

Depending on Session-Typed Processes

Bernardo Toninho and Nobuko Yoshida

Imperial College London, United Kingdom

Abstract. This work proposes a dependent type theory that combines functions and session-typed processes (with value dependencies) through a contextual monad, internalising typed processes in a dependently-typed λ -calculus. The proposed framework, by allowing session processes to depend on functions and vice-versa, enables us to specify and statically verify protocols where the choice of the next communication action can depend on specific values of received data. Moreover, the type theoretic nature of the framework endows us with the ability to internally describe and prove predicates on process behaviours. Our main results are type soundness of the framework, and a faithful embedding of the functional layer of the calculus within the session-typed layer, showcasing the expressiveness of dependent session types.

1 Introduction

Session types [25,14] are a typing discipline for communication protocols, whose simplicity provides an extensible framework that allows for integration with a variety of functional type features. One useful instance arising from the proof theoretic exploration of logical quantification is *value dependent session types* [26]. In this work, one can express properties of exchanged data in protocol specifications separately from communication, but *cannot* describe protocols where communication actions depend on the actual exchanged data (e.g. [17, § 2]). Moreover, it does not allow functions or values to depend on protocols (i.e. sessions) or communication, thus preventing reasoning about dependent process behaviours, exploring the proofs-as-programs paradigm of dependent type theory, e.g. [18,8].

Our work addresses the limitations of existing formulations of session types by proposing a type theory that integrates dependent functions *and* session types using a *contextual monad*. This monad internalises a session-typed calculus within a dependently-typed λ -calculus. By allowing session types to depend on λ -terms *and* λ -terms to depend on typed processes (using the monad), we are able to achieve heightened degrees of expressiveness. Exploiting the former direction, we enable writing actual data-dependent communication protocols. Exploiting the latter, we can define and *prove* properties of linearly-typed objects (i.e. processes) within our intuitionistic theory.

To informally demonstrate how our type theory goes beyond the state of the art in order to represent data-dependent protocols, consider the following session type (we write $\tau \wedge A$ for $\exists x:\tau.A$ where x does not occur in A and similarly $\tau \supset A$ for $\forall x:\tau.A$ when x is not free in A), $T \triangleq \text{Bool} \supset \oplus\{\mathbf{t} : \text{Nat} \wedge \mathbf{1}, \mathbf{f} : \text{Bool} \wedge \mathbf{1}\}$,

representable in existing session typing systems. The type T denotes a protocol which first, inputs a boolean and then either emits the label \mathbf{t} , which will be followed by an output of a natural number; or emits the label \mathbf{f} and a boolean. The intended protocol described by T is to take the \mathbf{t} branch if the received value is \mathbf{t} and the \mathbf{f} branch otherwise, which we can implement as Q with channel z typed by T as follows:

$$Q \triangleq z(x).\text{case } x \text{ of } (\text{true} \Rightarrow z.\mathbf{t}; z\langle 23 \rangle.\mathbf{0}, \text{false} \Rightarrow z.\mathbf{f}; z\langle \text{true} \rangle.\mathbf{0})$$

where $z(x).P$ denotes an input process, $z.\mathbf{t}$ is a process which selects label \mathbf{t} and $z\langle 23 \rangle.P$ is an output on z . However, since the specification is imprecise, process $z(x).\text{case } x \text{ of } (\text{false} \Rightarrow z.\mathbf{t}; z\langle 23 \rangle.\mathbf{0}, \text{true} \Rightarrow z.\mathbf{f}; z\langle \text{true} \rangle.\mathbf{0})$ is also a type-correct implementation of T that does not adhere to the intended protocol. Using our dependent type system, we can narrow the specification to guarantee that the desired protocol is precisely enforced. Consider the following definition of a session-type level conditional where we assume inductive definition and dependent pattern matching mechanisms (**stype** denotes the *kind* of session types):

$$\begin{aligned} \text{if} &:: \text{Bool} \rightarrow \text{stype} \rightarrow \text{stype} \rightarrow \text{stype} \\ \text{if true } AB &= A \quad \text{if false } AB = B \end{aligned}$$

The type-level function above case analyses the boolean and produces its first session type argument if the value is **true** and the second otherwise. We may now specify a session type that faithfully implements the protocol:

$$T' \triangleq \forall x:\text{Bool}.\text{if } x (\text{Nat} \wedge \mathbf{1}) (\text{Bool} \wedge \mathbf{1})$$

A process R implementing such a type on channel z is given below:

$$R \triangleq z(x).\text{case } x \text{ of } (\text{true} \Rightarrow z\langle 23 \rangle.\mathbf{0}, \text{false} \Rightarrow z\langle \text{true} \rangle.\mathbf{0})$$

Note that if we flip the two branches of the case analysis in R , the session is no longer typable with T' , ensuring that the protocol is implemented faithfully.

The example above illustrates a simple yet useful data-dependent protocol. When we further extend our dependent types with a *process* monad [30], where $\{c \leftarrow P \leftarrow \overline{u_j}; \overline{d_i}\}$ is a functional term denoting a process that may be *spawned* by other processes by instantiating the names in $\overline{u_j}$ and $\overline{d_i}$, we can provide more powerful reasoning on processes, enabling refined specifications through the use of type indices (i.e. type families) and an ability to internally specify and verify predicates on process behaviours. We also show that *all* functional types and terms can be faithfully embedded in the process layer using the dependently-typed sessions and process monads.

Contributions. § 2 introduces our dependent type theory, augmenting the example above by showing how we can reason about process behaviour using type families and dependently-typed functions (§ 2.3). We then establish the soundness of the theory (§ 2.4). § 3 develops a faithful embedding of the dependent function space in the process layer (Theorem 3.4). § 4 concludes with related work. Proofs, omitted definitions and additional examples can be found in [33].

Kinds	$K, K' ::= \text{type} \mid \text{stype} \mid \Pi x:\tau.K \mid \Pi t:K.K'$
Functional	$\tau, \sigma ::= \Pi x:\tau.\sigma \mid \lambda x:\tau.\sigma \mid \tau M \mid \{\overline{u_j:B_j}; \overline{d_i:A_i} \vdash c:A\} \mid \lambda t :: K.\tau \mid \tau \sigma$
Sessions	$A, B ::= !A \mid A \multimap B \mid A \otimes B \mid \forall x:\tau.A \mid \exists x:\tau.A \mid \mathbf{1}$ $\mid \&\{\overline{l_i:A_i}\} \mid \oplus\{\overline{l_i:A_i}\} \mid \lambda x:\tau.A \mid A M \mid \lambda t::K.A \mid AB$
Terms	$M, N ::= \lambda x:\tau.M \mid \{c \leftarrow P \leftarrow \overline{u_j}; \overline{d_i}\} \mid MN \mid x$
Processes	$P, Q ::= \overline{c}(d).P \mid (\nu c)P \mid c(x).P \mid c(M).P \mid !c(x).P$ $\mid c.\text{case}\{\overline{l_i} \Rightarrow P_i\} \mid c.l; P \mid [c \leftrightarrow d] \mid \mathbf{0} \mid c \leftarrow M \leftarrow \overline{u_j}; \overline{d_i}; Q$

Fig. 1. Syntax of Kinds, Types, Terms and Processes

2 A Dependent Type Theory of Processes

This section introduces our dependent type theory combining session-typed processes and functions. The theory is a generalisation of the line of work relating linear logic and session types [4,26,30], considering type-level functions and dependent kinds in an intensional type theory with full *mutual* dependencies between functions and processes. This generalisation enables us to express more sophisticated session types (such as those of § 1) and also to define and *prove* properties of processes expressed as type families with proofs as their inhabitants. We focus on the new rules and judgements, pointing the interested reader to [26,5,27] for additional details on the base theory.

2.1 Syntax

The calculus is stratified into two mutually dependent layers of processes and terms, which we often refer to as the *process* and *functional* layers, respectively. The syntax of the theory is given in Fig. 1 (we use x, y for variables ranging over terms and t for variables ranging over types).

Types and Kinds. The process layer is able to refer to terms of the functional layer via appropriate (dependently-typed) communication actions and through a *spawn* construct, allowing for processes encapsulated as functional values to be executed. Dually, the functional layer can refer to the process layer via a *contextual* monad [30] that internalises (open) typed processes as opaque functional values. This mutual dependency is also explicit in the type structure on several axes: process channel usages are typed by a language of session types, which specifies the communication protocols implemented on the used channels, extended with two dependent communication operations $\forall x:\tau.A$ and $\exists x:\tau.A$, where τ is a functional type and A is a session type in which x may occur. Moreover, we also extend the language of session types with type-level λ -abstraction over terms $\lambda x:\tau.A$ and session types $\lambda t::K.A$ (with the corresponding elimination forms AM and AB). As we show in § 1, the combination of these features allows for a new degree of expressiveness, enabling us to construct session types whose structure depends on previously communicated values.

The remaining session constructs are standard, following [5]: $!A$ denotes a *shared* session of type A that may be used an arbitrary (finite) number of times;

$A \multimap B$ represents a session offering to input a session of type A to then offer the session behaviour B ; $A \otimes B$ is the dual operator, denoting a session that outputs A and proceeds as B ; $\oplus\{\overline{l_i : A_i}\}$ and $\&\{\overline{l_i : A_i}\}$ represent internal and external labelled choice, respectively; $\mathbf{1}$ denotes the terminated session.

The functional layer is a λ -calculus with dependent functions $\Pi x:\tau.\sigma$, type-level λ -abstractions over terms and types (and respective type-level applications) and a *contextual monadic* type $\{u_j:\overline{B_j}; \overline{d_i:A_i} \vdash c:A\}$, denoting a (quoted) process offering session $c:A$ by using the *linear* sessions $\overline{d_i:A_i}$ and *shared* sessions $u_j:\overline{B_j}$ [30]. We often write $\{A\}$ for $\{\cdot; \vdash c:A\}$. The kinding system for our theory contains two base kinds **type** and **stype** of functional and session types, respectively. Type-level λ -abstractions require dependent kinds $\Pi x:\tau.K$ and $\Pi t::K.K'$, respectively. We note that the functional connectives form a standard dependent type theory [11,22].

Terms and Processes. Terms include the standard λ -abstractions $\lambda x:\tau.M$, applications MN and variables x . In order to internalise processes within the functional layer we make use of a monadic process wrapper, written $\{c \leftarrow P \leftarrow \overline{u_j}; \overline{d_i}\}$. In such a construct, the channels c , $\overline{u_j}$ and $\overline{d_i}$ are bound in P , where c is the session channel being offered and $\overline{u_j}$ and $\overline{d_i}$ are the session channels (linear and shared, respectively) being used. We write $\{c \leftarrow P \leftarrow \epsilon\}$ when P does not use any ambient channels, which we abbreviate to $\{P\}$.

The syntax of processes follows that of [5] extended with the monadic elimination form $c \leftarrow M \leftarrow \overline{u_j}; \overline{d_i}; Q$. Such a process construct denotes a term M that is to be evaluated to a monadic value of the form $\{c \leftarrow P \leftarrow \overline{u_j}; \overline{d_i}\}$ which will then be executed in parallel with Q , sharing with it a session channel c and using the provided channels $\overline{u_j}$ and $\overline{d_i}$. We write $c \leftarrow M \leftarrow \epsilon; Q$ when no channels are provided for the execution of M and often abbreviate this to $c \leftarrow M; Q$. The process $\overline{c}(d).P$ denotes the output of the *fresh* channel d along channel c with continuation P , which binds d ; $(\nu c)P$ denotes channel hiding, restricting the scope of c to P ; $c(x).P$ denotes an input along c , bound to x in P ; $c\langle M \rangle.P$ denotes the output of term M along c with continuation P ; $!c(x).P$ denotes a replicated input which spawns copies of P ; the construct $c.\text{case}\{\overline{l_i} \Rightarrow \overline{P_i}\}$ codifies a process that waits to receive some label l_j along c , with continuation P_j ; dually, $c.l; P$ denotes a process that emits a label l along c and continues as P ; $[c \leftrightarrow d]$ denotes a forwarder between c and d , which is operationally implemented as renaming; $P \mid Q$ denotes parallel composition and $\mathbf{0}$ the null process.

2.2 A Dependent Typing System

We now introduce our typing system, defined by a series of mutually inductive judgements, given in Fig. 2. We use Ψ to stand for a typing context for dependent λ -terms (i.e. assumptions of the form $x:\tau$ or $t :: K$, not subject to exchange), Γ for a typing context for *shared* sessions of the form $u:A$ (implicitly subject to weakening and contraction) and Δ for a linear context of sessions $x:A$. The context well-formedness judgments $\Psi \vdash$ and $\Psi; \Delta \vdash$ require that types and kinds (resp. session types) in Ψ (resp. Δ) are well-formed. The judgments $\Psi \vdash K$,

$\Psi \vdash$	Context Ψ is well-formed.
$\Psi; \Delta \vdash$	Context Δ is well-formed, under assumptions in Ψ .
$\Psi \vdash K$	K is a kind in context Ψ .
$\Psi \vdash \tau :: K$	τ is a (functional) type of kind K in context Ψ .
$\Psi \vdash A :: K$	A is a session type of kind K in context Ψ .
$\Psi \vdash M : \tau$	M has type τ in context Ψ .
$\Psi; \Gamma; \Delta \vdash P :: z:A$	P offers session $z:A$ when composed with processes offering sessions specified in Γ and Δ in context Ψ .
$\Psi \vdash K_1 = K_2$	Kinds K_1 and K_2 are equal.
$\Psi \vdash \tau = \sigma :: K$	Types τ and σ are equal of kind K .
$\Psi \vdash A = B :: K$	Session types A and B are equal of kind K .
$\Psi \vdash M = N : \tau$	Terms M and N are equal of type τ .
$\Psi \vdash \Delta = \Delta' :: \mathbf{stype}$	Contexts Δ and Δ' are equal, under the assumptions in Ψ .
$\Psi; \Gamma; \Delta \vdash P = Q :: z:A$	Processes P and Q are equal with typing $z:A$.

Fig. 2. Typing Judgements

$\Psi \vdash \tau :: K$ and $\Psi \vdash A :: K$ codify well-formedness of kinds, functional and session types (with kind K), respectively. Their rules are standard.

Typing. An excerpt of the typing rules for terms and processes is given in Fig. 3 and 4, respectively, noting that typing enforces types to be of base kind **type** (respectively **stype**). The rules for dependent functions are standard, including the type conversion rule which internalises definitional equality of types. We highlight the introduction rule for the monadic construct, which requires the appropriate session types to be well-formed and the process P to offer $c:A$ when provided with the appropriate session contexts.

In the typing rules for processes (Fig. 4), presented as a set of right and left rules (the former identifying how to *offer* a session of a given type and the latter how to use such a session), we highlight the rules for dependently-typed communication and monadic elimination (for type-checking purposes we annotate constructs with the respective dependent type – this is akin to functional type theories). To offer a session $c:\exists x:\tau.A$ we send a term M of type τ and then offer a session $c:A\{M/x\}$; dually, to use such a session we perform an input along c , bound to x in Q , warranting a use of c as a session of (open) type A . The rules for the universal are dual. Offering a session $c:\forall x:\tau.A$ entails receiving on c a term of type τ and offering $c:A$. Using a session of such a type requires sending along c a term M of type τ , warranting the use of c as a session of type $A\{M/x\}$.

The rule for the monadic elimination form requires that the term M be of the appropriate monadic type and that the provided channels \bar{u}_j and \bar{y}_i adhere to the typing specified in M 's type. Under these conditions, the process Q may then use the session c as session A . The type conversion rules reflect session type definitional equality in typing.

$$\begin{array}{c}
\text{(III)} \\
\frac{\Psi \vdash \tau :: \text{type} \quad \Psi, x:\tau \vdash M : \sigma}{\Psi \vdash \lambda x:\tau. M : \Pi x:\tau. \sigma} \\
\text{(IE)} \\
\frac{\Psi \vdash M : \Pi x:\tau. \sigma \quad \Psi \vdash N : \tau}{\Psi \vdash MN : \sigma\{N/x\}} \\
\text{({}I)} \\
\frac{\forall i, j. \Psi \vdash A_i, B_j :: \text{stype} \quad \Psi; \overline{u_j:B_j}; \overline{d_i:A_i} \vdash P :: c:A}{\Psi \vdash \{c \leftarrow P \leftarrow \overline{u_j}; \overline{d_i}\} : \{\overline{u_j:B_j}; \overline{d_i:A_i}\} \vdash c:A} \\
\text{(Conv)} \\
\frac{\Psi \vdash M : \tau \quad \Psi \vdash \tau = \sigma :: \text{type}}{\Psi \vdash M : \sigma}
\end{array}$$

Fig. 3. Typing for Terms (Excerpt – See [33])

$$\begin{array}{c}
\text{(\exists R)} \\
\frac{\Psi \vdash M:\tau \quad \Psi; \Gamma; \Delta \vdash P :: c:A\{M/x\}}{\Psi; \Gamma; \Delta \vdash c\langle M \rangle_{\exists x:\tau. A}. P :: c:\exists x:\tau. A} \\
\text{(\exists L)} \\
\frac{\Psi \vdash \tau :: \text{type} \quad \Psi, x:\tau; \Gamma; \Delta, c:A \vdash Q :: d:D}{\Psi; \Gamma; \Delta, c:\exists x:\tau. A \vdash c\langle x:\tau \rangle. Q :: d:D} \\
\text{(\forall R)} \\
\frac{\Psi \vdash \tau :: \text{type} \quad \Psi, x:\tau; \Gamma; \Delta \vdash P :: c:A}{\Psi; \Gamma; \Delta \vdash c\langle x:\tau \rangle. P :: c:\forall x:\tau. A} \\
\text{(\forall L)} \\
\frac{\Psi \vdash M:\tau \quad \Psi; \Gamma; \Delta, c:A\{M/x\} \vdash Q :: d:D}{\Psi; \Gamma; \Delta, c:\forall x:\tau. A \vdash c\langle M \rangle_{\forall x:\tau. A}. Q :: d:D} \\
\text{({}E)} \\
\frac{\Delta' = \overline{d_i : B_i} \quad \overline{u_j : C_j} \subseteq \Gamma \quad \Psi \vdash M : \{\overline{u_j : C_j}; \overline{d_i : B_i}\} \vdash c:A \quad \Psi; \Gamma; \Delta, c:A \vdash Q :: z:C}{\Psi; \Gamma; \Delta', \Delta \vdash c \leftarrow M \leftarrow \overline{u_j}; \overline{y_i}; Q :: z:C} \\
\text{(ConvR)} \\
\frac{\Psi; \Gamma; \Delta \vdash P :: z:A \quad \Psi \vdash A = B :: \text{stype}}{\Psi; \Gamma; \Delta \vdash P :: z:B} \\
\text{(ConvL)} \\
\frac{\Psi; \Gamma'; \Delta' \vdash P :: z:A \quad \Psi; \Gamma'; \Delta' = \Psi; \Gamma; \Delta}{\Psi; \Gamma; \Delta \vdash P :: z:A} \\
\text{(cut)} \\
\frac{\Psi; \Gamma; \Delta \vdash P :: c:A \quad \Psi; \Gamma; \Delta', c:A \vdash Q :: d:D}{\Psi; \Gamma; \Delta, \Delta' \vdash (\nu c)(P \mid Q) :: d:D}
\end{array}$$

Fig. 4. Typing for Processes (Excerpt – See [33])

Definitional Equality. The crux of any dependent type theory lies in its *definitional equality*. Type equality relies on equality of terms which, by including the monadic construct, necessarily relies on a notion of *process* equality.

Our presentation of an intensional definitional equality of terms follows that of [12], where we consider an intrinsically typed relation, including β and η conversion (similarly for type equality which includes β and η principles for the type-level λ -abstractions). An excerpt of the rules for term equality is given in Fig. 5. The remaining rules are congruence rules and closure under symmetry, reflexivity and transitivity. Rule (TMEq β) captures the β -reduction, identifying a λ -abstraction applied to an argument with the substitution of the argument in the function body (typed with the appropriately substituted type). We highlight rule (TMEq $\{\eta\}$), which codifies a general η -like principle for arbitrary terms of monadic type: We form a monadic term that applies the monadic elimination form to M , forwarding the result along the appropriate channel, which becomes a term equivalent to M .

$$\begin{array}{c}
\text{(TMEq}\beta\text{)} \\
\frac{\Psi \vdash \tau :: \text{type} \quad \Psi, x:\tau \vdash M : \sigma \quad \Psi \vdash N : \tau}{\Psi \vdash (\lambda x:\tau.M) N = M\{N/x\} : \sigma\{N/x\}} \quad \text{(TMEq}\eta\text{)} \\
\frac{\Psi \vdash M : \Pi x:\tau.\sigma \quad x \notin \text{fv}(M)}{\Psi \vdash \lambda x:\tau.M x = M : \Pi x:\tau.\sigma} \\
\text{(TMEq}\{\}\eta\text{)} \\
\frac{\Psi \vdash M : \{\overline{u_j:B_j}; \overline{d_i:A_i} \vdash c:A\}}{\Psi \vdash \{c \leftarrow (y \leftarrow M; \overline{u_j}; \overline{d_i}); [y \leftrightarrow c] \leftarrow \overline{u_j}; \overline{d_i}\} = M : \{\overline{u_j:B_j}; \overline{d_i:A_i} \vdash c:A\}}
\end{array}$$

Fig. 5. Definitional Equality of Terms (Excerpt – See [33])

$$\begin{array}{c}
\text{(PEqRed)} \frac{\Psi; \Gamma; \Delta \vdash P :: z:A \quad P \rightarrow Q \quad \Psi; \Gamma; \Delta \vdash Q :: z:A}{\Psi; \Gamma; \Delta \vdash P = Q :: z:A} \\
\text{(PEq}\forall\eta\text{)} \frac{}{\Psi; \Gamma; d:\forall x:\tau.A \vdash c(x).d(x).[d \leftrightarrow c] = [d \leftrightarrow c] :: c:\forall x:\tau.A} \\
\text{(PEqCC}\forall\text{)} \frac{\Psi; \Gamma; \Delta \vdash P :: d:B \quad \Psi, x:\tau; \Gamma; \Delta', d:B \vdash Q :: c:A}{\Psi; \Gamma; \Delta, \Delta' \vdash (\nu d)(P \mid c(x).Q) = c(x).(\nu d)(P \mid Q) :: c:\forall x:\tau.A}
\end{array}$$

Fig. 6. Definitional Equality of Processes (Excerpt – See [33])

Definitional equality of processes is summarised in Fig. 6. We rely on process reduction defined below. Definitional equality of processes consists of the usual congruence rules, (typed) reductions and the commuting conversions of linear logic and η -like principles, which allows for forwarding actions to be equated with the primitive syntactic forwarding construct. Commuting conversions amount to sound observational equivalences between processes [23], given that session composition requires name restriction (embodied by the (cut) rule): In rule (PEqCC \forall), either process can only be interacted with via channel c and so postponing actions of P to after the input on c (when reading the equality from left to right) cannot impact the process' observable behaviours. While P can in general interact with sessions in Δ (or with Q), these interactions are unobservable due to hiding in the (cut) rule.

Operational Semantics. The operational semantics for the λ -calculus is standard, noting that no reduction can take place inside monadic terms. The operational (reduction) semantics for processes is presented below where we omit closure under structural congruence and the standard congruence rules [4,26,30]. The last rule defines spawning a process in a monadic term.

$$\begin{array}{c}
c\langle M \rangle.P \mid c(x).Q \rightarrow P \mid Q\{M/x\} \quad \overline{c}\langle x \rangle.P \mid c(x).Q \rightarrow (\nu x)(P \mid Q) \\
!c(x).P \mid \overline{c}\langle x \rangle.Q \rightarrow !c(x).P \mid (\nu x)(P \mid Q) \quad c.\text{case}\{\overline{l_i} \Rightarrow \overline{P_i}\} \mid c.l_j; Q \rightarrow P_j \mid Q \quad (l_j \in \overline{l_i}) \\
(\nu c)(P \mid [c \leftrightarrow d]) \rightarrow P\{d/c\} \quad c \leftarrow \{c \leftarrow P \leftarrow \overline{u_j}; \overline{d_i}\} \leftarrow \overline{u_j}; \overline{d_i}; Q \rightarrow (\nu c)(P \mid Q)
\end{array}$$

2.3 Example – Reasoning about Processes using Dependent Types

The use of type indices (i.e. type families) in dependently typed frameworks adds information to types to produce more refined specifications. Our framework enables us to do this at the level of session types.

Consider a session type that “counts down” on a natural number (we assume inductive definitions and dependent pattern matching in the style of [22]):

$$\begin{aligned} \text{countDown} & \quad :: \Pi x:\text{Nat}.\text{stype} \\ \text{countDown}(\text{succ}(n)) & = \exists y:\text{Nat}.\text{countDown}(n) \\ \text{countDown } z & = \mathbf{1} \end{aligned}$$

The type family $\text{countDown}(n)$ denotes a session type that emits exactly n numbers and then terminates. We can now write a (dependently-typed) function that produces processes with the appropriate type, given a starting value:

$$\begin{aligned} \text{counter} & \quad : \Pi x:\text{Nat}.\{\text{countDown}(x)\} \\ \text{counter}(\text{succ}(n)) & = \{c \leftarrow c(\text{succ}(n)). d \leftarrow \text{counter}(n); [d \leftrightarrow c]\} \\ \text{counter } z & = \{c \leftarrow \mathbf{0}\} \end{aligned}$$

Note how the type of counter , through the type family countDown , allows us to specify exactly the number of times a value is sent. This is in sharp contrast with existing recursive (or inductive/coinductive [19,31]) session types, where one may only specify the general iterative nature of the behaviour (e.g. “send a number and then recurse or terminate”).

The example above relies on session type indexing in order to provide additional static guarantees about processes (and the functions that generate them). An alternative way is to consider “simply-typed” programs and then *prove* that they satisfy the desired properties, using the language itself. Consider a simply-typed version of the counter above described as an inductive session type:

$$\begin{aligned} \text{simpleCounterT} & \quad :: \text{stype} \\ \text{simpleCounterT} & = \oplus\{\text{dec} : \text{Nat} \wedge \text{simpleCounterT}, \text{done} : \mathbf{1}\} \end{aligned}$$

There are many processes that correctly implement such a type, given that the type merely dictates that the session outputs a natural number and recurses (modulo the dec and done messages to signal which branch of the internal choice is taken). A function that produces processes implementing such a session, mirroring those generated by the counter function above, is:

$$\begin{aligned} \text{simpleCounter} & \quad : \text{Nat} \rightarrow \{\text{simpleCounterT}\} \\ \text{simpleCounter}(\text{succ}(n)) & = \{c \leftarrow c.\text{dec}; (\nu d)(d(\text{succ}(n)).\mathbf{0} \mid d(x).c(x). \\ & \quad \quad \quad d \leftarrow \text{simpleCounter}(n); [d \leftrightarrow c])\} \\ \text{simpleCounter } z & = \{c \leftarrow c.\text{done}; \mathbf{0}\} \end{aligned}$$

The process generated by simpleCounter , after emitting the dec label, spawns a process in parallel that sends the appropriate number, which is received by the parallel thread and then sent along the session c . Despite its simplicity, this example embodies a general pattern where a computation is spawned in parallel (itself potentially spawning many other threads) and the main thread then waits for the result before proceeding.

While such a process is typable in most session typing frameworks, our theory enables us to *prove* that the counter implementation above indeed counts down

from a given number by defining an appropriate (inductive) type family, indexed by *monadic* values (i.e. processes):

```

corrCount ::  $\Pi x:\text{Nat}.\Pi y:\{\text{simpleCounterT}\}.\text{type}$ 
corrz    : corrCount z {c ← c.done; 0}
corrn    :  $\Pi n:\text{Nat}.\Pi P:\{\text{simpleCounterT}\}.\text{corrCount } n P \rightarrow$ 
           corrCount (succ(n)) {c ← c.dec; c(succ(n)).d ← P; [d ↔ c]}

```

The type family `corrCount`, indexed by a natural number and a monadic value implementing the session type `simpleCounter`, is defined via two constructors: `corrz`, which specifies that a correct 0 counter emits the `done` label and terminates; and `corrn`, which given a monadic value P that is a correct n -counter, defines that a correct $(n + 1)$ -counter emits $n + 1$ and then proceeds as P (modulo the label emission bookkeeping).

The proof of correctness of the `simpleCounter` function above is no more than a function of type $\Pi n:\text{Nat}.\text{corrCount } n$ (`simpleCounter(n)`), defined below:

```

prf          :  $\Pi n:\text{Nat}.\text{corrCount } n$  (simpleCounter(n))
prf z        = corrz
prf (succ(n)) = corrn n (simpleCounter(n)) (prf n)

```

Note that in this scenario, the processes that index the `corrCount` type family are not syntactically equal to those generated by `simpleCounter`, but rather *definitionally* equal.

Typically, the processes that index such correctness specifications tend to be distilled versions of the actual implementations, which often perform some additional internal computation or communication steps. Since our notion of definitional equality of processes includes reduction (and also commuting conversions which account for type-preserving shuffling of internal communication actions [27]), the type conversion mechanism allows us to use the techniques described above to generally reason about specification conformance.

2.4 Type Soundness of the Framework

The main goal of this section is to present type soundness of our framework through a subject reduction result. We also show that our theory guarantees progress for terms and processes. The development requires a series of auxiliary results (detailed in [33]) pertaining to the functional and process layers which are ultimately needed to produce the inversion properties necessary to establish subject reduction. We note that strong normalisation results for linear-logic based session processes are known in the literature [3,31,27], even in the presence of impredicative polymorphism, restricted corecursion and higher-order data. Such results are directly applicable to our work using appropriate semantics preserving type erasures.

In the remainder we often write $\Psi \vdash \mathcal{J}$ to stand for a well-formedness, typing or definitional equality judgment of the appropriate form. Similarly for $\Psi; \Gamma; \Delta \vdash \mathcal{J}$. We begin with the substitution property, which naturally holds for

both layers, noting that the dependently typed nature of the framework requires substitution in both contexts, terms and in types.

Lemma 2.1 (Substitution). *Let $\Psi \vdash M : \tau$:*

1. *If $\Psi, x:\tau, \Psi' \vdash \mathcal{J}$ then $\Psi, \Psi'\{M/x\} \vdash \mathcal{J}\{M/x\}$;*
2. *If $\Psi, x:\tau, \Psi'; \Gamma; \Delta \vdash \mathcal{J}$ then $\Psi, \Psi'\{M/x\}; \Gamma\{M/x\}; \Delta\{M/x\} \vdash \mathcal{J}\{M/x\}$*

Combining substitution with a form of functionality for typing (i.e. that substitution of equal terms in a well-typed term produces equal terms) and for equality (i.e. that substitution of equal terms in a definitional equality proof produces equal terms), we can establish validity for typing and equality, which is a form of internal soundness of the type theory stating that judgments are consistent across the different levels of the theory.

Lemma 2.2 (Validity for Typing). (1) *If $\Psi \vdash \tau :: K$ or $\Psi \vdash A :: K$ then $\Psi \vdash K$;* (2) *If $\Psi \vdash M : \tau$ then $\Psi \vdash \tau :: \text{type}$;* and (3) *If $\Psi; \Gamma; \Delta \vdash P :: z:A$ then $\Psi \vdash A :: \text{stype}$.*

Lemma 2.3 (Validity for Equality).

1. *If $\Psi \vdash M = N : \tau$ then $\Psi \vdash M : \tau, \Psi \vdash N : \tau$ and $\Psi \vdash \tau :: \text{type}$*
2. *If $\Psi \vdash \tau = \sigma :: K$ then $\Psi \vdash \tau :: K, \Psi \vdash \sigma :: K$ and $\Psi \vdash K$*
3. *If $\Psi \vdash A = B :: K$ then $\Psi \vdash A :: K, \Psi \vdash B :: K$ and $\Psi \vdash K$*
4. *If $\Psi \vdash K = K'$ then $\Psi \vdash K$ and $\Psi \vdash K'$*
5. *If $\Psi; \Gamma; \Delta \vdash P = Q :: z:A$ then $\Psi; \Gamma; \Delta \vdash P :: z:A, \Psi; \Gamma; \Delta \vdash Q :: z:A$ and $\Psi \vdash A :: \text{stype}$*

With these results we establish the appropriate inversion and injectivity properties which then enable us to show unicity of types (and kinds).

Theorem 2.4 (Unicity of Types and Kinds).

1. *If $\Psi \vdash M : \tau$ and $\Psi \vdash M : \tau'$ then $\Psi \vdash \tau = \tau' :: \text{type}$*
2. *If $\Psi \vdash \tau :: K$ and $\Psi \vdash \tau :: K'$ then $\Psi \vdash K = K'$*
3. *If $\Psi; \Gamma; \Delta \vdash P :: z:A$ and $\Psi; \Gamma; \Delta \vdash P :: z:A'$ then $\Psi \vdash A = A' :: \text{stype}$*
4. *If $\Psi \vdash A :: K$ and $\Psi \vdash A :: K'$ then $\Psi \vdash K = K'$*

All the results above, combined with the process-level properties established in [28,27,5] enable us to show the following:

Theorem 2.5 (Subject Reduction – Terms). *If $\Psi \vdash M : \tau$ and $M \rightarrow M'$ then $\Psi \vdash M' : \tau$*

Theorem 2.6 (Subject Reduction – Processes). *If $\Psi; \Gamma; \Delta \vdash P :: z:A$ and $P \rightarrow P'$ then $\exists Q$ such that $P' \equiv Q$ and $\Psi; \Gamma; \Delta \vdash Q :: z:A$*

Theorem 2.7 (Progress – Terms). *If $\Psi \vdash M : \tau$ then either M is a value or $M \rightarrow M'$*

As common in logical-based session type theories, typing enforces a strong notion of *global* progress which states that closed processes that are waiting to perform communication actions cannot get stuck (this relies on a notion of *live* process, defined as $\text{live}(P)$ iff $P \equiv (\nu \tilde{n})(\pi.Q \mid R)$ for some process R , sequence of names \tilde{n} and a non-replicated guarded process $\pi.Q$). We note that the restricted typing for P is without loss of generality, due to the (cut) rule.

Theorem 2.8 (Progress – Processes). *If $\Psi; \cdot; \cdot \vdash P :: c:1$ and $\text{live}(P)$ then $\exists Q$ such that $P \rightarrow Q$*

3 Embedding the Functional Layer in the Process Layer

Having introduced our type theory and showcased some of its informal expressiveness in terms of the ability to specify and *statically* verify true data dependent protocols, as well as the ability to prove properties of processes, we now develop a formal expressiveness result for our theory, showing that the process level type constructs are able to encode the dependently-typed functional layer, faithfully preserving type dependencies.

Specifically, we show that (1) the type-level constructs in the functional layer can be represented by those in the process layer combined with the contextual monad type, and (2) all term level constructs can be represented by session-typed processes that exchange monadic values. Thus, we show that both λ -abstraction and application can be eliminated while still preserving non-trivial type dependencies. Crucially, we note that the monadic construct *cannot* be fully eliminated due to the cross-layer nature of session type dependencies: In the process layer, simply-kinded dependent types (i.e. types with kind **stype**) are of the form $\forall x:\tau.A$ where τ is of kind **type** and A of kind **stype** (where x may occur). Operationally, such a session denotes an input of some term M of type τ with a continuation of type $A\{M/x\}$. Thus, to faithfully encode type dependencies we cannot represent such a type with a non-dependently typed input (e.g. a type of the form $A \multimap B$).

3.1 The Embedding

A first attempt. Given the observation above, a seemingly reasonable option would be to attempt an encoding that maintains monadic objects solely at the level of type indices and then exploits Girard’s encoding [9] of function types $\tau \rightarrow \sigma$ as $!\llbracket \tau \rrbracket \rightarrow \llbracket \sigma \rrbracket$, which is adequate for session-typed processes [29]. Thus a candidate encoding for the type $\Pi x:\tau.\sigma$ would be $\forall x:\{\llbracket \tau \rrbracket\}.!\llbracket \tau \rrbracket \multimap \llbracket \sigma \rrbracket$, where $\llbracket - \rrbracket$ denotes our encoding on types. If we then consider the encoding at the level of terms, typing dictates the following (we write $\llbracket M \rrbracket_z$ for the process encoding of $M : \tau$, where z is the session channel along which one may observe the “result” of the encoding, typed with $\llbracket \tau \rrbracket$):

$$\begin{aligned} \llbracket \lambda x:\tau.M \rrbracket_z &\triangleq z(x).z(x').\llbracket M \rrbracket_z \\ \llbracket M N \rrbracket_z &\triangleq (\nu x)(\llbracket M \rrbracket_x \mid x(\{\llbracket N \rrbracket_y\}).\bar{x}(x').(!x'(y).\llbracket N \rrbracket_y \mid [x \leftrightarrow z]) \end{aligned}$$

Kind:

$$\begin{array}{llll} \llbracket \text{type} \rrbracket & \triangleq \text{stype} & \llbracket \text{stype} \rrbracket & \triangleq \text{stype} \\ \llbracket \Pi x:\tau.K \rrbracket & \triangleq \Pi x:\{\llbracket \tau \rrbracket\}.\llbracket K \rrbracket & \llbracket \Pi t :: K_1.K_2 \rrbracket & \triangleq \Pi t::\llbracket K_1 \rrbracket.\llbracket K_2 \rrbracket \end{array}$$

Functional:

$$\begin{array}{llll} \llbracket \Pi x:\tau.\sigma \rrbracket & \triangleq \forall x:\{\llbracket \tau \rrbracket\}.\llbracket \sigma \rrbracket & \llbracket \{\overline{u_j:B_j}; \overline{d_i:B_i} \vdash c:A\} \rrbracket & \triangleq !\llbracket B_j \rrbracket \multimap \overline{\llbracket B_i \rrbracket} \multimap \llbracket A \rrbracket \\ \llbracket \lambda x:\tau.\sigma \rrbracket & \triangleq \lambda x:\{\llbracket \tau \rrbracket\}.\llbracket \sigma \rrbracket & \llbracket \tau M \rrbracket & \triangleq \llbracket \tau \rrbracket \{\llbracket M \rrbracket_c\} \\ \llbracket \lambda t::K.\tau \rrbracket & \triangleq \lambda t::\llbracket K \rrbracket.\llbracket \tau \rrbracket & \llbracket \tau \sigma \rrbracket & \triangleq \llbracket \tau \rrbracket \llbracket \sigma \rrbracket \end{array}$$

Session:

$$\begin{array}{llll} \llbracket \forall x:\tau.A \rrbracket & \triangleq \forall x:\{\llbracket \tau \rrbracket\}.\llbracket A \rrbracket & \llbracket \exists x:\tau.A \rrbracket & \triangleq \exists x:\{\llbracket \tau \rrbracket\}.\llbracket A \rrbracket \\ \llbracket \lambda x:\tau.A \rrbracket & \triangleq \lambda x:\{\llbracket \tau \rrbracket\}.\llbracket A \rrbracket & \llbracket A M \rrbracket & \triangleq \llbracket A \rrbracket \{\llbracket M \rrbracket_c\} \end{array}$$

Terms:

$$\begin{array}{ll} \llbracket \lambda x:\tau.M \rrbracket_z \triangleq z(x:\{\llbracket \tau \rrbracket\}).\llbracket M \rrbracket_z & \llbracket M N \rrbracket_z \triangleq (\nu x)(\llbracket M \rrbracket_x \mid x\langle\{\llbracket N \rrbracket_y\}\rangle.[x \leftrightarrow z]) \\ \llbracket x \rrbracket_z \triangleq y \leftarrow x; [y \leftrightarrow z] & \llbracket \{z \leftarrow P \leftarrow \overline{u_j}; \overline{d_i}\} \rrbracket_z \triangleq z(u_0).\dots.z(u_j).z(d_0).\dots.z(d_n).\llbracket P \rrbracket \end{array}$$

Processes:

$$\begin{array}{llll} \llbracket (\nu x)(P \mid Q) \rrbracket & \triangleq (\nu x)(\llbracket P \rrbracket \mid \llbracket Q \rrbracket) & \llbracket \mathbf{0} \rrbracket \triangleq \mathbf{0} & \llbracket \overline{x}\langle y \rangle.(P \mid Q) \rrbracket \triangleq \overline{x}\langle y \rangle.(\llbracket P \rrbracket \mid \llbracket Q \rrbracket) \\ \llbracket x(M).P \rrbracket & \triangleq x\langle\{\llbracket M \rrbracket_y\}\rangle.\llbracket P \rrbracket & \llbracket x(y).P \rrbracket \triangleq x(y).\llbracket P \rrbracket & \\ \llbracket c \leftarrow M \leftarrow \overline{u_j}; \overline{y_i}; Q \rrbracket & \triangleq (\nu c)(\llbracket M \rrbracket_c \mid \overline{c}\langle v_1 \rangle.(\overline{u_1}\langle a_1 \rangle.[a_1 \leftrightarrow v_1] \mid \dots \mid \\ & \overline{c}\langle d_1 \rangle.[y_1 \leftrightarrow d_1] \mid \dots \mid \overline{c}\langle d_n \rangle.[y_n \leftrightarrow d_n] \mid \llbracket Q \rrbracket) \dots) \end{array}$$

Fig. 7. An embedding of dependent functions into processes

However, this candidate encoding breaks down once we consider definitional equality. Specifically, compositionality (i.e. the relationship between $\llbracket M\{N/x\} \rrbracket_z$ and the encoding of N substituted in that of M) requires us to relate $\llbracket M\{N/x\} \rrbracket_z$ with $(\nu x)(\llbracket M \rrbracket_z \{\{\llbracket N \rrbracket_y\}/x\} \mid !x'(y).\llbracket N \rrbracket_y)$, which relies on reasoning up-to *observational equivalence* of processes, a much stronger relation than our notion of definitional equality. Therefore it is *fundamentally* impossible for such an encoding to preserve our definitional equality, and thus it cannot preserve typing in the general case.

A faithful embedding. We now develop our embedding of the functional layer into the process layer which is compatible with definitional equality. Our target calculus is reminiscent of a higher-order (in the sense of higher-order processes [24]) session calculus [20]. Our encoding $\llbracket - \rrbracket$ is inductively defined on kinds, types, session types, terms and processes. As usual in process encodings of the λ -calculus, the encoding of a term M is indexed by a result channel z , written $\llbracket M \rrbracket_z$, where the behaviour of M may be observed.

The embedding is presented in Fig. 7, noting that the encoding extends straightforwardly to typing contexts, where functional contexts $\Psi, x:\tau$ are mapped to $\{\llbracket \Psi \rrbracket\}, x:\{\llbracket \tau \rrbracket\}$. The mapping of base kinds is straightforward. Dependent kinds $\Pi x:\tau.K$ rely on the monad for well-formedness and are encoded as (session) kinds of the form $\Pi x:\{\llbracket \tau \rrbracket\}.\llbracket K \rrbracket$. The higher-kinded types in the functional layer are translated to the corresponding type-level constructs of the process layer where all objects that must be **type**-kinded rely on the monad to satisfy this constraint. For instance, $\lambda x:\tau.\sigma$ is mapped to the session-type abstraction $\lambda x:\{\llbracket \tau \rrbracket\}.\llbracket \sigma \rrbracket$ and the type-level application τM is translated to $\llbracket \tau \rrbracket \{\llbracket M \rrbracket_c\}$.

Given the observation above on embedding the dependent function type $\Pi x:\tau.\sigma$, we translate it directly to $\forall x:\llbracket\tau\rrbracket.\llbracket\sigma\rrbracket$, that is, functions from τ to σ are mapped to sessions that input *processes* implementing $\llbracket\tau\rrbracket$ and then behave as $\llbracket\sigma\rrbracket$ accordingly. The encoding for monadic types simply realises the contextual nature of the monad by performing a sequence of inputs of the appropriate types (with the shared sessions being of ! type).

The mutually dependent nature of the framework requires us to extend the mapping to the process layer. Session types are mapped homomorphically (e.g. $\llbracket A \multimap B \rrbracket \triangleq \llbracket A \rrbracket \multimap \llbracket B \rrbracket$) with the exception of dependent inputs and outputs which rely on the monad, similarly for type-level functions and application.

The encoding of λ -terms is guided by the embedding for types: the abstraction $\lambda x:\tau.M$ is mapped to an input of a term of type $\{\llbracket\tau\rrbracket\}$ with continuation $\llbracket M \rrbracket_z$; application $M N$ is mapped to the composition of the encoding of M on a fresh name x with the corresponding output of $\{\llbracket N \rrbracket_y\}$, which is then forwarded to the result channel z ; monadic expressions are translated to the appropriate sequence of inputs, as dictated by the translation of the monadic type; and, the translation of variables makes use of the monadic elimination form (since the encoding enforces variables to always be of monadic type) combined with forwarding to the appropriate result channel.

The mapping for processes is mostly homomorphic, using the monad constructor as needed. The only significant exception is the encoding for monadic elimination which must provide the encoded monadic term $\llbracket M \rrbracket_c$ with the necessary channels. Since the session calculus does not support communication of free names this is achieved by a sequence of outputs of fresh names combined with forwarding of the appropriate channel. To account for replicated sessions we must first trigger the replication via an output which is then forwarded accordingly.

We can illustrate our encoding via a simple example of an encoded function (we omit type annotations for conciseness):

$$\begin{aligned} \llbracket (\lambda x.x) (\lambda x.\lambda y.y) \rrbracket_z &= (\nu c)(\llbracket \lambda x.x \rrbracket_c \mid c\langle \{\llbracket \lambda x.\lambda y.y \rrbracket_w\} \rangle.[c \leftrightarrow z]) = \\ &(\nu c)(c(x).y \leftarrow x; [y \leftrightarrow c] \mid c\langle \{w(x).w(y).d \leftarrow y; [d \leftrightarrow w]\} \rangle.[c \leftrightarrow z]) \\ &\rightarrow^+ z(x).z(y).d \leftarrow y; [d \leftrightarrow z] = \llbracket \lambda x.\lambda y.y \rrbracket_z \end{aligned}$$

3.2 Properties of the Embedding

We now state the key properties satisfied by our embedding, ultimately resulting in type preservation and operational correspondence. For conciseness, in the statements below we list only the cases for terms and processes, omitting those for types and kinds (see [33]). The key property that is needed is a notion of compositionality, which unlike in the sketch above no longer falls outside of definitional equality.

Lemma 3.1 (Compositionality).

1. $\Psi; \Gamma; \Delta \vdash \llbracket M\{N/x\} \rrbracket_z = \llbracket M \rrbracket_z\{\{\llbracket N \rrbracket_y\}/x\} :: z:\llbracket A\{N/x\} \rrbracket$
2. $\Psi; \Gamma; \Delta \vdash \llbracket P\{M/x\} \rrbracket :: z:\llbracket A\{M/x\} \rrbracket$ iff $\Psi; \Gamma; \Delta \vdash \llbracket P \rrbracket\{\{\llbracket M \rrbracket_c\}/x\} :: z:\llbracket A \rrbracket\{\{\llbracket M \rrbracket_c\}/x\}$

Given the dependently typed nature of the framework, establishing the key properties of the encoding must be done simultaneously (relying on some auxiliary results – see [33]).

Theorem 3.2 (Preservation of Equality).

1. If $\Psi \vdash M = N : \tau$ then $\{\llbracket \Psi \rrbracket\}; \cdot; \cdot \vdash \llbracket M \rrbracket_z = \llbracket N \rrbracket_z :: z : \llbracket \tau \rrbracket$
2. If $\Psi; \Gamma; \Delta \vdash P = Q :: z : A$ then $\{\llbracket \Psi \rrbracket\}; \llbracket \Gamma \rrbracket; \llbracket \Delta \rrbracket \vdash \llbracket P \rrbracket = \llbracket Q \rrbracket :: z : \llbracket A \rrbracket$

Theorem 3.3 (Preservation of Typing).

1. If $\Psi \vdash M : \tau$ then $\{\llbracket \Psi \rrbracket\}; \cdot; \cdot \vdash \llbracket M \rrbracket_z :: z : \llbracket \tau \rrbracket$
2. If $\Psi; \Gamma; \Delta \vdash P :: z : A$ then $\{\llbracket \Psi \rrbracket\}; \llbracket \Gamma \rrbracket; \llbracket \Delta \rrbracket \vdash \llbracket P \rrbracket :: z : \llbracket A \rrbracket$

Theorem 3.4 (Operational Correspondence). *If $\Psi; \Gamma; \Delta \vdash P :: z : A$ and $\Psi \vdash M : \tau$ then:*

1. (a) If $P \rightarrow P'$ then $\llbracket P \rrbracket \rightarrow^+ Q$ with $\{\llbracket \Psi \rrbracket\}; \llbracket \Gamma \rrbracket; \llbracket \Delta \rrbracket \vdash Q = \llbracket P' \rrbracket :: z : \llbracket A \rrbracket$ and
(b) if $\llbracket P \rrbracket \rightarrow P'$ then $P \rightarrow^+ Q$ with $\{\llbracket \Psi \rrbracket\}; \llbracket \Gamma \rrbracket; \llbracket \Delta \rrbracket \vdash P' = \llbracket Q \rrbracket :: z : \llbracket A \rrbracket$
2. (a) If $M \rightarrow M'$ then $\llbracket M \rrbracket_z \rightarrow^+ N$ with $\{\llbracket \Psi \rrbracket\}; \cdot; \cdot \vdash N = \llbracket M' \rrbracket_z :: z : \llbracket \tau \rrbracket$ and
(b) if $\llbracket M \rrbracket_z \rightarrow P$ then $M \rightarrow N$ with $\{\llbracket \Psi \rrbracket\}; \cdot; \cdot \vdash \llbracket N \rrbracket_z = P :: z : \llbracket \tau \rrbracket$

In Theorem 3.4, (a) is commonly referred to as operational completeness, with (b) establishing soundness. As exemplified above, our encoding satisfies a very precise operational correspondence with the original λ -terms.

4 Related and Future Work

Enriching Session Types via Type Structure. Exploiting the linear logical foundations of session types, [26] considers a form of value dependencies where session types can state properties of exchanged data values, while the work [30] introduces the contextual monad in a simply-typed setting. Our development not only subsumes these two works, but goes beyond simple value dependencies by extending to a richer type structure and integrating dependencies with the contextual monad. Recently, [1] considers a non-conservative extension of linear logic-based session types with sharing, allowing true non-determinism. Their work includes dependent quantifications with shared channels, but their type syntax does *not* include free type variables, so the actual type dependencies do not arise (see [1, 37:8]). Thus none of the examples in this paper can be represented in [1]. The work [17] studies gradual session types. To the best of our knowledge, the main example in [17, § 2] is *statically* representable in our framework as in the example of § 1, where protocol actions depend on values that are communicated (or passed as function arguments).

In the context of multiparty session types, the theory of multiparty indexed session types is studied in [7], and implemented in a protocol description language [21]. The main aim of these works is to use indexed types to represent an arbitrary number of session *participants*. The work [32] extends [26] to multiparty sessions in order to treat value dependency across multiple participants. Extending our framework to multiparty [16] or non-logic based session types [15] is an interesting future topic.

Combining Linear and Dependent Types. Many works have studied the various challenges of integrating linearity in dependent functional type theories. We focus on the most closely related works. The work [6] introduced the Linear Logical Framework (LLF), integrating linearity with the LF [11] type theory, which was later extended to the Concurrent Logical Framework (CLF) [34], accounting for further linear connectives. Their theory is representable in our framework through the contextual monad (encompassing full intuitionistic linear logic), depending on linearly-typed processes that can express dependently typed functions (§ 3).

The work of [18] integrates linearity with type dependencies by extending LNL [2]. Their work is aimed at reasoning about imperative programs using a form of Hoare triples, requiring features that we do not study in this work such as proof irrelevance and computationally irrelevant quantification. Formally, their type theory is extensional which introduces significant technical differences from our intensional type theory, such as a realisability model in the style of NuPRL [10] to establish consistency.

Recently, [8] proposed an extension of LLF with first-class contexts (which may contain both linear and unrestricted hypotheses). While the contextual aspects of their theory are reminiscent of our contextual monad, their framework differs significantly from ours, since it is designed to enable higher-order abstract syntax (commonplace in the LF family of type theories), focusing on a type system for canonical LF objects with a meta-language that includes contexts and context manipulation. They do not consider additives since their integration with first-class contexts can break canonicity.

While none of the above works considers processes as primitive, their techniques should be useful for, e.g. developing algorithmic type-checking and integrating inductive and coinductive session types based on [27,31,19].

Dependent Types and Higher-Order π -calculus. The work [36] studies a form of dependent types where the type of processes takes the form of a mapping Δ from channels x to channel types T representing an interface of process P . The dependency is specified as $\Pi(x:T)\Delta$, representing a channel abstraction of the environment. This notion is extended to an existential channel dependency type $\Sigma(x:T)\Delta$ to address fresh name creation [35,13]. Combining our process monad with dependent types can be regarded as an “interface” which describes explicit channel usages for processes. The main differences are (1) our dependent types are more general, treating full dependent families including terms and processes in types, while [36,35,13] study only channel dependency to environments (i.e. neither terms nor processes appear in types, only channels); and (2) our calculus emits only fresh names, not needing to handle the complex scoping mechanism treated in [35,13]. In this sense, the process monad provides an elegant framework to handle higher-order computations and assign non-trivial types to processes.

Acknowledgements. The authors would like to thank the anonymous reviews for their comments and suggestions. This work is partially supported by EPSRC EP/K034413/1, EP/K011715/1, EP/L00058X/1, EP/N027833/1 and EP/N028201/1.

References

1. Balzer, S., Pfenning, F.: Manifest sharing with session types. *PACMPL* 1(ICFP), 37:1–37:29 (2017)
2. Benton, N.: A mixed linear and non-linear logic: Proofs, terms and models (extended abstract). In: *CSL*. pp. 121–135 (1994)
3. Caires, L., Pérez, J.A., Pfenning, F., Toninho, B.: Behavioral polymorphism and parametricity in session-based communication. In: *ESOP 2013*. pp. 330–349 (2013)
4. Caires, L., Pfenning, F.: Session types as intuitionistic linear propositions. In: *CONCUR 2010*. pp. 222–236 (2010)
5. Caires, L., Pfenning, F., Toninho, B.: Linear logic propositions as session types. *Mathematical Structures in Computer Science* 26(3), 367–423 (2016)
6. Cervesato, I., Pfenning, F.: A linear logical framework. *Inf. Comput.* 179(1), 19–75 (2002)
7. Deniérou, P., Yoshida, N., Bejleri, A., Hu, R.: Parameterised multiparty session types. *Logical Methods in Computer Science* 8(4) (2012), [http://dx.doi.org/10.2168/LMCS-8\(4:6\)2012](http://dx.doi.org/10.2168/LMCS-8(4:6)2012)
8. Georges, A.L., Murawska, A., Otis, S., Pientka, B.: LINCX: A linear logical framework with first-class contexts. In: *ESOP*. pp. 530–555 (2017)
9. Girard, J.: Linear logic. *Theor. Comput. Sci.* 50, 1–102 (1987)
10. Harper, R.: Constructing type systems over an operational semantics. *Journal of Symbolic Computation* 14(1), 71 – 84 (1992)
11. Harper, R., Honsell, F., Plotkin, G.D.: A framework for defining logics. *J. ACM* 40(1), 143–184 (1993)
12. Harper, R., Pfenning, F.: On equivalence and canonical forms in the LF type theory. *ACM Trans. Comput. Log.* 6(1), 61–101 (2005)
13. Hennessy, M., Rathke, J., Yoshida, N.: safeDpi: a language for controlling mobile code. *Acta Inf.* 42(4-5), 227–290 (2005)
14. Honda, K., Vasconcelos, V.T., Kubo, M.: Language primitives and type discipline for structured communication-based programming. In: *ESOP’98*. pp. 122–138 (1998)
15. Honda, K., Vasconcelos, V.T., Kubo, M.: Language primitives and type disciplines for structured communication-based programming. In: *ESOP’98*. vol. 1381, pp. 22–138 (1998)
16. Honda, K., Yoshida, N., Carbone, M.: Multiparty asynchronous session types. In: *POPL’08*. pp. 273–284 (2008)
17. Igarashi, A., Thiemann, P., Vasconcelos, V.T., Wadler, P.: Gradual session types. *PACMPL* 1(ICFP), 38:1–38:28 (2017)
18. Krishnaswami, N.R., Pradic, P., Benton, N.: Integrating linear and dependent types. In: *POPL’15*. pp. 17–30 (2015)
19. Lindley, S., Morris, J.G.: Talking bananas: structural recursion for session types. In: *ICFP 2016*. pp. 434–447 (2016)
20. Mostrous, D., Yoshida, N.: Two session typing systems for higher-order mobile processes. In: *TLCA07*. pp. 321–335 (2007)
21. Ng, N., Yoshida, N.: Pabble: parameterised Scribble. *Service Oriented Computing and Applications* 9(3-4), 269–284 (2015)
22. Norell, U.: Towards a practical programming language based on dependent type theory. Ph.D. thesis, Department of Computer Science and Engineering, Chalmers University of Technology (2007)

23. Pérez, J.A., Caires, L., Pfenning, F., Toninho, B.: Linear logical relations for session-based concurrency. In: ESOP. pp. 539–558 (2012)
24. Sangiorgi, D., Walker, D.: The pi-calculus: A theory of mobile processes. C.U.P (2001)
25. Takeuchi, K., Honda, K., Kubo, M.: An interaction-based language and its typing system. In: PARLE'94. pp. 398–413 (1994)
26. Toninho, B., Caires, L., Pfenning, F.: Dependent session types via intuitionistic linear type theory. In: PPDP'11. pp. 161–172 (2011)
27. Toninho, B.: A Logical Foundation for Session-based Concurrent Computation. Ph.D. thesis, Carnegie Mellon University and New University of Lisbon (2015)
28. Toninho, B., Caires, L., Pfenning, F.: Dependent session types via intuitionistic linear type theory. Tech. Rep. CMU-CS-11-139, School of Computer Science, Carnegie Mellon University (2011)
29. Toninho, B., Caires, L., Pfenning, F.: Functions as session-typed processes. In: FOSSACS 2012. pp. 346–360 (2012)
30. Toninho, B., Caires, L., Pfenning, F.: Higher-order processes, functions, and sessions: A monadic integration. In: ESOP. pp. 350–369 (2013)
31. Toninho, B., Caires, L., Pfenning, F.: Corecursion and non-divergence in session-typed processes. In: TGC 2014. pp. 159–175 (2014)
32. Toninho, B., Yoshida, N.: Certifying data in multiparty session types. *Journal of Logical and Algebraic Methods in Programming* 90(C), 61–83 (2017)
33. Toninho, B., Yoshida, N.: Depending on session-typed processes. Tech. Rep. TBD, TBD (2017)
34. Watkins, K., Cervesato, I., Pfenning, F., Walker, D.: A concurrent logical framework: The propositional fragment. In: TYPES'03. pp. 355–377 (2003)
35. Yoshida, N.: Channel dependent types for higher-order mobile processes. In: Proceedings of the 31st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2004, Venice, Italy, January 14-16, 2004. pp. 147–160 (2004)
36. Yoshida, N., Hennessy, M.: Assigning types to processes. *Inf. Comput.* 174(2), 143–179 (2002)