

University of Groningen

Typed Non-determinism in Functional and Concurrent Calculi

van den Heuvel, Bas; Paulus, Joseph W.N.; Nantes-Sobrinho, Daniele; Pérez, Jorge A.

Published in:
Programming Languages and Systems - 21st Asian Symposium, APLAS 2023, Proceedings

DOI:
[10.1007/978-981-99-8311-7_6](https://doi.org/10.1007/978-981-99-8311-7_6)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2023

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):
van den Heuvel, B., Paulus, J. W. N., Nantes-Sobrinho, D., & Pérez, J. A. (2023). Typed Non-determinism in Functional and Concurrent Calculi. In C.-K. Hur (Ed.), *Programming Languages and Systems - 21st Asian Symposium, APLAS 2023, Proceedings* (pp. 112-132). (Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics); Vol. 14405 LNCS). Springer Science and Business Media Deutschland GmbH. https://doi.org/10.1007/978-981-99-8311-7_6

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.




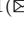

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.



Typed Non-determinism in Functional and Concurrent Calculi

Bas van den Heuvel¹, Joseph W. N. Paulus^{1,2},
Daniele Nantes-Sobrinho^{3,4}, and Jorge A. Pérez¹

¹ University of Groningen, Groningen, The Netherlands
`j.a.perez@rug.nl`

² University of Oxford, Oxford, UK

³ University of Brasília, Brasília, Brazil

⁴ Imperial College London, London, UK

Abstract. We study functional and concurrent calculi with non-determinism, along with type systems to control resources based on linearity. The interplay between non-determinism and linearity is delicate: careless handling of branches can discard resources meant to be used exactly once. Here we go beyond prior work by considering non-determinism in its standard sense: once a branch is selected, the rest are discarded. Our technical contributions are three-fold. First, we introduce a π -calculus with non-deterministic choice, governed by session types. Second, we introduce a resource λ -calculus, governed by intersection types, in which non-determinism concerns fetching of resources from bags. Finally, we connect our two typed non-deterministic calculi via a correct translation.

1 Introduction

In this paper, we present new formulations of typed programming calculi with *non-determinism*. A classical ingredient of models of computation, non-determinism brings flexibility and generality in specifications. In process calculi such as CCS and the π -calculus, one source of non-determinism is choice, which is typically *non-confluent*: that is, given $P + Q$, we have either $P + Q \longrightarrow P$ or $P + Q \longrightarrow Q$. Thus, committing to a branch entails discarding the rest.

We study non-determinism as a way of increasing the expressivity of typed calculi in which resource control is based on *linearity*. The interplay between non-determinism and linearity is delicate: a careless discarding of branches can jeopardize resources meant to be used exactly once. On the concurrent side, we consider the π -calculus, the paradigmatic model of concurrency [27]. We focus on π -calculi with *session types* [14, 15], in which linear logic principles ensure communication correctness: here the resources are names that perform session protocols; they can be *unrestricted* (used multiple times) and *linear* (used exactly once). To properly control resources, non-confluent non-determinism is confined to unrestricted names; linear names can only perform deterministic choices.

In this context, considering *confluent* forms of non-determinism can be appealing. Intuitively, such formulations allow all branches to proceed independently: given $P_1 \longrightarrow Q_1$ and $P_2 \longrightarrow Q_2$, then $P_1 + P_2 \longrightarrow Q_1 + P_2$ and

$P_1 + P_2 \longrightarrow P_1 + Q_2$. Because confluent non-determinism does not discard branches, it is compatible with a resource-conscious view of computation.

Confluent non-determinism has been studied mostly in the functional setting; it is present, e.g., in Pagani and Ronchi della Rocca’s resource λ -calculus [21] and in Ehrhard and Regnier’s differential λ -calculus [10]. In [21], non-determinism resides in the application of a term M to a *bag* of available resources C ; a β -reduction applies M to a resource *non-deterministically fetched* from C . Confluent non-deterministic choice is also present in the session-typed π -calculus by Caires and Pérez [5], where it expresses a choice between different implementations of the same session protocols, which are all *non-deterministically available*—they may be available but may also *fail*. In their work, a Curry-Howard correspondence between linear logic and session types (*propositions-as-sessions*’ [6, 31]) ensures confluence, protocol fidelity, and deadlock-freedom. Paulus et al. [22] relate functional and concurrent calculi with confluent non-determinism: they give a translation of a resource λ -calculus into the session π -calculus from [5], in the style of Milner’s *functions-as-processes*’ [17].

Although results involving confluent non-determinism are most significant, usual (non-confluent) non-determinism remains of undiscussed convenience in formal modeling; consider, e.g., specifications of distributed protocols [2, 20] in which commitment is essential. Indeed, non-confluent non-deterministic choice is commonplace in verification frameworks such as mCRL2 [12]. It is also relevant in functional calculi; a well-known framework is De’Liguoro and Piperno’s (untyped) non-deterministic λ -calculus [8] (see also [9] and references therein).

To further illustrate the difference between confluent and non-confluent non-determinism, we consider an example adapted from [5]: a movie server that offers a choice between buying a movie or watching its trailer. In $s\pi^+$, the typed π -calculus that we present in this paper, this server can be specified as follows:

$$\text{Server}_s = s.\text{case} \left\{ \begin{array}{l} \text{buy} : s(\text{title}); s(\text{paym}); \bar{s}[\text{movie}]; \bar{s}[] \\ \text{peek} : s(\text{title}); \bar{s}[\text{trailer}]; \bar{s}[] \end{array} \right\} \begin{array}{l} -\text{Server}_s^{\text{buy}} \\ -\text{Server}_s^{\text{peek}} \end{array}$$

where $s(-)$ and $\bar{s}[-]$ denote input and output prefixes on a name/channel s , respectively, and ‘movie’ and ‘trailer’ denote references to primitive data. Also, the free names of a process are denoted with subscripts. Process Server_s offers a choice on name s ($s.\text{case}\{-\}$) between labels **buy** and **peek**. If **buy** is received, process $\text{Server}_s^{\text{buy}}$ is launched: it receives the movie’s title and a payment method, sends the movie, and closes the session on $s(\bar{s}[])$. If **peek** is received, it proceeds as $\text{Server}_s^{\text{peek}}$: the server receives the title, sends the trailer, and closes the session.

Using the non-deterministic choice operator of $s\pi^+$, denoted ‘ \sharp ’, we can specify a process for a client Alice who is interested in the movie ‘Jaws’ but is undecided about buying the film or just watching its trailer for free:

$$\text{Alice}_s := \begin{array}{l} \bar{s}.\text{buy}; \bar{s}[\text{Jaws}]; \bar{s}[\text{mcard}]; s(\text{movie}); s(); \mathbf{0} - \text{Alice}_s^{\text{buy}} \\ \sharp \bar{s}.\text{peek}; \bar{s}[\text{Jaws}]; s(\text{trailer}); s(); \mathbf{0} - \text{Alice}_s^{\text{peek}} \end{array}$$

If Alice_s selects the label **buy** ($\bar{s}.\text{buy}$), process $\text{Alice}_s^{\text{buy}}$ is launched: it sends title and payment method, receives the movie, waits for the session to close ($s()$),

and then terminates $(\mathbf{0})$. If Alice_s selects peek , process $\text{Alice}_s^{\text{peek}}$ is launched: it sends a title, receives the trailer, waits for the session to close, and terminates. Then, process $\text{Sys} := (\nu s)(\text{Server}_s \mid \text{Alice}_s)$ denotes the composition of client and server, connected along s (using (νs)). Our semantics for $s\pi^+$, denoted \rightsquigarrow , enforces non-confluent non-determinism, as Sys can reduce to separate processes, as expected:

$$\text{Sys} \rightsquigarrow (\nu s)(\text{Server}_s^{\text{buy}} \mid \text{Alice}_s^{\text{buy}}) \quad \text{and} \quad \text{Sys} \rightsquigarrow (\nu s)(\text{Server}_s^{\text{peek}} \mid \text{Alice}_s^{\text{peek}})$$

In contrast, the confluent non-deterministic choice from [5], denoted \oplus , behaves differently: in their confluent semantics, Sys reduces to a *single* process including *both alternatives*, i.e., $\text{Sys} \longrightarrow (\nu s)(\text{Server}_s^{\text{buy}} \mid \text{Alice}_s^{\text{buy}}) \oplus (\nu s)(\text{Server}_s^{\text{peek}} \mid \text{Alice}_s^{\text{peek}})$.

Contributions. We study new concurrent and functional calculi with usual (non-confluent) forms of non-determinism. Framed in the typed (resource-conscious) setting, we strive for definitions that do not exert a too drastic discarding of branches (as in the non-confluent case) but also that do not exhibit a lack of commitment (as in the confluent case). Concretely, we present:

(Section 2) $s\pi^+$, a variant of the session-typed π -calculus in [5], now with non-confluent non-deterministic choice. Its semantics adapts to the typed setting the usual semantics of non-deterministic choice in the untyped π -calculus [27]. Well-typed processes enjoy type preservation and deadlock-freedom (Theorems 1 and 2).

(Section 3) λ_c^ℓ , a resource λ -calculus with non-determinism, enhanced with constructs for expressing resource usage and failure. Its non-idempotent intersection type system provides a quantitative measure of the need/usage of resources. Well-typed terms enjoy subject reduction and subject expansion (Theorems 3 and 4).

(Section 4) A typed translation of λ_c^ℓ into $s\pi^+$, which provides further validation for our non-deterministic calculi, and casts them in the context of ‘functions-as-processes’. We prove that our translation is *correct*, i.e., it preserves types and satisfies tight operational correspondences (Theorems 5 and 6).

Moreover, Sect. 5 closes by discussing related works. Omitted technical material can be found in the full version of the paper [13].

2 A Typed π -calculus with Non-deterministic Choice

We introduce $s\pi^+$, a session-typed π -calculus with non-deterministic choice. Following [5], session types express protocols to be executed along channels. These protocols can be *non-deterministic*: sessions may succeed but also fail. The novelty in $s\pi^+$ is the non-deterministic choice operator $\mathcal{P} \# \mathcal{Q}$, whose *lazily committing semantics* is compatible with linearity. We prove that well-typed processes satisfy two key properties: *type preservation* and *deadlock-freedom*.

$P, Q ::= \mathbf{0}$	inaction	$[x \leftrightarrow y]$	forwarder
$ (\nu x)(P Q)$	connect	$P \# Q$	non-determinism
$ \bar{x}[y]; (P Q)$	output	$x(y); P$	input
$ \bar{x}. \ell; P$	select	$x.\text{case}\{i : P\}_{i \in I}$	branch
$ \bar{x}[]$	close	$x(); P$	wait
$ x.\text{some}_{w_1, \dots, w_n}; P$	expect	$\bar{x}.\text{some}; P$	available
$ P Q$	parallel	$\bar{x}.\text{none}$	unavailable
$P \equiv P' [P \equiv_\alpha P'] \quad [x \leftrightarrow y] \equiv [y \leftrightarrow x] \quad P \mathbf{0} \equiv P$			
$(P Q) R \equiv P (Q R)$	$P Q \equiv Q P$	$(\nu x)(P Q) \equiv (\nu x)(Q P)$	
$P \# P \equiv P$	$P \# Q \equiv Q \# P$	$(P \# Q) \# R \equiv P \# (Q \# R)$	
$(\nu x)((P Q) R) \equiv (\nu x)(P R) Q$		$[x \notin \text{fn}(Q)]$	
$(\nu x)((\nu y)(P Q) R) \equiv (\nu y)((\nu x)(P R) Q)$		$[x \notin \text{fn}(Q), y \notin \text{fn}(R)]$	

Fig. 1. $s\pi^+$: syntax (top) and structural congruence (bottom).

2.1 Syntax and Semantics

We use P, Q, \dots to denote processes, and x, y, z, \dots to denote *names* representing channels. Figure 1 (top) gives the syntax of processes. $P\{y/z\}$ denotes the capture-avoiding substitution of y for z in P . Process $\mathbf{0}$ denotes inaction, and $[x \leftrightarrow y]$ is a forwarder: a bidirectional link between x and y . Parallel composition appears in two forms: while the process $P | Q$ denotes communication-free concurrency, process $(\nu x)(P | Q)$ uses restriction (νx) to express that P and Q implement complementary behaviors on x and do not share any other names.

Process $P \# Q$ denotes the non-deterministic choice between P and Q : intuitively, if one choice can perform a synchronization, the other option may be discarded if it cannot. Since $\#$ is associative, we often omit parentheses. Also, we write $\#_{i \in I} P_i$ for the non-deterministic choice between each P_i for $i \in I$.

Our output construct integrates parallel composition and restriction: process $\bar{x}[y]; (P | Q)$ sends a fresh name y along x and then continues as $P | Q$. The type system will ensure that behaviors on y and x are implemented by P and Q , respectively, which do not share any names—this separation defines communication-free concurrency and is key to ensuring deadlock-freedom. The input process $x(y); P$ receives a name z along x and continues as $P\{z/y\}$, which does not require the separation present in the output case. Process $x.\text{case}\{i : P_i\}_{i \in I}$ denotes a branch with labeled choices indexed by the finite set I : it awaits a choice on x with continuation P_j for each $j \in I$. The process $\bar{x}. \ell; P$ selects on x the choice labeled ℓ before continuing as P . Processes $\bar{x}[]$ and $x(); P$ are dual actions for closing the session on x . We omit replicated servers $!x(y); P$ and corresponding client requests $? \bar{x}[y]; P$, but they can be easily added (cf. [13]).

The remaining constructs define non-deterministic sessions which may provide a protocol or fail, following [5]. Process $\bar{x}.\text{some}; P$ confirms the availability of a session on x and continues as P . Process $\bar{x}.\text{none}$ signals the failure to provide the session on x . Process $x.\text{some}_{w_1, \dots, w_n}; P$ specifies a dependency on a non-deterministic session on x (names w_1, \dots, w_n implement sessions in P). This process can either (i) synchronize with a ' $\bar{x}.\text{some}$ ' and continue as P , or (ii) synchronize with a ' $\bar{x}.\text{none}$ ', discard P , and propagate the failure to w_1, \dots, w_n . To reduce eye strain, in writing $x.\text{some}$ we freely combine names and sets of names. This way, e.g., we write $x.\text{some}_{y, \text{fn}(P), \text{fn}(Q)}$ rather than $x.\text{some}_{\{y\} \cup \text{fn}(P) \cup \text{fn}(Q)}$.

Name y is bound in $(\nu y)(P \mid Q)$, $\bar{x}[y]; (P \mid Q)$, and $x(y); P$. The set $\text{fn}(P)$ includes the names in P that are not bound. We adopt Barendregt's convention.

Structural Congruence. Reduction defines the steps that a process performs on its own. It relies on *structural congruence* (\equiv), the least congruence relation on processes induced by the rules in Fig. 1 (bottom). Like the syntax of processes, the definition of \equiv is aligned with the type system (defined next), such that \equiv preserves typing (subject congruence, cf. Theorem 1). Differently from [5], we do not allow distributing non-deterministic choice over parallel and restriction. As shown in [13], the position of a non-deterministic choice in a process determines how it may commit, so changing its position affects commitment.

Reduction: Intuitions and Prerequisites. Barring non-deterministic choice, our reduction rules arise as directed interpretations of proof transformations in the underlying linear logic. We follow Caires and Pfenning [6] and Wadler [31] in interpreting cut-elimination in linear logic as synchronization in $\mathfrak{s}\pi^+$.

Before delving into our reduction rules (Fig. 2), it may be helpful to consider the usual reduction axiom for the (untyped) π -calculus (e.g., [19, 27]):

$$(\bar{x}[z]; P_1 + M_1) \mid (x(y); P_2 + M_2) \longrightarrow P_1 \mid P_2\{z/y\} \quad (1)$$

This axiom captures the interaction of two (binary) choices: it integrates the commitment of choice in synchronization; after the reduction step, the two branches not involved in the synchronization, M_1 and M_2 , are discarded. Our semantics of $\mathfrak{s}\pi^+$ is defined similarly: when a prefix within a branch of a choice synchronizes with its dual, that branch reduces and the entire process commits to it.

The key question at this point is: when and to which branches should we commit? In (1), a communication commits to a single branch. For $\mathfrak{s}\pi^+$, we define a *lazy semantics* that minimizes commitment as much as possible.

The intuitive idea is that multiple branches of a choice may contain the same prefix, and so all these branches represent possibilities for synchronization ("possible branches"). Other branches with different prefixes denote different possibilities ("impossible branches"). When one synchronization is chosen, the possible branches are maintained while the impossible ones are discarded.

Example 1. To distinguish possible and impossible branches, consider:

$$P := (\nu s)(s.\text{case}\{\text{buy} : \dots, \text{peek} : \dots\} \mid (\bar{s}.\text{buy}; \dots \# \bar{s}.\text{buy}; \dots \# \bar{s}.\text{peek}; \dots))$$

The branch construct (case) provides the context for the non-deterministic choice. When the case synchronizes on the ‘buy’ label, the two branches prefixed by ‘ $\bar{s}.\text{buy}$ ’ are possible, whereas the branch prefixed by ‘ $\bar{s}.\text{peek}$ ’ becomes impossible, and can be discarded. The converse occurs when the ‘peek’ label is selected. ∇

To formalize these intuitions, our reduction semantics (Fig. 2) relies on some auxiliary definitions. First, we define contexts.

Definition 1. We define ND-contexts (\mathbf{N}, \mathbf{M}) as follows:

$$\mathbf{N}, \mathbf{M} ::= [\cdot] \mid \mathbf{N} \mid P \mid (\nu x)(\mathbf{N} \mid P) \mid \mathbf{N} \# P$$

The process obtained by replacing $[\cdot]$ in \mathbf{N} with P is denoted $\mathbf{N}[P]$. We refer to ND-contexts that do not use the clause ‘ $\mathbf{N} \# P$ ’ as D-contexts, denoted \mathbf{C}, \mathbf{D} .

Using D-contexts, we can express that, e.g., $\prod_{i \in I} \mathbf{C}_i[\bar{x}]$ and $\prod_{j \in J} \mathbf{D}_j[x(); Q_j]$ should match. To account for reductions with impossible branches, we define a pre-congruence on processes, denoted \succeq_S , where the parameter S denotes the subject(s) of the prefix in the possible branches. Our semantics is closed under \succeq_S . Hence, e.g., anticipating a reduction on x , the possible branch $\mathbf{C}_1[x(y); P]$ can be extended with an impossible branch to form $\mathbf{C}_1[x(y); P] \# \mathbf{C}_2[z(); Q]$.

Before defining \succeq_S (Definition 3), we first define prefixes (and their subjects). Below, we write \tilde{x} to denote a finite tuple of names x_1, \dots, x_k .

Definition 2. Prefixes are defined as follows:

$$\alpha, \beta ::= \bar{x}[y] \mid x(y) \mid \bar{x}.\ell \mid x.\text{case} \mid \bar{x}[] \mid x() \mid \bar{x}.\text{some} \mid \bar{x}.\text{none} \mid x.\text{some}_{\bar{w}} \mid [x \leftrightarrow y]$$

The subjects of α , denoted $\text{subj}\{\alpha\}$, are $\{x, y\}$ in case of $[x \leftrightarrow y]$, or $\{x\}$. By abuse of notation, we write $\alpha; P$ even when α takes no continuation (as in $\bar{x}[]$, $\bar{x}.\text{none}$, and $[x \leftrightarrow y]$) and for $\bar{x}[y]$ which takes a parallel composition as continuation.

Definition 3. Let \bowtie denote the least relation on prefixes (Definition 2) defined by:

(i) $\bar{x}[y] \bowtie \bar{x}[z]$, (ii) $x(y) \bowtie x(z)$, and (iii) $\alpha \bowtie \alpha$ otherwise.

Given a non-empty set $S \subseteq \{x, y\}$, the pre-congruence $P \succeq_S Q$ holds when both following conditions hold:

1. $S = \{x\}$ implies

$$P = \left(\prod_{i \in I} \mathbf{C}_i[\alpha_i; P_i] \right) \# \left(\prod_{j \in J} \mathbf{C}_j[\beta_j; Q_j] \right) \text{ and } Q = \prod_{i \in I} \mathbf{C}_i[\alpha_i; P_i], \text{ where}$$

(i) $\forall i, i' \in I. \alpha_i \bowtie \alpha_{i'}$ and $\text{subj}\{\alpha_i\} = \{x\}$, and

(ii) $\forall i \in I. \forall j \in J. \alpha_i \not\bowtie \beta_j \wedge x \in \text{fn}(\beta_j; Q_j)$;

2. $S = \{x, y\}$ implies

$$P = \left(\prod_{i \in I} \mathbf{C}_i[[x \leftrightarrow y]] \right) \# \left(\prod_{j \in J} \mathbf{C}_j[[x \leftrightarrow z_j]] \right) \# \left(\prod_{k \in K} \mathbf{C}_k[\alpha_k; P_k] \right)$$

and $Q = \prod_{i \in I} \mathbf{C}_i[[x \leftrightarrow y]]$, where

(i) $\forall j \in J. z_j \neq y$, and (ii) $\forall k \in K. x \in \text{fn}(\alpha_k; P_k) \wedge \forall z. \alpha_k \not\bowtie [x \leftrightarrow z]$.

$$\begin{array}{l}
[\rightsquigarrow_{\text{ID}}] (\nu x) \left(\prod_{i \in I} \mathbf{C}_i [x \leftrightarrow y] \mid Q \right) \rightsquigarrow_{x,y} \prod_{i \in I} \mathbf{C}_i [Q\{y/x\}] \\
[\rightsquigarrow_{\otimes \otimes}] (\nu x) \left(\prod_{i \in I} \mathbf{C}_i [\bar{x}y_i]; (P_i \mid Q_i) \mid \prod_{j \in J} \mathbf{D}_j [x(z); R_j] \right) \\
\rightsquigarrow_x \prod_{i \in I} \mathbf{C}_i \left[(\nu x) (Q_i \mid (\nu w) (P_i \{w/y_i\} \mid \prod_{j \in J} \mathbf{D}_j [R_j \{w/z\}])) \right] \\
[\rightsquigarrow_{\oplus \&}] (\nu x) \left(\prod_{i \in I} \mathbf{C}_i [\bar{x}.k'; P_i] \mid \prod_{j \in J} \mathbf{D}_j [x.\text{case}\{k : Q_j^k\}_{k \in K}] \right) \\
\rightsquigarrow_x (\nu x) \left(\prod_{i \in I} \mathbf{C}_i [P_i] \mid \prod_{j \in J} \mathbf{D}_j [Q_j^{k'}] \right) \quad [k' \in K] \\
[\rightsquigarrow_{\mathbf{1} \perp}] (\nu x) \left(\prod_{i \in I} \mathbf{C}_i [\bar{x}[]] \mid \prod_{j \in J} \mathbf{D}_j [x(); Q_j] \right) \rightsquigarrow_x \prod_{i \in I} \mathbf{C}_i [\mathbf{0}] \mid \prod_{j \in J} \mathbf{D}_j [Q_j] \\
[\rightsquigarrow_{\text{some}}] (\nu x) \left(\prod_{i \in I} \mathbf{C}_i [\bar{x}.\text{some}; P_i] \mid \prod_{j \in J} \mathbf{D}_j [x.\text{some}_{w_1, \dots, w_n}; Q_j] \right) \\
\rightsquigarrow_x (\nu x) \left(\prod_{i \in I} \mathbf{C}_i [P_i] \mid \prod_{j \in J} \mathbf{D}_j [Q_j] \right) \\
[\rightsquigarrow_{\text{none}}] (\nu x) \left(\prod_{i \in I} \mathbf{C}_i [\bar{x}.\text{none}] \mid \prod_{j \in J} \mathbf{D}_j [x.\text{some}_{w_1, \dots, w_n}; Q_j] \right) \\
\rightsquigarrow_x \prod_{i \in I} \mathbf{C}_i [\mathbf{0}] \mid \prod_{j \in J} \mathbf{D}_j [\bar{w}_1.\text{none} \mid \dots \mid \bar{w}_n.\text{none}] \\
[\rightsquigarrow_{\succeq_S}] \frac{x \in S \quad P \succeq_S P' \quad Q \succeq_S Q' \quad (\nu x)(P' \mid Q') \rightsquigarrow_S R}{(\nu x)(P \mid Q) \rightsquigarrow_S R} \\
[\rightsquigarrow_{\nu \#}] \frac{(\nu x)(P \mid \mathbf{N}[\mathbf{C}[Q_1] \# \mathbf{C}[Q_2]]) \rightsquigarrow_S R}{(\nu x)(P \mid \mathbf{N}[\mathbf{C}[Q_1] \# Q_2]) \rightsquigarrow_S R} \\
[\rightsquigarrow_{\equiv}] \frac{P \equiv P' \quad P' \rightsquigarrow_S Q' \quad Q' \equiv Q}{P \rightsquigarrow_S Q} \quad [\rightsquigarrow_{\nu}] \frac{P \rightsquigarrow_S P'}{(\nu x)(P \mid Q) \rightsquigarrow_S (\nu x)(P' \mid Q)} \\
[\rightsquigarrow_{\mid}] \frac{P \rightsquigarrow_S P'}{P \mid Q \rightsquigarrow_S P' \mid Q} \quad [\rightsquigarrow_{\#}] \frac{P \rightsquigarrow_S P'}{P \# Q \rightsquigarrow_S P' \# Q}
\end{array}$$

Fig. 2. Reduction semantics for $\mathfrak{s}\pi^+$.

Intuitively, \otimes allows us to equate output/input prefixes with the same subject (but different object). The rest of Definition 3 accounts for two kinds of reduction, using S to discard “impossible” branches. In case S is $\{x\}$ (Item 1), it concerns a synchronization on x ; in case S is $\{x, y\}$, it concerns forwarding on x and y (Item 2). In both cases, P and Q contain matching prefixes on x , while P may contain additional branches with different or blocked prefixes on x ; x must appear in the hole of the contexts in the additional branches in P (enforced with $x \in \text{fn}(\dots)$), to ensure that no matching prefixes are discarded.

Example 2. Recall process P from Example 1. To derive a synchronization with the ‘buy’ alternative of the case, we can use \succeq_S to discard the ‘peek’ alternative, as follows: $\bar{s}.\text{buy}; \dots \# \bar{s}.\text{buy}; \dots \# \bar{s}.\text{peek}; \dots \succeq_s \bar{s}.\text{buy}; \dots \# \bar{s}.\text{buy}; \dots$ ∇

Reduction Rules. Figure 2 gives the rules for the (lazy) reduction semantics, denoted \rightsquigarrow_S , where the set S contains the names involved in the interaction. We

omit the curly braces in this annotation; this way, e.g., we write ‘ $\rightsquigarrow_{x,y}$ ’ instead of ‘ $\rightsquigarrow_{\{x,y\}}$ ’. Also, we write \rightsquigarrow_S^k to denote a sequence of $k \geq 0$ reductions.

The first six rules in Fig. 2 formalize forwarding and communication: they are defined on choices containing different D-contexts (cf. Definition 1), each with the same prefix but possibly different continuations; these rules preserve the non-deterministic choices. Rule $[\rightsquigarrow_{ID}]$ fixes S to the forwarder’s two names, and the other rules fix S to the one involved name. In particular, Rule $[\rightsquigarrow_{\otimes \otimes}]$ formalizes name communication: it involves multiple senders and multiple receivers (grouped in choices indexed by I and J , respectively). Because they proceed in lock-step, reduction leads to substitutions involving the same (fresh) name w ; also, the scopes of the choice and the contexts enclosing the senders is extended.

Rule $[\rightsquigarrow_{\succeq_s}]$ is useful to derive a synchronization that discards groups of choices. Rule $[\rightsquigarrow_{\nu \parallel}]$ allows inferring reductions when non-deterministic choices are not top-level: e.g., $(\nu x)(\bar{x}[] \mid (\nu y)((x()); Q_1 \parallel x(); Q_2) \mid R) \rightsquigarrow_x (\nu y)(Q_1 \mid R) \parallel (\nu y)(Q_2 \mid R)$. The last four rules formalize that reduction is closed under structural congruence, restriction, parallel composition, and non-deterministic choice.

As mentioned earlier, a key motivation for our work is to have non-deterministic choices that effectively enforce commitment, without a too drastic discarding of alternatives. Next we illustrate this intended form of *gradual commitment*.

Example 3. (A Modified Movie Server). Consider the following variant of the movie server from the introduction, where the handling of the payment is now modeled as a branch:

$$\text{NewServer}_s := s(\text{title}); s.\text{case} \left\{ \begin{array}{l} \text{buy} : s.\text{case} \left\{ \begin{array}{l} \text{card} : s(\text{info}); \bar{s}[\text{movie}]; \bar{s}[], \\ \text{cash} : \bar{s}[\text{movie}]; \bar{s}[] \end{array} \right\}, \\ \text{peek} : \bar{s}[\text{trailer}]; \bar{s}[] \end{array} \right\}$$

Consider a client, Eve, who cannot decide between buying ‘Oppenheimer’ or watching its trailer. In the former case, she has two options for payment method:

$$\text{Eve}_s := \bar{s}[\text{Oppenheimer}]; \left(\begin{array}{l} \bar{s}.\text{buy}; \bar{s}.\text{card}; \bar{s}[\text{visa}]; s(\text{movie}); s(); \mathbf{0} \\ \parallel \bar{s}.\text{buy}; \bar{s}.\text{cash}; s(\text{movie}); s(); \mathbf{0} \\ \parallel \bar{s}.\text{peek}; s(\text{link}); s(); \mathbf{0} \end{array} \right)$$

Let $\text{Sys}^* := (\nu s)(\text{NewServer}_s \mid \text{Eve}_s)$. After sending the movie’s title, Eve’s choice (buying or watching the trailer) enables gradual commitment. We have:

$$\text{Sys}^* \rightsquigarrow_s^2 (\nu s)(s.\text{case}\{\text{card} : \dots, \text{cash} : \dots\} \mid (\bar{s}.\text{card}; \dots \parallel \bar{s}.\text{cash}; \dots)) =: \text{Sys}_1^*$$

$$\text{and} \quad \text{Sys}^* \rightsquigarrow_s^2 (\nu s)(\bar{s}[\text{trailer}]; \dots \mid s(\text{trailer}); \dots) =: \text{Sys}_2^*$$

Process Sys_1^* represents the situation for Eve after selecting buy, in which case the third alternative $(\bar{s}.\text{peek}; \dots)$ can be discarded as an impossible branch. Process Sys_2^* represents the dual situation. From Sys_1^* , the selection of payment method completes the commitment to one alternative; we have: $\text{Sys}_1^* \rightsquigarrow_s (\nu s)(s(\text{info}); \dots \mid \bar{s}[\text{visa}]; \dots)$ and $\text{Sys}_1^* \rightsquigarrow_s (\nu s)(\bar{s}[\text{movie}]; \dots \mid s(\text{movie}); \dots)$. ∇

$[\text{T}_{\text{CUT}}] \frac{P \vdash \Gamma, x:A \quad Q \vdash \Delta, x:\bar{A}}{(\nu x)(P \mid Q) \vdash \Gamma, \Delta}$	$[\text{T}_{\text{MIX}}] \frac{P \vdash \Gamma \quad Q \vdash \Delta}{P \mid Q \vdash \Gamma, \Delta}$	
$[\text{T}_{\#}] \frac{P \vdash \Gamma \quad Q \vdash \Gamma}{P \# Q \vdash \Gamma}$	$[\text{T}_{\text{EMPTY}}] \frac{}{\mathbf{0} \vdash \emptyset}$	$[\text{T}_{\text{ID}}] \frac{}{[x \leftrightarrow y] \vdash x:A, y:\bar{A}}$
$[\text{T}_{\mathbf{1}}] \frac{}{\bar{x}[] \vdash x:\mathbf{1}}$	$[\text{T}_{\perp}] \frac{P \vdash \Gamma}{x(); P \vdash \Gamma, x:\perp}$	$[\text{T}_{\otimes}] \frac{P \vdash \Gamma, y:A \quad Q \vdash \Delta, x:B}{\bar{x}[y]; (P \mid Q) \vdash \Gamma, \Delta, x:A \otimes B}$
$[\text{T}_{\wp}] \frac{P \vdash \Gamma, y:A, x:B}{x(y); P \vdash \Gamma, x:A \wp B}$	$[\text{T}_{\oplus}] \frac{P \vdash \Gamma, x:A_j \quad j \in I}{\bar{x}.j; P \vdash \Gamma, x:\oplus\{i : A_i\}_{i \in I}}$	
$[\text{T}_{\&}] \frac{\forall i \in I. P_i \vdash \Gamma, x:A_i}{x.\text{case}\{i : P_i\}_{i \in I} \vdash \Gamma, x:\&\{i : A_i\}_{i \in I}}$	$[\text{T}_{\&\text{some}}] \frac{P \vdash \Gamma, x:A}{\bar{x}.\text{some}; P \vdash \Gamma, x:\&A}$	
$[\text{T}_{\&\text{none}}] \frac{}{\bar{x}.\text{none} \vdash x:\&A}$	$[\text{T}_{\oplus\text{some}}] \frac{P \vdash \&\Gamma, x:A}{x.\text{some}_{\text{dom}(\Gamma)}; P \vdash \&\Gamma, x:\oplus A}$	

Fig. 3. Typing rules for $\mathfrak{s}\pi^+$.

2.2 Resource Control for $\mathfrak{s}\pi^+$ via Session Types

We define a session type system for $\mathfrak{s}\pi^+$, following ‘propositions-as-sessions’ [6, 31]. As already mentioned, in a session type system, resources are names that perform protocols: the *type assignment* $x : A$ says that x should conform to the protocol specified by the session type A . We give the syntax of types:

$$A, B ::= \mathbf{1} \mid \perp \mid A \otimes B \mid A \wp B \mid \oplus\{i : A\}_{i \in I} \mid \&\{i : A\}_{i \in I} \mid \&A \mid \oplus A$$

The units $\mathbf{1}$ and \perp type closed sessions. $A \otimes B$ types a name that first outputs a name of type A and then proceeds as B . Similarly, $A \wp B$ types a name that inputs a name of type A and then proceeds as B . Types $\oplus\{i : A_i\}_{i \in I}$ and $\&\{i : A_i\}_{i \in I}$ are given to names that can select and offer a labeled choice, respectively. Then, $\&A$ is the type of a name that *may produce* a behavior of type A , or fail; dually, $\oplus A$ types a name that *may consume* a behavior of type A .

For any type A we denote its *dual* as \bar{A} . Intuitively, dual types serve to avoid communication errors: the type at one end of a channel is the dual of the type at the opposite end. Duality is an involution, defined as follows:

$$\begin{array}{llll} \bar{\mathbf{1}} = \perp & \overline{A \otimes B} = \bar{A} \wp \bar{B} & \overline{\oplus\{i : A_i\}_{i \in I}} = \&\{i : \bar{A}_i\}_{i \in I} & \overline{\&A} = \oplus \bar{A} \\ \bar{\perp} = \mathbf{1} & \overline{A \wp B} = \bar{A} \otimes \bar{B} & \overline{\&\{i : A_i\}_{i \in I}} = \oplus\{i : \bar{A}_i\}_{i \in I} & \overline{\oplus A} = \&\bar{A} \end{array}$$

Judgments are of the form $P \vdash \Gamma$, where P is a process and Γ is a context, a collection of type assignments. In writing $\Gamma, x : A$, we assume $x \notin \text{dom}(\Gamma)$. We write $\text{dom}(\Gamma)$ to denote the set of names appearing in Γ . We write $\&\Gamma$ to denote that $\forall x : A \in \Gamma. \exists A'. A = \&A'$.

Figure 3 gives the typing rules: they correspond to the rules in Curry-Howard interpretations of classical linear logic as session types (cf. Wadler [31]), with the rules for $\&A$ and $\oplus A$ extracted from [5], and the additional Rule $[T\ddagger]$ for non-confluent non-deterministic choice, which modifies the confluent rule in [5].

Most rules follow [31], so we focus on those related to non-determinism. Rule $[T\&\text{some}]$ types a process with a name whose behavior can be provided, while Rule $[T\&\text{none}]$ types a name whose behavior cannot. Rule $[T\oplus\text{some}]$ types a process with a name x whose behavior may not be available. If the behavior is not available, all the sessions in the process must be canceled; hence, the rule requires all names to be typed under the $\&A$ monad.

Rule $[T\ddagger]$ types our new non-deterministic choice operator; the branches must be typable under the same typing context. Hence, all branches denote the same sessions, which may be implemented differently. In context of a synchronization, branches that are kept are able to synchronize, whereas the discarded branches are not; nonetheless, the remaining branches still represent different implementations of the same sessions. Compared to the rule for non-determinism in [5], we do not require processes to be typable under the $\&A$ monad.

Example 4. Consider again process Eve_s from Example 3. The three branches of the non-deterministic choice give *different implementations of the same session*: assuming primitive, self-dual data types \mathbf{C} , \mathbf{M} , and \mathbf{L} , all three branches on s are typable by $\oplus\{\text{buy} : \oplus\{\text{card} : \mathbf{C} \otimes \mathbf{M} \wp \perp, \text{cash} : \mathbf{M} \wp \perp\}, \text{peek} : \mathbf{L} \wp \perp\}$. ∇

Example 5 (Unavailable Movies). Consider now a modified movie server, which offers movies that may not be yet available. We specify this server using non-deterministic choice and non-deterministically available sessions:

$$\text{BuyServ}_s := s(\text{title}); (\bar{s}.\text{none} \ddagger \bar{s}.\text{some}; s(\text{paym}); \bar{s}[\text{movie}]; \bar{s}[]) \vdash s : T \wp (\&(\mathbf{P} \wp \mathbf{M} \otimes \mathbf{1})),$$

where $T, \mathbf{P}, \mathbf{M}$ denote primitive, self-dual data-types. While the branch ‘ $\bar{s}.\text{none}$ ’ signals that the movie is not available, the branch ‘ $\bar{s}.\text{some}; \dots$ ’ performs the expected protocol. We now define a client Ada who buys a movie for Tim, using session s ; Ada only forwards it to him (using session u) if it is actually available:

$$\begin{aligned} \text{Ada}_{s,u} &:= \bar{s}[\text{Barbie}]; s.\text{some}_u; \bar{s}[\text{visa}]; s(\text{movie}); s(); \bar{u}.\text{some}; \bar{u}[\text{movie}]; \bar{u}[] \\ &\vdash s : T \otimes (\oplus(\mathbf{P} \otimes \mathbf{M} \wp \perp)), u : \&(\mathbf{M} \otimes \mathbf{1}) \end{aligned}$$

$$\text{Tim}_u := u.\text{some}; u(\text{movie}); u(); \mathbf{0} \vdash u : \oplus(\mathbf{M} \wp \mathbf{1})$$

Let $\text{BuySys} := (\nu s)(\text{BuyServ}_s \mid (\nu u)(\text{Ada}_{s,u} \mid \text{Tim}_u))$. Depending on whether the server has the movie ‘Barbie’ available, we have the following reductions:

$$\text{BuySys} \rightsquigarrow_s^2 (\nu u)(\bar{u}.\text{none} \mid \text{Tim}_u) \quad \text{or} \quad \text{BuySys} \rightsquigarrow_s^5 (\nu u)(\bar{u}.\text{some}; \dots \mid \text{Tim}_u)$$

∇

Our type system ensures *session fidelity* and *communication safety*, but not confluence: the former says that processes correctly follow their ascribed session protocols, and the latter that no communication errors/mismatches occur. Both properties follow from the fact that typing is consistent across structural congruence and reduction. See [13] for details.

$M, N, L ::= x$	variable	$M\langle\langle C/x \rangle\rangle$	intermediate subst.
$ (M C)$	application	$M\langle C/\tilde{x} \rangle$	explicit subst.
$ \lambda x.M$	abstraction	$\mathbf{fail}^{\tilde{x}}$	failure
$ M[\tilde{x} \leftarrow x]$	sharing		
$C, D ::= \mathbf{1} \mid \wr M \wr \cdot C$			bag
$\mathcal{C} ::= [\cdot] \mid \mathcal{C}[\tilde{x} \leftarrow x] \mid (C C) \mid \mathcal{C}\langle C/\tilde{x} \rangle$			context

Fig. 4. Syntax of $\lambda_{\mathcal{C}}^{\ell}$: terms, bags, and contexts.

Theorem 1 (Type Preservation). *If $P \vdash \Gamma$, then both $P \equiv Q$ and $P \rightsquigarrow_S Q$ (for any Q and S) imply $Q \vdash \Gamma$.*

Another important, if often elusive, property in session types is *deadlock-freedom*, which ensures that processes can reduce as long as they are not inactive. Our type system satisfies deadlock-freedom for processes with fully connected names, i.e., typable under the empty context. See [13] for details.

Theorem 2 (Deadlock-freedom). *If $P \vdash \emptyset$ and $P \not\equiv \mathbf{0}$, then there are Q and S such that $P \rightsquigarrow_S Q$.*

3 A Non-deterministic Resource λ -calculus

We present $\lambda_{\mathcal{C}}^{\ell}$, a resource λ -calculus with non-determinism and lazy evaluation. In $\lambda_{\mathcal{C}}^{\ell}$, non-determinism is non-confluent and *implicit*, as it arises from the fetching of terms from bags of *linear* resources. This is different from $\mathfrak{s}\pi^+$, where the choice operator ‘ \sharp ’ specifies non-determinism *explicitly*. A mismatch between the number of variable occurrences and the size of the bag induces *failure*.

In $\lambda_{\mathcal{C}}^{\ell}$, the *sharing* construct $M[x_1, \dots, x_n \leftarrow x]$, expresses that x may be used in M under “aliases” x_1, \dots, x_n . Hence, it atomizes n occurrences of x in M , via an explicit pointer to n variables. This way, e.g., the λ -term $\lambda x.(x x)$ is expressed in $\lambda_{\mathcal{C}}^{\ell}$ as $\lambda x.(x_1 \wr x_2 \wr [x_1, x_2 \leftarrow x])$, where $\wr x_2 \wr$ is a bag containing x_2 .

3.1 Syntax and Reduction Semantics

Syntax. We use x, y, z, \dots for variables, and write \tilde{x} to denote a finite sequence of pairwise distinct x_i ’s, with length $|\tilde{x}|$. Figure 4 gives the syntax of terms (M, N, L) and bags (C, D). The empty bag is denoted $\mathbf{1}$. We use C_i to denote the i -th term in C , and $\text{size}(C)$ denotes the number of elements in C . To ease readability, we often write, e.g., $\wr N_1, N_2 \wr$ as a shorthand notation for $\wr N_1 \wr \cdot \wr N_2 \wr$.

In $M[\tilde{x} \leftarrow x]$, we say that \tilde{x} are the *shared variables* and that x is the *sharing variable*. We require for each $x_i \in \tilde{x}$: (i) x_i occurs exactly once in M ; (ii) x_i is not a sharing variable. The sequence \tilde{x} can be empty: $M[\leftarrow x]$ means that x does not share any variables in M . Sharing binds the shared variables in the term.

$\frac{[\text{RS:Beta}]}{(\lambda x.M)C \longrightarrow M\langle\langle C/x \rangle\rangle}$	$\frac{[\text{RS:Ex-Sub}]}{\frac{\text{size}(C) = \tilde{x} \quad M \neq \text{fail}^{\tilde{y}}}{(M[\tilde{x} \leftarrow x])\langle\langle C/x \rangle\rangle \longrightarrow M\langle\langle C/\tilde{x} \rangle\rangle}}$	$\frac{[\text{RS : TCont}]}{M \longrightarrow N \quad C[M] \longrightarrow C[N]}$
$\frac{[\text{RS:Fetch}^\ell]}{\frac{\text{head}(M) = x_j \quad 0 < i \leq \text{size}(C)}{M\langle\langle C/\tilde{x}, x_j \rangle\rangle \longrightarrow M\{C_i/x_j\}\langle\langle C \setminus C_i/\tilde{x} \rangle\rangle}}$		
$\frac{[\text{RS:Fail}^\ell]}{\frac{\text{size}(C) \neq \tilde{x} \quad \tilde{y} = (\text{fv}(M) \setminus \{\tilde{x}\}) \cup \text{fv}(C)}{(M[\tilde{x} \leftarrow x])\langle\langle C/x \rangle\rangle \longrightarrow \text{fail}^{\tilde{y}}}}$	$\frac{[\text{RS:Cons}_1]}{\frac{\tilde{y} = \text{fv}(C)}{\text{fail}^{\tilde{x}} C \longrightarrow \text{fail}^{\tilde{x} \cup \tilde{y}}}}$	
$\frac{[\text{RS:Cons}_2]}{\frac{\text{size}(C) = \tilde{x} \quad \tilde{z} = \text{fv}(C)}{(\text{fail}^{\tilde{x} \cup \tilde{y}}[\tilde{x} \leftarrow x])\langle\langle C/x \rangle\rangle \longrightarrow \text{fail}^{\tilde{y} \cup \tilde{z}}}}$	$\frac{[\text{RS:Cons}_3]}{\frac{\tilde{z} = \text{fv}(C)}{\text{fail}^{\tilde{y} \cup \tilde{x}}\langle\langle C/\tilde{x} \rangle\rangle \longrightarrow \text{fail}^{\tilde{y} \cup \tilde{z}}}}$	
<p>where $\text{head}(M)$ is defined as follows:</p>		
$\begin{array}{lll} \text{head}(x) = x & \text{head}(\lambda x.M) = \lambda x.M & \text{head}((M C)) = \text{head}(M) \\ \text{head}(\text{fail}^{\tilde{x}}) = \text{fail}^{\tilde{x}} & \text{head}(M\langle\langle C/x \rangle\rangle) = M\langle\langle C/x \rangle\rangle & \text{head}(M\langle\langle C/\tilde{x} \rangle\rangle) = \text{head}(M) \\ \text{head}(M[\tilde{x} \leftarrow x]) = \begin{cases} x & \text{head}(M) = y \text{ and } y \in \tilde{x} \\ \text{head}(M) & \text{otherwise} \end{cases} \end{array}$		

Fig. 5. Reduction rules for λ_{C}^ℓ .

An abstraction $\lambda x.M$ binds occurrences of x in M . Application $(M C)$ is as usual. The term $M\langle\langle C/\tilde{x} \rangle\rangle$ is the *explicit substitution* of a bag C for \tilde{x} in M . We require $\text{size}(C) = |\tilde{x}|$ and for each $x_i \in \tilde{x}$: (i) x_i occurs in M ; (ii) x_i is not a sharing variable; (iii) x_i cannot occur in another explicit substitution in M . The term $M\langle\langle C/x \rangle\rangle$ denotes an intermediate explicit substitution that does not (necessarily) satisfy the conditions for explicit substitutions.

The term $\text{fail}^{\tilde{x}}$ denotes failure; the variables in \tilde{x} are “dangling” resources, which cannot be accounted for after failure. We write $\text{fv}(M)$ to denote the free variables of M , defined as expected. Term M is *closed* if $\text{fv}(M) = \emptyset$.

Semantics. Figure 5 gives the reduction semantics, denoted \longrightarrow , and the *head variable* of term M , denoted $\text{head}(M)$. Rule [RS : Beta] induces an intermediate substitution. Rule [RS : Ex-Sub] reduces an intermediate substitution to an explicit substitution, provided the size of the bag equals the number of shared variables. In case of a mismatch, the term evolves into failure via Rule [RS : Fail $^\ell$].

An explicit substitution $M\langle\langle C/\tilde{x} \rangle\rangle$, where the head variable of M is $x_j \in \tilde{x}$, reduces via Rule [R : Fetch $^\ell$]. The rule extracts a C_i from C (for some $0 < i \leq \text{size}(C)$) and substitutes it for x_j in M ; this is how fetching induces a non-

deterministic choice between $\text{size}(C)$ possible reductions. Rules $[\text{RS} : \text{Cons}_j]$ for $j \in \{1, 2, 3\}$ consume terms when they meet failure. Finally, Rule $[\text{RS} : \text{TCont}]$ closes reduction under contexts. The following example illustrates reduction.

Example 6. Consider the term $M_0 = (\lambda x.x_1 \lambda x_2 \lambda x_3 \mathbf{1} \int \int [\tilde{x} \leftarrow x]) \lambda \text{fail}^\emptyset, y, I \int$, where $I = \lambda x.(x_1[x_1 \leftarrow x])$ and $\tilde{x} = x_1, x_2, x_3$. First, M_0 evolves into an intermediate substitution (2). The bag can provide for all shared variables, so it then evolves into an explicit substitution (3):

$$M_0 \longrightarrow (x_1 \lambda x_2 \lambda x_3 \mathbf{1} \int \int [\tilde{x} \leftarrow x]) \langle \langle \lambda \text{fail}^\emptyset, y, I \int / x \rangle \rangle \quad (2)$$

$$\longrightarrow (x_1 \lambda x_2 \lambda x_3 \mathbf{1} \int \int) \langle \langle \lambda \text{fail}^\emptyset, y, I \int / \tilde{x} \rangle \rangle = M \quad (3)$$

Since $\text{head}(M) = x_1$, one of the three elements of the bag will be substituted. M represents a non-deterministic choice between the following three reductions:

$$\begin{array}{l} \nearrow (\text{fail}^\emptyset \lambda x_2 \lambda x_3 \mathbf{1} \int \int) \langle \langle \lambda y, I \int / x_2, x_3 \rangle \rangle = N_1 \\ M \longrightarrow (y \lambda x_2 \lambda x_3 \mathbf{1} \int \int) \langle \langle \lambda \text{fail}^\emptyset, I \int / x_2, x_3 \rangle \rangle = N_2 \\ \searrow (I \lambda x_2 \lambda x_3 \mathbf{1} \int \int) \langle \langle \lambda \text{fail}^\emptyset, y \int / x_2, x_3 \rangle \rangle = N_3 \end{array}$$

▽

3.2 Resource Control for λ_c^ℓ via Intersection Types

Our type system for λ_c^ℓ is based on non-idempotent intersection types. As in prior works [4, 21], intersection types account for available resources in bags, which are unordered and have all the same type. Because we admit the term $\text{fail}^{\tilde{x}}$ as typable, we say that our system enforces *well-formedness* rather than *well-typedness*. As we will see, well-typed terms form the sub-class of well-formed terms that does not include $\text{fail}^{\tilde{x}}$ (see the text after Theorem 3).

Strict types (σ, τ, δ) and multiset types (π, ζ) are defined as follows:

$$\sigma, \tau, \delta ::= \mathbf{unit} \mid \pi \rightarrow \sigma \quad \pi, \zeta ::= \bigwedge_{i \in I} \sigma_i \mid \omega$$

Given a non-empty I , multiset types $\bigwedge_{i \in I} \sigma_i$ are given to bags of size $|I|$. This operator is associative, commutative, and non-idempotent (i.e., $\sigma \wedge \sigma \neq \sigma$), with identity ω . Notation σ^k stands for $\sigma \wedge \dots \wedge \sigma$ (k times, if $k > 0$) or ω (if $k = 0$).

Judgments have the form $\Gamma \vDash M : \tau$, with contexts defined as follows:

$$\Gamma, \Delta ::= - \mid \Gamma, x : \pi \mid \Gamma, x : \sigma$$

where $-$ denotes the empty context. We write $\text{dom}(\Gamma)$ for the set of variables in Γ . For $\Gamma, x : \pi$, we assume $x \notin \text{dom}(\Gamma)$. To avoid ambiguities, we write $x : \sigma^1$ to denote that the assignment involves a multiset type, rather than a strict type. Given Γ , its *core context* Γ^\downarrow concerns variables with types different from ω ; it is defined as $\Gamma^\downarrow = \{x : \pi \in \Gamma \mid \pi \neq \omega\}$.

$\frac{[\text{FS:var}^\ell]}{x : \sigma \vDash x : \sigma}$	$\frac{[\text{FS:1}^\ell]}{- \vDash \mathbf{1} : \omega}$	$\frac{[\text{FS:bag}^\ell]}{\Gamma \vDash N : \sigma \quad \Delta \vDash C : \sigma^k}{\Gamma, \Delta \vDash \{N \} \cdot C : \sigma^{k+1}}$	$\frac{[\text{FS:fail}]}{\text{dom}(\Gamma^\downarrow) = \tilde{x}}{\Gamma^\downarrow \vDash \text{fail}^{\tilde{x}} : \tau}$
$\frac{[\text{FS:weak}]}{\Gamma \vDash M : \tau}{\Gamma, x : \omega \vDash M[\leftarrow x] : \tau}$	$\frac{[\text{FS:shar}]}{\Gamma, x_1 : \sigma, \dots, x_k : \sigma \vDash M : \tau \quad k \neq 0}{\Gamma, x : \sigma^k \vDash M[x_1, \dots, x_k \leftarrow x] : \tau}$		
$\frac{[\text{FS:abs-sh}]}{\Gamma, x : \sigma^k \vDash M[\tilde{x} \leftarrow x] : \tau}{\Gamma \vDash \lambda x. (M[\tilde{x} \leftarrow x]) : \sigma^k \rightarrow \tau}$	$\frac{[\text{FS:app}]}{\Gamma \vDash M : \sigma^j \rightarrow \tau \quad \Delta \vDash C : \sigma^k}{\Gamma, \Delta \vDash M C : \tau}$		
$\frac{[\text{FS:Esub}]}{\Gamma, x : \sigma^j \vDash M[\tilde{x} \leftarrow x] : \tau \quad \Delta \vDash C : \sigma^k}{\Gamma, \Delta \vDash (M[\tilde{x} \leftarrow x]) \langle\langle C/x \rangle\rangle : \tau}$	$\frac{[\text{FS:Esub}^\ell]}{\Gamma, x_1 : \sigma, \dots, x_k : \sigma \vDash M : \tau \quad \Delta \vDash C : \sigma^k}{\Gamma, \Delta \vDash M \langle C/x_1, \dots, x_k \rangle : \tau}$		

Fig. 6. Well-Formedness Rules for λ_c^ℓ .

Definition 4 (Well-formedness in λ_c^ℓ). A term M is well-formed if there exists a context Γ and a type τ such that the rules in Fig. 6 entail $\Gamma \vDash M : \tau$.

In Fig. 6, Rule $[\text{FS} : \text{var}^\ell]$ types variables. Rule $[\text{FS} : \mathbf{1}^\ell]$ types the empty bag with ω . Rule $[\text{FS} : \text{bag}^\ell]$ types the concatenation of bags. Rule $[\text{FS} : \text{fail}]$ types the term $\text{fail}^{\tilde{x}}$ with a strict type τ , provided that the domain of the core context coincides with \tilde{x} (i.e., no variable in \tilde{x} is typed with ω). Rule $[\text{FS} : \text{weak}]$ types $M[\leftarrow x]$ by weakening the context with $x : \omega$. Rule $[\text{FS} : \text{shar}]$ types $M[\tilde{x} \leftarrow x]$ with τ , provided that there are assignments to the shared variables in \tilde{x} .

Rule $[\text{FS} : \text{abs-sh}]$ types an abstraction $\lambda x. (M[\tilde{x} \leftarrow x])$ with $\sigma^k \rightarrow \tau$, provided that $M[\tilde{x} \leftarrow x] : \tau$ can be entailed from an assignment $x : \sigma^k$. Rule $[\text{FS} : \text{app}]$ types $(M C)$, provided that M has type $\sigma^j \rightarrow \tau$ and C has type σ^k . Note that, unlike usual intersection type systems, j and k may differ. Rule $[\text{FS} : \text{Esub}]$ types the intermediate substitution of a bag C of type σ^k , provided that x has type σ^j ; again, j and k may differ. Rule $[\text{FS} : \text{Esub}^\ell]$ types $M \langle C/\tilde{x} \rangle$ as long as C has type $\sigma^{|\tilde{x}|}$, and each $x_i \in \tilde{x}$ is of type σ .

Well-formed terms satisfy subject reduction (SR), whereas *well-typed* terms, defined below, satisfy also subject expansion (SE). See [13] for details.

Theorem 3 (SR in λ_c^ℓ). If $\Gamma \vDash M : \tau$ and $M \longrightarrow M'$, then $\Gamma \vDash M' : \tau$.

From our system for well-formedness we can extract a system for *well-typed* terms, which do not include $\text{fail}^{\tilde{x}}$. Judgments for well-typedness are denoted $\Gamma \vdash M : \tau$, with rules copied from Fig. 6 (the rule name prefix FS is replaced with TS), with the following modifications: (i) Rule $[\text{TS:fail}]$ is removed; (ii) Rules $[\text{TS:app}]$ and $[\text{TS:Esub}]$ are modified to disallow a mismatch between

$$\begin{aligned}
\llbracket x \rrbracket_u &= \bar{x}.\text{some}; [x \leftrightarrow u] \\
\llbracket \lambda x.M \rrbracket_u &= \bar{u}.\text{some}; u(x); \llbracket M \rrbracket_u \\
\llbracket (M C) \rrbracket_u &= (\nu v)(\llbracket M \rrbracket_v \mid v.\text{some}_{u, \text{fv}(C)}; \bar{v}[x]; (\llbracket C \rrbracket_x \mid [v \leftrightarrow u])) \\
\llbracket M \langle\langle C/x \rangle\rangle \rrbracket_u &= (\nu x)(\llbracket M \rrbracket_u \mid \llbracket C \rrbracket_x) \\
\llbracket \{ N \} \cdot C \rrbracket_x &= x.\text{some}_{\text{fv}(C), \text{fv}(N)}; x(y_i); x.\text{some}_{y_i, \text{fv}(C), \text{fv}(N)}; \bar{x}[z_i]; \\
&\quad (z_i.\text{some}_{\text{fv}(N)}; \llbracket N \rrbracket_{z_i} \mid \llbracket C \rrbracket_x \mid \bar{y}_i.\text{none}) \\
\llbracket \mathbf{1} \rrbracket_x &= x.\text{some}_\emptyset; x(y_n); (\bar{y}_n.\text{some}; \bar{y}_n \mid x.\text{some}_\emptyset; \bar{x}.\text{none}) \\
\llbracket M \langle\{ N_1, N_2 \} / x_1, x_2 \rangle \rrbracket_u &= (\nu z_1)(z_1.\text{some}_{\text{fv}(N_1)}; \llbracket N_1 \rrbracket_{z_1} \mid (\nu z_2)(z_2.\text{some}_{\text{fv}(N_2)}; \llbracket N_2 \rrbracket_{z_2} \\
&\quad \mid \bigoplus_{x_i \in \{x_1, x_2\}} \bigoplus_{x_j \in \{x_1, x_2 \setminus x_i\}} \llbracket M \rrbracket_u \{z_1/x_i\} \{z_2/x_j\})) \\
\llbracket M[\leftarrow x] \rrbracket_u &= \bar{x}.\text{some}; \bar{x}[y_i]; (y_i.\text{some}_{u, \text{fv}(M)}; y_i(); \llbracket M \rrbracket_u \mid \bar{x}.\text{none}) \\
\llbracket M[\tilde{x} \leftarrow x] \rrbracket_u &= \bar{x}.\text{some}; \bar{x}[y_i]; (y_i.\text{some}_\emptyset; y_i(); \mathbf{0} \mid \bar{x}.\text{some}; x.\text{some}_{u, \text{fv}(M) \setminus \tilde{x}}; \\
&\quad \bigoplus_{x_i \in \tilde{x}} x(x_i); \llbracket M[(\tilde{x} \setminus x_i) \leftarrow x] \rrbracket_u) \\
\llbracket \text{fail}^{x_1, \dots, x_k} \rrbracket_u &= \bar{u}.\text{none} \mid \bar{x}_1.\text{none} \mid \dots \mid \bar{x}_k.\text{none}
\end{aligned}$$

Fig. 7. Translation of λ_c^ℓ into $\mathfrak{s}\pi^+$.

variables and resources, i.e., multiset types should match in size. Well-typed terms are also well-formed, and thus satisfy SR. Moreover, as a consequence of adopting (non-idempotent) intersection types, they also satisfy SE:

Theorem 4 (SE in λ_c^ℓ). *If $\Gamma \vdash M' : \tau$ and $M \longrightarrow M'$, then $\Gamma \vdash M : \tau$.*

4 A Typed Translation of λ_c^ℓ into $\mathfrak{s}\pi^+$

While $\mathfrak{s}\pi^+$ features non-deterministic choice, λ_c^ℓ is a prototypical programming language in which implicit non-determinism implements fetching of resources. Resources are controlled using different type systems (session types in $\mathfrak{s}\pi^+$, intersection types in λ_c^ℓ). To reconcile these differences and illustrate the potential of $\mathfrak{s}\pi^+$ to precisely model non-determinism as found in realistic programs/protocols, we give a translation of λ_c^ℓ into $\mathfrak{s}\pi^+$. This translation preserves types (Theorem 5) and respects well-known criteria for dynamic correctness [11, 23, 24] (Theorem 6).

The Translation. Given a λ_c^ℓ -term M , its translation into $\mathfrak{s}\pi^+$ is denoted $\llbracket M \rrbracket_u$ and given in Fig. 7. As usual, every variable x in M becomes a name x in process $\llbracket M \rrbracket_u$, where name u provides the behavior of M . A peculiarity is that, to handle failures in λ_c^ℓ , u is a non-deterministically available session: the translated term can be available or not, as signaled by prefixes $\bar{u}.\text{some}$ and $\bar{u}.\text{none}$, respectively. As a result, reductions from $\llbracket M \rrbracket_u$ include synchronizations that codify M 's behavior but also synchronizations that confirm a session's availability.

At its core, our translation follows Milner's. This way, e.g., the process $\llbracket (\lambda x.M) C \rrbracket_u$ enables synchronizations between $\llbracket \lambda x.M \rrbracket_v$ and $\llbracket C \rrbracket_x$ along name v ,

$$\begin{array}{l}
 \llbracket \mathbf{unit} \rrbracket = \& \mathbf{1} \qquad \llbracket \sigma^k \rightarrow \tau \rrbracket = \& (\overline{\llbracket \sigma^k \rrbracket}_{(\sigma,i)}} \wp \llbracket \tau \rrbracket) \\
 \llbracket \sigma \wedge \pi \rrbracket_{(\tau,i)} = \oplus ((\& \mathbf{1}) \wp (\oplus \& ((\oplus \llbracket \sigma \rrbracket) \otimes (\llbracket \pi \rrbracket_{(\tau,i)})))) \\
 \llbracket \omega \rrbracket_{(\sigma,i)} = \begin{cases} \oplus ((\& \mathbf{1}) \wp (\oplus \& \mathbf{1})) & \text{if } i = 0 \\ \oplus ((\& \mathbf{1}) \wp (\oplus \& ((\oplus \llbracket \sigma \rrbracket) \otimes (\llbracket \omega \rrbracket_{(\sigma,i-1)})))) & \text{if } i > 0 \end{cases}
 \end{array}$$

Fig. 8. Translation of intersection types into session types (cf. Definition 5).

resulting in the translation of an intermediate substitution. The *key novelty* is the role and treatment of non-determinism. Accommodating non-confluent non-determinism is non-trivial, as it entails translating explicit substitutions and sharing in λ_c^ℓ using the non-deterministic choice operator \ddagger in $\mathfrak{s}\pi^+$. Next we discuss these novel aspects, while highlighting differences with respect to a translation by Paulus et al. [22], which is given in the confluent setting (see Section 5).

In Fig. 7, non-deterministic choices occur in the translations of $M\langle C/\tilde{x} \rangle$ (explicit substitutions) and $M[\tilde{x} \leftarrow x]$ (non-empty sharing). Roughly speaking, the position of \ddagger in the translation of $M\langle C/\tilde{x} \rangle$ represents the most desirable way of mimicking the fetching of terms from a bag. This use of \ddagger is a central idea in our translation: as we explain below, it allows for appropriate commitment in non-deterministic choices, but also for *delayed* commitment when necessary.

For simplicity, we consider explicit substitutions $M\langle C/\tilde{x} \rangle$ where $C = \wr N_1, N_2 \wr$ and $\tilde{x} = x_1, x_2$. The translation $\llbracket M\langle C/\tilde{x} \rangle \rrbracket_u$ uses the processes $\llbracket N_i \rrbracket_{z_i}$, where each z_i is fresh. First, each bag item confirms its behavior. Then, a variable $x_i \in \tilde{x}$ is chosen non-deterministically; we ensure that these choices consider all variables. Note that writing $\ddagger_{x_i \in \{x_1, x_2\}} \ddagger_{x_j \in \{x_1, x_2\} \setminus x_i}$ is equivalent to non-deterministically assigning x_i, x_j to each permutation of x_1, x_2 . The resulting choice involves $\llbracket M \rrbracket_u$ with x_i, x_j substituted by z_1, z_2 . Commitment here is triggered only via synchronizations along z_1 or z_2 ; synchronizing with $z_i.\mathbf{some}_{\text{fv}(N_i)}$; $\llbracket N_i \rrbracket_{z_i}$ then represents fetching N_i from the bag. The size of the translated term $\llbracket M\langle C/\tilde{x} \rangle \rrbracket_u$ is exponential with respect to the size of C .

The process $\llbracket M[\tilde{x} \leftarrow x] \rrbracket_u$ proceeds as follows. First, it confirms its behavior along x . Then it sends a name y_i on x , on which a failed reduction may be handled. Next, the translation confirms again its behavior along x and non-deterministically receives a reference to an $x_i \in \tilde{x}$. Each branch consists of $\llbracket M[(\tilde{x} \setminus x_i) \leftarrow x] \rrbracket_u$. The possible choices are permuted, represented by $\ddagger_{x_i \in \tilde{x}}$. Synchronizations with $\llbracket M[(\tilde{x} \setminus x_i) \leftarrow x] \rrbracket_u$ and bags delay commitment in this choice (we return to this point below). The process $\llbracket M[\leftarrow x] \rrbracket_u$ is similar but simpler: here the name x fails, as it cannot take further elements to substitute.

In case of a failure (i.e., a mismatch between the size of the bag C and the number of variables in M), our translation ensures that the confirmations of C will not succeed. This is how failure in λ_c^ℓ is correctly translated to failure in $\mathfrak{s}\pi^+$.

Translation Correctness. The translation is typed: intersection types in $\lambda_{\mathcal{C}}^{\ell}$ are translated into session types in $\mathfrak{s}\pi^+$ (Fig. 8). This translation of types abstractly describes how non-deterministic fetches are codified as non-deterministic session protocols. It is worth noting that this translation of types is the same as in [22]. This is not surprising: as we have seen, session types effectively abstract away from the behavior of processes, as all branches of a non-deterministic choice use the same typing context. Still, it is pleasant that the translation of types remains unchanged across different translations with our (non-confluent) non-determinism (in Fig. 7) and with confluent non-determinism (in [22]).

To state *static* correctness, we require the following definition:

Definition 5. *Let $\Gamma = x_1 : \sigma_1, \dots, x_m : \sigma_m, v_1 : \pi_1, \dots, v_n : \pi_n$ be a context. The translation $\llbracket \cdot \rrbracket_-$ in Fig. 8 extends to contexts as follows:*

$$\llbracket \Gamma \rrbracket = x_1 : \&\overline{\llbracket \sigma_1 \rrbracket}, \dots, x_m : \&\overline{\llbracket \sigma_m \rrbracket}, v_1 : \overline{\llbracket \pi_1 \rrbracket}_{(\sigma, i_1)}, \dots, v_n : \overline{\llbracket \pi_n \rrbracket}_{(\sigma, i_n)}$$

Well-formed terms translate into well-typed processes:

Theorem 5. *If $\Gamma \vDash M : \tau$, then $\llbracket M \rrbracket_u \vdash \llbracket \Gamma \rrbracket, u : \llbracket \tau \rrbracket$.*

To state *dynamic* correctness, we rely on established notions that (abstractly) characterize *correct translations*. A language $\mathcal{L} = (L, \rightarrow)$ consists of a set of terms L and a reduction relation \rightarrow on L . Each language \mathcal{L} is assumed to contain a success constructor \checkmark . A term $T \in L$ has *success*, denoted $T \Downarrow \checkmark$, when there is a sequence of reductions (using \rightarrow) from T to a term satisfying success criteria.

Given $\mathcal{L}_1 = (L_1, \rightarrow_1)$ and $\mathcal{L}_2 = (L_2, \rightarrow_2)$, we seek translations $\llbracket \cdot \rrbracket : L_1 \rightarrow L_2$ that are correct: they satisfy well-known correctness criteria [11, 23, 24]. We state the set of correctness criteria that determine the correctness of a translation.

Definition 6 (Correct Translation). *Let $\mathcal{L}_1 = (\mathcal{M}, \rightarrow_1)$ and $\mathcal{L}_2 = (\mathcal{P}, \rightarrow_2)$ be two languages. Let \simeq_2 be an equivalence over \mathcal{L}_2 . We use M, M' (resp. P, P') to range over terms in \mathcal{M} (resp. \mathcal{P}). Given a translation $\llbracket \cdot \rrbracket : \mathcal{M} \rightarrow \mathcal{P}$, we define:*

Completeness: *For every M, M' such that $M \rightarrow_1^* M'$, there exists P such that $\llbracket M \rrbracket \rightarrow_2^* P \simeq_2 \llbracket M' \rrbracket$.*

Weak Soundness: *For every M and P such that $\llbracket M \rrbracket \rightarrow_2^* P$, there exist M', P' such that $M \rightarrow_1^* M'$ and $P \rightarrow_2^* P' \simeq_2 \llbracket M' \rrbracket$.*

Success Sensitivity: *For every M , we have $M \Downarrow \checkmark$ if and only if $\llbracket M \rrbracket \Downarrow \checkmark$.*

Let us write Λ to denote the set of well-formed $\lambda_{\mathcal{C}}^{\ell}$ terms, and Π for the set of all well-typed $\mathfrak{s}\pi^+$ processes, both including \checkmark . We have our final result:

Theorem 6 (Translation correctness under \rightsquigarrow). *The translation $\llbracket \cdot \rrbracket_- : (\Lambda, \longrightarrow) \rightarrow (\Pi, \rightsquigarrow)$ is correct (cf. Definition 6) using equivalence \equiv (Fig. 1).*

The proof of Theorem 6 involves instantiating/proving each of the parts of Definition 6. Among these, *weak soundness* is the most challenging to prove. Prior work on translations of typed λ into π with confluent non-determinism [22] rely

critically on confluence to match a behavior in π with a corresponding behavior in λ . Because in our setting confluence is lost, we must resort to a different proof.

As already discussed, our translation makes the implicit non-determinism in a λ_C^ℓ -term M explicit by adding non-deterministic choices in key points of $\llbracket M \rrbracket_u$. Our reduction \rightsquigarrow preserves those branches that simultaneously have the same prefix available (up to \bowtie). In proving weak soundness, we exploit the fact that reduction entails delayed commitment. To see this, consider the following terms:

$$(\nu x)((\alpha_1; P_1 \# \alpha_2; P_2) \mid Q) \quad (4)$$

$$(\nu x)(\alpha_1; P_1 \mid Q) \# (\nu x)(\alpha_2; P_2 \mid Q) \quad (5)$$

In (4), commitment to a choice relies on whether $\alpha_1 \bowtie \alpha_2$ holds (cf. Definition 3). If $\alpha_1 \not\bowtie \alpha_2$, a choice is made; otherwise, commitment is delayed, and depends on P_1 and P_2 . Hence, in (4) the possibility of committing to either branch is kept open. In contrast, in (5) commitment to a choice is independent of $\alpha_1 \bowtie \alpha_2$.

Our translation exploits the delayed commitment of non-determinism illustrated by (4) to mimic commitment to non-deterministic choices in λ_C^ℓ , which manifests in fetching resources from bags. The fact that this delayed commitment preserves information about the different branches (e.g., P_1 and P_2 in (4)) is essential to establish weak soundness, i.e., to match a behavior in $s\pi^+$ with a corresponding step in λ_C^ℓ . In contrast, forms of non-determinism in $\llbracket N \rrbracket_u$ that resemble (5) are useful to characterize behaviors different from fetching.

5 Summary and Related Work

We studied the interplay between resource control and non-determinism in typed calculi. We introduced $s\pi^+$ and λ_C^ℓ , two calculi with non-confluent non-determinism, both with type systems for resource control. Inspired by the untyped π -calculus, non-determinism in $s\pi^+$ is lazy and explicit, with session types defined following ‘propositions-as-sessions’ [5]. In λ_C^ℓ , non-determinism arises in the fetching of resources, and is regulated by intersection types. A correct translation of λ_C^ℓ into $s\pi^+$ precisely connects their different forms of non-determinism.

Related Work. Integrating (non-confluent) non-determinism within session types is non-trivial, as carelessly discarding branches would break typability. Work by Caires and Pérez [5], already mentioned, develops a confluent semantics by requiring that non-determinism is only used inside the monad $\&\mathcal{A}$; our non-confluent semantics drops this requirement. This allows us to consider non-deterministic choices not possible in [5], such as, e.g., selections of different labels. We stress that linearity is not jeopardized: the branches of ‘ $\#$ ’ do not represent *different sessions*, but *different implementations* of the same sessions.

Atkey et al. [1] and Kokke et al. [16] extend ‘propositions-as-sessions’ with non-determinism. Their approaches are very different (conflation of the additives and bounded linear logic, respectively) and support non-determinism for unrestricted names only. Also, [1, 16] do not connect with typed λ -calculi, as we do.

Rocha and Caires also consider non-determinism, relying on confluence in [25] and on unrestricted names in [26]. Casal et al. [7, 30] develop a type system for *mixed sessions* (sessions with mixed choices), which can express non-determinism but does not ensure deadlock-freedom. Ensuring deadlock-freedom by typing is a key feature of the ‘propositions-as-sessions’ approach that we adopt for $\mathfrak{s}\pi^+$.

Our language λ_c^ℓ is most related to calculi by Boudol [3], Boudol and Laneve [4], and by Pagani and Ronchi Della Rocca [21]. Non-determinism in the calculi in [3, 4] is committing and implicit; their linear resources can be consumed *at most* once, rather than *exactly* once. The work [21] considers non-committing non-determinism that is both implicit (as in λ_c^ℓ) and explicit (via a sum operator on terms). Both [3, 21] develop (non-idempotent) intersection type systems to regulate resources. In our type system, all terms in a bag have the same type; the system in [21] does not enforce this condition. Unlike these type systems, our system for well-formedness can type terms with a lack or an excess of resources.

Boudol and Laneve [4] and Paulus et al. [22] give translations of resource λ -calculi into π . The translation in [4] is used to study the semantics induced upon λ -terms by a translation into π ; unlike ours, it does not consider types. As already mentioned in Sect. 4, Paulus et al. [22] relate calculi with *confluent* non-determinism: a resource λ -calculus with sums on terms, and the session π -calculus from [5]. Our translation of terms and that in [22] are very different: while here we use non-deterministic choice to mimic the sharing construct, the translation in [22] uses it to translate bags. Hence, our Theorem 6 cannot be derived from [22].

The last decade of work on ‘propositions-as-sessions’ has delivered insightful connections with typed λ -calculi—see, e.g., [28, 29, 31]. Excepting [22], already discussed, none of these works consider non-deterministic λ -calculi.

Acknowledgments. We are grateful to the anonymous reviewers for useful comments on previous versions of this paper. We are also grateful to Mariangiola Dezani for her encouragement and suggestions. This research has been supported by the Dutch Research Council (NWO) under project No. 016.Vidi.189.046 (‘Unifying Correctness for Communicating Software’) and the EPSRC Fellowship ‘VeTSpec: Verified Trustworthy Software Specification’ (EP/R034567/1).

References

1. Atkey, R., Lindley, S., Morris, J.G.: Conflation confers concurrency. In: Lindley, S., McBride, C., Trinder, P., Sannella, D. (eds.) *A List of Successes That Can Change the World*. LNCS, vol. 9600, pp. 32–55. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-30936-1_2
2. Berger, M., Honda, K.: The two-phase commitment protocol in an extended pi-calculus. In: Aceto, L., Victor, B. (eds.) *7th International Workshop on Expressiveness in Concurrency, EXPRESS 2000, Satellite Workshop of CONCUR 2000*, State College, PA, USA, 21 August 2000. *Electronic Notes in Theoretical Computer Science*, vol. 39, pp. 21–46. Elsevier (2000). [https://doi.org/10.1016/S1571-0661\(05\)82502-2](https://doi.org/10.1016/S1571-0661(05)82502-2)

3. Boudol, G.: The lambda-calculus with multiplicities. In: Best, E. (ed.) CONCUR 1993. LNCS, vol. 715, pp. 1–6. Springer, Heidelberg (1993). https://doi.org/10.1007/3-540-57208-2_1
4. Boudol, G., Laneve, C.: Lambda-calculus, multiplicities, and the pi-calculus. In: Proof, Language, and Interaction, Essays in Honour of Robin Milner, pp. 659–690 (2000)
5. Caires, L., Pérez, J.A.: Linearity, control effects, and behavioral types. In: Yang, H. (ed.) ESOP 2017. LNCS, vol. 10201, pp. 229–259. Springer, Heidelberg (2017). https://doi.org/10.1007/978-3-662-54434-1_9
6. Caires, L., Pfenning, F.: Session types as intuitionistic linear propositions. In: Gastin, P., Laroussinie, F. (eds.) CONCUR 2010. LNCS, vol. 6269, pp. 222–236. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-15375-4_16
7. Casal, F., Mordido, A., Vasconcelos, V.T.: Mixed sessions. *Theor. Comput. Sci.* **897**, 23–48 (2022). <https://doi.org/10.1016/j.tcs.2021.08.005>
8. de'Liguoro, U., Piperno, A.: Non deterministic extensions of untyped lambda-calculus. *Inf. Comput.* **122**(2), 149–177 (1995). <https://doi.org/10.1006/inco.1995.1145>
9. Dezani-Ciancaglini, M.: Logical semantics for concurrent lambda-calculus. Ph.D. thesis, Nijmegen University (1996). <https://www.di.unito.it/~dezani/papers/tesi.ps>
10. Ehrhard, T., Regnier, L.: The differential lambda-calculus. *Theor. Comput. Sci.* **309**(1–3), 1–41 (2003). [https://doi.org/10.1016/S0304-3975\(03\)00392-X](https://doi.org/10.1016/S0304-3975(03)00392-X)
11. Gorla, D.: Towards a unified approach to encodability and separation results for process calculi. *Inf. Comput.* **208**(9), 1031–1053 (2010). <https://doi.org/10.1016/j.ic.2010.05.002>
12. Groote, J.F., Mousavi, M.R.: Modeling and Analysis of Communicating Systems. MIT Press (2014). <https://mitpress.mit.edu/books/modeling-and-analysis-communicating-systems>
13. van den Heuvel, B., Paulus, J.W.N., Nantes-Sobrinho, D., Pérez, J.A.: Typed non-determinism in functional and concurrent calculi (extended version). CoRR abs/2205.00680 (2022). <https://doi.org/10.48550/arXiv.2205.00680>
14. Honda, K.: Types for dyadic interaction. In: Best, E. (ed.) CONCUR 1993. LNCS, vol. 715, pp. 509–523. Springer, Heidelberg (1993). https://doi.org/10.1007/3-540-57208-2_35
15. Honda, K., Vasconcelos, V.T., Kubo, M.: Language primitives and type discipline for structured communication-based programming. In: Hankin, C. (ed.) ESOP 1998. LNCS, vol. 1381, pp. 122–138. Springer, Heidelberg (1998). <https://doi.org/10.1007/BFb0053567>
16. Kokke, W., Morris, J.G., Wadler, P.: Towards races in linear logic. *Log. Meth. Comput. Sci.* **16**(4) (2020). [https://doi.org/10.23638/LMCS-16\(4:15\)2020](https://doi.org/10.23638/LMCS-16(4:15)2020)
17. Milner, R.: Functions as processes. Research Report 1154, INRIA, Sophia Antipolis (1990). Final version appeared as [18]
18. Milner, R.: Functions as processes. *Math. Struct. Comput. Sci.* **2**(2), 119–141 (1992). <https://doi.org/10.1017/S0960129500001407>
19. Milner, R., Parrow, J., Walker, D.: A calculus of mobile processes. I. *Inf. Comput.* **100**(1), 1–40 (1992). [https://doi.org/10.1016/0890-5401\(92\)90008-4](https://doi.org/10.1016/0890-5401(92)90008-4)
20. Nestmann, U., Fuzzati, R., Merro, M.: Modeling consensus in a process calculus. In: Amadio, R., Lugiez, D. (eds.) CONCUR 2003. LNCS, vol. 2761, pp. 399–414. Springer, Heidelberg (2003). https://doi.org/10.1007/978-3-540-45187-7_26

21. Pagani, M., della Rocca, S.R.: Solvability in resource lambda-calculus. In: Ong, L. (ed.) FoSSaCS 2010. LNCS, vol. 6014, pp. 358–373. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-12032-9_25
22. Paulus, J.W.N., Nantes-Sobrinho, D., Pérez, J.A.: Non-deterministic functions as non-deterministic processes. In: Kobayashi, N. (ed.) 6th International Conference on Formal Structures for Computation and Deduction, FSCD 2021, Buenos Aires, Argentina (Virtual Conference), 17–24 July 2021. LIPIcs, vol. 195, pp. 21:1–21:22. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2021). <https://doi.org/10.4230/LIPIcs.FSCD.2021.21>, Extended version on <https://arxiv.org/abs/2104.14759>
23. Peters, K.: Translational expressiveness. Comparing process calculi using encodings. Ph.D. thesis, Berlin Institute of Technology (2012). <https://doi.org/10.14279/depositonce-3416>
24. Peters, K.: Comparing process calculi using encodings. In: Pérez, J.A., Rot, J. (eds.) Proceedings Combined 26th International Workshop on Expressiveness in Concurrency and 16th Workshop on Structural Operational Semantics, EXPRESS/SOS 2019, Amsterdam, The Netherlands, 26th August 2019. EPTCS, vol. 300, pp. 19–38 (2019). <https://doi.org/10.4204/EPTCS.300.2>
25. Rocha, P., Caires, L.: Propositions-as-types and shared state. Proc. ACM Program. Lang. **5**(ICFP), 79:1–79:30 (2021). <https://doi.org/10.1145/3473584>
26. Rocha, P., Caires, L.: Safe session-based concurrency with shared linear state. In: Wies, T. (ed.) Programming Languages and Systems. LNCS, vol. 2072, pp. 421–450. Springer, Cham (2023). https://doi.org/10.1007/978-3-031-30044-8_16
27. Sangiorgi, D., Walker, D.: The Pi-Calculus - A Theory of Mobile Processes. Cambridge University Press, Cambridge (2001)
28. Toninho, B., Caires, L., Pfenning, F.: Functions as session-typed processes. In: Birkedal, L. (ed.) FoSSaCS 2012. LNCS, vol. 7213, pp. 346–360. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-28729-9_23
29. Toninho, B., Yoshida, N.: On polymorphic sessions and functions. In: Ahmed, A. (ed.) ESOP 2018. LNCS, vol. 10801, pp. 827–855. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-89884-1_29
30. Vasconcelos, V.T., Casal, F., Almeida, B., Mordido, A.: Mixed sessions. In: ESOP 2020. LNCS, vol. 12075, pp. 715–742. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-44914-8_26
31. Wadler, P.: Propositions as sessions. In: Thiemann, P., Findler, R.B. (eds.) ACM SIGPLAN International Conference on Functional Programming, ICFP 2012, Copenhagen, Denmark, 9–15 September 2012, pp. 273–286. ACM (2012). <https://doi.org/10.1145/2364527.2364568>