Towards Races in Linear Logic

Wen Kokke, J. Garrett Morris, Philip Wadler

► To cite this version:

Wen Kokke, J. Garrett Morris, Philip Wadler. Towards Races in Linear Logic. 21th International Conference on Coordination Languages and Models (COORDINATION), Jun 2019, Kongens Lyngby, Denmark. pp.37-53, 10.1007/978-3-030-22397-7_3 . hal-02365505

HAL Id: hal-02365505

<https://hal.inria.fr/hal-02365505>

Submitted on 15 Nov 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Distributed under a Creative Commons Attribution| 4.0 International License

Towards Races in Linear Logic

Wen Kokke¹, J. Garrett Morris², and Philip Wadler¹

¹ University of Edinburgh, Edinburgh, UK,
{wen.kokke@ed.ac.uk, wadler@inf.ed.ac.uk}

² University of Kansas, Lawrence, KS, USA,
garrett@ittc.ku.edu

Abstract Process calculi based in logic, such as π DILL and CP, provide a foundation for deadlock-free concurrent programming, but exclude non-determinism and races. HCP is a reformulation of CP which addresses a fundamental shortcoming: the fundamental operator for parallel composition from the π -calculus does not correspond to any rule of linear logic, and therefore not to any term construct in CP. We introduce HCP_{ND}^- , which extends HCP with a novel account of non-determinism. Our approach draws on bounded linear logic to provide a strongly-typed account of standard process calculus expressions of non-determinism. We show that our extension is expressive enough to capture many uses of non-determinism in untyped calculi, such as non-deterministic choice, while preserving HCP’s meta-theoretic properties, including deadlock freedom.

Keywords: π -calculus · linear logic · session types · non-determinism · deadlock freedom

1 Introduction

Consider the following scenario:

Ami and Boé are working from home one morning when they each get a craving for a slice of cake. Being denizens of the web, they quickly find the nearest store which does home deliveries. Unfortunately for them, they both order their cake at the *same* store, which has only one slice left. After that, all it can deliver is disappointment.

This is an example of a *race condition*. We can model this scenario in the π -calculus, where 🧑, 🧑 and 🏪 are processes modelling Ami, Boé and the store, and 🍰 and 🙅 are channels giving access to a slice of cake and disappointment, respectively. This process has two possible outcomes: either Ami gets the cake, and Boé gets disappointment, or vice versa.

$$\begin{aligned} & (x[\text{🍰}].x[\text{🙅}].\text{🏪} \mid x(y).\text{🧑} \mid x(z).\text{🧑}) \\ & \quad \Downarrow^* \\ & (\text{🏪} \mid \text{🧑}\{\text{🍰}/y\} \mid \text{🧑}\{\text{🙅}/z\}) \text{ or } (\text{🏪} \mid \text{🧑}\{\text{🙅}/y\} \mid \text{🧑}\{\text{🍰}/z\}) \end{aligned}$$

While Ami or Boé may not like all of the outcomes, it is the store which is responsible for implementing the online delivery service, and the store is happy with either outcome. Thus, the above is an interaction we would like to be able to model.

Now consider another scenario, which takes place *after* Ami has already bought the cake:

Boé is *really* disappointed when she finds out the cake has sold out. Ami, always looking to make some money, offers to sell the slice to her for a profit. Boé agrees to engage in a little bit of back-alley cake resale, but sadly there is no trust between the two. Ami demands payment first. Boé would rather get her slice of cake before she gives Ami the money.

This is an example of a *deadlock*. We can also model this scenario in the π -calculus, where € is a channel giving access to some adequate amount of money.

$$(x(z).y[\text{🍰}].\text{👤} \mid y(w).x[\text{€}].\text{👤}) \not\Rightarrow^*$$

The above process does not reduce. As both Ami and Boé would prefer the exchange to be made, this interaction is desired by *neither*. Thus, the above is an interaction we would like to exclude.

Session types [10] statically guarantee that concurrent programs, such as those above, respect communication protocols. Session-typed calculi with logical foundations, such as π DILL [8] and CP [16], obtain deadlock freedom as a result of a close correspondence with logic. These systems, however, also rule out non-determinism and race conditions. In this paper, we demonstrate that logic-inspired type systems need not rule out races.

We present HCP_{ND}^- , an extension of CP with a novel account of non-determinism and races. Inspired by bounded linear logic [9], we introduce a form of shared channels in which the type of a shared channel tracks how many times it is reused. As in the untyped π -calculus, sharing introduces the potential for non-determinism. We show that our approach is sufficient to capture practical examples of races, such as an online store, as well as other formal characterizations of non-determinism, such as non-deterministic choice. However, HCP_{ND}^- does not lose the meta-theoretical benefits of CP: we show that it enjoys termination and deadlock-freedom.

An important limitation of our work is that types in HCP_{ND}^- explicitly count the potential races on a channel. It works fine when there are two or three races, but not n for an arbitrary n . The latter case is obviously important, and we see the main value of our work as a stepping stone to this more general case.

HCP_{ND}^- is based on HCP [11,12]. HCP is a reformulation of CP which addresses a fundamental shortcoming: the fundamental operator for parallel composition from the π -calculus does not correspond to any rule of linear logic, and therefore not to any term construct in CP.

There are two versions of HCP: a version with delayed actions, introduced by Kokke, Montesi, and Peressotti [12]; and a version without delayed actions, introduced by Kokke, Montesi, and Peressotti [11], referred to as HCP^- . In

this work, we will base ourselves on the latter, as the former does not yet have reduction semantics.

This paper proceeds as follows. In section 2, we discuss recent approaches to non-determinism in logic-inspired session-typed process calculi. In section 3, we introduce a variant of CP and prove progress and preservation. In section 4, we introduce HCP_{ND}^- . In section 5, we discuss cuts with leftovers. Finally, in section 7, we conclude with a discussion of the work done in this paper and potential avenues for future work.

2 Non-determinism, Logic, and Session Types

Recent work extended πDILL and CP with operators for non-deterministic behaviour [1,6,7]. These extensions all implement an operator known as non-deterministic local choice. (This operator is written as $P + Q$, but should not be confused with input-guarded choice from the π -calculus [14].) Non-deterministic local choice can be summarised by the following typing and reduction rules:

$$\frac{P \vdash \Gamma \quad Q \vdash \Gamma}{P + Q \vdash \Gamma} \qquad \begin{array}{l} P + Q \Rightarrow P \\ P + Q \Rightarrow Q \end{array}$$

Local choice introduces non-determinism explicitly, by listing all possible choices. This is unlike the π -calculus, where non-determinism arises due to multiple processes communicating on shared channels. We can easily implement local choice in the π -calculus, using a nullary communication:

$$\begin{array}{c} (x[] . 0 \mid x().P \mid x().Q) \\ \Downarrow_* \\ (P \mid x().Q) \quad \text{or} \quad (x().P \mid Q) \end{array}$$

In this implementation, the process $x[] . 0$ will “unlock” either P or Q , leaving the other process deadlocked. Or we could use input-guarded choice:

$$(x[] . 0 \mid (x().P + x().Q))$$

However, there are many non-deterministic processes in the π -calculus that are awkward to encode using non-deterministic local choice. Let us recall our example:

$$\begin{array}{c} (x[\text{🍰}].x[\text{👋}].\text{📺} \mid x(y).\text{👤} \mid x(z).\text{👤}) \\ \Downarrow_* \\ (\text{📺} \mid \text{👤}\{\text{🍰}/y\} \mid \text{👤}\{\text{👋}/z\}) \quad \text{or} \quad (\text{📺} \mid \text{👤}\{\text{👋}/y\} \mid \text{👤}\{\text{🍰}/z\}) \end{array}$$

This non-deterministic interaction involves communication. If we wanted to write down a process which exhibited the same behaviour using non-deterministic local choice, we would have to write the following process:

$$\begin{aligned}
& (x[\text{🍰}].y[\text{👋}].\text{👤} \mid x(z).\text{👤} \mid y(w).\text{👤}) + (y[\text{🍰}].x[\text{👋}].\text{👤} \mid x(z).\text{👤} \mid y(w).\text{👤}) \\
& \quad \Downarrow \star \\
& (\text{👤} \mid \text{👤}\{\text{🍰}/y\} \mid \text{👤}\{\text{👋}/z\}) \text{ or } (\text{👤} \mid \text{👤}\{\text{👋}/y\} \mid \text{👤}\{\text{🍰}/z\})
\end{aligned}$$

In essence, instead of modelling a non-deterministic interaction, we are enumerating the resulting deterministic interactions. This means non-deterministic local choice cannot model non-determinism in the way the π -calculus does. Enumerating all possible outcomes becomes worse the more processes are involved in an interaction. Imagine the following scenario:

Three customers, Ami, Boé, and Cat, have a craving for cake. Should cake be sold out, however, well... a doughnut will do. They prepare to order their goods via an online store. Unfortunately, they all decide to use the same *shockingly* under-stocked store, which has only one slice of cake, and a single doughnut. After that, all it can deliver is disappointment.

We can model this scenario in the π -calculus, where 🧑, 🧑, 🧑, and 🛒 are four processes modelling Ami, Boé, Cat, and the store, and 🍰, 🍩, and 🙄 are three channels giving access to a slice of cake, a so-so doughnut, and disappointment, respectively.

$$\begin{aligned}
& (x[\text{🍰}].x[\text{🍩}].x[\text{🙄}].\text{🛒} \mid x(y).\text{🧑} \mid x(z).\text{🧑} \mid x(w).\text{🧑}) \\
& \quad \Downarrow \star \\
& (\text{🛒} \mid \text{🧑}\{\text{🍰}/y\} \mid \text{🧑}\{\text{🍩}/z\} \mid \text{🧑}\{\text{🙄}/w\}) \text{ or } (\text{🛒} \mid \text{🧑}\{\text{🍰}/y\} \mid \text{🧑}\{\text{🙄}/z\} \mid \text{🧑}\{\text{🍩}/w\}) \\
& (\text{🛒} \mid \text{🧑}\{\text{🍩}/y\} \mid \text{🧑}\{\text{🙄}/z\} \mid \text{🧑}\{\text{🍰}/w\}) \text{ or } (\text{🛒} \mid \text{🧑}\{\text{🍩}/y\} \mid \text{🧑}\{\text{🍰}/z\} \mid \text{🧑}\{\text{🙄}/w\}) \\
& (\text{🛒} \mid \text{🧑}\{\text{🙄}/y\} \mid \text{🧑}\{\text{🍰}/z\} \mid \text{🧑}\{\text{🍩}/w\}) \text{ or } (\text{🛒} \mid \text{🧑}\{\text{🙄}/y\} \mid \text{🧑}\{\text{🍩}/z\} \mid \text{🧑}\{\text{🍰}/w\})
\end{aligned}$$

With the addition of one process, modelling Cat, we have increased the number of possible outcomes enormously! In general, the number of outcomes for these types of scenarios is $n!$, where n is the number of processes. This means that if we wish to translate any non-deterministic process to one using non-deterministic local choice, we can expect a factorial growth in the size of the term.

3 Hypersequent Classical Processes

In this section, we introduce HCP [11,12], the basis for our calculus HCP_{ND}^- . The term language for HCP^- is a variant of the π -calculus [14].

Definition 1 (Terms).

$P, Q, R ::= x \leftrightarrow y$	<i>link</i>
0	<i>terminated process</i>
$(\nu x)P$	<i>name restriction, “cut”</i>
$(P \mid Q)$	<i>parallel composition, “mix”</i>
$x[y].P$	<i>output</i>
$x(y).P$	<i>input</i>
$x[] . P$	<i>halt</i>
$x().P$	<i>wait</i>
$x \triangleleft \text{inl}.P$	<i>select left choice</i>
$x \triangleleft \text{inr}.P$	<i>select right choice</i>
$x \triangleright \{\text{inl} : P; \text{inr} : Q\}$	<i>offer binary choice</i>
$x \triangleright \{\}$	<i>offer nullary choice</i>

The variables x, y, z and w range over channel names. Occasionally, we use a, b , and c to range over *free* channel names. The construct $x \leftrightarrow y$ links two channels [15,5], forwarding messages received on x to y and vice versa. The construct $(\nu x)P$ creates a new channel x , and the construct $P \mid Q$ composes two processes. In $x(y).P$ and $x[y].P$, round brackets denote input, square brackets denote output. We use bound output [15], meaning that both input and output bind a new name.

Terms in HCP^- are identified up to structural congruence.

Definition 2 (Structural congruence). *The structural congruence $\equiv$ is the congruence closure over terms which satisfies the following additional axioms:*

$(\leftrightarrow\text{-comm})$	$x \leftrightarrow y \equiv y \leftrightarrow x$	$(\nu\text{-comm})$	$(\nu x)(\nu y)P \equiv (\nu y)(\nu x)P$
$(\mid\text{-comm})$	$P \mid Q \equiv Q \mid P$	$(\mid\text{-assoc})$	$P \mid (Q \mid R) \equiv (P \mid Q) \mid R$
(halt)	$P \mid 0 \equiv P$	(scope-ext)	$(\nu x)(P \mid Q) \equiv P \mid (\nu x)Q$ if $x \notin P$

Channels in HCP^- are typed using a session type system which is a conservative extension of linear logic.

Definition 3 (Types).

$A, B, C ::= A \otimes B$	<i>independent channels</i>	$\mathbf{1}$	<i>unit for $\otimes$</i>
$A \wp B$	<i>interdependent channels</i>	$\perp$	<i>unit for $\wp$</i>
$A \oplus B$	<i>internal choice</i>	$\mathbf{0}$	<i>unit for $\oplus$</i>
$A \& B$	<i>external choice</i>	$\top$	<i>unit for $\&$</i>

Duality plays a crucial role in both linear logic and session types. In HCP^- , the two endpoints of a channel are assigned dual types. This ensures that, for instance, whenever a process *sends* across a channel, the process on the other end of that channel is waiting to *receive*. Each type A has a dual, written $A^\perp$. Duality $(\cdot^\perp)$ is an involutive function on types.

Definition 4 (Duality).

$(A \otimes B)^\perp = A^\perp \wp B^\perp$	$\mathbf{1}^\perp = \perp$	$(A \wp B)^\perp = A^\perp \otimes B^\perp$	$\perp^\perp = \mathbf{1}$
$(A \oplus B)^\perp = A^\perp \& B^\perp$	$\mathbf{0}^\perp = \top$	$(A \& B)^\perp = A^\perp \oplus B^\perp$	$\top^\perp = \mathbf{0}$

Environments associate channels with types. Names in environments must be unique, and environments Γ and Δ can only be combined (Γ, Δ) if $\text{cn}(\Gamma) \cap \text{cn}(\Delta) = \emptyset$, where $\text{cn}(\Gamma)$ denotes the set of channel names in Γ .

Definition 5 (Environments). $\Gamma, \Delta, \Theta ::= x_1 : A_1 \dots x_n : A_n$

HCP⁻ registers parallelism using hyper-environments. A hyper-environment is a multiset of environments. While names within environments must be unique, names may be shared between multiple environments in a hyper-environment. We write $\mathcal{G} \mid \mathcal{H}$ to combine two hyper-environments.

Definition 6 (Hyper-environments). $\mathcal{G}, \mathcal{H} ::= \emptyset \mid \mathcal{G} \mid \Gamma$

Typing judgements associate processes with collections of typed channels.

Definition 7 (Typing judgements). A typing judgement $P \vdash \Gamma_1 \mid \dots \mid \Gamma_n$ denotes that the process P consists of n independent, but potentially entangled processes, each of which communicates according to its own protocol Γ_i . Typing judgements can be constructed using the inference rules below.

Structural rules

$$\frac{}{x \leftrightarrow y \vdash x : A, y : A^\perp} \text{Ax} \quad \frac{P \vdash \mathcal{G} \mid \Gamma, x : A \mid \Delta, x : A^\perp}{(\nu x)P \vdash \mathcal{G} \mid \Gamma, \Delta} \text{CUT}$$

$$\frac{P \vdash \mathcal{G} \quad Q \vdash \mathcal{H}}{P \mid Q \vdash \mathcal{G} \mid \mathcal{H}} \text{H-MIX} \quad \frac{}{0 \vdash \emptyset} \text{H-MIX}_0$$

Logical rules

$$\frac{P \vdash \mathcal{G} \mid \Gamma, y : A \mid \Delta, x : B}{x[y].P \vdash \mathcal{G} \mid \Gamma, \Delta, x : A \otimes B} \otimes \quad \frac{P \vdash \mathcal{G} \mid \Gamma, y : A, x : B}{x(y).P \vdash \mathcal{G} \mid \Gamma, x : A \wp B} (\wp)$$

$$\frac{P \vdash \mathcal{G}}{x[].P \vdash \mathcal{G} \mid x : \mathbf{1}} \mathbf{1} \quad \frac{P \vdash \mathcal{G} \mid \Gamma}{x().P \vdash \mathcal{G} \mid \Gamma, x : \perp} (\perp)$$

$$\frac{P \vdash \mathcal{G} \mid \Gamma, x : A}{x \triangleleft \text{inl}.P \vdash \mathcal{G} \mid \Gamma, x : A \oplus B} (\oplus_1) \quad \frac{P \vdash \mathcal{G} \mid \Gamma, x : B}{x \triangleleft \text{inr}.P \vdash \mathcal{G} \mid \Gamma, x : A \oplus B} (\oplus_2)$$

$$\frac{P \vdash \Gamma, x : A \quad Q \vdash \Gamma, x : B}{x \triangleright \{\text{inl} : P; \text{inr} : Q\} \vdash \Gamma, x : A \& B} (\&)$$

$$(no \text{ rule for } 0) \quad \frac{}{x \triangleright \{\} \vdash \Gamma, x : \top} (\top)$$

Furthermore, each logical rule has the side condition that $x \notin \mathcal{G}$.

Reductions relate processes with their reduced forms.

Definition 8 (Reduction). Reductions are described by the smallest relation $\Rightarrow$ on process terms closed under the rules below:

$$\begin{array}{lll} (\leftrightarrow) & (\nu x)(w \leftrightarrow x \mid P) & \Rightarrow P\{w/x\} \\ (\beta \otimes \wp) & (\nu x)(x[y].P \mid x(y).R) & \Rightarrow (\nu x)(\nu y)(P \mid R) \\ (\beta \mathbf{1} \perp) & (\nu x)(x[].P \mid x().Q) & \Rightarrow P \mid Q \\ (\beta \oplus \&_1) & (\nu x)(x \triangleleft \text{inl}.P \mid x \triangleright \{\text{inl} : Q; \text{inr} : R\}) & \Rightarrow (\nu x)(P \mid Q) \\ (\beta \oplus \&_2) & (\nu x)(x \triangleleft \text{inr}.P \mid x \triangleright \{\text{inl} : Q; \text{inr} : R\}) & \Rightarrow (\nu x)(P \mid R) \end{array}$$

$$\begin{array}{c}
\frac{P \Rightarrow P'}{(\nu x)P \Rightarrow (\nu x)P'} (\gamma\nu) \quad \frac{P \Rightarrow P'}{P \mid Q \Rightarrow P' \mid Q} (\gamma\mid) \\
\\
\frac{P \equiv Q \quad Q \Rightarrow Q' \quad Q' \equiv P'}{P \Rightarrow P'} (\gamma\equiv)
\end{array}$$

3.1 Example

HCP⁻ uses hyper-sequents to structure communication, and it is this structure which rules out deadlocked interactions. Let us go back to our example of a deadlocked interaction from section 1. If we want to type this interaction in HCP⁻, we run into a problem: to communicate on x , we need to add name restrictions on x and y , *e.g.*,

$$(\nu x)(\nu y)(x(z).y[\text{🍰}].\text{👤} \mid y(w).x[\text{📺}].\text{👤}).$$

However, there is no typing derivation for this term. We illustrate this with the partial typing derivation below. In this derivation, there is no way to proceed and type the final name restriction. The CUT rule needs a hypersequent separator to eliminate, so that it only ever links up two independent processes, but the bottom-most sequent has none. Furthermore, the two occurrences of x appearing in the same environment make it ill-formed.

$$\begin{array}{c}
\frac{\frac{\text{👤} \vdash \Gamma, z:\$, \text{🍰}^\perp}{y[\text{🍰}].\text{👤} \vdash \Gamma, z:\$, \text{🍰}^\perp, y:\text{🍰}} \otimes \quad \frac{\text{👤} \vdash \Delta, \text{📺}^\perp, \$, w:\text{🍰}^\perp}{x[\text{📺}].\text{👤} \vdash \Delta, x:\$, w:\text{🍰}^\perp} \otimes}{\frac{x(z).y[\text{🍰}].\text{👤} \vdash \Gamma, x:\$, \text{🍰}^\perp, y:\text{🍰} \quad y(w).x[\text{📺}].\text{👤} \vdash \Delta, x:\$, y:\text{🍰}^\perp}{(x(z).y[\text{🍰}].\text{👤} \mid y(w).x[\text{📺}].\text{👤}) \vdash \Gamma, x:\$, y:\text{🍰}^\perp \mid \Delta, x:\$, y:\text{🍰}^\perp} \text{H-MIX}} \\
\text{CUT} \\
(\nu y)(x(z).y[\text{🍰}].\text{👤} \mid y(w).x[\text{📺}].\text{👤}) \vdash \Gamma, \Delta, x:\$, \text{🍰}^\perp, x:\$
\end{array}$$

3.2 Metatheory

HCP⁻ enjoys subject reduction, termination, and progress [11].

Lemma 9 (Preservation for $\equiv$). *If $P \equiv Q$, then $P \vdash \mathcal{G}$ iff $Q \vdash \mathcal{G}$.*

Proof. By induction on the derivation of $P \equiv Q$.

Theorem 10 (Preservation). *If $P \vdash \mathcal{G}$ and $P \Rightarrow Q$, then $Q \vdash \mathcal{G}$.*

Proof. By induction on the derivation of $P \Rightarrow Q$.

Definition 11 (Actions). *A process P acts on x whenever x is free in the outermost term constructor of P , e.g., $x[y].P$ acts on x but not on y , and $x \leftrightarrow y$ acts on both x and y . A process P is an action if it acts on some channel x .*

Definition 12 (Canonical forms). A process P is in canonical form if

$$P \equiv (\nu x_1) \dots (\nu x_n)(P_1 \mid \dots \mid P_{n+m+1}),$$

such that: no process P_i is a cut or a mix; no process P_i is a link acting on a bound channel x_i ; and no two processes P_i and P_j are acting on the same bound channel x_i .

Lemma 13. If a well-typed process P is in canonical form, then it is blocked on an external communication, i.e., $P \equiv (\nu x_1) \dots (\nu x_n)(P_1 \mid \dots \mid P_{n+m+1})$ such that at least one process P_i acts on a free name.

Proof. We have $P \equiv (\nu x_1) \dots (\nu x_n)(P_1 \mid \dots \mid P_{n+m+1})$, such that no P_i is a cut or a link acting on a bound channel, and no two processes P_i and P_j are acting on the same bound channel. The prefix of cuts and mixes introduces n channels. Each application of cut requires an application of mix, so the prefix introduces $n + m + 1$ processes. Therefore, at least $m + 1$ of the processes P_i must be acting on a free channel, i.e., blocked on an external communication.

Theorem 14 (Progress). If $P \vdash \Gamma$, then either P is in canonical form, or there exists a process Q such that $P \Rightarrow Q$.

Proof. We consider the maximum prefix of cuts and mixes of P such that $P \equiv (\nu x_1) \dots (\nu x_n)(P_1 \mid \dots \mid P_{n+m+1})$, and no P_i is a cut. If any process P_i is a link, we reduce by $(\leftrightarrow)$. If any two processes P_i and P_j are acting on the same channel x_i , we rewrite by $\equiv$ and reduce by the appropriate β -rule. Otherwise, P is in canonical form.

Theorem 15 (Termination). If $P \vdash \mathcal{G}$, then there are no infinite $\Rightarrow$ -reduction sequences.

Proof. Every reduction reduces a single cut to zero, one or two cuts. However, each of these cuts is smaller, measured in the size of the cut formula. Furthermore, each instance of the structural congruence preserves the size of the cut. Therefore, there cannot be an infinite $\Rightarrow$ -reduction sequence.

4 Shared Channels and Non-determinism

In this section, we will discuss our main contribution: an extension of HCP^- which allows for races while still excluding deadlocks. We have seen in section 3.1 how HCP^- excludes deadlocks, but how exactly does HCP^- exclude races? Let us return to our example from section 1, to the interaction between Ami, Boé and the store.

$$\begin{array}{c}
 (x[\text{🍰}].x[\text{👉}].\text{👤} \mid x(y).\text{👤} \mid x(z).\text{👤}) \\
 \Downarrow^* \\
 (\text{👤} \mid \text{👤}\{\text{🍰}/y\} \mid \text{👤}\{\text{👉}/z\}) \quad \text{or} \quad (\text{👤} \mid \text{👤}\{\text{👉}/y\} \mid \text{👤}\{\text{🍰}/z\})
 \end{array}$$

Races occur when more than two processes attempt to communicate simultaneously over the *same* channel. However, the CUT rule of HCP^- requires that *exactly two* processes communicate over each channel:

$$\frac{P \vdash \mathcal{G} \mid \Gamma, x:A \mid \Delta, x:A^\perp}{(\nu x)P \vdash \mathcal{G} \mid \Gamma, \Delta} \text{CUT}$$

We could attempt to write down a protocol for our example, stating that the store has a pair of channels $x, y : \text{🍰}$ with which it communicates with Ami and Bo, taking 🍰 to be the type of interactions in which cake *may* be obtained, i.e. of both 🍰 and 👋 , and state that the store communicates with Ami *and* Bo over a channel of type $\text{🍰} \wp \text{🍰}$. However, this *only* models interactions such as the following:

$$\frac{\frac{\frac{\text{👤} \vdash \Gamma, x:\text{🍰}^\perp \quad \text{👩} \vdash \Delta, y:\text{🍰}^\perp}{(\text{👤} \mid \text{👩}) \vdash \Gamma, x:\text{🍰}^\perp \mid \Delta, y:\text{🍰}^\perp} \text{H-MIX} \quad \frac{\frac{\text{👤} \vdash \Theta, x:\text{🍰}, y:\text{🍰}}{y(x).\text{👤} \vdash \Theta, y:\text{🍰} \wp \text{🍰}} \wp}{\frac{y[x].(\text{👤} \mid \text{👩}) \vdash \Gamma, \Delta, y:\text{🍰}^\perp \wp \text{🍰}^\perp \quad y(x).\text{👤} \vdash \Theta, y:\text{🍰} \wp \text{🍰}}{\frac{(y[x].(\text{👤} \mid \text{👩}) \mid y(x).\text{👤}) \vdash \Gamma, \Delta, y:\text{🍰}^\perp \wp \text{🍰}^\perp \mid \Theta, y:\text{🍰} \wp \text{🍰}}{(\nu y)(y[x].(\text{👤} \mid \text{👩}) \mid y(x).\text{👤}) \vdash \Gamma, \Delta, \Theta} \text{CUT}} \otimes \text{H-MIX}$$

In this interaction, Ami will get whatever the store decides to send on x , and Bo will get whatever the store decides to send on y . This means that this interaction gives the choice of who receives what *to the store*. This is not an accurate model of our original example, where the choice of who receives the cake is non-deterministic and depends on factors outside of any of the participants' control!

Modelling racy behaviour, such as that in our example, is essential to describing the interactions that take place in realistic concurrent systems. We would like to extend HCP^- to allow such races in a way which mirrors the way in which the π -calculus handles non-determinism. Let us return to our example:

$$(x[\text{🍰}].x[\text{👋}].\text{👤} \mid x(y).\text{👤} \mid x(z).\text{👩})$$

In this interaction, we see that the channel x is only used as a way to connect the various clients, Ami and Bo, to the store. The *real* communication, sending the slice of cake and disappointment, takes place on the channels 🍰 , 👋 , y and z . Inspired by this, we add two new constructs to the term language of HCP^- for sending and receiving on a *shared* channel. These actions are marked with a $\star$ to distinguish them from ordinary sending and receiving.

Definition 16 (Terms). We extend definition 1 as follows:

$$\begin{aligned} P, Q, R ::= & \dots \\ & \mid \star x[y].P \quad \text{client creation} \\ & \mid \star x(y).P \quad \text{server interaction} \end{aligned}$$

As before, round brackets denote input, square brackets denote output. Note that $\star x[y].P$, much like $x[y].P$, is a bound output: both client creation and server interaction bind a new name. The structural congruence, which identifies certain terms, is the same as definition 2.

In any non-deadlock interaction between a server and some clients, there must be *exactly* as many clients as there are server interactions. Therefore, we add two new *dual* types for client pools and servers, which track how many clients or server interactions they represent.

Definition 17 (Types). *We extend definition 3 as follows:*

$$\begin{array}{lcl} A, B, C & ::= & \dots \\ & | & !_n A \quad \text{pool of } n \text{ clients} \\ & | & ?_n A \quad n \text{ server interactions} \end{array}$$

The types $!_n A$ and $?_n A^\perp$ are dual. Duality remains an involutive function.

We have to add typing rules to associate our new client and server interactions with their types. The definition for environments will remain unchanged, but we will extend the definition for the typing judgement. To determine the new typing rules, we essentially answer the question “What typing constructs do we need to complete the following proof?”

$$\begin{array}{c} \text{👤} \vdash \Gamma, y : \text{🍰}^\perp \quad \text{👤} \vdash \Delta, y' : \text{🍰}^\perp \quad \text{🏪} \vdash \Theta, z : \text{🍰}, z' : \text{🍰} \\ \vdots \quad \quad \quad \vdots \quad \quad \quad \vdots \\ (\nu x)((\star x[y].\text{👤} \mid \star x[y'].\text{👤}) \mid \star x(z).\star x(z').\text{🏪}) \vdash \Gamma, \Delta, \Theta \end{array}$$

The constructs $\star x[y].P$ and $\star x(y).P$ introduce a single client or server action, respectively—hence, channels of type $!_1$ and $?_1$. However, when we cut, we want to cut on both interactions simultaneously. We need rules for the *contraction* of shared channel names.

4.1 Clients and Pooling

A client pool represents a number of independent processes, each wanting to interact with the same server. Examples of such a pool include Ami and Boé from our example, customers for online stores in general, and any number of processes which interact with a single, centralised server.

We introduce two new rules: one to construct clients, and one to pool them together. The first rule, $(!_1)$, interacts over a channel as a client. It does this by receiving a channel y over a *shared* channel x . The channel y is the channel across which the actual interaction will eventually take place. The second rule, $\text{CONT}_!$, allows us to contract shared channel names with the same type. When used together with H-MIX, this allows us to pool clients together.

$$\frac{P \vdash \mathcal{G} \mid \Gamma, y : A}{\star x[y].P \vdash \mathcal{G} \mid \Gamma, x : !_1 A} (!_1) \quad \frac{P \vdash \mathcal{G} \mid \Gamma, x : !_m A \mid \Delta, y : !_n A}{P\{x/y\} \vdash \mathcal{G} \mid \Gamma, \Delta, x : !_{m+n} A} \text{CONT}_!$$

Using these rules, we can derive the left-hand side of our proof by marking Ami and Boé as clients, and pooling them together.

$$\begin{array}{c}
\frac{\frac{\text{Ami} \vdash \Gamma, y : \text{cake}^\perp}{\star x[y].\text{Ami} \vdash \Gamma, y : !_1 \text{cake}^\perp} (!_1) \quad \frac{\frac{\text{Boé} \vdash \Delta, y' : \text{cake}^\perp}{\star x'[y'].\text{Boé} \vdash \Delta, x' : !_1 \text{cake}^\perp} (!_1)}{(\star x[y].\text{Ami} \mid \star x'[y'].\text{Boé}) \vdash \Gamma, x : !_1 \text{cake}^\perp \mid \Delta, x' : !_1 \text{cake}^\perp} \text{H-Mix} \\
\frac{(\star x[y].\text{Ami} \mid \star x'[y'].\text{Boé}) \vdash \Gamma, x : !_1 \text{cake}^\perp \mid \Delta, x' : !_1 \text{cake}^\perp}{(\star x[y].\text{Ami} \mid \star x'[y'].\text{Boé}) \vdash \Gamma, \Delta, x : !_2 \text{cake}^\perp} \text{CONT}_!
\end{array}$$

4.2 Servers and Sequencing

Dual to a pool of n clients in parallel is a server with n actions in sequence. Our interpretation of a server is a process which offers some number of interdependent interactions of the same type. Examples include the store from our example, which gives out slices of cake and disappointment, online stores in general, and any central server which interacts with some number of client processes.

We introduce two new rules to construct servers. The first rule, $(?_1)$, marks a interaction over some channel as a server interaction. It does this by sending a channel y over a *shared* channel x . The channel y is the channel across which the actual interaction will take place. The second rule, $\text{CONT}_?$, allows us to merge two (possibly interleaved) sequences of server interactions. This allows us to construct a server which has multiple interactions of the same type, across the same shared channel.

$$\frac{P \vdash \mathcal{G} \mid \Gamma, y : A}{\star x(y).P \vdash \mathcal{G} \mid \Gamma, x : ?_1 A} (?_1) \quad \frac{P \vdash \mathcal{G} \mid \Gamma, x : ?_m A, y : ?_n A}{P\{x/y\} \vdash \mathcal{G} \mid \Gamma, x : ?_{m+n} A} \text{CONT}_?$$

Using these rules, we can derive the right-hand side of our proof, by marking each of the store's interactions as server interactions, and then contracting them.

$$\begin{array}{c}
\frac{\frac{\text{Store} \vdash \Theta, z : \text{cake}, z' : \text{cake}}{\star x'(z').\text{Store} \vdash \Theta, z : \text{cake}, x' : ?_1 \text{cake}} (?_1)}{\frac{\star x(z).\star x'(z').\text{Store} \vdash \Theta, x : ?_1 \text{cake}, x' : ?_1 \text{cake}}{\star x(z).\star x'(z').\text{Store} \vdash \Theta, x : ?_2 \text{cake}} (?_1)} \text{CONT}_?
\end{array}$$

Thus, we complete the typing derivation of our example.

Definition 18 (Typing judgements). We extend definition 7 as follows:

$$\begin{array}{c}
\frac{P \vdash \mathcal{G} \mid \Gamma, y : A}{\star x[y].P \vdash \mathcal{G} \mid \Gamma, x : !_1 A} (!_1) \quad \frac{P \vdash \mathcal{G} \mid \Gamma, y : A}{\star x(y).P \vdash \mathcal{G} \mid \Gamma, x : ?_1 A} (?_1) \\
\frac{P \vdash \mathcal{G} \mid \Gamma, x : !_m A \mid \Delta, y : !_n A}{P\{x/y\} \vdash \mathcal{G} \mid \Gamma, \Delta, x : !_{m+n} A} \text{CONT}_! \quad \frac{P \vdash \mathcal{G} \mid \Gamma, x : ?_m A, y : ?_n A}{P\{x/y\} \vdash \mathcal{G} \mid \Gamma, x : ?_{m+n} A} \text{CONT}_?
\end{array}$$

4.3 Running Clients and Servers

Finally, we need to extend the reduction rules to allow for the reduction of client and server processes. The reduction rule we add is a variant of the reduction rule for $\otimes$ and $\wp$, $(\beta \otimes \wp)$.

Definition 19 (Reduction). *We extend definition 8 as follows:*

$$(\beta \star) \quad (\nu x)((\star x[y].P \mid \star x(z).Q) \mid R) \Longrightarrow (\nu x)((\nu y)(P \mid Q\{y/z\}) \mid R)$$

The difference between $(\beta \star)$ and $(\beta \otimes \wp)$ is that the former allows reduction to happen in the presence of an unrelated process R , which is passed along unchanged. This is necessary, as there may be other clients waiting to interact with the server on the shared channel x , which cannot be moved out of scope of the name restriction (νx) . When there is no unrelated process R , *i.e.*, when there is only a single client, we can rewrite by (halt) before and after applying $(\beta \star)$.

So where does the non-determinism in HCP_{ND}^- come from? Let us say we have a term of the following form:

$$(\nu x)((\star x[y_1].P_1 \mid \cdots \mid \star x[y_n].P_n) \mid \star x(y_1) \dots \star x(y_n).Q)$$

As parallel composition is commutative and associative, we can rewrite this term to pair any client in the pool with the server before applying $(\beta \star)$. Thus, like in the π -calculus, the non-determinism is introduced by the structural congruence.

Does this mean that, for an arbitrary client pool P in $(\nu x)(P \mid \star x(z).Q)$, every client in that pool is competing for the server interaction on x ? Not necessarily, as some portion of the clients can be blocked on an external communication. For instance, in the term below, clients $\star x[y_{n+1}].P_{n+1} \dots \star x[y_m].P_m$ are blocked on a communication on the external channel a :

$$\begin{aligned} & \nu x.((\star x[y_1].P_1 \mid \cdots \mid \star x[y_n].P_n) \\ & \quad \mid a().(\star x[y_{n+1}].P_{n+1} \mid \cdots \mid \star x[y_m].P_m)) \\ & \quad \mid \star x(y_1) \dots \star x(y_m).Q) \end{aligned}$$

If we reduce this term, then only the clients $\star x[y_1].P_1 \dots \star x[y_n].P_n$ will be assigned server interactions, and we end up with the following canonical form:

$$\begin{aligned} & \nu x.(a().(\star x[y_{n+1}].P_{n+1} \mid \cdots \mid \star x[y_m].P_m) \\ & \quad \mid \star x(y_{n+1}) \dots \star x(y_m).Q) \end{aligned}$$

This matches our intuition and the behaviour of the π -calculus.

Alternative syntax. If we choose to reuse the terms $x[y].P$ and $x(y).P$ for shared channels, we could replace $(\beta \otimes \wp)$ with $(\beta \star)$, using the latter rule for both cases.

4.4 Metatheory

HCP_{ND}^- enjoys subject reduction, termination, and progress.

Lemma 20 (Preservation for $\equiv$). *If $P \equiv Q$ and $P \vdash \mathcal{G}$, then $Q \vdash \mathcal{G}$.*

Proof. By induction on the derivation of $P \equiv Q$.

Theorem 21 (Preservation). *If $P \vdash \mathcal{G}$ and $P \Rightarrow Q$, then $Q \vdash \mathcal{G}$.*

Proof. By induction on the derivation of $P \Rightarrow Q$.

Definition 22 (Actions). *A process P acts on x whenever x is free in the outermost term constructor of P , e.g., $\star x(y).P$ acts on x but not on y , and $x \leftrightarrow y$ acts on both x and y . A process P is an action if it acts on some channel x . Two actions are dual when they introduce dual type constructors, e.g., $x[y].P$ is dual to $x(z).Q$, but $x \leftrightarrow y$ is not dual to any action.*

Definition 23 (Canonical forms). *A process P is in canonical form if*

$$P \equiv (\nu x_1) \dots (\nu x_n)(P_1 \mid \dots \mid P_{n+m+1}),$$

such that: no process P_i is a cut or a mix; no process P_i is a link acting on a bound channel x_i ; and no two processes P_i and P_j are acting on the same bound channel x_i with dual actions.

The new definition of canonical forms is slightly more precise than definition 12: we added the phrase “with dual actions”. With the addition of shared channels, it has become possible to have a process which cannot reduce, but in which two processes are waiting to act on the same channel, e.g., in $(\star x[y].\text{🐼} \mid \star x[y'].\text{🐼})$.

Lemma 24. *If a well-typed process P is in canonical form, then it is blocked on an external communication, i.e., $P \equiv (\nu x_1) \dots (\nu x_n)(P_1 \mid \dots \mid P_{n+m+1})$ such that at least one process P_i acts on a free name.*

Proof. We have $P \equiv (\nu x_1) \dots (\nu x_n)(P_1 \mid \dots \mid P_{n+m+1})$, such that no P_i is a cut or a link acting on a bound channel, and no two processes P_i and P_j are acting on the same bound channel with dual actions. The prefix of cuts and mixes introduces n channels. Each application of cut requires an application of mix, so the prefix introduces $n + m + 1$ processes. Each application of CONT! requires an application of mix, so there are at most m clients acting on the same bound channel. Therefore, at least one of the processes P_i must be acting on a free channel, i.e., blocked on an external communication.

Theorem 25 (Progress). *If $P \vdash \Gamma$, then either P is in canonical form, or there exists a process Q such that $P \Rightarrow Q$.*

Proof. We consider the maximum prefix of cuts and mixes of P such that $P \equiv (\nu x_1) \dots (\nu x_n)(P_1 \mid \dots \mid P_{n+m+1})$, and no P_i is a cut. If any process P_i is a link, we reduce by $(\leftrightarrow)$. If any two processes P_i and P_j are acting on the same channel x_i with dual actions, we rewrite by $\equiv$ and reduce by the appropriate β -rule. Otherwise, P is in canonical form.

Theorem 26 (Termination). *If $P \vdash \mathcal{G}$, then there are no infinite $\Rightarrow$ -reduction sequences.*

Proof. Every reduction reduces a single cut to zero, one or two cuts. However, each of these cuts is smaller, measured in the size of the cut formula. Furthermore, each instance of the structural congruence preserves the size of the cut. Therefore, there cannot be an infinite $\Rightarrow$ -reduction sequence.

4.5 HCP_{ND}^- and Non-deterministic Local Choice

In section 2, we discussed the non-deterministic local choice operator, which is used in several extensions of πDILL and CP [1,6,7]. This operator is admissible in HCP_{ND}^- . We can derive the non-deterministic choice $P + Q$ by constructing the following term:

$$\begin{aligned} & \nu x. ((\star x[y].y \triangleleft \text{inl}.y[] .0 \\ & \quad | \star x[z].z \triangleleft \text{inr}.z[] .0) \\ & \quad | \star x(y). \star x(z).y \triangleright \\ & \quad \{ \text{inl} : (\nu w)(z \triangleright \{ \text{inl} : z().w[] .0; \text{inr} : z().w[] .0 \} \mid w().P) \\ & \quad ; \text{inr} : (\nu w)(z \triangleright \{ \text{inl} : z().w[] .0; \text{inr} : z().w[] .0 \} \mid w().Q) \}) \end{aligned}$$

This term is a cut between two processes.

- On the left-hand side, we have a pool of two processes, $\star x[y].y \triangleleft \text{inl}.y[] .0$ and $\star x[z].z \triangleleft \text{inr}.z[] .0$. Each makes a choice: the first sends **inl**, and the second sends **inr**.
- On the right-hand side, we have a server with both P and Q . This server has two channels on which a choice is offered, y and z . The choice on y selects between P and Q . The choice on z does not affect the outcome of the process at all. Instead, it is discarded.

When these clients and the server are put together, the choices offered by the server will be non-deterministically lined up with the clients which make choices, and either P or Q will run.

While there is a certain amount of overhead involved in this encoding, it scales linearly in terms of the number of processes. The reverse—encoding the non-determinism present in HCP_{ND}^- using non-deterministic local choice—scales exponentially, see, *e.g.*, the examples in section 2.

5 Cuts with Leftovers

So far, our account of a non-determinism in client/server interactions only allows for interactions between equal numbers of clients and server interactions. A natural question is whether or not we can deal with the scenario in which there are more client than server interactions or vice versa, *i.e.*, whether or not the following rules are derivable:

$$\frac{\vdash \Gamma, !_{n+m}A \quad \vdash \Delta, ?_n A^\perp}{\vdash \Gamma, \Delta, !_m A} \quad \frac{\vdash \Gamma, !_n A \quad \vdash \Delta, ?_{n+m} A^\perp}{\vdash \Gamma, \Delta, ?_m A^\perp}$$

These rules are derivable using a link. For instance, we can derive the rule for the case in which there are more clients than servers as follows:

$$\frac{\begin{array}{c} \frac{Q \vdash \Delta, x : ?_n A^\perp \quad x \leftrightarrow w \vdash x' : ?_m A^\perp, w : !_m A}{(Q \mid x \leftrightarrow w) \vdash \Delta, x : ?_n A^\perp \mid x' : ?_m A^\perp, w : !_m A} \text{H-MIX} \\ \frac{P \vdash \Gamma, x : !_{n+m} A \quad (Q \mid x \leftrightarrow w) \vdash \Delta, x : ?_{n+m} A^\perp, w : !_m A}{(P \mid (Q \mid x \leftrightarrow w)) \vdash \Gamma, x : !_{n+m} A \mid \Delta, x : ?_{n+m} A^\perp, w : !_m A} \text{CONT!} \end{array}}{\frac{(P \mid (Q \mid x \leftrightarrow w)) \vdash \Gamma, x : !_{n+m} A \mid \Delta, x : ?_{n+m} A^\perp, w : !_m A}{(\nu x)(P \mid (Q \mid x \leftrightarrow w)) \vdash \Gamma, \Delta, w : !_m A} \text{H-MIX}} \text{CUT}$$

6 Relation to Manifest Sharing

In section 2, we mentioned related work which extends π DILL and CP with non-deterministic local choice [1,6,7], and contrasted these approaches with ours. In this section, we will contrast our work with the more recent work on manifest sharing [2].

Manifest sharing extends the session-typed language SILL with two connectives, $\uparrow_L^S A$ and $\downarrow_L^S A$, which represent the places in a protocol where a shared resource is acquired and released, respectively. In the resulting language, SILL_S, we can define a type for, *e.g.*, shared queues (using the notation for types introduced in this paper):

$$\text{queue } A ::= \uparrow_L^S (A^\perp \wp \downarrow_L^S (\text{queue } A)) \ \& \ (A \oplus \perp) \wp \downarrow_L^S (\text{queue } A)$$

The type $\text{queue } A$ types a shared channel which, after we acquire exclusive access, gives us the choice between enqueueing a value ($A^\perp$) and releasing the queue, or dequeuing a value if there is any ($A \oplus \perp$) and releasing the queue.

The language SILL_S is much more expressive than HCP_{ND}^- , as it has support for both shared channels and recursion. In fact, Balzer, Pfenning, and Toninho [3] show that SILL_S is expressive enough to embed the untyped asynchronous π -calculus. This expressiveness comes with a cost, as SILL_S processes are not guaranteed to be deadlock free, though recent work addresses this issue [4].

Despite the difference in expressiveness, there are some similarities between HCP_{ND}^- and SILL_S. In the former, shared channels represent (length-indexed) streams of interactions of the same type. In the latter, it is necessary for type preservation that shared channels are always released at the same type at which they were acquired, meaning that shared channels also represent (possibly infinite) streams of interactions of the same type. In fact, in HCP_{ND}^- , the type for queues (with n interactions) can be written as $!_n(A^\perp \ \& \ (A \oplus \perp))$.

One key difference between HCP_{ND}^- and SILL_S is that in SILL_S a server must finish interacting with one client before interacting with another, whereas in HCP_{ND}^- the server may interact with multiple clients simultaneously.

7 Discussion and Future Work

We presented HCP_{ND}^- , an extension of HCP^- which permits non-deterministic communication without losing the strong connection to logic. We gave proofs for preservation, progress, and termination for the term reduction system of HCP_{ND}^- . We showed that we can define non-deterministic local choice in HCP_{ND}^- .

Our formalism so far has only captured servers that provide for a fixed number of clients. More realistically, we would want to define servers that provide for arbitrary numbers of clients. This poses two problems: how would we define arbitrarily-interacting stateful processes, and how would we extend the typing discipline of HCP_{ND}^- to account for them without losing its static guarantees.

One approach to defining server processes would be to combine HCP_{ND}^- with structural recursion and corecursion, following the μCP extension of Lindley and Morris [13]. Their approach can express processes which produce streams of A channels. Such a process would expose a channel with the co-recursive type $\nu X. A \wp (1 \oplus X)$. Given such a process, it is possible to produce a channel of type $A \wp A \wp \dots \wp A$ for any number of A s, allowing us to satisfy the type $?_n A$ for an arbitrary n .

We would also need to extend the typing discipline to capture arbitrary use of shared channels. One approach would be to introduce resource variables and quantification. Following this approach, in addition to having types $?_n A$ and $!_n A$ for concrete n , we would also have types $?_x A$ and $!_x A$ for resource variables x . These variables would be introduced by quantifiers $\forall x A$ and $\exists x A$. Defining terms corresponding to $\forall x A$, and its relationship with structured recursion, presents an interesting area of further work.

Our account of HCP^- did not include the exponentials $?A$ and $!A$. The type $!A$ denotes arbitrarily many independent instances of A , while the type $?A$ denotes a concrete (if unspecified) number of potentially-dependent instances of A . Existing interpretations of linear logic as session types have taken $!A$ to denote A -servers, while $?A$ denotes A -clients. However, the analogy is imperfect: while we expect servers to provide arbitrarily many instances of their behaviour, we also expect those instances to be interdependent.

With quantification over resource variables, we can give precise accounts of both CP's exponentials and idealised servers and clients. CP exponentials could be embedded into this framework using the definitions $!A ::= \forall n !_n A$ and $?A ::= \exists n ?_n A$. We would also have types that precisely matched our intuitions for server and client behavior: an A server is of type $\forall n ?_n A$, as it serves an unbounded number of requests with the requests being interdependent, while a collection of A clients is of type $\exists n !_n A$, as we have a specific number of clients with each client being independent.

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: Essays Dedicated to Philip Wadler on the Occasion of His 60th Birthday. Lecture Notes in Computer Science (2016). https://doi.org/10.1007/978-3-319-30936-1_2
2. Balzer, S., Pfenning, F.: Manifest sharing with session types. Proceedings of the ACM on Programming Languages **1**(ICFP), 1–29 (Aug 2017). <https://doi.org/10.1145/3110281>
 3. Balzer, S., Pfenning, F., Toninho, B.: A Universal Session Type for Untyped Asynchronous Communication. In: Schewe, S., Zhang, L. (eds.) 29th International Conference on Concurrency Theory (CONCUR 2018). Leibniz International Proceedings in Informatics (LIPIcs), vol. 118, pp. 30:1–30:18. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany (2018). <https://doi.org/10.4230/LIPIcs.CONCUR.2018.30>
 4. Balzer, S., Toninho, B., Pfenning, F.: Manifest deadlock-freedom for shared session types. In: Programming Languages and Systems, pp. 611–639. Springer International Publishing (2019). https://doi.org/10.1007/978-3-030-17184-1_22
 5. Boreale, M.: On the expressiveness of internal mobility in name-passing calculi. Theoretical Computer Science **195**(2), 205–226 (3 1998). [https://doi.org/10.1016/s0304-3975\(97\)00220-x](https://doi.org/10.1016/s0304-3975(97)00220-x)
 6. Caires, L.: Types and logic, concurrency and non-determinism. In: Abadi, M., Gardner, P., Gordon, A., Mardare, R. (eds.) Essays for the Luca Cardelli Fest. Microsoft Research (9 2014)
 7. Caires, L., Pérez, J.A.: Linearity, control effects, and behavioral types. In: Programming Languages and Systems – 26th European Symposium on Programming, ESOP 2017, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2017, Uppsala, Sweden, April 22–29, 2017, Proceedings. Springer (4 2017)
 8. Caires, L., Pfenning, F.: Session Types as Intuitionistic Linear Propositions, pp. 222–236. Springer Berlin Heidelberg, Berlin, Heidelberg (2010). https://doi.org/10.1007/978-3-642-15375-4_16
 9. Girard, J.Y., Scedrov, A., Scott, P.J.: Bounded linear logic: A modular approach to polynomial-time computability. Theor. Comput. Sci. **97**(1), 1–66 (Apr 1992). [https://doi.org/10.1016/0304-3975\(92\)90386-T](https://doi.org/10.1016/0304-3975(92)90386-T)
 10. Honda, K.: Types for dyadic interaction. In: CONCUR'93, pp. 509–523. Springer Nature (1993). https://doi.org/10.1007/3-540-57208-2_35
 11. Kokke, W., Montesi, F., Peressotti, M.: Taking linear logic apart. In: Workshop on Linearity & TLLA at FloC'18 (Jul 2018)
 12. Kokke, W., Montesi, F., Peressotti, M.: Better late than never: A fully-abstract semantics for classical processes. PACMPL **3**(POPL) (Jan 2019)
 13. Lindley, S., Morris, J.G.: Talking bananas: Structural recursion for session types. In: Proceedings of the 21st ACM SIGPLAN International Conference on Functional Programming, pp. 434–447. ICFP 2016, ACM, New York, NY, USA (2016). <https://doi.org/10.1145/2951913.2951921>
 14. Milner, R., Parrow, J., Walker, D.: A calculus of mobile processes, II. Information and Computation **100**(1), 41–77 (Sep 1992). [https://doi.org/10.1016/0890-5401\(92\)90009-5](https://doi.org/10.1016/0890-5401(92)90009-5)
 15. Sangiorgi, D.: π -calculus, internal mobility, and agent-passing calculi. Theoretical Computer Science **167**(1-2), 235–274 (1996). [https://doi.org/10.1016/0304-3975\(96\)00075-8](https://doi.org/10.1016/0304-3975(96)00075-8)
 16. Wadler, P.: Propositions as sessions. In: Proceedings of the 17th ACM SIGPLAN International Conference on Functional Programming, pp. 273–286. ICFP '12, ACM, New York, NY, USA (2012). <https://doi.org/10.1145/2364527.2364568>