



HAL
open science

π

π with Leftovers: A Mechanisation in Agda

Uma Zalakain, Ornela Dardha

► **To cite this version:**

Uma Zalakain, Ornela Dardha.

π

π with Leftovers: A Mechanisation in Agda. 41th International Conference on Formal Techniques for Distributed Objects, Components, and Systems (FORTE), Jun 2021, Valletta, Malta. pp.157-174, 10.1007/978-3-030-78089-0_9 . hal-03740255

HAL Id: hal-03740255

<https://inria.hal.science/hal-03740255>

Submitted on 29 Jul 2022

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



This document is the original author manuscript of a paper submitted to an IFIP conference proceedings or other IFIP publication by Springer Nature. As such, there may be some differences in the official published version of the paper. Such differences, if any, are usually due to reformatting during preparation for publication or minor corrections made by the author(s) during final proofreading of the publication manuscript.

π with leftovers: a mechanisation in Agda ^{*}

Uma Zlakain¹[0000-0002-3268-9338] and Ornela Dardha²[0000-0001-9927-7875]

¹ University of Glasgow, Scotland `u.zalakain.1@research.gla.ac.uk`

² University of Glasgow, Scotland `ornela.dardha@glasgow.ac.uk`

Abstract. Linear type systems need to keep track of how programs use their resources. The standard approach is to use *context splits* specifying how resources are (disjointly) split across subterms. In this approach, context splits redundantly echo information which is already present within subterms. An alternative approach is to use *leftover typing* [24,2], where in addition to the usual (input) usage context, typing judgments have also an output usage context: the leftovers. In this approach, the leftovers of one typing derivation are fed as input to the next, threading through linear resources while avoiding context splits. We use leftover typing to define a type system for a resource-aware π -calculus [28,27], a process algebra used to model concurrent systems. Our type system is parametrised over a set of *usage algebras* [21,35] that are general enough to encompass *shared types* (free to reuse and discard), *graded types* (use exactly n number of times) and *linear types* (use exactly once). Linear types are important in the π -calculus: they ensure privacy and safety of communication and avoid race conditions, while graded and shared types allow for more flexible programming. We provide a framing theorem for our type system, generalise the weakening and strengthening theorems to include linear types, and prove subject reduction. Our formalisation is fully mechanised in about 1850 lines of Agda [38].

Keywords: Pi-calculus · Linear types · Leftover typing · Concurrency · Mechanisation · Agda

1 Introduction

The π -calculus [28,27] is a computational model for communication and concurrency that boils concurrent processing down to the sending and receiving of data over communication channels. Notably, it features channel mobility: channels themselves are first class values and can be sent and received. Kobayashi et al. [23] introduced a typed version of the π -calculus with linear channel types, where channels must be used *exactly* once. Linearity in the π -calculus guarantees privacy and safety of communication and avoids race conditions.

^{*} This work is supported by the EU HORIZON 2020 MSCA RISE project 778233 “Behavioural Application Program Interfaces” (BehAPI).

More broadly, linearity allows for resource-aware programming and more *efficient* implementations [36], and it inspired unique types (as in Clean [4]), and ownership types (as in Rust [25]). A linear type system must keep track of what resources are used in which parts of the program, and guarantee that they are *neither duplicated nor discarded*. To do so, the standard approach is to use context splits: typing rules for terms with multiple subterms add an extra side condition specifying what resources to allocate to each of the subterms. The typing derivations for the subterms must then use the entirety of their allocated resources. A key observation here is that each subterm already *knows* about the resources it needs. *Context splits contain usage information that is already present in the subterms*. Moreover, the subterms cannot be typed until the context splits have been defined. On top of that, using binary context splits means that typing rules with n subterms require $n - 1$ context splits, which considerably clutters the type system.

An alternative approach is *leftover typing*, a technique used to formulate intuitionistic linear logic [24] and to mechanise the linear λ -calculus [2]. Leftover typing changes the shape of the typing judgments and includes a second *leftover* output context that contains the resources that were left unused by the term. As a result, typing rules *thread* the resources *through* subterms without needing context splits: each subterm uses the resources it needs, and leaves the rest for its siblings. The first subterm in this chain of resources immediately knows what resources it has available.

In this paper, we use leftover typing to define for the first time a resource-aware type system for the π -calculus, and we fully mechanise our work in Agda [38]. All previous work on mechanisation of linear process calculi uses context splits instead [16,19,17,34,8]. We will further highlight the benefits of leftover typing as opposed to context splits in contributions and the rest of the paper.

Below we present two alternative typing rules for parallel composition in the linear π -calculus: the one on the left uses context splits, while the one on the right does not, and uses leftover typing instead:

$$\frac{\Gamma := \Delta \otimes \Xi \quad \Delta \vdash P \quad \Xi \vdash Q}{\Gamma \vdash P \parallel Q} \qquad \frac{\Gamma \vdash P \triangleright \Delta \quad \Delta \vdash Q \triangleright \Xi}{\Gamma \vdash P \parallel Q \triangleright \Xi}$$

Contributions and Structure of the Paper

1. **Leftover typing for resource-aware π -calculus.** Our type system uses leftover typing to model the resource-aware π -calculus (§ 4.3) and satisfies subject reduction (Theorem 5). In addition to making context splits unnecessary, leftover typing allows for a *framing* theorem (Theorem 1) to be stated and is naturally associative, making type safety properties considerably easier to reason about (§ 5). Thanks to leftover typing, we can now state *weakening* (Theorem 2) and *strengthening* (Theorem 3) for the whole framework, not just the shared fragment. This give a uniform and complete presentation of all the meta-theory for the resource-aware π -calculus.
2. **Shared, graded and linear unified π -calculus.** We generalise resource counting to a set of usage algebras that can be mixed within the same type

system. We do not instantiate our type system to only work with linear resources, instead we present an algebra-agnostic type system, and admit a mix of user-defined *resource aware* algebras [21,35] (§ 4.1). Any *partial commutative monoid* that is *decidable*, *deterministic*, *cancellative* and has a *minimal element* is a valid such algebra. Multiple algebras can be mixed in the type system — usage contexts keep information about what algebra to use for each type (§ 4.2). In particular, this allows for type systems combining linear (use exactly once), graded (exact number of n times) and shared (free to reuse and discard) types under the same framework.

3. **Full mechanisation in Agda.** The formalisation of the π -calculus with leftover typing, from the syntax to the semantics and the type system, has been fully mechanised in Agda in about 1850 lines of code, and is publicly available at [38]. We have fully mechanised all meta-theory and the details of a proof of subject reduction can be found in our extended paper [37] and repository [38].

We use type level de Bruijn indices [12,15] to define a syntax of π -calculus processes that is *well scoped by construction*: every free variable is accounted for in the type of the process that uses it (§ 2). We then provide an operational semantics for the π -calculus, prior to any typing (§ 3). This operational semantics is defined as a reduction relation on processes. The reduction relation tracks at the type level the channel on which communication occurs. This information is later used to state the subject reduction theorem. The reduction relation is defined modulo *structural congruence* — a relation defined on processes that acts as a quotient type to remove unnecessary syntactic minutiae introduced by the syntax of the π -calculus. We then define an interface for resource-aware algebras (§ 4.1) and use it to parametrise a type system based on leftover typing (§ 4.3). Finally, we present the meta theoretical properties of our type system in § 5.

Notation Data type definitions (\mathbb{N}) use double inference lines and index-free synonyms (NAT) as rule names for ease of reference. Constructors (0 and $1+$) are used as inference rule names. We maintain a close correspondence between the definitions presented in this paper and our mechanised definitions in Agda: inference rules become type constructors, premises become argument types and conclusions return types. Universe levels and universe polymorphism are omitted for brevity — all our types are of type SET . Implicit arguments are mentioned in type definitions but omitted by constructors.

$$\frac{\text{NAT}}{\mathbb{N} : \text{SET}} \qquad \frac{}{0 : \mathbb{N}} \qquad \frac{n : \mathbb{N}}{1+n : \mathbb{N}}$$

We use colours to further distinguish the different entities in this paper. **TYPES** are blue and uppercase, with indices as subscripts, **constructors** are orange, **functions** are teal, variables are black, and some constructor names are overloaded — and disambiguated by context.

2 Syntax

In order to mechanise the π -calculus syntax in Agda, we need to deal with bound names in continuation processes. Names are cumbersome to mechanise: they are not inherently well scoped, one has to deal with alpha-conversion, and inserting new variables into a context entails proving that their names differ from all other names in context. To overcome these challenges, we use de Bruijn indices [12], where a natural number n (aka *index*) is used to refer to the variable introduced n binders ago. That is, binders no longer introduce names; terms at different *depths* use different indices to refer to the same binding.

While de Bruijn indices are useful for mechanisation, they are not as readable as names. To overcome this difficulty and demonstrate the correspondence between a π -calculus that uses names and one that uses de Bruijn indices, we provide *conversion functions* in both directions and prove that they are inverses of each other up to α -conversion. Further details can be found in our extended paper [37] and repository [38].


Definition 1 (VAR and PROCESS). A variable reference occurring under n binders can refer to n distinct variables. We introduce the indexed family of types [15] VAR_n : for all naturals n , the type VAR_n has n distinct elements. We index processes according to their *depth*: for all naturals n , a process of type PROCESS_n contains free variables that can refer to n distinct elements. Every time we go under a binder, we increase the index of the continuation process, allowing the variable references within to refer to one more thing.

$$\begin{array}{c}
 \frac{n : \mathbb{N}}{\text{VAR}_n : \text{SET}} \text{VAR} \qquad \frac{n : \mathbb{N}}{0 : \text{VAR}_{1+n}} \qquad \frac{x : \text{VAR}_n}{1+x : \text{VAR}_{1+n}} \\
 \\
 \frac{n : \mathbb{N}}{\text{PROCESS}_n : \text{SET}} \text{PROCESS} \\
 \\
 \begin{array}{l}
 \text{PROCESS}_n ::= 0 \qquad \text{(inaction)} \\
 | \nu \text{PROCESS}_{1+n} \qquad \text{(restriction)} \\
 | \text{PROCESS}_n \parallel \text{PROCESS}_n \qquad \text{(parallel)} \\
 | \text{VAR}_n () \text{PROCESS}_{1+n} \qquad \text{(input)} \\
 | \text{VAR}_n \langle \text{VAR}_n \rangle \text{PROCESS}_n \qquad \text{(output)}
 \end{array}
 \end{array}$$

Process 0 denotes the terminated process, where no further communications can occur; process νP creates a new channel and binds it at index 0 in the continuation process P ; process $P \parallel Q$ composes P and Q in parallel; process $x () P$ receives data along channel x and makes that data available at index 0 in the continuation process P ; process $x \langle y \rangle P$ sends variable y over channel x and continues as process P .

Example 1 (The courier system).

We present a courier system that consists of three *roles*: a sender, who wants to send a package; a receiver, who receives the package sent by the sender; and a courier, who carries the package from the sender to the receiver.



Our courier system is defined by four π -calculus processes composed in parallel instantiating the above three roles: we have two sender processes, `send x` and `send y`, sending data over channels x and y , respectively; one receiver process, `recv z`, which receives over channel z the data sent from each of the senders – hence receives twice; and a courier process `carry x y z`, which synchronises communication among the senders and the receiver. The courier process first receives data from the two senders along its input channels x and y , and then sends the two received bits of data to the receiver along its output channel z .

The sender and receiver *roles* are defined below, parametrised by the channels on which they operate. The sender creates a new channel to be sent as data, and sends it over channel c , and then terminates. Processes `send x` and `send y` are an instantiation of `send c`. The receiver receives data *twice* on a channel c and then terminates. The receiver process `recv z` is an instantiation of `recv c`.

$$\text{send } c = \nu (1+c \langle 0 \rangle 0) \qquad \text{recv } c = c () (1+c) () 0$$

The courier role is defined below as `carry x y z`. It sequentially receives on the two input channels x and y , instantiated as $in0$ and $in1$, and then outputs the two pieces of received data on the output channel z , instantiated as out . Finally, we create three communication channels and compose all four processes together: the first channel is shared between the one sender and the courier, the second between the other sender and the courier, and the third between the receiver and the courier. The result is the courier `system` defined below.

$$\begin{aligned} \text{carry } in0 \ in1 \ out &= in0 () (1+ in1) () (1+1+ out) \langle 1+0 \rangle (1+1+ out) \langle 0 \rangle 0 \\ \text{system} &= \nu (\text{send } 0 \parallel \nu (\text{send } 0 \parallel \nu (\text{recv } 0 \parallel \text{carry } (1+1+0) (1+0) 0))) \end{aligned}$$

We continue this running example in §4.3, where we provide typing derivations for the above processes and use a mix of linear, graded and shared typing to type the courier `system`.

3 Operational Semantics

Thanks to our well-scoped grammar in §2, we now define the semantics of our language on the totality of the syntax.

Definition 2 (UNUSED). We consider a variable i to be unused in P ($\text{UNUSED}_i P$) if none of the inputs nor the outputs refer to it. $\text{UNUSED}_i P$ is defined as a recursive predicate on P , incrementing i every time we go under a binder, and using $i \neq x$ (which unfolds to the negation of propositional equality on VAR, i.e. $i \equiv x \rightarrow \perp$) to compare variables.

Definition 3 (STRUCTCONG). We define the base cases of a structural congruence relation \cong as follows:

$$\begin{array}{c}
\frac{}{P \cong Q : \text{SET}} \quad \text{STRUCTCONG} \quad \frac{}{\text{comp-assoc} : P \parallel (Q \parallel R) \cong (P \parallel Q) \parallel R} \\
\frac{}{\text{comp-sym} : P \parallel Q \cong Q \parallel P} \quad \frac{}{\text{comp-id} : P \parallel \mathbb{0}_n \cong P} \\
\frac{}{\text{scope-end} : \nu \mathbb{0}_{1+n} \cong \mathbb{0}_n} \quad \frac{uQ : \text{UNUSED}_0 Q}{\text{scope-ext} : \nu (P \parallel Q) \cong (\nu P) \parallel \text{lower}_0 Q \ uQ} \\
\frac{}{\text{scope-comm} : \nu \nu P \cong \nu \nu \text{exchange}_0 P}
\end{array}$$

The first three rules (comp-^*) state associativity, symmetry, and $\mathbb{0}$ as being the neutral element of parallel composition, respectively. The last three (scope-^*) state garbage collection, scope extrusion and commutativity of restrictions, respectively. In scope-ext the side condition $\text{UNUSED}_i Q$ makes sure that i is unused in Q (see Definition 2). The function $\text{lower}_i Q \ uQ$ traverses Q decrementing every index greater than i . In scope-comm the function $\text{exchange}_i P$ traverses P (of type PROCESS_{1+1+n}) and swaps variable references i and $1+i$. In all the above, i is incremented every time we go under a binder.

Definition 4 (EQUALS). We lift the relation $\text{STRUCTCONG} \cong$ and close it under equivalence and congruence in \simeq . This relation is structurally congruent under a context $\mathcal{C}[\cdot]$ [32] and is reflexive, symmetric and transitive.

Definition 5 (REDUCES). The operational semantics of the π -calculus is defined as a reduction relation \longrightarrow_c indexed by the channel c on which communication occurs. We keep track of channel c so we can state subject reduction (Theorem 5).

$$\begin{array}{c}
 \frac{n : \mathbb{N}}{\text{CHANNEL}_n : \text{SET}} \text{ CHANNEL} \qquad \frac{}{\text{internal} : \text{CHANNEL}_n} \\
 \frac{i : \text{VAR}_n}{\text{external } i : \text{CHANNEL}_n} \qquad \frac{c : \text{CHANNEL}_n \quad P \ Q : \text{PROCESS}_n}{P \longrightarrow_c Q : \text{SET}} \text{ REDUCES} \\
 \\
 \frac{i \ j : \text{VAR}_n \quad P : \text{PROCESS}_{1+n} \quad Q : \text{PROCESS}_n}{\text{comm} : i \langle \rangle P \parallel i \langle j \rangle Q \longrightarrow_{\text{external } i} \text{lower}_0 (P [0 \mapsto 1+j]) uP' \parallel Q} \\
 \\
 \frac{\text{red} : P \longrightarrow_c P'}{\text{par red} : P \parallel Q \longrightarrow_c P' \parallel Q} \qquad \frac{\text{red} : P \longrightarrow_c Q}{\text{res red} : \nu P \longrightarrow_{\text{dec } c} \nu Q} \\
 \\
 \frac{eq_1 : P \simeq P' \quad \text{red} : P' \longrightarrow_c Q' \quad eq_2 : Q' \simeq Q}{\text{struct eq red} : P \longrightarrow_c Q}
 \end{array}$$

We distinguish between channels that are created inside the process (**internal**), and channels that are created outside (**external** i), where i is the index of the channel variable. In rule **comm**, parallel processes reduce when they communicate over a common channel with index i . As a result of that communication, the continuation of the input process P has all the references to its most immediate variable substituted with references to $1+j$, the variable sent by the output process $i \langle j \rangle Q$. After this substitution, $P [0 \mapsto 1+j]$ is *lowered* — all variable references are decreased by one (and we derive the proof $\text{UNUSED}_0 (P [0 \mapsto 1+j])$). Reduction is closed under parallel composition (rule **par**), restriction (rule **res**) and structural congruence (rule **struct**) — notably, not under input nor output, as doing so would not preserve the sequencing of actions [32]. Rule **res** uses **dec** to decrement the index of channel c as we wrap processes P and Q inside a binder. It is defined as expected below:

$$\begin{array}{ll}
 \text{dec internal} & = \text{internal} \\
 \text{dec (external } 0) & = \text{internal} \\
 \text{dec (external } (1+n)) & = \text{external } n
 \end{array}$$

4 Resource-aware Type System

In § 4.1 we characterise a usage algebra for our type system. It defines how resources are *split* in parallel composition and *consumed* in input and output. We define typing and usage contexts in § 4.2. We provide a type system for a resource-aware π -calculus in § 4.3.

4.1 Multiplicities and Capabilities

In the linear π -calculus each channel has an input and an output *capability*, and each capability has a given *multiplicity* of 0 (exhausted) or 1 (available).

We generalise over this notion by defining an algebra for multiplicities [21,35] that is satisfied by linear, graded and shared types alike. We then use pairs of multiplicities as usage annotations for a channel's input and output capabilities.

Definition 6 (ALGEBRA). A *usage algebra* is a ternary relation $x := y \cdot z$ that is *partial* (as not any two multiplicities can be combined), *deterministic* and *cancellative* (to aid equational reasoning) and *associative* and *commutative* (following directly from subject congruence for parallel composition). In addition, we ask that the leftovers can be *computed* so that we can automatically update the usage context every time input and output occurs — this is purely for usability. It has a *neutral element* -0 that is absorbed on either side, and that is also *minimal* (so that new resources cannot arbitrarily spring into life). It has an element -1 that is used to count inputs and outputs. Below we define such an algebra as a record ALGEBRA_C on a carrier C . (We use \forall for universal quantification. The dependent product \exists uses the value of its first argument in the type of its second. The type $\text{DEC } P$ is a witness of either P or $P \rightarrow \perp$, where \perp is the empty type with no constructors.)

$- := - \cdot -$:	$C \rightarrow C \rightarrow C \rightarrow \text{SET}$
$-\text{unique}$:	$\forall x x' y z \rightarrow x' := y \cdot z \rightarrow x := y \cdot z \rightarrow x' \equiv x$
$-\text{unique}^1$:	$\forall x y y' z \rightarrow x := y' \cdot z \rightarrow x := y \cdot z \rightarrow y' \equiv y$
$-\text{assoc}$:	$\forall x y z u v \rightarrow x := y \cdot z \rightarrow y := u \cdot v \rightarrow \exists w (x := u \cdot w \times w := v \cdot z)$
$-\text{comm}$:	$\forall x y z \rightarrow x := y \cdot z \rightarrow x := z \cdot y$
$-\text{compute}^\tau$:	$\forall x y \rightarrow \text{DEC } (\exists z (x := y \cdot z))$
-0	:	C
$-\text{id}^1$:	$\forall x \rightarrow x := -0 \cdot x$
$-\text{min}^1$:	$\forall y z \rightarrow -0 := y \cdot z \rightarrow y \equiv -0$
-1	:	C

We sketch the implementation of linear, graded and shared types as instances of our usage algebra below. Their use in typing derivations is illustrated in Example 3.

	carrier	operation
linear	$0 : \text{Lin}$ $1 : \text{Lin}$	$0 := 0 \cdot 0$ $1 := 1 \cdot 0$ $1 := 0 \cdot 1$
graded	$0 : \text{Gra}$ $1+ : \text{Gra} \rightarrow \text{Gra}$	$\forall x y z$ $\rightarrow x \equiv y + z$ $\rightarrow x := y \cdot z$
shared	$\omega : \text{Sha}$	$\omega := \omega \cdot \omega$

4.2 Typing Contexts

We use indexed sets of usage algebras to allow several usage algebras to coexist in our type system with leftovers (§4.3).

Definition 7 (ALGEBRAS). An *indexed set of usage algebras* is a type **IDX** of indices that is nonempty (\exists **IDX**) together with an interpretation **USAGE** of indices into types, and an interpretation **ALGEBRAS** of indices into usage algebras of the corresponding type.

$$\begin{aligned} \mathbf{IDX} & : \mathbf{SET} \\ \exists \mathbf{IDX} & : \mathbf{IDX} \\ \mathbf{USAGE} & : \mathbf{IDX} \rightarrow \mathbf{SET} \\ \mathbf{ALGEBRAS} & : (idx : \mathbf{IDX}) \rightarrow \mathbf{ALGEBRA}_{\mathbf{USAGE}_{idx}} \end{aligned}$$

We keep typing contexts (**PRECTX**) and usage contexts (**CTX**) separate. The former are preserved throughout typing derivations; the latter are transformed as a result of input, output, and context splits.

Definition 8 (TYPE and PRECTX: types and typing contexts). A *type* is either a unit type (**1**), or a channel type (**C**[t ; x]).

$$\frac{}{\mathbf{TYPE} : \mathbf{SET}} \text{TYPE} \qquad \frac{}{\mathbf{1} : \mathbf{TYPE}} \qquad \frac{idx : \mathbf{IDX} \quad x : \mathbf{USAGE}_{idx}^2 \quad t : \mathbf{TYPE}}{\mathbf{C}[t;x] : \mathbf{TYPE}}$$

The unit type **1** serves as a base case for types. The type **C**[t ; x] of a channel determines what type t of data and what usage annotations x are sent over that channel — we use the notation \mathbf{C}^2 to stand for a $\mathbf{C} \times \mathbf{C}$ pair of input and output multiplicities, respectively. This channel notation aligns with $[t] \mathbf{chan}_{(i^y, o^z)}$, where y, z are the input and output multiplicities, respectively [22]. Henceforth, we use ℓ_\emptyset to denote the multiplicity pair $--0, --0$, ℓ_i for the pair $--1, --0$, ℓ_o for $--0, --1$, and $\ell_\#$ for $--1, --1$. This notation was originally used in the linear π -calculus [23,32]. A *typing context* \mathbf{PRECTX}_n is a length-indexed list of types that is either empty (\square) or the result of appending a type $t : \mathbf{TYPE}$ to an existing context (γ, t) .

Definition 9 (IDXs and CTX: contexts of indices and usage contexts). A context of indices \mathbf{IDXs}_n is a length-indexed list that is either empty (\square) or the result of appending an index $i : \mathbf{IDX}$ to an existing context $(idxs, i)$. A *usage context* is a context \mathbf{CTX}_{idxs} indexed by a context of indices $idxs : \mathbf{IDXs}_n$ that is either empty (\square) or the result of appending a usage annotation pair $u : \mathbf{USAGE}_{idx}^2$ with index $idx : \mathbf{IDX}$ to an existing context (Γ, u) .

4.3 Typing with Leftovers

We present a resource-aware type system for the π -calculus based on *leftover typing* [2], a technique that, in addition to the usual typing context \mathbf{PRECTX}_n and (input) usage context \mathbf{CTX}_{idxs} , adds an extra (*output*) usage context \mathbf{CTX}_{idxs} to the typing rules. This output context contains the *leftovers* (the unused multiplicities) of the process being typed. These leftovers can then be used as input to another typing derivation.

Leftover typing inverts the information flow of usage annotations so that it is the typing derivations of subprocesses which determine how resources are allocated. As a result, context split proofs are no longer necessary. Leftover typing also allows *framing* to be stated, and *weakening* and *strengthening* to cover linear types too.

Our type system is composed of two typing judgments: one for variable references (Definition 10) and one for processes (Definition 11). Both judgments are indexed by a typing context γ , an input usage context Γ , and an output usage context Δ (the leftovers). The **typing judgement for variables** $\gamma; \Gamma \ni_i t; y \triangleright \Delta$ asserts that “index i in typing context γ is of type t , and subtracting y at position i from input usage context Γ results in leftovers Δ ”. The **typing judgement for processes** $\gamma; \Gamma \vdash P \triangleright \Delta$ asserts that “process P is well typed under typing context γ , usage input context Γ and leftovers Δ ”.

Definition 10 (VARREF: typing variable references). The VARREF typing relation for variable references is presented below.

$$\frac{\begin{array}{c} t : \mathbf{TYPE} \\ idx : \mathbf{IDX} \quad idxs : \mathbf{IDX}_n \\ \gamma : \mathbf{PRECTX}_n \quad i : \mathbf{VAR}_n \quad y : \mathbf{USAGE}_{idx}^2 \quad \Gamma \Delta : \mathbf{CTX}_{idxs} \end{array}}{\gamma; \Gamma \ni_i t; y \triangleright \Delta : \mathbf{SET}} \text{VARREF}$$

$$\frac{x := y \cdot^2 z}{\mathbf{0} : \gamma, t; \Gamma, x \ni_0 t; y \triangleright \Gamma, z} \quad \frac{v : \gamma; \Gamma \ni_i t; x \triangleright \Delta}{\mathbf{1+} v : \gamma, t'; \Gamma, x' \ni_{\mathbf{1+}i} t; x \triangleright \Delta, x'}$$

We lift the operation $x := y \cdot z$ and its algebraic properties to an operation $(x_l, x_r) := (y_l, y_r) \cdot^2 (z_l, z_r)$ on pairs of multiplicities. The base case $\mathbf{0}$ splits the usage annotation x of type \mathbf{USAGE}_{idx}^2 into y and z (the leftovers). Note that the remaining context Γ is preserved unused as a leftover. This splitting $x := y \cdot^2 z$ is as per the usage algebra provided by the developer for the index idx . In our Agda implementation, $x := y \cdot^2 z$ is actually a trivially satisfiable implicit argument if $x := y \cdot^2 z$ is inhabited and an unsatisfiable argument otherwise. The inductive case $\mathbf{1+}$ appends the type t' to the typing context, and the usage annotation x' to both the input and output usage contexts.

Example 2 (Variable reference). `egVar` defines a variable reference $\mathbf{1+0}$ with type $\mathbf{C}[\mathbf{1}; \ell_i]$ and usage ℓ_i . We must show that this variable is well typed in an environment with a typing context $\gamma = \square, \mathbf{C}[\mathbf{1}; \ell_i], \mathbf{1}$ and a usage context $\Gamma = \square, \ell_{\#}, \ell_{\#}$. The VARREF constructors are completely determined by the variable index $\mathbf{1+0}$ in the type. The constructor $\mathbf{1+}$ steps under the outermost variable in the context, preserving its usage annotation $\ell_{\#}$ from input to output. The constructor $\mathbf{0}$ asserts that the next variable is of type $\mathbf{C}[\mathbf{1}; \ell_i]$, and that the usage annotation $\ell_{\#}$ can be split such that $\ell_{\#} := \ell_i \cdot \ell_o$ — using `--computer` to automatically fulfill the proof obligation.

$$\begin{array}{l} \text{egVar} : (\square, \mathbf{C}[\mathbf{1}; \ell_i], \mathbf{1}); (\square, \ell_{\#}, \ell_{\#}) \ni_{\mathbf{1+0}} \mathbf{C}[\mathbf{1}; \ell_i]; \ell_i \triangleright (\square, \ell_o, \ell_{\#}) \\ \text{egVar} = \mathbf{1+0} \end{array}$$

Definition 11 (TYPES: typing processes). The TYPES typing relation for the resource-aware π -calculus processes is presented below. For convenience, we reuse the constructor names introduced for the syntax in §2.

$$\begin{array}{c}
 \frac{\gamma : \text{PRECTX}_n \quad P : \text{PROCESS}_n \quad \frac{idxs : \text{IDXS}_n \quad \Gamma \Delta : \text{CTX}_{idxs}}{\text{TYPES}}}{\gamma ; \Gamma \vdash P \triangleright \Delta : \text{SET}} \\
 \\
 \frac{\frac{\mathbb{0} : \gamma ; \Gamma \vdash \mathbb{0} \triangleright \Gamma}{\text{}} \quad \frac{t : \text{TYPE} \quad x : \text{USAGE}_{idx}^2 \quad y : \text{USAGE}_{idx'} \quad cont : \gamma, \mathbb{C}[t; x]; \Gamma, (y, y) \vdash P \triangleright \Delta, \ell_\emptyset}{\nu t x y cont : \gamma ; \Gamma \vdash \nu P \triangleright \Delta}}{\mathbb{0} : \gamma ; \Gamma \vdash \mathbb{0} \triangleright \Gamma} \\
 \\
 \frac{\frac{chan : \gamma \quad ; \Gamma \quad \ni_i \mathbb{C}[t; x]; \ell_i \triangleright \Xi \quad cont : \gamma, t; \Xi, x \vdash P \quad \triangleright \Theta, \ell_\emptyset}{chan (\) cont : \gamma ; \Gamma \vdash i (\) P \triangleright \Theta} \quad \frac{chan : \gamma ; \Gamma \ni_i \mathbb{C}[t; x]; \ell_o \triangleright \Delta \quad loc : \gamma ; \Delta \ni_j t \quad ; x \triangleright \Xi \quad cont : \gamma ; \Xi \vdash P \quad \triangleright \Theta}{chan \langle loc \rangle cont : \gamma ; \Gamma \vdash i \langle j \rangle P \triangleright \Theta}}{\mathbb{0} : \gamma ; \Gamma \vdash \mathbb{0} \triangleright \Gamma} \\
 \\
 \frac{l : \gamma ; \Gamma \vdash P \triangleright \Delta \quad r : \gamma ; \Delta \vdash Q \triangleright \Xi}{l \parallel r : \gamma ; \Gamma \vdash P \parallel Q \triangleright \Xi}
 \end{array}$$

The inaction process in rule $\mathbb{0}$ does not change usage annotations. The scope restriction in rule ν expects three arguments: the type t of data being transmitted; the usage annotation x of what is being transmitted; and the multiplicity y given to the channel itself. This multiplicity y is used for both input and output, so that they are balanced. The continuation process P is provided with the new channel with usage annotation y, y , which it must completely exhaust. The input process in rule $(\)$ requires a channel $chan$ at index i with usage ℓ_i available, such that data with type t and usage x can be sent over it. Note that the index i is determined by the syntax of the typed process. We use the leftovers Ξ to type the continuation process, which is also provided with the received element — of type t and multiplicity x — at index $\mathbb{0}$. The received element x must be completely exhausted by the continuation process. Similarly to input, the output process in rule $\langle \ \rangle$ requires a channel $chan$ at index i with usage ℓ_o available, such that data with type t and usage x can be sent over it. We use the leftover context Δ to type the transmitted data, which needs an element loc at index j with type t and usage x , as per the type of the channel $chan$. The leftovers Ξ are used to type the continuation process. Note that both indices i and j are determined by the syntax of the typed process. Parallel composition in rule \parallel uses the leftovers of the left-hand process to type the right-hand process. Indeed, Theorem 4 shows that an alternative rule where the resources are first threaded through Q is admissible too.

Example 3 (Typing derivation (Continued)). We provide the typing derivation for the courier system defined in Example 1. For the sake of simplicity, we instantiate these processes with concrete variable references before typing them.

The receiver defined by the `recv` process receives data along the channel with index `0`, which needs to be of type $\mathbf{C}[t; u]$ for some t and u . After receiving twice, the process ends: we must not be left with any unused multiplicities, thus $u = \ell_\emptyset$. We will use graded types to keep track of the exact number of times communication happens. Whatever the input multiplicity of the channel, we will consume 2 of it and leave the remaining as leftovers. The sender defined by the `send` process sends data along the channel with index `0`, which needs to be of type $\mathbf{C}[t; u]$ for some t and u . We instantiate t (the type of data that the sender sends) to the trivial channel $\mathbf{C}[\mathbb{1}; \omega]$. As per the type of the process `recv`, $u = \ell_\emptyset$. We will transmit once, thus use $\mathbf{1+0}$ output multiplicity, and leave the rest as leftovers. Agda can uniquely determine the arguments required by the ν constructor.

```

recvwt :  $\gamma, \mathbf{C}[t; \ell_\emptyset]; \Gamma, (\mathbf{1+1+l}, r) \vdash \text{recv } 0 \triangleright \Gamma, (l, r)$ 
recvwt =  $0 \ () \ (\mathbf{1+0}) \ () \ 0$ 
sendwt :  $\gamma, \mathbf{C}[\mathbf{C}[\mathbb{1}; \omega]; \ell_\emptyset]; \Gamma, (l, \mathbf{1+r}) \vdash \text{send } 0 \triangleright \Gamma, (l, r)$ 
sendwt =  $\nu \ \_ \_ \_ 0 \ (\mathbf{1+0} \langle 0 \rangle 0)$ 

```

Dually, the courier defined by the `carry` process expects input multiplicities for the channels shared with `send` and output multiplicities for the channel shared with `recv`. We can now compose these processes in parallel and type the courier system.

```

carrywt :  $\gamma, \mathbf{C}[t; \ell_\emptyset], \mathbf{C}[t; \ell_\emptyset], \mathbf{C}[t; \ell_\emptyset]$ 
         ;  $\Gamma, (\mathbf{1+lx}, rx), (\mathbf{1+ly}, ry), (lz, \mathbf{1+1+rz})$ 
          $\vdash \text{carry } (\mathbf{1+1+0}) \ (\mathbf{1+0}) \ 0$ 
          $\triangleright \Gamma, (lx, rx), (ly, ry), (lz, rz)$ 
carrywt =  $(\mathbf{1+1+0}) \ () \ (\mathbf{1+1+0}) \ () \ (\mathbf{1+1+0}) \ \langle \mathbf{1+0} \rangle \ (\mathbf{1+1+0}) \ \langle 0 \rangle \ 0$ 
systemwt :  $[\ ]; [\ ] \vdash \text{system} \triangleright [\ ]$ 
systemwt =  $\nu \ \_ \_ \_ (\text{sendwt} \parallel \nu \ \_ \_ \_ (\text{sendwt} \parallel \nu \ \_ \_ \_ (\text{recvwt} \parallel \text{carrywt})))$ 

```

5 Meta-Theory

We have mechanised subject reduction for our π -calculus with leftovers in 850 lines of Agda code. The meta-theory of resource-aware type systems often needs to reason on typing derivations modulo associativity in the allocation of resources. For type systems using context splitting side conditions, this means applying associativity lemmas to recompute context splits; for type systems using leftover typing it does not. As an example, the proof that `comp-assoc` preserves typing proceeds by deconstructing the input derivation into $P \parallel (Q \parallel R)$ and reassembling it as $(P \parallel Q) \parallel R$ without the need of any extra reasoning.

All the reasoning carried out in our type safety proofs is based on the algebraic properties introduced in § 4.1 – the exception to this is `--computer`, only there for the user’s convenience. We lift the operation $x := y \cdot^2 z$ and its algebraic

properties to an operation $\Gamma := \Delta \otimes \Xi$ on usage contexts that have the same underlying context of indices. The algebraic properties of the algebras allow us to see a typing derivation $\gamma; \Gamma \vdash P \triangleright \Delta$ as a unique *arrow* from Γ to Δ , and to freely compose and reason with arrows with the same typing context and a matching output and input usage contexts.

Leftover typing also allows us to state a *framing* theorem showing that adding or subtracting arbitrary usage annotations to the input and output usage contexts preserves typing – one can understand a typing derivation independently from its unused resources. With framing one can show that **comp-comm** preserves typing: in $P \parallel Q$ the typing of P and Q is independent of one another.

Theorem 1 (Framing). Let $\gamma; \Gamma_l \vdash P \triangleright \Xi_l$. Let Δ be such that $\Gamma_l := \Delta \otimes \Xi_l$. Then for any Γ_r and Ξ_r where $\Gamma_r := \Delta \otimes \Xi_r$ it holds that $\gamma; \Gamma_r \vdash P \triangleright \Xi_r$.

Leftover typing allows *weakening* and *strengthening* to acquire a more general form where linear variables can freely be added or removed from context too – as long as they are added and removed to and from both the input and output contexts.

Theorem 2 (Weakening). Let ins_i insert an element into a context at position i . Let P be well typed in $\gamma; \Gamma \vdash P \triangleright \Xi$. Then, lifting every variable greater than or equal to i in P is well typed in $\text{ins}_i t \gamma; \text{ins}_i x \Gamma \vdash \text{lift}_i P \triangleright \text{ins}_i x \Xi$.

Theorem 3 (Strengthening). Let del_i delete the element at position i from a context. Let P be well typed in $\gamma; \Gamma \vdash P \triangleright \Xi$. Let i be a variable not in P , such that $uP : \text{UNUSED}_i P$. Then lowering every variable greater than i in P is well typed in $\text{del}_i \gamma; \text{del}_i \Gamma \vdash \text{lower}_i P uP \triangleright \text{del}_i \Xi$.

Subject congruence states that structural congruence (Definition 4) preserves the well-typedness of a process.

Theorem 4 (Subject Congruence). Let P and Q be processes. If $P \simeq Q$ and $\gamma; \Gamma \vdash P \triangleright \Xi$, then $\gamma; \Gamma \vdash Q \triangleright \Xi$.

Finally, subject reduction states that reducing on a channel c (Definition 5) preserves the well-typedness of a process — after consuming $\ell_\#$ from c if c is an **external** channel. Below we use $\Gamma \ni_i x \triangleright \Delta$ to stand for $\gamma; \Gamma \ni_i t; x \triangleright \Delta$ for some γ and t .

Theorem 5 (Subject Reduction). Let $\gamma; \Gamma \vdash P \triangleright \Xi$ and $P \rightarrow_c Q$. If c is **internal**, then $\gamma; \Gamma \vdash Q \triangleright \Xi$. If c is **external** i and $\Gamma \ni_i \ell_\# \triangleright \Delta$, then $\gamma; \Delta \vdash Q \triangleright \Xi$.

We refer to our extended paper [37] and repository [38] for a more detailed account of the mechanised proofs.

6 Conclusions, Related and Future Work

Extrinsic Encodings Extrinsic encodings define a syntax (often well-scoped) and a runtime semantics prior to any type system. This allows one to talk about ill-typed terms, and defers the proof of subject reduction to a later stage. To the best of our knowledge, leftover typing makes its appearance in 1994, when Ian Mackie first uses it to formulate intuitionistic linear logic [24]. Allais [2] uses leftover typing to mechanise in Agda a bidirectional type system for the linear λ -calculus. He proves type preservation and provides a decision procedure for type checking and type inference. In this paper, we follow Allais [2] and apply leftover typing to the π -calculus for the first time. We generalise the usage algebra, leading to linear, graded and shared type systems. Drawing from quantitative type theory (by McBride and Atkey [26,3]), in our work we too are able to talk about fully consumed resources — e.g., we can transmit ℓ_0 multiplicities of a fully exhausted channel. Recent years have seen an increase in the efforts to mechanise resource-aware process algebras, but one of the earliest works is the mechanisation of the linear π -calculus in Isabelle/HOL by Gay [16]. Gay encodes the π -calculus with linear and shared types using de Bruijn indices, a reduction relation and a type system posterior to the syntax. However, in his work typing rules demand user-provided context splits, and variables with consumed usage annotations are erased from context. We remove the demand for context splits, preserve the ability to talk about consumed resources, and adopt a more general usage algebra. Orchard et al. introduce Granule [29], a fully-fledged functional language with graded modal types, linear types, indexed types and polymorphism. Modalities include exact usages, security levels and intervals; resource algebras are pre-ordered semirings with partial addition. The authors provide bidirectional typing rules, and show the type safety of their semantics. The work by Goto et al. [19] is, to the best of our knowledge, the first formalisation of session types which comes along with a mechanised proof of type safety in Coq. The authors extend session types with polymorphism and pattern matching. They use a locally-nameless encoding for variable references, a syntax prior to types, and an LTS semantics that encodes session-typed processes into the π -calculus. Their type system uses reordering of contexts and extrinsic context splits, which are not needed in our work.

Intrinsic Encodings Intrinsic encodings merge syntax and type system. As a result, one can only ever talk about well-typed terms, and the reduction relation by construction carries a proof of subject reduction. Significantly, by merging the syntax and static semantics of the object language one can fully use the expressive power of the host language. Thiemann formalises in Agda the MicroSession (minimal GV [17]) calculus with support for recursion and subtyping [34]. As Gay does in [16], context splits are given extrinsically, and exhausted resources are removed from typing contexts altogether. The runtime semantics are given as an intrinsically typed CEK machine with a global context of session-typed channels. In their recent paper, Ciccone and Padovani mechanise a dependently-typed linear π -calculus in Agda [8]. Their intrinsic encoding allows them to

leverage Agda’s dependent types to provide a dependently-typed interpretation of messages — to avoid linearity violations the interpretation of channel types is erased. Message input is modeled as a dependent function in Agda, and as a result message predicates, branching, and variable-length conversations can be encoded. In contrast to our work, their algebra is on the multiplicities 0, 1, ω , and top-down context splitting proofs must be provided. In another recent work, Rouvoet et al. provide an intrinsic type system for a λ -calculus with session types [31]. They use proof relevant separation logic and a notion of a supply and demand *market* to make context splits transparent to the user. Their separation logic is based on a partial commutative monoid that need not be deterministic nor cancellative. Their typing rules preserve the balance between supply and demand, and are extremely elegant. They distill their typing rules even further by modelling the supply and demand market as a state monad.

Other Work Castro et al. [6] provide tooling for locally-nameless representations of process calculi in Coq, where de Bruijn indices are less popular than in Agda or Idris. They use their tool to help automate proofs of subject reduction for a type system with session types. Orchard and Yoshida [30] embed a small effectful imperative language into the session-typed π -calculus, showing that session types are expressive enough to encode effect systems. Based on contextual type theory, LINCX [18] extends the linear logical framework LLF [7] by internalising the notion of bindings and contexts. The result is a meta-theory in which HOAS encodings with both linear and dependent types can be described. The developer obtains for free an equational theory of substitution and decidable typechecking without having to encode context splits within the object language. Further work on mechanisation of the π -calculus [13,20,5,14,1], focuses on non-linear variations, differently from our range of linear, graded and shared types.

Conclusions and Future Work We provide a well-scoped syntax and a semantics for the π -calculus, extrinsically define a type system on top of the syntax capable of handling linear, graded and shared types under the same unified framework and show subject reduction. We avoid extrinsic context splits by defining a type system based on leftover typing [2]. As a result, theorems like framing, weakening and strengthening can now be stated also for the linear π -calculus. Our work is fully mechanised in around 1850 lines of code in Agda [38].

As future work we intend to expand our framework to include infinite behaviour by adding process replication, which is challenging, as to prove subject congruence one needs to uniquely determine the resources consumed by a process — e.g., by adding type annotations to the syntax. Orthogonally, we aim to investigate making our typing rules bidirectional which would allow us to provide a decision procedure for type checking processes in a given set of algebras. Finally, we will use our π -calculus with leftovers as an underlying framework on top of which we can implement session types, via their encodings into linear types [9,11,33] and other advanced type theories.

Acknowledgments We want to thank Erika Kreuter, Wen Kokke, James Wood, Guillaume Allais, Bob Atkey, and Conor McBride for their valuable suggestions.

References

1. Affeldt, R., Kobayashi, N.: A Coq Library for Verification of Concurrent Programs. *Electron. Notes Theor. Comput. Sci.* **199**, 17–32 (2008). <https://doi.org/10.1016/j.entcs.2007.11.010>
2. Allais, G.: Typing with Leftovers - A mechanization of Intuitionistic Multiplicative-Additive Linear Logic. In: *Types for Proofs and Programs, TYPES. LIPIcs*, vol. 104, pp. 1:1–1:22. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2017). <https://doi.org/10.4230/LIPIcs.TYPES.2017.1>
3. Atkey, R.: Syntax and Semantics of Quantitative Type Theory. In: *Logic in Computer Science, LICS*. pp. 56–65. ACM (2018). <https://doi.org/10.1145/3209108.3209189>
4. Barendsen, E., Smetsers, S.: Uniqueness Typing for Functional Languages with Graph Rewriting Semantics. *Math. Struct. Comput. Sci.* **6**(6), 579–612 (1996)
5. Bengtson, J.: The pi-calculus in nominal logic, vol. 2012 (2012), https://www.isa-afp.org/entries/Pi_Calculus.shtml
6. Castro, D., Ferreira, F., Yoshida, N.: EMTST: Engineering the Meta-theory of Session Types. In: *Tools and Algorithms for the Construction and Analysis of Systems, TACAS. Lecture Notes in Computer Science*, vol. 12079, pp. 278–285. Springer (2020). https://doi.org/10.1007/978-3-030-45237-7_17
7. Cervesato, I., Pfenning, F.: A Linear Logical Framework. In: *Logic in Computer Science, LICS*. pp. 264–275. IEEE Computer Society (1996). <https://doi.org/10.1109/LICS.1996.561339>
8. Ciccone, L., Padovani, L.: A Dependently Typed Linear π -Calculus in Agda. In: *PPDP '20: 22nd International Symposium on Principles and Practice of Declarative Programming*. pp. 8:1–8:14. ACM (2020). <https://doi.org/10.1145/3414080.3414109>
9. Dardha, O.: Recursive Session Types Revisited. In: Carbone, M. (ed.) *Workshop on Behavioural Types, BEAT. EPTCS*, vol. 162, pp. 27–34 (2014). <https://doi.org/10.4204/EPTCS.162.4>
10. Dardha, O., Giachino, E., Sangiorgi, D.: Session types revisited. In: *Principles and Practice of Declarative Programming, PPDP*. pp. 139–150. ACM (2012). <https://doi.org/10.1145/2370776.2370794>
11. Dardha, O., Giachino, E., Sangiorgi, D.: Session types revisited. *Inf. Comput.* **256**, 253–286 (2017). <https://doi.org/10.1016/j.ic.2017.06.002>, extended version of [10]
12. de Bruijn, N.G.: Lambda Calculus Notation with Nameless Dummies, a Tool for Automatic Formula Manipulation, with Application to the Church-Rosser Theorem. In: *Indagationes Mathematicae (Proceedings)*. vol. 75, pp. 381–392. Elsevier (1972)
13. Deransart, P., Smaus, J.: Subject Reduction of Logic Programs as Proof-Theoretic Property, vol. 2002 (2002), <http://danae.uni-muenster.de/lehre/kuchen/JFLP/articles/2002/S02-01/JFLP-A02-02.pdf>
14. Despeyroux, J.: A Higher-Order Specification of the pi-Calculus, *Lecture Notes in Computer Science*, vol. 1872. Springer (2000). https://doi.org/10.1007/3-540-44929-9_30

15. Dybjer, P.: Inductive Families. *Formal Asp. Comput.* **6**(4), 440–465 (1994). <https://doi.org/10.1007/BF01211308>
16. Gay, S.J.: A Framework for the Formalisation of Pi Calculus Type Systems in Isabelle/HOL. In: *Theorem Proving in Higher Order Logics, TPHOLs. Lecture Notes in Computer Science*, vol. 2152, pp. 217–232. Springer (2001). https://doi.org/10.1007/3-540-44755-5_16
17. Gay, S.J., Vasconcelos, V.T.: Linear type theory for asynchronous session types. *J. Funct. Program.* **20**(1), 19–50 (2010). <https://doi.org/10.1017/S0956796809990268>
18. Georges, A.L., Murawska, A., Otis, S., Pientka, B.: LINCX: A Linear Logical Framework with First-Class Contexts. In: *European Symposium on Programming, ESOP, Lecture Notes in Computer Science*, vol. 10201, pp. 530–555. Springer (2017). https://doi.org/10.1007/978-3-662-54434-1_20
19. Goto, M.A., Jagadeesan, R., Jeffrey, A., Pitcher, C., Riely, J.: An extensible approach to session polymorphism. *Math. Struct. Comput. Sci.* **26**(3), 465–509 (2016). <https://doi.org/10.1017/S0960129514000231>
20. Honsell, F., Miculan, M., Scagnetto, I.: pi-calculus in (Co)inductive-type theory. *Theor. Comput. Sci.* **253**(2), 239–285 (2001). [https://doi.org/10.1016/S0304-3975\(00\)00095-5](https://doi.org/10.1016/S0304-3975(00)00095-5)
21. Jung, R., Swasey, D., Sieczkowski, F., Svendsen, K., Turon, A., Birkedal, L., Dreyer, D.: Iris: Monoids and invariants as an orthogonal basis for concurrent reasoning. In: Rajamani, S.K., Walker, D. (eds.) *Symposium on Principles of Programming Languages, POPL 2015*. pp. 637–650. ACM (2015). <https://doi.org/10.1145/2676726.2676980>
22. Kobayashi, N.: *Type Systems for Concurrent Programs* (2007), <http://www.kb.ecei.tohoku.ac.jp/~koba/papers/tutorial-type-extended.pdf>
23. Kobayashi, N., Pierce, B.C., Turner, D.N.: Linearity and the Pi-Calculus. In: *Symposium on Principles of Programming Languages, POPL*. pp. 358–371. ACM Press (1996). <https://doi.org/10.1145/237721.237804>
24. Mackie, I.: Lilac: A Functional Programming Language Based on Linear Logic. *J. Funct. Program.* **4**(4), 395–433 (1994). <https://doi.org/10.1017/S0956796800001131>
25. Matsakis, N.D., II, F.S.K.: The rust language. In: *High integrity language technology, HILT*. pp. 103–104. ACM (2014). <https://doi.org/10.1145/2663171.2663188>
26. McBride, C.: I Got Plenty o’ Nuttin’. In: *A List of Successes That Can Change the World, Lecture Notes in Computer Science*, vol. 9600, pp. 207–233. Springer (2016). https://doi.org/10.1007/978-3-319-30936-1_12
27. Milner, R.: *Communicating and mobile systems - the Pi-calculus*. Cambridge University Press (1999)
28. Milner, R., Parrow, J., Walker, D.: A Calculus of Mobile Processes, Parts I and II. *Inf. Comput.* **100**(1) (1992). [https://doi.org/10.1016/0890-5401\(92\)90008-4](https://doi.org/10.1016/0890-5401(92)90008-4)
29. Orchard, D., Liepelt, V., III, H.E.: Quantitative program reasoning with graded modal types. *Proc. ACM Program. Lang.* **3**(ICFP), 110:1–110:30 (2019). <https://doi.org/10.1145/3341714>
30. Orchard, D.A., Yoshida, N.: Using session types as an effect system. In: Gay, S., Alglave, J. (eds.) *Programming Language Approaches to Concurrency- and Communication-centric Software, PLACES 2015. EPTCS*, vol. 203, pp. 1–13 (2015). <https://doi.org/10.4204/EPTCS.203.1>
31. Rouvoet, A., Poulsen, C.B., Krebbers, R., Visser, E.: Intrinsically-typed definitional interpreters for linear, session-typed languages. In: *Certified Programs and Proofs, CPP*. pp. 284–298. ACM (2020). <https://doi.org/10.1145/3372885.3373818>

32. Sangiorgi, D., Walker, D.: The Pi-Calculus - a theory of mobile processes. Cambridge University Press (2001)
33. Scalas, A., Dardha, O., Hu, R., Yoshida, N.: A Linear Decomposition of Multiparty Sessions for Safe Distributed Programming. In: European Conference on Object-Oriented Programming, ECOOP. LIPIcs, vol. 74, pp. 24:1–24:31. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2017). <https://doi.org/10.4230/LIPIcs.ECOOP.2017.24>
34. Thiemann, P.: Intrinsically-Typed Mechanized Semantics for Session Types pp. 19:1–19:15 (2019). <https://doi.org/10.1145/3354166.3354184>
35. Turon, A.J., Thamsborg, J., Ahmed, A., Birkedal, L., Dreyer, D.: Logical relations for fine-grained concurrency. In: Giacobazzi, R., Cousot, R. (eds.) Symposium on Principles of Programming Languages, POPL '13. pp. 343–356. ACM (2013). <https://doi.org/10.1145/2429069.2429111>
36. Wadler, P.: Linear Types can Change the World! In: Programming concepts and methods. p. 561. North-Holland (1990)
37. Zalakain, U., Dardha, O.: π with leftovers: a mechanisation in agda. CoRR **abs/2005.05902** (2020), <https://arxiv.org/abs/2005.05902>
38. Zalakain, U., Dardha, O.: Typing the Linear π -Calculus – Formalisation in Agda (2021), <https://github.com/umazalakain/typing-linear-pi>