

Non-Linear Communication via Graded Modal Session Types

Daniel Marshall^a, Dominic Orchard^{a,b}

^a*University of Kent, United Kingdom*

^b*University of Cambridge, United Kingdom*

Abstract

Session types provide guarantees about concurrent behaviour and can be understood through their correspondence with linear logic, with propositions as sessions and proofs as processes. However, a strictly linear setting is somewhat limiting since there exist various useful patterns of communication that rely on non-linear behaviours. For example, shared channels provide a way to repeatedly spawn a process with binary communication along a fresh linear session-typed channel. Non-linearity can be introduced in a controlled way in programming through the general concept of *graded modal types*, which are a framework encompassing various kinds of *coeffect* typing (describing how computations make demands on their context). This paper shows how graded modal types can be leveraged alongside session types to enable various well-known non-linear concurrency behaviours to be re-introduced in a precise manner in a type system with a linear basis. The ideas here are demonstrated using Granule, a functional programming language with linear, indexed, and graded modal types. We then define a core calculus capturing the requisite type features and our new graded modal session-typed primitives. We give an operational model and establish key properties following from the typing.

Keywords: session types, graded types, communication patterns

Email addresses: dm635@kent.ac.uk (Daniel Marshall), D.A.Orchard@kent.ac.uk (Dominic Orchard)

1. Introduction

Most programming languages ascribe a notion of *type* (dynamic or static) to data, classifying *what* data we are working with—integer, string, function, etc. Languages like Go and Rust among many others extend this to the realm of communication between two participants via *typed channels*, where any process with access to a channel endpoint can send and receive values of a given type across it. These typed channels only give us information about the nature of the data being communicated, however; to go beyond this we must consider *behavioural types*, which capture not just *what* data is, but *how* it is calculated.

One such behavioural type system for concurrency is *session types* [16, 42], adding more structure to typed channels by incorporating a protocol that describes the precise behaviour of the processes communicating over the channel. As an example of such a protocol, consider the standard example (with some added flavour) of a customer who wishes to purchase a cake from a bakery. In natural language, we could describe the protocol as follows: the customer names their desired cake (a string), the bakery provides the price (an integer), and then the customer decides if they are happy with said price. If so, they confirm delivery details (another string), but if not, they end the communication immediately. From the customer’s perspective, this protocol can be described as a session type, which would look something like:

```
Send String (Recv Int (Select (Send String End) End))
```

The particular constructs at play here will be explained in more depth later, but it is clear to see that this allows more information about the communication to be encoded in the type than simply the kinds of data that are being sent and received. We will also see later how the session type for the bakery can be understood as the *dual* to that of the customer, which helps to enforce certain communication properties such as deadlock freedom.

Session types naturally fit into the more general substructural discipline of *linear types* [14, 38, 40]. By treating data as a *resource* which must be used exactly once, linear types can capture various kinds of stateful protocols of interaction. Session types are inherently substructural: channels cannot in general be arbitrarily duplicated or discarded if they are to be used according to a precise sequence of operations (the protocol). This idea has allowed logical foundations to be developed for session types based on Curry-Howard correspondences with both intuitionistic [7, 35] and classical

linear logic [39, 20]. Furthermore, linearly-typed functional languages then provide an excellent basis for session-typed programming [12]; session types can be represented in the linear lambda calculus via an additional interface for channels without any other changes to the type system [12, 21].

Modern substructural type systems, however, allow us to go beyond the notion of linearity, classifying usage into more than just linear or non-linear. This idea originates from Bounded Linear Logic (BLL), which generalises the ! modality to a family of modalities indexed by a polynomial expressing an upper bound on usage [15]. For example, $!_{x \leq 2} A$ describes a proposition A which can be used at most twice.

From various directions this notion of BLL has been generalised further to develop *coeffect types* [6, 13, 28], which are a framework for capturing different analyses of a program’s use of the context, all within a single type system. *Graded modal types* [27] unify coeffects with their dual notion of *effects*, and are an expressive system allowing for the specification and verification of many behavioural properties of programs. In this paper, we demonstrate that combining session types with graded modal types (particularly for capturing coeffects) allows for various well-known non-linear behaviours of concurrent programs to be reintroduced in a linear setting in a controlled and precise way via the facilities of graded modalities.

A key part of this work is to reconcile the tension between three competing requirements: (1) side effects inherent in communication primitives, (2) non-linearity, and (3) the call-by-value semantics one might expect in most programming languages (in contrast to call-by-name, which is the basis of most theoretical explanations of linear and graded type systems). Requirements (1) and (2) can be satisfied more easily in a call-by-name setting, but (3) (CBV semantics) introduces unsoundness. Whilst a call-by-name semantics for a modal treatment of non-linearity captures the idea of a number of ‘copies’ of a value, a call-by-value semantics unifies these copies into just one value. In the context of concurrency, then, a non-linear channel in a CBV semantics points to a single channel and so multiple uses within a program can lead to deadlock or race-conditions. Our solution revolves around preventing resource allocating primitives (those that create channels) from being promoted to non-linear contexts, and representing multiplicity of channels through data structures. Section 4 describes in more detail the problem and our solution.

Granule is a research programming language built upon the core of the linear λ -calculus, extended with *graded modalities* as described above. A

previous iteration of session types for Granule was described by Orchard et al. [27]. However, the system described required a monadic interface to avoid some issues caused by CBV; we solve these problems here to provide a more general and powerful interface for session-typed programming, and also demonstrate the broad range of possible applications for this interface by introducing a suite of primitives which capture the non-linear behaviours of reuse, replication, and repetition (providing *multicasting*).

We show three common patterns and how they can be described by combining linear session types and graded modal types: *reusable channels* (Section 6), *replicated servers* (Section 7), and *multicast communication* (Section 8). All of these patterns are a variation on the “fork” primitive, with some amount of substructural behaviour introduced via graded modal types, and sometimes restrictions to protocols via *type predicates*. All of the ideas presented here are implemented in the latest versions of the Granule compiler; interested readers who wish to experiment for themselves should feel free to download and install the most recent release.¹ We develop the operational semantics of our approach, via a core calculus (Section 5), and establish that type safety is preserved (which encompasses communication safety and deadlock freedom) in the presence of the new non-linear concurrency primitives (Section 9).

Changelog. The present paper is an extended version of work appearing at the PLACES 2022 workshop, with the more voluminous title of *Replicate, Reuse, Repeat: Capturing Non-Linear Communication via Session Types and Graded Modal Types* [22]. The present paper greatly expands the technical development, introducing a formal calculus and its operational semantics for both the core language and the concurrency primitives introduced here. In this version, we prove that session types combined with graded modal types satisfy syntactic type safety. Furthermore, we have expanded the number of primitives considered, added more examples and fixed a number of technical errors from the original paper.

¹At time of writing, the latest build is available at <https://github.com/granule-project/granule/releases/tag/v0.9.0.1>.

2. Granule and Graded Modal Types Primer

Granule’s type system is based on the linear λ -calculus augmented with *graded modal types* [27] as well as polymorphic and indexed types (in the form of user-definable GADTs). With linear typing as the basis, we cannot write functions that discard or duplicate their inputs as in a standard functional programming language. However, we can introduce non-linearity via graded modalities and use these to represent such functions, exemplified by the following Granule functions:

```
1 drop :  $\forall \{a : \text{Type}\} . a [0] \rightarrow ()$ 
2 drop [x] = ()
3
4 copy :  $\forall \{a : \text{Type}\} . a [2] \rightarrow (a, a)$ 
5 copy [x] = (x, x)
```

The function arrow can be read as the type of linear functions (which consume their input exactly once; note that these are generally notated \multimap , but in Granule functions are linear by default, so we simply use \rightarrow), but a $[r]$ is a graded modal type capturing the capability to use the value ‘inside’ in a non-linear way as described by r , i.e., `drop` uses the value 0 times and `copy` uses it 2 times. The pattern match `[x]` on the left-hand side of each function eliminates the graded modality, binding `x` as a non-linear variable.

The central idea of graded modal types is to capture program structure via an indexed family of modalities where the indices have some algebraic structure which gives an abstract view of program structure. We focus on *semiring graded necessity* in this paper (written $\Box_r A$ in mathematical notation but `A [r]` in Granule) which generalises linear logic’s $!$ [14] and Bounded Linear Logic [15]. The semiring structure is also equipped with a preorder. The above example uses the natural number graded modality with equality as its preorder, thus providing and analysis which counts exactly how many times a value can be used.

Another useful graded modality has grades drawn from a semiring of *intervals* which allows us to give bounds on how a value might be used. To demonstrate, we define the classic `fromMaybe` function which allows a value that may or may not exist to be retrieved from a `Maybe` type.

```
1 data Maybe a = Just a | Nothing
2
3 fromMaybe :  $\forall \{a : \text{Type}\} . a [0..1] \rightarrow \text{Maybe } a \rightarrow a$ 
4 fromMaybe [_] (Just x) = x;
5 fromMaybe [d] Nothing = d
```

Note that without the graded modality, this function would be ill-typed in Granule, since values are linear by default and one of the cases requires discarding the default value (given by the first parameter). By giving this parameter a type of a `[0..1]`, we specify that it can be used *either* 0 times or 1 time (in other words, this value has *affine* behaviour rather than linear). There is only one total function in Granule which inhabits the type we give to `fromMaybe` here—linearity forbids defining an instance which always returns the default value.

In order to use `fromMaybe` we need to ‘promote’ its first input to be a graded modal value, which is written by wrapping a value in brackets, e.g., `fromMaybe [42]`. Promotion propagates any requirements implied by the graded modality to the free variables. For example, the following takes an input and shares the capabilities implied by its grade to two uses of `fromMaybe`:

```

1 fromMaybeIntPlus : Int [0..2] → Maybe Int → Maybe Int → Int
2 fromMaybeIntPlus [d] x y = fromMaybe [d] x + fromMaybe [d] y

```

The `0..1` usage implied by the first `fromMaybe [d]` and `0..1` usage by the second are added together to get the requirement that the incoming integer `d` is graded as `0..2`.

Lastly, Granule includes *indexed* types which offer a lightweight form of dependency, allowing for type-level access to information about data. For example, length-indexed vectors can be defined and used as follows:

```

1 data Vec (n : Nat) (a : Type) where
2   Nil  : Vec 0 a;
3   Cons : a → Vec n a → Vec (n + 1) a
4
5 append : ∀ {a : Type, n m : Nat} . Vec n a → Vec m a → Vec (n + m) a
6 append Nil ys          = ys;
7 append (Cons x xs) ys = Cons x (append xs ys)

```

Indexed types allow us to ensure at the type-level that when we append two vectors, the length of the output is equal to the sum of the lengths of the two inputs. Again, thanks to linearity, we gain further assurances from our type signatures—here we can guarantee that since we must use every element of the input vectors, these must also appear in the output, and so no information is being lost along the way.

Indexed natural numbers can be similarly defined to vectors, and are used later to provide lightweight type-dependence on natural number inputs:

```

1 data N (n : Nat) where
2   Z : N 0;

```

```
3 S : ∀ {n : Nat} . N n → N (n+1)
```

This definition gives a way to capture a runtime witness for a type level `Nat`.

Granule allows type indices to be used also as grades, thus we can define for example the functorial action on vectors which links the size of the vector with the usage of the function to be applied to each element:

```
1 vmap : ∀ {a b : Type, n : Nat} . (a → b) [n] → Vec n a → Vec n b
2 vmap [_] Nil = Nil;
3 vmap [f] (Cons x xs) = Cons (f x) (vmap [f] xs)
```

In Granule, predicates on types can be represented similarly to Haskell’s type class constraints with type signatures of the form:

```
functionName : ∀{typeVariables} . {constraints} ⇒ type
```

For example, we can define a much more general version of `copy` as:

```
1 copy : ∀ {a : Type, s : Semiring, r : s}
2       . {{{(1 + 1) : s} ≤ r} ⇒ a [r] → (a, a)
3 copy [x] = (x, x)
```

This is polymorphic on the semiring `s`, as well as including a constraint as a predicate that the grade `r` is approximated by `1 + 1` in this particular semiring. Types appearing in this constraint position have kind `Predicate`. We later introduce a number of type functions, returning a type of kind `Predicate`, in order to restrict protocols in various settings.

3. Core calculus, and linear session-typed channels

In order to later establish key properties of our approach, we replay the details of Granule’s core calculus from the work of Orchard et al. [27], which gives a linear type theory with graded modal types. We extend this core calculus with our session-typed concurrency primitives. We outline the syntax and static semantics (type system) in this section. Section 5 defines an operational semantics that incorporates the notion of processes and channels which was not previously given for Granule’s core calculus.

The core type theory we will focus on extends the linear λ -calculus with products, a multiplicative unit, and a *semiring-graded necessity modality* $\square_r A$, where for a pre-ordered semiring $(\mathcal{R}, *, 1, +, 0, \sqsubseteq)$ there exists a family of types $\{\square_r A\}_{r \in \mathcal{R}}$. Being a pre-ordered semiring requires that $+$ and $*$ must be monotone with respect to \sqsubseteq (but no other requirements beyond

the semiring axioms). This language represents a simplified monomorphic subset of Granule [27], but closely resembles other graded systems from the literature [9, 1, 4, 5, 41]; the session type system presented in this paper should also be compatible with these, to varying extents depending on which kinds of grading the system in question accommodates.

Syntax. Our syntax is that of the linear λ -calculus with multiplicative products and unit (first line of syntax below) with additional terms for introducing and eliminating the \square modality (second line):

$$t ::= x \mid \lambda x.t \mid t_1 t_2 \mid (t_1, t_2) \mid \mathbf{let} (x, y) = t_1 \mathbf{in} t_2 \mid \mathbf{unit} \mid \mathbf{let} \mathbf{unit} = t_1 \mathbf{in} t_2 \\ \mid [t] \mid \mathbf{let} [x] = t_1 \mathbf{in} t_2 \quad (\text{terms})$$

Following the syntax of variables, we group terms above into pairs of introduction and elimination forms, for functions, tensors, units, and the graded modal respectively. The meaning of these terms is explained in the next subsection with reference to their typing.

3.1. Type System

Typing judgements are of the form $\Gamma \vdash t : A$, assigning type A to a term t under context Γ . Types A are defined:

$$A, B ::= A \multimap B \mid A \otimes B \mid \mathbf{unit} \mid \square_r A \quad (\text{types})$$

Hence, our type syntax comprises linear function types $A \multimap B$, linear multiplicative products $A \otimes B$, a linear multiplicative unit \mathbf{unit} , and the graded modality $\square_r A$.

Contexts Γ contain both linear assumptions $x : A$ and graded assumptions $x : [A]_r$ which have originated from inside a graded modality. Syntax and typing are then given by the following rules:

$$\frac{}{x : A \vdash x : A} \text{VAR} \quad \frac{\Gamma, x : A \vdash t : B}{\Gamma \vdash \lambda x.t : A \multimap B} \multimap_i \quad \frac{\Gamma_1 \vdash t_1 : A \multimap B \quad \Gamma_2 \vdash t_2 : A}{\Gamma_1 + \Gamma_2 \vdash t_1 t_2 : B} \multimap_e$$

The VAR, \multimap_i , and \multimap_e rules are the standard rules of the linear λ -calculus, augmented with a notion of *contraction* captured by the $+$ operation on contexts coming from multiple sub-terms. This operation is only defined when contexts are disjoint with respect to linear assumptions; on overlapping

graded assumptions we add their grades, e.g. $(\Gamma_1, x : [A]_r) + (\Gamma_2, x : [A]_s) = (\Gamma_1 + \Gamma_2), x : [A]_{r+s}$. More explicitly, context addition is specified as follows:

$$\begin{aligned} (\Gamma, x : A) + \Gamma' &= (\Gamma + \Gamma'), x : A && \text{iff } x \notin |\Gamma'| && \emptyset + \Gamma = \Gamma \\ \Gamma + (\Gamma', x : A) &= (\Gamma + \Gamma'), x : A && \text{iff } x \notin |\Gamma| && \Gamma + \emptyset = \Gamma \\ (\Gamma, x : [A]_r) + (\Gamma', x : [A]_s) &= (\Gamma + \Gamma'), x : [A]_{(r+s)} \end{aligned}$$

Note that this is a declarative specification of $+$ rather than an algorithmic definition, since graded assumptions for the same variable may appear in different positions within the two contexts.

The rules involving graded modalities are perhaps more interesting:

$$\begin{array}{c} \frac{\Gamma \vdash t : A}{r * \Gamma \vdash [t] : \Box_r A} \text{ PR} \quad \frac{\Gamma_1 \vdash t_1 : \Box_r A \quad \Gamma_2, x : [A]_r \vdash t_2 : B}{\Gamma_1 + \Gamma_2 \vdash \mathbf{let} [x] = t_1 \mathbf{in} t_2 : B} \text{ ELIM} \\ \\ \frac{\Gamma, x : A \vdash t : B}{\Gamma, x : [A]_1 \vdash t : B} \text{ DER} \quad \frac{\Gamma \vdash t : B}{\Gamma, 0 * \Gamma' \vdash t : B} \text{ WEAK} \quad \frac{\Gamma, x : [A]_r, \Gamma' \vdash t : B \quad r \sqsubseteq s}{\Gamma, x : [A]_s, \Gamma' \vdash t : B} \text{ APPROX} \end{array}$$

The PR rule (promotion) introduces a graded modality with grade r and thus must scale by r all of the input assumptions Γ , where scaling is a partial operation defined:

$$r * \emptyset = \emptyset \quad r * (\Gamma, x : [A]_s) = (r * \Gamma), x : [A]_{r*s}$$

Thus, $r * \Gamma$ is undefined if Γ contains any linear assumptions and therefore implicitly in the PR rule we require that Γ contains only graded assumptions. We later denote a context which contains only graded variables as $[\Gamma]$.

The ELIM rule eliminates a graded modality and captures the idea that a requirement for x to be used in an r -like way in t_2 can be matched with the capability of t_1 described by its graded modal type. In Granule, this construct is folded into pattern matching (seen above): we can ‘unpack’ (eliminate) a graded modality via pattern matching to provide a non-linear (graded) variable in the body of the function (the analogue to t_2 in this rule).

The DER rule (‘dereliction’) connects linear typing to graded typing, showing that a requirement for a linear assumption is satisfied by an assumption graded by 1. Note that it is not the case that $A \cong [A]_1$ since dereliction only converts from a graded assumption to linear assumption; there is no inverse. Meanwhile, WEAK explains how we can *weaken* with variables graded by 0.

The last important rule is the APPROX rule which provides grade approximation; this allows a grade r to be converted to another grade s , provided that s *approximates* r where \sqsubseteq is the pre-order of the semiring in question.

Example 1. Some useful choices for the pre-ordered semiring (all of which are provided in Granule) include:

- The semiring of natural numbers, which we define as $(\mathbb{N}, \times, 1, +, 0, \equiv)$. Note that the *discrete ordering* \equiv here means that this particular semiring has no approximation, and so we are tracking the exact number of times a term has been used. We could instead use the standard \leq ordering on natural numbers which would then permit approximation; this would allow for an *upper bound* on a term's usage.
- A semiring of intervals over \mathbb{N} [27] (seen earlier), with the standard interval arithmetic definitions for $+$ and $*$, captures lower- and upper-bound usage, with $(l_1..u_1) \sqsubseteq (l_2..u_2)$ if $l_2 \leq l_1$ and $u_1 \leq u_2$.
- Intervals over the extended natural number semiring, with $\mathbb{N} \cup \{\infty\}$; these are additionally useful since constructing an interval with bounds $0..\infty$ allows us to capture a potentially arbitrary amount of use.
- The ‘none-one-tens’ semiring [25] with $\mathcal{R} = \{0, 1, \omega\}$ (in Granule, these are written `Zero`, `One`, `Many`) for tracking linear vs. non-linear usage, but with zero use identified as separate to non-linearity to capture irrelevance.
- The semiring of security levels [11], with $\mathcal{R} = \{\text{Lo}, \text{Hi}\}$ such that $\text{Lo} \sqsubseteq \text{Hi}$ and $1 = \text{Lo}$ and $0 = \text{Hi}$ with $*$ = \sqcap and $+$ = \sqcup .

Finally, we have the rules for introducing and eliminating tensor products and the multiplicative unit, which are standard, though it is of course important to remember that products are *linear* and so it is not possible to discard either side: both elements must be used when consuming a product.

$$\frac{\Gamma_1 \vdash t_1 : A \quad \Gamma_2 \vdash t_2 : B}{\Gamma_1 + \Gamma_2 \vdash (t_1, t_2) : A \otimes B} \otimes_i \quad \frac{\Gamma_1 \vdash t_1 : A \otimes A' \quad \Gamma_2, x : A, y : A' \vdash t_2 : B}{\Gamma_1 + \Gamma_2 \vdash \mathbf{let} (x, y) = t_1 \mathbf{in} t_2 : B} \otimes_e$$

$$\frac{}{\emptyset \vdash \mathbf{unit} : 1} \text{UNIT}_i \quad \frac{\Gamma_1 \vdash t_1 : 1 \quad \Gamma_2 \vdash t_2 : B}{\Gamma_1 + \Gamma_2 \vdash \mathbf{let} \mathbf{unit} = t_1 \mathbf{in} t_2 : B} \text{UNIT}_e$$

The following gives an example that combines together several of the above rules to show how a value of type $\Box_3 A$ presents the capability to use an A value three times, which is split into forming a pair with a single usage of A and the formation of a further graded modality:

$$\frac{\frac{\frac{}{x : A \vdash x : A} \text{VAR}}{x : [A]_1 \vdash x : A} \text{DER} \quad \frac{\frac{\frac{}{x : A \vdash x : A} \text{VAR}}{x : [A]_1 \vdash x : A} \text{DER}}{x : [A]_2 \vdash [x] : \Box_2 A} \text{PR}}{x : [A]_3 \vdash (x, [x]) : A \otimes \Box_2 A} \otimes_i \quad \frac{}{y : \Box_3 A \vdash y : \Box_3 A} \text{ELIM}}{y : \Box_3 A \vdash \mathbf{let} [x] = y \mathbf{in} (x, [x]) : A \otimes \Box_2 A} \text{ELIM}$$

Note that $\Box_r A$ has the structure of a *graded exponential comonad* [11] meaning it has the following operations, all of which can be derived from the typing rules here: $\delta : \Box_{(r*s)} A \multimap \Box_r(\Box_s A)$ and $\varepsilon : \Box_1 A \multimap A$, and $\mathbf{contr} : \Box_{(r+s)} A \multimap (\Box_r A) \otimes (\Box_s A)$ and $\mathbf{weak} : \Box_0 A \multimap 1$.

3.2. Linear Session Types in Granule

The implementation of session types in Granule is based on the GV calculus (originating from Gay and Vasconcelos [12], further developed by Wadler [39], for which we use Lindley and Morris' formulation [21]). The GV system extends the linear λ -calculus with a data type of channels $\mathbf{Chan}(S)$ parameterised by session types S [42], which capture the protocol of interaction allowed over the channel. Communication in GV is usually asynchronous: sending always succeeds, but receiving can block as usual, though Lindley and Morris consider a synchronous semantics [21].

Starting with a simple subset, session types can describe channels which send or receive a value of type T and then can be used according to session type S written $!T.S$ or $?T.S$ respectively, or can be closed written $\mathbf{end}_!$ or $\mathbf{end}_?$. There are then functions for sending or receiving a value on a channel, forking a process at one end of a channel returning the other, and waiting for a channel to be closed:

$$\begin{array}{ll} \mathbf{send} & : T \otimes \mathbf{Chan}(!T.S) \multimap \mathbf{Chan}(S) & \mathbf{fork} & : (\mathbf{Chan}(S) \multimap \mathbf{Chan}(\mathbf{end}_!)) \multimap \mathbf{Chan}(\overline{S}) \\ \mathbf{recv} & : \mathbf{Chan}(?T.S) \multimap T \otimes \mathbf{Chan}(S) & \mathbf{wait} & : (\mathbf{Chan}(\mathbf{end}_?) \multimap 1) \end{array}$$

where \overline{S} is the *dual* session type to S , defined $\overline{!T.S} = ?T.\overline{S}$, $\overline{?T.S} = !T.\overline{S}$, $\overline{\mathbf{end}_?} = \mathbf{end}_!$ and $\overline{\mathbf{end}_!} = \mathbf{end}_?$. The \mathbf{fork} combinator leverages the duality

operation, spawning a process which applies the parameter function with a fresh channel, thus returning the dual endpoint.

The core interface in Granule correspondingly has operations with the following types, where `Protocol` is the kind of session types whose constructors and associated operations we highlight in purple:

```

send      : ∀ {a : Type, p : Protocol} . LChan (Send a p) → a → LChan p
recv     : ∀ {a : Type, p : Protocol} . LChan (Recv a p) → (a, LChan p)
forkLinear : ∀ {p : Protocol}          . (LChan p → ()) → LChan (Dual p)
close    : LChan End → ()

```

The type-level function `Dual : Protocol → Protocol` provides the standard notion of duality for session types, defined:

```

Dual End = End
Dual (Send a p) = Recv a (Dual p)
Dual (Recv a p) = Send a (Dual p)
Dual (Select p1 p2) = Offer (Dual p1) (Dual p2)
Dual (Offer p1 p2) = Select (Dual p1) (Dual p2)

```

We can define a parallel composition combinator on thunked computations (i.e., of type `() → a`) using `forkLinear` to run one computation in parallel and then return its result via the associated channel:

```

1 -- `par` expects two thunks, `f` and `g`
2 -- and returns their resulting values in a pair
3 par : ∀ {a : Type, b : Type} . (() → a) → (() → b) → (a, b)
4 par f g = let
5     c      = forkLinear (λc' → close (send c' (f ()))));
6     b      = g ();    -- compute g in current thread
7     (a, c') = recv c; -- get result of f from other thread
8     ()     = close c'
9     in (a, b)

```

Choice. We also provide some utility operations for internal and external choice (also known as selection and offering); note that as described in the original formulation by Lindley and Morris [21] these can be defined in terms of the core functions given above, and so when it comes to our semantics we will not need to define explicit reductions for or prove properties relating to these operations, but we do provide them as primitives in the implementation for ease of use.

```

selectLeft  :  $\forall \{p1\ p2 : \text{Protocol}\} . \text{LChan } (\text{Select } p1\ p2) \rightarrow \text{LChan } p1$ 
selectRight :  $\forall \{p1\ p2 : \text{Protocol}\} . \text{LChan } (\text{Select } p1\ p2) \rightarrow \text{LChan } p2$ 
offer      :  $\forall \{p1\ p2 : \text{Protocol}, a : \text{Type}\}$ 
             .  $(\text{LChan } p1 \rightarrow a) \rightarrow (\text{LChan } p2 \rightarrow a) \rightarrow \text{LChan } (\text{Offer } p1\ p2) \rightarrow a$ 

```

The following gives a brief example putting all these primitives together:

```

1 server : LChan (Offer (Recv Int End) (Recv () End)) → Int [0..1] → Int
2 server c [d] = offer (λc → let (x, c) = recv c in let () = close c in x)
3               (λc → let ((), c) = recv c in let () = close c in d) c
4
5 client :  $\forall \{p : \text{Protocol}\} . \text{LChan } (\text{Select } (\text{Send } \text{Int } \text{End})\ p) \rightarrow ()$ 
6 client c = let c = selectLeft c;
7           c = send c 42
8           in close c
9
10 example : Int -- Evaluates to 42
11 example = server (forkLinear client) [100]

```

The server offers a choice between being sent an integer or a unit value. The second parameter (bound to `d`) is used as a default value in the case that a unit value is received by the server, where `Int [0..1]` denotes that this integer can be used 0 or 1 times (see Section 2 for more explanation of this grading). The client selects the left behaviour, sends an integer, then closes its side of the communication.

The final definition `example` spawns the client with a channel and connects the dual end of the channel to the server which returns the received value of 42 here.

Core calculus representation. In the core calculus we add a type of channels indexed by session types (protocols), with grammar:

$$\begin{aligned}
A &:= \dots \mid \mathbf{Chan } P && \text{(types, extended)} \\
P &:= \mathbf{Send } A P \mid \mathbf{Recv } A P \mid \mathbf{End} && \text{(session types)}
\end{aligned}$$

with the duality function \overline{P} on session types defined as expected:

$$\overline{(\mathbf{Send } A P)} = \mathbf{Recv } A \overline{P} \quad \overline{(\mathbf{Recv } A P)} = \mathbf{Send } A \overline{P} \quad \overline{\mathbf{End}} = \mathbf{End}$$

As mentioned above, we elide the offer and select types from the core calculus and the theoretical development in this paper, since they can also be emulated using send and receive and coproducts as in the FST presentation of GV [21].

The term language is extended with the following primitives and their typing:

$$\begin{aligned}
\text{send} &: A \otimes \mathbf{Chan} (\mathbf{Send} A P) \multimap \mathbf{Chan} P \\
\text{receive} &: \mathbf{Chan} (\mathbf{Recv} A P) \multimap A \otimes \mathbf{Chan} P \\
\text{close} &: \mathbf{Chan} \mathbf{End} \multimap 1 \\
\text{forkLinear} &: (\mathbf{Chan} P \multimap 1) \multimap \mathbf{Chan} \bar{P}
\end{aligned}$$

We assume asynchronous communication, which is captured in the operational model of Section 5.

4. The Relationship Between Grading, Call-by-Value, and Effects

The type system described thus far is based heavily on that of Orchard et al. [27], though with some simplifications where we elide concepts that are irrelevant to this paper. Orchard et al. also described an earlier version of session types in Granule based somewhat on the GV calculus, but the session types were not included in the language’s formal semantics at the time, which resulted in some cases where combining grading and session types led to unsoundness and the possible introduction of deadlocks. This work resolves these issues and we will present a full operational semantics in Section 5 as part of our formal theory of graded session types, but before we do this we first discuss the problems present in the original system and describe our solution.

The first and primary problem becomes apparent when promoting a primitive function which creates new channels. Consider the following example which is allowed by the type system described so far:

```

1 problematic : Int
2 problematic =
3   let [c] : ((LChan (Recv Int End)) [2]) = [forkLinear (λc → close (send c 42))];
4     (n, c') = recv c; () = close c';
5     (m, c') = recv c; () = close c'
6   in (n + m)

```

On line 3, the program forks a process that sends 42 on a channel, but under a promotion, with the type explaining that we want to use the resulting value twice (given by the explicit type signature here). This promotion then allows two uses of the channel on lines 4-5.

The typical semantics for effect-based calculi in the literature is call-by-name [6, 11, 28, 1]. Under a call-by-name semantics, which Granule allows via the extension `language CBN`, this program executes and produces the expected result of 84. The key is that call-by-name reduction substitutes the call to `forkLinear` into the two uses of variable `c` on lines 4 and 5, and thus we are receiving from two different channels. However, under the default call-by-value semantics (which was chosen in Granule for simplicity and to avoid complications resulting from effects), line 3 is fully evaluated (underneath the graded modal introduction), and so variable `c` on lines 4 and 5 refers to a single channel. This means that executing this program blocks indefinitely on line 5; we get an error from the underlying implementation’s concurrency primitives (written in Haskell) describing a `thread blocked indefinitely in an MVar operation`. This comes from Granule’s runtime which leverages these standard primitives.

In order to get around this problem with non-linear channels in a call-by-value setting, Orchard et al. [27] instead deal with channels monadically, so here we would end up with a channel of type `(Session (LChan ...)) [2]` (with the `Session` monad).² To make this example well-defined would then require a distributive law which maps from this to `Session ((LChan ...) [2])`, copying the channel. Providing such a distributive law is unsound—it would enable problematic and thus indefinite blocking—and fortunately it is also *not* derivable. However, we wish to avoid the monadic programming style as it is not required when working with just linear channels.

Our alternative solution is that we instead syntactically restrict promotion to terms which do not allocate resources, i.e., whose reduction does not create linear channels. In particular this includes reducible expressions containing the `forkLinear` primitive which causes the difficulty here. Non-linear channels are then re-introduced by additional primitives in Sections 6, 7 and 8 which allow for precise and carefully managed non-linear usage, allowing for grading to be combined with linear channels even in a call-by-value setting and without the extra overhead that was required when every channel had to be wrapped in the `Session` monad.

²In Granule such a channel would in fact have the type `(LChan ... <Session>) [2]`, with `<Session>` being an instance of the graded monadic `<r>` modalities which are dual to the graded comonadic `[r]` modalities. We elide discussing these further here since we will not need any monad other than `Session`.

Thus, we modify the promotion rule to be:

$$\frac{\Gamma \vdash t : A \quad \neg \text{resourceAllocator}(t)}{r * \Gamma \vdash [t] : \square_r A} \text{ PR}'$$

where $\text{resourceAllocator}(t)$ classifies those terms whose reduction to a value would trigger creation of a resource; such terms contain a use of the `forkLinear` primitive (or the other fork primitives introduced in later sections) in a reducible position. Note in particular that $\neg \text{resourceAllocator}(\lambda x. \text{forkLinear } t)$ since reduction does not happen underneath a λ -abstraction (Section 5 defines the reduction semantics). Definition 5 in Appendix A gives the formal definition. Generally, any term whose reduction semantics creates channels is classified as a resource allocator.

Promotion is not the only source of unsoundness in the session type system of Orchard et al., [27] however. Some additional restrictions on the primitives which combine grading with linear channels are required, in order to prevent possible problems that can occur when two ends of the same channel reach inconsistent communication states or are used at different grades. We will discuss these problems further in Sections 6, 7 and 8 where we will go on to define how exactly these new primitives must operate in order to properly preserve communication safety.

Linear Haskell. Note that since we developed this approach, other work has resolved the same problem relating to promotion through different techniques. Originally, the linear types extension to Haskell [4] got around this difficulty by only ever allocating new resources (such as linear channels or mutable arrays) inside a continuation, which is passed around at every step until deallocation to prevent the allocator itself from ever being promoted.

More recently, this strategy of requiring specialised restrictions on the interface for each resource allocator has been generalised by introducing a “linear constraint” [33] called `Linearly`. This constraint is assumed whenever a new resource is allocated, and since the assumption must be used exactly once this prevents the promotion of a linear resource constructor in much the same way. A continuation is still necessary when `Linearly` is initially assumed, but unlike before this same qualification can now be used generically across many types of resource.

5. Semantics

We define an operational model for our linear session types here, in order to make the connection between our type system and the underlying communication behaviours more concrete, and so that we can carefully prove that properties such as deadlock freedom will continue to hold after our extensions with additional primitives in Sections 6, 7 and 8. We give a semantics to terms in the context of runtime “configurations” of processes and channels. This allows us to show that deadlock freedom holds by proving that progress can always be made regardless of the current state of the global configuration, and proving preservation of typing of the overall configuration.

5.1. Operational Semantics

The semantics is call-by-value and at its core is based on the standard small-step operational semantics for the linear λ -calculus and its extension to graded modal types in a call-by-value setting given by Orchard et al. [27]. We extend this with the processes and channels.

Processes. A *process configuration* tracks the state of a number of processes, comprising ordinary terms, running in parallel:

$$\begin{aligned} \mathcal{P} &::= \cdot \mid (\mathcal{P} \mid t^{\mathcal{R}}) && \text{(process configurations)} \\ \mathcal{R} &::= \circ \mid \bullet && \text{(role)} \end{aligned}$$

(the parentheses are typically omitted). The role \circ marks a subordinate process and \bullet marks the main process. Typing of process configurations later establishes that a well-typed process configuration has exactly one main process (Section 5.2).

Configurations are subject to an equivalence relation capturing commutativity and associativity of the parallel composition in process configurations:

$$\mathcal{P}_1 \mid \mathcal{P}_2 \equiv \mathcal{P}_2 \mid \mathcal{P}_1 \quad (\mathcal{P}_1 \mid \mathcal{P}_2) \mid \mathcal{P}_3 \equiv \mathcal{P}_1 \mid (\mathcal{P}_2 \mid \mathcal{P}_3)$$

In subsequent relations (e.g., reduction) we work up-to process equivalence, i.e., we can freely apply the equivalence relation to rearrange process configurations. This relation is akin to *structural congruence* in process algebras.

Channels. Communication channels in the operational semantics are represented by pairs of dual *channel endpoints* to allow asynchronous duplex communication. Throughout, c ranges over channel endpoint names, which appear in ‘positive’ and ‘negative’ forms:

$$c ::= n^+ \mid n^- \quad (\text{channel endpoint names})$$

The *dual* of an endpoint name c is another endpoint name, denoted \bar{c} , where $\overline{n^+} = n^-$ and $\overline{n^-} = n^+$ and thus duality is an involution, i.e., $\overline{\bar{c}} \equiv c$.

Channel configurations record the contents of channel endpoints in the current environment by associating names with a queue \vec{v} of message values:

$$\mathcal{C} ::= \cdot \mid \mathcal{C}, c = \vec{v} \quad (\text{channel configurations})$$

We write $\langle \rangle$ for an empty queue, $\vec{v} : v$ for a queue with value v enqueued on the end, and $v : \vec{v}$ for a queue with v on the front (ready to be dequeued).

We will sometimes refer to c as simply a channel, and also as a reference to its contents (the queue) within a configuration.

Reduction. Single-step global reductions are then defined over pairs of process and channel configurations (which we refer to as *global configurations*) with reductions of the form:

$$(\mathcal{P}; \mathcal{C}) \rightsquigarrow (\mathcal{P}'; \mathcal{C}')$$

where $\mathcal{P}, \mathcal{P}'$ are (incoming and outgoing) process configurations and $\mathcal{C}, \mathcal{C}'$ are (incoming and outgoing) channel configurations respectively. The channel configuration may be updated if reduction involves applying a concurrency primitive, which could involve spawning two endpoints of a new channel or communicating a value from one end of a channel to the other.

Channels inside a channel configuration can either be *linear* or they can be *shared*, which means that rather than being used linearly they are used according to some grade r . We will discuss this in more detail when we describe primitives that can spawn graded channels, starting from Section 6.

The standard call-by-value small-step semantics for the linear λ -calculus embeds within this more complex global configuration semantics through the use of the following rule:

$$\frac{t \rightsquigarrow t'}{(\mathcal{P} \mid t^{\mathcal{R}}; \mathcal{C}) \rightsquigarrow (\mathcal{P} \mid t'^{\mathcal{R}}; \mathcal{C})} \text{REDUCE}$$

which states that if the term t can reduce to t' (without any need to interact with the global context) then we can make that same reduction within the global configuration, at any role, whilst threading all of the extra channel configuration information through so that it can come back into play once we reach a communication step. Going forwards, we elide the role in places where a reduction rule is agnostic to the role, assuming that the role of a process being reduced is preserved in its corresponding location in the output process configuration.

We will now briefly discuss how these standard reduction steps work in practice, to provide a foundation from which we can go on to explain the behaviour of our core session type primitives in this setting of process and channel configurations. We will explain some rules in detail and elide others which are very similar to those already introduced. All reduction rules for the entire semantics are collected together in Appendix A.

The core reduction rules for the (CBV) linear λ -calculus terms are of course β -reduction and function application, which are given as follows.

$$\frac{t_1 \rightsquigarrow t'_1}{t_1 t_2 \rightsquigarrow t'_1 t_2} \rightsquigarrow_{\text{APP}} \quad \frac{t \rightsquigarrow t'}{v t \rightsquigarrow v t'} \rightsquigarrow_{\text{prim}} \quad \overline{(\lambda x.t) v \rightsquigarrow t[v/x]} \rightsquigarrow_{\beta}$$

Our semantics is call-by-value, and left-to-right: we fully reduce the left-hand side of a function application to a value and then reduce on the right.

Perhaps more interesting are the rules for “boxed” terms which are those typed with a graded modality in our system; note that in many other Granule-related papers with call-by-need semantics any boxed term is a value by default, but here a boxed term is only a value if the term inside the box is itself a value. The reduction rules are:

$$\overline{\mathbf{let} [x] = [v] \mathbf{in} t \rightsquigarrow t[v/x]} \rightsquigarrow_{\square\beta} \quad \frac{t \rightsquigarrow t'}{[t] \rightsquigarrow [t']} \rightsquigarrow_{\square\text{reduce}}$$

$$\frac{t_1 \rightsquigarrow t'_1}{\mathbf{let} [x] = t_1 \mathbf{in} t_2 \rightsquigarrow \mathbf{let} [x] = t'_1 \mathbf{in} t_2} \rightsquigarrow_{\text{LET}\square}$$

Thus we can β -reduce an introduction-elimination pair for the graded modal values, eliminating the box modality through pattern matching. We reduce inside the box for terms that are not yet values, and we have a congruence rule for reducing the term in an elimination form. We will skip over the rules for tensor product and unit in this section, since they are very similar to the

rules for boxed terms, with rules for reducing terms inside a pair and also substitution and elimination for products and unit as above.

The real action happens when we reach an instance of application where the value being applied on the left is some (potentially already partially applied) communication primitive and the term on the right is the final parameter for said primitive, as this is where some communication behaviour needs to occur and so we will have to consider the state of the channel configuration. We begin by considering the creation of a new channel, which occurs through the use of the `forkLinear` primitive. For this we need to now extend the syntax of terms to include *runtime values*, which in this setting are just channel endpoint names c :

$$t ::= \dots \mid c \quad (\text{Runtime terms})$$

Furthermore, we also define the notion of a term context, denoted $Ctxt$, in which reduction can occur at a single “hole” in the syntax tree:

Definition 1 (Term contexts). Syntax trees with a single hole are defined:

$$\begin{array}{lcl}
 Ctxt & ::= & \text{Contexts} \\
 & | & - \\
 & | & Ctxt\ t \\
 & | & v\ Ctxt \\
 & | & [Ctxt] \\
 & | & \mathbf{let}\ [x] = Ctxt\ \mathbf{in}\ t \\
 & | & \mathbf{let}\ \mathbf{unit} = Ctxt\ \mathbf{in}\ t \\
 & | & \mathbf{let}\ (x, y) = Ctxt\ \mathbf{in}\ t \\
 & | & (Ctxt, t) \\
 & | & (v, Ctxt)
 \end{array}$$

where $-$ denotes a hole and we write $Ctxt\{t\}$ to denote the operation of filling the hole in $Ctxt$ with a term t .

The operational semantics of `forkLinear`, in some term context, is then:

$$\frac{\#c}{(\mathcal{P} \mid Ctxt\{\mathbf{forkLinear}\ (\lambda x.t)\}; \mathcal{C}) \rightsquigarrow (\mathcal{P} \mid t[c/x]^\circ \mid Ctxt\{\bar{c}\}; \mathcal{C}, \bar{c} = \langle \rangle, c = \langle \rangle)} \text{FORKLINEAR}$$

In this case, we select a fresh name c for the new channel (denoted by $\#c$ in the premise) and add both dual endpoints of the channel (c and \bar{c}) to

the channel configuration, each empty of values for now. As well as reducing the term by carrying out the expected substitution we also spawn a new process for communication on the fresh channel and add it to the process configuration, marked as a subordinate process.

In the above rule, we need not concern ourselves with whether the fresh endpoint name c is a positive channel name n^+ or negative n^- since the actual polarity is immaterial; whatever fresh name we have, the channel configuration is extended with a queue associated to this endpoint name c and a queue associated to its dual endpoint name \bar{c} .

The next cases we consider are sending and receiving information on one end of a channel, where the dual end will have to receive or provide said information in turn. Sending is asynchronous, given by:

$$\overline{(\mathcal{P} \mid \text{Ctx}\{\text{send } c \ v\}; \mathcal{C}, \bar{c} = \vec{v})} \rightsquigarrow (\mathcal{P} \mid \text{Ctx}\{c\}; \mathcal{C}, \bar{c} = \vec{v} : v) \quad \text{SEND}$$

The value v provided by the user is consumed, and the channel configuration is updated by appending v to the end of whatever queue of values is already present inside the *dual* end of the channel they sent on, where they will later be necessarily received by virtue of the session type guarantees. The process configuration contains the original end of the channel in context, which the sender will continue to use according to the next step of its protocol given by the session type. Conversely, receiving operates as follows:

$$\overline{(\mathcal{P} \mid \text{Ctx}\{\text{receive } c\}; \mathcal{C}, c = v : \vec{v})} \rightsquigarrow (\mathcal{P} \mid \text{Ctx}\{(v, c)\}; \mathcal{C}, c = \vec{v}) \quad \text{RECV}$$

In this case, the first value is popped from the queue of values contained in the end of the channel held by the receiver, and this is given to the user inside the context in the process configuration; they can again continue to use their end of the channel according to its protocol. This may involve receiving more values or sending values to the dual end depending on the particular session type that has been specified.

Finally, channels are closed as follows:

$$\overline{(\mathcal{P} \mid \text{Ctx}\{\text{close } c\}; \mathcal{C}, c = \langle \rangle)} \rightsquigarrow (\mathcal{P} \mid \text{Ctx}\{\mathbf{unit}\}; \mathcal{C}) \quad \text{CLOSE}$$

Note this reduction closes the referred-to endpoint of the channel c , rather than its dual, akin to how RECV receives from the referred to endpoint of the channel. This allows for asynchronous behaviour since the dual endpoint \bar{c} may still have values in its queue to be received.

5.2. Runtime typing

In order to later consider type safety properties, we extend typing to process and channel configurations, considering a notion of *runtime typing*.

For runtime typing of terms, the same typing rules apply but where the context also contains channel names and their typing. We refer to a *runtime context* by the letter Δ , defined:

$$\begin{aligned} \Delta &= \cdot \mid \Delta, c : \mathbf{Chan} P && \text{(Runtime contexts)} \\ \Gamma &= \dots \mid \Delta && \text{(Contexts (redefined))} \end{aligned}$$

where $c : \mathbf{Chan} P$ types a linear channel. Contexts Γ are then extended to include Δ , though it is sometimes useful to denote contexts that contain only runtime types (channel types) separately by using the Δ notation.

Context addition extends to Δ , treating the linear channel names just as linear variables. The variable rule of typing extends to channel names, i.e., we can derive runtime typing:

$$\frac{}{c : \mathbf{Chan} P \vdash c : \mathbf{Chan} P} \text{VAR}$$

Process configurations \mathcal{P} are typed by a judgment $\Gamma, \Delta \vdash \mathcal{P} : A$ where A is the type of the main process and Γ are the term variables and Δ the channel names in scope of the processes in \mathcal{P} . Here, the type \perp is a type that cannot be given to any term, but is used only for process configuration typing to denote a configuration without a main process.

$$\frac{}{\emptyset \vdash \cdot : \perp} \text{EMPTY}$$

$$\frac{\Gamma_1, \Delta_1 \vdash \mathcal{P} : \perp \quad \Gamma_2, \Delta_2 \vdash t : A}{\Gamma_1 + \Gamma_2, \Delta_1 + \Delta_2 \vdash \mathcal{P} \mid t^\bullet : A} \text{MAIN} \quad \frac{\Gamma_1, \Delta_1 \vdash \mathcal{P} : A \quad \Gamma_2, \Delta_2 \vdash t : 1}{\Gamma_1 + \Gamma_2, \Delta_1 + \Delta_2 \vdash \mathcal{P} \mid t^\circ : A} \text{SUB}$$

Empty process configurations are given the type \perp . The MAIN rule puts a term t in parallel with a configuration \mathcal{P} and marks it as the main process and thus the process configuration has the type of this process. The SUB rule puts a term t in parallel with a configuration \mathcal{P} , but this term is marked as a subordinate process with \circ (of which there may be many) and thus the return type of t must be **unit**. The process configuration's type is that of \mathcal{P} . Note that this type can itself be \perp , and so a process configuration with type \perp is not necessarily empty.

Channel configurations \mathcal{C} are typed by a judgment $\Delta \vdash \mathcal{C}$, defined:

$$\frac{}{\emptyset \vdash \cdot} \text{NIL}$$

$$\frac{\Delta \vdash \mathcal{C}}{\Delta, c : \mathbf{Chan} P \vdash \mathcal{C}, c = \langle \rangle} \text{EMPTY} \quad \frac{\Delta \vdash \mathcal{C}}{\Delta, c : \mathbf{Chan} \mathbf{End} \vdash \mathcal{C}, \bar{c} = \bar{v}} \text{END}$$

$$\frac{\Delta_1 \vdash v : A \quad \Delta_1, \Delta_2, c : \mathbf{Chan} P \vdash \mathcal{C}, \bar{c} = \bar{v} : v}{\Delta_1, \Delta_2, c : \mathbf{Chan} (\mathbf{Send} A P) \vdash \mathcal{C}, \bar{c} = \bar{v}} \text{SEND}$$

$$\frac{\Delta_1 \vdash v : A \quad \Delta_1, \Delta_2, c : \mathbf{Chan} P \vdash \mathcal{C}, c = \bar{v}}{\Delta_1, \Delta_2, c : \mathbf{Chan} (\mathbf{Recv} A P) \vdash \mathcal{C}, c = v : \bar{v}} \text{RECV}$$

The **EMPTY** rule says that a channel endpoint that is empty can have any type (it is in a ‘waiting state’). Meanwhile, the **END** rule says a channel can be typed at **Chan End** while its dual endpoint may contain any number of values. Reading the **SEND** rule bottom-up, the conclusion types a channel on which a value v of type A can be sent and thus we can type that channel with the continuation behaviour P in the premise when that value has been enqueued on the channel. Note that channel value is the dual to the typing, since sending is actioned by enqueueing on the dual endpoint. Reading the **RECV** rule bottom-up, the conclusion types a channel on which a value v of type A can now be received (from the head of the channel) with the continuation behaviour P in the premise after the value has been received (and hence removed) from the front of the channel.

5.2.1. Summary

So far we have established just the operational semantics for the core calculus with linear channels. We are now in a position to establish the key properties of well-typed programs containing graded modal session types but we delay this to Section 9.

Next, we introduce the new non-linear communication primitives which are the main contribution of this paper and expand the operational semantics to include their model, and associated definitions, which also includes expanding the definition of processes to include some intermediate runtime processes. In Section 9, we then consider the syntactic type safety properties for the full calculus with the non-linear extensions included.

6. Reusable Channels

A reusable or *non-linear* channel is one which can be shared and thus used repeatedly in a sound way. This contrasts with the notion seen in Section 4 of promoting a fresh linear channel to being non-linear in a call-by-value setting, which was unsound; a linear channel used in a shared way can easily lead to a deadlock where one user of the channel leaves it in a state which is then blocking for another user of the channel. The key problem with sharing a channel is ending up with inconsistent states across different shared usages. This can be avoided if the protocol allowed on the channel is restricted such that only a single ‘action’ (send, receive, choice, or offer) is allowed, and thus if the channel is used multiple times it can never be left in an inconsistent state across shared uses. This is captured by the idea that $(P \mid P) \neq P$ in general in process calculi (like CCS or the π -calculus) and so a replicated channel which has more than a single action cannot be split off into many parallel uses, e.g., $(ab)^* \mid (ab)^* \not\equiv (ab)^*$. However, if there is only a single action then multiple repeated parallel uses are consistent, e.g., $a^* \mid a^* \equiv a^*$.

We capture this idea via the following variant of the fork primitive with arbitrarily graded channels:

```
forkNonLinear :  $\forall$  {p : Protocol, s : Semiring, r : s}
  . {SingleAction p, ExactSemiring s}
   $\Rightarrow$  ((LChan p) [r]  $\rightarrow$  ())  $\rightarrow$  (LChan (Dual p)) [r]
```

where `ExactSemiring : Semiring \rightarrow Predicate` restricts acceptable semirings to only those that do not permit any approximation (so as far as semirings mentioned in this paper are concerned, the grade `r` is allowed to be a natural number, representing exact usage, but importantly is *not* allowed to be an interval, representing a range of possible usages) and `SingleAction : Protocol \rightarrow Predicate` is a type constraint that characterises only those protocols which comprise a single send, receive, choice, or offer, i.e.,

```
SingleAction End          SingleAction (Send a End)      SingleAction (Offer End End)
                          SingleAction (Recv a End)     SingleAction (Select End End)
```

As an example, consider a channel in a graded modality which says it can be used exactly `n` times. The following uses this channel to send every element of a vector of size `n` (using the `Vec` type of Section 2):

```
1 sendVec :  $\forall$  {n : Nat, a : Type} . (LChan (Send a End)) [n]  $\rightarrow$  Vec n a  $\rightarrow$  ()
2 sendVec [c] Nil = ();
3 sendVec [c] (Cons x xs) =
```



```

4     let c' = send c x;
5         () = close c'
6     in sendVec [c] xs

```

This code shows the powerful interaction between grading, linearity, and indexed types in Granule. Note that the above code is typeable without any of the new primitives described in this section, but would be unusable without promoting a use of `forkLinear`. However, we can complete this example with a dual process (`recvVec`) that is then connected to `sendVec` via `forkNonLinear`:

```

1  recvVec : ∀ {n : Nat, a : Type} . N n → (LChan (Recv a End)) [n] → Vec n a
2  recvVec Z [c] = Nil;
3  recvVec (S n) [c] = let (x, c') = recv c;
4                       () = close c'
5                       in Cons x (recvVec n [c])
6
7  length' : ∀ {n : Nat, a : Type} . Vec n a → (N n, Vec n a)
8  length' Nil = (Z, Nil);
9  length' (Cons x xs) = let (n, xs) = length' xs in (S n, Cons x xs)
10
11 example : ∀ {n : Nat, a : Type} . Vec n a → Vec n a
12 example xs = let (n, list) = length' xs
13              in recvVec n (forkNonLinear (λc → sendVec c list))
14
15 main : Vec 5 Int
16 main = example (Cons 1 (Cons 1 (Cons 2 (Cons 3 (Cons 5 Nil))))))

```

Since `forkNonLinear` give us a pair of channels whose semiring-usage is exact, we must consume all values produced by `sendVec` on the dual endpoint. Therefore, the receiver `recvVec` needs to know how many elements to receive. This information is passed separately (via an indexed natural number \mathbb{N} n) so that `recvVec` can induct on the natural number and leave no message unreceived. A system with dependent session types [34, 35] could avoid this by first sending the length, but this is not (yet) possible in Granule; the Gerty prototype language provides full dependent types and graded modal types which would be a good starting point [26].

6.1. Semantics

In the notation of the core calculus, the `forkNonLinear` primitive has its typing rendered as:

$$\frac{\text{SingleAction}(P) \quad \text{ExactSemiring}(S) \quad r : S}{\emptyset \vdash \text{forkNonLinear} : (\Box_r \mathbf{Chan} P \multimap 1) \multimap \Box_r(\overline{\mathbf{Chan} P})} \text{FORKNONLIN}$$

This primitive now introduces the notion of a non-linear, or *shared* channel, which we reflect into the run-time typing and channel configuration typing that was setup in Section 5.

As an example (partial) derivation here, consider the following term which uses `forkNonLinear` with a protocol **Send Int End**, and the natural numbers semiring with grade 2 to spawns a process reusing the channel twice to send two integer terms (3 then 5), yielding the dual channel which can be used twice to receive two integers:

$$\frac{\frac{\frac{y : [\mathbf{Chan}(\mathbf{Send\ Int\ End})]_2 \vdash \text{let } \mathbf{unit} = \text{close}(\text{send } y\ 3) \text{ in } \text{close}(\text{send } y\ 5) : 1}{x : \square_2(\mathbf{Chan}(\mathbf{Send\ Int\ End})) \vdash \text{let } [y] = x \text{ in } \text{let } \mathbf{unit} = \text{close}(\text{send } y\ 3) \text{ in } \text{close}(\text{send } y\ 5) : 1}^{\text{ELIM}}}{\emptyset \vdash \lambda x. \text{let } [y] = x \text{ in } \text{let } \mathbf{unit} = \text{close}(\text{send } y\ 3) \text{ in } \text{close}(\text{send } y\ 5) : 1}^{\text{--}_i}}{\emptyset \vdash \text{forkNonLinear}(\lambda x. \text{let } [y] = x \text{ in } \text{let } \mathbf{unit} = \text{close}(\text{send } y\ 3) \text{ in } \text{close}(\text{send } y\ 5)) : \square_2(\mathbf{Chan}(\mathbf{Recv\ Int\ End}))}^{\text{FORKNONLIN}}$$

Shared channel typing. We extend runtime-typing contexts Δ with a shared channel typing:

$$\Delta = \cdot \mid \Delta, c : \mathbf{Chan } P \mid \Delta, c :^* \mathbf{Chan } P \quad (\text{Runtime contexts})$$

where $c :^* \mathbf{Chan } P$ types a shared channel.

The predicate $[\Gamma]$ classifying graded contexts then also extends to include the rule that $[\Gamma] \implies [\Gamma, c :^* \mathbf{Chan } P]$, i.e., we treat shared channels akin to graded variables, as graded variables may be used non-linearly. Runtime context addition then needs to account for sharing:

Definition 2 (Runtime context addition). Runtime context addition is akin to context addition, where channels can be marked as linear or shared and addition is only defined for contexts that are disjoint in their linear channels and where shared channels may appear in either context:

$$\begin{aligned} \emptyset + \Delta_2 &= \Delta_2 \\ \Delta_1 + \emptyset &= \Delta_1 \\ (\Delta_1, c : \mathbf{Chan } P) + \Delta_2 &= (\Delta_1 + \Delta_2), c : \mathbf{Chan } P \quad \text{iff } c \notin \text{dom}(\Delta_2) \\ \Delta_1 + (\Delta_2, c : \mathbf{Chan } P) &= (\Delta_1 + \Delta_2), c : \mathbf{Chan } P \quad \text{iff } c \notin \text{dom}(\Delta_1) \\ (\Delta_1, c :^* \mathbf{Chan } P_1) + (\Delta_2, c :^* \mathbf{Chan } P_2) &= (\Delta_1 + \Delta_2), c :^* \mathbf{Chan } P \quad \text{where } P = P_1 \bowtie P_2 \end{aligned}$$

where $P_1 \bowtie P_2$ defines *compatibility of shared channel session types*, a symmetric partial operation defined:

$$\begin{aligned} P \bowtie P &= P \quad \text{if } \text{SingleAction}(P) \\ P \bowtie \mathbf{End} &= P \quad \text{if } \text{SingleAction}(P) \\ \mathbf{End} \bowtie P &= P \quad \text{if } \text{SingleAction}(P) \end{aligned}$$

where `SingleAction` is a built-in predicate on protocols:

`SingleAction(End)` `SingleAction(Recv A End)` `SingleAction(Send A End)`

Thus, two shared channels are compatible when they have the same protocol, or when one of the shared channels has been used and is in the **End** state. For example, consider three non-linear uses of a shared channel c whose protocol is a single action **Send A End**; after one communication is performed, one of the uses of c is typed as **End**, but this is still compatible with the two remaining non-linear uses of c typed at **Send A End**.

Thus, $+$ combines contexts which are disjoint in their linear channels, and allows shared channels to be included in both contexts as long as their session types are compatible, computing via \bowtie their compatible shared type.

Scalar multiplication of contexts extends smoothly to scalar multiplication of runtime contexts where $r * (\Delta, c : * \mathbf{Chan} P) = (r * \Delta), c : * \mathbf{Chan} P$.

The variable rule extends to shared channel names:

$$\frac{}{c : * \mathbf{Chan} P \vdash c : \mathbf{Chan} P} \text{VARCHANS}$$

The weakening rule also extends to shared channel names:

$$\frac{\Gamma \vdash t : B}{\Gamma, 0 * \Gamma', c : * \mathbf{Chan} P \vdash t : B} \text{WEAKCHANS}$$

Channel configurations and reduction. In channel configurations, we explicitly denote shared channels with assignments written $c^* = \vec{v}$. If we need to distinguish a linear channel in a channel configuration, we write $c^1 = \vec{v}$ though if the linear or shared nature is immaterial we write $c = \vec{v}$ as before. The typing of channel configurations behaves the same for shared channels as linear channels.

Send and receive reduction rules apply to shared channels exactly as linear channels. However, the `close` primitive has a different reduction semantics for shared channels:

$$\frac{}{(\mathcal{P} \mid \text{Ctx}\{\text{close } c\}; \mathcal{C}, c^* = \vec{v}) \rightsquigarrow (\mathcal{P} \mid \text{Ctx}\{\mathbf{unit}\}; \mathcal{C}, c^* = \vec{v})} \text{SHAREDCLOSE}$$

Note that when a shared channel is closed, we preserve it inside the channel configuration rather than removing it. Furthermore, we do not even require that it is empty since the channel may be shared and messages may have arrived for another interaction, e.g., particularly in the case of `forkNonLinear`.

Reduction semantics for forkNonLinear. Now that we have extended the runtime typing notions to shared channels we can consider the operational semantics of `forkNonLinear`.

The operational model of `forkNonLinear` follows the typing: its reduct starts with an application of this primitive to a function value `forkNonLinear` $(\lambda x.t)$ in some term context within a process reducing by creating a fresh shared channel endpoint c (and its dual endpoint \bar{c}) and spawning a new process t with x replaced by the graded modal lifting of the channel $[c]$ and returning the dual endpoint also inside a graded modal box:

$$\frac{\#c}{\begin{array}{l} (\mathcal{P} \mid \text{Ctx}\{\text{forkNonLinear}(\lambda x.t)\}; \mathcal{C}) \\ \rightsquigarrow (\mathcal{P} \mid \text{Ctx}\{\bar{c}\} \mid t[[c]/x]^\circ; \mathcal{C}, \bar{c}^* = \langle \rangle, c^* = \langle \rangle) \end{array}} \text{FORKNONLINEAR}$$

The spawned process based on the function body t (with the promoted channel substituted in) is marked as being a subordinate process with \circ . The newly created channels in the context are marked as being shared channels c^* and \bar{c}^* .

7. Replicated Servers

The π -calculus provides replication of a process P as the process $!P$, which session-typed π -calculus variants have refined into the more controlled idea of having a replicated ‘server’. Here, the spawning of replicated instances is controlled via a special receive [3, 32, 37], e.g., written like $*c(x).P$ meaning receive an x on channel c and then continue as P , whilst still providing the original process. From an operational semantics point of view this looks like $*c(x).P \mid \bar{c}(d).Q \rightarrow *c(x).P \mid P[d/x] \mid Q$ where $\bar{c}(d).Q$ sends the message d to the server which ‘spawns’ off a fresh copy of the server process P whilst also preserving the original server process for further clients to interact with.

This important pattern is not achievable with the reusable channels of Section 6 but the non-linearity inherent to replicated servers can be captured and explained again via the combination of graded and indexed types. We provide the replicated server functionality here as the fork variant `forkReplicate`:

$$\begin{array}{l} \text{forkReplicate} : \forall \{p : \text{Protocol}, n : \text{Nat}\} . \{\text{ReceivePrefix } p\} \\ \Rightarrow (\text{LChan } p \rightarrow ()) \ [0..n] \rightarrow \mathbb{N} \ n \rightarrow \text{Vec } n \ ((\text{LChan } (\text{Dual } p)) \ [0..1]) \end{array}$$

Here the grading is less general than in Section 6; we instead focus on a particular grading which says that given a server process $(\text{LChan } p \rightarrow ())$ that

can be used 0 to n times, then we get a vector of size n of dual channels which we can use to interact with the server. The predicate `ReceivePrefix p` classifies those protocols which start with a receive, i.e.:

`ReceivePrefix (Recv a p)` `ReceivePrefix (Offer p1 p2)`

Each returned client channel can itself be discarded due to the graded modality ... [0..1]. Thus, we can choose not to use any/all of the client channels, reflected in the dual side where the server can be used at most n times. This introduces some flexibility in the amount of usage.

The following example demonstrates `forkReplicate` to define a simple server that can receive requests to receive two further integers which it adds and sends back, or requests to compare a received integer for equality to 0:

```

1  addServer : LChan (Offer (Recv Int (Recv Int (Send Int End)))
2                      (Recv Int (Send Bool End))) → ()
3  addServer c =
4    offer
5      (λc → let (x, c) = recv c;
6              (y, c) = recv c;
7              in close (send c (x + y))
8
9      (λc → let (x, c) = recv c;
10           in close (send c (x == 0))) c

```

We then define two clients `client1` and `client2` that can interact with this server, exercising both parts of the offered choice:

```

1  client1 : ∀ {p : Protocol}
2           . LChan (Select (Send Int (Send Int (Recv Int End))) p) → Int
3  client1 c = let (x, c) = recv (send (send (selectLeft c) 10) 20);
4              () = close c
5              in x
6
7  client2 : ∀ {p : Protocol}
8           . LChan (Select p (Send Int (Recv Bool End))) → Bool
9  client2 c = let (b, c) = recv (send (selectRight c) 42);
10             () = close c
11             in b

```

We then compose the server and clients together using `forkReplicate` to get a vector of channels which we pattern match on immediately to get the channels to pass to each client (which are composed in parallel using `par` defined in Section 3.2):

```

1 import Parallel -- Provides the `par` combinator derived from `forkLinear`
2 main : (Int, Bool)
3 main =
4   let (Cons [c1] (Cons [c2] Nil)) = forkReplicate [addServer] (S (S Z))
5   in par (λ () → client1 c1) (λ () → client2 c2)

```

In the intuitionistic linear logical propositions of Caires and Pfenning [7], this same idea of a replicating fork is captured by a non-linear channel (in Dual Intuitionistic Linear Logical style) which yields a linear version when interacted with; this is akin to the idea of *shared channels* [42]. Our approach codifies the same principle using graded modalities, but only allows finite replication. Section 7.2 shows a variation of this primitive which allows unbounded replication.

7.1. Semantics

In the notation of the core calculus, the `forkReplicate` primitive has its typing rendered as:

$$\frac{\text{ReceivePrefix}(P)}{\emptyset \vdash \text{forkReplicate} : \square_{0..n}(\mathbf{Chan} P \multimap 1) \multimap \mathbf{N} n \multimap \mathbf{Vec} n (\square_{0..1}(\mathbf{Chan} \overline{P}))} \text{FORKREP}$$

We thus also extend the core calculus with indexed types for natural numbers and vectors, which previously were derived inside Granule using the syntax of GADTs (see Section 2):

$$A ::= \dots \mid \mathbf{N} n \mid \mathbf{Vec} n A$$

These indexed types are internally definable in Granule as a GADTs, but for simplicity we add them as primitive to the core calculus, with standard introduction rules:

$$\frac{}{\emptyset \vdash \mathbf{Z} : \mathbf{N} 0} \text{Z} \qquad \frac{\Gamma \vdash t : \mathbf{N} n}{\Gamma \vdash \mathbf{S} t : \mathbf{N} (n+1)} \text{S}$$

$$\frac{}{\emptyset \vdash \langle \rangle : \mathbf{Vec} 0 A} \text{NIL} \qquad \frac{\Gamma_1 \vdash t_1 : A \quad \Gamma_2 \vdash t_2 : \mathbf{Vec} n A}{\Gamma_1 + \Gamma_2 \vdash t_1 :: t_2 : \mathbf{Vec} (n+1) A} \text{CONS}$$

We elide their eliminators here for brevity, but their elimination is standard following usual eliminators for indexed types [29] (see also Granule’s type system [27], whose indexed type eliminators also cause specialisation in grades, allowing type indices and grades to interact). As a shorthand, we write a vector constructor with n elements as $\langle t_1, \dots, t_n \rangle$.

As an example (partial) derivation, the following uses `forkReplicate` to spawn a server that can be used at most three times to receive an integer and sends it back incremented:

$$\frac{\frac{\emptyset \vdash \lambda x. \text{let } (y, x) = \text{receive } x \text{ in close } (\text{send } x (y + 1)) : \mathbf{Chan} (\mathbf{Recv Int} (\mathbf{Send Int End})) \multimap 1}{\emptyset \vdash [\lambda x. \text{let } (y, x) = \text{receive } x \text{ in close } (\text{send } x (y + 1))] : \square_{0..3} (\mathbf{Chan} (\mathbf{Recv Int} (\mathbf{Send Int End})) \multimap 1)} \text{PR}}{\emptyset \vdash \text{forkReplicate } [\lambda x. \text{let } (y, x) = \text{receive } x \text{ in close } (\text{send } x (y + 1))] : \mathbf{N} \ 3 \multimap \mathbf{Vec} \ 3 (\square_{0..1} (\mathbf{Chan} (\mathbf{Send Int} (\mathbf{Recv Int End})))} \text{FORKREP}}$$

The operational model of `forkReplicate` splits into two rules. The first follows the typing and corresponds to a β -reduction of `forkReplicate`: its reduct has an application of this primitive to a promoted function application and a natural number term `forkReplicate` $[\lambda x.t] n$. The result is to spawn off a number of intermediate processes called “watchers” and to return to the originating process a vector of fresh channel endpoints (each promoted), of type $\mathbf{Vec} \ n (\square_{0..1} (\mathbf{Chan} \bar{P}))$:

$$\frac{\#c_1 \dots \#c_n}{(\mathcal{P} \mid \text{Ctxt}\{\text{forkReplicate } [\lambda x.t] n\}; \mathcal{C})} \text{FORKREPLICATE} \\ \rightsquigarrow (\mathcal{P} \mid \text{Ctxt}\{\langle [\bar{c}_1], \dots, [\bar{c}_n] \rangle\} \mid \mathbf{Wch} (c_1 \mapsto \lambda x.t) \mid \dots \mid \mathbf{Wch} (c_n \mapsto \lambda x.t) \\ ; \mathcal{C}, c_1^* = \langle \rangle, \dots, c_n^* = \langle \rangle, \bar{c}_1^* = \langle \rangle, \dots, \bar{c}_n^* = \langle \rangle)$$

Thus, n pairs of fresh channel endpoints are created and the `forkReplicate` reduct, in some term context *Ctxt*, is replaced by a vector of the dual endpoints (promoted to the graded modality). In parallel, n “watcher” processes are created, where $\mathbf{Wch} (c \mapsto t)$ is a runtime process waiting to receive a value on channel c and containing the continuation of the replicated process body: the suspended computation of process t , abstracted over the channel x .

Watcher processes (which are not ordinary terms) have the following typing, under the process typing judgment:

$$\frac{\Gamma_1, \Delta_1 \vdash \mathcal{P} : A \quad \Gamma_2, x : \mathbf{Chan} (\mathbf{Recv} A' P), \Delta_2 \vdash t : 1}{\Gamma_1 + (0..1 * \Gamma_2), \Delta_1 + \Delta_2, c : \mathbf{Chan} (\mathbf{Recv} A' P) \vdash \mathcal{P} \mid \mathbf{Wch} (c \mapsto \lambda x.t) : A} \text{WATCH}$$

Thus, given a t which has a free x which is a channel on which we can receive, then t can be placed into a watcher process, linked to the channel c on which a receive can occur. Note that in the conclusion here we scale the requirement of t by $0..1$ since t may never be evaluated as `forkReplicate` provides an upper bound on the number of possible replications.

The second part of the operational semantics for `forkReplicate` then involves reducing a watcher process into a regular term-based process upon

receiving an initial request:

$$\overline{(\mathcal{P} \mid \text{Wch}(c \mapsto \lambda x.t); \mathcal{C}, c = v : \vec{v})} \rightsquigarrow (\mathcal{P} \mid t[c/x]^{\circ}; \mathcal{C}, c = v : \vec{v}) \quad \text{WATCHER}$$

For a watcher on channel c , if there is a value v at the end head of the channel v then a new subordinate process is formed, substituting c for x . Note that the channel is left unchanged so that the spawned process can receive the value; the presence of the value v in the channel simply triggers the spawning of the process from the watcher.

7.2. Variations

In `forkReplicate`, each returned channel (in a vector) can be discarded due to the graded modality `... [0..1]`: we can choose not to use any/all of the client channels, which allows for some flexibility in the amount of usage. A more strict variant is given as:

```
forkReplicateExactly : ∀ {p : Protocol, n : Nat} . {ReceivePrefix p}
  ⇒ (LChan p → ()) [n] → N n → Vec n (LChan (Dual p))
```

meaning we have exactly n clients that *must* spawn the server n times.

A more powerful version of replication allows an unknown number of clients (possibly infinite), with an interface similar to `forkReplicate` but returning a lazy stream of client channels which can be affinely used:

```
data Stream a = Next a (Stream a)
```

```
forkReplicateForever : ∀ {p : Protocol} . {ReceivePrefix p}
  ⇒ (LChan p → ()) [0..∞] → (Stream (LChan (Dual p))) [0..1]
```

The result is a stream wrapped in the graded modality with grade `0..1`. Due to the typing of pattern matching in Granule, one can pattern match on the stream inside the graded modality, such that its components inherit the grade `0..1`; e.g., `example [Next c rest] = ...` binds the stream `rest` in the scope of the function body with the same grade `0..1` so it can be used once or not at all.

This last variation of replication captures more closely the replicated server concept found in session-typed π -calculi. The above two primitives are provided in the implementation but are not considered in the core calculus and metatheory here.

8. Multicast Sending

The final primitive we introduce here provides the notion of *multicast* (or *broadcast*) communication where messages on a channel can be received by multiple processes by copying the messages. This non-linearity is thus on the payload values being sent, with a graded modal type deployed to explain that the amount of non-linearity in the messages matches the number of receiving processes (clients). The `forkMulticast` primitive provides this behaviour, typed in Granule:

```
forkMulticast : ∀ {p : Protocol, n : Nat}
  . {Sends p}
  ⇒ (Chan (Graded n p) → ()) → N n → Vec n (Chan (Dual p))
```

where `Graded : Nat → Protocol → Protocol` is a built-in type function that adds a graded modality to payload types, defined:

```
Graded n End           = End
Graded n (Send a p)    = Send (a [n]) (Graded n p)
Graded n (Recv a p)    = Recv (a [n]) (Graded n p)
Graded n (Select p1 p2) = Select (Graded n p1) (Graded n p2)
Graded n (Offer p1 p2) = Offer (Graded n p1) (Graded n p2)
```

and `Sends : Protocol → Predicate` is a built-in predicate defined:

$$\frac{}{\text{Sends End}} \quad \frac{\text{Sends } p}{\text{Sends (Send } a \text{ } p)} \quad \frac{\text{Sends } p1 \quad \text{Sends } p2}{\text{Sends (Select } p1 \text{ } p2)}$$

Thus, as long as we are sending values that are wrapped in the graded modality such that they can be used n times, we can then broadcast these to n client processes via the `Vec n (Chan (Dual p))` channels returned by `forkMulticast`.

For example, in the following we have a broadcaster that takes a channel which expects an integer to be sent which can be used 3 times.

```
1 broadcaster : LChan (Send (Int [3]) End) → ()
2 broadcaster c = close (send c [42])
```

We can then broadcast these results with `forkMulticast broadcaster (S (S (S Z)))` producing a 3-vector of channels. Below we aggregate the results from these three receiver channels by applying `aggregateRecv` to the vector, giving us the result 126, i.e. $3 * 42$.

```
1 aggregateRecv : ∀ {n : Nat} . Vec n (LChan (Recv Int End)) → Int
2 aggregateRecv Nil           = 0;
3 aggregateRecv (Cons c cs) = let (x, c) = recv c;
```

```

4           () = close c
5           in x + aggregateRecv cs
6
7 main : Int -- Evaluates to 126
8 main = aggregateRecv (forkMulticast broadcaster (S (S (S Z))))

```

8.1. Semantics

In the notation of the core calculus, the `forkMulticast` primitive has its typing rendered as:

$$\frac{\text{Sends}(P)}{\emptyset \vdash \text{forkMulticast} : (\mathbf{Chan}(\text{Graded}_n P) \multimap 1) \multimap \mathbf{N} \ n \multimap \mathbf{Vec} \ n (\mathbf{Chan} \overline{P})} \text{FORKMULTI}$$

Similarly to `forkReplicate`, the operational model of `forkMulticast` comprises two rules; a β -reduction that generate some intermediate processes, and a rule actioning those processes in response to communication. The β -reduction rule follows the typing: its reduct has an application of the primitive to a function value and a natural number term: `forkMulticast` $(\lambda x.t) \ n$. The result is to spawn the process derived from t as well as a special “forwarding” process (explained below):

$$\frac{\#c \quad \#c_1 \dots \#c_n}{\begin{array}{l} (\mathcal{P} \mid \text{Ctx}\{\text{forkMulticast}(\lambda x.t) \ n\}; \mathcal{C}) \\ \rightsquigarrow (\mathcal{P} \mid \text{Ctx}\{\langle \overline{c}_1, \dots, \overline{c}_n \rangle\} \mid t[c/x]^\circ \mid \text{Fwd}(c \mapsto c_1, \dots, c_n) \\ \quad ; \mathcal{C}, c = \langle \rangle, \overline{c} = \langle \rangle, c_1 = \langle \rangle, \dots, c_n = \langle \rangle, \overline{c}_1 = \langle \rangle, \dots, \overline{c}_n = \langle \rangle) \end{array}} \text{FORKMULTICAST}$$

Thus, n pairs of fresh channel endpoints are created, $c_1 \dots c_n$, and the `forkMulticast` reduct is replaced by a vector of the dual endpoints $\langle \overline{c}_1, \dots, \overline{c}_n \rangle$. In parallel, the function body t from the reduct has c substituted for x to form a subordinate process $t[c/x]^\circ$, and a “forwarding” process is created. The forwarding process $\text{Fwd}(c \mapsto c_1, \dots, c_n)$ is a runtime process that receives a graded modal value from c , unboxes the value and forwards copies of this payload on the (dual end of) channels c_1 to c_n .

Forwarder processes are typed by the process typing judgment as:

$$\frac{\text{Sends}(P) \quad \Gamma, \Delta \vdash \mathcal{P} : A}{\Gamma, \Delta, c : \mathbf{Chan}(\text{Graded}_n \overline{P}), c_i : \mathbf{Chan} \ P \vdash \mathcal{P} \mid \text{Fwd}(c \mapsto c_1, \dots, c_n) : A} \text{FWD}$$

where $c_i : \mathbf{Chan} \ P$ is short hand for the typing $c_1 : \mathbf{Chan} \ P, \dots, c_n : \mathbf{Chan} \ P$. For example, if $P = \mathbf{Send} \ B \ P'$ (satisfying the `Sends` predicate) then $c : \mathbf{Chan}(\mathbf{Recv}(\square_n B)(\text{Graded}_n \overline{P}'))$ in the above.

The reduction of the forwarding process then occurs by receiving on channel c an n -graded value, which gets sent along n further channels:

$$\frac{}{(\mathcal{P} \mid \text{Fwd}(c \mapsto c_1, \dots, c_n); \mathcal{C}, c = [v] : \vec{v}, \overline{c_1} = \vec{v}_1, \dots, \overline{c_n} = \vec{v}_n)} \text{FORWARDER} \\ \rightsquigarrow (\mathcal{P} \mid \text{Fwd}(c \mapsto c_1, \dots, c_n); \mathcal{C}, c = \vec{v}, \overline{c_1} = \vec{v}_1 : v, \dots, \overline{c_n} = \vec{v}_n : v)$$

Note that there is no “garbage collection” step for the forwarding process, though this can be easily added as an implementational detail.

9. Properties

We consider the standard probes on the language design: whether typing guarantees progress of reduction (that we do not get stuck, or deadlocked) and preservation of types (that reduction does not conflict with the specifications of types and maintains its meaning throughout). Proofs of all the statements that follow are provided in the appendix.

We are now in a position to establish the key properties of well-typed programs containing graded modal session types. Our semantics defined over global configurations that was presented in Section 5 allows us to demonstrate that deadlock freedom holds by proving two main theorems: progress and preservation.

Lemma 1. Value lemma

Given $\Gamma \vdash v : A$ then, depending on the type, the shape of v can be inferred:

- $A = A' \multimap B$ then $v = \lambda x.v'$ or a partially applied primitive term p .
- $A = \square_r A'$ then $v = [v']$.
- $A = A' \otimes B$ then $v = (v_1, v_2)$.
- $A = 1$ then $v = \mathbf{unit}$.
- $A = \mathbf{Chan} P$ then $v = c$.

Note that as in other session-typed systems, typability of an expression is not sufficient to guarantee that it can immediately make progress at runtime [12]. For example, a channel with type $\mathbf{Chan}(\mathbf{Recv} A P)$ can only make progress if the dual channel has already sent the value that this channel needs to receive (as described by the \mathbf{RECV} rule). Our progress theorem applies to well-typed configurations that also satisfy the following property.

Definition 3. Given $\Delta \vdash t : A \wedge \Delta \vdash \mathcal{C}$ then we say that \mathcal{C} is *buffered* with respect to Δ , written $\text{buffered}(\mathcal{C}, \Delta)$, if for any channel c in \mathcal{C} with type $\text{Chan}(\text{Recv } AP)$, we have that c contains at least one value v .

Theorem 1 (Progress terms).

$$\begin{aligned}
& \Delta \vdash t : A \\
& \wedge \Delta \vdash \mathcal{C} \\
& \wedge \text{buffered}(\mathcal{C}, \Delta) \\
& \wedge \Delta \vdash \mathcal{P} : B \\
\implies & \text{value}(t) \vee \exists t', \mathcal{C}', \mathcal{P}'. (\mathcal{P} \mid t; \mathcal{C}) \rightsquigarrow (\mathcal{P}' \mid t'; \mathcal{C}')
\end{aligned}$$

A process configuration \mathcal{P} contains all values, defined by the predicate $\text{values}(\mathcal{P})$, if every process is either a value v or is a forwarder or watcher process.

Theorem 2 (Progress global config).

$$\begin{aligned}
& \Delta \vdash \mathcal{P} : A \wedge \Delta \vdash \mathcal{C} \wedge \text{buffered}(\mathcal{C}, \Delta) \\
\implies & \text{values}(\mathcal{P}) \vee \exists \mathcal{C}', \mathcal{P}'. (\mathcal{P}; \mathcal{C}) \rightsquigarrow (\mathcal{P}'; \mathcal{C}')
\end{aligned}$$

Proof. (Highlights) See Appendix C for the full proof.

Note that there are two separate progress theorems here; the first is akin to the standard progress theorem on terms, though some cases require additional information from the process and channel configurations and the overall result is that either the term t is a value or a global step can be taken which can reduce either or both of the two configurations along with the term. The second global progress theorem completes the deadlock freedom result by considering cases where we must reduce a process configuration without being able to focus on a term, for example if we need to reduce based on watcher and forwarder processes which only exist at runtime.

Most cases in the first progress theorem proceed exactly as they would in a standard proof of progress for a type system; in these cases the REDUCE rule from the global semantics allows us to say that if a term can be reduced in the standard way, then we can take a global step which does not affect the remainder of the process configuration or any of the channels which exist at runtime, simply reducing that term in place.

The interesting cases in this proof are the ones where a communication primitive is applied to a value, which by the value lemma (Lemma 1) must

be a channel. In these cases we cannot reduce by only considering the term; the rules we must use are the global semantic rules for the communication primitives, which take into account the channels that are present in the channel configuration and in some cases (such as FORKREPLICATE and FORKMULTICAST) also spawn additional processes.

The second proof is not too involved, and almost follows as a corollary of the first, but does require a little extra work for the cases of watcher and forwarder processes. Note in particular that it is possible for a forwarder process to still be present and for no reduction to be possible, but this is only in cases where all other terms are values, so the theorem does still hold here because $\text{values}(\mathcal{P})$ holds for a process configuration as long as all terms have been fully reduced. \square

Theorem 3 (Preservation).

$$\begin{aligned} & \Gamma, \Delta \vdash \mathcal{P} : A \wedge \Delta \vdash \mathcal{C} \wedge (\mathcal{P}; \mathcal{C}) \rightsquigarrow (\mathcal{P}'; \mathcal{C}') \\ \implies & \exists \Delta'. \Gamma, \Delta' \vdash \mathcal{P}' : A \wedge \Delta' \vdash \mathcal{C}' \end{aligned}$$

Proof. (Highlights) See Appendix D for the full proof.

The proof proceeds by induction on the structure of process configuration typing, in concert with channel configuration typing and then the inductive definition of global reductions that are possible at each typing.

Preservation for reduction of a single term comes from previous preservation results for Granule's core calculus [27].

For the concurrent primitives (both the communication primitives and forking primitives introduced here), their reductions involve terms in a reduction context. A lemma provides *factorisation* for a well-typed term in context, i.e., if $\text{Ctx}\{t\}$ is well-typed then $\text{Ctx}\{x\}$ is well-typed and t is well-typed separately (Lemma 6, p. 55). This lemma tells us that either x is a linear variable, i.e., the term t is substituted into a linear hole position, or x is a graded variable, i.e., the term t is substituted into a graded hole position, which for $\Gamma \vdash \text{Ctx}\{t\} : B$ then gives us:

$$\begin{aligned} & (\exists \Gamma_1, \Gamma_2, A, r. \neg \text{resourceAllocator}(t) \\ & \wedge \Gamma_1, x : [A]_r \vdash \text{Ctx}\{x\} : B \\ & \wedge [\Gamma_2] \vdash t : A \\ & \wedge \Gamma = \Gamma_1 + r * \Gamma_2) \end{aligned}$$

Critical here is that we know that this factorisation tells us that t is not a resource allocator. Given a global reduction which leads to a new term

t' , we then have the evidence to show that $\neg\text{resourceAllocator}(t')$ in these cases. *Context application* lemmas then give the dual result allowing us to reconstruct a well-typed term in context, i.e., $\text{Ctx}\{t'\}$ is well-typed (with preserved types).

Note that communication primitives (send, receive, close) are not resource allocators and can occur in graded contexts, in which case the factorisation lemma also states that their channels must therefore be shared, since $[\Gamma_2]$ classifies (runtime) contexts which are all graded or all shared channels. In these cases, the typing preservation is straightforward.

For the forking primitives, and auxiliary processes like forwarding and watching, the factorisation lemma tells us that these must necessarily be within a linear context as they are resource allocators, and thus given $\Gamma \vdash \text{Ctx}\{t\} : B$ we can factorise the typing into $\Gamma_1, x : A \vdash \text{Ctx}\{x\} : B$ and $\Gamma_2 \vdash t : A$ where $\Gamma = \Gamma_1 + \Gamma_2$. This shows us the effect of restricting terms to non-resource allocators in ensuring preservation. The type preservation then follows by the constructing the typing of the subject produced by reduction, which all satisfy the theorem statement.

For the `forkReplicate` case, the $\text{ReceivePrefix}(P)$ predicate is essential to type preservation: the watcher processes that `forkReplicate` reduces into rely on having receiving channels in order to continue reducing. Furthermore, the typing of a `forkReplicate` reduct has its first argument of type $\square_{0..n}(\mathbf{Chan} P \multimap 1)$. Since this argument is a value in the reduct, by inversion (the value lemma 1), it is necessarily of the form $[\lambda x.t]$ and we have $\Gamma_1, x : \mathbf{Chan} P, \Delta_1 \vdash t : 1$ and $0..n * (\Gamma_1, \Delta_1) \vdash [\lambda x.t] : (\mathbf{Chan} P \multimap 1)$ where Γ_1, Δ_1 is a graded context. Reduction yields n watchers, each of which holds a copy of t that by the typing of watchers promotes their resources by $0..1$ to $0..1 * (\Gamma_1, \Delta_1)$. Thus the n -wise parallel composition therefore (by distributivity of $*$ over $+$) results in preserving the $0..n * (\Gamma_1, \Delta_1)$ resources.

For the `forkMulticast` case, the $\text{Sends}(P)$ predicate is essential to type preservation also: the forwarding processes that `forkMulticast` reduce into rely on having a receiving channel endpoint and subsequent channels to send upon. Furthermore, the payload grading given by $\mathbf{Graded} n P$ ensures that the copying of the value to n channel endpoints is well-resourced. \square

10. Related work

Shared channels. Session type theories often provide non-linearity via *shared names* which provide repeated behaviours such as a replicable ‘server’ [42]

akin to our `forkReplicate` concept.

Particularly relevant in our context of linear types, Caires and Pfenning give a session-types meaning to intuitionistic linear logic in which propositions denote the types of (linear) channels and the exponential modality $!$ types the communication of shared servers which can be repeatedly spawned [7]; that is, $!A$ denotes the type of a non-linear shared channel which is used by a ‘server’ process to spawn new sessions conforming to the communication behaviour of A . The production of a new ‘client’ channel involves giving a semantic meaning to *derelection* where a non-linear (shared) channel has a linear copy ‘sharded’ off. In their formulation, $\Gamma; \Delta \vdash P :: T$ types a process P providing behaviour T along a channel, and using non-linear channels Γ and linear channels Δ . Access to the replicable server then involves the following process and typing:

$$\frac{\Gamma, u : A; \Delta, y : A \vdash P :: T}{\Gamma, u : A; \Delta \vdash (\nu y)u\langle y \rangle.P :: T}$$

i.e., access to the shared server via shared channel $u : A$ (essentially a $!A$ linear channel) is provided to P by the new linear channel $y : A$, amounting to contraction of u and derelection of u into y . This is coherent with a call-by-name semantics for linear logic where a value of type $!A$ is essentially an infinite number of copies of an A value, whereas in the call-by-value semantics $!A$ is just a single A value that can be used non-linearly. Therefore in our CBV approach, `forkReplicate` provides a vector of client channels to capture the notion of channel copies.

In the GV calculus, Lindley and Morris add replicated channels by first extending GV with an S4-style \Box modality (akin to the exponential modality) to provide “unlimited types” [21]. They derive a `replicate` primitive of the following type (where in GV, the type $!A.S$ denotes a channel on which we first send an A then behave as S , see Section 3.2):

$$\begin{aligned} \text{replicate} &: (!(\Box \bar{S}).\text{end}_!) \otimes (\Box(S \multimap \text{end}_!)) \multimap \text{end}_! \\ \text{replicate}(x, f) &= \text{send}([\text{let } [g] = f \text{ in fork } f], x) \end{aligned}$$

(we have used our notation for modal introduction and elimination to make the derivation more easily understood here).

Thus, replication takes a pair of (1) a channel x on which is sent a $\Box \bar{S}$ value, i.e., a non-linear channel which can behave as \bar{S} and (2) a non-linear function f which expects a channel behaving as S as input). The function

unboxes the non-linear function f into an instance g , and forks this, but underneath modal promotion, yielding a ‘boxed’ channel which is then sent on the input channel x . Whilst GV is call-by-value like Granule, the derivation above is only sound if \square behaves in a call-by-name way, i.e., we do not reduce under \square but instead capture the enclosing `fork` such that uses of it on the dual endpoint of x get a freshly forked function g each time. Thus, the semantics of \square in GV in this extension differs to Granule which is ‘pervasively’ call-by-value, evaluating inside of modalities.

For more detail, Appendix E gives a rendering of `replicate` in Granule, but using its call-by-name language extension for soundness.

Non-linear communication. Possibly the closest ideas to the ones described in our work that exist in the literature appear in the type system developed by Zalakain and Dardha, who use leftover typing to define a resource-aware variant of the π -calculus that can represent shared channels (which can be freely discarded or reused), graded channels (which must be used precisely n times for some n) and linear channels [43].

Applying their notion of ‘graded channel’ allows them to capture roughly the same behaviour as our `forkNonLinear` primitive, but their work does not consider more general forms of grading (though their mechanisation allows for user-defined resource algebras). We capture some of the same ideas inside our unified graded approach, showing how different amounts of sharing can be characterised precisely with the interaction of linear, indexed, and graded types; our setting is also more general, since we work in a pre-existing programming language rather than developing a calculus specifically focused on communication.

Another paper which describes some similar non-linear communication patterns to our work is the generalisation of session types via adjoint logic presented by Pruiksmá and Pfenning [30]. Here, they offer the capacity for both multicast communication and replicated servers, corresponding to our `forkMulticast` and `forkReplicate` primitives respectively. However, their system does not include any notion of grading, and so while they offer non-linear communication this is without the precise quantification allowed by grading in our system.

Classical session types. As well as the session-based interpretation of intuitionistic linear logic mentioned above [7], it is also possible to understand the correspondence between communication and logic by relying on *classical*

linear logic, as was first demonstrated by Wadler [39]. The differences between these two systems go beyond the superficial; it has been shown that there exists a class of processes which are typable following the classical interpretation but not via the intuitionistic interpretation, and that these are the processes which break the so-called *locality principle* for shared channels [36].

Recent work by Kokke, Morris and Wadler extends this interpretation through an approach which draws on notions from bounded linear logic [20]. While their notion of a shared channel corresponds closely to the one that we develop in our work (which may be unsurprising, since graded modal types are also built upon the foundation of bounded linear logic), their focus is more on accounting for non-determinism and races in a logical setting while preserving deadlock freedom. Their paper contrasts the system they develop with other related work on *manifest sharing* [2]; said work uses modal operators to describe the interaction between linear and shared processes, and has greater expressivity when it comes to non-linear communication, but this comes at the cost of reintroducing deadlocks in some circumstances.

Graded sessions in Linear Haskell. We explored combining grading with linear session types in Granule, but this could be extended to other settings. The linear types extension to Haskell is based on a calculus involving graded function arrows [4], and adding additional multiplicities to this system is a possibility. The Priority Sesh library provides a convenient embedding of the GV linear session calculus in Haskell [19], and extending this to make use of more precise information about channel usage could be valuable. This would also allow for experimenting with graded channels in a setting where grading can be implicit, rather than one where (like Granule) all grades are explicitly encoded via modalities.

11. Discussion and Conclusion

The key idea here is that graded types capture various standard non-linear forms of communication pattern atop the usual linear session types presentation. The demands of call-by-value required a different approach to past work on the interaction between graded types and session types, including the syntactic restriction discussed in Section 4 and vectors to capture multiplicity of channels. We have also proven type safety and various essential communication properties given the additional restriction on promoting channels and the other features introduced here, showing in particular that we can still enforce deadlock freedom.

Recursion and other combinators. A notable omission from our core session types calculus is the ability to define *recursive protocols*. However, some of the power of recursive session types is provided here; reusable channels (Section 6) are equivalent to linearly recursive session types (e.g., $\mu x.P.x$), although the expected usage has to be precisely specified in the grading since for example allowing the grade r to be instantiated to $0..∞$ to capture an arbitrary amount of use would break our guarantee of deadlock freedom. Further work is to integrate standard ideas on recursive session types and to explore their interaction with grading.

Note that combinations of the ideas here should also be possible; for example, combining multicast sending with reuse to get channels which we can repeatedly use to broadcast upon.

Applications of graded sessions. Using the linear channels already present in Granule it is possible to represent functions that are *sequentially realizable* [23]; a sequentially realizable function is one which has outwardly pure behaviour but relies on a notion of local side effects which are contained within the body of the function. Looking into which such behaviours may be more easily expressed by introducing non-linearity via graded channels would be an avenue for future work.

More recently, the GV session type calculus has been extended to allow for *multiparty* session types [18], allowing for concurrent communication between several participants at a time. It has also been shown that this multiparty system is strictly more general than the original binary GV system. It would be interesting to explore an implementation of this in Granule, particularly with a view to introducing grading into this setting to see if it can increase expressivity further.

Other resource allocators. As discussed in Section 4, it is necessary to restrict promotion of channels due to Granule’s default call-by-value semantics, since omitting this restriction allows for a linear channel to be used non-linearly. A very similar caveat applies if we consider mutable arrays. Recent work describes how to represent uniqueness of reference within Granule in addition to linearity [24]; if a unique array is promoted then we end up with multiple references to the same array, so we can no longer guarantee uniqueness for mutation. This is resolved by restricting promotion in much the same way as we have here for linear channels.

These are most likely not the only resource allocators for which such a restriction applies. Future work is to develop a better understanding of the

general case here, and to further explore the best way to embed interfaces for allocating resources in a graded type system such as Granule’s which can express multiple different evaluation strategies.

Uniqueness and graded sessions. Uniqueness in a concurrent setting was considered in various works by de Vries et al., but in particular we draw attention to their paper on resource management via unique and affine channels [10]. They allow for a notion of a channel which we can guarantee unique access to after i communication steps where i is some natural number; this is a form of grading, though quite different from the ideas we have discussed. It would be interesting to explore how we can represent this in Granule which already has both uniqueness and graded session types, and whether we gain further expressivity from doing so. This idea closely relates to recent work by Rocha and Caires on introducing mutable state into the classical linear logic interpretation of session types [31], through typing rules inspired by differential linear logic.

Acknowledgments. Thanks to the participants of PLACES 2022 for their comments on the workshop version of this paper and to the anonymous reviews of both versions of this paper. Thanks also to Paulo Torrens for his comments regarding resource allocation. Thanks to the rest of the Granule team for their feedback and overall work on the research vehicle that is the Granule programming language. This work was supported by an EPSRC Doctoral Training Award (Marshall) and EPSRC grant EP/T013516/1 (*Verifying Resource-like Data Use in Programs via Types*). The second author is also supported by the generosity of Eric and Wendy Schmidt by recommendation of the Schmidt Futures program.

References

- [1] A. Abel and J. Bernardy. A unified view of modalities in type systems. *Proc. ACM Program. Lang.*, 4(ICFP):90:1–90:28, 2020. doi: 10.1145/3408972.
- [2] S. Balzer and F. Pfenning. Manifest sharing with session types. *Proc. ACM Program. Lang.*, 1(ICFP), aug 2017. doi: 10.1145/3110281. URL <https://doi.org/10.1145/3110281>.

- [3] M. Berger, K. Honda, and N. Yoshida. Sequentiality and the π -calculus. In *International Conference on Typed Lambda Calculi and Applications*, pages 29–45. Springer, 2001. doi: 10.1007/3-540-45413-6_7.
- [4] J.-P. Bernardy, M. Boespflug, R. R. Newton, S. Peyton Jones, and A. Spiwack. Linear haskell: Practical linearity in a higher-order polymorphic language. *Proc. ACM Program. Lang.*, 2(POPL), dec 2017. doi: 10.1145/3158093.
- [5] E. C. Brady. *Idris 2: Quantitative type theory in practice*, 2021.
- [6] A. Brunel, M. Gaboardi, D. Mazza, and S. Zdancewic. A core quantitative effect calculus. In *Proceedings of the 23rd European Symposium on Programming Languages and Systems - Volume 8410*, page 351–370, Berlin, Heidelberg, 2014. Springer-Verlag. ISBN 9783642548321. doi: 10.1007/978-3-642-54833-8_19.
- [7] L. Caires and F. Pfenning. Session types as intuitionistic linear propositions. In *Proceedings of the 21st International Conference on Concurrency Theory, CONCUR’10*, page 222–236, Berlin, Heidelberg, 2010. Springer-Verlag. ISBN 3642153747. doi: 10.1007/978-3-642-15375-4_16.
- [8] M. Carbone and R. Neykova, editors. *Proceedings of the 13th International Workshop on Programming Language Approaches to Concurrency and Communication-cEntric Software, PLACES@ETAPS 2022, Munich, Germany, 3rd April 2022*, volume 356 of *EPTCS*, 2022. doi: 10.4204/EPTCS.356.
- [9] P. Choudhury, H. E. III, R. A. Eisenberg, and S. Weirich. A graded dependent type system with a usage-aware semantics. *Proc. ACM Program. Lang.*, 5(POPL):1–32, 2021. doi: 10.1145/3434331.
- [10] E. de Vries, A. Francalanza, and M. Hennessy. Uniqueness typing for resource management in message-passing concurrency. *Journal of Logic and Computation*, 24(3):531–556, 06 2012. ISSN 0955-792X. doi: 10.1093/logcom/exs022.
- [11] M. Gaboardi, S.-y. Katsumata, D. Orchard, F. Breuvar, and T. Uustalu. Combining effects and effects via grading. In *Proceedings of the 21st ACM SIGPLAN International Conference on Functional Programming, ICFP 2016*, page 476–489, New York, NY, USA,

2016. Association for Computing Machinery. ISBN 9781450342193. doi: 10.1145/2951913.2951939.
- [12] S. J. Gay and V. T. Vasconcelos. Linear type theory for asynchronous session types. *J. Funct. Program.*, 20(1):19–50, 2010. doi: 10.1017/S0956796809990268.
 - [13] D. R. Ghica and A. I. Smith. Bounded linear types in a resource semiring. In *Proceedings of the 23rd European Symposium on Programming Languages and Systems - Volume 8410*, page 331–350, Berlin, Heidelberg, 2014. Springer-Verlag. ISBN 9783642548321. doi: 10.1007/978-3-642-54833-8_18.
 - [14] J.-Y. Girard. Linear logic. *Theoretical Computer Science*, 50(1):1 – 101, 1987. ISSN 0304-3975. doi: 10.1016/0304-3975(87)90045-4. URL <http://www.sciencedirect.com/science/article/pii/0304397587900454>.
 - [15] J.-Y. Girard, A. Scedrov, and P. J. Scott. Bounded linear logic: a modular approach to polynomial-time computability. *Theoretical Computer Science*, 97(1):1–66, 1992. doi: 10.1007/978-1-4612-3466-1_11.
 - [16] K. Honda. Types for dyadic interaction. In E. Best, editor, *CONCUR’93*, pages 509–523, Berlin, Heidelberg, 1993. Springer Berlin Heidelberg. ISBN 978-3-540-47968-0. doi: 10.1007/3-540-57208-2_35.
 - [17] J. Hughes, D. Marshall, J. Wood, and D. Orchard. Linear Exponentials as Graded Modal Types. In *5th International Workshop on Trends in Linear Logic and Applications (TLLA 2021)*, Rome (virtual), Italy, June 2021. URL <https://hal-lirmm.ccsd.cnrs.fr/lirmm-03271465>.
 - [18] J. Jacobs, S. Balzer, and R. Krebbers. Multiparty gv: Functional multiparty session types with certified deadlock freedom. *Proc. ACM Program. Lang.*, 6(ICFP), aug 2022. doi: 10.1145/3547638.
 - [19] W. Kokke and O. Dardha. Deadlock-free session types in linear haskell. In *Proceedings of the 14th ACM SIGPLAN International Symposium on Haskell, Haskell 2021*, page 1–13, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450386159. doi: 10.1145/3471874.3472979.

- [20] W. Kokke, J. G. Morris, and P. Wadler. Towards races in linear logic. In *Coordination Models and Languages: 21st IFIP WG 6.1 International Conference, COORDINATION 2019, Held as Part of the 14th International Federated Conference on Distributed Computing Techniques, DisCoTec 2019, Kongens Lyngby, Denmark, June 17–21, 2019, Proceedings*, pages 37–53. Springer, 2019.
- [21] S. Lindley and J. G. Morris. A semantics for propositions as sessions. In *European Symposium on Programming Languages and Systems*, pages 560–584. Springer, 2015. doi: 10.1007/978-3-662-46669-8_23.
- [22] D. Marshall and D. Orchard. Replicate, reuse, repeat: Capturing non-linear communication via session types and graded modal types. In Carbone and Neykova [8], pages 1–11. doi: 10.4204/EPTCS.356.1.
- [23] D. Marshall and D. Orchard. How to Take the Inverse of a Type. In K. Ali and J. Vitek, editors, *36th European Conference on Object-Oriented Programming (ECOOP 2022)*, volume 222 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 5:1–5:27, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. ISBN 978-3-95977-225-9. doi: 10.4230/LIPIcs.ECOOP.2022.5. URL <https://drops.dagstuhl.de/opus/volltexte/2022/16233>.
- [24] D. Marshall, M. Vollmer, and D. Orchard. Linearity and uniqueness: An entente cordiale. In I. Sergey, editor, *Programming Languages and Systems*, pages 346–375, Cham, 2022. Springer International Publishing. ISBN 978-3-030-99336-8. doi: 10.1007/978-3-030-99336-8_13.
- [25] C. McBride. I got plenty o’ nuttin’. In S. Lindley, C. McBride, P. W. Trinder, and D. Sannella, editors, *A List of Successes That Can Change the World - Essays Dedicated to Philip Wadler on the Occasion of His 60th Birthday*, volume 9600 of *Lecture Notes in Computer Science*, pages 207–233. Springer, 2016. doi: 10.1007/978-3-319-30936-1_12. URL https://doi.org/10.1007/978-3-319-30936-1_12.
- [26] B. Moon, H. E. III, and D. Orchard. Graded modal dependent type theory. In N. Yoshida, editor, *Programming Languages and Systems - 30th European Symposium on Programming, ESOP 2021, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2021, Luxembourg City, Luxembourg, March 27 - April 1, 2021*,

- Proceedings*, volume 12648 of *Lecture Notes in Computer Science*, pages 462–490. Springer, 2021. doi: 10.1007/978-3-030-72019-3_17.
- [27] D. Orchard, V.-B. Liepelt, and H. Eades III. Quantitative program reasoning with graded modal types. *Proceedings of the ACM on Programming Languages*, 3(ICFP):1–30, 2019. doi: 10.1145/3341714.
- [28] T. Petricek, D. Orchard, and A. Mycroft. Coeffects: A calculus of context-dependent computation. In *Proceedings of the 19th ACM SIGPLAN International Conference on Functional Programming, ICFP '14*, page 123–135, New York, NY, USA, 2014. Association for Computing Machinery. ISBN 9781450328739. doi: 10.1145/2628136.2628160.
- [29] S. Peyton Jones, D. Vytiniotis, S. Weirich, and G. Washburn. Simple unification-based type inference for gadts. In *ACM SIGPLAN Notices*, volume 41, pages 50–61. ACM, 2006.
- [30] K. Pruikma and F. Pfenning. A message-passing interpretation of adjoint logic. *Electronic Proceedings in Theoretical Computer Science*, 291: 60–79, Apr 2019. ISSN 2075-2180. doi: 10.4204/eptcs.291.6.
- [31] P. Rocha and L. Caires. Propositions-as-types and shared state. *Proc. ACM Program. Lang.*, 5(ICFP), aug 2021. doi: 10.1145/3473584. URL <https://doi.org/10.1145/3473584>.
- [32] D. Sangiorgi and D. Walker. *The π -calculus: a Theory of Mobile Processes*. Cambridge University Press, 2003.
- [33] A. Spiwack, C. Kiss, J.-P. Bernardy, N. Wu, and R. A. Eisenberg. Linearly qualified types: Generic inference for capabilities and uniqueness. *Proc. ACM Program. Lang.*, 6(ICFP), aug 2022. doi: 10.1145/3547626.
- [34] B. Toninho and N. Yoshida. Depending on session-typed processes. In C. Baier and U. D. Lago, editors, *Foundations of Software Science and Computation Structures - 21st International Conference, FOSSACS 2018, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2018, Thessaloniki, Greece, April 14-20, 2018, Proceedings*, volume 10803 of *Lecture Notes in Computer Science*, pages 128–145. Springer, 2018. doi: 10.1007/978-3-319-89366-2_7.

- [35] B. Toninho, L. Caires, and F. Pfenning. A decade of dependent session types. In N. Veltri, N. Benton, and S. Ghilezan, editors, *PPDP 2021: 23rd International Symposium on Principles and Practice of Declarative Programming, Tallinn, Estonia, September 6-8, 2021*, pages 3:1–3:3. ACM, 2021. doi: 10.1145/3479394.3479398.
- [36] B. van den Heuvel and J. A. Pérez. Session type systems based on linear logic: Classical versus intuitionistic. *PLACES@ ETAPS*, 314:1–11, 2020.
- [37] V. T. Vasconcelos. Fundamentals of session types. *Information and Computation*, 217:52–70, 2012. ISSN 0890-5401. doi: <https://doi.org/10.1016/j.ic.2012.05.002>. URL <https://www.sciencedirect.com/science/article/pii/S0890540112001022>.
- [38] P. Wadler. Linear Types Can Change the World! In M. Broy and C. B. Jones, editors, *Programming Concepts and Methods: Proceedings of the IFIP Working Group 2.2/2.3 Working Conference on Programming Concepts and Methods, Sea of Galilee, Israel, 2–5 April, 1990*, pages 561–581, Amsterdam, 1990. North-Holland. URL <https://homepages.inf.ed.ac.uk/wadler/topics/linear-logic.html#linear-types>.
- [39] P. Wadler. Propositions as sessions. *Journal of Functional Programming*, 24(2-3):384–418, 2014. doi: 10.1017/S095679681400001X.
- [40] D. Walker. Substructural type systems. *Advanced Topics in Types and Programming Languages*, pages 3–44, 2005.
- [41] J. Wood and R. Atkey. A framework for substructural type systems. In I. Sergey, editor, *Programming Languages and Systems - 31st European Symposium on Programming, ESOP 2022, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2022, Munich, Germany, April 2-7, 2022, Proceedings*, volume 13240 of *Lecture Notes in Computer Science*, pages 376–402. Springer, 2022. doi: 10.1007/978-3-030-99336-8_14.
- [42] N. Yoshida and V. T. Vasconcelos. Language primitives and type discipline for structured communication-based programming revisited: Two systems for higher-order session communication. *Electronic Notes in Theoretical Computer Science*, 171(4):73–93, 2007. doi: 10.1016/j.entcs.2007.02.056.

- [43] U. Zalakain and O. Dardha. π with leftovers: A mechanisation in agda. In K. Peters and T. A. C. Willemse, editors, *Formal Techniques for Distributed Objects, Components, and Systems*, pages 157–174, Cham, 2021. Springer International Publishing. ISBN 978-3-030-78089-0. doi: 10.1007/978-3-030-78089-0_9.

Appendix A. Collected rules

Appendix A.1. Typing

$$\begin{array}{c}
\frac{}{x : A \vdash x : A} \text{VAR} \quad \frac{\Gamma, x : A \vdash t : B}{\Gamma \vdash \lambda x. t : A \multimap B} \multimap_i \quad \frac{\Gamma_1 \vdash t_1 : A \multimap B \quad \Gamma_2 \vdash t_2 : A}{\Gamma_1 + \Gamma_2 \vdash t_1 t_2 : B} \multimap_e \\
\\
\frac{\Gamma_1 \vdash t_1 : A \quad \Gamma_2 \vdash t_2 : B}{\Gamma_1 + \Gamma_2 \vdash (t_1, t_2) : A \otimes B} \otimes_i \quad \frac{\Gamma_1 \vdash t_1 : A \otimes A' \quad \Gamma_2, x : A, y : A' \vdash t_2 : B}{\Gamma_1 + \Gamma_2 \vdash \mathbf{let} (x, y) = t_1 \mathbf{in} t_2 : B} \otimes_e \\
\\
\frac{}{\emptyset \vdash \mathbf{unit} : 1} \text{UNIT}_i \quad \frac{\Gamma_1 \vdash t_1 : 1 \quad \Gamma_2 \vdash t_2 : B}{\Gamma_1 + \Gamma_2 \vdash \mathbf{let} \mathbf{unit} = t_1 \mathbf{in} t_2 : B} \text{UNIT}_e \\
\\
\frac{\Gamma, x : A \vdash t : B}{\Gamma, x : [A]_1 \vdash t : B} \text{DER} \quad \frac{\neg\text{resourceAllocator}(t) \quad [\Gamma] \vdash t : A}{r * \Gamma \vdash [t] : \square_r A} \text{PR} \quad \frac{\Gamma \vdash t : B}{\Gamma, 0 * \Gamma' \vdash t : B} \text{WEAK} \\
\\
\frac{\Gamma, x : [A]_r, \Gamma' \vdash t : B \quad r \sqsubseteq s}{\Gamma, x : [A]_s, \Gamma' \vdash t : B} \text{APPROX}
\end{array}$$

Definition 4 (Graded contexts). $[\Gamma]$ classifies those contexts which contain only graded variables:

$$\frac{}{[\emptyset]} \quad \frac{[\Gamma]}{[\Gamma, x : [A]_r]}$$

Definition 5 (Resource allocating terms). Predicate definition:

$$\begin{array}{c}
\frac{}{\text{resourceAllocator}(\text{forkLinear})} \quad \frac{}{\text{resourceAllocator}(\text{forkNonLinear})} \\
\frac{}{\text{resourceAllocator}(\text{forkReplicate})} \quad \frac{}{\text{resourceAllocator}(\text{forkMulticast})} \\
\frac{\text{resourceAllocator}(t)}{\text{resourceAllocator}([t])} \\
\frac{\text{resourceAllocator}(t_1)}{\text{resourceAllocator}(t_1 \ t_2)} \quad \frac{\text{resourceAllocator}(t_2)}{\text{resourceAllocator}(t_1 \ t_2)} \quad \frac{\text{resourceAllocator}(t_1)}{\text{resourceAllocator}((\lambda x. t_1) \ t_2)} \\
\frac{\text{resourceAllocator}(t_1)}{\text{resourceAllocator}(\text{let } [x] = t_1 \text{ in } t_2)} \quad \frac{\text{resourceAllocator}(t_2)}{\text{resourceAllocator}(\text{let } [x] = t_1 \text{ in } t_2)} \\
\frac{\text{resourceAllocator}(t_1)}{\text{resourceAllocator}(\text{let } \mathbf{unit} = t_1 \text{ in } t_2)} \quad \frac{\text{resourceAllocator}(t_2)}{\text{resourceAllocator}(\text{let } \mathbf{unit} = t_1 \text{ in } t_2)} \\
\frac{\text{resourceAllocator}(t_1)}{\text{resourceAllocator}(\text{let } (x, y) = t_1 \text{ in } t_2)} \quad \frac{\text{resourceAllocator}(t_2)}{\text{resourceAllocator}(\text{let } (x, y) = t_1 \text{ in } t_2)} \\
\frac{\text{resourceAllocator}(t_1)}{\text{resourceAllocator}((t_1, t_2))} \quad \frac{\text{resourceAllocator}(t_2)}{\text{resourceAllocator}((t_1, t_2))}
\end{array}$$

Primitives.

$$\begin{array}{c}
\frac{}{\emptyset \vdash \text{send} : A \otimes \mathbf{Chan Send } A P \multimap \mathbf{Chan } P} \text{ SEND} \\
\frac{}{\emptyset \vdash \text{receive} : \mathbf{Chan Recv } A P \multimap A \otimes \mathbf{Chan } P} \text{ RECV} \\
\frac{}{\emptyset \vdash \text{forkLinear} : (\mathbf{Chan } P \multimap 1) \multimap \mathbf{Chan } \overline{P}} \text{ FORK} \\
\frac{}{\emptyset \vdash \text{close} : \mathbf{Chan End } \multimap 1} \text{ CLOSE} \\
\frac{\text{SingleAction}(P) \quad \text{ExactSemiring}(S) \quad r : S}{\emptyset \vdash \text{forkNonLinear} : (\square_r \mathbf{Chan } P \multimap 1) \multimap \square_r(\mathbf{Chan } \overline{P})} \text{ FORKNONLIN} \\
\frac{\text{ReceivePrefix}(P)}{\emptyset \vdash \text{forkReplicate} : \square_{0..n}(\mathbf{Chan } P \multimap 1) \multimap \mathbf{N } n \multimap \mathbf{Vec } n (\square_{0..1}(\mathbf{Chan } \overline{P}))} \text{ FORKREP} \\
\frac{\text{Sends}(P)}{\emptyset \vdash \text{forkMulticast} : (\mathbf{Chan } (\text{Graded}_n P) \multimap 1) \multimap \mathbf{N } n \multimap \mathbf{Vec } n (\mathbf{Chan } \overline{P})} \text{ FORKMULTI}
\end{array}$$

Runtime typing.

$$\begin{array}{c}
\frac{}{c : \mathbf{Chan } P \vdash c : \mathbf{Chan } P} \text{ VAR} \quad \frac{}{c :^* \mathbf{Chan } P \vdash c : \mathbf{Chan } P} \text{ VARCHANS} \\
\frac{\Gamma \vdash t : B}{\Gamma, 0 * \Gamma', c :^* \mathbf{Chan } P \vdash t : B} \text{ WEAKCHANS}
\end{array}$$

In typing derivations in the appendix, we abstract over whether a channel is linear or shared by writing $c :^r \mathbf{Chan } P$.

Appendix A.2. Global configuration reduction rules

$$\begin{array}{c}
\frac{t \rightsquigarrow t'}{(\mathcal{P} \mid t^{\mathcal{R}}; \mathcal{C}) \rightsquigarrow (\mathcal{P} \mid t'^{\mathcal{R}}; \mathcal{C})} \text{REDUCE} \\
\frac{(\mathcal{P}; \mathcal{C}) \rightsquigarrow (\mathcal{P}'; \mathcal{C}')}{(\mathcal{P} \mid t; \mathcal{C}) \rightsquigarrow (\mathcal{P}' \mid t; \mathcal{C}')} \text{PAR} \\
\frac{}{(\mathcal{P} \mid \text{Ctxt}\{\text{send } c \ v\}; \mathcal{C}, \bar{c} = \vec{v}) \rightsquigarrow (\mathcal{P} \mid \text{Ctxt}\{c\}; \mathcal{C}, \bar{c} = \vec{v} : v)} \text{SEND} \\
\frac{}{(\mathcal{P} \mid \text{Ctxt}\{\text{receive } c\}; \mathcal{C}, c = v : \vec{v}) \rightsquigarrow (\mathcal{P} \mid \text{Ctxt}\{(v, c)\}; \mathcal{C}, c = \vec{v})} \text{RECV} \\
\frac{\#c}{(\mathcal{P} \mid \text{Ctxt}\{\text{forkLinear } (\lambda x.t)\}; \mathcal{C}) \rightsquigarrow (\mathcal{P} \mid t[c/x]^{\circ} \mid \text{Ctxt}\{\bar{c}\}; \mathcal{C}, \bar{c} = \langle \rangle, c = \langle \rangle)} \text{FORKLINEAR} \\
\frac{}{(\mathcal{P} \mid \text{Ctxt}\{\text{close } c\}; \mathcal{C}, c^* = \vec{v}) \rightsquigarrow (\mathcal{P} \mid \text{Ctxt}\{\mathbf{unit}\}; \mathcal{C}, c^* = \vec{v})} \text{SHARED}\text{CLOSE} \\
\frac{}{(\mathcal{P} \mid \text{Ctxt}\{\text{close } c\}; \mathcal{C}, c^1 = \langle \rangle) \rightsquigarrow (\mathcal{P} \mid \text{Ctxt}\{\mathbf{unit}\}; \mathcal{C})} \text{CLOSE} \\
\frac{\#c}{(\mathcal{P} \mid \text{Ctxt}\{\text{forkNonLinear } (\lambda x.t)\}; \mathcal{C}) \rightsquigarrow (\mathcal{P} \mid \text{Ctxt}\{[\bar{c}]\} \mid t[[c]/x]^{\circ}; \mathcal{C}, \bar{c}^* = \langle \rangle, c^* = \langle \rangle)} \text{FORKNONLINEAR} \\
\frac{\#c_1 \dots \#c_n}{(\mathcal{P} \mid \text{Ctxt}\{\text{forkReplicate } [\lambda x.t] \ n\}; \mathcal{C}) \rightsquigarrow (\mathcal{P} \mid \text{Ctxt}\{([\bar{c}_1], \dots, [\bar{c}_n])\} \mid \text{Wch}(c_1 \mapsto \lambda x.t) \mid \dots \mid \text{Wch}(c_n \mapsto \lambda x.t); \mathcal{C}, c_1 = \langle \rangle, \dots, c_n = \langle \rangle, \bar{c}_1 = \langle \rangle, \dots, \bar{c}_n = \langle \rangle)} \text{FORKREPLICATE} \\
\frac{\#c \quad \#c_1 \dots \#c_n}{(\mathcal{P} \mid \text{Ctxt}\{\text{forkMulticast } (\lambda x.t) \ n\}; \mathcal{C}) \rightsquigarrow (\mathcal{P} \mid \text{Ctxt}\{([\bar{c}_1], \dots, [\bar{c}_n])\} \mid t[c/x]^{\circ} \mid \text{Fwd}(c \mapsto c_1, \dots, c_n); \mathcal{C}, c = \langle \rangle, \bar{c} = \langle \rangle, c_1 = \langle \rangle, \dots, c_n = \langle \rangle, \bar{c}_1 = \langle \rangle, \dots, \bar{c}_n = \langle \rangle)} \text{FORKMULTICAST}
\end{array}$$

Appendix A.3. Standard small-step reduction rules

$$\begin{array}{c}
\frac{}{(\lambda x.t) v \rightsquigarrow t[v/x]} \rightsquigarrow_{\beta} \quad \frac{t_1 \rightsquigarrow t'_1}{t_1 t_2 \rightsquigarrow t'_1 t_2} \rightsquigarrow_{\text{APP}} \quad \frac{t \rightsquigarrow t'}{v t \rightsquigarrow v t'} \rightsquigarrow_{\text{prim}} \\
\\
\frac{}{\mathbf{let} (x, y) = (v_1, v_2) \mathbf{in} t_3 \rightsquigarrow t_3[v_1/x][v_2/y]} \rightsquigarrow_{\otimes\beta} \\
\\
\frac{t_1 \rightsquigarrow t'_1}{\mathbf{let} (x, y) = t_1 \mathbf{in} t_2 \rightsquigarrow \mathbf{let} (x, y) = t'_1 \mathbf{in} t_2} \rightsquigarrow_{\text{LET}\otimes} \\
\\
\frac{t_1 \rightsquigarrow t'_1}{(t_1, t_2) \rightsquigarrow (t'_1, t_2)} \rightsquigarrow_{\otimes\text{reduceL}} \quad \frac{t_2 \rightsquigarrow t'_2}{(v, t_2) \rightsquigarrow (v, t'_2)} \rightsquigarrow_{\otimes\text{reduceR}} \\
\\
\frac{}{\mathbf{let unit} = \mathbf{unit} \mathbf{in} t \rightsquigarrow t} \rightsquigarrow_{\beta\text{unit}} \quad \frac{t_1 \rightsquigarrow t'_1}{\mathbf{let unit} = t_1 \mathbf{in} t_2 \rightsquigarrow \mathbf{let unit} = t'_1 \mathbf{in} t_2} \rightsquigarrow_{\text{LETunit}} \\
\\
\frac{}{\mathbf{let} [x] = [v] \mathbf{in} t \rightsquigarrow t[v/x]} \rightsquigarrow_{\square\beta} \quad \frac{t_1 \rightsquigarrow t'_1}{\mathbf{let} [x] = t_1 \mathbf{in} t_2 \rightsquigarrow \mathbf{let} [x] = t'_1 \mathbf{in} t_2} \rightsquigarrow_{\text{LET}\square} \\
\\
\frac{t \rightsquigarrow t'}{[t] \rightsquigarrow [t']} \rightsquigarrow_{\square\text{reduce}}
\end{array}$$

Appendix B. Auxiliary results

Linear and graded substitution lemmas have been previously proved by Orchard et al. [27] for the core calculus used here:

Lemma 1 (Linear substitution). [27] *Given $\Gamma_1, x : A \vdash t_1 : B$ and $\Gamma_2 \vdash t_2 : A$ then $\Gamma_1 + \Gamma_2 \vdash t_1[t_2/x] : B$.*

Lemma 2 (Graded substitution). [27] *Given $\Gamma_1, x : [A]_r \vdash t_1 : B$ and $[\Gamma_2] \vdash t_2 : A$ and $\neg\text{resourceAllocator}(t_2)$ then $\Gamma_1 + r * \Gamma_2 \vdash t_1[t_2/x] : B$.*

Proof. Proof as in Orchard et al. [27], but where the additional premise of $\neg\text{resourceAllocator}(t_2)$ is deployed in the promotion case. \square

	$Ctxt$	$::=$		Contexts
			–	
			$Ctxt\ t$	
			$v\ Ctxt$	
			$[Ctxt]$	
Definition 6 (Term contexts).			let $[x] = Ctxt$ in t	
			let unit = $Ctxt$ in t	
			let $(x, y) = Ctxt$ in t	
			$(Ctxt, t)$	
			$(v, Ctxt)$	

Lemma 3 (Non-resource allocator context). $\neg\text{resourceAllocator}(Ctxt\{t\}) \implies \neg\text{resourceAllocator}(Ctxt\{x\})$

Proof. Trivial by the definition of $\text{resourceAllocator}(t)$ and since x is a variable and not a resource allocator. \square

Lemma 4 (Non-resource allocator composition). $\neg\text{resourceAllocator}(Ctxt\{x\}) \wedge \neg\text{resourceAllocator}(t) \implies \neg\text{resourceAllocator}(Ctxt\{t\})$

Proof. By induction on the definition of contexts and the definition of $\text{resourceAllocator}(t)$. \square

Lemma 5 (Contexts preserve resource allocator property).

$$\text{resourceAllocator}(t) \implies \text{resourceAllocator}(Ctxt\{t\})$$

Proof. By induction on the grammar of context terms: the structure of context terms matches the inductive term structure of resourceAllocator .

- $Ctxt = -$ by identity;
- $Ctxt = Ctxt'\ t'$. By induction on the first premise and modus ponens with the antecedent then $\text{resourceAllocator}(Ctxt'\{t\})$ and thus $\text{resourceAllocator}(Ctxt'\{t\}\ t')$ by the definition.
- $Ctxt = v\ Ctxt'$. By induction on the second premise and modus ponens with the antecedent then $\text{resourceAllocator}(Ctxt'\{t\})$ and thus $\text{resourceAllocator}(v\ Ctxt'\{t\})$ by the definition.
- $Ctxt = [Ctxt']$. By induction on the first premise and modus ponens with the antecedent then $\text{resourceAllocator}(Ctxt'\{t\})$ and thus $\text{resourceAllocator}([Ctxt'\{t\}])$ by the definition.

- $Ctxt = \mathbf{let} [x] = Ctxt' \mathbf{in} t_2$. By induction on the first premise and modus ponens with the antecedent then $\mathbf{resourceAllocator}(Ctxt'\{t\})$ and thus $\mathbf{resourceAllocator}(\mathbf{let} [x] = Ctxt'\{t\} \mathbf{in} t_2)$ by the definition.
- $Ctxt = \mathbf{let} \mathbf{unit} = Ctxt' \mathbf{in} t_2$. By induction on the first premise and modus ponens with the antecedent then $\mathbf{resourceAllocator}(Ctxt'\{t\})$ and thus $\mathbf{resourceAllocator}(\mathbf{let} \mathbf{unit} = Ctxt'\{t\} \mathbf{in} t_2)$ by the definition.
- $Ctxt = \mathbf{let} (x, y) = Ctxt' \mathbf{in} t_2$. By induction on the first premise and modus ponens with the antecedent then $\mathbf{resourceAllocator}(Ctxt'\{t\})$ and thus $\mathbf{resourceAllocator}(\mathbf{let} (x, y) = Ctxt'\{t\} \mathbf{in} t_2)$ by the definition.
- $Ctxt = (Ctxt', t')$. By induction on the first premise and modus ponens with the antecedent then $\mathbf{resourceAllocator}(Ctxt'\{t\})$ and thus $\mathbf{resourceAllocator}((Ctxt'\{t\}, t'))$ by the definition.
- $Ctxt = (v, Ctxt')$. By induction on the first premise and modus ponens with the antecedent then $\mathbf{resourceAllocator}(Ctxt'\{t\})$ and thus $\mathbf{resourceAllocator}((v, Ctxt'\{t\}))$ by the definition.

□

Corollary 1 (Non-resource-allocating contextual term).

$$\neg \mathbf{resourceAllocator}(Ctxt\{t\}) \implies \neg \mathbf{resourceAllocator}(t)$$

Proof. Dual of Lemma 5 (modus tollens)

□

Lemma 6 (Context factorisation).

$$\begin{aligned} & \Gamma \vdash Ctxt\{t\} : B \\ \implies & \quad (\exists \Gamma_1, \Gamma_2, A. \Gamma_1, x : A \vdash Ctxt\{x\} : B \\ & \quad \wedge \Gamma_2 \vdash t : A \\ & \quad \wedge \Gamma = \Gamma_1 + \Gamma_2) \\ \vee & \quad (\exists \Gamma_1, \Gamma_2, A, r. \neg \mathbf{resourceAllocator}(t) \\ & \quad \wedge \Gamma_1, x : [A]_r \vdash Ctxt\{x\} : B \\ & \quad \wedge [\Gamma_2] \vdash t : A \\ & \quad \wedge \Gamma = \Gamma_1 + r * \Gamma_2) \end{aligned}$$

Proof. By induction on $Ctxt$:

- $Ctxt = -$, i.e., the premise is that $\Gamma \vdash t : B$ and thus the result is given by $\Gamma_1 = \cdot$ and $\Gamma_2 = \Gamma$ with derivation:

$$\text{VAR } x : B \vdash x : B \quad (\text{ctxt})$$

and the substitutee derivation given by the premise. Thus, we prove this case by giving the first disjunct.

- $Ctxt = Ctxt' t'$ where $Ctxt\{t\} = Ctxt'\{t\} t'$ i.e., with premise derivation:

$$\frac{\Gamma'_1 \vdash Ctxt'\{t\} : A' \multimap B \quad \Gamma'_2 \vdash t' : A'}{\Gamma \vdash Ctxt'\{t\} t' : B} \multimap_e$$

where $\Gamma = \Gamma'_1 + \Gamma'_2$.

By induction on the first premise, then there are two cases:

- (first disjunct) We have Γ''_1, Γ''_2 such that $\Gamma''_1, x : A \vdash Ctxt'\{x\} : A' \multimap B$ and $\Gamma''_2 \vdash t : A$ such that $\Gamma'_1 = \Gamma''_1 + \Gamma''_2$.

We conclude with derivation:

$$\frac{\Gamma''_1, x : A \vdash Ctxt'\{x\} : A' \multimap B \quad \Gamma'_2 \vdash t' : A'}{(\Gamma''_1 + \Gamma'_2), x : A \vdash Ctxt'\{x\} t' : B} \multimap_e \quad (\text{ctxt})$$

i.e., where $\Gamma_1 = \Gamma''_1 + \Gamma'_2$ and $\Gamma_2 = \Gamma''_2$ where $\Gamma = \Gamma_1 + \Gamma_2 = \Gamma''_1 + \Gamma'_2 + \Gamma''_2 = \Gamma'_1 + \Gamma'_2$ from the induction and by commutativity and associativity of $+$.

The concluding substitutee $\Gamma''_2 \vdash t : A$ is from the induction.

- (second disjunct) We have Γ'_1, Γ''_2 such that $\neg\text{resourceAllocator}(t)$ and $\Gamma'_1, x : [A]_r \vdash Ctxt'\{x\} : A' \multimap B$ and $[\Gamma''_2] \vdash t : A$ such that $\Gamma'_1 = \Gamma''_1 + r * \Gamma''_2$.

We conclude with derivation:

$$\frac{\Gamma'_1, x : [A]_r \vdash Ctxt'\{x\} : A' \multimap B \quad \Gamma'_2 \vdash t' : A'}{(\Gamma''_1 + \Gamma'_2), x : A \vdash Ctxt'\{x\} t' : B} \multimap_e \quad (\text{ctxt})$$

i.e., where $\Gamma_1 = \Gamma''_1 + \Gamma'_2$ and where $\Gamma_2 = [\Gamma''_2]$ where:

$$\begin{aligned} \Gamma &= \Gamma''_1 + \Gamma'_2 && \{\text{premise}\} \\ &= \Gamma''_1 + r * \Gamma''_2 + \Gamma'_2 && \{\text{induction } \Gamma''_1\} \\ &= \Gamma''_1 + \Gamma'_2 + r * \Gamma''_2 && \{+ \text{ comm. and assoc.}\} \\ &= \Gamma_1 + r * \Gamma_2 && \{\text{concluding contexts}\} \end{aligned}$$

The concluding substitutee $[\Gamma_2''] \vdash t : A$ is from the induction and that $\neg\text{resourceAllocator}(t)$.

- $Ctxt = v Ctxt'$ where $Ctxt\{t\} = v Ctxt'\{t\}$ i.e., with premise derivation:

$$\frac{\Gamma_1' \vdash v : A' \multimap B \quad \Gamma_2' \vdash Ctxt'\{t\} : A'}{\Gamma \vdash v Ctxt'\{t\} : B} \multimap_e$$

where $\Gamma = \Gamma_1' + \Gamma_2'$.

By induction on the second premise, then there are two cases:

- (first disjunct) We have Γ_1'', Γ_2'' such that $\Gamma_1'', x : A \vdash Ctxt'\{x\} : A'$ and $\Gamma_2'' \vdash t : A$ such that $\Gamma_2'' = \Gamma_1'' + \Gamma_2''$. We can then conclude with:

$$\frac{\Gamma_1' \vdash v : A' \multimap B \quad \Gamma_1'', x : A \vdash Ctxt'\{x\} : A'}{\Gamma_1' + \Gamma_1'', x : A \vdash v Ctxt'\{x\} : B} \multimap_e \quad (\text{ctxt})$$

and substitutee $\Gamma_2'' \vdash t : A$ (from the induction) where $\Gamma_1 = \Gamma_1' + \Gamma_1''$ and $\Gamma_2 = \Gamma_2''$ where $\Gamma = \Gamma_1 + \Gamma_2 = \Gamma_1' + \Gamma_1'' + \Gamma_2''$ from the induction and by commutativity and associativity of $+$.

- (second disjunct) We have Γ_1'', Γ_2'' such that $\neg\text{resourceAllocator}(t)$ and $\Gamma_1'', x : [A]_r \vdash Ctxt'\{x\} : A'$ and $[\Gamma_2''] \vdash t : A$ such that $\Gamma_2'' = \Gamma_1'' + r * \Gamma_2''$.

We can then conclude with:

$$\frac{\Gamma_1' \vdash v : A' \multimap B \quad \Gamma_1'', x : [A]_r \vdash Ctxt'\{x\} : A'}{\Gamma_1' + \Gamma_1'', x : [A]_r \vdash v Ctxt'\{x\} : B} \multimap_e \quad (\text{ctxt})$$

and $\neg\text{resourceAllocator}(t)$ (from the induction) and substitutee $[\Gamma_2''] \vdash t : A$ (from the induction) where $\Gamma_1 = \Gamma_1' + \Gamma_1''$ and $\Gamma_2 = \Gamma_2''$ where:

$$\begin{aligned} \Gamma &= \Gamma_1' + \Gamma_2'' && \{\text{premise}\} \\ &= \Gamma_1' + \Gamma_1'' + r * \Gamma_2'' && \{\text{induction } \Gamma_1''\} \\ &= \Gamma_1 + r * \Gamma_2 && \{\text{concluding contexts}\} \end{aligned}$$

- $Ctxt = [Ctxt']$ where $Ctxt\{t\} = [Ctxt'\{t\}]$ i.e., with premise derivation:

$$\frac{\neg\text{resourceAllocator}(Ctxt'\{t\}) \quad [\Gamma_1'] \vdash Ctxt'\{t\} : A'}{\Gamma \vdash [Ctxt'\{t\}] : \square_r A'} \text{PR}'$$

where $\Gamma = r * \Gamma'_1$.

By induction on the second premise, there are two cases:

- (first disjunct) We have Γ''_1, Γ''_2 such that $\Gamma''_1, x : A \vdash \text{Ctx}'\{x\} : A'$ and $\Gamma''_2 \vdash t : A$ such that $\Gamma'_1 = \Gamma''_1 + \Gamma''_2$, and we know that $[\Gamma''_1]$ since the premise derivation gives us that $[\Gamma'_1]$.

We can then conclude with the derivation (giving the second disjunct):

$$\frac{\frac{\text{Lemma 3}}{\neg\text{resourceAllocator}(\text{Ctx}'\{x\})} \quad \frac{\frac{\text{ind.hyp.}}{[\Gamma''_1], x : A \vdash \text{Ctx}'\{x\} : A'}{\text{DER}}}{[\Gamma''_1], x : [A]_1 \vdash \text{Ctx}'\{x\} : A'}}{r * \Gamma''_1, x : [A]_r \vdash [\text{Ctx}'\{x\}] : \Box_r A'} \text{PR} \quad (\text{ctxt})$$

with substitutee $\Gamma''_2 \vdash t : A$ (from the induction) and with $\Gamma_1 = r * \Gamma''_1$ and $\Gamma_2 = \Gamma''_2$ where:

$$\begin{aligned} \Gamma &= r * \Gamma'_1 && \{\text{premise}\} \\ &= r * (\Gamma''_1 + \Gamma''_2) && \{\text{induction } \Gamma'_1\} \\ &= r * \Gamma''_1 + r * \Gamma''_2 && \{*\text{ dist.}\} \\ &= \Gamma_1 + r * \Gamma_2 && \{\text{concluding contexts}\} \end{aligned}$$

From the premise derivation $\neg\text{resourceAllocator}(\text{Ctx}'\{t\})$, by Corollary 1, we then have that $\neg\text{resourceAllocator}(t)$ as needed here.

- (second disjunct) We have $\neg\text{resourceAllocator}(t)$ and Γ''_1, Γ''_2 such that $\Gamma''_1, x : [A]_s \vdash \text{Ctx}'\{x\} : A'$ and $[\Gamma''_2] \vdash t : A$ such that $\Gamma'_1 = \Gamma''_1 + s * \Gamma''_2$, and we know that $[\Gamma''_1]$ since the premise derivation gives us that $[\Gamma'_1]$.

We can then conclude with the derivation (giving the second disjunct):

$$\frac{\frac{\text{Lemma 3}}{\neg\text{resourceAllocator}(\text{Ctx}'\{x\})} \quad \frac{\text{ind.hyp.}}{[\Gamma''_1], x : [A]_s \vdash \text{Ctx}'\{x\} : A'}}{r * \Gamma''_1, x : [A]_{r*s} \vdash [\text{Ctx}'\{x\}] : \Box_r A'} \text{PR} \quad (\text{ctxt})$$

with substitutee $[\Gamma_2''] \vdash t : A$ (from the induction) and with $\Gamma_1 = r * \Gamma_1''$ and $\Gamma_2 = \Gamma_2''$ where:

$$\begin{aligned}
\Gamma &= r * \Gamma_1' && \{\text{premise}\} \\
&= r * (\Gamma_1'' + s * \Gamma_2'') && \{\text{induction } \Gamma_1'\} \\
&= r * \Gamma_1'' + r * s * \Gamma_2'' && \{ * \text{ dist. } \} \\
&= \Gamma_1 + r * s * \Gamma_2 && \{\text{concluding contexts}\}
\end{aligned}$$

and $\neg\text{resourceAllocator}(t)$ from the induction, as required.

- $Ctxt = \mathbf{let} [y] = Ctxt' \mathbf{in} t_2$ where $Ctxt\{t\} = \mathbf{let} [y] = Ctxt'\{t\} \mathbf{in} t_2$.
i.e., with premise derivation:

$$\frac{\Gamma_1' \vdash Ctxt'\{t\} : \Box_r A' \quad \Gamma_2', y : [A']_r \vdash t_2 : B}{\Gamma_1' + \Gamma_2' \vdash \mathbf{let} [y] = Ctxt'\{t\} \mathbf{in} t_2 : B} \text{ELIM}$$

where $\Gamma = \Gamma_1' + \Gamma_2'$.

By induction on the first premise, then there are two case:

- (first disjunct) We have Γ_1'', Γ_2'' such that $\Gamma_1'', x : A \vdash Ctxt'\{x\} : A'$ and $\Gamma_2'' \vdash t : A$ such that $\Gamma_1' = \Gamma_1'' + \Gamma_2''$.

We can then conclude with:

$$\frac{\Gamma_1'', x : A \vdash Ctxt'\{x\} : \Box_r A' \quad \Gamma_2'', y : [A']_r \vdash t_2 : B}{(\Gamma_1'' + \Gamma_2''), x : A \vdash \mathbf{let} [y] = Ctxt'\{x\} \mathbf{in} t_2 : B} \text{ELIM}$$

and substitutee $\Gamma_2'' \vdash t : A$ (from the induction) where $\Gamma_1 = \Gamma_1'' + \Gamma_2''$ and $\Gamma_2 = \Gamma_2''$ therefore:

$$\begin{aligned}
\Gamma &= \Gamma_1'' + \Gamma_2'' && \{\text{premise}\} \\
&= \Gamma_1'' + \Gamma_2'' + \Gamma_2'' && \{\text{induction } \Gamma_1'\} \\
&= \Gamma_1'' + \Gamma_2'' + \Gamma_2'' && \{ + \text{ comm. and assoc. } \} \\
&= \Gamma_1 + \Gamma_2 && \{\text{concluding contexts}\}
\end{aligned}$$

- (second disjunct) We have $\neg\text{resourceAllocator}(t)$ and Γ_1'', Γ_2'' such that $\Gamma_1'', x : [A]_s \vdash Ctxt'\{x\} : A'$ and $[\Gamma_2''] \vdash t : A$ such that $\Gamma_1' = \Gamma_1'' + s * \Gamma_2''$.

We can then conclude with:

$$\frac{\Gamma_1'', x : [A]_s \vdash Ctxt'\{x\} : \Box_r A' \quad \Gamma_2'', y : [A']_r \vdash t_2 : B}{(\Gamma_1'' + \Gamma_2''), x : [A]_s \vdash \mathbf{let} [y] = Ctxt'\{x\} \mathbf{in} t_2 : B} \text{ELIM}$$

and substitutee $[\Gamma_2''] \vdash t : A$ (from the induction) where $\Gamma_1 = \Gamma_1'' + \Gamma_2'$ and $\Gamma_2 = \Gamma_2''$ therefore:

$$\begin{aligned}
\Gamma &= \Gamma_1' + \Gamma_2' && \{\text{premise}\} \\
&= \Gamma_1'' + s * \Gamma_2'' + \Gamma_2' && \{\text{induction } \Gamma_1'\} \\
&= \Gamma_1'' + \Gamma_2' + s * \Gamma_2'' && \{+ \text{ comm. and assoc.}\} \\
&= \Gamma_1 + s * \Gamma_2 && \{\text{concluding contexts}\}
\end{aligned}$$

and $\neg\text{resourceAllocator}(t)$ as required, from the induction.

- $Ctxt = \mathbf{let\ unit} = Ctxt' \mathbf{in} t_2$ where $Ctxt\{t\} = \mathbf{let\ unit} = Ctxt'\{t\} \mathbf{in} t_2$.
i.e., with premise derivation:

$$\frac{\Gamma_1' \vdash Ctxt'\{t\} : 1 \quad \Gamma_2' \vdash t_2 : B}{\Gamma_1 + \Gamma_2' \vdash \mathbf{let\ unit} = Ctxt'\{t\} \mathbf{in} t_2 : B} \text{ELIM}$$

where $\Gamma = \Gamma_1' + \Gamma_2'$.

By induction on the first premise, then there are two cases:

- (first disjunct) We have Γ_1'', Γ_2'' such that $\Gamma_1'', x : A \vdash Ctxt'\{x\} : 1$ and $\Gamma_2'' \vdash t : A$ such that $\Gamma_1' = \Gamma_1'' + \Gamma_2''$.

We can then conclude with:

$$\frac{\Gamma_1'', x : A \vdash Ctxt'\{x\} : 1 \quad \Gamma_2'' \vdash t_2 : B}{(\Gamma_1'' + \Gamma_2''), x : A \vdash \mathbf{let} [y] = Ctxt'\{x\} \mathbf{in} t_2 : B} \text{ELIM}$$

and substitutee $\Gamma_2'' \vdash t : A$ (from the induction) where $\Gamma_1 = \Gamma_1'' + \Gamma_2''$ and $\Gamma_2 = \Gamma_2''$ therefore:

$$\begin{aligned}
\Gamma &= \Gamma_1'' + \Gamma_2'' && \{\text{premise}\} \\
&= \Gamma_1'' + \Gamma_2'' + \Gamma_2'' && \{\text{induction } \Gamma_1'\} \\
&= \Gamma_1'' + \Gamma_2'' + \Gamma_2'' && \{+ \text{ comm. and assoc.}\} \\
&= \Gamma_1 + \Gamma_2 && \{\text{concluding contexts}\}
\end{aligned}$$

- (second disjunct) We have $\neg\text{resourceAllocator}(t)$ and Γ_1'', Γ_2'' such that $\Gamma_1'', x : [A]_s \vdash Ctxt'\{x\} : 1$ and $[\Gamma_2''] \vdash t : A$ such that $\Gamma_1' = \Gamma_1'' + s * \Gamma_2''$.

We can then conclude with:

$$\frac{\Gamma_1'', x : [A]_s \vdash Ctxt'\{x\} : 1 \quad \Gamma_2'' \vdash t_2 : B}{(\Gamma_1'' + \Gamma_2''), x : [A]_s \vdash \mathbf{let} [y] = Ctxt'\{x\} \mathbf{in} t_2 : B} \text{ELIM}$$

and substitutee $[\Gamma_2''] \vdash t : A$ (from the induction) where $\Gamma_1 = \Gamma_1'' + \Gamma_2'$ and $\Gamma_2 = \Gamma_2''$ therefore:

$$\begin{aligned}
\Gamma &= \Gamma_1' + \Gamma_2' && \{\text{premise}\} \\
&= \Gamma_1'' + s * \Gamma_2'' + \Gamma_2' && \{\text{induction } \Gamma_1'\} \\
&= \Gamma_1'' + \Gamma_2' + s * \Gamma_2'' && \{+ \text{ comm. and assoc.}\} \\
&= \Gamma_1 + s * \Gamma_2 && \{\text{concluding contexts}\}
\end{aligned}$$

and $\neg\text{resourceAllocator}(t)$ as required, from the induction.

- $Ctxt = \mathbf{let}(x, y) = Ctxt \mathbf{in} t_2$ the same as above, modulo the additional free variables for x and y in the scope of t_2 ;
- $Ctxt = (Ctxt', t')$ where $Ctxt\{t\} = (Ctxt'\{t\}, t')$. i.e., with premise derivation:

$$\frac{\Gamma_1' \vdash Ctxt'\{t\} : A' \quad \Gamma_2' \vdash t' : B'}{\Gamma_1' + \Gamma_2' \vdash (Ctxt'\{t\}, t') : A' \otimes B'} \otimes_i$$

where $\Gamma = \Gamma_1' + \Gamma_2'$.

By induction on the first premise, then there are two case:

- (first disjunct) We have Γ_1'', Γ_2'' such that $\Gamma_1'', x : A \vdash Ctxt'\{x\} : A'$ and $\Gamma_2'' \vdash t : A$ such that $\Gamma_1' = \Gamma_1'' + \Gamma_2''$.

We can then conclude with:

$$\frac{\Gamma_1'', x : A \vdash Ctxt'\{x\} : A' \quad \Gamma_2'' \vdash t : B'}{\Gamma_1'' + \Gamma_2'', x : A \vdash (Ctxt'\{t\}, t') : A' \otimes B'} \otimes_i$$

and substitutee $\Gamma_2'' \vdash t : A$ (from the induction) where $\Gamma_1 = \Gamma_1'' + \Gamma_2''$ and $\Gamma_2 = \Gamma_2''$ therefore:

$$\begin{aligned}
\Gamma &= \Gamma_1' + \Gamma_2' && \{\text{premise}\} \\
&= \Gamma_1'' + \Gamma_2'' + \Gamma_2' && \{\text{induction } \Gamma_1'\} \\
&= \Gamma_1'' + \Gamma_2' + \Gamma_2'' && \{+ \text{ comm. and assoc.}\} \\
&= \Gamma_1 + \Gamma_2 && \{\text{concluding contexts}\}
\end{aligned}$$

- (second disjunct) We have $\neg\text{resourceAllocator}(t)$ and Γ_1'', Γ_2'' such that $\Gamma_1'', x : [A]_s \vdash Ctxt'\{x\} : A'$ and $[\Gamma_2''] \vdash t : A$ such that $\Gamma_1' = \Gamma_1'' + s * \Gamma_2''$.

We can then conclude with:

$$\frac{\Gamma_1'', x : [A]_s \vdash \text{Ctx}'\{x\} : A' \quad \Gamma_2' \vdash t' : B'}{\Gamma_1'' + \Gamma_2', x : [A]_s \vdash (\text{Ctx}'\{t\}, t') : A' \otimes B'} \otimes_i$$

and substitutee $[\Gamma_2''] \vdash t : A$ (from the induction) where $\Gamma_1 = \Gamma_1'' + \Gamma_2'$ and $\Gamma_2 = \Gamma_2''$ therefore:

$$\begin{aligned} \Gamma &= \Gamma_1' + \Gamma_2' && \{\text{premise}\} \\ &= \Gamma_1'' + s * \Gamma_2'' + \Gamma_2' && \{\text{induction } \Gamma_1'\} \\ &= \Gamma_1'' + \Gamma_2' + s * \Gamma_2'' && \{+ \text{ comm. and assoc.}\} \\ &= \Gamma_1 + s * \Gamma_2 && \{\text{concluding contexts}\} \end{aligned}$$

and $\neg\text{resourceAllocator}(t)$ as required, from the induction.

- $\text{Ctx} = (v, \text{Ctx}')$ where $\text{Ctx}\{t\} = (v, \text{Ctx}'\{t\})$. i.e., with premise derivation:

$$\frac{\Gamma_1' \vdash v : A' \quad \Gamma_2' \vdash \text{Ctx}'\{t\} : B'}{\Gamma_1' + \Gamma_2' \vdash (v, \text{Ctx}'\{t\}) : A' \otimes B'} \otimes_i$$

where $\Gamma = \Gamma_1' + \Gamma_2'$.

By induction on the second premise, then there are two case:

- (first disjunct) We have Γ_1'', Γ_2'' such that $\Gamma_1'', x : A \vdash \text{Ctx}'\{x\} : A'$ and $\Gamma_2'' \vdash t : A$ such that $\Gamma_2' = \Gamma_1'' + \Gamma_2''$.

We can then conclude with:

$$\frac{\Gamma_1' \vdash v : A' \quad \Gamma_1'', x : A \vdash \text{Ctx}'\{x\} : B'}{(\Gamma_1' + \Gamma_1''), x : A \vdash (v, \text{Ctx}'\{x\}) : A' \otimes B'} \otimes_i$$

and substitutee $\Gamma_2'' \vdash t : A$ (from the induction) where $\Gamma_1 = \Gamma_1' + \Gamma_1''$ and $\Gamma_2 = \Gamma_2''$ therefore:

$$\begin{aligned} \Gamma &= \Gamma_1' + \Gamma_2' && \{\text{premise}\} \\ &= \Gamma_1' + \Gamma_1'' + \Gamma_2'' && \{\text{induction } \Gamma_1'\} \\ &= \Gamma_1 + \Gamma_2 && \{\text{concluding contexts}\} \end{aligned}$$

- (second disjunct) We have $\neg\text{resourceAllocator}(t)$ and Γ_1'', Γ_2'' such that $\Gamma_1'', x : [A]_s \vdash \text{Ctx}'\{x\} : A'$ and $[\Gamma_2''] \vdash t : A$ such that $\Gamma_2' = \Gamma_1'' + s * \Gamma_2''$.

We can then conclude with:

$$\frac{\Gamma'_1 \vdash v : A' \quad \Gamma''_1, x : [A]_s \vdash Ctxt'\{x\} : B'}{(\Gamma'_1 + \Gamma''_1), x : [A]_s \vdash (v, Ctxt'\{x\}) : A' \otimes B'} \otimes_i$$

and substitute $[\Gamma''_2] \vdash t : A$ (from the induction) where $\Gamma_1 = \Gamma'_1 + \Gamma''_1$ and $\Gamma_2 = \Gamma''_2$ therefore:

$$\begin{aligned} \Gamma &= \Gamma'_1 + \Gamma'_2 && \{\text{premise}\} \\ &= \Gamma'_1 + \Gamma''_1 + s * \Gamma''_2 && \{\text{induction } \Gamma'_1\} \\ &= \Gamma_1 + s * \Gamma_2 && \{\text{concluding contexts}\} \end{aligned}$$

and $\neg\text{resourceAllocator}(t)$ as required, from the induction. □

Lemma 7 (Context application (linear)).

$$\begin{aligned} &\Gamma_1, x : A \vdash Ctxt\{x\} : B \\ &\wedge \Gamma_2 \vdash t : A \\ \implies &\Gamma_1 + \Gamma_2 \vdash Ctxt\{t\} : B \end{aligned}$$

Proof. By induction on the structure of term contexts

- $Ctxt = -$, i.e., the premise is:

$$\frac{}{x : A \vdash x : A} \text{VAR} \quad (\text{ctxt})$$

(i.e., $\Gamma_1 = \cdot$ and $A = B$).

Then the conclusion follows by $Ctxt\{t\} = t$ therefore the second premise provides the conclusion.

- $Ctxt = Ctxt' t'$ where $Ctxt\{x\} = Ctxt'\{x\} t'$ i.e., with premise derivation:

$$\frac{\Gamma'_1, x : A \vdash Ctxt'\{x\} : A' \multimap B \quad \Gamma'_2 \vdash t' : A'}{(\Gamma'_1 + \Gamma'_2), x : A \vdash Ctxt'\{x\} t' : B} \multimap_e$$

By induction on the first premise, yielding: $\Gamma'_1 + \Gamma_2 \vdash Ctxt'\{t\} : A' \multimap B$ therefore the conclusion can be constructed as:

$$\frac{\Gamma'_1 + \Gamma_2 \vdash Ctxt'\{t\} : A' \multimap B \quad \Gamma'_2 \vdash t' : A'}{(\Gamma'_1 + \Gamma_2) + \Gamma'_2 \vdash Ctxt'\{t\} t' : B} \multimap_e$$

which by commutativity and associativity of $+$ matches the conclusion.

- $Ctxt = v Ctxt'$ where $Ctxt\{x\} = v Ctxt'\{x\}$ i.e., with premise derivation:

$$\frac{\Gamma'_1 \vdash v : A' \multimap B \quad \Gamma'_2, x : A \vdash Ctxt'\{x\} : A'}{(\Gamma'_1 + \Gamma'_2), x : A \vdash v Ctxt'\{x\} : B} \multimap_e$$

By induction on the second premise, yielding: $\Gamma'_2 + \Gamma_2 \vdash Ctxt'\{t\} : A'$ therefore the conclusion can be constructed as:

$$\frac{\Gamma'_1 \vdash v : A' \multimap B \quad \Gamma'_2 + \Gamma_2 \vdash Ctxt'\{t\} : A'}{(\Gamma'_1 + \Gamma'_2) + \Gamma_2 \vdash v Ctxt'\{t\} : B} \multimap_e$$

satisfying the goal of the lemma.

- $Ctxt = [Ctxt']$ where $Ctxt\{x\} = [Ctxt'\{x\}]$ i.e., with premise derivation:

$$\frac{[\Gamma'_1] \vdash Ctxt'\{x\} : A'}{r * \Gamma'_1 \vdash [Ctxt'\{x\}] : \Box_r A'} \text{PR}$$

where Γ'_1 must contain x . However, this leads to a contradiction since there cannot possibly be a linear assumption x in the conclusion due to the form of the typing. Therefore, the premise is false and this case is trivial.

- $Ctxt = \mathbf{let} [y] = Ctxt' \mathbf{in} t_2$ where $Ctxt\{x\} = \mathbf{let} [y] = Ctxt'\{x\} \mathbf{in} t_2$. i.e., with premise derivation:

$$\frac{\Gamma'_1, x : A \vdash Ctxt'\{x\} : \Box_r A' \quad \Gamma'_2, y : [A']_r \vdash t_2 : B}{(\Gamma'_1 + \Gamma'_2), x : A \vdash \mathbf{let} [y] = Ctxt'\{x\} \mathbf{in} t_2 : B} \text{ELIM}$$

By induction on the first premise, yielding: $\Gamma'_1 + \Gamma_2 \vdash Ctxt'\{t\} : \Box_r A'$. Therefore the conclusion can be constructed as:

$$\frac{\Gamma'_1 + \Gamma_2 \vdash Ctxt'\{t\} : \Box_r A' \quad \Gamma'_2, y : [A']_r \vdash t_2 : B}{(\Gamma'_1 + \Gamma_2) + \Gamma'_2 \vdash \mathbf{let} [y] = Ctxt'\{t\} \mathbf{in} t_2 : B} \text{ELIM}$$

which by commutativity and associativity of $+$ matches the goal.

- $Ctxt = \mathbf{let unit} = Ctxt' \mathbf{in} t_2$ where $Ctxt\{x\} = \mathbf{let unit} = Ctxt'\{x\} \mathbf{in} t_2$. i.e., with premise derivation:

$$\frac{\Gamma'_1, x : A \vdash Ctxt'\{x\} : 1 \quad \Gamma'_2 \vdash t_2 : B}{(\Gamma'_1 + \Gamma'_2), x : A \vdash \mathbf{let} [y] = Ctxt'\{x\} \mathbf{in} t_2 : B} \text{ELIM}$$

By induction on the first premise, yielding: $\Gamma'_1 + \Gamma_2 \vdash Ctxt'\{t\} : \square_r A'$.
Therefore the conclusion can be constructed as:

$$\frac{\Gamma'_1 + \Gamma_2 \vdash Ctxt'\{t\} : 1 \quad \Gamma'_2 \vdash t_2 : B}{(\Gamma'_1 + \Gamma_2) + \Gamma'_2 \vdash \mathbf{let} [y] = Ctxt'\{t\} \mathbf{in} t_2 : B} \text{ELIM}$$

which by commutativity and associativity of $+$ matches the goal.

- $Ctxt = \mathbf{let} (x, y) = Ctxt \mathbf{in} t_2$ the same as above, modulo the additional free variables for x and y in the scope of t_2 ;
- $Ctxt = (Ctxt', t')$ where $Ctxt\{x\} = (Ctxt'\{x\}, t')$. i.e., with premise derivation:

$$\frac{\Gamma'_1, x : A \vdash Ctxt'\{x\} : A' \quad \Gamma'_2 \vdash t' : B'}{(\Gamma'_1 + \Gamma'_2), x : A \vdash (Ctxt'\{x\}, t') : A' \otimes B'} \otimes_i$$

By induction on the first premise then $\Gamma'_1 + \Gamma_2 \vdash Ctxt'\{t\} : A'$. Therefore the goal can be constructed as:

$$\frac{\Gamma'_1 + \Gamma_2 \vdash Ctxt'\{t\} : A' \quad \Gamma'_2 \vdash t' : B'}{(\Gamma'_1 + \Gamma_2) + \Gamma'_2 \vdash (Ctxt'\{x\}, t') : A' \otimes B'} \otimes_i$$

satisfying the goal by commutativity and associativity of $+$.

- $Ctxt = (v, Ctxt')$ where $Ctxt\{x\} = (v, Ctxt'\{x\})$. i.e., with premise derivation:

$$\frac{\Gamma'_1 \vdash v : A' \quad \Gamma'_2, x : A \vdash Ctxt'\{x\} : B'}{(\Gamma'_1 + \Gamma'_2), x : A \vdash (v, Ctxt'\{x\}) : A' \otimes B'} \otimes_i$$

By induction on the second premise we have $\Gamma'_2 + \Gamma_2 \vdash Ctxt'\{t\} : B'$.
Therefore the goal can be constructed as:

$$\frac{\Gamma'_1 \vdash v : A' \quad \Gamma'_2 + \Gamma_2 \vdash Ctxt'\{t\} : B'}{\Gamma'_1 + (\Gamma'_2 + \Gamma_2) \vdash (v, Ctxt'\{x\}) : A' \otimes B'} \otimes_i$$

satisfying the goal by associativity of $+$.

□

Lemma 8 (Context application (graded)).

$$\begin{aligned}
& \Gamma_1, x : [A]_r \vdash \text{Ctx}\{x\} : B \\
& \wedge [\Gamma_2] \vdash t : A \\
& \wedge \neg \text{resourceAllocator}(t) \\
\implies & \Gamma_1 + r * \Gamma_2 \vdash \text{Ctx}\{t\} : B
\end{aligned}$$

Proof. By induction on the structure of term contexts and their typing.

Note that for all possible contexts, it is possible that typing is by (non-syntax directed) (approx) rule and thus we have:

$$\frac{\Gamma_1, x : [A]_s \vdash t' : B \quad s \sqsubseteq r}{\Gamma_1, x : [A]_r \vdash t' : B} \text{ APPROX}$$

Then by induction on the premise we have $\Gamma_1 + s * \Gamma_2 \vdash t' : B$.

Let $\Gamma_2 = x_i : [A_i]_{s_i}$ for $1 \leq i \leq |\Gamma_2|$ since by the premise every assumption is graded in the context. Then we can then construct the conclusion as follows:

$$\frac{\frac{\frac{}{s_i \sqsubseteq s_i} \text{ REFL} \quad s \sqsubseteq r}{s * s_i \sqsubseteq r * s_i} \text{ MONOTON.+UNIT}}{\Gamma_1 + s * \Gamma_2 \vdash t' : B} \text{ APPROX}^*$$

The rest is by induction now we have considered this possibility at each context term:

- $\text{Ctx} = -$.

There are therefore two possibilities for the derivation of the the premise:

- (var/der combination)

$$\frac{\frac{}{x : A \vdash x : A} \text{ VAR}}{x : [A]_1 \vdash x : A} \text{ DER}$$

(i.e., $\Gamma_1 = \cdot$ and $A = B$ and $r = 1$).

Since $1 * \Gamma_2 = \Gamma_2$ and since $\text{Ctx}\{t_2\} = t_2$ then we can conclude with the second premise $\Gamma_2 \vdash t_2 : B$

– (var/der/approx combination)

$$\frac{\frac{\overline{x : A \vdash x : A} \text{ VAR}}{x : [A]_1 \vdash x : A} \text{ DER} \quad 1 \sqsubseteq r}{x : [A]_r \vdash x : A} \text{ APPROX}$$

(i.e., $\Gamma_1 = \cdot$ and $A = B$).

Let $\Gamma_2 = x_i : [A]_{s_i}$ (since by the premise every assumption is graded in the context). Then we can then construct the conclusion as follows:

$$\frac{\overline{\Gamma_2 \vdash t : A} \text{ (SECOND PREMISE)} \quad \frac{\overline{s_i \sqsubseteq s_i} \text{ REFL} \quad 1 \sqsubseteq r}{s_i \sqsubseteq r * s_i} \text{ MONOTON. + UNIT}}{r * \Gamma_2 \vdash t : A} \text{ APPROX*}$$

- $Ctxt = Ctxt' t'$ where $Ctxt\{x\} = Ctxt'\{x\} t'$. i.e., with premise derivation:

$$\frac{\Gamma'_1, x : [A]_{s_1} \vdash Ctxt'\{x\} : A' \multimap B \quad \Gamma'_2, x : [A]_{s_2} \vdash t' : A'}{(\Gamma'_1 + \Gamma'_2), x : [A]_{s_1+s_2} \vdash Ctxt'\{x\} t' : B} \multimap_e$$

i.e., where $r = s_1 + s_2$. This derivation applies without loss of generality here for application since we can set $s_2 = 0$ in the case that $x \notin \text{FV}(t')$.

By induction on both premise we attain:

$$\Gamma'_1 + s_1 * \Gamma_2 \vdash Ctxt'\{t\} : A' \multimap B \quad (\text{ih1})$$

$$\Gamma'_2 + s_2 * \Gamma_2 \vdash t' : A' \quad (\text{ih2})$$

From which we construct the concluding judgment:

$$\frac{\Gamma'_1 + s_1 * \Gamma_2 \vdash Ctxt'\{t\} : A' \multimap B \quad \Gamma'_2 + s_2 * \Gamma_2 \vdash t' : A'}{\Gamma'_1 + s_1 * \Gamma_2 + \Gamma'_2 + s_2 * \Gamma_2 \vdash Ctxt'\{t\} t' : B} \multimap_e$$

which since $r = s_1 + s_2$, and by distributivity of $*$ over $+$, and commutativity and associativity of $+$ yields: $\Gamma'_1 + \Gamma'_2 + r * \Gamma_2 \vdash Ctxt'\{t\} t' : B$.

- $Ctxt = v Ctxt'$ where $Ctxt\{x\} = v Ctxt'\{x\}$ i.e., with premise derivation:

$$\frac{\Gamma'_1, x : [A]_{s_1} \vdash v : A' \multimap B \quad \Gamma'_2, x : [A]_{s_2} \vdash Ctxt'\{x\} : A'}{(\Gamma'_1 + \Gamma'_2), x : [A]_{s_1+s_2} \vdash v Ctxt'\{x\} : B} \multimap_e$$

This case proceeds as above by induction on both premises.

- $Ctxt = [Ctxt']$ where $Ctxt\{x\} = [Ctxt'\{x\}]$ i.e., with premise derivation:

$$\frac{\neg\text{resourceAllocator}(Ctxt'\{x\}) \quad [\Gamma'_1], x : [A]_{s_2} \vdash Ctxt'\{x\} : A'}{s_1 * (\Gamma'_1, x : [A]_{s_2}) \vdash [Ctxt'\{x\}] : \Box_r A'} \text{PR}'$$

i.e., where $r = s_1 * s_2$ and $\Gamma_1 = s_1 * \Gamma'_1$

By induction on the second premise we attain:

$$\Gamma'_1 + s_2 * \Gamma_2 \vdash Ctxt'\{t\} : A' \quad (\text{ih})$$

From which we construct the concluding judgment:

$$\frac{\text{Lemma 4} \quad \Gamma'_1 + s_2 * \Gamma_2 \vdash Ctxt'\{t\} : A'}{\neg\text{resourceAllocator}(Ctxt'\{t\}) \quad s_1 * (\Gamma'_1 + s_2 * \Gamma_2) \vdash [Ctxt'\{t\}] : \Box_r A'} \text{PR}'$$

which by distributivity of $*$ over $+$ and associativity of $*$ yields: $s_1 * \Gamma'_1 + (s_1 * s_2) * \Gamma_2 \vdash [Ctxt'\{t\}] : \Box_r A'$. and by $r = s_1 * s_2$ and $\Gamma_1 = s_1 * \Gamma'_1$ yields the goal: $\Gamma_1 + r * \Gamma_2 \vdash [Ctxt'\{t\}] : A'$

- $Ctxt = \mathbf{let} [y] = Ctxt' \mathbf{in} t_2$ where $Ctxt\{x\} = \mathbf{let} [y] = Ctxt'\{x\} \mathbf{in} t_2$. i.e., with premise derivation:

$$\frac{\Gamma'_1, x : A \vdash Ctxt'\{x\} : \Box_r A' \quad \Gamma'_2, y : [A]_r \vdash t_2 : B}{(\Gamma'_1 + \Gamma'_2), x : A \vdash \mathbf{let} [y] = Ctxt'\{x\} \mathbf{in} t_2 : B} \text{ELIM}$$

Proceeds by induction similarly to the (App-L) case, with no extra context manipulation necessary.

- $Ctxt = \mathbf{let unit} = Ctxt' \mathbf{in} t_2$ where $Ctxt\{x\} = \mathbf{let unit} = Ctxt'\{x\} \mathbf{in} t_2$. i.e., with premise derivation:

$$\frac{\Gamma'_1, x : A \vdash Ctxt'\{x\} : 1 \quad \Gamma'_2 \vdash t_2 : B}{(\Gamma'_1 + \Gamma'_2), x : A \vdash \mathbf{let} [y] = Ctxt'\{x\} \mathbf{in} t_2 : B} \text{ELIM}$$

Proceeds by induction similarly to the (App-L) case, with no extra context manipulation necessary.

- $Ctxt = \mathbf{let} (y, z) = Ctxt' \mathbf{in} t_2$ where $Ctxt\{x\} = \mathbf{let} (y, z) = Ctxt'\{x\} \mathbf{in} t_2$.
i.e., with premise derivation:

$$\frac{\Gamma'_1, x : A \vdash Ctxt'\{x\} : A_1 \otimes A_2 \quad \Gamma'_2, y : A_1, z : A_2 \vdash t_2 : B}{(\Gamma'_1 + \Gamma'_2), x : A \vdash \mathbf{let} (y, z) = Ctxt'\{x\} \mathbf{in} t_2 : B} \text{ELIM}$$

Proceeds by induction similarly to the (App-L) case, with no extra context manipulation necessary.

- $Ctxt = (Ctxt', t')$ where $Ctxt\{x\} = (Ctxt'\{x\}, t')$. i.e., with premise derivation:

$$\frac{\Gamma'_1, x : A \vdash Ctxt'\{x\} : A' \quad \Gamma'_2 \vdash t' : B'}{(\Gamma'_1 + \Gamma'_2), x : A \vdash (Ctxt'\{x\}, t') : A' \otimes B'} \otimes_i$$

Proceeds by induction similarly to the (App-L) case, with no extra context manipulation necessary.

- $Ctxt = (v, Ctxt')$ where $Ctxt\{x\} = (v, Ctxt'\{x\})$. i.e., with premise derivation:

$$\frac{\Gamma'_1 \vdash v : A' \quad \Gamma'_2, x : A \vdash Ctxt'\{x\} : B'}{(\Gamma'_1 + \Gamma'_2), x : A \vdash (v, Ctxt'\{x\}) : A' \otimes B'} \otimes_i$$

Proceeds by induction similarly to the (App-R) case, with no extra context manipulation necessary.

□

Lemma 9. *Value lemma*

Given $\Gamma \vdash v : A$ then, depending on the type, the shape of v can be inferred:

- $A = A' \multimap B$ then $v = \lambda x.v'$ or a partially applied primitive term p .
- $A = \square_r A'$ then $v = [v']$.
- $A = A' \otimes B$ then $v = (v_1, v_2)$.
- $A = 1$ then $v = \mathbf{unit}$.
- $A = \mathbf{Chan} P$ then $v = c$.

Proof. Recall that the value terms sub-grammar is:

$$v ::= (v_1, v_2) \mid \mathbf{unit} \mid [v] \mid \lambda x.v \mid c \mid p \quad (\text{value terms sub-grammar})$$

where p are partially-applied primitives. In other words,

$$\begin{aligned} p ::= & \text{close} \mid \text{send} \mid \text{send } c \mid \text{receive} \mid \text{forkLinear} \mid \text{forkNonLinear} \\ & \mid \text{forkReplicate} \mid \text{forkReplicate } c \\ & \mid \text{forkMulticast} \mid \text{forkMulticast } c \end{aligned}$$

We then proceed by case analysis on the type A to match the structure of the lemma. In each case we must consider what possible values can be assigned the type A and by which rules.

In all cases, there exist additional derivations based on dereliction, weakening and approximation, e.g., for the case where $A = A' \multimap B$:

$$\begin{aligned} & \frac{\Gamma, x : A' \vdash t : A' \multimap B}{\Gamma, x : \Box_1 A' \vdash t : A' \multimap B} \text{DER} \\ & \frac{\Gamma, x : A' \vdash t : A' \multimap B}{\Gamma, x : \Box_1 A', 0 * \Delta \vdash t : A' \multimap B} \text{DER} \\ & \frac{\Gamma, y : \Box_r A'', \Gamma', x : A' \vdash t : A' \multimap B \quad r \sqsubseteq s}{\Gamma, y : \Box_s A'', \Gamma', x : A' \vdash t : A' \multimap B} \text{DER} \end{aligned}$$

In all of these cases we can apply induction on the premise to get the result since the term is preserved between the premise and the conclusion.

We elide handling this separately each time in the cases that follow as the reasoning through dereliction is the same each time.

- $A = A' \multimap B$ then there are two classes of possible typing:

- Abstract term:

$$\frac{\Gamma, x : A' \vdash v' : B}{\Gamma \vdash \lambda x.v' : A' \multimap B} \multimap_i$$

thus $v = \lambda x.v'$ as in the lemma statement.

- Primitive term p formed by an application of zero or more values to a primitive operation, of which there are then ten possibilities:

1.

$$\overline{\emptyset \vdash \text{close} : \mathbf{Chan End} \multimap 1} \text{ CLOSE}$$

then $v = \text{close}$.

2.

$$\overline{\emptyset \vdash \text{send} : A \otimes \mathbf{Chan Send} A P \multimap \mathbf{Chan} P} \text{ SEND}$$

then $v = \text{send}$.

3.

$$\frac{\overline{c : \mathbf{Chan} P \vdash c : \mathbf{Chan} P} \text{ VAR} \quad \overline{\emptyset \vdash \text{send} : A \otimes \mathbf{Chan Send} A P \multimap \mathbf{Chan} P} \text{ SEND}}{\overline{c : \mathbf{Chan} P \vdash \text{send } c : A \multimap \mathbf{Chan} P} \multimap_e}$$

then $v = \text{send } c$.

4.

$$\overline{\emptyset \vdash \text{receive} : \mathbf{Chan Recv} A P \multimap A \otimes \mathbf{Chan} P} \text{ RECV}$$

then $v = \text{receive}$.

5.

$$\overline{\emptyset \vdash \text{forkLinear} : (\mathbf{Chan} P \multimap 1) \multimap \mathbf{Chan} \overline{P}} \text{ FORK}$$

then $v = \text{forkLinear}$.

6.

$$\overline{\emptyset \vdash \text{forkNonLinear} : (\square_r \mathbf{Chan} P \multimap 1) \multimap \square_r (\mathbf{Chan} \overline{P})} \text{ FORKNONLIN}$$

then $v = \text{forkNonLinear}$.

7.

$$\overline{\emptyset \vdash \text{forkReplicate} : \square_{0..n} (\mathbf{Chan} P \multimap 1) \multimap \mathbf{N} \ n \multimap \mathbf{Vec} \ n \ (\square_{0..1} (\mathbf{Chan} \overline{P}))} \text{ FORKREP}$$

then $v = \text{forkReplicate}$.

8.

$$\frac{\overline{c : \mathbf{Chan} P \vdash c : \mathbf{Chan} P} \text{ VAR} \quad \overline{c : \square_{0..1} \mathbf{Chan} P \vdash [c] : \square_{0..1} \mathbf{Chan} P} \text{ PR} \quad \overline{\emptyset \vdash \text{forkReplicate} : \square_{0..n} (\mathbf{Chan} P \multimap 1) \multimap \mathbf{N} \ n \multimap \mathbf{Vec} \ n \ (\square_{0..1} (\mathbf{Chan} \overline{P}))} \text{ FORKREP}}{\overline{c : \square_{0..1} \mathbf{Chan} P \vdash \text{forkReplicate} [c] : \mathbf{N} \ n \multimap \mathbf{Vec} \ n \ (\mathbf{Chan} \overline{P})} \multimap_e}$$

then $v = \text{forkReplicate} [c]$.

9.

$$\frac{\text{Sends}(P)}{\emptyset \vdash \text{forkMulticast} : (\mathbf{Chan}(\text{Graded}_n P) \multimap 1) \multimap \mathbf{N} \ n \multimap \mathbf{Vec} \ n \ (\mathbf{Chan} \overline{P})} \text{FORKMULTI}$$

then $v = \text{forkMulticast}$.

10.

$$\frac{\frac{c : \mathbf{Chan}(\text{Graded}_n P) \vdash c : \mathbf{Chan}(\text{Graded}_n P)}{\text{VAR}} \quad \frac{\text{Sends}(P)}{\emptyset \vdash \text{forkMulticast} : (\mathbf{Chan}(\text{Graded}_n P) \multimap 1) \multimap \mathbf{N} \ n \multimap \mathbf{Vec} \ n \ (\mathbf{Chan} \overline{P})} \text{FORKMULTI}}{c : \mathbf{Chan}(\text{Graded}_n P) \vdash \text{forkMulticast} \ c : \mathbf{N} \ n \multimap \mathbf{Vec} \ n \ (\mathbf{Chan} \overline{P})} \multimap_c$$

then $v = \text{forkMulticast} \ c$.

- $A = \square_r A'$ then there is only one possible non-dereliction typing of a value at that type:

$$\frac{\Gamma \vdash v' : A'}{r * \Gamma \vdash [v'] : \square_r A'} \text{PR}$$

thus $v = [v']$ as in the lemma statement.

- $A = A' \otimes B$ then there is only one possible non-dereliction typing of a value at that type:

$$\frac{\Gamma_1 \vdash v_1 : A' \quad \Gamma_2 \vdash v_2 : B}{\Gamma_1 + \Gamma_2 \vdash (v_1, v_2) : A' \otimes B} \otimes_i$$

thus $v = (v_1, v_2)$ as in the lemma statement.

- $A = 1$ then there is only one possible non-dereliction typing of a value at that type:

$$\frac{}{\emptyset \vdash \mathbf{unit} : 1} \text{UNIT}_i$$

thus $v = \mathbf{unit}$ as in the lemma statement.

- $A = \mathbf{Chan} \ P$ then the only possible typing that is a value is given by:

$$\frac{}{c : \mathbf{Chan} \ P \vdash c : \mathbf{Chan} \ P} \text{VAR}$$

thus $v = c$ as in the lemma statement, regardless of the protocol P .

□

Appendix C. Progress proof

Theorem 1 (Progress terms).

$$\begin{aligned}
& \Delta \vdash t : A \\
& \wedge \Delta \vdash \mathcal{C} \\
& \wedge \text{buffered}(\mathcal{C}, \Delta) \\
& \wedge \Delta \vdash \mathcal{P} : B \\
\implies & \text{value}(t) \vee \exists t', \mathcal{C}', \mathcal{P}'. (\mathcal{P} \mid t; \mathcal{C}) \rightsquigarrow (\mathcal{P}' \mid t'; \mathcal{C}')
\end{aligned}$$

Proof. We prove this by induction on the three judgments forming the assumptions.

First, we case split on the typing of process configurations: $\Delta \vdash \mathcal{P} : B$.

- (cons)

$$\frac{\Gamma_1, \Delta_1 \vdash \mathcal{P} : A \quad \Gamma_2, \Delta_2 \vdash t : B}{\Gamma_1 + \Gamma_2, \Delta_1 + \Delta_2 \vdash \mathcal{P} \mid t^{\mathcal{R}} : A} \text{CONSR}$$

In this case, we can write the process configuration \mathcal{P} as $\mathcal{P}' \mid t''$, or in other words as another process configuration \mathcal{P}' extended with some term t'' .

By induction, there are two possibilities here: either t'' is itself a value, or there exists some reduction $(\mathcal{P}'; \mathcal{C}) \rightsquigarrow (\mathcal{P}''; \mathcal{C}')$.

If there is such a reduction, then we can make a global reduction using the following rule:

$$\frac{(\mathcal{P}; \mathcal{C}) \rightsquigarrow (\mathcal{P}'; \mathcal{C}')}{(\mathcal{P} \mid t; \mathcal{C}) \rightsquigarrow (\mathcal{P}' \mid t; \mathcal{C}')} \text{PAR}$$

and the overall progress theorem is satisfied.

Otherwise, we must consider the possible reductions for the additional term t .

- (empty)

$$\frac{}{\emptyset \vdash \cdot : \perp} \text{EMPTY}$$

Since the process configuration is empty in this case, it cannot make any reductions, so again we only need to consider the possible reductions originating from the term t .

We then induct on the last step of the typing derivation $\Delta_1 \vdash t : A$. Note that in some cases the process and channel configurations do not come into play, in which case the proof proceeds roughly the same as the standard progress theorem on terms; either t is a value, in which case the theorem holds directly by induction, or we can take a step $t \rightsquigarrow t'$, in which case we will be able to take a global step via the REDUCE rule. In these cases we can generally ignore the presence of the runtime context Δ_1 .

- (var)

$$\frac{}{x : A \vdash x : A} \text{VAR}$$

Trivial since the variable x cannot be reduced and so is necessarily a value.

- (abs)

$$\frac{\Gamma, x : A \vdash t : B}{\Gamma \vdash \lambda x.t : A \multimap B} \multimap_i$$

Trivial since $(\lambda x.t)$ is already a value.

- (app)

$$\frac{\Gamma_1 \vdash t_1 : A \multimap B \quad \Gamma_2 \vdash t_2 : A}{\Gamma_1 + \Gamma_2 \vdash t_1 t_2 : B} \multimap_e$$

By induction on the first premise, there are two possibilities: either $t_1 = v$ for some value v or $t_1 \rightsquigarrow t'_1$ for some t'_1 .

In the first case, by the value lemma there are eleven cases to consider.

1. $t_1 = \lambda x.t'_1$. Here, there are two possibilities: either $t_2 = v$ for some value v or $t_2 \rightsquigarrow t'_2$ for some t'_2 .

In the first case, we can reduce by REDUCE using the reduction:

$$\frac{}{(\lambda x.t) v \rightsquigarrow t[v/x]} \rightsquigarrow_\beta$$

In the second case, since $\lambda x.t'_1$ is a value then by induction on the premise we can reduce by **REDUCE** using the reduction:

$$\frac{t \rightsquigarrow t'}{v t \rightsquigarrow v t'} \rightsquigarrow_{\text{prim}}$$

2. $t_1 = \text{close}$, where $A = \mathbf{ChanEnd}$. Therefore, we induct on the second argument. There are two cases.

* t_2 is a value. Then, by the value lemma, $t_2 = c$ and the runtime typing derivation has two possible (consequential) specialisations:

(a)

$$\frac{\frac{\cdot}{c : \mathbf{ChanEnd} \vdash c : \mathbf{ChanEnd}} \text{VAR} \quad \frac{\cdot}{\emptyset \vdash \text{close} : \mathbf{ChanEnd} \multimap 1} \text{CLOSE}}{c : \mathbf{ChanEnd} \vdash \text{close } c : 1} \multimap_e$$

Therefore, by Lemma 7 our overall succedent is of the form $\Delta'_2, c : \mathbf{ChanEnd} \vdash \text{Ctx}\{\text{close } c\} : A$. There is then a reduction as follows:

$$\frac{\cdot}{(\mathcal{P} \mid \text{Ctx}\{\text{close } c\}; \mathcal{C}, c^1 = \langle \rangle) \rightsquigarrow (\mathcal{P} \mid \text{Ctx}\{\mathbf{unit}\}; \mathcal{C})} \text{CLOSE}$$

(b)

$$\frac{\frac{\cdot}{c :^* \mathbf{ChanEnd} \vdash c : \mathbf{ChanEnd}} \text{VAR}_S \quad \frac{\cdot}{\emptyset \vdash \text{close} : \mathbf{ChanEnd} \multimap 1} \text{CLOSE}}{c :^* \mathbf{ChanEnd} \vdash \text{close } c : 1} \multimap_e$$

Therefore, by Lemma 7 our overall succedent is of the form $\Delta'_2, c :^* \mathbf{ChanEnd} \vdash \text{Ctx}\{\text{close } c\} : A$. There is then a reduction as follows:

$$\frac{\cdot}{(\mathcal{P} \mid \text{Ctx}\{\text{close } c\}; \mathcal{C}, c^* = \vec{v}) \rightsquigarrow (\mathcal{P} \mid \text{Ctx}\{\mathbf{unit}\}; \mathcal{C}, c^* = \vec{v})} \text{SHARED}_{\text{CLOSE}}$$

* t_2 is not a value, and has a reduction $t_2 \rightsquigarrow t'_2$ for some t'_2 . Then we can reduce by **REDUCE** using the reduction:

$$\frac{t_2 \rightsquigarrow t'_2}{\text{close } t_2 \rightsquigarrow \text{close } t'_2} \rightsquigarrow_{\text{prim}}$$

3. $t_1 = \mathbf{send}$, where $A = \mathbf{Chan}(\mathbf{Send} A' P) \multimap A' \multimap \mathbf{Chan} P$. Therefore, we induct on the second argument. There are two cases.

- * t_2 is a value. Then, by the value lemma, $t_2 = c$ and so $t_1 t_2 = \mathbf{send} c$ which is also a value.
- * t_2 is not a value, and has a reduction $t_2 \rightsquigarrow t'_2$ for some t'_2 . Then we can reduce by REDUCE using the reduction:

$$\frac{t_2 \rightsquigarrow t'_2}{\mathbf{send} t_2 \rightsquigarrow \mathbf{send} t'_2} \rightsquigarrow_{\text{prim}}$$

4. $t_1 = \mathbf{send} c$, where $A = A' \multimap \mathbf{Chan} P$. Therefore, we induct on the second argument. There are two cases.

- * $t_2 = x$ for some value x . The runtime typing derivation specialises to:

$$\frac{\frac{\dots}{c :^r \mathbf{Chan}(\mathbf{Send} A' P) \vdash \mathbf{send} c : A' \multimap \mathbf{Chan} P} \multimap_e \quad \frac{\cdot}{\emptyset \vdash v : A'} \text{VAR}}{c :^r \mathbf{Chan}(\mathbf{Send} A' P) \vdash \mathbf{send} c x : \mathbf{Chan} P} \multimap_e$$

Therefore, by Lemma 7 our overall succedent is of the form $\Delta'_2, c :^r \mathbf{Chan}(\mathbf{Send} A' P) \vdash \mathit{Ctx}\{\mathbf{send} c x\} : A$.

By inversion, channel configuration typing must be of the form:

$$\frac{\Delta_1 \vdash v : A \quad \Delta_1, \Delta_2, c :^r \mathbf{Chan} P \vdash \mathcal{C}, \bar{c} = \vec{v} : v}{\Delta_1, \Delta_2, c :^r \mathbf{Chan}(\mathbf{Send} A P) \vdash \mathcal{C}, \bar{c} = \vec{v}} \text{SEND}$$

with $\Delta = \Delta_1, \Delta_2$.

In this case, there is then a reduction as follows:

$$\frac{}{(\mathcal{P} \mid \mathit{Ctx}\{\mathbf{send} c v\}; \mathcal{C}, \bar{c} = \vec{v}) \rightsquigarrow (\mathcal{P} \mid \mathit{Ctx}\{c\}; \mathcal{C}, \bar{c} = \vec{v} : v)} \text{SEND}$$

- * t_2 is not a value, and has a reduction $t_2 \rightsquigarrow t'_2$ for some t'_2 . Then we can reduce by REDUCE using the reduction:

$$\frac{t_2 \rightsquigarrow t'_2}{\mathbf{send} c t_2 \rightsquigarrow \mathbf{send} c t'_2} \rightsquigarrow_{\text{prim}}$$

5. $t_1 = \text{receive}$, where $A = \mathbf{Chan}(\mathbf{Recv} A' P) \multimap A' \otimes \mathbf{Chan} P$. Therefore, we induct on the second argument. There are two cases.
- * t_2 is a value. Then, by the value lemma, $v_2 = c$ and the runtime typing derivation refines to:

$$\frac{\frac{\dots}{c :^r \mathbf{Chan}(\mathbf{Recv} A' P) \vdash c : \mathbf{Chan}(\mathbf{Recv} A' P)} \text{VAR} \quad \frac{}{\emptyset \vdash \text{receive} : \mathbf{Chan} \mathbf{Recv} A' P \multimap A \otimes \mathbf{Chan} P} \text{RECV}}{c :^r \mathbf{Chan}(\mathbf{Recv} A' P) \vdash \text{receive } c : A' \otimes \mathbf{Chan} P} \multimap_e$$

Therefore, by Lemma 7 our overall succedent is of the form $\Delta'_2, c :^r \mathbf{Chan}(\mathbf{Recv} A' P) \vdash \text{Ctx}\{\text{receive } c\} : A$. Now, since the channel configuration \mathcal{C} must be buffered with respect to Δ , the channel c cannot be empty, and so it must contain at least one value v . Hence, there is only one possibility for the channel configuration typing:

$$\frac{\Delta_1 \vdash v : A \quad \Delta_1, \Delta_2, c :^r \mathbf{Chan} P \vdash \mathcal{C}, c = \vec{v}}{\Delta_1, \Delta_2, c :^r \mathbf{Chan}(\mathbf{Recv} A' P) \vdash \mathcal{C}, c = v : \vec{v}} \text{RECV}$$

with $\Delta = \Delta_1, \Delta_2$.

Then there is a reduction as follows:

$$\frac{}{(\mathcal{P} \mid \text{Ctx}\{\text{receive } c\}; \mathcal{C}, c = v : \vec{v}) \rightsquigarrow (\mathcal{P} \mid \text{Ctx}\{(v, c)\}; \mathcal{C}, c = \vec{v})} \text{RECV}$$

- * t_2 is not a value, and has a reduction $t_2 \rightsquigarrow t'_2$ for some t'_2 . Then we can reduce by REDUCE using the reduction:

$$\frac{t_2 \rightsquigarrow t'_2}{\text{receive } t_2 \rightsquigarrow \text{receive } t'_2} \rightsquigarrow_{\text{prim}}$$

6. $t_1 = \text{forkLinear}$, where $A = (\mathbf{Chan} P \multimap 1) \multimap \mathbf{Chan} \bar{P}$. Therefore, we induct on the second argument. There are two cases.
- * t_2 is a value. Then, by the value lemma, $t_2 = \lambda x.t'_2$ for some t'_2 and we can construct the following runtime typing derivation:

$$\frac{\frac{x : \mathbf{Chan} P \vdash t'_2 : 1}{\emptyset \vdash \lambda x.t'_2 : \mathbf{Chan} P \multimap 1} \multimap_i \quad \frac{}{\emptyset \vdash \text{forkLinear} : (\mathbf{Chan} P \multimap 1) \multimap \mathbf{Chan} \bar{P}} \text{FORK}}{\emptyset \vdash \text{forkLinear } \lambda x.t'_2 : \mathbf{Chan} \bar{P}} \multimap_e$$

Therefore, by Lemma 7 our overall succedent is of the form $\Delta'_2 \vdash \text{Ctx}\{\text{forkLinear } \lambda x.t'_2\} : A$.
Then there is a reduction as follows:

$$\frac{\#c}{(\mathcal{P} \mid \text{Ctx}\{\text{forkLinear } (\lambda x.t)\}; \mathcal{C}) \rightsquigarrow (\mathcal{P} \mid t[c/x]^\circ \mid \text{Ctx}\{\bar{c}\}; \mathcal{C}, \bar{c} = \langle \rangle, c = \langle \rangle)} \text{FORKLINEAR}$$

where the channel configuration is typed $\Delta'_2 = \mathcal{C}$ and we select a fresh name c for the channel being forked.

- * t_2 is not a value, and has a reduction $t_2 \rightsquigarrow t'_2$ for some t'_2 .
Then we can reduce by REDUCE using the reduction:

$$\frac{t_2 \rightsquigarrow t'_2}{\text{forkLinear } t_2 \rightsquigarrow \text{forkLinear } t'_2} \rightsquigarrow_{\text{prim}}$$

7. $t_1 = \text{forkNonLinear}$, where $A = (\Box_r(\mathbf{Chan } P) \multimap 1) \multimap \Box_r(\mathbf{Chan } \bar{P})$.
Therefore, we induct on the second argument. There are two cases.

- * t_2 is a value. Then, by the value lemma, $t_2 = \lambda x.t'_2$ for some t'_2 and we can construct the following runtime typing derivation:

$$\frac{\frac{x : \Box_r(\mathbf{Chan } P) \vdash t'_2 : 1}{\emptyset \vdash \lambda x.t'_2 : \Box_r(\mathbf{Chan } P) \multimap 1} \multimap_i \quad \frac{\emptyset \vdash \text{forkNonLinear} : (\Box_r \mathbf{Chan } P \multimap 1) \multimap \Box_r(\mathbf{Chan } \bar{P})}{\emptyset \vdash \text{forkNonLinear } \lambda x.t'_2 : \Box_r(\mathbf{Chan } \bar{P})} \text{FORKNONLIN}}{\emptyset \vdash \text{forkNonLinear } \lambda x.t'_2 : \Box_r(\mathbf{Chan } \bar{P})} \multimap_e \quad (\text{C.1})$$

Note that this derivation cannot include any instances of the APPROX rule, because of the ExactSemiring constraint present in the type signature of `forkNonLinear`, and so the grade r must remain the same throughout the typing.

Therefore, by Lemma 7 our overall succedent is of the form $\Delta'_2 \vdash \text{Ctx}\{\text{forkNonLinear } \lambda x.t'_2\} : A$.
Then there is a reduction as follows:

$$\frac{\#c}{(\mathcal{P} \mid \text{Ctx}\{\text{forkNonLinear } (\lambda x.t)\}; \mathcal{C}) \rightsquigarrow (\mathcal{P} \mid \text{Ctx}\{\bar{c}\} \mid t[[c]/x]^\circ; \mathcal{C}, \bar{c} = \langle \rangle, c^* = \langle \rangle)} \text{FORKNONLINEAR}$$

where the channel configuration is typed $\Delta'_2 = \mathcal{C}$ and we select a fresh name c for the channel being forked.

- * t_2 is not a value, and has a reduction $t_2 \rightsquigarrow t'_2$ for some t'_2 .
Then we can reduce by REDUCE using the reduction:

$$\frac{t_2 \rightsquigarrow t'_2}{\text{forkNonLinear } t_2 \rightsquigarrow \text{forkNonLinear } t'_2} \rightsquigarrow^{\text{prim}}$$

8. $t_1 = \text{forkReplicate}$, where $A = \square_{0..1}(\mathbf{Chan } P \multimap 1) \multimap \mathbf{N } n \multimap \mathbf{Vec } n(\mathbf{Chan } \overline{P})$. Therefore, we induct on the second argument. There are two cases.

- * t_2 is a value. Then, by the value lemma, $t_2 = [\lambda c.t'_2]$ for some t'_2 and $t_1 t_2$ is $\text{forkReplicate } [\lambda c.t'_2]$ which is also a value.
- * t_2 is not a value, and has a reduction $t_2 \rightsquigarrow t'_2$ for some t'_2 .
Then we can reduce by REDUCE using the reduction:

$$\frac{t_2 \rightsquigarrow t'_2}{\text{forkReplicate } t_2 \rightsquigarrow \text{forkReplicate } t'_2} \rightsquigarrow^{\text{prim}}$$

9. $t_1 = \text{forkReplicate } c$, where $A = \mathbf{N } n \multimap \mathbf{Vec } n(\mathbf{Chan } \overline{P})$. Therefore, we induct on the second argument. There are two cases.
- * t_2 is a value. Then, by the value lemma, $t_2 = \lambda x.t'_2$ for some t'_2 and the runtime typing derivation refines to:

$$\frac{\frac{\frac{x : \mathbf{Chan } P \vdash t'_2 : 1}{\emptyset \vdash \lambda x.t'_2 : \mathbf{Chan } P \multimap 1} \multimap_c}{\emptyset \vdash [\lambda x.t'_2] : \square_{0..n}(\mathbf{Chan } P \multimap 1)} \text{PR} \quad \frac{\text{ReceivePrefix}(P)}{\emptyset \vdash \text{forkReplicate} : \square_{0..n}(\mathbf{Chan } P \multimap 1) \multimap \mathbf{N } n \multimap \mathbf{Vec } n(\square_{0..1}(\mathbf{Chan } \overline{P}))} \text{FORKREP}}{\frac{\emptyset \vdash \text{forkReplicate } [\lambda x.t'_2] : \mathbf{N } n \multimap \mathbf{Vec } n(\square_{0..1}(\mathbf{Chan } \overline{P}))}{\emptyset \vdash \text{forkReplicate } [\lambda x.t'_2] n : \mathbf{Vec } n(\square_{0..1}(\mathbf{Chan } \overline{P}))} \multimap_c} \multimap_c$$

Therefore, by Lemma 7 our overall succedent is of the form $\Delta'_2 \vdash \text{Ctx}\{\text{forkReplicate } [\lambda x.t'_2] n\} : A$.

Then there is a reduction as follows:

$$\frac{\#c_1 \dots \#c_n}{(\mathcal{P} \mid \text{Ctx}\{\text{forkReplicate } [\lambda x.t] n\}; \mathcal{C}) \rightsquigarrow (\mathcal{P} \mid \text{Ctx}\{\langle \overline{c_1} \rangle, \dots, \langle \overline{c_n} \rangle\} \mid \text{Wch}(c_1 \mapsto \lambda x.t) \mid \dots \mid \text{Wch}(c_n \mapsto \lambda x.t); \mathcal{C}, c_1 = \langle \rangle, \dots, c_n = \langle \rangle, \overline{c_1} = \langle \rangle, \dots, \overline{c_n} = \langle \rangle)} \text{FORKREPLICATE}$$

where the channel configuration is typed $\Delta'_2 = \mathcal{C}$ and we select fresh names c and $c_1 \dots c_n$ for the channels being forked.

- * t_2 is not a value, and has a reduction $t_2 \rightsquigarrow t'_2$ for some t'_2 .
Then we can reduce by REDUCE using the reduction:

$$\frac{t_2 \rightsquigarrow t'_2}{\text{forkReplicate } c t_2 \rightsquigarrow \text{forkReplicate } c t'_2} \rightsquigarrow^{\text{prim}}$$

10. $t_1 = \text{forkMulticast}$, where $A = (\mathbf{Chan}(\text{Graded}_n P) \multimap 1) \multimap \mathbf{N} \ n \multimap \mathbf{Vec} \ n(\mathbf{Chan} \bar{P})$. Therefore, we induct on the second argument. There are two cases.

- * t_2 is a value. Then, by the value lemma, $t_2 = \lambda c.t'_2$ for some t'_2 and $t_1 \ t_2$ is $\text{forkMulticast} \ \lambda c.t'_2$ which is also a value.
- * t_2 is not a value, and has a reduction $t_2 \rightsquigarrow t'_2$ for some t'_2 . Then we can reduce by REDUCE using the reduction:

$$\frac{t_2 \rightsquigarrow t'_2}{\text{forkMulticast} \ t_2 \rightsquigarrow \text{forkMulticast} \ t'_2} \rightsquigarrow_{\text{prim}}$$

11. $t_1 = \text{forkMulticast} \ c$, where $A = \mathbf{N} \ n \multimap \mathbf{Vec} \ n(\mathbf{Chan} \bar{P})$. Therefore, we induct on the second argument. There are two cases.

- * t_2 is a value. Then, by the value lemma, $t_2 = \lambda x.t'_2$ for some t'_2 and the runtime typing derivation refines to:

$$\frac{\frac{x : \mathbf{Chan}(\text{Graded}, P) \vdash t'_2 : 1}{\emptyset \vdash \lambda x.t'_2 : \mathbf{Chan}(\text{Graded}, P) \multimap 1} \multimap_i}{\emptyset \vdash \text{forkMulticast}(\lambda x.t'_2) : \mathbf{N} \ n \multimap \mathbf{Vec} \ n(\mathbf{Chan} \bar{P})} \multimap_c \quad \frac{\text{Sends}(P)}{\emptyset \vdash \text{forkMulticast} : (\mathbf{Chan}(\text{Graded}, P) \multimap 1) \multimap \mathbf{N} \ n \multimap \mathbf{Vec} \ n(\mathbf{Chan} \bar{P})} \text{FORKMULTI}}{\emptyset \vdash \text{forkMulticast}(\lambda x.t'_2) \ n : \mathbf{Vec} \ n(\mathbf{Chan} \bar{P})} \multimap_c$$

Therefore, by Lemma 7 our overall succedent is of the form $\Delta'_2 \vdash \text{Ctxt}\{\text{forkMulticast} \ \lambda x.t'_2 \ n\} : A$.

Then there is a reduction as follows:

$$\frac{\#c \quad \#c_1 \dots \#c_n}{(\mathcal{P} \mid \text{Ctxt}\{\text{forkMulticast}(\lambda x.t) \ n\}; C) \rightsquigarrow (\mathcal{P} \mid \text{Ctxt}\{\bar{c}_1, \dots, \bar{c}_n\} \mid t[c/x]^\circ \mid \text{Fwd}(c \mapsto c_1, \dots, c_n); C, c = \langle \rangle, \bar{c} = \langle \rangle, c_1 = \langle \rangle, \dots, c_n = \langle \rangle, \bar{c}_1 = \langle \rangle, \dots, \bar{c}_n = \langle \rangle)} \text{FORKMULTI}$$

where the channel configuration is typed $\Delta'_2 = \mathcal{C}$ and we select fresh names c and $c_1 \dots c_n$ for the channels being forked.

- * t_2 is not a value, and has a reduction $t_2 \rightsquigarrow t'_2$ for some t'_2 . Then we can reduce by REDUCE using the reduction:

$$\frac{t_2 \rightsquigarrow t'_2}{\text{forkMulticast} \ c \ t_2 \rightsquigarrow \text{forkMulticast} \ c \ t'_2} \rightsquigarrow_{\text{prim}}$$

In the second case, by induction on the premise we can reduce by REDUCE using the reduction:

$$\frac{t_1 \rightsquigarrow t'_1}{t_1 \ t_2 \rightsquigarrow t'_1 \ t_2} \rightsquigarrow_{\text{APP}}$$

– (pr)

$$\frac{\Gamma \vdash t : A}{r * \Gamma \vdash [t] : \square_r A} \text{ PR}$$

There are two cases: either $t = v$ for some value v , or $t \rightsquigarrow t'$ for some t' .

In the first case, $[v]$ is a value.

In the second case, we can reduce by REDUCE using the reduction:

$$\frac{t \rightsquigarrow t'}{[t] \rightsquigarrow [t']} \rightsquigarrow_{\square \text{reduce}}$$

– (elim)

$$\frac{\frac{\Gamma_1 \vdash t_1 : \square_r A \quad \Gamma_2, x : [A]_r \vdash t_2 : B}{\Gamma_1 + \Gamma_2 \vdash \mathbf{let} [x] = t_1 \mathbf{in} t_2 : B} \text{ ELIM}}{\Gamma_1 + \Gamma_2 \vdash \mathbf{let} [x] = t_1 \mathbf{in} t_2 : B} \text{ ELIM}$$

By induction on the first premise, there are two possibilities: either $t_1 = v$ for some value v or $t_1 \rightsquigarrow t'_1$ for some t'_1 .

In the first case, by the value lemma $t_1 = [v']$ for some value v' . This refines the typing as follows:

$$\frac{\frac{\emptyset \vdash v' : A}{\emptyset \vdash [v'] : \square_r A} \text{ PR} \quad x : [A]_r \vdash t_2 : B}{\emptyset \vdash \mathbf{let} [x] = [v'] \mathbf{in} t_2 : B} \text{ ELIM}$$

Therefore, we can reduce by REDUCE using the reduction:

$$\overline{\mathbf{let} [x] = [v] \mathbf{in} t} \rightsquigarrow_{\square \beta} \overline{\mathbf{let} [x] = [v] \mathbf{in} t[v/x]}$$

In the second case, by induction on the premise we can reduce by REDUCE using the reduction:

$$\frac{t_1 \rightsquigarrow t'_1}{\mathbf{let} [x] = t_1 \mathbf{in} t_2 \rightsquigarrow \mathbf{let} [x] = t'_1 \mathbf{in} t_2} \rightsquigarrow_{\text{LET}\square}$$

– (der)

$$\frac{\Gamma, x : A \vdash t : B}{\Gamma, x : [A]_1 \vdash t : B} \text{ DER}$$

Goal achieved immediately by induction on the premise.

– (weak)

$$\frac{\Gamma \vdash t : B}{\Gamma, 0 * \Gamma' \vdash t : B} \text{WEAK}$$

Goal achieved immediately by induction on the premise.

– (approx)

$$\frac{\Gamma, x : [A]_r, \Gamma' \vdash t : B}{\Gamma, x : [A]_s, \Gamma' \vdash t : B} \text{APPROX}$$

Goal achieved immediately by induction on the premise.

– (tensor)

$$\frac{\Gamma_1 \vdash t_1 : A \quad \Gamma_2 \vdash t_2 : B}{\Gamma_1 + \Gamma_2 \vdash (t_1, t_2) : A \otimes B} \otimes_i$$

There are three cases: either $t_1 = v_1$ and $t_2 = v_2$ for some pair of values (v_1, v_2) , $t_1 = v$ for some value v but $t_2 \rightsquigarrow t'_2$ for some t'_2 , or $t_1 \rightsquigarrow t'_1$ for some t'_1 .

In the first case, (v_1, v_2) is a value.

In the second case, we can reduce by REDUCE using the reduction:

$$\frac{t_2 \rightsquigarrow t'_2}{(v, t_2) \rightsquigarrow (v, t'_2)} \rightsquigarrow_{\otimes \text{reduceR}}$$

In the third case, we can reduce by REDUCE using the reduction:

$$\frac{t_1 \rightsquigarrow t'_1}{(t_1, t_2) \rightsquigarrow (t'_1, t_2)} \rightsquigarrow_{\otimes \text{reduceL}}$$

– (pairElim)

$$\frac{\Gamma_1 \vdash t_1 : A \otimes A' \quad \Gamma_2, x : A, y : A' \vdash t_2 : B}{\Gamma_1 + \Gamma_2 \vdash \mathbf{let} (x, y) = t_1 \mathbf{in} t_2 : B} \otimes_e$$

By induction on the first premise, there are two possibilities: either $t_1 = v$ for some value v or $t_1 \rightsquigarrow t'_1$ for some t'_1 .

In the first case, by the value lemma $t_1 = (v_1, v_2)$ for some pair of values (v_1, v_2) . This refines the typing as follows:

$$\frac{\frac{\emptyset \vdash v_1 : A \quad \emptyset \vdash v_2 : A'}{\emptyset \vdash (v_1, v_2) : A \otimes A'} \otimes_i \quad x : A, y : A' \vdash t_2 : B}{\emptyset \vdash \mathbf{let} (x, y) = (v_1, v_2) \mathbf{in} t_2 : B} \otimes_e$$

Therefore, we can reduce by REDUCE using the reduction:

$$\overline{\mathbf{let} (x, y) = (v_1, v_2) \mathbf{in} t_3} \rightsquigarrow_{\otimes\beta} t_3[v_1/x][v_2/y]$$

In the second case, by induction on the premise we can reduce by REDUCE using the reduction:

$$\frac{t_1 \rightsquigarrow t'_1}{\mathbf{let} (x, y) = t_1 \mathbf{in} t_2 \rightsquigarrow \mathbf{let} (x, y) = t'_1 \mathbf{in} t_2} \rightsquigarrow_{\text{LET}\otimes}$$

– (unit)

$$\overline{\emptyset \vdash \mathbf{unit} : 1} \text{UNIT}_i$$

A value.

– (unitElim)

$$\frac{\Gamma_1 \vdash t_1 : 1 \quad \Gamma_2 \vdash t_2 : B}{\Gamma_1 + \Gamma_2 \vdash \mathbf{let} \mathbf{unit} = t_1 \mathbf{in} t_2 : B} \text{UNIT}_e$$

By induction on the first premise, there are two possibilities: either $t_1 = v$ for some value v or $t_1 \rightsquigarrow t'_1$ for some t'_1 .

In the first case, by the value lemma $t_1 = \mathbf{unit}$ and so we can reduce by REDUCE using the reduction:

$$\overline{\mathbf{let} \mathbf{unit} = \mathbf{unit} \mathbf{in} t} \rightsquigarrow_{\beta\text{unit}} t$$

In the second case, by induction on the premise we can reduce by REDUCE using the reduction:

$$\frac{t_1 \rightsquigarrow t'_1}{\mathbf{let} \mathbf{unit} = t_1 \mathbf{in} t_2 \rightsquigarrow \mathbf{let} \mathbf{unit} = t'_1 \mathbf{in} t_2} \rightsquigarrow_{\text{LETunit}}$$

□

A process configuration \mathcal{P} contains all values, defined by the predicate $\text{values}(\mathcal{P})$, if every process is either a value v or is a forwarder or watcher process.

Theorem 2 (Progress global config).

$$\begin{aligned} & \Delta \vdash \mathcal{P} : A \wedge \Delta \vdash \mathcal{C} \wedge \text{buffered}(\mathcal{C}, \Delta) \\ \implies & \text{values}(\mathcal{P}) \vee \exists \mathcal{C}', \mathcal{P}'. (\mathcal{P}; \mathcal{C}) \rightsquigarrow (\mathcal{P}'; \mathcal{C}') \end{aligned}$$

Proof. By induction on the two judgments forming the assumptions.

First, we case split on the typing of process configurations: $\Delta \vdash \mathcal{P} : A$.

- (empty)

$$\frac{}{\emptyset \vdash \cdot : \perp} \text{EMPTY}$$

This case is trivial, as $\text{values}(\mathcal{P})$ holds by default for an empty process configuration.

- (cons)

$$\frac{\begin{array}{c} \Gamma_1, \Delta_1 \vdash \mathcal{P} : A \\ \Gamma_2, \Delta_2 \vdash t : B \end{array}}{\Gamma_1 + \Gamma_2, \Delta_1 + \Delta_2 \vdash \mathcal{P} \mid t^{\mathcal{R}} : A} \text{CONSR}$$

Here, we can decompose \mathcal{P} into $\mathcal{P}' \mid t$. First, focus on the particular term t and apply Theorem 1.

Then, either t is a value or there exists some possible reduction $(\mathcal{P} \mid t; \mathcal{C}) \rightsquigarrow (\mathcal{P}' \mid t'; \mathcal{C}')$.

If we can make such a reduction, then the theorem is satisfied. If not, then apply induction to the smaller process configuration \mathcal{P}' . We find that either $\text{values}(\mathcal{P}')$ holds or there exists a reduction $(\mathcal{P}'; \mathcal{C}) \rightsquigarrow (\mathcal{P}''; \mathcal{C}')$.

If $\text{values}(\mathcal{P}')$, then since t is also a value we have $\text{values}(\mathcal{P})$ and we are done.

Otherwise, we apply the following rule to find a reduction satisfying the theorem:

$$\frac{(\mathcal{P}; \mathcal{C}) \rightsquigarrow (\mathcal{P}'; \mathcal{C}')}{(\mathcal{P} \mid t; \mathcal{C}) \rightsquigarrow (\mathcal{P}' \mid t; \mathcal{C}')} \text{PAR}$$

- (watch)

$$\frac{\Gamma_1, \Delta_1 \vdash \mathcal{P} : A \quad \Gamma_2, x : \mathbf{Chan}(\mathbf{Recv} A' P), \Delta_2 \vdash t : 1}{\Gamma_1 + (0..1 * \Gamma_2), \Delta_1 + \Delta_2, c : \mathbf{Chan}(\mathbf{Recv} A' P) \vdash \mathcal{P} \mid \mathbf{Wch}(c \mapsto \lambda x.t) : A} \text{WATCH}$$

By induction on the first assumption, we have either that $\mathbf{values}(\mathcal{P})$ or that there exists some reduction $(\mathcal{P}; \mathcal{C}) \rightsquigarrow (\mathcal{P}'; \mathcal{C}')$.

If there is such a reduction, we can make a global reduction as follows:

$$\frac{(\mathcal{P}; \mathcal{C}) \rightsquigarrow (\mathcal{P}'; \mathcal{C}')}{(\mathcal{P} \mid \mathbf{Wch}(c \mapsto \lambda x.t); \mathcal{C}) \rightsquigarrow (\mathcal{P}' \mid \mathbf{Wch}(c \mapsto \lambda x.t); \mathcal{C}')} \text{PAR}$$

Otherwise, we must reduce on the watcher process itself. Note that due to the session typing of the channel in the runtime context, and the condition requiring that \mathcal{C} is buffered, the channel configuration must be typed as follows:

$$\frac{\Delta_1 + \Delta_2 \vdash \mathcal{C} \quad \Delta_1 + \Delta_2 \vdash v : A'}{\Delta_1 + \Delta_2, c : (\mathbf{Recv} A' P) \vdash \mathcal{C}, c = v : \vec{v}} \text{RECV} \quad (\text{C.2})$$

Hence, we can construct the following global reduction:

$$\frac{}{(\mathcal{P} \mid \mathbf{Wch}(c \mapsto \lambda x.t); \mathcal{C}, c = v : \vec{v}) \rightsquigarrow (\mathcal{P} \mid t[c/x]^\circ; \mathcal{C}, c = v : \vec{v})} \text{WATCHER}$$

- (forward)

$$\frac{\mathbf{Sends}(P) \quad \Gamma, \Delta \vdash \mathcal{P} : A}{\Gamma, \Delta, c : \mathbf{Chan}(\mathbf{Graded}_n \overline{P}), c_i : \mathbf{Chan} P \vdash \mathcal{P} \mid \mathbf{Fwd}(c \mapsto c_1, \dots, c_n) : A} \text{FWD}$$

By induction on the first assumption, we have either that $\mathbf{values}(\mathcal{P})$ or that there exists some reduction $(\mathcal{P}; \mathcal{C}) \rightsquigarrow (\mathcal{P}'; \mathcal{C}')$.

If there is such a reduction, we can make a global reduction as follows:

$$\frac{(\mathcal{P}; \mathcal{C}) \rightsquigarrow (\mathcal{P}'; \mathcal{C}')}{(\mathcal{P} \mid \text{Fwd}(c \mapsto c_1, \dots, c_n); \mathcal{C}) \rightsquigarrow (\mathcal{P}' \mid \text{Fwd}(c \mapsto c_1, \dots, c_n); \mathcal{C}')} \text{PAR}$$

Otherwise, we must reduce on the forwarder process itself. Note that due to the session typing of the channel in the runtime context, and the condition requiring that \mathcal{C} is buffered, then the channel configuration must be typed as follows:

$$\frac{\frac{\emptyset \vdash v : A'}{\emptyset \vdash [v] : \square_n A'} \text{PR} \quad \frac{\Delta, c_1 : \text{Chan } P, \dots, c_n : \text{Chan } P, c : \text{Chan } \bar{P} \vdash \mathcal{C}, \bar{c}_1 = \bar{v}_1 : v, \dots, \bar{c}_n = \bar{v}_n : v, c = \bar{v}}{\Delta, c_1 : \text{Chan}(\text{Send } A P), \dots, c_n : \text{Chan}(\text{Send } A P), c : \text{Chan } \bar{P} \vdash \mathcal{C}, \bar{c}_1 = \bar{v}_1, \dots, \bar{c}_n = \bar{v}_n, c = \bar{v}} \text{SEND}^n}{\Delta, c_1 : \text{Chan}(\text{Send } A P), \dots, c_n : \text{Chan}(\text{Send } A P), c : \text{Chan}(\text{Recv}(\square_n A') \bar{P}) \vdash \mathcal{C}, \bar{c}_1 = \bar{v}_1, \dots, \bar{c}_n = \bar{v}_n, c = [v] : \bar{v}} \text{RECV} \quad (\text{C.3})$$

Hence, we can construct the following global reduction:

$$\frac{}{(\mathcal{P} \mid \text{Fwd}(c \mapsto c_1, \dots, c_n); \mathcal{C}, c = [v] : \bar{v}, \bar{c}_1 = \bar{v}_1, \dots, \bar{c}_n = \bar{v}_n) \rightsquigarrow (\mathcal{P} \mid \text{Fwd}(c \mapsto c_1, \dots, c_n); \mathcal{C}, c = \bar{v}, \bar{c}_1 = \bar{v}_1 : v, \dots, \bar{c}_n = \bar{v}_n : v)} \text{FORWARDER}$$

□

Appendix D. Preservation proof

Theorem 3 (Preservation).

$$\begin{aligned} & \Gamma, \Delta \vdash \mathcal{P} : A \wedge \Delta \vdash \mathcal{C} \wedge (\mathcal{P}; \mathcal{C}) \rightsquigarrow (\mathcal{P}'; \mathcal{C}') \\ \implies & \exists \Delta'. \Gamma, \Delta' \vdash \mathcal{P}' : A \wedge \Delta' \vdash \mathcal{C}' \end{aligned}$$

Proof. By induction on the three judgments forming the assumptions here.

We first case split on the typing of process configurations: $\Gamma, \Delta \vdash \mathcal{P} : A$.

- (empty)

$$\frac{}{\emptyset \vdash \cdot : \perp} \text{EMPTY}$$

There is no reduction on the empty process configuration so this case is trivially true since the antecedent is false.

- (cons) We consider an abstraction over the two possibilities for extending a process configuration with a term which is either main or subordinate process:

$$\frac{\Gamma_1, \Delta_1 \vdash t : B \quad \Gamma_2, \Delta_2 \vdash \mathcal{P} : A}{\Gamma_1 + \Gamma_2, \Delta_1 + \Delta_2 \vdash \mathcal{P} \mid t^{\mathcal{R}} : A} \text{CONS}\mathcal{R}$$

where $\mathcal{R} = \bullet$ in which case $B = A$, otherwise $B = 1$.

We then case on the reduction of process configurations consider the role \mathcal{R} where pertinent.

- (send)

$$\frac{}{(\mathcal{P} \mid \text{Ctxt}\{\text{send } c v\}; \mathcal{C}, \bar{c} = \vec{v}) \rightsquigarrow (\mathcal{P} \mid \text{Ctxt}\{c\}; \mathcal{C}, \bar{c} = \vec{v} : v)} \text{SEND}$$

with process configuration typing:

$$\frac{\Gamma_1, \Delta_1, c :^r \mathbf{Chan}(\mathbf{Send } A P) \vdash \text{Ctxt}\{\text{send } c v\} : B \quad \Gamma_2, \Delta_2 \vdash \mathcal{P} : A}{\Gamma_1 + \Gamma_2, \Delta_1 + \Delta_2, c :^r \mathbf{Chan}(\mathbf{Send } A P) \vdash \mathcal{P} \mid \text{Ctxt}\{\text{send } c v\}^{\mathcal{R}} : A} \text{CONS}\mathcal{R} \quad (\text{D.1})$$

and channel configuration typing:

$$\frac{\Delta_3 \vdash v : A \quad \Delta_3, \Delta_4, c :^r \mathbf{Chan } P \vdash \mathcal{C}, \bar{c} = \vec{v} : v}{\Delta_3, \Delta_4, c :^r \mathbf{Chan}(\mathbf{Send } A P) \vdash \mathcal{C}, \bar{c} = \vec{v}} \text{SEND} \quad (\text{D.2})$$

where $(\Delta_3, \Delta_4) = \Delta_1 + \Delta_2$.

We can then construct the consequent channel configuration typing given by the premise of (D.2):

$$\Delta_1 + \Delta_2, c :^r \mathbf{Chan } P \vdash \mathcal{C}, \bar{c} = v : \vec{v}$$

From typing $\Gamma_1, \Delta_1 \vdash \text{Ctxt}\{\text{send } c v\} : B$, and by lemma (6) we have two cases:

1. Typing $\Gamma'_1, \Delta'_1 \vdash \text{send } c v : A''$ and $\Gamma'_2, y : A'', \Delta'_2 \vdash \text{Ctxt}\{y\} : B$ with $\Gamma_1 = \Gamma'_1 + \Gamma'_2$ and $\Delta_1 = \Delta'_1 + \Delta'_2$, which by inversion must therefore be of the form:

$$\frac{\dots \quad \frac{}{c : \mathbf{Chan}(\mathbf{Send } A' P) \vdash c : \mathbf{Chan}(\mathbf{Send } A' P)} \text{VAR}}{\emptyset, c : \mathbf{Chan}(\mathbf{Send } A' P) \vdash \text{send } c : A' \multimap \mathbf{Chan } P} \multimap_e \quad \emptyset \vdash v : A'}{\emptyset, c : \mathbf{Chan}(\mathbf{Send } A' P) \vdash \text{send } c v : \mathbf{Chan } P} \multimap_e$$

i.e., where $A'' = \mathbf{Chan} P$ and $\Gamma'_1 = \cdot$ and $\Delta'_1 = c : \mathbf{Chan} (\mathbf{Send} A' P)$.
Therefore, for the goal of this theorem we construct the typing:

$$\overline{\emptyset, c : \mathbf{Chan} P \vdash c : \mathbf{Chan} P}^{\text{VAR}}$$

By lemma 7 then $\Gamma'_2, \Delta'_2, c : \mathbf{Chan} P \vdash \mathit{Ctxt}\{c\} : B$ which is equal to $\Gamma_1, \Delta'_2, c : \mathbf{Chan} P \vdash \mathit{Ctxt}\{c\} : B$.

We then build the process configuration type:

$$\frac{\frac{\Gamma'_2, \Delta'_2, c : \mathbf{Chan} P \vdash \mathit{Ctxt}\{c\} : B \quad \Gamma_2, \Delta_2 \vdash \mathcal{P} : A}{\Gamma'_2 + \Gamma_2, \Delta'_2 + \Delta_2, c : \mathbf{Chan} P \vdash \mathcal{P} \mid \mathit{Ctxt}\{c\}^{\mathcal{R}} : A} \text{CONSR}}{\Gamma_1 + \Gamma_2, \Delta_1 + \Delta_2, c : \mathbf{Chan} P \vdash \mathcal{P} \mid \mathit{Ctxt}\{c\}^{\mathcal{R}} : A} \equiv \quad (\text{D.3})$$

where the equality step follows by $\Delta_1 = \Delta'_1 + \Delta'_2 = \Delta'_2$ and $\Gamma_1 = \Gamma'_1 + \Gamma'_2 = \Gamma'_2$.

2. Typing $[\Gamma'_1, \Delta'_1] \vdash \mathit{send} c v : A''$ and $\Gamma'_2, y : [A'']_r, \Delta'_2 \vdash \mathit{Ctxt}\{y\} : B$ with $\Gamma_1 = r * \Gamma'_1 + \Gamma'_2$ and $\Delta_1 = \Delta'_1 + \Delta'_2$, which by inversion must therefore be of the form:

$$\frac{\frac{\overline{c : * \mathbf{Chan} (\mathbf{Send} A' P) \vdash c : \mathbf{Chan} (\mathbf{Send} A' P)}}{c : * \mathbf{Chan} (\mathbf{Send} A' P) \vdash \mathit{send} c : A' \multimap \mathbf{Chan} P} \multimap_e \quad \emptyset \vdash v : A'}{c : * \mathbf{Chan} (\mathbf{Send} A' P) \vdash \mathit{send} c v : \mathbf{Chan} P} \multimap_e$$

i.e., where $A'' = \mathbf{Chan} P$ and $\Gamma'_1 = \cdot$ and $\Delta'_1 = c : * \mathbf{Chan} (\mathbf{Send} A' P)$.
Therefore, for the goal of this theorem we construct the typing:

$$\overline{c : * \mathbf{Chan} P \vdash c : \mathbf{Chan} P}^{\text{VARCHANS}}$$

By lemma 8 then $\Gamma'_2, \Delta'_2, c : * \mathbf{Chan} P \vdash \mathit{Ctxt}\{c\} : B$ which is equal to $\Gamma_1, \Delta'_2, c : * \mathbf{Chan} P \vdash \mathit{Ctxt}\{c\} : B$ (since $\Gamma_1 = r * \Gamma'_1 + \Gamma'_2 = r * \emptyset + \Gamma'_2 = \Gamma'_2$).

We then build the process configuration type:

$$\frac{\frac{\Gamma'_2, \Delta'_2, c : * \mathbf{Chan} P \vdash \mathit{Ctxt}\{c\} : B \quad \Gamma_2, \Delta_2 \vdash \mathcal{P} : A}{\Gamma'_2 + \Gamma_2, \Delta'_2 + \Delta_2, c : * \mathbf{Chan} P \vdash \mathcal{P} \mid \mathit{Ctxt}\{c\}^{\mathcal{R}} : A} \text{CONSR}}{\Gamma_1 + \Gamma_2, \Delta_1 + \Delta_2, c : * \mathbf{Chan} P \vdash \mathcal{P} \mid \mathit{Ctxt}\{c\}^{\mathcal{R}} : A} \equiv \quad (\text{D.4})$$

where the equality step follows by $\Delta_1 = \Delta'_1 + \Delta'_2 = \Delta'_2$ and $\Gamma_1 = r * \Gamma'_1 + \Gamma'_2 = \Gamma'_2$.

– (recv)

$$\frac{}{(\mathcal{P} \mid \text{Ctx}\{\text{receive } c\}; \mathcal{C}, c = v : \vec{v}) \rightsquigarrow (\mathcal{P} \mid \text{Ctx}\{(v, c)\}; \mathcal{C}, c = \vec{v})} \text{RECV}$$

with process configuration typing:

$$\frac{\Gamma_1, \Delta_1, c :^r \mathbf{Chan}(\mathbf{Recv } A P) \vdash \text{Ctx}\{\text{receive } c\} : B \quad \Gamma_2, \Delta_2 \vdash \mathcal{P} : A}{\Gamma_1 + \Gamma_2, \Delta_1 + \Delta_2, c :^r \mathbf{Chan}(\mathbf{Recv } A P) \vdash \mathcal{P} \mid \text{Ctx}\{\text{receive } c\}^{\mathcal{R}} : A} \text{CONSR} \quad (\text{D.5})$$

with channel configuration typing:

$$\frac{\Delta_3 \vdash v : A \quad \Delta_3, \Delta_4, c :^r \mathbf{Chan } P \vdash \mathcal{C}, c = \vec{v}}{\Delta_3, \Delta_4, c :^r \mathbf{Chan}(\mathbf{Recv } A P) \vdash \mathcal{C}, c = v : \vec{v}} \text{RECV} \quad (\text{D.6})$$

where $(\Delta_3, \Delta_4) = \Delta_1 + \Delta_2$. (note that **EMPTY** cannot be used in the root of this derivation as the reduction constraints the channel to start non-empty). We then construct the consequent channel configuration typing given by the premise of (D.6):

$$\Delta_1 + \Delta_2, c :^r \mathbf{Chan } P \vdash \mathcal{C}, c = \vec{v}$$

From typing $\Gamma_1, \Delta_1 \vdash \text{Ctx}\{\text{receive } c\} : B$, and by lemma 6 then there are two possible cases:

1. Typing $\Gamma'_1, \Delta'_1 \vdash \text{receive } c : A''$ and $\Gamma'_2, y : A'', \Delta'_2 \vdash \text{Ctx}\{y\} : B$ with $\Gamma_1 = \Gamma'_1 + \Gamma'_2$ and $\Delta_1 = \Delta'_1 + \Delta'_2$, which by inversion must therefore be of the form:

$$\frac{\frac{}{\emptyset, c : \mathbf{Chan}(\mathbf{Recv } A P) \vdash c : \mathbf{Chan}(\mathbf{Recv } A P)} \text{VAR}}{\emptyset, c : \mathbf{Chan}(\mathbf{Recv } A P) \vdash \text{receive } c : (A \otimes \mathbf{Chan } P)} \text{---}_e$$

where $A'' = (A \otimes \mathbf{Chan } P)$ and $\Gamma'_1 = \cdot$ and $\Delta'_1 = c : \mathbf{Chan}(\mathbf{Recv } A P)$. Therefore, for the goal of this theorem we construct the typing:

$$\frac{\frac{}{c : \mathbf{Chan } P \vdash c : \mathbf{Chan } P} \text{VAR} \quad \emptyset \vdash v : A' \text{ (by (D.6))}}{c : \mathbf{Chan } P \vdash (v, c) : (A \otimes \mathbf{Chan } P)} \otimes_i$$

By Lemma 7 then $\Gamma'_2, \Delta'_2, c : \mathbf{Chan} P \vdash \mathit{Ctxt}\{(v, c)\} : B$ which is equal to $\Gamma_1, \Delta'_2, c : \mathbf{Chan} P \vdash \mathit{Ctxt}\{(v, c)\} : B$

We then construct the preserved process configuration typing:

$$\frac{\frac{\Gamma'_2, \Delta'_2, c : \mathbf{Chan} P \vdash \mathit{Ctxt}\{(v, c)\} : B \quad \Gamma_2, \Delta_2 \vdash \mathcal{P} : A}{\Gamma'_2 + \Gamma_2, \Delta'_2 + \Delta_2, c : \mathbf{Chan} P \vdash \mathcal{P} \mid \mathit{Ctxt}\{(v, c)\}^{\mathcal{R}} : A} \text{CONSR}}{\Gamma_1 + \Gamma_2, \Delta_1 + \Delta_2, c : \mathbf{Chan} P \vdash \mathcal{P} \mid \mathit{Ctxt}\{(v, c)\}^{\mathcal{R}} : A} \equiv \quad (\text{D.7})$$

where the equality step follows by $\Delta_1 = \Delta'_1 + \Delta'_2 = \Delta'_2$ and $\Gamma_1 = \Gamma'_1 + \Gamma'_2 = \Gamma'_2$.

2. Typing $[\Gamma'_1, \Delta'_1] \vdash \text{receive } c : A''$ and $\Gamma'_2, y : [A'']_r, \Delta'_2 \vdash \mathit{Ctxt}\{y\} : B$ with $\Gamma_1 = r * \Gamma'_1 + \Gamma'_2$ and $\Delta_1 = \Delta'_1 + \Delta'_2$, which by inversion must therefore be of the form:

$$\frac{\overline{c : * \mathbf{Chan} (\mathbf{Recv} A P) \vdash c : \mathbf{Chan} (\mathbf{Recv} A P)}}{\text{VARCHAN S}} \quad \frac{\overline{c : * \mathbf{Chan} (\mathbf{Recv} A P) \vdash \text{receive } c : (A \otimes \mathbf{Chan} P)}}{-\circ_e}$$

where $A'' = (A \otimes \mathbf{Chan} P)$ and $\Gamma'_1 = \cdot$ and $\Delta'_1 = c : * \mathbf{Chan} (\mathbf{Recv} A P)$. Therefore, for the goal of this theorem we construct the typing:

$$\frac{\overline{c : * \mathbf{Chan} P \vdash c : \mathbf{Chan} P} \quad \text{VARCHAN S} \quad \emptyset \vdash v : A' \text{ (by (D.6))}}{c : * \mathbf{Chan} P \vdash (v, c) : (A \otimes \mathbf{Chan} P)} \otimes_i$$

By Lemma 8 then $\Gamma'_2, \Delta'_2, c : * \mathbf{Chan} P \vdash \mathit{Ctxt}\{(v, c)\} : B$ which is equal to $\Gamma_1, \Delta'_2, c : * \mathbf{Chan} P \vdash \mathit{Ctxt}\{(v, c)\} : B$ (since $\Gamma_1 = r * \Gamma'_1 + \Gamma'_2 = r * \emptyset + \Gamma'_2 = \Gamma'_2$).

We then construct the preserved process configuration typing:

$$\frac{\frac{\Gamma'_2, \Delta'_2, c : * \mathbf{Chan} P \vdash \mathit{Ctxt}\{(v, c)\} : B \quad \Gamma_2, \Delta_2 \vdash \mathcal{P} : A}{\Gamma'_2 + \Gamma_2, \Delta'_2 + \Delta_2, c : * \mathbf{Chan} P \vdash \mathcal{P} \mid \mathit{Ctxt}\{(v, c)\}^{\mathcal{R}} : A} \text{CONSR}}{\Gamma_1 + \Gamma_2, \Delta_1 + \Delta_2, c : * \mathbf{Chan} P \vdash \mathcal{P} \mid \mathit{Ctxt}\{(v, c)\}^{\mathcal{R}} : A} \equiv \quad (\text{D.8})$$

where the equality step follows by $\Delta_1 = \Delta'_1 + \Delta'_2 = \Delta'_2$ and $\Gamma_1 = r * \Gamma'_1 + \Gamma'_2 = \Gamma'_2$.

– (close)

$$\frac{}{(\mathcal{P} \mid \text{Ctxt}\{\text{close } c\}; \mathcal{C}, c^1 = \langle \rangle) \rightsquigarrow (\mathcal{P} \mid \text{Ctxt}\{\mathbf{unit}\}; \mathcal{C})} \text{CLOSE}$$

with process configuration typing:

$$\frac{\Gamma_1, \Delta_1, c : \mathbf{Chan\ End} \vdash \text{Ctxt}\{\text{close } c\} : B \quad \Gamma_2, \Delta_2 \vdash \mathcal{P} : A}{\Gamma_1 + \Gamma_2, \Delta_1 + \Delta_2, c : \mathbf{Chan\ End} \vdash \mathcal{P} \mid \text{Ctxt}\{\text{close } c\}^{\mathcal{R}} : A} \text{CONSR} \quad (\text{D.9})$$

(i.e., $c \notin \text{dom}(\Delta_2)$) with only one potential channel configuration typing:

$$\frac{\Delta_1 + \Delta_2 \vdash \mathcal{C}}{\Delta_1 + \Delta_2, c : \mathbf{Chan\ End} \vdash \mathcal{C}, c^1 = \langle \rangle} \text{EMPTY} \quad (\text{D.10})$$

The goal channel configuration $\Delta_1 + \Delta_2 \vdash \mathcal{C}$ is then given by the premise of (D.10).

From typing $\Gamma_1, \Delta_1 \vdash \text{Ctxt}\{\text{close } c\} : B$, and by lemma 6, and since we know here that c is a linear channel, then the typing factorises to the typing $\Gamma'_1, \Delta'_1 \vdash \text{close } c : A''$ and $\Gamma'_2, y : A'', \Delta'_2 \vdash \text{Ctxt}\{y\} : B$ with $\Gamma_1 = \Gamma'_1 + \Gamma'_2$ and $\Delta_1 = \Delta'_1 + \Delta'_2$, which by inversion must therefore be of the form:

$$\frac{c : \mathbf{Chan\ End} \vdash c : \mathbf{Chan\ End}}{c : \mathbf{Chan\ End} \vdash \text{close } c : 1} \text{--}\circ_e$$

where $A'' = 1$ and $\Gamma'_1 = \cdot$ and $\Delta'_1 = c : \mathbf{Chan\ End}$.

Therefore, for the goal of this theorem we construct the typing:

$$\frac{}{\emptyset \vdash \mathbf{unit} : 1} \text{UNIT}_i$$

By Lemma 7 then $\Gamma'_2, \Delta'_2 \vdash \text{Ctxt}\{\mathbf{unit}\} : B$.

We then construct the preserved process configuration typing:

$$\frac{\frac{\Gamma'_2, \Delta'_2 \vdash \text{Ctxt}\{\mathbf{unit}\} : B \quad \Gamma_2, \Delta_2 \vdash \mathcal{P} : A}{\Gamma'_2 + \Gamma_2, \Delta'_2 + \Delta_2 \vdash \mathcal{P} \mid \text{Ctxt}\{\mathbf{unit}\}^{\mathcal{R}} : A} \text{CONSR}}{\Gamma_1 + \Gamma_2, \Delta_1 + \Delta_2 \vdash \mathcal{P} \mid \text{Ctxt}\{\mathbf{unit}\}^{\mathcal{R}} : A} \equiv \quad (\text{D.11})$$

where the equality step follows by $\Delta_1 = \Delta'_1 + \Delta'_2 = \Delta'_2$ and $\Gamma_1 = \Gamma'_1 + \Gamma'_2 = \Gamma'_2$.

– (sharedClose)

$$\overline{(\mathcal{P} \mid \text{Ctxt}\{\text{close } c\}; \mathcal{C}, c^* = \vec{v}) \rightsquigarrow (\mathcal{P} \mid \text{Ctxt}\{\mathbf{unit}\}; \mathcal{C}, c^* = \vec{v})} \text{ SHARED_CLOSE}$$

with process configuration typing:

$$\frac{\Gamma_1, \Delta_1, c :^* \mathbf{Chan\ End} \vdash \text{Ctxt}\{\text{close } c\} : B \quad \Gamma_2, \Delta_2 \vdash \mathcal{P} : A}{\Gamma_1 + \Gamma_2, \Delta_1 + \Delta_2, c :^* \mathbf{Chan\ End} \vdash \mathcal{P} \mid \text{Ctxt}\{\text{close } c\}^{\mathcal{R}} : A} \text{ CONS}\mathcal{R} \quad (\text{D.12})$$

and channel configuration typing $\Delta_1 + \Delta_2, c :^* \mathbf{Chan\ End} \vdash \mathcal{C}, c^* = \vec{v}$.

Since the reduction semantics preserves the channel configuration, this channel configuration typing then also provides the goal configuration typing as-is. From typing $\Gamma_1, \Delta_1 \vdash \text{Ctxt}\{\text{close } c\} : B$, and by lemma 6, and since we know here that c is a non-linear channel, then the typing factorises to the typing $[\Gamma'_1, \Delta'_1] \vdash \text{close } c : A''$ and $\Gamma'_2, y : [A'']_r, \Delta'_2 \vdash \text{Ctxt}\{y\} : B$ with $\Gamma_1 = r * \Gamma'_1 + \Gamma'_2$ and $\Delta_1 = \Delta'_1 + \Delta'_2$, which by inversion must therefore be of the form:

$$\frac{c :^* \mathbf{Chan\ End} \vdash c : \mathbf{Chan\ End}}{c :^* \mathbf{Chan\ End} \vdash \text{close } c : 1} \text{ }^{-\circ_e}$$

where $A'' = 1$ and $\Gamma'_1 = \cdot$ and $\Delta'_1 = c :^* \mathbf{Chan\ End}$.

Therefore, for the goal of this theorem we construct the typing:

$$\frac{\overline{\emptyset \vdash \mathbf{unit} : 1} \text{ UNIT}_i}{c :^* \mathbf{Chan\ End} \vdash \mathbf{unit} : 1} \text{ WEAKCHANS}$$

By Lemma 8 then $\Gamma'_2, \Delta'_2, c :^* \mathbf{Chan\ End} \vdash \text{Ctxt}\{\mathbf{unit}\} : B$.

We then construct the preserved process configuration typing:

$$\frac{\frac{\Gamma'_2, \Delta'_2, c :^* \mathbf{Chan\ End} \vdash \text{Ctxt}\{\mathbf{unit}\} : B \quad \Gamma_2, \Delta_2 \vdash \mathcal{P} : A}{\Gamma'_2 + \Gamma_2, \Delta'_2 + \Delta_2, c :^* \mathbf{Chan\ End} \vdash \mathcal{P} \mid \text{Ctxt}\{\mathbf{unit}\}^{\mathcal{R}} : A} \text{ CONS}\mathcal{R}}{\Gamma_1 + \Gamma_2, \Delta_1 + \Delta_2, c :^* \mathbf{Chan\ End} \vdash \mathcal{P} \mid \text{Ctxt}\{\mathbf{unit}\}^{\mathcal{R}} : A} \equiv \quad (\text{D.13})$$

where the equality step follows by $\Delta_1 = \Delta'_1 + \Delta'_2 = \Delta'_2$ and $\Gamma_1 = r * \Gamma'_1 + \Gamma'_2 = \Gamma'_2$.

– (fork)

$$\frac{\#c}{(\mathcal{P} \mid \text{Ctx}\{\text{forkLinear}(\lambda x.t)\}; \mathcal{C}) \rightsquigarrow (\mathcal{P} \mid t[c/x]^\circ \mid \text{Ctx}\{\bar{c}\}; \mathcal{C}, \bar{c} = \langle \rangle, c = \langle \rangle)} \text{FORKLINEAR}$$

with process configuration typing:

$$\frac{\Gamma_1, \Delta_1 \vdash \text{Ctx}\{\text{forkLinear}(\lambda x.t)\} : B \quad \Gamma_2, \Delta_2 \vdash \mathcal{P} : A}{\Gamma_1 + \Gamma_2, \Delta_1 + \Delta_2 \vdash \mathcal{P} \mid \text{Ctx}\{\text{forkLinear}(\lambda x.t)\}^{\mathcal{R}} : A} \text{CONSR} \quad (\text{D.14})$$

with channel configuration typing:

$$\Delta_1 + \Delta_2 \vdash \mathcal{C} \quad (\text{D.15})$$

From typing $\Gamma_1, \Delta_1 \vdash \text{Ctx}\{\text{forkLinear}(\lambda x.t)\} : B$, and by lemma 6 there are two possibilities:

1. (second disjunct) (with $\text{Ctx}\{\text{forkLinear}(\lambda x.t)\}$ factorising into a graded binding) Gives $\neg\text{resourceAllocator}(\text{forkLinear}(\lambda x.t))$ which is a contradiction, thus we can conclude trivial by ex falso.
2. (first disjunct) Typing $\Gamma'_1, \Delta'_1 \vdash \text{forkLinear}(\lambda x.t) : A$, which by inversion must therefore be of the form:

$$\frac{\frac{\Gamma'_1, x : \mathbf{Chan} P, \Delta'_1 \vdash t : 1}{\Gamma'_1, \Delta'_1 \vdash (\lambda x.t) : \mathbf{Chan} P \multimap 1} \multimap_i}{\Gamma'_1, \Delta'_1 \vdash \text{forkLinear}(\lambda x.t) : \mathbf{Chan}(\bar{P})} \multimap_e \quad (\text{D.16})$$

where $A = \mathbf{Chan}(\bar{P})$ and $\Gamma''_1, x : \mathbf{Chan}(\bar{P}), \Delta''_1 \vdash \text{Ctx}\{x\} : B$ where $\Gamma_1 = \Gamma'_1 + \Gamma''_1$ and $\Delta_1 = \Delta'_1 + \Delta''_1$.

Therefore, for the goal of this theorem we construct the typings:

$$\overline{\emptyset, \bar{c} : \mathbf{Chan}(\bar{P}) \vdash \bar{c} : \mathbf{Chan}(\bar{P})} \text{VAR}$$

and

$$\overline{\emptyset, c : \mathbf{Chan} P \vdash c : \mathbf{Chan} P} \text{VAR} \quad (\text{D.17})$$

By Lemma 7 then $\Gamma''_1, \bar{c} : \mathbf{Chan}(\bar{P}) \vdash \text{Ctx}\{\bar{c}\} : B$.

We construct the goal channel configuration typing as:

$$\frac{\Delta_1 \vdash \mathcal{C}}{\Delta_1 + \Delta_2, c : \mathbf{Chan} P \vdash \mathcal{C}, c = \langle \rangle} \text{EMPTY} \quad \frac{}{\Delta_1 + \Delta_2, c : \mathbf{Chan} P, \bar{c} : \mathbf{Chan}(\bar{P}) \vdash \mathcal{C}, c = \langle \rangle, \bar{c} = \langle \rangle} \text{EMPTY}$$

with process configuration typing then given by:

$$\frac{\frac{\Gamma_1'', \Delta_1'', \bar{c} : \mathbf{Chan}(\bar{P}) \vdash \text{Ctx}\{\bar{c}\} : B \quad \Gamma_2, \Delta_2 \vdash \mathcal{P} : A}{\Gamma_1'' + \Gamma_2, \Delta_1'' + \Delta_2, \bar{c} : \mathbf{Chan}(\bar{P}) \vdash \mathcal{P} \mid \text{Ctx}\{\bar{c}\}^{\mathcal{R}} : A} \text{CONSR} \quad \frac{}{\Gamma_1', \Delta_1', c : \mathbf{Chan} P \vdash t[c/x] : 1} \text{SUBST} \quad \text{(D.16) (D.17)}}{\frac{\Gamma_1'' + \Gamma_2 + \Gamma_1', \Delta_1'' + \Delta_2 + \Delta_1', c : \mathbf{Chan} P, \bar{c} : \mathbf{Chan}(\bar{P}) \vdash \mathcal{P} \mid t[c/x]^\circ \mid \text{Ctx}\{\bar{c}\}^{\mathcal{R}} : A}{\Gamma_1 + \Gamma_2, \Delta_1 + \Delta_2, c : \mathbf{Chan} P, \bar{c} : \mathbf{Chan}(\bar{P}) \vdash \mathcal{P} \mid t[c/x]^\circ \mid \text{Ctx}\{\bar{c}\}^{\mathcal{R}} : A} \text{SUB}} \equiv$$

where disjointness conditions in the above for c and \bar{c} are satisfied by the fact that c is fresh as per the semantics, and disjointness of the constituent contexts comes from their original disjointness being defined (e.g., $\Delta_1 = \Delta_1' + \Delta_1''$).

– (forkNonLinear)

$$\frac{\#c}{(\mathcal{P} \mid \text{Ctx}\{\text{forkNonLinear}(\lambda x.t)\}; \mathcal{C}) \rightsquigarrow (\mathcal{P} \mid \text{Ctx}\{\bar{c}\} \mid t[c/x]^\circ; \mathcal{C}, \bar{c}^* = \langle \rangle, c^* = \langle \rangle)} \text{FORKNONLINEAR}$$

with process configuration typing:

$$\frac{\Gamma_1, \Delta_1 \vdash \text{Ctx}\{\text{forkNonLinear}(\lambda x.t)\} : B \quad \Gamma_2, \Delta_2 \vdash \mathcal{P} : A}{\Gamma_1 + \Gamma_2, \Delta_1 + \Delta_2 \vdash \mathcal{P} \mid \text{Ctx}\{\text{forkNonLinear}(\lambda x.t)\}^{\mathcal{R}} : A} \text{CONSR} \quad \text{(D.18)}$$

with channel configuration typing:

$$\Delta_1 + \Delta_2 \vdash \mathcal{C} \quad \text{(D.19)}$$

From typing $\Gamma_1, \Delta_1 \vdash \text{Ctx}\{\text{forkNonLinear}(\lambda x.t)\} : B$, and by lemma 6 there are two possibilities:

1. (second disjunct) (with $\text{Ctx}\{\text{forkNonLinear}(\lambda x.t)\}$ factorising into a graded binding) Gives $\neg \text{resourceAllocator}(\text{forkNonLinear}(\lambda x.t))$ which is a contradiction, thus we can conclude trivial by ex falso.
2. (first disjunct) There is a typing $\Gamma_1', \Delta_1' \vdash \text{forkNonLinear}(\lambda x.t) : A$, which by inversion must therefore be of the form:

$$\frac{\frac{\Gamma_1', x : \square_r(\mathbf{Chan} P), \Delta_1' \vdash t : 1}{\Gamma_1', \Delta_1' \vdash (\lambda x.t) : \square_r(\mathbf{Chan} P) \multimap 1} \multimap_i}{\Gamma_1', \Delta_1' \vdash \text{forkNonLinear}(\lambda x.t) : \square_r(\mathbf{Chan}(\bar{P}))} \multimap_e \quad \text{(D.20)}$$

where $A = \square_r(\mathbf{Chan}(\bar{P}))$ and $\Gamma_1'', x : \square_r(\mathbf{Chan}(\bar{P})), \Delta_1'' \vdash \mathit{Ctxt}\{x\} : B$ and where $\Gamma_1 = \Gamma_1' + \Gamma_1''$ and $\Delta_1 = \Delta_1' + \Delta_1''$.

For the goal of this theorem we construct the typings:

$$\frac{\overline{\emptyset, \bar{c} : * \mathbf{Chan}(\bar{P}) \vdash \bar{c} : \mathbf{Chan}(\bar{P})}^{\text{VAR}}}{\overline{\emptyset, \bar{c} : * \mathbf{Chan}(\bar{P}) \vdash [\bar{c}] : \square_r(\mathbf{Chan}(\bar{P}))}^{\text{PR}}} \quad (\text{D.21})$$

and

$$\frac{\overline{\emptyset, c : * \mathbf{Chan} P \vdash c : \mathbf{Chan} P}^{\text{VAR}}}{\overline{\emptyset, c : * \mathbf{Chan} P \vdash [c] : \square_r(\mathbf{Chan} P)}^{\text{PR}}} \quad (\text{D.22})$$

By Lemma 7 then $\Gamma_1'', \Delta_1'', \bar{c} : * \mathbf{Chan} \bar{P} \vdash \mathit{Ctxt}\{[\bar{c}]\} : B$.

We then construct the consequent channel configuration typing:

$$\frac{\overline{\Delta_1 + \Delta_2 \vdash \mathcal{C}}^{\text{(D.19)}}}{\overline{\Delta_1 + \Delta_2, c : * \mathbf{Chan} P \vdash \mathcal{C}, c^* = \langle \rangle}^{\text{EMPTY}}} \quad \text{EMPTY} \quad \frac{}{\overline{\Delta_1 + \Delta_2, c : * \mathbf{Chan} P, \bar{c} : * \mathbf{Chan} \bar{P} \vdash \mathcal{C}, c^* = \langle \rangle, \bar{c}^* = \langle \rangle}^{\text{EMPTY}}} \quad \text{EMPTY}$$

and process configuration:

$$\frac{\frac{\Gamma_1'', \Delta_1'', \bar{c} : * \mathbf{Chan} \bar{P} \vdash \mathit{Ctxt}\{[\bar{c}]\} : B \quad \Gamma_2, \Delta_2 \vdash \mathcal{P} : A}{\Gamma_1'' + \Gamma_2, \Delta_1'' + \Delta_2, \bar{c} : * \mathbf{Chan} \bar{P} \vdash \mathcal{P} \mid \mathit{Ctxt}\{[\bar{c}]\}^{\mathcal{R}} : A} \text{CONS}\mathcal{R}}{\Gamma_1', \Delta_1', c : * \mathbf{Chan} P \vdash t[[c]/x] : 1} \text{GRADED-SUBST}}{\frac{\Gamma_1' + \Gamma_1'' + \Gamma_2, \Delta_1' + \Delta_1'' + \Delta_2, c : * \mathbf{Chan} P, \bar{c} : * \mathbf{Chan} \bar{P} \vdash \mathcal{P} \mid \mathit{Ctxt}\{[\bar{c}]\}^{\mathcal{R}} \mid t[[c]/x]^{\circ} : A}{\Gamma_1 + \Gamma_2, \Delta_1 + \Delta_2, c : * \mathbf{Chan} P, \bar{c} : * \mathbf{Chan} \bar{P} \vdash \mathcal{P} \mid \mathit{Ctxt}\{[\bar{c}]\}^{\mathcal{R}} \mid t[[c]/x]^{\circ} : A} \text{SUB}}} \quad \equiv$$

where disjointness conditions in the above for c and \bar{c} are satisfied by the fact that c is fresh as per the semantics, and disjointness of the constituent contexts comes from their original disjointness being defined (e.g., $\Delta_1 = \Delta_1' + \Delta_1''$).

– (forkReplicate)

$$\frac{\#c_1 \dots \#c_n}{(\mathcal{P} \mid \mathit{Ctxt}\{\text{forkReplicate}[\lambda x.t] n\}; \mathcal{C}) \rightsquigarrow (\mathcal{P} \mid \mathit{Ctxt}\{([\bar{c}_1], \dots, [\bar{c}_n])\} \mid \text{Wch}(c_1 \mapsto \lambda x.t) \mid \dots \mid \text{Wch}(c_n \mapsto \lambda x.t); \mathcal{C}, c_1 = \langle \rangle, \dots, c_n = \langle \rangle, \bar{c}_1 = \langle \rangle, \dots, \bar{c}_n = \langle \rangle)} \text{FORKREPLICATE}$$

with process configuration typing:

$$\frac{\Gamma_1, \Delta_1 \vdash \mathit{Ctxt}\{\mathit{forkReplicate} [\lambda x.t] n\} : B \quad \Gamma_2, \Delta_2 \vdash \mathcal{P} : A}{\Gamma_1 + \Gamma_2, \Delta_1 + \Delta_2 \vdash \mathcal{P} \mid \mathit{Ctxt}\{\mathit{forkReplicate} [\lambda x.t] n\}^{\mathcal{R}} : A} \text{CONS}\mathcal{R} \quad (\text{D.23})$$

with channel configuration typing:

$$\Delta_1 + \Delta_2 \vdash \mathcal{C} \quad (\text{D.24})$$

From typing $\Gamma_1, \Delta_1 \vdash \mathit{Ctxt}\{\mathit{forkReplicate} [\lambda x.t] n\} : B$, and by lemma 6 there are two possibilities:

1. (second disjunct) (with $\mathit{Ctxt}\{\mathit{forkReplicate} [\lambda x.t] n\}$ factorising into a graded binding) Gives $\neg \text{resourceAllocator}(\mathit{forkReplicate} [\lambda x.t] n)$ which is a contradiction, thus we can conclude trivial by ex falso.
2. (first disjunct) then there is a typing $\Gamma_1'', \Delta_1'' \vdash \mathit{forkReplicate} [\lambda x.t] n : A$, which by inversion must therefore be of the form:

$$\frac{\frac{\frac{\Gamma_1'', x : \mathbf{Chan} P, \Delta_1'' \vdash t : 1}{\Gamma_1'', \Delta_1'' \vdash \lambda x.t : \mathbf{Chan} P \multimap 1} \multimap_i}{(0..n) * \Gamma_1'', \Delta_1'' \vdash [\lambda x.t] : \square_{0..n}(\mathbf{Chan} P \multimap 1)} \text{PR}}{(0..n) * \Gamma_1'', \Delta_1'' \vdash \mathit{forkReplicate} [\lambda x.t] : \mathbf{N} n \multimap \mathbf{Vec} n (\square_{0..1}(\mathbf{Chan}(\overline{P})))} \multimap_e}{(0..n) * \Gamma_1'', \Delta_1'' \vdash \mathit{forkReplicate} [\lambda x.t] n : \mathbf{Vec} n (\square_{0..1}(\mathbf{Chan}(\overline{P})))} \multimap_e \quad (\text{D.25})$$

where $A = \mathbf{Vec} n (\square_{0..1}(\mathbf{Chan}(\overline{P})))$ and $\Gamma_1', x : \mathbf{Vec} n (\square_{0..1}(\mathbf{Chan}(\overline{P})))$, $\Delta_1' \vdash \mathit{Ctxt}\{x\} : B$ and where $\Delta_1 = \Delta_1' + \Delta_1''$ and where $\Gamma_1 = \Gamma_1' + (0..n) * \Gamma_1''$. Furthermore by the $\text{ReceivePrefix}(P)$ constraint of $\mathit{forkReplicate}$ we know that $P = \mathbf{Recv} A' P'$.

For the goal of this theorem we construct the typings:

$$\frac{\dots}{\overline{c_1} : * \mathbf{Chan}(\mathbf{Send} A' P'), \dots, \overline{c_n} : * \mathbf{Chan}(\mathbf{Send} A' P') \vdash \langle [\overline{c_1}], \dots, [\overline{c_n}] \rangle : \mathbf{Vec} n (\square_{0..1}(\mathbf{Chan}(\mathbf{Send} A' P')))} \text{CONS}^n + \text{NIL}$$

By Lemma 7 then:

$$\Gamma_1', \Delta_1', \overline{c_1} : * \mathbf{Chan}(\mathbf{Send} A' P'), \dots, \overline{c_n} : * \mathbf{Chan}(\mathbf{Send} A' P') \vdash \mathit{Ctxt}\{\langle [\overline{c_1}], \dots, [\overline{c_n}] \rangle\} : B.$$

The consequent channel configuration has type derivation (and as a short hand we write $c_i : \mathbf{Chan} P_i$ for $c_1 : * \mathbf{Chan} P_1, \dots, c_n : * \mathbf{Chan} P_n$):

$$\frac{\Delta_1 + \Delta_2 \vdash \mathcal{C}}{\Delta_1 + \Delta_2, c_i : \mathbf{Chan}(\mathbf{Recv} A' P'), \overline{c_i} : * \mathbf{Chan}(\mathbf{Send} A' P') \vdash \mathcal{C}, c_1^* = \langle \rangle, \dots, c_n^* = \langle \rangle, \overline{c_1}^* = \langle \rangle, \dots, \overline{c_n}^* = \langle \rangle} \text{EMPTY}^*$$

We then construct the process configuration typing where $\Delta' = \Delta$, $c_i :^* \mathbf{Chan}(\mathbf{Recv} A' P')$, $\overline{c}_i :^* \mathbf{Chan}(\mathbf{Send} A' \overline{P}')$

$$\frac{\frac{\frac{\Gamma'_1, \Delta'_1, \overline{c}_i :^* \mathbf{Chan}(\mathbf{Send} A' \overline{P}') \vdash \mathit{Ctxt}\{\langle [\overline{c}_1], \dots, [\overline{c}_n] \rangle\} : B}{\Gamma_2, \Delta_2 \vdash \mathcal{P} : A}}{\Gamma'_1 + \Gamma_2, \Delta'_1 + \Delta_2, \overline{c}_i :^* \mathbf{Chan}(\mathbf{Send} A' \overline{P}') \vdash \mathcal{P} \mid \mathit{Ctxt}\{\langle [\overline{c}_1], \dots, [\overline{c}_n] \rangle\}^{\mathcal{R}} : A} \text{CONS}\mathcal{R}}{\frac{\Gamma''_1, x : \mathbf{Chan}(\mathbf{Recv} A' P'), \Delta''_1 \vdash t : 1}{\Gamma'_1 + (0..1 * \Gamma''_1) + \Gamma_2, \Delta'_1 + \Delta''_1 + \Delta_2, \Delta' \vdash \mathcal{P} \mid \mathit{Ctxt}\{\langle [\overline{c}_1], \dots, [\overline{c}_n] \rangle\}^{\mathcal{R}} \mid \mathbf{Wch}(c_1 \mapsto \lambda x.t) : A} \text{WATCH}}{\Gamma'_1 + (0..n * \Gamma''_1) + \Gamma_2, \Delta'_1 + \Delta''_1 + \Delta_2, \Delta' \vdash \mathcal{P} \mid \mathit{Ctxt}\{\langle [\overline{c}_1], \dots, [\overline{c}_n] \rangle\}^{\mathcal{R}} \mid \mathbf{Wch}(c_1 \mapsto \lambda x.t) \mid \dots \mid \mathbf{Wch}(c_n \mapsto \lambda x.t) : A} \text{WATCH}^*} \equiv \frac{\Gamma'_1 + (0..n * \Gamma''_1) + \Gamma_2, \Delta'_1 + \Delta''_1 + \Delta_2, \Delta' \vdash \mathcal{P} \mid \mathit{Ctxt}\{\langle [\overline{c}_1], \dots, [\overline{c}_n] \rangle\}^{\mathcal{R}} \mid \mathbf{Wch}(c_1 \mapsto \lambda x.t) \mid \dots \mid \mathbf{Wch}(c_n \mapsto \lambda x.t) : A}{\Gamma_1 + \Gamma_2, \Delta_1 + \Delta_2, \Delta' \vdash \mathcal{P} \mid \mathit{Ctxt}\{\langle [\overline{c}_1], \dots, [\overline{c}_n] \rangle\}^{\mathcal{R}} \mid \mathbf{Wch}(c_1 \mapsto \lambda x.t) \mid \dots \mid \mathbf{Wch}(c_n \mapsto \lambda x.t) : A} \equiv$$

since $0..1 + \dots + 0..1$ n-times is equal to $0..n$, satisfying the typing goal, and where disjointness conditions in the above for c_i and \overline{c}_i are satisfied by the fact that c_i are all fresh as per the semantics, and disjointness of the constituent contexts comes from their original disjointness being defined (e.g., $\Delta_1 = \Delta'_1 + \Delta''_1$).

– (watch)

$$\frac{(\mathcal{P} \mid \mathbf{Wch}(c \mapsto \lambda x.t); \mathcal{C}, c = v : \vec{v}) \rightsquigarrow (\mathcal{P} \mid t[c/x]^\circ; \mathcal{C}, c = v : \vec{v})}{\text{WATCHER}}$$

with process configuration typing:

$$\frac{\Gamma_1, x : \mathbf{Chan}(\mathbf{Recv} A' P), \Delta_1 \vdash t : 1 \quad \Gamma_2, \Delta_2 \vdash \mathcal{P} : A}{0..1 * \Gamma_1 + \Gamma_2, \Delta_1 + \Delta_2, c : \mathbf{Chan}(\mathbf{Recv} A' P) \vdash \mathcal{P} \mid \mathbf{Wch}(c \mapsto \lambda x.t) : A} \text{WATCH} \quad (\text{D.26})$$

and channel configuration typing:

$$\frac{\Delta_1 + \Delta_2 \vdash \mathcal{C} \quad \Delta_1 + \Delta_2 \vdash v : A'}{\Delta_1 + \Delta_2, c : (\mathbf{Recv} A' P) \vdash \mathcal{C}, c = v : \vec{v}} \quad (\text{D.27})$$

For the goal of this theorem we construct the typing:

$$\Gamma_1, x : \mathbf{Chan}(\mathbf{Recv} A' P), \Delta_1 \vdash t : 1 \quad (\text{D.32})$$

$$\frac{\frac{\frac{c : \mathbf{Chan}(\mathbf{Recv} A' P) \vdash c : \mathbf{Chan}(\mathbf{Recv} A' P)}{1..1 * \Gamma_1, \Delta_1, c : \mathbf{Chan}(\mathbf{Recv} A' P) \vdash t[c/x] : 1} \text{SUBST}}{0..1 * \Gamma_1, \Delta_1, c : \mathbf{Chan}(\mathbf{Recv} A' P) \vdash t[c/x] : 1} \text{APPROX}}{\text{VAR}} \quad (\text{D.28})$$

From this we can construct the channel configuration typing as just (D.27) and the process config typing:

$$\frac{\Gamma_2, \Delta_2 \vdash \mathcal{P} : A \quad 0..1 * \Gamma_1, \Delta_1, c : \mathbf{Chan}(\mathbf{Recv} A' P) \vdash t[c/x] : 1 \text{ (D.28)}}{0..1 * \Gamma_1 + \Gamma_2, \Delta_1 + \Delta_2, c : \mathbf{Chan}(\mathbf{Recv} A' P) \vdash \mathcal{P} \mid t[c/x]^\circ : A} \text{ WATCH}$$

and the premise channel configuration provides the consequent channel configuration.

– (forkMulticast)

$$\frac{(\mathcal{P} \mid \mathit{Ctxt}\{\mathbf{forkMulticast}(\lambda x.t) n\}; C) \rightsquigarrow (\mathcal{P} \mid \mathit{Ctxt}\{\bar{c}_1, \dots, \bar{c}_n\} \mid t[c/x]^\circ \mid \mathbf{Fwd}(c \mapsto c_1, \dots, c_n); \mathcal{C}, c = \langle \rangle, \bar{c} = \langle \rangle, c_1 = \langle \rangle, \dots, c_n = \langle \rangle, \bar{c}_1 = \langle \rangle, \dots, \bar{c}_n = \langle \rangle)}{\text{FORKMULTICAST}}$$

with process configuration typing: process configuration typing:

$$\frac{\Gamma_1, \Delta_1 \vdash \mathit{Ctxt}\{\mathbf{forkMulticast}(\lambda x.t) n\} : B \quad \Gamma_2, \Delta_2 \vdash \mathcal{P} : A}{\Gamma_1 + \Gamma_2, \Delta_1 + \Delta_2 \vdash \mathcal{P} \mid \mathit{Ctxt}\{\mathbf{forkMulticast}(\lambda x.t) n\}^{\mathcal{R}} : A} \text{ CONSR} \quad (\text{D.29})$$

with channel configuration typing:

$$\Delta_1 + \Delta_2 \vdash \mathcal{C} \quad (\text{D.30})$$

From typing $\Gamma_1, \Delta_1 \vdash \mathit{Ctxt}\{\mathbf{forkMulticast}(\lambda x.t) n\} : B$, and by lemma 6 there are two possibilities:

1. (second disjunct) (with $\mathit{Ctxt}\{\mathbf{forkReplicate}[\lambda x.t] n\}$ factorising into a graded binding) Gives $\neg \text{resourceAllocator}(\mathbf{forkReplicate}[\lambda x.t] n)$ which is a contradiction, thus we can conclude trivial by ex falso.
2. (first disjunct) Then there is a typing $\Gamma'_1, \Delta'_1 \vdash \mathbf{forkMulticast}(\lambda x.t) n : A$, which by inversion must therefore be of the form:

$$\frac{\frac{\Gamma'_1, x : \mathbf{Chan}(\mathbf{Graded}_n P'), \Delta'_1 \vdash t : 1}{\Gamma'_1, \Delta'_1 \vdash \lambda x.t : \mathbf{Chan}(\mathbf{Graded}_n P') \multimap 1} \multimap_i}{\frac{\Gamma'_1, \Delta'_1 \vdash \mathbf{forkMulticast}(\lambda x.t) : \mathbf{N} n \multimap \mathbf{Vec} n(\mathbf{Chan} \bar{P}')}{\Gamma'_1, \Delta'_1 \vdash \mathbf{forkMulticast}(\lambda x.t) n : \mathbf{Vec} n(\mathbf{Chan} \bar{P}')} \multimap_e} \multimap_e \quad (\text{D.31})$$

where $A = \mathbf{Vec} n(\mathbf{Chan} \bar{P})$ and $\Gamma''_1, x : \mathbf{Vec} n(\mathbf{Chan} \bar{P}), \Delta''_1 \vdash \mathit{Ctxt}\{x\} : B$ and $P = \mathbf{Graded} n P'$ and where $\Gamma_1 = \Gamma'_1 + \Gamma''_1$ and where $\Delta_1 = \Delta'_1 + \Delta''_1$.

For the goal of this theorem we construct the typing:

$$\frac{\dots}{\overline{c_1} : \mathbf{Chan} \overline{P}, \dots, \overline{c_n} : \mathbf{Chan} \overline{P} \vdash \langle \overline{c_1}, \dots, \overline{c_n} \rangle : \mathbf{Vec} \, n \, (\mathbf{Chan} \overline{P})} \text{CONS}^n + \text{NIL}$$

By Lemma 7 then $\Gamma_1'', \Delta_1'', \Delta' \vdash \text{Ctxt}\{\langle \overline{c_1}, \dots, \overline{c_n} \rangle\} : B$

The consequent channel configuration has type derivation where we write $\Delta' = c_1 : \mathbf{Chan} \, P, \dots, c_n : \mathbf{Chan} \, P$ and $\overline{\Delta'} = \overline{c_1} : \mathbf{Chan} \, \overline{P}, \dots, \overline{c_n} : \mathbf{Chan} \, \overline{P}$.

$$\frac{\Delta_1 + \Delta_2 \vdash \mathcal{C}}{\Delta_1 + \Delta_2, c : \mathbf{Chan} \, (\text{Graded}_n \, P), \overline{c} : \mathbf{Chan} \, (\overline{\text{Graded}_n \, P}), \Delta', \overline{\Delta'} \vdash \mathcal{C}, c = \langle \rangle, \overline{c} = \langle \rangle, c_1 = \langle \rangle, \dots, c_n = \langle \rangle, \overline{c_1} = \langle \rangle, \dots, \overline{c_n} = \langle \rangle} \text{EMPTY}^*$$

We then construct the process configuration typing:

$$\frac{\frac{\frac{\frac{\Gamma_1'', \Delta_1'', \Delta' \vdash \text{Ctxt}\{\langle \overline{c_1}, \dots, \overline{c_n} \rangle\} : B}{\Gamma_2, \Delta_2 \vdash \mathcal{P} : A}}{\Gamma_1'' + \Gamma_2, \Delta_1'' + \Delta_2, \Delta' \vdash \mathcal{P} \mid \text{Ctxt}\{\langle \overline{c_1}, \dots, \overline{c_n} \rangle\}^{\mathcal{R}} : A} \text{CONS}\mathcal{R}}{\frac{\Gamma_1', x : \mathbf{Chan} \, (\text{Graded}_n \, P'), \Delta_1' \vdash t : 1}{\Gamma_1', \Delta_1', c : \mathbf{Chan} \, (\text{Graded}_n \, P') \vdash t[c/x] : 1} \text{SUBST}}{\Gamma_1' + \Gamma_1'' + \Gamma_2, \Delta_1' + \Delta_1'' + \Delta_2, c : \mathbf{Chan} \, (\text{Graded}_n \, P'), \overline{\Delta'} \vdash \mathcal{P} \mid \text{Ctxt}\{\langle \overline{c_1}, \dots, \overline{c_n} \rangle\}^{\mathcal{R}} \mid t[c/x]^\circ : A} \text{CONS}\mathcal{R}}}{\frac{\Gamma_1' + \Gamma_1'' + \Gamma_2, \Delta_1' + \Delta_1'' + \Delta_2, \Delta', \overline{\Delta'}, c : \mathbf{Chan} \, (\text{Graded}_n \, P'), \overline{c} : \mathbf{Chan} \, (\overline{\text{Graded}_n \, P'}) \vdash \mathcal{P} \mid \text{Ctxt}\{\langle \overline{c_1}, \dots, \overline{c_n} \rangle\}^{\mathcal{R}} \mid t[c/x]^\circ \mid \mathbf{Fwd} \, (c \mapsto c_1, \dots, c_n) : A}{\Gamma_1 + \Gamma_2, \Delta_1 + \Delta_2, \Delta', \overline{\Delta'}, c : \mathbf{Chan} \, (\text{Graded}_n \, P'), \overline{c} : \mathbf{Chan} \, (\overline{\text{Graded}_n \, P'}) \vdash \mathcal{P} \mid \text{Ctxt}\{\langle \overline{c_1}, \dots, \overline{c_n} \rangle\}^{\mathcal{R}} \mid t[c/x]^\circ \mid \mathbf{Fwd} \, (c \mapsto c_1, \dots, c_n) : A} \text{FWD}} \equiv$$

satisfying the goal.

– (forward)

$$\frac{}{(\mathcal{P} \mid \mathbf{Fwd} \, (c \mapsto c_1, \dots, c_n); \mathcal{C}, c = [v] : \vec{v}, \overline{c_1} = \vec{v}_1, \dots, \overline{c_n} = \vec{v}_n) \rightsquigarrow (\mathcal{P} \mid \mathbf{Fwd} \, (c \mapsto c_1, \dots, c_n); \mathcal{C}, c = \vec{v}, \overline{c_1} = \vec{v}_1 : v, \dots, \overline{c_n} = \vec{v}_n : v)} \text{FORWARDER}$$

with process configuration typing:

$$\frac{\text{Sends}(P) \quad \Gamma, \Delta \vdash \mathcal{P} : A}{\Gamma, \Delta, c : \mathbf{Chan} \, (\text{Graded}_n \, \overline{P}), c_i : \mathbf{Chan} \, P \vdash \mathcal{P} \mid \mathbf{Fwd} \, (c \mapsto c_1, \dots, c_n) : A} \text{FWD} \quad (\text{D.32})$$

(specialisation of the FWD rule since $\text{Sends}(\overline{\mathbf{Recv}} \, A' \, \overline{P}) = \text{Sends}(\overline{P})$) and with channel configuration typing:

$$\frac{\frac{\Delta_0 \vdash v : A'}{\Delta_0 \vdash [v] : \square_n A'} \text{PR} \quad \frac{\Delta_0 \vdash v : A' \quad \Delta_0, \Delta_1, c_1 : \mathbf{Chan} \, P, \dots, c_n : \mathbf{Chan} \, P, c : \mathbf{Chan} \, \overline{P} \vdash \mathcal{C}, \overline{c_1} = \vec{v}_1 : v, \dots, \overline{c_n} = \vec{v}_n : v, c = \vec{v}}{\Delta_0, \Delta_1, c_1 : \mathbf{Chan} \, (\mathbf{Send} \, A \, P), \dots, c_n : \mathbf{Chan} \, (\mathbf{Send} \, A \, P), c : \mathbf{Chan} \, \overline{P} \vdash \mathcal{C}, \overline{c_1} = \vec{v}_1, \dots, \overline{c_n} = \vec{v}_n, c = \vec{v}} \text{SEND}^n}{\Delta_0, \Delta_1, c_1 : \mathbf{Chan} \, (\mathbf{Send} \, A \, P), \dots, c_n : \mathbf{Chan} \, (\mathbf{Send} \, A \, P), c : \mathbf{Chan} \, (\mathbf{Recv} \, (\square_n A') \, \overline{P}) \vdash \mathcal{C}, \overline{c_1} = \vec{v}_1, \dots, \overline{c_n} = \vec{v}_n, c = [v] : \vec{v}} \text{RECV}} \quad (\text{D.33})$$

where $\Delta = \Delta_0, \Delta_1$.

The goal channel configuration typing is then given by the premise of (D.33):

$$\frac{\Delta_0 \vdash v : A' \quad \Delta_1, c_1 : \mathbf{Chan} P, \dots, c_n : \mathbf{Chan} P, c : \mathbf{Chan} \bar{P} \vdash \mathcal{C}, \bar{c}_1 = \vec{v}_1 : v, \dots, \bar{c}_n = \vec{v}_n : v, c = \vec{v}}{\Delta_0, \Delta_1, c_1 : \mathbf{Chan} (\mathbf{Send} A P), \dots, c_n : \mathbf{Chan} (\mathbf{Send} A P), c : \mathbf{Chan} \bar{P} \vdash \mathcal{C}, \bar{c}_1 = \vec{v}_1, \dots, \bar{c}_n = \vec{v}_n, c = \vec{v}}^{\mathbf{SEND}^n}$$

From the definition of $\mathbf{Sends}(\overline{\mathbf{Recv} A' P}) = \mathbf{Sends}(\bar{P})$ and $\mathbf{Sends}(\mathbf{Send} A' P) = \mathbf{Sends}(P)$ therefore we can define the process configuration typing as:

$$\frac{\mathbf{Sends}(\bar{P}) \quad \mathbf{Sends}(P) \quad \Gamma, \Delta \vdash \mathcal{P} : A}{\Gamma, \Delta, \bar{c} : \mathbf{Chan} (\bar{P}), c_1 : \mathbf{Chan} P, \dots, c_n : \mathbf{Chan} P \vdash \mathcal{P} \mid \mathbf{Fwd}(c \mapsto c_1, \dots, c_n) : A}^{\mathbf{FWD}}$$

□

Appendix E. Related work further details; comparison with GV

The construction in GV of the `replicate` primitive can also be captured in Granule. Their construction [21, Appendix B] can be written in Granule with $\square A = A \text{ [Many]}$ (where $\mathcal{R} = \{\text{Zero}, \text{One}, \text{Many}\}$ is a semiring for capturing linear/non-linear behaviour [17]) but only if we both turn *off* the resource allocator restriction on promotion that we introduced in Section 4, which was critical for soundness in our call-by-value semantics, and use call-by-name semantics, both of which are provided by a language extension:

```

1  language CBN -- Call-by-Name semantics, also turns off the restriction
2                -- to non-resource allocating terms in promotion
3
4  -- GV `replicate`
5  replicate : ∀ {s : Protocol}
6              . (LChan (Send ((LChan (Dual s)) [Many]) End)
7                , (LChan s → ()) [Many]) → LChan End
8  replicate (x, f) =
9  let [g] = f in send x [forkLinear g]
10
11 -- Dual `request` primitive
12 request : ∀ {s : Protocol}
13           . LChan (Recv ((LChan s) [Many]) End) → LChan s
14 request s = let (w , s) = recv s
15             in let () = close s
16             in let [y] = w in y

```