

Asynchronous Functional Sessions: Cyclic and Concurrent

Bas van den Heuvel Jorge A. Pérez

University of Groningen, The Netherlands

{b.van.den.heuvel, j.a.perez} @ rug.nl

We present Concurrent GV (CGV), a functional calculus with message-passing concurrency governed by session types. With respect to prior calculi, CGV has increased support for concurrent evaluation and for cyclic network topologies. The design of CGV draws on APCP, a session-typed asynchronous π -calculus developed in prior work. Technical contributions are (i) the syntax, semantics, and type system of CGV; (ii) a correct translation of CGV into APCP; (iii) a technique for establishing deadlock-free CGV programs, by resorting to APCP’s priority-based type system.

1 Introduction

The goal of this paper is to introduce a new functional calculus with message-passing concurrency governed by linearity and session types. Our work contributes to a research line initiated by Gay and Vasconcelos [8], who proposed a functional calculus with sessions here referred to as λ^{sess} ; this line of work has received much recent attention thanks to Wadler’s GV calculus [26], which is a variation of λ^{sess} .

Our new calculus is dubbed Concurrent GV (CGV); with respect to previous work, it presents three intertwined novelties: asynchronous (buffered) communication; a highly concurrent reduction strategy; and thread configurations with cyclic topologies. The design of CGV rests upon a solid basis: an operationally correct translation into APCP (Asynchronous Priority-based Classical Processes), a session-typed π -calculus in which asynchronous processes communicate by forming cyclic networks [13].

We discuss the salient features of CGV by example, using a simplified syntax. As in λ^{sess} , communication in CGV is asynchronous: send operations place their messages in buffers, and receive operations read the messages from these buffers. Let us write $\text{send}(u, x)$ to denote the output of message u along channel x , and $\text{recv } y$ to denote an input on y . The following program expresses the parallel composition (\parallel) of two threads:

$$\begin{array}{l} \text{let } x = \text{send}(u, x) \text{ in} \\ \text{let } (v, y) = \text{recv } y \text{ in } () \end{array} \parallel \begin{array}{l} \text{let } y = \underline{\text{send}}(w, y) \text{ in } () \end{array}$$

In variants of λ^{sess} with *synchronous* communication, such as GV and Kokke and Dardha’s PGV [17, 18], this program is stuck: the send on y (underlined, on the right) cannot synchronize with the receive on y (on the left): it is blocked by the send on x (on the left), and there is no receive on x . In contrast, in CGV the send on x can be buffered after which the communication on y can take place.

In CGV, reduction is “more concurrent” than usual call-by-value or call-by-name strategies. Consider the following program:

$$\left(\begin{array}{l} \lambda x. \text{let } (u, y) = \text{recv } y \text{ in} \\ \text{let } x = \text{send}(u, x) \text{ in } () \end{array} \right) (\text{send}(v, z))$$

In λ^{sess} , reduction is call-by-value and so the function on x can only be applied on a value. However, the function’s parameter (send on z) is not a value, so it needs to be evaluated before the function on x can

be applied. Hence, this program can only be evaluated in one order: first the send on z , then the receive on y . In contrast, the semantics of CGV evaluates a function and its parameters *concurrently*: the send on z and the receive on y can be evaluated in any order. Note that asynchrony plays no role here: both buffering a message and synchronous communication entail a reduction in the function’s parameter.

The third novelty is cyclic thread configurations: threads can be connected by channels to form cyclic networks. Consider the following program:

$$\begin{array}{c} \text{let } (u, x) = \text{recv } x \text{ in} \\ \quad \text{let } y = \text{send } (u, y) \text{ in } () \end{array} \quad \parallel \quad \begin{array}{c} \text{let } x = \text{send } (v, x) \text{ in} \\ \quad \text{let } (w, y) = \text{recv } y \text{ in } () \end{array}$$

Here we have two threads connected on channels x and y , thus forming a cyclic thread configuration. Clearly, this program is deadlock-free. In λ^{sess} , the program is well-typed, but there is no deadlock-freedom guarantee: the type system of λ^{sess} admits deadlocked cyclic thread configurations. In GV and Fowler *et al.*’s EGV [7] (an extension of Fowler’s AGV [6]) there is a deadlock-freedom guarantee for well-typed programs; however, their type systems only support *tree-shaped* thread configurations—this limitation is studied in [4, 5]. Hence, the program above is not well-typed in GV and EGV.

These novelties are *intertwined*, in the following sense. Asynchronous communication reduces the synchronization points in programs (as output-like operations are non-blocking), therefore increasing concurrent evaluation. In turn, reduced synchronization points can streamline verification techniques for deadlock-freedom based on *priorities* [16, 22, 23, 3], which unlock the analysis of process networks with cyclic topologies. Indeed, in an asynchronous setting only input-like operations require priorities.

We endow CGV with a type system with functional types and session types; we opted for a design in which well-typed terms enjoy subject reduction / type preservation but not deadlock-freedom. To validate our semantic design and attain the three novelties motivated above, we resort to APCP. In our developments, APCP operates as a “low-level” reference programming model. We give a typed translation of CGV into APCP, which satisfies strong correctness properties, in the sense of Gorla [10]. In particular, it enjoys operational correspondence, which provides a significant sanity check to justify our key design decisions in CGV’s operational semantics. Interestingly, using our correct translation and the deadlock-freedom guarantees for well-typed processes in APCP, we obtain a technique for transferring the deadlock-freedom property to CGV programs. That is, given a CGV program C , we prove that if the APCP translation of C is typable (and hence, deadlock-free), then C itself is deadlock-free. This result thus delineates a class of deadlock-free CGV programs that includes cyclic thread configurations.

In summary, this paper presents the following technical contributions: (1) CGV, a new functional calculus with session-based asynchronous concurrency; (2) A typed translation of CGV into APCP, which is proven to satisfy well-studied encodability criteria; (3) A transference result for the deadlock-freedom property from APCP to CGV programs. An extended version contains omitted technical details [14].

2 Concurrent GV

2.1 Syntax and Semantics

The main syntactic entities in CGV are *terms*, *runtime terms*, and *configurations*. Intuitively, terms reduce to runtime terms; configurations correspond to the parallel composition of a main thread and several child threads, each executing a runtime term. Buffered messages are part of configurations. We define two reduction relations: one is on terms, which is then subsumed by reduction on configurations.

The syntax of terms (L, M, N) is given and described in Figure 1. We use x, y, \dots for *variables*; we write *endpoint* to refer to a variable used for session operations (send, receive, select, offer). Let

Terms (L, M, N) :	
x (variable)	new (create new channel)
$()$ (unit value)	$\text{spawn } M$ (execute pair M in parallel)
$\lambda x. M$ (abstraction)	$\text{send } (M, N)$ (send M along N)
$M N$ (application)	$\text{recv } M$ (receive along M)
(M, N) (pair construction)	$\text{select } \ell M$ (select label ℓ along M)
$\text{let } (x, y) = M \text{ in } N$ (pair deconstruction)	$\text{case } M \text{ of } \{i : M\}_{i \in I}$ (offer labels in I along M)
.....	
Runtime terms (L, M, N) and reduction contexts (\mathcal{R}) :	
$L, M, N ::= \dots \mid M\{N/x\} \mid \text{send}'(M, N)$	
$\mathcal{R} ::= [] \mid \mathcal{R} M \mid \text{spawn } \mathcal{R} \mid \text{send } \mathcal{R} \mid \text{recv } \mathcal{R} \mid \text{let } (x, y) = \mathcal{R} \text{ in } M$	
$\mid \text{select } \ell \mathcal{R} \mid \text{case } \mathcal{R} \text{ of } \{i : M\}_{i \in I} \mid \mathcal{R}\{M/x\} \mid M\{\mathcal{R}/x\} \mid \text{send}'(M, \mathcal{R})$	
.....	
Structural congruence for terms (\equiv_M) and term reduction (\longrightarrow_M) :	
SC-SUBEXT	$x \notin \text{fn}(\mathcal{R}) \Rightarrow (\mathcal{R}[M])\{N/x\} \equiv_M \mathcal{R}\{M\{N/x\}\}$
E-LAM	$(\lambda x. M) N \longrightarrow_M M\{N/x\}$
E-PAIR	$\text{let } (x, y) = (M_1, M_2) \text{ in } N \longrightarrow_M N\{M_1/x, M_2/y\}$
E-SUBSTNAME	$M\{y/x\} \longrightarrow_M M\{y/x\}$
E-NAMESUBST	$x\{M/x\} \longrightarrow_M M$
E-SEND	$\text{send } (M, N) \longrightarrow_M \text{send}'(M, N)$
E-LIFT	$M \longrightarrow_M N \Rightarrow \mathcal{R}[M] \longrightarrow_M \mathcal{R}[N]$
E-LIFTSC	$M \equiv_M M' \wedge M' \longrightarrow_M N' \wedge N' \equiv_M N \Rightarrow M \longrightarrow_M N$

Figure 1: The CGV term language.

$\text{fn}(M)$ denote the free variables of a term. All variables are free unless bound: $\lambda x. M$ binds x in M , and $\text{let } (x, y) = M \text{ in } N$ binds x and y in N . We introduce syntactic sugar for applications of abstractions: $\text{let } x = M \text{ in } N$ denotes $(\lambda x. N) M$. For $(\lambda x. M) N$, we assume $x \notin \text{fn}(N)$, and for $\text{let } (x, y) = M \text{ in } N$, we assume $x \neq y$ and $x, y \notin \text{fn}(M)$.

Figure 1 also gives the reduction semantics of CGV terms (\longrightarrow_M) , which relies on runtime terms (L, M, N) , reduction contexts (\mathcal{R}) , and structural congruence (\equiv_M) . Note that this semantics comprises the functional fragment of CGV; we define the concurrent semantics of CGV hereafter.

Runtime terms, whose syntax extends that of terms, guide the evaluation strategy of CGV; we discuss an example evaluation of a term using runtime terms after introducing the reduction rules (Example 2.1). Explicit substitution $M\{N/x\}$ enables the concurrent execution of a function and its parameters. The intermediate primitive $\text{send}'(M, N)$ enables N to reduce to an endpoint; the send primitive takes a pair of terms as an argument, inside which reduction is not permitted (cf. [8]). Reduction contexts define the

Markers (ϕ), messages (m, n), configurations (C, D, E), thread (\mathcal{F}) and configuration (\mathcal{G}) contexts:	
$\phi ::= \blacklozenge \mid \diamond$	$m, n ::= \mathbb{M} \mid \ell$
$C, D, E ::= \phi \mathbb{M} \mid C \parallel D \mid (\mathbf{v}x[\vec{m}]y)C \mid C\{\mathbb{M}/x\}$	
$\mathcal{F} ::= \phi \mathcal{R} \mid C\{\mathcal{R}/x\}$	$\mathcal{G} ::= [] \mid \mathcal{G} \parallel C \mid (\mathbf{v}x[\vec{m}]y)\mathcal{G} \mid \mathcal{G}\{\mathbb{M}/x\}$
.....	
Structural congruence for configurations (\equiv_c):	
SC-TERMSC	$\mathbb{M} \equiv_{\mathbb{M}} \mathbb{M}' \Rightarrow \phi \mathbb{M} \equiv_c \phi \mathbb{M}'$
SC-RESSWAP	$(\mathbf{v}x[\varepsilon]y)C \equiv_c (\mathbf{v}y[\varepsilon]x)C$
SC-RESCOMM	$(\mathbf{v}x[\vec{m}]y)(\mathbf{v}z[\vec{n}]w)C \equiv_c (\mathbf{v}z[\vec{n}]w)(\mathbf{v}x[\vec{m}]y)C$
SC-RESEXT	$x, y \notin \text{fn}(C) \Rightarrow (\mathbf{v}x[\vec{m}]y)(C \parallel D) \equiv_c C \parallel (\mathbf{v}x[\vec{m}]y)D$
SC-RESNIL	$x, y \notin \text{fn}(C) \Rightarrow (\mathbf{v}x[\varepsilon]y)C \equiv_c C$
SC-SEND'	$(\mathbf{v}x[\vec{m}]y)(\hat{\mathcal{F}}[\text{send}'(\mathbb{M}, x)] \parallel C) \equiv_c (\mathbf{v}x[\mathbb{M}, \vec{m}]y)(\hat{\mathcal{F}}[x] \parallel C)$
SC-SELECT	$(\mathbf{v}x[\vec{m}]y)(\mathcal{F}[\text{select } \ell.x] \parallel C) \equiv_c (\mathbf{v}x[\ell, \vec{m}]y)(\mathcal{F}[x] \parallel C)$
SC-PARNIL	$C \parallel \diamond() \equiv_c C$
SC-PARCOMM	$C \parallel D \equiv_c D \parallel C$
SC-PARASSOC	$C \parallel (D \parallel E) \equiv_c (C \parallel D) \parallel E$
SC-CONFSUBST	$\phi(\mathbb{M}\{\mathbb{N}/x\}) \equiv_c (\phi \mathbb{M})\{\mathbb{N}/x\}$
SC-CONFSUBSTEXT	$x \notin \text{fn}(\mathcal{G}) \Rightarrow (\mathcal{G}[C])\{\mathbb{M}/x\} \equiv_c \mathcal{G}\{C\{\mathbb{M}/x\}\}$
.....	
Configuration reduction (\longrightarrow_c):	
E-NEW	$\mathcal{F}[\text{new}] \longrightarrow_c (\mathbf{v}x[\varepsilon]y)(\mathcal{F}[(x, y)])$
E-SPAWN	$\hat{\mathcal{F}}[\text{spawn } (\mathbb{M}, \mathbb{N})] \longrightarrow_c \hat{\mathcal{F}}[\mathbb{N}] \parallel \diamond \mathbb{M}$
E-RECV	$(\mathbf{v}x[\vec{m}, \mathbb{M}]y)(\hat{\mathcal{F}}[\text{recv } y] \parallel C) \longrightarrow_c (\mathbf{v}x[\vec{m}]y)(\hat{\mathcal{F}}[(\mathbb{M}, y)] \parallel C)$
E-CASE	$j \in I \Rightarrow (\mathbf{v}x[\vec{m}, j]y)(\mathcal{F}[\text{case } y \text{ of } \{i : \mathbb{M}_i\}_{i \in I}] \parallel C) \longrightarrow_c (\mathbf{v}x[\vec{m}]y)(\mathcal{F}[\mathbb{M}_j y] \parallel C)$
E-LIFTC	$C \longrightarrow_c C' \Rightarrow \mathcal{G}[C] \longrightarrow_c \mathcal{G}[C']$
E-LIFTM	$\mathbb{M} \longrightarrow_{\mathbb{M}} \mathbb{M}' \Rightarrow \mathcal{F}[\mathbb{M}] \longrightarrow_c \mathcal{F}[\mathbb{M}']$
E-CONFLIFTSC	$C \equiv_c C' \wedge C' \longrightarrow_c D' \wedge D' \equiv_c D \Rightarrow C \longrightarrow_c D$

Figure 2: The CGV configuration language.

non-blocking parts of terms, where subterms may reduce. We write $\mathcal{R}[\mathbb{M}]$ to denote the runtime term obtained by replacing the hole $[]$ in \mathcal{R} by \mathbb{M} , and $\text{fn}(\mathcal{R})$ to denote $\text{fn}(\mathcal{R}[])$; we will use similar notation for other kinds of contexts later.

We discuss the reduction rules. The Structural congruence rule SC-SUBEXT allows the scope extrusion of explicit substitutions along reduction contexts. Rule E-LAM enforces application, resulting in an explicit substitution. Rule E-PAIR unpacks the elements of a pair into two explicit substitutions

(arbitrarily ordered, due to the syntactical assumptions introduced above). Rules E-SUBSTNAME and E-NAMESUBST convert explicit substitutions of or on variables into standard substitutions. Rule E-SEND reduces a `send` into a `send'` primitive. Rules E-LIFT and E-LIFTSC close term reduction under contexts and structural congruence, respectively. We write $M \xrightarrow{M}^k N$ to denote that M reduces to N in k steps.

Example 2.1. We illustrate the evaluation of terms using runtime terms through the following example, which contains a `send` primitive and nested abstractions and applications. In each reduction step, we underline the subterm that reduces and give the applied rule:

$$\begin{array}{ll}
(\lambda x. \underline{\text{send}}((), x)) ((\lambda y. y) z) & \text{(E-LAM)} \\
\longrightarrow_M (\underline{\text{send}}((), x)) \{((\lambda y. y) z)/x\} & \text{(E-SEND)} \\
\longrightarrow_M (\underline{\text{send}'((), x)}) \{((\lambda y. y) z)/x\} & \text{(E-LAM)} \\
\longrightarrow_M (\underline{\text{send}'((), x)}) \{(y\{z/y\})/x\} & \text{(SC-SUBEXT)} \\
\equiv_M \underline{\text{send}'((), x\{(y\{z/y\})/x\})} & \text{(E-NAMESUBST)} \\
\longrightarrow_M \underline{\text{send}'((), y\{z/y\})} & \text{(E-SUBSTNAME)} \\
\longrightarrow_M \underline{\text{send}'((), z)} &
\end{array}$$

Notice how the `send` primitive needs to reduce to a `send'` runtime primitive such that the explicit substitution of x can be applied. Also, note that the concurrency of CGV allows many more paths of reduction.

Note that the concurrent evaluation strategy of CGV may also be defined without explicit substitutions. In principle, this would require additional reduction contexts specific to applications on abstractions and pair deconstruction, as well as variants of Rules E-SUBSTNAME and E-NAMESUBST specific to these contexts. However, it is not clear how to define scope extrusion (Rule SC-SUBEXT) for such a semantics. Hence, we find that using explicit substitutions drastically simplifies the semantics of CGV.

Concurrency in CGV allows the parallel execution of terms that communicate through buffers. The syntax of configurations (C, D, E) is given in Figure 2. The configuration ϕM denotes a *thread*: a concurrently executed term. The thread marker helps to distinguish the *main* thread ($\phi = \blacklozenge$) from *child* threads ($\phi = \blacklozenge$). The configuration $C \parallel D$ denotes parallel composition. The configuration $(\nu x[\vec{m}]y)C$ denotes a *buffered restriction*: it connects the endpoints x and y through a buffer $[\vec{m}]$, binding x and y in C . The buffer's content, \vec{m} , is a sequence of messages (terms and labels). Buffers are directed: in $x[\vec{m}]y$, messages can be added to the front of the buffer on x , and they can be taken from the back of the buffer on y . We write $[\varepsilon]$ for the empty buffer. The configuration $C\{M/x\}$ lifts explicit substitution to the level of configurations: this allows spawning and sending terms under explicit substitution, such that the substitution can be moved to the context of the spawned or sent term.

The reduction semantics for configurations (\longrightarrow_C , also in Figure 2) relies on thread and configuration contexts (\mathcal{F} and \mathcal{G} , respectively) and structural congruence (\equiv_C). We write $\hat{\mathcal{F}}$ to denote a thread context in which the hole does not occur under explicit substitution, i.e. the context is not constructed using the clause $\mathcal{R}\{M/x\}$; this is used in rules for `send'`, `spawn`, and `recv`, effectively forcing the scope extrusion of explicit substitutions when terms are moved between contexts (cf. Example 2.4).

We comment on some of the congruences and reduction rules. Rule SC-RESSWAP allows to swap the direction of an empty buffer; this way, the endpoint that could read from the buffer before the swap can now write to it. Rule SC-RESCOMM allows to interchange buffers, and Rule SC-RESEXT allows to extrude their scope. Rule SC-RESNIL garbage collects buffers of closed sessions. Rule SC-CONFSUBST lifts explicit substitution at the level of terms to the level of threads, and

Rule SC-CONFSUBSTEXT allows the scope extrusion of explicit substitution along configuration contexts. Notably, putting messages in buffers is not a reduction: Rules SC-SEND' and SC-SELECT equate sends and selects on an endpoint x with terms and labels in the buffer for x , as asynchronous outputs are computationally equivalent to messages in buffers.

Reduction rule E-NEW creates a new buffer, leaving a reference to the newly created endpoints in the thread. Rule E-SPAWN spawns a child thread (the parameter pair's first element) and continues (as the pair's second element) inside the calling thread. Rule E-RECV retrieves a term from a buffer, resulting in a pair containing the term and a reference to the receiving endpoint. Rule E-CASE retrieves a label from a buffer, resulting in a function application of the label's corresponding branch to a reference to the receiving endpoint. There are no reduction rules for closing sessions, as they are closed silently. We write $C \xrightarrow{c}_C^k D$ to denote that C reduces to D in k steps. Also, we write \xrightarrow{c}_C^+ to denote the transitive closure of \xrightarrow{c}_C (i.e., reduction in at least one step).

We illustrate CGV's semantics by giving some examples. The following discusses a cyclic thread configuration which does not deadlock due to asynchrony:

Example 2.2. Consider configuration C_1 below, in which two threads are spawned and cyclically connected through two channels. One thread first sends on the first channel and then receives on the second, while the other thread first sends on the second channel and then receives on the first. Under synchronous communication, this would determine a configuration that deadlocks; however, under asynchronous communication, this is not the case (cf. the third example in Sec. 1). We detail some interesting reductions:

$$\begin{aligned}
C_1 &= \blacklozenge(\text{let } (f, g) = \text{new in let } (h, k) = \text{new in spawn} \left(\begin{array}{l} \text{let } f' = (\text{send } (u, f)) \text{ in} \\ \text{let } (v', h') = (\text{recv } h) \text{ in } (), \\ \text{let } k' = (\text{send } (v, k)) \text{ in} \\ \text{let } (u', g') = (\text{recv } g) \text{ in } () \end{array} \right)) \\
&\xrightarrow{c}_C^8 (\mathbf{v}x[\varepsilon]y)(\mathbf{v}w[\varepsilon]z)(\blacklozenge(\text{spawn} \left(\begin{array}{l} \text{let } f' = (\text{send } (u, x)) \text{ in} \\ \text{let } (v', h') = (\text{recv } w) \text{ in } (), \\ \text{let } k' = (\text{send } (v, z)) \text{ in} \\ \text{let } (u', g') = (\text{recv } y) \text{ in } () \end{array} \right))) \quad (1) \\
&\xrightarrow{c}_C (\mathbf{v}x[\varepsilon]y)(\mathbf{v}w[\varepsilon]z)(\blacklozenge \left(\begin{array}{l} \text{let } k' = (\text{send } (v, z)) \text{ in} \\ \text{let } (u', g') = (\text{recv } y) \text{ in } () \end{array} \right) \parallel \blacklozenge \left(\begin{array}{l} \text{let } f' = (\text{send } (u, x)) \text{ in} \\ \text{let } (v', h') = (\text{recv } w) \text{ in } () \end{array} \right))) \quad (2) \\
&\xrightarrow{c}_C^2 (\mathbf{v}x[\varepsilon]y)(\mathbf{v}w[\varepsilon]z)(\blacklozenge \left(\begin{array}{l} (\text{let } (u', g') = (\text{recv } y) \text{ in } ()) \\ \{\!\! \{ \text{send } (v, z) / k' \} \!\!\} \end{array} \right) \parallel \blacklozenge \left(\begin{array}{l} (\text{let } (v', h') = (\text{recv } w) \text{ in } ()) \\ \{\!\! \{ \text{send } (u, x) / f' \} \!\!\} \end{array} \right))) \quad (3) \\
&\xrightarrow{c}_C^2 (\mathbf{v}x[\varepsilon]y)(\mathbf{v}w[\varepsilon]z)(\blacklozenge \left(\begin{array}{l} (\text{let } (u', g') = (\text{recv } y) \text{ in } ()) \\ \{\!\! \{ \text{send}'(v, z) / k' \} \!\!\} \end{array} \right) \parallel \blacklozenge \left(\begin{array}{l} (\text{let } (v', h') = (\text{recv } w) \text{ in } ()) \\ \{\!\! \{ \text{send}'(u, x) / f' \} \!\!\} \end{array} \right))) \quad (4) \\
&\equiv (\mathbf{v}x[u]y)(\mathbf{v}z[v]w)(\blacklozenge((\text{let } (u', g') = (\text{recv } y) \text{ in } ())\{\!\! \{ z / k' \} \!\!\}) \parallel \blacklozenge((\text{let } (v', h') = (\text{recv } w) \text{ in } ())\{\!\! \{ x / f' \} \!\!\})) \quad (5) \\
&\xrightarrow{c}_C^2 (\mathbf{v}x[u]y)(\mathbf{v}z[v]w)(\blacklozenge(\text{let } (u', g') = (\text{recv } y) \text{ in } ()) \parallel \blacklozenge(\text{let } (v', h') = (\text{recv } w) \text{ in } ())) \\
&\xrightarrow{c}_C^2 (\mathbf{v}x[\varepsilon]y)(\mathbf{v}z[\varepsilon]w)(\blacklozenge(\text{let } (u', g') = (v, y) \text{ in } ()) \parallel \blacklozenge(\text{let } (v', h') = (v, w) \text{ in } ())) \xrightarrow{c}_C^4 \blacklozenge() \quad (6)
\end{aligned}$$

Intuitively, reduction (1) instantiates two buffers and assigns the endpoints through explicit substitutions. Reduction (2) spawns the left term as a child thread. Reduction (3) turns lets into explicit substitutions. Reduction (4) turns the sends into send's. Structural congruence (5) equates the send's with messages in the buffers. Reduction (6) retrieves the messages from the buffers. Note that many of these steps represent several reductions that may happen in any order.

The following example illustrates CGV's flexibility for communicating functions over channels:

Example 2.3. In the following configuration, a buffer and two threads have already been set up (cf. Example 2.2 for an illustration of such an initialization). The main thread sends an interesting term to the child thread: it contains the `send` primitive from which the main thread will subsequently receive from the child thread. We give the configuration's major reductions, with the reducing parts underlined:

$$\begin{aligned}
& (\mathbf{v}x[\varepsilon]y) \left(\diamond \left(\frac{\text{let } x' = \text{send } (\lambda z. \text{send } ((, z), x) \text{ in } ())}{\text{let } (v, x'') = \text{recv } x' \text{ in } v} \right) \parallel \diamond \left(\frac{\text{let } (w, y') = \text{recv } y \text{ in } ()}{\text{let } y'' = (w \ y') \text{ in } ()} \right) \right) \\
& \rightarrow_c^3 (\mathbf{v}x[\lambda z. \text{send } ((, z)]y) \left(\diamond \left(\text{let } (v, x'') = \text{recv } x \text{ in } v \right) \parallel \diamond \left(\frac{\text{let } (w, y') = \text{recv } y \text{ in } ()}{\text{let } y'' = (w \ y') \text{ in } ()} \right) \right) \\
& \rightarrow_c (\mathbf{v}y[\varepsilon]x) \left(\diamond \left(\text{let } (v, x'') = \text{recv } x \text{ in } v \right) \parallel \diamond \left(\frac{\text{let } (w, y') = (\lambda z. \text{send } ((, z), y) \text{ in } ())}{\text{let } y'' = (w \ y') \text{ in } ()} \right) \right) \\
& \rightarrow_c (\mathbf{v}y[\varepsilon]x) \left(\diamond \left(\text{let } (v, x'') = \text{recv } x \text{ in } v \right) \parallel \diamond \left(\frac{\text{let } y'' = (w \ y') \text{ in } ()}{\llbracket (\lambda z. \text{send } ((, z)) / w, y / y' \rrbracket} \right) \right) \\
& \rightarrow_c^2 (\mathbf{v}y[\varepsilon]x) \left(\diamond \left(\text{let } (v, x'') = \text{recv } x \text{ in } v \right) \parallel \diamond \left(\text{let } y'' = \left(\frac{\lambda z. \text{send } ((, z))}{y} \right) \text{ in } () \right) \right) \\
& \rightarrow_c^2 (\mathbf{v}y[\varepsilon]x) \left(\diamond \left(\text{let } (v, x'') = \text{recv } x \text{ in } v \right) \parallel \diamond \left(\text{let } y'' = \text{send } ((, y) \text{ in } () \right) \right) \\
& \rightarrow_c^3 (\mathbf{v}y[()]x) \left(\diamond \left(\text{let } (v, x'') = \text{recv } x \text{ in } v \right) \right) \rightarrow_c^4 \diamond ()
\end{aligned}$$

The following example illustrates why the restricted thread context $\hat{\mathcal{F}}$ is used:

Example 2.4. Consider the configuration $C = (\mathbf{v}x[\varepsilon]y) \left(\diamond \left((\text{send}'(z, x)) \llbracket v/z \rrbracket \right) \parallel D \right)$. Suppose Structural congruence rule SC-SEND' were defined on unrestricted thread contexts; then the rule applies under the explicit substitution of z : $C \equiv_c (\mathbf{v}x[z]y) \left(\diamond \left(x \llbracket v/z \rrbracket \right) \parallel D \right)$. Here, C and the right-hand-side are inconsistent with each other: in C , the variable z is bound by the explicit substitution, whereas z is free on the right-hand-side. With the restricted thread contexts we are forced to first extrude the scope of the explicit substitution before applying Rule SC-SEND', making sure that z remains bound:

$$C \equiv_c \left((\mathbf{v}x[\varepsilon]y) \left(\diamond \left(\text{send}'(z, x) \parallel D \right) \right) \llbracket v/z \rrbracket \equiv_c ((\mathbf{v}x[z]y) (\diamond x \parallel D)) \llbracket v/z \rrbracket.$$

2.2 Type System

We define a type system for CGV, with functional types for functions and pairs and session types for communication. The syntax and meaning of functional types (T, U) and session types (S) are as follows:

$$\begin{aligned}
T, U &::= T \times U && \text{(pair)} \mid T \multimap U && \text{(function)} \mid \mathbf{1} && \text{(unit)} \mid S && \text{(session)} \\
S &::= !T.S && \text{(output)} \mid ?T.S && \text{(input)} \mid \oplus\{i : T\}_{i \in I} && \text{(select)} \mid \&\{i : T\}_{i \in I} && \text{(case)} \mid \text{end}
\end{aligned}$$

Session type duality (\bar{S}) is defined as usual; note that only the continuations, and not the messages, of output and input types are dualized.

$$\overline{!T.S} = ?T.\bar{S} \quad \overline{?T.S} = !T.\bar{S} \quad \overline{\oplus\{i : S_i\}_{i \in I}} = \&\{i : \bar{S}_i\}_{i \in I} \quad \overline{\&\{i : S_i\}_{i \in I}} = \oplus\{i : \bar{S}_i\}_{i \in I} \quad \overline{\text{end}} = \text{end}$$

Typing judgments use typing environments $(\Gamma, \Delta, \Lambda)$ consisting of types assigned to variables ($x : T$). We write \emptyset to denote the empty environment; in writing Γ, Δ , we assume that the variables in Γ and Δ

$\frac{}{\Gamma \vdash_M x : T} \text{T-VAR}$	$\frac{\Gamma, x : T \vdash_M M : U}{\Gamma \vdash_M \lambda x. M : T \multimap U} \text{T-ABS}$	$\frac{\Gamma \vdash_M M : T \multimap U \quad \Delta \vdash_M N : T}{\Gamma, \Delta \vdash_M M N : U} \text{T-APP}$	$\frac{}{\emptyset \vdash_M () : \mathbf{1}} \text{T-UNIT}$	
$\frac{\Gamma \vdash_M M : T \quad \Delta \vdash_M N : U}{\Gamma, \Delta \vdash_M (M, N) : T \times U} \text{T-PAIR}$	$\frac{\Gamma \vdash_M M : T \times T' \quad \Delta, x : T, y : T' \vdash_M N : U}{\Gamma, \Delta \vdash_M \text{let}(x, y) = M \text{ in } N : U} \text{T-SPLIT}$		$\frac{}{\emptyset \vdash_M \text{new} : S \times \bar{S}} \text{T-NEW}$	
$\frac{\Gamma \vdash_M M : \mathbf{1} \times T}{\Gamma \vdash_M \text{spawn } M : T} \text{T-SPAWN}$	$\frac{\Gamma \vdash_M M : T}{\Gamma, x : \text{end} \vdash_M M : T} \text{T-ENDL}$	$\frac{}{\emptyset \vdash_M x : \text{end}} \text{T-ENDR}$	$\frac{\Gamma \vdash_M M : T \times !T.S}{\Gamma \vdash_M \text{send } M : S} \text{T-SEND}$	$\frac{\Gamma \vdash_M M : ?T.S}{\Gamma \vdash_M \text{recv } M : T \times S} \text{T-RECV}$
$\frac{\Gamma \vdash_M M : \oplus\{i : T_i\}_{i \in I} \quad j \in I}{\Gamma \vdash_M \text{select } j M : T_j} \text{T-SELECT}$		$\frac{\Gamma \vdash_M M : \&\{i : T_i\}_{i \in I} \quad \forall i \in I. \Delta \vdash_M N_i : T_i \multimap U}{\Gamma, \Delta \vdash_M \text{case } M \text{ of } \{i : N_i\}_{i \in I} : U} \text{T-CASE}$		
$\frac{\Gamma, x : T \vdash_M M : U \quad \Delta \vdash_M N : T}{\Gamma, \Delta \vdash_M M\{N/x\} : U} \text{T-SUB}$		$\frac{\Gamma \vdash_M M : T \quad \Delta \vdash_M N : !T.S}{\Gamma, \Delta \vdash_M \text{send}'(M, N) : S} \text{T-SEND'}$		
.....				
$\frac{}{\emptyset \vdash_B [\varepsilon] : S' > S'} \text{T-BUF}$	$\frac{\Gamma \vdash_M M : T \quad \Delta \vdash_B [\bar{m}] : S' > S}{\Gamma, \Delta \vdash_B [\bar{m}, M] : S' > !T.S} \text{T-BUFSIZE}$	$\frac{\Gamma \vdash_B [\bar{m}] : S' > S_j \quad j \in I}{\Gamma \vdash_B [\bar{m}, j] : S' > \oplus\{i : S_i\}_{i \in I}} \text{T-BUFSELECT}$		
.....				
$\frac{\Gamma \vdash_M M : T}{\Gamma \vdash_C^\diamond M : T} \text{T-MAIN}$	$\frac{\Gamma \vdash_M M : \mathbf{1}}{\Gamma \vdash_C^\diamond M : \mathbf{1}} \text{T-CHILD}$	$\frac{\Gamma \vdash_C^\diamond C : \mathbf{1} \quad \Delta \vdash_C^\diamond D : T}{\Gamma, \Delta \vdash_C^{\diamond+\phi} C \parallel D : T} \text{T-PARL}$	$\frac{\Gamma \vdash_C^\diamond C : T \quad \Delta \vdash_C^\diamond D : \mathbf{1}}{\Gamma, \Delta \vdash_C^{\phi+\diamond} C \parallel D : T} \text{T-PARR}$	
$\frac{\Gamma \vdash_B [\bar{m}] : S' > S \quad \Delta, x : S', y : \bar{S} \vdash_C^\diamond C : T}{\Gamma, \Delta \vdash_C^\diamond (\mathbf{v}x[\bar{m}]y)C : T} \text{T-RES}$		$\frac{\Gamma, y : \bar{S} \vdash_B [\bar{m}] : S' > S \quad \Delta, x : S' \vdash_C^\diamond C : T}{\Gamma, \Delta \vdash_C^\diamond (\mathbf{v}x[\bar{m}]y)C : T} \text{T-RESBUF}$		
$\frac{\Gamma, x : T \vdash_C^\diamond C : U \quad \Delta \vdash_M M : T}{\Gamma, \Delta \vdash_C^\diamond C\{M/x\} : U} \text{T-CONF SUB}$				

Figure 3: Typing rules for terms (top), buffers (center), and configurations (bottom).

are pairwise distinct. Figure 3 (top) gives the type system for (runtime) terms. Judgments are denoted $\Gamma \vdash_M M : T$ and have a *use-provide* reading: term M uses the variables in Γ to provide a behavior of type T (cf. Caires and Pfenning [1]). When a term provides type T , we often say that the term is of type T .

Typing rules T-VAR, T-ABS, T-APP, T-UNIT, T-PAIR, and T-SPLIT are standard. Rule T-NEW types a pair of dual session types $S \times \bar{S}$. Rule T-SPAWN types spawning a $\mathbf{1}$ -typed term as a child thread, continuing as a term of type T . Rules T-ENDL and T-ENDR type finished sessions. Rule T-SEND (resp. T-RECV) uses a term of type $!T.S$ (resp. $?T.S$) to type a send (resp. receive) of a term of type T , continuing as type S . Rule T-SELECT uses a term of type $\oplus\{i : T_i\}_{i \in I}$ to type selecting a label

$j \in I$, continuing as type T_j . Rule T-CASE uses a term of type $\&\{i: T_i\}_{i \in I}$ to type branching on labels $i \in I$, continuing as type U —each branch is typed $T_i \multimap U$. Rule T-SUB types an explicit substitution. Rule T-SEND' types sending directly, not requiring a pair but two separate terms.

Figure 3 (bottom) gives the typing rules for configurations. The typing judgments here are annotated with a thread marker: $\Gamma \vdash_{\mathbb{C}}^{\phi} C : T$. The thread marker serves to keep track of whether the typed configuration contains the main thread or not (i.e. $\phi = \blacklozenge$ if so, and $\phi = \diamond$ otherwise). When typing the parallel composition of two configurations, we thus have to combine the thread markers of their judgments. This combination of thread markers ($\phi + \phi'$) is defined as follows:

$$\blacklozenge + \diamond = \blacklozenge \quad \diamond + \blacklozenge = \blacklozenge \quad \diamond + \diamond = \diamond \quad (\blacklozenge + \blacklozenge \text{ is undefined})$$

Typing rules T-MAIN and T-CHILD turn a typed term into a thread, where child threads may only be of type **1**. Rules T-PARL and T-PARR compose configurations: one configuration must be of type **1** and have thread marker \diamond (i.e., it does not contain a main thread), providing the other configuration's type. Rule T-RES types buffered restriction, with output endpoint x and input endpoint y used in the configuration. It is possible to send the endpoint y on x , so there is also a Rule T-RESBUF where y is used in the buffer. Unlike with usual typing rules for restriction, the types S' of x and \bar{S} of y do not necessarily have to be duals. This is because the restriction's buffer may already contain messages sent on x but not yet received on y , such that the restricted configuration only needs to use x according to a continuation of S . To ensure that S' is indeed a continuation of S in accordance with the messages in the buffer, we have additional typing rules for buffers, which we explain hereafter. Finally, Rule T-CONFSUB types an explicit substitution on the level of configurations.

For typing buffers, in Figure 3 (center), we have judgments of the form: $\Gamma \vdash_{\mathbb{B}} [\bar{m}] : S' > S$. The judgment denotes that S' is a continuation of S , in accordance with the messages \bar{m} , which use the variables in Γ . The idea of the typing rules is that, starting with an empty buffer at the top of the typing derivation (Rule T-BUF) where $S' = S$, Rules T-BUFSEND and T-BUFSELECT add messages to the end of the buffer. Rule T-BUFSEND then prefixes S with an output of the sent term's type, and Rule T-BUFSELECT prefixes S with a selection such that the sent label's continuation is S .

Example 2.5. Figure 4 (top) shows the typing derivation of a configuration reduced from C_1 in Example 2.2 (following an alternative path after Reduction (5)). Figure 4 (bottom) shows the typing of the configuration $(\mathbf{v}x[M, \ell, ()]y)C$, which has some messages in a buffer; notice how the type of x in C is a continuation of the dual of the type of y .

In the configuration $(\mathbf{v}x[\text{let}(z, y) = \text{recv } y \text{ in } y])C$, the endpoint y is inside the buffer connecting it with x ; to type it, we need Rule T-RESBUF (omitting the derivation of the buffer):

$$\frac{y : ?\text{end} . \text{end} \vdash_{\mathbb{B}} [\text{let}(z, y) = \text{recv } y \text{ in } y] : \text{end} > !\text{end} . \text{end} \quad \Gamma, x : \text{end} \vdash_{\mathbb{C}}^{\phi} C : U}{\Gamma \vdash_{\mathbb{C}}^{\phi} (\mathbf{v}x[\text{let}(z, y) = \text{recv } y \text{ in } y])C : U}$$

Note that such buffers will always deadlock: the message in the buffer can never be received.

Type Preservation Well-typed CGV terms and configurations satisfy protocol fidelity and communication safety. These properties follow from type preservation: typing is consistent across structural congruence and reduction. In both cases the proof is by induction on the derivation of the congruence and reduction, respectively; we include full proofs in the extended version of this paper [14].

Theorem 2.6. If $\Gamma \vdash_{\mathbb{C}}^{\phi} C : T$ and $C \equiv_{\mathbb{C}} D$ or $C \longrightarrow_{\mathbb{C}} D$, then $\Gamma \vdash_{\mathbb{C}}^{\phi} D : T$.

Process syntax:					
$P, Q ::= x[y, z]$	(output)	$x(y, z) . P$	(input)		
$ x[z] \triangleleft i$	(selection)	$x(z) \triangleright \{i : P\}_{i \in I}$	(branching)	$(\mathbf{v}xy)P$	(restriction)
$ P Q$	(parallel)	$\mathbf{0}$	(inaction)	$x \leftrightarrow y$	(forwarder)
Structural congruence:					
$P \equiv P'$	(if $P \equiv_{\alpha} P'$)		$P Q \equiv Q P$	$x \leftrightarrow y \equiv y \leftrightarrow x$	
$P (Q R) \equiv (P Q) R$			$P \mathbf{0} \equiv P$	$(\mathbf{v}xy)x \leftrightarrow y \equiv \mathbf{0}$	
$P (\mathbf{v}xy)Q \equiv (\mathbf{v}xy)(P Q)$	(if $x, y \notin \text{fn}(P)$)		$(\mathbf{v}xy)\mathbf{0} \equiv \mathbf{0}$		
$(\mathbf{v}xy)(\mathbf{v}zw)P \equiv (\mathbf{v}zw)(\mathbf{v}xy)P$			$(\mathbf{v}xy)P \equiv (\mathbf{v}yx)P$		
Reduction:					
$\frac{z, y \neq x}{(\mathbf{v}yz)(x \leftrightarrow y P) \longrightarrow P\{x/z\}}$		$\xrightarrow{\text{ID}}$			
$\frac{j \in I}{(\mathbf{v}xy)(x[b] \triangleleft j y(z) \triangleright \{i : P_i\}_{i \in I}) \longrightarrow P_j\{b/z\}}$		$\xrightarrow{\oplus \&}$			
$\frac{P \longrightarrow Q}{(\mathbf{v}xy)P \longrightarrow (\mathbf{v}xy)Q}$		$\xrightarrow{\mathbf{v}}$			
$\frac{z, y \neq x}{(\mathbf{v}xy)(x[a, b] y(v, z) . P) \longrightarrow P\{a/v, b/z\}}$		$\xrightarrow{\otimes \&}$			
$\frac{P \equiv P' \quad P' \longrightarrow Q' \quad Q' \equiv Q}{P \longrightarrow Q}$		$\xrightarrow{\equiv}$			
$\frac{P \longrightarrow Q}{P R \longrightarrow Q R}$		$\xrightarrow{ }$			

Figure 5: Definition of APCP's process language.

Also, we write $P\{x/y\}$ to denote the capture-avoiding substitution of the free occurrences of y in P for x . We write sequences of substitutions $P\{x_1/y_1\} \dots \{x_n/y_n\}$ as $P\{x_1/y_1, \dots, x_n/y_n\}$.

The reduction relation for processes ($P \longrightarrow Q$) formalizes how complementary actions on connected endpoints may synchronize. As usual for π -calculi, reduction relies on *structural congruence* ($P \equiv Q$), which relates processes with minor syntactic differences; it is the smallest congruence on the syntax of processes (Fig. 5 (top)) satisfying the axioms in Figure 5 (center).

We define the reduction relation $P \longrightarrow Q$ by the axioms and closure rules in Figure 5 (bottom). Rule $\xrightarrow{\text{ID}}$ implements the forwarder as a substitution. Rule $\xrightarrow{\otimes \&}$ synchronizes an output and an input on connected endpoints and substitutes the message and continuation endpoints. Rule $\xrightarrow{\oplus \&}$ synchronizes a selection and a branch: the received label determines the continuation process, substituting the continuation endpoint appropriately. Rules $\xrightarrow{\equiv}$, $\xrightarrow{\mathbf{v}}$, and $\xrightarrow{|}$ close reduction under congruence, restriction, and parallel composition, respectively. We write \longrightarrow^* for the reflexive, transitive closure of \longrightarrow .

The Type System APCP types processes by assigning binary session types to channel endpoints. Following Curry-Howard interpretations, we present session types as linear logic propositions (cf. Caires *et al.* [2] and Wadler [26]) extended with *priority* annotations. Intuitively, actions typed with lower priority cannot be blocked by those with higher priority.

We write $\circ, \kappa, \pi, \rho, \dots$ to denote priorities, and ω to denote the ultimate priority that is greater than all other priorities and cannot be increased further. That is, $\forall \circ \in \mathbb{N}. \omega > \circ$ and $\forall \circ \in \mathbb{N}. \omega + \circ = \omega$.

$\frac{}{\mathbf{0} \vdash \emptyset} \text{EMPTY}$	$\frac{P \vdash \Gamma}{P \vdash \Gamma, x : \bullet} \bullet$	$\frac{}{x \leftrightarrow y \vdash x : \bar{A}, y : A} \text{ID}$	$\frac{P \vdash \Gamma \quad Q \vdash \Delta}{P Q \vdash \Gamma, \Delta} \text{MIX}$
$\frac{P \vdash \Gamma, x : A, y : \bar{A}}{(\mathbf{v}xy)P \vdash \Gamma} \text{CYCLE}$	$\frac{}{x[y, z] \vdash x : A \otimes^\circ B, y : \bar{A}, z : \bar{B}} \otimes$	$\frac{P \vdash \Gamma, y : A, z : B \quad \circ < \text{pr}(\Gamma)}{x(y, z) \cdot P \vdash \Gamma, x : A \wp^\circ B} \wp$	
$\frac{j \in I}{x[z] \triangleleft j \vdash x : \oplus^\circ \{i : A_i\}_{i \in I}, z : \bar{A}_j} \oplus$		$\frac{\forall i \in I. P_i \vdash \Gamma, z : A_i \quad \circ < \text{pr}(\Gamma)}{x(z) \triangleright \{i : P_i\}_{i \in I} \vdash \Gamma, x : \&^\circ \{i : A_i\}_{i \in I}} \&$	

Figure 6: The typing rules of APCP.

Definition 3.1. *The following grammar defines the syntax of session types A, B . Let $\circ \in \mathbb{N}$.*

$A, B ::= A \otimes^\circ B$ (output) $| A \wp^\circ B$ (input) $| \oplus^\circ \{i : A_i\}_{i \in I}$ (select) $| \&^\circ \{i : A_i\}_{i \in I}$ (branch) $| \bullet$ (end)

Note that type \bullet does not require a priority.

Duality, the cornerstone of session types and linear logic, ensures that the two endpoints of a channel have matching actions. Furthermore, dual types must have matching priority annotations.

Definition 3.2. *The dual of session type A , denoted \bar{A} , is defined inductively as follows:*

$$\begin{array}{lll} \overline{A \otimes^\circ B} := \bar{A} \wp^\circ \bar{B} & \overline{\oplus^\circ \{i : A_i\}_{i \in I}} := \&^\circ \{i : \bar{A}_i\}_{i \in I} & \bar{\bullet} := \bullet \\ \overline{A \wp^\circ B} := \bar{A} \otimes^\circ \bar{B} & \overline{\&^\circ \{i : A_i\}_{i \in I}} := \oplus^\circ \{i : \bar{A}_i\}_{i \in I} & \end{array}$$

The priority of a type is determined by the priority of the type's outermost connective:

Definition 3.3. *For session type A , $\text{pr}(A)$ denotes its priority:*

$$\text{pr}(A \otimes^\circ B) := \text{pr}(A \wp^\circ B) := \text{pr}(\oplus^\circ \{i : A_i\}_{i \in I}) := \text{pr}(\&^\circ \{i : A_i\}_{i \in I}) := \circ \quad \text{pr}(\bullet) := \omega$$

The priority of \bullet is ω : it denotes a “final” action of protocols without blocking behavior. Although associated with non-blocking behavior, \otimes and \oplus do have a non-constant priority: they are connected to \wp and $\&$, respectively, which denote blocking actions.

The typing rules of APCP ensure that actions with lower priority are not blocked by those with higher priority (cf. Dardha and Gay [3]). To this end, typing rules enforce the following laws:

1. An action with priority \circ must be prefixed only by inputs and branches with priority strictly smaller than \circ —this law does not hold for output and selection, as they are not prefixes;
2. dual actions leading to a synchronization must have equal priorities (cf. Def. 3.2).

Judgments are of the form $P \vdash \Gamma$, where P is a process and Γ is a context that assigns types to endpoints ($x : A$). A judgment $P \vdash \Gamma$ then means that P can be typed in accordance with the type assignments for names recorded in Γ . The context Γ obeys *exchange*: assignments may be silently reordered. Γ is *linear*, disallowing *weakening* (i.e., all assignments must be used) and *contraction* (i.e., assignments may not be duplicated). The empty context is written \emptyset . In writing $\Gamma, x : A$ we assume that $x \notin \text{dom}(\Gamma)$. We write $\text{pr}(\Gamma)$ to denote the least priority of all types in Γ (cf. Def. 3.3).

Figure 6 gives the typing rules. Rule EMPTY types an inactive process with no endpoints. Rule \bullet silently removes a closed endpoint from the typing context. Rule ID types forwarding between endpoints

$\begin{aligned} \llbracket T \times U \rrbracket &= (\llbracket T \rrbracket \wp \bullet) \otimes (\llbracket U \rrbracket \wp \bullet) \\ \llbracket T \multimap U \rrbracket &= (\overline{\llbracket T \rrbracket} \otimes \bullet) \wp \llbracket U \rrbracket \\ \llbracket \mathbf{1} \rrbracket &= \bullet \end{aligned}$	$\begin{aligned} \llbracket !T.S \rrbracket &= (\overline{\llbracket T \rrbracket} \otimes \bullet) \wp \llbracket S \rrbracket & \llbracket \oplus\{i : T_i\}_{i \in I} \rrbracket &= \&\{i : \llbracket T_i \rrbracket\}_{i \in I} \\ \llbracket ?T.S \rrbracket &= (\llbracket T \rrbracket \wp \bullet) \otimes \llbracket S \rrbracket & \llbracket \&\{i : T_i\}_{i \in I} \rrbracket &= \oplus\{i : \llbracket T_i \rrbracket\}_{i \in I} \\ \llbracket \text{end} \rrbracket &= \bullet \end{aligned}$
---	--

Figure 7: Translation of CGV types into session types.

of dual type. Rule MIX types the parallel composition of two processes that do not share assignments on the same endpoints. Rule CYCLE types a restriction, where the two restricted endpoints must be of dual type. Rule \otimes types an output action; this rule does not have premises to provide a continuation process, leaving the free endpoints to be bound to a continuation process using MIX and CYCLE. Similarly, Rule \oplus types an unbound selection action. Priority checks are confined to Rules \wp and $\&$, which type input and branching prefixes, respectively. In both cases, the used endpoint's priority must be lower than the priorities of the other types in the continuation's typing context, thus enforcing Law 1 above.

Well-typed processes satisfy protocol fidelity, communication safety, and deadlock-freedom. The first two properties follow from *type preservation*. Here we only state these results; see [13] for details.

Theorem 3.4 (Type Preservation). *If $P \vdash \Gamma$ and $P \equiv Q$ or $P \longrightarrow Q$, then $Q \vdash \Gamma$.*

Theorem 3.5 (Deadlock-freedom). *If $P \vdash \emptyset$, then either $P \equiv \mathbf{0}$ or $P \longrightarrow Q$ for some Q .*

4 Translating CGV into APCP

4.1 The Translation

In this section, we translate CGV into APCP. We translate entire typing derivations, following, e.g., Wadler [26]. Given the structure of CGV and its type system, the translation is defined in parts: for (run-time) terms, for configurations, and for buffers. The translation is defined on well-typed configurations which may be deadlocked, so our translation does not consider priority requirements. As we will see, typability in APCP will enable us to identify deadlock-free configurations in CGV (cf. Sec. 4.3).

The translation is informed by the semantics of CGV. It is crucial that subterms may only reduce when they occur in reduction contexts. For example, M_1 and M_2 may not reduce if they appear in a pair (M_1, M_2) . The translation must thus ensure that subterms are blocked when they do not occur in reduction contexts. Translations such as Wadler's hinge on blocking outputs and inputs; for example, the pair (M_1, M_2) is translated as an output that blocks the translations of M_1 and M_2 . However, outputs in APCP are non-blocking and so we use additional inputs to disable the reduction of subterms. For example, the translation of (M_1, M_2) adds extra inputs to block the translations of M_1 and M_2 .

Figure 7 gives the translation of CGV types into APCP types ($\llbracket T \rrbracket$), which already captures the operation of the translation: our translation is similar to the one by Wadler, but includes the aforementioned additional inputs. It may seem odd that this translation dualizes CGV session types (e.g., an output ' $!$ ' becomes an input ' \wp '). To understand this, consider that a variable x typed $!T.S$ represents access to a session which expects the user to send a term of type T and continue as S , but not the output itself. Hence, to translate an output on x into APCP, we need to connect the translation of x to an actual output. Since this actual output would be typed with \otimes , this means that the translation of x would need to be dually typed, i.e., typed with \wp . A more technical explanation is that the translation moves from two-sided CGV judgments to one-sided APCP judgments, which requires dualization (see, e.g., [9, 12]).

Importantly, the translation preserves duality of session types (by induction on their structure):

Proposition 4.1. *Given a CGV session type S , $\overline{\llbracket S \rrbracket} = \llbracket \overline{S} \rrbracket$.*

We extend the translation of types to typing environments, defined as expected. Similarly, we extend duality to typing environments: $\overline{\Gamma}$ denotes Γ with each type dualized. In this section, we give simplified presentations of the translations, showing only the conclusions of the source and target derivations; we include the translations with full derivations in the extended version of this paper [14].

A remark on notation. Some translated terms include annotated restrictions $(\overleftrightarrow{\mathbf{v}}xy)$. These so-called *forwarder-enabled* restrictions can be ignored in this subsection, but will be useful later when proving soundness (one of the correctness properties of the translation; cf. Section 4.2).

We define the translation of (the typing rules of) terms. Since a term has a provided type, the translation takes as a parameter a name on which the translation provides this type. Figure 8 gives the translation of terms, denoted $\llbracket \Gamma \vdash_{\mathbb{M}} M : T \rrbracket z$, where the type T is provided on z . By abuse of notation, we write $\llbracket M \rrbracket z$ to denote the process translation of the term M , and similarly for configurations and buffers. Notice the aforementioned additional inputs to block behavior of subterms in rules such as Rule T-PAIR. Before moving to buffers and configurations, we illustrate the translation of terms by an example:

Example 4.2. *Consider the following subterm from Example 2.3: $(\lambda z. \text{send}((\cdot), z)) y$. We gradually discuss how this term translates to APCP, and how the translation is set up to mimick the term's behavior.*

$$\llbracket (\lambda z. \text{send}((\cdot), z)) y \rrbracket q = (\mathbf{v}ab)(\llbracket \lambda z. \text{send}((\cdot), z) \rrbracket a \mid (\mathbf{v}cd)(b[c, q] \mid d(e, \cdot) \cdot \llbracket y \rrbracket e))$$

The function application translates the function on a , which is connected to b . The output on b serves to activate the function, which will subsequently activate the functions parameter ($\llbracket y \rrbracket e = y \leftrightarrow e$) by means of an output that will be received on d .

$$\llbracket \lambda z. \text{send}((\cdot), z) \rrbracket a = a(f, g) \cdot (\overleftrightarrow{\mathbf{v}}hz)((\mathbf{v}_)_)f[h, _] \mid \llbracket \text{send}((\cdot), z) \rrbracket g$$

The translation of the function is indeed blocked until it receives on a . It then outputs on f to activate the function's parameter (which receives on d), while the function's body appears in parallel.

$$\llbracket \text{send}((\cdot), z) \rrbracket g = (\mathbf{v}kl) \left(\llbracket ((\cdot), z) \rrbracket k \mid l(m, n) \cdot (\mathbf{v}op)((\mathbf{v}_)_)n[o, _] \mid (\mathbf{v}rs)(p[m, r] \mid s \leftrightarrow g) \right)$$

The translation of the `send` primitive connects the translation of the pair $((\cdot), z)$ on k to an input on l , receiving endpoints for the output term (m) and the output endpoint (n). Once activated by the input on l , the term representing the output endpoint is activated by means of an output on n . In parallel, the actual output (on p) sends the endpoint of the output term (m) and a fresh endpoint (r) representing the continuation channel after the message has been placed in a buffer (the forwarder $s \leftrightarrow g$).

$$\llbracket ((\cdot), z) \rrbracket k = (\mathbf{v}tu)(\mathbf{v}vw)(k[t, v] \mid u(a', _) \cdot \llbracket (\cdot) \rrbracket a' \mid w(b', _) \cdot \llbracket z \rrbracket b')$$

The translation of the pair outputs on k two endpoints for the two terms it contains (to be received by whatever intends to use the pair in the context, e.g., the `send` primitive on l). The translations of the two terms inside the pair ($\llbracket (\cdot) \rrbracket a' = \mathbf{0}$ and $\llbracket z \rrbracket b' = z \leftrightarrow b'$) are both guarded by an input, preventing the terms from reducing until the context explicitly activates them by means of outputs.

Analogously to the reductions from Example 2.3— $(\lambda z. \text{send}((\cdot), z)) y \rightarrow_{\mathbb{M}}^3 \text{send}'((\cdot), y)$ —we have

$$\llbracket (\lambda z. \text{send}((\cdot), z)) y \rrbracket q \rightarrow^5 \llbracket \text{send}'((\cdot), y) \rrbracket q.$$

T-VAR	$\llbracket x : T \vdash_M x : T \rrbracket z = x \leftrightarrow z \vdash x : \overline{\llbracket T \rrbracket}, z : \llbracket T \rrbracket$	T-UNIT	$\llbracket \emptyset \vdash_M () : \mathbf{1} \rrbracket z = \mathbf{0} \vdash z : \bullet$
T-ABS	$\llbracket \Gamma \vdash_M \lambda x. M : T \multimap U \rrbracket z = z(a, b) \cdot (\overleftrightarrow{\mathbf{v}cx})((\mathbf{v}ef)a[c, e] \mid \llbracket M \rrbracket b) \vdash \overline{\llbracket \Gamma \rrbracket}, z : (\overline{\llbracket T \rrbracket} \otimes \bullet) \wp \llbracket U \rrbracket$		
T-APP	$\llbracket \Gamma, \Delta \vdash_M MN : U \rrbracket z = (\mathbf{v}ab)(\llbracket M \rrbracket a \mid (\mathbf{v}cd)(b[c, z] \mid d(e, f) \cdot \llbracket N \rrbracket e)) \vdash \overline{\llbracket \Gamma \rrbracket}, \overline{\llbracket \Delta \rrbracket}, z : \llbracket U \rrbracket$		
T-PAIR	$\llbracket \Gamma, \Delta \vdash_M (M, N) : T \times U \rrbracket z = \begin{array}{l} (\mathbf{v}ab)(\mathbf{v}cd)(z[a, c] \mid b(e, f) \cdot \llbracket M \rrbracket e \mid d(g, h) \cdot \llbracket N \rrbracket g) \\ \vdash \overline{\llbracket \Gamma \rrbracket}, \overline{\llbracket \Delta \rrbracket}, z : (\llbracket T \rrbracket \wp \bullet) \otimes (\llbracket U \rrbracket \wp \bullet) \end{array}$		
T-SPLIT	$\llbracket \Gamma, \Delta \vdash_M \text{let } (x, y) = M \text{ in } N : U \rrbracket z = (\mathbf{v}ab)(\llbracket M \rrbracket a \mid b(c, d) \cdot (\overleftrightarrow{\mathbf{v}ex})(\overleftrightarrow{\mathbf{v}fy})(\mathbf{v}gh)c[e, g] \mid (\mathbf{v}kl)d[f, k] \mid \llbracket N \rrbracket z)) \vdash \overline{\llbracket \Gamma \rrbracket}, \overline{\llbracket \Delta \rrbracket}, z : \llbracket U \rrbracket$		
T-NEW	$\llbracket \emptyset \vdash_M \text{new} : S \times \overline{S} \rrbracket z = (\mathbf{v}ab)((\mathbf{v}cd)a[c, d] \mid b(e, f) \cdot (\mathbf{v}xy)\llbracket (x, y) \rrbracket z) \vdash z : (\llbracket S \rrbracket \wp \bullet) \otimes (\llbracket \overline{S} \rrbracket \wp \bullet)$		
T-SPAWN	$\llbracket \Gamma \vdash_M \text{spawn } M : T \rrbracket z = (\mathbf{v}ab)(\llbracket M \rrbracket a \mid b(c, d) \cdot ((\mathbf{v}ef)c[e, f] \mid (\mathbf{v}gh)d[z, g])) \vdash \overline{\llbracket \Gamma \rrbracket}, z : \llbracket T \rrbracket$		
T-ENDL	$\llbracket \Gamma, x : \text{end} \vdash_M M : T \rrbracket z = \llbracket M \rrbracket z \vdash \overline{\llbracket \Gamma \rrbracket}, x : \bullet, z : \llbracket T \rrbracket$	T-ENDR	$\llbracket \emptyset \vdash_M x : \text{end} \rrbracket z = \mathbf{0} \vdash z : \bullet$
T-SEND	$\llbracket \Gamma \vdash_M \text{send } M : S \rrbracket z = (\mathbf{v}ab)(\llbracket M \rrbracket a \mid b(c, d) \cdot (\mathbf{v}ef)((\mathbf{v}gh)d[e, g] \mid (\mathbf{v}kl)(f[c, k] \mid l \leftrightarrow z))) \vdash \overline{\llbracket \Gamma \rrbracket}, z : \llbracket S \rrbracket$		
T-RECV	$\llbracket \Gamma \vdash_M \text{recv } M : T \times S \rrbracket z = (\mathbf{v}ab)(\llbracket M \rrbracket a \mid b(c, d) \cdot (\mathbf{v}ef)(z[c, e] \mid f(g, h) \cdot d \leftrightarrow z)) \vdash \overline{\llbracket \Gamma \rrbracket}, z : (\llbracket T \rrbracket \wp \bullet) \otimes (\llbracket S \rrbracket \wp \bullet)$		
T-SELECT	$\llbracket \Gamma \vdash_M \text{select } jM : T_j \rrbracket z = (\mathbf{v}ab)(\llbracket M \rrbracket a \mid (\mathbf{v}cd)(b[c] \triangleleft j \mid d \leftrightarrow z)) \vdash \overline{\llbracket \Gamma \rrbracket}, z : \llbracket T_j \rrbracket$		
T-CASE	$\llbracket \Gamma, \Delta \vdash_M \text{case } M \text{ of } \{i : N_i\}_{i \in I} : U \rrbracket z = (\mathbf{v}ab)(\llbracket M \rrbracket a \mid b(c) \triangleright \{i : \llbracket N_i c \rrbracket z\}_{i \in I}) \vdash \overline{\llbracket \Gamma \rrbracket}, \overline{\llbracket \Delta \rrbracket}, z : \llbracket U \rrbracket$		
T-SUB	$\llbracket \Gamma, \Delta \vdash_M M \llbracket N/x \rrbracket : U \rrbracket z = (\overleftrightarrow{\mathbf{v}xa})(\llbracket M \rrbracket z \mid \llbracket N \rrbracket a) \vdash \overline{\llbracket \Gamma \rrbracket}, \overline{\llbracket \Delta \rrbracket}, z : \llbracket U \rrbracket$		
T-SEND'	$\llbracket \Gamma, \Delta \vdash_M \text{send}' (M, N) : S \rrbracket z = (\mathbf{v}ab)(a(c, d) \cdot \llbracket M \rrbracket c \mid (\mathbf{v}ef)(\llbracket N \rrbracket e \mid (\mathbf{v}gh)(f[b, g] \mid h \leftrightarrow z))) \vdash \overline{\llbracket \Gamma \rrbracket}, \overline{\llbracket \Delta \rrbracket}, z : \llbracket S \rrbracket$		

Figure 8: Translation of (runtime) term typing rules. See [14] for typing derivations.

Figure 9 (top) gives the translation of configurations, denoted $\llbracket \Gamma \vdash_C^\phi C : T \rrbracket z$. We omit the translation of Rule T-PARR. Noteworthy are the translations of buffered restrictions: the translation of $(\mathbf{v}x[\vec{m}]y)C$ relies on the translation of $\llbracket \vec{m} \rrbracket$, which is given the translation of C as its continuation.

The translation of buffers requires care: each message in the buffer is translated as an output in ACP, where the output of the following messages is on the former output's continuation endpoint. Once there are no more messages in the buffer, the translation uses a typed ACP process—a parameter of the translation—to provide the behavior of the continuation of the lastmost output. The translation has no requirements for the continuation process and its typing, except for the type of the buffer's endpoint. With this in mind, Figure 9 (bottom) gives the translation of the typing rules of buffers, denoted $\llbracket \Gamma \vdash_B \llbracket \vec{m} \rrbracket : S' > S \rrbracket_x^{P \vdash^* \Lambda, x : \overline{S'}}$, where x is the endpoint on which the buffer outputs, and P is the continuation of the buffer's last message. Note that we never use the typing rules for buffers by themselves: they always accompany the typing of endpoint restriction, of which the translation properly instantiates the continuation process.

Because CGV configurations may deadlock, the type preservation result of our translation holds up

T-MAIN/CHILD	$\llbracket \Gamma \vdash_{\mathcal{C}}^{\phi} \phi M : T \rrbracket z = \llbracket M \rrbracket z \vdash \overline{\llbracket \Gamma \rrbracket}, z : \llbracket T \rrbracket$
T-PARL	$\llbracket \Gamma, \Delta \vdash_{\mathcal{C}}^{\phi+\phi} C \parallel D : T \rrbracket z = (\mathbf{v}ab) \llbracket C \rrbracket a \mid \llbracket D \rrbracket z \vdash \overline{\llbracket \Gamma \rrbracket}, \overline{\llbracket \Delta \rrbracket}, z : \llbracket T \rrbracket$
T-RES/T-RESBUF	$\llbracket \Gamma, \Delta \vdash_{\mathcal{C}}^{\phi} (\mathbf{v}x[\vec{m}]y)C : T \rrbracket z = (\mathbf{v}xy) \llbracket [\vec{m}] \rrbracket_x^{\llbracket C \rrbracket z} \vdash \overline{\llbracket \Gamma \rrbracket}, \overline{\llbracket \Delta \rrbracket}, z : \llbracket T \rrbracket$
T-CONFSUB	$\llbracket \Gamma, \Delta \vdash_{\mathcal{C}}^{\phi} C \{M/x\} : U \rrbracket z = (\overleftrightarrow{\mathbf{v}}xa) (\llbracket C \rrbracket z \mid \llbracket M \rrbracket z) \vdash \overline{\llbracket \Gamma \rrbracket}, \overline{\llbracket \Delta \rrbracket}, z : \llbracket U \rrbracket$
T-BUF	$\llbracket \emptyset \vdash_{\mathcal{B}} [\varepsilon] : S' > S' \rrbracket_x^{P \vdash \Lambda, x : \overline{\llbracket S' \rrbracket}} = P \vdash \Lambda, x : \overline{\llbracket S' \rrbracket}$
T-BUFSEND	$\llbracket \Gamma, \Delta \vdash_{\mathcal{B}} [\vec{m}, M] : S' > !T.S \rrbracket_x^{P \vdash \Lambda, x : \overline{\llbracket S' \rrbracket}} = (\mathbf{v}ab)(\mathbf{v}cd)((\mathbf{v}gh)(x \leftrightarrow g \mid h[a, c] \mid b(e, f) \cdot \llbracket M \rrbracket e \mid \llbracket [\vec{m}] \rrbracket_d^{P\{d/x\}}) \vdash \overline{\llbracket \Gamma \rrbracket}, \overline{\llbracket \Delta \rrbracket}, \Lambda, x : (\llbracket T \rrbracket \wp \bullet) \otimes \overline{\llbracket S \rrbracket})$
T-BUFSELECT	$\llbracket \Gamma \vdash_{\mathcal{B}} [\vec{m}, j] : S' > \oplus \{i : S_i\}_{i \in I} \rrbracket_x^{P \vdash \Lambda, x : \overline{\llbracket S' \rrbracket}} = (\mathbf{v}ab)((\mathbf{v}cd)(x \leftrightarrow c \mid d[a] \triangleleft j \mid \llbracket [\vec{m}] \rrbracket_b^{P\{b/x\}}) \vdash \overline{\llbracket \Gamma \rrbracket}, \Lambda, x : \oplus \{i : \overline{\llbracket S_i \rrbracket}\}_{i \in I})$

Figure 9: Translation of configuration and buffer typing rules. See [14] for typing derivations.

to priority requirements. To formalize this, we have the following definition:

Definition 4.3. Let P be a process. We write $P \vdash^* \Gamma$ to denote that P is well-typed according to the typing rules in Figure 6 where Rules \wp and $\&$ are modified by erasing priority checks.

Hence, if $P \vdash \Gamma$ then $P \vdash^* \Gamma$ but the converse does not hold. Our translation correctly preserves the typing of terms, configurations, and buffers:

Theorem 4.4 (Type Preservation for the Translation).

- $\llbracket \Gamma \vdash_{\mathcal{M}} M : T \rrbracket z = \llbracket M \rrbracket z \vdash^* \overline{\llbracket \Gamma \rrbracket}, z : \llbracket T \rrbracket$
- $\llbracket \Gamma \vdash_{\mathcal{C}}^{\phi} C : T \rrbracket z = \llbracket C \rrbracket z \vdash^* \overline{\llbracket \Gamma \rrbracket}, z : \llbracket T \rrbracket$
- $\llbracket \Gamma \vdash_{\mathcal{B}} [\vec{m}] : S' > S \rrbracket_x^{P \vdash \Lambda, x : \overline{\llbracket S' \rrbracket}} = \llbracket [\vec{m}] \rrbracket_x^P \vdash^* \overline{\llbracket \Gamma \rrbracket}, \Lambda, x : \overline{\llbracket S \rrbracket}$

Example 4.5. Consider again the configuration $(\mathbf{v}x[M, \ell, ()]y)C$. We illustrate the translation of buffers into APCP by giving the translation of this configuration (writing $\langle x \rangle[a, b]$ to denote the forwarded output $(\mathbf{v}cd)(x \leftrightarrow c \mid d[a, b])$):

$$\begin{aligned} \llbracket (\mathbf{v}x[M, \ell, ()]y)C \rrbracket z &= (\mathbf{v}xy) \llbracket [M, \ell, ()] \rrbracket_x^{\llbracket C \rrbracket z} = (\mathbf{v}xy) (\mathbf{v}ab)(\mathbf{v}cx')(\langle x \rangle[a, c] \mid b(d, -) \cdot \mathbf{0} \\ &\mid (\mathbf{v}ex'')(\langle x' \rangle[e] \triangleleft \ell \mid (\mathbf{v}fg)(\mathbf{v}hx''')(\langle x'' \rangle[f, h] \mid g(k, -) \cdot \llbracket M \rrbracket k \mid \llbracket C \rrbracket z \{x'''/x\}))) \end{aligned}$$

Notice how the (forwarded) outputs are sequenced by continuation endpoints, and how the translation of C uses the last continuation endpoint x''' to interact with the buffer.

4.2 Operational Correctness

Following Gorla [10], we focus on *operational correspondence*: a translated configuration can reproduce all of the source configuration's reductions (completeness; Theorem 4.6), and any of the translated configuration's reductions can be traced back to reductions of the source configuration (soundness; Theorem 4.7). With the soundness result, our translation is stronger than related prior translations [20, 24, 19].

Our completeness result states that the reductions of a well-typed configuration can be mimicked by its translation in zero or more steps.

Theorem 4.6 (Completeness). *Given $\Gamma \vdash_{\mathcal{C}}^{\phi} C : T$, if $C \longrightarrow_{\mathcal{C}} D$, then $\llbracket C \rrbracket_z \longrightarrow^* \llbracket D \rrbracket_z$.*

Proof (Sketch). By induction on the derivation of the configuration’s reduction. In each case, we infer the shape of the configuration from the reduction and well-typedness. We then consider the translation of the configuration, and show that the resulting process reduces in zero or more steps to the translation of the reduced configuration. See the extended version of this paper [14] for a full proof. \square

Soundness states that any sequence of reductions from the translation of a well-typed configuration eventually leads to the translation of another configuration, which the initial configuration also reduces to. Asynchrony in APCP requires us to be careful, specifically concerning the semantics of variables in CGV. Variables can only cause reductions under specific circumstances. On the other hand, variables translate to forwarders in APCP, which reduce as soon as they are bound by restriction. This semantics for forwarders turns out to be too eager for soundness. As a result, soundness only holds for an alternative, so-called *lazy semantics* for APCP, denoted \longrightarrow_L , in which forwarders may only cause reductions under specific circumstances. It is here that the forwarder-enabled restrictions ($\overset{\leftrightarrow}{\mathbf{v}}xy$) anticipated in Section 4.1 come into play. As we will see in Section 4.3, this alternative semantics does not prevent us from identifying a class of deadlock-free CGV configurations through the translation into APCP. Due to space limitations, the definitions of the lazy semantics only appears in the extended version of this paper [14].

Theorem 4.7 (Soundness). *Given $\Gamma \vdash_{\mathcal{C}}^{\phi} C : T$, if $\llbracket C \rrbracket_z \longrightarrow_L^* Q$, then $C \longrightarrow_{\mathcal{C}}^* D$ and $Q \longrightarrow_L^* \llbracket D \rrbracket_z$ for some D .*

Proof (Sketch). By induction on the structure of C . In each case, we additionally apply induction on the number k of steps $\llbracket C \rrbracket_z \longrightarrow_L^k Q$. We then consider which reductions might occur from $\llbracket C \rrbracket_z$ to Q . Considering the structure of C , we then isolate a sequence of k' possible steps, such that $\llbracket C \rrbracket_z \longrightarrow_L^{k'} \llbracket D' \rrbracket_z$ for some D' where $C \longrightarrow_{\mathcal{C}} D'$. Since $\llbracket D' \rrbracket_z \longrightarrow_L^{k-k'} Q$, it then follows from the induction hypothesis that there exists D such that $D' \longrightarrow_{\mathcal{C}}^* D$ and $\llbracket D' \rrbracket_z \longrightarrow_L^* \llbracket D \rrbracket_z$.

Key here is the independence of reductions in APCP: if two or more reductions are enabled from a (well-typed) process, they must originate from independent parts of the process, and so they do not interfere with each other. This essentially means that the order in which independent reductions occur does not affect the resulting process. Hence, we can pick “desirable” sequences of reductions, postponing other possible reductions. See the extended version of this paper [14] for a full proof of soundness. \square

From the proof above we can deduce that if the translation takes at least one step, then so does the source:

Corollary 4.8. *Given $\Gamma \vdash_{\mathcal{C}}^{\phi} C : T$, if $\llbracket C \rrbracket_z \longrightarrow_L^+ Q$, then $C \longrightarrow_{\mathcal{C}}^+ D$ and $Q \longrightarrow_L^* \llbracket D \rrbracket_z$ for some D .*

4.3 Transferring Deadlock-freedom from APCP to CGV

In APCP, well-typed processes typable under empty contexts ($P \vdash \emptyset$) are deadlock-free. By appealing to the operational correctness of our translation, we transfer this result to CGV configurations. Each deadlock-free configuration in CGV obtained via transference satisfies two requirements:

- The configuration is typable $\emptyset \vdash_{\mathcal{C}}^{\diamond} C : \mathbf{1}$: it needs no external resources and has no external behavior.
- The typed translation of the configuration satisfies APCP’s priority requirements: it is well-typed under ‘ \vdash ’, not only under ‘ \vdash^* ’ (cf. Def. 4.3).

We rely on soundness (Theorem 4.7) to transfer deadlock-freedom to configurations. However, APCP’s deadlock-freedom (Theorem 3.5) considers standard semantics (\longrightarrow), whereas soundness considers the lazy semantics (\longrightarrow_L). Therefore, we first must show that if the translation of a configuration

	λ^{sess} [8]	GV [26]	EGV [7]	PGV [17, 18]	CGV (this paper)
Communication	Asynch.	Synch.	Asynch.	Synch.	Asynch.
Cyclic Topologies	Yes	No	No	Yes	Yes
Deadlock-Freedom	No	Yes (typing)	Yes (typing)	Yes (typing)	Yes (via APCP)

Table 1: The features of CGV compared to its predecessors.

satisfying the requirements above reduces under \longrightarrow , it also reduces under $\longrightarrow_{\text{L}}$; this is Theorem 4.9 below. The deadlock-freedom of these configurations (Theorem 4.10) then follows from Theorems 3.5 and 4.9. See the extended version of this paper [14] for detailed proofs of these results.

Theorem 4.9. *Given $\emptyset \vdash_{\mathcal{C}}^{\dagger} C : \mathbf{1}$, if $\llbracket C \rrbracket_z \vdash \Gamma$ for some Γ and $\llbracket C \rrbracket_z \longrightarrow Q$, then $\llbracket C \rrbracket_z \longrightarrow_{\text{L}} Q'$, for some Q' .*

Proof (Sketch). By inspecting the derivation of $\llbracket C \rrbracket_z \longrightarrow Q$. If the reduction is not derived from \rightarrow_{ID} , it can be directly replicated under $\longrightarrow_{\text{L}}$. Otherwise, we analyze the possible shapes of C and show that a different reduction under $\longrightarrow_{\text{L}}$ is possible. \square

Theorem 4.10 (Deadlock-freedom for CGV). *Given $\emptyset \vdash_{\mathcal{C}}^{\dagger} C : \mathbf{1}$, if $\llbracket C \rrbracket_z \vdash \Gamma$ for some Γ , then $C \equiv \blacklozenge()$ or $C \longrightarrow_{\mathcal{C}} D$ for some D .*

Proof (Sketch). By assumption and Theorem 4.4, $\llbracket C \rrbracket_z \vdash z : \bullet$. Then $(\mathbf{v}_{z-})\llbracket C \rrbracket_z \vdash \emptyset$. By Theorem 3.5, (i) $(\mathbf{v}_{z-})\llbracket C \rrbracket_z \equiv \mathbf{0}$ or (ii) $(\mathbf{v}_{z-})\llbracket C \rrbracket_z \longrightarrow Q$ for some Q . In case (i) it follows from the well-typedness and translation of C that $C \equiv_{\mathcal{C}} \blacklozenge()$. In case (ii) we deduce that the reduction of $(\mathbf{v}_{z-})\llbracket C \rrbracket_z$ cannot involve the endpoint z . Hence, $\llbracket C \rrbracket_z \longrightarrow Q_0$ for some Q_0 . By Theorem 4.9, then $\llbracket C \rrbracket_z \longrightarrow_{\text{L}} Q'$ for some Q' . Then, by Corollary 4.8, there exists D' such that $C \longrightarrow_{\mathcal{C}}^+ D'$. Hence, $C \longrightarrow_{\mathcal{C}} D$ for some D , proving the thesis. \square

As an example, using Theorem 4.10 we can show that C_1 from Example 2.2 is deadlock-free; see [14].

5 Conclusion

We have presented CGV, a new functional language with asynchronous session-typed communication. As illustrated in Section 1, CGV is strictly more expressive than its predecessors, thanks to a highly asynchronous semantics (compared to GV and PGV), its support for cyclic thread configurations (compared to EGV), and the ability to send whole terms and not just values (compared to all the mentioned calculi). Table 1 summarizes the features of CGV compared to its predecessors.

An operationally correct translation into APCP solidifies the design of CGV, and enables identifying a class of deadlock-free CGV programs. Interestingly, the asynchronous semantics of CGV is reminiscent of *future/promise* programming paradigms (see, e.g., [11, 21, 25]), which have been little studied in the context of session-typed communication.

The alternative to establishing deadlock-freedom in CGV via translation into APCP would be to enhance CGV's type system with priorities (in the spirit of, e.g., work by Padovani and Novara [23]). Another useful addition concerns recursion / recursive types. We leave these extensions to future work.

Acknowledgments Thanks to Simon Fowler and the anonymous reviewers for their helpful feedback. We gratefully acknowledge the support of the Dutch Research Council (NWO) under project No. 016.Vidi.189.046 (Unifying Correctness for Communicating Software).

References

- [1] Luís Caires & Frank Pfenning (2010): *Session Types as Intuitionistic Linear Propositions*. In Paul Gastin & François Laroussinie, editors: *CONCUR 2010 - Concurrency Theory*, Lecture Notes in Computer Science, Springer, Berlin, Heidelberg, pp. 222–236, doi:[10.1007/978-3-642-15375-4_16](https://doi.org/10.1007/978-3-642-15375-4_16).
- [2] Luís Caires, Frank Pfenning & Bernardo Toninho (2016): *Linear Logic Propositions as Session Types*. *Mathematical Structures in Computer Science* 26(3), pp. 367–423, doi:[10.1017/S0960129514000218](https://doi.org/10.1017/S0960129514000218).
- [3] Ornela Dardha & Simon J. Gay (2018): *A New Linear Logic for Deadlock-Free Session-Typed Processes*. In Christel Baier & Ugo Dal Lago, editors: *Foundations of Software Science and Computation Structures*, Lecture Notes in Computer Science, Springer International Publishing, pp. 91–109, doi:[10.1007/978-3-319-89366-2_5](https://doi.org/10.1007/978-3-319-89366-2_5).
- [4] Ornela Dardha & Jorge A. Pérez (2015): *Comparing Deadlock-Free Session Typed Processes*. *Electronic Proceedings in Theoretical Computer Science* 190, pp. 1–15, doi:[10.4204/EPTCS.190.1](https://doi.org/10.4204/EPTCS.190.1). arXiv:[1508.06707](https://arxiv.org/abs/1508.06707).
- [5] Ornela Dardha & Jorge A. Pérez (2022): *Comparing Type Systems for Deadlock Freedom*. *Journal of Logical and Algebraic Methods in Programming* 124, p. 100717, doi:[10.1016/j.jlamp.2021.100717](https://doi.org/10.1016/j.jlamp.2021.100717).
- [6] Simon Fowler (2019): *Typed Concurrent Functional Programming with Channels, Actors, and Sessions*. Ph.D. thesis, University of Edinburgh.
- [7] Simon Fowler, Sam Lindley, J. Garrett Morris & Sára Decova (2019): *Exceptional Asynchronous Session Types: Session Types without Tiers*. *Proceedings of the ACM on Programming Languages*, doi:[10.1145/3290341](https://doi.org/10.1145/3290341).
- [8] Simon J. Gay & Vasco T. Vasconcelos (2010): *Linear Type Theory for Asynchronous Session Types*. *Journal of Functional Programming* 20(1), pp. 19–50, doi:[10.1017/S0956796809990268](https://doi.org/10.1017/S0956796809990268).
- [9] Jean-Yves Girard (1993): *On the Unity of Logic*. *Annals of Pure and Applied Logic* 59(3), pp. 201–217, doi:[10.1016/0168-0072\(93\)90093-S](https://doi.org/10.1016/0168-0072(93)90093-S).
- [10] Daniele Gorla (2010): *Towards a Unified Approach to Encodability and Separation Results for Process Calculi*. *Information and Computation* 208(9), pp. 1031–1053, doi:[10.1016/j.ic.2010.05.002](https://doi.org/10.1016/j.ic.2010.05.002).
- [11] Robert H. Halstead (1985): *MULTILISP: A Language for Concurrent Symbolic Computation*. *ACM Transactions on Programming Languages and Systems* 7(4), pp. 501–538, doi:[10.1145/4472.4478](https://doi.org/10.1145/4472.4478).
- [12] Bas van den Heuvel & Jorge A. Pérez (2020): *Session Type Systems Based on Linear Logic: Classical versus Intuitionistic*. *Electronic Proceedings in Theoretical Computer Science* 314, pp. 1–11, doi:[10.4204/EPTCS.314.1](https://doi.org/10.4204/EPTCS.314.1). arXiv:[2004.01320](https://arxiv.org/abs/2004.01320).
- [13] Bas van den Heuvel & Jorge A. Pérez (2021): *Deadlock Freedom for Asynchronous and Cyclic Process Networks (Extended Version)*. arXiv:[2111.13091](https://arxiv.org/abs/2111.13091) [cs]. arXiv:[2111.13091](https://arxiv.org/abs/2111.13091). A short version appears in the Proceedings of ICE’21: arXiv:[2110.00146](https://arxiv.org/abs/2110.00146).
- [14] Bas van den Heuvel & Jorge A. Pérez (2022): *Asynchronous Functional Sessions: Cyclic and Concurrent (Extended Version)*, doi:[10.48550/arXiv.2208.07644](https://doi.org/10.48550/arXiv.2208.07644). arXiv:[2208.07644](https://arxiv.org/abs/2208.07644).
- [15] Kohei Honda (1993): *Types for Dyadic Interaction*. In Eike Best, editor: *CONCUR’93*, Lecture Notes in Computer Science, Springer, Berlin, Heidelberg, pp. 509–523, doi:[10.1007/3-540-57208-2_35](https://doi.org/10.1007/3-540-57208-2_35).
- [16] Naoki Kobayashi (2006): *A New Type System for Deadlock-Free Processes*. In Christel Baier & Holger Hermanns, editors: *CONCUR 2006 – Concurrency Theory*, Lecture Notes in Computer Science, Springer Berlin Heidelberg, pp. 233–247, doi:[10.1007/11817949_16](https://doi.org/10.1007/11817949_16).
- [17] Wen Kokke & Ornela Dardha (2021): *Prioritise the Best Variation*. In Kirstin Peters & Tim A. C. Willemse, editors: *Formal Techniques for Distributed Objects, Components, and Systems*, Lecture Notes in Computer Science, Springer International Publishing, Cham, pp. 100–119, doi:[10.1007/978-3-030-78089-0_6](https://doi.org/10.1007/978-3-030-78089-0_6).
- [18] Wen Kokke & Ornela Dardha (2021): *Prioritise the Best Variation*, doi:[10.48550/arXiv.2103.14466](https://doi.org/10.48550/arXiv.2103.14466). arXiv:[2103.14466](https://arxiv.org/abs/2103.14466).

- [19] Sam Lindley & J. Garrett Morris (2015): *A Semantics for Propositions as Sessions*. In Jan Vitek, editor: *Programming Languages and Systems*, Lecture Notes in Computer Science, Springer, Berlin, Heidelberg, pp. 560–584, doi:[10.1007/978-3-662-46669-8_23](https://doi.org/10.1007/978-3-662-46669-8_23).
- [20] Robin Milner (1989): *Communication and Concurrency*. Prentice Hall International Series in Computer Science, Prentice Hall, New York, USA.
- [21] Gerald K. Ostheimer & Antony J. T. Davie (1993): *Pi-Calculus Characterizations of Some Practical Lambda-Calculus Reduction Strategies*. Technical Report CS/93/14, Department of Computing Sciences, University of St Andrews.
- [22] Luca Padovani (2014): *Deadlock and Lock Freedom in the Linear π -Calculus*. In: *Proceedings of the Joint Meeting of the Twenty-Third EACSL Annual Conference on Computer Science Logic (CSL) and the Twenty-Ninth Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, CSL-LICS '14, ACM, New York, NY, USA, pp. 72:1–72:10, doi:[10.1145/2603088.2603116](https://doi.org/10.1145/2603088.2603116).
- [23] Luca Padovani & Luca Novara (2015): *Types for Deadlock-Free Higher-Order Programs*. In Susanne Graf & Mahesh Viswanathan, editors: *Formal Techniques for Distributed Objects, Components, and Systems*, Lecture Notes in Computer Science, Springer International Publishing, Cham, pp. 3–18, doi:[10.1007/978-3-319-19195-9_1](https://doi.org/10.1007/978-3-319-19195-9_1).
- [24] Davide Sangiorgi & David Walker (2003): *The Pi-Calculus: A Theory of Mobile Processes*. Cambridge University Press.
- [25] G. Tremblay & B. Malenfant (2000): *Lenient Evaluation and Parallelism*. *Computer Languages* 26(1), pp. 27–41, doi:[10.1016/S0096-0551\(01\)00007-8](https://doi.org/10.1016/S0096-0551(01)00007-8).
- [26] Philip Wadler (2012): *Propositions As Sessions*. In: *Proceedings of the 17th ACM SIGPLAN International Conference on Functional Programming, ICFP '12*, ACM, New York, NY, USA, pp. 273–286, doi:[10.1145/2364527.2364568](https://doi.org/10.1145/2364527.2364568).