

On Session Typed Contracts for Imperative Languages

Chuta Sano

CMU-CS-19-133

December 2019

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Thesis Committee:

Frank Pfenning, Chair

Iliano Cervesato

*Submitted in partial fulfillment of the requirements
for the degree of Master of Science.*

Copyright © 2019 Chuta Sano

Keywords: Session, Concurrency, Contracts

Meow.

Abstract

Session types prescribe the protocols for communication between concurrently executing processes. Following discoveries of the correspondence between intuitionistic linear logic and linear session types, there have been many related works ranging from practical implementations to theoretical extensions. In particular, linear session types, which assume a strong condition that there is only one client for a session, have been extended to shared session types, which introduce semantics for multiple clients. This extension was subsequently implemented in an imperative setting in the language Concurrent CO. In another direction, contracts, which are well-studied constructs in languages without session types, have been extended to monitors, or contracts for linear session types, in the usual functional setting. In this work, we formalize an imperative programming language based on previous work which implements both linear and shared session typed channels and adapt the monitors to the imperative setting. We further extend the notion of monitoring to shared session types and introduce its semantics. Finally, we introduce several case studies of linear and shared monitors.

Acknowledgments

I would like to express my appreciation to everyone who has been (and will hopefully remain) a part of my life – I am sure that even the smallest of interaction I have had with someone has in some way manifested itself to shape who I am. Of course, the ones with more noticeable impacts, who I should thank with greater weight, include the usual candidates such as my parents, grandparents, cats, friends, etc.

From an academic standpoint, I want to start by thanking my undergraduate advisor Xinwen Fu who gave me a wonderful introduction and impression of the so-called “academia” and Jay McCarthy for introducing me to programming language theory through his inspirational courses.

At Carnegie Mellon University, I was able to meet many inspirational people. I want to start by thanking Iliano Cervesato, who not only served as a committee member for me but also was a very enthusiastic source of answers for many technical questions I had. I then would like to thank my MSCS advisor Dave Eckhardt for many much-needed reminder and help on some of the administrative tasks that I had to do.

Mainly, I want to thank my advisor and committee chair Frank Pfenning for the three semesters of guidance. Since my prior research and technical experience was primarily unrelated to type theory and logic, I had assumed that there were many bridges that I had to cross to even begin having productive conversations, but I felt that he was able to level down the technicality and rigor of whatever topic was at hand for my sake to make me feel comfortable.

Contents

1	Introduction	1
1.1	C0 and Contracts	1
1.2	Linear Session Types	3
1.3	Shared Session Types	5
1.4	Monitors	7
1.5	Concurrent C0	8
1.6	Problem Statement	8
2	Type System and Syntax for MCC0	11
2.1	Types	11
2.1.1	Ordinary Types	11
2.1.2	Ordinary Type Constructors	11
2.1.3	Linear Session Types	12
2.1.4	Shared Session Types	14
2.1.5	Persistent Types	15
2.1.6	Function and Process Types	15
2.2	Structure	16
2.2.1	Expression	16
2.2.2	Statement	16
2.2.3	Declaration and Program	19
2.3	Abstract Syntax Summary	22

2.4	Final Remarks	23
2.5	Type Judgements	24
2.5.1	Assumptions	24
2.5.2	Expressions	24
2.5.3	Statements	26
2.5.4	Declarations	34
2.5.5	Program	34
2.6	Example	35
2.6.1	Remark	35
2.6.2	Parallel Fibonacci	36
3	Linear Monitors and Partial Identity	39
3.1	Introduction	39
3.2	Language	40
3.3	Type Judgements	41
3.3.1	Statements	41
3.3.2	Partial Identity	42
3.3.3	Declarations	48
3.4	Examples	48
3.5	Wrapper	51
4	Shared Monitors	55
4.1	Type Judgements	57
4.1.1	Declaration	57
4.1.2	Statement	57
4.2	Acknowledgement of Accept/Acquire	59
4.3	To Release Message	61
4.4	Examples	63
4.4.1	Extended Partial Identity	63

4.4.2	Acknowledgement and To Release	65
5	Monitors in Concurrent C0	67
6	Case Studies	69
6.1	Linear Monitors in CC0	69
6.1.1	Service Verification	69
6.1.2	Binary Search Tree	71
6.2	Shared Monitors in CC0	74
6.2.1	Shared Ordinary Variables	74
6.2.2	Meta Protocol Enforcement	74
7	Conclusion	81
A	Shared Integer with Monotonicity Monitor Code	83
B	Binary Search Tree with Monitor Code	85
C	Global Deadlock Detection Monitor Code	91
	Bibliography	95

List of Figures

1.1	Left: Linear sessions have exactly one provider and client. Middle: The intended semantics of shared sessions require one provider but allow multiple clients (c_1, c_2, c_3). Right: the two different operational interpretations of shared channels as described in this section. The dotted lines in manifest sharing indicate a blocked channel waiting to acquire. In the figure, c_2 currently has exclusive access to p and c_1, c_3 are blocked waiting for access to p until c_2 releases.	7
2.1	A demonstration of forward by process p . In the context of p , it is offering a channel c which is used by process u and also acting as a client to process q with a channel d . In this context, a forward $c = d$ means that p effectively disappears; all messages sent by u to p are now forwarded to q and vice versa.	18
3.1	A demonstration of a linear monitor m acting as an eavesdropper between a user u and its client channel provided by p	40
4.1	A shared process p being monitored by m . Note that although p now only has one client, it is still a shared process.	55
4.2	General shared monitors. Note that u_3 has unmonitored access to p	56
6.1	The meta protocol that a fixed node is required to follow across multiple acquire/release cycles.	77

List of Tables

2.1 Selected Syntactical Differences in MCC0	38
--	----

Chapter 1

Introduction

1.1 C0 and Contracts

C0 is a C-like language used in CMU’s Principles of Imperative Computation, a class typically taken by first year undergraduate computer science students. Aside from being a simpler subset of C Cre [2007], C0 also provides type and memory safety. Furthermore, C0 supports a rich variety of contracts, which help to enforce a disciplined style of programming Pfenning [2010], Arnold [2010].

To give an intuitive understanding of contracts Meyer [1992], we first consider one of the most common and primitive examples of contracts: assertions. In C0, an assertion is expressed as a statement of form `//@assert e` for some boolean expression e , where the semantics of this statement is that if e evaluates to true, the program continues, making the statement effectively a no-op, and on the other hand if e evaluates to false, then the program aborts with an error message.

Assertions are widely available across various imperative programming languages such as C, C++ Cpp [2017], and Java Gosling et al. [2014], and despite minor variances in syntax and failure consequences, all asserts can be suppressed during compile time, allowing the same source code to produce a slower debug build for developers and a faster assertion-free production build for releases.

Since assertions can be enabled and disabled based on compiler settings, the actual presence of assertions should not affect the state of the program assuming the particular assertion is satisfied. We do mention that certain languages do not explicitly check this requirement, such as C, where for example, one can technically write effectful assertions such as `assert(x++ > 0)`. This is considered bad practice and is heavily discouraged for good reason. In C0 however, this requirement is enforced by the compiler through a pass called the purity check.

More generally, contracts can be thought of as runtime checks and tests that programmers can write which are effectless given that the tests pass. Contracts should be suppressible, meaning that a compiler should be able to transform a program with contracts into an equivalent program without contracts (again, assuming that these runtime tests pass). In this paper, we capture the recurring notion of “assuming that the tests pass” as “partial.”

With this terminology, we can succinctly describe contracts as runtime tests that when part of a program, are partially equivalent to the same program with contracts disabled. An assertion is then a subset of a contract; it is a statement with a predicate which when executed is partially equivalent to a no-op. Since contracts help catch unintended program states and immediately halt the program on failure, they are not only useful in a teaching environment to train good coding discipline but also in industry, where code is updated frequently on large projects.

We also would like to point out that there are more considerations for contracts. For example, a very important topic regarding contracts is the content of the error messages on failure. Concepts like blame assignment, tracing, etc. are also very important aspects of contracts, but due to the preliminary nature of this work, we leave out many of these components and focus only on the fundamental idea behind contracts.

Including the aforementioned assertions, contracts in C0 come in four flavors: pre-conditions, post-conditions, loop invariants, and asserts. Pre-conditions (`@requires`) state requirements at function entry, for example properties of arguments, whereas post-conditions (`@ensures`) state requirements after a function call, for example properties of return values or invariants of parameters that are passed by reference. Loop invariants are assertions that must be true before the exit condition. See Listing 1.1 for an example of the four contracts.

Listing 1.1: An example of a function in C0 with contracts Pfenning [2010].

```
1 int log(int x)
2 //@requires x >= 1;
3 //@ensures \result >= 0;
4 //@ensures (1 << \result) <= x;
5 { int r = 0;
6   int y = x;
7   while (y > 1)
8     //@loop_invariant y >= 1 && r >= 0;
9     //@loop_invariant y * (1 << r) <= x;
10    {
11      y = y / 2;
12      //@assert r < x;
13      r = r + 1;
14    }
15  return r;
16 }
```

The pre-condition at line 2 requires that $x \geq 1$ on function call. The post-condition at line 3 demonstrates the use of the special variable `\result`, which allow post-condition contracts to reference the return value of the function. The post-condition at line 4 shows a usage of a parameter; this is only allowed because x is not mutated within the function definition (hence we copy the value of x to y at line 6). Inside the function body, two loop invariants are demonstrated at lines 8 and 9, and an assert is demonstrated at line 12.

1.2 Linear Session Types

Linear session types, initially coined as “session types,” prescribe bidirectional communication protocols between concurrent processes and guarantee strong properties like deadlock freedom Honda [1993], Honda et al. [1998]. The linearity condition requires that a given session must have exactly one client/user and server/provider. This type system was subsequently shown to correspond to intuitionistic linear logic Caires and Pfenning [2010].

We briefly introduce a modified portion of Honda’s formulation to provide an intuition for linear session types. A more complete and formal formulation is provided in Chapter 2. We first begin with some type T . Then we introduce the input and output type constructors $?T$ and $!T$ which operationally means to receive and send some value of type T respectively. These input and output types are called “actions”. A first attempt at defining a linear session type (A) can then be made as a list of these actions: ¹

$$A ::= \{a_1; a_2; \dots; a_n\}$$

with the meaning being that in this type, an action a_1 is performed and then an action a_2 is performed, and so on. The curly brackets are used to later distinguish between linear and shared session types.

Similar to how an ordinary type, such as `int`, is associated with a variable holding a value of type `int`, a linear session type is associated with a linear channel adhering to the session type. By convention, we type channels from the provider’s perspective; for example, if a protocol requires the provider (server) to send a value of type T , then the corresponding channel’s session type must be $!T$, even though a client must receive a value of type T .

For example, consider a server that when given an integer greater than 1, returns a single prime factor of that natural number. Since we type from the server’s perspective, the protocol can be written `?int; lint`. However, to satisfy this protocol, a client must first send a natural number and then receive a natural number; we see that the semantics of the session type from the client is component-wise “flipped.”

We next introduce branching constructors $A_1 \& A_2$ and $A_1 \oplus A_2$: $a_1 \& a_2$ offers two alternative choices while $A_1 \oplus A_2$ selects one of two alternative choices. This means that a server providing a channel of type $A_1 \& A_2$ must receive a decision between A_1 or A_2 by the client and then proceed as the corresponding session type, whereas a server providing a channel of type $A_1 \oplus A_2$ must send a decision between A_1 and A_2 to the client and then proceed as the corresponding session type.

For example, consider a server that offers a random number. Perhaps it allows a user to select between an `int` and `bool`, both cases providing a random value within the domain

¹The original formulation by Honda does not distinguish actions and ordinary types for type constructors. We present this distinction to better transition to our type system.

of the respective type. In this case, the server offers a choice between two actions `!int` and `!bool` since after the choice has been made by a user, it must send a value. The protocol can therefore be denoted as `!int & !bool`. Just as with actions, we see that the semantics of these branching constructors are flipped for the client as well – whereas the server receives a choice on `&`, a client must send a choice on `&`, and so on.

As shown, each session type when viewed from the server’s perspective gives rise to a dual session type when viewed from the client’s perspective. By convention, we choose to view session types from the server’s perspective, meaning when typechecking the client’s interactions, we implicitly infer the dual of the session type. We also would like to point out that due to channels being bidirectional, the direction of the message is in no way related to whether the process using the channel is providing the channel (server) or using the channel (client).

From an operational standpoint, similar to how a value of an ordinary type can be held by a variable, a single session adhering to some session type (a protocol) can be held by a channel. However, the operations one can perform on channels are not at all similar to ordinary variables since channels do not hold values; they instead hold an avenue of communication to some endpoint. Although expanded later, the informal idea is that we can either send to a channel or receive from a channel. As described so far, we can send or receive values or choices, but we have not covered recursive session types nor any notion of termination of a session.

1.3 Shared Session Types

Although linear session types can express many protocols, the strong assumption of linearity means that we are not able to model many other situations. We recall that what made linear session types “linear” was the assumption that a given channel consists of exactly one client and one provider (server). A shared session type is a weakening of that assumption; a given shared channel consists of potentially multiple clients but only one provider.

One approach to define such an extension was presented by Caires and Pfenning Caires and Pfenning [2010]. To motivate this construction, we start by a simple scenario. Consider a server that provides a single random integer. In the context of linear session

types, the protocol can then be informally expressed as `!int` (recall that we express session types from the provider/server’s perspective). However, it’s desirable for such a service to be available to multiple users with the intended use case being that multiple processes P_1, \dots, P_n can independently receive some random integer from the service. A more practical example is a typical web server, which certainly would like to serve multiple clients.

Caires and Pfenning resolved this problem by introducing the copy semantics shared session type. In the random number example, the protocol would instead be represented as some copiable session type: `copiable(!int)`. Unlike the linear channels, a channel of this session type can have multiple clients. The idea then is that a client can then transform a copiable session type into its underlying linear session type and then proceed like a linear session. From the server’s point of view, it creates a copy of itself as many times as there are clients where each copy operates independently and act linearly.

Balzer and Pfenning introduced another operational interpretation that fits the requirement of a channel serving multiple clients: manifest sharing Balzer and Pfenning [2017]. This new semantics of shared session types was motivated by the inability for the copy semantics version to implement shared resources or retain state; since copy semantics merely replicate a provider, the sessions across different clients are independent of each other, which cannot express mutable resources (if a client writes to a copied provider, then that information cannot be propagated to different copies).

The intuition behind manifest shared session types is that at any given time, only one client can acquire some shared channel. The client that successfully acquired the shared channel must then at some point release the channel, allowing other clients waiting for the shared channel to acquire the channel. In short, a client must follow an acquire, linear, then release pattern. A single iteration of this pattern is called an acquire/release cycle. On the other hand, a server must accept an acquire request, act linearly, and then wait for its client to detach itself from the channel. Similar to the client case, a single iteration of this pattern is called a accept/detach cycle.

Since multiple clients can be waiting to acquire the same shared channel, any release must return the shared channel to its original type. Balzer and Pfenning name this property “equi-synchronizing” Balzer and Pfenning [2017]. This version of shared session types will be formalized in detail in Chapter 2.

We summarize the three forms of session types: linear, copy shared, and manifest shared in Figure 1.1.




Figure 1.1: Left: Linear sessions have exactly one provider and client. Middle: The intended semantics of shared sessions require one provider but allow multiple clients ($c1, c2, c3$). Right: the two different operational interpretations of shared channels as described in this section. The dotted lines in manifest sharing indicate a blocked channel waiting to acquire. In the figure, $c2$ currently has exclusive access to p and $c1, c3$ are blocked waiting for access to p until $c2$ releases.

1.4 Monitors

An addition of channels and session types to a language means that boolean predicates, which we saw as seemingly sufficient, cannot express meaningful tests for a session since channels do not hold any notion of “values.” Any attempt at observing a channel inside the familiar assertions or any similar boolean predicates will cause an observable effect to the program, which seems to violate the recurring idea that contracts must partially have no effect.

Instead, contracts for channels, or monitors, are formulated as intermediary processes that attach to a channel and listen along the channel Gommerstadt et al. [2018]. That is, given some session from a process P to another process Q , a monitor M must attach itself in between the two processes, replacing the session from P to Q with two sessions from P to M and M to Q . Thus, M effectively acts as a middleman with local state that can listen along the communication and raise errors as necessary.

The requirement of partial equivalence under suppression or “no effect” can then be captured by imposing an additional requirement on the monitor; it must faithfully pass all messages it receives, unless a contract is violated. This restriction is called partial identity and is later developed in the context of an imperative language. Therefore, session typed contracts, or monitors, are just partial identity processes.

1.5 Concurrent C0

Concurrent C0, or CC0, was later developed which added inter-process asynchronous communication over linear session-typed channels to C0 Willsey et al. [2017], Gay and Vasconcelos [2010]. Shared session-types were then implemented over CC0 based on the aforementioned manifest sharing approach Balzer and Pfenning [2017]. Therefore in this paper, unless explicitly mentioned, we refer to manifest sharing as “sharing” (similarly, manifest shared session types as “shared session types”).

1.6 Problem Statement

Although CC0 introduced new features for the language, it does not support monitors, which are essentially contracts for session types. This causes a disconnect between the style of C0 and CC0, and although monitors for linear session types had been proposed in a functional setting Gommerstadt et al. [2018], the proposed system does not easily translate to the imperative setting of CC0. Furthermore, there is no known work of a corresponding system of monitors for shared session types.

Thus, our thesis has four main contributions: the adaptation of linear monitors from functional to imperative, the proposal of shared monitors (in an imperative setting) that correspond with the known semantics of linear monitors, thereby extending the notion

of partial identity to shared session types, a corresponding implementation of linear and shared monitors in $CC0$, and case studies which demonstrate the utility and appeal of monitors. We also note that we do not provide formal operational semantics for the proposed type system of linear and shared monitors (which can be understood in analogy with the functional version) and therefore do not prove the usual progress and preservation theorems; we instead rely on informal arguments appealing to intuition.

Chapter 2 introduces Minimalist $CC0$, or $MCC0$, a minimal imperative language with linear and shared session types closely resembling $CC0$, and its corresponding type system. Chapter 3 introduces linear monitors and the corresponding additions to the type system. Chapter 4 introduces shared monitors. Chapter 6 contains various case studies and examples showcasing both variants of monitors in $CC0$.

Chapter 2

Type System and Syntax for MCC0

2.1 Types

We first consider the type system of MCC0.

2.1.1 Ordinary Types

We begin with the ordinary types, τ . The ordinary types are the so-called primitive types. Since we are constructing a minimalist language, the only ordinary type in MCC0 is a `bool`, which is the usual boolean type with two possible values. Note that richer types such as `int`, `bool`, `char`, etc. are available in CC0 and can even be included in MCC0 without affecting most of the type system that we present.

2.1.2 Ordinary Type Constructors

Although we do not include any non-session type related type constructors for simplicity, we do want to briefly discuss the typical type constructors. Constructors such as `structs`, `arrays`, `pointers`, etc. fall in this category. For the most part, these constructors do not affect our type system similar to the addition of more ordinary types like `int`. That being said, we do want to specify that types that somehow reference memory, such as `pointers` (and `arrays` depending on the underlying implementation) would require some care since for example, it does not make sense to send a pointer through a channel; if sent across different machines, the pointer no longer makes sense, and even if sent across processes

with shared memory where it can technically make sense, it seems to go against the point of session types. Allowing pointers also further complicates the partial identity check in Chapter 3 where one verification step involves checking that certain variables are not mutated.

2.1.3 Linear Session Types

Action As introduced, there are two actions (denoted a); for all sendable types T (formally defined in Section 2.3), $!T$ means to send something of type T and $?T$ means to receive something of type T . Thus, at its most basic form, a linear session type is a list of actions denoted:

$$A ::= \{a_1; a_2; \dots ; a_n\}$$

However, this is not sufficient; in particular, it is not clear what happens after a_n . We next introduce a special “action” 1 , or the unit, denoting termination of a session. From the provider’s point of view, it must close a channel of type 1 and from the client’s point of view, it must wait on a channel of type 1 for it to be closed by the provider.

Thus, our second formulation of linear session types consist of a list of non-unit actions (from now on, actions are implicitly assumed to be non-unit) followed by the unit

$$A ::= \{a_1; \dots ; a_n; 1\}$$

where the implication of this session type is that after performing some action a_n , the provider must terminate, and similarly, the user must wait for termination.

External and Internal Choices Let an arbitrary linear session type be denoted as A . We recall the informally introduced binary operators denoting branching, $\&$ and \oplus , where $A_1 \& A_2$ means to offer a choice between A_1 and A_2 and $A_1 \oplus A_2$ means to select a choice between A_1 and A_2 . Note that unlike in its earlier introduction, we specifically require these operators to operate on linear session types and not individual actions. These $\&$ and \oplus are called external and internal choices respectively since $\&$ requires an external decision (user) while \oplus requires the provider to internally make a decision.

For usability and in the interest of more closely matching the type system of $CC0$, we generalize these binary choices to instead allow n -ary choices by the following syntax:

$$!\mathbf{choice}\{\overline{(l, A)}\}, ?\mathbf{choice}\{\overline{(l, A)}\},$$

where the former denotes an internal choice (since the provider “sends” a choice) whereas the latter denotes an external choice (provider “receives” a choice). The l (label) is a metavariable denoting some identifier, and the overline is the usual list notation. Thus, instead of a series of left or right choices like in the binary case, we send and receive a label corresponding to the session type we would like to proceed as.

Thus, our third formulation of linear session types is

$$A ::= \{\bar{a}; 1\} \mid \{\bar{a}; !\mathbf{choice}\{\dots\}\} \mid \{\bar{a}; ?\mathbf{choice}\{\dots\}\}$$

Recursive Choices Although finite session types can express many protocols, there are other protocols that are better represented as recursive protocols. For example, recall the random number service that offers a choice of two types and sends a random number of that type. In the type system that we have introduced so far, this can be typed as an external choice (since it receives a choice from the user):

$$?\mathbf{choice}\{(l_1, \tau_1; 1), (l_2, \tau_2; 1)\}$$

for some arbitrary types τ_1, τ_2 (such as **bool**) and labels l_1, l_2 .

This protocol implies that the service terminates after sending some value of type τ_1 or τ_2 depending on the choice, but it might make sense for the service to continue for as long as the user requests it. In this case, we would like to implement something of type

$$?\mathbf{choice}\{(l_1, \tau_1; ?\mathbf{choice}\{\dots\}), (l_2, \tau_2; ?\mathbf{choice}\{\dots\}), (done, 1)\}$$

where ideally, the inner $\mathbf{choice}\{\dots\}$ is the original type. The final *done* choice was added for flavor with the intention that a user might want to request for the service to terminate.

To resolve this, we can allow choices to be named and referenced. Thus, the final type of the protocol can be written as

$$?\mathbf{choice} \omega\{(l_1, \tau_1; ?\mathbf{choice} \omega), (l_2, \tau_2; ?\mathbf{choice} \omega), (done, 1)\}$$

This is reminiscent of the usual recursive type constructor $\mu\rho.\tau$ but restricted for session types. This construction follows CC0, where there were stylistic motivations to allow only choice naming. From a purely functionality standpoint, general recursive types can still be simulated by using a choice with only one option in a recursive manner.

Summary Although we have introduced named choices to justify recursive session types, in the following sections, we ignore the syntactic details of named choices and instead assume that choices are appropriately expanded as necessary during typechecking.

Therefore, a linear session type A is denoted as

$$A ::= \{\bar{a}; 1\} \mid \{\bar{a}; \mathbf{!choice}\{\overline{(l, A)}\}\} \mid \{\bar{a}; \mathbf{?choice}\{\overline{(l, A)}\}\}$$

2.1.4 Shared Session Types

Recall that in this paper, we ignore the copy shared session types and only focus on the manifest shared session types.

In a shared channel, a client must acquire exclusive access of the provider before proceeding and the provider must accept a single acquire request. This is represented with the shift action: $a ::= \dots \mid \#$

We assume that these session types are equi-synchronizing, which requires linear channels to be released (upshift) at the same type of the acquire (downshift) if it was originally a shared channel. We defer the formal definition and verification rules to the work by Balzer and Pfenning Balzer and Pfenning [2017], section 3.3¹.

Since a shared session type can only downshift to a linear session type and a linear session type can only upshift to a shared session type, instead of introducing two different shift operators, we introduce a new notation for shared session types and infer the direction of the shift (up/down) based on whether the session type is linear or shared.

¹Since equi-synchronizing is a property of session types and is independent of the underlying structure of the language, the fact that we have deviated from a functional setting where equi-synchronizing was proposed to an imperative setting is irrelevant, justifying this deferral.

A shared session type S is then of the following form (note the use of square brackets $[]$ instead of curly brackets $\{\}$):

$$S ::= [\#; \bar{a}; 1] \mid [\#; \bar{a}; !\mathbf{choice}\{\overline{(l, A)}\}] \mid [\#; \bar{a}; ?\mathbf{choice}\{\overline{(l, A)}\}]$$

Remark Since we defined $\#$ to be an action, we can include the initial $\#$ as part of \bar{a} , but we explicitly showed the initial shift to emphasize that a shared channel can only acquire/accept depending on user/provider.

2.1.5 Persistent Types

So far, we introduced three main categories of types: ordinary, linear session, and shared session. We define the persistent types, κ , to be either ordinary or a shared session. Since linear channels are non-copiable, it turns out that shared session types are more closely related to ordinary variables in some sense. Therefore, we group the ordinary types and shared session types in the same category for conciseness.

2.1.6 Function and Process Types

A function takes some inputs as arguments and returns a value of some ordinary type. On the other hand, a process is something that takes some inputs as arguments and provide a channel of some session type.

Therefore the type of a function returning a value of type τ is of form:

$$(\bar{\kappa}, \bar{A}) \rightarrow \tau$$

for persistent types κ and linear session types A .

Similarly, the type of a linear process providing a channel of type A is of form:

$$(\bar{\kappa}, \bar{A}) \rightarrow A$$

and the type of a shared process providing a channel of type S is of form:

$$(\bar{\kappa}) \rightarrow S$$

It is important to note that a shared process cannot depend on linear channels (as shown by the type structure). This is due to the independence principle as explained in Balzer and Pfenning [2017].

2.2 Structure

2.2.1 Expression

Expressions are simple mostly due to our intentionally small ordinary types. An expression e is either one of the boolean constants `true`, `false`, a persistent variable x (either an ordinary variable or a shared channel), or an ordinary function call $f(\bar{e})$, which is a specific function call that does not involve linear channels.

Expressions are intended to be independent of linear channels. Some forms that are typically seen as expressions are instead represented as statements.

2.2.2 Statement

On the other hand, statements are rich. In this section we introduce the interesting statements that mainly deal with linear and shared session types. Other more familiar statements such as variable declaration are listed in the summary table in Section 2.3.

Control Flow We begin with control flow related statements. We allow sequential composition of statements $s_1; s_2$ where the intended meaning is that after s_1 executes, s_2 executes. We also allow an if-then-else statement `if(e) { s_1 } { s_2 }` and a while loop statement `while(e) { s }` with the usual imperative language semantics. Finally, we include a pattern matching over choices which is later explained in this section along with labels in choices.

Shared Spawn and Aliasing Since shared channels are persistent, we group shared spawn and aliasing with ordinary assignments. In particular, a shared spawn involves invoking a shared channel and assigning its offering channel to an alias:

$$h = s(\bar{e}, \bar{A})$$

where h is a shared channel and s is the name of a shared process.

Shared aliasing on the other hand is a simple assignment of form:

$$h = h'$$

which aliases the channel referenced in h' . This is implicitly handled as part of a familiar variable assignment syntax:

$$h = e$$

for some expression e .

Linear Spawn and Rename/Forward We can spawn a new linear process and assign its offering channel to a fresh linear channel in a similar syntax to the shared spawn:

$$d = p(\bar{e}, \bar{A})$$

We can also rename linear channels:

$$d = d'$$

where the distinction here is that since linear channels only allow a single client, d' is no longer available and d simply takes the role of d' .

A special case of the syntactical renaming is a forward:

$$c = d$$

where c is an offering channel in the context of a process and d is some client channel to another process like in Figure 2.1.

In this case, $c = d$ is not a renaming but a signaling that the process will now continue as if it is the process that provides d . Semantically, this statement means that the offering channel c is to be provided by the client channel d .

Wait and Close From the client's perspective, if a linear channel is of type $\{1\}$, it must wait until the channel is closed by the provider. This is expressed as `wait(d)` for some channel d . On the other hand, from the provider's perspective, it must close the channel with the statement `close(c)` for its providing channel c .

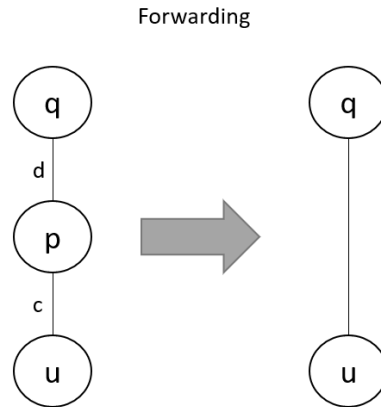


Figure 2.1: A demonstration of forward by process p . In the context of p , it is offering a channel c which is used by process u and also acting as a client to process q with a channel d . In this context, a forward $c = d$ means that p effectively disappears; all messages sent by u to p are now forwarded to q and vice versa.

Send and Receive We previously established the intuitive idea that we either receive or send values along linear channels. We therefore introduce two language primitives for these two operations:

$$\begin{aligned} & \mathbf{send}(d, e) \\ & \mathbf{send}(d, d') \\ & x = \mathbf{recv}(d) \\ & d' = \mathbf{recv}(d) \end{aligned}$$

where the first two send a value (thus a persistent type) or a linear channel respectively and the final two similarly receive from a channel and assign to a persistent typed variable x or a fresh linear channel d' respectively. We chose to make $\mathbf{recv}(\dots)$ not an expression since receiving from a linear channel changes its session type.

Send and Receive Labels Recall that choices consist of a list of pairs (l, A) . Thus, the side that is sending a choice (provider if internal choice and user if external choice) must send this choice through a label. A label l can be sent across a channel d with the syntax $d.l$.

To receive a label, we must be ready to account for all possible choices and provide an implementation of each possible case. This is accomplished using a switch-like statement:

$$\mathbf{switch}(d) \{\overline{l : s}\}$$

with d as the channel which we receive a label (a choice) from. The switch statement consists of a list of pairs consisting of labels and statements. Here, all labels as given by the type of d (recall that the type of d must be an internal or external choice) must be present among the branches.

Shift Operations We recall that given a shared session type, it can only be downshifted into its corresponding linear session type. Similarly, a linear session type with a shift ($\#$) action can only be upshifted into its corresponding shared session type. Operationally, a downshift requires a client to acquire exclusive access to the shared channel and a provider to accept the request. Similarly, an upshift requires a client to release its exclusive access and the server to then detach itself from the client.

$$\begin{aligned} d &= \mathbf{acq}(h) \\ h &= \mathbf{rel}(d) \\ c &= \mathbf{acc}(g) \\ g &= \mathbf{det}(c) \end{aligned}$$

An acquire therefore gives us a fresh linear channel d representing the downshifted channel and release restores the linear channel d back to its shared form. Similarly, an accept gives us a new offering linear channel c representing the shared process's transition to a linear form and a detach restores the linear c to the shared g .

The difference in variable names are reflections of our convention that c is a linear channel a process offers while d is a linear channel that it (or a function) uses and similarly g is a shared offering channel and h is a shared client channel.

2.2.3 Declaration and Program

A declaration dcl is either a function or a process definition.

A function declaration follows a C-like syntax with additional parameters for contracts:

$$\tau f(\overline{\kappa x}, \overline{A d})[\overline{\mathbf{req}(e)}, \overline{\mathbf{ens}(e)}]\{s\}$$

where τ is the return type, f is the function name, $\overline{\kappa x}$ are arguments of persistent types (ordinary and shared session), $\overline{A d}$ are arguments of linear session types, $\mathbf{req}(e)$ are pre-condition contracts, $\mathbf{ens}(e)$ are post-condition contracts, and finally, s is a statement consisting of the implementation of f .

A process declaration offers a channel and hence has a slightly different form:

$$\begin{aligned} &Acp(\overline{\kappa x}, \overline{A d})[\overline{\mathbf{req}(e)}, \overline{\mathbf{ens}(e)}]\{s\} \\ &Sgr(\overline{\kappa x})[\overline{\mathbf{req}(e)}, \overline{\mathbf{ens}(e)}]\{s\} \end{aligned}$$

where c is the name of the linear offering channel and g is the name of the shared offering channel by our convention. Shared processes cannot receive linear channels by the independence principle.

Finally, a program pg consists of a list of declarations dcl .

2.3 Abstract Syntax Summary

Sort	Abstract Form	Remarks
	c, d	Linear channel (offering, using respectively)
	g, h	Shared channel (offering, using respectively)
	x	Persistent variable
Metavariable	p, q	Linear process name
	r	Shared process name
	f	Function name
	l	Choice names (labels)
	z	Offering channel (either linear or shared)
	Program pg	$dcl\ pg$
Decl dcl	$\tau f(\overline{\kappa x}, \overline{A d})[\overline{\mathbf{req}(e)}, \overline{\mathbf{ens}(e)}]\{s\}$	Function with contracts
	$Ac p(\overline{\kappa x}, \overline{A d})[\overline{\mathbf{req}(e)}, \overline{\mathbf{ens}(e)}]\{s\}$	Linear process with contracts
	$Sgr(\overline{\kappa x})[\overline{\mathbf{req}(e)}, \overline{\mathbf{ens}(e)}]\{s\}$	Shared process with contracts
Stmt s	return e	Return ordinary value
	τx	Declare fresh ordinary variable
	$x = e$	Variable assignment
	$x = f(\overline{e}, \overline{d})$	Function call assignment
	$d' = d$	Alias/forward
	$d' = p(\overline{e}, \overline{d'})$	Spawn linear process
	$x = \mathbf{recv}(d)$	Receive ordinary value/shared channel
	$d' = \mathbf{recv}(d)$	Receive linear channel
	$d.l$	Send choice label
	send (d, e)	Send ordinary value or shared channel
	send (d, a)	Send linear channel
	wait (d)	Wait for channel to close
	close (c)	Close channel
	$d = \mathbf{acq}(h)$	Acquire (downshift by client)
	$h = \mathbf{rel}(d)$	Release (upshift by client)
	$d = \mathbf{acc}(h)$	Accept (downshift by provider)
	$h = \mathbf{det}(d)$	Detach (upshift by provider)
	$s_1; s_2$	Sequence of statements
	if (e) { s } { s }	If-then-else statement
	while (e) { s }	While statement
switch (d) { $l : s$ }	Pattern matching over choices	
assert (e)	Assertion (contract)	

Sort	Abstract Form	Remarks
Expression e	true, false	Boolean constants
	$f(\bar{e})$	Function call (no linear channels)
	x	Persistent variable
Persistent type κ	τ	Ordinary Type
	S	Shared session type
Ordinary Type τ	bool	Boolean
Signature type $sigtp$	$(\bar{\kappa}, \bar{A}) \rightarrow \tau$	Function type
	$(\bar{\kappa}, \bar{A}) \rightarrow A$	Linear process type
	$(\bar{\kappa}, \bar{A}) \rightarrow S$	Shared process type
Types T	κ	Persistent type
	A	Linear session type
Action a	$!T$	Provider sends value of T
	$?T$	Provider receives value of T
	$\#$	Shift point
Linear Session Type A	$\{\bar{a}; 1\}$	Ending with unit
	$\{\bar{a}; !\mathbf{choice}\{\overline{(l, A)}\}\}$	Ending with internal choice
	$\{\bar{a}; ?\mathbf{choice}\{\overline{(l, A)}\}\}$	Ending with external choice
Shared Session Type S	$[\#; \bar{a}; 1]$	Ending with unit
	$[\#; \bar{a}; !\mathbf{choice}\{\overline{(l, A)}\}]$	Ending with internal choice
	$[\#; \bar{a}; ?\mathbf{choice}\{\overline{(l, A)}\}]$	Ending with external choice
Session Type Z	A	linear session type
	S	shared session type

2.4 Final Remarks

The minimalist style of this language removes many convenient constructs from CC0 that are effectively sugaring such as for loops (variant of a while loop), if statements without an else clause (an if-then-else with an empty else branch), and simultaneous declaration assignments ($\tau x = e$). Furthermore, ordinary types are extremely limited (only **bool**) for simplicity as mentioned earlier.

Aside from these more obvious differences to CC0, there are some nuanced differences. For example, we are forbidden from referencing linear session variables inside expression level. In particular, `recv(···)` are now statements in the form of variable assignments and function calls that rely on any linear channels are now lifted to statements in the form of variable assignments. In practical terms, these changes disallow nesting receives and arbitrary functions, which are already undefined behaviors in CC0 due to a fixed evaluation order of function arguments not being fully implemented. We also group shared session types with the ordinary types instead of with linear session types since shared session types more closely resemble ordinary types than linear session types in most cases. We refer readers who are interested in the difference between MCC0 and CC0 to the C0 manual Pfenning [2018] for further reading ².

2.5 Type Judgements

2.5.1 Assumptions

Uniqueness of Identifiers Newly introduced identifiers, including function and process names, label names, and variable names, are all assumed to be unique in their respective groupings. That is, we assume no naming conflicts can occur.

Global Function/Process table We assume all functions and process types are globally available through Σ , which consists of a list of function and process identifiers and their corresponding signature types:

$$\Sigma ::= m : sigtp, \Sigma \mid \cdot$$

In particular, this means that all functions and processes are accessible from any context and forward declarations of signatures are not only unavailable but also unnecessary.

2.5.2 Expressions

As a result of an imperative language with its primary focus on linear session types, which are typed inside statement level, expression checking is concise. However, we do note that

²Note that a CC0 reference manual is not available, so this manual will only cover the non-session type related distinctions between MCC0 and CC0

adding practical programming language constructs that were intentionally omitted such as operators would mainly populate expression judgement rules, which of course will not interact with any of the session type checking done on statement level.

We first introduce the usual context consisting of the persistent types:

$$\Gamma ::= \cdot \mid \Gamma, x : \kappa$$

An important point is that we do not have any context for the linear channels when type-checking expressions. This is due to our constraint that expressions should be independent of linear session types and instead closely resemble expressions in typical imperative languages without session types. The fact that we allow some limited uses of shared channels inside expressions is fine because shared channels operate similarly to ordinary variables.

In general, an expression judgement is of form

$$\Gamma \vdash e : \kappa$$

where the judgement reads: in a context Γ , expression e is judged to be of type κ .

We first introduce rules for **true**, **false**, and variables.

$$\frac{}{\Gamma \vdash \mathbf{true} : \mathbf{bool}} \text{ (EXP-TRUE)}$$

$$\frac{}{\Gamma \vdash \mathbf{false} : \mathbf{bool}} \text{ (EXP-FALSE)}$$

$$\frac{}{\Gamma, x : \kappa \vdash x : \kappa} \text{ (EXP-VAR)}$$

Finally, function calls without linear channels requires the signature of f as found in the global Σ to match with its usage.

$$\frac{\Gamma \vdash \bar{e} : \bar{\kappa} \quad f : (\bar{\kappa}) \rightarrow \tau \in \Sigma}{\Gamma \vdash f(\bar{e}) : \tau} \text{ (EXP-FCALL)}$$

2.5.3 Statements

To typecheck statements, we maintain the same persistent context Γ as in the expression judgements:

$$\Gamma ::= \cdot \mid \Gamma, x : \kappa$$

and a linear context Δ :

$$\Delta ::= \cdot \mid \Delta, d : A$$

Linear Context Structurally, Δ and Γ are exactly the same except for the underlying types; Γ stores the typing of persistent variables while Δ similarly stores the typing of linear channels (linear session type variables). However, they operate very differently because of the distinction that the types of persistent variables are permanent whereas the types of linear channels are ephemeral.

Consider for example some client channel $d : \{\text{?bool}; a; \dots\}$. Since it is a client channel, the first action `?bool` requires us to send a value of type `bool`. After executing a statement that sends across d , for example

$$\text{send}(d, \text{true})$$

we say that d “progresses” – d is now of type $\{a; \dots\}$.

This property of linear channels differentiate it from the persistent variables which cannot change in type. Thus, we require a two-state formulation as used by Hodas et al. Hodas et al. [1996]:

$$\Delta \setminus \Delta'$$

where the intended meaning in the later introduced statement judgement rule is that Δ is the linear context before execution of some statement and Δ' is a potentially different linear context after execution of the statement. Γ , on the other hand, does not require a two-state formulation and therefore operates in the usual way.

Statements need to be aware of the context of its declaration. For example, to typecheck `return e`, we need to know that this statement is part of a function declaration and its return type τ . Similarly for processes, we need to know the offering channel name z and its type Z . Although the return type of a function does not change, the analogous offering linear session type of a linear process does change with the same reasoning as for the

linear context. As for shared processes, since an accept somehow “transforms” itself into an exclusive access mode where it acts linearly and changes its name, both the session type and the name of the offering channel changes. Therefore, we require a two-state formulation of the offering channel as well:

$$(z : Z \setminus z' : Z')$$

Next, we define a statement judgement to be in a terminated state if the end of control flow is reached, which means that operationally, no statements can further be executed. Statements that can cause termination include returns for functions and forwarding, tail-call, and closing for processes. This distinction becomes necessary when we typecheck control flow statements like the if-then-else statement; if one branch successfully enters a terminated state while the other does not, typechecking on subsequent statements should proceed from the state that the non-terminating branch ended on. We allow a statement causing termination to be expressed as part of the linear context and the offering channel (only linear context if function since it has no offering channel):

$$\begin{array}{c} \Delta \setminus \downarrow \\ (z : Z \setminus z' : \uparrow) \end{array}$$

We next introduce two meta-variables Δ_{\downarrow} and Z_{\uparrow} :

$$\begin{array}{l} \Delta_{\downarrow} ::= \Delta \mid \downarrow \\ Z_{\uparrow} ::= Z \mid \uparrow \end{array}$$

where \downarrow captures a terminated state through the linear context and \uparrow captures a terminated state through the offering channel.

A statement judgement therefore is of two possible forms: function and process:

$$\begin{array}{l} \Gamma ; \Delta \setminus \Delta'_{\downarrow} \vdash s :: \tau \\ \Gamma ; \Delta \setminus \Delta'_{\downarrow} \vdash s :: (z : Z \setminus z' : Z'_{\uparrow}) \end{array}$$

Because many judgement rules do not necessarily need to reference or change the right side of the double colon, we introduce a meta-variable Θ :

$$\Theta ::= \tau \mid (z : Z \setminus z' : Z'_\uparrow)$$

which can replace the two contexts of declaration. Therefore, a general statement judgement is of form:

$$\Gamma ; \Delta \setminus \Delta'_\downarrow \vdash s :: \Theta$$

We first introduce the two subsumption rules. As described when introducing termination, this is used in control flow operations; when one side of a branch (for example, due to an if statement) terminates, we allow the side of termination to construct an arbitrary state such that we can continue from the state of the other branch.

$$\frac{\Gamma ; \Delta \setminus \downarrow \vdash s :: \tau}{\Gamma ; \Delta \setminus \Delta' \vdash s :: \tau} \text{ (STM-TERMA)}$$

$$\frac{\Gamma ; \Delta \setminus \downarrow \vdash s :: (z : Z \setminus z : \uparrow)}{\Gamma ; \Delta \setminus \Delta' \vdash s :: (z : Z \setminus z' : Z')} \text{ (STM-TERMB)}$$

Since these rules allow inferring any Δ' (and $z : Z'$ for processes), we require considerable care when considering when statements are allowed to terminate.

Next we deal with assignment to persistent variables. This includes assignments on ordinary variables and also shared channels. Therefore, the two rules also account for aliasing and spawning on shared channels respectively. Note that upshifts are later expressed as special cases. Note that in (STM-ASG-F), linear channels that are transferred to the function f are no longer available in our context. This makes sense since linear channels only allow one client.

$$\frac{\Gamma, x : \kappa \vdash e : \kappa}{\Gamma, x : \kappa ; \Delta \setminus \Delta \vdash x = e :: \Theta} \text{ (STM-ASG-E)}$$

$$\frac{\Gamma, x : \tau \vdash \bar{e} : \bar{\kappa} \quad f : (\bar{\kappa}, \bar{A}) \rightarrow \tau \in \Sigma}{\Gamma, x : \tau ; \Delta, \bar{d} : \bar{A} \setminus \Delta \vdash x = f(\bar{e}, \bar{d}) :: \Theta} \text{ (STM-ASG-F)}$$

A return can only occur inside a function and the type must match the return type of the function it is in.

$$\frac{\Gamma \vdash e : \tau}{\Gamma ; \cdot \setminus \downarrow \vdash \mathbf{return} e :: \tau} \text{ (STM-RET)}$$

Since our language cannot have any side effects from a boolean expression due to expressions not being able to operate on channels, we do not need to check for any notion of “effectless” inside assertions. Thus, we only need to check that the inner expression is a predicate; that is, a **bool**. A more extensive language would need to further check that e is pure; it contains no side effects.

$$\frac{\Gamma \vdash e : \mathbf{bool}}{\Gamma ; \Delta \setminus \Delta \vdash \mathbf{assert}(e) :: \Theta} \text{ (STM-ASSERT)}$$

We finally move on to session types. We begin with the simplest rules. A process offering some linear channel c is only allowed to close itself if it is not a client to other linear processes, meaning the linear context Δ needs to be empty. Closing also puts itself in a termination state, represented by the up and down arrows in the respective locations.

$$\frac{}{\Gamma ; \cdot \setminus \downarrow \vdash \mathbf{close}(c) :: (c : \{1\} \setminus c : \uparrow)} \text{ (STM-CLOSE)}$$

On the other hand, a wait simply causes the channel d to disappear from the context after execution. The Θ on the right hand side means that this rule applies to both function and process contexts.

$$\frac{}{\Gamma ; \Delta, d : \{1\} \setminus \Delta \vdash \mathbf{wait}(d) :: \Theta} \text{ (STM-WAIT)}$$

Next, we introduce the send rules for the providing channel from a linear process. Since we view session types from a provider perspective, an action of form $!T$ on the offering channel means that we must send a value of that type. Note that when we send a channel (STM-SEND-CH1), the channel that we send is no longer available in our linear context – this is not surprising because sending a linear channel is effectively transferring the connection since we can only have one client for a given linear channel.

$$\frac{\Gamma \vdash e : \kappa}{\Gamma ; \Delta \setminus \Delta \vdash \mathbf{send}(c, e) :: (c : \{!\kappa; \bar{a}\} \setminus c : \{\bar{a}\})} \text{ (STM-SEND-PERS1)}$$

$$\frac{}{\Gamma ; \Delta, d : A \setminus \Delta \vdash \mathbf{send}(c, d) :: (c : \{!A; \bar{a}\} \setminus c : \{\bar{a}\})} \text{ (STM-SEND-CH1)}$$

On the other hand, an offering channel with an action of form $?T$ requires us to receive something of type T . Analogous to the above send rules, we have one rule each for the persistent types and linear session types.

$$\frac{}{\Gamma, x : \kappa ; \Delta \setminus \Delta \vdash x = \mathbf{recv}(c) :: (c : \{?\kappa; \bar{a}\} \setminus c : \{\bar{a}\})} \text{(STM-RECV-PERS1)}$$

$$\frac{}{\Gamma ; \Delta \setminus \Delta, d : A \vdash d = \mathbf{recv}(c) :: (c : \{?A; \bar{a}\} \setminus c : \{\bar{a}\})} \text{(STM-RECV-CH1)}$$

Sending a label through the offering channel requires the session type to be an internal choice. After selection, the offering channel continues as the linear session type corresponding to the label that it sent.

$$\frac{}{\Gamma ; \Delta \setminus \Delta \vdash c.l_i :: (c : \{!\mathbf{choice}\{\bar{l}, \bar{A}\}\} \setminus c : A_i)} \text{(STM-LABEL1)}$$

On the other hand, we can receive a label through the offering channel if the session type is an external choice. This introduces a pattern matching over all possible labels in the external choice. An important note is that the offering channel c begins as A_i in all the individual cases by the semantics of an external choice. The premise also requires that all branches end up with the same linear context and offering channel type; the subsumption rule will internally be used to satisfy this premise as necessary if any of the individual cases enter a terminated state. Note that the metavariables Δ'_\downarrow and A'_\uparrow are still being used since all branches can lead to termination, in which case we would like to propagate the information that the termination state has been reached.

$$\frac{\forall i (\Gamma ; \Delta \setminus \Delta'_\downarrow \vdash s_i :: (c : A_i \setminus c : A'_\uparrow))}{\Gamma ; \Delta \setminus \Delta'_\downarrow \vdash \mathbf{switch}(c) \{\bar{l} : \bar{s}\} :: (c : \{?\mathbf{choice}\{\bar{l}, \bar{A}\}\} \setminus c : A'_\uparrow)} \text{(STM-SWITCH1)}$$

The following rules dualize the previous rules on the offering channel to an arbitrary client channel inside our linear context. Recall that since we view session types from a provider perspective, the meaning of $!$ and $?$ are flipped from the provider's meaning when typechecking client channels.

$$\frac{\Gamma \vdash e : \kappa}{\Gamma ; \Delta, d : \{?\kappa; \bar{a}\} \setminus \Delta, d : \{\bar{a}\} \vdash \mathbf{send}(d, e) :: \Theta} \text{(STM-SEND-PERS2)}$$

$$\frac{}{\Gamma ; \Delta, d : \{?A; \bar{a}\}, d' : A \setminus \Delta, d : \{\bar{a}\} \vdash \mathbf{send}(d, d') :: \Theta} \text{(STM-SEND-CH2)}$$

$$\frac{}{\Gamma, x : \kappa ; \Delta, d : \{!\kappa; \bar{a}\} \setminus \Delta, d : \{\bar{a}\} \vdash x = \mathbf{recv}(d) :: \Theta} \text{(STM-RECV-PERS2)}$$

$$\frac{}{\Gamma ; \Delta, d : \{!A; \bar{a}\} \setminus \Delta, d : \bar{a}, d' : A \vdash d' = \mathbf{recv}(d) :: \Theta} \text{(STM-RECV-CH2)}$$

$$\frac{}{\Gamma ; \Delta, d : \{?\mathbf{choice}\{\bar{l}, \bar{A}\}\} \setminus \Delta, d : A_i \vdash d.l_i :: \Theta} \text{(STM-LABEL2)}$$

$$\frac{\forall i (\Gamma ; \Delta, d : A_i \setminus \Delta'_\downarrow \vdash s_i :: \tau)}{\Gamma ; \Delta, d : \{!\mathbf{choice}\{\bar{l}, \bar{A}\}\} \setminus \Delta'_\downarrow \vdash \mathbf{switch}(d) \{\bar{l} : s\} :: \Theta} \text{(STM-SWITCH2)}$$

In an imperative language, variable declarations are handled as special cases to statement sequences. We also highlight that whereas shared channels, being a persistent variable, need to be declared, linear channels do not.

$$\frac{\Gamma, x : \kappa ; \Delta \setminus \Delta_\downarrow \vdash s :: \Theta}{\Gamma ; \Delta \setminus \Delta_\downarrow \vdash \kappa x ; s :: \Theta} \text{(STM-VDECL)}$$

Sequences are typechecked component-wise. It is important to note that when checking $s_1; s_2$, we must check s_2 with the state after executing s_1 , meaning we cannot use Θ and instead separate out the process case to name the state after s_1 in (STM-SEQB).

$$\frac{\Gamma ; \Delta_1 \setminus \Delta_2 \vdash s_1 :: \tau \quad \Gamma ; \Delta_2 \setminus \Delta_\downarrow \vdash s_2 :: \tau}{\Gamma ; \Delta_1 \setminus \Delta_\downarrow \vdash s_1; s_2 :: \tau} \text{(STM-SEQA)}$$

$$\frac{\Gamma ; \Delta_1 \setminus \Delta_2 \vdash s_1 :: (z : Z_1 \setminus z' : Z_2) \quad \Gamma ; \Delta_2 \setminus \Delta_\downarrow \vdash s_2 :: (z' : Z_2 \setminus z'' : Z_\uparrow)}{\Gamma ; \Delta_1 \setminus \Delta_\downarrow \vdash s_1; s_2 :: (z : Z_1 \setminus z' : Z_\uparrow)} \text{(STM-SEQB)}$$

The requirement that the returning type τ must remain consistent in (STM-SEQA) reflects our meta-knowledge that statement judgements can only be applied within statements in the same function.

If-then-else statements are actually simpler forms of the pattern matching switch. We require that both branches lead to the same state. Again, if one side terminates but the other does not, the subsumption rules allow us to continue as the side that does not terminate. Similar to the switch rules, we propagate termination state if both branches terminate, which explains the use of the metavariables Δ'_\downarrow and Z'_\uparrow .

$$\frac{\Gamma \vdash e : \mathbf{bool} \quad \Gamma ; \Delta \setminus \Delta'_\downarrow \vdash s_1 :: \tau \quad \Gamma ; \Delta \setminus \Delta'_\downarrow \vdash s_2 :: \tau}{\Gamma ; \Delta \setminus \Delta'_\downarrow \vdash \mathbf{if}(e) \{s_1\} \{s_2\} :: \Theta} \text{(STM-IF)}$$

To typecheck a while statement, we first highlight the fact that the loop can execute zero times to arbitrary many times. Since we cannot statically infer this in the general sense, we require that a single execution of the loop to maintain the state. Thus, linear channels are in some sense loop invariant. Of course, a single execution of the loop can instead terminate the process – this is again allowed due to the subsumption rule. An important distinction from the previous control flow statements such as an if-then-else where if both branches terminate then the entire if-then-else statement is marked as terminating is that even if the inner loop statement can cause termination, we proceed with the original state since the inner loop statement in a while loop can potentially never execute.

Since a loop iteration does not allow linear session types to progress, it may seem natural to question its utility in the context of session types. The only way linear session types can be meaningfully used in while loops is when the underlying session type is recursive; in this case, it is possible for a single loop iteration to return the offering channel and any client channels to their respective original types.

$$\frac{\Gamma \vdash e : \mathbf{bool} \quad \Gamma ; \Delta \setminus \Delta \vdash s :: \tau}{\Gamma ; \Delta \setminus \Delta \vdash \mathbf{while}(e) \{s\} :: \tau} \text{ (STM-WHILEA)}$$

$$\frac{\Gamma \vdash e : \mathbf{bool} \quad \Gamma ; \Delta \setminus \Delta \vdash s :: (z : Z \setminus z : Z)}{\Gamma ; \Delta \setminus \Delta \vdash \mathbf{while}(e) \{s\} :: (z : Z \setminus z : Z)} \text{ (STM-WHILEB)}$$

Since a forwarding effectively ends further execution of the process, it is considered to enter a termination state. To forward, we must only have one linear channel as client similar to the condition imposed in (STM-CLOSE).

$$\frac{}{\Gamma ; d : A \setminus \downarrow \vdash c = d :: (c : A \setminus c : \uparrow)} \text{ (STM-FWD)}$$

A tailcall $c = p(\dots)$ is effectively a syntactic sugar for $d = p(\dots); c = d$. Hence, we require that we transfer all client channels to the process we continue as to mimic the requirement imposed in (STM-FWD).

$$\frac{\Gamma \vdash \bar{e} : \bar{\kappa} \quad p : (\bar{\kappa}, \bar{A}) \rightarrow A' \in \Sigma}{\Gamma ; \bar{d} : A \setminus \downarrow \vdash c = p(\bar{e}, \bar{d}) :: (c : A' \setminus c : \uparrow)} \text{ (STM-TAILCALL)}$$

Renaming a channel simply replaces itself with its new name in the linear context.

$$\frac{}{\Gamma ; \Delta, d' : A \setminus \Delta, d : A \vdash d = d' :: \Theta} \text{ (STM-RENAME)}$$

Spawning a process creates a fresh channel, and the typechecking is mostly similar to function calls. Note that linear channels that are transferred to the spawning process are no longer available in our context, similar to (STM-ASG-F). Unlike with shared channels, we do not need to forward declare linear channels and therefore, they simply appear in the linear context whenever spawned. Furthermore, since shared channels are not allowed to depend on linear channels by the independence principle, a process offering a shared channel cannot spawn (with the implication that once it “turns linear” after accepting, this restriction is removed). As a side note, we would also like to remind the reader that spawning shared processes is handled by the previously introduced (STM-ASG-E).

$$\frac{\Gamma \vdash \bar{e} : \bar{\kappa} \quad p : (\bar{\kappa}, \bar{A}) \rightarrow A' \in \Sigma}{\Gamma ; \Delta, \bar{d} : \bar{A} \setminus \Delta, d' : A' \vdash d' = p(\bar{e}, \bar{d}) :: \tau} \text{ (STM-SPAWN1)}$$

$$\frac{\Gamma \vdash \bar{e} : \bar{\kappa} \quad p : (\bar{\kappa}, \bar{A}) \rightarrow A' \in \Sigma}{\Gamma ; \Delta, \bar{d} : \bar{A} \setminus \Delta, d' : A' \vdash d' = p(\bar{e}, \bar{d}) :: (c : A_c \setminus c : A_c)} \text{ (STM-SPAWN2)}$$

Below we have the two shift statements for clients of shared channels. An acquire creates a fresh linear channel of corresponding type. And a release removes the linear channel from the context. Since acquire creates a fresh channel, we must be in a non-shared context.

$$\frac{}{\Gamma, h : [\#; \bar{a}] ; \Delta \setminus \Delta, d : \{\bar{a}\} \vdash d = \mathbf{acq}(h) :: \tau} \text{ (STM-ACQ1)}$$

$$\frac{}{\Gamma, h : [\#; \bar{a}] ; \Delta \setminus \Delta, d : \{\bar{a}\} \vdash d = \mathbf{acq}(h) :: (c : A_c \setminus c : A_c)} \text{ (STM-ACQ2)}$$

$$\frac{}{\Gamma, h : [\#; \bar{a}] ; \Delta, d : \{\#; \bar{a}\} \setminus \Delta \vdash h = \mathbf{rel}(d) :: \Theta} \text{ (STM-REL)}$$

For offering shared channels, an accept changes the offering channel to the linear version. Thus, typechecking then proceeds as if a linear process until it is restored through a detach, where at this point, there must be no linear channels in context.

$$\frac{}{\Gamma ; \cdot \setminus \cdot \vdash c = \mathbf{acc}(g) :: (g : [\#; \bar{a}] \setminus c : \{\bar{a}\})} \text{ (STM-ACC)}$$

$$\frac{}{\Gamma ; \cdot \setminus \cdot \vdash g = \mathbf{det}(c) :: (c : \{\#; \bar{a}\} \setminus g : [\#; \bar{a}])} \text{ (STM-DET)}$$

2.5.4 Declarations

Although function and process contracts appear as special attributes in the definition, they can instead be thought of as syntactic sugar for statement level asserts³. Since $\text{req}(e)$ is a pre-condition, it is checked as an $\text{assert}(e)$. For $\text{ens}(e)$, we similarly check it as an $\text{assert}(e)$ but also include a fixed variable *result* representing the return type of the function if appropriate. To save space, we will introduce the declaration rules without any declaration-level contracts. We describe a pass to convert all declarations with contracts into an equivalent set of declarations without contracts in section 3.5 which justifies this decision.

$$\frac{\overline{x : \kappa}; \overline{d : A} \setminus \downarrow \vdash s :: \tau}{\tau f(\overline{\kappa x}, \overline{A d})\{s\} \text{ decl}} \text{ (DECL-FUNC)}$$

$$\frac{\overline{x : \kappa}; \overline{d : A} \setminus \downarrow \vdash s :: (c : A' \setminus c : \uparrow)}{A' cp(\overline{\kappa x}, \overline{A d})\{s\} \text{ decl}} \text{ (DECL-LPROC)}$$

$$\frac{\overline{x : \kappa}; \overline{d : A} \setminus \downarrow \vdash s :: (g : S \setminus g' : \uparrow)}{S gr(\overline{\kappa x})\{s\} \text{ decl}} \text{ (DECL-SPROC)}$$

2.5.5 Program

Since we assume a global table, we do not require any context at program level. A program judgement simply requires that all declarations are valid declarations.

$$\frac{}{\cdot \text{ prog}} \text{ (PROG-EMPTY)}$$

$$\frac{dcl \text{ decl} \quad pg \text{ prog}}{dcl \text{ pg prog}} \text{ (PROG-DCL)}$$

³This is not totally true since function and process contracts can be specifically used to obtain blame assignment; a pre-condition violation can give an error message pointing to the malformed caller and so on. We ignore this distinction in this thesis since we do not deal with blame assignment on monitors.

2.6 Example

2.6.1 Remark

Although the formal rules will be written in terms of MCC0, many non-trivial examples of code will be introduced in CC0 for aesthetic reasons. In general, as per keeping the specification of MCC0 as simple as possible, many “shortcuts” that are desirable for programmers are not available, so examples in CC0 would be cleaner and more understandable. See Table 2.1 for a selected summary of syntactical and stylistic differences.

Listing 2.1: Parallel Fibonacci in CC0 with typechecking annotation. Instead of using the two state formation we annotate the result of the context and the succedent after execution of a given statement.

```

1 <!int;> $c fib(int n)
2 { // . |- s :: ($c : <!int;>)
3   if (n == 0 || n == 1) {
4     send($c, 1); // . |- s :: ($c : <>)
5     close($c); // . |- s :: ($c : uparrow)
6   } else {
7     <!int;> $f1 = fib(n-1); // $f1 : <!int;> |- s :: ($c : <!int;>)
8     <!int;> $f2 = fib(n-2); // $f1, $f2 : <!int;> |- ...
9     int x1 = recv($f1); // $f1 : <>, $f2 : <!int;> |- ...
10    wait($f1); // $f2 : <!int;> |- ...
11    int x2 = recv($f2); // $f2 : <> |- ...
12    wait($f2); // . |- ...
13    send($c, x1 + x2); // . |- s :: ($c : <>)
14    close($c); // . |- s :: ($c : uparrow)
15  }
16 }

```

2.6.2 Parallel Fibonacci

We will briefly walk through the typechecker for a process implementing parallel Fibonacci in Listing 2.1. Although the annotation gives reasonable detail in the inner state of the typechecker, the two state formation for Δ was not used due to space reasons. Therefore, the annotations in lines 7 to 12 should be interpreted as the result of Δ after the execution of the corresponding statement, or in other words, the right side of the two state formation.

At line 2, Γ solely consists of n of type `int`. Δ is empty because there are no linear channels in the argument. The process offers a channel of type $\{!int; 1\}$ (as a reminder, the final unit is inferred in CC0), meaning we expect to send an `int` and then terminate.

At line 3, we enter an `if` statement meaning that we now check that both cases terminate with the same state. At line 4 and 5, we send an `int` and then terminate. The termination is valid because $\$c$ is of type $\{1\}$ and Δ is empty. At this point, the linear context is \downarrow and

the offering is \uparrow .

At line 7 and 8, we locally spawn two processes and bind them to the linear channels $\$f1$ and $\$f2$. Therefore,

$$\Delta = \$f1 : \{!int; 1\}, \$f2 : \{!int; 1\}$$

At line 9, we bind a receive result from $\$f1$ to $x1$ and wait for termination of $\$f1$.

By line 13, both local processes have successfully terminated, and Δ is again empty. Γ on the other hand consists of $n, x1, x2$ all of type `int`. Finally on line 13 and 14, we send and close $\$c$ like on line 4 and 5. The termination is valid because Δ is empty.

Since both branches terminated, the `if` statement is also valid since the two states are equivalent. The `if` statement itself is then considered to have terminated since both sides terminate, meaning the process is successfully typechecked.

Note that if we remove termination from one branch, for instance the `if` branch on line 5, then the subsumption rule applied to the `else` branch (since it terminated) would allow us to continue the typechecking as if we are continuing on the `if` branch, meaning the typechecker would correctly require that after the `if else` statement, $\$c$ is closed.

Table 2.1: Selected Syntactical Differences in MCC0

In CC0	In MCC0	Notes
$\$d, \c	d, c	Identifiers of linear channels in CC0 are preceded by a dollar sign.
$\#d, \#c$	h, g	Identifiers of shared channels in CC0 are preceded by a hash. For this reason, the same letters c and d are often used in CC0.
$d = p(d, \dots)$	$d' = p(d, \dots); d = d'$	CC0 can infer that d is no longer available on the right hand side and thus allows an immediate assignment.
$L \$d = \dots$	$d = \dots$	CC0 requires type declarations when spawning linear channels.
$\kappa x = \dots$	$\kappa x; x = \dots$	CC0 allows an assignment on declaration.
typedef $T n$	N/A	CC0 allows a familiar typedef to bind types to names during compilation.
$\langle \dots; \rangle$	$\{\dots; 1\}$	The final unit in session types is inferred in CC0.
choice n { $A_1 \quad l_1;$ $A_2 \quad l_2;$ \dots $}$	choice $\{\dots\}$	In CC0, choices cannot be defined inline and instead are defined at declaration level. Thus, choices in CC0 reference defined choice names instead of constructing it inline.
$A \$c = (A)\#c$ $A \$d = (A)\#d$ $\#c = (S)\$c$ $\#d = (S)\$d$	$c = \mathbf{acc}(g)$ $d = \mathbf{acq}(h)$ $g = \mathbf{det}(c)$ $h = \mathbf{rel}(h)$	In CC0, the four shift operations on shared channels are written in the same syntax as type casts. A and S are assumed to follow the type rules as introduced in Section 2.5.3.

Chapter 3

Linear Monitors and Partial Identity

3.1 Introduction

So far, we have introduced ordinary contracts as statements that, aside from on contract violation, have no observable effect to the outcome of the program.

For ordinary variables, we can “look into” the variable to obtain its value and therefore have the ability to write productive contracts as boolean expressions. However for channels, there is no notion of “value”, and so it does not seem possible to express any meaningful contracts as mere boolean expressions.

As we noted previously, Gommerstadt et al. formulate linear session typed contracts as intermediary processes (monitor) that attach to a channel and listen along the channel Gommerstadt et al. [2018]. See Figure 3.1 for a visual representation of the monitor.

Like contracts, monitors cannot have any observable effect on the program, meaning they must faithfully pass all messages they receive, unless a contract is violated. When a process meets this requirement, it is said to have the property of partial identity. Thus, a valid monitor can be formally stated as a partial identity process and session type contracts are simply statements that attach a monitor to an existing channel.

At this point, it may be helpful to look at some of the examples in Section 3.4.

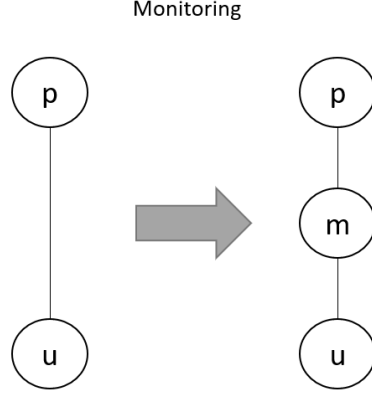


Figure 3.1: A demonstration of a linear monitor m acting as an eavesdropper between a user u and its client channel provided by p .

3.2 Language

We will now augment the language with monitors over linear channels:

$$\begin{aligned}
 dcl & ::= \tau f(\overline{\kappa x}, \overline{A d})[\overline{\mathbf{req}(e)}, \overline{\mathbf{ens}(e)}, \overline{\mathbf{mon}(d = p(d, \bar{e}))}] \{s\} \\
 & \quad | A c p(\overline{\kappa x}, \overline{A d})[\overline{\mathbf{req}(e)}, \overline{\mathbf{ens}(e)}, \overline{\mathbf{mon}(d = p(d, \bar{e}))}] \{s\} \\
 & \quad | S g r(\overline{\kappa x})[\overline{\mathbf{req}(e)}, \overline{\mathbf{ens}(e)}, \overline{\mathbf{mon}(d = p(d, \bar{e}))}] \{s\} \\
 & \quad | \mathbf{mon}(A c p(\overline{A d}, \overline{\tau x})[\overline{\mathbf{req}(e)}, \overline{\mathbf{ens}(e)}, \overline{\mathbf{mon}(d = p(d, \bar{e}))}] \{s\}) \\
 s & ::= \dots \\
 & \quad | \mathbf{mon}(d = p(d, \bar{e})) \\
 sigtp & ::= \dots \\
 & \quad | \mathbf{mon}((A, \overline{tp}) \rightarrow A)
 \end{aligned}$$

All declarations now support monitors over their parameters and providing channel if appropriate. Note that unlike ordinary contracts where there is a syntactic distinction between pre-condition (requires) and post-condition (ensures) contracts, there is no syntactic distinction between monitors that correspond to requires and monitors corresponding to ensures. This is discussed in Section 3.5. The final definition above allows construction of a linear monitor, which is a subset of linear processes. Note that there is only one linear channel as input; the rest are ordinary values (shared channels are not allowed). The addi-

tion of shared channels as parameters in monitors require a more sophisticated type system which we briefly discuss in Section 7. In practice, the explicit annotation of monitors can instead be inferred with a single pass over the program to detect any processes used as monitors.

Statements now support monitors similar to an assert. The syntax and usage is similar to those in declarations. A compiler should simply transform $\mathbf{mon}(d = m(d, \dots))$ to $d' = m(d, \dots); d = d'$ in a debug-build with contracts enabled which effectively inserts the monitor p as in figure 3.1.

Finally, a linear monitor’s typing signature will explicitly be marked as so. In practice, a global table with a listing of all valid linear monitors will serve the same purpose. Although in MCC0 monitors are treated differently from processes, in CC0 they are instead treated as processes. In particular, this means (in CC0) we can use monitors outside contracts to simulate tests that “always run.”

3.3 Type Judgements

3.3.1 Statements

First we add a rule for a statement that attaches a monitor to a channel. Since we explicitly mark monitors and check them separately, we simply assume monitors that are marked as such are valid to use; either the monitor has already been checked or it will be checked in the future.

$$\frac{\Gamma \vdash \bar{e} : \bar{\tau} \quad p : \mathbf{mon}((A, \bar{\tau}) \rightarrow A) \in \Sigma}{\Gamma ; \Delta, d : A \ \backslash \ \Delta, d : A \vdash \mathbf{mon}(d = p(d, \bar{e})) :: \Theta} \text{(STM-MONITOR)}$$

The semantics of this statement is that when the contract is enabled, this statement should locally be transformed into

$$d' = p(d, \dots); d = d'$$

where the first statement means the monitor p will attach itself onto d with the channel name d' and the second statement simply renames d' into the original name d so that future statements that reference d make sense.

3.3.2 Partial Identity

Recall that we introduced partial identity previously as a property of linear processes; a linear process with one providing channel and one client channel that satisfies partial identity is one that faithfully forwards all messages between the providing and client channels except on contract violation. Thus, partial identity check is a form of typechecking with much stronger conditions that only runs on monitors. In particular, the following rules make no assumptions on the well-typedness of any s .

Intuitively, partial identity can be checked by maintaining a queue of messages to be sent by either the providing channel c or the client d depending on the current direction of the messages and requiring the process to adhere to it. For example, if we receive two values consecutively from the client, then we remember that we need to send two consecutive values that we received to the provider and put it in a queue. Thus, we need to consider what types of messages we can receive in general; we can either receive persistent variables, linear channels, labels (for choices), or termination (only from client channel), so we define a member of the queue as follows:

$$q ::= \mathbf{Pers}(x) \mid \mathbf{Ch}(d) \mid \mathbf{Label}(l) \mid \mathbf{End}$$

Next, we introduce a tracking context which keeps track of the providing and the client end along with the current message direction:

$$tc ::= (c : A_{\uparrow} \Leftarrow d : A'_{\downarrow}) \mid (c : A_{\uparrow} \Rightarrow d : A'_{\downarrow})$$

where $(\dots \Leftarrow \dots)$ means that the message is flowing from d to c ; we are currently receiving from the client d and the queue of messages are messages we need to send to c . On the other hand, $(\dots \Rightarrow \dots)$ means that the message is flowing from c to d , meaning we are currently receiving from c and therefore our queue consist of messages we need to send to d .

We also introduce two metavariables:

$$\begin{aligned} \Leftrightarrow &::= \Leftarrow \mid \Rightarrow \\ \omega &::= (c : A \Leftrightarrow d : A') \end{aligned}$$

where \Leftrightarrow represents some direction (either left or right) and ω represents some tracking context.

Since both the queue and tracking context can change from a statement, they will be in a two state formation as with the linear context Δ .

In summary, a partial identity judgement is of form:

$$\begin{array}{l} \Gamma; \Delta \setminus \Delta'; ((c : A_c \Leftrightarrow d : A_d) \setminus (c : A'_c \Leftrightarrow d' : A'_d)); (\bar{q} \setminus \bar{q}') \vdash s \text{ pid} \\ \Gamma; \Delta \setminus \cdot; ((c : A_c \Leftrightarrow d : A_d) \setminus (\uparrow \Leftrightarrow \downarrow)); (\bar{q} \setminus \cdot) \vdash s \text{ pid} \end{array}$$

where Γ and Δ are analogous to the ones in statement judgements and \bar{q} is the message queue. Note that the second tracking context allows a different tracking client d' ; this can occur from renaming $d' = d$. The former denotes a typical partial identity judgement, and the latter denotes a partial identity judgement after entering a termination state. Just as in the statement judgements, we maintain a notion of termination state using the up and down arrows \uparrow and \downarrow . However, for stylistic reasons, we choose to use the tracking client $d : A_d$ to hold \downarrow instead of Δ . Thus, Δ in a terminated state is simply empty: \cdot .

The first rule is a subsumption rule similar to the statement judgements where a terminated state can allow inference of any arbitrary state. In partial identity checks, a state consists of a tracking context, queue, and the linear context. As in the statement judgements, these will be encountered in control flow statements like if or while to allow non-terminated branches to continue if appropriate.

$$\frac{\Gamma; \Delta \setminus \cdot; (\omega \setminus (\uparrow \Leftrightarrow \downarrow)); (\bar{q} \setminus \cdot) \vdash s \text{ pid}}{\Gamma; \Delta \setminus \Delta'; (\omega \setminus \omega'); (\bar{q} \setminus \bar{q}') \vdash s \text{ pid}} \text{ (PID-TERM)}$$

We begin with the partial identity judgement rules that do not involve using the tracking context. Thus, the following rules are almost identical to the statement judgements. The only significant difference is that we now require statements that can modify variables not to modify any members of the queue. This makes sense because the queue only captures what variable it ought to send; if we allow these variables to be mutated, then partial identity no longer becomes guaranteed. A more involved check for mutation becomes necessary if we allow pointers to MCC0. In CC0, where we do have pointers, we simply disallow referencing for members of the queue.

$$\frac{\Gamma, x : \kappa \vdash e : \kappa \quad \mathbf{Pers}(x) \notin \bar{q}}{\Gamma, x : \kappa; \Delta \setminus \Delta; (\omega \setminus \omega); (\bar{q} \setminus \bar{q}) \vdash x = e \text{ pid}} \text{ (PID-ASG-E)}$$

$$\frac{\Gamma, x : \tau \vdash e : \kappa \quad f : (\bar{\kappa}, \bar{A}) \rightarrow \tau \in \Sigma \quad \mathbf{Pers}(x) \notin \bar{q} \quad \overline{\mathbf{Ch}(d)} \notin \bar{q}}{\Gamma, x : \tau; \Delta, \bar{d} : \bar{A} \setminus \Delta; (\omega \setminus \omega); (\bar{q} \setminus \bar{q}) \vdash x = f(\bar{e}, \bar{d}) \text{ pid}} \text{ (PID-ASG-F)}$$

$$\frac{\Gamma \vdash e : \mathbf{bool}}{\Gamma; \Delta \setminus \Delta; (\omega \setminus \omega); (\bar{q} \setminus \bar{q}) \vdash \mathbf{assert}(e) \text{ pid}} \text{ (PID-ASSERT)}$$

$$\frac{\Gamma \vdash \bar{e} : \bar{\tau} \quad p : \mathbf{mon}((A, \bar{\tau}) \rightarrow A) \in \Sigma}{\Gamma; \Delta, d : A \setminus \Delta, d : A; (\omega \setminus \omega); (\bar{q} \setminus \bar{q}) \vdash \mathbf{mon}(d = p(d, \bar{e})) \text{ pid}} \text{ (PID-MONITOR1)}$$

$$\frac{\Gamma \vdash \bar{e} : \bar{\tau} \quad p : \mathbf{mon}((A, \bar{\tau}) \rightarrow A) \in \Sigma}{\Gamma; \Delta \setminus \Delta; ((c : A' \Leftrightarrow d : A) \setminus (c : A' \Leftrightarrow d : A)); (\bar{q} \setminus \bar{q}) \vdash \mathbf{mon}(d = p(d, \bar{e})) \text{ pid}} \text{ (PID-MONITOR2)}$$

At this point, we would like to recall that we do not allow forward declaration of linear channels and that newly introduced channel names are unique. Therefore, although an ordinary assignment $x = e$ requires that x is not in the queue, we do not impose any conditions on d in a renaming: $d = d'$. However, d' cannot be in the queue ¹.

$$\frac{\mathbf{Ch}(d') \notin \bar{q}}{\Gamma; \Delta, d' : A \setminus \Delta, d : A; (\omega \setminus \omega); (\bar{q} \setminus \bar{q}) \vdash d = d' \text{ pid}} \text{ (PID-RENAME1)}$$

$$\frac{\Gamma \vdash \bar{e} : \bar{\kappa} \quad p : (\bar{\kappa}, \bar{A}) \rightarrow A' \in \Sigma \quad \overline{\mathbf{Ch}(d)} \notin \bar{q}}{\Gamma; \Delta, \bar{d} : \bar{A} \setminus \Delta, d' : A'; (\omega \setminus \omega); (\bar{q} \setminus \bar{q}) \vdash d' = p(\bar{e}, \bar{d}) \text{ pid}} \text{ (PID-SPAWN1)}$$

$$\frac{\mathbf{Ch}(d) \notin \bar{q}}{\Gamma; \Delta, d : \{1\} \setminus \Delta; (\omega \setminus \omega); (\bar{q} \setminus \bar{q}) \vdash \mathbf{wait}(d) \text{ pid}} \text{ (PID-WAIT1)}$$

$$\frac{\Gamma \vdash e : \kappa \quad \mathbf{Ch}(d) \notin \bar{q}}{\Gamma; \Delta, d : \{?\kappa; \bar{a}\} \setminus \Delta, d : \{\bar{a}\}; (\omega \setminus \omega); (\bar{q} \setminus \bar{q}) \vdash \mathbf{send}(d, e) \text{ pid}} \text{ (PID-SEND-PERS1)}$$

$$\frac{\mathbf{Ch}(d) \notin \bar{q} \quad \mathbf{Ch}(d') \notin \bar{q}}{\Gamma; \Delta, d : \{?A; \bar{a}\}, d' : A \setminus \Delta, d : \{\bar{a}\}; (\omega \setminus \omega); (\bar{q} \setminus \bar{q}) \vdash \mathbf{send}(d, d') \text{ pid}} \text{ (PID-SEND-CH1)}$$

$$\frac{\mathbf{Pers}(x) \notin \bar{q} \quad \mathbf{Ch}(d) \notin \bar{q}}{\Gamma, x : \kappa; \Delta, d : \{!\kappa; \bar{a}\} \setminus \Delta, d : \{\bar{a}\}; (\omega \setminus \omega); (\bar{q} \setminus \bar{q}) \vdash x = \mathbf{recv}(d) \text{ pid}} \text{ (PID-RECV-PERS1)}$$

$$\frac{\mathbf{Ch}(d) \notin \bar{q}}{\Gamma; \Delta, d : \{!A; \bar{a}\} \setminus \Delta, d : \bar{a}, d' : A; (\omega \setminus \omega); (\bar{q} \setminus \bar{q}) \vdash d' = \mathbf{recv}(d) \text{ pid}} \text{ (PID-RECV-CH1)}$$

¹It is possible to relax this restriction by instead renaming d' to d in the queue, but we chose not to do so for simplicity.

$$\frac{\forall i (\Gamma; \Delta, d : A_i \setminus \Delta'; (\omega \setminus \omega'); (\bar{q} \setminus \bar{q}') \vdash s_i \mathbf{pid}) \quad \mathbf{Ch}(d) \notin \bar{q}}{\Gamma; \Delta, d : \{\mathbf{choice}\{\bar{l}, \bar{A}\}\} \setminus \Delta'; (\omega \setminus \omega'); (\bar{q} \setminus \bar{q}') \vdash \mathbf{switch}(d) \{\bar{l} : s\} \mathbf{pid}} \text{ (PID-SWITCH1)}$$

$$\frac{\Gamma; \Delta_1 \setminus \Delta_2; (\omega_1 \setminus \omega_2); (\bar{q}_1 \setminus \bar{q}_2) \vdash s_1 \mathbf{pid} \quad \Gamma; \Delta_2 \setminus \Delta_3; (\omega_2 \setminus \omega_3); (\bar{q}_2 \setminus \bar{q}_3) \vdash s_2 \mathbf{pid}}{\Gamma; \Delta_1 \setminus \Delta_3; (\omega_1 \setminus \omega_3); (\bar{q}_1 \setminus \bar{q}_3) \vdash s_1; s_2 \mathbf{pid}} \text{ (PID-SEQ)}$$

$$\frac{\Gamma \vdash e : \mathbf{bool} \quad \Gamma; \Delta \setminus \Delta'; (\omega \setminus \omega'); (\bar{q} \setminus \bar{q}') \vdash s_1 \mathbf{pid} \quad \Gamma; \Delta \setminus \Delta'; (\omega \setminus \omega'); (\bar{q} \setminus \bar{q}') \vdash s_2 \mathbf{pid}}{\Gamma; \Delta \setminus \Delta'; (\omega \setminus \omega'); (\bar{q} \setminus \bar{q}') \vdash \mathbf{if}(e) \{s_1\} \{s_2\} \mathbf{pid}} \text{ (PID-IF)}$$

$$\frac{\Gamma \vdash e : \mathbf{bool} \quad \Gamma; \Delta \setminus \Delta; (\omega \setminus \omega); (\bar{q} \setminus \bar{q}) \vdash s \mathbf{pid}}{\Gamma; \Delta \setminus \Delta; (\omega \setminus \omega); (\bar{q} \setminus \bar{q}) \vdash s \mathbf{pid}} \text{ (PID-WHILE)}$$

$$\frac{\mathbf{Pers}(h) \notin \bar{q}}{\Gamma, h : [\#; \bar{a}]; \Delta \setminus \Delta, d : \{\bar{a}\}; (\omega \setminus \omega); (\bar{q} \setminus \bar{q}) \vdash d = \mathbf{acq}(h) \mathbf{pid}} \text{ (PID-ACQ)}$$

$$\frac{\mathbf{Ch}(d) \notin \bar{q} \quad \mathbf{Pers}(h) \notin \bar{q}}{\Gamma, h : [\#; \bar{a}]; \Delta, d : \{\#; \bar{a}\} \setminus \Delta; (\omega \setminus \omega); (\bar{q} \setminus \bar{q}) \vdash h = \mathbf{rel}(d) \mathbf{pid}} \text{ (PID-REL)}$$

Remark The rules so far are for the most part identical to the statement judgements in Section 2.5.3 since the tracking context and the queue were not used. As an example, we compare (STM-SEND-PERS1) and (PID-SEND-PERS1) below:

$$\frac{\Gamma \vdash e : \kappa}{\Gamma; \Delta \setminus \Delta \vdash \mathbf{send}(c, e) :: (c : \{\!|\kappa; \bar{a}\} \setminus c : \{\bar{a}\})} \text{ (STM-SEND-PERS1)}$$

$$\frac{\Gamma \vdash e : \kappa \quad \mathbf{Ch}(d) \notin \bar{q}}{\Gamma; \Delta, d : \{?\kappa; \bar{a}\} \setminus \Delta, d : \{\bar{a}\}; (\omega \setminus \omega); (\bar{q} \setminus \bar{q}) \vdash \mathbf{send}(d, e) \mathbf{pid}} \text{ (PID-SEND-PERS1)}$$

The only significant difference between the two rules is that d , the channel that we send some persistent value to cannot be in the queue as represented by

$$\mathbf{Ch}(d) \notin \bar{q}$$

which reflects the requirement that members in the queue must be immutable to satisfy partial identity.

In general, when we ignore the additional bookkeeping of the tracking context ω and queue \bar{q} , the rules so far are essentially identical aside from the frequent \notin clauses ensure that variables in the queue are not mutated. For persistent variables, this is reflected in assignments, but for linear channels, this is reflected in any statement that may progress a channel (like PID-SEND-PERS1). The following rules consist of the more interesting parts of the partial identity check.

We first introduce the turn rules which allows message direction to change when the queue is empty.

$$\frac{\Gamma; \Delta \setminus \Delta; ((c : \{!t; \bar{a}\} \Leftarrow d : \{!t; \bar{a}\}) \setminus \omega); (\cdot \setminus \bar{q}) \vdash s \text{ pid}}{\Gamma; \Delta \setminus \Delta; ((c : \{!t; \bar{a}\} \Rightarrow d : \{!t; \bar{a}\}) \setminus \omega); (\cdot \setminus \bar{q}) \vdash s \text{ pid}} \text{ (PID-TURN-RL)}$$

$$\frac{\Gamma; \Delta \setminus \Delta; ((c : \{?t; \bar{a}\} \Rightarrow d : \{?t; \bar{a}\}) \setminus \omega); (\cdot \setminus \bar{q}) \vdash s \text{ pid}}{\Gamma; \Delta \setminus \Delta; ((c : \{?t; \bar{a}\} \Leftarrow d : \{?t; \bar{a}\}) \setminus \omega); (\cdot \setminus \bar{q}) \vdash s \text{ pid}} \text{ (PID-TURN-LR)}$$

where t denotes a generic entry in the session type, including any sendable type or a choice (the $!$ specifies that it is an internal choice and so on), in which case the following \bar{a} is empty, since an action cannot follow a choice.

We also add a turn rule to handle the unit type, which implicitly sends through the providing channel and therefore operates like a session type with $!t$:

$$\frac{\Gamma; \Delta \setminus \Delta; ((c : \{1\} \Leftarrow d : \{1\}) \setminus \omega); (\cdot \setminus \bar{q}) \vdash s \text{ pid}}{\Gamma; \Delta \setminus \Delta; ((c : \{1\} \Rightarrow d : \{1\}) \setminus \omega); (\cdot \setminus \bar{q}) \vdash s \text{ pid}} \text{ (PID-TURN-RL2)}$$

A wait and close use the message **End**. Analogous to the statement judgement, **close** requires the linear context to be empty and the queue to be a singleton **End**, since no more messages can be sent after termination.

$$\frac{}{\Gamma; \Delta \setminus \Delta; ((c : A \Leftarrow d : \{1\}) \setminus (c : A \Leftarrow \downarrow)); (\bar{q} \setminus \bar{q}, \mathbf{End}) \vdash \text{wait}(d) \text{ pid}} \text{ (PID-WAIT2)}$$

$$\frac{}{\Gamma; \cdot \setminus \cdot; ((c : \{1\} \Leftarrow \downarrow) \setminus (\uparrow \Leftarrow \downarrow)); (\mathbf{End} \setminus \cdot) \vdash \text{close}(c) \text{ pid}} \text{ (PID-CLOSE)}$$

recv and **send** record and consume $\text{Pers}(x)$ and $\text{Ch}(d')$ messages in the queue respectively.

$$\frac{\mathbf{Pers}(x) \notin \bar{q}}{\Gamma, x : \kappa; \Delta \setminus \Delta; ((c : A \Leftarrow d : \{!\kappa; \bar{a}\}) \setminus (c : A \Leftarrow d : \{\bar{a}\})); (\bar{q} \setminus \bar{q}, \mathbf{Pers}(x)) \vdash x = \mathbf{recv}(d) \mathbf{pid}} \text{ (PID-RECV-PERS2)}$$

$$\frac{}{\Gamma; \Delta \setminus \Delta, d' : A; ((c : A \Leftarrow d : \{!A'; \bar{a}\}) \setminus (c : A \Leftarrow d : \{\bar{a}\})); (\bar{q} \setminus \bar{q}, \mathbf{Ch}(d')) \vdash d' = \mathbf{recv}(d) \mathbf{pid}} \text{ (PID-RECV-CH2)}$$

$$\frac{}{\Gamma, x : \kappa; \Delta \setminus \Delta; ((c : A \Rightarrow d : \{?\kappa; \bar{a}\}) \setminus (c : A \Rightarrow d : \{\bar{a}\})); (\mathbf{Pers}(x), \bar{q} \setminus \bar{q}) \vdash \mathbf{send}(d, x) \mathbf{pid}} \text{ (PID-SEND-PERS2)}$$

$$\frac{}{\Gamma; \Delta, d' : A' \setminus \Delta; ((c : A \Rightarrow d : \{?A'; \bar{a}\}) \setminus (c : A \Rightarrow d : \{\bar{a}\})); (\mathbf{Ch}(d'), \bar{q} \setminus \bar{q}) \vdash \mathbf{send}(d, d') \mathbf{pid}} \text{ (PID-SEND-CH2)}$$

Corresponding rules for the providing channel can be obtained by dualizing as usual:

$$\frac{\mathbf{Pers}(x) \notin \bar{q}}{\Gamma, x : \kappa; \Delta \setminus \Delta; ((c : \{?\kappa; \bar{a}\} \Rightarrow d : A) \setminus (c : \{\bar{a}\} \Rightarrow d : A)); (\bar{q} \setminus \bar{q}, \mathbf{Pers}(x)) \vdash x = \mathbf{recv}(c) \mathbf{pid}} \text{ (PID-RECV-PERS3)}$$

$$\frac{}{\Gamma; \Delta \setminus \Delta, d' : A; ((c : \{?A'; \bar{a}\} \Rightarrow d : A) \setminus (c : \{\bar{a}\} \Rightarrow d : A)); (\bar{q} \setminus \bar{q}, \mathbf{Ch}(d')) \vdash d' = \mathbf{recv}(c) \mathbf{pid}} \text{ (PID-RECV-CH3)}$$

$$\frac{}{\Gamma, x : \kappa; \Delta \setminus \Delta; ((c : \{!\kappa; \bar{a}\} \Leftarrow d : A) \setminus (c : \{\bar{a}\} \Leftarrow d : A)); (\mathbf{Pers}(x), \bar{q} \setminus \bar{q}) \vdash \mathbf{send}(c, x) \mathbf{pid}} \text{ (PID-SEND-PERS3)}$$

$$\frac{}{\Gamma; \Delta, d' : A' \setminus \Delta; ((c : \{!A'; \bar{a}\} \Leftarrow d : A) \setminus (c : \{\bar{a}\} \Leftarrow d : A)); (\mathbf{Ch}(d'), \bar{q} \setminus \bar{q}) \vdash \mathbf{send}(c, d') \mathbf{pid}} \text{ (PID-SEND-CH3)}$$

Switch statements and sending labels again follow a similar pattern:

$$\frac{}{\Gamma; \Delta \setminus \Delta; ((c : A \Rightarrow d : \{\mathbf{choice}\{\bar{l}, \bar{A}\}) \setminus (c : A \Rightarrow d : A_i)); (\bar{q}, \mathbf{Label}(l_i) \setminus \bar{q}) \vdash d.l_i \mathbf{pid}} \text{ (PID-LABEL2)}$$

$$\frac{\forall i (\Gamma; \Delta \setminus \Delta'; ((c : A \Leftarrow d : A_i) \setminus \omega'); (\bar{q}, \mathbf{Label}(l_i) \setminus \bar{q}') \vdash s_i \mathbf{pid})}{\Gamma; \Delta \setminus \Delta'; ((c : A \Leftarrow d : \{\mathbf{choice}\{\bar{l}, \bar{A}\}) \setminus \omega'); (\bar{q} \setminus \bar{q}') \vdash \mathbf{switch}(d) \{\bar{l} : s\} \mathbf{pid}} \text{ (PID-SWITCH2)}$$

$$\frac{}{\Gamma; \Delta \setminus \Delta; ((c : \{\mathbf{choice}\{\bar{l}, \bar{A}\}) \Leftarrow d : A) \setminus (c : A_i \Leftarrow d : A)); (\bar{q}, \mathbf{Label}(l_i) \setminus \bar{q}) \vdash c.l_i \mathbf{pid}} \text{ (PID-LABEL3)}$$

$$\frac{\forall i (\Gamma; \Delta \setminus \Delta'; ((c : \{\mathbf{choice}\{\bar{l}, \bar{A}\}) \Rightarrow d : A) \setminus \omega'); (\bar{q}, \mathbf{Label}(l_i) \setminus \bar{q}') \vdash s_i \mathbf{pid})}{\Gamma; \Delta \setminus \Delta'; ((c : \{\mathbf{choice}\{\bar{l}, \bar{A}\}) \Rightarrow d : A) \setminus \omega'); (\bar{q} \setminus \bar{q}') \vdash \mathbf{switch}(c) \{\bar{l} : s\} \mathbf{pid}} \text{ (PID-SWITCH3)}$$

Forwarding the client simply renames the client in the tracking context, and forwarding the provider (with the client) requires the linear context and queue to be empty. As a trivial consequence of this, a monitor that immediately forwards is a valid partial identity monitor (though useless).

$$\frac{}{\Gamma; \Delta \setminus \Delta; ((c : A \Leftrightarrow d : A') \setminus (c : A \Leftrightarrow d' : A')); (\bar{q} \setminus \bar{q}) \vdash d' = d \mathbf{pid}} \text{ (PID-RENAME2)}$$

$$\frac{}{\Gamma; \cdot \setminus \cdot; ((c : A \Leftrightarrow d : A) \setminus (\uparrow \Leftrightarrow \downarrow)); (\cdot \setminus \cdot) \vdash c = d \mathbf{pid}} \text{ (PID-FWD)}$$

A tailcall on a monitor is a special case to allow monitors on recursive session types.

$$\frac{\Gamma \vdash \bar{e} : \bar{\tau} \quad p : \mathbf{mon}((A, \bar{\tau}) \rightarrow A) \in \Sigma}{\Gamma; \cdot \setminus \cdot; ((c : A \Leftarrow d : A) \setminus (\uparrow \Leftarrow \downarrow)); (\cdot \setminus \cdot) \vdash c = p(d, \bar{e}) \mathbf{pid}} \text{(PID-TAILCALL-MON)}$$

We reject passing the client tracking channel (d) to spawn or tailcall because it requires additional tracking states, adding more clutter to the already complicated inference rules. In particular, statements of form:

$$c = p(\dots, d, \dots)$$

$$d' = p(\dots, d, \dots)$$

are rejected (where c and d are the offering and client tracking channels respectively).

3.3.3 Declarations

Going back to the type judgements for declarations, declarations now contain monitors, which are somewhat analogous to $\mathbf{req}()$ and $\mathbf{ens}()$. Monitors over clients (the inputs) can be mapped to equivalent statement level monitors, similar to $\mathbf{req}()$. However, monitors over the providing channel require typechecking a wrapper process which we describe in Section 3.5. Therefore, we will assume that declarations do not have any contracts, as in Section 2.5.4.

$$\frac{\bar{x} : \bar{\tau}; \cdot \setminus \downarrow; ((c : A \Leftarrow d : A) \setminus (\uparrow \Leftarrow \downarrow)); (\cdot \setminus \cdot) \vdash s \mathbf{pid}}{\mathbf{mon}(A \text{ c p}(A \text{ d}, \bar{\tau} \bar{x}) \{s\}) \mathbf{decl}} \text{(DECL-LMON)}$$

Note that the initial direction of the tracking context does not matter; if the direction happened to be \Rightarrow , then the partial identity check will infer a valid change of directions since the queue is empty from (PID-TURN-LR).

3.4 Examples

Suppose the two processes are checked for partial identity. First, both processes have a client d and an offering c of the same type. They have no other channels as arguments and are therefore possible candidates for being valid monitors.

Listing 3.1: Simple Monitors in CC0

```

1 <!int;> $c is_pos(<!int;> $d)
2 {
3     int x = recv($d); // q = [x]
4     assert(x > 0);
5     send($c, x);      // q = []
6     $c = $d;
7 }
8
9 <!int;> $c is_pos2(<!int;> $d)
10 {
11     int x = recv($d); // q = [x]
12     assert(x > 0);
13     wait($d);         // q = [x, End]
14     send($c, x);     // q = [End]
15     close($c);      // q = []
16 }

```

Both processes begin similarly. On line 3 and 11, both processes bind the result of a receive on the client channel to x . At this point, the queue would consist of a single message $\mathbf{Pers}(x)$ with the message direction being \leftarrow . Line 4 and 12, which although uses a variable x in the queue, is allowed because there is no assignment, so x cannot be mutated.

We now focus on **is_pos**; in line 5, we send x through $\$c$, which is valid since $\mathbf{Pers}(x)$ is the head of the queue. The forwarding at line 6 is valid because the queue is empty. At this point, we are done with the verification, so **is_pos** is indeed a valid monitor.

Moving on to the remaining statements in **is_pos2**, in line 13, we wait on $\$d$. At this point,

$$\bar{q} = \mathbf{Pers}(x), \mathbf{End}$$

Line 14 is therefore valid because $\mathbf{Pers}(x)$ is at the head of the queue, and similarly, line 15 is valid because \mathbf{End} is now the only queue member. Therefore, **is_pos2** is also a valid monitor.

Listing 3.2: List Monitor in CC0

```
1 choice list {
2     <> Nil;
3     <!int; !choice list> Cons;
4 };
5 typedef <!choice list> list;
6
7 list $c nil() {
8     $c.Nil;
9     close($c);
10 }
11
12 list $c cons(int n, list $d) {
13     $c.Cons;
14     send($c, n);
15     $c = $d;
16 }
17
18 list $c allbigger_than(list $d, int threshold) {
19     switch ($d) {
20     case Nil: // q = [Nil]
21         wait($d); // q = [Nil, End]
22         $c.Nil; // q = [End]
23         close($c); // q = []
24     case Cons: { // q = [Cons]
25         int res = recv($d); // q = [Cons, res]
26         assert(res > threshold);
27         $c.Cons; // q = [res]
28         send($c, res); // q = []
29         $c = allbigger_than($d, threshold);
30     }
31     }
32 }
33
34 int sum(list $d)
35 //@monitor $d = allbigger_than($d, 10);
36 {
37     ...
38 }
```

Since line 35 uses the process **allbigger_than** as a monitor, the typechecker infers that it must be checked using partial identity. Again, we will briefly walk through the partial identity checking.

$list$ is a recursive session type: $\{!choice\{(Nil, \{1\}), (Cons, \{!int, list\})\}\}$. We first case over the choices in line 19. In line 20, we assume that we receive the label Nil from d , so the queue consists of $\mathbf{Label}(Nil)$. Lines 21 to 23 are trivial.

In the $Cons$ branch from line 24, we first bind the result of a receive to res . At this point, $\bar{q} = \mathbf{Label}(Cons), \mathbf{Pers}(res)$. In line 27, we send the label $Cons$, which is allowed because the head of the queue is $\mathbf{Label}(Cons)$. Similarly, we send res in line 28. By line 29, the queue is empty, so we are allowed to recurse using **allbigger_than**, which is a monitor.

3.5 Wrapper

So far, we have not formally accounted for contracts attached with functions or processes and instead explained that no specific judgement is required since they can be interpreted as syntactic sugar.

We first classify all declaration contracts into two categories before and after. For ordinary contracts, this is precisely requires and ensures respectively and are already differentiated syntactically. For monitors, we can infer the classification; before monitors include all monitors on the arguments (clients), whereas after monitors include all monitors on the offering channel. Thus, no after monitor exists on a function since it does not offer any channel.

Functions We first consider functions, which are of form:

$$\tau f(\overline{\kappa x}, \overline{A d})[\overline{\mathbf{req}(e_1)}, \overline{\mathbf{ens}(e_2)}, \overline{\mathbf{mon}(d = p(d, \bar{e}_b))}]\{s\}$$

We first rename this function to a fresh f' and strip its declaration contracts with the implication that f' can only be referenced by the inferred wrapper function which we will describe. The wrapper function f has the same form as f' :

$$\tau f(\overline{\kappa x}, \overline{A d})\{s_w\}$$

So what is s_w ? It must first begin with the contracts in the before category, which includes requires and all monitors on the arguments. It must then call f' (the original function), passing in all the arguments without any modification, and assign the result to a special variable $result$. Then, it must execute the contracts in the after category, where for functions, there are no after monitors. Finally, it must return $result$. We present s_w in a more readable C-like syntax below:

```
assert(...); ...; // all requires contracts as asserts
mon(d = ...); ...; // all monitors as statement-level monitors
tau result = f'(...); // pass in all arguments in correct order
assert(...); ...; // all ensures contracts as asserts
return result;
```

Processes Although we have not yet introduced shared monitors, for the sake of generalization, we will assume that shared monitors exist and are syntactically similar to the linear monitors. Therefore, we can monitor the offering channel of not only linear processes but also shared processes (the “after” monitors).

Linear and shared processes are of form:

$$Z z v(\overline{\kappa x}, \overline{A d})[\overline{\mathbf{req}(e_1)}, \overline{\mathbf{ens}(e_2)}, \overline{\mathbf{mon}(d = p(d, \overline{e_b})})}] \{s\}$$

Similar to the function case, we rename v to a fresh v' and strip its declaration contracts. Again, this means that v' cannot be referenced in any context but the wrapper process v . One subtle detail is that the wrapper function must offer a channel with a different name z' – this is because the wrapper cannot immediately tailcall after spawning v' and must instead act as a client to v' temporarily while the contracts in the after category execute.

We present the implementation of the wrapper in C-like syntax below:

```
assert(e1); ...; // all requires contracts as asserts
mon(d = ...); ...; // all before monitors as statement-level monitors
Z z = v'(...); // pass in all arguments in correct order
assert(e2); ...; // all ensures contracts as asserts
mon(z = ...); ...; // all after contracts as statement-level monitors
z' = z           // forward
```

As shown, the wrapper transforms all before contracts into statement contracts at the beginning, calls or spawns the original function or process, and then finally applies the after contracts. One key point is that the wrapper function should provide a different channel name such that after contracts can properly be called.

Chapter 4

Shared Monitors

We next introduce shared monitors. Similar to linear monitors, we view shared monitors as shared processes that satisfy an extended notion of partial identity. By analogy from the linear monitor, we can conceive of a situation where a shared process acts as a "middleman" with multiple clients. This attempt is shown in Figure 4.1.




Figure 4.1: A shared process p being monitored by m . Note that although p now only has one client, it is still a shared process.

However as it turns out, there can be more complicated scenarios where certain clients may have unmonitored access to the provider. This can be done since shared channels can

be aliased either locally through assignment $h' = h$ or by passing to functions or processes. In particular, if both h, h' point to the same shared process, then we can monitor only h , allowing clients using h' to either communicate with the shared process directly or even attach a different monitor to h' . This more general situation is shown in Figure 4.2.




Figure 4.2: General shared monitors. Note that $u3$ has unmonitored access to p .

As it turns out, partial identity for shared monitors at first seems fairly intuitive; similar to the recurring idea that a receive must record to a queue and send must remove elements from a queue, we can view operations on shared channels as variants of receiving and sending, which we will formulate later.

However, a naive code transformation of an otherwise valid monitor statement will not give intended results and therefore require a more involved transformation step. Hence, unlike for linear monitors, where partial identity and type checking were the main interest whereas the translation step was a mere side note, in this section, the translation step is the main interest whereas the type checking is quite trivial. Nevertheless, we begin with the typechecking.

$$\begin{aligned}
dcl & ::= \dots \\
& \quad | \mathbf{mon}(Sgr(Sh, \overline{tpx})[\overline{\mathbf{req}(e)}, \overline{\mathbf{ens}(e)}, \overline{\mathbf{mon}(d = p(d, \bar{e})})}]\{s\}) \\
s & ::= \dots \\
& \quad | \mathbf{mon}(g = p(h, \bar{e})) \\
sigtp & ::= \dots \\
& \quad | \mathbf{mon}((S, \overline{tp}) \rightarrow S)
\end{aligned}$$

The corresponding type rule for statement level shared monitors is reminiscent of the linear version.

$$\frac{\Gamma \vdash \bar{e} : \bar{\tau} \quad k : \mathbf{mon}((S, \bar{\tau}) \rightarrow S) \in \Sigma}{\Gamma, h : S ; \Delta \setminus \Delta \vdash \mathbf{mon}(h = k(h, \bar{e})) :: \Theta} \text{ (STM-SMONITOR)}$$

Similar to the linear case, this statement should locally be transformed to

$$Sh'; h' = k(h, \dots); h = h'$$

when contracts are enabled. Again, the intended meaning is that k monitors h in the second statement with the channel name h' , which is then aliased back to h so that later references to h make sense.

4.1 Type Judgements

4.1.1 Declaration

$$\frac{\overline{x : \tau}; \cdot \setminus \downarrow; ((g : S \Leftarrow h : S) \setminus (\uparrow \Leftarrow \downarrow)); (\cdot \setminus \cdot) \vdash_s \mathbf{pid}}{Sgr(Sh, \overline{\tau x})\{s\} \mathbf{decl}} \text{ (DECL-SMON)}$$

4.1.2 Statement

We want to reiterate the idea that monitors are effectively an expansion over the type structure of the tracking channels augmented with additional local computation and assertions. Partial identity checking for linear channels were complicated mostly because of the rich

structure of linear channels with many available expansion operations. However for shared channels, we can only acquire/release for the client and accept/detach for the provider. For ordinary statement judgements, we observed that after acquiring or accepting, the type checking was continued in the appropriate linear setting where it was assumed that either the linear channel terminated or was upshifted back appropriately. A similar idea applies for partial identity; any non-trivial shared monitor consists of accept and acquire, a linear section, and terminate with release and detach. Thus, checking a shared monitor at first seems to be a matter of checking for appropriate form and then checking the middle linear section, meaning we can defer much of the work to the partial identity for linear channels that we introduced in the previous section.

When we view shared operators as messages, an accept is effectively a receiving of a downshift message, and a detach is a sending of an upshift message. By duality, we then have that an acquire is sending a downshift while release is sending an upshift. Thus, we append two additional possible messages in our queue.

$$q ::= \dots | \mathbf{Dshift} | \mathbf{Ushift}$$

We first introduce the two turn rules. An accept is interpreted as a message from the user (an acquire request), and we will later explain why it makes sense for a detach to anticipate a message from the user as well. We formulate these rules as an analogy to the other turn rules, such as (PID-TURN-LR).

$$\frac{\Gamma; \Delta \setminus \Delta; ((g : [\#; \bar{a}] \Rightarrow h : [\#; \bar{a}]) \setminus \omega); (\cdot \setminus \bar{q}) \vdash s \mathbf{pid}}{\Gamma; \Delta \setminus \Delta; ((g : [\#; \bar{a}] \Leftarrow h : [\#; \bar{a}]) \setminus \omega); (\cdot \setminus \bar{q}) \vdash s \mathbf{pid}} \text{ (PID-TURND)}$$

$$\frac{\Gamma; \Delta \setminus \Delta; ((c : \{\#\; \bar{a}\} \Leftarrow d : \{\#\; \bar{a}\}) \setminus \omega); (\cdot \setminus \bar{q}) \vdash s \mathbf{pid}}{\Gamma; \Delta \setminus \Delta; ((c : \{\#\; \bar{a}\} \Rightarrow d : \{\#\; \bar{a}\}) \setminus \omega); (\cdot \setminus \bar{q}) \vdash s \mathbf{pid}} \text{ (PID-TURNU)}$$

A monitor involving shared channels consists of:

1. accept
2. acquire
3. linear portion
4. release

5. detach

where well-typed arbitrary local computation (statements that do not progress or consume the queue) can be inserted anywhere. Accept must precede acquire since otherwise, a shared monitor will acquire a shared resource and hold it, which can violate partial identity. The following rules are then reflections of this accept, acquire, linear, release, and finally detach paradigm.

$$\frac{}{\Gamma; \Delta \setminus \Delta; ((g : [\#; \bar{a}] \Rightarrow h : [\#; \bar{a}]) \setminus (c : \{\bar{a}\} \Rightarrow h : [\#; \bar{a}])); (\cdot \setminus \mathbf{Ushift}) \vdash c = \mathbf{acc}(g) \text{ pid}} \text{(PID-ACC)}$$

$$\frac{}{\Gamma; \Delta \setminus \Delta; ((c : \{\bar{a}\} \Rightarrow h : [\#; \bar{a}]) \setminus (c : \{\bar{a}\} \Leftarrow d : \{\bar{a}\})); (\mathbf{Ushift} \setminus \cdot) \vdash d = \mathbf{acq}(h) \text{ pid}} \text{(PID-ACQ2)}$$

$$\frac{}{\Gamma; \Delta \setminus \Delta; ((c : \{\#; \bar{a}\} \Leftarrow d : \{\#; \bar{a}\}) \setminus (c : \{\#; \bar{a}\} \Leftarrow h : [\#; \bar{a}])); (\cdot \setminus \mathbf{Dshift}) \vdash h = \mathbf{rel}(d) \text{ pid}} \text{(PID-REL2)}$$

$$\frac{}{\Gamma; \Delta \setminus \Delta; ((c : \{\#; \bar{a}\} \Leftarrow h : [\#; \bar{a}]) \setminus (g : [\#; \bar{a}] \Rightarrow h : [\#; \bar{a}])); (\mathbf{Dshift} \setminus \cdot) \vdash g = \mathbf{det}(c) \text{ pid}} \text{(PID-DET)}$$

4.2 Acknowledgement of Accept/Acquire

Unfortunately, the model of accept, acquire, linear code, release, and then detach is not sufficient for partial identity. This issue can be intuitively understood when viewing shared processes as instead shared resources. With this viewpoint, a client successfully acquiring a shared resource puts itself into a critical region. However, if the shared resource is monitored, we get an arbitrarily long window of delay between the accept and acquire; this can most notably occur when some client has an unmonitored access to the resource, which can be intentionally created through aliasing. Thus, a client successfully acquiring a monitored shared resource does not necessarily mean that it has actually acquired the shared resource, and it is not a mere timing issue.

To resolve this issue, we first introduce an acknowledgement of acceptance to be sent by the provider whenever it accepts. We therefore need to globally append each shift action by a `!bool` for all session types. We do not place any meaning in the value of the boolean and choose only to send true instead of false; therefore, the same effect can be achieved through a traditional unit type. However, we choose not to do so to keep the language minimal but also to avoid confusion between the unit action 1 which signals termination as we introduced and the unit type which is an ordinary type with only one value. In general, the session types

$$\{\#; \bar{a}\}$$
$$[\#; \bar{a}]$$

should respectively transform to

$$\{\#; !\mathbf{bool}; \bar{a}\}$$
$$[\#; !\mathbf{bool}; \bar{a}]$$

Then, outside a monitor, a statement of form

$$c = \mathbf{acc}(g)$$

should transform to

$$c = \mathbf{acc}(g); \mathbf{send}(c, \mathbf{true})$$

and similarly, a statement of form

$$d = \mathbf{acq}(h)$$

should transform to

$$d = \mathbf{acq}(h); x = \mathbf{recv}(d)$$

where x is fresh and therefore not referenced in the context at which this transformation occurs.

Inside a monitor with providing shared channel g and client h that is already checked for partial identity, we encounter the accept and acquire paradigm in the following format:

$$c = \mathbf{acc}(g); s_1; d = \mathbf{acq}(h); s_2$$

where s_1 is some local well-typed computation that does not affect the tracking context nor the queue. These sequences of statements should be transformed to

$$c = \mathbf{acc}(g); s_1; d = \mathbf{acq}(h); x = \mathbf{recv}(d); \mathbf{send}(c, x); s_2$$

Why do these transformations solve the problem? Going back to the original scenario where a client acquires a monitored shared resource, we now see that the client is then blocked from proceeding while waiting for an acknowledgement $x = \mathbf{recv}(d)$. Looking into the monitor, we see that the monitor sends an acknowledgement only after it has successfully acquired the resource.

As with the other introduced notions in the paper, we do not prove a desirable property that a well-typed program will remain well-typed under this global transformation. However, we hope that it is very intuitive why that would be the case; we modify all shared session types and their subsequent uses to add an acknowledgement, making sure no non-monitors that interact with the monitors do not have mismatched session types. For monitors, we make sure to receive an acknowledgement and then forward the result, maintaining partial identity as the queue must be empty at the point of receiving the acknowledgement.

4.3 To Release Message

Unfortunately, the acknowledgement change is a step in the right direction but is still not sufficient. Again, consider a client accessing an unmonitored shared resource. After successfully acquiring the resource, the client enters a critical region and then exits the critical region after it has released the resource. Now consider a monitored shared resource. With the acknowledgement changes, we have successfully fixed the issue where the client can enter a critical region before actually acquiring the resource. However, the monitor, can release the resource before the client has actually released the resource as long as the protocol allows for it. Thus, we still have an issue where the client can unknowingly and prematurely exit a critical region by the whim of the monitor.

This issue is again resolved through a so called “to release” message that is sent by the client. We first prepend all shifts with a `?bool` (again, we do not use the value of the boolean), meaning before an upshift, a boolean must be sent. Thus, we must yet again apply a global transformation on all non-monitors.

$$h = \mathbf{rel}(d)$$

transforms to

$$\mathbf{send}(d, \mathbf{true}); h = \mathbf{rel}(d)$$

and

$$g = \mathbf{det}(c)$$

transforms to

$$x = \mathbf{recv}(c); g = \mathbf{det}(c)$$

where again, x is fresh.

Inside a monitor with providing shared channel g and client h that is already checked for partial identity, we must encounter a release and detach with optional local computation:

$$h = \mathbf{rel}(d); s_1; g = \mathbf{det}(c); s_2$$

where s_1 is some local well-typed computation that does not affect the tracking context nor the queue. These sequences of statements should be transformed to

$$x = \mathbf{recv}(c); \mathbf{send}(d, x); h = \mathbf{rel}(d); s_1; g = \mathbf{det}(c); s_2$$

Going back to the issue we mentioned where, we see that the monitor will now be blocked from releasing the resource until signaled by the client. Thus, we have resolved the problem that the monitor can cause the client to unknowingly exit a critical region by prematurely releasing the resource. We do not prove the desirable property that this transformation still leads to a well-typed program although we believe the proof will mainly be mechanical. However, it is not clear how we can prove that this interpretation of shared monitors preserves a shared analog of the (linear) partial identity property.

Remark We see that enabling contracts for shared processes suddenly require global-level transformations to the code. Until now, all code transformations were local. These transformations were required because we allow both monitored and unmonitored access to shared channels through aliasing. For example, consider some shared channel h . We can then alias h : $h' = h$. Next, if we monitor h only: $\mathbf{mon}(h = \mathbf{mon}(h, \dots))$, then we have h' giving unmonitored access to the original resource and h giving monitored access. Another situation might be attaching two different monitors to the same shared channel with the intent that a certain subset of the clients must follow some invariants while another subset of the clients must follow other invariants. In these situations, we must somehow instrument the program such that acquiring a monitored resource is equivalent to acquiring the resource itself and so on, which due to the semantics of manifest sharing requires additional care in the form of these acknowledgement and to release messages.

From a practical point of view, this means that unmodifiable code, such as libraries must always be provided in a debug variant where shared session types respect the acknowledgement and to release messages and a release variant where shared session types do not.

Listing 4.1: Type Declarations for a shared integer.

```
1 typedef <?choice shd_prov> l_int;
2 typedef <#; l_int> s_int;
3 choice shd_prov {
4     <?int; s_int> Set;
5     <!int; s_int> Get;
6 };
```

4.4 Examples

Consider an implementation of a reference cell for an integer. It consists of a shared process offering a shared channel where after a downshift, it becomes an external choice consisting of the labels *Get* and *Set* where if *Get* is selected by a user, the provider must send an integer and then upshift back to the original type, and if *Set* is selected by a user, it must receive an integer to update the locally stored integer and then upshift back to the original type. The type declarations of this system written in CC0 are introduced in Listing 4.1.

The full code is listed in the Appendix A and the more general implementation of reference cells, or shared variables, is discussed in detail in Section 6.2.1.

4.4.1 Extended Partial Identity

In this section, we detail the partial identity check for a sample monitor which checks that set messages are monotonically increasing.

At lines 4 and 5, we see the only possible way shared monitors can begin (aside from local computations) – an accept followed by an acquire. At this point, our queue is empty and our tracking context consists of the linear channels

$$(\$c : l_int \Leftarrow \$d : l_int)$$

with l_int as defined in Listing 4.1.

Listing 4.2: Monitor code that enforces monotonicity on a shared integer.

```
1 // checks if Set messages are monotonically increasing
2 s_int #c mon(s_int #d, int prev)
3 {
4     l_int $c = (l_int)#c; // accept, q = [Dshift]
5     l_int $d = (l_int)#d; // acquire, q = []
6     switch($c)
7     {
8         case Set: { // q = [Set]
9             int n = recv($c); // q = [Set, n]
10            assert(n >= prev);
11            prev = n;
12            $d.Set; // q = [n]
13            send($d, n); // q = []
14            break;
15        }
16        case Get: { // q = [Get]
17            $d.Get; // q = []
18            int n = recv($d); // q = [n]
19            send($c, n); // q = []
20            break;
21        }
22    }
23    #d = (s_int)$d; // release // q = [Ushift]
24    #c = (s_int)$c; // detach // q = []
25    #c = mon(#d, prev);
26 }
```

Therefore, lines 6 to 22 correspond to the linear section, where we follow linear partial identity checking. Although we already showcase some partial identity examples in Section 3.4, we will still briefly discuss the linear partial identity check that occurs.

The switch statement at line 6 pattern matches over the external choice of **shd_prov**. Since there are two labels, there are correspondingly two cases in the switch statement.

At line 8, we assume we receive a *Set* label, meaning $\$c$ is now of type $\{?int; s_int\}$ and the queue consists of the label message **Label**(*Set*). We then receive an int from $\$c$, appending **Pers**(n) to the queue. At lines 12 and 13, we clear the queue by first sending the label *Set* to $\$d$ and then n . At this point, we finish the *Set* branch with an empty queue and both $\$c$ and $\$d$ ready to upshift.

For the other branch from line 16, we assume a *Get* label. We immediately clear the queue by sending *Get* to $\$d$ at line 17, receiving an integer from $\$d$ at line 18, and clearing the queue at line 19 by sending n to $\$c$. Similar to the *Set* branch, we break with an empty queue and both $\$c$ and $\$d$ ready to upshift, making the entire switch statement valid.

At lines 23 and 24, we detach and release. Finally, we recurse with a tailcall, which is handled by a shared monitor extended version of (PID-TAILCALL-MON).

4.4.2 Acknowledgement and To Release

In this section, we demonstrate the previously described code transformation that occurs after typechecking. Note that programmers would not be exposed to the transformation in typical situations. The annotated transformed code is shown in Listing 4.3 where lines 6 to 9 contain the acknowledgement transformation and lines 27 to 30 contain the to release transformation.

Listing 4.3: Transformed monitor code.

```
1 // checks if Set messages are monotonically increasing
2 s_int #c mon(s_int #d, int prev)
3 {
4     l_int $c = (l_int)#c; // accept
5     l_int $d = (l_int)#d; // acquire
6     // begin acknowledge
7     bool x = recv($d);
8     send($c, x); // notify client that we have successfully acquired
9     // end acknowledge
10    switch($c)
11    {
12        case Set: {
13            int n = recv($c);
14            assert(n >= prev);
15            prev = n;
16            $d.Set;
17            send($d, n);
18            break;
19        }
20        case Get: {
21            $d.Get;
22            int n = recv($d); // q = [n]
23            send($c, n);
24            break;
25        }
26    }
27    // begin to release
28    bool x = recv($c); // block until client notifies that we can release
29    send($d, x);
30    // end to release
31    #d = (s_int)$d; // release
32    #c = (s_int)$c; // detach
33    #c = mon(#d, prev);
34 }
```

Chapter 5

Monitors in Concurrent C0

As shown through some of the examples, the monitors that we introduced in MCC0 were also implemented in CC0. In this section we highlight some of the details of the implementation.

Left Values Although MCC0 allows only variables or channel as left values on assignments, CC0 also allows structs: $s \rightarrow \text{field1} = e$ where s is some pointer to a struct containing field1 . We did not generalize these forms of assignments for the partial identity check, and so CC0 only allows variable and channel assignments when receiving from one of the tracking channels in a partial identity check.

General Spawn and Tailcall In MCC0, we disallowed passing the tracking client d as an argument to a spawn or a tailcall. In CC0, we do allow these constructs; this can be done by remapping the contents of the queue to the context in the new process and continuing the partial identity check. This also requires some generalization of the termination conditions; for example, it was no longer the case that forwarding necessarily required an empty queue. However, it was instead necessary that for a given process, all terminations ended with the same queue and the same message direction, where checking that all terminations ended with an empty queue as in the system we presented is a simple case of this more general condition.

Shared Monitors In the CC0 implementation, we do not allow arbitrary local computation between acquire/accept and release/detach. Thus, a shared monitor in CC0 is slightly more restrictive in form than in MCC0. This came with the benefit that we did not need to

extend the queue messages with the two shift messages and instead check for the correct form.

Chapter 6

Case Studies

6.1 Linear Monitors in CC0

6.1.1 Service Verification

A typical service that requires input from the client can be modeled using a linear session type. For example, consider a service that when given an integer greater than 1, returns the prime factorization of it as a list. One way we express this protocol is for the server to receive an int and then proceed to send a (terminating) list of ints. See Listing 6.1 for the complete implementation of the monitor.

A monitor of such a service can first verify that the initial input is indeed a natural number. This requirement is locally captured in line 14. Next, the monitor can verify that the provided list of factors indeed multiply to the original number. This is done by a local variable y that is multiplied every time the monitor receives a factor as shown in line 32. When the monitor receives a Nil, signifying that no more factors exist, then y is required to match the original input x in line 24. The actual primality of the factors are intentionally left unverified as a design choice since that seems to be just as hard a problem as the service that we are monitoring. Technically, it is possible to verify that each p at line 31 is prime by locally running a factorization algorithm. A probabilistic test, for example Fermat's primality test, can instead sensibly be added, which would allow the monitor to claim great confidence that the factors are indeed prime factors.

Listing 6.1: An implementation of a service verification monitor.

```
1 choice list {
2     <> Nil;
3     <!int; !choice list> Cons;
4 };
5 typedef <!choice list> list;
6 typedef <?int; !choice list> serv;
7
8 // verifies that the initial input is an integer greater than 1
9 // and that the provided list multiplies to the original input
10 // however, it does not check that the list is all prime
11 serv $c ver_factor(serv $d)
12 {
13     int x = recv($c);
14     //@assert x > 1;
15     send($d, x);
16     // at this point, both $c and $d are lists
17     int y = 1; // cumulative product of the factors
18     while(true)
19     {
20         switch($d)
21         {
22             case Nil:
23             {
24                 //@assert y == x;
25                 $c.Nil;
26                 wait($d); close($c);
27             }
28             case Cons:
29             {
30                 $c.Cons;
31                 int p = recv($d);
32                 y = y * p;
33                 send($c, p);
34             }
35         }
36     }
37 }
```

Listing 6.2: Session type representing a binary search tree.

```
1 typedef <?choice bst> bst;
2 choice bst {
3     <?int; !bst; !bool; !bst; > split;
4     <?bst; bst> join;
5     <?int; ?int; bst> chk;
6 };
```

Although the shown service has a relatively simple protocol, this structure of monitoring can readily apply to more involved protocols.

6.1.2 Binary Search Tree

On the other hand, processes can be used to model complex data structures such as a binary search tree.

Consider a binary search tree using a split/join model where each linear process represents a node in the tree. The session type representing a binary search tree process is shown in Listing 6.2.

A binary search tree is an external choice with three labels *split*, *join*, and *chk*. After receiving *split*, the process receives an *int* (denoting the node at which the process should split at) and then proceeds to split the tree and send the left binary search tree, a boolean indicating whether the node was found or not, and then the right binary search tree. After receiving a *join*, the process receives a binary search tree and then remains as a binary search tree (where the implication is that the received binary search tree was merged with the current process). Finally, after receiving a *chk*, the process receives two integers representing the expected range of itself (and its children) and remains itself. While the *split* and *join* labels are required for implementation, the *chk* label is a proposition of some form of non-monitor contract which we discuss in this section. The full code is available in Appendix B.

Let us first consider a monitor **mon** as shown in Listing 6.3.

Listing 6.3: Monitor for a bst node.

```
1 bst $c mon(bst $d, int lo, int hi)
2 {
3     switch($c)
4     {
5     case split:
6     {
7         int k = recv($c);
8         $d.split;
9         send($d, k);
10        bst $l2 = recv($d);
11        bool b = recv($d);
12        bst $r2 = recv($d);
13        wait($d);
14        assert(!b || (hi > k && k > lo));
15        send($c, $l2);
16        send($c, b);
17        send($c, $r2);
18        close($c);
19    }
20    case join:
21    {
22        bst $t = recv($c);
23        //@monitor $t = mon($t, hi, 0x7FFFFFFF);
24        $d.join;
25        send($d, $t);
26        $c = mon($d, lo, 0x7FFFFFFF);
27    }
28    case chk:
29    {
30        $d.chk;
31        int l = recv($c);
32        int u = recv($c);
33        send($d, l);
34        send($d, u);
35        $c = mon($d, lo, hi);
36    }
37    }
38 }
```

Listing 6.4: *chk* case for the node implementation.

```
1  case chk:
2  {
3      int lower = recv($t);
4      int upper = recv($t);
5      assert(upper >= k && k >= lower);
6      $l.chk; // new bound for left children: [lower, k]
7      send($l, lower);
8      send($l, k);
9      $r.chk; // new bound for right children: [k, upper]
10     send($r, k);
11     send($r, upper);
12     $t = node($l, $r, k);
13 }
```

Ignoring the major details of implementation, the monitor sits between two nodes and locally stores a constraint on the range of the values that may be stored in the tree. These values are compared with a split request at line 21 with the meaning that whenever a split request is propagated to this monitor, then it must be in the range (otherwise, the tree should have been split somewhere in the ancestor).

This style of monitoring seems to be the best effort one can make at attempting to enforce a kind of global invariant that a binary search tree must be appropriately ordered across multiple interconnected session types. On the other hand, checking for this invariant is simple by sending a *chk* label. Listing 6.4 showcases a node's implementation of the aforementioned *chk* branch which checks that the tree is properly ordered.

However, this is clearly not a monitor and is instead dynamic checking that is coupled with the implementation of a node. Given our theory so far, it is not clear how one can even statically check that this *chk* label is independent to the implementation (that is, somehow removing it will still give a partially equivalent program). Thus, we see a limitation of monitors, which ultimately reduces to two issues – they can only sit in between two processes and by linearity, they cannot traverse and access information in the way we can with ordinary values. The fact that there are seemingly useful contracts one can write that is not expressible as monitors suggests opportunities for future work which are briefly discussed later.

6.2 Shared Monitors in CC0

6.2.1 Shared Ordinary Variables

In Section 4.4, we briefly introduced the shared session type representing a process that offers a reference cell for an integer (or a shared integer). Although the full code is available in Appendix A, we introduce a shortened version of the code without the monitor implementation in Listing 6.5.

The process definition is fairly self-explanatory; when a *Set* label is received, it sets the internal variable n , and when a *Get* label is received, it sends the internal variable n .

Although CC0 does not implement parametric polymorphism Reynolds [1983], it is easy to imagine that shared variables for any ordinary types (anything with a notion of value) except pointers can be arbitrarily constructed with this paradigm. Using a shared monitor (similar to the example in Appendix A), we can then apply predicates to the received value whenever a *Set* is received. Hence, when we implement shared variables as shared processes, shared monitors then become a natural way to refine the underlying variable type Freeman and Pfenning [1991]. This parallels the idea that refinement types of linear sessions can be implemented as monitors as introduced by Gommerstadt Gommerstadt [2019].

6.2.2 Meta Protocol Enforcement

We next investigate a CC0 implementation by Pfenning Pfenning of a centralized form of Mitchell and Merritt’s distributed deadlock detection algorithm Mitchell and Merritt [1984]. The algorithm assumes a distributed system with shared resources (implemented with shared session types) and nodes. Both nodes and resources are assumed to have unique identifiers represented by ints. From a node’s perspective, this implementation consists of a global shared process that acts as an overseer to some distributed system with all involving nodes and resources known ahead of time. The global overseer can detect (and with some modifications even recover) from deadlocks assuming that each node adheres to not only the protocol denoted by the session type shown in Listing 6.6 but also a sort of “meta protocol” across multiple acquire/release cycles.

Listing 6.5: An implementation of a shared integer.

```
1 typedef <?choice shd_prov> l_int;
2 typedef <#; l_int> s_int;
3 choice shd_prov {
4     <?int; s_int> Set;
5     <!int; s_int> Get;
6 };
7
8 s_int #c shd_int()
9 {
10     int n = 0; // default
11     while (true)
12     {
13         l_int $c = (l_int)#c; // accept
14         switch($c)
15         {
16             case Set: {
17                 n = recv($c);
18                 break;
19             }
20             case Get: {
21                 send($c, n);
22                 break;
23             }
24         }
25         #c = (s_int)$c; // detach
26     }
27 }
```

Listing 6.6: The protocol for the global overseer.

```
1 typedef <?choice overseer> l_overseer;
2 typedef <# ; l_overseer> s_overseer;
3
4 choice overseer {
5     <?int ; ?int ; s_overseer> tryacq; //node id, res id, upshift
6     <?int ; ?int ; s_overseer> didacq; //node id, res id, upshift
7     <?int ; ?int ; s_overseer> willrel; //node id, res id, upshift
8 };
```

Listing 6.7: A simple example showcasing a simple node adhering to the meta protocol.

```
1 <> $c node(int p, int r, srestp #sres, s_overseer #sa)
2 {
3     l_overseer $la = (l_overseer)#sa; // acquire #sa
4     $la.tryacq; // notify the overseer that an attempt at acquire will begin
5     send($la, p); send($la, r);
6     #sa = (s_overseer)$la; // release $la
7     lrestp $lres = (lrestp)#sres; // the actual attempt at acquiring the resource
8     // code procedes once #sres has been acquired
9
10    l_overseer $la = (l_overseer)#sa; // acquire #sa
11    $la.didacq; // notify the overseer that the previous acquire was completed
12    send($la, p); send($la, r);
13    #sa = (s_overseer)$la; // release $la
14
15    // critical region code using $lres
16
17    l_overseer $la = (l_overseer)#sa; // acquire #sa
18    $la.willrel; // notify the overseer that the resource will now be released
19    send($la, p); send($la, r);
20    #sa = (s_overseer)$la; // release $la
21
22    #sres = (srestp)$lres; // the actual release of the resource
23
24    close($c);
25 }
```

To illustrate the meta protocol, we consider a situation where some node with an id p is acquiring a shared resource $\#res$ with an id r . Because the aforementioned overseer is global, it is also available through some channel s of type $s_overseer$. When the node is ready to acquire the shared resource, it must first notify the global overseer before attempting the acquire through the *tryacq* choice. After the acquisition was successful, the node must then notify the global overseer that the acquire was successful through the *didacq* choice. Later, when the node is ready to release the shared resource, it can then notify the global overseer that it will release the resource before the actual release with the *willrel* choice. A minimum working example of such a process is showcased in Listing 6.7.

As shown in this simple case, a node must send three notifications through three individual acquire/release cycles for every shared resource it acquires. For some unique resource r , a node must send a *tryacq*, followed by a *didacq*, and finally a *willrel* for that particular r .

What if a node requires ownership of multiple shared resources at the same time? In this case, the final *willrel* notification can be delayed according to the system. However, a *tryacq* must always be followed by a *didacq*. Therefore, we have some notion of a meta protocol for this shared service – for a fixed node, we must begin with a *tryacq* followed by a *didacq*. At this point, we can either perform yet another *tryacq* followed by a *didacq* or instead send a *willrel*. If we sent a *willrel*, we can again either send a *tryacq* followed by a *didacq* or yet another *willrel*. See Figure 6.1 for a graphical representation of this meta protocol.




Figure 6.1: The meta protocol that a fixed node is required to follow across multiple acquire/release cycles.

Because shared session types cannot remember any information about any previous sessions due to the equi-synchronizing requirement, this sort of meta protocol cannot be statically enforced using the type system. Furthermore, this issue is worsened by the fact that this global overseer can be acquired and released by multiple processes, and since the meta protocol only applies to a fixed node, simply observing that a *tryacq* was sent twice in a row is not sufficient since two separate nodes could have sent a *tryacq*. However as it turns out, shared monitors can enforce this meta protocol by locally keeping track of each process’s last sent notification and upon receiving another notification, verifying that it is indeed an allowed message according to the meta protocol.

Listing 6.8: Monitor structure for the global overseer.

```

1 // should initially be attached with a null pointer
2 s_overseer #sa verify_protocol(s_overseer #sb, pstatus_list* pss)
3 {
4     l_overseer $la = (l_overseer)#sa; //acc
5     l_overseer $lb = (l_overseer)#sb; //acq
6     switch ($la) {
7         case tryacq: {
8             ...
9         }
10        case didacq: {
11            ...
12        }
13        case willrel: {
14            ...
15        }
16    }
17    #sb = (s_overseer)$lb;
18    #sa = (s_overseer)$la;
19    #sa = verify_protocol(#sb, pss);
20 }

```

Aside from this meta protocol, a monitor can verify that a *tryacq* and a subsequent *didacq* were indeed on the same resources (it makes no sense for a node to claim to be acquiring some resource r but then immediately after notify that it has acquired some other resource r'). It can also check that it is not claiming to release a resource it never acquired. Finally, it can also check that it is not attempting to acquire a resource it currently owns.

We layout the structure of the monitor with the helper functions in Listing 6.8. The monitor, named **verify_protocol**, keeps a local list of process statuses, where the type definition of a process status is shown in line 2. Some helper functions interacting with the local data structure are forward declared.

The implementations of the three individual cases (lines 35, 38, and 41) share redundant code and so in the interest of space, we show a sample implementation of the *didacq* case in Listing 6.9.

Listing 6.9: Implementation of the didacq case.

```
1 case didacq: {
2     int p = recv($la);
3     int r = recv($la);
4     process_status* ps = find(pss, p);
5     assert(ps != NULL); // ps should exist since tryacq must have been called
6     assert(ps->acquiring);
7     assert(ps->acquiring_res == r);
8     int_list* new_acq_res = alloc(struct int_list);
9     new_acq_res->n = r;
10    new_acq_res->next = ps->acquired_res;
11    ps->acquired_res = new_acq_res;
12    ps->acquiring = false;
13    $lb.didacq;
14    send($lb, p);
15    send($lb, r);
16    break;
17 }
```

In line 4, we search for the process status entry matching p , the process id we receive. The following three lines at lines 5 to 7 assert that the process status entry exists, a *tryacq* was previously called, and the previous *tryacq* call was on the same resource as r . At this point, we store the resource id r as a resource that is already owned by p to later be checked during *willrel* to verify the requirement that *willrel* must only be communicated on a resource the process owns. At line 12, we update our status to no longer acquiring. This means that a next iteration with the same process allows either a *tryacq* or a *willrel*.

The full implementation of this monitor is available in Appendix C.

Although the verifications are useful and interesting in their own right, we particularly found that this idea that monitors can verify these meta protocols that essentially require fixed clients to follow some protocol across multiple acquire/release cycles to be a very interesting part about shared monitors. This idea of a meta protocol seems to appear in many applications of shared channels, and so we can conclude that monitors can naturally check these meta protocols that cannot be statically typechecked.

We provide no formal description of how this meta protocol monitor can be implemented, but informally, a general implementation would involve a global transformation where some message indicating a unique identifier (say, an int) is inserted to all shared session types (and their corresponding linear session types) following a downshift. This id will then be used by the meta protocol monitor to locally construct some map of the protocols per unique id.

Chapter 7

Conclusion

In this thesis, we have introduced groundwork on the implementation of linear session types and their corresponding known model of contracts, (linear) monitors, to an imperative language. We further extend the type system with shared session types following the manifest sharing semantics where we extended linear monitors to its manifest shared variant and postulated its type rules and semantics. We also enriched the type system of CC0 and its compiler with these additions, which we used to review some practical and theoretical cases of these monitors. Some key observations that we would like to highlight include observations that linear monitors seem to lack expressiveness in certain ways (Section 6.1.2), shared monitors being able to express refinement of shared types (Section 6.2.1), and shared monitors seem to be capable of expressing meta protocols, which seem to appear in many uses of shared processes.

An obvious future work is to prove many of the claims we have made. In particular, although we intuitively believe that our proposed type rules and semantics for shared monitors indeed capture a shared extension to the linear monitors, it seems that a partial identity proof for shared monitors analogous to the proof for linear monitors Gommerstadt et al. [2018] is not trivial.

We also ignored blame assignment, an important component of contracts from a usability perspective. With a message passing model with potentially many forwardings, it seems subjective who to blame when some message is wrong, but either way, unlike with ordinary contracts where blame assignment can be implemented as a transformation on function arguments, with session types, it must be implemented as a transformation on the

protocol itself.

Another important direction we can move towards is to extend the recurring theme of suppressing computation between some debug build and a production build, known as ghost code or ghosting Filliâtre et al. [2016]. We observed that monitors seem more limited than ordinary contracts; we conjecture that an implementation and extension of ghosting to session types can fill the seeming gap that we noticed. In particular, we believe that blame assignment, which is fairly trivial on ordinary contracts, can be implemented for monitors using ghost code based on a preliminary investigation.

Appendix A

Shared Integer with Monotonicity Monitor Code

Listing A.1: Full code showcasing the monitor for a shared integer with its implementation.

```
1 typedef <?choice shd_prov> l_int;
2 typedef <#; l_int> s_int;
3 choice shd_prov {
4     <?int; s_int> Set;
5     <!int; s_int> Get;
6 };
7
8 // checks if Set messages are monotonically increasing
9 s_int #c mon(s_int #d, int prev)
10 {
11     l_int $c = (l_int)#c; // accept, q = [Dshift]
12     l_int $d = (l_int)#d; // acquire, q = []
13     switch($c)
14     {
15         case Set: { // q = [Set]
16             int n = recv($c); // q = [Set, n]
17             assert(n >= prev);
18             prev = n;
19             $d.Set; // q = [n]
20             send($d, n); // q = []
21             break;
```

```

22     }
23     case Get: {           // q = [Get]
24         $d.Get;           // q = []
25         int n = recv($d); // q = [n]
26         send($c, n);     // q = []
27         break;
28     }
29 }
30 #d = (s_int)$d; // release // q = [Ushift]
31 #c = (s_int)$c; // detach // q = []
32 #c = mon(#d, prev);
33 }
34
35 s_int #c shd_int()
36 //@monitor #c = mon(#c, 0x80000000);
37 {
38     int n = 0; // default
39     while (true)
40     {
41         l_int $c = (l_int)#c; // accept
42         switch($c)
43         {
44             case Set: {
45                 n = recv($c);
46                 break;
47             }
48             case Get: {
49                 send($c, n);
50                 break;
51             }
52         }
53         #c = (s_int)$c; // detach
54     }
55 }

```

Appendix B

Binary Search Tree with Monitor Code

The code is based on an implementation in SILL Pfenning and Griffith [2015], Toninho et al. [2013] by Gommerstadt Gommerstadt.

Listing B.1: Full code showcasing the split join based binary search tree.

```
1 typedef <?choice bst> bst;
2 choice bst {
3     <?int; !bst; !bool; !bst; > split;
4     <?bst; bst> join;
5     <?int; ?int; bst> chk;
6 };
7
8 bst $c mon(bst $d, int lo, int hi)
9 {
10     switch($c)
11     {
12     case split:
13     {
14         int k = recv($c);
15         $d.split;
16         send($d, k);
17         bst $l2 = recv($d);
18         bool b = recv($d);
19         bst $r2 = recv($d);
20         wait($d);
21         assert(!b || (hi > k && k > lo));
22         send($c, $l2);
```

```

23     send($c, b);
24     send($c, $r2);
25     close($c);
26 }
27 case join:
28 {
29     bst $t = recv($c);
30     //@monitor $t = mon($t, hi, 0x7FFFFFFF);
31     $d.join;
32     send($d, $t);
33     $c = mon($d, lo, 0x7FFFFFFF);
34 }
35 case chk:
36 {
37     $d.chk;
38     int l = recv($c);
39     int u = recv($c);
40     send($d, l);
41     send($d, u);
42     $c = mon($d, lo, hi);
43 }
44 }
45 }
46
47 bst $t node(bst $l, bst $r, int k)
48 {
49     switch($t)
50     {
51     case split:
52     {
53         int pivot = recv($t);
54         if (k > pivot)
55         {
56             $l.split;
57             send($l, pivot);
58             bst $l2 = recv($l);
59             bool b = recv($l);
60             bst $r2 = recv($l);
61             wait($l);
62             send($t, $l2);
63             send($t, b);

```

```

64     bst $r3 = node($r2, $r, k);
65     //@monitor $r3 = mon($r3, pivot, 0x7FFFFFFF);
66     send($t, $r3);
67     close($t);
68 }
69 else if (pivot == k)
70 {
71     send($t, $l);
72     send($t, true);
73     send($t, $r);
74     close($t);
75 }
76 else // pivot > k
77 {
78     $r.split;
79     send($r, pivot);
80     bst $l2 = recv($r);
81     bool b = recv($r);
82     bst $r2 = recv($r);
83     wait($r);
84     bst $l3 = node($l, $l2, k);
85     //@monitor $l3 = mon($l3, 0x80000000, pivot);
86     send($t, $l3);
87     send($t, b);
88     send($t, $r2);
89     close($t);
90 }
91 }
92 case join:
93 {
94     bst $t2 = recv($t);
95     $r.join;
96     send($r, $t2);
97     $t = node($l, $r, k);
98 }
99 case chk:
100 {
101     int lower = recv($t);
102     int upper = recv($t);
103     assert(upper >= k && k >= lower);
104     $l.chk; // new bound for left children: [lower, k]

```

```

105     send($l, lower);
106     send($l, k);
107     $r.chk; // new bound for right children: [k, upper]
108     send($r, k);
109     send($r, upper);
110     $t = node($l, $r, k);
111 }
112 }
113 }
114
115 bst $t leaf()
116 {
117     switch($t)
118     {
119     case split:
120     {
121         int k = recv($t);
122         bst $l = leaf();
123         bst $r = leaf();
124         send($t, $l);
125         send($t, false);
126         send($t, $r);
127         close($t);
128     }
129     case join:
130     {
131         bst $t2 = recv($t);
132         $t = $t2;
133     }
134     case chk: // trivially a valid bst
135     {
136         int u = recv($t);
137         int r = recv($t);
138         $t = leaf();
139     }
140     }
141 }
142
143 bst $t singleton(int a)
144 {
145     bst $l = leaf();

```



```

146     bst $r = leaf();
147     $t = node($l, $r, a);
148 }
149
150 bst $t insert(bst $t1, int a)
151 {
152     bst $m = singleton(a);
153     $t1.split;
154     send($t1, a);
155     bst $l = rcv($t1);
156     bool b = rcv($t1);
157     bst $r = rcv($t1);
158     wait($t1);
159     $m.join;
160     send($m, $r);
161     $l.join;
162     send($l, $m);
163     $t = $l;
164 }
165
166 bst $t delete(bst $t1, int a)
167 {
168     $t1.split;
169     send($t1, a);
170     bst $l = rcv($t1);
171     bool b = rcv($t1);
172     bst $r = rcv($t1);
173     wait($t1);
174     $l.join;
175     send($l, $r);
176     $t = $l;
177 }
178
179 // send bool then become bst
180 <!bool; bst> $res find(bst $t, int a)
181 {
182     $t.split;
183     send($t, a);
184     bst $l = rcv($t);
185     bool b = rcv($t);
186     bst $r = rcv($t);

```

```
187     wait($t);
188
189     send($res, b);
190     if (b)
191     {
192         $res = node($l, $r, a);
193     }
194     else
195     {
196         $l.join;
197         send($l, $r);
198         $res = $l;
199     }
200 }
```

Appendix C

Global Deadlock Detection Monitor Code

Listing C.1: Full code showcasing the monitor for the global overseer.

```
1 typedef struct process_status process_status;
2 struct process_status {
3     int pid; // process id
4     bool acquiring; // is it in the middle of acquiring?
5     int acquiring_res; // if acquiring, resource id being acquired
6     int_list* acquired_res; // list of acquired resources
7 };
8
9 typedef struct pstatus_list pstatus_list;
10 struct pstatus_list {
11     process_status* ps;
12     pstatus_list* next;
13 };
14
15 process_status* new_process_status(int p);
16
17 // true if list contains r
18 // should be called when will_release is received
19 bool found_res(int_list* acquired_res, int r);
20
21 // removes the first occurrence of x from int_list
22 int_list* remove_res(int_list* acquired_res, int x);
```

```

23
24 // returns pointer to a process status corresponding to p
25 // if it cannot find it, returns a null pointer
26 process_status* find(pstatus_list* pss, int p);
27
28 // should initially be attached with a null pointer
29 s_overseer #sa verify_protocol(s_overseer #sb, pstatus_list* pss)
30 {
31     l_overseer $la = (l_overseer)#sa; //acc
32     l_overseer $lb = (l_overseer)#sb; //acq
33     switch ($la) {
34         case tryacq: {
35             int p = recv($la);
36             int r = recv($la);
37             process_status* ps = find(pss, p);
38             if (ps == NULL)
39             {
40                 // add new ps
41                 ps = new_process_status(p);
42                 pstatus_list* new_pss = alloc(struct pstatus_list);
43                 new_pss->ps = ps;
44                 new_pss->next = pss;
45                 pss = new_pss;
46             }
47             assert(!ps->acquiring); // should not be acquiring
48             ps->acquiring = true;
49             ps->acquiring_res = r;
50             $lb.tryacq;
51             send($lb, p);
52             send($lb, r);
53             break;
54         }
55         case didacq: {
56             int p = recv($la);
57             int r = recv($la);
58             process_status* ps = find(pss, p);
59             assert(ps != NULL); // ps exists from earlier tryacq
60             assert(ps->acquiring);
61             assert(ps->acquiring_res == r);
62             int_list* new_acq_res = alloc(struct int_list);
63             new_acq_res->n = r;

```

```

64     new_acq_res->next = ps->acquired_res;
65     ps->acquired_res = new_acq_res;
66     ps->acquiring = false;
67     $lb.didacq;
68     send($lb, p);
69     send($lb, r);
70     break;
71 }
72 case willrel: {
73     int p = recv($la);
74     int r = recv($la);
75     process_status *ps = find(pss, p);
76     assert(ps != NULL);
77     assert(!ps->acquiring); // not be acquiring
78     assert(found_res(ps->acquired_res, r)); // only release acquired resources
79     int_list* new_acq_res = remove_res(ps->acquired_res, r);
80     ps->acquired_res = new_acq_res;
81     $lb.willrel;
82     send($lb, p);
83     send($lb, r);
84     break;
85 }
86 }
87 #sb = (s_overseer)$lb;
88 #sa = (s_overseer)$la;
89 #sa = verify_protocol(#sb, pss);
90 }

```

Bibliography

- Iso/iec 9899:1999. <http://www.open-std.org/jtc1/sc22/wg14/www/docs/n1256.pdf>, 09 2007. 1.1
- Working draft, standard for programming language c++. <http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2017/n4713.pdf>, 11 2017. 1.1
- Rob Arnold. *C₀, an Imperative Programming Language for Novice Computer Scientists*. M.s. thesis, Department of Computer Science, Carnegie Mellon University, December 2010. Available as Technical Report CMU-CS-10-145. 1.1
- Stephanie Balzer and Frank Pfenning. Manifest sharing with session types. In *International Conference on Functional Programming (ICFP)*, pages 37:1–37:29. ACM, September 2017. Extended version available as Technical Report CMU-CS-17-106R, June 2017. 1.3, 1.3, 1.5, 2.1.4, 2.1.6
- Luís Caires and Frank Pfenning. Session types as intuitionistic linear propositions. In *Proceedings of the 21st International Conference on Concurrency Theory (CONCUR 2010)*, pages 222–236, Paris, France, August 2010. Springer LNCS 6269. 1.2, 1.3
- Jean-Christophe Filliâtre, Léon Gondelman, and Andrei Paskevich. The spirit of ghost code. *Formal Methods in System Design*, 48(3):152–174, 2016. ISSN 1572-8102. doi: 10.1007/s10703-016-0243-x. 7
- Tim Freeman and Frank Pfenning. Refinement types for ML. In *Proceedings of the SIGPLAN '91 Symposium on Language Design and Implementation*, pages 268–277, Toronto, Ontario, June 1991. ACM Press. 6.2.1
- Simon J. Gay and Vasco T. Vasconcelos. Linear type theory for asynchronous session types. *Journal of Functional Programming*, 20(1):19–50, January 2010. 1.5
- Hannah Gommerstadt. personal communication. B

- Hannah Gommerstadt. *Session-Typed Concurrent Contracts*. PhD thesis, Carnegie Mellon University, September 2019. Available as Technical Report CMU-CS-19-119. 6.2.1
- Hannah Gommerstadt, Limin Jia, and Frank Pfenning. Session-typed concurrent contracts. In A. Ahmed, editor, *European Symposium on Programming (ESOP'18)*, pages 771–798, Thessaloniki, Greece, April 2018. Springer LNCS 10801. 1.4, 1.6, 3.1, 7
- James Gosling, Bill Joy, Guy L. Steele, Gilad Bracha, and Alex Buckley. *The Java Language Specification, Java SE 8 Edition*. Addison-Wesley Professional, 1st edition, 2014. ISBN 013390069X, 9780133900699. 1.1
- Joshua S. Hodas, Frank Pfenning, and Iliano Cervesato. Efficient Resource Management for Linear Logic Proof Search. In R. Dyckhoff, H. Herre, and P. Schröder-Heister, editors, *International Workshop on Extensions of Logic Programming — ELP'96*, pages 67–81, Leipzig, Germany, 03 1996. Springer-Verlag LNAI 1050. 2.5.3
- Kohei Honda. Types for dyadic interaction. In *4th International Conference on Concurrency Theory, CONCUR'93*, pages 509–523. Springer LNCS 715, 1993. 1.2
- Kohei Honda, Vasco T. Vasconcelos, and Makoto Kubo. Language primitives and type discipline for structured communication-based programming. In *7th European Symposium on Programming Languages and Systems (ESOP 1998)*, pages 122–138. Springer LNCS 1381, 1998. 1.2
- Bertrand Meyer. Applying "design by contract". *IEEE Computer*, 25(10):40–51, 1992. doi: 10.1109/2.161279. URL <https://doi.org/10.1109/2.161279>. 1.1
- Don P. Mitchell and Michael Merritt. A distributed algorithm for deadlock detection and resolution. In *Symposium on Principles of Distributed Computation (PODC 1984)*, pages 282–284, Vancouver, British Columbia, August 1984. ACM. 6.2.2
- Frank Pfenning. personal communication. 6.2.2
- Frank Pfenning. C0: Specification and verification in introductory computer science. <http://c0.typesafety.net>, 2010. 1.1, 1.1
- Frank Pfenning. C0 reference. <http://c0.typesafety.net/doc/c0-reference.pdf>, 03 2018. 2.4
- Frank Pfenning and Dennis Griffith. Polarized substructural session types. In A. Pitts, editor, *Proceedings of the 18th International Conference on Foundations of Software Science and Computation Structures (FoSSaCS 2015)*, pages 3–22, London, England, April 2015. Springer LNCS 9034. Invited talk. B

John C. Reynolds. Types, abstraction, and parametric polymorphism. In R.E.A. Mason, editor, *Information Processing 83*, pages 513–523. Elsevier, September 1983. 6.2.1

Bernardo Toninho, Luís Caires, and Frank Pfenning. Higher-order processes, functions, and sessions: A monadic integration. In M.Felleisen and P.Gardner, editors, *Proceedings of the European Symposium on Programming (ESOP'13)*, pages 350–369, Rome, Italy, March 2013. Springer LNCS 7792. B

Max Willsey, Rokhini Prabhu, and Frank Pfenning. Design and implementation of Concurrent C0. *Electronic Proceedings in Theoretical Computer Science*, 238:73–82, 01 2017. doi: 10.4204/EPTCS.238.8. 1.5