

Resource-Aware Session Types for Digital Contracts

Ankush Das

CMU-CS-21-112

May 2021

Computer Science Department
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Thesis Committee:

Jan Hoffmann, Chair

Frank Pfenning

Bryan Parno

Andrew Miller (University of Illinois Urbana-Champaign)

Shaz Qadeer (Facebook)

*Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy.*

Copyright © 2021 Ankush Das

This research was sponsored by the National Science Foundation award numbers CCF1812876 and CNS1801369. The views and conclusions contained in this document are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of any sponsoring institution, the U.S. government or any other entity.

Keywords: Programming Languages, Resource Analysis, Session Types, Language Design, Refinement Types, Smart Contracts, Blockchain

To Prof. Supratik Chakraborty, Prof. Akshay S, and Dr. Akash Lal for introducing me to the beautiful world of programming languages

Abstract

Programming distributed systems is already very challenging due to the presence of data races and deadlocks; bugs are difficult to detect and reproduce when they only arise in certain thread interleavings. The rise of modern distributed systems have introduced unique domain-specific challenges further complicating software development. Although program analysis tools exist for distributed systems, the most popular and usable tools are still centered around traditional programming languages. With the pervasive usage of distributed systems in software design, there is an urgent need for formal tools to help with the design, verification, and quantitative analysis of distributed software.

In response, this thesis designs novel resource-aware session types that serve as a sound and practical foundation for distributed systems with strong type-theoretic guarantees. Session types statically prescribe and enforce bidirectional communication protocols for message-passing processes. However, they cannot express quantitative properties of a distributed system, such as energy consumption, latency, response time, and throughput. This thesis addresses this limitation by designing two extensions to express the *work* and *span* of parallel computation. To compute work, the key innovation was that messages and processes both carry an abstract notion of *potential* which is consumed to perform work. To compute span, the key innovation was to introduce operators from temporal logic to capture the timing of message exchanges. Resource-aware session types combine session types with work and span extensions allowing programmers to reason about both qualitative and quantitative aspects of distributed systems.

The thesis further applies resource-aware session types to the blockchain domain. Blockchains allow execution of complex protocols between mutually distrusting parties through *smart contracts*. Programming smart contracts comes with unique challenges such as enforcing transaction protocols, computing their execution cost, and ensuring that assets are not accidentally duplicated or discarded. This thesis presents *Nomos*: a language for smart contracts based on resource-aware session types. Session types statically express contract protocols. Resource-aware types automatically infer the execution cost of transactions leveraging ideas from automatic amortized resource analysis. The built-in linear type system of session types provides a natural representation for assets ensuring that they are preserved across transactions. The *Nomos* type checker statically enforces the above requirements: protocols are enforced at runtime, bounds inferred are sound and precise, and assets used are neither duplicated nor discarded. *Nomos* also significantly develops the theory of programming languages: integrating session types with functional programming, linear-time type checking to prevent denial-of-service attacks, and an acquire-release discipline to rule out re-entrancy attacks. The thesis concludes with two complementary future directions: enhancing *Nomos* with refinement types for lightweight verification of smart contracts, and applying *Nomos* to other distributed domains e.g. cryptographic protocols.

Acknowledgments

Although I am its author, this dissertation is really a joint effort. First and foremost, I would like to thank my advisor Jan Hoffmann for his unwavering support. Jan, thank you for your patience with me while I learned the basics of programming language theory. Your vision, not just about PL, but about success in research will always stay with me. And a special thanks for your constant guidance during my job application process. I also want to express my sincere gratitude to Frank Pfening who, despite not being my official advisor, always treated me like his student. Frank, your vast experience and clarity of thought has taught me so much about PL research. I really enjoyed my collaborations and pair programming experiences with you; we have created some really impressive software together!

Before anyone else, I would like to take this opportunity to thank my partner Dr. Shrimai Prabhume. Without you, I would not have been able to successfully complete my PhD. The life principles I have learned from you have made me, not just a better researcher, but a better person.

I would also like to thank my committee members: Andrew Miller for introducing me to the world of security, privacy, and applied cryptography; Bryan Parno for providing amazing feedback on my application materials and dissertation; and especially Shaz Qadeer for enabling an awesome internship experience and teaching me how to become an independent researcher. I'd also like to thank all my collaborators: Henry DeYoung, Andreia Mordido, Di Wang, Stephanie Balzer, Surya Bakshi, and Shreyas Gandlur. And also to Ishani Santurkar and Stephen McIntosh for giving me the privilege to mentor them. I have thoroughly enjoyed each and every project I have worked on with you all.

Finally, I would like to thank friends and family for being by my side throughout this challenging journey! Shout out to Shruti Palaskar, Dheeraj Rajagopal, Vidhisha Balachandran, Venkat Prasath Perumal, Chaitanya Ahuja, Bhavya Balu, Priyank Lathwal, Samarth Gupta, and Anup Kumar. This would not have been possible without you!

Contents

1	Introduction	1
1.1	Programming Digital Contracts using Resource-Aware Session Types	3
1.2	Overview	5
2	Background	7
2.1	Session Types	7
2.1.1	Examples	11
2.1.2	Preservation and Progress	12
2.2	Resource Analysis	13
2.2.1	Manual Amortized Analysis	14
2.2.2	Automatic Amortized Analysis	14
I	Theory of Resource-Aware Session Types	17
3	Refinement Session Types	19
3.1	Introduction	19
3.2	Arithmetic Refinements	21
3.3	Undecidability of Type Equality	24
3.4	A Practical Algorithm for Type Equality	26
3.4.1	Soundness of the Type Equality Algorithm	28
3.4.2	Type Equality Declarations	29
3.5	Formal Description of the Rast Language	30
3.5.1	The Refinement Layer	32
3.6	Type Safety	34
3.7	Constraint Reconstruction	35
3.8	Implementation	39
3.9	Further Examples	41
3.10	Related Work	45
3.11	Conclusion	46
4	Work Analysis	47
4.1	Introduction	47
4.2	Overview	48

4.3	Operational Cost Semantics	52
4.4	Type System	55
4.5	Soundness	57
4.6	Case Study: Stacks and Queues	58
4.7	Related Work	60
4.8	Future Directions	61
5	Time Analysis	63
5.1	Introduction	63
5.2	The Temporal Modality Next ($\circ A$)	66
5.3	The Temporal Modalities Always ($\square A$) and Eventually ($\diamond A$)	69
5.4	Preservation and Progress	74
5.5	Further Examples	76
5.6	Related Work	80
5.7	Future Directions	82
II	Application of Resource-Aware Session Types	83
6	Programming Digital Contracts	85
6.1	Introduction	85
6.2	Nomos by Example	87
6.3	Sharing Contracts	90
6.4	Adding a Functional Layer	93
6.5	Tracking Resource Usage	96
6.6	Type Soundness	100
7	Implementation of Nomos	103
7.1	Overview of Nomos with an Auction Implementation	103
7.2	Mode Inference	105
7.3	Implementation	106
7.4	Blockchain Integration	110
	7.4.1 Blockchain-specific Features	113
	7.4.2 Exact Gas Computation (Facebook Internship)	115
7.5	Web Interface	116
7.6	Evaluation	118
7.7	Related Work	122
7.8	Future Directions	124
8	Conclusion	125
8.1	Summary of Contributions	125
8.2	Future Directions	126
	Bibliography	129

Chapter 1

Introduction

The design of safe, efficient, and secure software has always been a challenging task. With the proliferation of distributed systems, software development has become even more complex. In addition to the usual challenges, developers must also carefully

- ensure that no bugs manifest in any possible thread interleaving,
- avoid deadlocks and data races, and
- prevent performance bottlenecks in the system.

The rise of modern distributed systems such as server farms, cloud computing platforms, and blockchains have further complicated software design by introducing unique challenges. For instance, blockchain is a highly adversarial domain due to its transparency. The state of every decentralized application deployed on the blockchain can be publicly viewed and potentially exploited by malicious attackers. Security vulnerabilities in these applications have caused losses to the tune of several billions of dollars.

Advances in the design and principles of programming languages have significantly benefited the design of software, improving its performance, safety and security. Such advances have enabled programmers to analyze the *qualitative* aspects of software using verification techniques as well as the *quantitative* aspects using complexity analysis tools. Today, developers can utilize automatic verification tools such as Boogie [28] and Dafny [109], or program assistants such as Coq [29] and Isabelle [128] to express and verify that programs satisfy their specifications. They can exploit tools such as Fiat Crypto [67] and EverCrypt [134] to generate secure high-performance cryptographic functionalities. They can employ fully formally verified compilers such as CompCert [110] to compile these programs to efficient low-level machine code. Finally, they can take advantage of static analysis tools such as RaML [91, 92] and SPEED [84] to compute complexity bounds and Infer [44] to detect performance bugs.

Although analysis tools exist for distributed systems, most of the industrial strength support is still centered around traditional sequential programming languages. Programs implemented in sequential languages are easier to analyze and verify since all computation typically occurs on a single processor. However, with the advent of multi-processor systems, most of the industrial software systems today are no longer sequential. Distributed algorithms lie at the heart of commitment protocols, distributed consensus, leader election and broadcast mechanisms with a wide array of applications ranging from communication networks and database management

to industrial control systems, cloud computing and blockchains.

With the advent of such distributed algorithms and systems, there is an urgent need of formal tools to assist with the design, verification and quantitative analysis of distributed software. However, analysis of concurrent and distributed systems poses additional challenges. First, while implementing concurrent programs, developers must consider all possible thread interleavings and ensure that no bug manifests in any such interleaving. Second, concurrent programs are notoriously prone to deadlocks (where each process in a group is waiting on a resource held by another process), and data races (where one thread is writing to a shared resource that is simultaneously being read by another thread). Finally, unlike sequential programs, a compositional reasoning for concurrent programs is especially challenging since their behavior ultimately depends on interactions within the system.

In response, this thesis has designed novel resource-aware session types [61, 62] that serve as a sound and practical foundation for distributed systems with strong type-theoretic guarantees. A common theme in a distributed system is communication between its components. And often, the type, value, and direction of communication depends on the state of the system. Session types leverage this property by capturing the system state in the type, providing a structured way of prescribing communication protocols. Session types can also guarantee freedom from deadlocks and data races. They also possess a confluence property stating that a communicating system converges to the same final state under all possible thread interleavings. All these properties greatly benefit programmers and automatically prevent a large class of bugs that can result from communication mismatches.

However, simple session types cannot express quantitative properties of a distributed system, such as energy consumption, latency, response time, and throughput. Realizing this, this thesis proposes two extensions to express the *work* [61] and *span* [62] of parallel computation. Work is defined as the *total number of operations* executed as part of the computation. However, due to parallelism in the system, many of these operations execute simultaneously. Therefore, span is defined as the total *time* of computation assuming arbitrarily many processors, thus taking the maximum parallelism in the system into account.

But, work and span bounds often depend on the intrinsic sizes and values of data structures, e.g., size of a list, value of a bit, or height of a tree. Simple session types also do not have a mechanism to express sizes and values, thus limiting the expressivity of our resource-aware type systems to constant amortized bounds. Since our type system is general enough to handle higher-degree polynomial bounds, it is imperative to statically express the above intrinsic quantities. Our solution is to introduce *refinement session types* which index types with a natural number quantity. Resource annotations can then be expressed as a function of these refinements, enabling us to express more general bounds. We also introduce novel *type constraints* from Presburger arithmetic as type constructors allowing us to express lightweight properties about such systems. Thus, as a precursor, we develop the metatheory and implementation of these refinements and describe how resource-aware types employ them.

Work analysis is based on a linear type system that combines standard session types with type-based amortized analysis. We introduce two novel dual type constructors, \triangleright^r and \triangleleft^r to send and receive r *potential* units. This potential (in the sense of classical amortized analysis [149]) may either be spent by sending other messages in the process network, or stored in a process for future interactions. Since the process interface is characterized entirely by the resource-aware

session type of the channels it interacts with, this design provides a compositional resource specification. A conceptual challenge is to express symbolic bounds in a setting without static data structures and intrinsic sizes. Our innovation is that resource-aware session types describe bounds as functions of interactions on a channel. As a result, the type system derives parametric bounds on the resource usage of message-passing processes. Finally, a type safety theorem proves that the derived bounds are sound with respect to an operational cost semantics that tracks the total number of messages exchanged in a network of communicating processes.

In addition to work, the timing of messages is of central interest for analyzing parallel cost. We developed a type system that captures the parallel complexity of session-typed programs by adding *temporal modalities next* ($\circ A$), *always* ($\square A$) and *eventually* ($\diamond A$), interpreted over a linear model of time. When considered as types, the temporal modalities allow us to express properties of concurrent programs such as the *message rate* of a stream, the *latency* of a pipeline, the *response time* of a concurrent data structure, or the *span* of a fork-join parallel computation, all in the same uniform manner. The circle operator (\circ) expresses precision of message timings, while the box (\square) and diamond operators (\diamond) provide flexibility, allowing us to express the timing of a wide variety of standard session-typed programs. Finally, a type safety theorem establishes that the message timing expressed by the type system are realized by the timed operational semantics.

Resource-aware session types combine session types with work and span extensions allowing programmers to reason about both qualitative and quantitative aspects of distributed systems.

1.1 Programming Digital Contracts using Resource-Aware Session Types

Digital contracts are computer protocols that describe and enforce the execution of a contract. With the rise of blockchains and cryptocurrencies such as Bitcoin [122], Ethereum [161], and Tezos [79], digital contracts have become popular in the form of *smart contracts*, which provide potentially distrusting parties with programmable money and an enforcement mechanism that does not rely on third parties. Smart contracts have been used to implement auctions [1], investment instruments [121], insurance agreements [98], supply chain management [108], and mortgage loans [120]. In general, digital contracts hold the promise to reduce friction, lower cost, and broaden access to financial infrastructure.

Smart contracts are implemented using a high-level programming language such as Solidity [54], Rholang [6], and Liquidity [4]. It is then compiled down to bytecode and executed using a runtime environment (e.g. Ethereum Virtual Machine for the Ethereum blockchain). However, these languages have significant shortcomings as they do not accommodate the domain-specific requirements of digital contracts.

- Instead of centering contracts on their interactions with users, the high-level protocol of the intended interactions with a contract is buried in the implementation code, hampering understanding, formal reasoning, and trust.
- Resource (or *gas*) usage of digital contracts is of particular importance for transparency and consensus. However, obliviousness of resource usage in existing contract languages

makes it hard to predict the cost of executing a contract and prevent denial-of-service vulnerabilities.

- Existing languages fail to enforce linearity of assets, endangering the validity of a contract when assets are duplicated or deleted, accidentally or maliciously [118].

Such limitations of contract languages can lead to vulnerabilities in the implemented contracts which can be exploited by malicious users having direct financial consequences. A well-known example is the attack on The DAO [121], resulting in a multi-million dollar theft by exploiting a contract vulnerability. Maybe even more important is the potential erosion of trust as a result of such failures.

To address these challenges, I have designed the type-theoretic foundations of Nomos [56], a programming language whose genetics match directly with the domain-specific requirements to provide strong static guarantees that facilitate the design of correct contracts. To express and enforce the protocols underlying a contract, we base Nomos on resource-aware session types. Type checking is automatic and guarantees that Nomos programs follow the communication protocol expressed by the type.

Resource-aware types also make transaction cost in Nomos predictable and transparent, and prevents bugs resulting from excessive resource usage. Since resource-aware session types are parametric in the cost model, they can be instantiated to derive the gas bounds for Nomos programs. The type soundness theorem for Nomos guarantees that these bounds are both sound and precise. Other advantages of this type-based resource analysis are natural compositionality and reduction of bound inference to off-the-shelf LP solving.

To eliminate a class of bugs in which the internal state of a contract loses track of its assets, Nomos integrates a linear type system into a functional language. Linear type systems use the ideas of Girard’s linear logic [76] to represent certain resources ensuring they are never discarded or duplicated. Assets such as money and other commodities that can be exchanged between parties in a contract are typed using a linear channel. Type safety guarantees that processes maintain proper ownership of linear assets and do not terminate while holding access to a linear asset.

Since there exist multiple clients of a contract, we use a *shared* session type [25] to define the protocol of a contract. This ensures that clients interact with a contract in *mutual exclusion*. The type clearly demarcates the parts of the protocol that become a critical section using \uparrow_L^S modality marking its start and \downarrow_L^S modality marking its end. Programmatically, \uparrow_L^S translates into an acquire of the contract, while \downarrow_L^S into its release.

We complement the theory of Nomos with a stable open-source implementation [57] in OCaml. This implementation consists of a lexer, parser, type checker, inference engine, and interpreter. Thus, a programmer can implement contracts in the Nomos source language, verify that they indeed satisfy their session-typed protocol using the type checker, automatically compute its execution cost using the inference engine, and run transactions and observe how they modify the blockchain state using the interpreter.

To improve its usability, we enhance Nomos along two directions: *simplicity of programming* and *efficiency*. Crucial to type safety of Nomos are channel modes that describe the role of a process and the channels that it can depend on. However, computing and tagging each channel with its mode puts enormous burden on the programmer. Similarly, computing the

exact gas cost of transactions involves solving linear constraints which is also manually infeasible. To reduce this burden, we have designed an efficient inference engine that statically and automatically computes channel modes and gas bounds with minimal programmer interaction. To further enhance usability, we introduce (i) a built-in map data structure accompanied with syntactic sugar for ease of use which can be flexibly used both linearly to store assets or non-linearly, (ii) an abstract linear coin type to easily create Non-Fungible Tokens (NFTs), (iii) blockchain-specific expressions to obtain the transaction sender and number, and (iv) a bi-directional type checker with precise error messages.

To improve efficiency, we have carefully implemented a bi-directional Nomos type checker that is *linear-time* in the size of the program. In a blockchain setting, transaction validation which entails type checking is part of the attack surface. Thus, an inefficient type checker can be exploited by adversaries who can submit programs that are too slow to validate causing denial-of-service. We also introduce a novel `Nomos.deposit{r}` construct [60] that deposits r gas units in the transaction sender’s account. This helps us *equalize* the static gas cost in program branches, thus enabling the inference engine to produce an exact gas bound. At runtime, the `Nomos.deposit{r}` operations on the execution path safely return the leftover gas back to the user, thus eliminating the need for dynamic cost monitoring. This improves the overall hygiene and efficiency of the Nomos interpreter.

In addition, we describe how Nomos is integrated with an account model blockchain like Ethereum [161]. We describe the challenges in this integration, and highlight the main limitation of the language. We formalize how to represent blockchain states as session-typed configurations which can be manipulated by transaction programs. These configurations are conveniently stored in a file using OCaml S-expressions making them persistent. We conclude by evaluating Nomos on a variety of standard smart contracts, emphasizing the guarantees provided by Nomos and their advantages.

1.2 Overview

Thesis Statement *Resource-aware session types serve as a sound and practical type-theoretic foundation for digital contracts enforcing strong domain-specific guarantees while simplifying programming and without compromising efficiency.*

Chapter 2 sets up the necessary background for this thesis. It first provides the theory of basic session types, its type constructors, formal type systems, and semantics. Then, it explains the techniques of type-based automatic amortized resource analysis (AARA) in the context of a functional programming language. Chapter 3 describes our novel extension to refinement session types. We also discuss the problem of type equality, which remarkably, is undecidable, even though Presburger arithmetic itself is decidable. Chapter 4 employs refinement types to define potential type constructors and uses them to upper bound the work performed by the system. Chapter 5 employs refinements to define temporal annotations which are then used to express span bounds. Chapter 6 presents the design and type-theoretic foundation of the Nomos language which is proved to satisfy type preservation (session fidelity) and a limited form of progress (deadlock freedom). Chapter 7 describes its implementation and our efforts to enhance usability and efficiency. Finally, Chapter 8 concludes with broader applicability of

Nomos and resource-aware session types.

Chapter 2

Background

This chapter introduces two central concepts used in this thesis, namely session types and resource analysis. The first half of the chapter focuses on session types, defining them formally with static and dynamic semantics. The second half introduces type-based amortized analysis for functional programming languages. Resource-aware session types formalize the combination of these two techniques.

2.1 Session Types

Session types are a type discipline for communication-centric programming based on message passing via channels. Session-typed channels describe and enforce the protocol of communication among processes. The base system of session types is derived from a Curry-Howard interpretation of intuitionistic linear logic [42]. This chapter focuses on the linear fragment of SILL [129] that internalizes session-based concurrency. Session types were introduced by Honda [96].

Linear logic [76] is a substructural logic that enjoys *exchange* as its only structural property, i.e., it does not exhibit *weakening* or *contraction*. As a result, purely linear propositions can be viewed as resources that must be used *exactly once* in a proof. Here, I adopt the intuitionistic version of linear logic, yielding the following sequent

$$A_1, \dots, A_n \vdash C$$

where A_1, \dots, A_n are linear antecedents, while C is the linear succedent.

Under the Curry-Howard isomorphism for intuitionistic linear logic, propositions are related to session types, proofs to processes and cut reduction in proofs to communication. Appealing to this correspondence, a process term P is assigned to the above judgment and each hypothesis as well as the conclusion is labeled with a *channel*:

$$x_1 : A_1, \dots, x_n : A_n \vdash P :: (z : C)$$

The resulting judgment states that process P *provides* a service of session type C along channel z , *using* the services of session types A_1, \dots, A_n provided along channels x_1, \dots, x_n respectively. The assignment of a channel to the conclusion is convenient because, unlike functions,

Type	Provider Action	Session Continuation
$\oplus\{\ell : A_\ell\}_{\ell \in L}$	send label $k \in L$	A_k
$\&\{\ell : A_\ell\}_{\ell \in L}$	receive and branch on label $k \in L$	A_k
$\mathbf{1}$	send token close	<i>none</i>
$A \otimes B$	send channel $c : A$	B
$A \multimap B$	receive channel $c : A$	B

Table 2.1: Basic Session Types. Every provider action has a matching client action.

processes do not evaluate to a value but continue to communicate along their providing channel once they have been created until they terminate. For the judgment to be well-formed, all channel names have to be distinct. The antecedents are often abbreviated to Δ .

The balance between providing and using a session is established by the two fundamental rules of the sequent calculus that are independent of all logical connectives: *cut* and *identity*. Cut states that if P provides service A along channel x , then Q can use the service along the same channel at the same type. Identity states that a client of service A can directly provide A .

$$\frac{\Delta_1 \vdash P_x :: (x : A) \quad \Delta_2, x : A \vdash Q_x :: (z : C)}{\Delta_1, \Delta_2 \vdash x \leftarrow P_x ; Q_x :: (z : C)} \text{ cut} \quad \frac{}{y : A \vdash x \leftrightarrow y :: (x : A)} \text{ id}$$

Operationally, the process $x \leftarrow P_x ; Q_x$ creates a globally fresh channel c , spawns a new process $[c/x]P_x$ providing along c , and continues as $[c/x]Q_x$. Conversely, the process $c \leftrightarrow d$ forwards any message M that arrives along d to c and vice-versa. Because channels are used linearly, the forwarding process can then terminate, making sure to apply proper renaming.

The operational semantics are formalized as a system of *multiset rewriting rules* [48]. I introduce semantic objects $\text{proc}(c, P)$ and $\text{msg}(c, M)$ describing process P (or message M) providing service along channel c . Remarkably, in this formulation, a message is just a particular form of process, thereby not requiring any special rules for typing; it can be typed just as processes. The semantics rules for cut and id are presented below.

$$\begin{aligned} (\text{cut}C) \quad & \text{proc}(d, x \leftarrow P_x ; Q_x) \mapsto \text{proc}(c, [c/x]P_x), \text{proc}(d, [c/x]Q_x) \quad (c \text{ fresh}) \\ (\text{id}^+C) \quad & \text{msg}(d, M), \text{proc}(c, c \leftrightarrow d) \mapsto \text{msg}(c, [c/d]M) \\ (\text{id}^-C) \quad & \text{proc}(c, c \leftrightarrow d), \text{msg}(e, M(c)) \mapsto \text{msg}(e, [d/c]M(c)) \end{aligned}$$

Here, I adopt the convention to use x, y and z for channel *variables* and c, d and e for *channels*. Channels are created at runtime and substituted for channel variables in process terms. In the last rule, $M(c)$ indicates that c must occur in M , implying it is the sole client of c .

The Curry-Howard correspondence gives each linear logic connective an interpretation as a session type. This session type prescribes the kind of message that must be sent or received along a channel of this type. Table 2.1 summarizes the description of the type along with the provider action. I follow a detailed description of each session type operator.

Internal Choice A type A is said to describe a *session*, which is a particular sequence of interactions. As a first type construct, consider *internal choice* $\oplus\{\ell : A_\ell\}_{\ell \in L}$, an n -ary labeled

generalization of the linear logic connective $A \oplus B$. A process that provides $x : \oplus\{\ell : A_\ell\}_{\ell \in L}$ can send any label $k \in L$ along x and then continue by providing $x : A_k$. The corresponding process is written as $(x.k ; P)$, where P is the continuation that provides A_k . This typing is formalized by the *right rule* $\oplus R$ in linear sequent calculus. The corresponding client branches on the label received along x as specified by the *left rule* $\oplus L$.

$$\frac{(k \in L) \quad \Delta \vdash P :: (x : A_k)}{\Delta \vdash (x.k ; P) :: (x : \oplus\{\ell : A_\ell\}_{\ell \in L})} \oplus R$$

$$\frac{(\forall \ell \in L) \quad \Delta, (x : A_\ell) \vdash Q_\ell :: (z : C)}{\Delta, (x : \oplus\{\ell : A_\ell\}_{\ell \in L}) \vdash \text{case } x (\ell \Rightarrow Q_\ell)_{\ell \in L} :: (z : C)} \oplus L$$

Operationally, since communication is asynchronous, the process $(c.k ; P)$ sends a message k along c and continues as P without waiting for it to be received. As a technical device to ensure that consecutive messages on a channel arrive in order, the sender also creates a fresh continuation channel c' so that the message k is actually represented as $(c.k ; c \leftrightarrow c')$ (read: send k along c and continue as c'). The provider also substitutes c' for c , enforcing that the next message is sent on c' .

$$(\oplus S) \quad \text{proc}(c, c.k ; P) \mapsto \text{proc}(c', [c'/c]P), \text{msg}(c, c.k ; c \leftrightarrow c') \quad (c' \text{ fresh})$$

When the message k is received along c , the client selects branch k and also substitutes the continuation channel c' for c , thereby ensuring that it receives the next message on c' . This implicit substitution of the continuation channel ensures the ordering of the messages.

$$(\oplus C) \quad \text{msg}(c, c.k ; c \leftrightarrow c'), \text{proc}(d, \text{case } c (\ell \Rightarrow Q_\ell)_{\ell \in L}) \mapsto \text{proc}(d, [c'/c]Q_k)$$

External Choice The dual of internal choice is *external choice* $\&\{\ell : A_\ell\}_{\ell \in L}$, which is the n -ary labeled generalization of the linear logic connective $A \& B$. This dual operator simply reverses the role of the provider and client. The provider process of $x : \&\{\ell : A_\ell\}_{\ell \in L}$ branches on receiving a label $k \in L$ (described in $\&R$), while the client sends this label (described in $\&L$).

$$\frac{(\forall \ell \in L) \quad \Delta \vdash P_\ell :: (x : A_\ell)}{\Delta \vdash \text{case } x (\ell \Rightarrow P_\ell)_{\ell \in L} :: (x : \&\{\ell : A_\ell\}_{\ell \in L})} \&R$$

$$\frac{\Delta, (x : A_k) \vdash Q :: (z : C)}{\Delta, (x : \&\{\ell : A_\ell\}_{\ell \in L}) \vdash x.k ; Q :: (z : C)} \&L$$

The operational semantics rules are just the inverse of internal choice. The provider receives the branching label k sent by the provider. Both processes perform appropriate substitutions to ensure the order of messages sent and received is preserved.

$$(\& S) \quad \text{proc}(d, c.k ; Q) \mapsto \text{msg}(c', c.k ; c' \leftrightarrow c), \text{proc}(d, [c'/c]Q) \quad (c' \text{ fresh})$$

$$(\& C) \quad \text{proc}(c, \text{case } c (\ell \Rightarrow Q_\ell)_{\ell \in L}), \text{msg}(c', c.k ; c' \leftrightarrow c) \mapsto \text{proc}(c', [c'/c]Q_k)$$

Higher-Order Channels Session types allow channels to be *higher-order*, i.e., channels can be exchanged over channels. The session type corresponding to the linear logic connective $A \otimes B$ allows its provider to send a channel of type A and then continue with providing B . The corresponding process term $(\text{send } x \ w ; P)$ describes sending channel w over channel x and continuing with P . This typing is provided by the rule $\otimes R$. The client, on the other hand, receives this channel using the term $(y \leftarrow \text{recv } x ; Q)$ and binds it to a channel variable y , as described by $\otimes L$.

$$\frac{\Delta \vdash P :: (x : B)}{\Delta, (w : A) \vdash (\text{send } x \ w ; P) :: (x : A \otimes B)} \otimes R$$

$$\frac{\Delta, (y : A), (x : B) \vdash Q :: (z : C)}{\Delta, (x : A \otimes B) \vdash (y \leftarrow \text{recv } x ; Q) :: (z : C)} \otimes L$$

$$(\otimes S) \quad \text{proc}(c, \text{send } c \ e ; P) \mapsto \text{proc}(c', [c'/c]P), \text{msg}(c, \text{send } c \ e ; c \leftrightarrow c') \quad (c' \text{ fresh})$$

$$(\otimes C) \quad \text{msg}(c, \text{send } c \ e ; c \leftrightarrow c'), \text{proc}(d, x \leftarrow \text{recv } c ; Q) \mapsto \text{proc}(d, [c', e/c, x]Q)$$

The lolti (\multimap) operator is dual to \otimes . The provider and client invert their roles, i.e., the provider of $x : A \multimap B$ receives a channel of type A sent by its client.

$$\frac{\Delta, (y : A) \vdash P :: (x : B)}{\Delta \vdash (y \leftarrow \text{recv } x ; P) :: (x : A \multimap B)} \multimap R$$

$$\frac{\Delta, (x : B) \vdash Q :: (z : C)}{\Delta, (x : A \multimap B), (y : A) \vdash (\text{send } x \ w ; Q) :: (z : C)} \multimap L$$

$$(\multimap S) \quad \text{proc}(d, \text{send } c \ e ; Q) \mapsto \text{msg}(c', \text{send } c \ e ; c' \leftrightarrow c), \text{proc}(d, [c'/c]Q) \quad (c' \text{ fresh})$$

$$(\multimap C) \quad \text{proc}(c, x \leftarrow \text{recv } c), \text{msg}(c', \text{send } c \ e ; c' \leftrightarrow c) \mapsto \text{proc}(c', [c', d/c, x]P)$$

Termination The type $\mathbf{1}$, the multiplicative unit of linear logic, represents termination of a process, which (due to linearity) is not allowed to use any channels.

$$\frac{}{\cdot \vdash \text{close } x :: (x : \mathbf{1})} \mathbf{1}R \qquad \frac{\Delta \vdash Q :: (z : C)}{\Delta, (x : \mathbf{1}) \vdash (\text{wait } x ; Q) :: (z : C)} \mathbf{1}L$$

Operationally, a client has to wait for the corresponding closing message, which has no continuation since the provider terminates.

$$(\mathbf{1}S) \quad \text{proc}(c, \text{close } c) \mapsto \text{msg}(c, \text{close } c)$$

$$(\mathbf{1}C) \quad \text{msg}(c, \text{close } c), \text{proc}(d, \text{wait } c ; Q) \mapsto \text{proc}(d, Q)$$

Process Definitions Process definitions have the form $\Delta \vdash f = P :: (x : A)$ where f is the name of the process and P its definition. All definitions are collected in a fixed global signature Σ . Also, since process definitions are mutually recursive, it is required that every process in the signature is well-typed w.r.t. Σ , i.e. $\Sigma ; \Delta \vdash P :: (x : A)$. For readability of the examples,

I break a definition into two declarations, one providing the type and the other the process definition binding the variables x and those in Ω (generally omitting their types):

$$\begin{aligned} \Delta \vdash f &:: (x : A) \\ x \leftarrow f \Delta &= P \end{aligned}$$

A new instance of a defined process f can be spawned with the expression

$$x \leftarrow f \bar{y} ; Q$$

where \bar{y} is a sequence of variables matching the antecedents Δ . The newly spawned process will use all variables in \bar{y} and provide x to the continuation Q . The operational semantics reduces the spawn to a cut.

$$(\text{def}C) \quad \text{proc}(c, x \leftarrow f \leftarrow \bar{e} ; Q) \mapsto \text{proc}(a, [a/x, \bar{e}/\Delta]P), \text{proc}(c, [a/x]Q) \quad (a \text{ fresh})$$

where $\Delta \vdash f = P :: (x : A) \in \Sigma$. Here I write \bar{e}/Δ to denote substitution of the channels in \bar{e} for the corresponding variables in Δ .

Sometimes a process invocation is a *tail call*, written without a continuation as $x \leftarrow f \bar{y}$. This is a short-hand for $x' \leftarrow f \bar{y} ; x \leftrightarrow x'$ for a fresh variable x' , that is, a fresh channel is created and immediately identified with x (although it is generally implemented more efficiently).

Recursive Types Session types can be naturally extended to include recursive types. For this purpose I allow (possibly mutually recursive) type definitions $X = A$ in the signature, where I require A to be *contractive* [73]. This means here that A should not itself be a type name. The type definitions are *equi-recursive* so X can be silently replaced by A during type checking, and no explicit rules for recursive types are needed.

2.1.1 Examples

As a first example, consider a stream of bits defined recursively as

$$\text{bits} = \oplus\{\text{b0} : \text{bits}, \text{b1} : \text{bits}, \$: 1\}$$

When considering bits as representing natural numbers, the least significant bit is sent first. For example, a process *six* sending the number $6 = (110)_2$ would be

$$\cdot \vdash \text{six} :: (x : \text{bits})$$

$$x \leftarrow \text{six} = x.\text{b0} ; x.\text{b1} ; x.\text{b1} ; x.\$; \text{close } x$$

Executing $\text{proc}(c_0, c_0 \leftarrow \text{six})$ yields (with some fresh channels c_1, \dots, c_4)

$$\begin{aligned} \text{proc}(c_0, c_0 \leftarrow \text{six}) \mapsto^* & \text{msg}(c_4, \text{close } c_4), \\ & \text{msg}(c_3, c_3.\$; c_3 \leftrightarrow c_4), \\ & \text{msg}(c_2, c_2.\text{b1} ; c_2 \leftrightarrow c_3), \\ & \text{msg}(c_1, c_1.\text{b1} ; c_1 \leftrightarrow c_2), \\ & \text{msg}(c_0, c_0.\text{b0} ; c_0 \leftrightarrow c_1), \end{aligned}$$

As a first example of a recursive process definition, consider one that just copies the incoming bits on to the outgoing bits.

```

y : bits ⊢ copy :: (x : bits)
x ← copy y =
  case y (b0 ⇒ x.b0 ; x ← copy y   % received b0 on y, send b0 on x, recurse
        | b1 ⇒ x.b1 ; x ← copy y   % received b1 on y, send b1 on x, recurse
        | $ ⇒ x.$ ; wait y ; close x) % received $ on y, send $ on x, wait on y, close x

```

Note the occurrence of a (recursive) *tail call* to *copy*.

A last example: to increment a bit stream turn b0 to b1 but then forward the remaining bits unchanged ($x \leftrightarrow y$), or turn b1 to b0 but then increment the remaining stream ($x \leftarrow \text{plus1 } y$) to capture the effect of the carry bit.

```

y : bits ⊢ plus1 :: (x : bits)
x ← plus1 y =
  case y (b0 ⇒ x.b1 ; x ↔ y
        | b1 ⇒ x.b0 ; x ← plus1 y
        | $ ⇒ x.$ ; wait y ; close x)

```

2.1.2 Preservation and Progress

The main theorems that exhibit the deep connection between our type system and the operational semantics are the usual *type preservation* and *progress*, sometimes called *session fidelity* and *deadlock freedom*, respectively.

So far, I have only described individual processes. However, processes exist in a *configuration*. A process configuration is a multiset of semantic objects, $\text{proc}(c, P)$ and $\text{msg}(c, M)$, where any two offered channels are distinct. A key question is how to type these configurations. Since they consist of both processes and messages, they both *use* and *provide* a collection of channels. And even though a configuration is treated as a multiset, typing imposes a partial order on the processes and messages where a provider of a channel appears to the left of its client.

A configuration is typed w.r.t. a signature providing the type declaration of each process. A signature Σ is *well formed* if (a) every type definition $V = A_V$ is *contractive*, and (b) every process definition $\Delta \vdash f = P :: (x : A)$ in Σ is well typed according to the process typing judgment, i.e. $\Sigma ; \Delta \vdash P :: (x : A)$.

I use the following judgment to type a configuration.

$$\Sigma ; \Delta_1 \vDash \mathcal{S} :: \Delta_2$$

It states that Σ is well-formed and that the configuration \mathcal{S} uses the channels in the context Δ_1 and provides the channels in the context Δ_2 . The configuration typing judgment is defined using the rules presented in Figure 2.1. The rule *empty* defines that an empty configuration is well-typed. The rule *compose* composes two configurations \mathcal{S}_1 and \mathcal{S}_2 ; \mathcal{S}_1 provides service on the channels in Δ_1 while \mathcal{S}_2 uses the channels in Δ_2 . The *proc* rule creates a configuration out of a single process. Similarly, the *msg* rule creates a configuration out of a single message.

Theorem 1 (Type Preservation). *If $\Sigma ; \Delta' \vDash \mathcal{S} :: \Delta$ and $\mathcal{S} \mapsto \mathcal{S}'$, then $\Sigma ; \Delta' \vDash \mathcal{S}' :: \Delta$.*

Proof. By case analysis on the transition rule, applying inversion to the given typing derivation, and then assembling a new derivation of \mathcal{S}' . \square

$$\begin{array}{c}
\frac{}{\Sigma ; \Delta \vDash (\cdot) :: \Delta} \text{ empty} \qquad \frac{\Sigma ; \Delta_0 \vDash \mathcal{S}_1 :: \Delta_1 \quad \Sigma ; \Delta_1 \vDash \mathcal{S}_2 :: \Delta_2}{\Sigma ; \Delta_0 \vDash (\mathcal{S}_1 \mathcal{S}_2) :: \Delta_2} \text{ compose} \\
\\
\frac{\Sigma ; \Delta_1 \vdash P :: (c : A)}{\Sigma ; \Delta, \Delta_1 \vDash \text{proc}(c, P) :: (\Delta, (c : A))} \text{ proc} \qquad \frac{\Sigma ; \Delta_1 \vDash P :: (c : A)}{\Sigma ; \Delta, \Delta_1 \vDash \text{msg}(c, P) :: (\Delta, (c : A))} \text{ msg}
\end{array}$$

Figure 2.1: Typing rules for a configuration

A process or message is said to be *poised* if it is trying to communicate along the channel that it provides. A poised process is comparable to a value in a sequential language. A configuration is poised if every process or message in the configuration is poised. Conceptually, this implies that the configuration is trying to communicate externally, i.e. along one of the channels it provides. The progress theorem then shows that either a configuration can take a step or it is poised. To prove this I show first that the typing derivation can be rearranged to go strictly from right to left and then proceed by induction over this particular derivation.

Theorem 2 (Global Progress). *If $\cdot \vDash \mathcal{S} :: \Delta$ then either*

- (i) $\mathcal{S} \mapsto \mathcal{S}'$ for some \mathcal{S}' , or
- (ii) \mathcal{S} is poised.

Proof. By induction on the right-to-left typing of \mathcal{S} so that either \mathcal{S} is empty (and therefore poised) or $\mathcal{S} = (\mathcal{S}' \text{ proc}(c, P))$ or $\mathcal{S} = (\mathcal{S}' \text{ msg}(c, M))$. By induction hypothesis, \mathcal{S}' can either take a step (and then so can \mathcal{S}), or \mathcal{S}' is poised. In the latter case, I analyze the cases for P and M , applying multiple steps of inversion to show that in each case either \mathcal{S} can take a step or is poised. \square

2.2 Resource Analysis

The quality of software crucially depends on the amount of resources – time, memory and energy – that are required for its execution. Statically understanding and controlling resource usage continues to be a central issue in software development. Recent years have seen fast progress in developing tools and frameworks for statically reasoning about resource usage. The obtained *size change* information forms the basis for the computation of actual bounds on loop iterations and recursion depths; using counter instrumentation [85], ranking functions [10, 18, 40, 148], recurrence relations [11, 12] and abstract interpretation [47, 166]. Automatic resource analysis for functional programs are based on sized types [156], term-rewriting [23] and amortized resource analysis [91, 93, 100, 147].

Automatic amortized resource (AARA) was introduced for a strict first-order functional language with built-in data types [93]. Since then, AARA techniques have been applied to univariate polynomial bounds [89], multivariate bounds [91], higher-order functional programs [100] and user-defined data types [92]. In this section, I will mainly focus on linear bounds for first-order programs as it outlines the main ideas of AARA without complicating the type system.

2.2.1 Manual Amortized Analysis

Often the cost of an operation on a data structure depends on its state. Thus, it is natural to account for the total cost of a sequence of operations on such a data structure. To analyze such a sequence of operations, Sleator and Tarjan [149] proposed amortized analysis with the *potential method*.

The concept of potential is inspired by the notion of potential energy in physics. The idea is to define a potential function $\Phi : \mathcal{D} \rightarrow \mathbb{R}^{\geq 0}$ that maps data structure $D \in \mathcal{D}$ to a non-negative number. Operations on the data structure can then *increase* or *decrease* the potential. The amortized cost of an operation $\text{op}(D)$ is then defined as the sum of its actual cost K and the difference of the potential caused by op , i.e., $K + \Phi(\text{op}(D)) - \Phi(D)$. The sum of the amortized costs over a sequence of operations and the initial potential of D then furnishes an upper bound on the actual cost of the sequence.

A standard example that demonstrates the benefits of amortization is the analysis of a functional queue, represented as two lists L_{in} and L_{out} . Enqueuing an element simply adds it to the head of L_{in} , while dequeuing removes the element from the head of L_{out} . If L_{out} is empty, the elements from L_{in} are transferred to L_{out} , thereby reversing the order of the elements.

The cost of a dequeue operation for this queue depends on the state of L_{out} , whether its empty or not. In the worst case, when L_{out} is empty, the cost of dequeue is linear. However, we can introduce a potential $\Phi(L_{in}, L_{out}) = 2|L_{in}|$. Then, the amortized cost of enqueue is 3 – one to pay for consing to L_{in} , and two for the increase in potential. More importantly, the amortized cost of dequeue is 1. If L_{out} is not empty, the cost of detach is 1, while there is no change to the potential. While if L_{out} is empty, the potential stored in L_{in} is used to pay for the cost of transfer from L_{in} to L_{out} . Formally, the cost of transfer is $2|L_{in}|$, equal to the change in the potential. Hence, the amortized cost of dequeue remains 1, used to pay for the detach from L_{out} after the transfer. Thus, amortized analysis proves that the worst-case cost of both enqueue and dequeue operations is constant.

2.2.2 Automatic Amortized Analysis

The potential method from amortized analysis can be applied to statically analyze functional programs. The key idea here is that the arguments of a function store potential, which is consumed during function evaluation. The initial potential of the arguments therefore equals the sum of the resource cost of the function, and the potential of the return value. Thus, it acts as an upper bound on the resource cost. Since the excess potential is stored in the result, this technique is completely compositional.

Automation is a key requirement here since requiring the programmer to provide the potential functions will significantly increase their burden. To make automation feasible, it is necessary to restrict the space of possible potential functions. There is a precision-scalability trade-off here since increasing the space of potential functions will improve the precision of resource bounds, but will make automation more challenging.

I will restrict the potential functions to be *linear* and the language to be strict, first-order and functional. I attach the potential of the data structure to its type. Then, a sound type checking algorithm statically verifies that the potential is sufficient to pay for all operations that are

performed on this data structure during any possible evaluation of the program. Consider the *append* function that takes two lists and appends the second list to the first. The function is implemented as follows.

```
let rec append l1 l2 =
  match l1 with
  | [] -> l2
  | x::xs -> let ys = append xs l2 in x::ys
```

To understand the resource usage for *append*, we first need to fix the resource we are interested in counting. For this example, suppose we count the number of cons (::) operations. The above code suggests that the number of cons operations equals the length of l_1 . To understand the type-based analysis, consider the following type for *append*.

$$\text{append} : L^1(A) \times L^0(A) \xrightarrow{0/0} L^0(A)$$

Intuitively, this describes that a unit potential is attached to every element in l_1 , and no potential on l_2 and the result. In the nil branch of the match, the context is assigned the type $l_1 : L^1(A), l_2 : L^0(A)$. Since l_1 is nil, the total potential of the context is 0, and l_2 is directly returned. In the cons branch, the context is typed as $x : A, xs : L^1(A), l_2 : L^0(A)$. Remarkably, the recursive call to *append* utilizes the same type, since xs and l_2 have the same type as described in the signature. After the recursive call, the context becomes $x : A, ys : L^0(A), l_2 : L^0(A)$, and the unit potential stored in x is used to perform the cons operation.

The type inference algorithm assigns variable potential annotations to *append*.

$$\text{append} : L^{q_1}(A) \times L^{q_2}(A) \xrightarrow{r/s} L^q(A)$$

The inference algorithm then derives linear constraints on the annotations. For *append*, the constraints generated are $q_1 \geq q_2 + 1$, and $q_2 = q$, and $r \geq s$. These constraints are solved with an off-the-shelf linear-programming solver (LP solver) whose goal is to minimize the initial potential to derive the most precise bound. The LP solver recovers the annotation values described earlier, thus proving the exact bound $|l_1|$ for *append*.

Part I

Theory of Resource-Aware Session Types

Chapter 3

Refinement Session Types

Traditional session types prescribe bidirectional communication protocols for concurrent computations, where well-typed programs are guaranteed to adhere to the protocols. However, simple session types cannot capture properties beyond the basic type of the exchanged messages. In response, this chapter indexes session types with refinements from linear arithmetic, capturing intrinsic attributes of processes and data. These refinements then play a central role in describing sequential and parallel complexity bounds on session-typed programs (Chapters 4 and 5).

This chapter describes the metatheory of such indexed types. We show that, despite the decidability of Presburger arithmetic, type equality and therefore also type checking are now undecidable, which stands in contrast to analogous dependent refinement type systems from functional languages. We also present a practical incomplete algorithm for type equality and an algorithm for type checking which is complete relative to an oracle for type equality. Process expressions in this explicit language are rather verbose, so we also introduce an implicit form and a sound and complete algorithm for reconstructing explicit programs, borrowing ideas from the proof-theoretic technique of focusing. All the aforementioned ideas have been implemented in an open-source language named Rast [63]. We conclude by illustrating our systems and algorithms with a variety of examples that have been verified in the Rast implementation.

3.1 Introduction

Basic session types have limited expressivity. As a simple example, consider the session type offered by a queue data structure storing elements of type A .

$$\text{queue}_A = \&\{\mathbf{ins} : A \multimap \text{queue}_A, \\ \mathbf{del} : \oplus\{\mathbf{none} : \mathbf{1}, \\ \mathbf{some} : A \otimes \text{queue}_A\}\}$$

This type describes a queue interface supporting insertion and deletion. The *external choice* operator $\&$ dictates that the process providing this data structure accepts either one of two messages: the labels \mathbf{ins} or \mathbf{del} . In the case of the label \mathbf{ins} , it then receives an element of type A denoted by the \multimap operator, and then the type recurses back to queue_A . On receiving a \mathbf{del} request, the process can respond with one of two labels (\mathbf{none} or \mathbf{some}), indicated by

the *internal choice* operator \oplus . It responds with **none** and then *terminates* (indicated by **1**) if the queue is empty, or with **some** followed by the element of type A (expressed with the \otimes operator) and recurses if the queue is nonempty. However, the simple session type does not express the conditions under which the **none** and **some** branches must be chosen, which requires tracking the length of the queue.

We propose extending session types with simple arithmetic refinements to express, for instance, the size of a queue. The more precise type

$$\begin{aligned} \text{queue}_A[n] = \&\{\text{ins} : A \multimap \text{queue}_A[n + 1], \\ &\text{del} : \oplus\{\text{none} : ?\{n = 0\}. \mathbf{1}, \\ &\quad \text{some} : ?\{n > 0\}. A \otimes \text{queue}_A[n - 1]\}\} \end{aligned}$$

uses the index refinement n to indicate the size of the queue. In addition, we introduce a *type constraint* $?\{\phi\}. A$ which can be read as “*there exists a proof of ϕ* ” and is analogous to the *assertion* of ϕ in imperative languages. Here, the process providing the queue must (conceptually) send a proof of $n = 0$ after it sends **none**, and a proof of $n > 0$ after it sends **some**. It is therefore constrained in its choice between the two branches based on the value of the index n . Because the the index domain from which the propositions ϕ are drawn is Presburger arithmetic and hence decidable, no proof of ϕ will actually be sent, but we can nevertheless verify the constraint statically (which is the subject of this chapter) or dynamically (see [77, 78]). Although not used in this example, we also add the dual $!\{\phi\}. A$ (*for all proofs of ϕ* , analogous to the *assumption* of ϕ), and explicit quantifiers $\exists n. A$ and $\forall n. A$ that send and receive natural numbers, respectively.

Of course, arithmetic type refinements are not new and have been explored extensively in functional languages, for example, by Zenger [164], in DML [162], or in the form of *Liquid Types* [140]. Variants have been adapted to session types as well [78, 82, 165], generally with the implicit assumption that index refinements are somehow “orthogonal” to session types. In this chapter we show that, upon closer examination, this is *not* the case. In particular, unlike in the functional setting, session type equality and therefore type checking become undecidable. Remarkably, this is the case whether we treat session types equirecursively [73] or isorecursively [113], and even in the quantifier-free fragment. In response, we develop a new algorithm for type equality which, though incomplete, easily handles the wide variety of example programs we have tried. Moreover, it is naturally extensible through the additional assertion of type invariants should the need arise.

With a practically effective type equality algorithm in hand, we then turn our attention to *type checking*. It turns out that assuming an oracle for type equality, type checking is decidable because it can be reduced to checking the validity of propositions in Presburger arithmetic. We define *type checking* over a language where constructs related to arithmetic constraints ($\exists n. A$, $\forall n. A$, $?\{\phi\}. A$, and $!\{\phi\}. A$) have explicit communication counterparts. Despite the high theoretical complexity of deciding Presburger arithmetic, all our examples check very quickly using Cooper’s decision procedure [52] with two optimizations.

Many programs in this explicit language are unnecessarily verbose and therefore tedious for the programmer to write, because the process constructs pertaining to the refinement layer contribute only to verifying its properties, but not its observable computational outcomes. As is common for refinement types, we therefore also designed an *implicit* language for processes where most constructs related to index refinements are omitted. The problem of *reconstruction*

is then to map such an implicit program to an explicit one which is sound (the result type-checks) and complete (if there is a reconstruction, it can be found). Interestingly, the nature of Presburger arithmetic makes full reconstruction impossible. For example, the proposition $\forall n. \exists k. (n = 2k \vee n = 2k + 1)$ is true but the witness for k as a Skolem function of n (namely $\lfloor n/2 \rfloor$) cannot be expressed in Presburger arithmetic. Since witnesses are critical if we want to understand the work performed by a computation, we require that type quantifiers $\forall n. A$ and $\exists n. A$ have explicit witnesses in processes. We provide a sound and complete algorithm for the resulting reconstruction problem. This algorithm exploits proof-theoretic properties of the sequent calculus akin to focusing [20] to avoid backtracking and consequently provides precise error messages that we have found to be helpful.

We have implemented our language, named Rast, in SML, where a programmer can choose explicit or implicit syntax and the exact cost model for work analysis. The implementation consists of a lexer, parser, type checker, and reconstruction engine, with particular attention to providing precise error messages.

3.2 Arithmetic Refinements

Before we extend our language of types formally, we revisit the examples in order to motivate the specific constructs available. We write $V[\vec{e}]$ for a type indexed by a sequence of arithmetic expressions e . Since it has been appropriate for most of our examples, we restrict ourselves to natural numbers rather than arbitrary integers.

Example 1 (Queues, v2). *The provider of a queue should be constrained to answer **none** exactly if the queue contains no elements and **some** if it is nonempty. The queue type from Section 3.1 does not express this. This means a client may need to have some redundant branches to account for responses that should be impossible. Instead we use the refined type $\text{queue}_A[n]$ to stand for a queue with n elements.*

$$\begin{aligned} \text{queue}_A[n] = \&\{\mathbf{ins} : A \multimap \text{queue}_A[n + 1], \\ &\mathbf{del} : \oplus\{\mathbf{none} : ?\{n = 0\}. \mathbf{1}, \\ &\quad \mathbf{some} : ?\{n > 0\}. A \otimes \text{queue}_A[n - 1]\}\} \end{aligned}$$

*The first branch is easy to understand: if we add an element to a queue of length n , it subsequently contains $n + 1$ elements. In the second branch we constrain the arithmetic variable n to be equal to 0 if the provider sends **none** and positive if the provider sends **some**. In the latter case, we subtract one from the length after an element has been dequeued.*

Conceptually, the type $?\{\phi\}. A$ means that the provider must send a proof of ϕ , so it corresponds to $\exists p : \phi. A$. A characteristic of *type refinement*, in contrast to fully dependent types, is that the computation of A can only depend on the *existence* of a proof p , but not on its form. Since our index domain is also decidable no actual proof needs to be sent (since one can be constructed from ϕ automatically, if needed), just a token *asserting* its existence. There is also a dual constructor $!\{\phi\}. A$ that licenses the *assumption* of ϕ , which, conceptually, corresponds to *receiving* a proof of ϕ .

Example 2 (Binary Numbers). *We would like the indexed type $\text{bin}[n]$ to represent a binary number with value n . Because the least significant bit comes first, we expect, for example, that $\text{bin}[n] = \oplus\{\mathbf{b0} : ?\{2 \mid n\}. \text{bin}[n/2], \dots\}$. However, while divisibility is available in Presburger*

arithmetic, division itself is not; instead, we can express the constraint and the index of the recursive occurrence using quantification.

$$\begin{aligned} \text{bin}[n] = \oplus \{ & \mathbf{b0} : \exists k. ?\{n = 2 * k\}. \text{bin}[k], \\ & \mathbf{b1} : \exists k. ?\{n = 2 * k + 1\}. \text{bin}[k], \\ & \mathbf{e} : ?\{n = 0\}. \mathbf{1} \} \end{aligned}$$

As a further refinement, we could rule out leading zeros by adding the constraint $n > 0$ in the branch for $\mathbf{b0}$.

The type $\exists n. A$ means that the provider must send a natural number i and proceed at type $A[i/n]$, corresponding to existential quantification in arithmetic. The dual universal quantifier $\forall n. A$ requires the provider to receive a number i and proceed at type $A[i/n]$.

We now extend our definitions to account for these new constructs. Below, i represents a constant, while n represents a natural number variable.

Types	$A ::= \dots$		
	$?\{\phi\}. A$	assert ϕ	continue at type A
	$!\{\phi\}. A$	assume ϕ	continue at type A
	$\exists n. A$	send number i	continue at type $A[i/n]$
	$\forall n. A$	receive number i	continue at type $A[i/n]$
	$V[\bar{e}]$	variable instantiation	
Arith. Expressions	$e ::=$	$i \mid e + e \mid e - e \mid i \times e \mid (i \mid e) \mid n$	
Arith. Propositions	$\phi ::=$	$e = e \mid e > e \mid \top \mid \perp \mid \phi \wedge \phi \mid \phi \vee \phi \mid \neg \phi \mid \exists n. \phi \mid \forall n. \phi$	
Signature	$\Sigma ::=$	$\cdot \mid \Sigma, V[\bar{n} \mid \phi] = A$	

An indexed type definition $V[\bar{n} \mid \phi] = A$ requires every instance \bar{e} of the sequence of variables \bar{n} to satisfy $\phi[\bar{e}/\bar{n}]$. This is verified statically when a type signature is checked for validity, as defined below. We use \mathcal{V} for a collection of arithmetic variables and \mathcal{C} (to signify *constraints*) for an arithmetic proposition occurring among the antecedents of a judgment. We then have the following rules defining the validity of signatures ($\vdash \Sigma$ *signature*), declarations ($\vdash_{\Sigma} \Sigma'$ *valid*), and types ($\mathcal{V} ; \mathcal{C} \vdash_{\Sigma} A$ *valid*) where \mathcal{V} is a collection of arithmetic variables including all free variables in constraint \mathcal{C} and type A . We silently rename variables so that n does not already occur in \mathcal{V} in the $\exists V$ and $\forall V$ rules. We also call upon the semantic entailment judgment $\mathcal{V} ; \mathcal{C} \vDash \phi$ which means that $\forall \mathcal{V}. \mathcal{C} \supset \phi$ holds in arithmetic and $\vDash \phi$ abbreviates $\cdot ; \top \vDash \phi$.

$$\begin{array}{c} \frac{\vdash_{\Sigma} \Sigma \text{ valid}}{\vdash \Sigma \text{ signature}} \qquad \frac{}{\vdash_{\Sigma} (\cdot) \text{ valid}} \qquad \frac{\vdash_{\Sigma} \Sigma' \text{ valid} \quad \bar{n} ; \phi \vdash_{\Sigma} A \text{ valid} \quad A \neq V[\bar{e}']}{\vdash_{\Sigma} \Sigma', V[\bar{n} \mid \phi] = A \text{ valid}} \\ \\ \frac{\mathcal{V} ; \mathcal{C} \wedge \phi \vdash_{\Sigma} A \text{ valid}}{\mathcal{V} ; \mathcal{C} \vdash_{\Sigma} ?\{\phi\}. A \text{ valid}} ?V \qquad \frac{\mathcal{V} ; \mathcal{C} \wedge \phi \vdash_{\Sigma} A \text{ valid}}{\mathcal{V} ; \mathcal{C} \vdash_{\Sigma} !\{\phi\}. A \text{ valid}} !V \\ \\ \frac{\mathcal{V}, n ; \mathcal{C} \vdash_{\Sigma} A \text{ valid}}{\mathcal{V} ; \mathcal{C} \vdash_{\Sigma} \exists n. A \text{ valid}} \exists V^n \qquad \frac{\mathcal{V}, n ; \mathcal{C} \vdash_{\Sigma} A \text{ valid}}{\mathcal{V} ; \mathcal{C} \vdash_{\Sigma} \forall n. A \text{ valid}} \forall V^n \\ \\ \frac{V[\bar{n} \mid \phi] = A \in \Sigma \quad \mathcal{V} ; \mathcal{C} \vDash \phi[\bar{e}/\bar{n}]}{\mathcal{V} ; \mathcal{C} \vdash_{\Sigma} V[\bar{e}] \text{ valid}} \text{tdef} \end{array}$$

We elide the compositional rules for all the other type constructors. Since we like to work over natural numbers rather than integers, it is convenient to assume that every definition $V[\bar{n}] = A$ abbreviates $V[\bar{n} \mid \bar{n} \geq 0] = A$. This means that in valid signatures every occurrence $V[\bar{e}]$ is such that $\bar{e} \geq 0$ follows from the known constraints.

Example 3. *The declaration*

$$\begin{aligned} \text{queue}_A[n] = & \&\{\text{ins} : A \multimap \text{queue}_A[n+1], \\ & \text{del} : \oplus\{\text{none} : ?\{n=0\}. \mathbf{1}, \\ & \quad \text{some} : ?\{n>0\}. A \otimes \text{queue}_A[n-1]\}\} \end{aligned}$$

is valid because $n \geq 0 \models n+1 \geq 0$ and $n \geq 0 \wedge n > 0 \models n-1 \geq 0$.

Unfolding a definition must substitute for the arithmetic variables we abstract over.

Definition 1. $\text{unfold}_\Sigma(V[\bar{e}]) = A[\bar{e}/\bar{n}]$ if $V[\bar{n} \mid \phi] = A \in \Sigma$ and $\text{unfold}_\Sigma(A) = A$ otherwise.

We say that a type is *closed* if it contains no free arithmetic variables n .

Definition 2. A relation \mathcal{R} on types is a type bisimulation if $(A, B) \in \mathcal{R}$ implies that for $S = \text{unfold}_\Sigma(A)$, $T = \text{unfold}_\Sigma(B)$ we have

1. If $S = \oplus\{\ell : A_\ell\}_{\ell \in L}$ then $T = \oplus\{\ell : B_\ell\}_{\ell \in L}$ and $(A_\ell, B_\ell) \in \mathcal{R}$ for all $\ell \in L$.
2. If $S = \&\{\ell : A_\ell\}_{\ell \in L}$ then $T = \&\{\ell : B_\ell\}_{\ell \in L}$ and $(A_\ell, B_\ell) \in \mathcal{R}$ for all $\ell \in L$.
3. If $S = A_1 \otimes A_2$, then $T = B_1 \otimes B_2$ and $(A_1, B_1) \in \mathcal{R}$ and $(A_2, B_2) \in \mathcal{R}$.
4. If $S = A_1 \multimap A_2$, then $T = B_1 \multimap B_2$ and $(A_1, B_1) \in \mathcal{R}$ and $(A_2, B_2) \in \mathcal{R}$.
5. If $S = \mathbf{1}$ then $T = \mathbf{1}$.
6. If $S = ?\{\phi\}. A'$ then $T = ?\{\psi\}. B'$ and either (i) $\models \phi, \models \psi$, and $(A', B') \in \mathcal{R}$, or (ii) $\models \neg\phi$ and $\models \neg\psi$.
7. If $S = !\{\phi\}. A'$ then $T = !\{\psi\}. B'$ and either (i) $\models \phi, \models \psi$, and $(A', B') \in \mathcal{R}$, or (ii) $\models \neg\phi$ and $\models \neg\psi$.
8. If $S = \exists m. A'$ then $T = \exists n. B'$ and for all $i \in \mathbb{N}$, $(A'[i/m], B'[i/n]) \in \mathcal{R}$.
9. If $S = \forall m. A'$ then $T = \forall n. B'$ and for all $i \in \mathbb{N}$, $(A'[i/m], B'[i/n]) \in \mathcal{R}$.

Definition 3. We say that A is equal to B , written $A \equiv B$, if there is a type bisimulation \mathcal{R} such that $(A, B) \in \mathcal{R}$.

An interesting point here is provided by the cases (ii) in the clauses (6) and (7). Because the type must be closed, we know that ϕ and ψ will be either true or false. If both are false, no messages can be sent along a channel of either type and therefore the continuation types A' and B' are irrelevant when considering type equality.

Fundamentally, due to the presence of arbitrary recursion and therefore non-termination, we always view a type as a restriction of what a process *might* send or receive along some channel, but it is neither *required* to send a message nor *guaranteed* to receive one. This is similar to functional programming with unrestricted recursion where an expression may not return a value. The definition based on observability of messages is then somewhat strict, as exemplified by the next example.

Example 4. *Consider*

$$\begin{aligned} \text{bin}[n] = & \oplus\{\mathbf{b0} : \exists k. ?\{n = 2 * k\}. \text{bin}[k], \\ & \quad \mathbf{b1} : \exists k. ?\{n = 2 * k + 1\}. \text{bin}[k], \\ & \quad \mathbf{e} : ?\{n = 0\}. \mathbf{1}\} \end{aligned}$$

$$\text{zero} = \oplus\{\mathbf{b0} : \exists k. ?\{k = 0\}. \text{zero},$$

$e : \{0 = 0\}.1\}$

We might expect $\text{bin}[0] \equiv \text{zero}$, but this is not so. A process of type $\text{bin}[0]$ could send the label $\mathbf{b1}$ and maybe even, say, 0 for k and then just loop forever (because there is no proof of $0 = 1$). The type zero can not exhibit this behavior so the types are not equivalent.

In our implementation, missing branches for a choice in process definitions are reconstructed with a continuation that marks it as impossible, which is then verified by the type checker. We found this simple technique significantly limited the need for subtyping or explicit definition of types such as zero —instead, we just work with $\text{bin}[0]$.

The following properties of type equality are straightforward.

Lemma 1 (Properties of Type Equality). *The relation \equiv is reflexive, symmetric, transitive and a congruence on closed valid types.*

3.3 Undecidability of Type Equality

We prove the undecidability of type equality by exhibiting a reduction from an undecidable problem about two counter machines.

The type system allows us to simulate two counter machines [119]. Intuitively, arithmetic constraints allow us to model branching zero-tests available in the machine. This, coupled with recursion in the language of types, establishes undecidability. Remarkably, a small fragment of our language containing only type definitions, internal choice (\oplus) and assertions ($\{ \phi \}$. A) where ϕ just contains constraints $n = 0$ and $n > 0$ is sufficient to prove undecidability. Moreover, the proof still applies if we treat types isorecursively. In the remainder of this section we provide some details of the undecidability proof.

Definition 4 (Two Counter Machine). *A two counter machine \mathcal{M} is defined by a sequence of instructions $\iota_1, \iota_2, \dots, \iota_m$ where each instruction is one of the following.*

- “ $\text{inc}(c_j); \text{goto } k$ ” (increment counter j by 1 and go to instruction k)
- “ $\text{zero}(c_j)? \text{goto } k : \text{dec}(c_j); \text{goto } l$ ” (if the value of the counter j is 0, go to instruction k , else decrement the counter by 1 and go to instruction l)
- “halt” (stop computation)

A configuration C of the machine \mathcal{M} is defined as a triple (i, c_1, c_2) , where i denotes the number of the current instruction and c_j ’s denote the value of the counters. A configuration C' is defined as the successor configuration of C , written as $C \mapsto C'$ if C' is the result of executing the i -th instruction on C . If $\iota_i = \text{halt}$, then $C = (i, c_1, c_2)$ has no successor configuration. The computation of \mathcal{M} is the unique maximal sequence $\rho = \rho(0)\rho(1) \dots$ such that $\rho(i) \mapsto \rho(i+1)$ and $\rho(0) = (1, 0, 0)$. Either ρ is infinite, or ends in (i, c_1, c_2) such that $\iota_i = \text{halt}$ and $c_1, c_2 \in \mathbb{N}$.

The *halting problem* refers to determining whether the computation of a two counter machine \mathcal{M} with given initial values $c_1, c_2 \in \mathbb{N}$ is finite. Both the halting problem and its dual, the *non-halting problem*, are undecidable.

Theorem 3. *Given a valid signature Σ and two types A and B such that $m, n ; \top \vdash_{\Sigma} A, B$ valid. Then it is undecidable whether for concrete $i, j \in \mathbb{N}$ we have $A[i/m, j/n] \equiv B[i/m, j/n]$.*

Proof. Given a two counter machine, we construct a signature Σ and two types A and B with

free arithmetic variables m and n such that the computation of the machine starting with initial counter values i and j is infinite iff $A[i/m, j/n] \equiv B[i/m, j/n]$ in Σ .

We define types $T_{\text{inf}} = \oplus\{\ell : T_{\text{inf}}\}$ and $T'_{\text{inf}} = \oplus\{\ell' : T'_{\text{inf}}\}$ for *distinct* labels ℓ and ℓ' . Next, for every instruction ι_i , we define types T_i and T'_i based on the form of the instruction.

- Case ($\iota_i = \text{inc}(c_1); \text{goto } k$): We define

$$\begin{aligned} T_i[c_1, c_2] &= \oplus\{\text{inc}_1 : T_k[c_1 + 1, c_2]\} \\ T'_i[c_1, c_2] &= \oplus\{\text{inc}_1 : T'_k[c_1 + 1, c_2]\} \end{aligned}$$

- Case ($\iota_i = \text{inc}(c_2); \text{goto } k$): We define

$$\begin{aligned} T_i[c_1, c_2] &= \oplus\{\text{inc}_2 : T_k[c_1, c_2 + 1]\} \\ T'_i[c_1, c_2] &= \oplus\{\text{inc}_2 : T'_k[c_1, c_2 + 1]\} \end{aligned}$$

- Case ($\iota_i = \text{zero}(c_1)? \text{goto } k : \text{dec}(c_1); \text{goto } l$): We define

$$\begin{aligned} T_i[c_1, c_2] &= \oplus\{\text{zero}_1 : ?\{c_1 = 0\}. T_k[c_1, c_2], \text{dec}_1 : ?\{c_1 > 0\}. T_l[c_1 - 1, c_2]\} \\ T'_i[c_1, c_2] &= \oplus\{\text{zero}_1 : ?\{c_1 = 0\}. T'_k[c_1, c_2], \text{dec}_1 : ?\{c_1 > 0\}. T'_l[c_1 - 1, c_2]\} \end{aligned}$$

- Case ($\iota_i = \text{zero}(c_2)? \text{goto } k : \text{dec}(c_2); \text{goto } l$): We define

$$\begin{aligned} T_i[c_1, c_2] &= \oplus\{\text{zero}_2 : ?\{c_2 = 0\}. T_k[c_1, c_2], \text{dec}_2 : ?\{c_2 > 0\}. T_l[c_1, c_2 - 1]\} \\ T'_i[c_1, c_2] &= \oplus\{\text{zero}_2 : ?\{c_2 = 0\}. T'_k[c_1, c_2], \text{dec}_2 : ?\{c_2 > 0\}. T'_l[c_1, c_2 - 1]\} \end{aligned}$$

- Case ($\iota_i = \text{halt}$): We define

$$\begin{aligned} T_i[c_1, c_2] &= T_{\text{inf}} \\ T'_i[c_1, c_2] &= T'_{\text{inf}} \end{aligned}$$

We set type $A = T_1[m, n]$ and $B = T'_1[m, n]$. Now suppose, the counter machine \mathcal{M} is initialized in the state $(1, i, j)$. The type equality question we ask is whether $T_1[i, j] \equiv T'_1[i, j]$. The two types only differ at the halting instruction. If \mathcal{M} does not halt, the two types capture exactly the same communication behavior, since the halting instruction is never reached and they agree on all other instructions. If \mathcal{M} halts, the first type proceeds with label ℓ and the second with ℓ' and they are therefore not equal. Hence, the two types are equal iff \mathcal{M} does not halt. \square

We can easily modify this reduction for an isorecursive interpretation of types, by wrapping $\oplus\{\text{unfold} : __\}$ around the right-hand side of each type definition to simulate the unfold message. We also see that a host of other problems are undecidable, such as determining whether two types with free arithmetic variables are equal for all instances. This is the problem that arises while type-checking parametric process definitions.

3.4 A Practical Algorithm for Type Equality

Despite its undecidability, we have designed a coinductive algorithm for soundly approximating type equality. Similar to Gay and Hole’s algorithm [73], it proceeds by attempting to construct a bisimulation. Due to the undecidability of the problem, our algorithm can terminate in three different states: (1) we have succeeded in constructing a bisimulation, (2) we have found a counterexample to type equality by finding a place where the types may exhibit different behavior, or (3) we have terminated the search without a definitive answer. From the point of view of type-checking, both (2) and (3) are interpreted as a failure to type-check (but there is a recourse; see Section 3.4.2). Our algorithm is expressed as a set of inference rules where the execution of the algorithm corresponds to the bottom-up construction of a deduction. The algorithm is deterministic (no backtracking) and the implementation is quite efficient in practice (see Section 3.8).

One of the basic operations in Gay and Hole’s algorithm is *loop detection*, that is, we have to determine that we have already added an equation $A \equiv B$ to the bisimulation we are constructing. Since we must treat *open types*, that is, types with free arithmetic variables subject to some constraints, determining if we have considered an equation already becomes a difficult operation. To that purpose we make an initial pass over the given type and introduce fresh *internal names* abstracted over their free type variables and known constraints. In the resulting signature defined type variables and type constructor alternates and we can perform loop detection entirely on type definitions (whether internal or external).

Example 5 (Queues, v3). *After creating internal names $\%i$ for the type of queue we obtain the following signature (here parametric in A).*

$$\begin{aligned}
 \text{queue}_A[n] &= \&\{\mathbf{ins} : \%0[n], \mathbf{del} : \%1[n]\} \\
 \%0[n] &= A \multimap \text{queue}_A[n+1] & \%3 &= \mathbf{1} \\
 \%1[n] &= \oplus\{\mathbf{none} : \%2[n], \mathbf{some} : \%4[n]\} & \%4[n] &= ?\{n > 0\}. \%5[n] \\
 \%2[n] &= ?\{n = 0\}. \%3 & \%5[n \mid n > 0] &= A \otimes \text{queue}_A[n-1]
 \end{aligned}$$

Based on the invariants established by internal names, the algorithm only needs to compare two type variables or two structural types. The rules are shown in Figure 3.1. The judgment has the form $\mathcal{V} ; \mathcal{C} ; \Gamma \vdash A \equiv B$ where \mathcal{V} contains the free arithmetic variables in the constraints \mathcal{C} and the types A and B , and Γ is a collection of *closures* $\langle \mathcal{V}' ; \mathcal{C}' ; V_1'[\bar{e}_1] \equiv V_2'[\bar{e}_2] \rangle$. If a derivation can be constructed, all ground instances of all closures are included in the resulting bisimulation (see the proof of Theorem 4). A ground instance $V_1'[\bar{e}_1'[\sigma']] \equiv V_2'[\bar{e}_2'[\sigma']]$ is given by a substitution σ' over variables in \mathcal{V}' such that $\models \mathcal{C}'[\sigma']$.

The rules for type constructors simply compare the components. If the type constructors (or the label sets in the \oplus and $\&$ rules) do not match, then type equality fails (having constructed a counterexample to bisimulation) unless the \perp rule applies. This rule handles the case where the constraints are contradictory and no communication is possible.

The rule of reflexivity is needed explicitly here (but not in the version of Gay and Hole) because due to the incompleteness of the algorithm we may otherwise fail to recognize type variables with equal index expressions as equal.

Now we come to the key rules, *expd* and *def*. In the *expd* rule we expand the definitions of $V_1[\bar{e}_1]$ and $V_2[\bar{e}_2]$, and we also add the closure $\langle \mathcal{V} ; \mathcal{C} ; V_1[\bar{e}_1] \equiv V_2[\bar{e}_2] \rangle$ to Γ . Since the equality of $V_1[\bar{e}_1]$ and $V_2[\bar{e}_2]$ must hold for all its ground instances, the extension of Γ with the

$$\begin{array}{c}
\frac{\mathcal{V}; \mathcal{C}; \Gamma \vdash A_\ell \equiv B_\ell \quad (\forall \ell \in L)}{\mathcal{V}; \mathcal{C}; \Gamma \vdash \oplus\{\ell : A_\ell\}_{\ell \in L} \equiv \oplus\{\ell : B_\ell\}_{\ell \in L}} \oplus \quad \frac{\mathcal{V}; \mathcal{C}; \Gamma \vdash A_\ell \equiv B_\ell \quad (\forall \ell \in L)}{\mathcal{V}; \mathcal{C}; \Gamma \vdash \&\{\ell : A_\ell\}_{\ell \in L} \equiv \&\{\ell : B_\ell\}_{\ell \in L}} \& \\
\\
\frac{\mathcal{V}; \mathcal{C}; \Gamma \vdash A_1 \equiv B_1 \quad \mathcal{V}; \mathcal{C}; \Gamma \vdash A_2 \equiv B_2}{\mathcal{V}; \mathcal{C}; \Gamma \vdash A_1 \otimes A_2 \equiv B_1 \otimes B_2} \otimes \\
\\
\frac{\mathcal{V}; \mathcal{C}; \Gamma \vdash A_1 \equiv B_1 \quad \mathcal{V}; \mathcal{C}; \Gamma \vdash A_2 \equiv B_2}{\mathcal{V}; \mathcal{C}; \Gamma \vdash A_1 \multimap A_2 \equiv B_1 \multimap B_2} \multimap \quad \frac{}{\mathcal{V}; \mathcal{C}; \Gamma \vdash \mathbf{1} \equiv \mathbf{1}} \mathbf{1} \\
\\
\frac{\mathcal{V}; \mathcal{C} \vDash \phi \leftrightarrow \psi \quad \mathcal{V}; \mathcal{C} \wedge \phi; \Gamma \vdash A \equiv B}{\mathcal{V}; \mathcal{C}; \Gamma \vdash ?\{\phi\}. A \equiv ?\{\psi\}. B} ? \\
\\
\frac{\mathcal{V}; \mathcal{C} \vDash \phi \leftrightarrow \psi \quad \mathcal{V}; \mathcal{C} \wedge \phi; \Gamma \vdash A \equiv B}{\mathcal{V}; \mathcal{C}; \Gamma \vdash !\{\phi\}. A \equiv !\{\psi\}. B} ! \quad \frac{\mathcal{V}, k; \mathcal{C}; \Gamma \vdash A[k/m] \equiv B[k/n]}{\mathcal{V}; \mathcal{C}; \Gamma \vdash \exists m. A \equiv \exists n. B} \exists^k \\
\\
\frac{\mathcal{V}, k; \mathcal{C}; \Gamma \vdash A[k/m] \equiv B[k/n]}{\mathcal{V}; \mathcal{C}; \Gamma \vdash \forall m. A \equiv \forall n. B} \forall^k \quad \frac{\mathcal{V}; \mathcal{C} \vDash \perp}{\mathcal{V}; \mathcal{C}; \Gamma \vdash A \equiv B} \perp \\
\\
\frac{\mathcal{V}; \mathcal{C} \vDash e_1 = e'_1 \wedge \dots \wedge e_n = e'_n}{\mathcal{V}; \mathcal{C}; \Gamma \vdash V[\bar{e}] \equiv V[\bar{e}']} \text{ refl} \quad \frac{V_1[\bar{v}_1 \mid \phi_1] = A \in \Sigma \quad V_2[\bar{v}_2 \mid \phi_2] = B \in \Sigma \quad \gamma = \langle \mathcal{V}; \mathcal{C}; V_1[\bar{e}_1] \equiv V_2[\bar{e}_2] \rangle \quad \mathcal{V}; \mathcal{C}; \Gamma, \gamma \vdash A[\bar{e}_1/\bar{v}_1] \equiv B[\bar{e}_2/\bar{v}_2]}{\mathcal{V}; \mathcal{C}; \Gamma \vdash V_1[\bar{e}_1] \equiv V_2[\bar{e}_2]} \text{ expd} \\
\\
\frac{\langle \mathcal{V}'; \mathcal{C}'; V_1[\bar{e}'_1] \equiv V_2[\bar{e}'_2] \rangle \in \Gamma \quad \mathcal{V}; \mathcal{C} \vDash \exists \mathcal{V}'. \mathcal{C}' \wedge \bar{e}'_1 = \bar{e}_1 \wedge \bar{e}'_2 = \bar{e}_2}{\mathcal{V}; \mathcal{C}; \Gamma \vdash V_1[\bar{e}_1] \equiv V_2[\bar{e}_2]} \text{ def}
\end{array}$$

Figure 3.1: Algorithmic Rules for Type Equality

corresponding closure remembers exactly that.

In the def rule we close off the derivation successfully if all instances of the equation $V_1[\bar{e}_1] \equiv V_2[\bar{e}_2]$ are already instances of a closure in Γ . This is checked by the entailment in the second premise, $\mathcal{V}; \mathcal{C} \vDash \exists \mathcal{V}'. \mathcal{C}' \wedge \bar{E}_1 = \bar{e}_1 \wedge \bar{E}_2 = \bar{e}_2$. This entailment is verified as a closed $\forall \exists$ arithmetic formula, even if the original constraints \mathcal{C} and \mathcal{C}' do not contain any quantifiers. While for Presburger arithmetic we can decide such a proposition using quantifier elimination, other constraint domains may not permit such a decision procedure.

The algorithm so far is sound, but potentially nonterminating because when encountering variable/variable equations, we can use the expd rule indefinitely. To ensure termination, we restrict the expd rule to the case where *no* formula with the same type variables V_1 and V_2 is already present in Γ . This also removes the overlap between these two rules. Note that if type variables have no parameters, our algorithm specializes to Gay and Hole's (with the small optimizations of reflexivity and internal naming), which means our algorithm is sound and complete on unindexed types.

Example 6 (Integer Counter). An integer counter with increment (**inc**), decrement (**dec**) and sign-test (**sgn**) operations provides type $\text{intctr}[x, y]$, where the current value of the counter is $x - y$ for natural numbers x and y .

$$\begin{aligned} \text{intctr}[x, y] = \&\{\mathbf{inc} : \text{intctr}[x + 1, y], \\ &\mathbf{dec} : \text{intctr}[x, y + 1], \\ &\mathbf{sgn} : \oplus\{\mathbf{neg} : ?\{x < y\}. \text{intctr}[x, y], \\ &\quad \mathbf{zer} : ?\{x = y\}. \text{intctr}[x, y], \\ &\quad \mathbf{pos} : ?\{x > y\}. \text{intctr}[x, y]\}\} \end{aligned}$$

Under this definition our algorithm verifies, for example, that an increment followed by a decrement does not change the counter value. That is,

$$x, y ; \top ; \cdot \vdash \text{intctr}[x, y] \equiv \text{intctr}[x + 1, y + 1]$$

where we have elided the assumptions $x, y \geq 0$. When applying expd , we assume that $\gamma = \langle x', y' ; \top ; \text{intctr}[x', y'] \equiv \text{intctr}[x' + 1, y' + 1] \rangle$. Then, for example, in the first branch (for **inc**) we conclude $x, y ; \top ; \gamma \vdash \text{intctr}[x + 1, y] \equiv \text{intctr}[x + 2, y + 1]$ using the def rule and the entailment $x, y ; \top \vDash \exists x'. \exists y'. x' = x + 1 \wedge y' = y \wedge x' + 1 = x + 2 \wedge y' + 1 = y + 1$. The other branches are similar.

3.4.1 Soundness of the Type Equality Algorithm

We prove that the type equality algorithm is sound with respect to the definition of type equality. The soundness is proved by constructing a type bisimulation from a derivation of the algorithmic type equality judgment. We sketch the key points of the proofs.

The first gap we have to bridge is that the type bisimulation is defined only for closed types, because observations can only arise from communication along channels which, at runtime, will be of closed type. So, if we can derive $\mathcal{V} ; \mathcal{C} ; \cdot \vdash A \equiv B$ then we should interpret this as stating that for all ground substitutions σ over \mathcal{V} such that $\vDash \mathcal{C}[\sigma]$ we have $A[\sigma] \equiv B[\sigma]$.

Definition 5. Given a relation \mathcal{R} on valid ground types and two types A and B such that $\mathcal{V} ; \mathcal{C} \vdash A, B$ valid, we write $\forall \mathcal{V}. \mathcal{C} \Rightarrow A \equiv_{\mathcal{R}} B$ if for all ground substitutions σ over \mathcal{V} such that $\vDash \mathcal{C}[\sigma]$ we have $(A[\sigma], B[\sigma]) \in \mathcal{R}$.

Furthermore, we write $\forall \mathcal{V}. \mathcal{C} \Rightarrow A \equiv B$ if there exists a type bisimulation \mathcal{R} such that $\forall \mathcal{V}. \mathcal{C} \Rightarrow A \equiv_{\mathcal{R}} B$.

Note that if $\mathcal{V} ; \mathcal{C} \vDash \perp$, then $\forall \mathcal{V}. \mathcal{C} \Rightarrow A \equiv B$ is vacuously true, since there does not exist a ground substitution σ such that $\vDash \mathcal{C}[\sigma]$. A key lemma is the following, which is needed to show the soundness of the def rule.

Lemma 2. Suppose $\forall \mathcal{V}'. \mathcal{C}' \Rightarrow V_1[\bar{e}_1'] \equiv_{\mathcal{R}} V_2[\bar{e}_2']$ holds. Further assume that $\mathcal{V} ; \mathcal{C} \vDash \exists \mathcal{V}'. \mathcal{C}' \wedge \bar{e}_1' = \bar{e}_1 \wedge \bar{e}_2' = \bar{e}_2$ for some $\mathcal{V}, \mathcal{C}, \bar{e}_1, \bar{e}_2$. Then, $\forall \mathcal{V}. \mathcal{C} \Rightarrow V_1[\bar{e}_1] \equiv_{\mathcal{R}} V_2[\bar{e}_2]$ holds.

Proof. To prove $\forall \mathcal{V}. \mathcal{C} \Rightarrow V_1[\bar{e}_1] \equiv_{\mathcal{R}} V_2[\bar{e}_2]$, it is sufficient to show that $V_1[\bar{e}_1[\sigma]] \equiv_{\mathcal{R}} V_2[\bar{e}_2[\sigma]]$ for any substitution σ over \mathcal{V} such that $\vDash \mathcal{C}[\sigma]$. Applying this substitution to $\mathcal{V} ; \mathcal{C} \vDash \exists \mathcal{V}'. \mathcal{C}' \wedge \bar{e}_1' = \bar{e}_1 \wedge \bar{e}_2' = \bar{e}_2$, we infer $\exists \mathcal{V}'. \mathcal{C}' \wedge \bar{e}_1' = \bar{e}_1[\sigma] \wedge \bar{e}_2' = \bar{e}_2[\sigma]$ since $\vDash \mathcal{C}[\sigma]$. Thus, there exists σ' over \mathcal{V}' such that $\vDash \mathcal{C}'[\sigma']$ holds, and $\bar{e}_1'[\sigma'] = \bar{e}_1[\sigma]$ and $\bar{e}_2'[\sigma'] = \bar{e}_2[\sigma]$. And since $\forall \mathcal{V}'. \mathcal{C}' \Rightarrow V_1[\bar{e}_1'] \equiv_{\mathcal{R}} V_2[\bar{e}_2']$, we deduce that for any ground substitution (including the current

one) σ' over \mathcal{V}' , $V_1[\bar{e}_1'[\sigma']] \equiv_{\mathcal{R}} V_2[\bar{e}_2'[\sigma']]$ holds. This implies that $V_1[\bar{e}_1[\sigma]] \equiv_{\mathcal{R}} V_2[\bar{e}_2[\sigma]]$ since $\bar{e}_1'[\sigma'] = \bar{e}_1[\sigma]$ and $\bar{e}_2'[\sigma'] = \bar{e}_2[\sigma]$. \square

We construct the bisimulation from a derivation of $\mathcal{V} ; \mathcal{C} ; \Gamma \vdash A \equiv B$ by (i) collecting the conclusions of all the sequents, excepting only the def rule, and (ii) forming all ground instances from them.

Definition 6. Given a derivation \mathcal{D} of $\mathcal{V} ; \mathcal{C} ; \Gamma \vdash A \equiv B$, we define the set $\mathcal{S}(\mathcal{D})$ of closures. For each sequent $\mathcal{V}' ; \mathcal{C}' ; \Gamma' \vdash A' \equiv B'$ (except the conclusion of the def rule) we include the closure $\langle \mathcal{V}' ; \mathcal{C}' ; A' \equiv B' \rangle$ in $\mathcal{S}(\mathcal{D})$.

Theorem 4. If $\mathcal{V} ; \mathcal{C} ; \cdot \vdash A \equiv B$, then $\forall \mathcal{V}. \mathcal{C} \Rightarrow A \equiv B$.

Proof. We are given a derivation \mathcal{D}_0 of $\mathcal{V}_0 ; \mathcal{C}_0 ; \cdot \vdash A_0 \equiv B_0$. Construct $\mathcal{S}(\mathcal{D}_0)$ and define a relation \mathcal{R} on closed valid types as follows:

$$\mathcal{R} = \{(A[\sigma], B[\sigma]) \mid \langle \mathcal{V} ; \mathcal{C} ; A \equiv B \rangle \in \mathcal{S}(\mathcal{D}_0) \text{ and } \sigma \text{ over } \mathcal{V} \text{ with } \models \mathcal{C}[\sigma]\}$$

We first prove that \mathcal{R} is a type bisimulation. Then our theorem follows since the closure $\langle \mathcal{V}_0 ; \mathcal{C}_0 ; A_0 \equiv B_0 \rangle \in \mathcal{S}(\mathcal{D}_0)$.

Consider $(A[\sigma], B[\sigma]) \in \mathcal{R}$ where $\langle \mathcal{V} ; \mathcal{C} ; A \equiv B \rangle \in \mathcal{S}(\mathcal{D}_0)$ for some σ over \mathcal{V} and $\models \mathcal{C}[\sigma]$.

First, consider the case where $\mathcal{V} ; \mathcal{C} \models \perp$. Under such a constraint, $\mathcal{V} ; \mathcal{C} ; \cdot \vdash A \equiv B$ holds true due to the \perp rule. Furthermore, $\forall \mathcal{V}. \mathcal{C} \Rightarrow A \equiv B$ holds vacuously, and the algorithm is sound. For the remaining cases, we case analyze on the structure of $A[\sigma]$ and assume that there exists a ground substitution σ such that $\models \mathcal{C}[\sigma]$.

Consider the case where $A = \oplus\{\ell : A_\ell\}_{\ell \in L}$. Since A and B are both structural, $B = \oplus\{\ell : B_\ell\}_{\ell \in L}$. Since $\langle \mathcal{V} ; \mathcal{C} ; A \equiv B \rangle \in \mathcal{S}(\mathcal{D}_0)$, by definition of $\mathcal{S}(\mathcal{D}_0)$, we get $\langle \mathcal{V} ; \mathcal{C} ; A_\ell \equiv B_\ell \rangle \in \mathcal{S}(\mathcal{D}_0)$ for all $\ell \in L$. By the definition of \mathcal{R} , we get that $(A_\ell[\sigma], B_\ell[\sigma]) \in \mathcal{R}$. Also, $A[\sigma] = \oplus\{\ell : A_\ell[\sigma]\}_{\ell \in L}$ and similarly, $B[\sigma] = \oplus\{\ell : B_\ell[\sigma]\}_{\ell \in L}$. Hence, \mathcal{R} satisfies the appropriate closure condition for a type bisimulation.

Next, consider the case where $A = ?\{\phi\}. A'$. Since A and B are both structural, $B = ?\{\psi\}. B'$. Since $\langle \mathcal{V} ; \mathcal{C} ; A \equiv B \rangle \in \mathcal{S}(\mathcal{D}_0)$, we obtain $\mathcal{V} ; \mathcal{C} \models \phi \leftrightarrow \psi$ and $\langle \mathcal{V} ; \mathcal{C} \wedge \phi ; A' \equiv B' \rangle \in \mathcal{S}(\mathcal{D}_0)$. Thus, for any substitution σ such that $\models \mathcal{C}[\sigma] \wedge \phi[\sigma]$, we get that $(A'[\sigma], B'[\sigma]) \in \mathcal{R}$ with $A[\sigma] = ?\{\phi[\sigma]\}. A'[\sigma]$ and $B[\sigma] = ?\{\psi[\sigma]\}. B'[\sigma]$. Since $\models \phi[\sigma]$ and $\mathcal{V} ; \mathcal{C} \models \phi \leftrightarrow \psi$ we also obtain $\models \psi[\sigma]$ and the closure condition is satisfied.

Next, consider the case where $A = \exists m. A'$. Since A and B are both structural, $B = \exists n. B'$. Since $\langle \mathcal{V} ; \mathcal{C} ; A \equiv B \rangle \in \mathcal{S}(\mathcal{D}_0)$, we get that $\langle \mathcal{V}, k ; \mathcal{C} ; A'[k/m] \equiv B'[k/n] \rangle \in \mathcal{S}(\mathcal{D}_0)$. Since k was chosen fresh and does not occur in \mathcal{C} , we obtain that for any $i \in \mathbb{N}$ we have $\models \mathcal{C}[\sigma, i/k]$ and therefore $(A'[\sigma, i/k], B'[\sigma, i/k]) \in \mathcal{R}$ for all $i \in \mathbb{N}$ and the closure condition is satisfied.

The only case where a conclusion is not added to $\mathcal{S}(\mathcal{D}_0)$ is the def rule. In this case, adding $(\forall \mathcal{V}. \mathcal{C} \Rightarrow V_1[\bar{e}_1] \equiv V_2[\bar{e}_2])$ is redundant: Lemma 2 states that $V_1[\bar{e}_1[\sigma]] \equiv_{\mathcal{R}} V_2[\bar{e}_2[\sigma]]$ which implies $(V_1[\bar{e}_1[\sigma]], V_2[\bar{e}_2[\sigma]]) \in \mathcal{R}$. \square

3.4.2 Type Equality Declarations

Even though the type equality algorithm in Section 3.4 is incomplete, we have yet to find a natural example where it fails after we added reflexivity as a general rule. But since we cannot

see a simple reason why this should be so, we made our type equality algorithm extensible by the programmer via an additional form of declaration

$$\forall \mathcal{V}. \mathcal{C} \Rightarrow V_1[\bar{e}_1] \equiv V_2[\bar{e}_2]$$

in signatures. Let Γ_Σ denote the set of all such declarations. Then we check

$$\mathcal{V} ; \mathcal{C} ; \Gamma_\Sigma \vdash V_1[\bar{e}_1] \equiv V_2[\bar{e}_2]$$

for each such declaration, seeding the construction of a bisimulation with all the given equations. Then, when type-checking has to decide the equality of two types, it starts not with the empty context Γ but with Γ_Σ . Our soundness proof can easily accommodate this more general algorithm.

3.5 Formal Description of the Rast Language

In this section, we give a formal description of the Rast language. We have already seen the rules for statics and dynamics for the basic session types in Chapter 2. Hence, this section will primarily present the rules for the refinement layer. We have already described the grammar for types in Section 3.2. We now present the grammar for process expressions in Rast.

$$\begin{array}{l} \text{Procs } P, Q ::= x.k ; P \mid \text{case } x (l \Rightarrow P)_{l \in L} \\ \quad \mid \text{send } x y ; P \mid y \leftarrow \text{recv } x ; P \\ \quad \mid \text{close } x \mid \text{wait } x ; P \\ \quad \mid x \leftrightarrow y \mid x \leftarrow f y ; P \\ \quad \mid \text{assert } x \{ \phi \} ; P \mid \text{assume } x \{ \phi \} ; P \\ \quad \mid \text{send } x \{ e \} ; P \mid \{ n \} \leftarrow \text{recv } x ; P \end{array}$$

The typing judgment has the form of a sequent

$$\mathcal{V} ; \mathcal{C} ; \Delta \vdash_\Sigma P :: (x : A)$$

where \mathcal{V} are index variables n , \mathcal{C} are constraints over these variables expressed as a single proposition, Δ are the linear antecedents $x_i : A_i$, P is a process expression, and $x : A$ is the linear succedent. We propose and maintain that the x_i 's and x are all distinct, and that all free index variables in \mathcal{C} , Δ , P , and A are contained among \mathcal{V} . Finally, Σ is a fixed signature containing type and process definitions. Because it is fixed, we elide it from the presentation of the rules. In addition we write $\mathcal{V} ; \mathcal{C} \vDash \phi$ for semantic entailment (proving ϕ assuming \mathcal{C}) in the constraint domain where \mathcal{V} contains all arithmetic variables in \mathcal{C} and ϕ . Table 3.1 reviews the basic session types their associated process terms, their continuation (both in types and terms) and operational description.

We formalize the operational semantics as a system of *multiset rewriting rules* [48]. We introduce semantic objects $\text{proc}(c, P)$ and $\text{msg}(c, M)$ which mean that process P or message M provide along channel c . A process configuration is a multiset of such objects, where any two channels provided are distinct (formally described in Section 3.4.1).

Type	Cont.	Process Term	Cont.	Description
$c : \oplus\{\ell : A_\ell\}_{\ell \in L}$	$c : A_k$	$c.k ; P$ $\text{case } c (\ell \Rightarrow Q_\ell)_{\ell \in L}$	P Q_k	send label k along c branch on received label along c
$c : \&\{\ell : A_\ell\}_{\ell \in L}$	$c : A_k$	$\text{case } c (\ell \Rightarrow P_\ell)_{\ell \in L}$ $c.k ; Q$	P_k Q	branch on received label along c send label k along c
$c : A \otimes B$	$c : B$	$\text{send } c w ; P$ $y \leftarrow \text{recv } c ; Q$	P $Q[w/y]$	send channel $w : A$ along c receive channel $w : A$ along c
$c : A \multimap B$	$c : B$	$y \leftarrow \text{recv } c ; P$ $\text{send } c w ; Q$	$P[w/y]$ Q	receive channel $w : A$ along c send channel $w : A$ along c
$c : \mathbf{1}$	—	close c wait $c ; Q$	— Q	send <i>close</i> along c receive <i>close</i> along c

Table 3.1: Basic session types with operational description

Process Definitions Process definitions have the form $\Delta \vdash f[\bar{n}] = P :: (x : A)$ where f is the name of the process and P its definition. In addition, \bar{n} is a sequence of arithmetic variables that Δ , P and A can refer to. All definitions are collected in a fixed global signature Σ . For a *well-formed signature*, we require that $\bar{n} ; \top ; \Delta \vdash P :: (x : A)$ for every definition, thereby allowing definitions to be mutually recursive. A new instance of a defined process f can be spawned with the expression $x \leftarrow f[\bar{e}] \bar{y} ; Q$ where \bar{y} is a sequence of channels matching the antecedents Δ and $[\bar{e}]$ is a sequence of arithmetic expression matching the variables $[\bar{n}]$. The newly spawned process will use all variables in \bar{y} and provide x to the continuation Q .

$$\frac{\overline{y' : B} \vdash f[\bar{n}] = P_f :: (x' : A) \in \Sigma \quad \Delta' = \overline{(y : B)[\bar{e}/\bar{n}]} \quad \mathcal{V} ; \mathcal{C} ; \Delta, (x : A[\bar{e}/\bar{n}]) \vdash Q :: (z : C)}{\mathcal{V} ; \mathcal{C} ; \Delta, \Delta' \vdash (x \leftarrow f[\bar{e}] \bar{y} ; Q) :: (z : C)} \text{ def}$$

The declaration of f is looked up in the signature Σ (first premise), and \bar{e} is substituted for \bar{n} while matching the types in Δ' and \bar{y} (second premise). Similarly, the freshly created channel x has type A from the signature with \bar{e} substituted for \bar{n} . The corresponding semantics rule also performs a similar substitution (a fresh).

$$(\text{def}C) : \text{proc}(c, x \leftarrow f[\bar{e}] \bar{d} ; Q) \mapsto \text{proc}(a, P_f[a/x, \bar{d}/\bar{y}', \bar{e}/\bar{n}]), \text{proc}(c, Q[a/x])$$

where $\overline{y' : B} \vdash f[\bar{n}] = P_f :: (x' : A) \in \Sigma$.

Sometimes a process invocation is a tail call, written without a continuation as $x \leftarrow f[\bar{e}] \bar{y}$. This is a short-hand for $x' \leftarrow f[\bar{e}] \bar{y} ; x \leftrightarrow x'$ for a fresh variable x' , that is, we create a fresh channel and immediately identify it with x .

Type Definitions As our queue example already showed, session types can be defined recursively, departing from a strict Curry-Howard interpretation of linear logic, analogous to the way pure ML or Haskell depart from a pure interpretation of intuitionistic logic. For this purpose we allow (possibly mutually recursive) type definitions $V[\bar{n} \mid \phi] = A$ in the signature Σ . Here, \bar{n} denotes a sequence of arithmetic variables. Again, for a well-formed signature, we require

Type	Cont.	Process Term	Cont.	Description
$c : \exists n. A$	$c : A[i/n]$	$\text{send } c \{e\} ; P$ $\{n\} \leftarrow \text{recv } c ; Q$	P $Q[i/n]$	provider sends value i of e along c client receives number i along c
$c : \forall n. A$	$c : A[i/n]$	$\{n\} \leftarrow \text{recv } c ; P$ $\text{send } c \{e\} ; Q$	$P[i/n]$ Q	provider receives number i along c client sends value i of e along c
$c : ?\{\phi\}. A$	$c : A$	$\text{assert } c \{\phi\} ; P$ $\text{assume } c \{\phi\} ; Q$	P Q	provider asserts ϕ on channel c client assumes ϕ on c
$c : !\{\phi\}. A$	$c : A$	$\text{assume } c \{\phi\} ; P$ $\text{assert } c \{\phi\} ; Q$	P Q	provider assumes ϕ on channel c client asserts ϕ on c

Table 3.2: Refined session types with operational description

A to be *contractive* [73] meaning A should not itself be a type name. Our type definitions are *equirecursive* so we can silently replace type names $V[\bar{e}]$ indexed with arithmetic refinements by $A[\bar{e}/\bar{n}]$ during type checking, and no explicit rules for recursive types are needed.

All types in a signature must be *valid*, formally denoted with the judgment $\mathcal{V} ; \mathcal{C} \vdash A$ *valid*, which requires that all free arithmetic variables of \mathcal{C} and A are contained in \mathcal{V} , and that for each arithmetic expression e in A we can prove $\mathcal{V}' ; \mathcal{C}' \vdash e : \text{nat}$ for the constraints \mathcal{C}' known at the occurrence of e (implicitly proving that $e \geq 0$).

3.5.1 The Refinement Layer

We describe quantifiers ($\exists n. A, \forall n. A$) and constraints ($?\{\phi\}. A, !\{\phi\}. A$) [59]. An overview of the types, process expressions, and their operational meaning can be found in Table 3.2.

Quantification The provider of $(c : \exists n. A)$ should send a witness i along channel c and then continue as $A[i/n]$. The witness is specified by an arithmetic expression e which, since it must be closed at runtime, can be evaluated to a number i (following standard evaluation rules of arithmetic). From the typing perspective, we just need to check that the expression e denotes a natural number, using only the permitted variables in \mathcal{V} . This is represented with the auxiliary judgment $\mathcal{V} ; \mathcal{C} \vdash e : \text{nat}$ (implicitly proving that $e \geq 0$ under constraint \mathcal{C}).

$$\frac{\mathcal{V} ; \mathcal{C} \vdash e : \text{nat} \quad \mathcal{V} ; \mathcal{C} ; \Delta \vdash P :: (x : A[e/n])}{\mathcal{V} ; \mathcal{C} ; \Delta \vdash \text{send } x \{e\} ; P :: (x : \exists n. A)} \exists R$$

$$\frac{\mathcal{V}, n ; \mathcal{C} ; \Delta, (x : A) \vdash Q :: (z : C) \quad (n \text{ fresh})}{\mathcal{V} ; \mathcal{C} ; \Delta, (x : \exists n. A) \vdash \{n\} \leftarrow \text{recv } x ; Q :: (z : C)} \exists L$$

Statically, the client adds n to \mathcal{V} to ensure that Q and A are closed w.r.t. \mathcal{V} . Operationally, the provider sends the arithmetic expression with the continuation channel as a message that the client receives and appropriately substitutes.

$(\exists S) : \text{proc}(c, \text{send } c \{e\} ; P) \mapsto \text{proc}(c', P[c'/c]), \text{msg}(c, \text{send } c \{e\} ; c \leftrightarrow c')$

$(\exists C) : \text{msg}(c, \text{send } c \{e\} ; c \leftrightarrow c'), \text{proc}(d, \{n\} \leftarrow \text{recv } c ; Q) \mapsto \text{proc}(d, Q[e/n][c'/c])$

The dual type $\forall n. A$ reverses the role of the provider and client. The client sends (the value of) an arithmetic expression e which the provider receives and binds to n .

$$\frac{\mathcal{V}, n; \mathcal{C}; \Delta \vdash P_n :: (x : A)}{\mathcal{V}; \mathcal{C}; \Delta \vdash \{n\} \leftarrow \text{recv } x; P_n :: (x : \forall n. A)} \forall R$$

$$\frac{\mathcal{V}; \mathcal{C} \vdash e : \text{nat} \quad \mathcal{V}; \Delta, (x : A[e/n]) \vdash Q :: (z : C)}{\mathcal{V}; \mathcal{C}; \Delta, (x : \forall n. A) \vdash \text{send } x \{e\}; Q :: (z : C)} \forall L$$

$$(\forall S) : \text{proc}(d, \text{send } c \{e\}; P) \mapsto \text{msg}(c', \text{send } c \{e\}; c' \leftrightarrow c), \text{proc}(d, [c'/c]P)$$

$$(\forall C) : \text{proc}(d, \{n\} \leftarrow \text{recv } c; Q), \text{msg}(c', \text{send } c \{e\}; c' \leftrightarrow c) \mapsto \text{proc}(d, [e/n][c'/c]Q)$$

Constraints Refined session types also allow constraints over index variables. As we have already seen in the examples, these critically govern permissible messages. From the message-passing perspective, the provider of $(c : ?\{\phi\}. A)$ should send a proof of ϕ along c and the client should receive such a proof. However, since the index domain is decidable and future computation cannot depend on the form of the proof (what is known in type theory as *proof irrelevance*) such messages are not actually exchanged. Instead, it is the provider's responsibility to ensure that ϕ holds, while the client is permitted to assume that ϕ is true. Therefore, and in an analogy with imperative languages, we write $\text{assert } c \{\phi\}; P$ for a process that *asserts* ϕ for channel c and continues with P , while $\text{assume } c \{\phi\}; Q$ *assumes* ϕ and continues with Q .

Thus, the typing rules for this new type constructor are

$$\frac{\mathcal{V}; \mathcal{C} \vDash \phi \quad \mathcal{V}; \mathcal{C}; \Delta \vdash P :: (x : A)}{\mathcal{V}; \mathcal{C}; \Delta \vdash \text{assert } x \{\phi\}; P :: (x : ?\{\phi\}. A)} ?R$$

$$\frac{\mathcal{V}; \mathcal{C} \wedge \phi; \Delta, (x : A) \vdash Q :: (z : C)}{\mathcal{V}; \mathcal{C}; \Delta, (x : ?\{\phi\}. A) \vdash \text{assume } x \{\phi\}; Q :: (z : C)} ?L$$

Notice how the provider must verify the truth of ϕ given the currently known constraints \mathcal{C} (the premise $\mathcal{V}; \mathcal{C} \vDash \phi$), while the client assumes ϕ by adding it to \mathcal{C} .

Operationally, the provider creates a message containing the constraint that is received by the client (c' fresh).

$$(?S) : \text{proc}(c, \text{assert } c \{\phi\}; P) \mapsto \text{proc}(c', [c'/c]P), \text{msg}(c, \text{assert } c \{\phi\}; c \leftrightarrow c')$$

$$(?C) : \text{msg}(c, \text{assert } c \{\phi\}; c \leftrightarrow c'), \text{proc}(d, \text{assume } c \{\phi'\}; Q) \mapsto \text{proc}(d, [c'/c]Q)$$

In well-typed configurations (which arise from executing well-typed processes) the constraint ϕ in these rules will always be closed and true so there is no need to check this explicitly.

The dual operator $!\{\phi\}. A$ reverses the role of provider and client. The provider of $x : !\{\phi\}. A$ may assume the truth of ϕ , while the client must verify it. The dual rules are

$$\frac{\mathcal{V}; \mathcal{C} \wedge \phi; \Delta \vdash P :: (x : A)}{\mathcal{V}; \mathcal{C}; \Delta \vdash \text{assume } x \{\phi\}; P :: (x : !\{\phi\}. A)} !R$$

$$\frac{\mathcal{V}; \mathcal{C} \vDash \phi \quad \mathcal{V}; \mathcal{C}; \Delta, (x : A) \vdash Q :: (z : C)}{\mathcal{V}; \mathcal{C}; \Delta, (x : !\{\phi\}. A) \vdash \text{assert } x \{\phi\}; Q :: (z : C)} !L$$

$$\begin{array}{c}
\frac{}{\Delta \Vdash (\cdot) :: \Delta} \text{ emp} \qquad \frac{\Delta_1 \Vdash \mathcal{S}_1 :: \Delta_2 \quad \Delta_2 \Vdash \mathcal{S}_2 :: \Delta_3}{\Delta_1 \Vdash (\mathcal{S}_1, \mathcal{S}_2) :: \Delta_3} \text{ comp} \\
\frac{\cdot ; \top ; \Delta \vdash P :: (x : A)}{\Delta \Vdash \text{proc}(x, P) :: (x : A)} \text{ proc} \qquad \frac{\cdot ; \top ; \Delta \vdash M :: (x : A)}{\Delta \Vdash \text{msg}(x, M) :: (x : A)} \text{ msg}
\end{array}$$

Figure 3.2: Typing rules for a configuration

The remaining issue is how to type-check a branch that is impossible due to unsatisfiable constraints. For example, if a client sends a `del` request to a provider along $c : \text{queue}_A[0]$, the type then becomes

$$c : \oplus \{ \mathbf{none} : ?\{0=0\}. \mathbf{1}, \mathbf{some} : ?\{0>0\}. A \otimes \text{queue}_A[0-1] \}$$

The client would have to branch on the label received and then assume the constraint asserted by the provider

$$\text{case } c \left(\begin{array}{l} \mathbf{none} \Rightarrow \text{assume } c \{0 = 0\} ; P_1 \\ \mathbf{some} \Rightarrow \text{assume } c \{0 > 0\} ; P_2 \end{array} \right)$$

but what could we write for P_2 in the `some` branch? Intuitively, computation should never get there because the provider can not assert $0 > 0$. Formally, we use the process expression ‘impossible’ to indicate that computation can never reach this spot:

$$\text{case } c \left(\begin{array}{l} \mathbf{none} \Rightarrow \text{assume } c \{0 = 0\} ; P_1 \\ \mathbf{some} \Rightarrow \text{assume } c \{0 > 0\} ; \text{impossible} \end{array} \right)$$

In implicit syntax (see Section 3.7) we could omit the `some` branch altogether and it would be reconstructed in the form shown above. Abstracting away from this example, the typing rule for impossibility simply checks that the constraints are indeed unsatisfiable

$$\frac{\mathcal{V} ; \mathcal{C} \vDash \perp}{\mathcal{V} ; \mathcal{C} ; \Delta \vdash \text{impossible} :: (x : A)} \text{ unsat}$$

There is no operational rule for this scenario since in well-typed configurations the process expression ‘impossible’ is dead code and can never be reached.

3.6 Type Safety

The main theorems that establish the deep connection between our refined type system and operational semantics are the usual *type preservation* and *progress*, also referred as *session fidelity* and *deadlock freedom*. At runtime, a program is represented using a set of semantic objects, i.e. processes and messages together defined as a *configuration*.

$$\mathcal{S} ::= \cdot \mid \mathcal{S}, \mathcal{S}' \mid \text{proc}(c, P) \mid \text{msg}(c, M)$$

We say that $\text{proc}(c, P)$ (or $\text{msg}(c, M)$) provide channel c . We stipulate that no two distinct semantic objects provide the same channel.

Type Preservation A key question then is how to *type configurations*? We define a well-typed configuration using the judgment $\Delta_1 \Vdash_{\Sigma} \mathcal{S} :: \Delta_2$ denoting that configuration \mathcal{S} uses channels Δ_1 and provides channels Δ_2 . The rules for typing a configuration are defined in Figure 3.2. A configuration is always typed w.r.t. a *well-formed signature* Σ , requiring that all (i) all type definitions are valid and contractive, and (ii) all process definitions are well-typed. Since the signature Σ is fixed, we elide it from the presentation.

The rule `emp` defines that an empty configuration provides all the channels Δ that it uses. The `comp` rule composes two configurations \mathcal{S}_1 and \mathcal{S}_2 ; \mathcal{S}_1 provides channels Δ_2 while \mathcal{S}_2 uses channels Δ_2 . The rule `proc` creates a configuration out of a single process. Configurations only exist at runtime where all arithmetic expressions in process terms are closed, i.e. they do not refer to any free variables. Hence, we use $\mathcal{V} = \cdot$ and $\mathcal{C} = \top$ when typing process P (premise in `proc` rule). Similar to `proc`, the rule `msg` creates a configuration out of a single message (where a message is also represented as a process).

Global Progress To state progress, we need the notion of a *poised process* [129]. A process `proc(c, P)` is poised if it is trying to receive a message on c . Dually, a message `msg(c, M)` is poised if it is sending along c . A configuration is poised if every message or process in the configuration is poised. Conceptually, this means that the configuration is trying to communicate *externally* along one of the channels it uses or provides.

Theorem 5 (Type Safety). *For a well-typed configuration $\Delta_1 \Vdash_{\Sigma} \mathcal{S} :: \Delta_2$:*

- (i) (Preservation) *If $\mathcal{S} \mapsto \mathcal{S}'$, then $\Delta_1 \Vdash_{\Sigma} \mathcal{S}' :: \Delta_2$*
- (ii) (Progress) *Either \mathcal{S} is poised, or $\mathcal{S} \mapsto \mathcal{S}'$.*

Proof. The proof of preservation proceeds by case analysis on the rules of operational semantics, applying inversion to the given typing derivation of \mathcal{S} , and then assembling a new derivation of \mathcal{S}' . Progress is proved by induction on the right-to-left typing of \mathcal{S} so that either \mathcal{S} is empty (and therefore poised) or $\mathcal{S} = (\mathcal{D}, \text{proc}(c, P))$ or $\mathcal{S} = (\mathcal{D}, \text{msg}(c, M))$. By induction hypothesis, \mathcal{D} can either take a step (and then so can \mathcal{S}), or \mathcal{D} is poised. In the latter case, we analyze the cases for P and M , applying multiple steps of inversion to show that in each case either \mathcal{S} can take a step or is poised. \square

3.7 Constraint Reconstruction

The process expressions introduced so far in the language follow simple syntax-directed typing rules. This means they are immediately amenable to be interpreted as an algorithm for type-checking, calling upon a decision procedure where arithmetic entailments and type equalities need to be verified. However, this requires the programmer to write a significant number of explicit process constructs pertaining to the refinement layer in their code. Relatedly, this hinders reuse: we are unable to provide multiple types to the same program so that it can be used in different contexts.

This section introduces an *implicit type system* in which the source program never contains the `assume` and `assert` expressions, i.e. constructs corresponding to proof constraints. Moreover,

$$\begin{array}{c}
\frac{\mathcal{V}; \mathcal{C} \vDash \phi \quad \mathcal{V}; \mathcal{C}; \Delta \vdash_i P :: (x : A)}{\mathcal{V}; \mathcal{C}; \Delta \vdash_i P :: (x : ?\{\phi\}. A)} \text{?R} \qquad \frac{\mathcal{V}; \mathcal{C} \wedge \phi; \Delta, (x : A) \vdash_i Q :: (z : C)}{\mathcal{V}; \mathcal{C}; \Delta, (x : ?\{\phi\}. A) \vdash_i Q :: (z : C)} \text{?L} \\
\\
\frac{\mathcal{V}; \mathcal{C} \wedge \phi; \Delta \vdash_i P :: (x : A)}{\mathcal{V}; \mathcal{C}; \Delta \vdash_i P :: (x : !\{\phi\}. A)} \text{!R} \qquad \frac{\mathcal{V}; \mathcal{C} \vDash \phi \quad \mathcal{V}; \mathcal{C}; \Delta, (x : A) \vdash_i Q :: (z : C)}{\mathcal{V}; \mathcal{C}; \Delta, (x : !\{\phi\}. A) \vdash_i Q :: (z : C)} \text{!L}
\end{array}$$

Figure 3.3: Implicit Typing Rules

impossible branches may be omitted from case expressions. The missing branches and other constructs are restored by a type-directed process of *reconstruction*.

In the first phase, a case expression with a missing branch for label ℓ is transformed into a branch $\ell \Rightarrow \text{impossible}$ so that type checking later verifies that the omitted branch is indeed impossible. Then assumes and asserts are inserted according to a reconstruction algorithm described in this section.

Following branch reconstruction, the resulting process expression is checked with the implicit typing judgment $\mathcal{V}; \mathcal{C}; \Delta \vdash_i P :: (x : A)$. The implicit system differs from the explicit system in only one way: for the implicit constructs related to constraints (!R, !L, ?R, ?L), the process expression does not change on application of these rules. Selected typing rules are described in Figure 3.3 and illustrate that expressions P and Q are unchanged in the premise and conclusion. For the remaining rules pertaining to base session types and quantifiers ($\exists R, \exists L, \forall R, \forall L$), no reconstruction is involved and the implicit rules exactly match the explicit rules.

The implicit rules are sound and complete with respect to the explicit system, since from an implicit typing derivation we can read off the corresponding explicit process expression and vice versa. The rules are also manifestly decidable since the types in the premise are smaller than the conclusion for all the rules presented.

However, the implicit type system is highly nondeterministic. Since the process expressions do not change on the application of implicit rules in Figure 3.3, they can be applied in many different orders. And *each valid order corresponds to a different explicit program*, intuitively changing the order in which constraints are sent and received. Thus, an implicit source program may correspond to many different explicit programs. The necessary backtracking would greatly complicate error messages and would also be exponential and severely inefficient.

To solve this problem, we introduce a novel *forcing calculus* which enforces an order among these implicit constructs. The core idea of this calculus is to follow the structure of each type, but within that *assume should be inserted as early as possible, and assert should be inserted as late as possible*. This reasoning is sound since the constraints obey a *monotonicity property*: if a constraint is true at a program point, it will always be true later in the program. Thus, eagerly assuming and lazily asserting constraints is sound: if a constraint can be proved now, it can be proved later. It is also complete under the mild assumption that the types can be polarized (explained below). Logically, the !R, ?L rules are invertible, and are applied eagerly while their dual rules are applied lazily.

This strategy is formally realized in the forcing calculus using the judgment $\mathcal{V}; \mathcal{C}; \Delta; \Omega \vdash P :: (x : A)$. The context is split into two: the linear context Δ contains stable propositions on

which the invertible left rules have been applied, while the ordered context Ω stores channels on which invertible rules can possibly still be applied to. First, we assign polarities to the type operators with implicit expressions, a notion borrowed from focusing [20] with a similar function here. Type definitions are unfolded in order to determine their polarity, which is always possible since type definitions are contractive. The types that involve communication are called *structural* and represented by S .

$$\begin{aligned}
A^+ &::= S \mid ?\{\phi\}. A^+ \\
A^- &::= S \mid !\{\phi\}. A^- \\
A &::= A^+ \mid A^- \\
S &::= \oplus\{\ell : A\}_{\ell \in L} \mid \&\{\ell : A\}_{\ell \in L} \mid A \otimes A \mid \mathbf{1} \mid A \multimap A \mid \exists n. A \mid \forall n. A
\end{aligned}$$

Not all types can be polarized in this manner, particularly types containing alternating proof constraints e.g., $!\{\phi\}. ?\{\psi\}. A$. When checking the validity of types before performing reconstruction we reject such types with alternating polarities. We also require that all process declarations contain only structural types at the top-level. Both these restrictions turn out to be mild in practice and can be resolved by introducing additional communications.

Thus, the $?$ operator is positive, while $!$ is negative. The structural types, denoted by S are considered neutral. In the forcing calculus, the invertible rules are applied first.

$$\begin{aligned}
&\frac{\mathcal{V} ; \mathcal{C} \wedge \phi ; \Delta^- ; \Omega \vdash P :: (x : A^-)}{\mathcal{V} ; \mathcal{C} ; \Delta^- ; \Omega \vdash P :: (x : !\{\phi\}. A^-)} !R \\
&\frac{\mathcal{V} ; \mathcal{C} \wedge \phi ; \Delta^- ; \Omega \cdot (x : A^+) \vdash P :: (z : C^+)}{\mathcal{V} ; \mathcal{C} ; \Delta^- ; \Omega \cdot (x : ?\{\phi\}. A^+) \vdash P :: (z : C^+)} ?L
\end{aligned}$$

If a negative type is encountered in the ordered context, it is considered stable (invertible rules applied) and moved to Δ^- .

$$\frac{\mathcal{V} ; \mathcal{C} ; \Delta^-, (x : A^-) ; \Omega \vdash P :: (z : C^+)}{\mathcal{V} ; \mathcal{C} ; \Delta^- ; \Omega \cdot (x : A^-) \vdash P :: (z : C^+)} \text{move}$$

The ordered context Ω imposes an order on the channels on which these invertible rules are applied.

Once all the invertible rules are applied, we reach a stable sequent which has the form $\mathcal{V} ; \mathcal{C} ; \Delta^- ; \cdot \vdash P :: (x : A^+)$, i.e., the ordered context is empty and the provided type A^+ is positive. A stable sequent implies that all constraints have been received. We send a constraint lazily, i.e., just before communicating on that channel. We realize this by *forcing* the channel just before communicating on it. As an example, while sending (or receiving) a label on channel x , we force it effectively sending any pending constraints.

$$\begin{aligned}
&\frac{\mathcal{V} ; \mathcal{C} ; \Delta^- ; \cdot \vdash x.k ; P :: [x : A^+]}{\mathcal{V} ; \mathcal{C} ; \Delta^- ; \cdot \vdash x.k ; P :: (x : A^+)} \oplus F_R \\
&\frac{\mathcal{V} ; \mathcal{C} ; \Delta, [x : A^-] ; \cdot \vdash \text{case } x (\ell \Rightarrow Q_\ell)_{\ell \in L} :: (z : C^+)}{\mathcal{V} ; \mathcal{C} ; \Delta, (x : A^-) ; \cdot \vdash \text{case } x (\ell \Rightarrow Q_\ell)_{\ell \in L} :: (z : C^+)} \oplus F_L
\end{aligned}$$

The square brackets $[\cdot]$ indicates that the channel is forced, indicating that a communication is about to happen on it. If there are assert constructs pending on the forced channel, they are applied now.

$$\frac{\mathcal{V}; \mathcal{C} \vDash \phi \quad \mathcal{V}; \mathcal{C}; \Delta^-; \cdot \vdash P :: [x : A^+]}{\mathcal{V}; \mathcal{C}; \Delta^-; \cdot \vdash P :: [x : ?\{\phi\}. A^+]} ?R$$

$$\frac{\mathcal{V}; \mathcal{C} \vDash \phi \quad \mathcal{V}; \mathcal{C}; \Delta^-, [x : A^-]; \cdot \vdash P :: (z : C^+)}{\mathcal{V}; \mathcal{C}; \Delta^-, [x : !\{\phi\}. A^-]; \cdot \vdash P :: (z : C^+)} !L$$

Finally, if a forced channel has a structural type, we apply the corresponding structural rule and *lose the forcing*. Again, as an example, we consider the internal choice operator.

$$\frac{(k \in L) \quad \mathcal{V}; \mathcal{C}; \Delta^-; \cdot \vdash P :: (x : A_k)}{\mathcal{V}; \mathcal{C}; \Delta^-; \cdot \vdash (x.k; P) :: [x : \oplus\{\ell : A_\ell\}_{\ell \in L}]} \oplus R_k$$

$$\frac{(\forall \ell \in L) \quad \mathcal{V}; \mathcal{C}; \Delta; (x : A_\ell) \vdash Q_\ell :: (z : C^+)}{\mathcal{V}; \mathcal{C}; \Delta, [x : \oplus\{\ell : A_\ell\}]; \cdot \vdash \text{case } x (\ell \Rightarrow Q_\ell) :: (z : C^+)} \oplus L$$

In either case, applying the structural rule creates a possibly unstable sequent, thereby restarting the inversion phase.

Remarkably, *the forcing calculus is sound and complete with respect to the implicit type system*, assuming types can be polarized. Since every rule in the forcing calculus is also present in the implicit system, it is trivially sound. Moreover, applying assume eagerly, and assert lazily also turns out to be complete due to the monotonicity property of constraints.

Theorem 6 (Soundness and Completeness). *For (valid) polarized types A and contexts Δ we have:*

1. *If $\mathcal{V}; \mathcal{C}; \Delta \vdash_i P :: (x : A)$, then $\mathcal{V}; \mathcal{C}; \cdot; \Delta \vdash P :: (x : A)$.*
2. *If $\mathcal{V}; \mathcal{C}; \cdot; \Delta \vdash P :: (x : A)$, then $\mathcal{V}; \mathcal{C}; \Delta \vdash_i P :: (x : A)$.*

Proof. Part (1) of Theorem 6 corresponds to soundness. The proof of soundness follows by induction on the implicit typing judgment. Intuitively, soundness follows from the simple observation that every rule in the forcing calculus is also valid in the implicit typing judgment. Theorem 6 part (2) corresponds to completeness whose proof proceeds by induction on the forcing judgment. The proof relies on two *key lemmas*: (i) the rules $!R$ and $?L$ are invertible, and (ii) if $\mathcal{V}; \mathcal{C}; \Delta^-; \Omega \vdash (x : A^+)$ and $\mathcal{V}; \mathcal{C} \vDash \phi$, then $\mathcal{V}; \mathcal{C}; \Delta^-; \Omega \vdash (x : ?\{\phi\}. A^+)$, i.e. asserting a constraint ϕ on a channel can be done at any program point where ϕ holds assuming \mathcal{C} , and thus, can be delayed. \square

If a process is well-typed in the implicit system, it is well-typed using the forcing calculus. Once the typing derivation, i.e., ordering of the typing rules is fixed by the forcing calculus, *a unique explicit program is constructed by applying the explicit typing rules to the derivation*. Thus, if a reconstruction is possible, the forcing calculus will find it! We use this calculus to reconstruct the explicit program, which is then typechecked using the explicit typing system.

Abstract Types	Concrete Types	Abstract Syntax	Concrete Syntax
$\oplus\{l : A, \dots\}$	$+\{l : A, \dots\}$	$x.k$	$x.k$
$\&\{l : A, \dots\}$	$\&\{l : A, \dots\}$	$\text{case } x (\ell \Rightarrow P)_{\ell \in L}$	$\text{case } x (l \Rightarrow P \mid \dots)$
$A \otimes B$	$A * B$	$\text{send } x w$	$\text{send } x w$
$A \multimap B$	$A \multimap B$	$y \leftarrow \text{recv } x$	$y \leftarrow \text{recv } x$
1	1	$\text{close } x$	$\text{close } x$
		$\text{wait } x$	$\text{wait } x$
$\exists n. A$	$?n. A$	$\text{send } x \{e\}$	$\text{send } x \{e\}$
$\forall n. A$	$!n. A$	$\{n\} \leftarrow \text{recv } x$	$\{n\} \leftarrow \text{recv } x$
$?\{\phi\}. A$	$?\{\text{phi}\}. A$	$\text{assert } x \{\phi\}$	$\text{assert } x \{\text{phi}\}$
$!\{\phi\}. A$	$!\{\text{phi}\}. A$	$\text{assume } x \{\phi\}$	$\text{assume } x \{\text{phi}\}$
$V[\bar{e}]$	$V\{e1\}\{e2\} \dots$	$x \leftrightarrow y$	$x \leftrightarrow y$
		$x \leftarrow f x_1 \dots x_n$	$x \leftarrow f x_1 \dots x_n$

Table 3.3: Abstract and Corresponding Concrete Syntax for Types and Expressions

3.8 Implementation

We have implemented a prototype for Rast in Standard ML (8100 lines of code). This implementation contains a lexer and parser (1200 lines), reconstruction engine (900 lines), an arithmetic solver (1200 lines), a type checker (2500 lines), pretty printer (400 lines), and an interpreter (200 lines). The source code is well-documented and available open-source [63].

Syntax Table 3.3 describes the syntax for Rast programs. Each row presents the abstract and concrete representation of a session type, and its corresponding providing expression. A program contains a series of mutually recursive type and process declarations and definitions.

```

type v{n}... = A
decl f{n}... : (x1 : A1) ... (xn : An) |- (x : A)
proc x <- f {n}... x1 ... xn = P

```

Listing 3.1: Top-Level Declarations

The first line is a *type definition*, where v is the type name with index variables \bar{n} and A is its definition. The second line is a *process declaration*, where f is the process name, $(x_1 : A_1) \dots (x_n : A_n)$ are the used channels and corresponding types, while the provided channel is x of type A . Finally, the last line is a *process definition* for the same process f defined using the process expression P . In addition, f can be parameterized by index variables \bar{n} . We use a hand-written lexer and shift-reduce parser to read an input file and generate the corresponding abstract syntax tree of the program. The reason to use a hand-written parser instead of a parser generator is to anticipate the most common syntax errors that programmers make and respond with the best possible error messages.

Validity Checking Once the program is parsed and its abstract syntax tree is extracted, we perform a validity check on it. We check that all index refinements, potentials, and delay op-

erators are non-negative. We also check that all index expressions are closed with respect to the the index variables in scope, and similarly for type expressions. To simplify and improve the efficiency of the type equality algorithm, we also assign internal names to type subexpressions [59, 64] parameterized over their free type and index variables. These internal names are not visible to the programmer.

Reconstruction and Type Checking The programmer can use a flag in the program file to indicate whether they are using *explicit* or *implicit* syntax. If the syntax is explicit, the reconstruction engine performs no program transformation. However, if the syntax is implicit, we use the implicit type system to approximately type-check the program. Once completed, we use the forcing calculus to insert assert and assume constructs.

The implementation takes some care to provide constructive and precise error messages, in particular as session types are likely to be unfamiliar. One technique is staging: first check approximate type correctness, ignoring refinements and only if that check passes perform reconstruction and strict type checking. Another particularly helpful technique has been *type compression*. Whenever the type checker expands a type $V[\bar{e}]$ with $V[\bar{n}] = B$ to $B[\bar{e}/\bar{n}]$, we record a reverse mapping back to $V[\bar{e}]$. When printing types for error messages this mapping is consulted, and complex types may be compressed to much simpler forms, greatly aiding readability of error messages. This is feasible in part because all intermediate subexpressions have an explicit (internal) definition, simplifying the lookup. Finally, our implementation uses a bi-directional [59, 64] type checking algorithm which reconstructs intermediate types for each channel. This helps localize the source of the error message as the program point where reconstruction fails. We designed the abstract syntax tree to also contain the relevant source code location information which is utilized while generating the error message.

Arithmetic Solver To determine the validity of arithmetic propositions that is used by our refinement layer, we use a straightforward implementation of Cooper’s decision procedure [52] for Presburger arithmetic. We found a small number of optimizations were necessary, but the resulting algorithm has been quite efficient and robust.

1. We eliminate constraints of the form $x = e$ (where x does not occur in e) by substituting e for x in all other constraints to reduce the total number of variables.
2. We exploit that we are working over natural numbers so all solutions have a natural lower bound, i.e., 0.

We also extend our solver to handle non-linear constraints. Since non-linear arithmetic is undecidable, in general, we use a normalizer which collects coefficients of each term in the multinomial expression.

1. To check $e_1 = e_2$, we normalize $e_1 - e_2$ and check that each coefficient of the normal form is 0.
2. To check $e_1 \geq e_2$, we normalize $e_1 - e_2$ and check that each coefficient is non-negative.
3. If we know that $x \geq c$, we substitute $y + c$ for x in the constraint that we are checking with the knowledge that the fresh $y \geq 0$.

4. We try to find a quick counterexample to validity by plugging in 0 and 1 for the index variables.

If the constraint does not fall in the above two categories, we print the constraint and trust that it holds. A user can then view these constraints manually and confirm their validity. At present, all of our examples pass without having to trust unsolvable constraints with our set of heuristics beyond Presburger arithmetic.

Interpreter The current version of the interpreter pursues a sequential schedule following a prior proposal [136]. We only execute programs that have no free type or index variables and only one externally visible channel, namely the one provided. When the computation finishes, the messages that were asynchronously sent along this distinguished channel are shown, while running processes waiting for input are displayed simply as a dash '-'.

The interpreter is surprisingly fast. For example, using a linear prime sieve to compute the status (prime or composite) or all number in the range [2, 257] takes 27.172 milliseconds using MLton during our experiments (see machine specifications below).

3.9 Further Examples

We present several different kinds of examples from varying domains illustrating different features of the type system and algorithms. Table 3.4 describes the results: iLOC describes the lines of source code in implicit syntax, eLOC describes the lines of code after reconstruction (which inserts implicit constructs), #Defs shows the number of process definitions, R (ms) and T (ms) show the reconstruction and type-checking time in milliseconds respectively. Note that reconstruction is faster than type-checking since reconstruction does not involve solving any arithmetic propositions. The experiments were run on an Intel Core i5 2.7 GHz processor with 16 GB 1867 MHz DDR3 memory.

1. **arithmetic**: natural numbers in unary and binary representation indexed by their value and processes implementing standard arithmetic operations.
2. **integers**: an integer counter represented using two indices x and y with value $x - y$.
3. **linlam**: expressions in the linear λ -calculus indexed by their size.
4. **list**: lists indexed by their size, and some standard operations such as *append*, *reverse*, *map*, *fold*, etc. Also provides and implementation of stacks and queues using lists.
5. **primes**: the sieve of Eratosthenes to classify numbers as prime or composite.
6. **segments**: the type $\text{seg}[n] = \forall k. \text{list}[k] \multimap \text{list}[n + k]$ representing partial lists with a constant-work append operation.
7. **ternary**: natural numbers and integers represented in balanced ternary form with digits 0, 1, -1, indexed by their value, and a few standard operations on them. This example is noteworthy since it is the only one stressing the arithmetic decision procedure.
8. **theorems**: processes representing valid circular [66] proofs of simple theorems such as $n(k + 1) = nk + n, n + 0 = n, n * 0 = 0$, etc.

Module	iLOC	eLOC	#Defs	R (ms)	T (ms)
arithmetic	395	619	29	0.959	5.732
integers	90	125	8	0.488	0.659
linlam	88	112	10	0.549	1.072
list	341	642	37	3.164	4.637
primes	118	164	11	0.289	4.580
segments	48	76	8	0.183	0.225
ternary	270	406	20	0.947	140.765
theorems	79	156	13	0.182	1.095
tries	243	520	13	2.122	6.408
Total	1672	2820	149	8.883	165.173

Table 3.4: Case Studies

9. **tries**: a trie data structure to store multisets of binary numbers, with constant amortized work insertion and deletion verified with ergonomic types.

We highlight interesting examples from some case studies showcasing the invariants that can be proved using arithmetic refinements and nested polymorphism.

Linear λ -Calculus We implemented the linear λ -calculus with evaluation (weak head normalization) of terms. We use higher-order abstract syntax, representing linear abstraction in the object language by a process receiving a message corresponding to its argument.

```
type exp = +{ lam : exp -o exp ,
             app : exp * exp }
```

We would like evaluation to return a value (a λ -abstraction), so we take advantage of the structural nature of types (allowing us to reuse the label **lam**) to define the value type.

```
type val = +{ lam : exp -o exp }
```

Rast can infer that **val** is a subtype of **exp**. We can derive constructors *apply* for expressions and *lambda* for values (we do not need the corresponding constructor for expressions).

```
decl apply : (e1 : exp) (e2 : exp) |- (e : exp)
proc e <- apply e1 e2 =
  e.app ; send e e1 ; e <-> e2
```

```
decl lambda : (f : exp -o exp) |- (v : val)
proc v <- lambda f = v.lam ; v <-> f
```

As a simple example, we follow the representation of a combinator that swaps the arguments to a function.

```
(* swap = \f. \x. \y. (f y) x *)
decl swap : . |- (e : exp)
proc e <- swap =
  e.lam ; f <- recv e ;
```

```

e.lam ; x <- recv e ;
e.lam ; y <- recv e ;
fy <- apply f y ;
e <- apply fy x

```

Evaluation is now the following very simple process.

```

decl eval : (e : exp) |- (v : val)
proc v <- eval e =
  case e ( lam => v <- lambda e
        | app => e1 <- recv e ;      % e = e2
              v1 <- eval e1 ;
              case v1 ( lam => send v1 e ;
                    v <- eval v1 ) )

```

If e sends a **lam** label, we just rebuild the expression as a value. If e sends an **app** label then e represents a linear application $e_1 e_2$ and the continuation has type $\text{exp} \otimes \text{exp}$. This means we *receive* a channel representing e_1 and the continuation (still called e) behaves like e_2 . We note this with a comment in the source. We then evaluate e_1 which exposes a λ -expression along the channel v_1 . We send e along v_1 , carrying out the reduction via communication. The result of this (still called v_1) is evaluated to yield the final value v . This program is available in the repository at `examples/linlam.rast`.

We would now like to prove that the value of a linear λ -expression is smaller than or equal to the original expression. At the same time we would like to rule out a class of so-called *exotic terms* in the representation, which are possible due to the presence of recursion in the meta-language. We achieve this by indexing the types `exp` and `val` with their *size*. For an application, this is easy: the size is one more than the sum of the sizes of the subterms.

```

type exp{n} = +{ lam : ...
               app : ?n1. ?n2. ?{n=n1+n2+1}. exp{n1} * exp{n2}}

```

The size $n_2 + 1$ of a λ -expression is one more than the size n_2 of its body, but what is that in our higher-order representation? The body is a linear function takes an expression of size n_1 and then behaves like an expression of size $n_1 + n_2$. Solving for n_2 then gives use the following type definitions and types for the constructor processes.

```

type exp{n} = +{lam : ?{n > 0}. !n1.exp{n1} -o exp{n1+n-1},
               app : ?n1. ?n2. ?{n=n1+n2+1}. exp{n1} * exp{n2}}

```

```

type val{n} = +{ lam : ?{n > 0}. !n1.exp{n1} -o exp{n1+n-1} }

```

```

decl apply{n1}{n2} :
  (e1 : exp{n1}) (e2 : exp{n2}) |- (e : exp{n1+n2+1})
decl lambda{n2} :
  (f : !n1. exp{n1} -o exp{n1+n2}) |- (v : val{n2+1})

```

The universal quantification over n_1 in the type `lam` is important, because a linear λ -expression may be applied to an argument of any size. We also cannot predict the size of the result of evaluation, so we have to use existential quantification: The value of an expression of size n will have size k for some $k \leq n$.

```
decl eval{n} : (e : exp{n}) |- (v : ?k. ?{k <= n}. val{k})
```

Because witnesses for quantifiers are not reconstructed, the evaluation process has to send and receive suitable sizes.

```
proc v <- eval{n} e =
  case e ( lam => send v {n} ;
          v <- lambda{n-1} e
        | app => {n1} <- recv e ;
                 {n2} <- recv e ;
                 e1 <- recv e ;
                 v1 <- eval{n1} e1 ;
                 {k2} <- recv v1 ;
                 case v1 ( lam => send v1 {n2} ;
                           send v1 e ;
                           v2 <- eval{n2+k2-1} v1 ;
                           {l} <- recv v2 ;
                           send v {l} ; v <-> v2))
```

Type-checking now verifies that if evaluation terminates, the resulting value is smaller than the expression (or of equal size if the expression is already a value). The repository contains the implementation in the file `examples/linlam-size.rast`.

Trie Data Structure We now implement multisets of natural numbers (in binary form). One of the key questions is how to maintain linearity in the design of the data structure and interface. For example, should we be able to delete an element from the trie, not knowing a priori if it is even *in* the trie? To avoid exceedingly complex types to account for these situations, the process maintaining a trie offers an interface with two operations: insert (label `ins`) and delete (label `del`). We index the type `trie{n}` with the number of elements in the trie, so inserting an element always increases n by 1. If the element is already present, we just add 1 to its multiplicity. Deleting an element actually removes all copies of it and returns its multiplicity m . If the element is *not* in the trie, we just return a multiplicity of $m = 0$. In either case, the trie contains $n - m$ elements afterwards.

```
type trie{n} =
  &{ins : !k. bin{k} -o trie{n+1},
    del : !k. bin{k} -o ?m. ?{m <= n}. bin{m} * trie{n-m}}
```

This type requires universal quantification over k , (written `!k`) which is the value of the number inserted into or deleted from the trie on each interaction (which is arbitrary).

The basic idea of the implementation is that each bit in the number $x : \text{bin}\{k\}$ addresses a subtrie: if it is `b0` we descend into the left subtrie, if it is `b1` we descent into the right subtrie. If it is `e` we have found (or constructed) the node corresponding to x and we either increase its multiplicity (for insert), or respond with its multiplicity and set the new multiplicity to zero (for delete). We have two forms of processes: a *leaf* with zero elements and an interior node with $n_0 + m + n_1$ elements (where n_0 and n_1 and the number of elements in the left and right subtrees, and m is the multiplicity of the number corresponding to this node in the trie).

```

decl leaf : . |- (t : trie{0})
decl node{n0}{m}{n1} :
  (l : trie{n0})(c : ctr{m})(r : trie{n1}) |- (t : trie{n0+m+n1})

```

The code is somewhat repetitive, so we only show the code for inserting an element into an interior node.

```

proc t <- node{n0}{m}{n1} l c r =
  case t (
    ins => {k} <- recv t ;
          x <- recv t ;
          case x ( b0 => {k'} <- recv x ;
                   l.ins ; send l {k'} ; send l x ;
                   t <- node{n0+1}{m}{n1} l c r
                | b1 => {k'} <- recv x ;
                   r.ins ; send r {k'} ; send r x ;
                   t <- node{n0}{m}{n1+1} l c r
                | e => wait x ;
                   c.inc ;
                   t <- node{n0}{m+1}{n1} l c r )
    | del => ...)

```

What does type-checking verify in this case? It shows that the number of elements in the trie increases and decreases as expected for each insert and delete operation. On the other hand, it does not verify that the *correct* multiplicities are incremented or decremented, which is beyond the reach of the current type system. The source code is at `examples/trie-work.rast`.

3.10 Related Work

The literature on session types is by now vast, so we focus our review of related work on *binary session types* (rather than multiparty session types) with *implementations* (rather than theoretical foundations). Among them, we can distinguish those that offer a library or embedding to a pre-existing language, and those that may be considered stand-alone language designs.

Libraries There are a number of libraries for session types. Such libraries tend to have a very different flavor from Rast because they focus on practical usability in the context of a general-purpose language. As such, the challenge usually is how to encode session types so programs can be statically checked against them and how to achieve the expected dynamic behavior. Among them we find libraries for Haskell [112, 125], Scala [143], OCaml [126], and Rust [99]. Noteworthy is the embedding of session types in ATS [163] because, unlike the others, ATS supports arithmetic indexing similar to Rast. The most recent library for Rust [49] is perhaps the closest to Rast in that it extends the exact basic system of session types from Chapter 2 with shared types [25]. While some of these libraries permit limited polymorphism, none of them support ergometric or temporal types.

Languages Designing complete languages like Rast frees the researcher from the limitations and idiosyncrasies of the host language as they explore the design space. A relatively early effort was the object-oriented language MOOL [157] which distinguishes linear and nonlinear channels.

A different style of language is SePi [24, 72] based on the π -calculus. It supports *linear* refinements in terms of uninterpreted propositions (which may reference integers) in addition to assert and assume primitives on them. They are not intended to capture internal properties of data structures of processes; instead, they allow the programmer to express some security properties.

The CO₂ middleware language [30, 31] supports *binary timed session types*. The notion of time here is *external*. As such, it does not measure work or span based on a cost model like Rast, but specifies interaction time windows for processes that can be enforced dynamically via monitors.

Concurrent C0 [160] is an implementation of linear and shared session types as an extension of C0, a small type-safe and memory-safe subset of C. It integrates the basic session types from Chapter 2 with shared session types [25] in the context of an imperative language.

Links [71, 113, 114] is a language aimed at developing web applications. While based on a different foundations, it is related to SILL [81, 153] in that both integrate traditional functional types with linear session types. As such, they can express many (nonlinear) programs that Rast cannot, but they support neither arithmetic refinements nor ergometric or temporal types.

Context-free session types [19, 151] generalize ordinary session types with sequential composition as well as permitting some polymorphism. The linear sublanguage of context-free session types can be modeled in Rast with nested polymorphism [64].

3.11 Conclusion

This chapter describes the Rast programming language. In particular, we focused on the concrete syntax, type checking and equality, and the refinement layer [58, 59]. The refinements rely on an arithmetic solver based on Cooper’s algorithm [52]. The interpreter uses the shared memory semantics introduced in recent work [136]. We concluded with several examples demonstrating the efficacy of the refined type system in expressing and verifying properties about data structure sizes and values. All our examples have been verified with our system, and are available in an open-source repository [63].

In the future, we plan to address some limitations of the Rast language. One goal of Rast was to explore the boundaries of purely linear programming with general recursion. Often, this imposes a certain programming discipline and can be inconvenient if we need to drop or duplicate channels. Recent work on adjoint logic [135] uniformly integrates different logical layers into a unified language by assigning modes to communication. We plan to utilize this adjoint formulation to support shared [25] and unrestricted channels. Prior work on SILL [81] has demonstrated such an integration is helpful in general-purpose programming. With respect to refinements, we intend to pursue richer constraint domains such as non-linear arithmetic, particularly SMT.

Chapter 4

Work Analysis

Chapter 3 introduces refinement session types and describes its metatheory and implementation. This chapter crucially relies on refined types to express the *work* bounds of concurrent programs. The chapter introduces the novel type system and cost semantics, proves its soundness, and evaluates the system on a series of standard benchmarks.

4.1 Introduction

This chapter studies the foundations of worst-case resource analysis for session-typed programs. The key idea here is to rely on *resource-aware session types* to describe the resource bounds for inter-process communication. We extend session types to not only exchange messages, but also potential along a channel. The potential (in the sense of classical amortized analysis) may be spent by sending other messages as part of the network of interacting processes, or maintained locally for future interactions. Resource analysis is static, using the type system, and the runtime behavior of programs is not affected.

Here, I mainly focus on bounds on the total work performed by a system, counting the number of messages that are exchanged. While this alone does not yet account for the concurrent nature of message-passing programs, it constitutes a significant and necessary first step. The derived bounds are also useful in their own right. For example, the information can be used in scheduling decisions, to derive the number of messages that are sent along a specific channel, or to statically decide whether we should spawn a new thread of control or execute sequentially when possible. Additionally, bounds on the work of a process also serve as input to a Brent-style theorem [39] that relates the complexity of the execution of a program on a k -processor machine to the program's work (this chapter) and span (next chapter).

The analysis is based on a linear type system that extends standard session types with two new type constructors, one to receive potential (\langle^r) and one to send potential (\triangleright^r). The superscript r declares the amount of potential that must be transferred (conceptually!). Since the interface to a process is characterized entirely by the resource-aware session types of the channels it interacts with, this design provides a compositional resource specification. For closed programs (which evolve into a closed network of interacting processes), the bound becomes a single constant.

A conceptual challenge is to express symbolic bounds in a setting without static data structures and intrinsic sizes. The innovation is that resource-aware session types describe bounds as functions of interactions (messages sent) on a channel. A major technical challenge is to account for the global number of messages sent with local derivation rules: operationally, local message counts are forwarded to a parent process when a sub-process terminates. As a result, local message counts are incremented by sub-processes in a non-local fashion. My solution is that messages and processes carry potential to amortize the cost of a terminating sub-process proactively as a side-effect of the communication.

The main contributions are as follows. I present the first session type system for deriving parametric bounds on the resource usage of message-passing processes. I also prove the non-trivial soundness of the type system with respect to an operational cost semantics that tracks the total number of messages exchanged in a network of communicating processes. I also demonstrate the effectiveness of the technique by deriving tight bounds for some standard examples of amortized analysis from the literature on session types. I also show how resource-aware session types can be used to specify and compare the performance characteristics of different implementations of the same protocol. The analysis is currently manual, with automation left for future work.

4.2 Overview

This section will motivate and informally introduce resource-aware session types and show how they can be used to analyze the resource usage of message-passing processes. I describe an implementation of a counter and use resource-aware session types to analyze its resource usage. Like in the rest of this chapter, the resource we are interested in is the total number of messages sent along all channels in the system.

As a first simple example, I consider natural numbers in binary form. A process *providing* a natural number sends a stream of bits starting with the least significant bit. These bits are represented by messages zero and one, eventually terminated by \$.

$$\text{bits} = \oplus\{\text{zero} : \text{bits}, \text{one} : \text{bits}, \$: \mathbf{1}\}$$

For instance, the number $6 = (110)_2$ would be represented by the sequence of messages zero, one, one, \$, *close*. A client of a channel $c : \text{bits}$ has to branch on whether it receives zero, one, or \$. As a second example, I describe the interface to a counter. A client can repeatedly send *inc* messages to a counter, until they want to read its value and send *val*. At that point the counter will send a stream of bits representing its value as prescribed by the type *bits*.

$$\text{ctr} = \&\{\text{inc} : \text{ctr}, \text{val} : \text{bits}\}$$

A well-known example of amortized analysis counts the number of bits that must be flipped to increment a counter. It turns out the amortized cost per increment is 2, so that n increments require at most $2n$ bits to be flipped. This is observed by introducing a potential of 1 for every bit that is 1 and using this potential to *pay* for the expensive case in which an increment triggers many flips. When the lowest bit is zero, it is flipped to one (costing 1) and a remaining potential




Figure 4.1: Binary counter system representing $5 = (101)_2$ with messages triggered when inc message is received on s_4 .

of 1 is also stored with this bit. When the lowest bit is one, the stored potential is used to flip the bit back to zero (with no stored potential) and the remaining potential of 2 is passed along for incrementing the higher bits.

A binary counter is modeled as a chain of processes where each process represents a single bit (process b_0 or b_1) with a final process e at the end. Each of the processes in the chain *provides* a channel of the ctr type, and each (except the last) also *uses* a channel of this type representing the higher bits. For example, in the first chain in Figure 4.1, the process b_0 offers along channel s_3 (indicated by \bullet between b_0 and s_3) and uses channel s_2 . This is formally written as

$$\begin{array}{c} \cdot \vdash e :: (s_1 : \text{ctr}) \quad s_1 : \text{ctr} \vdash b_1 :: (s_2 : \text{ctr}) \\ s_2 : \text{ctr} \vdash b_0 :: (s_3 : \text{ctr}) \quad s_3 : \text{ctr} \vdash b_1 :: (s_4 : \text{ctr}) \end{array}$$

The *definitions* of e , b_0 , and b_1 can be found in Figures 4.3 and 4.4. The only channel visible to an outside client (not shown) is s_4 . Figure 4.1 shows the messages triggered if an increment message is received along s_4 .

Expressing resource bounds. My basic approach is that *messages carry potential* and *processes store potential*. This means the sender has to pay not just 1 unit for sending the message, but whatever additional units to amortize future costs. In the amortized analysis of the counter, each bit flip corresponds exactly to an `inc` message, because that is what triggers a bit to be flipped. My cost model focuses on messages as prescribed by the session type and does not count other operations, such as spawning a new process or terminating a process. This choice is not essential to the approach, but convenient here.

To capture the informal analysis we need to express *in the type* that we have to send 1 unit of potential right after the label `inc`. We do this using the \triangleleft operator indicating the required potential with the superscript, postponing the discussion of `val`.

$$\text{ctr} = \&\{\text{inc} : \triangleleft^1 \text{ctr}, \text{val} : \triangleleft^? \text{bits}\}$$

When types are assigned to processes, we use the more expressive resource-aware session types. We indicate the potential stored in a particular process as a superscript on the turnstile.

$$t : \text{ctr} \vdash^0 b0 :: (s : \text{ctr}) \quad (4.1)$$

$$t : \text{ctr} \vdash^1 b1 :: (s : \text{ctr}) \quad (4.2)$$

$$\cdot \vdash^0 e :: (s : \text{ctr}) \quad (4.3)$$

These typing constraints can be verified using the typing rules of the system, using the definitions of $b0$, $b1$, and e . Informally, the reason is as follows:

$b0$: After $b0$ receives `inc` it receives 1 unit of potential. It continues as $b1$ (which requires no communication) which stores this 1 unit (as prescribed from the type of $b1$ in Equation 4.2).

$b1$: After $b1$ receives `inc` it receives 1 unit of potential which, when combined with the stored one, makes 2 units. It sends an `inc` message which consumes 1 unit, followed by sending a unit potential, thereby consuming the 2 units. It has no remaining potential, which is sufficient because it transitions to $b0$ which stores no potential (inferred from the type of $b0$ in Equation 4.1).

e : After e receives `inc` it receives 1 unit of potential. It spawns a new process e and continues as $b1$. Spawning a process is free, and e requires no potential, so it can store the potential it received with $b1$ as required.

How do we handle the type annotation $\text{val} : \triangleleft^? \text{bits of the label val}$? Recall that $\text{bits} = \oplus\{\text{zero} : \text{bits}, \text{one} : \text{bits}, \$: 1\}$. In our implementation, upon receiving a `val` message, a $b0$ or $b1$ process will first respond with `zero` or `one` respectively. It then sends `val` along the channel it uses (representing the higher bits of the number) and terminates by *forwarding* further communication to the higher bits in the chain. Figure 4.2 demonstrates the messages triggered when `val` message is received along s_4 . The e process will just send `$` and `close`, indicating the empty stream of bits.

There will be enough potential to carry out the required send operations if each process ($b0$, $b1$, and e) carries an additional 2 units of potential. These could be imparted with the `inc` and `val` messages by sending 2 more units with `inc` and 2 units with `val`. That is, the following type is valid:

$$\begin{aligned} \text{bits} &= \oplus\{\text{zero} : \text{bits}, \text{one} : \text{bits}, \$: 1\} \\ \text{ctr} &= \&\{\text{inc} : \triangleleft^3 \text{ctr}, \text{val} : \triangleleft^2 \text{bits}\} \end{aligned}$$

However, this type is a gross over-approximation! The processes of a counter of value n , would carry $2n$ additional potential while only $2 \lceil \log(n + 1) \rceil + 2$ are needed. To obtain this more precise bound, we need to define a more *refined* type.

A more precise analysis. This requires that, *in the type*, either the number of bits in the representation of a number or its value can be referred to. This form of internal measure is needed only for type-checking purposes, not at runtime. It is also not intrinsically tied to a property of a representation, the way the length of a list in a functional language is tied to its




Figure 4.2: Binary counter system representing $5 = (101)_2$ with messages triggered when val message is received on s_4 .

memory requirements. We have already set the stage for indexing types using refinements in Chapter 3. We now employ refinement type $\text{ctr}[n]$ to denote a counter of value n . Following the reasoning above, we obtain the following type:

$$\begin{aligned} \text{bits} &= \oplus\{\text{zero} : \text{bits}, \text{one} : \text{bits}, \$: \mathbf{1}\} \\ \text{ctr}[n] &= \&\{\text{inc} : \triangleleft^1 \text{ctr}[n + 1], \text{val} : \triangleleft^{2 \lceil \log(n+1) \rceil + 2} \text{bits}\} \end{aligned}$$

To check the types of our implementation, we need to revisit and refine the typing of the b_0 , b_1 and e processes.

$$\begin{aligned} b_0[n] &:: (t : \text{ctr}[n]) \vdash^0 (s : \text{ctr}[2n]) \\ b_1[n] &:: (t : \text{ctr}[n]) \vdash^1 (s : \text{ctr}[2n + 1]) \\ e &:: \cdot \vdash^0 (s : \text{ctr}[0]) \end{aligned}$$

The type system verifies these types against the implementation of b_0 , b_1 , and e (see Figures 4.3 and 4.4, potential annotations marked in red). I will briefly explain the type derivation of b_0 , as shown in Figure 4.3 after the $\%$. After receiving the inc message, the b_0 process receives a unit potential on s using the \triangleleft^1 type constructor. This constructor is accompanied by the get construct (line 4) which receives the unit potential which is stored in the process, as indicated by the number on the turnstile. The type on the right exactly matches b_1 's type in the signature, thereby making the call to b_1 valid. Similarly, b_0 receives $2 \lceil \log(2n + 1) \rceil + 2$ units of potential after receiving the val .

The cost model of interest in this chapter counts the total number of messages exchanged in the system. This is realized formally by consuming a unit of potential before every message sent. The corresponding construct is $\text{work}\{1\}$, as indicated in line 7. This consumes a unit of potential, as indicated by the type on the right. A unit potential is similarly consumed on line 9 before sending the val message on t (line 10). The dual to the get construct is pay . This is used to send potential on a channel, as indicated on line 11, consuming potential stored in the process. Finally, the b_0 process remains with no potential and can successfully terminate by forwarding. Note that a process is not allowed to terminate while it stores potential as that would violate

```

1:  $b0[n] :: (t : \text{ctr}[n]) \vdash^0 (s : \text{ctr}[2n])$ 
2:    $s \leftarrow b0[n] t =$ 
3:   case  $s$ 
4:     (inc  $\Rightarrow$  get  $s \{1\}$  ;                                $\% (t : \text{ctr}[n]) \vdash^1 s : \text{ctr}[2n + 1]$ 
5:        $s \leftarrow b1[n] t$ 
6:     | val  $\Rightarrow$  get  $s \{2 \lceil \log(2n + 1) \rceil + 2\}$  ;  $\% (t : \text{ctr}[n]) \vdash^{2 \lceil \log(2n+1) \rceil + 2} s : \text{bits}$ 
7:       work  $\{1\}$  ;                                        $\% (t : \text{ctr}[n]) \vdash^{2 \lceil \log(2n+1) \rceil + 2 - 1} s : \text{bits}$ 
8:        $s.\text{zero}$  ;                                        $\% (t : \text{ctr}[n]) \vdash^{2 \lceil \log(2n+1) \rceil + 1} s : \text{bits}$ 
9:       work  $\{1\}$  ;                                        $\% (t : \text{ctr}[n]) \vdash^{2 \lceil \log(2n+1) \rceil + 1 - 1} s : \text{bits}$ 
10:       $t.\text{val}$  ;                                          $\% (t : \text{ctr}[n]) \vdash^{2 \lceil \log(2n+1) \rceil} s : \text{bits}$ 
11:      pay  $t \{2 \lceil \log(n + 1) \rceil + 2\}$  ;  $\% (t : \text{bits}) \vdash^{2 \lceil \log(2n+1) \rceil - 2 \lceil \log(n+1) \rceil - 2} s : \text{bits}$ 
12:       $s \leftrightarrow t$                                         $\% (t : \text{bits}) \vdash^0 s : \text{bits}$ 

```

Figure 4.3: Implementation for $b0$ process with its type derivation.

the linearity constraint on the potential. The derivations for $b1$ and e are similar and described in Figure 4.4.

The typing rules reduce the well-typedness of these processes to arithmetic inequalities which can be solved by hand, for example, using that $\log(2n) = \log(n) + 1$. The intrinsic measure n and the precise potential annotations are not automatically derived, but come from our insight about the nature of the algorithms.

The typing derivation provides a proof certificate on the resource bound for a process. For closed processes typed as

$$\cdot \vdash^p Q :: (c : \mathbf{1})$$

the number p provides a worst case bound on the number of messages sent during computation of Q , which always ends with the process sending *close* along c , indicating termination.

4.3 Operational Cost Semantics

The cost semantics for standard session types is augmented to track the total work performed by the system. The work is tracked by the local counter w in $\text{proc}(c, w, P)$ and $\text{msg}(c, w, M)$ propositions. For process P , w maintains the total work performed by P so far. When a process executes the work $\{c\}$ construct, its counter w is incremented by c . When a process terminates, the respective predicate is removed from the configuration, but its work done is preserved. A process can terminate either by sending a close message, or by forwarding. In either case, the process' work is conveniently preserved in the msg predicate to pass it on to the client process.

The cost semantics is parametric in the cost model. That is, the programmer can specify the resource they intend to measure. This cost model is then realized by inserting a work construct before the respective expressions. For instance, inserting a work $\{1\}$ before each send will count the total number of messages exchanged.

The semantics is defined in Figure 4.5. Each rule consumes the propositions to the left of \mapsto and produces the proposition to its right. The rules $\text{cut}C$ and $\text{def}C$ spawn a new process

```

13:  $b1[n] :: (t : \text{ctr}[n]) \vdash^1 (s : \text{ctr}[2n + 1])$ 
14:  $s \leftarrow b1[n] t =$ 
15:   case  $s$ 
16:     (inc  $\Rightarrow$  get  $s \{1\}$  ;  $\% (t : \text{ctr}[n]) \vdash^2 s : \text{ctr}[2n + 2]$ 
17:       work  $\{1\}$  ;  $\% (t : \text{ctr}[n]) \vdash^{2-1} s : \text{ctr}[2n + 2]$ 
18:        $t.\text{inc}$  ;  $\% (t : \triangleleft^1 \text{ctr}[n + 1]) \vdash^1 s : \text{ctr}[2n + 2]$ 
19:       pay  $t \{1\}$  ;  $\% (t : \text{ctr}[n + 1]) \vdash^{1-1} s : \text{ctr}[2n + 2]$ 
20:        $s \leftarrow b0[n + 1] t$ 
21:     | val  $\Rightarrow$  get  $s \{2 \lceil \log(2n + 2) \rceil + 2\}$  ;  $\% (t : \text{ctr}[n]) \vdash^{2 \lceil \log(2n+2) \rceil + 2} s : \text{bits}$ 
22:       work  $\{1\}$  ;  $\% (t : \text{ctr}[n]) \vdash^{2 \lceil \log(2n+2) \rceil + 2 - 1} s : \text{bits}$ 
23:        $s.\text{one}$  ;  $\% (t : \text{ctr}[n]) \vdash^{2 \lceil \log(2n+2) \rceil + 1} s : \text{bits}$ 
24:       work  $\{1\}$  ;  $\% (t : \text{ctr}[n]) \vdash^{2 \lceil \log(2n+2) \rceil + 1 - 1} s : \text{bits}$ 
25:        $t.\text{val}$  ;  $\% (t : \triangleleft^{2 \lceil \log(n+1) \rceil + 2} \text{bits}) \vdash^{2 \lceil \log(2n+2) \rceil} s : \text{bits}$ 
26:       pay  $t \{2 \lceil \log(n + 1) \rceil + 2\}$  ;  $\% (t : \text{bits}) \vdash^{2 \lceil \log(2n+2) \rceil - 2 \lceil \log(n+1) \rceil - 2} s : \text{bits}$ 
27:        $s \leftrightarrow t$  ;  $\% (t : \text{bits}) \vdash^0 s : \text{bits}$ 
28:  $e :: \cdot \vdash^0 (s : \text{ctr}[0])$ 
29:  $s \leftarrow e =$ 
30:   case  $s$ 
31:     (inc  $\Rightarrow$  get  $s \{1\}$  ;  $\% \cdot \vdash^1 s : \text{ctr}[0 + 1]$ 
32:        $t \leftarrow e$  ;  $\% (t : \text{ctr}[0]) \vdash^1 s : \text{ctr}[1]$ 
33:        $s \leftarrow b1[0] t$ 
34:     | val  $\Rightarrow$  get  $s \{2 \lceil \log(0 + 1) \rceil + 2\}$  ;  $\% \cdot \vdash^{2 \lceil \log(0+1) \rceil + 2} s : \text{bits}$ 
35:       work  $\{1\}$  ;  $\% \cdot \vdash^{2-1} s : \text{bits}$ 
36:        $s.\$$  ;  $\% \cdot \vdash^1 s : \mathbf{1}$ 
37:       work  $\{1\}$  ;  $\% \cdot \vdash^{1-1} s : \mathbf{1}$ 
38:       close  $s$  ;  $\% \cdot \vdash^0 s : \mathbf{1}$ 

```

Figure 4.4: Implementations for $b1$ and e processes with their type derivations.

with 0 work (as it has not done any work so far), while Q_c continues with the same amount of work. A forwarding process transfers its work to a corresponding message and terminates after identifying the channels, as described in rules id^+C and id^-C . All other communication rules create a message with work 0, which is then later received by its recipient, thereby transferring the work done by the message (which it gathered by possibly interacting with forwarding processes). The standard semantics rules can be obtained by simply deleting the work counters.

Work counter can be incremented only by executing the work construct.

$$(\text{work}) \quad \text{proc}(c, w, \text{work } \{w'\} ; P) \mapsto \text{proc}(c, w + w', P)$$

Finally, the two type constructors \triangleright and its dual \triangleleft are used to exchange potential. The potential

$$\begin{array}{l}
(\text{cut}C) \quad \text{proc}(d, w, x \leftarrow P_x ; Q_x) \mapsto \text{proc}(c, 0, [c/x]P_x), \text{proc}(d, w, [c/x]Q_x) \quad (c \text{ fresh}) \\
(\text{def}C) \quad \text{proc}(d, w, x \leftarrow f \leftarrow \bar{e} ; Q) \mapsto \\
\quad \text{proc}(c, 0, [c/x, \bar{e}/\Delta]P), \text{proc}(d, w, [c/x]Q) \quad (c \text{ fresh}) \\
(\text{id}^+C) \quad \text{msg}(d, w, M), \text{proc}(c, w', c \leftrightarrow d) \mapsto \text{msg}(c, w + w', [c/d]M) \\
(\text{id}^-C) \quad \text{proc}(c, w, c \leftrightarrow d), \text{msg}(e, w', M(c)) \mapsto \text{msg}(e, w + w', [d/c]M(c)) \\
(\oplus S) \quad \text{proc}(c, w, c.k ; P) \mapsto \text{proc}(c', w, [c'/c]P), \text{msg}(c, 0, c.k ; c \leftrightarrow c') \quad (c' \text{ fresh}) \\
(\oplus C) \quad \text{msg}(c, w, c.k ; c \leftrightarrow c'), \text{proc}(d, w', \text{case } c (\ell \Rightarrow Q_\ell)_{\ell \in L}) \mapsto \\
\quad \text{proc}(d, w + w', [c'/c]Q_k) \\
(\& S) \quad \text{proc}(d, w, c.k ; Q) \mapsto \text{msg}(c', 0, c.k ; c' \leftrightarrow c), \text{proc}(d, w, [c'/c]Q) \quad (c' \text{ fresh}) \\
(\& C) \quad \text{proc}(c, w, \text{case } c (\ell \Rightarrow Q_\ell)_{\ell \in L}), \text{msg}(c', w', c.k ; c' \leftrightarrow c) \mapsto \\
\quad \text{proc}(c', w + w', [c'/c]Q_k) \\
(\otimes S) \quad \text{proc}(c, w, \text{send } c e ; P) \mapsto \\
\quad \text{proc}(c', w, [c'/c]P), \text{msg}(c, 0, \text{send } c e ; c \leftrightarrow c') \quad (c' \text{ fresh}) \\
(\otimes C) \quad \text{msg}(c, w, \text{send } c e ; c \leftrightarrow c'), \text{proc}(d, w', x \leftarrow \text{recv } c ; Q) \mapsto \\
\quad \text{proc}(d, w + w', [c', e/c, x]Q) \\
(-\circ S) \quad \text{proc}(d, w, \text{send } c e ; Q) \mapsto \\
\quad \text{msg}(c', 0, \text{send } c e ; c' \leftrightarrow c), \text{proc}(d, w, [c'/c]Q) \quad (c' \text{ fresh}) \\
(-\circ C) \quad \text{proc}(c, w, x \leftarrow \text{recv } c), \text{msg}(c', w', \text{send } c e ; c' \leftrightarrow c) \mapsto \\
\quad \text{proc}(c', w + w', [c', d/c, x]P) \\
(\mathbf{1}S) \quad \text{proc}(c, w, \text{close } c) \mapsto \text{msg}(c, w, \text{close } c) \\
(\mathbf{1}C) \quad \text{msg}(c, w, \text{close } c), \text{proc}(d, w', \text{wait } c ; Q) \mapsto \text{proc}(d, w + w', Q)
\end{array}$$

Figure 4.5: Cost semantics tracking total work for programs

is only a theoretical construct, and potentials have no role to play at runtime.

$$\begin{array}{l}
(\triangleright S) \quad \text{proc}(c, w, \text{pay } c \{r\} ; P) \mapsto \\
\quad \text{proc}(c', w, [c'/c]P), \text{msg}(c, 0, \text{pay } c \{r\} ; c \leftrightarrow c') \quad (c' \text{ fresh}) \\
(\triangleright C) \quad \text{msg}(c, w, \text{pay } c \{r\} ; c \leftrightarrow c'), \text{proc}(d, w', \text{get } c \{r\} ; Q) \mapsto \\
\quad \text{proc}(d, w + w', [c'/c]Q) \\
(\triangleleft S) \quad \text{proc}(d, w, \text{pay } c \{r\} ; Q) \mapsto \\
\quad \text{msg}(c', 0, \text{pay } c \{r\} ; c' \leftrightarrow c), \text{proc}(d, w, [c'/c]Q) \quad (c' \text{ fresh}) \\
(\triangleleft C) \quad \text{proc}(c, w, \text{get } c \{r\}), \text{msg}(c', w', \text{pay } c \{r\} ; c' \leftrightarrow c) \mapsto \\
\quad \text{proc}(c', w + w', [c'/c]P)
\end{array}$$

The rules of the cost semantics are successively applied to a configuration until the configuration becomes empty or the configuration is stuck and none of the rules can be applied. At any point in this local stepping, the total work performed by the system can be obtained by summing the local counters w for each predicate in the configuration.

$$\frac{\mathcal{V}; \mathcal{C} \vDash q = 0}{\mathcal{V}; \mathcal{C}; y : A \vdash^q x \leftrightarrow y :: (x : A)} \text{id}$$

$$\frac{\mathcal{V}; \mathcal{C} \vDash r = p + q \quad \Delta \vdash^p f[\bar{n}] = P :: (x : A) \in \Sigma \quad \Delta_1 =_\alpha \Delta[\bar{e}/\bar{n}] \quad \mathcal{V}; \mathcal{C}; \Delta_2, (x : A) \vdash^q Q_x :: (z : C)}{\mathcal{V}; \mathcal{C}; \Delta_1, \Delta_2 \vdash^r (x \leftarrow f \bar{y} = Q_x) :: (z : C)} \text{spawn}$$

Figure 4.6: Typing rules for resource-aware session types for id, spawn

4.4 Type System

The typing judgment has the form

$$\mathcal{V}; \mathcal{C}; \Delta \vdash_\Sigma^q P :: (x : A)$$

Intuitively, the judgment describes a process in state P using the context Δ and signature Σ and providing service along channel x of type A . In other words, P is the provider for channel $x : A$, and a client for all the channels in Δ . The resource annotation q is a natural number and defines the potential stored in the process P .

Σ defines the signature containing type and process definitions. It is defined as a finite set of type definitions $V = A$ and process definitions $\Delta \vdash^q f[\bar{n}] = P :: (x : A)$. The equation $V = A$ is used to define the type variable V as type expression A . We treat such definitions *equirecursively*. The process definition $\Delta \vdash^q f[\bar{n}] = P :: (x : A)$ defines a (possibly recursive) process named f parameterized by index variables \bar{n} implemented by P providing along channel $x : A$ and using the channels Δ as a client, storing potential q . Because the signature is fixed, it is elided from the presentation of the rules.

Figures 4.7, 4.6 describe the usual typing rules for our system. The interesting rules here are spawn and id. The spawn splits the potential $r = p + q$, and provides potential p to the spawned process, and q to the continuation. A forwarding process $x \leftrightarrow y$ must be typed with no potential as it is about to terminate. The rest of the rules are standard and I am omitting their discussion. Deleting the potential annotation from the process typing judgment recovers the typing rules for standard session types. Messages are typed exactly as processes.

In addition, the language has explicit rules for consuming and transfer of potential. Executing the work $\{w\}$ construct consumes w (non-negative) units from the potential stored in a process. Thus, a process must have at least w units of potential to execute this construct. This is expressed in the rule with the annotation $q + w$ in the conclusion.

$$\frac{\mathcal{V}; \mathcal{C}; \Delta \vdash^q P :: (x : A)}{\mathcal{V}; \mathcal{C}; \Delta \vdash^{q+w} \text{work } \{w\}; P :: (x : A)} \text{work}$$

Similarly, executing a $\text{pay } x \{r\}$ consumes r units from the process potential, while $\text{get } x \{r\}$ provides r units to the process potential. The main innovation here is the introduction of the two dual type operators, \triangleright and \triangleleft . The \triangleright operator expresses that the provider must pay potential

$$\begin{array}{c}
\frac{\mathcal{V}; \mathcal{C}; \Delta \Vdash P :: (x : A_k) \quad (k \in L)}{\mathcal{V}; \mathcal{C}; \Delta \Vdash (x.k; P) :: (x : \oplus\{\ell : A_\ell\}_{\ell \in L})} \oplus R \\
\\
\frac{\mathcal{V}; \mathcal{C}; \Delta, (x : A_\ell) \Vdash Q_\ell :: (z : C) \quad (\forall \ell \in L)}{\mathcal{V}; \mathcal{C}; \Delta, (x : \oplus\{\ell : A_\ell\}_{\ell \in L}) \Vdash \text{case } x (\ell \Rightarrow Q_\ell)_{\ell \in L} :: (z : C)} \oplus L \\
\\
\frac{\mathcal{V}; \mathcal{C}; \Delta \Vdash P_\ell :: (x : A_\ell) \quad (\forall \ell \in L)}{\mathcal{V}; \mathcal{C}; \Delta \Vdash \text{case } x (\ell \Rightarrow P_\ell)_{\ell \in L} :: (x : \&\{\ell : A_\ell\}_{\ell \in L})} \& R \\
\\
\frac{\mathcal{V}; \mathcal{C}; \Delta, (x : A_k) \Vdash Q :: (z : C)}{\mathcal{V}; \mathcal{C}; \Delta, (x : \&\{\ell : A_\ell\}_{\ell \in L}) \Vdash x.k; Q :: (z : C)} \& L \\
\\
\frac{\mathcal{V}; \mathcal{C}; \Delta, (y : A) \Vdash P_y :: (x : B)}{\mathcal{V}; \mathcal{C}; \Delta \Vdash (y \leftarrow \text{recv } x; P_y) :: (x : A \multimap B)} \multimap R \\
\\
\frac{\mathcal{V}; \mathcal{C}; \Delta, (x : B) \Vdash Q :: (z : C)}{\mathcal{V}; \mathcal{C}; \Delta, (w : A), (x : A \multimap B) \Vdash (\text{send } x w; Q) :: (z : C)} \multimap L \\
\\
\frac{\mathcal{V}; \mathcal{C}; \Delta \Vdash P :: (x : B)}{\mathcal{V}; \mathcal{C}; \Delta, (w : A) \Vdash \text{send } x w; P :: (x : A \otimes B)} \otimes R \\
\\
\frac{\mathcal{V}; \mathcal{C}; \Delta, (y : A), (x : B) \Vdash Q_y :: (z : C)}{\mathcal{V}; \mathcal{C}; \Delta, (x : A \otimes B) \Vdash y \leftarrow \text{recv } x; Q_y :: (z : C)} \otimes L \\
\\
\frac{\mathcal{V}; \mathcal{C} \models q = 0}{\mathcal{V}; \mathcal{C}; \cdot \Vdash \text{close } x :: (x : \mathbf{1})} \mathbf{1}R \quad \frac{\mathcal{V}; \mathcal{C}; \Delta \Vdash Q :: (z : C)}{\mathcal{V}; \mathcal{C}; \Delta, (x : \mathbf{1}) \Vdash \text{wait } x; Q :: (z : C)} \mathbf{1}L
\end{array}$$

Figure 4.7: Typing rules for structural resource-aware session types

which is received by its client. Dually, the \triangleleft operator requires that the provider receives potential paid by the client. The type guarantees that the potential paid by the sender equals what is gained by the recipient, thereby preserving the total potential of a configuration.

$$\begin{array}{c}
\frac{\mathcal{V}; \mathcal{C} \models q \geq r \quad \mathcal{V}; \mathcal{C}; \Delta \Vdash^{q-r} P :: (x : A)}{\mathcal{V}; \mathcal{C}; \Delta \Vdash \text{pay } x \{r\}; P :: (x : \triangleright^r A)} \triangleright R \\
\\
\frac{\mathcal{V}; \mathcal{C}; \Delta, (x : A) \Vdash^{q+r} Q :: (z : C)}{\mathcal{V}; \mathcal{C}; \Delta, (x : \triangleright^r A) \Vdash \text{get } x \{r\}; Q :: (z : C)} \triangleright L \\
\\
\frac{\mathcal{V}; \mathcal{C}; \Delta \Vdash^{q+r} P :: (x : A)}{\mathcal{V}; \mathcal{C}; \Delta \Vdash \text{get } x \{r\}; P :: (x : \triangleleft^r A)} \triangleleft R \\
\\
\frac{\mathcal{V}; \mathcal{C} \models q \geq r \quad \mathcal{V}; \mathcal{C}; \Delta, (x : A) \Vdash^{q-r} Q :: (z : C)}{\mathcal{V}; \mathcal{C}; \Delta, (x : \triangleleft^r A) \Vdash \text{pay } x \{r\}; Q :: (z : C)} \triangleleft L
\end{array}$$

$$\begin{array}{c}
\frac{}{\Delta \vDash^0 (\cdot) :: \Delta} \text{empty} \qquad \frac{\Delta \vDash^E \mathcal{S} :: \Delta' \quad \Delta' \vDash^{E'} \mathcal{S}' :: \Delta''}{\Delta \vDash^{E+E'} (\mathcal{S} \mathcal{S}') :: \Delta''} \text{compose} \\
\\
\frac{\cdot ; \top ; \Delta_1 \vdash^q P :: (x : A)}{\Delta, \Delta_1 \vDash^{q+w} (\text{proc}(x, w, P)) :: (\Delta, (x : A))} \text{proc} \\
\\
\frac{\cdot ; \top ; \Delta_1 \vdash^q M :: (x : A)}{\Delta, \Delta_1 \vDash^{q+w} (\text{msg}(x, w, M)) :: (\Delta, (x : A))} \text{msg}
\end{array}$$

Figure 4.8: Typing rules for a configuration

4.5 Soundness

This section demonstrates the soundness of the resource-aware type system with respect to the operational cost semantics. So far, we have analyzed and type-checked processes in isolation. However, as our cost semantics indicates, processes always exist in a configuration interacting with other processes. Thus, we need to extend the typing rules to arbitrary configurations.

Configuration Typing At runtime, a program evolves into a set of processes interacting via typed channels. Such a configuration is typed w.r.t. a well-formed signature. A signature Σ is *well formed* if (a) every type definition $V = A$ is contractive, and (b) every process definition $\Delta \vdash^q f[\bar{n}] = P :: (x : A)$ in Σ is well typed according to the process typing judgment, i.e. $\bar{n} ; \top ; \Delta \vdash^q P :: (x : A)$.

I use the following judgment to type a configuration.

$$\Delta_1 \vDash_{\Sigma}^E \mathcal{S} :: \Delta_2$$

It states that Σ is well-formed and that the configuration \mathcal{S} uses the channels in the context Δ_1 and provides the channels in the context Δ_2 . The natural number E denotes the sum of the total potential and work done by the system. I call E the energy of the configuration. The configuration typing judgment is defined using the rules presented in Figure 4.8. The rule *empty* defines that an empty configuration is well-typed with energy 0. The rule *compose* composes two configurations \mathcal{S} and \mathcal{S}' ; \mathcal{S} provides service on the channels in Δ' while \mathcal{S}' uses the channels in Δ' . The energy of the composed configuration $\mathcal{S} \mathcal{S}'$ is obtained by summing up their individual energies. The rule *proc* creates a configuration out of a single process. The energy of this singleton configuration is obtained by adding the potential of the process and the work performed by it. Similarly, the rule *msg* creates a configuration out of a single message.

Soundness Theorem 7 is the main theorem of the chapter. It is a stronger version of a classical type preservation theorem and the usual type preservation is a direct consequence. Intuitively,

it states that the energy of a configuration never increases during an evaluation step, i.e. the energy remains conserved.

Theorem 7 (Soundness). *Consider a well-typed configuration \mathcal{S} w.r.t. a well-formed signature Σ such that $\Delta_1 \stackrel{E}{\vDash}_{\Sigma} \mathcal{S} :: \Delta_2$. If $\mathcal{S} \mapsto \mathcal{S}'$, then $\Delta_1 \stackrel{E}{\vDash}_{\Sigma} \mathcal{S}' :: \Delta_2$.*

The proof of the soundness theorem is achieved by a case analysis on the cost semantics, followed by an inversion on the typing of a configuration. The preservation theorem is a corollary since soundness implies that the configuration \mathcal{S}' is well-typed.

The soundness implies that the energy of an initial configuration is an upper bound on the energy of any configuration it will ever step to. In particular, if a configuration starts with 0 work, the initial energy (equal to initial potential) is an upper bound on the total work performed by an evaluation starting in that configuration.

Corollary 1 (Upper Bound). *If $\Delta_1 \stackrel{E}{\vDash}_{\Sigma} \mathcal{S} :: \Delta_2$, and $\mathcal{S} \mapsto^* \mathcal{S}'$, then $E \geq W'$, where W' is the total work performed by the configuration \mathcal{S}' , i.e. the sum of the work performed by each process and message in \mathcal{S}' . In particular, if the work done by the initial configuration \mathcal{S} is 0, then the potential P of the initial configuration satisfies $P \geq W'$.*

Proof. Applying the Soundness theorem successively, we get that if $\mathcal{S} \mapsto^* \mathcal{S}'$ and $\Delta_1 \stackrel{E}{\vDash} \mathcal{S} :: \Delta_2$, then $\Delta_1 \stackrel{E}{\vDash} \mathcal{S}' :: \Delta_2$. Also, $E = P' + W'$, where P' is the total potential of \mathcal{S}' , while W' is the total work performed so far in \mathcal{S}' . Since $P' \geq 0$, we get that $W' \leq P' + W' = E$. In particular, if $W = 0$, we get that $P = P + W = E \geq W'$, where P and W are the potential and work of the initial configuration respectively. \square

The progress theorem is a direct consequence of progress in SILL [153]. Our cost semantics are a cost observing semantics, i.e. it is just annotated with counters observing the work. Hence, any runtime step that can be taken by a program in SILL can be taken in this language.

4.6 Case Study: Stacks and Queues

As an illustration of the type system, I present a case study on stacks and queues. Stacks and queues have the same interaction protocol: they store elements of a variable type A and support inserting and deleting elements. They only differ in their implementation and resource usage. Their common interface type is expressed as the simple session type store_A (parameterized by type variable A and size n).

$$\begin{aligned} \text{store}_A[n] = \&\{ \text{ins} : A \multimap \text{store}_A[n+1], \\ &\text{del} : \oplus\{\text{none} : ?\{n=0\}. \mathbf{1}, \\ &\quad \text{some} : ?\{n>0\}. A \otimes \text{store}_A[n-1]\} \} \end{aligned}$$

The session type dictates that a process providing a service of type $\text{store}_A[n]$ gives a client the choice to either insert (ins) or delete (del) an element of type A . Upon receipt of the label ins, the providing process expects to receive a channel of type A to be enqueued and recurses. Upon receipt of the label del, the providing process either indicates that the queue is empty (none), in which case it terminates, or that there is an element stored in the queue (some), in which case it deletes this element, sends it to the client, and recurses.

To account for the resource cost, I add potential annotations leading to the store_A type to obtain two different resource-aware types for stacks and queues. The cost model again counts the total number of messages exchanged.

Stacks The type for stacks is defined as follows.

$$\begin{aligned} \text{stack}_A[n] = \&\{ \text{ins} : A \multimap \text{stack}_A[n+1], \\ &\text{del} : \triangleleft^2 \oplus \{ \text{none} : ?\{n=0\}. \mathbf{1}, \\ &\quad \text{some} : ?\{n>0\}. A \otimes \text{stack}_A[n-1] \} \} \end{aligned}$$

A stack is implemented using a sequence of *elem* processes terminated by an *empty* process. Each *elem* process stores an element of the stack, while *empty* denotes the end of stack.

Inserting an element simply spawns a new *elem* process (which has no cost in our cost model), thus having no resource cost. Deleting an element terminates the *elem* process at the head. Before termination, it sends two messages back to the client, either none followed by close or some followed by element. Thus, deletion has a resource cost of 2. This is reflected in the stack_A type, where *ins* and *del* are annotated with none and 2 units of potential respectively.

Queues The queue interface is achieved by using the same store_A type and annotating it with a different potential. The tight potential bound depends on the number of elements stored in the queue. Hence, a precise resource-aware type needs access to this internal measure in the type. A type $\text{queue}_A[n]$ intuitively defines a queue of size n (for $n > 0$).

$$\begin{aligned} \text{queue}_A[n] = \&\{ \text{ins} : \triangleleft^{2n} (A \multimap \text{queue}_A[n+1]), \\ &\text{del} : \triangleleft^2 \oplus \{ \text{none} : ?\{n=0\}. \mathbf{1}, \\ &\quad \text{some} : ?\{n>0\}. A \otimes \text{queue}_A[n-1] \} \} \end{aligned}$$

A queue is also implemented by a sequence of *elem* processes, connected via channels, terminated by the *empty* process, similar to a stack.

For each insertion, the *ins* label along with the element travels to the end of the queue. There, it spawns a new *elem* process that holds the inserted element. Hence, the resource cost of each insertion is $2n$ where n is the size of the queue. On the other hand, deletion is similar to that of stack and has a resource cost of 2. Again, this is reflected in the queue_A type, where *ins* and *del* are annotated with $2n$ and 2 units of potential respectively.

The resource-aware types show that stacks are more efficient than queues. The label *ins* is annotated with no potential for stack_A and with $2n$ for queue_A . The label *del* has the same annotation in both types. Hence, an efficiency comparison can be performed by simply observing the resource-aware session types.

Queues as two stacks In a functional language, a queue is often implemented with two lists. The idea is to enqueue into the first list and to dequeue from the second list. If the second list is empty, the the first list is copied over to the second list , thereby reversing its order. Since the cost of the dequeue operation varies drastically between the dequeue operations, amortized analysis is again instrumental in the analysis of the worst-case behavior and shows that the worst-case amortized cost for deletion is actually a constant. The type of the queue is

$$\begin{aligned} \text{queue}_A[n] = \&\{ \text{ins} : \triangleleft^6 (A \multimap \text{queue}_A[n+1]), \\ &\text{del} : \triangleleft^2 \oplus \{ \text{none} : ?\{n=0\}. \mathbf{1}, \end{aligned}$$

$$\text{some} : \{n > 0\}. A \otimes \text{queue}_A[n - 1]\}$$

Resource-aware session types enable us to translate the amortized analysis to the distributed setting. The type prescribes that an insertion has an amortized cost of 6 while the deletion has an amortized cost of 2. The main idea here is that the elements are inserted with a constant potential in the first list. While deleting, if the second list is empty, then this stored potential in the first list is used to pay for copying the elements over to the second list. As demonstrated from the resource-aware type, this implementation is more efficient than the previous queue implementation, which has a linear resource cost for insertion.

4.7 Related Work

In the context of process calculi, capabilities [150] and static analyses [103] have been used to statically restrict communication for controlling buffer sizes in languages without session types. For session-typed communication, upper bounding the size of message queues is simpler and studied in the compiler for Concurrent C0 [160]. In contrast to capabilities, our potential annotations do not control buffer sizes but provide a symbolic description of the number of messages exchanged at runtime. It is not clear how capabilities could be used to perform such an analysis.

Type systems for static resource bound analysis for sequential programs have been extensively studied (e.g., [50, 106]). The work is based on type-based amortized resource analysis. Automatic amortized resource analysis (AARA) has been introduced as a type system to automatically derive linear [93] and polynomial bounds [92] for sequential functional programs. It can also be integrated with program logics to derive bounds for imperative programs [21, 45]. Moreover, it has been used to derive bounds for term-rewrite systems [95] and object-oriented programs [94]. A recent work also considers bounds on the parallel evaluation cost (also called *span*) of functional programs [90]. The innovation of our work is the integration of AARA and session types and the analysis of message-passing programs that communicate with the outside world. Instead of function arguments, our bounds depend on the messages that are sent along channels. As a result, the formulation and proof of the soundness theorem is quite different from the soundness of sequential AARA.

I am only aware of a couple of other works that study resource bounds for concurrent programs. Gimenez et al. [75] introduced a technique for analyzing the parallel and sequential space and time cost of evaluating interaction nets. While it also based on linear logic and potential annotations, the flavor of the analysis is quite different. Interaction nets are mainly used to model parallel evaluation while session types focus on the interaction of processes. A main innovation of our work is that processes can exchange potential via messages. It is not clear how we can represent the examples we consider in this article as interaction nets. Albert et al. [13, 15] have studied techniques for deriving bounds on the cost of concurrent programs that are based on the actor model. While the goals of the work are similar to ours, the used technique and considered examples are dissimilar. A major difference is that our method is type-based and compositional. A unique feature of our work is that types describe bounds as functions of the messages that are sent along a channel.

4.8 Future Directions

I briefly mention some important future directions with regard to work analysis.

Work Inference To completely automate the type system, it is crucial to infer the work bounds, not just check them. This entails inserting pay or get constructors with every continuation with a parametric value and then obtaining constraints on these parameters. Then a solver tries to solve these constraints while minimizing the number of such constructors that need to be inserted. If an algorithmic version of this inference is implemented, a programmer will simply write the original simple session-typed program, and the inference engine will infer the resource-aware type, along with the resource bound.

Process Scheduling Inferring work bounds has several applications. One such direction to explore is the use of resource bounds in process scheduling. For instance, oracle schedulers [9] can use a priori knowledge of the runtime of each parallel thread to calculate thread creation overheads and enhance efficiency. Thus, resource-aware session types can be used to design an efficient scheduling algorithm that maximizes throughput.

Chapter 5

Time Analysis

Chapter 3 set the stage for refined session types. Chapter 4 employed refined types to express work bounds that computed the total number of operations executed by a set of concurrent processes. However, in a distributed setting, not all these operations would execute sequentially; some of these operations could occur in parallel. Thus, this chapter complements the previous chapter by exploring the inherent parallelism in the system by computing the total *time* of execution. We compute these time bounds assuming maximal parallelism; everything that can happen in parallel will happen in parallel. We present a cost semantics to measure this time, and a type system that can express such bounds.

5.1 Introduction

Analyzing the complexity of concurrent message-passing processes poses additional challenges over sequential programs. To begin with, we need information about the possible interactions between processes to enable compositional and local reasoning about concurrent cost. In addition to the structure of communication, the timing of messages is of central interest for analyzing concurrent cost. With information on message timing we may analyze not only properties such as the rate or latency with which a stream of messages can proceed through a pipeline, but also the span of a parallel computation, which can be defined as the time of the final response message assuming maximal parallelism.

There are several possible ways to enrich session types with timing information. A challenge is to find a balance between precision and flexibility. We would like to express precise times according to a global clock as in synchronous data flow languages whenever that is possible. However, sometimes this will be too restrictive. For example, we may want to characterize the response time of a concurrent queue where enqueue and dequeue operations arrive at unpredictable intervals.

In this chapter, I develop a type system that captures the parallel complexity of session-typed message-passing programs by adding *temporal modalities* *next* ($\circ A$), *always* ($\square A$), and *eventually* ($\diamond A$), interpreted over a linear model of time. When considered as types, the temporal modalities express properties of concurrent programs such as the *message rate* of a stream, the *latency* of a pipeline, the *response time* of concurrent data structure, or the *span* of a fork/join

parallel program, all in the same uniform manner. The results complement my prior work on expressing the *work* of session-typed processes in the same base language [61]. Together, they form a foundation for analyzing the parallel complexity of session-typed processes.

The type system is constructed conservatively over the base language of session types, which makes it quite general and easily able to accommodate various concrete cost models. The language contains standard session types and process expressions, and their typing rules remain unchanged. They correspond to processes that do not induce cost and send all messages at the constant time 0.

To model computation cost, a new syntactic form **delay** is introduced, which advances time by one step. A particular cost semantics is specified by taking an ordinary, non-temporal program and adding delays capturing the intended cost. For example, if only the blocking operations should cost one unit of time, a delay is added before the continuation of every receiving construct. If sends should have unit cost as well, a delay is added immediately after each send operation. Processes that contain delays cannot be typed using standard session types.

To type processes with non-zero cost, I first introduce the type $\circ A$, which is inhabited only by the process expression (**delay** ; P). This forces time to advance on all channels that P can communicate along. The resulting types prescribe the *exact* time a message is sent or received and sender and receiver are precisely synchronized.

As an example, consider a stream of bits terminated by $\$$, expressed as the recursive type

$$\text{bits} = \oplus\{\text{b0} : \text{bits}, \text{b1} : \text{bits}, \$: \mathbf{1}\}$$

where \oplus stands for *internal choice* and $\mathbf{1}$ for *termination*, ending the session. A simple cost model for asynchronous communication prescribes a cost of one unit of time for every receive operation. A stream of bits then needs to delay every continuation to give the recipient time to receive the message, expressing a *rate* of one. This can be captured precisely with the temporal modality $\circ A$:

$$\text{bits} = \oplus\{\text{b0} : \circ\text{bits}, \text{b1} : \circ\text{bits}, \$: \circ\mathbf{1}\}$$

A transducer *neg* that negates each bit it receives along channel x and passes it on along channel y would be typed as

$$x : \text{bits} \vdash \text{neg} :: (y : \circ\text{bits})$$

expressing a *latency* of one. A process *negneg* that puts two negations in sequence has a latency of two, compared with *copy* which passes on each bit, and *id* which terminates and identifies the channel y with the channel x , short-circuiting the communication.

$$x : \text{bits} \vdash \text{negneg} :: (y : \circ\circ\text{bits}) \quad x : \text{bits} \vdash \text{copy} :: (y : \circ\text{bits}) \quad x : \text{bits} \vdash \text{id} :: (y : \text{bits})$$

All these processes have the same extensional behavior, but different latencies. They also have the same rate since after the pipelining delay, the bits are sent at the same rate they are received, as expressed in the common type bits used in the context and the result.

While precise and minimalistic, the resulting system is often too precise for typical concurrent programs such as pipelines or servers. I therefore introduce the dual type formers $\diamond A$ and $\square A$ to talk about varying time points in the future. Remarkably, even if part of a program is typed using these constructs, we can still make precise and useful statements about other aspects.

For example, consider a transducer *compress* that shortens a stream by combining consecutive 1 bits so that, for example, 00110111 becomes 00101. For such a transducer, we cannot

bound the latency statically, even if the bits are received at a constant rate like in the type bits. So we have to express that after seeing a 1 bit we will *eventually* see either another bit or the end of the stream. For this purpose, we introduce a new type sbits with the same message alternatives as bits, but different timing. In particular, after sending b1, either the next bit or end-of-stream is *eventually* sent (\diamond sbits), rather than immediately.

$$\text{sbits} = \oplus\{\text{b0} : \circ\text{sbits}, \text{b1} : \circ\diamond\text{sbits}, \$: \circ\mathbf{1}\}$$

$$x : \text{bits} \vdash \text{compress} :: (y : \circ\text{sbits})$$

We write $\circ\diamond\text{sbits}$ instead of $\diamond\text{sbits}$ for the continuation type after b1 to express that there will always be a delay of at least one; to account for the unit cost of receive in this particular cost model.

The dual modality, $\square A$, is useful to express, for example, that a server providing A is *always* ready, starting from “now”. As an example, consider the following temporal type of an interface to a process of type $\square\text{queue}_A$ with elements of type $\square A$. It expresses that there must be at least four time units between successive enqueue operations and that the response to a dequeue request is immediate, only one time unit later ($\&$ stands for external choice, the dual to internal choice).

$$\text{queue}_A = \&\{\text{enq} : \circ(\square A \multimap \circ^3\square\text{queue}_A), \\ \text{deq} : \circ\oplus\{\text{none} : \circ\mathbf{1}, \text{some} : \circ(\square A \otimes \circ\square\text{queue}_A)\}\}$$

As an example of a *parametric* cost analysis, the following type can be given to a process that appends inputs l_1 and l_2 to yield l , where the message rate on all three lists is $r + 2$ units of time (that is, the interval between consecutive list elements needs to be at least 2).

$$l_1 : \text{list}_A[n], l_2 : \circ^{(r+4)n+2} \text{list}_A[k] \vdash \text{append}[n, k, r] :: (l : \circ\circ\text{list}_A[n+k])$$

It expresses that *append* has a latency of two units of time and that it inputs the first message from l_2 after $(r + 4)n + 2$ units of time, where n is the number of elements sent along l_1 .

To analyze the span of a fork/join parallel program, we capture the time at which the (final) answer is sent. For example, the type $\text{tree}[h]$ describes the span of a process that computes the parity of a binary tree of height h with boolean values at the leaves. The session type expresses that the result of the computation is a single boolean that arrives at time $5h + 3$ after the parity request.

$$\text{tree}[h] = \&\{\text{parity} : \circ^{5h+3} \text{bool}\}$$

In summary, the main contributions of the chapter are (1) a generic framework for parallel cost analysis of asynchronously communicating session-typed processes rooted in a novel combination of temporal and linear logic, (2) a soundness proof of the type system with respect to a timed operational semantics, showing progress and type preservation (3) instantiations of the framework with different cost models, e.g. where either just receives, or receives and sends, cost one time unit each, and (4) examples illustrating the scope of my method. My technique for proving progress and preservation does not require dependency graphs and may be of independent interest. I further provide decidable systems for *time reconstruction* and *subtyping* that greatly simplify the programmer’s task. They also enhance modularity by allowing the same program to be assigned temporally different types, depending on the context of use.

5.2 The Temporal Modality Next ($\circ A$)

This section introduces *actual cost* by explicitly advancing time. Remarkably, all the rules presented so far in Chapter 2 remain literally unchanged. They correspond to the cost-free fragment of the language in which time never advances. In addition, I have a new type construct $\circ A$ (read: *next A*) with a corresponding process construct ($\text{delay} ; P$), which advances time by one unit. In the corresponding typing rule

$$\frac{\Delta \vdash P :: (x : A)}{\circ\Delta \vdash (\text{delay} ; P) :: (x : \circ A)} \text{ } \circ LR$$

I abbreviate $y_1:\circ A_1, \dots, y_m:\circ A_m$ by $\circ(y_1:A_1, \dots, y_m:A_m)$. Intuitively, when $(\text{delay} ; P)$ idles, time advances on *all* channels connected to P . Computationally, I delay the process for one time unit without any external interactions. To understand this computation, I introduce semantic objects $\text{proc}(c, t, P)$ and $\text{msg}(c, t, M)$ which mean that process P or message M provide along channel c and are at an integral time t .

$$(\circ C) \quad \text{proc}(c, t, \text{delay} ; P) \mapsto \text{proc}(c, t + 1, P)$$

There is a subtle point about forwarding: A process $\text{proc}(c, t, c \leftarrow d)$ may be ready to forward a message *before* a client reaches time t while in all other rules the times must match exactly. We can avoid this mismatch by transforming uses of forwarding $x \leftarrow y$ at type $\circ^n S$ where $S \neq \circ(-)$ to $(\text{delay}^n ; x \leftarrow y)$. In this discussion I have used the following notation which will be useful later:

$$\begin{array}{ll} \circ^0 A & = A & \text{delay}^0 ; P & = P \\ \circ^{n+1} A & = \circ \circ^n A & \text{delay}^{n+1} ; P & = \text{delay} ; \text{delay}^n ; P \end{array}$$

Modeling a Cost Semantics

My system allows us to represent a variety of different abstract cost models in a straightforward way. I will mostly use two different abstract cost models. In the first, called \mathcal{R} , I assign unit cost to every receive (or wait) action while all other operations remain cost-free. We may be interested in this since receiving a message is the only blocking operation in the asynchronous semantics. A second one, called \mathcal{RS} and considered in Section 5.5, assigns unit cost to both send and receive actions.

To capture \mathcal{R} , I take a source program and insert a delay operation before the continuation of every receive. I write this delay as **tick** in order to remind the reader that it arises systematically from the cost model and is never written by the programmer. In all other respects, **tick** is just a synonym for **delay**.

For example, the copy process would become

$$\text{bits} = \oplus \{ \text{b0} : \text{bits}, \text{b1} : \text{bits}, \$: 1 \}$$

$$y : \text{bits} \vdash \text{copy} :: (x : \text{bits}) \quad \% \text{ No longer correct!}$$

$$x \leftarrow \text{copy} \leftarrow y =$$

$$\begin{array}{l} \text{case } y \text{ (} \text{b0} \Rightarrow \text{tick} ; x.\text{b0} ; x \leftarrow \text{copy} \leftarrow y \\ \quad | \text{b1} \Rightarrow \text{tick} ; x.\text{b1} ; x \leftarrow \text{copy} \leftarrow y \\ \quad | \$ \Rightarrow \text{tick} ; x.\$; \text{wait } y ; \text{tick} ; \text{close } x \text{)} \end{array}$$

As indicated in the comment, the type of *copy* is now no longer correct because the bits that arrive along *y* are delayed by one unit before they are sent along *x*. We can observe this concretely by starting to type-check the first branch

```

y : bits ⊢ copy :: (x : bits)
x ← copy ← y =
  case y ( b0 ⇒                               % y : bits ⊢ x : bits
           tick ; ... )

```

We see that the delay `tick` does not type-check, because neither *x* nor *y* have a type of the form $\circ(-)$. We need to redefine the type `bits` so that the continuation type after every label is delayed by one, anticipating the time it takes to receive the label `b0`, `b1`, or `$`. Similarly, we capture in the type of *copy* that its *latency* is one unit of time.

```
bits = ⊕{b0 : ○bits, b1 : ○bits, $ : ○1}
```

```
y : bits ⊢ copy :: (x : ○bits)
```

With these declarations, we can now type-check the definition of *copy*. I show the intermediate type of the used and provided channels after each interaction.

```

x ← copy ← y =
  case y ( b0 ⇒                               % y : ○bits ⊢ x : ○bits
           tick ;                               % y : bits ⊢ x : bits
           x.b0 ;                               % y : bits ⊢ x : ○bits
           x ← copy ← y                         % well-typed by type of copy
        | b1 ⇒                               % y : ○bits ⊢ x : ○bits
           tick ;                               % y : bits ⊢ x : bits
           x.b1 ;                               % y : bits ⊢ x : ○bits
           x ← copy ← y
        | $ ⇒                               % y : ○1 ⊢ x : ○bits
           tick ;                               % y : 1 ⊢ x : bits
           x.$ ;                               % y : 1 ⊢ x : ○1
           wait y ;                             % · ⊢ x : ○1
           tick ;                               % · ⊢ x : 1
           close x )

```

Armed with this experience, we now consider the increment process *plus1*. Again, we expect the latency of the increment to be one unit of time. Since we are interested in detailed type-checking, I show the transformed program, with a delay `tick` after each receive.

```
bits = ⊕{b0 : ○bits, b1 : ○bits, $ : ○1}
```

```
y : bits ⊢ plus1 :: (x : ○bits)
```

```

x ← plus1 ← y =
  case y ( b0 ⇒ tick ; x.b1 ; x ← y             % type error here!
           | b1 ⇒ tick ; x.b0 ; x ← plus1 ← y
           | $ ⇒ tick ; x.$ ; wait y ; tick ; close x )

```

The branches for `b1` and `$` type-check as before, but the branch for `b0` does not. I make the types at the crucial point explicit:

```
x ← plus1 ← y =
```

```

case y ( b0 ⇒ tick ; x.b1 ;                % y : bits ⊢ x : ○bits
        x ← y                               % ill-typed, since bits ≠ ○bits
  | ... )

```

The problem here is that identifying x and y removes the delay mandated by the type of $plus1$. A solution is to call $copy$ to reintroduce the latency of one time unit.

```

y : bits ⊢ plus1 :: (x : ○bits)
x ← plus1 ← y =
  case y ( b0 ⇒ tick ; x.b1 ; x ← copy ← y
          | b1 ⇒ tick ; x.b0 ; x ← plus1 ← y
          | $ ⇒ tick ; x.$ ; wait y ; tick ; close x )

```

In order to write $plus2$ as a pipeline of two increments we need to delay the second increment explicitly in the program and stipulate, in the type, that there is a latency of two.

```

y : bits ⊢ plus2 :: (x : ○○bits)
x ← plus2 ← y =
  z ← plus1 ← y ;           % z : ○bits ⊢ x : ○○bits
  delay ;                   % z : bits ⊢ x : ○bits
  x ← plus1 ← z

```

Programming with so many explicit delays is tedious, but fortunately a source program without all these delay operations (but explicitly temporal session types) can be transformed automatically in two steps: (1) insert the delays mandated by the cost model (here: a `tick` after each receive), and (2) perform *time reconstruction* to insert the additional delays so the result is temporally well-typed or issue an error message if this is impossible (see [62]).

The Interpretation of a Configuration

Let us reconsider the program to produce the number $6 = (110)_2$ under the cost model \mathcal{R} where each receive action costs one unit of time. There are no receive operations in this program, but time reconstruction must insert a delay after each send in order to match the delays mandated by the type `bits`.

```

bits = ⊕{b0 : ○bits, b1 : ○bits, $ : ○1}
· ⊢ six :: (x : bits)
x ← six = x.b0 ; delay ; x.b1 ; delay ; x.b1 ; delay ; x.$ ; delay ; close x

```

Executing $\text{proc}(c_0, 0, c_0 \leftarrow \text{six})$ then leads to the following configuration

```

msg(c4, 4, close c4),
msg(c3, 3, c3.$ ; c3 ← c4),
msg(c2, 2, c2.b1 ; c2 ← c3),
msg(c1, 1, c1.b1 ; c1 ← c2),
msg(c0, 0, c0.b0 ; c0 ← c1)

```

These messages are at increasing times, which means any client of c_0 will have to immediately (at time 0) receive `b0`, then (at time 1) `b1`, then (at time 2) `b1`, etc. In other words, the time stamps on messages predict *exactly* when the message will be received. Of course, if there is a client in parallel this state may never be reached because, for example, the first `b0` message along channel

c_0 may be received before the continuation of the sender produces the message $b1$. So different configurations may be reached depending on the *scheduler* for the concurrent processes. It is also possible to give a time-synchronous semantics in which all processes proceed *in parallel* from time 0 to time 1, then from time 1 to time 2, etc.

5.3 The Temporal Modalities Always ($\Box A$) and Eventually ($\Diamond A$)

The strength and also the weakness of the system so far is that its timing is very precise. Consider a process *compress* that combines runs of consecutive 1's to a single 1. For example, compressing 11011100 should yield 10100. First, in the cost-free the process is defined as

$$\text{bits} = \oplus\{\text{b0} : \text{bits}, \text{b1} : \text{bits}, \$: \mathbf{1}\}$$

$$y : \text{bits} \vdash \text{compress} :: (x : \text{bits})$$

$$y : \text{bits} \vdash \text{skip1s} :: (x : \text{bits})$$

$$x \leftarrow \text{compress} \leftarrow y =$$

$$\begin{array}{l} \text{case } y \text{ (} \text{b0} \Rightarrow x.\text{b0} ; x \leftarrow \text{compress} \leftarrow y \\ \quad | \text{b1} \Rightarrow x.\text{b1} ; x \leftarrow \text{skip1s} \leftarrow y \\ \quad | \$ \Rightarrow x.\$; \text{wait } y ; \text{close } x \text{)} \end{array}$$

$$x \leftarrow \text{skip1s} \leftarrow y =$$

$$\begin{array}{l} \text{case } y \text{ (} \text{b0} \Rightarrow x.\text{b0} ; x \leftarrow \text{compress} \leftarrow y \\ \quad | \text{b1} \Rightarrow x \leftarrow \text{skip1s} \leftarrow y \\ \quad | \$ \Rightarrow x.\$; \text{wait } y ; \text{close } x \text{)} \end{array}$$

The problem is that program cannot be typed under the cost mode \mathcal{R} , where every receive takes one unit of time. Actually worse: there is no way to insert next-time modalities into the type and additional delays into the program so that the result is well-typed. This is because if the input stream is unknown we cannot predict how long a run of 1's will be, but the length of such a run will determine the delay between sending a bit 1 and the following bit 0.

The best we can say is that after a bit 1 *compress* will *eventually* send either a bit 0 or the end-of-stream token $\$$. This is the purpose of the type $\Diamond A$. We capture this timing in the type *sbits* (for *slow bits*).

$$\text{bits} = \oplus\{\text{b0} : \circ\text{bits}, \text{b1} : \circ\text{bits}, \$: \circ\mathbf{1}\}$$

$$\text{sbits} = \oplus\{\text{b0} : \circ\text{sbits}, \text{b1} : \circ\Diamond\text{sbits}, \$: \circ\mathbf{1}\}$$

$$y : \text{bits} \vdash \text{compress} :: (x : \circ\text{sbits})$$

$$y : \text{bits} \vdash \text{skip1s} :: (x : \circ\Diamond\text{sbits})$$

The next section introduces the process constructs and typing rules so that *compress* and *skip1s* programs can be revised to have the right temporal semantics.

Eventually A

A process providing $\Diamond A$ promises only that it will eventually provide A . There is a somewhat subtle point here: since not every action may require time and because we do not check termina-

tion separately, $x : \diamond A$ expresses only that *if the process providing x terminates* it will eventually provide A . Thus, it expresses non-determinism regarding the (abstract) *time* at which A is provided, rather than a strict liveness property. Therefore, $\diamond A$ is somewhat weaker than one might be used to from LTL [131]. When restricted to a purely logical fragment, without unrestricted recursion, the usual meaning is fully restored so I feel the terminology is justified. Imposing termination, for example along the lines of Fortier and Santocanale [70] or Toninho et al. [154] is an interesting item for future work but not necessary for our present purposes.

When a process offering $c : \diamond A$ is ready, it will send a **now!** message along c and then continue at type A . Conversely, the client of $c : \diamond A$ will have to be ready and waiting for the **now!** message to arrive along c and then continue at type A . I use (**when?** c ; Q) for the corresponding client. These explicit constructs are a conceptual device and may not need to be part of an implementation. They also make type-checking processes entirely syntax-directed and trivially decidable.

The typing rules for **now!** and **when?** are somewhat subtle.

$$\frac{\Delta \vdash P :: (x : A)}{\Delta \vdash (\mathbf{now!} \ x ; P) :: (x : \diamond A)} \diamond R \quad \frac{\Delta \text{ delayed}^\square \quad \Delta, x : A \vdash Q :: (z : C) \quad C \text{ delayed}^\diamond}{\Delta, x : \diamond A \vdash (\mathbf{when?} \ x ; Q) :: (z : C)} \diamond L$$

The $\diamond R$ rule just states that, without constraints, we can at any time decide to communicate along $x : \diamond A$ and then continue the session at type A . The $\diamond L$ rule expresses that the process must be ready to receive a **now!** message along $x : \diamond A$, but there are two further constraints. Because the process (**when?** x ; Q) may need to wait an indefinite period of time, the rule must make sure that communication along z and any channel in Δ can also be postponed an indefinite period of time. The predicate $C \text{ delayed}^\diamond$ describes that C must have the form $\circ^* \diamond C'$ to require that C may be delayed a fixed finite number of time steps and then must be allowed to communicate at an arbitrary time in the future. Similarly, for every channel $y : B$ in Δ , B must have the form $\circ^* \square B$, where \square (as the dual of \diamond) is introduced in Section 5.3.

In the operational semantics, the central restriction is that **when?** is ready *before* the **now!** message arrives so that the continuation can proceed immediately as promised by the type.

$$\begin{aligned} (\diamond S) \quad & \text{proc}(c, t, \mathbf{now!} \ c ; P) \mapsto \text{proc}(c', t, [c'/c]P), \text{msg}(c, t, \mathbf{now!} \ c ; c \leftarrow c') \quad (c' \text{ fresh}) \\ (\diamond C) \quad & \text{msg}(c, t, \mathbf{now!} \ c ; c \leftarrow c'), \text{proc}(d, s, \mathbf{when?} \ c ; Q) \mapsto \text{proc}(d, t, [c'/c]Q) \quad (t \geq s) \end{aligned}$$

To rewrite the *compress* process in our cost model \mathcal{R} , I first insert **tick** before all the actions that must be delayed according to our cost model. Then I insert appropriate additional **delay**, **when?**, and **now!** actions. While *compress* turns out to be straightforward, *skip1s* creates a difficulty after it receives a **b1**:

$$\begin{aligned} \text{bits} &= \oplus\{\mathbf{b0} : \circ \text{bits}, \mathbf{b1} : \circ \text{bits}, \$: \circ \mathbf{1}\} \\ \text{sbits} &= \oplus\{\mathbf{b0} : \circ \text{sbits}, \mathbf{b1} : \circ \diamond \text{sbits}, \$: \circ \mathbf{1}\} \end{aligned}$$

$$\begin{aligned} y : \text{bits} \vdash \text{compress} &:: (x : \circ \text{sbits}) \\ y : \text{bits} \vdash \text{skip1s} &:: (x : \circ \diamond \text{sbits}) \end{aligned}$$

$$\begin{aligned} x \leftarrow \text{compress} \leftarrow y &= \\ \text{case } y \quad & (\mathbf{b0} \Rightarrow \mathbf{tick} ; x.\mathbf{b0} ; x \leftarrow \text{compress} \leftarrow y \\ & | \mathbf{b1} \Rightarrow \mathbf{tick} ; x.\mathbf{b1} ; x \leftarrow \text{skip1s} \leftarrow y \end{aligned}$$


```

    | $ ⇒ tick ; x.$ ; wait y ; tick ; close x )
x ← skip1s ← y =
  case y ( b0 ⇒ tick ; now! x ; x.b0 ; x ← compress ← y
    | b1 ⇒ tick ;                               % y : bits ⊢ x : ◇sbits
            x' ← skip1s ← y ;                   % x' : ○◇sbits ⊢ x : ◇sbits
            x ← idle ← x'                       % with x' : ○◇sbits ⊢ idle :: (x : ◇sbits)
    | $ ⇒ tick ; now! x ; x.$ ; wait y ; tick ; close x )

```

At the point where I would like to call *skip1s* recursively, I have

```

y : bits ⊢ x : ◇sbits
but y : bits ⊢ skip1s :: (x : ○◇sbits)

```

which prevents a tail call since $\circ\Diamond\text{sbits} \neq \Diamond\text{sbits}$. Instead *skip1s* is called to obtain a new channel x' and then use another process called *idle* to go from $x' : \circ\Diamond\text{sbits}$ to $x : \Diamond\text{sbits}$. Intuitively, it should be possible to implement such an idling process: $x : \Diamond\text{sbits}$ expresses *at some time in the future, including possibly right now* while $x' : \circ\Diamond\text{sbits}$ says *at some time in the future, but not right now*.

To type the idling process, the $\circ LR$ rule needs to be generalized to account for the interactions of $\circ A$ with $\square A$ and $\Diamond A$. After all, they speak about the same underlying model of time.

Interactions of $\circ A$ and $\Diamond A$

Recall the left/right rule for \circ :

$$\frac{\Delta \vdash P :: (x : A)}{\circ\Delta \vdash (\text{delay} ; P) :: (x : \circ A)} \circ LR$$

If the succedent were $x : \Diamond A$ instead of $x : \circ A$, we should still be able to delay since we can freely choose when to interact along x . We could capture this in the following rule (superseded later by a more general form of $\circ LR$):

$$\frac{\Delta \vdash P :: (x : \Diamond A)}{\circ\Delta \vdash (\text{delay} ; P) :: (x : \Diamond A)} \circ\Diamond$$

I keep $\Diamond A$ as the type of x since I want to retain the full flexibility of using x at any time in the future after the initial delay. I will generalize the rule once more in the next section to account for interactions with $\square A$.

With this, I can define and type the idling process parametrically over A :

```

x' : ○◇A ⊢ idle :: (x : ◇A)
x ← idle ← x' = delay ; x ← x'

```

This turns out to be an example of subtyping (see [62]), which means that the programmer actually will not have to explicitly define or even reference an idling process. The programmer simply writes the original *skip1s* process (without referencing the *idle* process) and our subtyping algorithm will use the appropriate rule to typecheck it successfully.

Always A

The last temporal modality, written as $\Box A$ (read: *always A*), is dual to $\Diamond A$. If a process P provides $x : \Box A$ it means it is ready to receive a **now!** message along x at any point in the future. In analogy with the typing rules for $\Diamond A$, but flipped to the other side of the sequent, we obtain

$$\frac{\Delta \text{ delayed}^{\Box} \quad \Delta \vdash P :: (x : A)}{\Delta \vdash (\text{when? } x ; P) :: (x : \Box A)} \Box R \qquad \frac{\Delta, x : A \vdash Q :: (z : C)}{\Delta, x : \Box A \vdash (\text{now! } x ; Q) :: (z : C)} \Box L$$

The operational rules just reverse the role of provider and client from the rules for $\Diamond A$.

$$\begin{aligned} (\Box S) \quad & \text{proc}(d, t, \text{now! } c ; Q) \mapsto \text{msg}(c', t, \text{now! } c ; c' \leftarrow c), \text{proc}(d, t, [c'/c]Q) \quad (c' \text{ fresh}) \\ (\Box C) \quad & \text{proc}(c, s, \text{when? } c ; P), \text{msg}(c', t, \text{now! } c ; c' \leftarrow c) \mapsto \text{proc}(c', t, [c'/c]P) \quad (s \leq t) \end{aligned}$$

As an example for the use of $\Box A$, and also to introduce a new kind of example, I specify and implement a counter process that can receive `inc` and `val` messages. When receiving an `inc` it will increment its internally maintained counter, when receiving `val` it will produce a finite bit stream representing the current value of the counter. In the cost-free setting the type is

`bits = $\oplus\{\text{b0} : \text{bits}, \text{b1} : \text{bits}, \$: \mathbf{1}\}$`

`ctr = $\&\{\text{inc} : \text{ctr}, \text{val} : \text{bits}\}$`

A counter is implemented by a chain of processes, each holding one bit (either *bit0* or *bit1*) or signaling the end of the chain (*empty*). For this purpose we implement three processes:

`$d : \text{ctr} \vdash \text{bit0} :: (c : \text{ctr})$`

`$d : \text{ctr} \vdash \text{bit1} :: (c : \text{ctr})$`

`$\cdot \vdash \text{empty} :: (c : \text{ctr})$`

`$c \leftarrow \text{bit0} \leftarrow d =$`

`case c (inc $\Rightarrow c \leftarrow \text{bit1} \leftarrow d$ % increment by continuing as bit1
 | val $\Rightarrow c.\text{b0} ; d.\text{val} ; c \leftarrow d$ % send b0 on c , send val on d , identify c and d)`

`$c \leftarrow \text{bit1} \leftarrow d =$`

`case c (inc $\Rightarrow d.\text{inc} ; c \leftarrow \text{bit0} \leftarrow d$ % send inc (carry) on d , continue as bit1
 | val $\Rightarrow c.\text{b1} ; d.\text{val} ; c \leftarrow d$ % send b1 on c , send val on d , identify c and d)`

`$c \leftarrow \text{empty} =$`

`case c (inc $\Rightarrow e \leftarrow \text{empty} ;$ % spawn a new empty process with channel e
 $c \leftarrow \text{bit1} \leftarrow e$ % continue as bit1
 | val $\Rightarrow c.\$; \text{close } c$ % send $\$$ on c and close c)`

Using our standard cost model \mathcal{R} there is a problem: the *carry bit* (the $d.\text{inc}$ message sent in the *bit1* process) is sent only on every other increment received because *bit0* continues as *bit1* *without* a carry, and *bit1* continues as *bit0* *with* a carry. So it will actually take 2^k increments received at the lowest bit of the counter (which represents the interface to the client) before an increment reaches the k th process in the chain. This is not a constant number, so the behavior cannot be characterized exactly using only the next time modality. Instead, I require, from a certain point on, a counter is always ready to receive either an `inc` or `val` message.

`bits = $\oplus\{\text{b0} : \text{Obits}, \text{b1} : \text{Obits}, \$: \text{O}\mathbf{1}\}$`

`ctr = $\Box\&\{\text{inc} : \text{Octr}, \text{val} : \text{Obits}\}$`

In the program, the ticks are mandated by our cost model and some additional **delay**, **when?**, and **now!** actions are present to satisfy the stated types. The two marked lines may look incorrect, but are valid based on the generalization of the $\circ LR$ rule in Section 5.3.

```

d :  $\circ$ ctr  $\vdash$  bit0 :: (c : ctr)
d : ctr  $\vdash$  bit1 :: (c : ctr)
·  $\vdash$  empty :: (c : ctr)
c  $\leftarrow$  bit0  $\leftarrow$  d =
  when? c ;                               % d :  $\circ$ ctr  $\vdash$  c : &\{\dots\}
  case c (inc  $\Rightarrow$  tick ;                   % d : ctr  $\vdash$  c : ctr
           c  $\leftarrow$  bit1  $\leftarrow$  d
         | val  $\Rightarrow$  tick ;               % d : ctr  $\vdash$  c : bits
           c.b0 ;                          % d : ctr  $\vdash$  c :  $\circ$ bits
           now! d ; d.val ;                 % d :  $\circ$ bits  $\vdash$  c :  $\circ$ bits
           c  $\leftarrow$  d)

c  $\leftarrow$  bit1  $\leftarrow$  d =
  when? c ;                               % d : ctr  $\vdash$  c : &\{\dots\}
  case c (inc  $\Rightarrow$  tick ;                   % d : ctr  $\vdash$  c : ctr (see Section 5.3)
           now! d ; d.inc ;                 % d :  $\circ$ ctr  $\vdash$  c : ctr
           c  $\leftarrow$  bit0  $\leftarrow$  d
         | val  $\Rightarrow$  tick ;               % d : ctr  $\vdash$  c : bit (see Section 5.3)
           c.b1 ;                          % d : ctr  $\vdash$  c :  $\circ$ bits
           now! d ; d.val ;                 % d :  $\circ$ bits  $\vdash$  c :  $\circ$ bits
           c  $\leftarrow$  d)

c  $\leftarrow$  empty =
  when? c ;                               % ·  $\vdash$  c : &\{\dots\}
  case c (inc  $\Rightarrow$  tick ;                   % ·  $\vdash$  c : ctr
           e  $\leftarrow$  empty ;               % e : ctr  $\vdash$  c : ctr
           c  $\leftarrow$  bit1  $\leftarrow$  e
         | val  $\Rightarrow$  tick ; c.$ ;         % ·  $\vdash$  c :  $\circ$ 1
           delay ; close c)

```

Interactions Between Temporal Modalities

Just as $\circ A$ and $\diamond A$ interacted in the rules since their semantics is based on the same underlying notion of time, so do $\circ A$ and $\square A$. Executing a delay allows any channel of type $\square A$ that is used and leaves its type unchanged because we are not obligated to communicate along it at any particular time. To cover all the cases, I introduce a new notation, writing $[A]_L^{-1}$ and $[A]_R^{-1}$ on types and extend it to contexts. Depending on one's point of view, this can be seen as stepping forward or backward by one unit of time.

$$\begin{array}{lll}
[\circ A]_L^{-1} = A & [\circ A]_R^{-1} = A & [x : A]_L^{-1} = x : [A]_L^{-1} \\
[\square A]_L^{-1} = \square A & [\square A]_R^{-1} = \text{undefined} & [x : A]_R^{-1} = x : [A]_R^{-1} \\
[\diamond A]_L^{-1} = \text{undefined} & [\diamond A]_R^{-1} = \diamond A & [\cdot]_L^{-1} = \cdot \\
[S]_L^{-1} = \text{undefined} & [S]_R^{-1} = \text{undefined} & [\Delta, \Delta']_L^{-1} = [\Delta]_L^{-1}, [\Delta']_L^{-1}
\end{array}$$

$$\begin{array}{c}
\frac{[\Delta]_L^{-1} \vdash P :: [x : A]_R^{-1}}{\Delta \vdash (\text{delay} ; P) :: (x : A)} \circ LR \quad \overline{\circ^* \Box A \text{ delayed}^\Box} \quad \overline{\circ^* \Diamond A \text{ delayed}^\Diamond} \\
\\
\frac{\Delta \vdash P :: (x : A)}{\Delta \vdash (\text{now! } x ; P) :: (x : \Diamond A)} \Diamond R \quad \frac{\Delta \text{ delayed}^\Box \quad \Delta, x:A \vdash Q :: (z : C) \quad C \text{ delayed}^\Diamond}{\Delta, x:\Diamond A \vdash (\text{when? } x ; Q) :: (z : C)} \Diamond L \\
\\
\frac{\Delta \text{ delayed}^\Box \quad \Delta \vdash P :: (x : A)}{\Delta \vdash (\text{when? } x ; P) :: (x : \Box A)} \Box R \quad \frac{\Delta, x:A \vdash Q :: (z : C)}{\Delta, x:\Box A \vdash (\text{now! } x ; Q) :: (z : C)} \Box L
\end{array}$$

Figure 5.1: Explicit Temporal Typing Rules

Here, S stands for any basic session type constructor as in Table 2.1. We use this notation in the general rule $\circ LR$ which can be found in Figure 5.1 together with the final set of rules for $\Box A$ and $\Diamond A$. In conjunction with the rules in Chapter 2 this completes the system of temporal session types where all temporal actions are explicit. The rule $\circ LR$ only applies if both $[\Delta]_L^{-1}$ and $[x : A]_R^{-1}$ are defined.

A type A is called *patient* if it does not force communication along a channel $x : A$ at any particular point in time. Because the direction of communication is reversed between the two sides of a sequent, a type A is patient if it has the form $\circ^* \Box A'$ if it is among the antecedents, and $\circ^* \Diamond A'$ if it is in the succedent. The judgments $A \text{ delayed}^\Box$ and $A \text{ delayed}^\Diamond$ are a shorthand for patient types. Further, $A \text{ delayed}^\Box$ is extended to contexts $\Delta \text{ delayed}^\Box$ if for every declaration $(x : A) \in \Delta$, $A \text{ delayed}^\Box$ holds.

5.4 Preservation and Progress

The main theorems that exhibit the deep connection between our type system and the timed operational semantics are the usual *type preservation* and *progress*, sometimes called *session fidelity* and *deadlock freedom*, respectively.

Configuration Typing

A key question is how we type configurations \mathcal{C} . Configurations consist of multiple processes and messages, so they both *use* and *provide* a collection of channels. And even though we treat a configuration as a multiset, typing imposes a partial order on the processes and messages where a provider of a channel appears to the left of its client.

$$\text{Configuration } \mathcal{C} ::= \cdot \mid \mathcal{C} \mathcal{C}' \mid \text{proc}(c, t, P) \mid \text{msg}(c, t, M)$$

The predicates $\text{proc}(c, t, P)$ and $\text{msg}(c, t, M)$ *provide* c . I stipulate that no two distinct processes or messages in a configuration provide the same channel c . Also recall that messages M are simply processes of a particular form and are typed as such. The possible messages (of which

there is one for each type constructor) can be read of from the operational semantics. They are summarized here for completeness.

$$M ::= (c.k ; c \leftarrow c') \mid (c.k ; c' \leftarrow c) \mid \text{close } c \mid (\text{send } c \ d ; c' \leftarrow c) \mid (\text{send } c \ d ; c \leftarrow c')$$

The typing judgment has the form $\Delta' \vDash C :: \Delta$ meaning that if composed with a configuration that provides Δ' , the result will provide Δ .

$$\frac{}{\Delta \vDash (\cdot) :: \Delta} \text{ empty} \qquad \frac{\Delta_0 \vDash C_1 :: \Delta_1 \quad \Delta_1 \vDash C_2 :: \Delta_2}{\Delta_0 \vDash (C_1 \ C_2) :: \Delta_2} \text{ compose}$$

To type processes and messages, I begin by considering *preservation*: I would like to achieve that if $\Delta' \vDash C :: \Delta$ and $C \mapsto C'$ then still $\Delta' \vDash C' :: \Delta$. Without the temporal modalities, this is guaranteed by the design of the sequent calculus: the right and left rules match just so that cut reduction (which is the basis for reduction in the operational semantics) leads to a well-typed deduction. The key here is what happens with time. Consider the special case of **delay**. When we transition from **delay** ; P to P we strip one \circ modality from Δ and A , but because we also advance time from t to $t + 1$, the \circ modality is restored, keeping the interface type invariant.

When we also consider types $\square A$ and $\diamond A$ the situation is a little less straightforward because of their interaction with \circ . I reuse the idea of the solution, allowing the subtraction of time from a type, possibly stopping when I meet a \square or \diamond .

$$\begin{aligned} [A]_L^{-0} &= A & [A]_R^{-0} &= A \\ [A]_L^{-(t+1)} &= [[A]_L^{-t}]_L^{-1} & [A]_R^{-(t+1)} &= [[A]_R^{-t}]_R^{-1} \end{aligned}$$

This is extended to channel declarations in the obvious way. Additionally, the imprecision of $\square A$ and $\diamond A$ may create temporal gaps in the configuration that need to be bridged by a weak form of subtyping $A <: B$,

$$\frac{m \leq n}{\circ^m \square A <: \circ^n \square A} \square_{\text{weak}} \qquad \frac{m \geq n}{\circ^m \diamond A <: \circ^n \diamond A} \diamond_{\text{weak}} \qquad \frac{}{A <: A} \text{ refl}$$

This relation is specified to be reflexive and clearly transitive. I extend it to contexts Δ in the obvious manner. In the final rules, I also account for some channels that are not used by P or M but just passed through.

$$\frac{\Delta' <: \Delta \quad [\Delta]_L^{-t} \vdash P :: [c : A]_R^{-t} \quad A <: A'}{\Delta_0, \Delta' \vDash \text{proc}(c, t, P) :: (\Delta_0, c : A')} \text{ proc}$$

$$\frac{\Delta' <: \Delta \quad [\Delta]_L^{-t} \vdash M :: [c : A]_R^{-t} \quad A <: A'}{\Delta_0, \Delta' \vDash \text{msg}(c, t, M) :: (\Delta_0, c : A')} \text{ msg}$$

Type Preservation

With the four rules for typing configurations (empty, compose, proc and msg), type preservation is relatively straightforward. We need some standard lemmas about being able to split a configuration and be able to move a provider (whether process or message) to the right in a typing derivation until it rests right next to its client. Regarding time shifts, we need the following properties.

Lemma 3 (Time Shift).

- (i) If $[A]_L^{-t} = [B]_R^{-t}$ and both are defined then $A = B$.
- (ii) $[[A]_L^{-t}]_L^{-s} = [A]_L^{-(t+s)}$ and if either side is defined, the other is as well.
- (iii) $[[A]_R^{-t}]_R^{-s} = [A]_R^{-(t+s)}$ and if either side is defined, the other is as well.

Theorem 8 (Type Preservation). If $\Delta' \vDash \mathcal{C} :: \Delta$ and $\mathcal{C} \mapsto \mathcal{D}$ then $\Delta' \vDash \mathcal{D} :: \Delta$.

Proof. By case analysis on the transition rule, applying inversion to the given typing derivation, and then assembling a new derivation of \mathcal{D} . \square

Type preservation on basic session types is a simple special case of this theorem.

Global Progress

A process or message is said to be *poised* if it is trying to communicate along the channel that it provides. A poised process is comparable to a value in a sequential language. A configuration is poised if every process or message in the configuration is poised. Conceptually, this implies that the configuration is trying to communicate externally, i.e. along one of the channel it provides. The progress theorem then shows that either a configuration can take a step or it is poised. To prove this I show first that the typing derivation can be rearranged to go strictly from right to left and then proceed by induction over this particular derivation.

The question is how can we prove that processes are either at the same time (for most interactions) or that the message recipient is ready before the message arrives (for **when?**, **now!**, and some forwards)? The key insight here is in the following lemma.

Lemma 4 (Time Inversion).

- (i) If $[A]_R^{-s} = [A]_L^{-t}$ and either side starts with a basic session type constructor then $s = t$.
- (ii) If $[A]_L^{-t} = \square B$ and $[A]_R^{-s} \neq \circ(-)$ then $s \leq t$ and $[A]_R^{-s} = \square B$.
- (iii) If $[A]_R^{-t} = \diamond B$ and $[A]_L^{-s} \neq \circ(-)$ then $s \leq t$ and $[A]_L^{-s} = \diamond B$.

Theorem 9 (Global Progress). If $\cdot \vDash \mathcal{C} :: \Delta$ then either

- (i) $\mathcal{C} \mapsto \mathcal{C}'$ for some \mathcal{C}' , or
- (ii) \mathcal{C} is poised.

Proof. By induction on the right-to-left typing of \mathcal{C} so that either \mathcal{C} is empty (and therefore poised) or $\mathcal{C} = (\mathcal{D} \text{ proc}(c, t, P))$ or $\mathcal{C} = (\mathcal{D} \text{ msg}(c, t, M))$. By induction hypothesis, \mathcal{D} can either take a step (and then so can \mathcal{C}), or \mathcal{D} is poised. In the latter case, we analyze the cases for P and M , applying multiple steps of inversion and Lemma 4 to show that in each case either \mathcal{C} can take a step or is poised. \square

5.5 Further Examples

This section presents example analyses of some of the properties that we can express in the type system, such as the response time of concurrent data structures and the span of a fork/join parallel program.

In some examples I use parametric definitions, both at the level of types and processes. For example, stack_A describes stacks parameterized over a type A , $\text{list}_A[n]$ describes lists of n elements, and $\text{tree}[h]$ describes binary trees of height h . Process definitions are similarly parameterized. They exist as families of ordinary definitions and calculated accordingly, at the metalevel, which is justified since they are only implicitly quantified across whole definitions. This common practice (for example, in work on interaction nets [75]) avoids significant syntactic overhead, highlighting conceptual insight. It is of course possible to internalize such parameters (see, for example, work on refinement of session types [82] or explicitly polymorphic session types [43, 81]).

Response Times: Stacks and Queues

To analyze response times, I present concurrent stacks and queues. A stack data structure provides a client with a choice between a push and a pop. After a push, the client has to send an element, and the provider will again behave like a stack. After a pop, the provider will reply either with the label *none* and terminate (if there are no elements in the stack), or send an element and behave again like a stack. In the cost-free model, this is expressed in the following session type.

$$\text{stack}_A = \&\{ \text{push} : A \multimap \text{stack}_A, \\ \text{pop} : \oplus\{ \text{none} : \mathbf{1}, \text{some} : A \otimes \text{stack}_A \} \}$$

A stack is implemented as a chain of processes. The bottom to the stack is defined by the process *empty*, while a process *elem* holds a top element of the stack as well as a channel with access to the top of the remainder of the stack.

$$x : A, t : \text{stack}_A \vdash \text{elem} :: (s : \text{stack}_A) \\ \cdot \vdash \text{empty} :: (s : \text{stack}_A)$$

The cost model I would like to consider here is \mathcal{RS} where both receives and sends cost one unit of time. Because a receive costs one unit, every continuation type must be delayed by one tick of the clock, which is denoted by prefixing continuations by the \circ modality. This delay is not an artifact of the implementation, but an inevitable part of the cost model—one reason I have distinguished the synonyms *tick* (delay of one, due to the cost model) and *delay* (delay of one, to correctly time the interactions). In this section of examples I will make the same distinction for the next-time modality: I write $\prime A$ for a step in time mandated by the cost model, and $\circ A$ for a delay necessitated by a particular set of process definitions.

As a first approximation,

$$\text{stack}_A = \&\{ \text{push} : \prime(A \multimap \prime \text{stack}_A), \\ \text{pop} : \prime \oplus\{ \text{none} : \prime \mathbf{1}, \text{some} : \prime(A \otimes \prime \text{stack}_A) \} \}$$

There are several problems with this type. The stack is a data structure and has little or no control over *when* elements will be pushed onto or popped from the stack. Therefore a type $\square \text{stack}_A$ should be used to indicate that the client can choose the times of interaction with the stack. While the elements are held by the stack time advances in an indeterminate manner. Therefore, the elements stored in the stack must also have type $\square A$, not A (so that they are always available).

$$\text{stack}_A = \&\{ \text{push} : \prime(\square A \multimap \prime \square \text{stack}_A),$$

$$\text{pop} : ' \oplus \{ \text{none} : ' \mathbf{1}, \text{some} : ' (\square A \otimes ' \square \text{stack}_A) \} \}$$

$$x : \square A, t : \square \text{stack}_A \vdash \text{elem} :: (s : \square \text{stack}_A)$$

$$\cdot \vdash \text{empty} :: (s : \square \text{stack}_A)$$

This type expresses that the data structure is very efficient in its response time: there is no additional delay after it receives a push and then an element of type $\square A$ before it can take the next request, and it will respond immediately to a pop request. It may not be immediately obvious that such an efficient implementation actually exists in the \mathcal{RS} cost model, but it does. I use the implicit form from [62] omitting the `tick` constructs after each receive and send, and also the `when?` before each case that goes along with type $\square A$.

$$s \leftarrow \text{elem} \leftarrow x \ t =$$

$$\text{case } s \ (\text{push} \Rightarrow y \leftarrow \text{recv } s ;$$

$$\quad s' \leftarrow \text{elem} \leftarrow x \ t ; \quad \% \text{ previous top of stack, holding } x$$

$$\quad s \leftarrow \text{elem} \leftarrow y \ s' \quad \% \text{ new top of stack, holding } y$$

$$\quad | \text{pop} \Rightarrow s.\text{some} ;$$

$$\quad \text{send } s \ x ; \quad \% \text{ send channel } x \text{ along } s$$

$$\quad s \leftarrow t) \quad \% s \text{ is now provided by } t, \text{ via forwarding}$$

$$s \leftarrow \text{empty} =$$

$$\text{case } s \ (\text{push} \Rightarrow y \leftarrow \text{recv } s ;$$

$$\quad e \leftarrow \text{empty} ; \quad \% \text{ new bottom of stack}$$

$$\quad s \leftarrow \text{elem} \leftarrow y \ e$$

$$\quad | \text{pop} \Rightarrow s.\text{none} ;$$

$$\quad \text{close } s)$$

The specification and implementation of a queue is very similar. The key difference in the implementation is that when a new element is received, it is passed along the chain of processes until it reaches the end. So instead of

$$s' \leftarrow \text{elem} \leftarrow x \ t ; \quad \% \text{ previous top of stack, holding } x$$

$$s \leftarrow \text{elem} \leftarrow y \ s' \quad \% \text{ new top of stack, holding } y$$

I write

$$t.\text{enq} ;$$

$$\text{send } t \ y ; \quad \% \text{ send } y \text{ to the back of the queue}$$

$$s \leftarrow \text{elem} \leftarrow x \ t$$

in the push branch of `elem` process. These two send operations take two units of time, which must be reflected in the type: after a channel of type $\square A$ has been received, there is a delay of an additional two units of time before the provider can accept the next request.

$$\text{queue}_A = \&\{ \text{enq} : ' (\square A \multimap ' \circ\circ \square \text{queue}_A),$$

$$\quad \text{deq} : ' \oplus \{ \text{none} : ' \mathbf{1}, \text{some} : ' (\square A \otimes ' \square \text{queue}_A) \} \}$$

$$x : \square A, t : \circ\circ \square \text{queue}_A \vdash \text{elem} :: (s : \square \text{queue}_A)$$

$$\cdot \vdash \text{empty} :: (s : \square \text{queue}_A)$$

Time reconstruction will insert the additional delays in the `empty` process through subtyping, using $\square \text{queue}_A \leq \circ\circ \square \text{queue}_A$. I have syntactically expanded the tail call so the second use of subtyping is more apparent.

$$s \leftarrow \text{empty} =$$

case s (enq $\Rightarrow y \leftarrow \text{recv } s$;	$\% y : \Box A \vdash s : \circ\circ\Box\text{queue}_A$
$e \leftarrow \text{empty}$;	$\% y : \Box A, e : \Box\text{queue}_A \vdash s : \circ\circ\Box\text{queue}_A$
$s' \leftarrow \text{elem} \leftarrow y e$;	$\% \Box\text{queue}_A \leq \circ\circ\Box\text{queue}_A$ (on e)
$s \leftarrow s'$	$\% \Box\text{queue}_A \leq \circ\circ\Box\text{queue}_A$ (on s')
deq $\Rightarrow s.\text{none}$;	
close s)	

The difference between the *response times* of stacks and queues in the cost model is minimal: both are constant, with the queue being two units slower. This is in contrast to the total work [61] which is constant for the stack but linear in the number of elements for the queue.

This difference in response times can be realized by typing clients of both stacks and queues. Compare clients S_n and Q_n that insert n elements into a stack and queue, respectively, send the result along channel d , and then terminate. I show only their type below, omitting the implementations.

$x_1 : \Box A, \dots, x_n : \Box A, s : \Box\text{stack}_A \vdash S_n :: (d : \circ^{2n} (\Box\text{stack}_A \otimes '1))$
 $x_1 : \Box A, \dots, x_n : \Box A, s : \Box\text{queue}_A \vdash Q_n :: (d : \circ^{4n} (\Box\text{queue}_A \otimes '1))$

The types demonstrate that the total execution time of S_n is only $2n + 1$, while it is $4n + 1$ for Q_n . The difference comes from the difference in response times. Note that we can infer precise execution times, even in the presence of the \Box modality in the stack and queue types.

Span Analysis: Trees

I use trees to illustrate an example that is typical for fork/join parallelism and computation of *span*. In order to avoid integers, I just compute the parity of a binary tree of height h with boolean values at the leaves. I do not show the obvious definition of *xor*, which in the \mathcal{RS} cost model requires a delay of four from the first input.

$\text{bool} = \oplus\{ b0 : '1, b1 : '1 \}$

$a : \text{bool}, b : \circ^2 \text{bool} \vdash \text{xor} :: (c : \circ^4 \text{bool})$

In the definition of *leaf* and *node* I have explicated the delays inferred by time reconstruction, but not the *tick* delays. The type of $\text{tree}[h]$ gives the *span* of this particular parallel computation as $5h + 2$. This is the time it takes to compute the parity under maximal parallelism, assuming that *xor* takes 4 cycles as shown in the type above.

$\text{tree}[h] = \&\{ \text{parity} : \circ^{5h+2} \text{bool} \}$

$\cdot \vdash \text{leaf} :: (t : \text{tree}[h])$

$t \leftarrow \text{leaf} =$

case t (parity \Rightarrow	$\% \cdot \vdash t : \circ^{5h+2} \text{bool}$
$\% \text{delay}^{5h+2}$	$\% \cdot \vdash t : \text{bool}$
$t.\text{b0}$;	$\% \cdot \vdash t : \mathbf{1}$
close t)	

$l : \circ^1 \text{tree}[h], r : \circ^3 \text{tree}[h] \vdash \text{node} :: (t : \text{tree}[h + 1])$

$t \leftarrow \text{node} \leftarrow l r =$

case t (parity \Rightarrow	$\% l : \text{tree}[h], r : \circ^2 \text{tree}[h] \vdash t : \circ^{5(h+1)+2} \text{bool}$
---------------------------------	---

```

l.parity ;           % l :  $\circ^{5h+2}$  bool, r :  $\circ^1$ tree[h]  $\vdash$  t :  $\circ^{5(h+1)+1}$  bool
% delay             % l :  $\circ^{5h+1}$  bool, r : tree[h]  $\vdash$  t :  $\circ^{5h+5}$  bool
r.parity ;         % l :  $\circ^{5h}$  bool, r :  $\circ^{5h+2}$  bool  $\vdash$  t :  $\circ^{5h+4}$  bool
% delay5h        % l : bool, r :  $\circ^2$  bool  $\vdash$  t :  $\circ^4$  bool
t  $\leftarrow$  xor  $\leftarrow$  l r )

```

The type $l : \circ \text{tree}[h]$ comes from the fact that, after receiving a parity request, it is first sent out the parity request to the left subtree l . The type $r : \circ^3 \text{tree}[h]$ is determined from the delay of 2 between the two inputs to xor . The magic number 5 in the type of tree was derived in reverse from setting up the goal of type-checking the *node* process under the constraints already mentioned. It can also be thought of as 4+1, where 4 is the time to compute the exclusive or at each level and 1 as the time to propagate the parity request down each level.

As is often done in abstract complexity analysis, I can also impose an alternative cost model. For example, I may count only the number of calls to xor while all other operations are cost free. Then I would have

```

a : bool, b : bool  $\vdash$  xor :: (c :  $\circ$ bool)      ·  $\vdash$  leaf :: (t : tree[h])
tree[h] = &{ parity :  $\circ^h$  bool }             l : tree[h], r : tree[h]  $\vdash$  node :: (t : tree[h + 1])

```

with the same code but different times and delays from before. The reader is invited to reconstruct the details.

5.6 Related Work

Most closely related is work on space and time complexity analysis of interaction nets by Gimenez and Moser [75], which is a parallel execution model for functional programs. While also inspired by linear logic and, in particular, proof nets, it treats only special cases of the additive connectives and recursive types and does not have analogues of the \square and \diamond modalities. It also does not provide a general source-level programming notation with a syntax-directed type system. On the other hand it incorporates sharing and space bounds, which are beyond the scope of this work.

Session types and process calculi. Another related thread is the research on timed multiparty session types [38] for modular verification of real-time choreographic interactions. Their system is based on explicit global timing interval constraints, capturing a new class of communicating timed automata, in contrast to our system based on binary session types in a general concurrent language. Therefore, their system has no need for general \square and \diamond modalities, the ability to pass channels along channels, or the ability to identify channels via forwarding. Their work is complemented by an expressive dynamic verification framework in real-time distributed systems [124], which I do not consider. Semantics counting communication costs for work and span in session-typed programs were given by Silva et al. [146], but no techniques for analyzing them were provided.

In addition to the work on timed multiparty session types, time has been introduced into the π -calculus (see, for example, Saeedloei and Gupta [142]) or session-based communication primitives (see, for example, López et al. [115]) but generally these works do not develop a type system. Kobayashi [102] extends a (synchronous) π -calculus with means to count parallel

reduction steps. He then provides a type system to verify time-boundedness. This is more general in some dimension than our work because of a more permissive underlying type and usage system, but it lacks internal and external choice, genericity in the cost model, and provides bounds rather than a fine gradation between exact and indefinite times. Session types can also be derived by a Curry-Howard interpretation of *classical linear logic* [159] but I am not aware of temporal extensions. I conjecture that there is a classical version of our system where \square and \diamond are dual and \circ is self-dual.

Reactive programming. Synchronous data flow languages such as Lustre [86], Esterel [34], or Lucid Synchrone [133] are time-synchronous with uni-directional flow and thus may be compared to the fragment of our language with internal choice (\oplus) and the next-time modality ($\circ A$), augmented with existential quantification over basic data values like booleans and integers (which we have omitted here only for the sake of brevity). The global clock would map to our underlying notion of time, but data-dependent local clocks would have to be encoded at a relatively low level using streams of option type, compromising the brevity and elegance of these languages. Furthermore, synchronous data flow languages generally permit sharing of channels, which, although part of many session-typed languages [25, 42], require further investigation with temporal modalities. On the other hand, I support a number of additional types such as external choice ($\&$) for bidirectional communication and higher-order channel-passing ($A \multimap B$, $A \otimes B$). In the context of functional reactive programming, a Nakano-style [123] temporal modality has been used to ensure productivity [104]. A difference in my work is that I consider concurrent processes and that the types prescribe the timing of messages.

Computational interpretations of $\circ A$. A first computational interpretation of the next-time modality under a proofs-as-programs paradigm was given by Davies [65]. The basis is natural deduction for a (non-linear!) intuitionistic linear-time temporal logic with only the next-time modality. Rather than capturing cost, the programmer could indicate *staging* by stipulating that some subexpressions should be evaluated “at the next time”. The natural operational semantics then is a logically-motivated form of *partial evaluation* which yields a residual program of type $\circ A$. This idea was picked up by Feltman et al. [69] to instead *split* the program statically into two stages where results from the first stage are communicated to the second. Again, neither linearity (in the sense of linear logic), nor any specific cost semantics appears in this work.

Other techniques. Computing the cost of concurrent programs is a fundamental problem in resource analysis. Hoffmann and Shao [90] introduced the first automatic analysis for deriving bounds on the worst-case evaluation cost of parallel first-order functional programs. Their main limitation is that they can only handle parallel computation; they don’t support message-passing or shared memory based concurrency. Blelloch and Reid-Miller [37] utilize pipelining [127] to improve the complexity of parallel algorithms. However, they use futures [87], a parallel language construct to implement pipelining without the programmer having to specify them explicitly. The runtime of algorithms is determined by analyzing the work and depth in a language-based cost model. The work relates to ours in the sense that pipelines can have delays,

which can be data dependent. However, the algorithms they analyze have no message-passing concurrency or other synchronization constructs. Albert et al. [14] devised a static analysis for inferring the parallel cost of distributed systems. They first perform a block-level analysis to estimate the serial cost, then construct a distributed flow graph (DFG) to capture the parallelism and then obtain the parallel cost by computing the maximal cost path in the DFG. However, the bounds they produce are modulo a points-to and serial cost analysis. Hence, an imprecise points-to analysis will result in imprecise parallel cost bounds. Moreover, since their technique is based on static analysis, it is not compositional and a whole program analysis is needed to infer bounds on each module. Recently, a bounded linear typing discipline [74] modeled in a semiring was proposed for resource-sensitive compilation. It was then used to calculate and control execution time in a higher-order functional programming language. However, this language did not support recursion.

5.7 Future Directions

I have presented a system of temporal session types that can accommodate and analyze concurrent programs with respect to a variety of different cost models. Types can vary in precision, based on desired and available information, and includes latency, rate, response time, and span of computations. It is constructed in a modular way, on top of a system of basic session types, and therefore lends itself to easy generalization. I have illustrated the type system through a number of simple programs on streams of bits, binary counters, lists, stacks, queues, and trees. I mention some of the further challenges that need to be addressed in this domain of temporal session types.

Inference Inference of time bounds will make resource-aware session types more practical and usable. The most severe difficulty here is the \circ operator. Computing the number of \circ operators to insert at a program point is non-trivial. The idea here would be similar to work inference. The inference engine first inserts a parametric amount of delay and then the type-checker determines the constraints on the inserted delays. A solver then tries to determine the value of said delays. The \square and \diamond operators do not involve any parameters and they should be easier to handle.

Dependent Types Time bounds are often dependent on the type refinements applied to session types. However, these dependencies bring along their own challenges. They require an arithmetic solver engine in the type checker, along with a system that allows parameters in types and process definitions. Moreover, they exacerbate the non-determinism in subtyping.

Part II

Application of Resource-Aware Session Types

Chapter 6

Programming Digital Contracts

The first part of this thesis develops the theory of resource-aware session types. The second major contribution of this thesis is to design a language based on resource-aware session types and apply it to the domain of digital contracts. This chapter describes how resource-aware session types are a perfect fit for smart contracts and cleanly addresses the core challenges developers face in this domain. This chapter also presents the type theory of this language, its various components, and type soundness theorems.

6.1 Introduction

Digital contracts are computer protocols that describe and enforce the execution of a contract. With the rise of blockchains and cryptocurrencies such as Bitcoin [122], Ethereum [161], and Tezos [79], digital contracts have become popular in the form of smart contracts, which provide potentially distrusting parties with programmable money and an enforcement mechanism that does not rely on third parties. Smart contracts have been used to implement auctions [1], investment instruments [121], insurance agreements [98], supply chain management [108], and mortgage loans [120]. In general, digital contracts hold the promise to reduce friction, lower cost, and broaden access to financial infrastructure.

Smart contracts have not only shed light on the benefits of digital contracts but also on their potential risks. Like all software, smart contracts can contain bugs and security vulnerabilities [22], which can have direct financial consequences. A well-known example, is the attack on The DAO [121], resulting in a multi-million dollar theft by exploiting a contract vulnerability. Maybe even more important than the direct financial consequences is the potential erosion of trust as a result of such failures.

Contract languages today are derived from existing general-purpose programming languages like JavaScript (Ethereum’s Solidity [54]), Go (in the Hyperledger project [41]), or OCaml (Tezos’ Liquidity [4]). While this makes contract languages look familiar to software developers, it is inadequate to accommodate the domain-specific requirements of digital contracts.

- Instead of centering contracts on their interactions with users, the high-level protocol of the intended interactions with a contract is buried in the implementation code, hampering understanding, formal reasoning, and trust.

- Resource (or *gas*) usage of digital contracts is of particular importance for transparency and consensus. However, obliviousness of resource usage in existing contract languages makes it hard to predict the cost of executing a contract and prevent denial-of-service vulnerabilities.
- Existing languages fail to enforce linearity of assets, endangering the validity of a contract when assets get duplicated or deleted, accidentally or maliciously [118].

As a result, developing a correct smart contract is no easier than developing bug-free software in general. Additionally, vulnerabilities are harder to fix, because changes in the code may proliferate into changes in the contract itself.

This chapter presents the *type-theoretic foundations* of Nomos, a programming language for digital contracts whose genetics match the domain-specific requirements to provide strong static guarantees that facilitate the design of correct contracts. In particular, Nomos’ type system makes explicit the protocols governing a contract, provides static bounds on the resource cost of interacting with a contract, and enforces a linear treatment of a contract’s assets.

To express and enforce the protocols underlying a contract, Nomos is based on *resource-aware session types* [61]. The types describe the protocol of interaction between users and contracts and serve as a high-level description of the functionality of the contract. Type checking can be automated and guarantees that Nomos programs follow the given protocol. In this way, the key functionality of the contract is visible in the type, and contract development is centered on the interaction of the contract with the world. In addition to the interaction, resource-aware types also make the transaction cost visible in the type. This makes transactions transparent in their resource usage, and type checking again guarantees that the transactions do not exceed the resource usage prescribed by the type.

To eliminate a class of bugs in which the internal state of a contract loses track of its assets or performs unintended transactions, Nomos integrates a linear type system [158] into a functional language. Linear type systems use the ideas of Girard’s linear logic [76] to ensure that certain data is neither duplicated nor discarded by a program. Programming languages such as Rust [7] have demonstrated that substructural type systems are practical in industrial-strength languages. Moreover, linear types are compatible with session types, which are themselves based on linear logic [25, 42, 129, 153, 159].

In addition to the design of the Nomos language, this chapter makes the following technical contributions.

1. Linear session types that support controlled sharing [25, 26] have been integrated into a conventional functional type system. To leave the logical foundation intact, the integration is achieved by a *contextual monad* [153] (Section 6.4) that gives process expressions first-class status in the functional language. Moreover, shared session types [25] are recast to accommodate the explicit notions of *contracts* and *clients* (Section 6.3).
2. I prove the type soundness of Nomos with respect to a novel asynchronous cost semantics using progress and preservation (Section 6.6).

6.2 Nomos by Example

This section provides an overview of the main features of Nomos based on a simple auction contract.

Explicit Protocols of Interaction Digital contracts, like traditional contracts, follow a *pre-defined protocol*. For instance, an auction contract distinguishes a bidding phase, where bidders submit their bids, possibly multiple times, from a subsequent collection phase, where the highest bidder receives the lot while all other bidders receive their bids back. In Solidity [54], the bidding phase of an auction is typically implemented as the bid function below. This function receives a bid (`msg.value`) from a bidder (`msg.sender`) and adds it to the bidder’s total previous bids (`bidValue`).

```
function bid() public payable {
  require (status == running);
  bidder = msg.sender; bid = msg.value;
  bidValue[bidder] = bidValue[bidder] + bid; }
```

To guarantee that a bid can only be placed in the bidding phase, the contract uses the state variable `status` to track the different phases of a contract. The `require` statement tests whether the auction is still running and thus accepts bids. It is checked at run-time and aborts the execution if the condition is not met. It is the responsibility of the programmer to define state variables, update them, and introduce corresponding guards.

Rather than burying the contract’s interaction protocol in implementation code by means of state variables and run-time checks, Nomos allows the explicit expression and static enforcement of protocols with *session types*. The auction’s protocol amounts to the below session type:

$$\text{auction} = \uparrow_{\mathbb{L}}^{\mathbb{S}} \triangleleft^{22} \oplus \{ \text{running} : \& \{ \text{bid} : \text{id} \rightarrow \text{money} \multimap \downarrow_{\mathbb{L}}^{\mathbb{S}} \text{auction}, \quad \% \text{rcv bid from client} \\ \text{cancel} : \triangleright^{21} \downarrow_{\mathbb{L}}^{\mathbb{S}} \text{auction} \}, \quad \% \text{client canceled} \\ \text{ended} : \& \{ \text{collect} : \text{id} \rightarrow \oplus \{ \text{won} : \text{lot} \otimes \downarrow_{\mathbb{L}}^{\mathbb{S}} \text{auction}, \quad \% \text{client won} \\ \text{lost} : \text{money} \otimes \triangleright^7 \downarrow_{\mathbb{L}}^{\mathbb{S}} \text{auction} \}, \quad \% \text{client lost} \\ \text{cancel} : \triangleright^{21} \downarrow_{\mathbb{L}}^{\mathbb{S}} \text{auction} \} \}$$

We first focus on how the session type defines the main interactions of a contract with a bidder and ignore the operators $\uparrow_{\mathbb{L}}^{\mathbb{S}}$, $\downarrow_{\mathbb{L}}^{\mathbb{S}}$, \triangleleft , and \triangleright for now. To distinguish the two main phases an auction can be in, the session type uses an internal choice (\oplus), leading the contract to either send the label `running` or `ended`, depending on whether the auction still accepts bids or not, respectively. Dual to an internal choice is an external choice ($\&$), which leaves the choice to the client (i.e., bidder) rather than the provider (i.e., contract). For example, in case the auction is running, the client can choose between placing a bid (label `bid`) or backing out (`cancel`). In the former case, the client indicates their identifier (type `id`), followed by a payment (type `money`). Nomos session types allow transfer of both non-linear (e.g., `id`) and linear assets (e.g., `money`), using the operators arrow (\rightarrow) and (\multimap), respectively. Should the auction have ended, the client can choose to check their outcome (label `collect`) or back out (`cancel`). In the case of `collect`, the auction will answer with either `won` or `lost`. In the former case, the auction will send the lot, in the latter case, it will return the client’s bid. The linear product (\otimes) is dual to \multimap and denotes

the transfer of a linear value from the contract to the client. The auction type guarantees that a client cannot collect during the running phase, while they cannot bid during the ended phase.

Nomos uses *shared* session types [25] to guarantee that bidders interact with the auction in mutual exclusion from each other and that the sequences of actions are executed *atomically*. To demarcate the parts of the protocol that become a *critical section*, the above session type uses the \uparrow_L^S and \downarrow_L^S modalities. The \uparrow_L^S modality denotes the beginning of a critical section, the \downarrow_L^S modality denotes its end. Programmatically, \uparrow_L^S translates into an *acquire* of the auction session and \downarrow_L^S into its *release*, which is only sound if the protocol behaves like an auction afterwards (*equi-synchronizing* type).

Contracts are implemented by *processes*, revealing the concurrent, message-passing nature of session-typed languages. The process *run* below implements the auction's running phase. Line 2 gives the process' signature, indicating that it offers a shared session of type *auction* along the channel *sa* and uses a linear hash map $b : \text{hashmap}_{\text{id}, \text{bid}}$ of bids indexed by id and a linear lot *l*. The bid session type (line 1) can be queried for the stored identifier and bid value, and is offered by a process (not shown) that internally stores this identifier and money. Line 4 onward list the process body. Line 1 defines session types *bid* and *bids*, respectively.

```

1: stype bid = &\{addr : id × bid, val : money\},    stype bids = hashmapid,bid
2: (b : bids), (l : lot) ⊢ run :: (sa : auction)      % syntax for process declaration
3:   sa ← run b l =                                  % syntax for process definition
4:     la ← accept sa ;                               % accept a client acquire request
5:     la.running ;                                   % auction is running
6:     case la ( bid ⇒ r ← recv la ;                  % receive identifier r : id
7:               m ← recv la ;                        % receive bid m : money
8:               sa ← detach la ;                    % detach from client
9:               b' ← addbid r b m ;                 % store bid internally
10:              sa ← check b' l                     % check if threshold reached
11:             | cancel ⇒ sa ← detach la ;          % detach from client
12:             sa ← run b l)                         % recurse

```

The contract process first *accepts* an acquire request by a bidder (line 4) and then sends the message *running* (line 5), indicating the auction status and waiting for the bidder's choice. Should the bidder choose to make a bid, the process waits to receive the bidder's identifier (line 6) followed by money equivalent to the bidder's bid (line 7). After this linear exchange, the process leaves the critical section by issuing a *detach* (line 8), matching the bidder's release request. Internally, the process stores the pair of the bidder's identifier and bid in the data structure *bids* (line 9). The ended protocol of the contract is governed by a different process (not shown), responsible for distributing the bids back to the clients. The contract transitions to the ended state when the number of bidders reaches a threshold (stored in *auction*). This is achieved by the *check* process (line 10) which checks if the threshold has been reached and makes this transition, or calls *run* otherwise. Should the bidder choose to cancel, the contract simply detaches and recurses (lines 11,12).

Re-Entrancy Vulnerabilities A contract function is re-entrant if, once called by a user, it can potentially be called again before the previous call has completed. As an illustration, con-

sider the following collect function of the auction contract in Solidity where the funds are transferred to the bidder before the hash map is updated to reflect this change.

```
function collect() public payable {
    require (status == ended);
    bidder = msg.sender; bid = bidValue[bidder];
    bidder.send(bid); bidValue[bidder] = 0; }

function () payable {
    // 'auction' variable stores the
    // address to auction contract
    auction.collect(); }
```

A bidder can now cause re-entrancy by creating a dummy contract with an unnamed *fallback* function (on the bottom) that calls the auction's collect function. This call is triggered when collect calls send (last line in collect), leading to an infinite recursive call to collect, depleting all funds from the auction. The message-passing framework of session types eliminates this vulnerability. While session types provide multiple clients access to a contract, the acquire-release discipline ensures that clients interact with the contract in mutual exclusion. To attempt re-entrancy, a bidder will need to acquire the auction contract twice without releasing it.

Linear Assets Nomos integrates a linear type system that tracks the assets stored in a process. The type system enforces that assets are never duplicated, but only exchanged between processes. Moreover, the type system prevents a process from terminating while it holds linear assets. For example, the auction contract treats money and lot as linear assets, which is witnessed by the use of the linear operators \multimap and \otimes for their exchange. In contrast, no provisions to handle assets linearly exist in Solidity, allowing such assets to be created out of thin air, duplicated, or discarded. In the above bid function, for instance, the language does not prevent the programmer from writing `bidValue[bidder] = bid` instead, losing the bidder's previous bid.

Resource Cost Another important aspect of digital contracts is their *resource usage*. On a blockchain, executing a contract function, or *transaction*, requires new blocks to be added to the blockchain. In existing blockchains like Ethereum, this is done by *miners* who charge a fee based on the *gas* usage of the transaction, indicating the cost of its execution. Precisely computing this cost statically is important because the sender of a transaction must pay this fee to the miners along with sending the transaction. If the sender does not pay a sufficient amount, the transaction will be rejected by the miners and the sender's fee is lost!

Nomos uses resource-aware session types [61] to statically analyze the resource cost of a transaction. They operate by assigning an initial *potential* to each process. This potential is consumed by each operation that the process executes or can be transferred between processes to share and amortize cost. The cost of each operation is defined by a cost model. If the cost model assigns a cost to each operation as equivalent to their gas cost during execution, the potential consumed during a transaction reflects an upper bound on the gas usage.

Resource-aware session types express the potential as part of the session type using the operators \triangleleft and \triangleright . The \triangleleft operator prescribes that the client must send potential to the contract,

with the amount of potential indicated as a superscript. Dually, \triangleright prescribes that the contract must send potential to the client. In the case of the auction contract, we require the client to pay potential for the operations that the contract must execute, both while placing and collecting their bids. If the cost model assigns a cost of 1 to each contract operation, then the maximum cost of an auction session is 22 (taking the max number of operations in all branches). Thus, we require the client to send 22 units of potential at the start of a session using \triangleleft^{22} . In the lost branch of the auction type, on the other hand, the contract returns 7 units of potential to the client using \triangleright^7 . This simulates gas usage in smart contracts, where the sender initiates a transaction with some initial gas, and the leftover gas at the end of the transaction is returned to the sender. In contrast to existing smart contract languages like Solidity, which provide no support for analyzing the cost of a transaction, Nomos’ type checker has automatically inferred these potential annotations and guarantees that well-typed transactions cannot run out of gas.

Bringing It All Together Combining all these features soundly in one language is challenging. In Nomos, we achieve this by using different *typing judgments* and *modes*, identifying the role of the process offered along that channel. The mode R denotes *purely linear processes* for linear assets or private data structures, such as b and l in the auction. The modes S and L denote *sharable processes*, i.e., contracts, that are either in their shared or linear phase such as sa and la , respectively. The mode T denotes a *transaction process* that can refer to shared and linear processes and is issued by a user, such as bidder in the auction. The mode assignment carries over into the process typing judgments imposing invariants (Definition 7) that are key to type safety. The mode annotations are automatically inferred by the type checker relieving programmers from this burden.

Typing Judgment For typing Nomos processes, we use the judgment $\Psi ; \Gamma ; \Delta \vdash^g P :: (x : A)$. As we introduce each concept, the role of each symbol will become clear. Henceforth, we indicate the current concepts from the judgment in black while the concepts that will be introduced in later sections are marked in blue.

6.3 Sharing Contracts

Multi-user support is fundamental to digital contract development. Linear session types, as defined in Chapter 3, unfortunately preclude such sharing because they restrict processes to exactly one client; only one bidder for the auction, for instance (who will always win!). To support multi-user contracts, we base Nomos on *shared* session types [25]. Shared session types impose an acquire-release discipline on shared processes to guarantee that multiple clients interact with a contract in *mutual exclusion* of each other. When a client acquires a shared contract, it obtains a private linear channel along which it can communicate with the contract undisturbed by any other clients. Once the client releases the contract, it loses its private linear channel and only retains a shared reference to the contract.

A key idea of shared session types is to lift the acquire-release discipline to the type level. Generalizing the idea of type *stratification* [33, 129, 139], session types are stratified into a linear and shared layer with two *adjoint modalities* going back and forth between them:

$$\begin{aligned}
A_R & ::= \oplus\{\ell : A_R\}_{\ell \in L} \mid \&\{\ell : A_R\}_{\ell \in L} \mid \mathbf{1} \mid A_m \multimap_m A_R \mid A_m \otimes_m A_R \\
& \quad \mid \tau \rightarrow A_R \mid \tau \times A_R \\
A_L & ::= \oplus\{\ell : A_L\}_{\ell \in L} \mid \&\{\ell : A_L\}_{\ell \in L} \mid \mathbf{1} \mid A_m \multimap_m A_L \mid A_m \otimes_m A_L \\
& \quad \mid \tau \rightarrow A_L \mid \tau \times A_L \mid \downarrow_L^S A_S \\
A_S & ::= \uparrow_L^S A_L \\
A_T & ::= A_R
\end{aligned}$$

Figure 6.1: Grammar for shared session types

$$\begin{aligned}
A_S & ::= \uparrow_L^S A_L && \text{shared session type} \\
A_L & ::= \dots \mid \downarrow_L^S A_S && \text{linear session types}
\end{aligned}$$

The \uparrow_L^S type modality translates into an *acquire*, while the dual \downarrow_L^S type modality into a *release*. Whereas mutual exclusion is one key ingredient to guarantee session fidelity (a.k.a. type preservation) for shared session types, the other key ingredient is the requirement that a session type is *equi-synchronizing*. A session type is equi-synchronizing if it imposes the invariant on a process to be released back to the same type at which the process was previously acquired. This is also the key behind eliminating *re-entrancy vulnerabilities* since it prevents a user from interrupting an ongoing session in the middle and initiating a new one.

Recall the process typing judgment in Nomos $\Psi ; \Gamma ; \Delta \vdash^q P :: (x_m : A)$ denoting a process P offering service of type A along channel x at mode m . The contexts Γ and Δ store the shared and linear channels that P can refer to, respectively (Ψ and q are explained later and thus marked in blue in Figure 6.2). The stratification of channels into layers arises from a difference in structural properties that exist for types at a mode. Shared propositions exhibit weakening, contraction and exchange, thus can be discarded or duplicated, while linear propositions only exhibit exchange.

Allowing Contracts to Rely on Linear Assets As exemplified by the auction contract, a digital contract typically amounts to a process that is shared at the outset, but oscillates between shared and linear to interact with clients, one at a time. Crucial for this pattern is the ability of a contract to maintain its linear assets (e.g., money or lot for the auction) regardless of its mode. Unfortunately, current shared session types [25] do not allow a shared process to rely on any linear channels, requiring any linear assets to be consumed before becoming shared. This precaution was logically motivated [137] and also crucial for type preservation.

A key novelty of our work is to lift this restriction while *maintaining type preservation*. The main concern regarding preservation is to prevent a process from acquiring its client, which would result in a cycle in the linear process tree. To this end, we factor the process typing judgment according to the *three roles* that arise in digital contract programs: *contracts*, *transactions*, and *linear assets*. Since contracts are shared and thus can oscillate between shared and linear, we get 4 sub-judgments for typing processes, each characterized by the mode of the channel being offered.

Definition 7 (Process Typing). *The judgment $\Psi ; \Gamma ; \Delta \vdash^q P :: (x_m : A)$ is categorized according to mode m . This factorization imposes certain invariants on the judgment outlined below. $\mathbf{L}(A)$ denotes the language generated by the grammar of A .*

1. If $m = R$, then (i) Γ is empty, (ii) for all $d_k \in \Delta \implies k = R$, and (iii) $A \in \mathbf{L}(A_R)$.
2. If $m = S$, then (i) for all $d_k \in \Delta \implies k = R$, and (ii) $A \in \mathbf{L}(A_S)$.
3. If $m = L$, then (i) for all $d_k \in \Delta \implies k = R \vee k = L$, and (ii) $A \in \mathbf{L}(A_L)$.
4. If $m = T$, then $A \in \mathbf{L}(A_T)$.

Figure 6.1 shows the session type grammar in Nomos. The first sub-judgment in Definition 7 is for typing linear assets. These type a purely linear process P using a purely linear context Δ (types belonging to grammar A_R in Figure 6.1) and offering a purely linear type A along channel x_R . The mode R of the channel indicates that a purely linear session is offered. The second and third sub-judgments are for typing contracts. The second sub-judgment shows the type of a contract process P using a shared context Γ and a purely linear channel context Δ (judgment Δ purelin) and offering shared type A on the shared channel x_S . Once this shared channel is acquired by a user, the shared process transitions to its linear phase, whose typing is governed by the third sub-judgment. The offered channel transitions to linear mode L , while the linear context may now contain channels at modes R or L . Finally, the fourth typing judgment types a linear process, corresponding to a *transaction* holding access to shared channels Γ and linear channels Δ , and offering at mode T .

This novel factorization upholds preservation while allowing shared contract processes to rely on linear resources. The modes impose the ordering $R < S < L < T$ among the linear channels in the configuration. A process (offering a channel) at a certain mode is allowed to rely only on processes at the same or lower mode. These are exactly the conditions imposed by Definition 7. This introduces an implicit ordering among the linear processes depending on their mode, thus eliminating cycles in the process tree. Relatedly, shared processes can only refer to shared channels (at mode S) or purely linear channels (at mode R) as exemplified by the judgment Δ purelin in Figure 6.2. Formally, Δ purelin denotes that for all $d_k \in \Delta \implies k = R$. This ensures that a shared contract must release all processes it has acquired before itself being released. This further enforces an ordering in which the channels are acquired and released, thus *allowing contracts to interact with other contracts without compromising type safety*.

Shared session types introduce new typing rules into our system, concerning the *acquire-release* constructs (see Figure 6.2). In rule $\uparrow_L^S L$, an acquire is applied to the shared channel $x_S : \uparrow_L^S A_L$ in Γ and yields a linear channel x_L added to Δ when successful. A contract process can *accept* an acquire request along its offering shared channel x_S . After the accept is successful, the shared contract process transitions to its linear phase, now offering along the linear channel x_L (rule $\uparrow_L^S R$).

The synchronous dynamics of the *acquire-accept* pair is

$$(\uparrow_L^S C) : \text{proc}(a_S, w', x_L \leftarrow \text{accept } a_S ; P_{x_L}), \text{proc}(c_m, w, x_L \leftarrow \text{acquire } a_S ; Q_{x_L}) \mapsto \\ \text{proc}(a_L, w', P_{a_L}), \text{proc}(c_m, w, Q_{a_L})$$

This rule exploits the invariant that a contract process' providing channel a can come at two different modes, a linear one a_L , and a shared one a_S . The linear channel a_L is substituted for the channel variable x_L occurring in the process terms P and Q .

The dual to acquire-accept is *release-detach*. A client can *release* linear access to a contract process, while the contract process *detaches* from the client. The corresponding typing rules are presented in Figure 6.2. The effect of releasing the linear channel x_L is that the continuation Q loses access to x_L , while a new reference to x_S is made available in the shared context Γ . The

$\boxed{\Psi ; \Gamma ; \Delta \Vdash^q P :: (x_m : A)}$ Process P uses shared channels in Γ and offers A along x .

$$\frac{\Psi ; \Gamma ; \Delta, (x_L : A_L) \Vdash^q Q :: (z_m : C)}{\Psi ; \Gamma, (x_S : \uparrow_L^S A_L) ; \Delta \Vdash^q x_L \leftarrow \text{acquire } x_S ; Q :: (z_m : C)} \uparrow_L^S L$$

$$\frac{\Delta \text{ purelin} \quad \Psi ; \Gamma ; \Delta \Vdash^q P :: (x_L : A_L)}{\Psi ; \Gamma ; \Delta \Vdash^q x_L \leftarrow \text{accept } x_S ; P :: (x_S : \uparrow_L^S A_L)} \uparrow_L^S R$$

$$\frac{\Psi ; \Gamma, (x_S : A_S) ; \Delta \Vdash^q Q :: (z_m : C)}{\Psi ; \Gamma ; \Delta, (x_L : \downarrow_L^S A_S) \Vdash^q x_S \leftarrow \text{release } x_L ; Q :: (z_m : C)} \downarrow_L^S L$$

$$\frac{\Delta \text{ purelin} \quad \Psi ; \Gamma ; \Delta \Vdash^q P :: (x_S : A_S)}{\Psi ; \Gamma ; \Delta \Vdash^q x_S \leftarrow \text{detach } x_L ; P :: (x_L : \downarrow_L^S A_S)} \downarrow_L^S R$$

Figure 6.2: Typing rules corresponding to the shared layer.

contract, on the other hand, detaches from the client by transitioning its offering channel from linear mode x_L back to the shared mode x_S . Both right rules $\uparrow_L^S R$ and $\downarrow_L^S R$ require Δ purelin ensuring that a shared process releases all shared channels before themselves being released. Operationally, the release-detach rule is inverse to the acquire-accept rule.

$$(\downarrow_L^S C) : \text{proc}(a_L, w', x_S \leftarrow \text{detach } a_L ; P_{x_S}), \text{proc}(c_m, w, x_S \leftarrow \text{release } a_L ; Q_{x_S}) \mapsto \text{proc}(a_S, w', P_{a_S}), \text{proc}(c_m, w, Q_{a_S})$$

6.4 Adding a Functional Layer

To support general-purpose programming patterns, Nomos combines linear channels with conventional data structures, such as integers, lists, or dictionaries. To reflect and track different classes of data in the type system, we take inspiration from prior work [129, 153] and incorporate processes into a functional core via a *linear contextual monad* that isolates session-based concurrency. To this end, we introduce a separate functional context to the typing of a process. The linear contextual monad encapsulates open concurrent computations, which can be passed in functional computations but also transferred between processes in the form of *higher-order processes*, providing a uniform integration of higher-order functions and processes.

The types are separated into a functional and concurrent part, mutually dependent on each other. The functional types τ are given by the type grammar below.

$$\begin{aligned} \tau ::= & \tau \rightarrow \tau \mid \tau + \tau \mid \tau \times \tau \mid \text{int} \mid \text{bool} \mid L^q(\tau) \\ & \mid \{A_R \leftarrow \overline{A_R}\}_R \mid \{A_S \leftarrow \overline{A_S} ; \overline{A_R}\}_S \mid \{A_T \leftarrow \overline{A_S} ; \overline{A}\}_T \end{aligned}$$

The types are standard, except for the potential annotation $q \in \mathbb{N}$ in list type $L^q(\tau)$, which we explain in Section 6.5, and the contextual monadic types in the last line, which are the topic of this section. The expressivity of the types and terms in the functional layer are not important

for the development in this paper. Thus, we do not formally define functional terms M but assume that they have the expected term formers such as function abstraction and application, type constructors, and pattern matching. We also define a standard type judgment for the functional part of the language.

$\Psi \Vdash^p M : \tau$ term M has type τ in functional context Ψ (potential p discussed later)

Contextual Monad The main novelty in the functional types are the three type formers for contextual monads, denoting the type of a process expression. The type $\{A_R \leftarrow \overline{A_R}\}_R$ denotes a process offering a *purely linear* session type A_R and using the purely linear vector of types $\overline{A_R}$. The corresponding introduction form in the functional language is the monadic value constructor $\{c_R \leftarrow P \leftarrow \overline{d_R}\}$, denoting a runnable process offering along channel c_R that uses channels $\overline{d_R}$, all at mode R. The corresponding typing rule for the monad is (ignore the blue portions)

$$\frac{\Delta = \overline{d_R : D} \quad \Psi ; \cdot ; \Delta \Vdash^g P :: (x_R : A)}{\Psi \Vdash^g \{x_R \leftarrow P \leftarrow \overline{d_R}\} : \{A \leftarrow \overline{D}\}_R} \{\} I_R$$

The monadic *bind* operation implements process composition and acts as the elimination form for values of type $\{A_R \leftarrow \overline{A_R}\}_R$. The bind operation, written as $c_R \leftarrow M \overline{d_R} ; Q_c$, composes the process underlying the monadic term M , which offers along channel c_R and uses channels $\overline{d_R}$, with Q_c , which uses c_R . The typing rule for the monadic bind is rule $\{\} E_{RR}$ in Figure 6.3. The linear context is split between the monad M and continuation Q , enforcing linearity. Similarly, the potential in the functional context is split using the sharing judgment (\forall), explained in Section 6.5. The shared context Γ is empty in accordance with the invariants established in Definition 7 (i), since the mode of offered channel x is R. The effect of executing a bind is the spawn of the purely linear process corresponding to the monad M , and the parent process continuing with Q . The corresponding operational semantics rule (named spawn_{RR}) is given as follows:

$$\text{proc}(d_R, w, x_R \leftarrow \{x'_R \leftarrow P_{x'_R, \overline{y}} \leftarrow \overline{y}\} \overline{a} ; Q) \mapsto \text{proc}(c_R, 0, P_{c_R, \overline{a}}), \text{proc}(d_R, w, [c_R/x_R]Q)$$

The above rule spawns the process P offering along a globally fresh channel c_R , and using channels \overline{a} . The continuation process Q acts as a client for this fresh channel c_R . The other two monadic types correspond to spawning a shared process $\{A_S \leftarrow \overline{A_S} ; \overline{A_R}\}_S$ and a transaction process $\{A_T \leftarrow \overline{A_S} ; \overline{A}\}_T$ at mode S and T, respectively. Their rules are analogous to $\{\} I_R$ and $\{\} E_{RR}$.

Value Communication Communicating a *value* of the functional language along a channel is expressed at the type level by adding the following two session types.

$$A ::= \dots \mid \tau \rightarrow A \mid \tau \times A$$

The type $\tau \rightarrow A$ prescribes receiving a value of type τ with continuation type A , while its dual $\tau \times A$ prescribes sending a value of type τ with continuation A . The corresponding typing rules for arrow ($\rightarrow R, \rightarrow L$) are given in Figure 6.3 (rules for \times are inverse). Receiving a value adds it to the functional context Ψ , while sending it requires proving that the value has type τ . Semantically, sending a value $M : \tau$ creates a message predicate along a fresh channel c_m^+ containing the value:

$\boxed{\Psi ; \Gamma ; \Delta \vdash^g P :: (x_m : A)}$ Process P uses functional values in Ψ .

$$\frac{\Psi_1 \Vdash^p M : \{A \leftarrow \overline{D}\} \quad \Delta = \overline{d_R : D} \quad \Psi \Downarrow (\Psi_1, \Psi_2) \quad \Psi_2 ; \cdot ; \Delta', (x_R : A) \vdash^g Q :: (z_R : C)}{\Psi ; \cdot ; \Delta, \Delta' \vdash^r x_R \leftarrow M \overline{d_R} ; Q :: (z_R : C)} \{ \} E_{RR}$$

$$\frac{\Psi, (y : \tau) ; \Gamma ; \Delta \vdash^g P :: (x_m : A)}{\Psi ; \Gamma ; \Delta \vdash^g y \leftarrow \text{recv } x_m ; P :: (x_m : \tau \rightarrow A)} \rightarrow R$$

$$\frac{r = p + q \quad \Psi \Downarrow (\Psi_1, \Psi_2) \quad \Psi_1 \Vdash^p M : \tau \quad \Psi_2 ; \Gamma ; \Delta, (x_m : A) \vdash^g Q :: (z_k : C)}{\Psi ; \Gamma ; \Delta, (x_m : \tau \rightarrow A) \vdash^r \text{send } x_m M ; Q :: (z_k : C)} \rightarrow L$$

Figure 6.3: Typing rules corresponding to the functional layer.

$$(\rightarrow S) : \text{proc}(d_k, w, \text{send } c_m M ; P) \mapsto \text{msg}(c_m^+, 0, \text{send } c_m M ; c_m^+ \leftrightarrow c_m),$$

$$\text{proc}(d_k, w, [c_m^+/c_m]P)$$

The recipient process substitutes M for x , and continues to offer along the fresh continuation channel received by the message. This ensures that messages are received in the order they are sent. The rule is formalized below.

$$(\rightarrow C) : \text{proc}(c_m, w', x \leftarrow \text{recv } c_m ; Q), \text{msg}(c_m^+, w, \text{send } c_m M ; c_m^+ \leftrightarrow c_m) \mapsto$$

$$\text{proc}(c_m^+, w + w', [c_m^+/c_m][M/x]Q)$$

Tracking Linear Assets As an illustration, consider the type money introduced in the auction example (Section 6.2). The type is an abstraction over funds stored in a process and is described as

```
money =
  &{value : int × money,           % send value
    add : money →R money,         % receive money and add it
    subtract : int → ⊕{sufficient : money ⊗R money, % receive int, send money
                      insufficient : money}      % funds insufficient to subtract
    coins : listcoin}
```

The type supports querying for value, and addition and subtraction. The type also expresses insufficiency of funds in the case of subtraction. The provider process only supplies money to the client if the requested amount is less than the current balance, as depicted in the **sufficient** label. The type is implemented by a *wallet* process that internally stores a linear list of coins and an integer representing its value. Since linearity is only enforced on the list of coins in the linear context, we trust the programmer updates the integer in the functional context correctly during transactions. The process is typed and implemented as (modes of channels l and m is R , skipped in the definition for brevity)

$$1: (n : \text{int}) ; (l_R : \text{list}_{\text{coin}}) \vdash \text{wallet} :: (m_R : \text{money})$$

```

2:   $m \leftarrow \text{wallet } n \ l =$ 
3:      case  $m$                                      % case analyze on label received on  $m$ 
4:          (value  $\Rightarrow$  send  $m \ n$  ;           % receive value, send  $n$ 
5:               $m \leftarrow \text{wallet } n \ l$ 
6:          | add  $\Rightarrow m' \leftarrow \text{rcv } m$  ;     % receive  $m'$  : money to add
7:               $m'.\text{value}$  ;                         % query value of  $m'$ 
8:               $v \leftarrow \text{rcv } m'$  ;
9:               $m'.\text{coins}$  ;                          % extract list of coins stored in  $m'$ 
10:              $k \leftarrow \text{append } l \ m'$  ;        % append list received to internal list
11:              $m \leftarrow \text{wallet } (n + v) \ k$ 
12:          | subtract  $\Rightarrow n' \leftarrow \text{rcv } m$  ; % receive int to subtract
13:             if ( $n' > n$ ) then
14:                  $m.\text{insufficient}$  ;                % funds insufficient
15:                  $m \leftarrow \text{wallet } n \ l$ 
16:             else
17:                  $m.\text{sufficient}$  ;                % funds sufficient
18:                  $l' \leftarrow \text{remove } n' \ l$  ;    % remove  $n'$  coins from  $l$ 
19:                  $k \leftarrow \text{rcv } l'$  ;          % and create its own list
20:                  $m' \leftarrow \text{wallet } n' \ k$  ;    % new wallet process for subtracted funds
21:                 send  $m \ m'$  ;                    % send new money channel to client
22:                  $m \leftarrow \text{wallet } (n - n') \ l'$ 
23:          | coins  $\Rightarrow m \leftrightarrow l$ )

```

If the *wallet* process receives the message value, it sends back the integer n , and recurses (lines 4 and 5). If it receives the message add followed by a channel of type money (line 6), it queries the value of the received money m' (line 7), stores it in v (line 8), extracts the coins stored in m' (line 9), and appends them to its internal list of coins (line 10). Similarly, if the *wallet* process receives the message subtract followed by an integer, it compares the requested amount against the stored funds. If the balance is insufficient, it sends the corresponding label, and recurses (lines 14 and 15). Otherwise, it removes n' coins using the *remove* process (line 18), creates a money abstraction using the *wallet* process (line 20), sends it (line 21) and recurses. Finally, if the *wallet* receives the message coins, it simply forwards its internal list along the offered channel.

6.5 Tracking Resource Usage

Resource usage is particularly important in digital contracts: Since multiple parties need to agree on the result of the execution of a contract, the computation is potentially performed multiple times or by a trusted third party. This immediately introduces the need to prevent denial of service attacks and to distribute the cost of the computation among the participating parties.

The predominant approach for smart contracts on blockchains like Ethereum is not to restrict the computation model but to introduce a cost model that defines the *gas* consumption of low level operations. Any transaction with a smart contract needs to be executed and validated

before adding it to the global distributed ledger, i.e., blockchain. This validation is performed by *miners*, who charge fees based on the gas consumption of the transaction. This fee has to be estimated and provided by the sender prior to the transaction. If the provided amount does not cover the gas cost, the money falls to the miner, the transaction fails, and the state of the contract is reverted back. Overestimates bear the risk of high losses if the contract has flaws or vulnerabilities.

It is not trivial to decide on the right amount for the fee since the gas cost of the contract does not only depend on the requested transaction but also on the (a priori unknown) state of the blockchain. Thus, precise and static estimation of gas cost facilitates transactions and reduces risks. We discuss our approach of tracking resource usage, both at the functional and process layer. The former is handled using an adoption of automatic amortized resource analysis (AARA) while the latter is handled using potential annotations introduced in Chapter 4.

Functional Layer Numerous techniques have been proposed to statically derive resource bounds for functional programs [23, 46, 55, 105, 138]. In Nomos, we adapt the work on AARA [91, 93] that has been implemented in Resource Aware ML (RaML) [92]. RaML can automatically derive worst-case resource bounds for higher-order polymorphic programs with user-defined inductive types. The derived bounds are multivariate resource polynomials of the size parameters of the arguments. AARA is parametric in the resource metric and can deal with non-monotone resources like memory that can become available during the evaluation.

As an illustration, consider the function *applyInterest* that iterates over a list of balances and applies interest on each element, multiplying them by a constant c . We use *tick* annotations to define the resource usage of an expression in this article. We have annotated the code to count the number of multiplications. The resource usage of an evaluation of *applyInterest* b is $|b|$.

```
let applyInterest balances =
  match balances with
  | [] -> []
  | hd::tl -> tick(1); (c*hd)::(applyInterest tl)
      (* consume unit potential for tick *)
```

The idea of AARA is to decorate base types with potential annotations that define a potential function as in amortized analysis. The typing rules ensure that the potential before evaluating an expression is sufficient to cover the cost of the evaluation and the potential defined by the return type. This posterior potential can then be used to pay for resource usage in the continuation of the program. For example, we can derive the following resource-annotated type.

$$\mathit{applyInterest} : L^1(\text{int}) \xrightarrow{0/0} L^0(\text{int})$$

The type $L^1(\text{int})$ denotes a list of integers assigning a unit potential to each element in the list. The return value, on the other hand, has no potential. The annotation on the function arrow indicates that we do not need any potential to call the function and that no constant potential is left after the function call has returned.

In a larger program, we might want to call the function *applyInterest* again on the result of a call to the function. In this case, we would need to assign the type $L^1(\text{int})$ to the resulting list and require $L^2(\text{int})$ for the argument. In general, the type for the function can be described with symbolic annotations with linear constraints between them. To derive a worst-case bound

for a function the constraints can be solved by an off-the-shelf LP solver, even if the potential functions are polynomial [91, 92].

In Nomos, we simply adopt the standard typing judgment of AARA for functional programs.

$$\Psi \Vdash^q M : \tau$$

It states that under the resource-annotated functional context Ψ , with constant potential q , the expression M has the resource-aware type τ .

The operational *cost* semantics is defined by the judgment

$$M \Downarrow V \mid \mu$$

which states that the closed expression M evaluates to the value V with cost μ . The type soundness theorem states that if $\cdot \Vdash^q M : \tau$ and $M \Downarrow V \mid \mu$ then $q \geq \mu$.

More details about AARA can be found in the literature [92, 93] and the Nomos supplementary material.

Process Layer To bound the resource usage of a process, Nomos features previously introduced resource-aware session types (Chapter 4) for work analysis. Resource-aware session types describe resource contracts for inter-process communication. The type system supports amortized analysis by assigning potential to both messages and processes. The derived resource bounds are functions of interactions between processes. As an illustration, consider the following resource-aware list interface from prior work [61].

$$\text{list}_A = \oplus\{\text{nil}^0 : \mathbf{1}^0, \text{cons}^1 : A \otimes^0 \text{list}_A\}$$

The type prescribes that the provider of a list must send one unit of potential with every `cons` message that it sends. Dually, a client of this list will receive a unit potential with every `cons` message. All other type constructors are marked with potential 0, and exchanging the corresponding messages does not lead to transfer of potential.

While resource-aware session types in Nomos are equivalent to the existing formulation [61], our version is simpler and more streamlined. Instead of requiring every message to carry a potential (and potentially tagging several messages with 0 potential), we introduce two new type constructors for exchanging potential.

$$A ::= \dots \mid \triangleright^r A \mid \triangleleft^r A$$

The type $\triangleright^r A$ requires the provider to pay r units of potential which are transferred to the client. Dually, the type $\triangleleft^r A$ requires the client to pay r units of potential that are received by the provider. Thus, the reformulated list type becomes

$$\text{list}_A = \oplus\{\text{nil} : \mathbf{1}, \text{cons} : \triangleright^1(A \otimes \text{list}_A)\}$$

The reformulation is more compact since we need to account for potential in only the typing rules corresponding to $\triangleright^r A$ and $\triangleleft^r A$.

With all aspects introduced, the process typing judgment

$$\Psi ; \Gamma ; \Delta \Vdash^q P :: (x_m : A)$$

denotes a process P accessing functional variables in Ψ , shared channels in Γ , linear channels in Δ , offers service of type A along channel x at mode m and stores a non-negative constant potential q . Similarly, the expressing typing judgment

$$\Psi \Vdash^p M : \tau$$

$\boxed{\Psi ; \Gamma ; \Delta \vdash^q P :: (x_m : A)}$ Process P has potential q .

$$\frac{p = q + r \quad \Psi ; \Gamma ; \Delta \vdash^p P :: (x_m : A)}{\Psi ; \Gamma ; \Delta \vdash^q \text{get } x_m \{r\} ; P :: (x_m : \triangleleft^r A)} \triangleleft R$$

$$\frac{q = p + r \quad \Psi ; \Gamma ; \Delta, (x_m : A) \vdash^p P :: (z_k : C)}{\Psi ; \Gamma ; \Delta, (x_m : \triangleleft^r A) \vdash^q \text{pay } x_m \{r\} ; P :: (z_k : C)} \triangleleft L$$

$$\frac{q = p + r \quad \Psi ; \Gamma ; \Delta \vdash^p P :: (x_m : A)}{\Psi ; \Gamma ; \Delta \vdash^q \text{tick}(r) ; P :: (x_m : A)} \text{tick}$$

Figure 6.4: Selected typing rules corresponding to potential.

denotes that expression M has type τ in the presence of functional context Ψ and potential p .

Figure 6.4 shows the rules that interact with the potential annotations. In the rule $\triangleleft R$, process P storing potential q receives r units along the offered channel $x_m : \triangleleft^r A$ using the *get* construct and the continuation executes with $p = q + r$ units of potential. In the dual rule $\triangleleft L$, a process storing potential $q = p + r$ sends r units along the channel $x_m : \triangleleft^r A$ in Δ using the *pay* construct, and the continuation remains with p units of potential. The typing rules for the dual constructor $\triangleright^r A$ are the exact inverse. Finally, executing the *tick*(r) construct consumes r potential from the stored process potential q , and the continuation remains with $p = q - r$ units, as described in the *tick* rule.

The *tick* construct is used to simulate a cost model in Nomos. If an operation (e.g., sending a message, calling a function, etc.) has a cost of r , this cost is simulated by inserting *tick*(r) just before the operation. Then, the *tick* operations are the only ones that cost potential, thus simplifying the type system. These *tick* operations are automatically inserted by the Nomos type checker, using a predefined cost model that assigns a constant cost to each operation. In addition, our implementation provides some standard cost models, for instance, that assign a unit cost to each internal operation and sending a message.

Integration Since both AARA for functional programs and resource-aware session types are based on the integration of the potential method into their type systems, their combination is natural. The two points of integration of the functional and process layer are (i) spawning a process, and (ii) sending/receiving a value from the functional layer. Recall the spawn rule $\{\} E_{RR}$ from Figure 6.3. A process storing potential $r = p + q$ can spawn a process corresponding to the monadic expression M , if M needs p units of potential to evaluate, while the continuation needs q units of potential to execute. Moreover, the functional context Ψ is shared in the two premises as Ψ_1 and Ψ_2 using the judgment $\Psi \curlywedge (\Psi_1, \Psi_2)$. This judgment, already explored in prior work [92] describes that the base types in Ψ are copied to both Ψ_1 and Ψ_2 , but the potential is split up. For instance, $L^{q_1+q_2}(\tau) \curlywedge (L^{q_1}(\tau), L^{q_2}(\tau))$. The rule $\rightarrow L$ in Figure 6.3 follows a similar pattern. A process Q storing $r = p + q$ potential sends a monadic expression M needing p units of potential to evaluate, and the continuation remains with q units of potential to execute. The p units of potential are consumed to evaluate M to a value before sending

since only values are exchanged at runtime. Thus, the combination of the two type systems is smooth, assigning a uniform meaning to potential, both for the functional and process layer. Remarkably, this technical device of exchanging functional values can be used to exchange non-constant potential with messages. For instance, exchanging a list $l : L^q(\tau)$ will exchange $q \cdot n$ units of potential where n is the size of the list l .

Operational Cost Semantics The resource usage of a process (or message) is tracked in semantic objects $\text{proc}(c, w, P)$ and $\text{msg}(c, w, N)$ using the local counters w . This signifies that the process P (or message N) has performed *work* w so far. The rules of semantics that explicitly affect the work counter are

$$\frac{M \Downarrow V \mid \mu}{\text{proc}(c_m, w, P[M]) \mapsto \text{proc}(c_m, w + \mu, P[V])} \text{ internal}$$

This rule describes that if an expression M evaluates to V with cost μ , then the process $P[M]$ depending on monadic expression M steps to $P[V]$, while the work counter increments by μ , denoting the total number of internal steps taken by the process. At the process layer, the work increments on executing a *tick* operation.

$$\text{proc}(c_m, w, \text{tick}(\mu) ; P) \mapsto \text{proc}(c_m, w + \mu, P)$$

A new process (or message) is spawned with $w = 0$, and a terminating process transfers its work to the corresponding message it interacts with before termination, thus preserving the total work performed by the system.

6.6 Type Soundness

The main theorems that exhibit the connections between our type system and the operational cost semantics are the usual *type preservation* and *progress*. First, Definition 7 asserts certain invariants on process typing judgment depending on the mode of the channel offered by a process. This mode, remains invariant, as the process evolves. This is ensured by the process typing rules, which remarkably preserve these invariants despite being parametric in the mode.

Lemma 5 (Invariants). *The typing rules on the judgment $\Psi ; \Gamma ; \Delta \vdash^g (x_m : A)$ preserve the invariants outlined in Definition 7, i.e., if the conclusion satisfies the invariant, so do all the premises.*

Configuration Typing At run-time, a program evolves into a number of processes and messages, represented by proc and msg predicates. This multiset of predicates is referred to as a *configuration* (abbreviated as Ω).

$$\Omega ::= \cdot \mid \Omega, \text{proc}(c, w, P) \mid \Omega, \text{msg}(c, w, N)$$

A key question is how to type these configurations because a configuration both uses and provides a number of channels. The solution is to have the typing impose a partial order among the processes and messages, requiring the provider of a channel to appear before its client. We stipulate that no two distinct processes or messages in a well-formed configuration provide the same channel c .

The typing judgment for configurations has the form $\Sigma ; \Gamma_0 \stackrel{E}{\Vdash} \Omega :: (\Gamma ; \Delta)$ defining a configuration Ω providing shared channels in Γ and linear channels in Δ . Additionally, we need to track the mapping between the shared channels and their linear counterparts offered by a contract process, switching back and forth between them when the channel is acquired or released respectively. This mapping, along with the type of the shared channels, is stored in Γ_0 . E is a natural number and stores the sum of the total potential and work as recorded in each process and message. We call E the energy of the configuration. The supplement details the configuration typing rules.

Finally, Σ denotes a signature storing the type and function definitions. A signature is well-formed if (i) every type definition $V = A_V$ is *contractive* [73] and (ii) every function definition $f = M : \tau$ is well-typed according to the expression typing judgment $\Sigma ; \cdot \Vdash^p M : \tau$. The signature does not contain process definitions; every process is encapsulated inside a function using the contextual monad.

Theorem 10 (Type Preservation).

- If a closed well-typed expression $\cdot \Vdash^q M : \tau$ evaluates to a value, i.e., $M \Downarrow V \mid \mu$, then $q \geq \mu$ and $\cdot \Vdash^{q-\mu} V : \tau$.
- Consider a closed well-formed and well-typed configuration Ω such that $\Sigma ; \Gamma_0 \stackrel{E}{\Vdash} \Omega :: (\Gamma ; \Delta)$. If the configuration takes a step, i.e. $\Omega \mapsto \Omega'$, then there exist Γ'_0, Γ' such that $\Sigma ; \Gamma'_0 \stackrel{E}{\Vdash} \Omega' :: (\Gamma' ; \Delta)$, i.e., the resulting configuration is well-typed. Additionally, $\Gamma_0 \subseteq \Gamma'_0$ and $\Gamma \subseteq \Gamma'$.

The preservation theorem is standard for expressions [92]. For processes, we proceed by induction on the operational cost semantics and inversion on the configuration and process typing judgment.

To state progress, we need the notion of a *poised* process [129]. A process $\text{proc}(c_m, w, P)$ is poised if it is trying to receive a message on c_m . Dually, a message $\text{msg}(c_m, w, N)$ is poised if it is sending along c_m . A configuration is poised if every message or process in the configuration is poised. Intuitively, this means that the configuration is trying to interact with the outside world along a channel in Γ or Δ . Additionally, a process can be *blocked* [25] if it is trying to acquire a contract process that has already been acquired by some process. This can lead to the possibility of deadlocks.

Theorem 11 (Progress). Consider a closed well-formed and well-typed configuration Ω such that $\Gamma_0 \stackrel{E}{\Vdash} \Omega :: (\Gamma ; \Delta)$. Either Ω is poised, or it can take a step, i.e., $\Omega \mapsto \Omega'$, or some process in Ω is blocked along a_S for some shared channel a_S and there is a process $\text{proc}(a_L, w, P) \in \Omega$.

The progress theorem is weaker than that for binary linear session types, where progress guarantees deadlock freedom due to absence of shared channels.

Chapter 7

Implementation of Nomos

Chapter 6 described the theory of Nomos, along with its static and dynamic semantics and type safety theorems. This chapter complements the previous chapter by describing the implementation of Nomos along with specific features making it easier to write smart contracts.

7.1 Overview of Nomos with an Auction Implementation

We highlight the main features of the Nomos language using the implementation of an auction contract in concrete syntax. An auction operates in two phases: a *running* phase where bidders bid into the auction followed by an *ended* phase where bidders collect their earnings. If a bidder wins the auction, they receive the *lot*, otherwise they receive their bids back.

The first key idea behind Nomos is to express and enforce the contract protocols like the auction via *session types*. The auction session type is defined as

```
type auction =
  /\ <{20} +{running : money -o |{3}> \/ auction ,
      ended : +{won : lot * |{5}> \/ auction ,
              lost : money * |{0}> \/ auction}}
```

We first ignore the operators $\langle\{q\}|$ and $|{q}\rangle$ for natural numbers q (described later) and describe the remaining type.

The type initiates with $/\$ (\uparrow_L^S in abstract syntax) indicating that auction is a *shared* session type [25] that must be acquired by a bidder to interact with the contract. Shared session types guarantee that bidders interact with the auction in *mutual exclusion* and their interaction with the auction executes atomically. Once the action contract is acquired, it replies either with *running* or *ended* indicating the phase of the auction. In the former case, the auction receives money using the $-o$ constructor ($-o$ in abstract syntax) followed by the bidder releasing the contract matching the $\backslash/$ (\downarrow_L^S in abstract syntax, dual to \uparrow_L^S) constructor in the type. In the latter case, the auction determines if the bidder won or lost the election. If the bidder wins, the auction sends the *won* label followed by sending the *lot* using the $*$ constructor (\otimes in abstract syntax, dual to $-o$). If the bidder loses, the auction sends the *lost* label followed by returning the bidder's money back. In either case, the type recurses back to auction after a $\backslash/$ indicating that the bidder must release the auction.

The second key feature of Nomos is that it statically enforces that assets are never duplicated nor discarded, but only transferred between processes. Nomos’ type system relies on session types [42] that are rooted in linear logic [76]. This linear type system tracks the assets stored in a process. For instance, the auction contract treats money and lot as linear assets, which is witnessed by the use of the linear logic operators \multimap and \otimes for their exchange.

Finally, an important aspect of smart contracts is their *execution cost*. Blockchains such as Ethereum [161] charge users a fee proportional to the execution (aka gas) cost of their transaction. A unique feature of Nomos is that it uses resource-aware session types [61] to statically analyze the execution cost of a transaction. They operate by assigning an initial *potential* to each process. This potential is consumed by each operation that the process executes or can be transferred between processes to share and amortize cost. The cost of each operation is defined by a cost model.

Resource-aware session types express the potential as part of the session type using the operators $\langle\{q\}|$ and $|\{q\}\rangle$ (\triangleleft^q and \triangleright^q in abstract syntax). The $\langle\{q\}|$ operator prescribes that the client must send q potential to the contract, with the amount of potential indicated as a superscript. Dually, $|\{q\}\rangle$ prescribes that the contract must send q potential to the client. In case of the auction contract, we require the client to pay potential for the operations that the contract must execute, both while placing and collecting their bids. If the cost model assigns a cost of 1 to each contract operation, then the maximum cost of an auction session is 20 (taking the maximum execution cost all branches). Thus, we require the client to send 20 units of potential at the start of a session using $\langle\{20\}|$. In the `won` branch of the auction type, on the other hand, the contract returns 5 units of potential to the client using $|\{q\}\rangle$. This mirrors gas usage in smart contracts, where the sender initiates a transaction with some initial gas, and the leftover gas at the end of the transaction is returned to the sender. All the above potential annotations have been automatically inferred by the Nomos type checker that internally relies on an LP solver to compute gas bounds.

Process Implementations As a final set of illustrations, we describe the main parts of the auction contract program. Since the auction operates in two phases, we have two main processes: `running_auction` for the running phase, and `ended_auction` for the ended phase. The type and definition of `running_auction` process is presented below.

```
contract running_auction :
($bm : Map<address, money>), ($l : lot) |- (#a : auction) =
  $la <- accept #a ;           % accept acquire request
  get $la {20} ;              % get 20 potential units
  $la.running ;               % send ‘running’ label
  $m <- recv $la ;            % receive money from bidder
  pay $la {3} ;               % pay leftover potential
  #a <- detach $la ;          % detach from bidder
  let addr = Nomos.GetTxnSender() ; % get bidder’s address
  $bm.insert(addr, $m) ;      % insert bid into bidmap
  #a <- run_or_end $bm $l     % call ‘run_or_end’ process
```

The process uses two linear channels: `$bm` represents the mapping from address to money, and `$l` represents the lot. On the other hand, it offers the shared channel `#a` that connects the

auction contract to the bidder process. To syntactically distinguish session-typed channels from functional variables, Nomos prefixes linear channels with $\$$ (e.g. $\$m$) and shared channels with $\#$ (e.g. $\#a$). The process closely follows the session-typed protocol described by the auction. It first accepts the acquire request from the bidder followed by receiving 20 potential units. Since this process represents the running phase of auction, it sends the **running** label, receives the money from the bidder in $\$m$, pays the leftover 3 potential units, and detaches from the bidder. Internally, the process then computes the sender’s address using the built-in `GetTxnSender` function, and inserts key-value pair $(\text{addr}, \$m)$ into the $\$m$ map. Maps have a built-in session type and can be used as such, but we simplify programming by introducing syntactic sugar construct `\$m.insert(addr, \$m)`. Finally, the process calls the `run_or_end` process which decides whether to call `running_auction` or `ended_auction` (maybe if the number of bidders reaches a certain threshold).

The bidder process lies on the other end of the auction channel and is responsible for bidding in the auction. We outline the process definition.

```
transaction bid_proc : ( $\$m$  : money), ( $\#a$  : auction) |- ( $\$d$  : 1) =
   $\$la$  <- acquire  $\#a$  ;                               % acquire auction contract
  pay  $\$la$  {20} ;                                       % pay 20 potential units
  case  $\$la$  (                                           % branch on label received
    running => send  $\$la$   $\$m$  ;                          % send money to contract
              get  $\$la$  {3} ;                             % get leftover potential
               $\#a$  <- release  $\$la$  ;                     % release the auction contract
              close  $\$d$                                  % terminate the transaction
  | ended => abort )                                    % abort if auction has ended
```

The process uses linear channel $\$m$ representing the bid, and the shared channel $\#a$ that connects to the auction contract. It offers on the channel $\$d$ of type 1. We mandate all transaction processes to offer type 1 for simplicity (more details in Section 7.3). The process initiates with acquiring the auction contract, pays 20 potential units, and case analyzes on the response. If the response is **running** (indicating that the auction is running), the process sends the money, gets the leftover potential, releases the contract, and terminates the transaction. If the response is **ended**, we simply abort the transaction. Note how the `running_auction` and `bid_proc` processes perform matching dual actions on the auction channel, as governed by the auction session type.

7.2 Mode Inference

Combining all these features in a single language is challenging. To achieve this integration without compromising type safety, Nomos introduces *channel modes*: R for linear asset channels; S for shared channels; L for shared channels in linear phase when acquired; and T for transaction channels. Nomos assigns a mode to *every* channel. For instance, in the `bid_proc` process, $\$m$ has mode R and $\#a$ has mode S. Once acquired, $\$la$ has mode L, and $\$d$ has mode T.

Practically however, annotating every channel with a mode can be a burden for the programmer. To address this, Nomos automatically infers the mode of every channel automatically.

Intuitively, we require the programmer to annotate the process with a *role*. We employ three roles: *asset* for linear assets, *contract* (e.g. `running_auction`), and *transaction* (e.g. `bid_proc`). Nomos then uses these process roles, written before the process name, to assign modes to all channels in that process.

First, based on the mode m of the channel offered by a process, Nomos asserts mode invariants on the shared and linear channels that the process uses. Definition 7 details those invariants.

Definition 8 (Process Typing). *Given judgment $\Psi ; \Delta \vdash^g P :: (x_m : A)$, and an arbitrary channel $y_k \in \Delta$,*

1. *If $m = R$, then $k = R$.*
2. *If $m = S$, then $k = R \vee k = S$.*
3. *If $m = L$, then $k = R \vee k = S \vee k = L$.*
4. *If $m = T$, then $k = R \vee k = S \vee k = L \vee k = T$.*

Intuitively, the above invariants impose a modal hierarchy $R < S < L < T$ and enforce that a process at mode m only uses channels at mode m' if $m' \leq m$. This hierarchy prevents cycles in the process dependency tree at runtime and is crucial to proving type safety [56].

Relevant to the tool implementation, the above invariants are central to inferring the channel modes automatically. First, Nomos uses the process roles to infer the mode m of the offered channel. We use the following rule:

- For role *asset*: $m = R$.
- For role *contract*: $m = S$.
- For role *contract*: $m = T$.
- (we do not allow defining processes at mode L)

Once we know the offered mode, we use the invariants from Definition 7 to generate constraints for mode k of each channel used by a process. In addition, we have three additional constraints.

- The process expression `$lc <- acquire #sc` imposes that the mode of linear channel `$lc` must be L, and the mode of shared channel `#sc` must be S.
- Dually, the process expression `#sc <- release $lc` imposes that mode of `$lc` must be L, and the mode of shared channel `#sc` must be S.
- Finally, if a channel c has a shared type, we conclude that the mode of c must be S.

The Nomos implementation generates and collects these constraints, and ships them to the LP solver. The solver, in turn, solves the constraints, computes the mode of each channel, and substitutes them back into the program. Section 7.3 provides more details on the LP solver.

7.3 Implementation

We have developed an open-source Nomos implementation [57] in OCaml (8469 lines of code). The implementation contains a lexer and parser (594 lines), a type checker (3486 lines), a pretty printer (531 lines), a cost inference engine with an LP solver interface (969 lines of code), and an interpreter (1942 lines). The program is first implemented in the concrete Nomos syntax and

then parsed into an abstract syntax tree. The program is then instrumented with work constructs to realize the cost model. Finally, the program is type checked to verify whether it implements the interface described by its session type. In this type checking phase, we also generate LP constraints on the potential annotations, and ship these constraints to the LP solver. The solver then minimizes the total potential while solving these constraints, computes a satisfying assignment, which is substituted back into the program, thus computing the gas cost of the program. Finally, the interpreter runs the program against the current blockchain state to obtain the new blockchain state. We follow a brief description of each aspect of the implementation.

Lexing and Parsing The Nomos lexer and parser have been implemented in Menhir [132], an LR(1) parser generator for OCaml. A Nomos program is a list of mutually recursive type and process definitions. The syntax for definitions is

```
type v = A
<role> f : (x1 : T), (#c2 : A2), ... |{q}- ($c : A) = M
```

The first line describes a type definition: A is the type expression that stands for the definition of type name v (e.g. `auction` type in Section 7.1). Since Nomos treats types equi-recursively, we can silently replace type name v with its definition A . The second line describes a process declaration and definition. We write the process role, followed by its name, its context and offered channel and type. The process role assigns a mode to the offered channel: *asset*, *contract* or *transaction*, assigning respective modes R , S and T to the offered channel. The modes for all other channels are inferred automatically (explained in Section 7.2). The context contains both functional variables and session-typed channel variables: $x1 : T$ defines a functional variable $x1$ of type T ; $\#c2 : A2$ defines a channel $\#c2$ with type $A2$; the process offers channel $\$c$ with type A . The expression M stands for the process definition. To visually separate out functional variables from session-typed channels, we require that shared channels are prefixed by $\#$, while linear channels are prefixed by $\$$. This avoids confusion between the two, both for the programmer and the parser. Finally, the potential $\{q\}$ of a process is marked on the turnstile in the declaration.

As a reference, Table 7.1 provides the abstract and concrete syntax of the session types and their corresponding process constructs in Nomos.

Cost Instrumentation Once a program has been parsed and converted into an abstract syntax tree, we instrument it with work constructs based on the cost model. The cost model intuitively defines the execution cost of each construct. The instrumentation engine takes the program and the cost model as input and produces a program with work constructs inserted at appropriate places. We use the following rule ($\llbracket P \rrbracket$ denotes the work-instrumented version of process P):

$$\llbracket S ; P \rrbracket ::= \text{work } \{C_S\} ; S ; \llbracket P \rrbracket$$

Here, S is a process construct with P as its continuation, and C_S denotes the cost of construct S according to the given cost model. We use this rule to add work annotations throughout the program, thus realizing the execution cost of the program. This instrumentation simplifies the cost analysis which can simply assign a cost of c to `work {c}` and cost 0 to all other process expressions.

Abstract Types	Concrete Types	Abstract Syntax	Concrete Syntax
$\oplus\{l : A, \dots\}$	$+ \{l : A, \dots\}$	$\$x.k$	$\$x.k$
$\&\{l : A, \dots\}$	$\&\{l : A, \dots\}$	$\text{case } \$x (\ell \Rightarrow P)_{\ell \in L}$	$\text{case } \$x (l \Rightarrow P \mid \dots)$
$A \otimes B$	$A * B$	$\text{send } \$x \w	$\text{send } \$x \w
$A \multimap B$	$A \multimap B$	$\$y \leftarrow \text{recv } \x	$\$y \leftarrow \text{recv } \x
$\mathbf{1}$	$\mathbf{1}$	$\text{close } \$x$	$\text{close } \$x$
$\uparrow_L^S A$	$\wedge A$	$\text{wait } \$x$	$\text{wait } \$x$
		$\$y \leftarrow \text{accept } x$	$\$y \leftarrow \text{accept } \#x$
		$\$y \leftarrow \text{acquire } x$	$\$y \leftarrow \text{acquire } \#x$
$\downarrow_L^S A$	$\vee A$	$\#y \leftarrow \text{detach } \x	$\#y \leftarrow \text{detach } \x
		$\#y \leftarrow \text{release } \x	$\#y \leftarrow \text{release } \x
$t \times A$	$t \wedge A$	$\text{send } \$x M$	$\text{send } \$x M$
$t \rightarrow A$	$t \multimap A$	$v \leftarrow \text{recv } \x	$v = \text{recv } \$x$
$\triangleright^r A$	$\mid\{r\}\rangle A$	$\text{pay } \$x \{r\}$	$\text{pay } \$x \{r\}$
$\triangleleft^r A$	$\langle\{r\}\mid A$	$\text{get } \$x \{r\}$	$\text{get } \$x \{r\}$

Table 7.1: Abstract and Corresponding Concrete Syntax for Nomos Types and Expressions

Type Checking The Nomos type checker is based on bi-directional type checking [130]. Intuitively, the programmer provides the initial type of each variable and channel in the declaration and the definition is checked against it, while reconstructing the intermediate types. This helps localize the source of a type error as the point where type reconstruction fails. As an illustration, recall the implementation of the running phase of the auction.

```

type auction =
  /\ <{20} +{running : money -o |{3}> \/ auction,
     ended : +{won : lot * |{5}> \/ auction,
     lost : money * |{0}> \/ auction}}

contract running_auction :
($bm : Map<address, money>), ($l : lot) |{0}- (#a : auction) =
  @% ($bm : Map<address, money>), ($l : lot)
  | - (#a : auction)@
  $la <- accept #a ;
  @% ($bm : Map<address, money>), ($l : lot)
  | - ($la : <{20}| +{running : ..., ended : ...})@
  get $la {20} ;
  @% ($bm : Map<address, money>), ($l : lot)
  | - ($la : +{running : ..., ended : ...})@
  $la.running ;
  @% ($bm : Map<address, money>), ($l : lot)
  | - ($la : money -o |{3}> \/ auction)@
  $m <- recv $la ;
  @% ($bm : Map<address, money>), ($l : lot), ($m : money)
  | - ($la : |{3}> \/ auction)@

```

```

pay $la {3} ;
  @% ($bm : Map<address, money>), ($l : lot), ($m : money)
    |- ($la : \ / auction)@
#a <- detach $la ;
  @% ($bm : Map<address, money>), ($l : lot), ($m : money)
    |- (#a : auction)@
let addr = Nomos.GetTxnSender() ;
  @% (addr : address), ($bm : Map<address, money>), ($l : lot)
    ($m : money) |- (#a : auction)@
$bm.insert(addr, $m) ;
  @% (addr : address), ($bm : Map<address, money>), ($l : lot)
    |- (#a : auction)@
#a <- run_or_end $bm $l

```

If the programmer forgets to write `$bm.insert(addr, $m)`, reconstruction would fail. There would be an extra `($m : money)` left in the process context, and since Nomos employs a linear type system, the type checker would report an error. In effect, the type checker *forces the programmer to consume* the channel `$m` to ensure linearity.

Type equality is restricted to reflexivity (constant time), although we have also implemented the standard co-inductive algorithm [73] which is quadratic in the size of type definitions. For all our examples, the reflexive notion of equality was sufficient. If type equality is restricted to constant-time reflexive notion, *type checking is linear time in the size of the program*. This property is quite relevant in the blockchain setting, since type checking can be part of the attack surface. If type checking is too slow, malicious users can issue transaction programs that take too long to type check, effectively forcing a denial-of-service attack.

LP Solver for Potential and Mode Inference The potential and mode annotations are the most interesting aspects of the Nomos type system. Since modes are associated with each channel, they are tedious to write. Similarly, the exact potential annotations depend on the cost assigned to each operation and is difficult to predict statically. Thus, we implemented an automatic inference algorithm for both these annotations by relying on an off-the-shelf LP solver.

Using ideas from existing techniques for type inference for AARA [92, 93], we reduce the reconstruction of potential annotations to linear optimization. To this end, Nomos’ inference engine uses the Coin-Or LP solver. In a Nomos program, the programmer can indicate unknown potential using `*`. Thus, resource-aware session types can be marked with \triangleright^* and \triangleleft^* , list types can be marked as $L^*(\tau)$ and process definitions can be marked with $\{|*\}$ – on the turnstile. The mode of all the channels is marked as ‘unknown’ while parsing.

As an example, consider the auction session type. The programmer writes the following type:

```

type auction =
  /\ <{*} +{running : money -o |{*}> \ / auction,
      ended : +{won : lot * |{*}> \ / auction,
              lost : money * |{*}> \ / auction}}

```


Using `*` annotations minimizes the programmer burden who does not need to compute exact potential annotations and execution cost.

The inference engine then iterates over the program and substitutes the star annotations with potential variables and ‘unknown’ with mode variables. For the auction type, this reduces to

```
type auction =
  /\ <{v0}> +{running : money -o |{v1}> \/ auction,
      ended : +{won : lot * |{v2}> \/ auction,
      lost : money * |{v3}> \/ auction}}
```

In the first phase of type checking, we ignore the potential and mode annotations and approximately type check the remaining program. This phase rules out type errors resulting from structural session types, i.e., protocol or linearity violation. In the second phase, we apply the rules for potential constructors (see Figures 6.3, 6.4) to generate linear constraints on the potential variables. For the above example, we generate the constraints

$$\begin{array}{lll}
 v_0 - v_1 \geq 17 & v_0 - v_2 \geq 15 & v_0 - v_3 \geq 20 \\
 v_0 \geq 0 & v_1 \geq 0 & v_2 \geq 0 & v_3 \geq 0 \\
 \min(v_0 + v_1 + v_2 + v_3)
 \end{array}$$

Finally, these constraints are shipped to the LP solver, which minimizes the value of the potential annotations to achieve tight bounds. The LP solver either returns that the constraints are infeasible, or returns a satisfying assignment, which is then substituted into the program. For the above example, we obtain the solution: $v_0 = 20, v_1 = 3, v_2 = 5, v_3 = 0$. The final program is pretty printed for the programmer to view and verify the potential and mode annotations.

7.4 Blockchain Integration

To integrate Nomos with a blockchain, we need a mechanism to (i) represent contracts and their addresses in the current blockchain state, (ii) create and execute transactions, and (iii) construct the global distributed ledger. This section addresses these challenges and also highlights the main limitation of the language: deadlocks in the transaction programs.

Nomos on a Blockchain We assume working in a blockchain setting similar to Ethereum with a standard account model. At any given time, the blockchain contains a set of Nomos contracts: C_1, \dots, C_n with their type information: $\cdot ; \Gamma^i ; \Delta_R^i \vdash^{q_i} C_i :: (x_S^i : A_S^i)$. The shared context Γ^i types the shared contracts that C_i refers to, and the linear context Δ_R^i types the contract’s linear assets. The channel name x_S^i of a contract is its address and has to be globally unique. Our implementation contains a deterministic mechanism to generate unique fresh names. We allow contracts to carry potential given by the annotation q_i , which can be used to share and amortize gas cost across transactions. It is also straightforward to modify the blockchain setup to suppress the potential stored in the contracts.

These contracts together form a stuck configuration (a valid virtual blockchain state) typed as

$$\Gamma \vDash^E \text{proc}(x_S^1, w_1, C_1) \dots \text{proc}(x_S^n, w_n, C_n) :: (\Gamma ; \cdot)$$

where $\Gamma = (x_S^1 : A_S^1), \dots, (x_S^n : A_S^n)$ and $E = \sum_{i=1}^n q_i + w_i$ is the total energy of the configuration, that is, the sum of the stored potential and previously performed work. The actual implementation stores some additional metadata such as mapping from linear channels to their shared counterpart, linear channels to their continuation, and a mapping from shared channels to their types. The mapping from linear channels to their counterpart is necessary for the $\downarrow_L^S C$ rule where a linear channel is released. To release a linear channel a_L to its corresponding shared channel a_S , we utilize this mapping that stores $a_L \mapsto a_S$. But since fresh channels are created with every communication, if a shared channel a_S is acquired to create a_L , then the linear channel to be released to a_S might be different from a_L , particularly if some communication occurred on a_L . The need for mapping shared channels to their types is explained in the paragraph on the Nomos interpreter.

To perform a transaction with a contract, a user submits a transaction script Q (a process) that is well-typed with respect to the existing contracts:

$$\cdot ; \Gamma ; \cdot \vdash^q Q :: (x_T : \mathbf{1})$$

We mandate that the transaction offers along a channel of type $\mathbf{1}$ and terminates by sending a close message on its offered channel. Intuitively, this enforces that the transaction, at termination, leaves the blockchain in a well-formed state. This transaction process is added to the set of contracts and the new (closed) configuration is typed as

$$\Gamma \vDash^{E+q} \text{proc}(x_S^1, w_1, C_1) \dots \text{proc}(x_S^n, w_n, C_n) \text{proc}(x_T, 0, Q) :: (\Gamma ; (x_T : \mathbf{1}))$$

This configuration then steps according to the Nomos semantics. A transaction can either create new contracts, or update the state of existing contracts. In the former case, new contracts are added to the blockchain state, making them visible in the type of the configuration for subsequent transactions to access. The type safety of Nomos ensures that transaction execution will be successful terminating in the following configuration

$$\Gamma \vDash^{E'} \text{proc}(x_S^1, w'_1, C'_1) \dots \text{proc}(x_S^m, w'_m, C'_m) \text{msg}(x_T, 0, \text{close } x_T) :: (\Gamma ; \cdot)$$

where m may be greater than n since the transaction can create additional contracts. At this point, we remove the close message from the configuration, resulting in the stuck configuration

$$\Gamma \vDash^{E'} \text{proc}(x_S^1, w'_1, C'_1) \dots \text{proc}(x_S^m, w'_m, C'_m) :: (\Gamma ; \cdot)$$

This stuck configuration represents a valid blockchain state guaranteeing that the transaction execution has successfully terminated, and we may initiate a new transaction on this new blockchain state.

Gas Accounts With this blockchain setup, one may reasonably wonder about the origin of potential. In our integration, we allow users to create *gas accounts* that store potential. To simplify matters, users can only store gas in these accounts; to store other private data and assets, they must create contracts where this additional information is collected. Gas accounts only contain a username and gas balance. We allow users to create new accounts and deposit gas into their accounts. When a user submits a transaction, the gas cost of the transaction is inferred by the LP solver, and this gas is automatically deducted from the account balance of that user. If the account balance is insufficient to pay for the transaction cost, the transaction is simply aborted.

Subsynchronizing Session Types A transaction program usually proceeds by *acquiring* existing contracts, exchanges data and assets with them, and subsequently *releasing* them. In Nomos, we require session types to be *equi-synchronizing* [25], i.e., contracts must be released at the same type that they are acquired. This requirement is crucial to type safety due to the shared nature of contracts. Since multiple users interact with a shared contract, it is critical that they observe a common shared type for the contract. If one client alters this type after interaction, other clients are not notified of this change, which will break type safety.

This equi-synchronizing constraint imposes a strong restriction: contracts must maintain a common type despite the phase they operate in. For instance, the auction contract operates over two phases: an *open* phase where bidders bid into the auction; and a *closed* phase when bidders withdraw their bids from the auction. However, these phases are not reflected in the type. Recall the auction type.

```
type auction =
  /\ <{20} +{running : money -o |{3}> \/ auction,
      ended : +{won : lot * |{5}> \/ auction,
              lost : money * |{0}> \/ auction}}
```

As we observe, the auction is acquired and released at the common type: auction.

In Nomos, we can relax this restriction by allowing *sub-synchronizing* types. Here, a contract can be released at a subtype of the type it was acquired at. For instance, we can introduce two mutually recursive auction types.

```
type open_auction =
  /\ <{20} +{running : money -o |{3}> \/ open_auction,
      ended : +{won : lot * |{5}> \/ closed_auction,
              lost : money * |{0}> \/ closed_auction}}

type closed_auction =
  /\ <{20} +{ended : +{won : lot * |{5}> \/ closed_auction,
                    lost : money * |{0}> \/ closed_auction}}
```

The `open_auction` type describes the open auction phase: it can either stay open (if it sends the running message), or transition to closed (if it sends the ended) message. As exemplified by the type, the phase of the auction is now visible in its type.

This relaxation, however, comes at a cost. Type checking in this relaxed setting needs to reason about subtyping, instead of type equality. Subtyping, unlike type equality, is more gen-

eral than reflexivity, hence cannot be decided in constant time. Thus, in the worst case, type checking would no longer be linear in this relaxed environment. Therefore, blockchains where efficiency is a particular concern should stick to reflexive equality.

Deterministic Execution Since blockchains rely on consensus among the miners, it is important to ensure deterministic execution of transactions. However, Nomos semantics has one source of non-determinism: the *acquire-accept* rule where an accepting contract latches on to any acquiring transaction. The Nomos implementation resolves this non-determinism by mandating that the contract must interact with the transaction with the lowest channel number. This simple heuristic, although not the cleanest, is easy to implement and sufficient for our purposes. Another promising approach is *record-and-replay* [111, 141]. The miner who mines the transaction records the order in which they resolved the acquire-accept non-determinism. All other miners validating the blockchain state must replay the same order, thus obtaining the same blockchain state after execution.

Interpreter The Nomos implementation provides two key functionalities: *inference* and *execution*. The inference engine takes a transaction program as input, infers the potential and mode annotations and outputs a well-typed program which is verified by the type checker. Next, the execution engine takes the well-typed transaction program and a valid blockchain state as input, executes the transaction against the state and outputs a valid blockchain state.

The Nomos interpreter uses OCaml S-expressions to represent blockchain states. The interpreter has read/write functionality which converts blockchain state to S-expression and vice-versa. This helps persist the blockchain state across transactions. The interpreter takes an input file, reads the S-expression from it, converts it to a blockchain state, executes the transaction against the state, and writes the output blockchain state to an output file. Internally, the interpreter is based on the semantics rules presented in Chapter 6.

Deadlocks The only language-specific reason a transaction can fail is a deadlock in the transaction code. Our progress theorem accounts for this possibility of deadlocks. We currently employ dynamic deadlock detection techniques internally in the implementation. Intuitively, since a valid blockchain state represents a stuck configuration of a particular form (only shared contracts in the configuration), we verify that the execution terminates with the configuration in this form. If not, we conclude that a deadlock occurred during the execution, and we simply abort the whole transaction. We maintain snapshots of the configuration after every transaction execution, so we simply revert to the previous valid blockchain state. It is the user’s responsibility to issue a new transaction that does not deadlock. In the future, we plan to employ deadlock prevention techniques [27] to statically rule out deadlocks.

7.4.1 Blockchain-specific Features

To further simplify programming, we enhance Nomos with blockchain-specific features. This section provides a brief overview of these features.

Map Data Structure The most widely used data structure in smart contracts is maps. It is often used in contracts to store a mapping from users to their balance. The auction example uses it to map users to their bids. We provide surface syntax to make it easier for programmers to interact with maps.

The first step is to distinguish between linear and non-linear maps. Although the two maps differ in their statics and semantics, we want to provide a unified syntax for ease of programming. We first describe the session type for a non-linear map with key type `kt` and value type `vt`.

```
type Map<kt, vt> = &{insert : kt -> vt -> Map<kt, vt>,
                    delete : kt -> vt option ^ Map<kt, vt>,
                    size : int ^ Map<kt, vt>,
                    close : 1}
```

The map is implemented with a recursive session type initiating with an external choice. It accepts one of three messages: `insert`, `delete`, `size` or `close`. In the case of `insert`, the map receives a key and value and inserts the pair in the dictionary. If the key already exists, the existing value is overwritten. In the case of `delete`, the map receives a key and returns an optional value depending on whether the key exists in the map or not. In the case of `size`, the map returns an integer corresponding to its size. In each of these three cases, the type then recurses back to `Map<kt, vt>`. Finally, in the case of `close`, the map simply terminates with a `close` message.

The case of a linear map is only slightly different. The type is as follows:

```
type Map<kt, vt> = &{insert : kt -> vt -o Map<kt, vt>,
                    delete : kt -> vt option * Map<kt, vt>,
                    size : int ^ Map<kt, vt>,
                    close : +{empty : 1,
                             nonempty: Map<kt, vt>}}
```

The type differs from a non-linear map in a few key ways. First, `vt` is a linear type, hence we use `-o` and `*` constructors to exchange them. Second, closing a linear map is only allowed when it is empty. Hence, on receiving a `close` message, a linear map will respond with either the empty message followed by termination, or with the `nonempty` message followed by recursing back to its original type.

Finally, we provide surface syntax to ease programming with maps.

- `$m <- new Map<kt, vt>`: for creating a new map `$m` of key type `kt` and value type `vt`
- `$m.insert(k, v)`: for inserting key `k` and value `v` into map `$m`. If the map is linear, the value is replaced by a channel `$v`.
- `v = $m.delete(k)`: for deleting key `k` from map `$m`. If the map is linear, the expression changes to `$v <- $m.delete(k)`.
- `n = $m.size`: for obtaining the size of map `$m` and storing it in variable `n`.
- `$m.close`: for closing the map.

Blockchain-specific Expressions To keep track of blockchain users, we introduce a built-in type called `address`. We use the expression `Nomos.GetTxnSender()` as an introduction form

that returns a value of type `address`. There are no elimination forms for this type.

We also introduce the expression `Nomos.GetTxnNum()` to track the current transaction number. Our implementation uses a sequential mode of execution: every new transaction is automatically assigned the number $n + 1$ where n is the number of the previous transaction. The first transaction is assigned the number 0.

Custom Coins Finally, Nomos provides a built-in abstract coin type. There are no introduction or elimination forms of this type. Thus, there are no ways to delete or duplicate a channel of this type; they can only be transferred across different contracts. This type can be used to truly enforce linearity of assets on the blockchain.

These coins can further be used to encode fancier tokens. For example, we can create a type `coin2` to represent 2 coins.

```
type coin2 = coin * coin
```

We can even store more information inside the coin, e.g. the owner's address or the history of all owners in the past. For instance, consider the `UTxO` coin.

```
type utxo = &{history : address list ^ utxo,
             delete : coin}
```

This type, when queried, can return the address list of all its owners. But we can also delete its history, and turn it into a simple coin. The Nomos language is thus, very general, and we are not tied to a particular system-specific coin.

7.4.2 Exact Gas Computation (Facebook Internship)

I wanted to make a special mention of the aspect of exact gas computation since this is a result of my summer internship at Facebook. So far, we have only discussed upper bound gas computation by the Nomos inference engine. Unfortunately, upper gas bounds are inadequate. At runtime, if a user provides excess gas units, the leftover gas needs to be returned to the user. Thus, in existing blockchains such as Ethereum and Diem, a monitor function known as *dynamic gas meter* tracks the gas cost during execution. If the execution runs out of gas, the meter raises an *out-of-gas exception*, otherwise it returns the excess gas back to the user. Thus, despite the benefits of static gas analysis, blockchains still need to meter gas at runtime. Moreover, dynamic gas metering has its own limitations. First, it creates an execution overhead, inadvertently increasing the transaction gas cost. For the Diem blockchain, this overhead is about 20% of execution time [32]. Second, if the transaction runs out of gas, it does not provide any feedback to the user for transaction resubmission.

Upper gas bounds can also be unfair to miners. Miners are usually paid in proportion to the gas cost of a transaction. As a result, they often accept transactions with a high gas limit, hoping that transactions with a high gas limit will have a high gas cost. However, a malicious entity can trick this system by submitting transactions with a high gas limit but a low gas cost. Miners would accept such transactions only to discover that their compensation would be low and most of the gas is returned back to the user. Thus, there is a need to *provide miners with a trusted exact gas bound that can be verified efficiently* before accepting transactions.

In response, we describe a technique to compute *exact* gas bounds so as to eliminate dynamic metering. The main challenge in computing exact bounds is *branching* since (i) gas cost may vary along different branches, and (ii) the branch taken during execution is difficult to predict statically. To compute exact bounds in the presence of branching, we need to ensure that branches have equal gas cost. We establish this by introducing a special operation `Nomos.deposit{n}` which deposits n gas units in the transaction sender's account at runtime. We augment the less costly branch with such an expression with n being the difference in the gas cost of both branches. We further illustrate that this mechanism is sufficient to produce exact gas bounds and eliminates the need for gas metering, improving the overall hygiene of the virtual machine.

To further reduce programmer overhead, we automatically augment both branches with the expression `Nomos.deposit{*}`. The program is then shipped to the LP solver which computes the value of all such $*$ annotations. Then, all the `Nomos.deposit{*}` expressions that are substituted with 0 value are simply deleted from the program. This completely eliminates the need for dynamic gas metering. At runtime, these expressions safely return the leftover gas in the sender's account. During my internship, I also created a tool called GasBoX [60] that employed similar ideas to compute exact bounds for Move programs.

7.5 Web Interface

To further enhance accessibility and usability, Nomos comes with a well-equipped web interface hosted at www.nomos-lang.org. The web interface has been carefully designed to simulate a blockchain execution with minimal user overhead. Programmers can utilize the web interface to type check transactions, automatically infer their gas usage, submit them and observe how they update the overall blockchain state.

Figure 7.1 provides a screenshot of the web interface in action. On the top left is the editor box where the transaction code can be written. As we can observe in line 23, the $*$ annotation communicates the user intention that the type checker should infer the gas cost. On the top right, we can create new user accounts and assign them an initial balance. We have already created an account in the name of Ankush with an account balance of 10,000 gas units. On the bottom left, we can view the blockchain state which contains a list of smart contracts currently deployed: their channel names, channel session type, state of the processes contained in that contract, and the work done and gas stored inside them. Finally, on the bottom right, we can view the transaction history which contains the list of all transactions ever executed along with their sender, code, and gas bound.

Figure 7.1 shows a transaction that creates an auction contract, and then places a bid in the auction. But, the programmer has made an error and accidentally sent a `collect` message (line 28) instead of `bid`. If the programmer submits this program, the type checker complains about the same as illustrated in Figure 7.2 in the popup in the bottom right. Once the programmer fixes that error, the type checker accepts the transaction program, computes the gas bound of 11 automatically (Figure 7.3). Once the transaction is posted, we observe that a new contract was created (bottom right of Figure 7.3) and programmers can view the type along with other contract information.

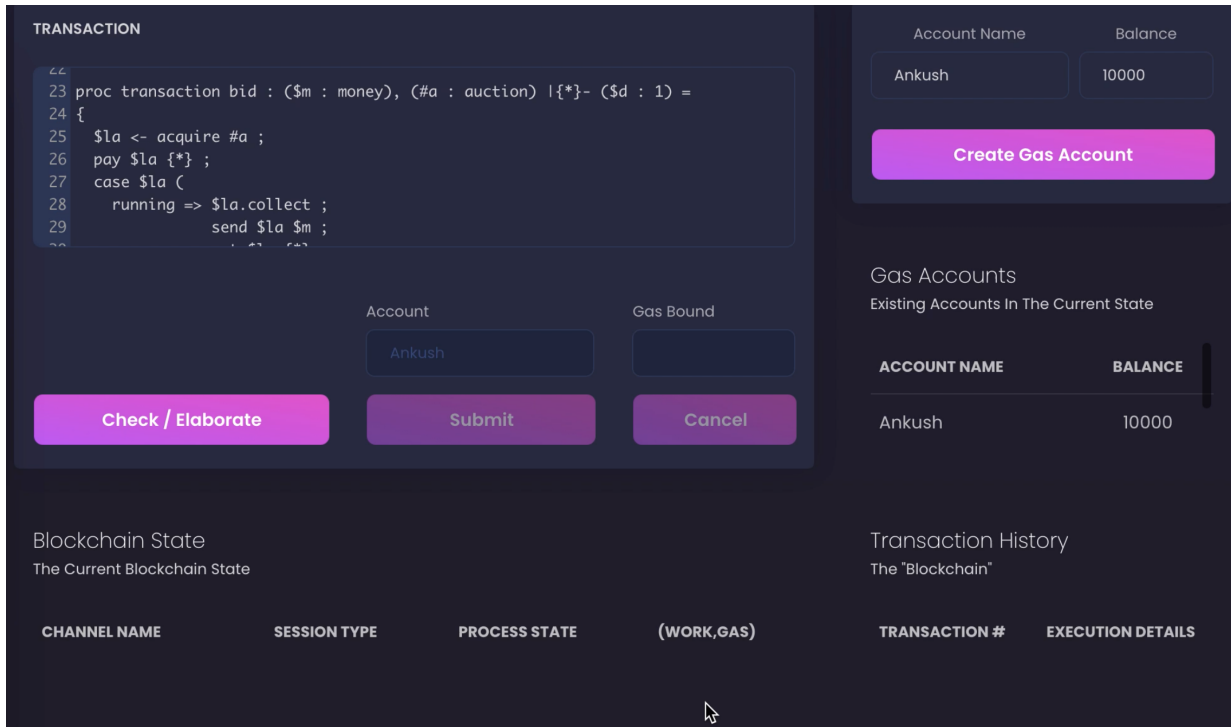


Figure 7.1: Screenshot of the Nomos Web Interface

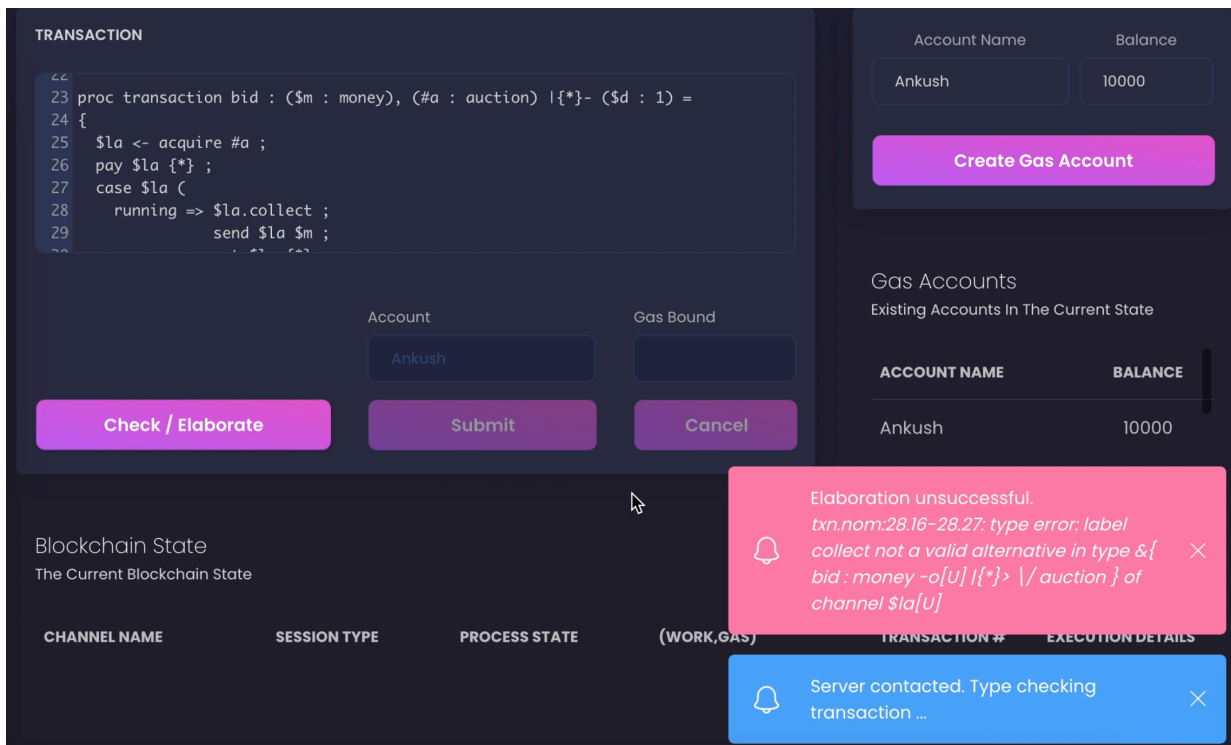


Figure 7.2: Error in Transaction Program

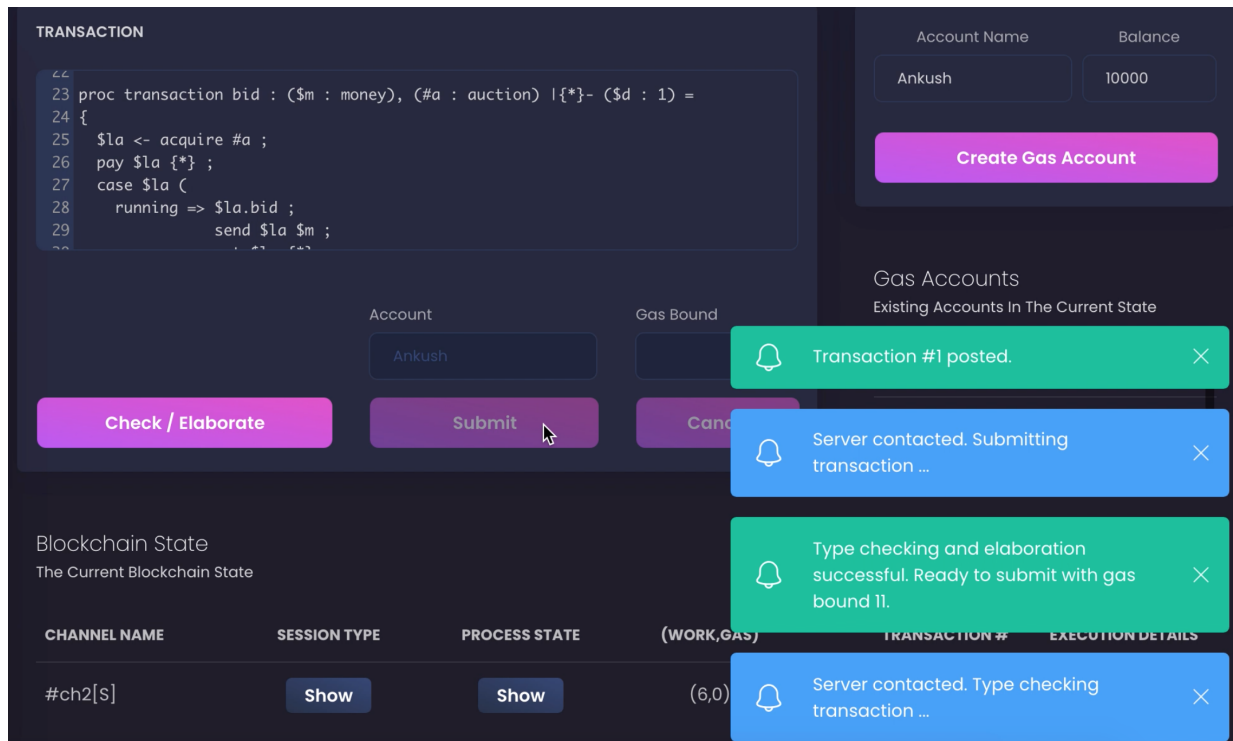


Figure 7.3: Correct Transaction Program

7.6 Evaluation

We evaluate the design of Nomos by implementing several smart contract applications and discussing the typical issues that arise. All the contracts are implemented and type checked in the prototype implementation and the potential and mode annotations are derived automatically by the inference engine. The cost model used for these examples assigns 1 unit of cost to every atomic internal computation and sending of a message. We show the contract types from the implementation with the following ASCII format: i) \wedge for \uparrow_L^S , ii) \vee for \downarrow_L^S , iii) $\langle\{q\}\rangle$ for \triangleleft^q , iv) $\mid\{q\}\rangle$ for \triangleright^q , v) \wedge for \times , vi) $*[m]$ for \otimes_m , vii) $-o[m]$ for \multimap_m .

ERC-20 Token Standard ERC-20 [68] is a technical standard for smart contracts on the Ethereum blockchain that defines a common list of standard functions that a token contract has to implement. The majority of tokens on the Ethereum blockchain are ERC-20 compliant.

The ERC-20 token contract implements the following session type in Nomos:

```

stype erc20token = /\ <{11}\ | &{
  totalSupply : int ^ |{9}\ > \ / erc20token ,
  balanceOf : id -> int ^ |{8}\ > \ / erc20token ,
  transfer : id -> id -> int -> |{0}\ > \ / erc20token ,
  approve : id -> id -> int -> |{6}\ > \ / erc20token ,
  allowance : id -> id -> int ^ |{6}\ > \ / erc20token }

```


The type ensures that the token implements the protocol underlying the ERC-20 standard. To query the total number of tokens in supply, a client sends the `totalSupply` label, and the contract sends back an integer. If the contract receives the `balanceOf` label followed by the owner's identifier, it sends back an integer corresponding to the owner's balance. A balance transfer can be initiated by sending the `transfer` label to the contract followed by sender's and receiver's identifier, and the amount to be transferred. If the contract receives `approve`, it receives the two identifiers and the value, and updates the allowance internally. Finally, this allowance can be checked by issuing the `allowance` label, and sending the owner's and spender's identifier.

The design of Nomos is orthogonal to the concrete representation of money or currency in the language. The Nomos implementation provides a simple built-in abstract coin type of a unit value. Our implementation of the `erc20token` session type relies on these abstract coins used exclusively for exchanges among the private accounts. Coins are treated linearly as no operations are allowed on primitive types. As a result, coins cannot be created or discarded.

It is straightforward to add features by using more sophisticated abstract coin types or by providing built-in operations that are executed by the runtime system. For example, we can add coins with unique identifiers or coins of different denominations by changing the underlying session type of coins. Similarly, we can add operations for minting (creating) or burning (discarding) coins if users have the respective privileges. Such operations could be, for instance, implemented in an abstract contract that is an interface to the runtime system. Finally, there can be operations for exchanging coins and gas at rates that are fixed when type-checking transactions.

It is also possible to allow programmers to define their own abstract types with their individual introduction and elimination forms to use them in an implementation of a session type like `erc20token`.

Hacker Gold (HKG) Token The HKG token is one particular implementation of the ERC-20 token specification. Recently, a vulnerability was discovered in the HKG token smart contract based on a typographical error leading to a re-issuance of the entire token [2]. When updating the receiver's balance during a transfer, instead of writing `balance+=value`, the programmer mistakenly wrote `balance+=value` (semantically meaning `balance=value`). Nomos' type system would have caught the linearity violation in the latter statement that drops the existing balance in the recipient's account.

Puzzle Contract This contract, taken from prior work [116] rewards users who solve a computational puzzle and submit the solution. The contract allows two functions, one that allows the owner to update the reward, and the other that allows a user to submit their solution and collect the reward.

In Nomos, this contract is implemented to offer the type

```
stype puzzle = /\ <{14}| &{
  update : id -> money -o[R] |{0}> \/ puzzle ,
  submit : int ^ &{
  success : int -> money *[R] |{5}> \/ puzzle ,
  failure : |{9}> \/ puzzle } }
```

The contract still supports the two transactions. To update the reward, it receives the update label and an identifier, verifies that the sender is the owner, receives money from the sender, and acts like a puzzle again. The transaction to submit a solution has a *guard* associated with it. First, the contract sends an integer corresponding to the reward amount, the user then verifies that the reward matches the expected reward (the guard condition). If this check succeeds, the user sends the success label, followed by the solution, receives the winnings, and the session terminates. If the guard fails, the user issues the failure label and immediately terminates the session. Thus, the contract implementation guarantees that the user submitting the solution receives their expected winnings.

Voting The voting contract provides a ballot type.

```
stype ballot = /\ <{16}| +{
  open : id -> +{ vote : id -> |{0}> \/ ballot ,
                novote : |{9}> \/ ballot },
  closed : id ^ |{13}> \/ ballot }
```

This contract allows voting when the election is **open** by sending the candidate's *id*, and prevents double voting by checking if the voter has already voted (the **novote** label). Once the election closes, the contract can be acquired to check the winner. We use two implementations for the contract: the first stores a counter for each candidate that is updated after each vote is cast (voting in Table 7.2); the second does not use a counter but stores potential inside the vote list that is consumed for counting the votes at the end (voting-aa in Table 7.2). This stored potential is provided by the voter to amortize the cost of counting. The type above shows the potential annotations corresponding to the latter.

Insurance Nomos has been carefully designed to allow inter-contract communication without compromising type safety. We illustrate this feature using an insurance contract that processes flight delay insurance claims after verifying them with a trusted third party. The insurer and third party verifier are implemented as separate contracts providing the following session types.

```
stype insurer = /\ <{6}| &{
  submit : claim -> +{
  success : money *[R] |{0}> \/ insurer ,
  failure : |{1}> \/ insurer } }

stype verifier = /\ <{3}| &{
  verify : claim -> +{
  valid : |{0}> \/ verifier ,
  invalid : |{0}> \/ verifier } }
```

The insurer type provides the option to **submit** a claim by receiving it and responds with **success** or **failure** depending upon verification of the claim. If the claim is successful, the insurer sends over the reimbursement in the form of money. The verifier type provides the option to **verify** a claim by receiving it and responding with **valid** or **invalid** depending on the validity of the claim.

Contract	LOC	T (ms)	Vars	Cons	I (ms)	Gap
auction	176	0.558	229	730	5.225	3
ERC 20	136	0.579	161	561	4.317	6
puzzle	108	0.410	126	389	8.994	8
voting	101	0.324	109	351	3.664	0
voting-aa	101	0.346	140	457	3.926	0
escrow	85	0.404	95	321	3.816	3
insurance	56	0.299	76	224	8.289	0
bank	147	0.663	173	561	4.549	0
wallet	30	0.231	32	102	3.224	0

Table 7.2: Evaluation of Nomos with Case Studies. LOC = lines of code; T (ms) = the type checking time in ms; Vars = #variables generated during type inference; Cons = #constraints generated during type inference; I (ms) = type inference time in ms; Gap = maximal gas bound gap.

The insurer, upon receiving a claim, acquires the verifier and sends it the claim details. If the claim is valid, then it responds with **success**, sends the money and detaches from its client. If the claim is invalid, it responds with **failure** and immediately detaches from its client.

Experimental Evaluation We describe the 8 case studies we implemented in Nomos. We have already discussed auction (Section 7.1), ERC 20, puzzle, voting, and insurance. The other case studies are:

- A bank account that allows users to register, make deposits and withdrawals and check the balance.
- An escrow to exchange bonds between two parties.
- A wallet allowing users to store money on the blockchain.

Table 7.2 contains a compilation of our experiments with the case studies and the prototype implementation. The experiments were run on an Intel Core i5 2.7 GHz processor with 16 GB 1867 MHz DDR3 memory. It presents the contract name, its lines of code (LOC), the type checking time (T (ms)), number of potential and mode variables introduced (Vars), number of potential and mode constraints that were generated while type checking (Cons) and the time the LP solver took to infer their values (I (ms)). The last column describes the maximal gap between the static gas bound inferred and the actual runtime gas cost. It accounts for the difference in the gas cost in different program paths. However, this waste is clearly marked in the program by explicit *tick* instructions so the programmer is aware of this runtime gap, based on the program path executed.

The evaluation shows that the type-checking overhead is less than a millisecond for case studies. This indicates that Nomos is applicable to settings like distributed blockchains in which type checking could add significant overhead and could be part of the attack surface. Type inference is also efficient but an order of magnitude slower than type checking. This is acceptable since inference is only performed once during deployment and can be carried out off-chain.

Gas bounds are tight in most cases. Loose gas bounds are caused by conditional branches with different gas cost. In practice, this is not a major concern since the Nomos semantics tracks the exact gas cost, and a user will not be overcharged for their transaction. However, Nomos' type system can be easily modified to only allow contracts with tight bounds.

Our implementation experience revealed that describing the session type of a contract crystallizes the important aspects of its protocol. Often, subtle aspects of a contract are revealed while defining the protocol as a session type. Once the type is defined, the implementation simply *follows* the type protocol. The error messages from the type checker were helpful in ensuring linearity of assets and adherence to the protocol. Using $*$ for potential annotations meant we could remain unaware of the exact gas cost of operations. The syntactic sugar constructs reduced the programming overhead and the size of the contract implementations.

7.7 Related Work

We classify the related work into 3 categories - i) new programming languages for smart contracts, ii) static analysis techniques for existing languages and bytecode, and iii) session-typed and type-based resource analysis systems technically related to Nomos.

Smart Contract Languages Existing smart contracts on Ethereum are predominantly implemented in Solidity [54], a statically typed object-oriented language influenced by Python and JavaScript. Languages like Vyper [8] address resource usage by disallowing recursion and infinite-length loops, thus making estimation of gas usage decidable. However, both languages still suffer from re-entrancy vulnerabilities. Bamboo [3], on the other hand, makes state transitions explicit and avoids re-entrance by design. In contrast to our work, none of these languages use linear type systems to track assets stored in a contract.

Domain-specific languages have been designed for other blockchains apart from Ethereum as well. Typecoin [53] uses affine logic to solve the peer-to-peer affine commitment problem using a generalization of Bitcoin where transactions deal in types rather than numbers. Although Typecoin does not provide a mechanism for expressing protocols, it also uses a linear type system to prevent resources from being discarded or duplicated. Rholang [6] is formally modeled by the ρ -calculus, a reflective higher-order extension of the π -calculus. Michelson [5] is a purely functional stack-based language that has no side effects. However, none of these languages describe and enforce communication protocols statically. Scilla [145] is an intermediate-level language where contracts are structured as communicating automata providing a continuation-passing style computational model to the language semantics. Scilla does not use session types or linearity but features static gas bounds. A difference is that Nomos' bounds are not asymptotic and are proved sound with respect to a cost semantics. Closer to Nomos, Obsidian [51] employs *typestate* to express state transition during transaction execution. It also utilizes linear types to track and preserve assets. The Move programming language [36] is a flexible language based on Rust [101] to implement contracts on the Libra blockchain. Similar to Nomos, it provides the ability to define custom linear types to represent assets. However, neither provides support to express or control gas usage.

Static Analysis Analysis of smart contracts has received substantial attention [83, 152] recently due to their security vulnerabilities [22, 155]. KEVM [88] creates a program verifier based on reachability logic that given an EVM program and specification, tries to automatically prove the corresponding reachability theorems. However, the verifier requires significant manual intervention, both in specification and proof construction. Oyente [116] is a symbolic execution tool that checks for 4 kinds of security bugs in smart contracts, transaction-order dependence, timestamp dependence, mishandled exceptions and re-entrancy vulnerabilities. MadMax [80] automatically detects gas-focused vulnerabilities with high confidence. The analysis is based on a decompiler that extracts control and data flow information from EVM bytecode, and a logic-based analysis specification that produces a high-level program model. Ethereum contracts are also translated to F^* [35] to prove runtime safety and functional correctness, although they do not support all syntactic features. VERISOL [107] is a highly-automated formal verifier for Solidity that can produce proofs as well as counterexamples and proves semantic conformance of smart contracts against a state machine model with access-control policy. However, in contrast to Nomos, where guarantees are proved by a soundness proof of the type system, static analysis techniques often do not explore all program paths, can report false positives that need to be manually filtered, and miss bugs due to timeouts and other sources of incompleteness.

Several static analysis tools also focus specifically on computing gas bounds. GASTAP [17] infers gas bounds on contracts implemented in Solidity [54] or EVM bytecode in terms of size of the input parameters, contract state and gas consumption. The inference procedure requires construction of control-flow graphs, decompilation to a high-level representation, inferring size relations, generating and solving gas equations. GASOL [16] is an extension to GASTAP which offers a variety of cost models to measure the cost of, for e.g., only storage opcodes, selected family of gas-consumption opcodes, selected program line, etc. It further detects under-optimized storage patterns and automatic optimization of such patterns. Marescotti et. al. [117] employ symbolic model checking to modularly enumerate all gas consumption paths based on unwinding loops to a limit. For each path, it then computes the environment state to force that path and simulates the transaction under the state to obtain an exact worst-case gas bound. Nomos differs from these tools in its goal of providing miners with a trusted exact gas bound which can be verified in linear time and eliminating dynamic gas metering.

Session types and Resource analysis Session types were introduced by Honda [96] as a typed formalism for inter-process dyadic interaction. They have been integrated into a functional language in prior work [153]. However, this integration does not account for resource usage or sharing. Sharing in session types has also been explored in prior work [25], but with the strong restriction that shared processes cannot rely on linear resources that we lift in Nomos. Shared session types were also never integrated with a functional layer or tracked for resource usage. While we consider binary session types that express local interactions, global protocols can be expressed using multi-party session types [97, 144]. Automatic amortized resource analysis (AARA) has been introduced as a type system to derive linear [93] and polynomial bounds [92] for functional programming languages. Resource usage has also previously been explored separately for the purely linear process layer [61], but were never combined with shared session types or integrated with the functional layer.

7.8 Future Directions

The strong safety guarantees of Nomos only hold in a *restricted attacker model* where both contracts and clients are well-typed. In practice, however, not all contracts and clients would be implemented in the Nomos language. This raises the concern of how contracts implemented in Nomos can safely interact with clients or contracts that are not implemented in Nomos. The contracts also contain no error raising and handling mechanisms. One approach here can be to introduce *runtime monitoring*. Monitors are wrapped around the contracts and observe the type of the data that is being exchanged on the channels. If they catch a type mismatch, they abort the transaction. We can also allow the programmer to define more sophisticated mechanisms of handling such type mismatches.

Another promising future direction is to explore refinement session types for lightweight verification of smart contracts. We also plan to explore open questions regarding integration with a blockchain system. These include the exact cost model, fluctuation of gas prices, and potential compilation to a lower-level language. Finally, we would also like to support parallel execution of transactions since Nomos naturally provides a concurrent semantics.

Chapter 8

Conclusion

This chapter concludes my thesis with a summary of contributions, a brief look at possible future directions, and some final thoughts.

8.1 Summary of Contributions

This thesis makes two major contributions: the design and theoretical foundation of resource-aware session types, and the application of resource-aware session types for a safe smart contract programming language called Nomos.

Chapter 3 begins with introducing novel arithmetically refined session types. The key innovation here was to index types with natural numbers that represent data structure sizes and values. We also allow session-typed processes to exchange linear arithmetic constraints on these indices to constrain process behavior. Our first, rather surprising, theoretical result was that type equality, and therefore, type checking is undecidable for refinement session types, even though Presburger arithmetic itself is decidable. But as a recourse, we devised an algorithm for approximating type equality which, despite its incompleteness, works very well in practice. Since refinement constructs introduce verbosity overhead to programs, we also devised a novel forcing algorithm to insert refinement constructs automatically. This assists with code reuse by greatly reducing programmer burden. Finally, the language including all the above algorithms are implemented in a type-safe system called Rast which is evaluated on standard session-typed benchmarks.

Chapter 4 crucially employs refinements to express work bounds for session-typed processes. To compute work, our key innovation was to introduce an abstract notion of potential. This potential can be stored inside processes and exchanged using special messages but, most importantly, it must be consumed to take an execution step. Intuitively, potential gets converted into work at runtime. This implies that the initial potential stands as an upper bound on the work bound of a concurrent computation. This is exactly what is formalized by the soundness theorem. As an application, we compare the efficiency of standard stack and queue implementations.

Chapter 5 employs refinements to express time bounds for programs. The key innovation here was to augment the timing information of message exchanges in the session type. In

particular, we introduce the \circ^r operator that stands for a delay of r units of time. For instance, a channel of type $\circ^r A$ behaves as A after exactly r time units. To capture unknown timings, we also introduce \square and \diamond operators which describe an eventual message exchange. We combine these three operators to flexibly express time bounds on standard concurrent data structures.

Session types, together with work and span extensions, are dubbed as *resource-aware session types*.

Chapter 6 is dedicated to the type-theoretic foundation of the Nomos language that provides 3 domain-specific features. First, the communication protocols that are statically prescribed by session types are instrumental in expressing and enforcing transaction behavior. Second, the linear characteristic of session types enables programmers to track assets preventing their accidental duplication or deletion. Finally, resource-aware session types automatically infer the execution cost of transactions relieving programmers of this burden too. In addition, the chapter also describes how session types are integrated into a functional language via a linear contextual monad. The chapter concludes with the usual progress and preservation theorems.

Chapter 7 complements the previous chapter by presenting the implementation of Nomos. The aim of this chapter is to demonstrate the practicality and wide applicability of Nomos. More concretely, the chapter describes our efforts in reducing programmer burden. First, verbose channel modes are automatically inferred using an LP solver. Second, the implementation provides several blockchain-specific features such as easy use of map data structures, transaction-related expressions, and gas accounts to send transactions. Third, the chapter also describes how Nomos integrates into an account model blockchain. Finally, Nomos is evaluated on a variety of standard smart contracts with a detailed description of the guarantees it can provide.

8.2 Future Directions

I conclude my thesis with a few broad future directions that closely align with resource-aware session types and the Nomos language. I have classified these directions in 3 broad projects.

- (i) design of optimal scheduling policies based on execution cost
- (ii) automatic synthesis of distributed programs from their specification
- (iii) implementation and analysis of cryptographic systems

Cost Analysis for Optimal Scheduling Cost analysis can assist developers in designing optimal scheduling policies for their applications. The sequential complexity bounds from Rast can be used to determine whether a new computation needs to be executed in the current or a freshly spawned thread. The parallel complexity bounds from Rast implicitly determine data dependency between threads and can be used to decide the order of thread execution. We can implement complexity-driven scheduling policies in Rast and evaluate their impact on performance.

Synthesis of Distributed Programs One effective way of assisting developers is by writing programs for them! Refinement session types can naturally express program specifications.

Programs can then be synthesized from their specifications by applying data-driven deep learning techniques. Training data can be randomly generated by first randomly creating a session-typed declaration for a process, and then synthesizing a process that follows the declaration. Refinements can also be utilized for semi-automatic synthesis of smart contracts. I believe refinement session types can propel program synthesis to distributed systems.

Rust for Cryptographic Protocols In a recent collaboration, I observed that refinement session types can neatly represent security protocols. With the recent growth in the complexity of such protocols, developers can greatly benefit from language support while building secure systems. We can also employ type-based techniques to model ill-behaved adversaries and formally verify cryptographic protocols. Furthermore, resource-aware types can provide computational security by modeling adversaries who are capable of only polynomial-time computation. We can also use programming language techniques to demonstrate *universal composability*, which entails that security properties of cryptographic protocols are preserved even when arbitrarily composed with other protocols.

More broadly, session types carry the safety guarantees that type systems can provide to the distributed domain. Resource-aware session types augment session types with cost information to provide computational guarantees along with safety. Together, they can inspire the design of next-generation concurrent programming languages. On the practical side, blockchains hold the potential of providing financial infrastructure access to underprivileged communities. Safe smart contract languages are the first step towards increase our trust in financial systems and improving their broader applicability. Nomos can inspire the design of future smart contract languages. In conclusion, programming language tools and techniques hold the power of improving software design and development and making technologies safer, faster, and more reliable!

Bibliography

- [1] Solidity by example. <https://solidity.readthedocs.io/en/v0.3.2/solidity-by-example.html>, march 2016. Accessed: 2018-11-04. 1.1, 6.1
- [2] Ether.camp’s hkg token has a bug and needs to be reissued. <https://www.ethnews.com/ethercamps-hkg-token-has-a-bug-and-needs-to-be-reissued>, January 2017. Accessed: 2019-02-25. 7.6
- [3] Bamboo. <https://github.com/cornellblockchain/bamboo>, August 2018. Accessed: 2018-11-04. 7.7
- [4] Welcome to liquidity’s documentation! <http://www.liquidity-lang.org/doc/index.html>, August 2018. Accessed: 2018-11-04. 1.1, 6.1
- [5] The michelson language. <https://www.michelson-lang.com/>, August 2018. Accessed: 2018-11-04. 7.7
- [6] Rholang. <https://github.com/rchain/Rholang>, August 2018. Accessed: 2018-11-04. 1.1, 7.7
- [7] The rust programming language. <https://doc.rust-lang.org/book>, 2018. 6.1
- [8] Vyper. <https://vyper.readthedocs.io/en/latest/index.html>, August 2018. Accessed: 2018-11-04. 7.7
- [9] Umut A Acar, Arthur Charguéraud, and Mike Rainey. Oracle-Guided Scheduling for Controlling Granularity in Implicitly Parallel Languages. *J. Funct. Programming*, 2016. 4.8
- [10] Elvira Albert, Puri Arenas, Samir Genaim, and Germán Puebla. Closed-Form Upper Bounds in Static Cost Analysis. *Journal of Automated Reasoning*, pages 161–203, 2011. 2.2
- [11] Elvira Albert, Puri Arenas, Samir Genaim, Miguel Gómez-Zamalloa, and Germán Puebla. Automatic Inference of Resource Consumption Bounds. In *18th Conference on Logic for Programming, Artificial Intelligence, and Reasoning (LPAR’12)*, 2012. 2.2
- [12] Elvira Albert, Puri Arenas, Samir Genaim, German Puebla, and Damiano Zanardini. Cost Analysis of Object-Oriented Bytecode Programs. *Theor. Comput. Sci.*, 413(1):142 – 159, 2012. 2.2
- [13] Elvira Albert, Puri Arenas, Jesús Correas, Samir Genaim, Miguel Gómez-Zamalloa, Enrique Martin-Martin, Germán Puebla, and Guillermo Román-Díez. Resource analysis: From sequential to concurrent and distributed programs. In *FM’15*, 2015. 4.7

- [14] Elvira Albert, Jesús Correas, Einar Broch Johnsen, and Guillermo Román-Díez. Parallel cost analysis of distributed systems. In Sandrine Blazy and Thomas Jensen, editors, *Static Analysis*, pages 275–292, Berlin, Heidelberg, 2015. Springer Berlin Heidelberg. ISBN 978-3-662-48288-9. [5.6](#)
- [15] Elvira Albert, Antonio Flores-Montoya, Samir Genaim, and Enrique Martin-Martin. May-Happen-in-Parallel Analysis for Actor-Based Concurrency. *ACM Trans. Comput. Log.*, 2016. [4.7](#)
- [16] Elvira Albert, Jesús Correas, Pablo Gordillo, Guillermo Román-Díez, and Albert Rubio. Gasol: Gas analysis and optimization for ethereum smart contracts, 2019. [7.7](#)
- [17] Elvira Albert, Pablo Gordillo, Albert Rubio, and Ilya Sergey. Running on fumes. In Pierre Ganty and Mohamed Kaâniche, editors, *Verification and Evaluation of Computer and Communication Systems*, pages 63–78, Cham, 2019. Springer International Publishing. ISBN 978-3-030-35092-5. [7.7](#)
- [18] Christophe Alias, Alain Darte, Paul Feautrier, and Laure Gonnord. Multi-dimensional Rankings, Program Termination, and Complexity Bounds of Flowchart Programs. In *17th Int. Static Analysis Symposium (SAS’10)*, 2010. [2.2](#)
- [19] Bernardo Almeida, Andreia Mordido, and Vasco T. Vasconcelos. Deciding the bisimilarity of context-free session types. In A. Biere and D. Parker, editors, *16th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS 2020)*, pages 39–56, Dublin, Ireland, April 2020. Springer LNCS 12079. [3.10](#)
- [20] Jean-Marc Andreoli. Logic programming with focusing proofs in linear logic. *Journal of Logic and Computation*, 2:297–347, 1992. [3.1](#), [3.7](#)
- [21] Robert Atkey. Amortised Resource Analysis with Separation Logic. In *19th Euro. Symp. on Prog. (ESOP’10)*, 2010. [4.7](#)
- [22] Nicola Atzei, Massimo Bartoletti, and Tiziana Cimoli. A survey of attacks on ethereum smart contracts (sok). In *Principles of Security and Trust - 6th International Conference, POST 2017*, pages 164–186, 2017. doi: 10.1007/978-3-662-54455-6_8. URL https://doi.org/10.1007/978-3-662-54455-6_8. [6.1](#), [7.7](#)
- [23] Martin Avanzini, Ugo Dal Lago, and Georg Moser. Analysing the complexity of functional programs: Higher-order meets first-order. In *Proceedings of the 20th ACM SIGPLAN International Conference on Functional Programming, ICFP 2015*, pages 152–164, New York, NY, USA, 2015. ACM. ISBN 978-1-4503-3669-7. doi: 10.1145/2784731.2784753. URL <http://doi.acm.org/10.1145/2784731.2784753>. [2.2](#), [6.5](#)
- [24] Pedro Baltazar, Dimitris Mostrous, and Vasco T. Vasconcelos. Linearly refined session types. In S. Alves and I. Mackie, editors, *International Workshop on Linearity (LINEARITY 2012)*, pages 38–49, Tallinn, Estonia, April 2012. EPTCS 101. doi: 10.4204/eptcs.101.4. [3.10](#)
- [25] Stephanie Balzer and Frank Pfenning. Manifest sharing with session types. *Proceedings of the ACM on Programming Languages (PACMPL)*, 1(ICFP):37:1–37:29, 2017. [1.1](#), [3.10](#), [3.10](#), [3.11](#), [5.6](#), [6.1](#), [1](#), [6.2](#), [6.3](#), [6.3](#), [6.6](#), [7.1](#), [7.4](#), [7.7](#)
- [26] Stephanie Balzer, Frank Pfenning, and Bernardo Toninho. A Universal Session Type

- for Untyped Asynchronous Communication. In Sven Schewe and Lijun Zhang, editors, *29th International Conference on Concurrency Theory (CONCUR 2018)*, volume 118 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 30:1–30:18, Dagstuhl, Germany, 2018. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. ISBN 978-3-95977-087-3. doi: 10.4230/LIPIcs.CONCUR.2018.30. URL <http://drops.dagstuhl.de/opus/volltexte/2018/9568>. 1
- [27] Stephanie Balzer, Bernardo Toninho, and Frank Pfenning. Manifest deadlock-freedom for shared session types. In Luís Caires, editor, *Programming Languages and Systems*, pages 611–639, Cham, 2019. Springer International Publishing. ISBN 978-3-030-17184-1. 7.4
- [28] Mike Barnett, Bor-Yuh Evan Chang, Robert DeLine, Bart Jacobs, and K. Rustan M. Leino. Boogie: A modular reusable verifier for object-oriented programs. In Frank S. de Boer, Marcello M. Bonsangue, Susanne Graf, and Willem-Paul de Roever, editors, *Formal Methods for Components and Objects*, pages 364–387, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg. ISBN 978-3-540-36750-5. 1
- [29] Bruno Barras, Samuel Boutin, Cristina Cornes, Judicaël Courant, Jean-Christophe Filliâtre, Eduardo Giménez, Hugo Herbelin, Gérard Huet, César Muñoz, Chetan Murthy, Catherine Parent, Christine Paulin-Mohring, Amokrane Saïbi, and Benjamin Werner. The Coq Proof Assistant Reference Manual : Version 6.1. Research Report RT-0203, INRIA, May 1997. URL <https://hal.inria.fr/inria-00069968>. Projet COQ. 1
- [30] Massimo Bartoletti, Tiziana Cimoli, Maurizio Murgia, Alessandro Sebastian Podda, and Livio Pompianu. A contract-oriented middleware. In C. Braga and P. Ölveczky, editors, *Formal Aspects of Component Software (FACS 2015)*, pages 86–104. Springer LNCS 9539, 2015. 3.10
- [31] Massimo Bartoletti, Tiziana Cimoli, Maurizio Murgia, Alessandro Sebastian Podda, and Livio Pompianu. Compliance and subtyping in timed session types. In S. Graf and M. Viswanathan, editors, *Formal Techniques for Distributed Objects, Components, and Systems (FORTE 2015)*, pages 161–177, Grenoble, France, June 2015. Springer LNCS 9039. 3.10
- [32] Mathieu Baudet, Avery Ching, Andrey Chursin, George Danezis, François Garillot, Zekun Li, Dahlia Malkhi, Oded Naor, Dmitri Perelman, and Alberto Sonnino. State machine replication in the libra blockchain, 2019. URL <https://developers.libra.org/docs/assets/papers/libra-consensus-state-machine-replication-in-the-libra-blockchain.pdf>. 7.4.2
- [33] P. N. Benton. A mixed linear and non-linear logic: Proofs, terms and models. In *8th International Workshop on Computer Science Logic (CSL)*, volume 933 of *Lecture Notes in Computer Science*, pages 121–135. Springer, 1994. An extended version appeared as Technical Report UCAM-CL-TR-352, University of Cambridge. 6.3
- [34] Gérard Berry and Georges Gonthier. The estereel synchronous programming language: Design, semantics, implementation. *Sci. Comput. Program.*, 19(2):87–152, November 1992. ISSN 0167-6423. doi: 10.1016/0167-6423(92)90005-V. URL <http://dx.doi.org/10.>

- [35] Karthikeyan Bhargavan, Antoine Delignat-Lavaud, Cédric Fournet, Anitha Gollamudi, Georges Gonthier, Nadim Kobeissi, Natalia Kulatova, Aseem Rastogi, Thomas Sibut-Pinote, Nikhil Swamy, and Santiago Zanella-Béguelin. Formal verification of smart contracts: Short paper. In *Proceedings of the 2016 ACM Workshop on Programming Languages and Analysis for Security*, PLAS '16, pages 91–96, New York, NY, USA, 2016. ACM. ISBN 978-1-4503-4574-3. doi: 10.1145/2993600.2993611. URL <http://doi.acm.org/10.1145/2993600.2993611>. 7.7
- [36] Sam Blackshear, David L. Dill, Shaz Qadeer, Clark W. Barrett, John C. Mitchell, Oded Padon, and Yoni Zohar. Resources: A safe language abstraction for money, 2020. 7.7
- [37] Guy E. Blelloch and Margaret Reid-Miller. Pipelining with futures. In *Proceedings of the Ninth Annual ACM Symposium on Parallel Algorithms and Architectures*, SPAA '97, pages 249–259, New York, NY, USA, 1997. ACM. ISBN 0-89791-890-8. doi: 10.1145/258492.258517. URL <http://doi.acm.org/10.1145/258492.258517>. 5.6
- [38] Laura Bocchi, Weizhen Yang, and Nobuko Yoshida. Timed multiparty session types. In Paolo Baldan and Daniele Gorla, editors, *CONCUR 2014 – Concurrency Theory*, pages 419–434, Berlin, Heidelberg, 2014. Springer Berlin Heidelberg. ISBN 978-3-662-44584-6. 5.6
- [39] Richard P Brent. *Algorithms for minimization without derivatives*. Courier Corporation, 2013. 4.1
- [40] Marc Brockschmidt, Fabian Emmes, Stephan Falke, Carsten Fuhs, and Jürgen Giesl. Alternating Runtime and Size Complexity Analysis of Integer Programs. In *20th Int. Conf. on Tools and Alg. for the Constr. and Anal. of Systems (TACAS'14)*, 2014. 2.2
- [41] Christian Cachin. Architecture of the hyperledger blockchain fabric. In *Workshop on Distributed Cryptocurrencies and Consensus Ledgers*, volume 310, 2016. 6.1
- [42] Luís Caires and Frank Pfenning. Session types as intuitionistic linear propositions. In *21st International Conference on Concurrency Theory (CONCUR)*, pages 222–236. Springer, 2010. 2.1, 5.6, 6.1, 7.1
- [43] Luís Caires, Jorge A. Pérez, Frank Pfenning, and Bernardo Toninho. Behavioral polymorphism and parametricity in session-based communication. In M.Felleisen and P.Gardner, editors, *Proceedings of the European Symposium on Programming (ESOP'13)*, pages 330–349, Rome, Italy, March 2013. Springer LNCS 7792. 5.5
- [44] Cristiano Calcagno and Dino Distefano. Infer: An automatic program verifier for memory safety of C programs. In Mihaela Bobaru, Klaus Havelund, Gerard J. Holzmann, and Rajeev Joshi, editors, *NASA Formal Methods*, pages 459–465, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg. ISBN 978-3-642-20398-5. 1
- [45] Quentin Carbonneaux, Jan Hoffmann, Thomas Reps, and Zhong Shao. Automated Resource Analysis with Coq Proof Objects. In *29th International Conference on Computer-Aided Verification (CAV'17)*, 2017. 4.7
- [46] Ezgi Çiçek, Gilles Barthe, Marco Gaboardi, Deepak Garg, and Jan Hoffmann. Relational Cost Analysis. In *44th Symposium on Principles of Programming Languages (POPL'17)*,

2017. 6.5

- [47] Pavol Cerný, Thomas A. Henzinger, Laura Kovács, Arjun Radhakrishna, and Jakob Zwirchmayr. Segment Abstraction for Worst-Case Execution Time Analysis. In *24th European Symposium on Programming (ESOP'15)*, 2015. 2.2
- [48] Iliano Cervesato and Andre Scedrov. Relating state-based and process-based concurrency through linear logic (full-version). *Information and Computation*, 207(10):1044 – 1077, 2009. ISSN 0890-5401. doi: <https://doi.org/10.1016/j.ic.2008.11.006>. URL <http://www.sciencedirect.com/science/article/pii/S089054010900100X>. Special issue: 13th Workshop on Logic, Language, Information and Computation (WoLLIC 2006). 2.1, 3.5
- [49] Ruofei Chen and Stephanie Balzer. Ferrite: A judgmental embedding of session types in Rust. *CoRR*, abs/2009.13619, 2020. URL <http://arxiv.org/abs/2009.13619>. 3.10
- [50] Ezgi Çiçek, Deepak Garg, and Umut Acar. Refinement types for incremental computational complexity. In Jan Vitek, editor, *Programming Languages and Systems*, pages 406–431, Berlin, Heidelberg, 2015. Springer Berlin Heidelberg. ISBN 978-3-662-46669-8. 4.7
- [51] Michael J. Coblenz, Reed Oei, Tyler Etzel, Paulette Koronkevich, Miles Baker, Yannick Bloem, Brad A. Myers, Joshua Sunshine, and Jonathan Aldrich. Obsidian: Typestate and assets for safer blockchain programming. *CoRR*, abs/1909.03523, 2019. URL <http://arxiv.org/abs/1909.03523>. 7.7
- [52] David C Cooper. Theorem proving in arithmetic without multiplication. *Machine Intelligence*, 7(91-99):300, 1972. 3.1, 3.8, 3.11
- [53] Karl Crary and Michael J. Sullivan. Peer-to-peer affine commitment using bitcoin. In *Proceedings of the 36th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '15*, pages 479–488, New York, NY, USA, 2015. ACM. ISBN 978-1-4503-3468-6. doi: 10.1145/2737924.2737997. URL <http://doi.acm.org/10.1145/2737924.2737997>. 7.7
- [54] Chris Dannen. *Introducing Ethereum and Solidity*. Springer, 2017. 1.1, 6.1, 6.2, 7.7, 7.7
- [55] Norman Danner, Daniel R. Licata, and Ramyaa Ramyaa. Denotational cost semantics for functional languages with inductive types. In *Proceedings of the 20th ACM SIGPLAN International Conference on Functional Programming, ICFP 2015*, pages 140–151, New York, NY, USA, 2015. ACM. ISBN 978-1-4503-3669-7. doi: 10.1145/2784731.2784749. URL <http://doi.acm.org/10.1145/2784731.2784749>. 6.5
- [56] A. Das, S. Balzer, J. Hoffmann, F. Pfenning, and I. Santurkar. Resource-aware session types for digital contracts. In *2021 IEEE 34th Computer Security Foundations Symposium (CSF)*, pages 111–126, Los Alamitos, CA, USA, jun 2021. IEEE Computer Society. doi: 10.1109/CSF51468.2021.00004. URL <https://doi.ieeecomputersociety.org/10.1109/CSF51468.2021.00004>. 1.1, 7.2
- [57] Ankush Das. Nomos implementation. <https://github.com/ankushdas/Nomos>, 2019. Accessed: 2019-11-11. 1.1, 7.3

- [58] Ankush Das and Frank Pfenning. Session types with arithmetic refinements. In I. Konnov and L. Kovács, editors, *31st International Conference on Concurrency Theory (CONCUR 2020)*, volume 171, pages 13:1–13:18. LIPIcs 171, 2020. doi: 10.4230/LIPIcs.CONCUR.2020.13. URL <https://drops.dagstuhl.de/opus/volltexte/2020/12825>. 3.11
- [59] Ankush Das and Frank Pfenning. Verified linear session-typed concurrent programming. In *Proceedings of the 22nd International Symposium on Principles and Practice of Declarative Programming (PPDP 2020)*, pages 7:1–7:15, Bologna, Italy, September 2020. ACM. 3.5.1, 3.8, 3.8, 3.11
- [60] Ankush Das and Shaz Qadeer. Exact and linear-time gas-cost analysis. In David Pichardie and Mihaela Sighireanu, editors, *Static Analysis*, pages 333–356, Cham, 2020. Springer International Publishing. ISBN 978-3-030-65474-0. 1.1, 7.4.2
- [61] Ankush Das, Jan Hoffmann, and Frank Pfenning. Work analysis with resource-aware session types. In *33rd ACM/IEEE Symposium on Logic in Computer Science (LICS’18)*, 2018. 1, 5.1, 5.5, 6.1, 6.2, 6.5, 7.1, 7.7
- [62] Ankush Das, Jan Hoffmann, and Frank Pfenning. Parallel complexity analysis with temporal session types. In *23rd International Conference on Functional Programming (ICFP’18)*, 2018. 1, 5.2, 5.3, 5.5
- [63] Ankush Das, Farzaneh Derakhshan, and Frank Pfenning. Rast Implementation. <https://bitbucket.org/fpfenning/rast/src/master/>, 2019. Accessed: 2019-11-11. 3, 3.8, 3.11
- [64] Ankush Das, Henry DeYoung, Andreia Mordido, and Frank Pfenning. Nested session types. In Nobuko Yoshida, editor, *Programming Languages and Systems*, pages 178–206, Cham, 2021. Springer International Publishing. ISBN 978-3-030-72019-3. 3.8, 3.8, 3.10
- [65] Rowan Davies. A temporal logic approach to binding-time analysis. In E. Clarke, editor, *Proceedings of the Eleventh Annual Symposium on Logic in Computer Science*, pages 184–195, New Brunswick, New Jersey, July 1996. IEEE Computer Society Press. URL <http://www.cs.cmu.edu/afs/cs/user/rowan/www/papers/multbta.ps.Z>. 5.6
- [66] Farzaneh Derakhshan and Frank Pfenning. Circular proofs as session-typed processes: A local validity condition. *CoRR*, abs/1908.01909, August 2019. URL <http://arxiv.org/abs/1908.01909>. 8
- [67] A. Erbsen, J. Philipoom, J. Gross, R. Sloan, and A. Chlipala. Simple high-level code for cryptographic arithmetic - with proofs, without compromises. In *2019 IEEE Symposium on Security and Privacy (SP)*, pages 1202–1219, 2019. 1
- [68] Vitalik Buterin Fabian Vogelsteller. Erc20 token standard. <https://eips.ethereum.org/EIPS/eip-20>, december 2018. Accessed: 2018-02-027. 7.6
- [69] Nicolas Feltman, Carlo Angiuli, Umut Acar, and Kayvon Fatahalian. Automatically splitting a two-stage lambda calculus. In P. Thiemann, editor, *Proceedings of the 25th European Symposium on Programming (ESOP)*, pages 255–281, Eindhoven, The Netherlands, April 2016. Springer LNCS 9632. 5.6
- [70] Jérôme Fortier and Luigi Santocanale. Cuts for circular proofs: Semantics and cut elimi-

- nation. In *22nd Conference on Computer Science Logic*, volume 23 of *LIPICs*, pages 248–262, 2013. 5.3
- [71] Simon Fowler, Sam Lindley, J. Garrett Morris, and Sára Decova. Exceptional asynchronous session types: Session types without tiers. *Proc. ACM Program. Lang.*, 3(POPL), January 2019. doi: 10.1145/3290341. URL <https://doi.org/10.1145/3290341>. 3.10
- [72] Juliana Franco and Vasco T. Vasconcelos. A concurrent programming language with refined session types. In S. Counsell and M. Núñez, editors, *Software Engineering and Formal Methods (SEFM 2013)*, pages 15–28, Madrid, Spain, September 2013. Springer LNCS 8368. 3.10
- [73] Simon Gay and Malcolm Hole. Subtyping for session types in the pi calculus. *Acta Informatica*, 42(2):191–225, Nov 2005. ISSN 1432-0525. doi: 10.1007/s00236-005-0177-z. URL <https://doi.org/10.1007/s00236-005-0177-z>. 2.1, 3.1, 3.4, 3.5, 6.6, 7.3
- [74] Dan R. Ghica and Alex I. Smith. Bounded linear types in a resource semiring. In *Proceedings of the 23rd European Symposium on Programming Languages and Systems - Volume 8410*, pages 331–350, New York, NY, USA, 2014. Springer-Verlag New York, Inc. ISBN 978-3-642-54832-1. doi: 10.1007/978-3-642-54833-8_18. URL http://dx.doi.org/10.1007/978-3-642-54833-8_18. 5.6
- [75] Stéphane Gimenez and Georg Moser. The complexity of interaction. In *POPL’16*, 2016. 4.7, 5.5, 5.6
- [76] Jean-Yves Girard. Linear logic. *Theoretical Computer Science*, 50:1–102, 1987. 1.1, 2.1, 6.1, 7.1
- [77] Hannah Gommerstadt. *Session-Typed Concurrent Contracts*. PhD thesis, Carnegie Mellon University, September 2019. Available as Technical Report CMU-CS-19-119. 3.1
- [78] 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. 3.1
- [79] L.M Goodman. Tezos — a self-amending crypto-ledger. https://tezos.com/static/papers/white_paper.pdf, 2014. 1.1, 6.1
- [80] Neville Grech, Michael Kong, Anton Jurisevic, Lexi Brent, Bernhard Scholz, and Yannis Smaragdakis. Madmax: Surviving out-of-gas conditions in ethereum smart contracts. *Proc. ACM Program. Lang.*, 2(OOPSLA):116:1–116:27, October 2018. ISSN 2475-1421. doi: 10.1145/3276486. URL <http://doi.acm.org/10.1145/3276486>. 7.7
- [81] Dennis Griffith. *Polarized Substructural Session Types*. PhD thesis, University of Illinois at Urbana-Champaign, April 2016. 3.10, 3.11, 5.5
- [82] Dennis Griffith and Elsa L. Gunter. LiquidPi: Inferrable dependent session types. In *Proceedings of the NASA Formal Methods Symposium*, pages 186–197. Springer LNCS 7871, 2013. 3.1, 5.5
- [83] Ilya Grishchenko, Matteo Maffei, and Clara Schneidewind. Foundations and tools for the static analysis of ethereum smart contracts. In Hana Chockler and Georg Weissenbacher, editors, *Computer Aided Verification*, pages 51–78, Cham, 2018. Springer International

Publishing. ISBN 978-3-319-96145-3. [7.7](#)

- [84] Sumit Gulwani. Speed: Symbolic complexity bound analysis. In Ahmed Bouajjani and Oded Maler, editors, *Computer Aided Verification*, pages 51–62, Berlin, Heidelberg, 2009. Springer Berlin Heidelberg. ISBN 978-3-642-02658-4. [1](#)
- [85] Sumit Gulwani, Krishna K. Mehra, and Trishul M. Chilimbi. SPEED: Precise and Efficient Static Estimation of Program Computational Complexity. In *36th ACM Symp. on Principles of Prog. Langs. (POPL’09)*, 2009. [2.2](#)
- [86] N. Halbwachs, P. Caspi, P. Raymond, and D. Pilaud. The synchronous data flow programming language lustre. *Proceedings of the IEEE*, 79(9):1305–1320, Sep 1991. ISSN 0018-9219. doi: 10.1109/5.97300. [5.6](#)
- [87] Robert H. Halstead, Jr. Multilisp: A language for concurrent symbolic computation. *ACM Trans. Program. Lang. Syst.*, 7(4):501–538, October 1985. ISSN 0164-0925. doi: 10.1145/4472.4478. URL <http://doi.acm.org/10.1145/4472.4478>. [5.6](#)
- [88] Everett Hildenbrandt, Manasvi Saxena, Xiaoran Zhu, Nishant Rodrigues, Philip Daian, Dwight Guth, Brandon Moore, Yi Zhang, Daejun Park, Andrei Stefanescu, and Grigore Rosu. Kevm: A complete semantics of the ethereum virtual machine. In *2018 IEEE 31st Computer Security Foundations Symposium*, pages 204–217. IEEE, 2018. [7.7](#)
- [89] Jan Hoffmann and Martin Hofmann. Amortized Resource Analysis with Polynomial Potential. In *19th European Symposium on Programming (ESOP’10)*, 2010. [2.2](#)
- [90] Jan Hoffmann and Zhong Shao. Automatic Static Cost Analysis for Parallel Programs. In *24th European Symposium on Programming (ESOP’15)*, 2015. [4.7](#), [5.6](#)
- [91] Jan Hoffmann, Klaus Aehlig, and Martin Hofmann. Multivariate Amortized Resource Analysis. In *38th ACM Symp. on Principles of Prog. Langs. (POPL’11)*, 2011. [1](#), [2.2](#), [6.5](#), [6.5](#)
- [92] Jan Hoffmann, Ankush Das, and Shu-Chun Weng. Towards Automatic Resource Bound Analysis for OCaml. In *44th Symposium on Principles of Programming Languages (POPL’17)*, 2017. [1](#), [2.2](#), [4.7](#), [6.5](#), [6.5](#), [6.5](#), [6.6](#), [7.3](#), [7.7](#)
- [93] Martin Hofmann and Steffen Jost. Static Prediction of Heap Space Usage for First-Order Functional Programs. In *30th ACM Symp. on Principles of Prog. Langs. (POPL’03)*, 2003. [2.2](#), [4.7](#), [6.5](#), [6.5](#), [7.3](#), [7.7](#)
- [94] Martin Hofmann and Steffen Jost. Type-Based Amortised Heap-Space Analysis. In *15th Euro. Symp. on Prog. (ESOP’06)*, 2006. [4.7](#)
- [95] Martin Hofmann and Georg Moser. Multivariate Amortised Resource Analysis for Term Rewrite Systems. In *13th International Conference on Typed Lambda Calculi and Applications (TLCA’15)*, 2015. [4.7](#)
- [96] Kohei Honda. Types for dyadic interaction. In *4th International Conference on Concurrency Theory (CONCUR)*, pages 509–523. Springer, 1993. [2.1](#), [7.7](#)
- [97] Kohei Honda, Nobuko Yoshida, and Marco Carbone. Multiparty asynchronous session types. In *35th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, pages 273–284. ACM, 2008. [7.7](#)

- [98] Blockchain Insurance Industry Initiative. B3i. 2008. [1.1](#), [6.1](#)
- [99] Thomas Bracht Laumann Jespersen, Philip Munksgaard, and Ken Friis Larsen. Session types for Rust. In P. Bahr and S. Erdweg, editors, *Proceedings of the 11th Workshop on Generic Programming (WGP 2015)*, pages 13–22, Vancouver, Canada, August 2015. ACM. [3.10](#)
- [100] Steffen Jost, Kevin Hammond, Hans-Wolfgang Loidl, and Martin Hofmann. Static Determination of Quantitative Resource Usage for Higher-Order Programs. In *37th ACM Symp. on Principles of Prog. Langs. (POPL '10)*, 2010. [2.2](#)
- [101] Steve Klabnik and Carol Nichols. *The Rust Programming Language*. No Starch Press, USA, 2018. ISBN 1593278284. [7.7](#)
- [102] Naoki Kobayashi. A type system for lock-free processes. *Information and Computation*, 177:122–159, 2002. [5.6](#)
- [103] Naoki Kobayashi, Motoki Nakade, and Akinori Yonezawa. Static analysis of communication for asynchronous concurrent programming languages. In *SAS'95*, 1995. [4.7](#)
- [104] Neelakantan R. Krishnaswami and Nick Benton. Ultrametric Semantics of Reactive Programs. In *26th IEEE Symposium on Logic in Computer Science, (LICS'11)*, pages 257–266, 2011. [5.6](#)
- [105] Ugo Dal Lago and Marco Gaboardi. Linear Dependent Types and Relative Completeness. In *26th IEEE Symp. on Logic in Computer Science (LICS'11)*, 2011. [6.5](#)
- [106] Ugo Dal Lago and Barbara Petit. The Geometry of Types. In *40th ACM Symp. on Principles Prog. Langs. (POPL '13)*, 2013. [4.7](#)
- [107] Shuvendu K. Lahiri, Shuo Chen, Yuepeng Wang, and Isil Dillig. Formal specification and verification of smart contracts for azure blockchain. *CoRR*, abs/1812.08829, 2018. URL <http://arxiv.org/abs/1812.08829>. [7.7](#)
- [108] Angwei Law. *Smart contracts and their application in supply chain management*. PhD thesis, Massachusetts Institute of Technology, 2017. [1.1](#), [6.1](#)
- [109] K. Rustan M. Leino. Dafny: An automatic program verifier for functional correctness. In Edmund M. Clarke and Andrei Voronkov, editors, *Logic for Programming, Artificial Intelligence, and Reasoning*, pages 348–370, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg. ISBN 978-3-642-17511-4. [1](#)
- [110] Xavier Leroy. Formal verification of a realistic compiler. *Commun. ACM*, 52(7):107–115, July 2009. ISSN 0001-0782. doi: 10.1145/1538788.1538814. URL <https://doi.org/10.1145/1538788.1538814>. [1](#)
- [111] Christopher Lidbury and Alastair F. Donaldson. Sparse record and replay with controlled scheduling. In *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2019*, pages 576–593, New York, NY, USA, 2019. ACM. ISBN 978-1-4503-6712-7. doi: 10.1145/3314221.3314635. URL <http://doi.acm.org/10.1145/3314221.3314635>. [7.4](#)
- [112] Sam Lindley and J. Garrett Morris. Embedding session types in Haskell. In G. Mainland, editor, *Proceedings of the 9th International Symposium on Haskell (Haskell 2016)*, pages

- 133–145, Nara, Japan, September 2016. ACM. 3.10
- [113] Sam Lindley and J. Garrett Morris. Talking bananas: Structural recursion for session types. In *Proceedings of the 21st ACM SIGPLAN International Conference on Functional Programming*, ICFP 2016, page 434–447. ACM, 2016. ISBN 9781450342193. doi: 10.1145/2951913.2951921. URL <https://doi.org/10.1145/2951913.2951921>. 3.1, 3.10
- [114] Sam Lindley and J. Garrett Morris. Lightweight functional session types. In S. Gay and A. Ravara, editors, *Behavioural Types: from Theory to Tools*, chapter 12, pages 265–286. River Publishers, June 2017. doi: <https://doi.org/10.13052/rp-9788793519817>. 3.10
- [115] Hugo A. López, Carlos Olarte, and Jorge A. Pérez. Towards a unified framework for declarative structure communications. In A. Beresford and S. Gay, editors, *Proceedings of the Workshop on Programming Language Approaches to Concurrency and Communication-Centric Software (PLACES)*, pages 1–15. EPTCS 17, March 2009. 5.6
- [116] Loi Luu, Duc-Hiep Chu, Hrishi Olickel, Prateek Saxena, and Aquinas Hobor. Making smart contracts smarter. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, CCS ’16, pages 254–269, New York, NY, USA, 2016. ACM. ISBN 978-1-4503-4139-4. doi: 10.1145/2976749.2978309. URL <http://doi.acm.org/10.1145/2976749.2978309>. 7.6, 7.7
- [117] Matteo Marescotti, Martin Blicha, Antti E. J. Hyvärinen, Sepideh Asadi, and Natasha Sharygina. Computing exact worst-case gas consumption for smart contracts. In *International Symposium on Leveraging Applications of Formal Methods ISOLA 2018: Leveraging Applications of Formal Methods, Verification and Validation. Industrial Practice*, Cyprus, 2018. Springer, Springer. URL https://link.springer.com/chapter/10.1007%2F978-3-030-03427-6_33. 7.7
- [118] Lucius Gregory Meredith. Linear types can change the blockchain. *arXiv preprint arXiv:1506.01001*, 2015. 1.1, 6.1
- [119] Marvin L. Minsky. *Computation: Finite and Infinite Machines*. Prentice-Hall, Inc., USA, 1967. ISBN 0131655639. 3.3
- [120] Vincenzo Morabito. Smart contracts and licensing. In *Business Innovation Through Blockchain*, pages 101–124. Springer, 2017. 1.1, 6.1
- [121] David Z. Morris. Blockchain-based venture capital fund hacked for \$ 60 million. <http://fortune.com/2016/06/18/blockchain-vc-fund-hacked>, June 2016. 1.1, 6.1
- [122] Satoshi Nakamoto. Bitcoin: A Peer-to-Peer Electronic Cash System, 2008. 1.1, 6.1
- [123] Hiroshi Nakano. A Modality for Recursion. In *15th IEEE Symposium on Logic in Computer Science (LICS’00)*, pages 255–266, 2000. 5.6
- [124] Romyana Neykova, Laura Bocchi, and Nobuko Yoshida. Timed runtime monitoring for multiparty conversations. In *3rd International Workshop on Behavioural Types (BEAT 2014)*, 2014. 5.6
- [125] Dominic A. Orchard and Nobuko Yoshida. Effects as sessions, sessions as effects. In R. Bodik and R. Majumdar, editors, *Proceedings of the 43rd Annual Symposium on Principles of Programming Languages (POPL 2016)*, pages 568–581, St. Petersburg, Florida, January

2016. ACM. 3.10

- [126] Luca Padovani. A simple library implementation of binary sessions. *Journal of Functional Programming*, 27(e4), 2017. 3.10
- [127] W. Paul, U. Vishkin, and H. Wagoner. Parallel dictionaries on 2–3 trees. In Josep Diaz, editor, *Automata, Languages and Programming*, pages 597–609, Berlin, Heidelberg, 1983. Springer Berlin Heidelberg. ISBN 978-3-540-40038-7. 5.6
- [128] Lawrence C Paulson. *Isabelle: A generic theorem prover*, volume 828. Springer Science & Business Media, 1994. 1
- [129] Frank Pfenning and Dennis Griffith. Polarized substructural session types. In *18th International Conference on Foundations of Software Science and Computation Structures (FoSSaCS)*, pages 3–22. Springer, 2015. 2.1, 3.6, 6.1, 6.3, 6.4, 6.6
- [130] Benjamin C. Pierce and David N. Turner. Local type inference. *ACM Trans. Program. Lang. Syst.*, 22(1):1–44, January 2000. ISSN 0164-0925. doi: 10.1145/345099.345100. URL <http://doi.acm.org/10.1145/345099.345100>. 7.3
- [131] Amir Pnueli. The temporal logic of programs. In *Proceedings of the 18th Symposium on Foundations of Computer Science (FOCS’77)*, pages 46–57. IEEE Computer Society, 1977. 5.3
- [132] Francois Pottier and Yann Régis-Gianas. *Menhir Reference Manual*, 2019. 7.3
- [133] Marc Pouzet. Lucid synchrone release, version 3.0 tutorial and reference manual, 2006. 5.6
- [134] J. Protzenko, B. Parno, A. Fromherz, C. Hawblitzel, M. Polubelova, K. Bhargavan, B. Beurdouche, J. Choi, A. Delignat-Lavaud, C. Fournet, N. Kulatova, T. Ramananandro, A. Rastogi, N. Swamy, C. M. Wintersteiger, and S. Zanella-Beguelin. Evercrypt: A fast, verified, cross-platform cryptographic provider. In *2020 IEEE Symposium on Security and Privacy (SP)*, pages 983–1002, Los Alamitos, CA, USA, may 2020. IEEE Computer Society. doi: 10.1109/SP40000.2020.00114. URL <https://doi.ieeecomputersociety.org/10.1109/SP40000.2020.00114>. 1
- [135] Klaas Pruiksma and Frank Pfenning. A message-passing interpretation of adjoint logic. In F. Martins and D. Orchard, editors, *Workshop on Programming Language Approaches to Concurrency and Communication-Centric Software (PLACES)*, pages 60–79, Prague, Czech Republic, April 2019. EPTCS 291. 3.11
- [136] Klaas Pruiksma and Frank Pfenning. Back to futures. *CoRR*, abs/2002.04607, February 2020. URL <http://arxiv.org/abs/2002.04607>. 3.8, 3.11
- [137] Klaas Pruiksma, William Chargin, Frank Pfenning, and Jason Reed. Adjoint logic. Technical report, Carnegie Mellon University, April 2018. 6.3
- [138] Ivan Radiček, Gilles Barthe, Marco Gaboardi, Deepak Garg, and Florian Zuleger. Monadic Refinements for Relational Cost Analysis. *Proc. ACM Program. Lang.*, 2(POPL), 2017. 6.5
- [139] Jason Reed. A judgmental deconstruction of modal logic. Unpublished manuscript, January 2009. URL <http://www.cs.cmu.edu/~jcreed/papers/jdml.pdf>. 6.3

- [140] Patrick M. Rondon, Ming Kawaguci, and Ranjit Jhala. Liquid types. In *Proceedings of the 29th ACM SIGPLAN Conference on Programming Language Design and Implementation*, PLDI '08, page 159–169, New York, NY, USA, 2008. Association for Computing Machinery. ISBN 9781595938602. doi: 10.1145/1375581.1375602. URL <https://doi.org/10.1145/1375581.1375602>. 3.1
- [141] Michiel Ronsse and Koen De Bosschere. Replay: A fully integrated practical record/replay system. *ACM Trans. Comput. Syst.*, 17(2):133–152, May 1999. ISSN 0734-2071. doi: 10.1145/312203.312214. URL <http://doi.acm.org/10.1145/312203.312214>. 7.4
- [142] Neda Saeedloei and Gopal Gupta. Timed π -calculus. In *8th International Symposium on Trustworthy Global Computing - Volume 8358*, TGC 2013, pages 119–135, New York, NY, USA, 2014. Springer-Verlag New York, Inc. ISBN 978-3-319-05118-5. doi: 10.1007/978-3-319-05119-2_8. URL https://doi.org/10.1007/978-3-319-05119-2_8. 5.6
- [143] Alceste Scalas and Nobuko Yoshida. Lightweight session programming in Scala. In *Proceedings of the 30th European Conference on Object-Oriented Programming (ECOOP 2016)*, pages 21:1–21:28, Rome, Italy, July 2016. LIClcs 56. 3.10
- [144] Alceste Scalas and Nobuko Yoshida. Less is more: Multiparty session types revisited. *Proc. ACM Program. Lang.*, 3(POPL):30:1–30:29, January 2019. ISSN 2475-1421. doi: 10.1145/3290343. URL <http://doi.acm.org/10.1145/3290343>. 7.7
- [145] Ilya Sergey, Amrit Kumar, and Aquinas Hobor. Scilla: a smart contract intermediate-level language. *CoRR*, abs/1801.00687, 2018. URL <http://arxiv.org/abs/1801.00687>. 7.7
- [146] Miguel Silva, Mário Florido, and Frank Pfenning. Non-blocking concurrent imperative programming with session types. In *Fourth International Workshop on Linearity*, June 2016. 5.6
- [147] Hugo R. Simões, Pedro B. Vasconcelos, Mário Florido, Steffen Jost, and Kevin Hammond. Automatic Amortised Analysis of Dynamic Memory Allocation for Lazy Functional Programs. In *17th Int. Conf. on Funct. Prog. (ICFP'12)*, 2012. 2.2
- [148] Moritz Sinn, Florian Zuleger, and Helmut Veith. A Simple and Scalable Approach to Bound Analysis and Amortized Complexity Analysis. In *Computer Aided Verification - 26th Int. Conf. (CAV'14)*, 2014. 2.2
- [149] Robert Endre Tarjan. Amortized computational complexity. *SIAM Journal on Algebraic Discrete Methods*, 6(2):306–318, 1985. 1, 2.2.1
- [150] Tachio Terauchi and Adam Megacz. Inferring channel buffer bounds via linear programming. In *ESOP'08*, 2008. 4.7
- [151] Peter Thiemann and Vasco T. Vasconcelos. Context-free session types. In *Proceedings of the 21st International Conference on Functional Programming (ICFP 2016)*, pages 462–475, Nara, Japan, September 2016. ACM. 3.10
- [152] S. Tikhomirov, E. Voskresenskaya, I. Ivanitskiy, R. Takhaviev, E. Marchenko, and Y. Alexandrov. Smartcheck: Static analysis of ethereum smart contracts. In *2018 IEEE/ACM 1st International Workshop on Emerging Trends in Software Engineering for Blockchain (WETSEB)*, pages 9–16, May 2018. 7.7

- [153] Bernardo Toninho, Luís Caires, and Frank Pfenning. Higher-order processes, functions, and sessions: a monadic integration. In *22nd European Symposium on Programming (ESOP)*, pages 350–369. Springer, 2013. [3.10](#), [4.5](#), [6.1](#), [1](#), [6.4](#), [7.7](#)
- [154] Bernardo Toninho, Luís Caires, and Frank Pfenning. Corecursion and non-divergence in session-typed processes. In M. Maffei and E. Tuosto, editors, *Proceedings of the 9th International Symposium on Trustworthy Global Computing (TGC 2014)*, pages 159–175, Rome, Italy, September 2014. Springer LNCS 8902. [5.3](#)
- [155] Petar Tsankov, Andrei Dan, Dana Drachler-Cohen, Arthur Gervais, Florian Bünzli, and Martin Vechev. Securify: Practical security analysis of smart contracts. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, CCS '18*, page 67–82, New York, NY, USA, 2018. Association for Computing Machinery. ISBN 9781450356930. doi: 10.1145/3243734.3243780. URL <https://doi.org/10.1145/3243734.3243780>. [7.7](#)
- [156] Pedro Vasconcelos. *Space Cost Analysis Using Sized Types*. PhD thesis, School of Computer Science, University of St Andrews, 2008. [2.2](#)
- [157] Vasco T. Vasconcelos. Session, from types to programming languages. *Bulletin of the EATCS*, 103:53–73, 2011. [3.10](#)
- [158] Philip Wadler. Linear types can change the world! In *IFIP TC 2 Working Conference on Programming Concepts and Methods*, pages 546–566. North, 1990. [6.1](#)
- [159] Philip Wadler. Propositions as sessions. In *17th ACM SIGPLAN International Conference on Functional Programming (ICFP)*, pages 273–286. ACM, 2012. [5.6](#), [6.1](#)
- [160] Max Willsey, Rohini Prabhu, and Frank Pfenning. Design and implementation of Concurrent C0. In *Fourth International Workshop on Linearity*, June 2016. [3.10](#), [4.7](#)
- [161] Gavin Wood. Ethereum: A secure decentralized transaction ledger, 2014. [1.1](#), [6.1](#), [7.1](#)
- [162] Hongwei Xi and Frank Pfenning. Dependent types in practical programming. In *Proceedings of the 26th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '99*, pages 214–227, New York, NY, USA, 1999. ACM. ISBN 1-58113-095-3. doi: 10.1145/292540.292560. URL <http://doi.acm.org/10.1145/292540.292560>. [3.1](#)
- [163] Hongwei Xi, Zhiqiang Ren, Hanwen Wu, and William Blair. Session types in a linearly typed multi-threaded lambda-calculus. *CoRR*, abs/1603.03727, 2016. URL <http://arxiv.org/abs/1603.03727>. [3.10](#)
- [164] Christoph Zenger. Indexed types. *Theor. Comput. Sci.*, 187(1–2):147–165, November 1997. ISSN 0304-3975. doi: 10.1016/S0304-3975(97)00062-5. URL [https://doi.org/10.1016/S0304-3975\(97\)00062-5](https://doi.org/10.1016/S0304-3975(97)00062-5). [3.1](#)
- [165] Fangyi Zhou, Francisco Ferreira, Rumyana Neykova, and Nobuko Yoshida. Fluid Types: Statically Verified Distributed Protocols with Refinements. In *11th Workshop on Programming Language Approaches to Concurrency and Communication-Centric Software*, 2019. [3.1](#)
- [166] Florian Zuleger, Moritz Sinn, Sumit Gulwani, and Helmut Veith. Bound Analysis of Imperative Programs with the Size-change Abstraction. In *18th Int. Static Analysis Symp.*

(SAS'11), 2011. [2.2](#)