TECHNISCHE
UNIVERSITAT
DARMSTADT

# Type-Safe Data Plane Programming 

vom Fachbereich Informatik<br>der Technischen Universität Darmstadt genehmigte<br>Dissertation<br>zur Erlangung des akademischen Grades eines<br>Doktor-Ingenieurs (Dr.-Ing.)<br>vorgelegt von<br>M.Sc. Matthias Eichholz<br>geboren in Oberhausen<br>Gutachter<br>Prof. Dr.-Ing. Mira Mezini<br>Technische Universität Darmstadt<br>Prof. Nate Foster<br>Cornell University, USA

Matthias Eichholz: Type-Safe Data Plane Programming
Darmstadt, Technische Universität Darmstadt
Tag der mündlichen Prüfung: 12.10.2022
Jahr der Veröffentlichung der Dissertation auf TUprints: 2022
URN: urn:nbn:de:tuda-tuprints-228736
Veröffentlicht unter CC BY-SA 4.0 International
https://creativecommons.org/licenses/by-sa/4.0/


#### Abstract

Since the mid-1990s, there have been efforts to enable more flexible processing of network packets by making packet processing programmable. With the advent of software-defined networking (SDN), this idea has now become a reality. Early approaches initially focused on control plane programming, with the goal of implementing centralized network policies at a high level of abstraction without having to use low-level, device-specific configuration mechanisms. For this purpose, various network programming languages have been developed, which provide correctness guarantees and make the formal verification of network policies possible.

More recently, it is also possible to program the network data plane. Being able to define the structure of network packet headers freely, opens up a whole new range of applications, from implementing new network protocols up to moving application logic directly into the network. Until today, the $\mathrm{P}_{4}$ language has become the de facto standard for programming data planes. While $\mathrm{P}_{4}$ provides declarative abstractions for programming data planes, $\mathrm{P}_{4}$ lacks basic safety guarantees to help avoid errors and implement correct applications for the data plane.

Modern programming languages use static type systems to provide languages with basic safety guarantees that completely eliminate the occurrence of entire categories of errors. Surprisingly, however, the use of type systems in the field of network programming has hardly been investigated. This dissertation investigates what appropriate type systems must look like in order to provide data plane programming languages-in particular, $\mathrm{P}_{4}$-with static correctness guarantees. As a first step, we present SafeP4, a domain-specific language for programmable data planes that is equipped with a static type system that guarantees that all headers that are read or written are valid, which is a common cause of errors. We then present $\Pi_{4}$, whose type system is based on dependent types and is thus able to bridge the gap in terms of expressiveness between $\mathrm{SAFEP}_{4}$ and full-fledged verification tools. At the same time, $\Pi_{4}$ enables modular verification of programs.

Our evaluation using open source programs confirms that accessing invalid packet headers is a common source of errors in practice and that the SafeP4's type system is capable of identifying buggy programs. Using case studies, we show that $\Pi$ 's type system is capable of expressing and verifying a variety of real-world correctness properties.


## Zusammenfassung

Seit Mitte der 1990er Jahre gibt es Bestrebungen, eine flexiblere Verarbeitung von Netzwerkpaketen zu ermöglichen, indem Paketverarbeitung programmierbar wird. Mit dem Aufkommen von Software-Defined-Networking (SDN) ist diese Idee nun Realität geworden. Frühe Ansätze haben sich zunächst auf die Programmierung der Control-Plane konzentriert, mit dem Ziel, auf hohem Abstraktionsniveau zentrale Netzwerkrichtlinien zu realisieren, ohne systemnahe, gerätespezifische Konfigurationsmechanismen nutzen zu müssen. Hierfür entstanden diverse Programmiersprachen, die Korrektheitsgarantien gewähren und die formale Verifikation der Netzwerkrichtlinien ermöglichen.

In jüngster Zeit ist es zudem auch möglich die Netzwerk-Data-Plane zu programmieren. Durch die Möglichkeit die Struktur von Netzwerkpaketheadern frei zu definieren, eröffnen sich eine Vielzahl neuer Anwendungsmöglichkeiten, von der Implementierung neuartiger Netzwerkprotokolle bis hin zur Auslagerung von Anwendungsfunktionalität direkt in das Netzwerk. Bis heute hat sich die Sprache $\mathrm{P}_{4}$ als Standard für die Programmierung von Data-Planes durchgesetzt. Zwar bietet P4 deklarative Abstraktionen für die Programmierung von Data-Planes, allerdings fehlen $\mathrm{P}_{4}$ grundlegende Sicherheitsgarantien, die dabei helfen Fehler zu vermeiden und korrekte Anwendungen für die Data-Plane zu implementieren.

Modernen Programmiersprachen nutzen statische Typsysteme, um Sprachen mit grundlegenden Sicherheitsgarantien auszustatten, die das Auftreten ganzer Fehlerkategorien vollständig ausschließen. Überraschenderweise wurde der Einsatz von Typsysteme im Bereich der Netzwerkprogrammierung jedoch bislang kaum untersucht. Diese Dissertation untersucht wie geeignete Typsysteme aussehen müssen, um Data-Plane-Programmiersprachen - insbesondere P4 - mit statischen Korrektheitsgarantien auszustatten. Im ersten Schritt präsentieren wir SafeP4, eine domänenspezifische Sprache für programmierbare Data-Planes, die mit einem statischen Typsystem ausgestattet ist, das garantiert, dass alle Header, die gelesen oder geschrieben werden gültig sind, was eine häufige Ursache für Fehler ist. Im zweiten Schritt präsentieren wir $\Pi_{4}$, dessen Typsystem auf Dependent-Types basiert und damit in der Lage ist, die Lücke hinsichtlich der Ausdrucksstärke zwischen SAFEP4 und vollwertigen Verifikationswerkzeugen zu schließen. Gleichzeitig ermöglicht $\Pi_{4}$ die modulare Verifikation von Programmen.

Unsere Auswertung anhand von Open-Source-Programmen bestätigt, dass der Zugriff auf ungültige Paketheader in der Praxis eine häufige Fehlerquelle ist und das SafeP4s Typsystem in der Lage ist, fehlerhafte Programme zu identifizieren. Anhand von Fallstudien zeigen wir zudem, dass $\Pi_{4 s}$ Typsystem imstande ist eine Vielzahl praktisch relevanter Korrektheitseigenschaften auszudrücken und zu verifizieren.

## Acknowledgements

First and foremost, I thank my advisor Mira Mezini, who made it possible for me to pursue a Phd in the first place and who helped me to ultimately achieve this goal. Thank you, Mira, for the freedom to follow my own research ideas. However, I am at least as grateful to Nate Foster, without whom many aspects of this work would probably have developed quite differently. Thank you, Nate, I really learned a lot from you about what it needs to be a PL researcher.

I also like to thank the remaining members of my PhD committee Reiner Hähnle, Zsolt István, and Marie-Christine Jakobs. Furthermore, I like to thank many others whom I met and worked with during my time as a Phd student. In particular, I thank Eric Hayden Campbell for his contributions to the papers that form the basis of this thesis, and for the many discussions about formalizations and proofs. I thank my co-authors and fellow PhD students Marcel Blöcher, Matthias Krebs, Johannes Krude, Katharina Keller, and Artur Sterz. I thank all members of the Software Technology Group especially my long-time office colleagues Ragnar Mogk, Pascal Weisenburger and Mirko Köhler for the numerous conversations, discussions and the occasional, very entertaining rants. Thanks to Gudrun Harris and Claudia Roßmann, who always had your back when it came to bureaucracy. I also thank former group member Guido Salvaneschi for his support at the beginning of my time as a PhD student.

I thank my family, especially my parents and my sister for always being there for me and believing in me. Last but not least, I like to thank my wife Christin and my little son Jonas. Thank you, Christin, for accompanying and supporting me all these years through all the ups and downs-it is finally done, but now for real. Thank you, Jonas, for always brightening up my day with your little smile. I love you.

## Contents

I Prologue ..... 1
1 Introduction ..... 3
1.1 Problem Statement ..... 4
1.2 State of the Art ..... 5
1.3 The Thesis in a Nutshell ..... 6
1.3.1 $\mathrm{SafeP}_{4}$ ..... 7
1.3.2 $\Pi_{4}$ ..... 7
1.4 Contributions ..... 8
1.5 List of Publications ..... 9
1.6 Structure of the Thesis ..... 10
2 Background ..... 11
2.1 Programmable Packet Processing ..... 11
2.2 The P4 Language ..... 13
2.2.1 Header Types and Header Instances ..... 14
2.2.2 Metadata ..... 15
2.2.3 Parsers ..... 15
2.2.4 Tables and Actions ..... 16
2.2.5 Control ..... 18
2.2.6 Deparser ..... 19
2.2.7 Externs ..... 19
2.2.8 $\quad \mathrm{P}_{4}$ Language Versions ..... 19
2.3 Chapter Summary ..... 20
3 Common Header Validity Bugs ..... 23
3.1 Parser Bugs ..... 23
3.2 Control Bugs ..... 25
3.3 Table Reads Bugs ..... 26
3.4 Table Action Bugs ..... 27
3.5 Default Action Bugs ..... 29
3.6 Chapter Summary ..... 29
II Typed Data Plane Programming ..... 31
4 A Typing Discipline to Ensure Header Validity ..... 33
4.1 Design ..... 34
4.2 Syntax ..... 37
4.3 Static Semantics ..... 39
4.3.1 Operations on header types ..... 40
4.3.2 Typing rules ..... 42
4.4 Dynamic Semantics ..... 45
4.5 Safety ..... 47
4.6 Related Work ..... 49
4.7 Chapter Summary ..... 50
5 Dependently-Typed Data Plane Programming ..... 53
5.1 An Overview of $\Pi_{4}$ ..... 54
5.2 Design ..... 55
5.3 Syntax ..... 57
5.4 Well-formedness ..... 58
5.5 Dynamic Semantics ..... 59
5.6 Static Semantics ..... 62
5.7 Chomp ..... 68
5.7.1 Single-bit Chomp ..... 68
5.7.2 Instance Refinement ..... 71
5.7.3 Correctness of Chomp ..... 71
5.8 Safety ..... 71
5.9 Related Work ..... 73
5.10 Chapter Summary ..... 74
6 An Implementation of $\Pi_{4}$ ..... 75
6.1 Algorithmic Typing Rules ..... 75
6.2 Decidability ..... 76
6.3 SMT Encoding ..... 79
6.4 Optimizations ..... 84
6.4.1 Optimizing the SMT Encoding ..... 84
6.4.2 Reducing the Number of SMT Solver Invocations ..... 91
6.5 P4 Frontend ..... 93
6.6 Chapter Summary ..... 94
III Evaluation ..... 97
7 Header Validity Bugs in Real-world Programs ..... 99
7.1 Detecting and Repairing Bugs ..... 100
7.2 Overhead ..... 103
7.3 Chapter Summary ..... 104
8 Expressivity of $\Pi_{4}$ ..... 105
8.1 Survey ..... 105
8.2 Checking Network Invariants ..... 107
8.2.1 Protocol conformance ..... 108
8.2.2 Determined Forwarding ..... 110
8.2.3 Parser-Deparser Compatibility ..... 111
8.2.4 Mutual Exclusion of Headers ..... 111
8.3 Designing for Modularity ..... 114
8.3.1 Specifying Invariants ..... 115
8.3.2 Checking Customer Programs ..... 115
8.4 Chapter Summary ..... 118
9 Performance Evaluation ..... 121
9.1 Checking Header Validity ..... 122
9.2 Effects of Optimizations on Runtime ..... 123
9.3 Effects of the MTU on Runtime ..... 124
9.4 Modular Verification ..... 126
9.5 Chapter Summary ..... 127
IV Epilogue ..... 129
10 Conclusion and Future Work ..... 131
Bibliography ..... 135
A Proofs ..... 145
A. 1 SafeP4 ..... 145
A.1.1 Operations on Header Types ..... 145
A.1.2 Safety ..... 153
А. $2 \Pi_{4} \ldots \ldots$ ..... 164
A.2.1 Safety ..... 164
A.2.2 Algorithmic Typing Correctness ..... 205
A.2.3 Decidability of Typechecking ..... 218
A.2.4 Type Equivalences ..... 226

## Part I

## Prologue

## Introduction

For more than a quarter of a century, the idea has existed to overcome the limitations imposed by the static nature of networks by providing programmability inside the network [Smi+96; FRZ13]. Until today, with a major focus on interoperability, new developments such as new network protocols usually have to undergo years of standardization [Ten+97; WGT98; TWo7; Calo6] before being adopted by hardware vendors and being usable in practical deployments, which overall slows down innovation in the field of networking. The functionality of most network devices is strongly linked to the underlying hardware and the available functionality is solely dictated by hardware vendors. Implementations are usually proprietary, and network administrators have only limited configuration mechanisms available to adapt the functionality of the network device to their needs. As such, there has long been a desire to implement new network services or tailor packet processing functionality to the needs of applications running on top of the network, but the original approaches have not caught on.

For a little over a decade, the way how we can configure networks has changed significantly, and today this idea is finally becoming a reality. There is a shift towards more flexible platforms, which allow to specify the behavior of the network in software. These platforms are based on the idea of breaking the tight coupling between deciding where to send packets and actually forwarding packets [Yan+04; Lak+04], which results in a logically centralized control plane and a separate data plane.

Early efforts related to software-defined networking (SDN) focused on the control plane software [Cas+07; $\mathrm{McK}+08$ ]. These approaches made it possible to write programs that determine how routes through the network are computed, load is balanced or how security policies are enforced. The data plane was modeled as a simple pipeline that operates on a fixed set of packet formats. In practice, it has been found that a fixed number of available packet headers that can be accessed is not sufficient [Bos+14]. To overcome the limitations imposed by a fixed data plane, there is a recent interest in allowing the functionality of the data plane itself to be specified as a program. The goal is to provide flexible mechanisms that allow arbitrary packet headers to be extracted and processed. This opens up a whole new range of applications since it allows network programmers to implement new network protocols, make more efficient use of hardware resources or even relocate application-level functionality into the network [Jin+17; Jin+18].

In particular, the $\mathrm{P}_{4}$ language is becoming the de facto standard for programming data planes. It enables the functionality of the data plane to be programmed in terms of declarative abstractions such as header types, packet parsers, match-action tables, and structured control flow that a compiler maps down to an underlying target device.

### 1.1 Problem Statement

Today, computer networks play a more important role than ever before, since they provide the communication fabric for nearly all modern software systems. Unfortunately networks are still programmed using low-level languages that lack basic safety guarantees. Unsurprisingly this results in networks being unreliable and remarkably insecure-e.g., the first step in a cyberattack often involves compromising a router or other network device [KG16; OCo+18].

While a number of P 4 's features were clearly inspired by designs found in modern languages, the central abstraction for representing packet data-header types-lacks basic safety guarantees. Header types are used to describe the structure of packet headers, which-as a first approximation-can be thought of as a record, with one entry for each header field. For example, the header type for an IPv4 packet would have a 4-bit version field, an 8-bit time-to-live field, two 32 -bit fields for the source and destination addresses, and so on. According to the $\mathrm{P}_{4}$ language specification [ P 416 ], an instance of a header type may either be valid or invalid. If the instance is valid, then all operations reading or writing the header instance produce a defined value, but if it is invalid, then operations yield undefined results. In practice, programs that manipulate invalid headers can exhibit a variety of faults including dropping the packet when it should be forwarded or even leaking information from one packet to the next. In addition, such programs are also not portable, since their behavior can vary when executed on different targets.

The choice to model the semantics of header types in an unsafe way was intended to make the language easier to implement on high-speed routers, which have often limited amounts of memory. A typical $\mathrm{P}_{4}$ program might specify the behavior for several dozen different protocols, but any particular packet is likely to contain only a small handful of headers. Consequently, if the compiler only needs to represent the valid headers at run-time, then memory requirements can be reduced. However, while it may have benefits for language implementers, the design is a disaster for programmers. It repeats Hoare's "mistake" [Hoao9] and bakes an unsafe feature deep into the design of a language that has the potential to become the de facto standard in a multi-billion-dollar industry.

As programmable data planes become more prevalent and more complex applications are implemented inside the network, the more the risk of bugs increases. In the past, various verification tools have been developed to statically detect errors resulting from insecure access of header instances. We believe that basic safety guarantees should be part of the programming language instead. Modern programming languages offer features such as type systems, structured control flow, objects, modules, etc. that make it possible to express rich computations in terms of high-level abstractions rather than machine-level code. Increasingly, many languages also offer fundamental safety guarantees-e.g., well-typed programs do not go wrong [Mil78]-that make entire categories of programming errors simply impossible.

Type systems are a lightweight and compositional way to establish program properties. Types for individual components document assumptions about the components they rely upon as well as the guarantees they offer. However, the use of type systems in the realm of network programming has barely been investigated until now. For simple
properties, such as handing the access to uninitialized memory locations, no additional program annotations are necessary in contrast to existing verification tools.

The compositionality inherent to type systems, enables modular verification of programs. Existing verification tools are monolithic and do not provide any support for modular verification of programs, which was not a problem in the past, since $\mathrm{P}_{4}$ programs were also mostly monolithic and mostly still are. Although pre-processors are used to separate the program code for individual sections of the packet processing pipeline, $\mathrm{P}_{4}$ programs still contain a complete description of the processing pipeline. Increasingly, however, attempts are being made to make data plane programs modular [Gao+20; Son+20], with the aim of creating libraries for programming data planes from which data plane programs can be composed and that functionality can be extended in a modular fashion, such that programmers can describe the intended behavior in terms of high-level abstractions instead of low-level, platform-specific language constructs.

In summary, the $\mathrm{P}_{4}$ language is used to program critical infrastructure, although it lacks basic safety guarantees. Thus, with the steady proliferation of programmable data planes, the risk of network devices exhibiting unexpected errors is increasing. Future versions of $\mathrm{P}_{4}$ or new languages for programming data planes should therefore be equipped with suitable type systems that provide the programmer with necessary safety guarantees already at compile time.

### 1.2 State of the Art

In order to provide a better understanding of the proposed solution, we first provide an overview of the current state of network programming and network program verification.

## Network Programming Languages

With OpenFlow [McK+o8], the idea of a network operating system [Gud+o8] emerged that provides abstractions for the resources of the underlying network in a manner similar to conventional operating systems providing access to system resources. The goal was to provide programmers with higher abstractions for configuring networks so that network policies no longer had to be realized by configuring individual devices using low-level configuration mechanisms. On this basis, a multitude of programming languages emerged, which raised the level of abstraction even further.

Various declarative programming paradigms such as logic programming [Hin+o9], functional reactive programming [VH11; VKF12] or tierless programming [Nel+14] were used, with the goal of avoiding errors through complex interactions between packet-handling rules, for example, when composing network policies [Mon+13], for stateful packet processing [Ara+16], consistent network policy updates [McC+16] or to enable formal reasoning [And+14; Kim+15].

In contrast, data planes are currently programmed primarily with low-level languages [Bos+14; Bro19]. Both $\mu \mathrm{P} 4$ [Son +20 ] and Lyra [Gao +20 ] aim to provide highlevel abstractions for the data plane, similar to the high-level languages for the control plane, to make code portable between architectures and to enable composition. However, neither approach is suitable for specifying and verifying correctness properties for data planes.

## Verification of Data Plane Programs

Until now, only dedicated verification tools have been used for the verification of data planes. Various verification tools have been developed in the recent years to statically check the correctness of data planes, using a number of different verification techniques. p4v [Liu+18] applies classical techniques based on predicate transformer semantics. Vera $[S t o+18]$ and Assert-P 4 [Fre $+18 ; \mathrm{Nev}+18]$ are symbolic execution engines for P 4 . p4pktgen [Nöt+18] uses symbolic execution to automatically generate test cases for $\mathrm{P}_{4}$ programs. The bf4 tool [Dum+2o] follows the approach pioneered in p 4 v , but also attempts to infer control-plane constraints that are sufficiently strong to establish correctness, and offers heuristics for repairing programs when verification fails. P4K [KR18] provides a formal semantics of $\mathrm{P}_{4}$ in the K framework [RȘ10] and thus, can make use of the verification tools provided by the $K$ framework. P4AIG [Nou+19] statically verifies programmable data planes using sequential circuits-a hardware verification technique. In contrast, $\mathrm{P}_{4}$ RL [Shu+19] uses a dynamic approach-fuzz testing-for the verification of P 4 programs.

Having to resort to a separate verification tool increases the entry barrier for verification compared to when verification mechanisms are part of the programming language itself, as it is the case with type systems. In addition, none of the previously mentioned verification tools allows to modularly reason about data planes.

## Type Systems for Networks

Surprisingly, the use of type systems in the context of network programming has hardly been investigated so far, although the use of strongly typed languages for programming packet processing systems was already investigated in the 1990s in the context of Active Networks [Hic +98 ; Ale +98 ], with the aim of providing the programmer with static correctness guarantees for programs executed inside the network.

PacLang [ESMo4] is an imperative, concurrent language for expressing packet processing applications for Network Processors [All+o3]. It uses a linear type system to ensure that a packet is never processed by multiple threads simultaneously and can only be processed by multiple threads if it was transferred between threads beforehand.

Muthukrishnan et al. proposed strongly typed networking [Mut+10], where packets carry additional type information that describes how the receiver of the data will interpret it. Network entities can then reject traffic based on this information if not enough context is provided to correctly carry out its functionality.

In a position paper [GGW15], Gaboardi et al. propose the use of a simple type system to ensure that rules installed by an SDN controller are compatible with the underlying match-action table. For other properties they recognize the need for refinement type systems or dependent type systems without providing a concrete solution. However, all the above systems use simple type systems that are not suitable for providing static correctness guarantees for data plane programs written in $\mathrm{P}_{4}$.

### 1.3 The Thesis in a Nutshell

For programming data planes, we lack programming languages that provide programmers even with basic safety guarantees. The development of a variety of verification tools clearly shows that there is a need to statically verify the correctness of data planes. The goal of this thesis is therefore to develop suitable type systems that are able to verify
the correctness of data planes, starting from basic security properties up to applicationspecific properties.

## Thesis Statement

Type systems are well suited to equip data plane programming languages with safety guarantees, which can be used to verify a wide range of safety properties.

To validate this thesis, we develop two type systems for $\mathrm{P}_{4}$, currently the most widely used programming language for data planes. Our first contribution is SAFEP4, a domain-specific language (DSL) that models the main abstractions of $\mathrm{P}_{4}$ and comes with a static type system that guarantees header validity. While header validity is a common cause of safety bugs, it is not sufficient to match the expressive power of fullfledged verification tools. As a secondary contribution, we therefore develop $\Pi_{4}$, a dependently-typed version of P 4 , which allows us to express and verify rich network properties while retaining the compositionality inherent to type systems.

### 1.3.1 SafeP4

To address header validity bugs, we design $\mathrm{SAFEP}_{4}$, a domain-specific language for programmable data planes, which has a static type system that guarantees that all headers read or written by the program are guaranteed to be valid. SAFEP4 models all the essential features of the $\mathrm{P}_{4}$ language, but prunes away unnecessary complexity, which results in a minimal calculus that is easy to reason about, but can still express numerous real-world data plane programs. The type system has an expressive algebra of so-called header types, by means of which it is possible to describe precisely which headers are valid on a certain program path. The type checker can automatically reject all programs that attempt to read or write headers that are not valid on all program paths leading to a certain program statement.

The challenge in designing the type system is that header validity is actually a dynamic property, because the set of valid headers can be modified at runtime. Data plane programs do not fully describe the functionality of the data plane. Instead, the functionality of the data plane is largely determined by the control plane, in the form of packet processing rules used to populate match action tables. Depending on the installed rules and the contents of the network packet that is processed, packet headers can thus be added or removed. To still be able to achieve static safety, SafeP4 uses a path-sensitive type system that incorporates information from conditional statements, forwarding tables and the control plane to precisely track header validity.

The type system was implemented in the form of P4CHECK, $^{2}$ a static analysis tool that allows $\mathrm{P}_{4}$ programs to be analyzed without having to re-implement them using the DSL provided by SAFEP4. Our evaluation shows that SAFEP4's type system is capable of discovering numerous unsafe header accesses in real-world $\mathrm{P}_{4}$ programs, both from academia and the industry.

### 1.3.2 $\Pi_{4}$

With the design of $\Pi_{4}$, we pursue two goals: on the one hand, we want to close the gap between $\mathrm{SAFEP}_{4}$ and full-fledged data plane verification tools, and on the other
hand, we want to use the compositionality inherent to type systems to enable modular verification of data planes.

Even though validity bugs are a common source of errors, there are a variety of other safety properties that SAFEP4 cannot express. The main reason is that SAFEP4's type system is not able to capture values of individual fields, which is an essential requirement for a multitude of safety properties, since network packets often rely on so-called type-length-value encoding where the first bits determine the type, length, and structure of subsequent bits.

For this reason, we resort to a more powerful typing discipline in ח4's type systemdependent types-which allow us to define types based on program values. $\Pi_{4}$ extends SAFEP4's header types to heap types, which not only capture which instances are valid, but also the shape of header instances and the incoming and outgoing network packet down to the bit-level. Overall $\Pi_{4}$ features a combination of dependent function types, dependent pairs, refinement types, union types and explicit substitutions, which enables precise typing in the presence of domain-specific features that combine packet serialization and de-serialization with imperative control flow.

In the design of $\Pi_{4}$ 's type system, we manage to balance the tradeoffs between expressiveness and decidability. By encoding types into the effectively propositional fragment of first-order logic over fixed-width bit vectors, we achieve automated subtyping and equivalence checks. This relieves the programmer from the burden of writing manual proofs, which is often the case for dependent type systems.

### 1.4 Contributions

The main contribution of this thesis is to demonstrate that type systems are well-suited to verify the correctness of data planes and that it is not necessary to sacrifice safety guarantees in favor of a simpler implementation. On the contrary, P4's constrained programming model lends itself perfectly to the use of expressive typing disciplines. As part of the design, implementation and evaluation of SAFEP4 and $\Pi_{4}$, this thesis makes the following individual contributions.

- We formalize SAFEP4, a core calculus that models the core features of $\mathrm{P}_{4}$ and is equipped with a type system based on header types-a limited form of regular types-that statically guarantees header validity and prove its type system sound.
- With the formalization of SafeP4's type system, we show how the dynamic behavior of the control plane can be approximated and thus the dynamic property of header validity can be made statically verifiable.
- We formalize $\Pi_{4}$, a dependently-typed version of the $\mathrm{P}_{4}$ language that combines dependent function types with heap types-a combination of refinement types, a limited form of regular types and explicit substitutions. We show how precise typing can be enabled in the presence of domain-specific language features that combine packet serialization and de-serialization with imperative control flow. Again, we prove soundness for $\Pi_{4}$ 's type system and in addition, we prove that type checking is decidable.
- With our chomp operator that computes the type that remains after extracting bits from a packet buffer, we demonstrate how we can compute derivatives of regular types in the presence of dependent types.
- We provide an implementation of $\Pi_{4}$, which provides automated subtyping and equivalence checks by encoding heap types into a decidable theory of first-order logic.
- We provide a classification of which of P4's language constructs are susceptible to validity bugs and exemplify how to avoid corresponding bugs.
- We evaluate $\mathrm{SafeP}_{4}$ by checking a set of open-source $\mathrm{P}_{4}$ programs, both from industry and academia, against SAFEP4's typing rules. Our results confirm that header validity bugs are a common source of errors, more than $70 \%$ of programs examined contained at least one such bug. Our evaluation further shows that fixing header validity bugs usually entails only low overhead for the programmer.
- We evaluate the expressiveness of $\Pi_{4}$ with a set of case studies, which demonstrate that $\Pi_{4}$ is capable of expressing and checking properties that were also addressed by other data plane verification tools. In an additional case study, we further demonstrate how $\Pi_{4}$ can be used to reason about modular data plane programs. Finally, we evaluate the runtime performance of $\Pi_{4}$ 's type checker on a set of open-source P4 programs.


### 1.5 List of Publications

The contributions of this thesis appeared previously in the following publications at peer-reviewed conferences. Parts of them are used verbatim.
[Eic+19] Matthias Eichholz et al. "How to Avoid Making a Billion-Dollar Mistake: Type-Safe Data Plane Programming with SafeP4". In: 33rd European Conference on Object-Oriented Programming (ECOOP 2019). Ed. by Alastair F. Donaldson. Vol. 134. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2019, 12:1-12:28. DOI: 10.4230/LIPIcs .ECOOP . 2019. 12.
[Eic+22] Matthias Eichholz et al. "Dependently-Typed Data Plane Programming". In: Proceedings of the ACM on Programming Languages 6.POPL (2022). Doi: $10.1145 / 3498701$.

I furthermore (co)authored the following peer-reviewed conference publications and workshop papers, which are not part of this thesis.
[Eic16] Matthias Eichholz. "Language Support for Verifiable SDNs". In: Companion Proceedings of the 2016 ACM SIGPLAN International Conference on Systems, Programming, Languages and Applications: Software for Humanity. SPLASH Companion 2016. Amsterdam, Netherlands: Association for Computing Machinery, 2016, pp. 9-11.
[ESM18] Matthias Eichholz, Guido Salvaneschi, and Mira Mezini. "Towards Safe Modular Composition of Network Functions". In: Conference Companion of the 2nd International Conference on Art, Science, and Engineering of Programming. Programming' 18 Companion. Nice, France: Association for Computing Machinery, 2018, pp. 81-86.
[Blö+19] Marcel Blöcher et al. "GRASS: Generic Reactive Application-Specific Scheduling". In: Proceedings of the 6th ACM SIGPLAN International Workshop on Reactive and Event-Based Languages and Systems. REBLS 2019. Athens, Greece: Association for Computing Machinery, 2019, pp. 21-30.
$\begin{array}{ll}\text { [Kru+19a] } & \begin{array}{l}\text { Johannes Krude et al. "Online Reprogrammable Multi Tenant Switches". } \\ \text { In: Proceedings of the 1st ACM CoNEXT Workshop on Emerging In-Network }\end{array} \\ & \begin{array}{l}\text { Computing Paradigms. ENCP '19. Orlando, FL, USA: Association for Com- } \\ \text { puting Machinery, 2019, pp. 1-8. }\end{array} \\ {[\text { Kru+19b] }} & \begin{array}{l}\text { Johannes Krude et al. "Optimizing Data Plane Programs for the Network". } \\ \text { In: Proceedings of the ACM SIGCOMM 2019 Workshop on Networking } \\ \text { and Programming Languages. NetPL'19. Beijing, China: Association for }\end{array} \\ \text { Computing Machinery, 2019, p. 1. }\end{array}$

### 1.6 Structure of the Thesis

Although the work presented in this thesis is my own, research is still ultimately a collaborative process and I have been fortunate to work with several co-authors, who have helped to shape and improve the texts, concepts and solutions presented. In the following I will state which parts of this thesis are based on which publication. In addition, I will mention which contributions were mainly made by my co-authors.

Chapter 2 provides background on software-defined networking and network programmability. We summarize the central language abstractions of $\mathrm{P}_{4}$ and show what differences exist between the two main language versions $\mathrm{P}_{14}$ and $\mathrm{P}_{416}$, which are mainly of syntactic nature. Chapter 3 classifies common bugs that arise in $\mathrm{P}_{4}$ programs caused by accessing invalid header instances and shows examples of how to fix these bugs. The chapter is based on [Eic+19] where the classification was contributed by Eric Hayden Campbell. Chapter 4 formalizes SAFEP4 and proves soundness for its type system. It is based on the work published at ECOOP [Eic+19]. Chapter 5 provides the formalization of $\Pi_{4}$, including the definition of our chomp operator and proves $\Pi_{4}$ 's type system to be safe and decidable. This chapter is based on the work published at POPL [Eic+22]. Chapter 6 describes the central insights underlying the implementation of $\Pi_{4}$, including the algorithmic typing, the SMT encoding and necessary optimizations. This chapter is in parts also based on the work published at POPL [Eic+22]. Chapter 7 provides evaluation results that show that bugs due to invalid header references arise in various industrial and academic data plane programs and is based on the work published at ECOOP [Eic+22]. The implementation of the prototype was done by Nate Foster, since he had access to a proprietary frontend for the $\mathrm{P}_{4}$ language and the evaluation was mostly performed by Eric Hayden Campbell. Chapter 8 is also based on [Eic+22] and describes several case studies that show that $\Pi_{4}$ is well-suited to express a wide range of network properties also addressed by verification efforts from the networking community. Chapter 9 evaluates $\Pi_{4}$ 's runtime performance when checking real $\mathrm{P}_{4}$ programs. Chapter 10 concludes this thesis and discusses future research directions.

## Background

This chapter introduces how programmability inside the network has emerged, leading to the current state of programmable data planes. We then show how the behavior of packet processing pipelines can be specified using the language abstractions provided by the $\mathrm{P}_{4}$ language.

### 2.1 Programmable Packet Processing

The functionality of network devices is based on a division of labor between two components, the control plane and the data plane. The control plane is responsible for deciding where to forward packets, for example by calculating routes, distributing load or enforcing security policies, while the data plane comprises specialized hardware capable of efficiently processing network packets at line rate. Until today, the majority of network devices exhibit a close tie between the control and data plane, as vendors control and closely coordinate both the hardware and software. This complicates a variety of network management tasks. For example, despite the distributed nature of network control software, the majority of these devices must be configured individually via vendor-specific configuration mechanisms-which usually differ between vendors and sometimes even between devices from the same vendor-making it difficult to debug network-wide configurations.

Moreover, the control software is usually limited to standardized network protocols to ensure interoperability between devices from different vendors. Customizing the behavior of network devices is only possible to a limited extent, but it is not readily possible to simply replace the software, for example, to implement new functions such as new network protocols. As a result, while network hardware continues to evolve, innovation at the control software level is slowed down. Before new developments are adopted by vendors, it usually takes several years for them to go through standardization processes. For example, the standardization of IPv6 started as early as 1998 [DH98] and was not completed until mid-2017 [DH17].

In order to accelerate innovations within the network, the idea of programmable networks therefore first emerged in the mid-1990s. With the growing success of the Internet, researchers wanted to address the limitation that implementing and testing new network protocols was not possible in real networks, but instead limited to small lab


Figure 2.1: Control plane and data plane in traditional networks (left) in comparison to a logically centralized control plane used in SDN (right).
deployments or simulations. Active Networks [BCZ97; Ale +98 ; TW07] therefore aimed at making packet processing on network nodes adaptable, such that new protocols could be implemented with little effort and packet processing could even be adapted to the requirements of individual applications. This was achieved by either providing network switches with a set of functions that a packet could execute at runtime as needed or by sending the code to be executed directly with the network packets. With the Packet Language for Active Networks (PLAN) [Hic+98], there were even efforts to establish safety guarantees as a fundamental part of the language by means of static and dynamic typing. For many the proposed approaches were too radical and as a result the vision of programmable networks did not catch on at the time.

The topic of programmable networks first gained considerable attention in the late 2000s with the advent of software-defined networks and, in particular, the development of OpenFlow [McK+o8]. OpenFlow managed to find a balance between the vision of programmable networks and practical deployability by building on existing hardware, which lead to a fast adoption both in academia and industry. As visualized in Figure 2.1, SDN is based on a fundamental change in network architecture and builds on two key ideas developed in previous years: the separation between the control and data plane and the consolidation of the previously distributed control plane into a logically centralized control plane. By removing the tight coupling between the control plane and the data plane, it became possible to program the control plane software independently. Henceforth, the control plane software runs on general-purpose machines while the data plane consists of dumb network switches that interact with the control plane via well-defined APIs such as OpenFlow. At the heart of an OpenFlow switch are matchaction tables, whose entries form a set of rules that determine how network packets are processed. Each rule consists of (1) a pattern that identifies a set of packets based on a fixed set of packet headers, (2) a list of actions that are applied to packets for which the pattern matches as well as (3) a priority for disambiguating rules and (4) counters for tracking the total size of packets processed using the respective rule. Depending on the set of installed rules, the switch can fulfill different roles, for example that of a router, a firewall or a load balancer. Via a dedicated communication channel between the switch and the controller, the controller can install new rules at runtime, but switches can also issue packets to be processed on the controller, for example, in case that none of the installed rules matches.

The development of OpenFlow was closely followed by the idea of a network operating system, which provides the programmer with a unified interface that abstracts over the resources of the network (e.g. switches), thus enabling network-wide policies


Figure 2.2: P4's abstract forwarding model.
and making low-level, device-specific configurations obsolete. The task of network operating system was taken over by a variety of controller platforms that were developed in the following period. Furthermore, a multitude of programming languages based on declarative programming paradigms such as logic programming [Hin+o9], functional reactive programming [VH11; VKF12] or tierless programming [Nel+14] emerged, which raised the level of abstraction even further. The goal of these languages was to avoid errors through complex interactions between packet-handling rules, for example, when composing network policies [Mon+13], for stateful packet processing [Ara+16], consistent network policy updates [ $\mathrm{McC}+16$ ] or to enable formal reasoning [And+14; Kim+15].

As OpenFlow became more widely adopted, it became apparent that matching on a set of predefined packet headers is not sufficient in practice to satisfy all application domains, which initially led to the standard being extended several times to provide new packet headers and eventually to efforts to make the data plane programmable as well. While packet processing is still based on match-action tables, a central component in these systems is a programmable parser, which allows defining new packet headers as well as the order in which these are read from the packet being processed. This opens up a variety of new applications, from new network protocols to more efficient use of hardware resources to offloading application logic into the network [Jin+17; Jin+18]. The $\mathrm{P}_{4}$ language $[\mathrm{Bos}+14]$ has emerged as the de facto standard for programming data planes.

### 2.2 The $P_{4}$ Language

$\mathrm{P}_{4}$ is a domain-specific language for specifying the behavior of network data planes. It provides declarative abstractions to describe how network devices process packets, i.e., arbitrary sequence of bits that can be divided into (1) a set of pre-determined headers that determine how the packet will be forwarded through the network, and (2) a payload that encodes application-level data. $\mathrm{P}_{4}$ is designed to be protocol-independent, which means that it handles packets with standard header formats (e.g., Ethernet, IP, TCP, etc.) as well as packets with custom header formats defined by the programmer.

By now, different kinds of devices can be programmed with P4, including PISA switches [Bos+13], FPGAs [Iba+19; Wan+17] or software devices, e.g., eBPF [Høi+18]. Because individual devices can differ significantly in their internal structure-for example, a programmable network interface card uses a different processing pipeline than a switch- $\mathrm{P}_{4}$ uses the concept of architectures to abstract over hardware details. A P4 architecture constitutes the programming model. It determines the programmable blocks and which functions are available to the programmer to interact with the hardware. The architecture is specified by the hardware vendor, which can either be a standard architecture such as the Portable Switch Architecture (PSA) [PSA16], or it can be a custom

```
header ethernet_t {
    bit<48> dstAddr;
    bit<48> srcAddr;
    bit<16> etherType;
}
```

Figure 2.3: Header type declarations in $\mathrm{P}_{16}$.

```
struct headers {
    ethernet_t ethernet;
    ethernet_t inner_ethernet;
    vlan_t[2] vlan;
}
```

Figure 2.4: Declaration of header instances in $\mathrm{P}_{416}$.
one. The PSA, for example, provides for a total of six programmable blocks, one parser, one match action stage and one deparser each for both the ingress and egress processing stages.

Simplified, $\mathrm{P}_{4}$ programs follow the abstract forwarding model shown in Figure 2.2, which generalizes how packets are processed in different forwarding devices. A P4 program first parses the headers in the input packet into a typed representation, which together with various metadata forms a global, per-packet state that is shared between the pipeline stages. Next, it uses a match-action pipeline to compute a transformation on those headers-e.g., modifying fields, adding headers, or removing them. Finally, a deparser serializes the headers back into a packet, which can be output to the next device. The rest of this section describes P4's typed representation, how the parsers, and deparsers convert between packets and this typed representation, and how control flows through the match-action pipeline.

### 2.2.1 Header Types and Header Instances

The packet headers on which a $\mathrm{P}_{4}$ program operates are referred to as header instances while header types specify the internal representation of packet data. For example, the code snippet shown in Figure 2.3 declares a header type (ethernet_t) for the Ethernet header with fields for the destination (dstAddr) and source ( $s r c A d d r$ ) addresses and the so-called EtherType (etherType). Each field in a header type declaration is annotated with a type. The type bit<W> indicates that the field is a bit vector of size $W$.

The code snippet in Figure 2.4 shows how headers are instantiated. Lines 2 and 3 define two header instances ethernet and inner_ethernet that share the same structure defined by header type ethernet_t. As shown, multiple header instances can be derived from the same header type, but since header instances in $\mathrm{P}_{4}$ are globally scoped, instance names must be unique. Besides header instance declarations, $\mathrm{P}_{4}$ also allows defining fixed-size arrays of headers with the same type, so-called header stacks. On Line 4 in Figure 2.4, a header stack vlan that can hold two instances of type vlan_t is defined. Both header stack instances can be referenced using the notation vlan [0] and $v \operatorname{lan}$ [1] respectively. They both are equivalent to a non-stack header instance.

In addition to fields that have been explicitly declared as part of a header type declaration, all header instances automatically include an additional Boolean validity field. Fields with Boolean type are implemented as bit vectors of size one, which is why we speak of the validity bit of a header instance. If the value of the field is true, we say that the header instance is valid, otherwise it is invalid. By default, header instances are invalid. The validity bit can be manipulated by the parser or by explicitly marking a header instance as valid or invalid. Each header stack instance has its own validity bit and adding or removing elements from the stack does only change the number of valid headers in the stack but not the number of headers.

### 2.2.2 Metadata

Metadata is per-packet state that is generated during the execution of a $\mathrm{P}_{4}$ program. Metadata behaves like header instances, i.e., individual fields can be read and written, but metadata has no validity bit and is therefore always valid. A distinction is made between user-defined metadata and intrinsic metadata. The structure of user-defined metadata is determined by the programmer and can be thought of as a set of temporary variables that can be modified by the program. In contrast, intrinsic metadata is provided by the architecture and holds information about the incoming packet such as the input port a packet arrives at. In addition, intrinsic metadata performs the task of a control register, i.e. by setting certain fields the programmer determines how the hardware processes the packet. For example, in the architecture used by $\mathrm{P}_{1_{14}}$, the programmer determines on which port the packet will be forwarded by setting the egress_spec field.

### 2.2.3 Parsers

A P4 parser specifies the order in which header instances are extracted from the input packet using a simple abstraction based on finite state machines. Figure 2.5 shows an implementation of a parser for three common headers Ethernet, VLAN and IPv4. Every parser contains at least one start state named start and two final states, accept and reject. The final state accept indicates successful parsing while the state reject indicates a parsing failure. For example, the parser in Figure 2.5 transitions to the acceptance state if none of the explicitly specified cases matches (Lines 7 and 14) or after parsing the IPv4 header (Line 19). When extracting into a header instance, bits are copied from the input packet into the header instance, which is as a result marked as valid. The select statement is P4's version of a switch statement and allows to transition to a different parser state based on the contents of a previously extracted header.

The parser first extracts the instance ethernet, optionally followed by a vlan instance, or an ipv4 instance, or both. Figure 2.6 depicts the headers accepted by the parser. If the EtherType field of the Ethernet header contains any value except oxo8oo or ox8100, no other headers except for Ethernet are extracted from the packet. If the EtherType is oxo8oo, the IPv4 header is also be extracted. Similarly, if the EtherType is ox8100, the VLAN header is extracted instead. In this case, the value of the EtherType field of the VLAN header determines whether the IPv4 header is additionally extracted or not. The corresponding parse graph is shown in Figure 2.7, where the final state on the right represents the accept state. When the parser reaches the accept state, processing transitions into the ingress pipeline, which begins the match-action processing.

```
parser Parser(packet_in packet, out headers hdr, ...) {
    state start {
        packet.extract(hdr.ethernet);
        transition select(hdr.ethernet.etherType) {
            0x0800: parse_ipv4;
            0x8100: parse_vlan;
            default: accept;
        }
    }
    state parse_vlan {
        packet.extract(hdr.vlan);
            transition select(hdr.vlan.etherType) {
            0x0800: parse_ipv4;
            default: accept;
        }
    }
    state parse_ipv4 {
        packet.extract(hdr.ipv4);
        transition accept;
    }
}
```

Figure 2.5: $\mathrm{P}_{4}$ code implementing a parser in $\mathrm{P}_{416}$.

| ethernet | $*$ |  |  |  |
| ---: | :---: | :---: | :---: | :---: |
| ethernet | $0 \times 0800$ | ipv4 |  |  |
| ethernet | $0 \times 0810$ | vlan | $*$ |  |
|  |  |  |  |  |
| ethernet | $0 x 0810$ | vlan | $0 \times 0800$ |  |
|  |  |  |  |  |

Figure 2.6: Headers accepted by the parsers in Figures 2.5 and 2.13 .

### 2.2.4 Tables and Actions

The bulk of the processing for each packet in a $\mathrm{P}_{4}$ program is performed using matchaction tables, a central data structure the match-action pipeline relies on. Match-action tables encode conditional processing, more specifically, the table first looks up the values being tested against a list of possible entries, and then executes a further piece of code depending on which entry (if any) matched. However, unlike standard conditionals, the entries in a match-action table are not known at compile-time. Rather, they are inserted and removed at run-time by the control plane, which may be logically centralized (as in a software-defined network), or it may operate as a distributed protocol (as in a conventional network).

As shown in Figure 2.8, a table is defined in terms of (1) the data it reads to determine a matching entry (if any) (2) the actions it may execute, and (3) an optional default action it executes if no matching entry is found. In $\mathrm{P}_{16}$ the data read by the table is specified


Figure 2.7: State machine of the parsers described by code in Figures 2.5 and 2.13.

```
table forward {
    key = {
        hdr.ipv4.isValid(): exact;
        hdr.vlan.isValid(): exact;
        hdr.ipv4.dstAddr: ternary;
    }
    actions = {
        nop;
        next_hop;
        remove;
    }
    default_action = nop;
}
```

Figure 2.8: Table declaration in $\mathrm{P}_{416}$. forward reads the validity of the ipv4 and vlan header instances and the dstAddr field of the ipv4 header instance, and calls one of its actions: nop, next_hop, or remove.
using the keyword key. Table forward shown in Figure 2.8 reads the validity of the ipv4 and vlan header instances and the dstAddr field of the ipv4 header instance.

A table also specifies the match-kind that describes how each header field should match with the patterns provided by the control plane. The $\mathrm{P}_{4_{16}}$ core library provides three predefined match-kinds, exact, ternary and $l p m$. An exact match requires the bits in the packet be exactly equivalent to the bits in the controller-installed pattern. A ternary match allows wildcards in arbitrary positions, so the controller-installed pattern $0 *$ would match bit sequences 00 and 01 . A longest prefix match (lpm) selects the table entry where the largest number of leading bits of the specified header field match those in the table entry. For example, let us assume that a table performs a longest prefix match on the $\mathrm{IPv}_{4}$ destination address and the table contains two entries with addresses 192.168.178.80/28 and 192.168.178.0/24. Both IP address ranges each contain the address 192.168.178.88, but due to the longer matching prefix, the table entry with the address 192.168.178.80/28 is selected for a packet containing the destination address 192.168.178.88.

The behavior of a table depends on the entries installed at run-time by the controlplane. Each table entry contains a match pattern, an action, and action data. Intuitively, the match pattern specifies the bits that should be used to match values, the action is the name of a pre-defined function (such as the ones in Figure 2.10), and the action data are the arguments to that function. Operationally, to process a packet, a table first

| Pattern |  |  | Action |  |
| :--- | :--- | :--- | :--- | :--- |
| ipv4 | vlan | ipv4.dstAddr | Name | Data |
| 1 | 0 | $10.0 .0 . *$ | next_hop | d,p |
| 0 | 1 | $*$ | remove |  |

Figure 2.9: Runtime contents of forward.

```
action next_hop(bit<48> dst, bit<9> port) {
    hdr.ethernet.srcAddr = hdr.ethernet.dstAddr;
    hdr.ethernet.dstAddr = dst;
    hdr.ipv4.ttl = hdr.ipv4.ttl - 1;
    standard_meta.egress_spec = port;
}
action remove() {
    hdr.ethernet.etherType = hdr.vlan.etherType;
    hdr.vlan.setInvalid();
}
```

Figure 2.10: $\mathrm{P}_{4}$ actions.
scans its entries to locate the first matching entry. If such a matching entry is found, the packet is said to "hit" in the table, and the associated action is executed. Otherwise, if no matching entry is found, the packet is said to "miss" in the table, and the default action (which is a no-op if unspecified) is executed.

Figure 2.9 provides an example of how the entries of the forward table populated by the control plane might look light. The first rule tests whether ipv4 is valid, vlan is invalid, and the first 24 bits of ipv4.dstAddr equal 10.0.0, and then applies next_hop with arguments $d$ and $p$, which stand for destination address and port. The second rule checks that ipv4 is invalid, then that vlan is valid, and skips evaluating the value of ipv4.dstAddr (since it is wildcarded), to finally apply the remove action.

Actions are functions containing sequences of primitive commands that perform operations such as adding and removing headers, assigning values to fields, etc. For example, Figure 2.10 depicts two actions next_hop and remove. The next_hop action first updates the Ethernet source with the current Ethernet destination, then updates the Ethernet destination with action data from the controller, then decrements the IPv4 TTL (time-to-live) field and finally sets the outgoing port, also with data provided by the controller. The remove action copies the data contained in the EtherType field from the VLAN header instance to the Ethernet header instance and invalidates the VLAN header.

### 2.2.5 Control

In $\mathrm{P}_{416}$, tables and actions are declared as part of control blocks. In addition, control blocks define in which order and under which conditions packet headers are transformed by match action tables. The body (apply) of a control block is an imperative program that uses standard control flow constructs to describe a pipeline of match action tables. The ingress control block usually begins to execute as soon as the parser completes. The

```
control Ingress(inout headers hdr, ...) {
    // table and declarations
    apply {
        if(hdr.ipv4.isValid() || hdr.vlan.isValid()) {
            forward.apply();
        }
    }
}
```

Figure 2.11: Ingress control block in $\mathrm{P}_{416}$.

```
control Deparser(packet_out packet, in headers hdr) {
    packet.emit(hdr.ethernet);
    packet.emit(hdr.vlan);
    packet.emit(hdr.ipv4);
}
```

Figure 2.12: $\mathrm{P}_{416}$ deparser.
ingress control shown in Figure 2.11 conditionally executes the table forward (Line 6) if one of IPv4 or VLAN is valid (Line 5).

### 2.2.6 Deparser

The deparser reassembles the final output packet after all processing has been done by serializing each valid header instance in some order. In $\mathrm{P}_{4}$, a control block is used to describe the deparsed headers and their order as shown in Figure 2.12. In the example, the deparser produces a packet with Ethernet, VLAN, and IPv4, in that order. Every emit statement contains an implicit validity check, so headers are only emitted if they are actually valid, otherwise emit behaves like a no-op instruction.

### 2.2.7 Externs

In addition to the core language, $\mathrm{P}_{4}$ provides so-called externs, a foreign-function interface into the hardware. Externs allow $\mathrm{P}_{4}$ programs to use functionality that cannot be expressed with $\mathrm{P}_{4}$ itself. $\mathrm{P}_{4}$ programs can interact with externs via well-defined APIs and change their internal state, but the internal behavior of an extern itself is fixed and dictated by the hardware. Which externs are available depends on the architecture of the hardware. Examples for externs are checksum units, or stateful elements such as counters, meters and registers that allow to save state between packets.

### 2.2.8 $\quad P_{4}$ Language Versions

There are two major versions of the $\mathrm{P}_{4}$ language, $\mathrm{P}_{414}$ and the newer revision $\mathrm{P}_{416}$. With $\mathrm{P}_{416}$, the language designers have attempted to reduce the complexity of the language into a core language that provides fewer distinguished language constructs. In

```
parser start {
    extract(ethernet);
    return select(latest.etherType) {
        0x0800: parse_ipv4;
        0x8100: parse_vlan;
        default: ingress;
    }
}
parser parse_vlan {
    extract(vlan);
    return select(latest.etherType) {
        0x0800: parse_ipv4;
        default: ingress;
    }
}
parser parse_ipv4 {
    extract(ipv4);
    return ingress;
}
```

Figure 2.13: P 4 code implementing a parser in $\mathrm{P}_{414}$.
explaining the central language abstractions, we so far concentrated on $\mathrm{P}_{416}$, however, since Chapters 3 and 4 refer to language version 14, we will now briefly point out the main syntactic differences.

Figure 2.13 shows the same parser as Figure 2.5, but this time implemented in $\mathrm{P}_{4_{14}}$. Parser states are specified using the keyword parser and parser states do not transition to a next state but rather return the next parser to be executed. Accordingly, in $\mathrm{P}_{414}$ the end of the parsing phase is not expressed by a transition to an acceptance state, but by the parser returning as the next state the name of a control block to be executed afterwards. For example, Lines 6, 13 and 18 in Figure 2.13 indicate that after the parsing phase, the ingress control should be executed.

The code in Figure 2.14 shows how the table from Figure 2.8 is declared in $\mathrm{P}_{4_{14}}$, which is overall analogous to $\mathrm{P}_{416}$. The major syntactic difference is that match keys are specified using the keyword reads. $\mathrm{P}_{14}$ also provides additional match-kinds such as valid for matching on the validity bit of header instances. Tables and actions are declared outside of control blocks, but otherwise control blocks are similar to $\mathrm{P}_{416}$ as shown in Figure 2.15.

In $\mathrm{P}_{14}$ the compiler automatically generated the deparser from the parser, and while the concept of externs was already present, stateful elements were part of the core language.

### 2.3 Chapter Summary

This chapter outlined how the idea of programmable networks started in the mid-1990s with Active Networks and how modern software-defined networks simplify network

```
table forward {
    reads {
        ipv4: valid;
        vlan: valid;
        ipv4.dstAddr: ternary;
    }
    actions {
        nop;
        next_hop;
        remove;
    }
    default_action : nop;
}
```

Figure 2.14: Table declaration in $\mathrm{P}_{1_{14}}$.

```
control ingress {
    if(valid(ipv4) or valid(vlan)) {
        apply(forward);
    }
}
```

Figure 2.15: Control block implemented in $\mathrm{P}_{414}$.
management by separating the network control plane from the data plane. While early efforts focused on the programmability of the control plane, more recently it is also possible to program the data plane, for which the $\mathrm{P}_{4}$ language has established itself as the de facto standard. The second part of the chapter provided a deeper insight into the $\mathrm{P}_{4}$ language. Starting from $\mathrm{P}_{4}$ 's abstract forwarding model consisting of the parser, the match action pipeline and the deparser, we showed how these can be programmed with the abstractions provided by $\mathrm{P}_{4}$.

## Common Header Validity Bugs

While a number of P 's features were inspired by designs found in modern languages, the central abstraction for representing packet data-header types-lacks basic safety guarantees. As discussed in the previous chapter, an instance of a header type may either be valid or invalid. According to the $\mathrm{P}_{4}$ language specification, if the instance is valid, then all operations produces a defined value, but if it is invalid, then reading or writing a field yields an undefined result.

In practice, this has serious consequences for $\mathrm{P}_{4}$ programmers, since they must be careful not to read or write invalid headers, because programs that manipulate invalid headers can exhibit a variety of faults including dropping the packet when it should be forwarded, or even leaking information from one packet to the next. In addition, such programs are also not portable, since their behavior can vary when executed on different targets.

In this chapter we present five categories of bugs found in open-source $\mathrm{P}_{4}$ programs that arise due to reading and writing invalid headers. To identify the bugs we surveyed a benchmark suite of 15 research and industrial $\mathrm{P}_{4}$ programs that are publicly available on GitHub and compile to the BMv2 [Net18] backend. We categorize bugs based on the following syntactic constructs: (1) parsers, (2) controls, (3) table reads, (4) table actions, and (5) default actions. For each bug, we also present a possible type-safe fix.

For this survey, we exclusively considered programs implemented in $\mathrm{P}_{4_{14}}$ because only few $\mathrm{P} 4_{16}$ programs were publicly available at that time, however, the same bug categories persist to exist in $\mathrm{P}_{416}$.

### 3.1 Parser Bugs

The first class of errors is caused by the parser being too conservative about dropping malformed packets, which increases the set of headers that may be invalid in the control pipeline. In most programs, the parser chooses which headers to extract based on the fields of previously-extracted headers using P4's select statement. Programmers often fail to handle packets falling through to the default case of these select statements.

Figure 3.1 illustrates this bug using an example from the codebase of the research tool NetHCF [Bai+18; Li+19]-a tool designed to combat TCP spoofing. After extracting the

```
/* Unsafe */
parser parse_ethernet {
    extract(ethernet);
    return select(latest.etherType) {
        0x0800: parse_ipv4;
        default: ingress;
    }
}
parser parse_ipv4 {
    extract(ipv4);
    return select(latest.protocol) {
        6: parse_tcp;
        default: ingress;
    }
}
parser parse_tcp {
    extract(tcp);
    return ingress;
}
```

```
control ingress {
    if(tcp.syn == 1 and ...) {...}
}
```

Figure 3.1: Parser bug example in NeтHCF: parser (top) and ingress control (bottom).

Ethernet and IPv4 headers, the parser handles TCP packets in parse_ipv4 (Line 12) based on the IPv4 protocol field and redirects all other packets to the ingress control, exemplary shown in the bottom of Figure 3.1. However, the ingress control does not check whether TCP is valid before accessing tcp. syn to check whether it is equal to 1 . This is unsafe since tcp is not guaranteed to be valid even though it is required to be valid in the ingress control.

A possible fix of this bug is to throw an error during parsing, to stop the processing of the current packet and drop it instead. In $\mathrm{P}_{414}$, this can be achieved by defining a parser exception (Line 2 in Figure 3.2) with a handler that drops packets (Line 3 in Figure 3.2). This handler protects the ingress control from having to handle unexpected packets as shown in Lines 10 and 18.

In $\mathrm{P}_{416}$ we have to resort to a different mechanism, since parser exceptions are no longer part of the language. Instead, we can use the verify statement, a simple form of error handling, as shown on Line 7 in Figure 3.3. If the boolean expression used as the first argument evaluates to false, the parser immediately transitions to the reject state, which causes parsing to terminate immediately and sets the error state of the parser to the error used as second argument (Unsupported). However, the actual runtime behavior of this implementation highly depends on the target it is executed on. It is up to the target's implementation whether a rejected packet is actually dropped or just an error flag is set, and therefore the code is not portable. In the latter case, to achieve an

```
/* Safe */
parser_exception unsupported {
    parser_drop;
}
parser parse_ethernet {
    extract(ethernet);
    return select(latest.etherType) {
        0x0800: parse_ipv4;
        default:
                parser_error unsupported;
    }
}
parser parse_ipv4 {
    extract(ipv4);
    return select(latest.protocol) {
        6: parse_tcp;
        default:
            parser_error unsupported;
    }
}
```

Figure 3.2: Fixing parser bugs in $\mathrm{P}_{14}$ using parser exceptions.
equivalent runtime behavior to the $\mathrm{P}_{414}$ fix, we could check for an error in the ingress and manually drop the packet.

In general, however, this fix might not be the best solution, because it alters the original behavior of the program. However, without knowing the programmer's intention, it is generally not possible to automatically repair a program with undefined behavior.

### 3.2 Control Bugs

Another common bug occurs when a table is executed in a context in which the instances referenced by that table are not guaranteed to be valid. This bug can be seen in the source code of NetCache [Jin18; Jin+17], a system that uses P4 to implement a load-balancing cache. The parser reserves a specific port (8888) to handle special-purpose traffic. If UDP traffic arrives at port 8888, the parser extracts the NetCache-specific header nc_hdr. Otherwise, it performs standard L2 and L3 routing.

As shown in the top-left of Figure 3.4, in the ingress control, the process_cache control block is executed, which itself is shown in the bottom of Figure 3.4. The latter defines and applies table check_cache_exist, which reads field nc_hdr. key as part of the match key. However, it is never checked that nc_hdr is actually valid. The invocation of the process_value table (not shown) contains another instance of the same bug. As shown in the top-right of Figure 3.4, to fix these bugs, we can wrap the calls to process_cache and process_value in a conditional that checks the validity of the header nc_hdr. This ensures that nc_hdr is valid when process_cache refers to it.

```
/* Safe */
error { Unsupported }
parser Parser(packet_in packet, out headers hdr, ...) {
    state start {
        packet.extract(hdr.ethernet);
            verify(hdr.ethernet.etherType == 0x0800, error.
                Unsupported)
            transition select(hdr.ethernet.etherType) {
            0x0800: parse_ipv4;
            default: accept;
        }
    }
    state parse_ipv4 {
        packet.extract(hdr.ipv4);
        verify(hdr.ipv4.protocol == 6, error.Unsupported)
        transition select(hdr.ipv4.protocol) {
            6: parse_tcp;
            default: accept;
        }
    }
}
```

Figure 3.3: Fixing parser bugs in $\mathrm{P}_{416}$ using the verify statement.

### 3.3 Table Reads Bugs

A similar bug arises in programs that contain tables that first match on the validity of certain header instances before matching on the fields of those instances. The advantage of this approach is that multiple types of packets can be processed in a single table, which saves memory. However, if implemented incorrectly, this programming pattern can lead to a bug, in which the match key reads bits from a header that may not be valid. An example of this bug is exhibited by switch.p4, a "realistic production switch" developed by Barefoot Networks, meant to be used "as-is, or as a starting point for more advanced switches" [Kod15].

Table port_vlan_mapping, shown in the top of Figure 3.5, shows an archetypal example of a table reads bug. This table is invoked in a context where it is unknown which of the VLAN tags is valid, despite containing references to both vlan_tag_ [0] and vlan_tag_[1] in the match key declaration. The references to header fields vlan_tag_[i].vid are guarded with keys that test the validity of vlan_tag_[i], for $i=0,1$. However, as written, it is impossible for the control plane to install a rule that will always avoid reading the value of an invalid header. The first match will check whether instance vlan_tag_ [0] is invalid, which is safe. However, the very next match will try to read the value of field vlan_tag_[0].vid, even when the instance is invalid. This attempt to access an invalid header results in undefined behavior, and

```
/* Unsafe */
/* Safe */
control ingress {
control ingress {
    if(valid(nc_hdr)) {
    process_cache();
        process_cache();
    process_value();
        process_value();
    }
    apply(ipv4_route);
    apply(ipv4_route);
}
}
```

```
table check_cache_exist {
    reads { nc_hdr.key : exact }
    actions { ... }
}
control process_cache {
    apply(check_cache_exist);
}
```

Figure 3.4: Control bug example in NetCache (top-left), type-safe fix (top-right) and common code (bottom).
is therefore a bug. It is worthy to note that this code is not actually buggy on some targets-in particular, on targets where invalid headers are initialized with o. However, o-initialization is not prescribed by the language specification, and therefore this code is not portable across targets.

The naive solution to fix this bug is to refactor the table into four different tables (one for each combination of validity bits) and then check the validity of each header before the tables are invoked. While this fix is perfectly safe, it can result in a combinatorial blowup in the number of tables, which is clearly undesirable both for efficiency reasons and because it requires modifying the control plane. Fortunately, rather than factoring the table into four tables, we can replace the exact match-kinds with ternary matchkinds as shown in the bottom of Figure 3.5, which permit matching with wildcards. In particular, the control plane can install rules that match invalid instances using an all-wildcard patterns, which is safe. In order for this solution to be an actual fix, we also need to assume that the control plane is well-behaved-i.e. that it will install wildcards for the ternary matches whenever the header is invalid.

### 3.4 Table Action Bugs

Another common bug arises when distinct actions in a table require different (and possible mutually exclusive) headers to be valid. This can lead to two problems: (1) the control plane can populate the table with unsafe match-action rules, and (2) there may be no validity checks that we can add to the control to make all the actions type-check.

Table fabric_ingress_dst_lkp shown in the top of Figure 3.6 provides an example of this bug ${ }^{1}$. It reads the value of header field fabric_hdr.dstDevice and

[^0]```
/* Unsafe */
table port_vlan_mapping {
    reads {
        vlan_tag_[0]: valid;
        vlan_tag_[0].vid: exact;
        vlan_tag_[1]: valid;
        vlan_tag_[1].vid: exact;
    }
}
/* Safe */
table port_vlan_mapping {
    reads {
        vlan_tag_[0]: valid;
        vlan_tag_[0].vid: ternary;
        vlan_tag_[1]: valid;
        vlan_tag_[1].vid: ternary;
    }
}
```

Figure 3.5: Table reads bug in switch.p4 (top) and type-safe fix (bottom).
then invokes one of several actions: (1) term_cpu_packet (2) term_fabric_unica st_packet, or (3) term_fabric_multicast_packet. These actions require that the headers (1) fabric_hdr_cpu, (2) fabric_hdr_unicast, and (3) fabric_hdr_ multicast respectively are valid. However, the validity of these headers is mutually exclusive.

Since all three headers are mutually exclusive, there is no single context that makes this table safe. The only facility the table provides to determine which action should be called is fabric_hdr.dstDevice. However, the P4 program doesn't establish a relationship between the value of fabric_hdr.dstDevice and the validity of any of these three header instances. So, the behavior of this table is only well-defined when the input packets are well-formed, an unreasonable expectation for real switches, which may receive any sequence of bits "on the wire."

We fix this bug by including validity matches in the match key, as shown in the bottom of Figure 3.6. Similar to the fix presented in Section 3.3, this solution avoids combinatorial blowup and extensive control plane refactoring. Again, we need to make an assumption about the way the control plane will populate the table. Concretely, if an action $a$ is only safe to execute if a header $h$ is valid, and $h$ is not necessarily valid when the table is applied, we assume that the control plane will only call $a$ if $h$ is matched as valid. For example, fabric_hdr_cpu is not known to be valid when (the fixed version of) fabric_ingress_dst_lkp is applied, so we assume that the control plane will only call action term_cpu_packet when fabric_hdr_cpu is matched as valid.

```
/* Unsafe */
table fabric_ingress_dst_lkp {
    reads {
        h.fabric_hdr.dstDevice: exact;
    }
    actions {
        term_cpu_packet;
        term_fabric_unicast_packet;
        term_fabric_multicast_packet;
    }
}
/* Safe */
table fabric_ingress_dst_lkp {
    reads {
        h.fabric_hdr.dstDevice: exact;
        h.fabric_hdr_cpu: valid;
        h.fabric_hdr_unicast: valid;
        h.fabric_hdr_multicast: valid;
    }
    actions {
        term_cpu_packet;
        term_fabric_unicast_packet;
        term_fabric_multicast_packet;
    }
}
```

Figure 3.6: Table action bug in switch.p4 (top) and type-safe fix (bottom).

### 3.5 Default Action Bugs

Default action bugs occur when the programmer incorrectly assumes that a table performs some action when a packet misses. The implementation of NetCache, which we introduced in Section 3.2, exhibits an example of this bug. The bug is shown in the top of Figure 3.7, where table add_value_header_1 is expected to make the nc_value_1 header valid, which is done in the add_value_header_1_act action. The control plane may refuse to add any rules to the table, which would cause all packets to miss, meaning that the add_value_header_1_act action would never be called and header nc_value_1 may not be valid. To fix this error, as shown in the bottom of Figure 3.7, we simply set the default action for the table to add_value_header_1_act, which will force the table to add the header no matter what rules the controller installs.

### 3.6 Chapter Summary

In this chapter, we presented five of P 4 's language constructs that are potentially vulnerable to validity bugs and what possible bug fixes for these bugs might look like. We

```
/* Unsafe */
table add_value_header_1 {
    actions {
        add_value_header_1_act;
    }
}
/* Safe */
table add_value_header_1 {
    actions {
        add_value_header_1_act;
    }
    default_action = add_value_header_1_act();
}
```

Figure 3.7: Default action bug in NetCache: unsafe code missing a default action (top) and type-safe fix (bottom).
do not claim that our taxonomy is complete and, moreover, the proposed bug fixes may change the behavior of the original program, but this cannot be avoided without knowing the programmer's intentions.

Header validity bugs can be caused by the parser when packages that should not be considered are handled inadequately. Additionally, they can also occur in various places related to match action tables. Invalid headers can be accessed if headers a table matches on are not guaranteed to be valid in the respective context, but also if in such a case, the programmer wants to ensure validity by means of an additional validity match, but which is not implemented correctly. Furthermore, bugs can occur when actions require that certain-possibly mutually exclusive-headers are valid without the program making a connection between the validity of these headers and the readsexpressions of the table defining the actions, or when it is incorrectly assumed that a header is made valid by an action without any guarantee that this action will ever be executed.

## Part II

## Typed Data Plane Programming

## A Typing Discipline to Ensure Header Validity

While the way headers are represented in $\mathrm{P}_{4}$ has advantages for language implementers, the design is a disaster for programmers. As we saw in the previous chapter, a variety of subtle bugs can creep in at various points in a $\mathrm{P}_{4}$ program when invalid headers are accessed, which makes it challenging for programmers to write correct data plane programs. The fact that accessing invalid headers returns undefined values is comparable to the existence of null references in various general-purpose programming languages. Computer scientist Tony Hoare once called his invention of the null reference a billiondollar mistake [Hoao9]. Computer networks provide the foundation for the majority of today's software systems. By embedding an insecure language feature deep into the design of a language that is about to become the standard in a billion-dollar industry, we are about to repeat Hoare's "mistake".

In this chapter, we look at how we can incorporate header safety guarantees into the language, thus enabling a correct-by-construction approach. We present SafeP4, a domain-specific language for programmable data planes in which all packet data is guaranteed to have a well-defined meaning and satisfy essential safety guarantees. We equip $S_{A F E P} 4$ with a static type system that statically guarantees header validity, which relies on an expressive algebra of so-called header types that tracks validity information at a fine level of granularity. One of the main challenges here is that $\mathrm{P}_{4}$ programs do not completely specify the behavior of data planes. Part of the behavior is determined by the match-action rules installed by the control plane at runtime. Since the available rules can be altered at runtime, and as a consequence the set of valid headers can dynamically change, header validity becomes a dynamic program property. As a result, SAFEP4's type system employs a form of path-sensitive reasoning that tracks dynamic information from conditional statements, routing tables and the control plane. We formalize the syntax and semantics of $\mathrm{SAFEP}_{4}$ in a core calculus and prove that the type system is sound.

### 4.1 Design

We made four key design decisions when designing SAFEP4. First, we represent only the core of $\mathrm{P}_{4}$, i.e., the features relevant to packet processing. Second, we simplify and generalize certain aspects to avoid unnecessary complications in the calculus. Third, we use a set of sets of headers to capture valid header instances per program path to avoid an overly restrictive type system and fourth, we assume that the control plane is well-behaved.

Core calculus In the design of SAFEP4, we draw inspiration from Featherweight Java [IPWo1], i.e., we model the essential features of $\mathrm{P}_{4}$, but prune away unnecessary complexity. The result is a minimal calculus that is easy to reason about, but can still express numerous real-world data plane programs. Our calculus is protocol-independent by allowing the programmer to specify the types of packet headers and their order in the bit stream. Also, SAFEP4 mimics P4's use of tables to interface with the control-plan and decide which actions to execute at run-time.

Simplification and generalization We omit a number of constructs that are secondary to how packets are processed, including parser exceptions, counters, meters, etc., however, it would be relatively straightforward to add these to the calculus. Also, compared to $\mathrm{P}_{4}, \mathrm{SAFEP}_{4}$ does not enforce a strict separation between the parsing phase and the control phase. Rather than unnecessarily complicating the syntax of SAFEP4, we allow the syntactic objects that represent parsers and controls to be freely mixed. Similarly, we only enforce informally which primitive commands can be invoked within actions, e.g., field modifications but not conditionals.

We also deviate from $\mathrm{P}_{4}$ with regard to the add command. In $\mathrm{P}_{416}$, the analogous operation setValid (respectively add_header in $\mathrm{P} 4_{14}$ ) only modifies the validity bit of a header instances. Accordingly, accessing an added header instance returns a non-deterministic value if the header instance's fields have not been manually initialized beforehand. Instead of complicating our type system by additionally capturing whether header fields are defined, we define the semantics of our $\operatorname{add}(h)$ command to initialize each field of header instance $h$ with a default value. We assume that in addition to our type constants there exists a function init that accepts a header type $\eta$ and returns a header instance of type $\eta$ with all fields set to their default value.

Last, we want to model the core behavior of both $\mathrm{P}_{4_{14}}$ and $\mathrm{P}_{416}$, however, both use different type systems and evaluation behaviors for expressions. We therefore abstract away expression typing and syntax variants by assuming that we are given a set of $n$ ary constants $k$ that can represent values like o or true (o-ary), or unary and binary operators such as negation or logical conjunction, etc. We further assume that these operators are assigned sound types. With these features in hand, one can instantiate our type system over arbitrary constants.

Granularity of types To analyze the validity of header instances, our type system needs a way to capture which header instances are valid. Naively, we can keep track of a set of headers, which are guaranteed to be valid on all program paths and reject all programs that reference headers not included in this set. However, this coarse-grained approach would generate a large number of false positives.

For example, consider the parser implementation in Figure 4.1. The parser extracts an Ethernet header and then either parses a VLAN header or proceeds to the ingress.

```
parser start {
    extract(ethernet);
    return select(latest.etherType) {
            0x8100 : parse_vlan;
            default: ingress;
    }
}
parser parse_vlan {
    extract(vlan);
    return ingress;
}
```

Figure 4.1: Parser extracting Ethernet and optionally VLAN.

```
control ingress {
    if(valid(vlan)) {
        modify_field(ethernet.etherType, vlan.etherType);
        remove_header(vlan);
    }
}
```

Figure 4.2: Ingress program

Hence, at the beginning of the ingress, only Ethernet is guaranteed to be valid. However, it is certainly safe to write an ingress program that references the VLAN header after checking that it is valid as shown in Figure 4.2.

To reflect this in the type system, we introduce a special construct valid $(h) c_{1}$ else $c_{2}$, which executes $c_{1}$ if $h$ is valid and $c_{2}$ otherwise. When we type check this command, following previous work on occurrence typing [TF1o], we check $c_{1}$ with the additional fact that $h$ is valid, and we check $c_{2}$ with the additional fact that $h$ is not valid. Despite this enhancement, our type system would still be overly restrictive. For example, the parser shown in Figure 4.3, first extracts Ethernet and then boots into ingress or extracts IPv4. If IPv4 is valid, the parser optionally extracts TCP or UDP.

Now, suppose that we have an ingress control that defines a table tcp_table that refers to both IPv4 and TCP in its key expression and that only checks the validity of TCP before applying the table as shown in Figure 4.4. In general, the validity of TCP also implies the validity of IPv4, so it should be safe to apply table tcp_table after checking only the validity of TCP. However, with the representation of valid headers as a set, the type checker would still reject the program, because at the point where tcp_table is applied (Line 11), only Ethernet and TCP would be guaranteed to be valid. Instead, we would need to explicitly check that both IPv4 and TCP are valid.

We solve the problem by using a more fine-grained type representation-namely a set of sets of headers instead of a set-to capture header instances guaranteed to be valid and their dependencies. Each inner set contains all headers that might be valid at the current program point for some program path. For a given header reference to be safe, it must be a member of all possible sets of headers, i.e., it must be valid on all paths

```
parser start {
    extract(ethernet);
    return select(latest.etherType) {
        0x0800: parse_ipv4;
        default: ingress;
    }
}
parser parse_ipv4 {
    extract(ipv4);
    return select(latest.protocol) {
        0x6: parse_tcp;
        0x11: parse_udp;
        default: ingress;
    }
}
parser parse_tcp { ... }
parser parse_udp { ... }
```

Figure 4.3: Parser extracting Ethernet, optionally followed by $\mathrm{IPv}_{4}$ and if $\mathrm{IPv}_{4}$ is valid, optionally followed by TCP or UDP.

```
table tcp_table {
    reads {
        ipv4.dstAddr: lpm;
        tcp.port: exact;
    }
    actions { ... }
}
control ingress {
    if(valid(tcp)) {
        apply(tcp_table);
    }
}
```

Figure 4.4: Ingress control that applies table tcp_table, which reads headers ipv4 and $t c p$.
through the program that reach the reference.

Control plane In the formalization of SAFEP4, we model the control plane as a function that-given a table and the currently valid headers-returns the action to call and the (possibly empty) action data arguments. Also, a second function analyzes the table and produces for each action a set of valid headers that can be safely assumed valid when the entries are populated by the control plane. From the table declaration
and the header instances that can be assumed valid, based on the match-kinds, we can derive a list of match key expression that must be evaluated when the table is invoked. Together, these functions model the run-time interface between the switch and the controller. We assume that the control plane interface satisfies three simple correctness properties: (1) the control plane can safely install table entries that never read invalid headers, (2) the action data provided by the control plane has the types expected by the action, and (3) the control plane will only assume valid headers for an action that are valid for a given packet.

### 4.2 Syntax

The syntax of SAFEP4 is shown in Figure 4.5. To lighten the notation, we write $\bar{x}$ as shorthand for a (possibly empty) sequence $x_{1}, \ldots, x_{n}$. A SAFEP4 program consists of a sequence of declarations $\bar{d}$ and a command $c$. Declarations include tables, header types and header instances.

A table declaration $t(\bar{h}, \overline{(e, m)}, \bar{a})$ is defined in terms of a sequence of valid-match header instances $\bar{h}$, a sequence of match-key expressions $\overline{(e, m)}$, where $e$ is an expression and $m$ is the match-kind ${ }^{1}$ used to match this expression and a sequence of actions $\bar{a}$. The notation $t$.valids denotes the valid-match instances, t.reads denotes the expressions, and t.actions denotes the actions. For example, the following table declaration

$$
\begin{aligned}
& \text { forward }\left((i p v 4, \text { vlan }),\left(\left(i p v_{4} . d s t A d d r, \text { ternary }\right)\right),( \right. \\
& \text { skip, } \\
& \lambda s \text {, d.ethernet.srcAddr }=s ; \text { ethernet.dstAddr }=d \text {, } \\
& \text { ethernet.etherType }=\text { vlan.etherType; remove }(\text { vlan }) \\
& \text { )) }
\end{aligned}
$$

corresponds to the following P4 table, assuming action nop is-as the name suggests-a no-op instruction, action next_hop sets the Ethernet source and destination addresses to values provided by the control plane and action remove populates the EtherType field of the Ethernet header with the value contained in EtherType field of the VLAN header and ultimately removes the VLAN header.

```
table forward {
    reads {
        ipv4 : valid;
        vlan : valid;
        ipv4.dstAddr: ternary;
    }
    actions {
        nop;
        next_hop;
        remove;
    }
}
```

[^1]
## Commands

| $c::=$ |  |
| :---: | :---: |
| \| extract(h) | EXtraction |
| \| emit(h) | DEPARSING |
| $c_{1} ; c_{2}$ | SEQUENCE* |
| if $(e) c_{1}$ else $c_{2}$ | CONDITIONAL |
| $\operatorname{valid}(h) c_{1}$ else $c_{2}$ | validity |
| t.apply () | application |
| skip | SKIP |
| $\operatorname{add}(h)$ | ADDITION* |
| remove( $h$ ) | Removal* |
| $h . f=e$ | MODIFICATION* |
| Actions |  |
| $a::=\lambda \bar{x} . c$ | ACTION |
| Expressions |  |
| $e::=$ |  |
| $v$ | values |
| $h . f$ | HEADER FIELD |
| $x$ | variable |
| $k^{n}$ | CONSTANT |

## Declarations


$\mid t(\bar{h} \overline{(e, m)}, \bar{a})$ TABLE
$m \in\{$ exact, ternary $\} \quad k^{n} \in K$
Program
Values
$v \in V$

## Header Types

$\Theta$ ::=

Action Types Expression Types
$\alpha::=\bar{\tau} \rightarrow \Theta \quad \tau::=$ Bool
$\left\lvert\, \begin{aligned} & \bar{\tau} \rightarrow \tau \\ & \ldots\end{aligned}\right.$

Figure 4.5: Syntax of SAFEP4

Actions are written as (uncurried) $\lambda$-abstractions. An action $\lambda \bar{x} . c$ declares a (possibly empty) sequence of parameters, drawn from a fresh set of names, which are in scope for the command $c$. The run-time arguments for actions (action data) are provided by the control plane. We artificially restrict the commands that can be called in the body of the action to addition, removal of headers, modification of header fields and sequence. These commands are identified with an asterisk in Figure 4.5.

Header type declarations describe the format of individual headers and are defined in terms of a name and a sequence of field declarations. The notation " $f: \tau$ " indicates that field $f$ has type $\tau$. We let $\eta$ range over header types. A header instance declaration assigns a name $h$ to a header type $\eta$. The map $\mathcal{H} \mathcal{T}$ encodes the (global) mapping between header instances and header types.

The calculus provides commands for extracting (extract), creating (add), removing (remove), and modifying ( $h . f=e$ ) header instances. The emit command is used in the deparser and serializes a header instance back into a bit sequence. The if-statement conditionally executes one of two commands based on the value of a boolean condition. Similarly, the valid-statement branches on the validity of $h$. Table application commands (t.apply()) are used to invoke a table $t$ in the current state. The skip command is a no-op.

The only built-in expressions in SAFEP4 are variables $x$ and header fields, written $h . f$. We let $v$ range over values and assume a collection of $n$-ary constant operators $k^{n} \in K$. For simplicity, we assume that every header referenced in an expression has a corresponding instance declaration. We also assume that header instance names $h$, header type names $\eta$, variable names $x$, and table names $t$ are drawn from disjoint sets

$$
\begin{aligned}
\llbracket \Theta \rrbracket & \subseteq \mathcal{P}(\text { Header }) \\
\llbracket 0 \rrbracket & =\{ \} \\
\llbracket 1 \rrbracket & =\{\{ \}\} \\
\llbracket h \rrbracket & =\{\{h\}\} \\
\llbracket \Theta_{1} \cdot \Theta_{2} \rrbracket & =\llbracket \Theta_{1} \rrbracket \bullet \llbracket \Theta_{2} \rrbracket \\
\llbracket \Theta_{1}+\Theta_{2} \rrbracket & =\llbracket \Theta_{1} \rrbracket \cup \llbracket \Theta_{2} \rrbracket
\end{aligned}
$$

Figure 4.6: Semantics of header types.
of names $\mathrm{H}, \mathrm{E}, \mathrm{V}$, and T respectively and that each name is declared only once.
$S_{A F E P} 4$ provides three kinds of types, header types $\Theta$, expression types $\tau$ and action types $\alpha$. We assume that the set of expression types includes booleans (for conditionals) as well as tuples and function types.

### 4.3 Static Semantics

SafeP4 uses a path-sensitive analysis, coupled with occurrence typing to keep track of which headers are guaranteed to be available at any program point. The type system rejects programs that reference headers that might be uninitialized, thus, preventing all references to invalid headers.

Semantics of types A header type $\Theta$ represents a set of header instances that may be valid at the same time. The type o denotes the empty set. This type arises when there are unsatisfiable assumptions about which headers are valid. The type 1 denotes the empty singleton set of headers. It describes the initial state of the program where no headers are valid. The type $h$ denotes a singleton set, $\{\{h\}\}$, i.e., states where only $h$ is valid. The type $\Theta_{1} \cdot \Theta_{2}$ denotes the set obtained by combining headers from $\Theta_{1}$ and $\Theta_{2}$, i.e., a product or concatenation $\left(H_{1} \bullet H_{2}=\left\{h_{1} \cup h_{2} \mid h_{1} \in H_{1} \wedge h_{2} \in H_{2}\right\}\right)$. Finally, the type $\Theta_{1}+\Theta_{2}$ denotes the union of $\Theta_{1}$ or $\Theta_{2}$, which intuitively represents an alternative.

The semantics of header types, $\llbracket \Theta \rrbracket$, is defined by the equations in Figure 4.6. Intuitively, each subset represents one alternative set of headers that may be valid. For example, the header type eth $\cdot(\operatorname{ipv} 4+1)$ denotes the set $\{\{e t h, i p v 4\},\{e t h\}\}$.

Typing judgment We use different typing judgments for command typing, expression typing and action typing. The typing judgment for commands has the form $\Gamma \vdash c$ : $\Theta_{1} \mapsto \Theta_{2}$, which means that in variable context $\Gamma$, if $c$ is executed in the header context $\Theta_{1}$, then a header instance type $\Theta_{2}$ is assigned. Intuitively, $\Theta_{1}$ encodes the sets of headers that may be valid when type checking a command. $\Gamma$ is a standard type environment which maps variables $x$ to type $\tau$. If there exists $\Theta_{2}$ such that $\Gamma \vdash c: \Theta_{1} \mapsto \Theta_{2}$, we say that $c$ is well-typed in $\Theta_{1}$.

The typing judgment for expressions and actions has the form $\Gamma ; \Theta \vdash e: \tau$ and $\Gamma ; \Theta \vdash a: \bar{\tau} \rightarrow \Theta$ respectively, meaning that expression $e$ has type $\tau$ respectively that action $a$ has type $\bar{\tau} \rightarrow \Theta$ in variable context $\Gamma$ and header context $\Theta$.

### 4.3.1 Operations on header types

To formulate the typing rules for $\mathrm{SAFEP}_{4}$, we first define a set of operations on header types, namely restriction, negated restriction, inclusion, removal and emptiness. In the following we assume that $S$ ranges over elements of the domain $\mathcal{P}(\mathcal{P}$ (Headers)).

Restriction The operator Restrict $\Theta h$ recursively traverses $\Theta$ and keeps only those choices in which $h$ is contained, zeroing out the others. Semantically this has the effect of throwing out the subsets of $\llbracket \Theta \rrbracket$ that do not contain $h$, i.e., we define restriction semantically as $\left.S\right|_{h} \triangleq\{h s \mid h s \in S \wedge h \in h s\}$. Syntactically we define restriction by induction on $\Theta$ as follows:

```
        Restrictoh\triangleqo
        Restrict 1 h 气 o
```



```
    Restrict ( }\mp@subsup{\Theta}{1}{}\cdot\mp@subsup{\Theta}{2}{})h\triangleq((\mathrm{ Restrict }\mp@subsup{\Theta}{1}{}h)\cdot\mp@subsup{\Theta}{2}{})+(\mp@subsup{\Theta}{1}{}\cdot(\mathrm{ Restrict }\mp@subsup{\Theta}{2}{}h)
Restrict ( }\mp@subsup{\Theta}{1}{}+\mp@subsup{\Theta}{2}{})h\triangleq(\mathrm{ Restrict }\mp@subsup{\Theta}{1}{}h)+(\mathrm{ Restrict }\mp@subsup{\Theta}{2}{}h
```

Lemma 4.1 captures the equivalence of the syntactic and the semantic definition.
Lemma 4.1 (Restrict Equal). $\left.\llbracket \Theta \rrbracket\right|_{h}=\llbracket$ Restrict $\Theta h \rrbracket$

Proof. By induction on $\Theta$.

Negated Restriction Dually to the restrict operator, NegRestrict $\Theta h$ produces only those subsets where $h$ is invalid. Semantically, negated restriction is defined as $\left.S\right|_{\neg h} \triangleq\{h s \mid h s \in S \wedge h \notin h s\}$. Syntactically we define negated restriction by induction on $\Theta$ :

$$
\begin{aligned}
& \text { NegRestricto } h \triangleq \text { o } \\
& \text { NegRestrict } 1 h \stackrel{ }{\vartheta} 1 \\
& \text { NegRestrict } g h \triangleq \begin{cases}0 & \text { if } g=h \\
g & \text { otherwise }\end{cases} \\
& \text { NegRestrict }\left(\Theta_{1} \cdot \Theta_{2}\right) h \triangleq\left(\text { NegRestrict } \Theta_{1} h\right) \cdot\left(\text { NegRestrict } \Theta_{2} h\right) \\
& \text { NegRestrict }\left(\Theta_{1}+\Theta_{2}\right) h \triangleq\left(\text { NegRestrict } \Theta_{1} h\right)+\left(\text { NegRestrict } \Theta_{2} h\right)
\end{aligned}
$$

Lemma 4.2 captures the equivalence of the syntactic and the semantic definition.
Lemma 4.2. (NegRestrict Equal) $\left.\llbracket \Theta \rrbracket\right|_{\neg h}=\llbracket$ NegRestrict $\Theta h \rrbracket$

Proof. By induction on $\Theta$.

Inclusion Includes $\Theta h$ traverses $\Theta$ and checks if $h$ is valid in every path. Semantically this says that $h$ is a member of every element of $\llbracket \Theta \rrbracket$, i.e., $h \sqsubset S \triangleq \forall h s \in S . h \in h s$. Syntactically we define inclusion by induction on $\Theta$ :

$$
\begin{gathered}
\text { Includes } 0 h \triangleq \text { true } \\
\text { Includes } 1 h \triangleq \text { false } \\
\text { Includes } g h \triangleq \begin{cases}\text { true } & \text { if } g=h \\
\text { false } & \text { otherwise }\end{cases} \\
\text { Includes }\left(\Theta_{1} \cdot \Theta_{2}\right) h \triangleq\left(\text { Includes } \Theta_{1} h\right) \vee\left(\text { Includes } \Theta_{2} h\right) \\
\text { Includes }\left(\Theta_{1}+\Theta_{2}\right) h \triangleq\left(\text { Includes } \Theta_{1} h\right) \wedge\left(\text { Includes } \Theta_{2} h\right)
\end{gathered}
$$

Lemma A. 3 captures the equivalence of the syntactic and the semantic definition.
Lemma 4.3. (Includes Equal) $\forall h s \in \llbracket \Theta \rrbracket . h \in h s=$ Includes $\Theta h$
Proof. By induction on $\Theta$.
Removal Remove $\Theta h$ removes $h$ from every path, which means, semantically that it removes $h$ from every element of $\llbracket \Theta \rrbracket$, i.e., $S \backslash h \triangleq\{h s \mid h s \in S \wedge h s \backslash\{h\}\}$. Syntactically we define removal by induction on $\Theta$ :

$$
\left.\begin{array}{rl}
\text { Remove o } h \triangleq 0 \\
\text { Remove } 1 h \triangleq 1
\end{array}\right] \begin{array}{ll}
\text { Remove } g h \triangleq \begin{cases}1 & \text { if } g=h \\
g & \text { otherwise }\end{cases} \\
\text { Remove }\left(\Theta_{1} \cdot \Theta_{2}\right) h \triangleq\left(\text { Remove } \Theta_{1} h\right) \cdot\left(\text { Remove } \Theta_{2} h\right) \\
\text { Remove }\left(\Theta_{1}+\Theta_{2}\right) h \triangleq\left(\operatorname{Remove} \Theta_{1} h\right)+\left(\text { Remove } \Theta_{2} h\right)
\end{array}
$$

Lemma A. 4 captures the equivalence of the syntactic and the semantic definition.
Lemma 4.4. (Remove Equal) $\llbracket \Theta \rrbracket \backslash h=\llbracket$ Remove $\Theta h \rrbracket$
Proof. By induction on $\Theta$.
Emptiness Empty $\Theta$ checks if $\Theta$ is semantically empty. Syntactically we define emptiness by induction on $\Theta$ :

$$
\begin{aligned}
& \text { Empty } \mathrm{o} \triangleq \text { true } \\
& \text { Empty } 1 \triangleq \text { false } \\
& \text { Empty } h \triangleq \text { false } \\
& \text { Empty }\left(\Theta_{1} \cdot \Theta_{2}\right) \triangleq \text { Empty } \Theta_{1} \vee \operatorname{Empty} \Theta_{2} \\
& \text { Empty }\left(\Theta_{1}+\Theta_{2}\right) \triangleq \text { Empty } \Theta_{1} \wedge \text { Empty } \Theta_{2}
\end{aligned}
$$

The equivalence of the syntactic and semantic definition is captured by Lemma A.5.

$$
\begin{aligned}
\mathcal{F}\left(h, f_{i}\right) & =\tau_{i} \\
\mathcal{A}(a) & =\lambda \bar{x}: \bar{\tau} . c \\
\mathcal{H}(e) & =\bar{h} \\
\mathcal{C} \mathcal{A}(t, H) & =\left(a_{i}, \bar{v}\right) \\
\mathcal{C V}(t) & =\bar{S}
\end{aligned}
$$

Field type lookup
Action lookup
Referenced Header instances
Control-plane actions
Control-plane validity

$$
\begin{aligned}
\text { maskable }(t, e, \text { exact }) & \triangleq \text { false } \\
\text { maskable }(t, e, \text { ternary }) & \triangleq \mathcal{H}(e) \subseteq t . \text { valids }
\end{aligned}
$$

Figure 4.7: Auxiliary functions.

Lemma 4.5. $\llbracket \Theta \rrbracket==\{ \}$ if and only if Empty $\Theta$.
Proof. By induction on $\Theta$.

### 4.3.2 Typing rules

SAFEP4's typing rules rely on several auxiliary functions shown in Figure 4.7. The field type lookup function $\mathcal{F}\left(h, f_{i}\right)$ returns the type assigned to a field $f_{i}$ in header instance $h$ by first looking up the corresponding header type $\eta$ from the global header table and then looking up the field type from the header type declaration. The action lookup function $\mathcal{A}(a)$ returns the action definition $\lambda \bar{x}: \bar{\tau}$. $c$ for action $a$ and $\mathcal{H}(e)$ returns the header instances $\bar{h}$ referenced by expression $e$. Given a table $t$ and the currently valid headers $H$, the function $\mathcal{C} \mathcal{A}(t, H)$ returns the run-time action to call $a_{i}$ and the (possibly empty) action data arguments $\bar{v}$. The function $\mathcal{C} \mathcal{V}(t)$ analyzes table $t$ and produces a list of sets of valid headers $\bar{S}$-one set for each action-that can be safely assumed valid when the entries are populated by the control plane. Both, $\mathcal{C A}(t, H)$ and $\mathcal{C} \mathcal{V}(t)$ are assumed to be instantiated by the control plane.

The function maskable $(t, e, m)$ takes in a table $t$, a match key expression $e$ and a match-kind $m$ and checks whether expression $e$ must be evaluated when table $t$ is invoked. If maskable evaluates to true, it means that the expression does not need to be evaluated. If the match-kind is exact, $e$ always needs to be evaluated and if the match-kind is ternary, $e$ only needs to be evaluated if it references at least one header instance that is not part of the valid-match header instances (t.valids).

Command typing The typing rules for commands are presented in Figure 4.8. The rule T-Zero gives a command an arbitrary output type if the input type is empty. The rules T-Skip and T-Seq are standard. The rule T-If a path-sensitive union type between the type computed for each branch. The rule T-If VaLid is similar, but leverages knowledge about the validity of $h$. So the true branch $c_{1}$ is checked in the context Restrict $\Theta h$, and the false branch $c_{2}$ is checked in the context NegRestrict $\Theta h$. The top-level output type is the union of the resulting output types for $c_{1}$ and $c_{2}$. The rule T-Mod checks that $h$ is guaranteed to be valid using the Includes operator, and uses the auxiliary function $\mathcal{F}$ to obtain the type assigned to $h$.f. The set of valid headers does not change when evaluating an assignment, so the output and input types are identical.

T-Mod
Includes $\Theta h$
T-Extr

$$
\frac{\mathcal{F}(h, f)=\tau_{i} \quad \Gamma ; \Theta \vdash e: \tau_{i}}{\Gamma \vdash h \cdot f=e: \Theta \mapsto \Theta}
$$

$$
\overline{\Gamma \vdash \operatorname{extract}(h): \Theta \mapsto \Theta \cdot h}
$$

T-Add

$$
\overline{\Gamma \vdash \operatorname{add}(h): \Theta \Leftrightarrow \Theta \cdot h}
$$

T-Rem

$$
\overline{\Gamma \vdash \operatorname{remove}(h): \Theta \mapsto \operatorname{Remove} \Theta h}
$$

T-Apply

$$
\begin{gathered}
\mathcal{C V}(t)=\bar{S} \quad \text { t.actions }=\bar{a} \\
\bar{e}=\left\{e_{j} \mid\left(e_{j}, m_{j}\right) \in \text { t.reads } \wedge \neg \text { maskable }\left(t, e_{j}, m_{j}\right)\right\}
\end{gathered}
$$

$$
\text { Т-Еміт } \quad ; \Theta \vdash e_{j}: \tau_{j} \text { for } e_{j} \in \bar{e}
$$

$$
\cdot ; \text { Restrict } \Theta S_{i} \vdash a_{i}: \bar{\tau}_{i} \rightarrow \Theta_{i}^{\prime} \text { for } a_{i} \in \bar{a}
$$

$$
\overline{\Gamma \vdash e m i t}(h): \Theta \Leftrightarrow \Theta
$$

$$
\Gamma \vdash t . \operatorname{apply}(): \Theta \mapsto\left(\sum_{a_{i} \in \bar{a}} \Theta_{i}^{\prime}\right)
$$

Figure 4.8: Command typing rules for SAFEP4.

The rules T-Extr and T-Add assign header extractions and header additions the type $\Theta \cdot h$, reflecting the fact that $h$ is valid after the command executes. Emitting packet headers does not change the set of valid headers, which is captured by rule T-Emit. The typing rule T-Rem uses the Remove operator to remove $h$ from the input type $\Theta$.

Finally, the rule T-Apply checks table applications. To understand how it works, let us first consider a simpler, but less precise, typing rule:

$$
\begin{array}{cc}
\text { t.reads }=\bar{e} & ; \Theta \vdash e_{i}: \tau_{i} \text { for } e_{i} \in \bar{e} \\
\text { t.actions }=\bar{a} & ; \Theta \vdash a_{i}: \overline{\tau_{i}} \rightarrow \Theta_{i}^{\prime} \text { for } a_{i} \in \bar{a} \\
\cdot \vdash t . \operatorname{apply}(): \Theta \mapsto\left(\sum \Theta_{i}^{\prime}\right)
\end{array}
$$

Intuitively, this rule says that to type check a table application, we check each expression it reads and each of its actions. The final header type is the union of the types computed for the actions. This rule models table application as a non-deterministic choice between its actions. However, while this rule is sound, it is overly conservative. In particular, it does not model the fact that the control plane often uses header validity bits to control which actions are executed.

Hence, the actual typing rule, T-Apply, makes use of function $\mathcal{C V}(t)$ to first obtain for each action $a_{i}$, a set of headers $S_{i}$ that can be assumed valid when type checking

$$
\begin{aligned}
& \begin{array}{ll}
\begin{array}{l}
\text { T-Zero } \\
\text { Empty } \Theta_{1}
\end{array} & \text { T-SKIP } \\
\Gamma \vdash c: \Theta_{1} \Leftrightarrow \Theta_{2}
\end{array} \quad \begin{array}{l}
\text { T-SEQ } \\
\Gamma \vdash s k i p: \Theta \Leftrightarrow \Theta
\end{array} \frac{\begin{array}{l}
\Gamma \vdash c_{1}: \Theta \mapsto \Theta_{1} \quad \Gamma \vdash c_{2}: \Theta_{1} \Leftrightarrow \Theta_{2} \\
\Gamma \vdash c_{1} ; c_{2}: \Theta \mapsto \Theta_{2}
\end{array}}{l} \\
& \begin{array}{l}
\text { T-IF } \quad \Gamma ; \Theta \vdash e: \text { Bool } \\
\begin{array}{l}
\Gamma \vdash c_{1}: \Theta \Leftrightarrow \Theta_{1} \quad \Gamma \vdash c_{2}: \Theta \mapsto \Theta_{2} \\
\Gamma \vdash \text { if }(e) c_{1} \text { else } c_{2}: \Theta \mapsto \Theta_{1}+\Theta_{2}
\end{array}
\end{array} \\
& \text { T-If Valid } \\
& \Gamma \vdash c_{1} \text { : Restrict } \Theta h \Leftrightarrow \Theta_{1} \\
& \Gamma \vdash c_{2} \text { : NegRestrict } \Theta h \Leftrightarrow \Theta_{2} \\
& \overline{\Gamma \vdash \operatorname{valid}(h) c_{1} \text { else } c_{2}: \Theta \mapsto \Theta_{1}+\Theta_{2}}
\end{aligned}
$$

```
table port_vlan_mapping {
    reads {
        vlan[0] : valid;
        vlan[0].vid : ternary;
        vlan[1] : valid;
        vlan[1].vid : ternary;
    } ...
}
```

Figure 4.9: Expressions vlan[0].vid and vlan[1].vid must be wildcarded when vlan [0] respectively vlan [1] are invalid. The typechecker can warn the programmer about these assumptions for the table to be safe.


Figure 4.10: Expression typing rules for $\mathrm{SAFEP}_{4}$.
$a_{i}$. From the match key expressions of the table declaration and the headers assumed valid, we can derive a subset of the expressions read by the table, e.g., excluding expressions that can be wildcarded when certain validity bits are false using the function maskable $(t, e, m)$. Consequently, we only need to type check these expressions. For example, given the table in Figure 4.9, if an action $a_{j}$ is matched by the rule ( $0, *, 0, *$ ), both $S_{j}$ and $e_{j}$ are empty. Just like the simplified rule, the final header type is the union of the types computed for the actions, however, we additionally restrict the header context with the header instances assumed to be valid. Here we lift the Restrict operator to sets of header instances:

$$
\begin{aligned}
\text { Restrict } \Theta\} \triangleq \Theta \\
\text { Restrict } \left.\Theta\left\{h_{1}, \ldots, h_{n}\right\} \triangleq \text { Restrict (Restrict (Restrict } \Theta h_{1}\right) \ldots \text { ) } h_{n}
\end{aligned}
$$

Expression typing The typing rules for expressions are shown in Figure 4.10. Constants are type-checked according to rule T-Constant, as long as each expression that is passed as an argument to the constant $k$ has the type required by the typeof function. The rule T-Var is standard. Rule T-Field checks that header instance $h$ is guaranteed to be valid and assigns the type obtained from the field type lookup function $\mathcal{F}$.

Action typing Given a variable context $\Gamma$ and header type $\Theta$, an action $\lambda \bar{x}: \bar{\tau} . c$ encodes a function of type $\bar{\tau} \rightarrow \Theta^{\prime}$, so long as the body $c$ is well-typed in the context where $\Gamma$ is extended with $x_{i}: \tau_{i}$ for every $i$.

$$
\frac{\Gamma, \bar{x}: \bar{\tau} \vdash c: \Theta \Leftrightarrow \Theta^{\prime}}{\Gamma ; \Theta \vdash \lambda \bar{x}: \bar{\tau} . c: \bar{\tau} \rightarrow \Theta^{\prime}} \quad \text { (T-Action) }
$$

Figure 4.11: Action typing rule for SAFEP4.

### 4.4 Dynamic Semantics

We define the operational semantics for commands in terms of four-tuples $\langle I, O, H, c\rangle$, where $I$ is the input bit stream (which is assumed to be infinite for simplicity), $O$ is the output bit stream, $H$ is a map that associates each valid header instance with a records containing the values of each field, and $c$ is the command to be evaluated.

For the definition of the operational semantics, we assume that for each declared header type $\eta$ there exists a deserialization function ( deserialize $_{\eta}$ ), a serialization function ( serialize $_{\eta}$ ) and an initialization function ( init $_{\eta}$ ). The function $\operatorname{deserialize~}_{\eta}(I)=$ ( $v, I^{\prime}$ ) takes in the input bit stream $I$ and returns a header value $v$ populated with bits from $I$ as well as the rest of the input bit stream. For example, assuming the header type $\eta=\{f: \operatorname{bit}\langle 3\rangle ; g: \operatorname{bit}\langle 2\rangle ;\}$ has two fields $f$ and $g$ and $I=11000 B$ where $B$ is the rest of the bit stream following, then deserialize $_{\eta}(I)=(\{f=110 ; g=o 0 ;\}, B)$. The serialize function is the corresponding counterpart. It takes the bit values of the fields of a header value and concatenates them in the order in which they are defined to produce a single bit sequence. For example, calling serialize on the header value $\{f=110 ; g=00 ;\}$ returns the bit string 11000 . Finally, the function init $_{\eta}$ returns a header instance of type $\eta$ with all fields set to a default value.

Semantics of commands The reduction rules are presented in Figure 4.12 and Figure 4.13. The command $\operatorname{extract}(h)$ evaluates via the rule E-Extr, which looks up the header type in the global header table $\mathcal{H T}$ and then invokes the corresponding deserialization function. The deserialized header value $v$ is added to the map of valid header instances, $H$ and evaluation continues with the remaining input bit stream $I^{\prime}$. The rules for validity checks step to the true branch if $h \in \operatorname{dom}(H)$ (E-If ValidTrue) and to the false branch otherwise (E-If ValidFalse). The rule E-Rem removes the header from the map $H$. If a header $h$ is already invalid, removing it has no effect. Modification of header fields is evaluated according to rules E-Mod and E-Modr. If the assigned expression is fully reduced, the respective field in the header value is updated with the new value in $H$. Otherwise, the assigned expression is reduced first. Table application commands are evaluated according to rule E-Apply. We first invoke the control plane function $\mathcal{C} \mathcal{A}(t, H)$ to determine an action $a_{i}$ and action data $v$. Then we use $\mathcal{A}$ to look up the definition of $a_{i}$, yielding $\lambda \bar{x}: \bar{\tau} . c_{i}$ and step to $c_{i}[\bar{v} / \bar{x}]$. Note that for simplicity, we model the evaluation of expressions read by the table using the control-plane function $\mathcal{C A}$. The rule E-ADD evaluates addition commands $\operatorname{add}(h)$. Similar to header extraction, we first obtain the header type $\eta$ of the instance $h$ and then use the function init $_{\eta}$ function to obtain a header instance $v$ of type $\eta$ with all fields set to a default value and extend the map $H$ with $h \mapsto v$. Note that according to E-AddValid, if the header instance is already valid, $\operatorname{add}(h)$ does nothing. The rule E-Emit serializes a header instance $h$ back into a bit stream. It first looks up the corresponding header type and header value in the header table $\mathcal{H} \mathcal{T}$ and the map of valid headers respectively. The header value is then passed to the serialization function for the header type to produce a bit sequence

$$
\begin{aligned}
& \text { E-Extr } \\
& \frac{\mathcal{H} \mathcal{T}(h)=\eta \quad \text { deserialize }_{\eta}(I)=\left(v, I^{\prime}\right)}{\langle I, O, H, \operatorname{extract}(h)\rangle \rightarrow\left\langle I^{\prime}, O, H[h \mapsto v], \text { skip }\right\rangle} \\
& \text { E-If ValidTrue } \\
& \frac{h \in \operatorname{dom}(H)}{\left\langle I, O, H, \operatorname{valid}(h) c_{1} \text { else } c_{2}\right\rangle \rightarrow\left\langle I, O, H, c_{1}\right\rangle} \\
& \text { E-If ValidFalse } \\
& \frac{h \notin \operatorname{dom}(H)}{\left\langle I, O, H, \operatorname{valid}(h) c_{1} \text { else } c_{2}\right\rangle \rightarrow\left\langle I, O, H, c_{2}\right\rangle} \\
& \text { E-Rem } \\
& \overline{\langle I, O, H, \text { remove }(h)\rangle \rightarrow\langle I, O, H \backslash h, s k i p\rangle} \\
& \text { E-Mod } \\
& \begin{array}{c}
H(h)=r \quad r^{\prime}=\{r \text { with } f=v\} \\
\langle I, O, H, h . f=v\rangle \rightarrow\left\langle I, O, H\left[h \mapsto r^{\prime}\right], \text { skip }\right\rangle
\end{array}
\end{aligned}
$$

| E-Modi $\quad\langle H, e\rangle \rightarrow e^{\prime}$ | E-Apply $\mathcal{C} \mathcal{A}(t, H)=\left(a_{i}, \bar{v}\right) \quad \mathcal{A}\left(a_{i}\right)=\lambda \bar{x} \cdot c_{i}$ |
| :---: | :---: |
| $\overline{\langle I, O, H, h . f=e\rangle \rightarrow\left\langle I, O, H, h . f=e^{\prime}\right\rangle}$ | $\overline{\langle I, O, H, t . a p p l y}()\rangle \rightarrow\left\langle I, O, H, c_{i}[\bar{v} / \bar{x}]\right\rangle$ |
| E-Add | E-AddValid |
| $\mathcal{H} \mathcal{T}(h)=\eta \quad$ init $_{\eta}=v$ | $h \in \operatorname{dom}(H)$ |
| $\overline{\langle I, O, H, \operatorname{add}(h)\rangle \rightarrow\langle I, O, H[h \mapsto v], s k i p\rangle}$ | $\overline{\langle I, O, H, \operatorname{add}(h)\rangle \rightarrow\langle I, O, H, s k i p\rangle}$ |
| E-Emit | E-Emitinvalid |
| $\mathcal{H} \mathcal{T}(h)=\eta \quad$ serialize $_{\eta}(H(h))=\bar{B}$ | $h \notin \operatorname{dom}(H)$ |
| $\overline{\langle I, O, H, e m i t}(h)\rangle \rightarrow\langle I, O . \bar{B}, H$, skip $\rangle$ | $\overline{\langle I, O, H, e m i t}(h)\rangle \rightarrow\langle I, O, H$, skip $\rangle$ |

Figure 4.12: Operational semantics of commands.
that is appended to the output bit stream. We adopt the semantics of $\mathrm{P}_{4}$ with respect to emitting invalid headers. Emitting an invalid header instance-i.e., a header instance which has not been added or extracted-has no effect on the output bit stream (rule E-Emitinvalid). Notice also that the header remains unchanged in $H$.

Sequential composition (cf. Figure 4.13) reduces left to right, i.e., the left command needs to be reduced to skip (rule E-SEQ1) before the right command can be reduced (rule E-Seq). The evaluation of conditionals (rules E-If, E-If True, E-If False) is standard. The expression in the condition is first evaluated to a value. If the condition is true, the conditional steps to command $c_{1}$ otherwise to command $c_{2}$.

Semantics of expressions The semantics for expressions is defined in Figure 4.14. It is defined in terms of tuples $\langle H, e\rangle$, where $H$ is the same map used in the semantics of commands and $e$ is the expression to evaluate. The rule E-Field reduces header field expressions to the value stored in the heap $H$ for the respective field. To evaluate

$$
\begin{aligned}
& \begin{array}{ll}
\text { E-SEQ } & \begin{array}{l}
\text { E-SEQ1 } \\
\\
\hline\left\langle I, O, H, s k i p ; c_{2}\right\rangle \rightarrow\left\langle I, O, H, c_{2}\right\rangle
\end{array}
\end{array} \begin{array}{l}
\left\langle I, O, H, c_{1}\right\rangle \rightarrow\left\langle I^{\prime}, O^{\prime}, H^{\prime}, c_{1}^{\prime}\right\rangle
\end{array} \\
& \text { E-If } \\
& \frac{\langle H, e\rangle \rightarrow e^{\prime}}{\left\langle I, O, H, \text { if }(e) \text { then } c_{1} \text { else } c_{2}\right\rangle \rightarrow\left\langle I, O, H, \text { if }\left(e^{\prime}\right) \text { then } c_{1} \text { else } c_{2}\right\rangle} \\
& \text { E-If True } \\
& \overline{\left\langle I, O, H, \text { if (true) then } c_{1} \text { else } c_{2}\right\rangle \rightarrow\left\langle I, O, H, c_{1}\right\rangle} \\
& \text { E-If FALSE } \\
& \overline{\left\langle I, O, H, \text { if }(\text { false }) \text { then } c_{1} \text { else } c_{2}\right\rangle \rightarrow\left\langle I, O, H, c_{2}\right\rangle}
\end{aligned}
$$

Figure 4.13: Operational semantics of commands (continued).

$$
\begin{aligned}
& \begin{array}{l}
\text { E-FIELD } \\
\begin{array}{c}
H(h)=\left\{f_{1}: n_{1}, \ldots, f_{k}: n_{k}\right\} \\
\left\langle H, h . f_{i}\right\rangle \rightarrow n_{i}
\end{array}
\end{array} \begin{array}{c}
\text { E-CONST } \\
\left\lfloor k \rrbracket\left(v_{1}, \ldots, v_{n}\right)=v\right.
\end{array} \\
& \begin{array}{l}
\text { E-CONST-CONG } \\
\frac{\left\langle H, k\left(v_{1}, \ldots, v_{n}\right)\right\rangle \rightarrow v}{\left\langle H, k\left(v_{1}, \ldots, v_{i-1}, e_{i}, \ldots, e_{n}\right)\right\rangle \rightarrow k\left(v_{1}, \ldots, v_{i-1}, e_{i}^{\prime}, \ldots, e_{n}\right)}
\end{array}
\end{aligned}
$$

Figure 4.14: Operational semantics for expressions.
constants via the rule E-Const, we assume that there is an evaluation function for constants $\llbracket k \rrbracket(\bar{v})=v$ that is well-behaved-i.e., if typeof $(k)=\bar{\tau} \rightarrow \tau^{\prime}$ and $\overline{v: \tau}$, then $\because \cdot \vdash \llbracket k \rrbracket(\bar{v}): \tau^{\prime}$. Arguments passed to constants are evaluated left to right (rule E-Const-Cong).

### 4.5 Safety

We prove safety in terms of progress and preservation [WF94]. Both theorems make use of the relation $H \vDash \Theta$ which intuitively holds if $H$ is described by $\Theta$. The formal definition, as given in Figure 4.15, satisfies $H \vDash \Theta$ if and only if $\operatorname{dom}(H) \in \llbracket \Theta \rrbracket$.

The empty header instance map only entails the empty header instance type 1 (Rule Ent-Емpty). If a header instance $h$ is contained in the map of valid header instances $H, H$ entails the header instance type $h$ (Rule Ent-Inst). The sequence type $\Theta_{1} \cdot \Theta_{2}$ is entailed by the distinct union of the maps entailing $\Theta_{1}$ and $\Theta_{2}$ respectively (Rule Ent-SEQ) and the choice type $\Theta_{1}+\Theta_{2}$ is entailed either by the map entailing $\Theta_{1}$ or the map entailing $\Theta_{2}$ (Rules Ent-ChoiceL and Ent-ChoiceR).

We prove progress and preservation only for commands. For expressions, we formulate these properties as additional lemmas (Lemmas 4.6 and 4.7). The respective proofs

$$
\begin{aligned}
& \begin{array}{lll}
\text { Ent-Empty } & \begin{array}{l}
\text { Ent-Inst } \\
\operatorname{dom}(H)=\{h\} \\
\cdot \vDash 1
\end{array} & \begin{array}{l}
\text { Ent-Seq }
\end{array} \\
& \frac{H_{1} \vDash \Theta_{1} \quad H_{2} \vDash \Theta_{2}}{H_{1} \cup H_{2} \vDash \Theta_{1} \cdot \Theta_{2}}
\end{array} \\
& \begin{array}{lc}
\begin{array}{c}
\text { Ent-ChoiceL } \\
H \vDash \Theta_{1}
\end{array} & \begin{array}{c}
\text { Ent-ChoiceR } \\
H \vDash \Theta_{2} \\
H \vDash \Theta_{1}+\Theta_{2}
\end{array}
\end{array}
\end{aligned}
$$

Figure 4.15: The Entailment relation between header instances and header instance types
are straightforward for our system.
Lemma 4.6 (Expression Progress). If $\cdot ; \Theta \vdash e: \tau$ and $H \vDash \Theta$, then either $e$ is a value or $\exists e^{\prime} .\langle H, e\rangle \rightarrow e^{\prime}$.

Lemma 4.7 (Expression Preservation). If $\Gamma ; \Theta \vdash e: \tau$ and $H \vDash \Theta$ and $\langle H, e\rangle \rightarrow e^{\prime}$ then $\Gamma ; \Theta \vdash e^{\prime}: \tau$.

To prove both theorems we also need the following propositions that model our assumptions about the functions $\mathcal{C A}$ and $\mathcal{C V}$ modeling the control plane.

Proposition 1 (Control Plane Reads). If $H \vDash \Theta$ and $\mathcal{C V}(t)=\bar{S}$ and $\bar{e}=\left\{e_{j} \mid\left(e_{j}, m_{j}\right) \epsilon\right.$ t.reads ()$\wedge \neg$ maskable $\left.\left(t, e_{j}, m_{j}\right)\right\}$ and $\Gamma ; \Theta \vdash e_{j}: \tau_{j}$ for $e_{j} \in \bar{e}$ then $\mathcal{C A}(t, H)=\left(a_{i}, \bar{v}\right)$.

Proposition 2 (Control Plane Action Data). If $H \vDash \Theta$ and $\mathcal{C A}(t, H)=\left(a_{i}, \bar{v}\right)$ and $\mathcal{A}\left(a_{i}\right)=\lambda \bar{x}: \bar{\tau} . c_{i}$ then $\cdot ; \cdot \vdash \bar{v}: \bar{\tau}$

Proposition 3 (Control Plane Assumptions). If $H \vDash \Theta$ and $\mathcal{C A}(t, H)=\left(a_{i}, \bar{v}\right)$ and $\mathcal{C} \mathcal{V}(t)=\bar{S}$ then $H \vDash$ Restrict $\Theta S_{i}$.

Proposition 1 captures the first assumption-i.e., the control plane can safely install table entries that never read invalid headers, Proposition 2 captures the second assumption-i.e., the action data provided by the control plane has the types expected by the action and Proposition 3 captures the third assumption-i.e., the control plane will only assume valid headers for an action that are valid for a given packet.

Theorem 4.8 (Progress). If $\cdot \vdash c: \Theta_{1} \Leftrightarrow \Theta_{2}$ and $H \vDash \Theta_{1}$, then either, $c=s k i p$, or $\exists\left\langle I^{\prime}, O^{\prime}, H^{\prime}, c^{\prime}\right\rangle .\langle I, O, H, c\rangle \rightarrow\left\langle I^{\prime}, O^{\prime}, H^{\prime}, c^{\prime}\right\rangle$.

Proof. By induction on the typing derivation.
Intuitively, progress states that a well-typed command is fully reduced or can take a step.

Theorem 4.9 (Preservation). If $\Gamma \vdash c: \Theta_{1} \Leftrightarrow \Theta_{2}$ and $\langle I, O, H, c\rangle \rightarrow\left\langle I^{\prime}, O^{\prime}, H^{\prime}, c^{\prime}\right\rangle$, where $H \vDash \Theta_{1}$, then $\exists \Theta_{1}^{\prime}, \Theta_{2}^{\prime}$. $\Gamma \vdash c: \Theta_{1}^{\prime} \Leftrightarrow \Theta_{2}^{\prime}$ where $H^{\prime} \vDash \Theta_{1}^{\prime}$ and $\Theta_{2}^{\prime}<\Theta_{2}$.

Proof. By induction on a derivation of $\Gamma \vdash c: \Theta_{1} \Leftrightarrow \Theta_{2}$, with a case analysis on the last rule used.

Preservation says that if a command $c$ is well-typed with input type $\Theta_{1}$ and output type $\Theta_{2}$ and $c$ evaluates to $c^{\prime}$ in a single step, then there exists an input type $\Theta_{1}^{\prime}$ and an output type $\Theta_{2}^{\prime}$ that make $c^{\prime}$ well-typed, and $\Theta_{1}^{\prime}$ describes the same maps of header instances $H$ as $\Theta_{1}$, and $\Theta_{2}^{\prime}$ is semantically contained in $\Theta_{2}$. We define syntactic containment to be $\Theta_{1}<\Theta_{2} \xlongequal{\triangleq} \llbracket \Theta_{1} \rrbracket \subseteq \llbracket \Theta_{2} \rrbracket$.

### 4.6 Related Work

Formal Reasoning for $\mathbf{P}_{4}$ Programs With respect to verifying the correctness of $\mathrm{P}_{4}$ programs, probably the most closely related work to SAFEP4 is p4v [Liu+18]. Unlike SAFEP4, which is based on a static type system, p4v uses Dijkstra's approach to program verification based on predicate transformer semantics. To model the behavior of the control plane, p4v uses first-order annotations. SAFEP4's typing rule for table application is inspired by this idea, but adopts simple heuristics-e.g., we only assume that the control plane is well-behaved-rather than requiring logical annotations. Both p4v and $\mathrm{P}_{4}$ Сheск can be used to verify safety properties of data planes modelled in $\mathrm{P}_{4}-$ e.g., that no read or write operations are possible on an invalid header. As it is often the case when comparing approaches based on types to those based on program verification, p 4 v can check more complex properties, including architectural invariants and programspecific properties-e.g., that the $\operatorname{IPv} 4$ time-to-live field is correctly decremented on every packet. However, in general, it requires annotating the program with formal specifications both for the correctness property itself and to model the behavior of the control plane. McKeown et al. developed an operational semantics for $\mathrm{P}_{4}[\mathrm{McK}+16]$, which is translated to Datalog to verify safety properties and to check program equivalence. An operational semantics for $\mathrm{P}_{4}$ was also developed in the K framework [ $\mathrm{R} S_{1} 10$ ], yielding a symbolic model checker and deductive verification tool [KR18]. Vera [Sto +18 ] models the semantics of $\mathrm{P}_{4}$ by translation to SymNet [Sto+16], and develops a symbolic execution engine for verifying a variety of properties, including header validity. Compared to SAFEP4, these approaches do not use their formalization of $\mathrm{P}_{4}$ as a foundation for defining a type system that addresses common bugs. To the best of our knowledge, $\mathrm{SAFEP}_{4}$ is the first formal calculus for a $\mathrm{P}_{4}$-like packet processing language that provides correct-by-construction guarantees of header safety properties.

Ensuring Null-Safety Other languages have used type systems to rule our safety problems due to null references. For example, NullAway [Sri18] analyzes all Java programs annotated with special annotations, making path-sensitive deductions about which references may be null. Similar to the validity checks in SAFEP4, NullAway analyzes conditionals for null checks of the form var != null using data flow analysis.

Packet Processing Languages Looking further afield, PacLang [ESMo4] is a concurrent packet-processing language that uses a linear type system to allow multiple references to a given packet within a single thread. PacLang and SAFEP4 share the use of a type system for verifying safety properties, but they differ in the kind of properties they address and, hence, the kind of type system they employ for this purpose. In addition, the primary focus in PacLang is on efficient compilation whereas $\mathrm{SAFEP}_{4}$ is concerned with ensuring safety of header data. Domino [Siv+16] is a domain-specific language for data plane algorithms supporting packet transactions-i.e., blocks of code that are guaranteed to be atomic and isolated from other transactions. In Domino, the programmer defines the operations needed for each packet without worrying about
other in-flight packets. In case of success, the compiler guarantees performance at the line rate supported on programmable switches. Overall, Domino focuses on transactional guarantees and concurrency rather than header safety properties. BPF+ [BMG99] and eEBPF [Cor14] are packet-processing frameworks that can be used to extend the kernel networking stack with custom functionality. The modern eBPF framework is based on machine-level programming model, but it uses a virtual machine and code verifier to ensure a variety of basic safety properties. Much of the recent work on eBPF focuses on techniques such as just-in-time compilation to achieve good performance. SNAP [Ara+16] is a language for stateful packet processing based on $\mathrm{P}_{4}$. It offers a programming model with global state registers that are distributed across many physical switches while optimizing for various criteria, such as minimizing congestion. More specifically, the compiler analyzes read/write dependencies to automatically optimize the placement of state and the routing of traffic across the underlying physical topology.

Taint Analysis and Formal Calculi While our approach to track validity is networkspecific, it is similar to taint analysis [VIS96; HOMo6; HDM14], which attempts to identify secure program parts that can be safely accessed. Of course, there is a long tradition of formal calculi that aim to capture some aspect of computation and make it amenable for mathematical reasoning. The design of SAFEP4 is directly inspired by Featherweight Java [IPWo1], which stands out for its elegant formalization of a real-world language in an extensible core calculus.

### 4.7 Chapter Summary

In this chapter we introduced $\mathrm{SAFEP}_{4}$, a domain-specific language for program-mable data planes, which is equipped with a static type system that guarantees that all headers read or written are guaranteed to be valid. SAFEP4 models the essential features of $\mathrm{P}_{4}$ but prunes away language constructs that are secondary to how packets are processed. To be compatible with both language versions $\mathrm{P}_{4_{14}}$ and $\mathrm{P}_{416}$, SAFEP4 abstracts away expression typing and evaluation behaviors for expressions in terms of $n$-ary constants that are assumed to have sound types.

SAFEP4 introduces the notion of header types, which allow the type system to statically capture which header instances are valid at a specific point in the program. To ensure that capturing valid headers is not too restrictive, SAFEP4 employs occurrence typing, which allows the type checker to use more precise types depending on whether explicit validity checks succeed or fail. Furthermore, using a fine-grained representation of valid header instances based on sets of sets of headers, $\mathrm{SAFEP}_{4}$ is able to capture valid header instances together with their dependencies per program path.

The problem of header validity being a dynamic property due to the interaction with the control plane is addressed by modeling the runtime interface between the switch and the controller using three functions. Given a table and the currently valid headers, the first function returns the action to call and the possibly empty action data arguments. The second function produces for each action of a table a set of valid headers that can be safely assumed valid when the entries are populated by the control plane. From the table declaration and the header instances that can be assumed valid, the third function derives a list of match key expressions that must be evaluated when the table is invoked. In addition, the control plane interface relies on three basic correctness properties that are assumed to hold: (1) the control plane can safely install table entries that never read invalid headers, (2) the action data provided by the control plane has the types expected
by the action, and (3) the control plane will only assume valid headers for an action that are valid for a given packet. These correctness properties allowed us to prove safety for SAFEP4's type system.

# Dependently-Typed Data Plane Programming 

In Chapter 4, we introduced SafeP4, a domain-specific language that uses a correct-byconstruction approach to statically eliminate a variety of errors based on header validity. A major advantage of the presented approach is that no complex program annotations are required, as is the case with other data plane verification tools. However, because header validity is deeply baked into the system as a central correctness property, $\mathrm{SafeP}_{4}$ is limited in its ability to verify richer properties compared to other verification tools. In particular, $\mathrm{SafeP}_{4}$ is not able to capture individual values of the program state. For example, it is not possible to guarantee that headers are only ever accessed on mutually exclusive program paths (e.g., IPv4 and IPv6) or to track dependencies between headers based on header field values.

While type system-based approaches are so far limited in their expressiveness, they have one key advantage over many existing data plane verification tools, which is compositionality. Type systems are designed to enable compositional reasoning-i.e., the types for individual components document assumptions about the components they rely upon as well as the guarantees they offer. So far, there has been little effort in writing modular $\mathrm{P}_{4}$ code in which individual parts of the code can be reused, which has long been common practice in general-purpose programming languages. However, recently, there have been efforts to enable modular designs of data plane programs [Gao+20; Son+20], paving the way for an "open-world" model in which third-party components are embedded into existing programs.

The question therefore arises whether it is possible to combine the expressive power of fully-fledged verification tools with the compositional checking inherent to type systems. Dependently-typed languages [XP99; Con+07; RKJo8; Vaz+14] are increasingly blurring the line between type checking and theorem proving. For instance, Liquid Haskell [RKJo8; Vaz+14] allows programmers to smoothly shift from properties that can be checked with traditional typing disciplines to more sophisticated ones. Under the hood, an SMT solver automatically discharges the formulas generated during type checking without requiring manual proofs. So far, the dependently-typed approach has not yet been explored in the context of network programming.

```
parser MyParser(packet_in pkt, out headers hdr, ...) {
    state start {
        pkt.extract(hdr.ethernet);
        transition select(hdr.ethernet.etherType) {
            0x0800: parse_ipv4;
            default: accept;
        }
    }
    state parse_ipv4 {
        pkt.extract(hdr.ipv4);
        transition accept;
    }
}
control MyIngress(inout headers hdr, ...) {
    apply {
        if (hdr.ipv4.isValid()) {
            if (hdr.ipv4.src == 10.10.10.10) {
                drop();
            }
        }
    }
}
```

Figure 5.1: Safe program according to $\mathrm{SaFEP}_{4}$, using a validity check in the ingress.

In this chapter, we present $\Pi_{4}$, a dependently-typed core of the $\mathrm{P}_{4}$ language. For $\Pi_{4}$ 's type system, we extend SafeP4's header types to heap types, which also capture the shape of valid packet headers as well as the shape of the incoming and outgoing packet. We formalize its syntax and semantics and show how we can enable precise typing in the presence of domain-specific features that combine packet serialization and deserialization operations with imperative control flow. For example, our novel chomp operator allows to precisely capture the effect of packet deserialization, i.e., it allows to compute the type that remains after extracting bits from a packet buffer. Finally, we prove safety for $\Pi_{4}$ 's type system.

### 5.1 An Overview of $\Pi_{4}$

To get a first understanding of how dependent types allow checking richer properties, this section provides a high-level overview of $\Pi_{4}$. Let us consider the program shown in Figure 5.1. The parser (lines 1 to 13) extracts Ethernet, optionally followed by IPv4. The ingress uses an explicit validity check to make sure that IPv4 is valid before conditionally dropping the packet depending on the value of the IPv4 source address (lines 17-19).

This program will pass SAFEP4's type checker. On the other hand, if we use the same parser but the implementation of the ingress control shown in Figure 5.2, SafeP4's type checker will reject the program. This might be surprising, because the parser guarantees

```
control MyIngress(inout headers hdr, ...) {
    apply {
            if (hdr.ethernet.etherType == 0x0800) {
                if (hdr.ipv4.src == 10.10.10.10) {
                    drop();
                }
            }
        }
    }
```

Figure 5.2: Ingress control using a data-dependent check to guarantee only valid headers are accessed.
that if the EtherType is equal to oxo8oo, the IPv 4 header is also valid, thus, the program is actually safe. However, the type computed for the parser ethernet • (ipv4 + 1) does not guarantee that $\mathrm{IPv}_{4}$ is valid on all program paths and also the path-sensitive reasoning does not add additional assumptions about the validity of $\mathrm{IPv}_{4}$ in the ingress.

To address this problem, $\Pi_{4}$ employs a dependent type system [ $\mathrm{XP}_{99}$ ], in which we can compute a precise type for the program after parsing:

$$
\begin{aligned}
& \left(x:\left\{y: \epsilon| | \mathrm{y} \cdot \mathrm{pkt} \mathrm{t}_{\mathrm{in}} \mid>272\right\}\right) \rightarrow \\
& \qquad\left(\begin{array}{l}
\sum y: \text { ether. }\{z: \text { ipv4 } \mid y . \text { ether.etherType }==\mathrm{oxo80o}\}+ \\
\{y: \text { ether } \mid y . \text { ether.etherType } \neq \text { oxo8oo }\}
\end{array}\right.
\end{aligned}
$$

Intuitively, this type says that, starting with the empty heap $(y: \epsilon)$ with at least enough bits to extract both the Ethernet and the IPv4 header $\left(\left|y \cdot p k t_{i n}\right|>272\right)$, the parser ends in one of two possible states (denoted by +): (1) both Ethernet and $\mathrm{IPv}_{4}$ are valid ( $\Sigma y$ : ether. $\{z: \operatorname{ipv} 4 \mid \ldots\}$ ), if the EtherType is equal to oxo8oo (note how $z$ : ipv4 is conditioned by $y$.ether.etherType $==0 \times 0800$ ) or (2) just Ethernet is valid, if EtherType is not equal to oxo8oo. When checking the ingress control, the type checker can use the predicate ether.etherType $==0 \times 0800$ on the conditional to derive the set of valid header instances, which in this case includes IPv4. Thus, accessing the $\mathrm{IPv}_{4}$ source address is safe, and the program correctly passes the type checker.

While the output type is admittedly notationally heavy-a common feature in precise type systems - the programmer is not forced to write down the most precise type. $\Pi_{4}$ only requires the annotated type to be sufficiently precise to capture basic safety guarantees and other desired invariants. For example, in a program where only the Ethernet header is needed to be valid at the end of the parser, we can use the type $\left(x:\left\{y: \epsilon| | y \cdot\right.\right.$ pkt $\left.\left._{\text {in }} \mid>272\right\}\right) \rightarrow$ ether $_{\approx}$, which indicates that at least Ethernet is valid but possibly others, too.

### 5.2 Design

In designing $\Pi_{4}$, our primary goal is to enable data plane programmers to make use of dependent types to verify useful program properties in a compositional way and without having to write manual proofs. We want to show that dependent types are a good match for data plane programming.

Deviation from P4 Similar to SAFEP4, $\Pi_{4}$ focuses on the unique aspects of the $\mathrm{P}_{4}$ programming language, which benefit from dependent types, (e.g., parsing, deparsing, validity, and control flow) and omits features that would simply add clutter (e.g., externs, registers, checksums, hashing, and pipelines). We can even get away without explicitly modeling tables. Following p4v [Liu+18], $\Pi_{4}$ uses ghost state and conditionals to encode tables, which is discussed in detail in Section 8.3.2. Consequently, $\Pi_{4}$ is a loop-free ${ }^{1}$ imperative language with a few domain-specific primitive commands: extract $(\iota)$, $\operatorname{add}(\iota)$, remove $(\iota)$, remit $(\iota)$, and reset.

P4's emit primitive serializes a header instance $\iota$ into a series of bits and prepends it to the outgoing packet payload, only if $\iota$ is valid, otherwise it does nothing. To simplify typing rules and semantics, $\Pi_{4}$ provides the primitive remit $(\iota)$, which really emits $\iota$ if it is valid, and otherwise gets stuck. Hence, emit ( $\iota$ ) can be expressed using the command if ( $\iota . v a l i d)$ remit ( $\iota$ ) else skip.

We follow the design decisions made in $\mathrm{SAFEP}_{4}$ with respect to explicitly validating header instances. Again, to avoid dealing with undefined values from reads to uninitialized header instances, rather than forcing the programmer to manually write default values, the $\operatorname{add}(\iota)$ command sets $\iota$.valid to true, and assigns instance $\iota$ a pre-determined default value (say o). If required, P4's behavior could be encoded using an extra 1-bit header to independently track the validity of the instance and initialization of its fields.

We also introduce a new primitive, reset, which models the behavior of $\mathrm{P}_{4}$ between pipeline stages. In many switch architectures [Bos+13], packets are deparsed and then reparsed between pipelines-e.g., after ingress and before egress. The reset command encodes the behavior of the inner step: it combines the deparsed bits with the packet's unparsed payload and passes it along as the input to the next stage.

We model header field accesses as direct bit-slices into the instance (to avoid another layer of indirection in our semantics)-i.e., eth.srcAddr is written eth[48:96].

Type system $\quad$ 's's type system is intended to promote modular reasoning, so we need a way to annotate and modularly check programs with types. We annotate a program $c$ with a type $\sigma$ using an ascription operator $c$ as $\sigma$. The ascription has no effect on the runtime behavior of the code (i.e., $c$ as $\sigma$ always just steps to $c$ ). It does, however, indicate a program point where type checking should occur. Hence, we can independently type check $c$ with type $\sigma$ and then use $\sigma$ when checking the rest of the program.

We always assign a dependent function type $\left(x: \tau_{1}\right) \rightarrow \tau_{2}$ to a command $c$, where $x$ may occur in $\tau_{2}$. This design allows us to relate the input and output heaps of commands described by heap types $\tau_{1}$ and $\tau_{2}$ respectively. For example, we may want to ensure that the Ethernet header has the same value after being deparsed, followed by a reset, and then being parsed again. To express equations like this, we use refinement types $\{y: \tau \mid \varphi\}$, where $\varphi$ is a formula in the logic of variable-width bit vectors with concatenation and length operators. In this example, we could say that the Ethernet header is unchanged by using the type $\left\{y: \tau_{2} \mid x\right.$.eth $=y$.eth $\}$.

Granularity of types Our goal is to equip $\Pi_{4}$ 's type system with the ability to capture information down to the bit level. In particular, we must ensure that the input type and output type remain consistent after bits have been shuffled around by a command, which is especially challenging when parsing. For example, given an input type $\left\{x: T\left|x . p k t_{i n}[0: 8]=0 x 04 \wedge\right| x . p k t_{i n} \mid>160\right\}$, where $x . p k t_{i n}$ represents an incoming

[^2]```
\(\tau \quad::=\varnothing|\top| \Sigma x: \tau . \tau|\tau+\tau|\{x: \tau \mid \varphi\} \mid \tau[x \mapsto \tau] \quad\) (heap types)
\(\sigma \quad::=\mathbb{N}|\mathbb{B}| \operatorname{BV} \mid(x: \tau) \rightarrow \tau \quad\) (base types)
\(\varphi \quad::=e=e|e>e| \varphi \wedge \varphi|\neg \varphi|\) x.ı.valid \(\mid\) true \(\mid\) false \(\quad\) (formulae)
\(e \quad::=n|b v||x . p||e+e| e @ e|x . p| x . p[l: r] \mid x . l[l: r] \quad\) (expressions)
\(b v \quad::=\langle \rangle|o:: b v| 1:: b v \mid \mathrm{b}_{\mathrm{n}}:: b v\)
\(p \quad::=p k t_{\text {in }} \mid p k t_{\text {out }}\)
\(c \quad::=\operatorname{extract}(\iota)|\operatorname{add}(\iota)| \operatorname{remove}(\iota)|\iota . f:=e| \operatorname{remit}(\iota) \mid\)
    reset \(\mid\) if \((\varphi)\) c else \(c|c ; c|\) skip \(\mid c\) as \((x: \tau) \rightarrow \tau\)
\(d \quad::=\eta\{\overline{f: B V}\} \mid \iota \mapsto \eta\)
\(P \quad::=(\bar{d}, c)\)
    (bit vectors)
    (packets)
(commands)
(declarations)
    (programs)
```

Figure 5.3: Syntax of $\Pi_{4}$
packet and the command extract(ipv4), the output type should reflect that the ipv4 header instance is now valid, that ipv4[0:8] is 0x04, and that $x . p k t_{i n}$ may have no more bits remaining. $\Pi_{4}$ accomplishes this using two key mechanisms: (1) a dependent sum type $\Sigma x: \tau_{1} . \tau_{2}$ that computes the disjoint union of the valid instances in $\tau_{1}$ and $\tau_{2}$ and concatenates the incoming and outgoing packets and (2) a refinement transformer, chomp, that manipulates input refinements to be consistent with the extraction operation.

### 5.3 Syntax

Figure 5.3 shows the syntax of $\Pi_{4}$. Boolean formulae $\varphi$ include expression equality ( $e_{1}=e_{2}$ ), expression comparison ( $e_{1}>e_{2}$ ), conjunction ( $\wedge$ ), negation ( $\neg$ ), validity of instances (x.ı.valid) and boolean literals true and false. Expressions $e$ include naturals, bit vectors, packet length $|x . p|$, addition (+), concatenation ( $e_{1} @ e_{2}$ ), packet access $x . p$ and slices of packets $(x . p[l: r])$ and instances $(x . \iota[l: r])$. Packet accesses refer either to the input packet $\left(p k t_{i n}\right)$ or the output packet $\left(p k t_{\text {out }}\right)$.

To ease the notation, we write $x . l[l]$ instead of $x . l[l: l+1]$ for bit-wise access, $x . \iota . f$ instead of $x . \iota[l: r]$ for ranges matching header instance fields, $x . \iota$ instead of $x . \iota[0: \operatorname{sizeof}(\iota)]$, and similarly for the corresponding expressions involving packet variables $x . p$. We use a list-like encoding of bit vectors. A bit vector is either the empty bit vector $\left\rangle\right.$ or a concatenation of bits. We assume that bit variables $b_{n}$ are not part of the surface syntax and are only used internally. For singleton bit vectors, we write $\langle b\rangle$ instead of $b::\langle \rangle$.

We write $x \equiv y$ (respectively $x \equiv, y$ ) as syntactic sugar for the boolean predicates capturing strict equality (respectively instance equality) between the heaps bound to $x$ and $y$. Strict equality requires that both the input and output packets are equivalent as well as all instances contained in the heap. It is formally defined as follows:

$$
\begin{aligned}
& x \equiv y \triangleq x \cdot p k t_{\text {in }}=y \cdot p k t_{\text {in }} \wedge x \cdot p k t_{\text {out }}=y \cdot p k t_{\text {out }} \wedge \\
& \bigwedge_{\operatorname{dom}(\mathcal{H} \mathcal{T})}(x . \iota . \text { valid } \wedge y . \iota . \text { valid } \wedge x . \iota=y . \iota) \vee(\neg x . \iota . \text { valid } \wedge \neg y . \iota . \text { valid })
\end{aligned}
$$

In contrast, instance equality only requires that the instances are equivalent in both heaps according to the following definition:

$$
x \equiv, y \triangleq \bigwedge_{\iota \in \operatorname{dom}(\mathcal{H} \mathcal{T})}(x . \iota . v a l i d \wedge y . \iota . v a l i d \wedge x . \iota=y . \iota) \vee(\neg x . \iota . \text { valid } \wedge \neg y . \iota . \text { valid })
$$

A program consists of a sequence of declarations $\bar{d}$ and a command $c$, where $\bar{x}$ is a shorthand for a possibly empty sequence $x_{1}, \ldots, x_{n}$. Declarations $d$ include header type declarations $\eta\{\overline{f: \mathrm{BV}}\}$ and header instance declarations $\iota \mapsto \eta$. Header type declarations specify the format of network packet headers. They are defined in terms of a name and a sequence of field declarations, where each field is itself defined in terms of a field name and a bit vector type. We write $f: \mathrm{BV}$ to denote that field $f$ has a bit vector type BV . With $\eta$ ranging over header types, the instance declaration $\iota \mapsto \eta$ assigns the name $\iota$ to header type $\eta$. The global mapping between header instances and header types is stored in the so-called header table $\mathcal{H} \mathcal{T}$. We assume that names of header instances and header types are drawn from disjoint sets of names and that each name is declared only once.
$\Pi_{4}$ provides commands for parsing (extract), creating (add), removing (remove) and modifying ( $\iota . f:=e$ ) header instances. The remit command serializes a header instance into a bit sequence. The reset command resets the program state-in particular, the packet buffers and all assumptions about header validity. The if-command conditionally executes one out of two commands based on the value of the boolean formulae $\varphi$. Commands can be sequentially composed $\left(c_{1} ; c_{2}\right)$, skip is a no-op, and with type ascription ( $c$ as $(x: \tau) \rightarrow \tau$ ) it is possible to explicitly assign a type to a command which it is assumed to have at the current point in the program. We assume that every header referenced in a program has a corresponding instance declaration-a property that could be enforced using a simple static analysis.

Both heap types and commands share the same syntactic categories of formulas and expressions. However, since there are no binders at the level of commands, we implicitly assume that formulas used as the condition for an if-statement as well as expressions assigned to header fields are implicitly prefixed with a variable named heap. We usually omit these binders in the surface syntax, for example, we write if (ethernet. valid) ...else... instead of if (heap.ethernet.valid)...else....
$\Pi_{4}$ provides two categories of types, base types $\sigma$ and heap types $\tau$. Base types include natural numbers $(\mathbb{N})$, booleans $(\mathbb{B})$, bit vectors $(\mathrm{BV})$ and dependent function types $((x: \tau) \rightarrow \tau)$. Heap types include the bottom type ( $\varnothing$ ), the top type ( $T$ ) and dependent pairs $\Sigma x: \tau_{1} . \tau_{2}$, where $x$ may occur in $\tau_{2}$. $\Pi_{4}$ also supports fine-grained path-dependent reasoning via union types $\left(\tau_{1}+\tau_{2}\right)$. Refinement types $\{x: \tau \mid \varphi\}$ allow to endow a type with a boolean predicate $\varphi$ which is assumed to hold for all heaps described by type $\tau$. We also often need to reference intermediate types, which is achieved with substitution types $\tau_{2}\left[x \mapsto \tau_{1}\right]$, where $x$ may occur in $\tau_{2}$. In such a type, $\tau_{1}$ may represent the type at any earlier point in the program.

To ease the notation, we additionally define the types $\epsilon, \iota$ and $\iota_{\approx}$. These types are only syntactic sugar and thus can be expressed by combinations of the previously described heap types. The type $\epsilon \triangleq\left\{x: \top \mid \bigwedge_{t \in \operatorname{dom}(\mathcal{H} \mathcal{T})} \neg x\right.$.l.valid $\}$ describes the empty heap on which no header instances are valid. The type $\iota \triangleq\{x: \top \mid x . \iota . v a l i d \wedge$ $\Lambda_{\iota^{\prime} \in \operatorname{dom}(\mathcal{H} \mathcal{T}), \iota^{\prime} \neq \iota} \neg x . \iota^{\prime}$.valid $\}$ describes the heap on which only instance $\iota$ is valid, while $\iota_{\approx} \triangleq\{x: \top \mid x . \iota . v a l i d\}$ describes the heap on which at least instance $\iota$ is guaranteed to be valid.

### 5.4 Well-formedness

We assume that all types, formulae and expressions are well-formed, i.e., they satisfy the following basic syntactic properties as defined in Figure 5.4, Figure 5.5 and Figure 5.6.

1. There are no free variables in types


Figure 5.4: Well-formedness of types.

$$
\begin{aligned}
& \begin{array}{llll}
\text { WF-True } & \text { WF-False } & \text { WF-VALid } & \text { WF-NeG } \\
\overline{\Gamma \vdash \mathrm{Wf}_{\varphi} \text { true }} & \frac{x \in \Gamma}{\Gamma \vdash \mathrm{Wf}_{\varphi} \text { false }} & \frac{\Gamma \vdash \mathrm{wf}_{\varphi} \varphi}{\Gamma \vdash \mathrm{wf}_{\varphi} x . l . v a l i d} & \frac{\Gamma \vdash \mathrm{wf}_{\varphi} \neg \varphi}{\Gamma}
\end{array} \\
& \text { Wf-Conj Wf-EQ } \\
& \frac{\Gamma \vdash \mathrm{wf}_{\varphi} \varphi_{1} \quad \Gamma \vdash \mathrm{wf}_{\varphi} \varphi_{2}}{\Gamma \vdash \mathrm{wf}_{\varphi} \varphi_{1} \wedge \varphi_{2}} \quad \frac{\Gamma \vdash \mathrm{wf}_{e} e_{1} \quad \Gamma \vdash \mathrm{wf}_{e} e_{2}}{\Gamma \vdash \mathrm{wf}_{\varphi} e_{1}=e_{2}} \\
& \text { Wf-GT } \\
& \frac{\Gamma \vdash \mathrm{wf}_{e} e_{1} \quad \Gamma \vdash \mathrm{wf}_{e} e_{2}}{\Gamma \vdash \mathrm{wf}_{\varphi} e_{1}>e_{2}}
\end{aligned}
$$

Figure 5.5: Well-formedness of formulae.
2. The bounds of bit vector slices are positive and describe at least a range of length one.
3. The bounds of instance slices respect the statically known size of the instance

### 5.5 Dynamic Semantics

The operational semantics of $\Pi_{4}$ is in many aspects similar to the operational semantics of SAFEP4 (cf. Section 4.4). It is also defined in terms of a four-tuple $\langle I, O, H, c\rangle$, where $I$ is the bit stream of the incoming packet, $O$ is the bit stream of the outgoing packet, $H$ is a map that relates instance names to records containing the field values, and $c$ is a command. The rules of the operational semantics of $\Pi_{4}$ are shown in Figure 5.8.

The semantics of the extract command (rule E-ExTract) is identical to the semantics of the command in SAFEP4. Again, we assume the existence of a deserialization function deserialize $_{\eta}$ (as defined in Section 4.4) that copies the appropriate number of bits from the input bit stream into the deserialized representation of the instance $v$ leaving the remainder of the input bit stream $I^{\prime}$. The deserialized value is added to the map of valid header instances $H$.

| Wf-Num | Wf-Bitvec | Wf-Length $x \in \Gamma$ | $\begin{aligned} & \text { WF-PLus } \\ & \Gamma \vdash \mathrm{wf}_{e} e_{1} \end{aligned}$ | $\Gamma \vdash \mathrm{wf}_{e} e_{2}$ |
| :---: | :---: | :---: | :---: | :---: |
| $\overline{\Gamma \vdash \mathrm{wf}_{e} n}$ | $\overline{\Gamma \vdash \mathrm{wf}_{e} b v}$ | $\overline{\Gamma \vdash \mathrm{wf}_{e}\|x \cdot p\|}$ | $\Gamma \vdash \mathrm{wf}_{e} e_{1}+e_{2}$ |  |
| Wf-Concat <br> $\Gamma \vdash \mathrm{wf}_{e} e_{1}$ |  | $\begin{gathered} \text { Wf-Packet } \\ x \in \Gamma \end{gathered}$ | Wf-PacketSlice |  |
|  | $\Gamma \vdash \mathrm{wf}_{e} e_{2}$ |  | $x \in \Gamma$ | $l<r$ |
| $\Gamma \vdash w$ | $e_{1}$ @ $e_{2}$ | $\Gamma \vdash \mathrm{wf}_{e} x . p$ | $\Gamma \vdash \mathrm{wf}$ | [l:r] |

Wf-InstanceSlice

$$
\frac{x \in \Gamma \quad l \geq 0 \quad l<r \quad r \leq \operatorname{sizeof}(\iota)}{\Gamma \vdash \mathrm{wf}_{e} x . l[l: r]}
$$

Figure 5.6: Well-formedness of expressions.

$$
\begin{array}{ll}
\begin{array}{l}
\text { E-MoD } \\
H(\iota)=r
\end{array} \quad r^{\prime} \triangleq\{r \text { with } f=v\} \\
\langle I, O, H, \iota . f:=v\rangle \rightarrow\left\langle I, O, H\left[\iota \mapsto r^{\prime}\right], s k i p\right\rangle
\end{array} \quad \text { E-MoD1 } \quad \begin{aligned}
& \langle I, O, H, t\rangle \rightarrow t^{\prime} \\
& \langle I, O, H, h . f=t\rangle \rightarrow\left\langle I, O, H, h . f=t^{\prime}\right\rangle
\end{aligned}
$$

E-Remit

$$
\begin{array}{cl}
\iota \in \operatorname{dom}(H) \\
\operatorname{serialize}_{\eta}(H(\iota))=b v
\end{array} \quad \begin{aligned}
& \text { E-Reset } \\
& \frac{I^{\prime}=O @ I}{\langle I, O, H, \operatorname{remit}(\iota)\rangle \rightarrow\langle I, O:: b v, H, s k i p\rangle}
\end{aligned} \quad \begin{aligned}
& \langle I, O, H, \text { reset }\rangle \rightarrow\left\langle I^{\prime},\langle \rangle,[], s k i p\right\rangle
\end{aligned}
$$

Figure 5.7: Small-step operational semantics of $\Pi_{4}$.

Command add evaluates by rule E-ADD if the instance is not yet valid. The evaluation is similar to rule E-Extract, except that no bits are taken from the input bit stream. Again, we assume that there exists an initialization function init $_{\eta}$ for every heap type $\eta$ that initializes all fields of an instance to a fixed value (cf. Section 4.4). If an instance is already valid, the program gets stuck.

Removing a header instance (rule E-Remove) requires that the header instance is valid-i.e., it is contained in $H$. While the input and output packet are not affected, the instance is removed from the map of valid header instances $H$, which is denoted by $H \backslash i$.

Command remit( $\iota$ ) (rule E-Remit) also requires that the header instance is valid.

$$
\begin{aligned}
& \text { E-Extract } \\
& \frac{\mathcal{H} \mathcal{T}(\iota)=\eta \quad \text { deserialize }_{\eta}(I)=\left(v, I^{\prime}\right)}{\langle I, O, H, \operatorname{extract}(\iota)\rangle \rightarrow\left\langle I^{\prime}, O, H[\iota \mapsto v], \text { skip }\right\rangle} \\
& \text { E-Add } \\
& \frac{\iota \notin \operatorname{dom}(H) \quad \mathcal{H} \mathcal{T}(\iota)=\eta \quad \text { init }_{\eta}=v}{\langle I, O, H, \operatorname{add}(\iota)\rangle \rightarrow\langle I, O, H[\iota \mapsto v], \text { skip }\rangle} \\
& \text { E-Remove } \\
& \frac{\iota \in \operatorname{dom}(H)}{\langle I, O, H, \text { remove }(\iota)\rangle \rightarrow\langle I, O, H \backslash \iota, \text { skip }\rangle}
\end{aligned}
$$

| E-SEQ | E-SEQ1 <br> $\left\langle I, O, H, c_{1}\right\rangle \rightarrow\left\langle I^{\prime}, O^{\prime}, H^{\prime}, c_{1}^{\prime}\right\rangle$ |
| :--- | :--- |
| $\left\langle I, O, H, s k i p ; c_{2}\right\rangle \rightarrow\left\langle I, O, H, c_{2}\right\rangle$ | $\frac{\left\langle I, O, H, c_{1} ; c_{2}\right\rangle \rightarrow\left\langle I^{\prime}, O^{\prime}, H^{\prime}, c_{1}^{\prime} ; c_{2}\right\rangle}{\left\langle I, O, H, c_{1}\right.}$ |

E-Ascribe

$$
\overline{\left\langle I, O, H, c \text { as }\left(x: \tau_{1}\right) \rightarrow \tau_{2}\right\rangle \rightarrow\langle I, O, H, c\rangle}
$$

E-IF

$$
\frac{\langle I, O, H, \varphi\rangle \rightarrow \varphi^{\prime}}{\left\langle I, O, H, \text { if }(\varphi) \text { then } c_{1} \text { else } c_{2}\right\rangle \rightarrow\left\langle I, O, H, \text { if }\left(\varphi^{\prime}\right) \text { then } c_{1} \text { else } c_{2}\right\rangle}
$$

E-If True
$\overline{\left\langle I, O, H, \text { if (true) then } c_{1} \text { else } c_{2}\right\rangle \rightarrow\left\langle I, O, H, c_{1}\right\rangle}$
E-If False
$\overline{\left\langle I, O, H, \text { if }(\text { false }) \text { then } c_{1} \text { else } c_{2}\right\rangle \rightarrow\left\langle I, O, H, c_{2}\right\rangle}$

Figure 5.8: Additional rules of the small-step operational semantics of $\Pi_{4}$.

Again, we assume there is a serialization function for every heap type (cf. Section 4.4), which turns a record representing the instance back into a bit sequence. The serialized bit sequence is appended to the end of the outgoing packet. Both the input packet and the set of valid headers remain unchanged. In contrast to $\mathrm{SAFEP}_{4}$, the program gets stuck, if we try to emit a header instance that is currently not valid.

Rule E-Mod defines the semantics of assigning a value to a header field. Assuming $r$ is the record storing the values of the fields, an updated record $r^{\prime}$ with the modified field value is stored in $H$. The input and output packets remain unchanged. If the assigned expression is not a value, it is reduced first (rule E-Mod1).

Rule E-Reset defines the semantics of the command reset. It would be invoked between the ingress and egress pipelines, when the packet emitted by the ingress becomes the input packet for the egress. Operationally, the bits contained in the output packet are prepended to the bits of the input packet. This concatenated bit sequence serves as the new input packet. The output packet is emptied and all valid header instances are discarded.

An ascribed command $c$ as $\sigma$ (rule E-Ascribe) evaluates to $c$ trivially, without modifying the heap. The rules for sequencing (E-SEQ, E-SEQ1) are standard. Sequences of commands evaluate from left to right-i.e., the left-hand command is reduced to skip before the right-hand command is evaluated. The evaluation rules for conditionals are also standard. Conditionals reduce the boolean formula first (rule E-If) and if it evaluates to true, the then-branch (rule E-IfTrue) otherwise the else-branch (rule E-If False) is evaluated.

$$
\begin{aligned}
& \llbracket \tau \rrbracket_{\mathcal{E}} \subseteq \mathcal{P}(\mathcal{H}) \\
& \llbracket \varnothing \rrbracket_{\mathcal{E}}=\{ \} \\
& \llbracket\rceil \rrbracket_{\mathcal{E}}=\mathcal{H} \\
& \llbracket \tau_{1}+\tau_{2} \rrbracket_{\mathcal{E}}=\llbracket \tau_{1} \rrbracket_{\mathcal{E}} \cup \llbracket \tau_{2} \rrbracket_{\mathcal{E}} \\
& \llbracket \Sigma x: \tau_{1} \cdot \tau_{2} \rrbracket_{\mathcal{E}}=\left\{h_{1}+h_{2} \mid h_{1} \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}} \wedge h_{2} \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}\left[x \leftrightarrow h_{1}\right]}\right\} \\
& \llbracket \tau_{1}\left[x \mapsto \tau_{2}\right] \rrbracket_{\mathcal{E}}=\left\{h \mid h_{2} \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}} \wedge h \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}\left[x \mapsto h_{2}\right]}\right\} \\
& \llbracket\{x: \tau \mid e\} \rrbracket_{\mathcal{E}}=\left\{h \mid h \in \llbracket \tau \rrbracket_{\mathcal{E}} \wedge \llbracket e \rrbracket_{\mathcal{E}[x \mapsto h]}=\text { true }\right\}
\end{aligned}
$$

Figure 5.9: Semantics of heap types

$$
\left[\begin{array}{lll}
a & \mapsto & 1011 \\
p k t_{i n} & \mapsto & 1101 \\
p k t_{\text {out }} & \mapsto & \rangle
\end{array}\right]+\left[\begin{array}{lll}
b & \mapsto & 11 \\
p k t_{\text {in }} & \mapsto & 0 \\
p k t_{\text {out }} & \mapsto & 0000
\end{array}\right]=\left[\begin{array}{lll}
a & \mapsto & 1011 \\
b & \mapsto & 11 \\
p k t_{\text {in }} & \mapsto & 11010 \\
p k t_{\text {out }} & \mapsto & 0000
\end{array}\right]
$$

Figure 5.10: Example of a heap concatenation. The resulting heap contains the union of header instances while the input packet respectively output packet of the second heap is appended to the input packet respectively output packet of the first heap.

### 5.6 Static Semantics

$\Pi$ 4's type system is capable of capturing bit-level dependencies between header instances and the incoming and outgoing packet at any given program point. A heap $h$ in the set of heaps $\mathcal{H}$ describes such a possible system state, consisting of the incoming and outgoing packet and the set of valid header instances. We model heaps as maps from names to bit vectors. A heap contains for every valid header instance a mapping from the instance name to a bit vector, as well as two special entries $p k t_{i n}$ and $p k t_{\text {out }}$ representing the incoming and outgoing packet buffers.

Semantics of heap types Heap types $\tau$ represent sets of heaps, where each element in the set describes a different program path. Heap types are evaluated in an environment $\mathcal{E}$, which maps variables $x, y, z$ to heaps and bit variables $\mathrm{b}_{0}, \mathrm{~b}_{1}, \mathrm{~b}_{2}$ to single bits. The environment models other heaps available in the current scope upon which the current header type may depend.

The semantics of types is shown in Figure 5.9. The type $\varnothing$ denotes the empty set. It is used in situations where there are unsatisfiable assumptions involving the header instances or the incoming and outgoing packet buffers. The top type $T$ denotes the set of all possible heaps. The choice type $\tau_{1}+\tau_{2}$ denotes the union of the sets of heaps represented by $\tau_{1}$ and $\tau_{2}$. The dependent pair $\Sigma x: \tau_{1} . \tau_{2}$ denotes the concatenation of heaps from $\tau_{1}$ and $\tau_{2}$, where heaps described by $\tau_{2}$ may depend on heaps from $\tau_{1}$. The concatenation $h=h_{1}++h_{2}$ of two heaps $h_{1}$ and $h_{2}$ requires that header instances contained in $h_{1}$ and $h_{2}$ are disjoint. The resulting heap contains all instances from $h_{1}$ and from $h_{2}$, with $p k t_{\text {in }}$ and $p k t_{\text {out }}$ being the concatenation of respective bit vectors in $h_{1}$ and $h_{2}$ as exemplified in Figure 5.10.

$$
\begin{aligned}
\llbracket \varphi \rrbracket_{\mathcal{E}} & \in \mathbb{B} \\
\llbracket e_{1}=e_{2} \rrbracket_{\mathcal{E}} & =\llbracket e_{1} \rrbracket_{\mathcal{E}}=\llbracket e_{2} \rrbracket_{\mathcal{E}} \\
\llbracket e_{1}>e_{2} \rrbracket_{\mathcal{E}} & =\llbracket e_{1} \mathbb{I}_{\mathcal{E}}>\mathbb{e}_{2} \mathbb{Z}_{\mathcal{E}} \\
\llbracket \varphi_{1} \wedge \varphi_{2} \rrbracket_{\mathcal{E}} & =\llbracket \varphi_{1} \rrbracket_{\mathcal{E}} \wedge \varphi_{2} \rrbracket_{\mathcal{E}} \\
\llbracket \neg \varphi \rrbracket_{\mathcal{E}} & =\llbracket \varphi \rrbracket_{\mathcal{E}} \\
\llbracket x \text { x.।.valid } \rrbracket_{\mathcal{E}} & =\iota \in \operatorname{dom}(\mathcal{E}(x)) \\
\llbracket \text { true } \rrbracket & =\text { true } \\
\llbracket \text { false } \rrbracket & =\text { false }
\end{aligned}
$$

Figure 5.11: Semantics of formulae.

The explicit substitution $\tau_{1}\left[x \mapsto \tau_{2}\right]$ denotes the set of heaps obtained by evaluating $\tau_{1}$ for every heap described by $\tau_{2}$. Finally, the refinement type $\{x: \tau \mid \varphi\}$ denotes the set of heaps described by $\tau$ for which the predicate $\varphi$ holds. The semantics of the additional heap types we have defined as syntactic sugar is defined as follows:

$$
\begin{aligned}
\llbracket \epsilon \rrbracket_{\mathcal{E}} & =\{h \mid \forall \iota . \iota \notin \operatorname{dom}(h)\} \\
\llbracket \iota \rrbracket_{\mathcal{E}} & =\left\{h \mid \iota \in \operatorname{dom}(h) \wedge \forall \iota^{\prime} \in \operatorname{dom}(\mathcal{H} \mathcal{T}) . \iota^{\prime} \neq \iota \rightarrow \iota^{\prime} \notin \operatorname{dom}(h)\right\} \\
\llbracket \iota \approx \rrbracket_{\mathcal{E}} & =\{h \mid \iota \in \operatorname{dom}(h)\}
\end{aligned}
$$

Semantics of formulae The semantics of formulae is defined in Figure 5.11. Refinement predicates $\varphi$ are evaluated in the same type of environment as heap types. Formulae evaluate to a boolean value, i.e., $\llbracket \varphi \rrbracket_{\mathcal{E}} \in \mathbb{B}$. The semantics of expression equality ( $e_{1}=$ $e_{2}$ ) is defined as the semantic equality between expressions $e_{1}$ and $e_{2}$. Similarly, the semantics of expression comparison ( $e_{1}>e_{2}$ ) is defined as the semantic comparison between expressions $e_{1}$ and $e_{2}$. Instance validity (x.ı.valid) evaluates to true, if header instance $\iota$ is contained in the heap bound to $x$ in environment $\mathcal{E}$, otherwise it evaluates to false. The semantics of conjunction, negation and the boolean literals true and false are standard.

Semantics of expressions The semantics of expressions is defined in Figure 5.12. Expressions evaluate to either a bit vector or a natural number. Again, we use the same evaluation environment $\mathcal{E}$. The semantics of naturals and bit vectors is standard, except for bit variables $b_{n}$. In addition to bit literals $o$ and 1 , bit vectors can contain bit variables, which are looked up from the environment during evaluation. The semantic of addition is also standard. To evaluate the length of a packet $|x . p|$, we compute the length of the bit vector of $p k t_{\text {in }}$ or $p k t_{o u t}$ respectively in the heap bound to $x$ in the environment. The semantics of bit vector concatenation is as expected. If $\llbracket e_{1} \rrbracket \mathcal{E}=\left\langle b_{0}, \ldots, b_{n}\right\rangle$ and $\llbracket e_{2} \rrbracket_{\mathcal{E}}=\left\langle b_{n+1}, \ldots, b_{m}\right\rangle$, then $\llbracket e_{1} @ e_{2} \rrbracket=\left\langle b_{0}, \ldots, b_{n}, b_{n+1}, \ldots, b_{m}\right\rangle$.

A packet access $x . p$ looks up the respective entry from the heap bound to variable $x$ in $\mathcal{E}$. A packet slice $x . p[l: r]$ is evaluated in the same way, but additionally the designated slice is obtained from the bit vector. The semantics of instance slices $x . \iota[l: r]$ is defined similarly, but the lookup occurs on header instance $\iota$. Since the size and order of header fields is statically known, a field access $x . ו . f$ is just a named slice on a header instance. We interpret slices as half-open intervals, where the left bound is included and the right bound is excluded. For example, given a bit vector $b v=1010$ we have $b v[1: 4]=010$.

$$
\begin{aligned}
\llbracket e_{\mathcal{E}} & \in \mathrm{BV} \cup \mathbb{N} \\
\llbracket n \rrbracket_{\mathcal{E}} & =n \\
\llbracket x . p \rrbracket_{\mathcal{E}} & = \begin{cases}0 & \text { if } \mathcal{E}(x)(p)=\langle \rangle \\
n & \text { if } \mathcal{E}(x)(p)=\left\langle b_{1}, \ldots, b_{n}\right\rangle\end{cases} \\
\llbracket e_{1}+e_{2} \rrbracket_{\mathcal{E}} & =\llbracket e_{1} \rrbracket_{\mathcal{E}}+\llbracket e_{2} \rrbracket_{\mathcal{E}} \\
\llbracket e_{1} @ e_{2} \rrbracket_{\mathcal{E}} & =\llbracket e_{1} \rrbracket_{\mathcal{E}} \llbracket e_{2} \rrbracket_{\mathcal{E}} \\
\llbracket x . p \rrbracket_{\mathcal{E}} & =\left\{\left\langle b_{1}, \ldots, b_{n}\right\rangle\right.
\end{aligned} \quad \text { if } \mathcal{E}(x)(p)=\left\langle b_{1}, \ldots, b_{n}\right\rangle,
$$

$$
\left.\begin{array}{rlrl}
\llbracket b v \rrbracket_{\mathcal{E}} & \in \mathrm{BV} & & \\
\llbracket\left\rangle \rrbracket_{\mathcal{E}}\right. & =\langle \rangle & & \\
\llbracket 0:: b v \rrbracket_{\mathcal{E}} & =\left\langle\mathrm{o}, b_{1}, \ldots, b_{n}\right\rangle & & \text { if } \llbracket b v \rrbracket_{\mathcal{E}}
\end{array}=\left\langle b_{1}, \ldots, b_{n}\right\rangle\right)
$$

Figure 5.12: Semantics of expressions (top) and bit vectors (bottom).

Operations on heap types We define two semantic operations on heap types: inclusion and exclusion of instances. The first, Includes $\Gamma \tau \iota$, traverses $\tau$ and checks that instance $\iota$ is valid in every heap. Semantically this says that $\iota$ is a member of every element of $\llbracket \tau \rrbracket \mathcal{E}$-i.e., if $\mathcal{E} \vDash \Gamma$, then $\forall h \in \llbracket \tau \rrbracket_{\mathcal{E} . ~} \in \operatorname{dom}(h)$. The second, Excludes $\Gamma \tau \iota$, traverses $\tau$ and checks that instance $\iota$ is invalid in every heap. Semantically this says that $\iota$ is no member of every element of $\llbracket \tau \rrbracket_{\mathcal{E}}$-i.e., if $\mathcal{E} \vDash \Gamma$, then $\forall h \in \llbracket \tau \rrbracket \mathcal{E} . \iota \notin \operatorname{dom}(h)$.

Typing Judgment The typing judgement has the form $\Gamma \vdash c:\left(x: \tau_{1}\right) \rightarrow \tau_{2}$. Intuitively, type $\tau_{1}$ describes the input heap and $\tau_{2}$ describes the output heap obtained after the execution of command $c . \Gamma$ is a variable context that maps variable names to heap types and is used to capture additional dependencies of the input type. If a command typechecks in a context where $y$ maps to $\tau$ (i.e., $\left.\Gamma, y: \tau \vdash c:\left(x: \tau_{1}\right) \rightarrow \tau_{2}\right)$ it means that given some heap described by type $\tau$ on which the input heap might depend, executing $c$ on the input heap described by $\tau_{1}$ will result in a heap described by $\tau_{2}$.

Subtyping We write $\Gamma \vdash \tau_{1}<: \tau_{2}$ to denote the subtyping check between $\tau_{1}$ and $\tau_{2}$. Context $\Gamma$ captures external dependencies of $\tau_{1}$ and $\tau_{2}$ respectively. We define subtyping semantically as follows:

$$
\Gamma \vdash \tau_{1}<: \tau_{2} \triangleq \forall \mathcal{E} \cdot \mathcal{E} \vDash \Gamma \Longrightarrow \llbracket \tau_{1} \rrbracket_{\mathcal{E}} \subseteq \llbracket \tau_{2} \rrbracket_{\mathcal{E}}
$$

Type $\tau_{1}$ is a subtype of type $\tau_{2}$ in context $\Gamma$, if and only if for any environment $\mathcal{E}$ such that environment $\mathcal{E}$ entails context $\Gamma$, the set of heaps described by $\tau_{1}$ is a subset of the set of heaps described by $\tau_{2}$ both evaluated in environment $\mathcal{E}$.

$$
\begin{aligned}
& \begin{array}{lll}
\text { Ent-Top } & \begin{array}{l}
\text { Ent-ChoiceL } \\
(I, O, H) \vDash \mathcal{E} \tau_{1}
\end{array} & \begin{array}{l}
\text { Ent-ChoiceR } \\
(I, O, H) \vDash \mathcal{E} \tau_{2}
\end{array} \\
\hline(I, O, H) \vDash_{\mathcal{E}} \top & \begin{array}{ll}
(I, O, H) \vDash \mathcal{E} \tau_{1}+\tau_{2} & (I, O, H) \vDash_{\mathcal{E}} \tau_{1}+\tau_{2}
\end{array}
\end{array} \\
& \text { Ent-Refine Ent-Sigma } \\
& (I, O, H) \vDash_{\mathcal{E}} \tau \\
& \frac{\llbracket e \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]}=\text { true }}{(I, O, H) \vDash \mathcal{E}\{x: \tau \mid e\}} \quad \frac{\left(I_{2}, O_{2}, H_{2}\right) \vDash \mathcal{E}\left[x \leftrightarrow\left(I_{1}, O_{1}, H_{1}\right)\right] \tau_{2}}{\left(I_{1} @ I_{2}, O_{1} @ O_{2}, H_{1} \cup H_{2}\right) \vDash \mathcal{E} \sum x: \tau_{1} . \tau_{2}} \\
& \text { Ent-Subst } \\
& \left(I_{2}, O_{2}, H_{2}\right) \vDash_{\mathcal{E}} \tau_{2} \\
& \frac{(I, O, H) \vDash_{\mathcal{E}\left[x \mapsto\left(I_{2}, O_{2}, H_{2}\right)\right]} \tau_{1}}{(I, O, H) \vDash_{\mathcal{E}} \tau_{1}\left[x \mapsto \tau_{2}\right]}
\end{aligned}
$$

Figure 5.13: Entailment between heaps and heap types.

The entailment between environment $\mathcal{E}$ and typing context $\Gamma$ is formally defined as follows:

$$
\mathcal{E} \vDash \Gamma \triangleq \forall x_{i} \in \operatorname{dom}(\Gamma) \cdot \mathcal{E}\left(x_{i}\right)=h_{i} \wedge h_{i} \vDash \mathcal{E} \Gamma\left(x_{i}\right)
$$

An environment $\mathcal{E}$ entails a context $\Gamma$, iff for every mapping from a variable name $x_{i}$ to some heap type $\tau_{i}$ in $\Gamma$ there exists a mapping from variable $x_{i}$ to some heap $h_{i}$ in environment $\mathcal{E}$ and that heap $h_{i}$ entails type $\tau_{i}$. The entailment relation between a heap and a type is defined in Figure 5.13. A heap $H\left[p k t_{\text {in }} \mapsto I, p k t_{o u t} \mapsto O\right]$, in short $(I, O, H)$ entails a type $\tau$, if the heap $(I, O, H)$ is contained in the semantics of $\tau$.

Command typing The typing rules for commands are shown in Figure 5.14 and Figure 5.15. The typing rule T-Extract first captures that in order to execute an extract command, the input packet must provide enough bits to populate the instance. The predicate sizeof ${ }_{p k t_{i n}}(\tau) \geq n$ holds, if and only if the input packet in any heap described by type $\tau$ contains at least $n$ bits, i.e., sizeof ${ }_{p k t_{i n}}(\tau) \geq n$ iff $\forall \mathcal{E}, h \in \llbracket \tau \rrbracket \mathcal{E} .\left|h\left(p k t_{i n}\right)\right| \geq n$.

Further, rule T-Extract also captures that instance $\iota$ must be valid after an extract command is executed. When a packet header is extracted, the first $n$ bits-where $n$ is the number of bits contained in the header instance-are removed from the input packet and copied into the instance. We need to reflect this change accordingly. This is the task of the chomp refinement transformer, which transforms our input type to obtain a new type, which reflects this change. The chomp operator will be discussed in detail in Section 5.7.

Rule T-ADD first checks that the instance is not yet included in the type and assigns an output type that reflects that all information from the input type $\tau_{1}$ is retained and just instance $\iota$ is added, which is initialized with the value provided by initialization function init $\mathcal{H T}_{\mathcal{T}(\imath)}$.

To typecheck a modification of an instance field, the typing rule T-Mod first checks if the instance to be modified is guaranteed to be valid in the input type. Helper function $\mathcal{F}(\iota, f)$ returns the bit vector type for field $f$ in instance $\iota$. While in our formalization we do not explicitly distinguish different bit vector types, this allows for a more fine-grained check if the assigned expression is compatible with the header field. Instead of just

T-Extract
$\Gamma \vdash \operatorname{sizeof}_{p k t_{i n}}\left(\tau_{1}\right) \geq \operatorname{sizeof}(\iota) \quad \varphi_{1} \triangleq z . p k t_{\text {in }}=z . p k t_{\text {out }}=\langle \rangle$
$\frac{\varphi_{2} \triangleq y . \iota @ z . p k t_{\text {in }}=x . p k t_{\text {in }} \wedge z . p k t_{\text {out }}=x . p k t_{\text {out }} \wedge z \equiv, x}{\Gamma \vdash \operatorname{extract}(\iota):\left(x: \tau_{1}\right) \rightarrow \Sigma y:\left\{z: \iota \mid \varphi_{1}\right\} \cdot\left\{z: \operatorname{chomp}\left(\tau_{1}, \iota, y\right) \mid \varphi_{2}\right\}}$
T-Add

$$
\begin{gathered}
\text { Excludes } \Gamma \tau_{1} \iota \quad \text { init }_{\mathcal{H} \mathcal{T}(\iota)}=v \\
\varphi \triangleq z . p k t_{\text {in }}=z \cdot p k t_{\text {out }}=\langle \rangle \wedge z . \iota=v \\
\Gamma \vdash \operatorname{add}(\iota):\left(x: \tau_{1}\right) \rightarrow \Sigma y:\left\{z: \tau_{1} \mid z \equiv x\right\} .\{z: \iota \mid \varphi\}
\end{gathered}
$$

T-Mod

$$
\begin{gathered}
\text { Includes } \Gamma \tau_{1} \iota \quad \mathcal{F}(\iota, f)=\mathrm{BV} \quad \Gamma ; \tau_{1} \vdash e: \mathrm{BV} \\
\varphi_{p k t} \triangleq y \cdot p k t_{\text {in }}=x . p k t_{i n} \wedge y . p k t_{\text {out }}=x \cdot p k t_{\text {out }} \\
\varphi_{\iota} \triangleq \forall \kappa \in \operatorname{dom}(\mathcal{H T}) \cdot \kappa \neq \iota \Longrightarrow y . \kappa=x . \kappa \wedge \\
\varphi_{f} \triangleq \forall g \in \operatorname{dom}(\mathcal{H \mathcal { T }}(\iota)) \cdot g \neq f \Longrightarrow y . \iota . g=x . \iota . g \\
\Gamma \vdash \iota . f:=e:\left(x: \tau_{1}\right) \rightarrow\left\{y: \top \mid \varphi_{p k t} \wedge \varphi_{\iota} \wedge \varphi_{f} \wedge y . \iota . f=e[x / \text { heap }]\right\}
\end{gathered}
$$

T-Remove

$$
\begin{gathered}
\text { Includes } \Gamma \tau_{1} \iota \\
\varphi_{\iota} \triangleq \forall \kappa \in \operatorname{dom}(\mathcal{H T}) \cdot \kappa \neq \iota \Longrightarrow y \cdot \kappa=x . \kappa \\
\varphi_{p k t} \triangleq y \cdot p k t_{\text {in }}=x \cdot p k t_{\text {in }} \wedge y . p k t_{\text {out }}=x . p k t_{\text {out }} \\
\Gamma \vdash \text { remove }(\iota):\left(x: \tau_{1}\right) \rightarrow\left\{y: \top \mid \varphi_{\iota} \wedge \varphi_{p k t} \wedge \neg y . \iota . \text { valid }\right\}
\end{gathered}
$$

T-Remit

$$
\frac{\text { Includes } \Gamma \tau_{1} \iota \quad \varphi \triangleq z . p k t_{\text {in }}=\langle \rangle \wedge z \cdot p k t_{\text {out }}=x . \iota}{\Gamma \vdash \operatorname{remit}(\iota):\left(x: \tau_{1}\right) \rightarrow \Sigma y:\left\{z: \tau_{1} \mid z \equiv x\right\} \cdot\{z: \epsilon \mid \varphi\}}
$$

T-Reset

$$
\begin{gathered}
\varphi_{1} \triangleq z \cdot p k t_{\text {out }}=\langle \rangle \wedge z \cdot p k t_{\text {in }}=x \cdot p k t_{\text {out }} \\
\varphi_{2} \triangleq z \cdot p k t_{\text {out }}=\langle \rangle \wedge z \cdot p k t_{\text {in }}=x \cdot p k t_{\text {in }} \\
\Gamma \vdash \text { reset }:\left(x: \tau_{1}\right) \rightarrow \Sigma y:\left\{z: \epsilon \mid \varphi_{1}\right\} \cdot\left\{z: \epsilon \mid \varphi_{2}\right\}
\end{gathered}
$$

Figure 5.14: Typing rules for domain-specific commands.
checking if the assigned expression is of type bit vector, we can additionally check that the bit-size of the assigned expression matches the bit-size of the header field we are assigning to. The output type is similar to the strongest post-condition of the input type: everything in the output type is the same as in $x$, except for the modified instance field $y . l . f$, which must be equal to $e[x /$ heap $]$. Predicate $\varphi_{p k t}$ ensures that the input and output packet remain unchanged, predicate $\varphi_{\iota}$ ensures that all instances beside instance $\iota$ remain unchanged and predicate $\varphi_{f}$ ensures that all fields except for field $f$ on instance $\iota$ remain unchanged. Rule T-Remove first checks that the instance to be removed is valid in the input type. The output type reflects that all other instances besides instance $\iota$ as well as the input packet and output packet remain unchanged. To typecheck the command T-Remit, we check whether the instance to be emitted is guaranteed to be valid in the input type. The assigned output type ensures that emitting a header instance appends the value of the instance to the end of the outgoing packet (second projection of the assigned $\Sigma$-type) but leaves the input packet and all other

T-If

$$
\begin{gathered}
\Gamma ; \tau \vdash \varphi: \mathbb{B} \quad \Gamma \vdash c_{1}:\left(x:\left\{y: \tau_{1} \mid \varphi[y / \text { heap }]\right\}\right) \rightarrow \tau_{12} \\
\Gamma \vdash c_{2}:\left(x:\left\{y: \tau_{1} \mid \neg \varphi[y / \text { heap }]\right\}\right) \rightarrow \tau_{22} \\
\Gamma \vdash \operatorname{if}(\varphi) c_{1} \text { else } c_{2}:\left(x: \tau_{1}\right) \rightarrow\left\{y: \tau_{12} \mid \varphi[x / \text { heap }]\right\}+\left\{y: \tau_{22} \mid \neg \varphi[x / \text { heap }]\right\}
\end{gathered}
$$

T-SEQ
$\frac{\Gamma \vdash c_{1}:\left(x: \tau_{1}\right) \rightarrow \tau_{12} \quad \Gamma, x: \tau_{1} \vdash c_{2}:\left(y: \tau_{12}\right) \rightarrow \tau_{22}}{\Gamma \vdash c_{1} ; c_{2}:\left(x: \tau_{1}\right) \rightarrow \tau_{22}\left[y \mapsto \tau_{12}\right]}$

T-Skip $\frac{\tau_{2} \triangleq\left\{y: \tau_{1} \mid y \equiv x\right\}}{\Gamma \vdash \operatorname{skip}:\left(x: \tau_{1}\right) \rightarrow \tau_{2}}$

$$
\begin{aligned}
& \text { T-SUB } \\
& \quad \Gamma \vdash \tau_{1}<: \tau_{3}
\end{aligned}
$$

$$
\text { T-Ascribe } \quad \Gamma, x: \tau_{1} \vdash \tau_{4}<: \tau_{2}
$$

$$
\frac{\Gamma \vdash c: \sigma}{\Gamma \vdash c \text { as } \sigma: \sigma} \quad \frac{\Gamma \vdash c:\left(x: \tau_{3}\right) \rightarrow \tau_{4}}{\Gamma \vdash c:\left(x: \tau_{1}\right) \rightarrow \tau_{2}}
$$

Figure 5.15: Additional command typing rules.
validity information unchanged (first projection of the assigned $\Sigma$-type). Rule T-Reset resets all assumptions about header validity, empties the output packet $p k t_{\text {out }}$ and sets the input packet $p k t_{i n}$ to be the concatenation of $p k t_{\text {out }}$ and $p k t_{i n}$ of the input type. In the output type, we use a $\Sigma$-type to model the concatenation.

Rule T-IF typechecks each branch of the conditional with the additional assumption that the condition $\varphi$ holds respectively does not hold. The resulting type is a pathsensitive union type, which includes the types of both paths. By default, all variables in formula $\varphi$ in the command are bound to heap (cf. Section 5.3) and since we want to use the formula as a refinement in the type, we have to adjust the binders accordingly. To turn $\varphi$ into a refinement on a type, we substitute every occurrence of heap with the respective binder of the type we want to refine. We write $\varphi[x /$ heap $]$ to denote the formula obtained from $\varphi$ in which all variables heap are substituted with $x$. For example, if the command is if $($ ethernet.etherType $=0 x 0800)$ extract $(i p v 4)$ else skip, we typecheck the then-branch with type $\left(x:\left\{y: \tau_{1} \mid y\right.\right.$.ethernet.etherType $\left.\left.=0 \times 0800\right\}\right) \rightarrow \tau_{12}$. The full command is checked with type $\left(x: \tau_{1}\right) \rightarrow\left\{y: \tau_{12} \mid x\right.$.ethernet.etherType $=$ ox0800 $\}+\left\{y: \tau_{22} \mid \neg x\right.$.ethernet.etherType $\left.=0 \times 0800\right\}$.

The typing rule for sequencing T-SEQ is mostly standard, with one peculiarity: because our typing judgement assigns dependent function types to commands, the result type $\tau_{22}$ of command $c_{2}$ might depend on its input type $\tau_{12}$-i.e., variable $y$ might appear free in $\tau_{22}$. Hence, we must also capture the type $\tau_{12}$ in the result type. The typing rule T-SKip is also standard, except that it strictly enforces that the heaps described by the output type and input type respectively are equivalent. The typing rule for ascription T-Ascribe is standard. The typing rule for subsumption T-Sub is also standard. Since $\tau_{4}$ can depend on the input type $\tau_{3}$ and similarly, since $\tau_{2}$ can depend on $\tau_{1}$, we need to extend the typing context with the respective types. Because of the subtype relation between $\tau_{1}$ and $\tau_{3}$, all heaps that $\tau_{4}$ can depend on are also described by type $\tau_{1}$, thus for the subtyping check between types $\tau_{4}$ and $\tau_{2}$ we extend the type typing context with $\tau_{1}$.

$$
\begin{gathered}
\text { chomp : } \tau \times \iota \times \mathcal{X} \rightarrow \tau \\
\operatorname{chomp}(\tau, \iota, x) \triangleq \operatorname{chompRec}(\tau, \operatorname{sizeof}(\iota), x, \iota)
\end{gathered}
$$

where

$$
\begin{gathered}
\text { chompRec : } \tau \times \mathbb{N} \times \mathcal{X} \times \iota \rightarrow \tau \\
\operatorname{chompRec}(\tau, n, x, \iota) \triangleq \begin{cases}\tau & \text { if } n=\mathrm{o} \\
\text { let } \tau^{\prime}=\operatorname{heapRef}_{1}( & \text { o/w }\left(\mathrm{b}_{\mathrm{n}} \text { fresh }\right) \\
\operatorname{chomp} \\
\left.\left.\operatorname{chompec}\left(\tau, \mathrm{b}_{\mathrm{n}}\right), \mathrm{b}_{\mathrm{n}}, \iota, n\right) \text { in }, x, \iota\right)\end{cases}
\end{gathered}
$$

Figure 5.16: Definition of chomp.

$$
\begin{aligned}
& \operatorname{chomp}_{1}: \tau \times B_{n} \rightarrow \tau \\
& \operatorname{chomp}_{1}\left(\sum x: \tau_{1} \cdot \tau_{2}, \mathrm{~b}_{\mathrm{n}}\right) \triangleq \sum x: \operatorname{chomp}_{1}\left(\tau_{1}, \mathrm{~b}_{\mathrm{n}}\right) \cdot \operatorname{chompRef}_{1}\left(\tau_{2}, x, \mathrm{~b}_{\mathrm{n}}\right)+ \\
& \sum x:\left\{y: \tau_{1}| | y \cdot p k t_{\text {in }} \mid=\mathrm{o}\right\} \cdot \operatorname{chomp}_{1}\left(\tau_{2}, \mathrm{~b}_{\mathrm{n}}\right) \\
& \operatorname{chomp}_{1}\left(\tau_{1}+\tau_{2}, \mathrm{~b}_{\mathrm{n}}\right) \triangleq \operatorname{chomp}_{1}\left(\tau_{1}, \mathrm{~b}_{\mathrm{n}}\right)+\operatorname{chomp}_{1}\left(\tau_{2}, b_{\circ}\right) \\
& \operatorname{chomp}_{1}\left(\{x: \tau \mid e\}, \mathrm{b}_{\mathrm{n}}\right) \triangleq\left\{x: \operatorname{chomp}_{1}\left(\tau, \mathrm{~b}_{\mathrm{n}}\right) \mid \operatorname{chomp}_{1}^{\varphi}\left(e, x, \mathrm{~b}_{\mathrm{n}}\right)\right\} \\
& \operatorname{chomp}_{1}\left(\tau_{1}\left[x \mapsto \tau_{2}\right], \mathrm{b}_{\mathrm{n}}\right) \triangleq \operatorname{chomp}_{1}\left(\tau_{1}, \mathrm{~b}_{\mathrm{n}}\right)\left[x \mapsto \tau_{2}\right] \\
& \operatorname{chomp}_{1}\left(\tau, \_\right) \triangleq \tau
\end{aligned}
$$

Figure 5.17: Definition of chomp ${ }_{1}$.

### 5.7 Chomp

When an instance $\iota$ is extracted, sizeof $(\iota)$ bits are moved from the input bit stream to the instance. We want this process, which we call chomping, to be reflected in the type that we assign to an extract command. For example, given a header instance $y . A$ of type $A_{\eta}=\{f: 2\}$, and heap type $\tau=\left\{x: \epsilon \mid x . p k t_{i n}[0: 2]=11\right\}$, moving the first two bits of the input packet into instance $A$ is captured by the type $\{x: \epsilon \mid y \cdot A[0: 2]=11\}$. Chomping is a syntactic transformation of a heap type, which is defined as a bitwise operation consisting of (1) chomp $p_{1}$, the consumption of exactly one bit from $p k t_{i n}$ and (2) heapRef ${ }_{1}$, the refinement of the consumed bit to the extracted instance. Finally, chomp (c.f. Figure 5.16) lifts the pairwise application of chomp ${ }_{1}$ and heapRef ${ }_{1}$ to be applicable to whole instances.

### 5.7.1 Single-bit Chomp

For chomping off one bit from a heap type, we need to update references to the length, as well as to the first bit of $p k t_{i n}$. The removal of one bit from $p k t_{i n}$ through chomp ${ }_{1}$ resembles the computation of a Brzozowski derivative [Brz64]. We define 1 -bit chomping for heap types ( chomp $_{1}$ ), expressions (chomp ${ }_{1}^{e}$ ) and formulae (chomp $p_{1}^{\varphi}$ ).

As shown in Figure 5.17, function chomp $1_{1}$ takes two arguments, a heap type $\tau$ and a bit variable $b_{n}$ out of the set of bit variables $B_{n}$ and returns a heap type. Intuitively,

$$
\begin{aligned}
& \text { chompRef }: \tau \times \mathcal{X} \times B_{n} \rightarrow \tau \\
& \operatorname{chompRef}_{1}\left(\Sigma x: \tau_{1} \cdot \tau_{2}, x, \mathrm{~b}_{\mathrm{n}}\right) \triangleq \Sigma x: \operatorname{chompRef}_{1}\left(\tau_{1}, x, \mathrm{~b}_{\mathrm{n}}\right) \cdot \operatorname{chompRef}_{1}\left(\tau_{2}, x, \mathrm{~b}_{\mathrm{n}}\right) \\
& \operatorname{chompRef}_{1}\left(\tau_{1}+\tau_{2}, x, \mathrm{~b}_{\mathrm{n}}\right) \triangleq \operatorname{chompRef}_{1}\left(\tau_{1}, x, \mathrm{~b}_{\mathrm{n}}\right)+\operatorname{chompRef}_{1}\left(\tau_{2}, x, \mathrm{~b}_{\mathrm{n}}\right) \\
& \operatorname{chompRef}_{1}\left(\{x: \tau \mid \varphi\}, x, \mathrm{~b}_{\mathrm{n}}\right) \triangleq\left\{x: \operatorname{chompRef}_{1}\left(\tau, x, \mathrm{~b}_{\mathrm{n}}\right) \mid \operatorname{chomp}_{1}^{\varphi}\left(\varphi, x, \mathrm{~b}_{\mathrm{n}}\right)\right\} \\
& \operatorname{chompRef}_{1}\left(\tau_{1}\left[y \mapsto \tau_{2}\right], x, \mathrm{~b}_{\mathrm{n}}\right) \triangleq \operatorname{chompRef}_{1}\left(\tau_{1}, x, \mathrm{~b}_{\mathrm{n}}\right)\left[y \mapsto \operatorname{chompRef}\left(\tau_{2}, x, \mathrm{~b}_{\mathrm{n}}\right)\right] \\
&\left.\operatorname{chompRef}_{1}\left(\tau,,_{-}\right)\right) \triangleq \tau
\end{aligned}
$$

Figure 5.18: Definition of chompRef ${ }_{1}$.
chomp ${ }_{1}$ removes the first bit of the input packet and replaces all references to this bit with bit variable $b_{n}$. Syntactically, when chomping a heap type $\tau$ we need to update each occurrence of $p k t_{i n}$ in a refinement, if that occurrence describes the first bit of $p k t_{i n}$ of a heap in the semantics of $\tau$.

Types $\varnothing$ and T , are not affected by chomping. For a choice type $\tau=\tau_{1}+\tau_{2}$ chomp ${ }_{1}$ is applied to both types $\tau_{1}$ and $\tau_{2}$ individually, as each branch of the choice type is describing isolated heaps of $\tau$. In the substitution type $\tau=\tau_{1}\left[x \mapsto \tau_{2}\right]$ only $\tau_{1}$ is chomped, as $\tau_{2}$ only captures information relevant for evaluating refinements. In the refinement type $\tau=\left\{x: \tau_{1} \mid \varphi\right\} \tau_{1}$ is chomped as well as the formula $\varphi$. The binder $x$ is used in the chomping of the expression, as $x$ is the latest leftmost binder for $p k t_{i n}$.

Chomping a sigma type $\tau=\Sigma x: \tau_{1} \cdot \tau_{2}$ is a bit more involved than the previous cases, because of the concatenation semantics of $p k t_{i n}$. Since we only chomp off exactly one bit, we need to consider two cases. First, the input packet in the left projection $\tau_{1}$ of the $\Sigma$-type contains at least one bit and second, the input packet of the first projection is empty, i.e., the length is equal to zero. In the first case, a single bit is removed from $\tau_{1}$, which requires that all references to $x . p k t_{i n}$ in $\tau_{2}$ must be updated as $\tau_{1}$ is bound to $x$ in $\tau_{2}$. This is the responsibility of function chompRef. Otherwise, chomping could cause contradictions between refinements referencing the same component.

For example, let us consider the heap type $\Sigma x:\left\{y: \epsilon| | y . p k t_{i n} \mid=1\right\} .\{z: \epsilon \mid$ $\left.\left|x . p k t_{i n}\right|=1\right\}$, where both refinements reference the first projection. To obtain an updated type that reflects the removal of the first bit of the input packet, chomp must update both refinements accordingly, i.e., chomp $_{1}\left(\Sigma x:\left\{y: \epsilon| | y \cdot p k t_{i n} \mid=1\right\} .\{z: \epsilon \mid\right.$ $\left.\left.\left|x \cdot p k t_{i n}\right|=1\right\}, \mathrm{b}_{0}\right)=\Sigma x:\left\{y: \epsilon| | y \cdot p k t_{i n} \mid+1=1\right\} .\left\{z: \epsilon| | x \cdot p k t_{i n} \mid+1=1\right\}$.

In the second case, we remove the bit from the input packet of the second projection. When we chomp in $\tau_{2}$ we use a refinement to assert the input packet $p k t_{i n}$ of $\tau_{1}$ is actually empty (e.g. $\operatorname{chomp}_{1}\left(\Sigma x: \epsilon .\left\{y: \epsilon| | x \cdot p k t_{i n} \mid=0\right\}, \mathrm{b}_{0}\right)=\Sigma x:\left\{y: \epsilon| | y \cdot p k t_{i n} \mid=\right.$ $\left.o\} .\left\{y: \epsilon| | x \cdot p k t_{i n} \mid=o\right\}\right)$. Finally, we combine both possibilities with a choice type.

As defined in Figure 5.18, the function chompRef ${ }_{1}$ takes three arguments, a heap type $\tau$, a variable $x$ out of the set of all variables $\mathcal{X}$ and a bit variable $b_{n}$. Intuitively, chompRef ${ }_{1}$ updates all references to $x . p k t_{i n}$ in $\tau$ and replaces all references to this bit with bit variable $b_{0}$. It is defined recursively on all heap types and for refinement types it passes on the execution to chomp ${ }_{1}^{\varphi}$.

Function chomp ${ }_{1}^{\varphi}$ (cf. Figure 5.19) takes three arguments, a formula $\varphi$, a variable $x$ and a bit variable $b_{n}$. Intuitively, chomp $p_{1}^{\varphi}$ updates via function chomp ${ }_{1}^{e}$ all expressions referencing $x . p k t_{i n}$ in formula $\varphi$ such that references to the first bit of $x . p k t_{i n}$ are replaced with bit variable $b_{0}$ and references to the length of $x . p k t_{i n}$ are increased by one. Formulae true, false and $x$.ı.valid are not affected by chomping and negation $\neg e$

$$
\begin{aligned}
& \operatorname{chomp}_{1}^{\varphi}: \varphi \times \mathcal{X} \times B_{n} \rightarrow \varphi \\
& \operatorname{chomp}_{1}^{\varphi}\left(e_{1}=e_{2}, x, \mathrm{~b}_{\mathrm{n}}\right) \triangleq \operatorname{chomp}_{1}^{e}\left(e_{1}, x, \mathrm{~b}_{\mathrm{n}}\right)=\operatorname{chomp}_{1}^{e}\left(e_{2}, x, \mathrm{~b}_{\mathrm{n}}\right) \\
& \operatorname{chomp}_{1}^{\varphi}\left(e_{1}>e_{2}, x, \mathrm{~b}_{\mathrm{n}}\right) \triangleq \operatorname{chomp}_{1}^{e}\left(e_{1}, x, \mathrm{~b}_{\mathrm{n}}\right)>\operatorname{chomp}_{1}^{e}\left(e_{2}, x, \mathrm{~b}_{\mathrm{n}}\right) \\
& \operatorname{chomp}_{1}^{\varphi}\left(\varphi_{1} \wedge \varphi_{2}, x, \mathrm{~b}_{\mathrm{n}}\right) \triangleq \operatorname{chomp}_{1}^{\varphi}\left(\varphi_{1}, x, \mathrm{~b}_{\mathrm{n}}\right) \wedge \operatorname{chomp}^{\varphi}\left(\varphi_{2}, x, \mathrm{~b}_{\mathrm{n}}\right) \\
& \operatorname{chomp}_{1}^{\varphi}\left(\neg \varphi, x, \mathrm{~b}_{\mathrm{n}}\right) \triangleq \neg \operatorname{chomp}_{1}^{\varphi}\left(\varphi, x, \mathrm{~b}_{\mathrm{n}}\right) \\
& \operatorname{chomp}_{1}^{\varphi}\left(\varphi,,_{-}\right) \triangleq \varphi
\end{aligned}
$$

Figure 5.19: Definition of chomp ${ }_{1}^{\varphi}$.

$$
\begin{aligned}
& \text { chomp }_{1}^{e}: e \times \mathcal{X} \times B_{n} \rightarrow e \\
& \operatorname{chomp}_{1}^{e}\left(\left|x \cdot p k t_{i n}\right|, y,_{-}\right) \triangleq \begin{cases}\left|x \cdot p k t_{i n}\right|+1 & \text { if } x=y \\
\left|x \cdot p k t_{i n}\right| & \text { otherwise }\end{cases} \\
& \operatorname{chomp}_{1}^{e}\left(x \cdot p k t_{i n}, y, \mathrm{~b}_{\mathrm{n}}\right) \triangleq \begin{cases}\left\langle\mathrm{b}_{\mathrm{n}}\right\rangle @ x . p k t_{i n} & \text { if } x=y \\
x \cdot p k t_{i n} & \text { otherwise }\end{cases} \\
& \operatorname{chomp}_{1}^{e}\left(x \cdot p k t_{i n}[l: r], y, \mathrm{~b}_{\mathrm{n}}\right) \triangleq \begin{cases}x \cdot p k t_{i n}[l: r] & \text { if } x \neq y \\
\left\langle\mathrm{~b}_{\mathrm{n}}\right\rangle & \text { if } x=y \wedge r \leq 1 \\
\left\langle\mathrm{~b}_{n}\right\rangle @ x \cdot p k t_{i n}[\mathrm{o}: r-1] & \text { if } x=y \wedge l=0 \\
x . p k t_{i n}[l-1: r-1] & \text { if } x=y \wedge l \neq 0\end{cases} \\
& \operatorname{chomp}_{1}^{e}\left(n+m, y, \mathrm{~b}_{\mathrm{n}}\right) \triangleq \operatorname{chomp}_{1}^{e}\left(n, y, \mathrm{~b}_{\mathrm{n}}\right)+\operatorname{chomp}_{1}^{e}\left(m, y, \mathrm{~b}_{\mathrm{n}}\right) \\
& \operatorname{chomp}_{1}^{e}\left(b v_{1} @ b v_{2}, y, \mathrm{~b}_{\mathrm{n}}\right) \triangleq \operatorname{chomp}_{1}^{e}\left(b v_{1}, y, \mathrm{~b}_{\mathrm{n}}\right) @ \operatorname{chomp}_{1}^{e}\left(b v_{2}, y, \mathrm{~b}_{\mathrm{n}}\right) \\
& \operatorname{chomp}_{1}^{e}\left(e,,_{-}\right) \triangleq e
\end{aligned}
$$

Figure 5.20: Definition of chomp ${ }_{1}^{e}$.
and conjunction $e_{1} \wedge e_{2}$ are standard congruence rules. For expression equality $e_{1}=e_{2}$ and expression comparison $e_{1}>e_{2}$, chomp ${ }_{1}^{e}$ is applied to both expressions individually.

Figure 5.20 defines how expressions are transformed as part of chomp ${ }_{1}^{e}$. Function chomp ${ }_{1}^{e}$ takes three arguments, an expression $e$, a variable $x$ and a bit variable $\mathrm{b}_{\mathrm{n}}$. Intuitively, chomping expressions with chomp ${ }_{1}^{e}\left(e, x, \mathrm{~b}_{n}\right)$ replaces references to the first bit of $x . p k t_{i n}$ in $e$ with $b_{n}$ and references to the length of $x . p k t_{i n}$ in $e$ are incremented by one. Expressions $e$ that do not reference $x . p k t_{i n}$ or $\left|x . p k t_{i n}\right|$, like $x . p k t_{o u t}$, numerals $n$ and bit vectors $b v$ are not affected by chomping, thus chomp ${ }_{1}^{e}\left(e, x, \mathrm{~b}_{\mathrm{n}}\right)=e$. If the expression is $x . p k t_{i n}$ or $\left|x . p k t_{i n}\right|$, but the binder does not match with the variable passed to chomp $1_{1}^{e}$, the expression remains unchanged. The cases $e_{1} @ e_{2}$ and $e_{1}+e_{2}$ are standard congruence rules. We increment expressions $\left|x . p k t_{i n}\right|$ by one if the binder matches with the variable argument passed to chomp ${ }_{1}^{e}$. If $e=x . p k t_{i n}$, we prepend bit variable $\mathrm{b}_{\mathrm{n}}$, which effectively removes the first bit. For input packet slices (x.pktin $[l: r]$ ), the result of chomp ${ }_{1}^{e}$ depends on the bounds $l$ and $r$. If the right bound is less or equal to one, the slice contains only a single bit, so we replace the whole slice with a bit vector only containing bit variable $\mathrm{b}_{\mathrm{n}}$ (e.g., chomp $\left._{1}^{e}\left(x . p k t_{i n}[\mathrm{o}: 1], x, \mathrm{~b}_{\mathrm{n}}\right)=\left\langle\mathrm{b}_{\mathrm{n}}\right\rangle\right)$. If the left bound is zero, the bit variable $b_{n}$ is prepended and the size of the slice is reduced by
one to retain the overall size of the expression (e.g., $\operatorname{chomp}_{1}^{e}\left(x . p k t_{i n}[0: 8], x, \mathrm{~b}_{\mathrm{n}}\right)=$ $\left.\left\langle\mathrm{b}_{\mathrm{n}}\right\rangle @ x . p k t_{i n}[\mathrm{o}: 7]\right)$. If the left bound is greater than zero, both bounds are decremented by one (e.g., chomp $\left.{ }_{1}^{e}\left(x . p k t_{i n}[4: 8], x, \mathrm{~b}_{0}\right)=x . p k t_{i n}[3: 7]\right)$.

Example Given type $\tau=\left\{x: \iota| | x \cdot p k t_{i n} \mid=8 \wedge x \cdot p k t_{i n}[0: 8]=x . \iota[4: 12]\right\}$, $\operatorname{chomp}_{1}\left(\tau, \mathrm{~b}_{0}\right)=\left\{x:\left\{y: \iota| | y \cdot p k t_{i n} \mid+1=8\right\} \mid \mathrm{b}_{0}:: x . p k t_{i n}[\mathrm{o}: 7]=x . \iota[4: 12]\right\}$

### 5.7.2 Instance Refinement

With heapRef $_{1}$ we replace the placeholder bits introduced by chomp ${ }_{1}$ with references to the extracted instance. As presented in Figure 5.21, the function heapRef takes five $^{\text {ta }}$ arguments, a heap type $\tau$ (we define the same function also for formulae and expressions), a bit variable $b_{n}$, a variable $x$, a header instance $\iota$ and a number $n$. Intuitively, heapRef $_{1}\left(\tau, \mathrm{~b}_{\mathrm{n}}, x, l, n\right)$ replaces $\mathrm{b}_{\mathrm{n}}$ in $\tau$ with $x . \iota[m: m+1]$, where $m=\operatorname{sizeof}(\iota)-n$.

The only place where a bit variable can occur according to the syntax of $\Pi_{4}$ is in the bit vector construction. Each case except for expressions $b:: b v$ will either be the identity if the input is a value or a congruence rule if the input has sub-nodes. For a expression $b:: b v(1)$ if $b \neq \mathrm{b}_{\mathrm{n}}$, we keep $b$ and continue searching in $b v$ or (2) if $b=\mathrm{b}_{\mathrm{n}}$, we replace $b_{n}$ with $x . l[\operatorname{sizeof}(\iota)-n]$ and continue searching in $b v$.

Example Given header instance $A$ of type $A_{\eta}=\{f: 2\}$ and heap type $\tau=\{x: \epsilon \mid$ $\left.\left\langle b_{0}\right\rangle @ x . p k t_{i n}[\mathrm{o}]=10\right\}$, $\operatorname{heapRef}_{1}\left(\tau, b_{0}, y, A, 2\right)=\left\{x: \epsilon \mid y \cdot A[\mathrm{o}] @ x . p k t_{i n}[\mathrm{o}]=10\right\}$.

### 5.7.3 Correctness of Chomp

We can prove that our definition of the chomp operator (cf. Figure 5.16) has the desired semantics. We first define a semantic chomp operation chomp ${ }^{\Downarrow}$ that-given a heap $h$ and a number $n$-removes the first $n$ bits from the input packet in heap $h$ :

$$
\operatorname{chomp}^{\Downarrow}(h, n) \triangleq h\left[p k t_{i n} \mapsto h\left(p k t_{i n}\right)[n:]\right]
$$

We can then prove the following lemma, which states that-given some heap $h \in$ $\llbracket \tau \rrbracket \mathcal{E}$-there exists a corresponding heap $h^{\prime}$ in the semantics of the chomped type that is equivalent to heap obtained after applying chomp ${ }^{\Downarrow}$ to heap $h$. Since chomp adds a refinement on $x . l$, we have to evaluate the chomped type in an environment, where $x$ maps to the heap in which $\iota$ contains the first sizeof $(\iota)$ bits from $h\left(p k t_{i n}\right)$. This corresponds with the intuition that chomp populates the header instance $\iota$ with the first sizeof $(\iota)$ bits from the input packet.

Lemma 5.1 (Semantic Chomp). If $x$ does not appear free in $\tau$, then for all heaps $h \in$ $\llbracket \tau \rrbracket \mathcal{E}$ where $\left|h\left(p k t_{i n}\right)\right| \geq \operatorname{sizeof}(t)$, there exists $h^{\prime} \in \llbracket \operatorname{chomp}(\tau, l, x) \rrbracket_{\mathcal{E}^{\prime}}$ such that $h^{\prime}=$ $\operatorname{chomp}^{\Downarrow}(h$, sizeof $(\iota))$ where $\mathcal{E}^{\prime}=\mathcal{E}\left[x \mapsto\left(\langle \rangle,\langle \rangle,\left[\iota \mapsto h\left(p k t_{\text {in }}\right)[\mathrm{o}: \operatorname{sizeof}(\iota)]\right]\right)\right]$.

Proof. By unfolding the definition of chomp and by induction on the number of bits consumed. The full proof can be found in Appendix A.2.1 (Lemma A.44).

### 5.8 Safety

We prove safety of $\Pi_{4}$ in terms of standard progress and preservation theorems. That is, well-typed programs do not get stuck and when well-typed programs are evaluated,

$$
\begin{aligned}
& \text { heapRef }_{1}: \tau \times \mathrm{b}_{\mathrm{n}} \times \mathcal{X} \times \iota \times \mathbb{N} \rightarrow \tau \\
& \operatorname{heapRef}_{1}\left(\Sigma x: \tau_{1} \cdot \tau_{2}, \mathrm{~b}_{\mathrm{n}}, y, l, n\right) \triangleq \Sigma x: \operatorname{heapRef}_{1}\left(\tau_{1}, \mathrm{~b}_{\mathrm{n}}, y, l, n\right) \text {. } \\
& \text { heapRef }_{1}\left(\tau_{2}, \mathrm{~b}_{\mathrm{n}}, y, l, n\right) \\
& \operatorname{heapRef}_{1}\left(\tau_{1}+\tau_{2}, \mathrm{~b}_{\mathrm{n}}, y, \iota, n\right) \triangleq \operatorname{heapRef}_{1}\left(\tau_{1}, \mathrm{~b}_{\mathrm{n}}, y, \iota, n\right)+ \\
& \text { heapRef }_{1}\left(\tau_{2}, \mathrm{~b}_{\mathrm{n}}, y, \iota, n\right) \\
& \operatorname{heapRef}_{1}\left(\{x: \tau \mid \varphi\}, \mathrm{b}_{n}, y, l, n\right) \triangleq\left\{x: \operatorname{heapRef}_{1}\left(\tau, \mathrm{~b}_{n}, y, l, n\right) \mid\right. \\
& \text { heapRef } \left._{1}\left(\varphi, \mathrm{~b}_{\mathrm{n}}, y, \iota, n\right)\right\} \\
& \text { heapRef }_{1}\left(\tau_{1}\left[x \mapsto \tau_{2}\right], b_{n}, y, l, n\right) \triangleq \operatorname{heapRef}_{1}\left(\tau_{1}, b_{n}, y, l, n\right)[x \mapsto \\
& \text { heapRef } \left._{1}\left(\tau_{2}, \mathrm{~b}_{n}, y, l, n\right)\right] \\
& \operatorname{heapRef}_{1}(\tau, \ldots,,-,-) \triangleq \tau \\
& \text { heapRef }{ }_{1}: \varphi \times \mathrm{b}_{\mathrm{n}} \times \mathcal{X} \times \iota \times n \rightarrow \varphi \\
& \operatorname{heapRef}_{1}\left(e_{1}=e_{2}, \mathrm{~b}_{\mathrm{n}}, x, l, n\right) \triangleq \operatorname{heapRef}_{1}\left(e_{1}, \mathrm{~b}_{\mathrm{n}}, x, l, n\right)=\operatorname{heapRef}_{1}\left(e_{2}, \mathrm{~b}_{\mathrm{n}}, x, l, n\right) \\
& \operatorname{heapRef}_{1}\left(e_{1}>e_{2}, b_{n}, x, l, n\right) \triangleq \operatorname{heapRef}_{1}\left(e_{1}, \mathrm{~b}_{\mathrm{n}}, x, l, n\right)>\operatorname{heapRef}_{1}\left(e_{2}, \mathrm{~b}_{\mathrm{n}}, x, l, n\right) \\
& \operatorname{heapRef}_{1}\left(\varphi_{1} \wedge \varphi_{2}, \mathrm{~b}_{\mathrm{n}}, x, \iota, n\right) \triangleq \operatorname{heapRef}_{1}\left(\varphi_{1}, \mathrm{~b}_{\mathrm{n}}, x, l, n\right) \wedge \operatorname{heapRef}_{1}\left(\varphi_{2}, \mathrm{~b}_{\mathrm{n}}, x, l, n\right) \\
& \operatorname{heapRef}_{1}\left(\neg \varphi, \mathrm{~b}_{\mathrm{n}}, x, \iota, n\right) \triangleq \neg \operatorname{heapRef}_{1}\left(\varphi, \mathrm{~b}_{\mathrm{n}}, x, l, n\right) \\
& \operatorname{heapRef}_{1}(\varphi,-,-,,-) \triangleq \varphi
\end{aligned}
$$

$$
\begin{aligned}
& \text { heapRef }_{1}: e \times b_{n} \times \mathcal{X} \times \iota \times n \rightarrow e \\
& \operatorname{heapRef}_{1}\left(b:: b v, \mathrm{~b}_{\mathrm{n}}, x, \iota, n\right) \triangleq\left\{\begin{array}{cl}
x . \iota[\operatorname{sizeof}(\iota)-n: \operatorname{sizeof}(\iota)-n+1] @ & \\
\operatorname{heapRef}_{1}\left(b v, \mathrm{~b}_{\mathrm{n}}, x, \iota, n\right) & \text { if } b=\mathrm{b}_{\mathrm{n}} \\
\langle b\rangle @ \text { heapRef } \\
\left(b v, \mathrm{~b}_{\mathrm{n}}, x, \iota, n\right) & \text { otherwise }
\end{array}\right. \\
& \operatorname{heapRef}_{1}\left(e,,_{-},-\right) \triangleq e
\end{aligned}
$$

Figure 5.21: Definition of heapRef ${ }_{1}$.
they remain well typed. Both theorems make use of the entailment relation defined in Figure 5.13.

Theorem 5.2 (Progress). If. $\vdash c:\left(x: \tau_{1}\right) \rightarrow \tau_{2}$ and $(I, O, H) \vDash \tau_{1}$, then either $c=$ skip or $\exists\left\langle I^{\prime}, O^{\prime}, H^{\prime}, c^{\prime}\right\rangle .\langle I, O, H, c\rangle \rightarrow\left\langle I^{\prime}, O^{\prime}, H^{\prime}, c^{\prime}\right\rangle$.

Proof. By induction on the typing derivation. The full proof can be found in Appendix A.2.1 (Theorem A.35).

The progress theorem states that if a command is well-typed, it is either skip or it can take a step.
Theorem 5.3 (Preservation). If $\Gamma \vdash c:\left(x: \tau_{1}\right) \rightarrow \tau_{2}$ and $\langle I, O, H, c\rangle \rightarrow\left\langle I^{\prime}, O^{\prime}, H^{\prime}, c^{\prime}\right\rangle$ and $\mathcal{E} \vDash \Gamma$ and $(I, O, H) \vDash \mathcal{E} \tau_{1}$, then there exists $\Gamma^{\prime}, \mathcal{E}^{\prime}, x^{\prime}, \tau_{1}^{\prime}, \tau_{2}^{\prime}$, such that $\Gamma^{\prime} \vdash c^{\prime}$ :
$\left(x^{\prime}: \tau_{1}^{\prime}\right) \rightarrow \tau_{2}^{\prime}$ and $\mathcal{E}^{\prime} \vDash \Gamma^{\prime}$ and $\Gamma \subseteq \Gamma^{\prime}$ and $\mathcal{E} \subseteq \mathcal{E}^{\prime}$ and $\left(I^{\prime}, O^{\prime}, H^{\prime}\right){\vDash \mathcal{E}^{\prime}} \tau_{1}^{\prime}$ and $\left.\llbracket \tau_{2}^{\prime}\right]_{\mathcal{E}}{ }^{\prime}\left[x^{\prime} \mapsto\left(I^{\prime}, O^{\prime}, H^{\prime}\right)\right] \subseteq \llbracket \tau_{2} \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]}$

Proof. By induction on the typing derivation. The full proof can be found in Appendix A.2.1 (Theorem A.54).

Preservation says that if a command $c$ is well-typed, command $c$ can step to $c^{\prime}$, and the starting heap entails input type $\tau_{1}$, then $c^{\prime}$ is well-typed with type $\left(x^{\prime}: \tau_{1}^{\prime}\right) \rightarrow \tau_{2}^{\prime}$ for some $x^{\prime}, \tau_{1}^{\prime}$ and $\tau_{2}^{\prime}$ s.t. the stepped heap entails $\tau_{1}^{\prime}$ and the set of heaps described by $\tau_{2}^{\prime}$ is a subset of the heaps described by $\tau_{2}$.

### 5.9 Related Work

Formal Reasoning for $\mathbf{P}_{4}$ Programs A number of different approaches for verifying properties of P 4 programs have been proposed in recent years. p4v [Liu+18] applies classical techniques based on predicate transformer semantics to achieve monolithic verification of $\mathrm{P}_{4}$ programs. In contrast, $\Pi_{4}$ uses dependent types and offers compositional verification. Vera $[\mathrm{Sto}+18]$ and Assert- $\mathrm{P} 4[\mathrm{Fre}+18 ; \mathrm{Nev}+18]$ are symbolic execution engines for $\mathrm{P}_{4}$. The bf4 tool [Dum+2o] follows the approach pioneered in p4v, but also attempts to infer control-plane constraints that are sufficiently strong to establish correctness, and offers heuristics for repairing programs when verification fails. P4K [KR18] provides a formal semantics of $\mathrm{P}_{4}$ in the K framework $\left[\mathrm{R} \mathrm{S}_{1} 10\right.$ ] and thus, can make use of the verification tools provided by the K framework. Petr4 [Doe+21] develops a formal semantics for $\mathrm{P}_{4}$ but does not itself offer verification tools. In contrast, P4RL [Shu+19] uses a dynamic approach-fuzz testing-for the verification of $\mathrm{P}_{4}$ programs.

Dependent Types There is a long history of using dependent types to capture properties of low-level code. Early work by Xi and Pfenning [XP99] showed how dependent types could be used to eliminate run-time safety checks-e.g., array bounds checks in imperative programs. Xanadu [Xioo] adds dependent typing to imperative programming, but does not capture the effect of mutations in the type. Xi and Harper later showed how dependent types could be applied to assembly code [XHor].

Deputy [Con+07] used dependent types to reason about complex, heap-allocated data structures. Similar to Deputy, П4's typing rule for modification of header fields is also inspired by the Hoare axiom for assignment. $\Pi_{4}$ is different in that type checking has no effect on the run time, and it also supports path-sensitive reasoning.

Similar to $\Pi_{4}$, Hoare Type Theory (HTT) [NMBo6] statically tracks how the heap evolves during execution. Typing of computations in HTT is similar to the dependent function types $\Pi_{4}$ uses for commands. The type captures the state before and after execution, possibly relating the output type with the input type. In our domain, this requires bit-by-bit transformations on the input type, provided by chomp. Other type systems like Ynot [Nan+o8], FCSL [Nan+14], and F* [Swa+16] provide dependent types for low-level imperative programming. While these type systems target general functional verification and often require manual programs-as-proofs to do so, our type system is designed with domain-specific properties of network programming in mind and is fully automatic.

Solver-Aided Tools A key focus of recent work on dependently-typed language has been on automation. This work builds on recent advances in SAT/SMT solvers and is designed to make dependent types usable by ordinary programs. A prominent example
is Liquid Haskell [RKJo8]. It extends Haskell with refinement types, but imposes restrictions to ensure the refinements remain decidable. Under the hood, all proof obligations generated by Liquid Haskell during type checking are handled by an SMT solver in a way that is transparent to the programmer. $\Pi_{4}$ draws inspiration from Liquid Haskell's decidable refinement types. However, our SMT encoding and our proof of correctness and decidability are novel. Liquid Haskell stipulates that its refinements must be in the theory of quantifier-free integer linear arithmetic in order to be decidable. We encode types into the effectively propositional fragment of first-order logic over bit vectors, which facilitates automatic subtyping and equivalence checks.

The Prototype Verification System (PVS) [ORS92; Owr+95] is an interactive theorem prover that similar to $\Pi_{4}$ combines dependent types and refinement types in its specification language. Also, PVS automatically extracts and solves proof obligations during type checking, but in contrast to $\Pi_{4}$, type checking is undecidable, which requires the user to get involved in complex cases.

Formalizing Protocols Another line of work has developed language-based specifications of protocols. CMU's FoxNet project used SML to specify the behavior of an entire networking stack [Bia+94]. McCann and Chandra used a type-based approach to give abstract specifications of protocols [MCoo]. Grammar-based tools such as PADS [FGo5], Narcissus[Del+19], and Yakker [JMW1o], enable specifying the syntax of complex, dependent formats including network protocols, and also provide tools for serializing and deserializing data.

### 5.10 Chapter Summary

In this chapter, we introduced $\Pi_{4}$, a dependently-typed version of the $\mathrm{P}_{4}$ language that aims at closing the gap between type-system-based approaches and full-fledged verification tools while enabling modular verification.

Featuring a combination of expressive types including refinement types, dependent pairs, union types and explicit substitutions, $\Pi_{4}$ is able to capture precise assumptions about valid header instances and the incoming and outgoing packet as well as dependencies between them down to the bit level. Dependent function types make it possible to statically capture the effect of executing a command in a certain program state and to relate both program states before and after the execution. $\Pi_{4}$ addresses the challenge of retaining precise types in the presence of packet deserialization, one central domainspecific feature of the $\mathrm{P}_{4}$ language, by resorting to the novel chomp operator, which computes a type that remains after extracting bits from a packet buffer.
$\Pi_{4}$ relies on type ascription to enable modular checking of programs. By ascribing dependent function types, it is possible to describe the requirements a command has on the context in which it is executed as well as the guarantees the command provides to the outside via the input type and output type respectively. Each ascribed type indicates a program point where type checking should occur, thus, ascribed commands can be independently type-checked. It is then sufficient to rely on the ascribed output type to check the rest of the program. As a result, as we will see in Chapter 8, $\Pi_{4}$ is capable of expressing and verifying a variety of rich network properties in a modular way.

## CHAPTER

 6
## An Implementation of $\Pi_{4}$

We have built a prototype implementation of $\Pi_{4}$ in OCaml and $Z_{3}$ [MBo8], which comprises approximately 7600 lines of OCaml code. Our implementation provides two frontends, one based on Menhir, an LR(1) parser generator for the OCaml programming language and one based on Petr4's parser [Doe+21]. While the first allows the programmer to directly use the syntax described in Section 5.3, the latter allows processing $\mathrm{P}_{4}$ programs leveraging the built-in annotation mechanism. The programmer can annotate $\mathrm{P}_{4}$ programs using @pi4 annotations to add custom type annotations to $\mathrm{P}_{4}$ code blocks that are then checked by $\Pi_{4}$ 's type checker. Under the hood, our implementation uses an encoding of $\Pi_{4}$ 's types into a decidable theory of first-order logic, facilitating an SMT solver to automatically discharge the various side conditions that arise during type checking.

We start this chapter by presenting the algorithmic type system of $\Pi_{4}$. Next, we show how to encode $\Pi_{4}$ 's types into a decidable theory of first-order logic, exploiting the fact that in practice, the size of network packets are bound by the Maximum Transmission Unit (MTU). With the algorithmic typing rules in place, we then formally prove that given an input type that respects the bounds imposed by the MTU, we can also prove a bound on the computed output types, which ultimately allows us to prove type checking to be decidable. Finally, we discuss how the algorithmic typing rules can be optimized to obtain a more efficient encoding and provide a brief overview of our $\mathrm{P}_{4}$ frontend.

### 6.1 Algorithmic Typing Rules

For our implementation, we define an algorithmic version of our type system (cf. Figures 6.1 and 6.2 ) whose rules are mostly identical to the rules from our declarative type system. The first difference to the declarative type system is that the semantic conditions that must be checked during type checking are encoded as subtype constraints. For example, when we type-check the command $\operatorname{add}(\iota)$, we must check that the newly added instance is not already valid in the input type $\tau_{1}$, i.e., Excludes $\Gamma \tau_{1} \iota$. As shown by rule T-Add-Algo in Figure 6.1, Excludes $\Gamma \tau_{1} \iota$ becomes the subtype check $\Gamma \vdash \tau_{1}<:\{x: \top \mid \neg x$.ı.valid $\}$. Similarly, rule T-Mod and T-Remit require Includes $\Gamma \tau_{1} \iota$, which becomes $\Gamma \vdash \tau_{1}<:\{x: \top \mid x . \iota$. valid $\}$ in T-Mod-Algo and T-

$$
\begin{aligned}
& \text { T-Extract-ALGO } \\
& \begin{array}{r}
\Gamma \vdash \tau_{1}<:\left\{x: \top| | x . p k t_{\text {in }} \mid \geq \operatorname{sizeof}(\iota)\right\} \quad \varphi_{1} \triangleq z \cdot p k t_{\text {in }}=z \cdot p k t_{\text {out }}=\langle \rangle \\
\quad \varphi_{2} \triangleq y . \iota @ z \cdot p k t_{\text {in }}=x \cdot p k t_{\text {in }} \wedge z \cdot p k t_{\text {out }}=x \cdot p k t_{\text {out }} \wedge z \equiv, x \\
\Gamma \vdash \operatorname{extract}(\iota):\left(x: \tau_{1}\right) \leadsto \Sigma y:\left\{z: \iota \mid \varphi_{1}\right\} \cdot\left\{z: \operatorname{chomp}\left(\tau_{1}, \iota, y\right) \mid \varphi_{2}\right\}
\end{array}
\end{aligned}
$$

T-Add-Algo

$$
\frac{\Gamma \vdash \tau_{1}<:\{x: \top \mid \neg x . \iota . v a l i d\} \quad \text { init }_{\mathcal{H} \mathcal{T}(\iota)}=v}{\Gamma \vdash \operatorname{add}(\iota):\left(x: \tau_{1}\right) \leadsto \Sigma y:\left\{z: \tau_{1} \mid z \equiv x\right\} .\left\{z: \iota \mid z . p k t_{\text {in }}=z . p k t_{\text {out }}=\langle \rangle \wedge z . \iota=v\right\}}
$$

T-Remove-Algo

$$
\begin{gathered}
\Gamma \vdash \tau_{1}<: \iota_{\approx} \\
\varphi_{\iota} \triangleq \forall \kappa \in \operatorname{dom}(\mathcal{H T}) \cdot \kappa \neq \iota \Longrightarrow y \cdot \kappa=x \cdot \kappa \\
\varphi_{p k t} \triangleq y \cdot p k t_{\text {in }}=x \cdot p k t_{\text {in }} \wedge y \cdot p k t_{\text {out }}=x \cdot p k t_{\text {out }} \\
\Gamma \vdash \operatorname{remove}(\iota):\left(x: \tau_{1}\right) \leadsto\left\{y: \top \mid \varphi_{\iota} \wedge \varphi_{p k t} \wedge \neg y . . \iota . v a l i d\right\}
\end{gathered}
$$

T-Mod-Algo

$$
\begin{gathered}
\Gamma \vdash \tau<: \iota_{\approx} \quad \mathcal{F}(\iota, f)=\mathrm{BV} \quad \Gamma ; \tau_{1} \vdash e: \mathrm{BV} \\
\varphi_{p k t} \triangleq y \cdot p k t_{\text {in }}=x \cdot p k t_{\text {in }} \wedge y \cdot p k t_{\text {out }}=x \cdot p k t_{\text {out }} \\
\varphi_{\iota} \triangleq \forall \kappa \in \operatorname{dom}(\mathcal{H} \mathcal{T}) . \iota \neq \kappa \Rightarrow y . \kappa=x . \kappa \\
\varphi_{f} \triangleq \forall g \in \operatorname{dom}(\mathcal{H T}(\iota)) . f \neq g \Rightarrow y . \iota . g=x . \iota . g \\
\Gamma \vdash \iota . f:=e:\left(x: \tau_{1}\right) \leadsto\left\{y: T \mid \varphi_{p k t} \wedge \varphi_{\iota} \wedge \varphi_{f} \wedge y . \iota . f=e[x / \text { heap }]\right\}
\end{gathered}
$$

T-Remit-Algo

$$
\begin{gathered}
\frac{\Gamma \vdash \tau_{1}<: \iota_{\approx} \quad \varphi \triangleq z \cdot p k t_{\text {in }}=\langle \rangle \wedge z \cdot p k t_{\text {out }}=x \cdot l}{\Gamma \vdash \operatorname{remit}(\iota):\left(x: \tau_{1}\right) \leadsto \Sigma y:\left\{z: \tau_{1} \mid z \equiv x\right\} \cdot\{z: \epsilon \mid \varphi\}} \\
\text { T-Reset-ALGO } \\
\varphi_{1} \triangleq z \cdot p k t_{\text {out }}=\langle \rangle \wedge z \cdot p k t_{\text {in }}=x \cdot p k t_{\text {out }} \\
\frac{\varphi_{2} \triangleq z \cdot p k t_{\text {out }}=\langle \rangle \wedge z \cdot p k t_{\text {in }}=x \cdot p k t_{\text {in }}}{\Gamma \vdash \operatorname{reset}:\left(x: \tau_{1}\right) \leadsto \Sigma y:\left\{z: \epsilon \mid \varphi_{1}\right\} \cdot\left\{z: \epsilon \mid \varphi_{2}\right\}}
\end{gathered}
$$

Figure 6.1: Algorithmic typing rules for domain-specific commands.

Remit-Algo respectively. The check sizeof ${ }_{p k t_{i n}}(\tau) \geq \operatorname{sizeof}(\iota)$ required by T-Extract becomes $\Gamma \vdash \tau_{1}<:\left\{x: T| | x . p k t_{i n} \mid \geq \operatorname{sizeof}(\iota)\right\}$ in rule T-Extract-Algo.

The second major difference is the rule for type ascription T-Ascribe-Algo. In our implementation we check if the input type $\tau_{1}$ is a subtype of the ascribed input type $\hat{\tau}_{1}$. We then use the ascribed input type to compute an output type $\tau_{c}$. Finally, we check if the computed output type $\tau_{c}$ is a subtype of the ascribed output type $\hat{\tau}_{2}$. Note, that our type checking algorithm can be used to obtain a weak form of type inference. Given an input type that describes the state before the execution, our algorithm computes an output type, which describes the state after the execution of the program.

### 6.2 Decidability

An essential prerequisite to prove decidability of ח4's type checking is that subtyping checks are decidable. Since we encode types into bit vectors, this follows by finite enumeration, if we can show that the bit vectors are finite. Unfortunately, the $p k t_{i n}$ and

T-Seq-Algo

$$
\begin{aligned}
& \Gamma \vdash c_{1}:\left(x: \tau_{1}\right) \leadsto \tau_{12} \quad \text { T-Skip-Algo } \\
& \frac{\Gamma,\left(x: \tau_{1}\right) \vdash c_{2}:\left(y: \tau_{12}\right) \leadsto \tau_{22}}{\Gamma \vdash c_{1} ; c_{2}:\left(x: \tau_{1}\right) \leadsto \tau_{22}\left[y \mapsto \tau_{12}\right]} \quad \frac{\tau_{2} \triangleq\left\{y: \tau_{1} \mid y \equiv x\right\}}{\Gamma \vdash \operatorname{skip}:\left(x: \tau_{1}\right) \leadsto \tau_{2}} \\
& \text { T-Ascribe-Algo } \\
& \Gamma \vdash c:\left(x: \hat{\tau}_{1}\right) \leadsto \tau_{c} \quad \Gamma \vdash \tau_{1}<: \hat{\tau}_{1} \\
& \frac{\Gamma, x: \hat{\tau}_{1} \vdash \tau_{c}<: \hat{\tau}_{2}}{\Gamma \vdash c \operatorname{as}\left(x: \hat{\tau}_{1}\right) \rightarrow \hat{\tau}_{2}:\left(x: \tau_{1}\right) \leadsto \hat{\tau}_{2}}
\end{aligned}
$$

T-If-Algo

$$
\begin{gathered}
\Gamma ; \tau_{1} \vdash \varphi: \mathbb{B} \quad \Gamma \vdash c_{1}:\left(x:\left\{y: \tau_{1} \mid \varphi[y / \text { heap }]\right\}\right) \leadsto \tau_{12} \\
\Gamma \vdash c_{2}:\left(x:\left\{y: \tau_{1} \mid \neg \varphi[y / \text { heap }]\right\}\right) \leadsto \tau_{22} \\
\Gamma \vdash \text { if }(\varphi) c_{1} \text { else } c_{2}:\left(x: \tau_{1}\right) \leadsto\left\{y: \tau_{12} \mid \varphi[x / \text { heap }]\right\}+\left\{y: \tau_{22} \mid \neg \varphi[x / \text { heap }]\right\}
\end{gathered}
$$

Figure 6.2: Additional algorithmic typing rules.
$p k t_{\text {out }}$ entries describe infinite bit vectors. However, in practice, packets are bounded by the Maximum Transmission Unit (MTU) and therefore, we can bound the number of bits required to represent heaps described by a certain heap type. More formally, we say that a heap type $\tau$ is bounded by $N$ in a context $\Gamma$, written $\Gamma \vdash \tau \leq N$, iff for every heap described by type $\tau$, the length of the input packet and output packet together is less or equal to $N$, which is captured by the following definition:

$$
\Gamma \vdash \tau \leq N \triangleq \forall \mathcal{E} \vDash \Gamma, \forall h \in \llbracket \tau \rrbracket \mathcal{E},\left|h\left(p k t_{\text {in }}\right)\right|+\left|h\left(p k t_{\text {out }}\right)\right| \leq N
$$

It is not sufficient to just require that $\left|h\left(p k t_{\text {in }}\right)\right| \leq N$ and $\left|h\left(p k t_{\text {out }}\right)\right| \leq N$, because as seen in case of the reset command, the resulting input packet is the concatenation of the input and output packet of the input heap, i.e., $h\left(p k t_{\text {out }}\right) @ h\left(p k t_{i n}\right)$, which might violate the constraint.

As a first step, we want to prove that if a program typechecks with a bound on its input type, then we can compute that maximum number of bits that we need to encode the output type. Ideally, this would be the same bound, however, it is possible for a program to emit more bits from the incoming packet than is allowed by the MTU. So we define a helper function emit $(c) \in \mathbb{N}$ that over-approximates the maximum number of bits that could be emitted along any program path in $c$, which is defined in Figure 6.3. Command remit ( $\iota$ ) emits $n$ bits, where $n$ is the size of the header instance $t$. The number of emitted bits of conditionals is at most the maximum of the bits emitted in the then-branch and the else-branch. Sequences emit the sum of emitted bits from both commands and for ascribed commands we take the number of bits emitted by the command itself. All other commands do not affect the number of emitted bits.

Theorem 6.1 captures the idea that given an algorithmic typing judgement on a program $c$, for which the input type and all ascribed types in $c$ respect the MTU $N$, the output type will require no more than $N+\operatorname{emit}(c)$ bits. Note that even though the input type is constrained by the MTU, intermediate states may require more than just $N$ bits. This theorem shows that $N+\mathrm{emit}(c)$ suffices as the maximum combined width of $p k t_{i n}$ and $p k t_{\text {out }}$.

```
emit(extract(l))}\triangleq
emit (if (b) c. else c}\mp@subsup{c}{2}{})\triangleq\operatorname{max}(\operatorname{emit}(\mp@subsup{c}{1}{}),\operatorname{emit}(\mp@subsup{c}{2}{})
emit}(\mp@subsup{c}{1}{};\mp@subsup{c}{2}{})\quad\triangleq\operatorname{emit}(\mp@subsup{c}{1}{})+\operatorname{emit}(\mp@subsup{c}{2}{}
emit}(l.f:=e)\quad\triangleq 
emit}(\operatorname{remit}(\iota))\quad\triangleq\operatorname{sizeof}(\iota
emit(skip) \triangleq O
emit(reset) \triangleq 0
emit(remove(l))}\triangleq
emit}(\operatorname{add}(\iota))\triangleq
emit (c as \sigma) }\quad\mathrm{ emit (c)
```

Figure 6.3: $\operatorname{emit}(c) \in \mathbb{N}$ computes the maximum number of bits that can be emitted along any path in c .

Theorem 6.1 (MTU-Bound). For every $\Gamma, c, x, \tau_{1}, \tau_{2}$, and $N$, if $\Gamma \vdash \tau_{1} \leq N$ and $\Gamma \vdash c:\left(x: \tau_{1}\right) \leadsto \tau_{2}$ and every ascribed type in $c$ is also bounded by $N$, then $\Gamma,\left(x: \tau_{1}\right) \vdash$ $\tau_{2} \leq N+\operatorname{emit}(c)$.

Proof. By induction on the typing derivation. For details, see Theorem A. 78 in Appendix A.2.3.

Theorem 6.2 establishes the correctness of the algorithmic typing relation. It states that a program $c$ typechecks in the declarative system with type $\left(x: \tau_{1}\right) \rightarrow \tau_{2}$ if and only if it also typechecks in the algorithmic system with type $\left(x: \tau_{1}\right) \leadsto \tau_{2}^{\prime}$ and the output type of the algorithmic system $\tau_{2}^{\prime}$ is a subtype of the output type $\tau_{2}$ of the declarative system.

Theorem 6.2 (Algorithmic Typing Correctness). For all $\Gamma, c, x, \tau_{1}$, and $\tau_{2}$, where $x$ is not free in $\tau_{1}, \Gamma \vdash c:\left(x: \tau_{1}\right) \rightarrow \tau_{2}$ if and only if there is some $\tau_{2}^{\prime}$ such that $\Gamma \vdash c:\left(x: \tau_{1}\right) \leadsto \tau_{2}^{\prime}$, and $\Gamma,\left(x: \tau_{1}\right) \vdash \tau_{2}^{\prime} \ll \tau_{2}$.

Proof. By induction on the typing derivation. For details, see Theorem A. 73 in Appendix A.2.3.

With Theorems 6.1 and 6.2 in hand, it is straightforward to show the decidability of the declarative type system, i.e., that $\Gamma \vdash c:\left(x: \tau_{1}\right) \rightarrow \tau_{2}$ is decidable, which is stated by Theorem 6.3. Theorem 6.2 allows us to equivalently show that type checking the command in the algorithmic type system-i.e., $\Gamma \vdash c:\left(x: \tau_{1}\right) \leadsto \tau_{2}^{\prime}-$ and that checking $\Gamma,\left(x: \tau_{1}\right) \vdash \tau_{2}^{\prime}<: \tau_{2}$ are decidable. We can prove the former by induction on the typing derivation, while the latter follows by finite enumeration using the bounds guaranteed by Theorem 6.1.

Theorem 6.3 (Decidability). If $\Gamma, \tau_{1}, \tau_{2}$ and every ascribed type in $c$ are bounded by the MTU $N$, then $\Gamma \vdash c:\left(x: \tau_{1}\right) \rightarrow \tau_{2}$ is decidable.

Proof. Proof by Algorithmic Typing Correctness, MTU-Bound and by induction on the typing derivation. For details, see Theorem A.79.

```
        \(\bigwedge_{\epsilon \in \operatorname{dom}(\Gamma)} \operatorname{smt}(y, \Gamma(y)) \wedge \operatorname{smt}\left(s, \tau_{1}\right) \wedge\)
\(\forall\) consts \(\left(\mathrm{t}, \tau_{2}\right) \cdot \operatorname{smt}\left(t, \tau_{2}\right) \Longrightarrow \neg\) equal \((s, t)\)
```

Figure 6.4: Encoding of subtyping check $\Gamma \vdash \tau_{1}<: \tau_{2}$

### 6.3 SMT Encoding

In our implementation, we convert every subtyping check $\Gamma \vdash \tau_{1}<: \tau_{2}$ into a formula in the theory of uninterpreted functions and (fixed-width) bit vectors (UFBV) [WHM13]. Again, we exploit the fact that the MTU limits the number of bits network switches can receive or transmit. So in our encoding, the MTU determines the size of the bit vectors.

Our encoding is based on the semantic notion of subtyping as defined in Section 5.6. Intuitively, type $\tau_{1}$ is a subtype of $\tau_{2}$, i.e., $\Gamma \vdash \tau_{1}<: \tau_{2}$, if the set of heaps described by $\tau_{1}$ is a subset of the set of heaps described by $\tau_{2}$. Let $S$ be the set of heaps described by $\tau_{1}$ and $T$ the set of heaps described by $\tau_{2}$. Set $S$ is a subset of set $T(S \subseteq T)$ iff for all $s \in S$ there exists a $t \in T$ such that $s=t$-i.e., $\forall s \in S . \exists t \in T . s=t$. We want the formula to hold for all assignments of $s$ and $t$, so we check if the formula is valid by checking that the negation of the formula $\exists s \in S . \forall t \in T . s \neq t$ is not satisfiable (UNSAT).

Accordingly, the subtype check $\Gamma \vdash \tau_{1}<: \tau_{2}$ is encoded into the formula shown in Figure 6.4. Similar to the elements in the set of heaps described by subtype $\tau_{1}$, heap types bound to variables in context $\Gamma$ respectively are also existentially quantified. However, in the encoding, we can eliminate the existential quantifiers by skolemization. We declare a set of global constants for all heap types (function consts $(x, \tau)$ ), which we use to model the heaps described by the types. We declare two constants for every header instance in the header table $\mathcal{H} \mathcal{T}$, a boolean constant capturing the validity of the header instance (e.g. x.ethernet.valid) and a bit vector constant for the instance data (e.g. $x$. ethernet), where the size of the bit vector corresponds to the size of the instance. We declare two bit vector constants each for the input and output packet, one capturing the length of the packet (e.g., x.pkt_in. length) and one capturing the contents of the packet (e.g. x.pkt_in). The length of both bit vectors is determined by the MTU. The length bit vector has as many bits as needed to store the value of the MTU, the length of the packet bit vector is equal to the MTU.

In the encoding, we decompose heap types according to their recursive structure. We reference the respective sub-heaps via the binders introduced by the types, which is the reason why all declared constants are prefixed with a variable. For example, assuming subtype $\tau_{1}$ has the form $\Sigma x:\{y: \top \mid \ldots\} .\{z: \top \mid \ldots\}$, the encoding will declare the following set of constants:

- $\left\{\mathrm{s}, \mathrm{x}_{1}, \mathrm{x}_{\mathrm{r}}, \mathrm{y}, \mathrm{z}\right\} \cdot \mathrm{l}_{0}, \ldots,\left\{\mathrm{~s}, \mathrm{x}_{1}, \mathrm{x}_{\mathrm{r}}, \mathrm{y}, \mathrm{z}\right\} \cdot \mathrm{l}_{\mathrm{n}}$
- $\left\{\mathrm{s}, \mathrm{x}_{1}, \mathrm{x}_{\mathrm{r}}, \mathrm{y}, \mathrm{z}\right\} . \iota_{0} \cdot \mathrm{valid}, \ldots,\left\{\mathrm{s}, \mathrm{x}_{1}, \mathrm{x}_{\mathrm{r}}, \mathrm{y}, \mathrm{z}\right\} . \mathrm{l}_{\mathrm{n}} . \mathrm{valid}$
- $\left\{s, x_{1}, x_{r}, y, z\right\}$.pkt_in, $\left\{s, x_{1}, x_{r}, y, z\right\}$.pkt_in.length
- $\left\{\mathrm{s}, \mathrm{x}_{1}, \mathrm{x}_{\mathrm{r}}, \mathrm{y}, \mathrm{z}\right\}$.pkt_out, $\left\{\mathrm{s}, \mathrm{x}_{1}, \mathrm{x}_{\mathrm{r}}, \mathrm{y}, \mathrm{z}\right\}$.pkt_out.length

Note, we introduce an additional binder $s$ for the subtype and $t$ for the supertype to be able to refer to the top-level heap, which, for example, is necessary if the top-level type

$$
\begin{aligned}
\operatorname{smt}(x, \varnothing) \triangleq & \text { false } \\
\operatorname{smt}(x, \top) \triangleq & \operatorname{true} \\
\operatorname{smt}\left(x, \tau_{1}+\tau_{2}\right) \triangleq & \operatorname{smt}\left(x, \tau_{1}\right) \vee \operatorname{smt}\left(x, \tau_{2}\right) \\
\operatorname{smt}\left(x, \Sigma y: \tau_{1} \cdot \tau_{2}\right) \triangleq & \operatorname{smt}\left(y_{l}, \tau_{1}\right) \wedge \operatorname{smt}\left(y_{r}, \tau_{2}\right) \wedge \\
& \operatorname{pktbounds}(x) \wedge \\
& \operatorname{pktbounds}\left(y_{l}\right) \wedge \\
& \operatorname{pktbounds}\left(y_{r}\right) \wedge \\
& \operatorname{append}\left(x, y_{l}, y_{r}\right) \\
\operatorname{smt}\left(x_{0}, \tau_{\circ}\left[x_{1} \mapsto \tau_{1}\right]\right) \triangleq & \operatorname{smt}\left(x_{0}, \tau_{0}\right) \wedge \operatorname{smt}\left(x_{1}, \tau_{1}\right) \wedge \\
& \operatorname{pktbounds}\left(x_{0}\right) \wedge \\
& \operatorname{pktbounds}\left(x_{1}\right) \\
\operatorname{smt}(x,\{y: \tau \mid e\}) \triangleq & \operatorname{smt}(y, \tau) \wedge \operatorname{smt} \text { form }(e) \wedge \\
& \operatorname{pktbounds}(x) \wedge \\
& \operatorname{pktbounds}(y) \wedge \\
& \text { equal }(x, y)
\end{aligned}
$$

Figure 6.5: SMT encoding of heap types.
is a choice type. For $\Sigma$-types we introduce binders for both projections, in the example $x_{l}$ and $x_{r}{ }^{1}$.

Encoding of heap types As defined in Figure 6.5, the function $\operatorname{smt}(x, \tau)$ encodes heap type $\tau$ into a first-order logic formula that describes the heap identified by binder $x$, i.e., the formula adds assertions for constants prefixed with variable $x$. The types $\varnothing$ and T are encoded into boolean literals false and true respectively. A choice type $\tau_{1}+\tau_{2}$ is encoded into the disjunction between formulas for $\tau_{1}$ and $\tau_{2}$. The encoding of a $\Sigma$-type $\Sigma y: \tau_{1} \cdot \tau_{2}$ returns a conjunction between the formulas for $\tau_{1}$ and $\tau_{2}$, constraints on the input and output packet (function pktbounds) as well as the encoding of the concatenation of heaps described by $\tau_{1}$ and $\tau_{2}$ (function append). The auxiliary functions pktbounds and append are defined in Figure 6.6. Function pktbounds $(x)$ asserts for the input and output packet of heap $x$ that the length is less than the MTU $\mathcal{M}$ and that the value of the input and output packet is constrained by the packet length. Function append $\left(x_{0}, x_{1}, x_{2}\right)$ returns a formula that describes the concatenation of heaps $x_{1}$ and $x_{2}$, where heap $x_{0}$ is the resulting heap. An instance $t$ is valid in $x_{0}$, if it is either valid in $x_{1}$ or $x_{2}$. Depending on the validity of instances, function hdreq asserts that instances in $x_{0}$ are equal to instances in $x_{1}$ and $x_{2}$ respectively. Function append_pkt $\left(x_{0}, x_{1}, x_{2}, p\right)$ asserts that the packet $p$ in $x_{0}$ is the concatenation of packets $p$ in $x_{1}$ and $x_{2}$. The packet length is the sum of the packet lengths in $x_{1}$ and $x_{2}$, but at most the MTU $\mathcal{M}$. If the total length is greater than zero but the packet in $x_{1}$ is empty, the resulting packet is determined by $x_{2}$. If the packet in $x_{1}$ is not empty, the result of the concatenation is obtained by computing the bitwise OR between the packet in $x_{1}$ and the packet

[^3]\[

$$
\begin{aligned}
& \operatorname{hdreq}(x, y, \iota) \triangleq y . \iota . v a l i d \Longrightarrow x . \iota=y . \iota \\
& \text { equal }(x, y) \triangleq \bigwedge_{t \in \operatorname{dom}(\mathcal{H T})} x . l=y . \iota \wedge x . \iota . v a l i d=y . \iota . v a l i d \wedge \\
& \bigwedge_{p \in\left\{p k t_{t_{\text {in }}}, p k t_{\text {out }}\right\}} x \cdot p=y \cdot p \wedge x . p . \text { length }=y . p . \text { length } \\
& \operatorname{pktbounds}(x) \triangleq \bigwedge_{p \in\left\{p k t_{\text {in }}, p k t_{\text {out }}\right\}}(x . p . \text { length } \leq \mathcal{M} \wedge \\
& \left.\left(x . p . \text { length }=0 \vee x . p<2^{x . p . l e n g t h}\right)\right) \\
& \text { append_pkt }\left(x_{0}, x_{1}, x_{2}, p\right) \triangleq x_{0} \cdot p . \text { length }=\min \left(x_{1} \cdot p . \text { length }+x_{2} \cdot p . \text { length, } \mathcal{M}\right) \wedge \\
& \left(x_{0} \cdot p . \text { length }=0 \vee\left(x_{1} \cdot p . \text { length }=0 \wedge x_{0} \cdot p=x_{2} \cdot p\right) \vee\right. \\
& \left.\left(x_{1} \cdot p . \text { length }>0 \wedge x_{0} \cdot p=x_{1} \cdot p \mid\left(x_{2} \cdot p \ll x_{1} \cdot p . \text { length }\right)\right)\right) \\
& \operatorname{append}\left(x_{0}, x_{1}, x_{2}\right) \triangleq \bigwedge_{p \in\left\{p k t_{\text {in }}, p k t_{\text {out }}\right\}} \text { append_pkt }\left(x_{0}, x_{1}, x_{2}, p\right) \wedge \\
& \wedge\left(x_{0} . . . \text { valid }=\left(x_{1} . . . \text { valid } \oplus x_{2} . t . \text {.valid }\right) \wedge\right. \\
& \iota \in \operatorname{dom}(\mathcal{H} \mathcal{T}) \\
& \left.\operatorname{hdreq}\left(x_{0}, x_{1}, l\right) \wedge \operatorname{hdreq}\left(x_{0}, x_{2}, l\right)\right)
\end{aligned}
$$
\]

Figure 6.6: Helper functions for SMT encoding.
in $x_{2}$ shifted by the length of the packet in $x_{1}$. For example, if $x_{1} \cdot p k t_{i n}$.length $=3$ and $x_{1} \cdot p k t_{i n}=000101$ and $x_{2} \cdot p k t_{i n}=000110, x_{0} \cdot p k t_{i n}=000101 \mid(000110 \ll 3)=$ ooo101|110000 $=110101$. Note, in the formalization of $\Pi_{4}$, the least significant bit of a bit vector is the left-most bit, while in $\mathrm{Z}_{3}$, the right-most bit is the least significant one-i.e, in the formalization the concatenation of bit vectors $x_{1} \cdot p k t_{i n}=101$ and $x_{2} \cdot p k t_{i n}=011$ results in $x_{1} . p k t_{i n} @ x_{2} . p k t_{i n}=101011$.

The encoding of a substitution type $\tau_{0}\left[x_{1} \mapsto \tau_{1}\right]$ returns a conjunction of the formulas for $\tau_{0}$ and $\tau_{1}$ as well as packet constraints on $x_{0}$ and $x_{1}$ The encoding of a refinement type $\{y: \tau \mid \varphi\}$ returns the conjunction between the formula for heap type $\tau$ and the formula for $\Pi_{4}$ formula $\varphi$, as well as constraints on the input and output packet and the equality between heaps $x$ and $y$ (function equal $(x, y)$ ) with regard to instances and input and output packets.

Encoding of formulae The encoding of formulae is defined in Figure 6.7. Boolean literals, conjunctions and negations are encoded into their respective counterparts. The encoding of instance validity checks simply returns the respective constant. In the encoding of expression equality and expression comparison, we distinguish between arithmetic expressions and bit vector expressions. The comparison of arithmetic expressions is defined in Figure 6.8. When comparing two arithmetic expressions $a_{1}$ and $a_{2}$, we first compute the maximum value (function max_value) these expressions can have. For numerals, this corresponds to their value, the packet length can have at most the value of the MTU $\mathcal{M}$, and the maximum value of an addition is the sum of both maximum values. Next, we compute the minimum number $n$ of bits required to represent the maximum value. Both terms are encoded into a fixed-width bit vector of size $n$ (function arith). We write $[m]_{n}$ to denote a bit vector of size $n$ with value $m$.

```
            smt_form(true) \triangleqtrue
            smt_form(false) \triangleq false
            smt_form}(\mp@subsup{\varphi}{1}{}\wedge\mp@subsup{\varphi}{2}{})\triangleq\operatorname{smt}\mathrm{ _form ( }\mp@subsup{\varphi}{1}{})\wedge\operatorname{smt_form}(\mp@subsup{\varphi}{2}{}
            smt_form}(\neg\varphi)\triangleq\negsmt_form(\varphi
                smt_form(x.l.valid ) \triangleqx.l.valid
            smt_form}(\mp@subsup{a}{1}{}=\mp@subsup{a}{2}{})\triangleq\textrm{cmp_arith}(\mp@subsup{a}{1}{},\mp@subsup{a}{2}{},=
            smt_form( }\mp@subsup{a}{1}{}>\mp@subsup{a}{2}{})\triangleq\textrm{cmp_arith}(\mp@subsup{a}{1}{},\mp@subsup{a}{2}{},>
                smt_form}(b\mp@subsup{v}{1}{}=b\mp@subsup{v}{2}{})\triangleqcmp_bv(b\mp@subsup{v}{1}{},b\mp@subsup{v}{2}{},=
```



Figure 6.7: SMT encoding of expressions.

$$
\begin{aligned}
& \operatorname{arith}(m, n) \triangleq[m]_{n} \\
& \operatorname{arith}(|x . p|, n) \triangleq\left[x . p . \text { length }_{n}\right. \\
& \operatorname{arith}\left(a_{1}+a_{2}, n\right) \triangleq \operatorname{arith}\left(a_{1}, n\right)+\operatorname{arith}\left(a_{2}, n\right) \\
& \text { max_value }(m) \triangleq m \\
& \text { max_value }(|x . p|) \triangleq \mathcal{M} \\
& \max \_ \text {value }(n+m) \triangleq \max \_ \text {value }(n)+\max \_ \text {value }(m) \\
& \text { cmp_arith }\left(a_{1}, a_{2}, o p\right) \triangleq \text { let } m_{1}=\text { max_value }\left(a_{1}\right) \text { in } \\
& \text { let } m_{2}=\text { max_value }\left(a_{2}\right) \text { in } \\
& \text { let len }=\min \_ \text {bit_width }\left(\max \left(m_{1}, m_{2}\right)\right) \text { in } \\
& \operatorname{arith}\left(a_{1}, \text { len }\right) \text { op } \operatorname{arith}\left(a_{2}, l e n\right)
\end{aligned}
$$

Figure 6.8: SMT encoding of arithmetic terms and comparison of arithmetic terms.

Since packet length terms are encoded into the respective constant, it might happen that the target size is larger than the pre-declared size of the constant's bit vector type. In this case, we zero-pad the declared constant to produce a bit vector of the desired size. The comparison itself is straightforwardly encoded into a comparison of fixed-width bit vectors

Similarly, we define the comparison of bit vectors in Figure 6.9. The main challenge in the encoding of bit vector expressions is the concatenation of two bit vectors, especially, if a reference to either the input packet or the output packet is involved, e.g., $x . p k t_{i n} @ 1010$. While the input packet (just as the output packet) is represented by a fixed-width bit vector in the encoding, the bit vector representing the packet length actually determines how many bits of the first actually contain meaningful data. For example, let us assume that the input packet is represented by an 8-bit bit vector, but only three bits are used, i.e., x.pkt_in $=00000111$ and x.pkt_in.length $=3$. We call the first the static size

```
        static_size \((\rangle) \triangleq 0\)
    \(\operatorname{static} \_\)size \((b:: b v) \triangleq 1+\operatorname{static} \_\operatorname{size}(b v)\)
static_size \(\left(b v_{1} @ b v_{2}\right) \triangleq \min \left(\right.\) static_size \(\left.\left(b v_{1}\right)+\operatorname{static} \_\operatorname{size}\left(b v_{2}\right), \mathcal{M}\right)\)
static_size \((x . p[l: r]) \triangleq r-l\)
\(\operatorname{static} \_\operatorname{size}(x . l[l: r]) \triangleq r-l\)
    static_size \((x . p) \triangleq \mathcal{M}\)
        \(\operatorname{bv} 2 \operatorname{smt}(b:: b v) \triangleq \operatorname{concat} \operatorname{bv2smt}(b v)[b]_{1}\)
            \(\operatorname{bits}(b v, n) \triangleq[\operatorname{bv} 2 \operatorname{smt} b v]_{n}, n\)
            \(\operatorname{bits}(x . p, n) \triangleq[x . p]_{n}, x . p . l e n g t h\)
    \(\operatorname{bits}(x . p[l: r], n) \triangleq[(\operatorname{extract}(r-1) l) x . p]_{n}, r-l\)
        \(\operatorname{bits}(x . l[l: r], n) \triangleq[(\operatorname{extract}(r-1) l) x . l]_{n}, r-l\)
        \(\operatorname{bits}\left(b v_{1} @ b v_{2}, n\right) \triangleq\) let \(v_{1}, n_{1}=\operatorname{bits}\left(b v_{1}, n\right)\) in
                let \(v_{2}, n_{2}=\operatorname{bits}\left(b v_{2}, n\right)\) in
                \(\begin{cases}v_{2} & \text { if } n_{1}=\mathrm{o} \\ {\left[v_{1} \mid\left(v_{2} \ll n_{1}\right)\right]_{n}, n_{1}+n_{2}} & \text { otherwise }\end{cases}\)
\(\mathrm{cmp} \_\mathrm{bv}\left(b v_{1}, b v_{2}, o p\right) \triangleq\) let len \(=\max \left(\operatorname{static\_ size}\left(b v_{1}\right)\right.\), static_size \(\left.\left(b v_{2}\right)\right)\) in
let \(v_{1}, n_{1}=\operatorname{bits}\left(b v_{1}\right.\), len \()\) in
let \(v_{2}, n_{2}=\operatorname{bits}\left(b v_{2}\right.\),len \()\) in
\(n_{1}=n_{2} \wedge\left(n_{1}=o \vee v_{1} o p v_{2}\right)\)
```

Figure 6.9: SMT encoding of bit vector terms
and the latter the dynamic size of $\mathrm{x} . \mathrm{pkt}$ _in. Inconsiderately concatenating both bit vectors results in 101000000111, but instead we want the bits 1010 to follow directly after the three bits embodying the contents of x.pkt_in, i.e., 01010111.

In the encoding of a comparison of two bit vectors $b v_{1}$ and $b v_{2}$ (function cmp_bv), we therefore first compute the maximum static size of both bit vectors len. We then encode both $b v_{1}$ and $b v_{2}$ into an SMT bit vector expression of size len and finally the formula for bit vector comparison asserts that the dynamic sizes of both bit vectors must be the same and the size is either zero or the comparison relation holds. As defined in the top of Figure 6.9, the static size of a bit vector value corresponds to the number of bits. The static size of a concatenation of two bit vectors is the sum of the static sizes of both bit vectors. The static size of a slice is determined by the number of bits the slice comprises and the static size of a packet bit vector is the MTU $\mathcal{M}$.

The encoding of bit vector terms is defined by function bits $(e, n)$ that encodes a bit vector expression $e$ into a bit vector of size $n$ and additionally returns the dynamic size of the respective expression. Bit vector constants are encoded into bit vectors of size $n$. The only peculiarity is that we have to reverse the order of bits during the encoding because, as stated above, $\Pi_{4}$ 's formalization uses a different bit order than $Z_{3}$. Note, because $Z_{3}$ does not support empty bit vectors, we cannot straightforwardly encode
the empty bit vector $\rangle$. In the encoding we therefore handle cases dealing with empty bit vectors using assertions on the packet length. The encoding of packet references simply returns the respective constant, possibly zero-extended to match size $n$. The dynamic size is determined by the respective packet length constant. Slices on packets and instances are encoded into a bit vector extraction. The dynamic size in these cases is statically known and results from the size of the interval. For the concatenation of two bit vectors $b v_{1}$ and $b v_{2}$ we first encode both bit vectors into an SMT expression. If the dynamic size of $b v_{1}$ is zero, the result is the result of encoding $b v_{2}$. Otherwise, the concatenated bit vector is obtained by shifting the result of encoding $b v_{2}$ by the dynamic size of the first bit vector to the left and applying the bitwise inclusive OR operator to the resulting bit vector and the result of encoding $b v_{1}$.

### 6.4 Optimizations

The performance of the type checker depends mainly on how fast the SMT solver is able to perform subtyping checks. This results in two levers that can be used to optimize the performance: (1) the encoding itself and (2) the number of invocations of the SMT solver during the type checking process. In the following, we will look at a total of three such optimizations.

### 6.4.1 Optimizing the SMT Encoding

First, we consider two optimizations to improve the SMT encoding. We discuss how we can treat $\Sigma$-types as syntactic sugar to obtain a more efficient encoding, and also how we can inline substitution types to reduce the overall complexity of generated SMT queries.

## Type Equivalences

The typing rules for parsing, adding and deparsing header instances as well as the rules for commands reset and skip have one significant drawback both in the declarative and in the algorithmic type system. The output type is obtained by extending the input type, mostly by means of a $\Sigma$-type with new header instances or bits in the output packet, which are added to the heap by the execution of the command. By carrying over the input type, the computed type grows larger and larger the longer the program is that is checked, and thus, subtyping checks become more expensive in terms of the time $Z_{3}$ needs to solve formulae. Furthermore, the encoding of $\Sigma$-types presented above has been shown to incur significant overhead in practice, i.e., subtyping checks that include $\Sigma$-types take significantly longer.

Fortunately, we can assign equivalent types, which allows us to eliminate $\Sigma$-types and even the chomp-operator. To eliminate $\Sigma$-types, we exploit the fact that $\Sigma$-types can be expressed using refinement types and substitution types, which is captured by Lemma 6.4. We write $\Gamma \vdash \tau_{1} \doteq \tau_{2}$ to denote the equivalence of types $\tau_{1}$ and $\tau_{2}$ in context $\Gamma$, more formally:

$$
\Gamma \vdash \tau_{1} \doteq \tau_{2} \triangleq \forall \mathcal{E} \vDash \Gamma \cdot \llbracket \tau_{1} \rrbracket_{\mathcal{E}}=\llbracket \tau_{2} \rrbracket_{\mathcal{E}}
$$

Lemma 6.4 (Rewriting Sigma Types). In any context $\Gamma$,

Proof. Proof each direction separately.
$(\Rightarrow)$ Let $\mathcal{E} \vDash \Gamma$ and let $h \in \llbracket \Sigma x: \tau_{1} \cdot \tau_{2} \rrbracket_{\mathcal{E}}$. By the semantics of heap types, we know there exists $h_{1}$ and $h_{2}$ such that $h=h_{1}++h_{2}$ and $h_{1} \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}}$ and $h_{2} \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}\left[x \mapsto h_{1}\right]}$. By definition of heap concatenation, $h\left(p k t_{i n}\right)=h_{1}\left(p k t_{i n}\right) @ h_{2}\left(p k t_{i n}\right)$ and also $h\left(p k t_{\text {out }}\right)=h_{1}\left(p k t_{\text {out }}\right) @ h_{2}\left(p k t_{\text {out }}\right)$. Further, $\operatorname{dom}(h)$ is the disjoint union of $\operatorname{dom}\left(h_{1}\right)$ and $\operatorname{dom}\left(h_{2}\right)$ such that if $\iota \in \operatorname{dom}\left(h_{i}\right), h(\iota)=h_{i}(\iota)$ for each $i=1,2$ and each $\iota \in \operatorname{dom}(\mathcal{H} \mathcal{T})$. The result follows by definition of the semantics.
$(\Leftarrow)$ Let $\mathcal{E} \vDash \Gamma$. By the definition of the semantics, it suffices to show, for $h_{1} \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}}$, and $h_{2} \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}\left[x \mapsto h_{1}\right]}$, and $h \in \llbracket \top \rrbracket_{\mathcal{E}\left[x \mapsto h_{1}, r \mapsto h_{2}\right]}$ such that the above refinement holds for $h$, that $h \in \llbracket \Sigma x: \tau_{1}, \tau_{2} \rrbracket \mathcal{E}$. By the semantics of heap types, it suffices to show that $h=h_{1}++h_{2}$. The refinement tells us that

$$
\text { - } h\left(p k t_{i n}\right)=h_{1}\left(p k t_{i n}\right) @ h_{2}\left(p k t_{i n}\right) \text { and }
$$

- $h\left(p k t_{\text {out }}\right)=h_{1}\left(p k t_{\text {out }}\right) @ h_{2}\left(p k t_{\text {out }}\right)$.

Further, $\operatorname{dom}(h)$ is the disjoint union of $\operatorname{dom}\left(h_{1}\right)$ and $\operatorname{dom}\left(h_{2}\right)$ such that if $\iota \in$ $\operatorname{dom}\left(h_{i}\right), h(\iota)=h_{i}(\iota)$ for each $i=1,2$ and each $\iota \in \operatorname{dom}(\mathcal{H} \mathcal{T})$.

Since the $\Sigma$-types used in the typing rules are even more specific due to various refinements, we are able to specify equivalent types that are even further simplified and do not need to introduce additional substitution types. Lemma 6.5 shows exemplarily for rule T-Extract, how we can construct a refinement type that is equivalent to the $\Sigma$-type used for the output type. Similarly, equivalent types can be specified for the add, remit, and reset commands.

Lemma 6.5 (Rewrite Sigma Extract). For all $\Gamma, x, \tau$ and $\iota$, if $\Gamma \vdash \operatorname{sizeof}_{p k t_{i n}}(\tau) \geq$ sizeof $(\iota)$ and $x$ does not occur free in $\tau$, then

$$
\begin{aligned}
& \Sigma y:\left\{\begin{array}{l|l}
z: \iota & \begin{array}{l}
z . p k t_{\text {in }}=\langle \rangle \wedge \\
z \cdot p k t_{\text {out }}=\langle \rangle
\end{array}
\end{array}\right\} \cdot\left\{\begin{array}{ll}
\Gamma, x: \tau \vdash \\
z \operatorname{chomp}(\tau, \iota, y) & \begin{array}{l}
y . \iota @ z \cdot p k t_{\text {in }}=x \cdot p k t_{\text {in }} \wedge \\
z \cdot p k t_{\text {out }}=x \cdot p k t_{\text {out }} \wedge \\
z \equiv \iota x
\end{array}
\end{array}\right\} \\
& \left\{\begin{array}{l|l}
y: \top & \begin{array}{l}
\text { y.ı.valid } \wedge \wedge_{\kappa \in \operatorname{dom}(\mathcal{H} \mathcal{T}) \wedge \kappa \neq \iota} y \cdot \kappa=x . \kappa \wedge \\
\text { y.l@y.pkt } t_{\text {in }}=x . p k t_{\text {in }} \wedge y . p k t_{\text {out }}=x . p k t_{\text {out }}
\end{array}
\end{array}\right\}
\end{aligned}
$$

Proof. Again, we prove both directions separately. The result follows by the semantics of heap types and the relation between the semantics of chomped types and semantic chomp (chomp ${ }^{\Downarrow}$ ). The full proof can be found in Appendix A.2.4 (Lemma A.8o)

Lemma 6.6 shows that an equivalent type also exists for the output type of the rule T-Skip, which does not copy the input type into the output type, but instead uses $T$ as the base type. Figure 6.10 summarizes all optimized algorithmic typing rules.

Lemma 6.6 (Rewriting Refinement Types). For $\Gamma, \tau, t, x, y$, such that $x$ and $y$ do not occur free in $\tau$,

$$
\Gamma,(x: \tau) \vdash\{y: \tau \mid x \equiv y\} \doteq\{y: \top \mid x \equiv y\}
$$

Proof. Prove each direction separately.
$(\Rightarrow)$ Let $\mathcal{E} \vDash \Gamma,(x: \tau)$. We know $\mathcal{E}=\mathcal{E}^{\prime}\left[x \mapsto h_{1}\right]$ such that $h_{1} \in \llbracket \tau \rrbracket_{\mathcal{E}^{\prime}}$. Let $h_{2} \in \llbracket\{y: \tau \mid x \equiv y\} \rrbracket_{\mathcal{E}^{\prime}\left[x \mapsto h_{1}\right]}$. Then $h_{2} \in \llbracket \tau \rrbracket_{\mathcal{E}^{\prime}\left[x \mapsto h_{1}\right]}$ and $\llbracket x \equiv y \rrbracket_{\mathcal{E}^{\prime}\left[x \mapsto h_{1}, y \mapsto h_{2}\right]}=$ true. From the latter, we can conclude that $h_{2}=h_{1}$. To show $h_{2} \in \llbracket\{y: T \mid$ $y \equiv x\} \rrbracket_{\mathcal{E}^{\prime}\left[x \mapsto h_{1}\right]}$, we have to show that $h_{2} \in \llbracket \top \rrbracket_{\mathcal{E}^{\prime}\left[x \mapsto h_{1}\right]}$, which is immediate, and $\llbracket x \equiv y \rrbracket_{\mathcal{E}^{\prime}\left[x \mapsto h_{1}, y \mapsto h_{2}\right]}=$ true, which immediately follows by the fact that $h_{2}=h_{1}$.
$(\Leftarrow)$ Let $\mathcal{E} \vDash \Gamma,(x: \tau)$. We know $\mathcal{E}=\mathcal{E}^{\prime}\left[x \mapsto h_{1}\right]$ such that $h_{1} \in \llbracket \tau \rrbracket_{\mathcal{E}^{\prime}}$. Let $h_{2} \in$ $\llbracket\{y: \top \mid x \equiv y\} \rrbracket_{\mathcal{E}^{\prime}\left[x \mapsto h_{1}\right]}$. Then $h_{2} \in \llbracket T \rrbracket_{\mathcal{E}^{\prime}\left[x \leftrightarrow h_{1}\right]}$, and $\llbracket x \equiv y \rrbracket_{\mathcal{E}\left[x \mapsto h_{1}, y \mapsto h_{2}\right]}=$ true. Observe that $h_{1}=h_{2}$. To show that $h_{2} \in \llbracket\{y: \tau \mid y \equiv x\} \rrbracket_{\mathcal{E}^{\prime}\left[x \mapsto h_{1}\right]}$, we must show that $h_{2} \in \llbracket \tau \rrbracket_{\mathcal{E}^{\prime}\left[x \mapsto h_{1}\right]}$ and $\llbracket x \equiv y \rrbracket_{\mathcal{E}\left[x \mapsto h_{1}, y \mapsto h_{2}\right]}=$ true. The first follows by assumption that $h_{1} \in \llbracket \tau \rrbracket \mathcal{E}^{\prime}$ and the fact that $h_{2}=h_{1}$. The second immediately follows from $h_{2}=h_{1}$.

## Substitution Inlining

Another possibility for optimization arises with respect to the typing rule for sequencing T-SEQ(-Algo). As a reminder, since the output type $\tau_{22}$ of command $c_{2}$ might depend on the output type $\tau_{12}$ of $c_{1}$, we need to memorize $\tau_{12}$ for the final output type. In the formalization we achieve this in terms of a substitution type, which is elegant for the formalization, but in practice again leads to the complexity of the computed types growing with the length of the program. This in turn leads to an increase in the complexity of the SMT queries, since as shown at the beginning of Section 6.3, additional constants are generated for each binder, for which the SMT solver must solve additional constraints.

As it turns out, in most cases it is possible to eliminate the explicit substitution while obtaining an equivalent heap type. For example, let us consider the simple $\Pi_{4}$ program shown in the following code listing and let us assume this program typechecks with some type $\left(x: \tau_{\text {in }}\right) \rightarrow \tau_{\text {out }}$.

```
1 extract(ether);
2 skip
```

Let us further assume that the header table only contains header instance ether. Given type $\tau_{i n}$, the type checker computes the following type $\tau_{c}$, where-starting from the input heap bound to variable $x$-the refinement type bound to $z$ describes the intermediate heap obtained after executing command extract(ether).

T-Extract-Algo-Opt

$$
\Gamma \vdash \tau_{1}<:\left\{x: \top| | x \cdot p k t_{i n} \mid \geq \operatorname{sizeof}(\iota)\right\}
$$

$$
\varphi_{\iota} \triangleq y . t . v a l i d \wedge y . \iota=x . p k t_{i n}[0: \operatorname{sizeof}(\iota)] \wedge \bigwedge_{\kappa \in \operatorname{dom}(\mathcal{H T}) \wedge \kappa \neq \iota} y . \kappa=x . \kappa
$$

$\frac{\varphi_{p k t} \triangleq y . \iota @ y . p k t_{\text {in }}=x . p k t_{\text {in }} \wedge y \cdot p k t_{\text {out }}=x . p k t_{\text {out }}}{\Gamma \vdash \operatorname{extract}(\iota):\left(x: \tau_{1}\right) \leadsto\left\{y: \top \mid \varphi_{\imath} \wedge \varphi_{p k t}\right\}}$

T-Remit-Algo-Opt
$\frac{\Gamma \vdash \tau_{1}<: \iota_{\approx} \varphi_{p k t} \triangleq y \cdot p k t_{\text {in }}=x \cdot p k t_{\text {in }} \wedge y \cdot p k t_{\text {out }}=x . p k t_{\text {out }} @ x . \iota}{\Gamma \vdash \operatorname{remit}(\iota):\left(x: \tau_{1}\right) \leadsto\left\{y: \mathrm{T} \mid y \equiv, ~ x \wedge \varphi_{p k t}\right\}}$

> T-Reset-Algo-Opt

$$
\frac{\varphi_{\text {pkt }} \triangleq y \cdot p k t_{\text {in }}=x . p k t_{\text {out }} @ x . p k t_{\text {in }} \wedge y . p k t_{\text {out }}=\langle \rangle}{\Gamma \vdash \text { reset }:\left(x: \tau_{1}\right) \leadsto\left\{y: \top \mid \bigwedge_{t \in \operatorname{dom}(\mathcal{H} \mathcal{T})} \neg y \cdot \iota . \text { valid } \wedge \varphi_{p k t}\right\}}
$$

T-Add-Algo-Opt

$$
\Gamma \vdash \tau_{1}<:\{x: \top \mid \neg x . \iota . \text { valid }\}
$$

init $_{\mathcal{H} \mathcal{T}(t)}=v \quad \varphi_{t} \triangleq y . \iota . v a l i d \wedge y . l=v \wedge \bigwedge_{\kappa \in \operatorname{dom}(\mathcal{H} \mathcal{T}) \wedge \kappa \neq \iota} y . \kappa=x . \kappa$
$\frac{\varphi_{p k t} \triangleq y \cdot p k t_{\text {in }}=x \cdot p k t_{\text {in }} \wedge y . p k t_{\text {out }}=x . p k t_{\text {out }}}{\Gamma \vdash \operatorname{add}(\iota):\left(x: \tau_{1}\right) \leadsto\left\{y: \mathrm{T} \mid \varphi_{\imath} \wedge \varphi_{p k t}\right\}}$
T-Skip-Algo-Opt

$$
\overline{\Gamma \vdash s k i p:(x: \tau) \leadsto\{y: \top \mid y \equiv x\}}
$$

Figure 6.10: Optimized algorithmic typing rules for $\Pi_{4}$.

$$
\begin{gathered}
\tau_{c}=\left\{y: \top| | y \cdot p k t_{\text {in }}\left|=\left|z \cdot p k t_{\text {in }}\right| \wedge\right.\right. \\
y \cdot p k t_{\text {in }}=z \cdot p k t_{\text {in }} \wedge \\
\left|y \cdot p k t_{\text {out }}\right|=\left|z \cdot p k t_{\text {out }}\right| \wedge \\
y . p k t_{\text {out }}=z . p k t_{\text {out }} \wedge \\
((\neg y . \text { ether.valid } \wedge \neg z . \text { ether.valid }) \vee \\
(y . \text { ether.valid } \wedge z . \text { ether.valid } \wedge y . \text { ether }=z . \text { ether }))\}[z \mapsto \\
\{v: \top \mid v . e t h e r . v a l i d \wedge \\
v . e t h e r @ v . p k t_{\text {in }}=x . p k t_{\text {in }} \wedge \\
\left|v . p k t_{\text {in }}\right|+112=\left|x . p k t_{\text {in }}\right| \wedge \\
\left|v . p k t_{\text {out }}\right|=\left|x . p k t_{\text {out }}\right| \wedge \\
\left.\left.v . p k t_{\text {out }}=x . p k t_{\text {out }}\right\}\right]
\end{gathered}
$$

Suppose we want to convince ourselves what length of the input packet this type allows. First, we see that the length of the input packet is determined by the refinement $\left|y . p k t_{i n}\right|=\left|z . p k t_{i n}\right|$. So intuitively, in the next step we replace the information about the length of the input packet that we have available in the type bound to variable $z$. In
the example, the refinement $\left|v \cdot p k t_{i n}\right|+112=\left|x \cdot p k t_{i n}\right|$ expresses that the input packet is 112 bits shorter than it is in the input heap $(x)$. Accordingly, we can also use this information directly to describe the length of $y \cdot p k t_{i n}$, i.e. $\left|y \cdot p k t_{i n}\right|+112=\left|x \cdot p k t_{i n}\right|$. We can repeat this step for all refinements referring to the substituted type $z$ to ultimately end up with the following equivalent type.

$$
\begin{aligned}
\Gamma, x: \tau_{\text {in }} \vdash \tau_{c} \doteq\{y: \top \mid & \mid \\
& x \cdot p k t_{\text {in }}\left|+112=\left|x \cdot p k t_{\text {in }}\right| \wedge\right. \\
& \left|y \cdot p k t_{\text {out }}\right|=\left|x \cdot p k t_{\text {out }}\right| \wedge \\
& y . p k t_{\text {out }}=x \cdot p k t_{\text {out }} \wedge \\
& y . \text { ether.valid } \wedge \\
& \left.y . \text { ether }=x . p k t_{\text {in }}[\mathrm{o}: 112]\right\}
\end{aligned}
$$

With regard to the SMT encoding, only 12 constants are generated for this type instead of 24 constants for which the SMT solver must solve constraints. If we consider longer, non-trivial programs, this difference becomes even bigger, since the calculated substitution types grow with the program length. If we directly inline the resulting substitution types during type checking of command sequences, the size of computed types remains constant.

However, there are also substitution types that cannot be inlined. For example, let us assume we are in a context, where variable $x$ binds to type $\left\{x: \epsilon| | x . p k t_{i n} \mid>\right.$ $\left.55 \wedge\left|x . p k t_{\text {out }}\right|>55\right\}$, i.e., the length of the input and output packet together is at least 112 bits.

$$
\begin{aligned}
& \left\{y: \top \mid y . \text { ether.valid } \wedge y . \text { ether }=z . p k t_{\text {in }}[\mathrm{o}: 112]\right\}[z \mapsto \\
& \left.\quad\left\{v: \top \mid v \cdot p k t_{\text {in }}=x . p k t_{\text {out }} @ x . p k t_{\text {in }}\right\}\right]
\end{aligned}
$$

If we follow a similar approach before, we end up with the following equivalent type:

$$
\left\{y: \top \mid y . \text { ether.valid } \wedge y . \text { ether }=\left(x . p k t_{\text {out }} @ x . p k t_{\text {in }}\right)[0: 112]\right\}
$$

Unfortunately, $\Pi_{4}$ 's current syntax prohibits expressing generic slices on bit vectors. Currently, only slices on header instances and on the input or output packet are supported. Since the type bound to $x$ does not specify how the 112 bits distributed over the input and output packet, there exists no equivalent type by means of packet slices. As a consequence, with $\Pi_{4}$ 's current syntax it is not possible to come up with an algorithm that is able to inline arbitrary substitution types. We therefore outline below how we can design an inlining algorithm specifically adapted to the types computed by $\Pi_{4}$ 's type checker.

We now discuss based on the commands available in $\Pi_{4}$ and their respective typing rules $^{2}$ which peculiarities arise for the inlining. Note, while we assume in the typing rules that we can derive the length of $p k t_{i n}$ and $p k t_{\text {out }}$ from the respective bit vectors, we have to handle them explicitly in our implementation and therefore consider them explicitly with respect to inlining. All typing rules adhere to the same structure, i.e., the respective output types define (1) which part of the heap remains unchanged and (2) which part of the heap has changed. The former consists of simple equalities such as $y \cdot p k t_{i n}=x \cdot p k t_{i n}$, which can be easily replaced syntactically, as we saw in the initial

[^4]example. We therefore focus mainly on the refinements that describe the changes to the heap.
reset The output type of rule T-Reset-Algo-Opt manifests the example which we have discussed before that cannot be inlined given the current syntax of $\Pi_{4}$. We therefore keep the substitution type around and only inline preceding and subsequent commands. For example, if we have a program $p ;$ reset; $s$, the resulting heap type is of the form $\left(\tau_{s}\left[y \mapsto \tau_{r}\right]\right)\left[z \mapsto \tau_{p}\right]$, with $\tau_{s}$ being the inlined heap type for all subsequent commands, $\tau_{r}$ being the inline heap type for command reset and $\tau_{p}$ being the inlined heap type of all preceding commands.
$c_{1}$ as $\left(x: \tau_{1}\right) \rightarrow \tau_{2}$ As the ascribed output type $\tau_{2}$ can be an arbitrary heap type, it may also contain statements that are not inlinable. Since there is currently no trivial solution to inline ascribed types as this would require a generic inlining approach, we skip the inlining for such commands. However, for a command $c_{1}$ as $\left(x: \tau_{1}\right) \rightarrow \tau_{2} ; c_{2}$ it is still worth to inline the heap type for $c_{1}$, even though it will ultimately be replaced by the ascribed type $\tau_{2}$.
skip Since rule T-Skip-Algo-Opt only asserts that the heap remains unchanged, we can straightforwardly inline all equalities as described before.
$\operatorname{add}(\iota)$ The changes introduced by typing rule T-Add-Algo-Opt are captured by the following two refinements: $y . \iota . v a l i d$ and $y . \iota=v$, where $v$ is a bit vector initialized with a statically known value such as $\langle 0 . . .0\rangle$. We can inline command sequences of the form $\operatorname{add}(\iota) ; c_{2}$ by replacing validity checks referencing the substitution with the boolean literal true, and references to instance $\iota$ can be replaced with the value of $v$. If preceded by another command, i.e. $c_{1} ; \operatorname{add}(\iota)$, nothing needs to be inlined because the two expressions do not depend on the input type.
remove ( $\iota$ ) Analogous to the inlining of command $\operatorname{add}(\iota)$, we can inline the refinement $\neg y$.ı.valid by replacing validity checks referencing the substitution with the boolean literal false.
extract ( $\iota$ ) If we leave out all refinements capturing the part of the heap that remains unchanged, the following type describes the changes introduced by command $\operatorname{extract}(\iota)$, where $r=\operatorname{sizeof}(\iota)$.
\[

$$
\begin{aligned}
\{y: \top \mid & y . \iota . v a l i d \wedge \\
& y . \iota=x \cdot p k t_{i n}[\mathrm{o}: r] \wedge \\
& x \cdot p k t_{i n}[\mathrm{o}: r] @ y \cdot p k t_{i n}=x \cdot p k t_{i n} \wedge \\
& \left.\left|y \cdot p k t_{i n}\right|+r=\left|x . p k t_{i n}\right|\right\}
\end{aligned}
$$
\]

As before, references to $y$.ı.valid are replaced by boolean literal true. References to $y . \iota$ are replaced by sub-slices on $x . p k t_{i n}[\mathrm{o}: r]$, depending on the fields defined for instance $\iota$. As extract $(\iota)$ is the only command except for reset that alters $p k t_{i n}$, and since we do not inline reset commands, the only interesting case arises from inlining two consecutive extract $(\iota)$ commands, for example, extract $\left(t_{1}\right) ; \operatorname{extract}\left(t_{2}\right)$. With $r_{1}=\operatorname{sizeof}\left(\iota_{1}\right), r_{2}=\operatorname{sizeof}\left(t_{2}\right)$, and binder $x$ referencing the input type, the resulting substitution type with only the refinements on the input packet looks as
follows.

$$
\begin{aligned}
& \left\{y: \mathrm{T} \mid z \cdot p k t_{i n}\left[\mathrm{o}: r_{2}\right] @ y . p k t_{i n}=z \cdot p k t_{i n} \wedge\right. \\
& \left.\left|y \cdot p k t_{i n}\right|+r_{2}=\left|z \cdot p k t_{i n}\right|\right\}[z \mapsto \\
& \left\{v: \top \mid x \cdot p k t_{i n}\left[\mathrm{o}: r_{1}\right] @ v \cdot p k t_{i n}=x \cdot p k t_{i n} \wedge\right. \\
& \left.\left.\left|v \cdot p k t_{i n}\right|+r_{1}=\left|x \cdot p k t_{i n}\right|\right\}\right]
\end{aligned}
$$

The inlined heap type captures that after executing the two consecutive extract commands, $x . p k t_{i n}[r:]$ remains left of the input packet, which is captured by the following type:

$$
\begin{gathered}
\left\{y: \top \mid x \cdot p k t_{i n}\left[\mathrm{o}: r_{1}\right] @ x . p k t_{i n}\left[r_{1}: r_{1}+r_{2}\right] @ y \cdot p k t_{i n}=x \cdot p k t_{i n} \wedge\right. \\
\left.\left|y \cdot p k t_{i n}\right|+r_{1}+r_{2}=\left|x \cdot p k t_{i n}\right|\right\}
\end{gathered}
$$

with $r=r_{1}+r_{2}$

$$
\doteq\left\{y: \top\left|x \cdot p k t_{i n}[\mathrm{o}: r] @ y \cdot p k t_{i n}=x . p k t_{i n} \wedge\right| y \cdot p k t_{i n}\left|+r=\left|x . p k t_{i n}\right|\right\}\right.
$$

remit ( $\iota$ ) We can argue similarly as before for command remit $(\iota)$, which is the only command besides reset that modifies the output packet. Accordingly, the most interesting case results from a sequence of remit commands. Again, the following substitution type captures the changes to the heap with respect to the output packet after executing $\operatorname{remit}\left(\iota_{1}\right) ; \operatorname{remit}\left(\iota_{2}\right)$.

$$
\begin{aligned}
& \left\{y: \top\left|y \cdot p k t_{\text {out }}=z . p k t_{\text {out }} @ z . \iota_{2} \wedge\right| y \cdot p k t_{\text {out }}\left|=\left|z \cdot p k t_{\text {out }}\right|+\operatorname{sizeof}\left(t_{2}\right)\right\}[z \mapsto\right. \\
& \quad\left\{v: \top\left|v \cdot p k t_{\text {out }}=x \cdot p k t_{\text {out }} @ x \cdot \iota_{1} \wedge\right| v \cdot p k t_{\text {out }}\left|=\left|x \cdot p k t_{\text {out }}\right|+\operatorname{sizeof}\left(t_{1}\right)\right\}\right]
\end{aligned}
$$

After inlining, we obtain the following type, where again $r=r_{1}+r_{2}$.

$$
\left\{y: \top\left|y . p k t_{\text {out }}=x . p k t_{\text {out }} @ x . \iota_{1} @ x . \iota_{2} \wedge\right| y . p k t_{\text {out }}\left|=\left|x . p k t_{\text {out }}\right|+r\right\}\right.
$$

ı. $f:=e \quad$ Field assignments are the only commands that operate at the level of header fields. As captured by typing rule T-Mod-Algo, in the output heap only the value of field $f$ of instance $t$ is changed, while the rest of the heap remains unchanged. This is captured by the formula $y \cdot l[l: r]=e[x /$ heap $]$, where $l$ and $r$ are the respective field bounds. The inlining itself is mostly straightforward, because as we have seen before, we can simply substitute all references to the field respectively instance slice with the assigned expression. Note, however, that we must decompose refinements that refer to the full instance. For example, given header instance $h$ with shape $\{a: 4 ; b: 4 ; c: 2\}$ and program $\mathrm{h} . \mathrm{b}:=0 b_{1011 ;}$ skip, we encounter the following substitution type.

$$
\{y: \top \mid y . \mathrm{h}=z . \mathrm{h} \wedge \ldots\}\left[z \mapsto\left\{v: \mathrm{T} \mid v \cdot \mathrm{~h}[4: 8]=\mathrm{o} b_{1011} \wedge v \cdot \mathrm{~h}[\mathrm{o}: 4]=x . \mathrm{h}[\mathrm{o}: 4] \ldots\right\}\right]
$$

To obtain the inlined type, we must decompose $z$.h properly to end up with type $\{y: \mathrm{T} \mid y . \mathrm{h}=x . \mathrm{h}[\mathrm{o}: 4] @ \mathrm{ob} 1011 @ x . \mathrm{h}[8: 10]\}$.
if $(\varphi) c_{1}$ else $c_{2}$ Regardless of whether we consider programs $c ;$ if $(\varphi) c_{1}$ else $c_{2}$ or alternatively $i f(\varphi) c_{1}$ else $c_{2} ; c$, we must inline the type for command $c$ in a union type (i.e., $\left.\left(\tau_{\text {then }}+\tau_{\text {else }}\right)\left[z \mapsto \tau_{c}\right]\right)$ or we must inline a union type in the type for command $c$ (i.e. $\tau_{c}\left[z \mapsto\left(\tau_{\text {then }}+\tau_{\text {else }}\right)\right]$ ). In both cases, both branches of the union type can be inlined independently as described before. The end result in both cases is again a union type.

```
@pi4("(Parser;Ingress) as (x:{y:\epsilon|y.pkt_in.length > 304}) ->
        {z:T|!z.vlan.valid}")
parser Parser(packet_in packet, out headers hdr, ...) {
    state start {
        packet.extract(hdr.ethernet);
        transition select(hdr.ethernet.etherType) {
            0x0800: parse_ipv4;
                0x8100: parse_vlan;
            default: accept;
        }
    }
    state parse_vlan {
            packet.extract(vlan);
            transition select(hdr.vlan.etherType) {
                0x0800: parse_ipv4;
                default: accept;
        }
    }
    state parse_ipv4 {
        packet.extract(hdr.ipv4);
        transition accept;
    }
}
```

Figure 6.11: VLAN decapsulation example.

### 6.4.2 Reducing the Number of SMT Solver Invocations

If we look at the reasons for invoking the SMT solver during type checking, the most common reason is to check the validity of headers. By strictly following the formalization with the implementation, many of these checks become redundant. For example, let us consider a somewhat realistic $\mathrm{P}_{4}$ program shown in Figures 6.11 and 6.12. The parser first parses the Ethernet header, optionally followed by VLAN and IPv4. If the IPv4 header is valid in the ingress, the program then performs a basic forwarding operation and if the VLAN header is valid, it removes the VLAN header. As indicated by the @pi4 annotation in Line 1 of Figure 6.11, we want to check the whole program, i.e., the parser followed by the ingress with type $\left(x:\left\{y: \epsilon| | y \cdot p k t_{i n} \mid>304\right\}\right) \rightarrow\{z: \top \mid$ $\neg z$.vlan.valid $\}$. This type asserts that starting in a heap where no headers are valid, and the input packet provides enough bits to extract all three header instances, we always end up in a heap for which it is guaranteed that the VLAN header is not present. The annotation mechanism provided by our $\mathrm{P}_{4}$ frontend will be explained in more detail in the next section.

Translating the $\mathrm{P}_{4}$ program into $\Pi_{4}$ 's syntax and type checking the program with the annotated type, takes about 12 seconds ${ }^{3}$, of which roughly $68 \%$ are used to check the validity of headers. A total of 18 validity checks are performed, each one resulting

[^5]```
control Ingress(inout headers hdr, ...) {
    action ipv4_forward(bit<48> dstAddr, bit<9> port) {
        standard_metadata.egress_spec = port;
        hdr.ethernet.srcAddr = hdr.ethernet.dstAddr;
        hdr.ethernet.dstAddr = dstAddr;
        hdr.ipv4.ttl = hdr.ipv4.ttl - 1;
    }
    action vlan_decap() {
        hdr.ethernet.etherType = hdr.vlan.etherType;
        hdr.vlan.setInvalid();
    }
    table ipv4_lpm {
        key = { hdr.ipv4.dstAddr: lpm; }
        actions = { ipv4_forward; }
    }
    apply {
        if (hdr.ipv4.isValid()) {
                ipv4_lpm.apply();
            }
            if (hdr.vlan.isValid()) {
                vlan_decap();
            }
    }
}
```

Figure 6.12: VLAN decapsulation example continued.
in a call to the SMT solver. If we look at the program more closely, it is noticeable that many of the checks are indeed redundant. For example, the parser guarantees that the Ethernet header is valid. If we cache this information, we save a call to the SMT solver for each subsequent read or write access to the Ethernet header, which are six in total for this example program. In total, we can save all but two solver calls for the example program this way.

Implementing this cache is mostly straightforward. If we successfully typecheck an extract or add command, we store the information that on the current program path the respective header instance is valid. After a remove command, we invalidate the instance in the cache. For conditionals, we apply a simple heuristic to the condition. If it consists of simple validity checks such as hdr.ipv4.isValid(), we add this information to the cache used for the then and else branches accordingly. Since the validity assumptions can diverge across both branches, we have to merge them for subsequent commands of conditionals. If both branches provide the same validity information for a specific instance, we keep the entry in the cache, otherwise we remove entries. After checking a reset command, we store the information that all header instances are invalid, and type ascriptions make it necessary to clear the entire cache. The reason for the latter is that when applying a type ascription, we check that the computed output type $\tau_{c}$ is a subtype of the ascribed output type $\tau_{a s c}$. This in turn means that the ascribed output type may

$$
\begin{array}{lr}
\alpha::=\beta \text { as }(x: \tau) \rightarrow \tau & \text { (annotations) } \\
\beta::=P|C| \text { reset }|\beta ; \beta| \beta \text { as }\left(x: \tau_{1}\right) \rightarrow \tau_{2} & \text { (annotation bodies) }
\end{array}
$$

Figure 6.13: Syntax of @pi4 annotations.

```
@pi4("(MyParser;MyIngress;MyDeparser) as (x: }\mp@subsup{\tau}{1}{})->\mp@subsup{\tau}{2}{\prime")
parser MyParser(packet_in p, ...) { ... }
control MyIngress(inout headers hdr, ...) { ... }
control MyDeparser(packet_out packet, ...) { ... }
```

Figure 6.14: The annotation on line 1 instructs the frontend to check the full pipeline consisting of the parser, ingress control and deparser with type $\left(x: \tau_{1}\right) \rightarrow \tau_{2}$.
be less specific than the actual type and as such we cannot transfer our knowledge about instance validity in $\tau_{c}$ to $\tau_{\text {asc }}$. When performing the validity checks, we first check the cache for an entry. If no entry is available, we have to resort to the SMT solver and store the computed result in the cache.

## 6.5 $P_{4}$ Frontend

With our $\mathrm{P}_{4}$ frontend we are able to automatically translate a part of the $\mathrm{P}_{4}$ language into $\Pi_{4}$ 's syntax. In order to be able to flexibly check $\mathrm{P}_{4}$ programs for different properties, we use the annotation mechanism of the language to specify which parts of the program should be checked with which type. For this we use the custom annotation @pi4(" $\alpha$ "), whose syntax is defined in Figure 6.13. An annotation $\alpha$ always consists of an annotation body $\beta$ ascribed with a dependent function type. The annotation body $\beta$ is either the name of a parser $P$ or a control $C$, reset, a sequence of two annotation bodies, or again an ascribed annotation body. For example, the annotation in Line 1 of Figure 6.14 expresses that the complete pipeline described by the parser MyParser, the ingress control MyIngress, and the deparser MyDeparser should be checked with type $\left(x: \tau_{1}\right) \rightarrow \tau_{2}$. We allow the programmer to use certain unicode symbols in annotations, e.g., $T$ or $\epsilon$, but also more convenient notations, for example, \&\& and || for conjunction and disjunction (instead of $\wedge$ and $\vee$ ) respectively. The parser and control blocks are translated into a $\Pi_{4}$ command $c$ and likewise reset, $\beta ; \beta$ and $\beta$ as $\left(x: \tau_{1}\right) \rightarrow \tau_{2}$ are translated into the corresponding $\Pi_{4}$ commands.

Since $\Pi_{4}$ is an imperative language, we unroll parsers during translation. We model match-action tables using an encoding similar to the one used in p4v [Liu+18], where we create an extra header that captures the match keys and selected action and action parameters. Figure 6.15 shows how table routing_v6 declared in control Forwarding is translated. The table header type declaration in the lower half of Figure 6.15 contains a one bit field act for modeling the choice of the actions, since a single bit is sufficient to represent two actions. Field set_next_id_next_id represents the action parameter next_id for action set_next_id and field ipv6_dst_addr_key models the match key of table routing_v6. The table is then translated into a set of conditional commands, where the conditions encode the matching on the key and the non-deterministic choice of actions. There are limitations in the translation particularly with regard to

```
control Forwarding(...) {
    action set_next_id(bit<32> next_id) {
        meta.next_id = next_id;
    }
    table routing_v6 {
            key = { hdr.ipv6.dst_addr: exact; }
            actions = {
                set_next_id; NoAction;
            }
            default_action = NoAction();
    }
}
```

header_type Forwarding_routing_v6_table_t \{
act: 1;
set_next_id_next_id: 32;
ipv6_dst_addr_key: 128;
\}
header Forwarding_routing_v6_table :
Forwarding_routing_v6_table_t
if(Forwarding_routing_v6_table.ipv6_dst_addr_key == ipv6.
dst_addr) \{
if(Forwarding_routing_v6_table.act == 0b0) \{
meta.next_id := Forwarding_routing_v6_table.
set_next_id_next_id
\} else \{
if(Forwarding_routing_v6_table.act == 0b1) \{
skip
\}
\}
\}

Figure 6.15: Translation of $\mathrm{P}_{4}$ tables into $\Pi_{4}: \mathrm{P}_{4}$ table declaration (top) and result of translation (bottom)
externs such as stateful elements or hash computations, since these cannot currently be modeled with the syntax of $\Pi_{4}$.

### 6.6 Chapter Summary

This chapter outlined the underlying details of the implementation of $\Pi_{4}$ 's type checker. We started by defining the algorithmic version of ח4's type system. We exploited the fact that network switches are limited by the Maximum Transmission Unit (MTU) in the size of packets they can receive and transmit, solving the problem that our heap types describe possibly unbounded bit vectors. We were able to prove an upper bound on the
size of computed output types under the assumption that the input type as well as all ascribed types are bound by the MTU as well. Ultimately, this allowed us to prove our type checking algorithm to be decidable. The MTU also allows us to encode subtyping checks into a formula in the theory of uninterpreted functions and fixed-width bit vectors (UFBV), facilitating automated subtyping checks. Finally, we discussed several optimizations to speed up solving SMT queries in terms of simpler encodings, inlining substitution types and by reducing the number of SMT solver invocations by caching assumptions about header validity.

A formal proof that our encoding is correct is left for future work. Since the performance of the SMT solver also depends heavily on the size of the bit vectors used, a possible further optimization is to adapt the number of bits to the size of heap types used to check a program. However, this requires being able to compute a bound for any heap type. Currently, it is unclear if that is generally possible. In case it is, this would also allow us to remove the constraint from our decidability theorem that every ascribed type must be bounded.

## Part III

## Evaluation

## Header Validity Bugs in Real-world Programs

In Chapter 3, we identified five syntactic constructs of the $\mathrm{P}_{4}$ language that are vulnerable to header validity bugs. In Chapter 4, we then introduced SAFEP4, a type system that statically ensures that all headers that are read or written, are valid. In this chapter, we want to answer whether SAFEP4's type system is able to identify the different types of header validity bugs in real-world $\mathrm{P}_{4}$ programs. We also investigate how often these bugs occur in practice and how much effort is required by a programmer to fix header validity bugs. Accordingly, we formulate the following three research questions:

RQ1 Is SAFEP4 able to detect header validity bugs in real-world programs?
RQ2 How often do the different kinds of bugs occur in practice?
RQ3 What is the overhead for fixing these bugs?
To evaluate the type system proposed in Chapter 4, we implemented P $_{4}$ Снеск, a tool that automatically checks whether P 4 programs comply with the rules presented in Section 4.3 and reports the respective violations if not. Our implementation uses the frontend provided by $\mathrm{p} 4 \mathrm{v}[\mathrm{Liu}+18]$ and is able to cover the whole $\mathrm{P}_{4_{14}}$ language. We used $\mathrm{P}_{4}$ Сheck to check 15 open source programs that differ in size and complexity, ranging from 143 lines of code to 9600 lines of code. We selected these programs based on the following criteria: (1) each program must be open-source, (2) publicly available on GitHub, (3) compile without errors and (4) contain either industrial or academic code that implements standard or novel network functionality, i.e., we excluded programs primarily used for teaching purposes.

Out of the 15 programs, only four did not violate any of our typing rules and passed the checker. These were primarily implementations of simple routers or DDoS mitigation mechanisms, only consisting of a few lines of code (188-635 lines) and only accepting a few packet types. For the remaining 11 programs, our tool found a total of 418 violations. The high number of violations stems from the fact that each individual violation is counted, even if multiple violations share the same root cause. For example, if a single


Figure 7.1: Proportional frequencies of each bug type per-program. The raw number of bugs for each program and category is reported on each stacked bar.
action modifies the source address and destination address field of an IPv 4 header in a context that cannot prove that IPv 4 is valid, then both references will be reported as violations, even though they are due to the same control bug. We therefore use another metric inspired by the metric proposed in $[\mathrm{Kle}+18]$. We relate the number of bugs in each program to the number of bug fixes required for the program to successfully type-check. Using this metric, we counted 58 bugs.

We classified the bugs according to the categories presented in Chapter 3. Figure 7.1 shows which bug type occurs how often per program and Figure 7.2 shows how often the bugs from the different categories occur overall. It is particularly noticeable that although table action bugs occur most often numerically ( 22 times), there are only occurrences in one program (switch.p4). The reason for this is that switch.p4 heavily relies on correct control plane configurations. In contrast, there were only nine respectively eight occurrences of parser bugs and table read bugs across five programs. Compared to p4v [Liu+18], P4CHECK was unable to detect any default action bugs in switch.p4, while p4v reported many of such bugs, which has two reasons. First, p4v allows programmers to specify complex properties that involve fine-grained conditions on tables and the relationships between tables. In contrast, SAFEP4 makes various assumptions that rule out a variety of bugs, including some default action bugs. Second, our repairs are often coarse-grained, potentially forcing stronger guarantees on the program than necessary. In contrast, p 4 v uses first-order logic annotations, which allows the programmer to formulate the weakest and hence more complex assumptions.

### 7.1 Detecting and Repairing Bugs

We will now look at how $\mathrm{P}_{4}$ СНеск detects bugs from the categories identified in Chapter 3 and how the repairs affect the computed types.

Parser bugfix We have shown an example of a parser bug in Figure 3.1. The bug occurs because the ingress expects packets where the IPv4 header and the TCP header are valid, but the parser does not guarantee that only such packets are successfully parsed. Instead, program execution switches to the ingress control, when the parser encounters an unexpected header.

Our type system is able to detect this bug because the ingress expects packets of type ethernet•ipv4•tcp, i.e., it must be guaranteed that all three packet headers Ethernet,


Figure 7.2: Frequency of each bug across all programs. The raw number of bugs in each category is reported to the right of the bar.

```
./h.p4, line 350, cols 12-21: error tcp not guaranteed to be valid
./h.p4, line 118, cols 8-16: error ipv4 not guaranteed to be valid
./h.p4, line 101, cols 42-50: error ipv4 not guaranteed to be valid
./h.p4, line 320, cols 8-15: error tcp not guaranteed to be valid
./h.p4, line 362, cols 12-19:error tcp not guaranteed to be valid
./h.p4, line 362, cols 29-36: error tcp not guaranteed to be valid
./h.p4, line 295, cols 60-69: error tcp not guaranteed to be valid
./h.p4, line 107, cols 8-16: error ipv4 not guaranteed to be valid
./h.p4, line 101, cols 42-50: error ipv4 not guaranteed to be valid
./h.p4, line 163, cols 8-16: error ipv4 not guaranteed to be valid
./h.p4, line 101, cols 42-50: error ipv4 not guaranteed to be valid
```

```
./h.p4, line 350, cols 12-21: error tcp not guaranteed to be valid
./h.p4, line 320, cols 8-15: error tcp not guaranteed to be valid
./h.p4, line 362, cols 12-19: error tcp not guaranteed to be valid
./h.p4, line 362, cols 29-36: error tcp not guaranteed to be valid
./h.p4, line 295, cols 60-69: error tcp not guaranteed to be valid
```

Figure 7.3: Curated output from Р4Снеск for the parser bug in NetHCF before (above) and after (below) modifying parse_ethernet

IPv4 and TCP are valid on all program paths. However, the parser only produces packets of type ethernet $\cdot(1+i p v 4 \cdot(1+t c p))$, which means that only Ethernet is guaranteed to be valid, while the IPv4 header and the TCP header are both optional. Accordingly, $\mathrm{P}_{4}$ Check reports each reference to the IPv 4 header and TCP header as a violation of the type system as shown in the top half of Figure 7.3. The ubiquity of the reports intimates a mismatch between the parsing and the control types, which gives the programmer a hint as how to fix the problem.

If we fix the program step-by-step and replace in parse_ethernet the default clause with a parser exception as shown in Line 10 of Figure 3.2, the parser henceforth guarantees that Ethernet and IPv4 are valid in the ingress, i.e, parsed packets have type ethernet • ipv4•(1+tcp). If we run the tool again, all violations related to the IPv4 header are removed from the output, as shown in the bottom of Figure 7.3. Additionally applying the second fix (cf. Line 18 of Figure 3.2) causes P4CHECK to output no violations, since the type upon entering the ingress control is ethernet $\cdot \mathrm{ipv} 4 \cdot \mathrm{tcp}$, resulting in
port.p4, line 248, cols 8-24: warning: assuming either vlan_tag_[0] matched as valid or vlan_tag_[0].vid wildcarded
port.p4, line 250, cols 8-24: warning: assuming either vlan_tag_[1] matched as valid or vlan_tag_[1].vid wildcarded
fabric.p4 line 42 , cols $41-67:$ warning: assuming fabric_header_cpu
matched as valid for rules with action terminate_cpu_packet
fabric.p4, line 57, cols $17-54:$ warning: assuming
$\quad$ fabric_header_unicast matched as valid for rules with action
terminate_fabric_unicast_packet
fabric.p4, line 81, cols $17-56:$ warning: assuming
$\quad$ fabric_header_multicast matched as valid for rules with action
terminate_fabric_multicast_packet
Figure 7.4: Warnings printed after fixing switch.p4's reads bug (top), and its actions bug (bottom)
all subsequent references to the IPv 4 and TCP headers being safe.

Control bugfix Recall that a control bug occurs when the incoming type presents a choice between two instances, but subsequent code expects one of the instances to be valid. The control bug example in Figure 3.4 uses a parser that produces the type $\Theta=$ ethernet $\cdot\left(1+\operatorname{ipv} 4 \cdot\left(1+\mathrm{udp} \cdot\left(1+\mathrm{nc} \_\right.\right.\right.$hdr $\left.\left.\left.\cdot \tau\right)+\mathrm{tcp}\right)\right)$, where $\tau$ is a type for caching operations. Especially, this type suggests that Includes $\Theta$ nc_hdr does not hold, however, controls process_cache and process_value only typecheck in contexts where Includes $\Theta$ nc_hdr is true.

Accordingly, $\mathrm{P}_{4} \mathrm{CHECK}$ reports type violations at every reference to nc_hdr. Fixing this error is simply a matter of wrapping both calls to process_cache() and process_value( ) in a validity check as demonstrated in the top right of Figure 3.4. As a result, the type inside the validity check becomes $\Theta=$ ethernet•ipv4•udp•nc_hdr $\cdot \tau$ and thus Includes $\Theta$ nc_hdr is always true. As NetCache handles TCP and UDP packets as well as its special-purpose packets, we can't include the application of the IPv4 routing table in the validity check. This is another instance of a parser bug, as the type does not guarantee Includes $\Theta$ ipv4, which is required by table ipv4_route.

Table reads bugfix As shown in Figure 3.5, table reads errors occur when a header $h$ is included in the reads declaration of a table $t$ with match kind $k$, and $h$ is not guaranteed to be valid at the call site of $t$, and if $h \notin$ valid_reads $(t)$ or the match-kind of $k \neq$ ternary. In the case of the port_vlan_mapping table in Figure 3.6, there is a valid bit for both vlan_tag_[0] and vlan_tag_[1], both of which are followed by exact matches. We fixed this bug by using a ternary match-kind instead, which allows the use of wildcard matching, since when a field is matched with a wildcard, the table does not attempt to compute the value of the reads expression.

However, this fix assumes that the controller is well-behaved and fills the table entry for vlan_tag_[0].vid with a wildcard whenever vlan_tag_[0] is matched as invalid (and similarly for vlan_tag_[1]). This also what SAFEP4's type system does, with its maskable checks in typing rule T-Apply. P4СНеск prints warnings describing these assumptions to the programmer as shown in the top of Figure 7.4, giving them guidelines against which to check their control plane implementation.

Table action bugfix As exemplified with table fabric_ingress_dst_lkp in Figure 3.6, table actions bugs occur when at least one action cannot be safely executed in all scenarios. In the example, the parser parses exactly one of the three headers (1) fabric_hdr_cpu, (2) fabric_hdr_unicast and (3) fabric_hdr_multicast. Thus, when the table is applied at type $\Theta$, exactly one of Includes $\Theta$ fabric_hdr_i for $i \in\{$ cpu, unicast, multicast $\}$ will hold. Now, the action term_cpu_packet type-checks only with the (nonempty) type Restrict $\Theta$ fabric_hdr_cpu, and the actions term_fabric_i_packet only typecheck with the (nonempty) types Restrict $\Theta$ term_fabric_i_packet for $i=$ unicast, multicast
$\mathrm{P}_{4}$ Снеск suggests that this is the cause of the bug since it reports type violations for all references to these three headers in the control paths following from the application of fabric_ingress_dst_lkp. As shown in the bottom of Figure 3.6, the optimal fix is to augment the reads declaration to include a validity check for each contentious header. We then assume that the controller is well-behaved enough to only call actions when their required headers are valid, allowing us to typecheck each action in the appropriate type restriction. Again, $\mathrm{P}_{4}$ CHECK alerts the programmer whenever it makes such an assumption. We show these warnings for the fixed version of fabric_ingress_dst_lkp in the bottom part of Figure 7.4

Default action bugfix Recall that default action bugs, such as the one shown in Figure 3.7, occur when a programmer creates a wrapper table for an action that modifies the type, and forgets to force the table to call that action when the packet misses. Table add_value_header_1 wraps action add_value_header_1_act, which itself executes add_header(nc_value_1). The default action, when left unspecified, is nop, which means that if the pre-application type was $\Theta$, then the type after applying the table is $\Theta^{\prime}=\Theta+\Theta \cdot n c \_v a l u e \_1$. For this type Includes $\Theta^{\prime}$ nc_value_1 is false, hence, P4CHECK reports every subsequent reference (on this code path) to nc_header_1 to be a type violation. Fixing this bug by setting the default action to add_value_1 makes the post-application type $\Theta \cdot n c \_v a l u e \_1+\Theta \cdot n c \_v a l u e \_1=\Theta \cdot n c \_v a l u e \_1$, and therefore Includes $\Theta^{\prime}$ nc_value_1 is true, which allows the subsequent code to typecheck.

### 7.2 Overhead

It is important to evaluate two kinds of overhead when considering a static type system: overhead on programmers and on the underlying implementation. Typically, adding a static type system to a dynamic type system requires more work for the programmerthe field of gradual typing is devoted breaking this task into smaller commit-sized chunks [Cam+17]. Surprisingly, in our experience, migrating real-world $\mathrm{P}_{4}$ code to pass the $\mathrm{SaFEP}_{4}$ type system only required modest programmer effort.

To qualitatively evaluate the effort required to change an unsafe program into a safe one using our type system, we manually fixed all detected bugs. The programs that had bugs required us to edit between $0.10 \%$ and $1.4 \%$ of the lines of code. The one exception was PPPoE_using_P4, which was a 143 line program that required 6 line-edits (4\%), all of which were validity checks. Conversely, switch.p4 required 34 line edits, the greatest observed number, but this only accounted for $0.37 \%$ of the total lines of code in the program. Each class of bugs has a simple one-to-two line fix, as described in Section 7.1: adding a validity check, adding a default action, or slightly modifying the parser. Each of these changes was straightforward to identify and simple to make.

Another possible concern is that extending tables with extra read expressions, or adding run-time validity checks to controls, might impose a heavy cost on implementations, especially on hardware. Although we have not yet performed an extensive study of the impact on compiled code, based on the size and complexity of the annotations we added, we believe the additional cost should be quite low. We were able to compile our fixed version of switch.p4 program to the Tofino architecture [Bos18] with only a modest increase in resource usage. Overall, given the large number of potential bugs located by $\mathrm{P}_{4}$ Снеск, we believe the assurance one gains about safety properties by using a static type system makes the costs well worth it.

### 7.3 Chapter Summary

In this chapter, we have shown how we can use SAFEP4's type system in practice to check real $\mathrm{P}_{4}$ programs for header validity bugs according to the taxonomy presented in Chapter 3 ( $R Q_{1}$ ). For this purpose we have shown exemplarily which effects buggy programs as well as the repaired versions of these programs have on the computed types. Validity bugs were present in the majority of the programs we examined, which clearly indicates that validity bugs are ubiquitous in practice. Number-wise, table-action bugs occurred most frequently, but only in a single program, which was significantly larger than the others in terms of the number of lines of code. If we look at how often certain bug categories appear in different programs, control and parser bugs predominate ( RQ 2 ). Our evaluation has shown that validity bugs can usually be fixed with little effort, which makes the effort for the programmer manageable. A first evaluation indicates that the additional code required to fix the header validity bugs does not result in significantly higher hardware resource consumption ( $R_{2} 3$ ).

## Expressivity of $\Pi_{4}$

One of our main goals in designing $\Pi_{4}$ was to create a type system that is capable of verifying relevant network properties beyond header validity but that retains at the same time the modularity inherent to type systems. In this chapter, we evaluate to what extent $\Pi_{4}$ can bridge the gap between $\mathrm{SAFEP}_{4}$ and full-fledged verification tools in terms of expressiveness, and whether $\Pi_{4}$ is actually able to verify $\mathrm{P}_{4}$ programs in a modular fashion. We thereby answer the following two research questions:

RQ1 Can $\Pi_{4}$ be used to verify practically relevant properties beyond header validity?
$\mathbf{R Q}_{2}$ Can $\Pi_{4}$ be used to verify $\mathrm{P}_{4}$ programs in a modular fashion?
We start this chapter with a brief overview of properties discussed by other $\mathrm{P}_{4}$ verification tools ranging from basic safety properties to advanced safety properties. We then show how the most common of these properties and other practically relevant properties can be expressed using our heap types and verified by $\Pi_{4}$ 's type checker. Finally, based on a case study, we investigate how we can facilitate $\Pi_{4}$ 's type system to modularly verify $\mathrm{P}_{4}$ programs.

### 8.1 Survey

We surveyed the publications on recent $\mathrm{P}_{4}$ verification tools with respect to the network properties discussed. Table 8.1 lists seven properties that were discussed in the context of at least two tools. Other properties discussed are mainly application-specific and are therefore not considered further. For each of the verification tools including $\Pi_{4}$, a checkmark indicates whether the respective property can be verified or not.

Header validity As already motivated in Chapter 3, accessing an invalid header instance in $\mathrm{P}_{4}$ yields undefined values, which in turn can result in a variety of subtle bugs. This problem has already been recognized, and as shown in Table 8.1, most $\mathrm{P}_{4}$ verification tools provide the ability to detect invalid headers accesses. As it was already the case with SAFEP4, header validity is also a central safety property enforced by $\Pi_{4}$ 's type system. The type checker rejects any program that attempts to access invalid headers.

|  |  | U ت̈ 0 0 0 0 0 0 0 0 0 0 | $\mathscr{U}$ <br> 0 <br> $\ddot{U}$ <br> 0 <br> 0 <br> 0 <br> 0 <br> 0 <br> 0 |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| ASSERT-P4 [Fre+18] | $\checkmark$ | $\checkmark$ |  |  |  |  |  |
| p4v [Liu+18] | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ |
| Vera [Sto+18] | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ |  | $\checkmark$ |
| P4RL [Shu+19] |  | $\checkmark$ |  |  |  | $\checkmark$ |  |
| $\mathrm{bf}_{4}$ [Dum+20] | $\checkmark$ |  | $\checkmark$ |  | $\checkmark$ | $\checkmark$ |  |
| $\Pi_{4}$ | $\checkmark$ | $\checkmark$ |  |  | $\checkmark$ | $\checkmark$ | $\checkmark$ |

Table 8.1: Common safety properties examined by other $\mathrm{P}_{4}$ verification tools.

The main difference is that $\Pi_{4}$ is able to ensure header validity even in the presence of data-dependent validity checks. A corresponding example where the validity of the IPv4 header is checked depending on the EtherType field of the Ethernet header, has already been discussed in Section 5.1.

Protocol conformance To ensure interoperability between network devices of different vendors, network protocols are usually standardized. For example, these standards specify which values certain fields can contain, which other protocols are encapsulated in the payload of the packet depending on certain header fields, or how packets must be processed in certain situations. Protocol conformance is an umbrella term for a variety of properties that ensure that protocols are implemented correctly. A frequently considered example is the time-to-live (TTL) of an IPv4 packet, which limits how often a packet can be forwarded before it must be discarded by the switch. We show in Section 8.2.1 how we can express and verify protocol conformance such as the mentioned example using $\Pi_{4}$ 's types.

Out-of-bound accesses P4's header stacks represent fixed-size arrays of headers. Accordingly, errors typical for arrays can occur, such as out-of-bounds accesses, which a programmer usually wants to rule out. Currently, $\Pi_{4}$ lacks the support for header stacks. However, since this is not a fundamental limitation and avoiding out-of-bounds access is a prime example of the use of dependent types, it should be possible to extend $\Pi_{4}$ accordingly in the future.

Arithmetic overflow Arithmetic operations in $\mathrm{P}_{4}$ do not detect overflows or underflows and should therefore be eliminated. Large values not fitting into a fixed number of bits will be cut off, which might lead to unexpected behavior of the program. Currently, it is not possible to check arithmetic overflows using $\Pi_{4}$ 's type system. Since our heap
types capture values down to the bit level, it seems possible to support this property in the future.

Determined forwarding Another basic safety property, is determined forwarding. Typical $\mathrm{P}_{4}$ programs contain thousands of paths on which a packet can be processed. To avoid situations where packets are dropped unexpectedly, a desirable invariant is that each program path contains an explicit forwarding decision-i.e., packets are either forwarded on some switch port or dropped. We show in Section 8.2.2 how we can encode this invariant with $\Pi_{4}$ 's heap types.

Parser-Deparser compatibility A P4 program typically defines the parser, controls for ingress and egress pipelines, and the deparser. The main reason for this four-phase structure is that separate ingress and egress pipelines allow packet processing to occur both before and after packets are scheduled, typically using one or more queues. In practice, parsing and deparsing may also happen between the ingress and egress stagesi.e., the deparser code is additionally executed at the end of the ingress followed by the parser code, before the egress. In such cases, it is important to ensure that data intended to be carried from ingress to egress is serialized and deserialized correctly. Otherwise, headers may be unexpectedly removed from the packet. We discuss this invariant in more detail in Section 8.2.3 and show how this invariant can be checked using $\Pi_{4}$.

Read-only fields In $\mathrm{P}_{4}$ there are certain metadata fields that are read-only (e.g. the egress_port), yet the compiler does not rule out writes but instead silently ignores them. To prevent unexpected behavior of the program, it is therefore desirable to prohibit all write access to read-only fields. We show in Section 8.3 how an instance of this invariant can be checked using $\Pi_{4}$.

### 8.2 Checking Network Invariants

We now show how $\Pi_{4}$ 's type system can be used to check real network protocol invariants and verify a variety of basic and advanced safety properties as discussed before. In most $\mathrm{P}_{4}$ programs, the packet-forwarding behavior of the device is specified using a predefined record of type standard_metadata_t. In particular, the field egress_spec is used to instruct the switch to forward the packet on a specific port. We assume that the field is initialized to oxoo, indicating that no forwarding decision has been made, and that by setting the field to oxiFF (i.e., the largest unsigned integer that can be encoded into 9 bits, which is the width of the field), the switch can be instructed to drop the packet. For simplicity, we treat $\mathrm{P}_{4}$ metadata as an ordinary header instance.

There is a general pattern that can be used to encode invariants in types. Given a program that typechecks with some type $\left(x: \tau_{i n}\right) \rightarrow \tau_{\text {out }}$, we can instead typecheck the program with type $\left(x: \tau_{\text {in }}\right) \rightarrow\left\{y: \tau_{\text {out }} \mid \varphi_{\text {inv }}\right\}$, which refines the output type with an expression describing an invariant ( $\varphi_{i n v}$ ). By doing so, we effectively filter out all heaps for which the invariant does not hold. If the type computed by the type checker is not a subtype of the annotated output type, there must be some heap allowed by the program, for which the invariant does not hold.

```
/* Unsafe */
if(ipv4.valid) {
    stdmeta.egress_spec := 0x1 ;
    ipv4.ttl := ipv4.ttl - 1
}
```

```
/* Safe */
if(ipv4.valid) {
    if(ipv4.ttl == 0) {
        stdmeta.egress_spec := 0xlff
    } else {
        stdmeta.egress_spec := 0x1;
        ipv4.ttl := ipv4.ttl - 1
    }
}
```

```
(x:{y:ipv4}||y.stdmeta.valid}) ) ,
    {y:ipv4}\mp@subsup{|}{|}{|}y.\mathrm{ stdmeta.valid }
    (y.ipv4.ttl =o\Longrightarrowy.stdmeta.egress_spec = 0x1ff)}
```

Figure 8.1: IPv4 TTL example. Top: property violated; middle: property holds; bottom: $\Pi_{4}$ type encoding the TTL invariant.

### 8.2.1 Protocol conformance

We start with examples showing how $\Pi_{4}$ 's type system can be used to ensure that a program conforms with standard network protocols.

IPv4-Time To Live For Internet Protocol (IP) packets, the time to live (TTL) limits how often a packet can be forwarded from one network switch to another. Every time a packet is forwarded, the TTL is decremented; when the TTL is zero before the packet has reached its destination, forwarding halts to eliminate the risk of infinite loops. ${ }^{1}$

In the code snippet in the top of Figure 8.1, the intended behavior is violated, because the packet is always forwarded on the same port while the TTL is decremented. We can detect this violation by checking the program with the type shown in the bottom of Figure 8.1, which reads as follows: starting in a heap where at least IPv4 and the standard metadata is valid, after executing the ingress code, still at least IPv 4 and the standard metadata is valid and if the IPv4 TTL is zero, the value of egress_spec indicates that the packet will be dropped. The middle of Figure 8.1 shows a program that successfully typechecks with the given type.

IPv4 Options The standard IPv4 header consists of at least 160 bits, but it may also carry additional data in optional fields. The Internet Header Length (IHL) field specifies the length of the header as multiples of 32 and indicates whether additional data is

[^6]```
/* Unsafe */
extract(ethernet);
if(ethernet.etherType == 0x0800) {
    extract(ipv4)
}
```

```
/* Safe */
extract(ethernet);
if(ethernet.etherType == 0x0800) {
    extract(ipv4);
    if(ipv4.ihl != 0x5) {
        extract(ipv4_opt)
    }
}
```

    \(\left(x:\left\{y: \epsilon| | y \cdot p k t_{i n} \mid>592\right\}\right) \rightarrow\)
        \(\left\{y: \mathrm{T} \mid\left(\left(y . i p v_{4} . v a l i d \wedge y . i p v 4 . i h l \neq 5\right) \Longrightarrow\right.\right.\) y.ipv4_opt.valid \() \wedge\)
            \(\left((y . i p v 4 . v a l i d \wedge y . i p v 4 . i h l=5) \Longrightarrow \neg y . i p v 4 \_\right.\)opt.valid \(\left.)\right\}\)
    Figure 8.2: IPv4 Options example. Top: property violated; middle: property holds; bottom: $\Pi_{4}$ type encoding the IPv4-Option specification.
available. The minimum IHL is $5(5 * 32=160)$ and the maximum value is 15 . Due to their flexibility, IP options are notoriously difficult to parse, and many real-world network devices handle them incorrectly.

We can use $\Pi_{4}$ 's type system to ensure that we also extract the $\operatorname{IPv} 4$ options from the input packet, whenever IPv 4 is valid and $\mathrm{IHL}>5$. Figure 8.2 provides one example where this property is violated (top) and one where it holds (middle). To rule out violations, we can check the respective programs with the type at the bottom of Figure 8.2. This type states that executing the parser in the empty heap where enough bits are available to extract Ethernet, IPv4 and IPv4 options, produces a heap satisfying the constraint that if IPv4 is valid, then either IHL is 5 and IPv4 options are not valid, or IHL $>5$ and IPv4 options are valid.

Header Dependencies Most protocols have some way of keeping track of what other protocols are encapsulated in the payload of a packet-i.e., which header follows next. The correspondence between field values and protocols is typically defined as part of the protocol standard. For example, an Ethernet frame uses the EtherType field (written ethernet. etherType) for this purpose: a value of oxo8oo indicates that the next header is an $\mathrm{IPv}_{4}$ header, while, for example, a value of ox86DD indicates that the next header is an IPv6 header. The code snippet in the top of Figure 8.3 violates the dependency between the $\mathrm{IPv}_{4}$ header and the EtherType field of the Ethernet header. Our type checker detects this violation by checking that executing the parser in an empty heap, with enough bits to extract both Ethernet and IPv4, produces a heap with either an invalid IPv4 header or a valid IPv4 header and an EtherType value of oxo8oo, which is captured by the type shown in the bottom of Figure 8.3. An example of a program

```
/* Unsafe */
extract(ethernet);
extract(ipv4)
/* Safe */
extract(ethernet);
if(ethernet.etherType == 0x0800) {
    extract(ipv4)
}
```

$$
\begin{aligned}
(x: & \left.\left\{y: \epsilon| | y \cdot p k t_{i n} \mid>272\right\}\right) \rightarrow \\
& \{y: \mathrm{T} \mid y . \mathrm{ipv} 4 . v a l i d \Longrightarrow y . e t h e r n e t . e t h e r T y p e==0 x 0800\}
\end{aligned}
$$

Figure 8.3: Header dependency example. Top: property violated; middle: property holds; bottom: $\Pi_{4}$ type encoding IPv4's dependency on Ethernet.

```
/* Unsafe */
if(ipv4.valid) {
    if(ipv4.dstAddr != 0x0a0a0a0a) {
        stdmeta.egress_spec := 0x1
    }
}
```

```
/* Safe */
```

/* Safe */
if(ipv4.valid) {
if(ipv4.valid) {
if(ipv4.dstAddr != 0x0a0a0a0a) {
if(ipv4.dstAddr != 0x0a0a0a0a) {
stdmeta.egress_spec := 0x1
stdmeta.egress_spec := 0x1
} else {
} else {
stdmeta.egress_spec := 0x1ff
stdmeta.egress_spec := 0x1ff
}
}
}

```
}
```

$\left(x:\left\{y: \operatorname{ipv}^{*} \mid y\right.\right.$. stdmeta.valid $\left.\}\right) \rightarrow\{y: \top \mid y$.stdmeta.egress_spec $\neq \mathrm{o} x \mathrm{o}\}$
Figure 8.4: Determined forwarding example. Top: property violated; middle: property holds; bottom: П4 type encoding the determined forwarding specification.
that successfully typechecks with the given type is shown in the middle of Figure 8.3.

### 8.2.2 Determined Forwarding

Our type checker is able to detect violations in programs that do not make an explicit forwarding decision. For example, for the type shown at the bottom of Figure 8.4, the program shown in the top of Figure 8.4 violates this property, while the property holds for the program shown in the middle of the same figure. The type states that starting
in a heap where it is guaranteed that at least IPv4 and the standard metadata are valid, after executing the program, we end in a state, where for all program paths the field stdmeta.egress_spec does not contain the initial value anymore, which indicates that a forwarding decision was made.

### 8.2.3 Parser-Deparser Compatibility

We now show how $\Pi_{4}$ can guarantee that deparsed packets can be correctly re-parsed without losing or corrupting information contained in packet headers. For example, assuming that the parser shown in Figure 8.5 successfully parses the Ethernet and IPv4 headers from the input packet, but not a VLAN header, from the code we can conclude that EtherType must be oxo8oo. Further, assuming that the programmer intends the ingress control shown in the middle right of Figure 8.5, after parsing, the switch checks if a VLAN header is present. If a VLAN header was already parsed from the input packet, no changes are made. Otherwise, a VLAN header is added (line 3) and the EtherType of the Ethernet header is updated accordingly. If an IPv4 header is present, the EtherType field of the VLAN header must also be updated (line 6) to obtain a protocol-conformant packet.

On the other hand, if the programmer forgot the statement on Line 4, i.e., they didn't update the ethernet. etherType field, serializing and deserializing the parsed headers will produce a corrupted packet. This unsafe example is shown on the middle left of Figure 8.5. After running the deparser at the end of ingress, all three headers are serialized: the first 112 bits correspond to the Ethernet header, followed by 32 bits of the VLAN header, and another 160 bits of the IPv4 header. However, since the programmer forgot to update the EtherType, bits 96 to 112 still contain the value oxo8oo. Hence, if the parser is run with this bit stream as the input, it will first parse the Ethernet header, then look at the etherType and given the value oxo8oo, it will continue to parse the $\mathrm{IPv}_{4}$ header. As a result, the bits of the VLAN header are parsed as an $\operatorname{IPv} 4$ header, leading to a corrupted packet.

To avoid such errors, we want to enforce the invariant that all instances valid at the end of ingress are equivalent to those obtained after deparsing and re-parsing. The code of the full pipeline is shown in the bottommost code snippet in Figure 8.5. The ascribed input type starting in Line 3 captures the assumptions about the state after executing the parser followed by the ingress control, i.e., both Ethernet and VLAN are guaranteed to be valid and the validity of IPv4 is indicated by the EtherType field of the VLAN header. The ascribed output type (Line 7) specifies the actual invariant, namely that executing the deparser followed by a reset statement and the parser produces a heap that is equivalent to the heap obtained after executing the parser followed by the ingress. We instruct our type checker to verify the property by checking the program with the type shown at the bottom of Figure 8.5, which ensures that there are enough bits to parse all possible headers.

### 8.2.4 Mutual Exclusion of Headers

Mutual exclusion is another property of interest that arises from the fact that the contents of specific header fields indicate which protocol header follows next and that each packet at runtime usually specifies only one such a protocol. In an implementation, we can take advantage of this property and use the same memory to store mutually exclusive headers. An example is the parser shown in the top of Figure 8.6 that conditionally parses either $\mathrm{IPv}^{2}$ or IPv6. Because only one of the paths is taken at runtime, it should

```
Parser \triangleq
    extract(ether);
    if(ether.etherType == 0x8100) {
        extract(vlan);
        if(vlan.etherType == 0x0800) { extract(ipv4) }
    } else {
        if(ether.etherType == 0x0800) { extract(ipv4) }
    }
```

```
UnsafeIngress =
    if(!vlan.valid) {
        add(vlan);
        if(ipv4.valid) {
            vlan.etherType := 0
                x0800
        }
    }
```

```
SafeIngress
    if(!vlan.valid) {
        add(vlan);
        ether.etherType := 0
            x8100;
        if(ipv4.valid) {
            vlan.etherType := 0
                x0800
        }
    }
```

```
Deparser 气
    if(ether.valid) { remit(ether) };
    if(vlan.valid) { remit(vlan) };
    if(ipv4.valid) { remit(ipv4) }
```

    /* Ingress is either UnsafeIngress or SafeIngress */
    Parser; Ingress;
(Deparser; reset; Parser) as (x:\{z:ether~|
z.ether.etherType $==0 \times 8100$ \&\& z.vlan.valid $\& \&$
z.ipv4.valid <=> z.vlan.etherType $==0 x 0800$ \&\&
z.pkt_out.length $==0$ \&\& z.pkt_in.length > 0\}) ->
\{y:T|x === y\}

$$
\left(x:\left\{y: \epsilon| | y \cdot p k t_{\text {out }} \mid=0 \wedge y \cdot p k t_{\text {in }} . \text { length }>304\right\}\right) \rightarrow T
$$

Figure 8.5: Roundtripping example. Common parser followed by the unsafe ingress code (left) and safe ingress code (right), followed by the deparser. The last code snippet shows the full pipeline, and type at the bottom is used to check the pipeline.
never happen that both instances are valid at the same time. In this small example, it is easy to see that this invariant holds, but in larger programs it is difficult to track which header instances are valid on which execution paths. We can check that the property continues to hold in the ingress shown in Line 10 using the type in the annotation on line 11. This type ensures that starting from a heap in which Ethernet and optionally either IPv4 or IPv6 are valid, at the end of the ingress, it is still guaranteed that IPv4

```
(extract(ether);
if(ether.etherType == 0x86dd) {
    extract(ipv6)
} else {
    if(ether.etherType == 0x0800) {
        extract(ipv4)
    }
}) as (x:{y:\epsilon|y.pkt_in.length > 432}) -> {y:ether~|!(y.ipv4.
        valid && y.ipv6.valid)};
Ingress /* Can be SafeIngress or UnsafeIngress*/
    as (x:{y:ether~|!(y.ipv4.valid && y.ipv6.valid)}) ->
            {y:ether~|!(y.ipv4.valid && y.ipv6.valid)};
if(ether.valid) {
    remit(ether)
};
if(ipv4.valid) {
    remit(ipv4)
};
if(ipv6.valid) {
    remit(ipv6)
}
```

```
UnsafeIngress 气
    add(ipv6);
    ether.etherType := 0x86dd
```

```
SafeIngress \triangleq
```

SafeIngress \triangleq
if(!ipv4.valid) {
if(!ipv4.valid) {
add(ipv6);
add(ipv6);
ether.etherType = 0
ether.etherType = 0
x86dd
x86dd
}

```
    }
```

$$
\left(x:\left\{y: \epsilon| | y \cdot p k t_{i n} \mid>432\right\}\right) \rightarrow \top
$$

Figure 8.6: Mutual exclusion example: IPv4 and IPv6 should never be simultaneously valid. Top: common pipeline; middle left: unsafe ingress code; middle right: safe ingress code; bottom: whole program type.
and IPv6 are not valid at the same time. The ingress code shown in the middle left of Figure 8.6 exemplifies a violation of the property. If a packet enters the control block with a valid $\mathrm{IPv}_{4}$ header, it will leave with both a valid IPv4 and a valid IPv6 header; a violation of our property. The code in the middle right is safe because it includes a conditional that explicitly checks the validity of IPv4 before adding IPv6.

Summary $\Pi_{4}$ 's types are expressive enough to express a variety of properties beyond header validity, from protocol conformance to basic safety properties to complex proper-


Figure 8.7: Hybrid switch architecture: fixed-function data plane in which the customer can execute custom code at pre-defined program points.
ties such as mutual exclusion of headers or the compatibility of the parser and deparser. This enables us to close the gap between previous type system-based approaches and existing verification tools for $\mathrm{P}_{4}\left(R Q_{1}\right)$.

### 8.3 Designing for Modularity

One key advantage of $\mathrm{P}_{4}$ programmable devices over traditional network devices is that the functionality of the data plane can be tailored to the application needs. However, at the same time, this advantage presents new challenges for network administrators, for example, the entire functionality of such devices usually must be implemented from scratch, without the possibility to rely on proven implementations provided by vendors. Therefore, an emerging design pattern for data plane switches is partial programmability, such as Cisco's daPIPE [Bal19], which is designed for the Nexus 3400 switch [Cis21]. The idea is that a device vendor provides a partially-implemented pipeline together with a set of program points where the customer can inject custom code as shown in Figure 8.7. Since the $\mathrm{P}_{4}$ language was not designed for incremental programming, where functionality is added to an existing data plane program, several challenges arise. Most importantly, the base data plane program provided by the vendor should not be affected by the injected code. Because of trade secrets or too much effort required to understand the base program in detail, it is usually not an option that the vendor shares the implementation of the base program such that the programmer can understand what his program is allowed to do. A better approach is that vendors require that customer programs satisfy certain properties, but in current architectures, these properties are not automatically checked.

For example, consider a deployment of the customizable pipeline in a campus network, where network engineers want to experiment with in-band network telemetry (INT) without perturbing the VLAN tag, which is used to enforce security policies. Let us assume that there are four classes of traffic, Visitor, Student, Faculty, and Staff, each with unique VLAN identifiers. We want to ensure that no matter how the customer-programmable part of the pipeline is instantiated, it cannot cause students and visitors to acquire the privileges of faculty or staff.

With $\Pi_{4}$, we can design a modular system that checks invariants on customer programs to be integrated into vendor pipelines statically. Practically, we can ensure that the VLAN tag is not changed, by checking that the customer's code has a type like: $(x: \tau) \rightarrow\left\{y: \tau^{\prime} \mid x\right.$.vlan.vid $\left.=y . v l a n . v i d\right\}$, where $\tau$ and $\tau^{\prime}$ are appropriate for the specific pipeline. We check, once-and-for-all, that the surrounding switch code composes with this type, and incrementally check that the customer code has this type (for an appropriate $\tau$ ).

```
extract(vlan);
Ingress /* Customer specified code: Default, Overwrite,
    Table, or UnsafeActions */
    as (x:\Sigmay:stdmeta.vlan) -> {y:\Sigmaz:stdmeta~.vlan~|y.vlan ==
        x.vlan};
remit(vlan)
```

Figure 8.8: Instantiation of modular router design; the parser and deparser are provided by the vendor, the ingress is provided by the customer.

### 8.3.1 Specifying Invariants

In the following we will consider a very simple implementation of the example described above, show in Figure 8.8. We assume that the header instance vlan is the standard VLAN header ( 32 bits), including a 12 -bit vlan tag field vlan. vid. We further assume that instance stdmeta, which is initially valid, provides access to the standard metadata ( 325 bits) used in the $\mathrm{P}_{4}$ switch model, including a 9 -bit egress specification field meta.egress_spec.

The control flow simply extracts the VLAN instance, executes the modular ingress control, and then emits the VLAN header. The overall behavior of the program is captured by type $\left(x:\left\{y:\right.\right.$ stdmeta $\left.\left.| | y . p k t_{i n} \mid>32\right\}\right) \rightarrow \Sigma y:$ stdmeta $_{\sim} . v$ lan $_{\sim}$. This type expresses that before executing the program stdmeta is valid and the input packet provides at least 32 bits, i.e., enough bits to extract the VLAN header. After the execution at least the metadata and VLAN headers are valid, but possibly also other headers added in the ingress. However, this type is too coarse-grained to guarantee that the ingress code will not change the parsed VLAN header.

Since both parser and deparser are provided by the vendor and stay unchanged, it is also not necessary to re-check that the whole pipeline is well-typed every time the customer implementation changes. Instead, we check once that the parser (Line 1 of Figure 8.8 ) is compatible with type $\left(x:\left\{y:\right.\right.$ stdmeta $\left.\left.| | y \cdot p k t_{i n} \mid>32\right\}\right) \rightarrow \Sigma y$ : stdmeta.vlan and that the deparser (Line 4 of Figure 8.8) is compatible with type $(x: \Sigma y:$ stdmeta.vlan $) \rightarrow \Sigma y:$ stdmeta $_{\approx}$. vlan $_{\sim}$. Of course, we could also provide a more specific output type for the deparser, for example, asserting that the content of the last 32 bits of the output packet is equal to the content of the VLAN instance.

When we swap in different implementations for the ingress control, we only need to check the ascribed type on Line 3 of Figure 8.8, without rechecking the surrounding code. This type defines the requirements the implementation of the ingress control must fulfill in order to be compatible with the rest of the pipeline, i.e, possibly new header instances can be added but the VLAN header must remain unchanged. With the infrastructure $\Pi_{4}$ 's type system provides, network engineers can make changes to their experimental module Ingress and check its compatibility with the switch without re-checking the feasibility of the whole switch in a modular fashion.

### 8.3.2 Checking Customer Programs

We now consider a collection of customer programs that an engineer may want to install into the switch and how $\Pi_{4}$ prevents security vulnerabilities by ensuring the customer code has the type annotated on Line 3 of Figure 8.8.

```
Default 气
    skip
```

```
Overwrite ^
```

Overwrite ^
vlan.vid := Faculty
vlan.vid := Faculty
Table 气
Table 气
add(_vlan_table);
add(_vlan_table);
if(_vlan_table.vid_key == vlan.vid) {
if(_vlan_table.vid_key == vlan.vid) {
if(_vlan_table.act == 0b0) {
if(_vlan_table.act == 0b0) {
stdmeta.egress_spec := 0x1ff
stdmeta.egress_spec := 0x1ff
} else {
} else {
stdmeta.egress_spec := 0x1
stdmeta.egress_spec := 0x1
}
}
}

```
    }
```

```
UnsafeActions 气
    add(_vlan_table);
    if(_vlan_table.vid_key == vlan.vid) {
        if(_vlan_table.act == 0b0) {
            vlan.vid := Faculty
        } else {
            vlan.vid := Staff
        }
    } else {
        vlan.vid := Visitor
    }
```

Figure 8．9：A collection of safe and unsafe customer implementations for the Ingress module from Figure 8．8．Top Left：Default；Top Right：Overwrite；Bottom Left：Table； Bottom Right：Unsafe Actions

Default Consider the empty program，shown in the top of Figure 8．9，which would surely be the default behavior when the programmer has not written any code yet．To typecheck this no－op module，we check that command skip has the annotated type：

$$
\cdot \vdash \operatorname{skip}:(x: \Sigma y: \text { stdmeta.vlan }) \rightarrow\left\{y: \Sigma z: \text { stdmeta }_{\approx} . \mathrm{vlan}_{\approx} \mid y . \mathrm{vlan}=x . \mathrm{vlan}\right\}
$$

The command skip typechecks with this type by rules T－SKip and T－Sub，since

$$
\begin{aligned}
(x: \Sigma y: \text { stdmeta.vlan }) \vdash & \{y: \Sigma z: \text { stdmeta.vlan } \mid y \equiv x\}<: \\
& \left\{y: \Sigma z: \text { stdmeta }_{\sim} . \mathrm{vlan}_{\approx} \mid y . \mathrm{vlan}=x . v \operatorname{van}\right\}
\end{aligned}
$$

Overwrite Conversely，if the customer were to install an obviously incorrect program， such as the second one in Figure 8．9，which always overwrites the VLAN tag with the identifier reserved for faculty members，the type system complains that the following subtyping check fails：

```
control Ingress(...) {
    action drp() {
        stdmeta.egress_spec = 0x1FF;
    }
    action fwd() {
        stdmeta.egress_spec = 1;
    }
    table vlan {
        key = { vlan.vid : exact; }
        actions = { drp; fwd; }
    }
    apply {
        vlan.apply();
    }
}
```

Figure 8.10: P 4 table encoded by the program Table from Figure 8.9.

$$
\begin{aligned}
& (x: \Sigma y: \text { stdmeta.vlan }) \vdash\left\{y: \Sigma z: \text { meta }_{\approx} . v \operatorname{van}_{\approx} \mid y . v l a n . v i d=F a c u l t y\right\}<: \\
& \left\{y: \Sigma z: \text { stdmeta }_{\approx} . \mathrm{vlan}_{\approx} \mid y . \mathrm{vl} \mathrm{an}^{2} . \mathrm{vid}=x . \mathrm{vlan} . \mathrm{vid}\right\}
\end{aligned}
$$

For example, if the VLAN tag of the incoming packet (x.vlan.vid) is Student, the two types denote disjoint sets of heaps.

Table The third program (Table) in Figure 8.9 encodes the $\mathrm{P}_{4}$ table shown in Figure 8.10, which matches on the value of header field vlan.vid and selects one of two actions: (1) drp, which sets the outgoing port (egress_spec) to oxiFF, and (2) fwd, which sets it to oxı.

As discussed in Section 6.5, to encode this table, we create a new header_vlan_table with a 12 -bit field vid_key and a 1-bit field act. The field _vlan_table.vid_key represents the match key of the table, while the field _vlan_table. act encodes the different actions. Since in this example the table only provides two actions, a single bit is sufficient to encode both alternatives. Actions drp and fwd are represented by _vlan_table.act = o and _vlan_table.act = 1 respectively. This program will typecheck since no branch of the code modifies the vlan.vid field, and _vlan_table is permitted to be valid.

Unsafe actions Finally, let us consider the last program in Figure 8.9, which encodes the $\mathrm{P}_{4}$ table shown in Figure 8.11. This table provides three actions fac, stf and vst, each modifying the VLAN tag. The $\mathrm{P}_{4}$ annotation @defaultonly indicates that action vst can only be used as default action and never in the table.

Whenever the table is applied, the VLAN tag is overwritten, however, without respecting the VLAN tag of the incoming packet. This clearly violates the requirement

```
control Ingress(...) {
    action fac() {
        vlan.vid = Faculty;
    }
    action stf() {
        vlan.vid = Staff;
    }
    action vst() {
        vlan.vid = Visitor;
    }
    table vlan {
        key = { vlan.vid : exact; }
        actions = {
            fac; stf;
            vst @defaultonly;
        };
        default_action = vst;
    }
    apply {
        vlan.apply();
    }
}
```

Figure 8.11: $\mathrm{P}_{4}$ table encoded by the program UnsafeActions from Figure 8.9.
that the VLAN tag must be unchanged after executing the ingress program, triggering a violation of the subset check just as in the Overwrite example.

Summary This case study has shown how $\Pi_{4}$ can be used to verify data plane programs in a modular way ( $R Q_{2}$ ). Type annotations allow expressing requirements that other modules must fulfill. If the implementation of modules changes, it is sufficient to recheck these modules in isolation according to their external requirements without having to re-check the entire pipeline.

### 8.4 Chapter Summary

In this chapter, we showed that $\Pi_{4}$ is indeed expressive enough to verify a wide range of real-world network properties. Thus, $\Pi_{4}$ is able to bridge the gap between type system-based approaches such as $\mathrm{SAFEP}_{4}$ and full-fledged verification tools. At the same time, our approach allows programs to be verified modularly, which will likely benefit efforts to modularize $\mathrm{P}_{4}$ code in the future. There are still properties that $\Pi_{4}$ does not yet support. As we have shown, these include, properties like array bound checks and arithmetic overflows, however we do not think this is a fundamental limitation since eliminating array bound checks is one of the standard examples discussed in the
dependent typing literature [ XP 98 ] and arithmetic overflows can be handled similarly. On the other hand, the verification of language features such as hash functions, externs, or registers poses a bigger challenge and is left for future work.

## CHAPTER

## Performance Evaluation

After exemplifying in the previous chapter that $\Pi_{4}$ 's types are suitable for expressing and checking a variety of practically relevant properties, in this chapter we evaluate the runtime performance of $\Pi_{4}$ using several real $\mathrm{P}_{4}$ programs. We thereby address the following research questions.

RQ1 How long does it take to verify that a certain property holds?
RQ2 What impact do the optimizations described in Section 6.4 have?
RQ3 What impact does the MTU and as such the size of bit vectors used in the encoding have?

RQ4 What is the impact of modular verification?
For our evaluation we used a collection of open-source programs and programs written by ourselves, which are summarized in Table 9.1. We used our $\mathrm{P}_{4}$ frontend to automatically translate $\mathrm{P}_{4}$ programs into $\Pi_{4}$ 's syntax and then used the type checker to check annotated types. We were therefore limited to programs that can be translated into the syntax of $\Pi_{4}$, for example, we exclude programs using persistent state (e.g. registers or counters). In some cases we have adapted the programs so that they could be translated, or we only translated parts automatically and parts by hand, while trying to preserve their semantics. For example, the program $n g s d n$ is a slightly adapted version of the next generation SDN platform ${ }^{1}$ tutorial, where we manually adjusted the control flow and removed the segment routing over IPv6 feature because the parser could not be fully translated. Similarly, we adapted the parser for program fabric-provided by the Open Network Operating System (ONOS) ${ }^{2}$-since it contains an infinite loop. In practice, the program relies on packets containing a specific header value that terminates the loop. However, this makes it impossible for us to unroll the parser, which is necessary for our translation to $\Pi_{4}$. In addition, we only included the filtering and forwarding features and removed all remaining. To provide a better intuition about the different programs, Table 9.1 provides for each program (1) how many lines of P 4 code it consists

[^7]| Program | LoC P4 | Parser <br> states | Tables | Total <br> header <br> length <br> (bytes) | $\Pi_{4}$ <br> Commands |
| :---: | :---: | :---: | :---: | :---: | :---: |
| multicast | 116 | 2 | 1 | 38 | 23 |
| basic | 120 | 3 | 1 | 34 | 33 |
| ecn | 134 | 3 | 1 | 38 | 39 |
| qos | 157 | 3 | 1 | 38 | 41 |
| vlan_decap | 105 | 3 | 1 | 38 | 45 |
| roundtrip | 93 | 3 | 0 | 38 | 48 |
| load_balance | 170 | 4 | 3 | 38 | 49 |
| basic_tunnel | 156 | 4 | 2 | 38 | 55 |
| ngsdn | 421 | 10 | 6 | 88 | 239 |
| fabric | 537 | 15 | 6 | 86 | 541 |

Table 9.1: P 4 programs used for the performance evaluation ordered by the number of $\Pi_{4}$ commands after translation.
of, (2) how many states the parser comprises, (3) how many tables the program uses, (4) what is the longest sequence of packet headers that is parsed (in bytes), and (5) how many $\Pi_{4}$ commands the translated program consists of.

### 9.1 Checking Header Validity

To address our first research question, as suggested by Liu et al. [Liu+18], we check all programs for header validity, as this property requires reasoning about all control-flow paths. Since header validity is a fundamental safety guarantee of $\Pi_{4}$ 's type system, it is sufficient to check that the programs produce any valid output heap. We can capture this with the following type.

$$
\left(x:\left\{y: \text { meta } \cdot \text { standard_metadata }| | y \cdot p k t_{i n}|>\operatorname{LEN} \wedge| y \cdot p k t_{\text {out }}=0 \mid\right\}\right) \rightarrow T
$$

This type states that, given an initial heap where both user-defined metadata and the intrinsic metadata provided by the target are valid ${ }^{3}$, and enough bits are available in the input packet to extract the longest sequence of headers, and additionally assuming that the output packet is empty, the output heap is in the set of all possible heaps. For example, for program fabric, we use LEN $=(86 * 8)-1$ to assert that the input packet must at least provide 86 bytes $=688$ bits.

We conducted all experiments on a workstation equipped with an Intel Core i76700 K CPU and 32 GiB of RAM. Since we quickly ran out of memory while checking program $n g s d n$, we repeated the experiment on a server equipped with an AMD EPYC 7542 CPU and 512 GiB of RAM. In all other experiments, we obtained very similar results on both machines, and surprisingly, our workstation was even able to check programs faster in several cases.

For our first experiment, we used an MTU of 1500 bits, which is only one-eighth of the standard MTU, but is sufficient for our programs because the length of the input

[^8]

Figure 9.1: Time needed to check header validity for the programs listed in Table 9.1, with substitution inlining and validity caching enabled (default), without validity caching enabled, without substitution inlining enabled and with both optimizations disabled.
packet plus the length of the output packet—which is twice the total header length reported in Table 9.1-does not exceed this value. The results are given by the blue bars (default) in Figure 9.1. Since there are huge spans between the measured values, we use a logarithmic scale for the $y$-axis. However, since this distorts the actual size differences in the plot, we additionally give the absolute value above each bar. Most notably are the results for programs $n g s d n$ and fabric. Intuitively, both programs take the longest to verify, as they are also the largest programs. Surprisingly, it takes almost 6 hours ( $5: 51$ ) to check $n g s d n$, while fabric with more than twice as many commands takes only about 50 seconds. The reason for this is that our type checker found a header validity bug in the implementation of fabric, and thus can terminate early.

Figure 9.2 shows a stripped-down excerpt from fabric.p4 showing the offensive code. The ingress control first applies the filtering control and afterwards applies the forwarding control. The first control applies table fwd_classifier-shown in the top of Figure 9.2-which matches on the Ethernet type field and sets the forwarding type (fwd_type) accordingly. If the forwarding type is equal to the constant named FWD_IPV6_UNICAST, the forwarding control applies table routing_v6, which reads the IPv6 header. However, this program does not guarantee that the IPv6 header is actually valid in this case. After repairing the bug by guarding the invocation of table routing_ipv6 with an additional validity check, we ran our checker again on the program, but-even on our server-eventually ran out of memory before the check was completed.

### 9.2 Effects of Optimizations on Runtime

In the next step, we evaluated whether the optimizations described in Section 6.4 actually have the desired effect, thus addressing $R Q_{2}$. As a reminder, as a first optimization we reduce the size of the computed types by eliminating explicit type substitutions and the

```
action set_forwarding_type(fwd_type_t fwd_type) {
    meta.fwd_type = fwd_type;
}
table fwd_classifier {
    key = { hdr.eth_type.value: ternary; }
    actions = { set_forwarding_type; }
}
```

```
table routing_v6 {
    key = { hdr.ipv6.dst_addr: lpm; }
}
apply {
    if (meta.fwd_type == FWD_IPV6_UNICAST) {
        routing_v6.apply();
    }
}
```

Figure 9.2: Header validity bug in fabric.p4. Excerpt from control Filtering (top) and control Forwarding (bottom). The validity of ipv6 in control Forwarding depends on correct control-plane entries in table fwd_classifier.
second optimization reduces the number of SMT solver calls by caching assumptions about header validity. The results can also be seen in Figure 9.1. Except for program $n g s d n$, we ran all programs respectively without the header validity cache but with substitution inlining, without substitution inlining but with the header validity cache, and without any optimization. Program $n g s d n$ could not be checked at all without the optimizations.

As can be seen with program multicast, substitution inlining may introduce additional overhead for programs that can be checked in a short time. For smaller programs, the validity cache has a bigger impact on the runtime than substitution inlining, however, this reverses for lager programs as can be seen with program fabric.

### 9.3 Effects of the MTU on Runtime

Next, we evaluate the impact of MTU on the runtime of the type checker ( $R Q_{3}$ ). Again, we look at the time it takes our type checker to check header validity. We start with an MTU of 1500 bits and then increase the MTU in steps of 1500 bits, up to a value of 12000 bits $=1500$ bytes. As can be seen in Figure 9.3, the runtime of the type checker increases with increasing MTU. For larger programs, the runtime increases significantly more with larger MTU values than for comparatively smaller programs. The program roundtrip is noticeable, which proportionally shows the largest increase in runtime. The reason is that the program uses the reset command, which prohibits the calculated types being completely inlined (cf. Section 6.4.1), which has a significant effect on the runtime of the SMT solver.

To better quantify the respective increase in runtime, Figure 9.4 additionally shows the percentage runtime increase compared to the next smaller MTU, i.e., how much does the runtime increase when we increase the MTU from 1500 bits to 3000 bits, or


Figure 9.3: Overall time needed to check header validity with varying MTUs.


Figure 9.4: Percentage increase in runtime compared to the previous MTU value, starting from an MTU of 1500 bits. For example, checking program roundtrip with an MTU of 3000 bits increases the type checking time by roughly $130 \%$ compared to an MTU of 1500 bits.
from 3000 bits to 4500 bits, and so on. Overall, the runtime increases with a larger MTU value, but no direct conclusions can be drawn from the selected MTU to the increase in runtime. Not every increase in MTU results in the same percentage runtime increase. Also, the runtime does not necessarily increase to a greater extent with larger MTU values. In a few cases, a larger MTU value even leads to a reduced runtime compared to the next smaller MTU value.

| Program | Parser | Ingress | Egress | Deparser | Complete |
| :---: | :---: | :---: | :---: | :---: | :---: |
| ngsdn | 2.99 | 19.85 | 0.31 | 2.01 | 201.93 |
| fabric | 15.24 | 21.81 | - | 743.50 | 355.94 |

Table 9.2: Time in seconds needed to modularly verify header validity for programs $n g s d n$ and fabric. Parser, Ingress, Egress and Deparser are the required times to check the respective pipeline stages in isolation. Complete indicates the time for the entire pipeline, where we ascribed each pipeline stage with a type that captures the requirements on the input heap and the guarantees for the resulting heap.

### 9.4 Modular Verification

A key aspect of $\Pi_{4}$ is its support for modular verification. Therefore, we will now evaluate to which extent we can exploit modularity to make verification of large programs in particular achievable (RQ4). For this purpose we use program $n g s d n$ and the fixed version of fabric without any validity bugs. Table 9.2 shows the results for checking each pipeline stage in isolation and for the whole pipeline. However, when checking the complete pipeline, we ascribe a type to every pipeline stage, which exactly captures the requirements it has on the input heap and the guarantees it offers with respect to the output heap.

For example, we use the following type to describe the parser of program ngsdn.

$$
\begin{aligned}
& \left(x:\left\{y: \text { standard_metadata } \cdot \text { meta }| | y \cdot p k t_{i n} \mid>703\right\}\right) \rightarrow \\
& \left\{y: \top \mid y . s t a n d a r d \_m e t a d a t a . v a l i d \wedge\right. \\
& y \text {.meta.valid } \wedge \\
& y \text {.ethernet.valid } \wedge \\
& \neg y . c p u \_i n . v a l i d \wedge \\
& \text { y.icmpv6.valid } \Longrightarrow \text { ( } y \text {.ipv6.valid } \wedge \\
& \text { ( } y \text {.icmpv6.type }=0 x 87 \vee y . \text { icmpv6.type }=0 \times 88) \Longrightarrow \\
& y \text {.ndp.valid) } \wedge \\
& \left.\left|y \cdot p k t_{i n}\right| \geq 0\right\}
\end{aligned}
$$

This type states that given an input packet that provides enough bits to extract all consecutive packet headers, the parser guarantees that afterwards the ethernet header will be parsed, but not the cpu_in header. Furthermore, if the icmpv6 header is parsed, it is guaranteed that the ipv6 header was also parsed and if the type field of header icmpv6 contains the value ox87 or ox88 the ndp header will be additionally parsed. This type captures exactly the requirements of the ingress control to also successfully typecheck. We can proceed similarly for the egress and the deparser. For example the egress has the requirement that the header cpu_in must not be valid at the beginning, which is guaranteed by both the parser and the ingress. Overall, this approach allowed us to check the complete program ngsdn in about 3.5 minutes, which is a significant improvement over the 6 hours previously measured.

By proceeding analogously, we can also significantly reduce the time required to check program fabric. Note, the table does not provide a value for the egress because we have initially limited this program to a portion of the ingress control. Since the functionality of the ingress is described by two separate control blocks, we could also
check them in isolation. For example, using this approach, it took us only about 5 seconds to check the filtering part of the ingress control. It is surprising that it takes more than twice as long to check the deparser in isolation compared to the full pipeline. However, a detailed root cause analysis is left for future work.

### 9.5 Chapter Summary

In this chapter, we showed that $\Pi_{4}$ 's SMT solver-based approach is suitable for checking real $\mathrm{P}_{4}$ applications. However, especially with regard to larger programs, it is essential to rely on $\Pi_{4}$ 's capability to modularly check programs, but still our approach is currently not able to compete with tools like p4v [Liu+18] or Aquila [Tia+21]. It is crucial for an implementation of our SMT-based approach that the complexity of the SMT queries as well as the number of SMT solver invocations is reduced as much as possible, for example, by applying the optimizations we have proposed. Furthermore, a crucial factor with regard to the complexity of the SMT queries is the MTU used, which determines the maximum size of the encoded bit vectors. An interesting question that has remained open is whether it is possible to automatically compute the minimum required MTU for each program and annotated type while preserving the properties of our type system, thus avoiding unnecessary overhead.

## Part IV

## Epilogue

## Conclusion and Future Work

This dissertation confirms our initial hypothesis that type systems are well suited to equip data plane programming languages-in particular the $\mathrm{P}_{4}$ language-with safety guarantees, which makes it possible to verify a rich set of safety properties.

We started with the basic property of header validity where we were able to show that several of P4's language features are susceptible to bugs due to accessing invalid header instances. Based on these findings, we then designed SAFEP4, a domain-specific language for programmable data planes whose static type system guarantees that all headers accessed are guaranteed to be valid. We were faced with the challenge that due to the interaction with the control plane, header validity becomes a dynamic property, which we addressed both by employing path-sensitive typing that incorporates information from forwarding table declarations and by assuming that the control plane satisfies three basic safety properties. Our evaluation showed that all the error categories we identified occur in real-world programs of varying sizes and that SAFEP4's type system can be utilized to detect them without having to annotate the source code beforehand, which is a major difference to existing verification tools. In the second step we addressed the issue of $\mathrm{SAFEP}_{4}$ being limited to checking header validity.

With $\Pi_{4}$, we managed to close the gap between simple approaches such as SAFEP4 and full-fledged verification tools with respect to the expressive power by resorting to the more powerful typing discipline of dependent types. Since type systems are a compositional way to establish program properties, our approach offers for the first time the possibility to verify rich correctness properties for data plane programs in a modular way. At the same time, we managed to automate subtype checks by encoding them into SMT queries in the theory of fixed-width bit vectors, relieving the programmer from writing manual proofs, which is common for dependently-typed systems.

With this dissertation we have laid the foundation for a wider adoption of type systems in the field of network programming-in particular for programmable data planes. This results in a variety of interesting directions in which our work can be expanded.

Integration into $\mathrm{P}_{4}$ The first direction in which our work can be extended is making our dependent type system practically usable. We have equipped our prototype implementation with a frontend based on Petr 4 that allows programs written in the

P4 language to be parsed so that they can be annotated with types and type-checked. However, our frontend is not complete and requires two major enhancements.

The first is with respect to the support of more advanced language features. There are a few $\mathrm{P}_{4}$ features that our current prototype does not support, mostly because they pose challenges to SMT-based approaches to verification. The unpredictability of hash functions is difficult to verify. Besides over-approximating their behavior by representing them as uninterpreted functions, we can resort to a more fine-grained approach such as concolic verification [GKSo5]. Registers are on-switch state that can be modified by the packet or the controller and persists between packets. Representing persistent state in our current semantics is tricky, since it involves distributed computing concerns and needs further investigation if one does not simply want to over-approximate their behavior, e.g., by making havoc of the values every time the register is read. More general, to enable typing of externs, the work on typing foreign function interfaces [FFo8] might serve a starting point.

The second improvement concerns the integration of our type system into the $\mathrm{P}_{4}$ language. To guide the design of impactful systems for modular verification of data plans, it is necessary to provide a gradual transition from untyped to fully typed code. Gradual Typing [STo6; TFo6] would allow $\mathrm{P}_{4}$ programmers to statically type parts of their code as needed. Approaches such as migrational typing [Cam+17] or the use of type inference [SVo8] could support a gradual migration. How type inference, such as the one used by Liquid Haskell [VTV18], can be realized for $\Pi_{4}$ is also future work.

In addition, the type-checking performance must be further optimized. Even though we have already implemented various optimizations, the overall time needed to check even simple programs is not sufficient for quick feedback during the development process. For this purpose, it can be examined whether special SMT tactics or an alternative encoding allow a more efficient solving of our subtyping constraints. An alternative could be an embedding of our type checker into a dependently-typed language that does not need to rely on an SMT solver at all.

Extend guarantees offered by types A second direction in which we can extend our work is to extend our type system to also cover the interaction between the data plane and the control plane. Even though the functionality of the data plane is not specified exclusively by the $\mathrm{P}_{4}$ program and depends to a large extent on the control plane, there are no mechanisms to ensure that no faulty rules-e.g., causing inconsistencies between the control plane and data plane [Shu+20]-are installed by the control plane.

It would therefore be interesting to investigate whether a specification of the data plane using the dependent types we have developed can be used to restrict the set of rules that can be installed by the control plane to conform to the network policy being implemented. For example, if the parser never extracts a particular packet header, this could result in the control plane program being forced to install a rule that ensures that appropriate packets are processed by the controller.

This would, however, require that both the data plane and the control plane are programmed together. Approaches from the field of tierless programming [Nel+14; RV18] might be suitable starting points. From the common description, a compiler could then generate platform-specific code, both for the data plane and the controller that implements the necessary guarantees.

Exploring Applications of chomp Another direction in which our work can be extended is with respect to our chomp operator. Here, it would be interesting to explore
what applications arise in other domains for our verified approach to parsing using derivatives. The domain of verified serializers and deserializers like EverParse [Ram+19] and Narcissus [Del+19] could be interesting. So far these approaches do not statically capture what remains after parsing parts of a certain input.

## Bibliography

| [Ale+98] | D. S. Alexander et al. "The SwitchWare Active Network Architecture". In: IEEE Network 12.3 (May 1998), pp. 29-36. Doi: 10. 1109/65. 690959. |
| :---: | :---: |
| [All+o3] | J. R. Allen et al. "IBM PowerNP Network Processor: Hardware, Software, and Applications". In: IBM Journal of Research and Development 47.2-3 (2003), pp. 177-193. DOI: 10.1147/rd. 472.0177. |
| [And+14] | Carolyn Jane Anderson et al. "NetKAT: Semantic Foundations for Networks". In: Proceedings of the 41 st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages. POPL '14. San Diego, California, USA: Association for Computing Machinery, 2014, pp. 113-126. Doi: 10.1145/2535838. 2535862. |
| [Ara+16] | Mina Tahmasbi Arashloo et al. "SNAP: Stateful Network-Wide Abstractions for Packet Processing". In: Proceedings of the 2016 ACM SIGCOMM Conference. SIGCOMM '16. Florianopolis, Brazil: ACM, 2016, pp. 29-43. doi: 10.1145/2934872. 2934892. |
| [Bai+18] | Jiasong Bai et al. "Filtering Spoofed IP Traffic Using Switching ASICs". In: Proceedings of the ACM SIGCOMM 2018 Conference on Posters and Demos. ACM. 2018, pp. 51-53. |
| [Balı9] | M. Baldi. "daPIPE - A Data Plane Incremental Programming Environment". In: 2019 ACM/IEEE Symposium on Architectures for Networking and Communications Systems (ANCS). 2019, pp. 1-6. doI: 10.1109/ANCS . 2019.8901893. |
| [BCZ97] | Samrat Bhattacharjee, Kenneth L. Calvert, and Ellen W. Zegura. "An Architecture for Active Networking". In: High Performance Networking VII: IFIP TC6 Seventh International Conference on High Performance Networks (HPN '97), 28th April - 2nd May 1997, White Plains, New York, USA. Ed. by Ahmed Tantawy. Boston, MA: Springer US, 1997, pp. 265-279. DoI: 10.1007/978-0-387-35279-4_17. |
| [Bia+94] | Edoardo Biagioni et al. "Signatures for a Network Protocol Stack: A Systems Application of Standard Ml". In: Proceedings of the 1994 ACM Conference on LISP and Functional Programming. LFP '94. Orlando, Florida, USA: Association for Computing Machinery, 1994, pp. 55-64. Doi: 10. 1145/182409. 182431. |

[BMG99] Andrew Begel, Steven McCanne, and Susan L. Graham. "BPF+: Exploiting Global Data-flow Optimization in a Generalized Packet Filter Architecture". In: Proceedings of the Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication. SIGCOMM '99. Cambridge, Massachusetts, USA: ACM, 1999, pp. 123-134. DoI: 10.1145/ 316188. 316214.
[Bos +13 Pat Bosshart et al. "Forwarding metamorphosis: Fast programmable matchaction processing in hardware for SDN". In: ACM SIGCOMM Computer Communication Review 43.4 (2013), pp. 99-110.
[Bos+14] Pat Bosshart et al. "P4: Programming Protocol-independent Packet Processors". In: SIGCOMM Comput. Commun. Rev. 44.3 (July 2014), pp. 87-95. Doi: $10.1145 / 2656877.2656890$.
[Bos18] Patrick Bosshart. "Programmable Forwarding Planes at Terabit/s Speeds". In: 2018 IEEE Hot Chips 30 Symposium (HCS). IEEE. 2018.
[Bro19] Broadcom Inc. Network Programming Language. https : / /nplang . org/. Accessed: 11.03.2022. 2019.
[Brz64] Janusz A. Brzozowski. "Derivatives of Regular Expressions". In: Journal of the ACM 11.4 (Oct. 1964), pp. 481-494. DoI: 10. 1145/321239. 321249.
[Calo6] Ken Calvert. "Reflections on Network Architecture: An Active Networking Perspective". In: SIGCOMM Comput. Commun. Rev. 36.2 (2006), pp. 2730. DOI: 10.1145/1129582. 1129590.
[Cam+17] John Peter Campora et al. "Migrating Gradual Types". In: Proceedings of the ACM on Programming Languages 2.POPL (2017), p. 15.
[Cas+07] Martin Casado et al. "Ethane: Taking Control of the Enterprise". In: Proceedings of the 2007 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications. SIGCOMM 'o7. Kyoto, Japan: Association for Computing Machinery, 2007, pp. 1-12. Dor: 10.1145/1282380. 1282382.
[Cis21] Cisco. Cisco Nexus 30oo Series Switches. https://www.cisco.com/ c/en/us/products/switches/nexus-3000-series-switches/ index.html. Jan. 2021.
[Con+07] Jeremy Condit et al. "Dependent Types for Low-level Programming". In: Proceedings of the 16th European Symposium on Programming. ESOP'o7. Braga, Portugal: Springer-Verlag, 2007, pp. 520-535.
[Cor14] Jonathan Corbet. BPF: The Universal In-kernel Virtual Machine. Available at https://lwn.net/Articles/599755/, May 2014.
[Del+19] Benjamin Delaware et al. "Narcissus: Correct-by-construction Derivation of Decoders and Encoders from Binary Formats". In: Proc. ACM Program. Lang. 3.ICFP (July 2019). DoI: $10.1145 / 3341686$.
[DH17] S. Deering and R. Hinden. Internet Protocol, Version 6 (IPv6) Specification. Tech. rep. 8200. RFC Editor, July 2017. URL: https://www. rfc-editor. org/rfc/rfc8200.txt.
[DH98] S. Deering and R. Hinden. Internet Protocol, Version 6 (IPv6) Specification. RFC 2460. RFC Editor, Dec. 1998. URL: https://www.rfc-editor. org/rfc/rfc2460.txt.
[Doe+21] Ryan Doenges et al. "Petr4: Formal Foundations for P4 Data Planes". In: Proc. ACM Program. Lang. 5.POPL (Jan. 2021). DoI: 10.1145/3434322.
[Dum+2o] Dragos Dumitrescu et al. "Bf4: Towards Bug-free P4 Programs". In: Proceedings of the Annual Conference of the ACM Special Interest Group on Data Communication on the Applications, Technologies, Architectures, and Protocols for Computer Communication. SIGCOMM '2o. Virtual Event, USA: Association for Computing Machinery, 2020, pp. 571-585.
[Eic+19] Matthias Eichholz et al. "How to Avoid Making a Billion-Dollar Mistake: Type-Safe Data Plane Programming with SafeP4". In: 33rd European Conference on Object-Oriented Programming (ECOOP 2019). Ed. by Alastair F. Donaldson. Vol. 134. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2019, 12:1-12:28. DOI: 10.4230/LIPIcs.ECOOP . 2019.12.
[Eic+22] Matthias Eichholz et al. "Dependently-Typed Data Plane Programming". In: Proceedings of the ACM on Programming Languages 6.POPL (2022). Doi: 10.1145/3498701.
[ESMo4] Robert Ennals, Richard Sharp, and Alan Mycroft. "Linear Types for Packet Processing". In: Programming Languages and Systems. Ed. by David Schmidt. Berlin, Heidelberg: Springer Berlin Heidelberg, 2004, pp. 204-218.
[FFo8] Michael Furr and Jeffrey S. Foster. "Checking Type Safety of Foreign Function Calls". In: ACM Trans. Program. Lang. Syst. 30.4 (Aug. 2008). Doi: 10.1145/1377492.1377493.
[FGo5] Kathleen Fisher and Robert Gruber. "Pads: A Domain-specific Language for Processing Ad Hoc Data". In: Proceedings of the 2005 ACM SIGPLAN Conference on Programming Language Design and Implementation. PLDI '05. Chicago, IL, USA: Association for Computing Machinery, 2005, pp. 295-304. DoI: 10.1145/1065010. 1065046.
[Fre+18] Lucas Freire et al. "Uncovering Bugs in P4 Programs with Assertion-based Verification". In: Proceedings of the Symposium on SDN Research. SOSR '18. Los Angeles, CA, USA: Association for Computing Machinery, 2018. Doi: 10.1145/3185467. 3185499.
[FRZ13] Nick Feamster, Jennifer Rexford, and Ellen Zegura. "The Road to SDN: An Intellectual History of Programmable Networks". In: Queue 11.12 (2013), pp. 20-40. Doi: 10.1145/2559899. 2560327.
[Gao+2o] Jiaqi Gao et al. "Lyra: A Cross-Platform Language and Compiler for Data Plane Programming on Heterogeneous ASICs". In: Proceedings of the Annual Conference of the ACM Special Interest Group on Data Communication on the Applications, Technologies, Architectures, and Protocols for Computer Communication. SIGCOMM '2o. Virtual Event, USA: Association for Computing Machinery, 2020, pp. 435-450. DoI: 10. 1145/ 3387514.3405879.
[GGW15] Marco Gaboardi, Michael Greenberg, and David Walker. Type Systems for SDN Controllers. https://www.cs.princeton.edu/~dpw/papers/ typed-controllers-plvnet-2015.pdf. 2015.
[GKSo5] Patrice Godefroid, Nils Klarlund, and Koushik Sen. "DART: Directed Automated Random Testing". In: Conference on Programming Language Design and Implementation (PLDI). 2005, pp. 213-223.
[Gud+o8] Natasha Gude et al. "NOX: Towards an Operating System for Networks". In: SIGCOMM Comput. Commun. Rev. 38.3 (July 2008), pp. 105-110. DoI: 10.1145/1384609. 1384625.
[HDM14] Wei Huang, Yao Dong, and Ana Milanova. "Type-Based Taint Analysis for Java Web Applications". In: Proceedings of the 17th International Conference on Fundamental Approaches to Software Engineering - Volume 8411. New York, NY, USA: Springer-Verlag New York, Inc., 2014, pp. 140-154. DOI: 10.1007/978-3-642-54804-8_10.
[Hic+98] Michael Hicks et al. "PLAN: A Packet Language for Active Networks". In: Proceedings of the Third ACM SIGPLAN International Conference on Functional Programming. ICFP '98. Baltimore, Maryland, USA: Association for Computing Machinery, 1998, pp. 86-93. Doi: 10.1145/289423. 289431.
[Hin+o9] Timothy L. Hinrichs et al. "Practical Declarative Network Management". In: Proceedings of the 1st ACM Workshop on Research on Enterprise Networking. WREN 'o9. Barcelona, Spain: Association for Computing Machinery, 2009, pp. 1-10. DOI: 10. 1145/1592681. 1592683.
[Hoao9] Tony Hoare. Null References: The Billion Dollar Mistake. https://www. infoq. com/presentations/Null-References-The-Billion-Dollar-Mistake-Tony-Hoare/. Aug. 2009.
[Høi+18] Toke Høiland-Jørgensen et al. "The EXpress Data Path: Fast Programmable Packet Processing in the Operating System Kernel". In: Proceedings of the 14th International Conference on Emerging Networking EXperiments and Technologies. CoNEXT '18. Heraklion, Greece: Association for Computing Machinery, 2018, pp. 54-66. doI: 10.1145/3281411. 3281443.
[HOMo6] William G. J. Halfond, Alessandro Orso, and Panagiotis Manolios. "Using Positive Tainting and Syntax-aware Evaluation to Counter SQL Injection Attacks". In: Proceedings of the 14th ACM SIGSOFT International Symposium on Foundations of Software Engineering. SIGSOFT '06/FSE-14. Portland, Oregon, USA: ACM, 2006, pp. 175-185. DOI: 10.1145/1181775. 1181797.
[Iba+19] Stephen Ibanez et al. "The P4->NetFPGA Workflow for Line-Rate Packet Processing". In: Proceedings of the 2019 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays. FPGA '19. Seaside, CA, USA: Association for Computing Machinery, 2019, pp. 1-9. DoI: 10 . 1145 / 3289602.3293924.
[IPWo1] Atsushi Igarashi, Benjamin C. Pierce, and Philip Wadler. "Featherweight Java: A Minimal Core Calculus for Java and GJ". In: ACM Trans. Program. Lang. Syst. 23.3 (May 2001), pp. 396-450. Doi: 10.1145/503502. 503505.
[Jin+17] Xin Jin et al. "NetCache: Balancing Key-Value Stores with Fast In-Network Caching". In: Proceedings of the 26th Symposium on Operating Systems Principles. SOSP '17. Shanghai, China: Association for Computing Machinery, 2017, pp. 121-136. DOI: 10. 1145/3132747. 3132764.
[Jin+18] Xin Jin et al. "NetChain: Scale-Free Sub-RTT Coordination". In: 15th USENIX Symposium on Networked Systems Design and Implementation (NSDI 18). Renton, WA: USENIX Association, Apr. 2018, pp. 35-49.
[Jin18] Xin Jin. BMV2-based implementation of NetCache. https://github. com/netx-repo/netcache-p4. Mar. 2018.
[JMW1o] Trevor Jim, Yitzhak Mandelbaum, and David Walker. "Semantics and Algorithms for Data-dependent Grammars". In: Proceedings of the 37th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages. POPL '1o. Madrid, Spain: Association for Computing Machinery, 2010, pp. 417-430. doi: 10.1145/1706299. 1706347.
[KG16] Rahul Kumar and B. B. Gupta. "Stepping Stone Detection Techniques: Classification and State-of-the-art". In: Proceedings of the international conference on recent cognizance in wireless communication \& image processing. Springer. 2016, pp. 523-533.
[Kim+15] Hyojoon Kim et al. "Kinetic: Verifiable Dynamic Network Control". In: 12th USENIX Symposium on Networked Systems Design and Implementation (NSDI 15). Oakland, CA: USENIX Association, May 2015, pp. 5972.
[Kle+18] George T. Klees et al. "Evaluating Fuzz Testing". In: Proceedings of the ACM Conference on Computer and Communications Security (CCS). Oct. 2018.
[Kod15] Chaitanya Kodeboyina. An open-source P4 switch with SAI support. https : / / p4 . org / p4 / an - open - source - p4 - switch - with - sai support.html. June 2015
[KR18] Ali Kheradmand and Grigore Roşu. P4K: A Formal Semantics of $P_{4}$ and Applications. Tech. rep. https://arxiv.org/abs/1804.01468. University of Illinois at Urbana-Champaign, Apr. 2018.
[Lak+o4] T. V. Lakshman et al. "The SoftRouter Architecture". In: ACM HOTNETS. ACM, 2004. URL: https : / / www . microsoft . com / en -us/research/publication/the-softrouter-architecture/.
[Li+19] Guanyu Li et al. "NETHCF: Enabling Line-rate and Adaptive Spoofed IP Traffic Filtering". In: 2019 IEEE 27th International Conference on Network Protocols (ICNP). 2019, pp. 1-12. Doi: 10.1109/ICNP. 2019.8888057.
[Liu+18] Jed Liu et al. "P4V: Practical Verification for Programmable Data Planes". In: Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication. SIGCOMM '18. Budapest, Hungary: ACM, 2018, pp. 490-503. DOI: 10.1145/3230543. 3230582.
[MBo8] Leonardo de Moura and Nikolaj Bjørner. "Z3: An Efficient SMT Solver". In: Tools and Algorithms for the Construction and Analysis of Systems. Ed. by C. R. Ramakrishnan and Jakob Rehof. Berlin, Heidelberg: Springer Berlin Heidelberg, 2008, pp. 337-340.
[MCoo] Peter J. McCann and Satish Chandra. "Packet Types: Abstract Specification of Network Protocol Messages". In: Proceedings of the Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication. SIGCOMM 'oo. Stockholm, Sweden: Association for Computing Machinery, 2000, pp. 321-333.
[McC+16] Jedidiah McClurg et al. "Event-Driven Network Programming". In: Proceedings of the 37th ACM SIGPLAN Conference on Programming Language Design and Implementation. PLDI '16. Santa Barbara, CA, USA: Association for Computing Machinery, 2016, pp. 369-385. DoI: 10.1145/2908080. 2908097.
[McK+o8] Nick McKeown et al. "OpenFlow: Enabling Innovation in Campus Networks". In: SIGCOMM Comput. Commun. Rev. 38.2 (Mar. 2008), pp. 69-74. doi: 10.1145/1355734. 1355746.
[McK+16] Nick McKeown et al. Automatically Verifying Reachability and Well-formedness in P4 Networks. Tech. rep. Sept. 2016. Url: https : / / www . microsoft . com / en - us / research / publication / automatically-verifying-reachability-well-formedness-p4-networks/.
[Mily8] Robin Milner. "A Theory of Type Polymorphism in Programming". In: Journal of Computer and System Sciences 17.3 (Dec. 1978), pp. 348-375.
[Mon+13] Christopher Monsanto et al. "Composing Software Defined Networks". In: 1oth USENIX Symposium on Networked Systems Design and Implementation (NSDI 13). Lombard, IL: USENIX Association, Apr. 2013, pp. 1-13.
[Mut+1o] Chitra Muthukrishnan et al. "Using Strongly Typed Networking to Architect for Tussle". In: Proceedings of the 9th ACM SIGCOMM Workshop on Hot Topics in Networks. Hotnets-IX. Monterey, California: Association for Computing Machinery, 2010. DoI: 10.1145/1868447. 1868456.
[Nan+o8] Aleksandar Nanevski et al. "Ynot: Dependent Types for Imperative Programs". In: Proceedings of the 13 th ACM SIGPLAN International Conference on Functional Programming. ICFP 'o8. Victoria, BC, Canada: Association for Computing Machinery, 2008, pp. 229-240. DOI: 10. 1145/1411204. 1411237.
[Nan+14] Aleksandar Nanevski et al. "Communicating State Transition Systems for Fine-Grained Concurrent Resources". In: Programming Languages and Systems. Ed. by Zhong Shao. Berlin, Heidelberg: Springer Berlin Heidelberg, 2014, pp. 290-310.
[Nel+14] Tim Nelson et al. "Tierless Programming and Reasoning for SoftwareDefined Networks". In: Proceedings of the 11th USENIX Conference on Networked Systems Design and Implementation. NSDI'14. Seattle: USENIX Association, 2014, pp. 519-531.
[Net18] Barefoot Networks. Behavioral Model. Dec. 2018. UrL: https://github. com/p4lang/behavioral-model.
[Nev+18] Miguel Neves et al. "Verification of P4 Programs in Feasible Time Using Assertions". In: Proceedings of the 14th International Conference on Emerging Networking EXperiments and Technologies. CoNEXT '18. Heraklion, Greece: Association for Computing Machinery, 2018, pp. 73-85. doi: 10.1145/3281411. 3281421.
[NMB06] Aleksandar Nanevski, Greg Morrisett, and Lars Birkedal. "Polymorphism and Separation in Hoare Type Theory". In: Proceedings of the Eleventh ACM SIGPLAN International Conference on Functional Programming. ICFP 'o6. Portland, Oregon, USA: Association for Computing Machinery, 2006, pp. 62-73. DOI: 10.1145/1159803.1159812.
[Nöt+18] Andres Nötzli et al. "P4pktgen: Automated Test Case Generation for P4 Programs". In: Proceedings of the Symposium on SDN Research. SOSR '18. Los Angeles, CA, USA: Association for Computing Machinery, 2018. doi: 10.1145/3185467.3185497.
[Nou+19] Mohammad A. Noureddine et al. "P4AIG: Circuit-Level Verification of P4 Programs". In: 2019 49th Annual IEEE/IFIP International Conference on Dependable Systems and Networks - Supplemental Volume (DSN-S). 2019, pp. 21-22. DOI: 10.1109/DSN-S. 2019. 00016.
[OCo+18] T. J. OConnor et al. "Pivotwall: SDN-based Information Flow Control". In: Proceedings of the Symposium on SDN Research. ACM. 2018, p. 3.
[ORS92] S. Owre, J. M. Rushby, and N. Shankar. "PVS: A prototype verification system". In: Automated Deduction-CADE-11. Ed. by Deepak Kapur. Berlin, Heidelberg: Springer Berlin Heidelberg, 1992, pp. 748-752.
[Owr+95] S. Owre et al. "Formal verification for fault-tolerant architectures: prolegomena to the design of PVS". In: IEEE Transactions on Software Engineering 21.2 (1995), pp. 107-125. DOI: 10. 1109/32 . 345827.
[P416] The P4 Language Consortium. $P_{416}$ Language Specification, Version 1.2.2. Tech. rep. Available at https://p4.org/specs/, 2021. URL: https: //p4 . org/p4-spec/docs/P4-16-v1.2.2.pdf.
[PSA16] The P4.org Architecture Working Group. ${ }_{44_{16}}$ Portable Switch Architecture (PSA). Tech. rep. 2018.
[Ram+19] Tahina Ramananandro et al. "EverParse: Verified Secure Zero-Copy Parsers for Authenticated Message Formats". In: 28th USENIX Security Symposium (USENIX Security 19). Santa Clara, CA: USENIX Association, Aug. 2019, pp. 1465-1482.
[RKJo8] 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 'o8. Tucson, AZ, USA: Association for Computing Machinery, 2008, pp. 159-169. DoI: 10.1145/ 1375581. 1375602.
[RŞ1o] Grigore Roşu and Traian Florin Şerbănuță. "An Overview of the K Semantic Framework". In: Journal of Logic and Algebraic Programming 79.6 (2010), pp. 397-434. Doi: 10.1016/j.j lap. 2010.03.012.
[RV18] Gabriel Radanne and Jérôme Vouillon. "Tierless Web Programming in the Large". In: Companion Proceedings of the The Web Conference 2018. WWW '18. Lyon, France, 2018, pp. 681-689. doi: 10.1145/3184558. 3185953.
[Shu+19] Apoorv Shukla et al. "Runtime Verification of $\mathrm{P}_{4}$ Switches with Reinforcement Learning". In: Proceedings of the 2019 Workshop on Network Meets AI \& ML. NetAI'19. Beijing, China: Association for Computing Machinery, 2019, pp. 1-7. DOI: 10.1145/3341216. 3342206.
[Shu+20] Apoorv Shukla et al. "P4Consist: Toward Consistent P4 SDNs". In: IEEE Journal on Selected Areas in Communications 38.7 (2020), pp. 1293-1307. doi: 10.1109/JSAC. 2020. 2999653.
[Siv+16] Anirudh Sivaraman et al. "Packet Transactions: High-Level Programming for Line-Rate Switches". In: Proceedings of the 2016 ACM SIGCOMM Conference. SIGCOMM '16. Florianopolis, Brazil: ACM, 2016, pp. 15-28. Doi: 10.1145/2934872. 2934900.
[Smi+96] Jonathan M. Smith et al. "SwitchWare: Accelerating Network Evolution (White Paper)". In: University of Pennsylvania Department of Computer and Information Science Technical Report No. MS-CIS-96-38 (1996).
[Son +20 ] Hardik Soni et al. "Composing Dataplane Programs with $\mu \mathrm{P} 4$ ". In: Proceedings of the Annual Conference of the ACM Special Interest Group on Data Communication on the Applications, Technologies, Architectures, and Protocols for Computer Communication. SIGCOMM '20. Virtual Event, USA: Association for Computing Machinery, 2020, pp. 329-343. DoI: 10.1145/3387514.3405872.
[Sri18] Manu Sridharan. Engineering NullAway, Uber's Open Source Tool for Detecting NullPointerExceptions on Android. Dec. 2018. URL: https://eng. uber. com/nullaway/.
[STo6] Jeremy G. Siek and Walid Taha. "Gradual Typing for Functional Languages". In: IN SCHEME AND FUNCTIONAL PROGRAMMING WORKSHOP. 2006, pp. 81-92.
[Sto+16] Radu Stoenescu et al. "SymNet: Scalable Symbolic Execution for Modern Networks". In: ACM SIGCOMM. Florianopolis, Brazil: ACM, 2016, pp. 314327. DOI: 10.1145/2934872. 2934881.
[Sto+18] Radu Stoenescu et al. "Debugging P4 Programs with Vera". In: ACM SIGCOMM. Budapest, Hungary: ACM, 2018, pp. 518-532. DoI: 10.1145/ 3230543.3230548.
[SVo8] Jeremy G. Siek and Manish Vachharajani. "Gradual Typing with UnificationBased Inference". In: Proceedings of the 2008 Symposium on Dynamic Languages. DLS 'o8. Paphos, Cyprus: Association for Computing Machinery, 2008. Doi: 10. 1145/1408681. 1408688.
[Swa+16] Nikhil Swamy et al. "Dependent Types and Multi-Monadic Effects in $\mathrm{F}^{* \prime}$. In: Proceedings of the 43 rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages. POPL '16. St. Petersburg, FL, USA: Association for Computing Machinery, 2016, pp. 256-270. DoI: 10.1145/ 2837614.2837655.
[Ten+97] D.L. Tennenhouse et al. "A survey of active network research". In: IEEE Communications Magazine 35.1 (1997), pp. 80-86. Doi: 10 . 1109/35. 568214.
[TF06] Sam Tobin-Hochstadt and Matthias Felleisen. "Interlanguage Migration: From Scripts to Programs". In: Companion to the 21st ACM SIGPLAN Symposium on Object-Oriented Programming Systems, Languages, and Applications. OOPSLA 'o6. Portland, Oregon, USA: Association for Computing Machinery, 2006, pp. 964-974. DoI: 10. 1145/1176617. 1176755.
[TF1o] Sam Tobin-Hochstadt and Matthias Felleisen. "Logical Types for Untyped Languages". In: Proceedings of the 15th ACM SIGPLAN International Conference on Functional Programming. ICFP '1o. Baltimore, Maryland, USA: ACM, 2010, pp. 117-128. DoI: 10.1145/1863543. 1863561.
[Tia+21] Bingchuan Tian et al. "Aquila: A Practically Usable Verification System for Production-Scale Programmable Data Planes". In: Proceedings of the 2021 ACM SIGCOMM 2021 Conference. SIGCOMM '21. Virtual Event, USA: Association for Computing Machinery, 2021, pp. 17-32. DoI: 10.1145/ 3452296.3472937.
[TWo7] David L. Tennenhouse and David J. Wetherall. "Towards an Active Network Architecture". In: SIGCOMM Comput. Commun. Rev. 37.5 (Oct. 2007), pp. 81-94. DoI: 10.1145/1290168. 1290180.
[Vaz+14] Niki Vazou et al. "Refinement Types for Haskell". In: ICFP. 2014, pp. 269282.
[VH11] Andreas Voellmy and Paul Hudak. "Nettle: Taking the Sting Out of Programming Network Routers". In: Practical Aspects of Declarative Languages. Ed. by Ricardo Rocha and John Launchbury. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 235-249.
[VIS96] Dennis Volpano, Cynthia Irvine, and Geoffrey Smith. "A Sound Type System for Secure Flow Analysis". In: J. Comput. Secur. 4.2-3 (Jan. 1996), pp. 167-187.
[VKF12] Andreas Voellmy, Hyojoon Kim, and Nick Feamster. "Procera: A Language for High-Level Reactive Network Control". In: Proceedings of the First Workshop on Hot Topics in Software Defined Networks. HotSDN '12. Helsinki, Finland: Association for Computing Machinery, 2012, pp. 43-48. DOI: 10.1145/2342441. 2342451.
[VTV18] Niki Vazou, Éric Tanter, and David Van Horn. "Gradual Liquid Type Inference". In: Proceedings of the ACM on Programming Languages (PACMPL) 2.OOPSLA (Oct. 2018). Doi: 10.1145/3276502.
[Wan+17] Han Wang et al. "P4FPGA: A Rapid Prototyping Framework for P4". In: Proceedings of the Symposium on SDN Research. SOSR '17. Santa Clara, CA, USA: Association for Computing Machinery, 2017, pp. 122-135. DoI: 10.1145/3050220. 3050234.
[WF94] A.K. Wright and M. Felleisen. "A Syntactic Approach to Type Soundness". In: Inf. Comput. 115.1 (Nov. 1994), pp. 38-94. Doi: 10.1006/inco. 1994. 1093.
[WGT98] D.J. Wetherall, J.V. Guttag, and D.L. Tennenhouse. "ANTS: a toolkit for building and dynamically deploying network protocols". In: 1998 IEEE Open Architectures and Network Programming. 1998, pp. 117-129. DoI: 10.1109/OPNARC. 1998. 662048.
[WHM13] Christoph M. Wintersteiger, Youssef Hamadi, and Leonardo de Moura. "Efficiently solving quantified bit-vector formulas". In: Formal Methods in System Design 42.1 (2013), pp. 3-23. Doi: 10.1007/s10703-012-01562.
[XHo1] Hongwei Xi and Robert Harper. "A Dependently Typed Assembly Language". In: Proceedings of the Sixth ACM SIGPLAN International Conference on Functional Programming. ICFP 'o1. Florence, Italy: Association for Computing Machinery, 2001, pp. 169-18o. Doi: 10.1145/507635. 507657.
[Xioo] Hongwei Xi. "Imperative programming with dependent types". In: Proceedings Fifteenth Annual IEEE Symposium on Logic in Computer Science (Cat. No.99CB36332). 2000, pp. 375-387. DOI: 10.1109/LICS . 2000. 855785.
[XP98] Hongwei Xi and Frank Pfenning. "Eliminating Array Bound Checking through Dependent Types". In: Proceedings of the ACM SIGPLAN 1998 Conference on Programming Language Design and Implementation. PLDI '98. Montreal, Quebec, Canada: Association for Computing Machinery, 1998, pp. 249-257. DOI: 10.1145/277650. 277732.
[XP99] 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. San Antonio, Texas, USA: Association for Computing Machinery, 1999, pp. 214-227.
[Yan+04] L. Yang et al. Forwarding and Control Element Separation (ForCES) Famework. RFC 3746. RFC Editor, Apr. 2004. URL: https : / / www . rfc editor.org/rfc/rfc3746.txt.

## APPENDIX A

## Proofs

## A. 1 SAFEP4

## A.1.1 Operations on Header Types

Lemma A.1. $\left.\llbracket \Theta \rrbracket\right|_{h}=\llbracket$ Restrict $\Theta h \rrbracket$.

Proof. By induction on $\Theta$.

## Case $\Theta=0$ :

$$
\begin{aligned}
& \llbracket \mathrm{o} \rrbracket \mid h \\
= & \} \mid h \\
= & \{h s \mid h s \in\{ \} \wedge h \in \\
= & \} \\
= & \llbracket \mathrm{o} \rrbracket \\
= & \llbracket \text { Restrict } o h \rrbracket
\end{aligned}
$$

$$
=\{ \} \mid h \quad \text { by definition of } \llbracket \cdot \rrbracket
$$

$$
=\{h s \mid h s \in\{ \} \wedge h \in h s\} \quad \text { by definition of } . \mid h
$$

$$
\text { by definition of Restrict } . h
$$

Case $\Theta=1$ :

$$
\begin{aligned}
& \llbracket 1 \rrbracket \mid h \\
= & \{\}\} \mid h \\
= & \{h s \mid h s \in\{\{ \}\} \wedge h \in h s\} \\
= & \} \\
= & \llbracket \mathrm{o} \rrbracket
\end{aligned}
$$

by definition of $\llbracket . \rrbracket$
by definition of .|h
by set theory
by definition of $\llbracket . \rrbracket$

$$
=\llbracket \text { Restrict } 1 h \rrbracket \quad \text { by definition of Restrict } . h
$$

## Case $\Theta=g$ :

## $\llbracket g \rrbracket \mid h$

$$
\begin{array}{ll}
=\{\{g\}\} \mid h & \\
=\{h s \mid h s \in\{\{g\}\} \wedge h \in h s\} & \\
=\text { by definition of } \mathbb{\llbracket} . \rrbracket \\
& \text { by definition of } . \mid h
\end{array}
$$

Subcase $h=g$

| $=\{\{g\}\}$ | by set theory |
| :--- | ---: |
| $=\llbracket g \rrbracket$ | by definition of $\llbracket . \rrbracket$ |

by def. of Restrict . $h$ and by assumption $h=g$
$=\llbracket$ Restrict $g h \rrbracket$
Subcase $h \neq g$

$$
\begin{array}{lr}
=\{ \} & \text { by set theory } \\
=\llbracket o \rrbracket & \text { by definition of } \llbracket . \rrbracket \\
\text { by def. of Restrict . hand by assumption } h \neq g & \\
=\llbracket \text { Restrict } g h \rrbracket &
\end{array}
$$

by set theory by definition of $\llbracket . \rrbracket$

Case $\Theta=\Theta_{1} \cdot \Theta_{2}$ :

$$
\begin{aligned}
& \llbracket \Theta_{1} \cdot \Theta_{2} \rrbracket \mid h \\
& =\left\{h s_{1} \cup h s_{2} \mid h s_{1} \in \llbracket \Theta_{1} \rrbracket \wedge h s_{2} \in \llbracket \Theta_{2} \llbracket\right\} \mid h \\
& =\left\{h s_{1} \cup h s_{2} \mid h s_{1} \in \llbracket \Theta_{1} \rrbracket \wedge h s_{2} \in \llbracket \Theta_{2} \rrbracket \wedge h \in\left(h s_{1} \cup h s_{2}\right)\right\} \\
& =\left\{h s_{1} \cup h s_{2} \mid h s_{1} \in \llbracket \Theta_{1} \rrbracket \wedge h s_{2} \in \llbracket \Theta_{2} \rrbracket \wedge h \in h s_{1}\right\} \cup \\
& \left\{h s_{1} \cup h s_{2} \mid h s_{1} \in \llbracket \Theta_{1} \rrbracket \wedge h s_{2} \in \llbracket \Theta_{2} \rrbracket \wedge h \in h s_{2}\right\} \\
& \text { by def. of } \llbracket . \rrbracket \\
& \text { by def. of . } h \\
& \text { by set theory } \\
& \text { by logic and set theory } \\
& =\left\{h s_{1} \cup h s_{2} \mid\left(h s_{1} \in \llbracket \Theta_{1} \rrbracket \wedge h \in h s_{1}\right) \wedge h s_{2} \in \llbracket \Theta_{2} \rrbracket\right\} \cup \\
& \left\{h s_{1} \cup h s_{2} \mid h s_{1} \in \llbracket \Theta_{1} \rrbracket \wedge\left(h s_{2} \in \llbracket \Theta_{2} \rrbracket \wedge h \in h s_{2}\right)\right\} \\
& =\left\{h s_{1} \mid h s_{1} \in \llbracket \Theta_{1} \rrbracket \wedge h \in h s_{1}\right\} \bullet\left\{h s_{2} \mid h s_{2} \in \llbracket \Theta_{2} \rrbracket\right\} \cup \quad \text { by def. of } S_{1} \bullet S_{2} \\
& \left\{h s_{1} \mid h s_{1} \in \llbracket \Theta_{1} \rrbracket\right\} \bullet\left\{h s_{2} \mid h s_{2} \in \llbracket \Theta_{2} \rrbracket \wedge h \in h s_{2}\right\} \\
& =\left\{h s_{1} \mid h s_{1} \in \llbracket \Theta_{1} \rrbracket\right\} \mid h \bullet\left\{h s_{2} \mid h s_{2} \in \llbracket \Theta_{2} \rrbracket\right\} \cup \quad \text { by def. of . } \mid h \\
& \left\{h s_{1} \mid h s_{1} \in \llbracket \Theta_{1} \rrbracket\right\} \bullet\left\{h s_{2} \mid h s_{2} \in \llbracket \Theta_{2} \rrbracket\right\} \mid h \\
& =\llbracket \Theta_{1} \rrbracket\left|h \bullet \llbracket \Theta_{2} \rrbracket \cup \llbracket \Theta_{1} \rrbracket \bullet \llbracket \Theta_{2} \rrbracket\right| h \\
& =\llbracket \text { Restrict } \Theta_{1} h \rrbracket \bullet \llbracket \Theta_{2} \rrbracket \cup \llbracket \Theta_{1} \rrbracket \bullet \llbracket \text { Restrict } \Theta_{2} h \rrbracket \\
& \text { by def. of } \llbracket . \rrbracket \\
& \text { by def. of } S_{1} \bullet S_{2} \text { and } \llbracket . \rrbracket \\
& =\llbracket \text { Restrict } \Theta_{1} h \cdot \Theta_{2}+\Theta_{1} \cdot \text { Restrict } \Theta_{2} h \rrbracket \\
& \text { by def. of Restrict. } h \\
& =\llbracket \text { Restrict }\left(\Theta_{1} \cdot \Theta_{2}\right) h \rrbracket
\end{aligned}
$$

Case $\Theta=\Theta_{1}+\Theta_{2}$ :

$$
\begin{array}{rlr} 
& \llbracket \Theta_{1}+\Theta_{2} \rrbracket \mid h & \\
= & \left(\llbracket \Theta_{1} \rrbracket \cup \llbracket \Theta_{2} \rrbracket\right) \mid h & \text { by definition of } \llbracket . \rrbracket \\
= & \left\{h s \mid h s \in\left(\llbracket \Theta_{1} \rrbracket \cup \llbracket \Theta_{2} \rrbracket\right) \wedge h \in h s\right\} & \text { by definition of . } \mid h \\
= & \left\{h s_{1} \mid h s_{1} \in \llbracket \Theta_{1} \rrbracket \wedge h \in h s_{1}\right\} \cup\left\{h s_{2} \mid h s_{2} \in \llbracket \Theta_{2} \rrbracket \wedge h \in h s_{2}\right\} & \text { by set theory } \\
= & \llbracket \Theta_{1} \rrbracket\left|h \cup \llbracket \Theta_{2} \rrbracket\right| h & \text { by definition of . } \mid h \\
\text { by induction hypothesis } & \\
= & \llbracket \text { Restrict } \Theta_{1} h \rrbracket \cup \llbracket \text { Restrict } \Theta_{2} h \rrbracket & \\
= & \llbracket \text { Restrict } \Theta_{1} h+\text { Restrict } \Theta_{2} h \rrbracket & \text { by definition of } \llbracket . \rrbracket \\
\text { by definition of Restrict. } h & \\
= & \llbracket \text { Restrict }\left(\Theta_{1}+\Theta_{2}\right) h \rrbracket &
\end{array}
$$

Lemma A.2. $\left.\llbracket \Theta \rrbracket\right|_{\neg h}=\llbracket$ NegRestrict $\Theta h \rrbracket$.

Proof. By induction on $\Theta$.

Case $\Theta=0$ :

$$
\begin{array}{rlr} 
& \llbracket \mathrm{o} \rrbracket \mid \neg h & \\
= & \} \mid \neg h & \text { by definition of } \llbracket \cdot \rrbracket \\
= & \{h s \mid h s \in\{ \} \wedge h \notin h s\} & \text { by definition of } . \mid \neg h \\
= & \} & \text { by set theory } \\
= & \llbracket \mathrm{o} \rrbracket & \text { by definition of } \llbracket \cdot \rrbracket \\
= & \llbracket \text { NegRestrict o } h \rrbracket & \text { by definition of NegRestrict } . h
\end{array}
$$

Case $\Theta=1$ :

$$
\begin{array}{rlr} 
& \llbracket 1 \rrbracket \mid \neg h & \\
= & \{\}\} \mid \neg h & \text { by definition of } \llbracket . \rrbracket \\
= & \{h s \mid h s \in\{\{ \}\} \wedge h \notin h s\} & \text { by definition of } . \mid \neg h \\
= & \{\}\} & \text { by set theory } \\
= & \llbracket 1 \rrbracket & \text { by definition of } \llbracket . \rrbracket \\
= & \llbracket \text { NegRestrict } 1 h \rrbracket & \text { by definition of NegRestrict. } h
\end{array}
$$

Case $\Theta=g$ :

$$
\begin{array}{rlr} 
& \llbracket g \rrbracket \mid \neg h & \\
= & \{\{g\}\} \mid \neg h & \text { by definition of } \llbracket . \rrbracket \\
= & \{h s \mid h s \in\{\{g\}\} \wedge h \notin h s\} & \text { by definition of } . \mid \neg h \\
\text { Subcase } h=g & \\
= & \} & \text { by set theory } \\
= & \llbracket o \rrbracket & \text { by definition of } \llbracket . \rrbracket
\end{array}
$$

by assumption $h=g$ and by definition of NegRestrict.$h$
$=\llbracket$ NegRestrict o $h \rrbracket$
Subcase $h \neq g$

$$
\begin{array}{lr}
=\{\{g\}\} & \text { by set theory } \\
=\llbracket g \rrbracket & \text { by definition of } \llbracket . \rrbracket
\end{array}
$$

by assumption $h \neq g$ and by definition of NegRestrict . $h$
$=\llbracket$ NegRestrict $g h \rrbracket$

Case $\Theta=\Theta_{1} \cdot \Theta_{2}$ :

$$
\llbracket \Theta_{1} \cdot \Theta_{2} \rrbracket \mid \neg h
$$

by definition of $S_{1} \bullet S_{2}$
$=\left(\llbracket \Theta_{1} \rrbracket \bullet \llbracket \Theta_{2} \rrbracket\right) \mid \neg h$
by definition of $\llbracket . \rrbracket$
$=\left\{h s_{1} \cup h s_{2} \mid h s_{1} \in \llbracket \Theta_{1} \rrbracket \wedge h s_{2} \in \llbracket \Theta_{2} \llbracket\right\} \mid \neg h$
by definition of. $\mid \neg h$
$=\left\{h s_{1} \cup h s_{2} \mid h s_{1} \in \llbracket \Theta_{1} \rrbracket \wedge h s_{2} \in \llbracket \Theta_{2} \rrbracket \wedge h \notin\left(h s_{1} \cup h s_{2}\right)\right\}$
by set theory and logic
$=\left\{h s_{1} \cup h s_{2} \mid h s_{1} \in \llbracket \Theta_{1} \rrbracket \wedge h s_{2} \in \llbracket \Theta_{2} \rrbracket \wedge h \notin h s_{1} \wedge h \notin h s_{2}\right\}$
by set theory and logic
$=\left\{h s_{1} \cup h s_{2} \mid\left(h s_{1} \in \llbracket \Theta_{1} \rrbracket \wedge h \notin h s_{1}\right) \wedge\left(h s_{2} \in \llbracket \Theta_{2} \rrbracket \wedge h \notin h s_{2}\right)\right\}$
by definition of $S_{1} \bullet S_{2}$
$=\left\{h s_{1} \mid h s_{1} \in \llbracket \Theta_{1} \rrbracket \wedge h \notin h s_{1}\right\} \bullet\left\{h s_{2} \mid h s_{2} \in \llbracket \Theta_{2} \rrbracket \wedge h \notin h s_{2}\right\}$
by definition of. $\mid \neg h$
$=\llbracket \Theta_{1} \rrbracket\left|\neg h \bullet \llbracket \Theta_{2} \rrbracket\right| \neg h$
by induction hypothesis
$=\llbracket$ NegRestrict $\Theta_{1} h \rrbracket \bullet \llbracket$ NegRestrict $\Theta_{2} h \rrbracket$
By definition of $\llbracket . \rrbracket$
$=\llbracket\left(\right.$ NegRestrict $\left.\Theta_{1} h\right) \cdot\left(\right.$ NegRestrict $\left.\Theta_{2} h\right) \rrbracket$
by definition of NegRestrict. $h$
$=\llbracket$ NegRestrict $\left(\Theta_{1} \cdot \Theta_{2}\right) h \rrbracket$

Case $\Theta=\Theta_{1}+\Theta_{2}$ :

$$
\llbracket \Theta_{1}+\Theta_{2} \rrbracket \mid \neg h
$$

by definition of $\llbracket . \rrbracket$

$$
=\left(\llbracket \Theta_{1} \rrbracket \cup \llbracket \Theta_{2} \rrbracket\right) \mid \neg h
$$

by definition of . $\mid \neg h$
$=\left\{h s \mid h s \in\left(\llbracket \Theta_{1} \rrbracket \cup \llbracket \Theta_{2} \rrbracket\right) \wedge h \notin h s\right\}$
by set theory
$=\left\{h s_{1} \mid h s_{1} \in \llbracket \Theta_{1} \rrbracket \wedge h \notin h s_{1}\right\} \cup\left\{h s_{2} \mid h s_{2} \in \llbracket \Theta_{2} \rrbracket \wedge h \notin h s_{2}\right\}$
by definition of $\mid \neg h$
$=\llbracket \Theta_{1} \rrbracket\left|\neg h \cup \llbracket \Theta_{2} \rrbracket\right| \neg h$
by induction hypothesis
$=\llbracket$ NegRestrict $\Theta_{1} h \rrbracket \cup \llbracket$ NegRestrict $\Theta_{2} h \rrbracket$
by definition of $\llbracket$.】
$=\llbracket$ NegRestrict $\Theta_{1} h+$ NegRestrict $\Theta_{2} h \rrbracket$
by definition of NegRestrict. $\neg h$
$=\llbracket$ NegRestrict $\left(\Theta_{1}+\Theta_{2}\right) h \rrbracket$

Lemma A.3. $\forall h s \in S . h \in h s==$ Includes $\Theta h$.

Proof. By induction on $\Theta$.

Case $\Theta=0$ :

$$
\begin{aligned}
& h \sqsubset \llbracket \mathrm{o} \rrbracket \\
= & \bigwedge(h s \in\} \wedge h \in h s) \\
= & \text { true } \quad \text { by definition of } \llbracket . \rrbracket \\
= & \text { Includes o } h
\end{aligned} \quad \text { by logic and set theory } 0 \text { by definition of Includes } . h
$$

Case $\Theta=1$ :

$$
h \sqsubset \llbracket 1 \rrbracket
$$

$$
=\bigwedge(h s \in\{\{ \}\} \wedge h \in h s) \quad \text { by definition of } \llbracket . \rrbracket
$$

$$
=\text { false } \quad \text { by logic and set theory }
$$

$$
=\text { Includes } 1 h \quad \text { by definition of Includes } . h
$$

## Case $\Theta=g$ :

$$
\begin{aligned}
& \quad h \sqsubset \llbracket g \rrbracket \\
& =\bigwedge(h s \in\{\{g\}\} \wedge h \in h s) \quad \text { by definition of } \llbracket \cdot \rrbracket \\
& \text { Subcase } h=g \quad \\
& =\text { true by logic and set theory } \\
& \text { by definition of (Includes } . h) \text { and assumption } h=g \\
& =\text { Includes } g h \\
& \begin{array}{l}
\text { Subcase } h \neq g \\
= \\
\text { balse } \\
\text { by definition of }(\text { Includes } . h) \text { and assumption } h \neq g \\
=
\end{array} \\
& \text { Includes } g h \quad \text { by logic and set theory }
\end{aligned}
$$

Case $\Theta=\Theta_{1} \cdot \Theta_{2}$ :

$$
\begin{aligned}
& \quad h \sqsubset \llbracket \Theta_{1} \cdot \Theta_{2} \rrbracket \\
& \text { by definition of } S_{1} \bullet S_{2} \\
& =h \sqsubset\left(\llbracket \Theta_{1} \rrbracket \bullet \llbracket \Theta_{2} \rrbracket\right) \\
& \text { by definition of } \llbracket \cdot \rrbracket \\
& =h \sqsubset\left\{h s_{1} \cup h s_{2} \mid h s_{1} \in \llbracket \Theta_{1} \rrbracket \wedge h s_{2} \in \llbracket \Theta_{2} \rrbracket\right\} \\
& \text { by set theory and logic } \\
& =h \sqsubset\left\{h s_{1} \mid h s_{1} \in \llbracket \Theta_{1} \rrbracket\right\} \vee h \sqsubset\left\{h s_{2} \mid h s_{2} \in \llbracket \Theta_{2} \rrbracket\right\} \\
& \text { by definition of } h \sqsubset . \\
& =\bigwedge\left(h s_{1} \in \llbracket \Theta_{1} \rrbracket \wedge h \in h s_{1}\right) \vee \bigwedge\left(h s_{2} \in \llbracket \Theta_{2} \rrbracket \wedge h \in h s_{2}\right) \\
& \text { by definition of } \llbracket \cdot \rrbracket \\
& =h \sqsubset \llbracket \Theta_{1} \rrbracket \vee h \sqsubset \llbracket \Theta_{2} \rrbracket \\
& \text { by induction hypothesis } \\
& =\left(\text { Includes } \Theta_{1} h\right) \vee\left(\text { Includes } \Theta_{2} h\right) \\
& \text { by definition of Includes } . h \\
& =\left(\text { Includes } \Theta_{1} \cdot \Theta_{2} h\right)
\end{aligned}
$$

Case $\Theta=\Theta_{1}+\Theta_{2}$ :

$$
h \sqsubset \llbracket \Theta_{1}+\Theta_{2} \rrbracket
$$

by definition of $\llbracket . \rrbracket$
$=h \check{ }\left(\llbracket \Theta_{1} \rrbracket \cup \llbracket \Theta_{2} \rrbracket\right)$
by definition of $h \sqsubset$.

$$
=\wedge\left(h s \in\left(\llbracket \Theta_{1} \rrbracket \cup \llbracket \Theta_{2} \rrbracket\right) \wedge h \in h s\right)
$$

by set theory and logic
$=\bigwedge\left(h s_{1} \in \llbracket \Theta_{1} \rrbracket \wedge h \in h s_{1} \wedge h s_{2} \in \llbracket \Theta_{2} \rrbracket \wedge h \in h s_{2}\right)$
by set theory and logic
$=\bigwedge\left(h s_{1} \in \llbracket \Theta_{1} \rrbracket \wedge h \in h s_{1}\right) \wedge \bigwedge\left(h s_{2} \in \llbracket \Theta_{2} \rrbracket \wedge h \in h s_{2}\right)$
by definition of $h \sqsubset$.
$=h \sqsubset \llbracket \Theta_{1} \rrbracket \wedge h \sqsubset \llbracket \Theta_{2} \rrbracket$
by induction hypothesis
$=\left(\right.$ Includes $\left.\Theta_{1} h\right) \wedge\left(\right.$ Includes $\left.\Theta_{2} h\right)$
by definition of Includes. $h$
$=\left(\right.$ Includes $\left.\left(\Theta_{1}+\Theta_{2}\right) h\right)$

Lemma A.4. $\llbracket \Theta \rrbracket \backslash h==\llbracket$ Remove $\Theta h \rrbracket$.

Proof. By induction on $\Theta$.

Case $\Theta=0$ :

$$
\llbracket o \rrbracket \backslash h
$$

$$
=\{ \} \backslash h \quad \text { by definition of } \llbracket . \rrbracket
$$

$$
=\{h s \mid h s \in\{ \} \wedge h s \backslash\{h\}\} \quad \text { by definition of } . \backslash h
$$

$$
=\{ \} \quad \text { by set theory }
$$

$$
=\llbracket \mathrm{o} \rrbracket \quad \text { by definition of } \llbracket . \rrbracket
$$

$$
=\llbracket \text { Remove o } h \rrbracket \quad \text { by definition of Remove . } h
$$

Case $\Theta=1$ :

$$
\begin{aligned}
& \llbracket 1 \rrbracket \backslash h \\
= & \{\}\} \backslash \\
= & \{h s \mid h s \\
= & \{\}\} \\
= & \llbracket 1 \rrbracket
\end{aligned}
$$

$$
=\{\{ \}\} \backslash h \quad \text { by definition of } \llbracket . \rrbracket
$$

$$
=\{h s \mid h s \in\{\{ \}\} \wedge h s \backslash\{h\}\} \quad \text { by definition of } . \backslash h
$$

$$
=\llbracket \text { Remove } 1 h \rrbracket \quad \text { by definition of Remove } . h
$$

Case $\Theta=g$ :

|  | $\llbracket g \rrbracket \backslash h$ |  |
| ---: | :--- | ---: |
| $=$ | $\{\{g\}\} \backslash h$ | by definition of $\llbracket . \rrbracket$ |
| $=$ | $\{h s \mid h s \in\{\{g\}\} \wedge h s \backslash\{h\}\}$ | by definition of . $h$ |
| Subcase $h=g$ |  |  |
| $=$ | $\{\}\}$ | by set theory |
| $=$ | $\llbracket 1 \rrbracket$ | by definition of $\llbracket . \rrbracket$ |
| $=$ | $\llbracket$ Remove $1 h \rrbracket$ | by definition of Remove.$h$ |

Subcase $h \neq g$
$=\{\{g\}\}$
$=\llbracket g \rrbracket$
by set theory
by definition of $\llbracket . \rrbracket$
$=\llbracket$ Remove $g h \rrbracket \quad$ by definition of Remove.$h$

Case $\Theta=\Theta_{1} \cdot \Theta_{2}$ :

$$
\llbracket \Theta_{1} \cdot \Theta_{2} \rrbracket \backslash h
$$

by definition of $S_{1} \bullet S_{2}$
$=\left(\llbracket \Theta_{1} \rrbracket \bullet \llbracket \Theta_{2} \rrbracket\right) \backslash h$
by definition of $\llbracket . \rrbracket$
$=\left\{h s_{1} \cup h s_{2} \mid h s_{1} \in \llbracket \Theta_{1} \rrbracket \wedge h s_{2} \in \llbracket \Theta_{2} \llbracket\right\} \backslash h$
by definition of . $\backslash h$
$=\left\{h s_{1} \cup h s_{2} \mid h s_{1} \in \llbracket \Theta_{1} \rrbracket \wedge h s_{2} \in \llbracket \Theta_{2} \rrbracket \wedge\left(h s_{1} \cup h s_{2}\right) \backslash h\right\}$
by set theory and logic
$=\left\{h s_{1} \cup h s_{2} \mid h s_{1} \in \llbracket \Theta_{1} \rrbracket \wedge h s_{2} \in \llbracket \Theta_{2} \rrbracket \wedge h s_{1} \backslash h \wedge h s_{2} \backslash h\right\}$
by set theory and logic
$=\left\{h s_{1} \cup h s_{2} \mid\left(h s_{1} \in \llbracket \Theta_{1} \rrbracket \wedge h s_{1} \backslash h\right) \wedge\left(h s_{2} \in \llbracket \Theta_{2} \rrbracket \wedge h s_{2} \backslash h\right)\right\}$
by definition of $S_{1} \bullet S_{2}$
$=\left\{h s_{1} \mid h s_{1} \in \llbracket \Theta_{1} \rrbracket \wedge h s_{1} \backslash h\right\} \bullet\left\{h s_{2} \mid h s_{2} \in \llbracket \Theta_{2} \rrbracket \wedge h s_{2} \backslash h\right\}$
by definition of . $h$
$=\llbracket \Theta_{1} \rrbracket \backslash h \bullet \llbracket \Theta_{2} \rrbracket \backslash h$
by induction hypothesis
$=\llbracket$ Remove $\Theta_{1} h \rrbracket \bullet \llbracket$ Remove $\Theta_{2} h \rrbracket$
By definition of $\llbracket . \rrbracket$
$=\llbracket\left(\right.$ Remove $\left.\Theta_{1} h\right) \cdot\left(\right.$ Remove $\left.\Theta_{2} h\right) \rrbracket$
by definition of Remove . $h$
$=\llbracket \operatorname{Remove}\left(\Theta_{1} \cdot \Theta_{2}\right) h \rrbracket$

Case $\Theta=\Theta_{1}+\Theta_{2}$ :

$$
\begin{aligned}
& \llbracket \Theta_{1}+\Theta_{2} \rrbracket \backslash h \\
= & \left(\llbracket \Theta_{1} \rrbracket \cup \llbracket \Theta_{2} \rrbracket\right) \backslash h \\
= & \left\{h s \mid h s \in\left(\llbracket \Theta_{1} \rrbracket \cup \llbracket \Theta_{2} \rrbracket\right) \wedge h s \backslash\{h\}\right\} \\
= & \left\{h s \mid h s \in\left(\llbracket \Theta_{1} \rrbracket \cup \llbracket \Theta_{2} \rrbracket\right) \wedge h s \backslash\{h\}\right\} \\
= & \left\{h s_{1} \mid h s_{1} \in \llbracket \Theta_{1} \rrbracket \wedge h s_{1} \backslash\{h\}\right\} \cup \\
& \left\{h s_{2} \mid h s_{2} \in \llbracket \Theta_{2} \rrbracket \wedge h s_{2} \backslash\{h\}\right\} \\
= & \llbracket \Theta_{1} \rrbracket \backslash h \cup \llbracket \Theta_{2} \rrbracket \backslash h \\
= & \llbracket \operatorname{Remove} \Theta_{1} h \rrbracket \cup \llbracket \text { Remove } \Theta_{2} h \rrbracket \\
= & \llbracket\left(\text { Remove } \Theta_{1} h\right) \cdot\left(\llbracket \operatorname{Remove} \Theta_{2} h\right) \rrbracket \\
= & \llbracket \operatorname{Remove}\left(\Theta_{1} \cdot \Theta_{2}\right) h \rrbracket
\end{aligned}
$$

by definition of $\llbracket . \rrbracket$ by definition of . $h$ by definition of . $\backslash h$ by logic and set theory
by definition of . $h$
by induction hypothesis
by definition of $\llbracket . \rrbracket$ by definition of Remove. $h$

Lemma A.5. $\llbracket \Theta \rrbracket==\{ \}$ if and only if Empty $\Theta$.
Proof. By induction on $\Theta$.
Case $\Theta=\mathrm{o}$ : We have $\llbracket \mathrm{o} \rrbracket=\{ \}$ and Empty $\mathrm{o}=$ true.
Case $\Theta=1$ : We have $\llbracket 1 \rrbracket \neq\{ \}$ and Empty $1=$ false.
Case $\Theta=h$ : We have $\llbracket h \rrbracket \neq\{ \}$ and Empty $h=$ false.
Case $\Theta=\Theta_{1} \cdot \Theta_{2}$ : By definition we have $\llbracket \Theta_{1} \cdot \Theta_{2} \rrbracket=\llbracket \Theta_{1} \rrbracket \bullet \llbracket \Theta_{2} \rrbracket$ which is equal to $\left\{S_{1} \cup S_{2} \mid S_{1} \in \llbracket \Theta_{1} \rrbracket \wedge S_{2} \in \llbracket \Theta_{2} \rrbracket\right\}$. It follows that $\llbracket \Theta_{1} \cdot \Theta_{2} \rrbracket=\{ \}$ iff $\llbracket \Theta_{1} \rrbracket=$ $\left\}\right.$ or $\llbracket \Theta_{2} \rrbracket=\{ \}$. By induction hypothesis, we have $\llbracket \Theta_{1} \rrbracket=\{ \}$ if and only if Empty $\Theta_{1}=$ true, and $\llbracket \Theta_{2} \rrbracket=\{ \}$ if and only if Empty $\Theta_{2}=$ true. The result follows as Empty $\left(\Theta_{1} \cdot \Theta_{2}\right)=$ Empty $\Theta_{1} \vee \operatorname{Empty} \Theta_{2}$.

Case $\Theta=\Theta_{1}+\Theta_{2}$ : By definition we have $\llbracket \Theta_{1}+\Theta_{2} \rrbracket=\llbracket \Theta_{1} \rrbracket \cup \llbracket \Theta_{2} \rrbracket$. It follows that $\llbracket \Theta_{1} \cdot \Theta_{2} \rrbracket \neq\{ \}$ iff $\llbracket \Theta_{1} \rrbracket \neq\{ \}$ and $\llbracket \Theta_{2} \llbracket \neq\{ \}$. By induction hypothesis, we have $\llbracket \Theta_{1} \rrbracket \neq\{ \}$ if and only if Empty $\Theta_{1}=$ true, and $\llbracket \Theta_{2} \rrbracket \neq\{ \}$ if and only if Empty $\Theta_{2}=$ true. The result follows as Empty $\left(\Theta_{1}+\Theta_{2}\right)=$ Empty $\Theta_{1} \wedge$ Empty $\Theta_{2}$.

## A.1. 2 Safety

Lemma A. 6 (Expression Substitution). If $\Gamma, x: \tau ; \Theta \vdash e: \tau^{\prime}$ and $\because \cdot \vdash \bar{v}: \bar{\tau}$ then $\Gamma ; \Theta \vdash e[\bar{v} / \bar{x}]: \tau^{\prime}$

Lemma A. 7 (Entailment is Type Alternative). If $H \vDash \Theta$ then $\operatorname{dom}(H) \in \llbracket \Theta \rrbracket$.
Proof. By induction on $\Theta$.
Case $\Theta=0$ : The case immediately holds as $H \vDash \mathrm{o}$ is a contradiction.
Case $\Theta=1$ : By inversion of Entailment, $H=\cdot$, and so $\operatorname{dom}(H)=\{ \} \in \llbracket 1 \rrbracket=\{\{ \}\}$.

Case $\Theta=h$ : By inversion of Entailment, $\operatorname{dom}(H)=\{h\} \in \llbracket h \rrbracket=\{\{h\}\}$.
Case $\Theta=\Theta_{1} \cdot \Theta_{2}$ : By inversion of Entailment, $H=H_{1} \cup H_{2}, H_{1} \vDash \Theta_{1}, H_{2} \vDash \Theta_{2}$. By induction hypothesis, $\operatorname{dom}\left(H_{1}\right) \in \llbracket \Theta_{1} \rrbracket$ and $\operatorname{dom}\left(H_{2}\right) \in \llbracket \Theta_{2} \rrbracket$. By set theory, $\operatorname{dom}(H)=\operatorname{dom}\left(H_{1}\right) \cup \operatorname{dom}\left(H_{2}\right)$ By definition of $\llbracket . \rrbracket$ and $(\bullet), \llbracket \Theta \rrbracket=\llbracket \Theta_{1} \rrbracket \bullet \llbracket \Theta_{2} \rrbracket=\left\{h s_{1} \cup h s_{2} \mid h s_{1} \in \llbracket \Theta_{1} \rrbracket \wedge h s_{2} \in\right.$ $\left.\llbracket \Theta_{2} \rrbracket\right\}$ and therefore $\operatorname{dom}\left(H_{1}\right) \cup \operatorname{dom}\left(H_{2}\right) \in\left\{h s_{1} \cup h s_{2} \mid h s_{1} \in \llbracket \Theta_{1} \rrbracket \wedge h s_{2} \in \llbracket \Theta_{2} \rrbracket\right\}$, i.e., $\operatorname{dom}(H) \in \llbracket \Theta \rrbracket$.

Case $\Theta=\Theta_{1}+\Theta_{2}$ : By inversion of Entailment, either $H \vDash \Theta_{1}$ or $H \vDash \Theta_{2}$.
Subcase $H \vDash \Theta_{1}$ : By the induction hypothesis, $\operatorname{dom}(H) \in \llbracket \Theta_{1} \rrbracket$ and by set theory $\operatorname{dom}(H) \in \llbracket \Theta_{1} \rrbracket \cup \llbracket \Theta_{2} \rrbracket$
Subcase $H \vDash \Theta_{2}$ : Symmetric to the previous subcase.

Lemma A. 8 (Included Instances in Domain). If $H \vDash \Theta$ and Includes $\Theta h$, then $h \in \operatorname{dom}(H)$.

Proof. By induction on $\Theta$.
Case $\Theta=0$ : The case immediately holds as $H \vDash \mathrm{o}$ is a contradiction.
Case $\Theta=1$ :
By inversion of Entailment, $H=\boldsymbol{\cdot}$. The case immediately holds, as Includes $\Theta h$ is a contradiction.

Case $\Theta=g$ :
By inversion of Entailment, $\operatorname{dom}(H)=\{g\}$. By assumption Includes $\Theta h, h=g$ and thus $h \in \operatorname{dom}(H)$.

Case $\Theta=\Theta_{1} \cdot \Theta_{2}$ : By inversion of Entailment, $H=H_{1} \cup H_{2}, H_{1} \vDash \Theta_{1}, H_{2} \vDash \Theta_{2}$. By set theory $\operatorname{dom}(H)=\operatorname{dom}\left(H_{1}\right) \cup \operatorname{dom}\left(H_{2}\right)$. By definition of Inclusion and by assumption Includes $\Theta h$, Includes $\Theta_{1} h \vee$ Includes $\Theta_{2} h$.

Subcase Includes $\Theta_{1} h$ : By induction hypothesis, $h \in \operatorname{dom}\left(H_{1}\right)$ and by assumption $\operatorname{dom}\left(H_{1}\right) \subseteq \operatorname{dom}(H)$, we can conclude $h \in \operatorname{dom}(H)$.
Subcase Includes $\Theta_{2} h$ : Symmetric to the previous subcase.
Case $\Theta=\Theta_{1}+\Theta_{2}$ : By inversion of Entailment, either $H \vDash \Theta_{1}$ or $H \vDash \Theta_{2}$. By definition of Inclusion and by assumption Includes $\Theta h$, Includes $\Theta_{1} h$ and Includes $\Theta_{2} h$.

Subcase $H \vDash \Theta_{1}$ : By induction hypothesis, we can conclude $h \in \operatorname{dom}(H)$.
Subcase $H \vDash \Theta_{2}$ : Symmetric to the previous subcase.

## Progress

Theorem A. 9 (Progress). If $\cdot \vdash c: \Theta \Leftrightarrow \Theta^{\prime}$ and $H \vDash \Theta$, then either $c=$ skip or $\exists\left\langle I^{\prime}, O^{\prime}, H^{\prime}, c^{\prime}\right\rangle .\langle I, O, H, c\rangle \rightarrow\left\langle I^{\prime}, O^{\prime}, H^{\prime}, c^{\prime}\right\rangle$

Proof. By induction on typing derivations of $\cdot \vdash c: \Theta \Leftrightarrow \Theta^{\prime}$.
Case T-Skip: $c=s k i p$
Immediate.
Case T-Extr: $c=\operatorname{extract}(h)$
Let $\left(I^{\prime}, v\right)=\operatorname{deserialize}_{\eta}(I)$ and $O^{\prime}=O$ and $H^{\prime}=H[h \mapsto v]$ and $c^{\prime}=s k i p$. The result follows by E-Extract.

Case T-Emit: $c=\operatorname{emit}(h)$
If $h \notin \operatorname{dom}(H)$, let $I^{\prime}=I$ and $O^{\prime}=O$ and $H^{\prime}=H$, and $c^{\prime}=s k i p$. The result follows by E-EmitInvalid. Otherwise, $h \in \operatorname{dom}(H)$. Let $H(h)=v$ and $\bar{B}=\operatorname{serialize}_{\eta}(v)$ and $I^{\prime}=I$ and $O^{\prime}=O \cdot \bar{B}$ and $H^{\prime}=H$ and $c^{\prime}=s k i p$. The result follows by E-Emit.

Case T-SEQ: $c=c_{1} ; c_{2}$ and $\cdot \vdash c_{1}: \Theta \Leftrightarrow \Theta_{1}$ and $\cdot \vdash c_{2}: \Theta_{1} \Leftrightarrow \Theta_{2}$
By induction hypothesis, $c_{1}$ is either skip or there is some $\left\langle I^{\prime}, O^{\prime}, H^{\prime}, c_{1}^{\prime}\right\rangle$, such that $\left\langle I, O, H, c_{1}\right\rangle \rightarrow\left\langle I^{\prime}, O^{\prime}, H^{\prime}, c_{1}^{\prime}\right\rangle$.
If $c_{1}=s k i p$, let $I^{\prime}=I$ and $O^{\prime}=O$ and $H^{\prime}=H$ and $c^{\prime}=c_{2}$. The result follows by E-Seq. Otherwise, the result follows by E-Seq1.

Case T-IF: $c=$ if $(e)$ then $c_{1}$ else $c_{2}$ and $; \Theta \vdash e:$ Bool and $\cdot \vdash c_{1}: \Theta \Leftrightarrow \Theta_{1}$ and $\cdot \vdash c_{2}: \Theta \Leftrightarrow \Theta_{2}$
By the progress theorem for expressions, we have that $e$ is either true, false, or there is some $e^{\prime}$ such that $\langle H, e\rangle \rightarrow e^{\prime}$.

Subcase $e=$ true: Let $I^{\prime}=I$ and $O^{\prime}=O$ and $H^{\prime}=H$ and $c^{\prime}=c_{1}$. The result follows by E-If True.
Subcase $e=$ false: Symmetric to the previous case.
Subcase $\langle H, e\rangle \rightarrow e^{\prime}$ : Let $I^{\prime}=I$ and $O^{\prime}=O$ and $H^{\prime}=H$ and $c^{\prime}=$ if $\left(e^{\prime}\right) c_{1} c_{2}$. The result follows by E-IF.

Case T-IF Valid: $c=\operatorname{valid}(h) c_{1}$ else $c_{2}$
If $h \in \operatorname{dom}(H)$, let $I^{\prime}=I$ and $O^{\prime}=O$ and $H^{\prime}=H$ and $c^{\prime}=c_{1}$. The result follows by E-If ValidTrue Otherwise, $h \notin \operatorname{dom}(H)$. Let $I^{\prime}=I$ and $O^{\prime}=O$ and $H^{\prime}=H$ and $c^{\prime}=c_{2}$. The result follows by E-If ValidFalse

Case T-Apply: $c=t . \operatorname{apply}()$
By Proposition 1, we have $\mathcal{C} \mathcal{A}(t, H)=\left(a_{i}, \bar{v}\right)$. Let $\mathcal{A}\left(a_{i}\right)=\lambda \bar{x}: \bar{\tau} . c_{i}$. Let $I^{\prime}=I$ and $O^{\prime}=O$ and $H^{\prime}=H$ and $c^{\prime}=c_{i}[\bar{v} / \bar{x}]$. The result follows by E-Apply.

Case T-Add: $c=\operatorname{add}(h)$
If $h \in \operatorname{dom}(H)$, let $I^{\prime}=I$ and $O^{\prime}=O$ and $H^{\prime}=H$ and $c^{\prime}=s k i p$. The result follows by E-AddValid. Otherwise, $h \notin \operatorname{dom}(H)$. Let $v=$ init $_{\eta}$ and $I^{\prime}=I$ and $O^{\prime}=O$ and $H^{\prime}=H[h \mapsto v]$ and $c^{\prime}=s k i p$. The result follows by E-ADD

Case T-Remove: $c=$ remove $(h)$
Let $I^{\prime}=I$ and $O^{\prime}=O$ and $H^{\prime}=H \backslash h$ and $c^{\prime}=s k i p$. The result follows by E-Remove.

Case T-MoD: $c=h . f=e$ and Includes $\Theta h$ and $\mathcal{F}(h, f)=\tau_{i}$ and $\cdot ; \Theta \vdash e: \tau_{i}$
By the progress rule for expressions, either $e$ is a value or there is some $e^{\prime}$ such that $\langle H, e\rangle \rightarrow e^{\prime}$.

Subcase $e=v$ : By Lemma A.8: $h \in \operatorname{dom}(H)$. Let $r=H(h)$ and $r^{\prime}=\{r$ with $f=$ $v\}$. Also let $I^{\prime}=I$ and $O^{\prime}=O$ and $H^{\prime}=H\left[h \mapsto r^{\prime}\right]$ and $c^{\prime}=s k i p$. The result follows by E-Mod.
Subcase $\langle H, e\rangle \rightarrow e^{\prime}$ : Let $I^{\prime}=I$ and $O^{\prime}=O$ and $H^{\prime}=H$ and $c^{\prime}=h . f=e^{\prime}$. The result follows by E-Modi.

Case T-Zero: Empty $\Theta_{1}$
By Lemma A.7, we have $\operatorname{dom}(H) \in \llbracket \Theta_{1} \rrbracket$. By Lemma A.5, we have $\llbracket \Theta_{1} \rrbracket=\{ \}$, which is a contradiction.

## Preservation

Lemma A. 10 (Restriction Entailed). If $H \vDash \Theta$ and $h \in \operatorname{dom}(H)$, then $H \vDash$ Restrict $\Theta h$.

Proof. By induction on $\Theta$.
Case $\Theta=0$ : The case immediately holds as $H \vDash \mathrm{o}$ is a contradiction.
Case $\Theta=1$ : By inversion of Entailment, $H=\cdot$. The case immediately holds as $h \in$ $\operatorname{dom}(\cdot)$ is a contradiction.

Case $\Theta=g$ : By inversion of Entailment, dom $(H)=\{g\}$, and so $h=g$. By definition Restrict $\Theta h=$ Restrict $g g=g$. By Ent-Inst $H \vDash g$, i.e., $H \vDash \Theta$.

Case $\Theta=\Theta_{1} \cdot \Theta_{2}$ : By inversion of Entailment $H=H_{1} \cup H_{2}, H_{1} \vDash \Theta_{1}, H_{2} \vDash \Theta_{2}$. By $h \in \operatorname{dom}(H)$, either $h \in \operatorname{dom}\left(H_{1}\right)$ or $h \in \operatorname{dom}\left(H_{2}\right)$.

Subcase $h \in \operatorname{dom}\left(H_{1}\right)$ :
By the induction hypothesis, we have $H_{1} \vDash$ Restrict $\Theta_{1} h$. By EntSeQ, we have $H_{1} \cup H_{2} \vDash$ Restrict $\Theta_{1} h \cdot \Theta_{2}$. By Ent-ChoiceL, we have $H_{1} \cup H_{2} \vDash\left(\right.$ Restrict $\left.\Theta_{1} h \cdot \Theta_{2}\right)+\left(\Theta_{1} \cdot\right.$ Restrict $\left.\Theta_{2} h\right)$ which concludes the case.
Subcase $h \in \operatorname{dom}\left(H_{2}\right)$ : Symmetric to the previous subcase.
Case $\Theta=\Theta_{1}+\Theta_{2}$ : By inversion of Entailment, either $H \vDash \Theta_{1}$ or $H \vDash \Theta_{2}$.
Subcase $H \vDash \Theta_{1}$ : By the induction hypothesis, we have $H \vDash$ Restrict $\Theta_{1} h$. By Ent-ChoiceL, $H \vDash$ Restrict $\Theta_{1} h+$ Restrict $\Theta_{2} h$.
Subcase $H \vDash \Theta_{2}$ : Symmetric to the previous subcase.

Lemma A. 11 (NegRestriction Entailed). If $H \vDash \Theta$ and $h \notin \operatorname{dom}(H)$, then $H \vDash$ NegRestrict $\Theta h$.

Proof. By induction on $\Theta$.
Case $\Theta=0$ : The case immediately holds as $H \vDash \mathrm{o}$ is a contradiction.

Case $\Theta=1$ : By inversion of Entailment, $H=\cdot$. By definition of Negated Restriction, NegRestrict $\Theta h=$ NegRestrict $1 h=1$. By Ent-Empty • $\vDash$ 1, i.e., $H \vDash$ NegRestrict $\Theta h$.

Case $\Theta=g$ : By inversion of Entailment, $\operatorname{dom}(H)=\{g\}$. By assumption $h \neq g$. By definition of Restriction NegRestrict $\Theta h=$ NegRestrict $g h=g$. By EntInst $H \vDash g$, i.e., $H \vDash$ NegRestrict $\Theta h$.

Case $\Theta=\Theta_{1} \cdot \Theta_{2}$ : By inversion of Entailment, $H=H_{1} \cup H_{2}, H_{1} \vDash \Theta_{1}, H_{2} \vDash \Theta_{2}$. By $h \notin \operatorname{dom}(H), h \notin \operatorname{dom}\left(H_{1}\right)$ and $h \notin \operatorname{dom}\left(H_{2}\right)$. By the induction hypothesis, $H_{1} \vDash$ NegRestrict $\Theta_{1} h$ and $H_{2} \vDash$ NegRestrict $\Theta_{2} h$. By Ent-Seq, $H_{1} \cup H_{2} \vDash$ NegRestrict $\Theta_{1} h \cdot$ NegRestrict $\Theta_{2} h$ which finishes the case.

Case $\Theta=\Theta_{1}+\Theta_{2}$ : By inversion of Entailment, either $H \vDash \Theta_{1}$ or $H \vDash \Theta_{2}$.
Subcase $H \vDash \Theta_{1}$ :
By the induction hypothesis, we have $H \vDash$ NegRestrict $\Theta_{1} h$.
By Ent-ChoiceL, $H \vDash$ NegRestrict $\Theta_{1} h+$ NegRestrict $\Theta_{2} h$.
Subcase $H \vDash \Theta_{2}$ : Symmetric to the previous subcase.

Lemma A. 12 (Entailment Congruence). If $H \vDash \Theta$ and $\operatorname{dom}(H)=\operatorname{dom}\left(H^{\prime}\right)$ then $H^{\prime} \vDash \Theta$.

Proof. By induction on $\Theta$.
Case $\Theta=\mathrm{o}$ : The case immediately holds as $H \vDash \mathrm{o}$ is a contradiction.
Case $\Theta=1$ : By inversion of Entailment $H=\cdot \cdot$. By assumption $\operatorname{dom}(H)=\operatorname{dom}\left(H^{\prime}\right)$, $H^{\prime}=\cdot$ and by Ent-Емpty, $H^{\prime} \vDash \Theta$.

Case $\Theta=g$ : By inversion of Entailment, $\operatorname{dom}(H)=\{g\}$. By assumption $\operatorname{dom}(H)=$ $\operatorname{dom}\left(H^{\prime}\right)$ and by Ent-Inst, $H^{\prime} \vDash \Theta$.

Case $\Theta=\Theta_{1} \cdot \Theta_{2}$ : By inversion of Entailment, $H=H_{1} \cup H_{2}, H_{1} \vDash \Theta_{1}, H_{2} \vDash \Theta_{2}$. By set theory, $\operatorname{dom}(H)=\operatorname{dom}\left(H_{1}\right) \cup \operatorname{dom}\left(H_{2}\right)$. By induction hypothesis if $\operatorname{dom}\left(H_{1}^{\prime}\right)=$ $\operatorname{dom}\left(H_{1}\right)$ and $\operatorname{dom}\left(H_{2}^{\prime}\right)=\operatorname{dom}\left(H_{2}\right)$, then $H_{1}^{\prime} \vDash \Theta_{1}$, and $H_{2}^{\prime} \vDash \Theta_{2}$. By Ent-Seq, $H^{\prime}=H_{1}^{\prime} \cup H_{2}^{\prime} \vDash \Theta_{1} \cdot \Theta_{2}$.

Case $\Theta=\Theta_{1}+\Theta_{2}$ : By inversion of Entailment, either $H \vDash \Theta_{1}$ or $H \vDash \Theta_{2}$.
Subcase $H \vDash \Theta_{1}$ : By induction hypothesis, we have $H^{\prime} \vDash \Theta_{1}$. By Ent-ChoiceL, $H^{\prime} \vDash \Theta_{1}+\Theta_{2}$.
Subcase $H \vDash \Theta_{2}$ : Symmetric to the previous subcase.

Lemma A. 13 (Substitution). If $\Gamma, x: \tau \vdash c: \Theta \Leftrightarrow \Theta^{\prime}$ and $\because \cdot \vdash v: \tau$ then $\Gamma \vdash c[v / x]$ : $\Theta \Leftrightarrow \Theta^{\prime}$.

Proof. By straightforward induction on the derivation $\Gamma, x: \tau \vdash c: \Theta \Leftrightarrow \Theta^{\prime}$.
Lemma A. 14 (Entails Subsumption). If $H \vDash \Theta$ then $H[h \mapsto v] \vDash \Theta \cdot h$

Proof. We analyze two cases.
Case $h \in \operatorname{dom}(H)$ : By the assumption of the case, we have $\operatorname{dom}(H)=\operatorname{dom}(H[h \mapsto v])$. Let $H_{1}=H[h \mapsto v]$ and $H_{2}=\{h \mapsto v\}$. Observe that $H[h \mapsto v]=H_{1} \cup H_{2}$. By Lemma A.12, we have that $H_{1} \vDash \Theta$. By Ent-Inst we have $H_{2} \vDash h$. By Ent-Seq we have $H[h \mapsto v] \vDash \Theta \cdot h$.

Case $h \notin \operatorname{dom}(H)$ : Let $H_{1}=H$ and $H_{2}=\{h \mapsto v\}$. Observe that $H[h \mapsto v]=H_{1} \cup H_{2}$. By assumption we have $H_{1} \vDash \Theta$. By Ent-Inst we have $H_{2} \vDash h$. By Ent-Seq we have $H[h \mapsto v] \vDash \Theta \cdot h$.

Lemma A. 15 (Entails Removal). If $H \vDash \Theta$ then $H \backslash h \vDash$ Remove $\Theta h$.
Proof. By induction on $\Theta$.
Case $\Theta=0$ : The case immediately holds, as $H \vDash \mathrm{o}$ is a contradiction.
Case $\Theta=1$ : By inversion of Entailment, $H=\cdot \cdot$. By set theory, $\bullet \backslash h=\cdot$ and Remove $1 h=1$. By Ent-Empty, ・ト 1.

Case $\Theta=g$ : By inversion of Entailment, $\operatorname{dom}(H)=\{g\}$.
Subcase $g=h$ : By set theory $H \backslash h=\cdot$. By definition of Remove, Remove $\Theta h=1$. By Еnt-Емрту, $\cdot \vDash 1$, which concludes the case.

Subcase $g \neq h$ : By set theory $H \backslash h=H$. By definition of Remove, Remove $\Theta h=$ $g$. By assumption, $H \vDash \Theta$, which concludes the case.

Case $\Theta=\Theta_{1} \cdot \Theta_{2}$ : By inversion of Entailment, $H=H_{1} \cup H_{2}, H_{1} \vDash \Theta_{1}, H_{2} \vDash \Theta_{2}$. By induction hypothesis, $H_{1} \backslash h \vDash$ Remove $\Theta_{1} h$ and $H_{2} \backslash h \vDash$ Remove $\Theta_{2} h$. By set theory, $H_{1} \backslash h \cup H_{2} \backslash h=\left(H_{1} \cup H_{2}\right) \backslash h$. By definition of Removal, Remove $\Theta_{1} h \cdot \operatorname{Remove} \Theta_{2} h=\operatorname{Remove}\left(\Theta_{1} \cdot \Theta_{2}\right) h$. By Ent-SeQ, $\left(H_{1} \cup H_{2}\right) \backslash h \vDash$ Remove $\left(\Theta_{1} \cdot \Theta_{2}\right) h$.

Case $\Theta=\Theta_{1}+\Theta_{2}$ : By inversion of Entailment, either $H \vDash \Theta_{1}$ or $H \vDash \Theta_{2}$.

Subcase $H \vDash \Theta_{1}$ : By induction hypothesis, $H \backslash h \vDash$ Remove $\Theta_{1} h$. By EntChoicel, applied to $H \backslash h \vDash$ Remove $\Theta_{1} h$, and Remove $\Theta_{2} h$, we can conclude $H \backslash h \vDash$ Remove $\Theta_{1} h+$ Remove $\Theta_{2} h$. By definition of Removal, Remove $\Theta_{1} h+$ Remove $\Theta_{2} h=\operatorname{Remove}\left(\Theta_{1}+\Theta_{2}\right) h$.
Subcase $H \vDash \Theta_{2}$ : Symmetric to previous subcase.

Lemma A. 16 (Order Extend). If $\Theta_{1}^{\prime}<\Theta_{1}$ then $\Theta_{1}^{\prime} \cdot h<\Theta_{1} \cdot h$.
Proof. By assumption $\Theta_{1}^{\prime}<\Theta_{1}$ and the definition of $<$ follows $\llbracket \Theta_{1}^{\prime} \rrbracket \subseteq \llbracket \Theta_{1} \rrbracket$ By $\llbracket . \rrbracket$ follows $\llbracket \Theta_{1}^{\prime} \cdot h \rrbracket==\llbracket \Theta_{1}^{\prime} \rrbracket \bullet\{\{h\}\}$ and $\llbracket \Theta_{1} \cdot h \rrbracket==\llbracket \Theta_{1} \rrbracket \bullet\{\{h\}\}$.

Let $S \in \llbracket \Theta_{1}^{\prime} \rrbracket \bullet\{\{h\}\}$. By definition of $\bullet S=S^{\prime} \cup\{h\}$, where $S^{\prime} \in \llbracket \Theta_{1}^{\prime} \rrbracket$. By $\llbracket \Theta_{1}^{\prime} \rrbracket \subseteq \llbracket \Theta_{1} \rrbracket$, follows that $S^{\prime} \in \llbracket \Theta_{1} \rrbracket$. By set theory, $S^{\prime} \cup\{h\} \in \llbracket \Theta_{1} \rrbracket \bullet\{\{h\}\}$. Then $\llbracket \Theta_{1}^{\prime} \rrbracket \bullet\{\{h\}\} \subseteq \llbracket \Theta_{1} \rrbracket \bullet\{\{h\}\}$.

Lemma A. 17 (Order Remove). If $\Theta_{1}^{\prime}<\Theta_{1}$ then $\llbracket$ Remove $\Theta_{1}^{\prime} h \rrbracket \subseteq \llbracket$ Remove $\Theta_{1} h \rrbracket$.
Proof. Since $\llbracket$ Remove $\Theta_{1}^{\prime} h \rrbracket==\llbracket \Theta_{1}^{\prime} \rrbracket \backslash h$ and $\llbracket$ Remove $\Theta_{1} h \rrbracket==\llbracket \Theta_{1} \rrbracket \backslash h$ by Lemma A.4, we can equivalently show that $\llbracket \Theta_{1}^{\prime} \rrbracket \backslash h \subseteq \llbracket \Theta_{1} \rrbracket \backslash h$, which follows from set theory.

Lemma A. 18 (Order Restrict). If $\Theta_{1}^{\prime}<\Theta_{1}$ then $\llbracket$ Restrict $\Theta_{1}^{\prime} h \rrbracket \subseteq \llbracket$ Restrict $\Theta_{1} h \rrbracket$
Proof. By Lemma 4.1, $\llbracket$ Restrict $\Theta_{1}^{\prime} h \rrbracket==\llbracket \Theta_{1}^{\prime} \rrbracket \mid h$ and $\llbracket$ Restrict $\Theta_{1} h \rrbracket==\llbracket \Theta_{1} \rrbracket \mid h$. By set theory, $\llbracket \Theta_{1}^{\prime} \rrbracket\left|h \subseteq \llbracket \Theta_{1} \rrbracket\right| h$ when $\llbracket \Theta_{1}^{\prime} \rrbracket \subseteq \llbracket \Theta_{1} \rrbracket$, so we are done.

Lemma A. 19 (Order NegRestrict). If $\Theta_{1}^{\prime}<\Theta_{1}$ then
$\llbracket$ NegRestrict $\Theta_{1}^{\prime} h \rrbracket \subseteq \llbracket$ NegRestrict $\Theta_{1} h \rrbracket$
Proof. By Lemma 4.2, $\llbracket$ NegRestrict $\Theta_{1}^{\prime} h \rrbracket==\llbracket \Theta_{1}^{\prime} \rrbracket \mid \neg h$ and $\llbracket$ NegRestrict $\Theta_{1} h \rrbracket==$ $\llbracket \Theta_{1} \rrbracket \mid \neg h$. By set theory, $\llbracket \Theta_{1}^{\prime} \rrbracket\left|\neg h \subseteq \llbracket \Theta_{1} \rrbracket\right| \neg h$ when $\llbracket \Theta_{1}^{\prime} \rrbracket \subseteq \llbracket \Theta_{1} \rrbracket$, so we are done.

Lemma A. 20 (Order Include). If $\Theta_{1}^{\prime}<\Theta_{1}$ and Includes $\Theta_{1} h$ then Includes $\Theta_{1}^{\prime} h$.
Proof. By Lemma A.3, Includes $\Theta_{1}^{\prime} h=h \sqsubset \Theta_{1}^{\prime}$. By the same lemma, Includes $\Theta_{1} h=$ $h \sqsubset \Theta_{1}$. Let $S \in \llbracket \Theta_{1}^{\prime} \rrbracket$ to show $h \in S$ and hence $h \check{ } \Theta_{1}^{\prime}$. Since $\llbracket \Theta_{1}^{\prime} \rrbracket \subseteq \llbracket \Theta_{1} \rrbracket$, by assumption and definition of $<$, then $S \in \llbracket \Theta_{1} \rrbracket$. Since $h \sqsubset \Theta_{1}$, conclude $h \in S$ and we are done.

Lemma A. 21 (Order Empty). If $\Theta_{1}^{\prime}<\Theta_{1}$ and Empty $\Theta_{1}$ then Empty $\Theta_{1}^{\prime}$.
Proof. By definition of $<, \llbracket \Theta_{1}^{\prime} \rrbracket \subseteq \llbracket \Theta_{1} \rrbracket$. By Lemma A. 5 and assumption Empty $\Theta_{1}$ follows $\llbracket \Theta_{1} \rrbracket=\{ \}$. By set theory $\llbracket \Theta_{1}^{\prime} \rrbracket=\{ \}$. The result follows by Lemma A. 5 .

Lemma A. 22 (Order Choice). If $\Theta_{a}^{\prime}<\Theta_{a}$ and $\Theta_{b}^{\prime}<\Theta_{b}$ then $\Theta_{a}^{\prime}+\Theta_{b}^{\prime}<\Theta_{a}+\Theta_{b}$.
Proof. We have to show that $\llbracket \Theta_{a}^{\prime}+\Theta_{b}^{\prime} \rrbracket \subseteq \llbracket \Theta_{a}+\Theta_{b} \rrbracket$ when $\Theta_{a}^{\prime}<\Theta_{a}$ and $\Theta_{b}^{\prime}<\Theta_{b}$. By definition of $\llbracket . \rrbracket$ we can equally show that $\llbracket \Theta_{a}^{\prime} \rrbracket \cup \llbracket \Theta_{b}^{\prime} \rrbracket \subseteq \llbracket \Theta_{a} \rrbracket \cup \llbracket \Theta_{b} \rrbracket$, which follows from set theory.

Lemma A. 23 (Expression Type Bounds). If $\Gamma ; \Theta \vdash e: \tau$ and $\Theta^{\prime}<\Theta$, then $\Gamma ; \Theta^{\prime} \vdash e: \tau$.
Proof. By induction on the typing derivation.
Case T-Constant: We know $e=k(\bar{e})$, and $\Gamma ; \Theta \vdash e_{i}: \tau_{i}$ for all $i$, and typeof $(k)=$ $\bar{\tau} \rightarrow \tau$ and $\Theta^{\prime}<\Theta$. By induction hypothesis, $\Gamma ; \Theta^{\prime} \vdash e_{i}: \tau_{i}$ for all $i$ and we are done by T-Constant.

Case T-Var: We know $e=x$, and $x: \tau \in \Gamma$, and $\Theta^{\prime}<\Theta$. We are done by T-Var.
Case T-Field: We know $e=h . f$, and Includes $\Theta h$ and $\Theta^{\prime}<\Theta$. By Lemma A.2o, we know Includes $\Theta^{\prime} \mathrm{h}$ and the result follows by T-Field.

Lemma A. 24 (Action Type Bounds). If $\Gamma ; \Theta_{1} \vdash a: \bar{\tau} \rightarrow \Theta_{2}$ and $\Theta_{1}^{\prime}<\Theta_{1}$, then $\exists \Theta_{2}^{\prime} \cdot \Gamma ; \Theta_{1}^{\prime} \vdash a: \bar{\tau} \rightarrow \Theta_{2}^{\prime}$ and $\Theta_{2}^{\prime}<\Theta_{2}$.

Proof. Rule T-Action is the only rule by which we have concluded that $\Gamma ; \Theta_{1} \vdash a: \bar{\tau} \rightarrow$ $\Theta_{2}$. This rule gives us two facts: we know $a=\lambda \bar{x}: \bar{\tau} . c$, and $\Gamma, \bar{x}: \bar{\tau} \vdash c: \Theta_{1} \mapsto \Theta_{2}$.

Since this $c$ is an action command, is only generated by the add, remove, modification and sequence commands. So we perform a limited induction on the structure of $c$ :

Case $c=\operatorname{add}(h)$ : The only typing rule that applies is T-ADd, so we know $\Theta_{2}=\Theta_{1} \cdot h$. Now let $\Theta_{2}^{\prime}=\Theta_{1}^{\prime} \cdot h$. Then T-Add shows $\Gamma, \bar{x}: \bar{\tau} \vdash \operatorname{add}(h): \Theta_{1}^{\prime} \Leftrightarrow \Theta_{1}^{\prime} \cdot h$ and $\Theta_{1}^{\prime} \cdot h<\Theta_{1} \cdot h$ follows by Lemma A.16, and we are done.

Case $c=\operatorname{remove}(h)$ : The only typing rule that could have applied is T-Remove, so we know that $\Theta_{2}=$ Remove $\Theta_{1} h$. Let $\Theta_{2}^{\prime}=$ Remove $\Theta_{1}^{\prime} h$. Then T-Remove shows $\Gamma, \bar{x}: \bar{\tau} \vdash \operatorname{remove}(h): \Theta_{1}^{\prime} \Leftrightarrow$ Remove $\Theta_{1} h$ and Remove $\Theta_{1}^{\prime} h<$ Remove $\Theta_{1} h$ follows by Lemma A. 17 and by definition of $<$.

Case $c=h . f=v$ : The only typing rule that could have applied is T-Mod, so we know that $\Theta_{2}=\Theta_{1}$. Let $\Theta_{2}^{\prime}=\Theta_{1}^{\prime}$, which proves $\Theta_{2}^{\prime}<\Theta_{2}$ by assumption. We know by our case assumption that $\Gamma, \bar{x}: \bar{\tau} ; \Theta_{1} \vdash e: \mathcal{F}(h, f)$ and Includes $\Theta_{1} h$.
By T-Mod, we only need to show that
(1) $\Gamma, \bar{x}: \bar{\tau} ; \Theta_{1}^{\prime} \vdash e: \mathcal{F}(h, f)$, which follows by Lemma A. 23 and
(2) Includes $\Theta_{1}^{\prime} h$, which follows by Lemma A.20.

Case $c=c_{1} ; c_{2}$ : The only rule that could have applied is T-SEQ, so we know that $\Gamma, \bar{x}$ : $\bar{\tau} \vdash c_{1}: \Theta_{1} \Leftrightarrow \Theta_{11}$, and $\Gamma, \bar{x}: \bar{\tau} \vdash c_{2}: \Theta_{11} \Leftrightarrow \Theta_{2}$.
The inductive hypothesis on $c_{1}$ gives us a $\Theta_{11}^{\prime}<\Theta_{11}$ such that $\Gamma, \bar{x}: \bar{\tau} \vdash c_{1}: \Theta_{1}^{\prime} \Leftrightarrow$ $\Theta_{11}^{\prime}$.
The inductive hypothesis on $c_{2}$ gives us a $\Theta_{2}^{\prime}<\Theta_{2}$ such that $\Gamma, \bar{x}: \bar{\tau} \vdash c_{2}: \Theta_{11}^{\prime} \mapsto$ $\Theta_{2}^{\prime}$. The result follows by T-SEQ.

Lemma A. 25 (Control Type Bounds). If $\Gamma \vdash c: \Theta_{1} \Leftrightarrow \Theta_{2}$ and $\Theta_{1}^{\prime}<\Theta_{1}$, then $\exists \Theta_{2}^{\prime}$. $\Gamma \vdash$ $c: \Theta_{1}^{\prime} \Leftrightarrow \Theta_{2}^{\prime}$ and $\Theta_{2}^{\prime}<\Theta_{2}$.

Proof. By induction on a derivation of $\Gamma \vdash c: \Theta_{1} \Leftrightarrow \Theta_{2}$. We refer to the general assumptions as follows:
(A) $\Gamma \vdash c: \Theta_{1} \Leftrightarrow \Theta_{2}$ and
(B) $\Theta_{1}^{\prime}<\Theta_{1}$

Similarly, we refer to the proof goals as
(1) $\exists \Theta_{2}^{\prime} \cdot \Gamma \vdash c: \Theta_{1}^{\prime} \Leftrightarrow \Theta_{2}^{\prime}$
(2) $\Theta_{2}^{\prime}<\Theta_{2}$

Case T-Zero: We know Empty $\Theta_{1}$. By assumption (B) and Lemma A. 21 we have Empty $\Theta_{1}^{\prime}$. Let $\Theta_{2}^{\prime}=\Theta_{2}$. We have $\Gamma \vdash c: \Theta_{1}^{\prime} \Leftrightarrow \Theta_{2}$ by T-Zero, proving ( 1 ), and and $\Theta_{2}^{\prime}<\Theta_{2}$ by reflexivity, proving (2).

Case T-Skip: We know $c=s k i p$ and $\Theta_{2}=\Theta_{1}$. Let $\Theta_{2}^{\prime}=\Theta_{1}^{\prime}$. Then by assumption (B) $\Theta_{2}^{\prime}=\Theta_{1}^{\prime}<\Theta_{1}=\Theta_{2}$, proving (2) and $\Gamma \vdash$ skip: $\Theta_{1}^{\prime} \Leftrightarrow \Theta_{1}^{\prime}$ by T-SKIP, proving (1).

Case T-Emit: We know $c=\operatorname{emit}(h)$ and $\Theta_{2}=\Theta_{1}$. Let $\Theta_{2}^{\prime}=\Theta_{1}^{\prime}$. Then by assumption (B), $\Theta_{2}^{\prime}=\Theta_{1}^{\prime}<\Theta_{1}=\Theta_{2}$, proving (2) and $\Gamma \vdash \operatorname{emit}(h): \Theta_{1}^{\prime} \Leftrightarrow \Theta_{1}^{\prime}$ by T-Eміт, proving (1).

Case T-ADD: We know $c=\operatorname{add}(h)$ and $\Theta_{2}=\Theta_{1} \cdot h$. Let $\Theta_{2}^{\prime}=\Theta_{1}^{\prime} \cdot h$. (1) follows since we can prove $\Gamma \vdash \operatorname{add}(h): \Theta_{1}^{\prime} \Leftrightarrow \Theta_{1}^{\prime} \cdot h$ by T-AdD. (2), i.e., $\Theta_{1}^{\prime} \cdot h<\Theta_{1} \cdot h$, follows from Lemma A.16.

Case T-Extr: We know $c=\operatorname{extract}(h)$ and $\Theta_{2}=\Theta_{1} \cdot h$. Let $\Theta_{2}^{\prime}=\Theta_{1}^{\prime} \cdot h$. (1) follows since we can prove $\Gamma \vdash \operatorname{extract}(h): \Theta_{1}^{\prime} \Leftrightarrow \Theta_{1}^{\prime} \cdot h$ by T-Extract. (2), i.e., $\Theta_{1}^{\prime} \cdot h<\Theta_{1} \cdot h$, follows from Lemma A.16.

Case T-Rem: We know $c=\operatorname{remove}(h)$ and $\Theta_{2}=$ Remove $\Theta_{1} h$. Let $\Theta_{2}^{\prime}=$ Remove $\Theta_{1}^{\prime} h$. (1) follows by T-Rem. For (2) we have to show that Remove $\Theta_{1}^{\prime} h<$ Remove $\Theta_{1} h$, which follows from Lemma A.17.

Case T-MoD: We know $c=h . f=e$ and $\Theta_{2}=\Theta_{1}$ and Includes $\Theta_{1} h$, and $\mathcal{F}(h, f)=$ $\tau$ and $\Gamma ; \Theta_{1} \vdash e: \tau$. Let $\Theta_{2}^{\prime}=\Theta_{1}^{\prime}$. (1) follows by T-MoD, if we can show Includes $\Theta_{1}^{\prime} h$, which follows by assumption (B) and Lemma A.20. (2) follows by assumption (B).

Case T-SEQ: We know $c=c_{1} ; c_{2}$ and $\Gamma \vdash c_{1}: \Theta_{1} \Leftrightarrow \Theta_{11}$ and $\Gamma \vdash c_{2}: \Theta_{11} \Leftrightarrow \Theta_{2}$. By induction hypothesis, $\exists \Theta_{11}^{\prime} . \Gamma \vdash c_{1}: \Theta_{1}^{\prime} \Leftrightarrow \Theta_{11}^{\prime}$ and $\Theta_{11}^{\prime}<\Theta_{11}$. Again, by induction hypothesis, $\exists \Theta_{2}^{\prime} . \Gamma \vdash c_{2}: \Theta_{11}^{\prime} \mapsto \Theta_{2}^{\prime}$ and $\Theta_{2}^{\prime}<\Theta_{2}$ (proving (2)). (1) follows by T-SEQ, which concludes the case.

Case T-If Valid: We know $c=\operatorname{valid}(h) c_{1}$ else $c_{2}$ and $\Gamma \vdash c_{1}$ : Restrict $\Theta_{1} h \mapsto$ $\Theta_{t}, \Gamma \vdash c_{2}$ : NegRestrict $\Theta_{1} h \Leftrightarrow \Theta_{f}, \Theta_{2}=\Theta_{t}+\Theta_{f}$. Let $\Theta_{2}^{\prime}=$ Restrict $\Theta_{1}^{\prime} h+$ NegRestrict $\Theta_{1}^{\prime} h$. (1) is immediate from T-IF Valid. (2) follows from Lemmas A.18, A. 19 and A. 22 .

Case T-IF: We know $c=$ if $(e) c_{1}$ else $c_{2}$, and $\Gamma \vdash c_{1}: \Theta_{1} \Leftrightarrow \Theta_{11}$, and $\Gamma \vdash c_{2}: \Theta_{1} \Leftrightarrow \Theta_{12}$, and $\Gamma ; \Theta_{1} \vdash e$ : Bool.
By induction hypothesis, there exists $\Theta_{11}^{\prime}$ such that (1a) $\Gamma \vdash c_{1}: \Theta_{1}^{\prime} \Leftrightarrow \Theta_{11}^{\prime}$ and (2a) $\Theta_{11}^{\prime}<\Theta_{11}$. Also by induction hypothesis, there exists $\Theta_{12}^{\prime}$ such that (1b) $\Gamma \vdash c_{2}: \Theta_{1}^{\prime} \Leftrightarrow \Theta_{12}^{\prime}$ and (2b) $\Theta_{12}^{\prime}<\Theta_{12}$. Let $\Theta_{2}^{\prime}=\Theta_{11}^{\prime}+\Theta_{12}^{\prime}$. By Lemma A. 23 follows $\Gamma ; \Theta_{1}^{\prime} \vdash e$ : Bool. (1) then follows by T-IF with (1a) and (1b). (2) follows by Lemma A. 22 with (2a) and (2b).

Case T-Apply: We know

- $c=t . \operatorname{apply}()$
- $\Theta_{2}=\Theta_{11}+\Theta_{12}+\ldots+\Theta_{1 n}$
- t.actions $=a_{1}+a_{2}+\ldots+a_{n}$
- ; $\Theta_{1} \vdash e_{j}: \tau_{j}$ for $j=1, \ldots, m$
- $\mathcal{C} \mathcal{V}(t)=\left(S_{1} \ldots S_{n}\right)$
- $\left(e_{1} \ldots e_{m}\right)=\left\{e_{i} \mid\left(e_{i}, m_{i}\right) \in\right.$ t.reads ()$\wedge \neg$ maskable $\left.\left(t, e_{i}, m_{i}\right)\right\}$
- ; Restrict $\Theta_{1} S_{i} \vdash a_{i}: \bar{\tau}_{i} \rightarrow \Theta_{1 i}$

We want to construct $\Theta_{2}^{\prime}<\Theta_{2}$ such that $\Gamma \vdash t$.apply ()$: \Theta_{1}^{\prime} \mapsto \Theta_{2}^{\prime}$. By repeated application of Lemma A.18, Restrict $\Theta_{1}^{\prime} S_{i}<\operatorname{Restrict} \Theta_{1} S_{i}$. For every $i$ apply Lemma A.24, which gives us $\Gamma$; Restrict $\Theta_{1}^{\prime} S_{i} \vdash a: \bar{\tau} \rightarrow \Theta_{1_{i}}$ and $\Theta_{1 i}^{\prime}<\Theta_{1 i}$. Also for every $j$ apply Lemma A.23, which gives us $\cdot ; \Theta_{1}^{\prime} \vdash e_{j}: \tau_{j}$ for $j=1, \ldots, m$ Let $\Theta_{2}^{\prime}=\sum_{i} \Theta_{1 i}^{\prime}$. (2) follows by T-Apply. To show (1), i.e., $\Theta_{2}^{\prime}=\sum_{i} \Theta_{1 i}^{\prime}<\sum_{i} \Theta_{1 i}=$ $\Theta_{2}$. We know $\Theta_{1 i}^{\prime}<\Theta_{1 i}$ for all $i$. The result follows by repeated application of Lemma A. 22 .

Theorem A. 26 (Preservation). If $\Gamma \vdash c: \Theta_{1} \Leftrightarrow \Theta_{2}$ and $\langle I, O, H, c\rangle \rightarrow\left\langle I^{\prime}, O^{\prime}, H^{\prime}, c^{\prime}\right\rangle$, where $H \vDash \Theta_{1}$, then $\exists \Theta_{1}^{\prime}, \Theta_{2}^{\prime}$. $\Gamma \vdash c^{\prime}: \Theta_{1}^{\prime} \Leftrightarrow \Theta_{2}^{\prime}$ where $H^{\prime} \vDash \Theta_{1}^{\prime}$ and $\Theta_{2}^{\prime}<\Theta_{2}$.

Proof. By induction on a derivation of $\Gamma \vdash c: \Theta_{1} \Leftrightarrow \Theta_{2}$, with a case analysis on the last rule used.

Case T-Skip: $c=s k i p$ and $\Theta_{2}=\Theta_{1}$
Vacuously holds as there is no $c^{\prime}$ such that $\langle I, O, H, c\rangle \rightarrow\left\langle I^{\prime}, O^{\prime}, H^{\prime}, c^{\prime}\right\rangle$.
Case T-ExTR: $c=\operatorname{extract}(h)$ and $\Theta_{2}=\Theta_{1} \cdot h$
The only evaluation rule that applies to $c$ is E-Extr, so we also have $c^{\prime}=s k i p$ and $\mathcal{H} \mathcal{T}(h)=\eta$ and $H^{\prime}=H[h \mapsto v]$ where deserialize ${ }_{\eta}(I)=\left(v, I^{\prime}\right)$.
Let $\Theta_{1}^{\prime}=\Theta_{2}^{\prime}=\Theta_{2}$. We have $\Gamma \vdash c^{\prime}: \Theta_{1}^{\prime} \Leftrightarrow \Theta_{2}^{\prime}$ by T-SKIP, we have $H^{\prime} \vDash \Theta_{2}^{\prime}$ by Lemma A.14, and we have $\Theta_{2}^{\prime}<\Theta_{2}$ by reflexivity.

Case T-Emit: $c=e m i t(h)$ and $\Theta_{2}=\Theta_{1}$.
There are two evaluation rules that apply to $c$, E-Emit and E-EmitInvalid. In either case, $c^{\prime}=$ skip and $H^{\prime}=H$.
Let $\Theta_{1}^{\prime}=\Theta_{2}^{\prime}=\Theta_{1}$. We have $\Gamma \vdash c^{\prime}: \Theta_{1}^{\prime} \Leftrightarrow \Theta_{2}^{\prime}$ by T-Skip, we have $H^{\prime} \vDash \Theta_{1}^{\prime}$ by assumption, and we have $\Theta_{2}^{\prime}<\Theta_{2}$ by reflexivity.

Case T-SEQ: $c=c_{1} ; c_{2}$ and $\Gamma \vdash c_{1}: \Theta_{1} \Leftrightarrow \Theta_{12}$ and $\Gamma \vdash c_{2}: \Theta_{12} \Leftrightarrow \Theta_{2}$
There are two evaluation rules that apply to $c$, E-SEQ1 and E-SEQ.
Subcase E-Seq: $c^{\prime}=c_{2}$ and $H^{\prime}=H$
By inversion of $\Gamma \vdash c_{1}: \Theta_{1} \Leftrightarrow \Theta_{12}$ we have $\Theta_{12}=\Theta_{1}$. Let $\Theta_{1}^{\prime}=\Theta_{1}$ and $\Theta_{2}^{\prime}=\Theta_{2}$. We have $\Gamma \vdash c^{\prime}: \Theta_{1}^{\prime} \Leftrightarrow \Theta_{2}^{\prime}$ by assumption, we have $H \vDash \Theta_{1}^{\prime}$ also by assumption, and $\Theta_{2}^{\prime}<\Theta_{2}$ by reflexivity.
Subcase E-SEQ1: $c^{\prime}=c_{1}^{\prime} ; c_{2}$ and $\left\langle I, O, H, c_{1}\right\rangle \rightarrow\left\langle I^{\prime}, O^{\prime}, H^{\prime}, c_{1}^{\prime}\right\rangle$.
By IH we have $\Gamma \vdash c_{1}: \Theta_{1}^{\prime} \Leftrightarrow \Theta_{12}^{\prime}$ such that $H^{\prime} \vDash \Theta_{1}^{\prime}$ and $\Theta_{12}^{\prime}<\Theta_{12}$. By Lemma A. 25 we have $\Gamma \vdash c_{2}: \Theta_{12}^{\prime} \Leftrightarrow \Theta_{2}^{\prime}$ for some $\Theta_{2}^{\prime}<\Theta_{2}$. We have $\Gamma \vdash c_{1} ; c_{2}: \Theta_{1}^{\prime} \Leftrightarrow \Theta_{2}^{\prime}$ by T-SEQ, which finishes the case.

Case T-IF: $c=$ if $(e) c_{1}$ else $c_{2}$ and $\Gamma ; \Theta_{1} \vdash e:$ Bool and $\Gamma \vdash c_{1}: \Theta_{1} \mapsto \Theta_{12}$ and $\Gamma \vdash c_{2}: \Theta_{1} \Leftrightarrow \Theta_{22}$ and $\Theta_{2}=\Theta_{12}+\Theta_{22}$.
There are three evaluation rules that apply to $c$, E-If, E-If True, and E-If False.
Subcase E-If: $c^{\prime}=$ if $\left(e^{\prime}\right) c_{1}$ else $c_{2}$ and $H^{\prime}=H$
Let $\Theta_{1}^{\prime}=\Theta_{1}$ and $\Theta_{2}^{\prime}=\Theta_{2}$. We have $\Gamma \vdash$ if $(e) c_{1}$ else $c_{2}: \Theta_{1}^{\prime} \Leftrightarrow \Theta_{2}^{\prime}$ by T-IF, we have $H \vDash \Theta_{1}$ by assumption, and we have $\Theta_{2}<\Theta_{2}^{\prime}$ by reflexivity.
Subcase E-If True: $c^{\prime}=c_{1}$ and $H^{\prime}=H$.
Let $\Theta_{1}^{\prime}=\Theta_{1}$ and $\Theta_{2}^{\prime}=\Theta_{12}$. We have $\Gamma \vdash c^{\prime}: \Theta_{1}^{\prime} \Leftrightarrow \Theta_{2}^{\prime}$ by assumption, we have $H \vDash \Theta_{1}^{\prime}$ also by assumption, and we have $\Theta_{2}^{\prime}<\Theta_{2}$ by the definition of < and the semantics of types.
Subcase E-If False: $c^{\prime}=c_{2}$ and $H^{\prime}=H$.
Symmetric to the previous case.
Case T-If Valid: $c=\operatorname{valid}(h) c_{1}$ else $c_{2}$ and $\Gamma \vdash c_{1}$ : Restrict $\Theta_{1} h \Leftrightarrow \Theta_{12}$ and $\Gamma \vdash c_{2}$ : NegRestrict $\Theta_{1} h \Leftrightarrow \Theta_{22}$ and $\Theta_{2}=\Theta_{12}+\Theta_{22}$. There are two evaluation rules that apply to $c$, E-If ValidTrue and E-If ValidFalse

Subcase E-If ValidTrue: $c^{\prime}=c_{1}$ and $h \in \operatorname{dom}(H)$ and $H^{\prime}=H$.
Let $\Theta_{1}^{\prime}=$ Restrict $\Theta_{1} h$ and $\Theta_{2}^{\prime}=\Theta_{12}$. We have $\Gamma \vdash c^{\prime}: \Theta_{1}^{\prime} \Leftrightarrow \Theta_{2}^{\prime}$ by assumption, we have $H \vDash \Theta_{1}^{\prime}$ by Lemma A.10, and we have $\Theta_{2}^{\prime}<\Theta_{2}$ by the definition of < and semantics of types.
Subcase E-If ValidFalse: $c^{\prime}=c_{2}$ and $h \notin \operatorname{dom}(H)$ and $H^{\prime}=H$.
Symmetric to the previous case.
Case T-Apply: $c=t . a p p l y()$ and $\mathcal{C} \mathcal{V}(t)=\bar{S}$ and t.actions $=\bar{a}$ and $\bar{e}=\left\{e_{j} \mid\left(e_{j}, m_{j}\right) \in\right.$ $t . \operatorname{reads}() \wedge \neg$ maskable $\left.\left(t, e_{j}, m_{j}\right)\right\}$ and $; \Theta \vdash e_{i}: \tau_{i}$ for $e_{i} \in \bar{e}$ and $\Theta_{2}=\sum\left(\Theta_{i}^{\prime}\right)$ and $\cdot ;$ Restrict $\Theta_{1} S_{i} \vdash a_{i}: \bar{\tau}_{i} \rightarrow \Theta_{i}^{\prime}$ for $a_{i} \in a$
There is only one evaluation rule that applies to $c$, E-Apply.
It follows that $\mathcal{C} \mathcal{A}(t, H)=\left(a_{i}, \bar{v}\right)$, and $c^{\prime}=c_{i}[\bar{v} / \bar{x}]$ where $\mathcal{A}\left(a_{i}\right)=\lambda \bar{x} . c_{i}$. Next, inverting T-Action, we have $\cdot \bar{x}: \bar{\tau}_{i} \vdash c_{i}$ : Restrict $\Theta S_{i} \Leftrightarrow \Theta_{i}^{\prime}$. By Proposition 2, we have $\because \cdot \vdash \bar{v}: \bar{\tau}_{i}$. Hence, by Lemma A.13, we have $\cdot \vdash c_{i}[\bar{v} / \bar{x}]$ : Restrict $\Theta S_{i} \Leftrightarrow \Theta_{i}^{\prime}$.
Let $\Theta_{1}^{\prime}=$ Restrict $\Theta S_{i}$ and $\Theta_{2}^{\prime}=\Theta_{i}^{\prime}$. We have already shown that $\Gamma \vdash c^{\prime}$ : $\Theta_{1}^{\prime} \Leftrightarrow \Theta_{2}^{\prime}$. We have that $H^{\prime} \vDash \Theta_{1}^{\prime}$ by Proposition 3, and we have $\Theta_{2}^{\prime}<\Theta_{2}$ by the definition of < and the semantics of union types.

Case T-Add: $c=\operatorname{add}(h)$ and $\Theta_{2}=\Theta_{1} \cdot h$
There are two evaluation rules that apply to $c$, E-Add and E-AddValid.
Subcase E-ADD: $c^{\prime}=s k i p$ and $\mathcal{H} \mathcal{T}(h)=\eta$ and init $_{\eta}=v$ and $H^{\prime}=H[h \mapsto v]$ Let $\Theta_{1}^{\prime}=\Theta_{2}^{\prime}=\Theta_{2}$. We have $\Gamma \vdash c^{\prime}: \Theta_{1}^{\prime} \Leftrightarrow \Theta_{2}^{\prime}$ by T-Skip, we have $H^{\prime} \vDash \Theta_{1}^{\prime}$ by Lemma A.14, and we have $\Theta_{2}^{\prime}<\Theta_{2}$ by reflexivity.
Subcase E-AddValid: $c^{\prime}=s k i p$ and $H^{\prime}=H$
Let $\Theta_{1}^{\prime}=\Theta_{2}^{\prime}=\Theta_{2}$. We have $\Gamma \vdash c^{\prime}: \Theta_{1}^{\prime} \Leftrightarrow \Theta_{2}^{\prime}$ by T-Sкip, We have $H^{\prime} \vDash \Theta_{1}^{\prime}$ by Lemma A. 12 and Lemma A. 14 since $\operatorname{dom}\left(H^{\prime}\right)=\operatorname{dom}(H[h \mapsto v])$ for any $v$, and we have $\Theta_{2}^{\prime}<\Theta_{2}$ by reflexivity.

Case T-Rem: $c=\operatorname{remove}(h)$ and $\Theta_{2}=$ Remove $\Theta_{1} h$
There is only one evaluation rule that applies to $c$, E-Rem, so we have $c^{\prime}=s k i p$ and $H^{\prime}=H \backslash h$. Let $\Theta_{1}^{\prime}=\Theta_{2}^{\prime}=$ Remove $\Theta_{1} h$. We have $\Gamma \vdash c^{\prime}: \Theta_{1}^{\prime} \mapsto \Theta_{2}^{\prime}$ by T-SkIP, we have $H^{\prime} \vDash \Theta_{1}^{\prime}$ by Lemma A.15, and we have $\Theta_{2}^{\prime}<\Theta_{2}$ by reflexivity.

Case T-Mod: $c=h . f=e$ and Includes $\Theta_{1} h$ and $\mathcal{H} \mathcal{T}(h, f)=\tau_{i}$ and $\cdot ; \Theta_{1} \vdash e: \tau_{i}$ and $\Theta_{2}=\Theta_{1}$
There are two evaluation rules that applies to $c$, E-Modi and E-Mod.
Subcase E-Mod1: $c^{\prime}=h . f=e^{\prime}$ and $e \rightarrow e^{\prime}$ and $H^{\prime}=H$
By preservation for expressions we have $; \Theta_{1} \vdash e^{\prime}: \tau_{i}$. Let $\Theta_{1}^{\prime}=\Theta_{2}^{\prime}=\Theta_{1}$. We have $\Gamma \vdash c^{\prime}: \Theta_{1}^{\prime} \Leftrightarrow \Theta_{2}^{\prime}$ by T-Mod, we have $H^{\prime} \vDash \Theta_{1}^{\prime}$ by assumption, and we have $\Theta_{2}^{\prime}<\Theta_{2}$ by reflexivity.
Subcase E-Mod: $c^{\prime}=\operatorname{skip}$ and $\operatorname{dom}\left(H^{\prime}\right)=\operatorname{dom}(H)$
Let $\Theta_{1}^{\prime}=\Theta_{2}^{\prime}=\Theta_{1}$. We have $\Gamma \vdash c^{\prime}: \Theta_{1}^{\prime} \Leftrightarrow \Theta_{2}^{\prime}$ by T-SKIP, we have $H^{\prime} \vDash \Theta_{1}^{\prime}$ by Lemma A.12, and we have $\Theta_{2}^{\prime}<\Theta_{2}$ by reflexivity.

Case T-Zero: Empty $\Theta_{1}$
By Lemma A.7, we have $\operatorname{dom}(H) \in \llbracket \Theta_{1} \rrbracket$. By Lemma A.5, we have $\llbracket \Theta_{1} \rrbracket=\{ \}$, which is a contradiction.

## A. $2 \Pi_{4}$

## A.2. 1 Safety

Lemma A. 27 (Semantic Entailment). If $(I, O, H) \vDash \mathcal{E}^{\tau}$, then $(I, O, H) \in \llbracket \tau \rrbracket_{\mathcal{E}}$
Proof. By induction on $\tau$.
Case $\tau=\varnothing$ : Immediate, since $(I, O, H) \vDash_{\mathcal{E}} \varnothing$ is a contradiction.
Case $\tau=\mathrm{T}$ : Immediate, since $\llbracket \uparrow \rrbracket_{\mathcal{E}}=\mathcal{H}$.
Case $\tau=\Sigma x: \tau_{1} . \tau_{2}:$ By inversion of entailment, we get
(A1) $(I, O, H)=\left(I_{1} @ I_{2}, O_{1} @ O_{2}, H_{1} \cup H_{2}\right)$ and
(A2) $\left(I_{1}, O_{1}, H_{1}\right) \vDash_{\mathcal{E}} \tau_{1}$ and
(A3) $\left(I_{2}, O_{2}, H_{2}\right) \vDash_{\mathcal{E}\left[x \mapsto\left(I_{1}, O_{1}, H_{1}\right)\right]} \tau_{2}$
By ( $\mathrm{A}_{2}$ ) respectively $\left(\mathrm{A}_{3}\right)$ and the induction hypothesis, we get
(A4) $\left(I_{1}, O_{1}, H_{1}\right) \in \llbracket \tau_{1} \rrbracket \mathcal{E}$ and
(A5) $\left(I_{2}, O_{2}, H_{2}\right) \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}\left[x \mapsto\left(I_{1}, O_{1}, H_{1}\right)\right]}$.
To show that $(I, O, H) \in \llbracket \Sigma x: \tau_{1}, \tau_{2} \rrbracket_{\mathcal{E}}=\left\{h_{1}+h_{2} \mid h_{1} \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}} \wedge h_{2} \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}\left[x \mapsto h_{1}\right]}\right\}$, we have to show that $(I, O, H)$ is the concatenation of two heaps $h_{1}$ and $h_{2}$, where $h_{1} \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}}$ and $h_{2} \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}\left[x \mapsto h_{1}\right]}$, which follows from (A1), (A4) and (A5).

Case $\tau=\tau_{1}+\tau_{2}$ : By inversion of entailment, either $(I, O, H) \vDash_{\mathcal{E}} \tau_{1}$ or $(I, O, H) \vDash_{\mathcal{E}} \tau_{2}$. To show that $(I, O, H) \in \llbracket \tau_{1}+\tau_{2} \rrbracket_{\mathcal{E}}=\llbracket \tau_{1} \rrbracket_{\mathcal{E}} \cup \llbracket \tau_{2} \rrbracket_{\mathcal{E}}$, we have to show that $(I, O, H) \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}}$ or $(I, O, H) \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}}$.

Subcase $(I, O, H) \vDash_{\mathcal{E}} \tau_{1}$ : By induction hypothesis, $(I, O, H) \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}}$. We can conclude $(I, O, H) \in \llbracket \tau_{1}+\tau_{2} \rrbracket_{\mathcal{E}}$.

Subcase $(I, O, H) \vDash_{\mathcal{E}} \tau_{2}$ : Symmetric to previous subcase.
Case $\tau=\left\{x: \tau_{1} \mid \varphi\right\}$ : By inversion of entailment, we get
(A1) $(I, O, H) \vDash_{\mathcal{E}} \tau$ and
(A2) $\llbracket \varphi \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]}=$ true
To show that $(I, O, H) \in \llbracket\{x: \tau \mid \varphi\} \rrbracket_{\mathcal{E}}=\left\{h \mid h \in \llbracket \tau \rrbracket_{\mathcal{E}} \wedge \llbracket \varphi \rrbracket_{\mathcal{E}[x \mapsto h]}\right\}$, we have to show that $(I, O, H) \in \llbracket \tau \rrbracket_{\mathcal{E}}$ and that $\llbracket \varphi \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]}=$ true. The first follows by induction hypothesis and ( $\mathrm{A}_{1}$ ) and the latter by ( $\mathrm{A}_{2}$ ).

Case $\tau=\tau_{1}\left[x \mapsto \tau_{2}\right]$ : By inversion of entailment, we get
(A1) $\left(I_{2}, O_{2}, H_{2}\right) \vDash \mathcal{E} \tau_{2}$ for some $I_{2}, O_{2}, H_{2}$ and
(A2) $(I, O, H) \vDash_{\mathcal{E}\left[x \mapsto\left(I_{2}, O_{2}, H_{2}\right)\right]} \tau_{1}$
To show that $(I, O, H) \in \llbracket \tau_{1}\left[x \mapsto \tau_{2}\right] \rrbracket_{\mathcal{E}}=\left\{h \mid h_{2} \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}} \wedge h \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}\left[x \mapsto h_{2}\right]}\right\}$, we have to show that $(I, O, H) \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}\left[x \mapsto h_{2}\right]}$ where $h_{2} \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}}$. By induction hypothesis and (A2) follows that $(I, O, H) \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}\left[x \rightarrow\left(I_{2}, O_{2}, H_{2}\right)\right]} .\left(I_{2}, O_{2}, H_{2}\right) \in$ $\llbracket \tau_{2} \rrbracket_{\mathcal{E}}$ follows by induction hypothesis and (A1), which concludes this case.

Lemma A. 28 (Semantic Containment Entails). If $(I, O, H) \in \llbracket \tau \rrbracket_{\mathcal{E}}$, then $(I, O, H) \vDash_{\mathcal{E}} \tau$.
Proof. By induction on $\tau$.
Case $\tau=\varnothing$ : Immediate, since there is no heap in $\llbracket \varnothing \rrbracket_{\mathcal{E}}$.
Case $\tau=\mathrm{T}$ : Result directly follows by Ent-Top.
Case $\tau=\Sigma x: \tau_{1} \cdot \tau_{2}$ : By the semantics of heap types, all heaps $h \in \llbracket \Sigma x: \tau_{1} . \tau_{2} \rrbracket$ have the form $h=h_{1}++h_{2}$, where $h_{1}=\left(I_{1}, O_{1}, H_{1}\right) \in \llbracket \tau_{1} \rrbracket \mathcal{E}$ and $h_{2}=\left(I_{2}, O_{2}, H_{2}\right) \in$ $\llbracket \tau_{2} \rrbracket_{\mathcal{E}\left[x \mapsto h_{1}\right]}$. By applying the induction hypothesis, we get $\left(I_{1}, O_{1}, H_{1}\right) \vDash \mathcal{E} \tau_{1}$ and $\left(I_{2}, O_{2}, H_{2}\right) \vDash_{\mathcal{E}\left[x \mapsto h_{1}\right]} \tau_{2}$. The result directly follows by Ent-Sigma.

Case $\tau=\tau_{1}+\tau_{2}$ : By the semantics of heap types, for any $h \in \llbracket \tau_{1}+\tau_{2} \rrbracket \mathcal{E}$ holds that either $h \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}}$ or $h \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}}$.

Subcase $h \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}}$ : By induction hypothesis, $h \vDash \mathcal{E}^{\tau_{1}}$. The result directly follows by Ent-ChoiceL.
Subcase $h \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}}$ : Symmetric to previous subcase.
Case $\tau=\left\{y: \tau_{1} \mid \varphi\right\}$ : By the semantics of heap types, $h \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}}$ and $\llbracket \varphi \rrbracket_{\mathcal{E}[x \mapsto h]}=$ true . By induction hypothesis, $h \not \mathcal{E} \tau$. The result directly follows by Ent-Refine.

Case $\tau=\tau_{1}\left[x \mapsto \tau_{2}\right]$ : By the semantics of heap types, $\left.h \in \llbracket \tau_{1}\right]_{\mathcal{E}\left[x \mapsto h_{2}\right]}$ where $h_{2} \in \llbracket \tau_{2} \rrbracket \mathcal{E}$. By induction hypothesis, $h_{2} \vDash_{\mathcal{E}} \tau_{2}$ and $h \vDash_{\mathcal{E}\left[x \mapsto h_{2}\right]} \tau_{1}$. The result directly follows by Ent-Subst.

Lemma A. 29 (Subtype Entailment). If $(I, O, H) \vDash \mathcal{E} \tau_{1}$ and $\mathcal{E} \vDash \Gamma$ and $\Gamma \vdash \tau_{1}<: \tau_{2}$, then $(I, O, H) \vDash \mathcal{E} \tau_{2}$.

Proof. By Lemma A. $27,(I, O, H) \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}}$. With $\mathcal{E} \vDash \Gamma$ and by definition of subtyping, $(I, O, H) \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}}$. The result follows by Lemma A.28.

Lemma A. 30 (Extended Environment Entails). If $\mathcal{E} \vDash \Gamma$ and $(I, O, H) \vDash \mathcal{E} \tau$ and $x \notin \operatorname{dom}(\mathcal{E})$, then $\mathcal{E}[x \mapsto(I, O, H)] \vDash \Gamma, x: \tau$.

Proof. By definition of entailment between environments and typing contexts and by assumptions.

## Progress

Lemma A. 31 (Included Instances in Domain). If $(I, O, H) \vDash_{\mathcal{E}} \tau$ and Includes $\Gamma \tau \iota$, then $\iota \in \operatorname{dom}(H)$.

Proof. By Lemma A. $27,(I, O, H) \in \llbracket \tau \rrbracket_{\mathcal{E}}$. By assumption Includes $\Gamma \tau \iota$ and by definition of inclusion, $\forall h \in \llbracket \tau \rrbracket_{\mathcal{E} . l} \in \operatorname{dom}(h)$, we can conclude that $\iota \in \operatorname{dom}(H)$.

Lemma A. 32 (Excluded Instances not in Domain). If $(I, O, H) \vDash_{\mathcal{E}} \tau$ and Excludes $\Gamma \tau \iota$, then $\iota \notin \operatorname{dom}(H)$.

Proof. By Lemma A. $27,(I, O, H) \in \llbracket \tau \rrbracket_{\mathcal{E}}$. By assumption Excludes $\Gamma \tau \iota$ and by definition of exclusion, $\forall h \in \llbracket \tau \rrbracket_{E . l} \notin \operatorname{dom}(h)$, we can conclude that $\iota \notin \operatorname{dom}(H)$.

Lemma A. 33 (Expression Progress). If $\Gamma ; \tau \vdash e: \sigma$ and $\mathcal{E} \vDash \Gamma$ and $(I, O, H) \vDash \mathcal{E} \tau$, then either $e$ is a value or $\exists e^{\prime} .\langle I, O, H, e\rangle \rightarrow e^{\prime}$.

Lemma A. 34 (Formulae Progress). If $\Gamma ; \tau \vdash \varphi: \mathbb{B}$ and $\mathcal{E} \vDash \Gamma$ and $(I, O, H) \vDash \mathcal{E} \tau$, then either $\varphi$ is a value or $\exists \varphi^{\prime} .\langle I, O, H, \varphi\rangle \rightarrow \varphi^{\prime}$
Theorem A. 35 (Progress). If $\Gamma \vdash c:\left(x: \tau_{1}\right) \rightarrow \tau_{2}$ and $\mathcal{E} \vDash \Gamma$ and $(I, O, H) \vDash \mathcal{E} \tau_{1}$, then either $c=$ skip or there exists $\left\langle I^{\prime}, O^{\prime}, H^{\prime}, c^{\prime}\right\rangle$ such that $\langle I, O, H, c\rangle \rightarrow\left\langle I^{\prime}, O^{\prime}, H^{\prime}, c^{\prime}\right\rangle$.
Proof. By induction on typing derivations of $\Gamma \vdash c:\left(x: \tau_{1}\right) \rightarrow \tau_{2}$.
Case T-SKIP: $c=s k i p$
The result is immediate.
Case T-Extract: $c=\operatorname{extract}(\iota)$ and $\Gamma \vdash \operatorname{sizeof}_{p k t_{i n}}\left(\tau_{1}\right) \geq \operatorname{sizeof}(\iota)$
By inversion of $(I, O, H) \vDash \mathcal{E} \tau_{1}$, we know that $I$ contains enough bits such that deserialize $_{\eta}(I)$ does not fail. Let $\left(v, I^{\prime}\right)=\operatorname{deserialize}_{\eta}(I)$ and $O^{\prime}=O$ and $H^{\prime}=$ $H[\iota \mapsto v]$ and $c^{\prime}=s k i p$. The result follows by E-Extract.

Case T-Reset: $c=$ reset Let $I^{\prime}=O @ I, O^{\prime}=\langle \rangle, H^{\prime}=[]$ and $c^{\prime}=s k i p$. The result follows by E-Reset.

Case T-Remit: $c=\operatorname{remit}(\iota)$ and Includes $\Gamma \tau_{1} \iota$
By Lemma A. 31 we know $\iota \in \operatorname{dom}(H)$. Let $I^{\prime}=I, O^{\prime}=O @ \operatorname{serialize}_{\eta}(H(\iota)), H^{\prime}=$ $H$ and $c^{\prime}=s k i p$. The result follows by E-Remit

Case T-Remove: $c=\operatorname{remove}(\iota)$ and Includes $\Gamma \tau_{1} \iota$
By Lemma A.31, we can conclude that $\iota \in \operatorname{dom}(H)$. Let $I^{\prime}=I$ and $O^{\prime}=O$ and $H^{\prime}=H \backslash \iota$. The result follows by E-Remove.

Case T-Mod: $c=\iota . f:=e$ and Includes $\Gamma \tau_{1} \iota$ and $\mathcal{F}(\iota, f)=\mathrm{BV}$ and $\Gamma ; \tau_{1} \vdash e: \mathrm{BV}$ and $\tau_{2}=\left\{y: \top \mid \varphi_{p k t} \wedge \varphi_{t} \wedge \varphi_{f} \wedge y . \iota . f=e[x /\right.$ heap $\left.]\right\}$
By Lemma A.33, either $e$ is a value or there is some $e^{\prime}$ such that $\langle I, O, H, e\rangle \rightarrow e^{\prime}$.
Subcase $e=v$ : By Lemma A.31, $\iota \in \operatorname{dom}(H)$. Let $r=H(\iota)$ and $r^{\prime}=\{r$ with $f=$ $v\}$. Let $I^{\prime}=I, O^{\prime}=O, H^{\prime}=H[\iota \mapsto r]$ and $c^{\prime}=s k i p$. The result follows by E-Mod.
Subcase $\langle I, O, H, e\rangle \rightarrow e^{\prime}$ : Let $I^{\prime}=I, O^{\prime}=O, H^{\prime}=H^{\prime}$ and $c^{\prime}=\iota . f:=t^{\prime}$. The result follows by E-Modi.

Case T-SEQ: $c=c_{1} ; c_{2}$ and $\Gamma \vdash c_{1}:\left(x: \tau_{1}\right) \rightarrow \tau_{1}^{\prime}$ and $\Gamma \vdash c_{2}:\left(x: \tau_{1}^{\prime}\right) \rightarrow \tau_{2}$
By induction hypothesis, $c_{1}$ is either skip or there is some $\left\langle I^{\prime}, O^{\prime}, H^{\prime}, c_{1}^{\prime}\right\rangle$, such that $\left\langle I, O, H, c_{1}\right\rangle \rightarrow\left\langle I^{\prime}, O^{\prime}, H^{\prime}, c_{1}^{\prime}\right\rangle$. If $c_{1}=s k i p$, let $I^{\prime}=I, O^{\prime}=O, H^{\prime}=H$ and $c^{\prime}=c_{2}$. The result follows by E-SEQ. Otherwise, the result follows by E-SEQ1.

Case T-IF: $c=i f(\varphi) c_{1}$ else $c_{2}$ and $\Gamma ; \tau_{1} \vdash e: \mathbb{B}$
By Lemma A.34, we have that $\varphi$ is either true, false or there is some $\varphi^{\prime}$ such that $\langle I, O, H, \varphi\rangle \rightarrow \varphi^{\prime}$.
Subcase $\varphi=$ true: Let $I^{\prime}=I, O^{\prime}=O, H^{\prime}=H$ and $c^{\prime}=c_{1}$. The result follows by E-If True.

Subcase $\varphi=$ false: Symmetric to previous subcase.
Subcase $\langle I, O, H, \varphi\rangle \rightarrow \varphi^{\prime}$ : Let $I^{\prime}=I$ and $O^{\prime}=O$ and $H^{\prime}=H$. Further, let $c^{\prime}=$ if $\left(\varphi^{\prime}\right) c_{1}$ else $c_{2}$. The result follows by E-IF.

Case T-Add: $c=\operatorname{add}(\iota)$ and Excludes $\Gamma \tau \iota$.
By Lemma A.32, $\downarrow \notin \operatorname{dom}(H)$. The result follows by E-Add.
Case T-Ascribe: $c=c_{a}$ as $\left(x: \tau_{a_{1}}\right) \rightarrow \tau_{a_{2}}$. Let $I^{\prime}=I, O^{\prime}=O, H^{\prime}=H$ and $c^{\prime}=c_{a}$. The result follows by E-Ascribe.

Case T-Sub: $\Gamma \vdash \tau_{1}<: \tau_{3}$ and $\Gamma, x: \tau_{1} \vdash \tau_{4}<: \tau_{2}$ and $\Gamma \vdash c:\left(x: \tau_{3}\right) \rightarrow \tau_{4}$. By Lemma A.29, $(I, O, H) \vDash \mathcal{E} \tau_{3}$. By IH, $c=s k i p$ or there exists $I^{\prime}, O^{\prime}, H^{\prime}, c^{\prime}$ s.t. $\langle I, O, H, c\rangle \rightarrow\left\langle I^{\prime}, O^{\prime}, H^{\prime}, c^{\prime}\right\rangle$. The result follows directly.

## Preservation

Lemma A. 36 (Semantic Chomp Expression). For all expressions $e$, heaps $h$ and $h^{\prime}$, environments $\mathcal{E}$ and $\mathcal{E}^{\prime}$ and variables $x$, if $h^{\prime}=$ chomp $^{\Downarrow}(h, 1)$, and $\mathcal{E}^{\prime}=\mathcal{E}[x \mapsto(\langle \rangle,\langle \rangle,[\iota \mapsto$ $v])]$ and, if $x \in \operatorname{dom}(\mathcal{E}), v=\mathcal{E}(x)(\iota) @ h\left(p k t_{\text {in }}\right)[0: 1]$ and $\mathcal{E}(x)\left(p k t_{i n}\right)=\langle \rangle$ and $\mathcal{E}(x)\left(p k t_{\text {out }}\right)=\langle \rangle$, and otherwise $v=h\left(p k t_{\text {in }}\right)[0: 1]$ and $x$ not free in $e$, then

$$
\llbracket e \rrbracket_{\mathcal{E}[y \mapsto h]}=\llbracket \operatorname{heapRef}_{1}\left(\operatorname{chomp}_{1}^{e}\left(e, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]}
$$

Proof. Proof by induction on $e$. We only consider expressions referencing $p k t_{i n}$. All other expressions are not affected by chomping, and therefore the semantic is unchanged.
Case $e=z \cdot p k t_{\text {in }}[l: r]: \quad$ Case distinction on $z=y:$
Subcase $z \neq y$ :

$$
\begin{aligned}
& \llbracket \text { heapRef }_{1}\left(\operatorname{chomp}_{1}^{e}\left(z \cdot p k t_{i n}[l: r], y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} \\
= & \llbracket \operatorname{heapRef}_{1}\left(z \cdot p k t_{i n}[l: r], \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}}\left[y \mapsto h^{\prime}\right] \\
= & \llbracket z \cdot p k t_{i n}[l: r] \rrbracket_{\mathcal{E}^{\prime}}\left[y \mapsto h^{\prime}\right]
\end{aligned}
$$

If $z \neq x, z$ binds to some heap in $\mathcal{E}$, which must also be contained in $\mathcal{E}^{\prime}$ unchanged. If $z=x$, by assumption, $x . p k t_{i n}$ maps to the empty bit vector, both in $\mathcal{E}$ and $\mathcal{E}^{\prime}$.

$$
=\llbracket z . p k t_{i n}[l: r] \rrbracket_{\mathcal{E}[y \mapsto h]}
$$

Subcase $z=y, r \leq 1$ :

$$
\begin{aligned}
& \llbracket \text { heapRef }_{1}\left(\mathrm{chomp}_{1}^{e}\left(y \cdot p k t_{i n}[\mathrm{o}: 1], y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}}\left[y \mapsto h^{\prime}\right] \\
= & \llbracket \operatorname{heapRef}_{1}\left(\operatorname{chomp}_{1}^{e}\left(y \cdot p k t_{i n}[\mathrm{o}: 1], y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} \\
= & \llbracket \operatorname{heapRef}_{1}\left(\mathrm{~b}_{0}::\langle \rangle, \mathrm{b}_{0}, x, \iota, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} \\
= & \llbracket x . l[\operatorname{sizeof}(\iota)-1: \operatorname{sizeof}(\iota)-1+1] @\langle \rangle \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} \\
= & \llbracket x . \iota[|v|-1:|v|] \rrbracket \rrbracket_{\mathcal{E}^{\prime}}\left[y \mapsto h^{\prime}\right] \\
= & \mathcal{E}^{\prime}(x)(\iota)[|v|-1] \\
= & h\left(p k t_{i n}\right)[\mathrm{o}: 1] \\
= & \llbracket y \cdot p k t_{i n}[\mathrm{o}: 1] \rrbracket \mathcal{E}[y \mapsto h]
\end{aligned}
$$

Subcase $z=y, l=0$ :

$$
\begin{aligned}
& \text { 【heapRef }{ }_{1}\left(\operatorname{chomp}_{1}^{e}\left(y \cdot p k t_{i n}[\mathrm{o}: r], y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, t, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} \\
& =\llbracket \operatorname{heapRef}_{1}\left(\operatorname{chomp}_{1}^{e}\left(y \cdot p k t_{i n}[\mathrm{o}: r], y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \rightarrow h^{\prime}\right]} \\
& =\llbracket \text { heapRef }_{1}\left(\mathrm{~b}_{0}:: y . p k t_{i n}[\mathrm{o}: r-1], \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} \\
& =\llbracket x . \iota[\operatorname{sizeof}(\iota)-1: \operatorname{sizeof}(\iota)-1+1] @ y \cdot p k t_{i n}[0: r-1] \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} \\
& \left.=\llbracket x . \iota[|v|-1:|v|] @ y \cdot p k t_{i n}[\mathrm{o}: r-1]\right]_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} \\
& =\mathcal{E}^{\prime}(x)(\iota)[|v|-1:|v|] @ h^{\prime}\left(p k t_{i n}\right)[0: r-1] \\
& \text { with } v=\mathcal{E}(x)(\iota) @ h\left(p k t_{i n}\right)[\mathrm{o}] \text { follows } \\
& =h\left(p k t_{i n}\right)[0: 1] @ h^{\prime}\left(p k t_{i n}\right)[0: r-1] \\
& \text { with } h^{\prime}=\text { chomp }^{\Downarrow}(h, 1) \text { follows } \\
& =h\left(p k t_{i n}\right)[\mathrm{o}: 1] @ h\left(p k t_{i n}\right)[1: r] \\
& =h\left(p k t_{i n}\right)[\mathrm{o}: r] \\
& =\llbracket y \cdot p k t_{i n}[\mathrm{o}: r] \rrbracket_{\mathcal{E}[y \mapsto h]}
\end{aligned}
$$

Subcase $z=y, l \neq 0$ :

$$
\begin{aligned}
& \llbracket \operatorname{heapRef}_{1}\left(\operatorname{chomp}_{1}^{e}\left(y \cdot p k t_{i n}[l: r], y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} \\
= & {\left[\operatorname{heapRef}_{1}\left(\operatorname{chomp}_{1}^{e}\left(y \cdot p k t_{i n}[l: r], y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]}\right.} \\
= & \llbracket \operatorname{heapRef}_{1}\left(y \cdot p k t_{i n}[l-1: r-1], \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} \\
= & \llbracket y \cdot p k t_{i n}[l-1: r-1] \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} \\
= & h^{\prime}\left(p k t_{i n}\right)[l-1: r-1]
\end{aligned}
$$

with $h^{\prime}=$ chomp $^{\Downarrow}(h, 1)$ follows

$$
\begin{aligned}
& =h\left(p k t_{i n}\right)[l: r] \\
& =\llbracket y \cdot p k t_{i n}[l: r] \rrbracket_{\mathcal{E}[y \mapsto h]}
\end{aligned}
$$

Case $e=z . p k t_{i n}$ : Case distinction on $z=y$ :
Subcase $z \neq y$ : Symmetric to first subcase of previous case.
Subcase $z=y$ :

$$
\begin{aligned}
& \llbracket \text { heap } \operatorname{Ref}_{1}\left(\operatorname{chomp}_{1}^{e}\left(y \cdot p k t_{i n}, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} \\
= & \llbracket \operatorname{heapRef}_{1}\left(\mathrm{~b}_{0}:: y \cdot p k t_{\text {in }}, \mathrm{b}_{0}, x, \iota, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} \\
= & \llbracket x . \iota[\operatorname{sizeof}(\iota)-1: \operatorname{sizeof}(\iota)-1+1] @ y \cdot p k t_{i n} \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} \\
= & \llbracket x . \iota[|v|-1:|v|] @ y \cdot p k t_{\text {in }} \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} \\
= & h\left(p k t_{i n}\right)[\mathrm{o}: 1] @ h^{\prime}\left(p k t_{\text {in }}\right)
\end{aligned}
$$

with $h^{\prime}=$ chomp ${ }^{\Downarrow}(h, 1)$ follows

$$
\begin{aligned}
& =h\left(p k t_{i n}\right) \\
& =\llbracket y \cdot p k t_{i n} \rrbracket_{\mathcal{E}[y \mapsto h]}
\end{aligned}
$$

Case $e=\left|z \cdot p k t_{i n}\right|:$ Case distinction on $z=y$ :
Subcase $z \neq y$ :

$$
\begin{aligned}
& \llbracket \text { heapRef }_{1}\left(\text { chomp }_{1}^{e}\left(\left|z \cdot p k t_{i n}\right|, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, \iota, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} \\
= & \llbracket\left|z \cdot p k t_{i n}\right| \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]}
\end{aligned}
$$

If $z=x$, the length of $x \cdot p k t_{i n}=\mathrm{o}$ in both environments and otherwise, $z . p k t_{i n}$ refers to the same heap in both $\mathcal{E}$ and $\mathcal{E}^{\prime}$.

$$
=\llbracket\left|z \cdot p k t_{i n}\right| \rrbracket_{\mathcal{E}[y \mapsto h]}
$$

Subcase $z=y$ :

$$
\begin{aligned}
& \llbracket \text { heapRef }_{1}\left(\text { chomp }_{1}^{e}\left(\left|y \cdot p k t_{i n}\right|, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, \iota, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} \\
= & \llbracket\left|y \cdot p k t_{i n}\right|+1 \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]}
\end{aligned}
$$

with $h^{\prime}=$ chomp $^{\Downarrow}(h, 1)$ follows

$$
=\llbracket\left|y \cdot p k t_{i n}\right| \rrbracket_{\mathcal{E}[y \mapsto h]}
$$

Case $e=b:: b v$ :

$$
\begin{aligned}
& \llbracket \operatorname{heapRef}_{1}\left(\operatorname{chomp}_{1}^{e}\left(b:: b v, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} \\
= & \llbracket \operatorname{heapRef}_{1}\left(b:: \operatorname{chomp}_{1}^{e}\left(b v, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, t, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} \\
= & \llbracket b:: \operatorname{heap}^{2} \operatorname{Ref}_{1}\left(\operatorname{chomp}_{1}^{e}\left(b v, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, t, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} \\
= & \llbracket b \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]}:: \llbracket \operatorname{heapRef}_{1}\left(\operatorname{chomp}_{1}^{e}\left(b v, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} \\
= & \llbracket b \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]:: \llbracket b v \rrbracket_{\mathcal{E}[y \mapsto h]}}
\end{aligned}
$$

by IH follows

$$
=\llbracket b \rrbracket_{\mathcal{E}[y \mapsto h]}:: \llbracket b v \rrbracket_{\mathcal{E}[y \mapsto h]}
$$

since $b$ is either o or 1

$$
=\llbracket b:: b v \rrbracket_{\mathcal{E}[y \mapsto h]}
$$

Case $e=b v_{1} @ b v_{2}$ :

$$
\begin{aligned}
& \llbracket \text { heapRef }_{1}\left(\operatorname{chomp}_{1}^{e}\left(b v_{1} @ b v_{2}, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, t, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} \\
= & \llbracket \text { heapRef }_{1}\left(\operatorname{chomp}_{1}^{e}\left(b v_{1}, y, \mathrm{~b}_{0}\right) @ \operatorname{chomp}_{1}^{e}\left(b v_{2}, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, t, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} \\
= & \llbracket \text { heapRef }_{1}\left(\operatorname{chomp}_{1}^{e}\left(b v_{1}, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, t, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} @ \\
& \llbracket \text { heapRef }_{1}\left(\operatorname{chomp}_{1}^{e}\left(b v_{1}, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, t, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} \\
= & \llbracket b v_{1} \rrbracket_{\mathcal{E}[y \mapsto h]} @ b v_{2} \rrbracket_{\mathcal{E}[y \mapsto h]}
\end{aligned}
$$

by IH

$$
=\llbracket b v_{1} @ b v_{2} \rrbracket_{\mathcal{E}[y \mapsto h]}
$$

Case $e=n+m$ : Symmetric to previous case.

Lemma A. 37 (Semantic Chomp Formulae). For all formulae $\varphi$, heaps $h$ and $h^{\prime}$, environments $\mathcal{E}$ and $\mathcal{E}^{\prime}$ and variables $x$, if $h^{\prime}=$ chomp $^{\Downarrow}(h, 1)$, and $\mathcal{E}^{\prime}=\mathcal{E}[x \mapsto(\langle \rangle,\langle \rangle,[\iota \mapsto v])]$, and if $x \in \operatorname{dom}, v=\mathcal{E}(x)(\iota) @ h\left(p k t_{\text {in }}\right)[0: 1]$ and $\mathcal{E}(x)\left(p k t_{\text {in }}\right)=\langle \rangle$ and $\mathcal{E}(x)\left(p k t_{\text {out }}\right)=$ $\left\rangle\right.$ and otherwise $v=h\left(p k t_{i n}\right)[0: 1]$ and $x$ not free in $\varphi$, then

$$
\llbracket \varphi \rrbracket_{\mathcal{E}[y \mapsto h]}=\llbracket \operatorname{heapRef}_{1}\left(\operatorname{chomp}_{1}^{\varphi}\left(\varphi, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]}
$$

Proof. By induction on $\varphi$.
Case $\varphi=e_{1}=e_{2}$ :

$$
\begin{aligned}
& \llbracket \text { heapRef }_{1}\left(\operatorname{chomp}_{1}^{\varphi}\left(e_{1}=e_{2}, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, \iota\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} \\
= & \llbracket \operatorname{heapRef}_{1}\left(\operatorname{chomp}_{1}^{\varphi}\left(e_{1}, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l\right)= \\
& \text { heapRef }_{1}\left(\operatorname{chomp}_{1}^{\varphi}\left(e_{1}, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} \\
= & \llbracket \operatorname{heapRef}_{1}\left(\operatorname{chomp}_{1}^{\varphi}\left(e_{1}, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]}= \\
& \llbracket \operatorname{heapRef}_{1}\left(\operatorname{chomp}_{1}^{\varphi}\left(e_{1}, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]}
\end{aligned}
$$

by Lemma A. 36 follows

$$
\begin{aligned}
& =\llbracket e_{1} \rrbracket_{\mathcal{E}[y \rightarrow h]}=\llbracket e_{2} \rrbracket_{\mathcal{E}[y \mapsto h]} \\
& =\llbracket e_{1}=e_{2} \rrbracket_{\mathcal{E}[y \mapsto h]}
\end{aligned}
$$

$\operatorname{Case} \varphi=\varphi_{1} \wedge \varphi_{2}:$

$$
\begin{aligned}
& \llbracket \text { heapRef }_{1}\left(\operatorname{chomp}_{1}^{\varphi}\left(\varphi_{1} \wedge \varphi_{2}, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} \\
= & \boxed{\text { heapRef }_{1}\left(\operatorname{chomp}_{1}^{\varphi}\left(\varphi_{1}, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \wedge} \\
& \text { heapRef }_{1}\left(\operatorname{chomp}_{1}^{\varphi}\left(\varphi_{2}, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, t, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} \\
= & \llbracket \operatorname{heapRef}_{1}\left(\operatorname{chomp}_{1}^{\varphi}\left(\varphi_{1}, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, \iota, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} \wedge \\
& \llbracket \text { heapRef }_{1}\left(\operatorname{chomp}_{1}^{\varphi}\left(\varphi_{2}, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, t, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]}
\end{aligned}
$$

by IH follows

$$
\begin{aligned}
& =\llbracket \varphi_{1} \rrbracket_{\mathcal{E}[y \mapsto h]} \wedge \llbracket \varphi_{2} \rrbracket_{\mathcal{E}[y \mapsto h]} \\
& =\llbracket \varphi_{1} \wedge \varphi_{2} \rrbracket_{\mathcal{E}[y \mapsto h]}
\end{aligned}
$$

Case $\varphi=\neg \varphi_{1}$ :

$$
\begin{aligned}
& \llbracket \text { heapRef }_{1}\left(\operatorname{chomp}_{1}^{\varphi}\left(\neg \varphi_{1}, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket \mathcal{E}^{\prime}[y \mapsto h] \\
= & \llbracket \neg \text { heapRef }_{1}\left(\operatorname{chomp}_{1}^{\varphi}\left(\varphi_{1}, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}[y \mapsto h]} \\
= & \neg \operatorname{heapRef}_{1}\left(\operatorname{chomp}_{1}^{\varphi}\left(\varphi_{1}, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}[y \mapsto h]}
\end{aligned}
$$

by IH follows

$$
\begin{aligned}
& =\neg \llbracket \varphi \rrbracket_{\mathcal{E}[y \mapsto h]} \\
& =\llbracket \neg \varphi \rrbracket_{\mathcal{E}[y \mapsto h]}
\end{aligned}
$$

Case $\varphi=z . \iota^{\prime}$. valid $:$

$$
\begin{aligned}
& \llbracket \text { heapRef }\left(\operatorname{chomp}_{1}^{\varphi}\left(z . \iota^{\prime} . v a l i d, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} \\
= & \llbracket z . \iota^{\prime} . \operatorname{valid} \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} \\
= & \iota^{\prime} \in \operatorname{dom}\left(\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right](z)\right) \\
= & \iota^{\prime} \in \operatorname{dom}(\mathcal{E}[y \mapsto h](z))
\end{aligned}
$$

by definition of $\mathcal{E}^{\prime}$ and $h^{\prime}$ follows

$$
=\llbracket z . \iota^{\prime} . \text { valid } \rrbracket_{\mathcal{E}[y \mapsto h]}
$$

Case $\varphi=$ true:

$$
\begin{aligned}
& \llbracket \text { heapRef } \\
1 & \left(\text { chomp }_{1}^{\varphi}\left(\text { true }, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}[y \mapsto h]} \\
= & \text { 【true } \rrbracket_{\mathcal{E}^{\prime}[y \mapsto h]} \\
= & \text { true } \\
= & \llbracket \text { true } \rrbracket_{\mathcal{E}[y \mapsto h]}
\end{aligned}
$$

Case $\varphi$ = false: Symmetric to previous case.

Lemma A. 38 (Semantic Chomp Refinement). For all heap types $\tau$, heaps $h$ and $h^{\prime}$, environments $\mathcal{E}$ and $\mathcal{E}^{\prime}$ and variables $x$, if $h^{\prime}=\operatorname{chomp}^{\Downarrow}(h, 1)$, and $\mathcal{E}^{\prime}=\mathcal{E}[x \mapsto(\langle \rangle,\langle \rangle,[\iota \mapsto$ $v])]$, and, if $x \in \operatorname{dom}, v=\mathcal{E}(x)(\iota) @ h\left(p k t_{i n}\right)[0: 1]$ and $\mathcal{E}(x)\left(p k t_{i n}\right)=\langle \rangle$ and $\mathcal{E}(x)\left(p k t_{\text {out }}\right)=\langle \rangle$, and otherwise $v=h\left(p k t_{\text {in }}\right)[0: 1]$ and $x$ not free in $\tau$, then

$$
\llbracket \tau \rrbracket_{\mathcal{E}[y \mapsto h]}=\llbracket \operatorname{heapRef}_{1}\left(\operatorname{chompRef}_{1}\left(\tau, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]}
$$

Proof. Proof by induction on $\tau$.
Case $\tau=\varnothing$ :

$$
\begin{aligned}
& \llbracket \text { heapRef }_{1}\left(\operatorname{chompRef}_{1}\left(\varnothing, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} \\
= & \llbracket \varnothing \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} \\
= & \} \\
= & \llbracket \tau \rrbracket_{\mathcal{E}[y \mapsto h]}
\end{aligned}
$$

## Case $\tau=\mathrm{T}$ :

$$
\begin{aligned}
& \llbracket \text { heapRef }_{1}\left(\operatorname{chompRef}_{1}\left(\mathrm{~T}, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket \mathbb{\mathcal { E }}^{\prime}\left[y \mapsto h^{\prime}\right] \\
= & \llbracket \top \rrbracket \mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right] \\
= & H \\
= & \llbracket T \rrbracket_{\mathcal{E}}
\end{aligned}
$$

Case $\tau=\Sigma z: \tau_{1} \cdot \tau_{2}$ :

$$
\begin{aligned}
& \llbracket \text { heapRef }_{1}\left(\operatorname{chompRef}_{1}\left(\Sigma z: \tau_{1}, \tau_{2}, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} \\
= & \llbracket \Sigma z: \text { heapRef }_{1}\left(\operatorname{chompRef}_{1}\left(\tau_{1}, x, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) . \\
& \quad \operatorname{heapRef}_{1}\left(\operatorname{chompRef}_{1}\left(\tau_{2}, x, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} \\
= & \left\{h_{1}^{\prime}++h_{2}^{\prime} \mid h_{1}^{\prime} \in \llbracket \operatorname{heapRef}_{1}\left(\operatorname{chompRef}_{1}\left(\tau_{1}, x, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} \wedge\right. \\
& \left.h_{2}^{\prime} \in \llbracket \operatorname{heapRef}_{1}\left(\operatorname{chompRef}_{1}\left(\tau_{2}, x, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}, z \mapsto h_{1}^{\prime}\right]}\right\} \\
= & \left\{h_{1}^{\prime}++h_{2}^{\prime} \mid h_{1}^{\prime} \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}[y \mapsto h]} \wedge h_{2}^{\prime} \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}\left[y \mapsto h, z \mapsto h_{1}^{\prime}\right]}\right\} \\
= & \llbracket \Sigma z: \tau_{1} . \tau_{2} \rrbracket_{\mathcal{E}[y \mapsto h]}
\end{aligned}
$$

Case $\tau=\tau_{1}+\tau_{2}$ :

$$
\begin{aligned}
& \llbracket \text { heapRef }\left(\operatorname{chompRef}_{1}\left(\tau_{1}+\tau_{2}, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} \\
= & \llbracket \text { heapRef }_{1}\left(\operatorname{chompRef}_{1}\left(\tau_{1}, x, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right)+ \\
& \text { heapRef }_{1}\left(\operatorname{chompRef}_{1}\left(\tau_{2}, x, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}}\left[y \mapsto h^{\prime}\right] \\
= & \llbracket \text { heapRef }_{1}\left(\operatorname{chompRef}_{1}\left(\tau_{1}, x, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} \cup \\
& \llbracket \text { heapRef }_{1}\left(\operatorname{chompRef}_{1}\left(\tau_{2}, x, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} \\
= & \llbracket \tau_{1} \rrbracket \mathcal{E}[y \mapsto h] \\
= & \llbracket \tau_{1}+\tau_{1} \rrbracket \mathcal{E}_{2} \rrbracket_{\mathcal{E}[y \mapsto h \mapsto}[y]
\end{aligned}
$$

$\operatorname{Case} \tau=\left\{z: \tau_{1} \mid \varphi\right\}:$

$$
\begin{aligned}
& \llbracket \text { heapRef }_{1}\left(\operatorname{chompRef}_{1}\left(\left\{z: \tau_{1} \mid \varphi\right\}, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} \\
= & \llbracket\left\{z: \text { heapRef }_{1}\left(\operatorname{chompRef}_{1}\left(\tau_{1}, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \mid\right. \\
& \left.\quad \operatorname{heapRef}_{1}\left(\operatorname{chomp}_{1}^{\varphi}\left(\varphi, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, \iota, 1\right)\right\} \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} \\
= & \left\{h \mid h \in \llbracket \operatorname{heapRef}_{1}\left(\operatorname{chompRef}_{1}\left(\tau_{1}, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} \wedge\right. \\
& \left.\quad \llbracket \operatorname{heapRef}_{1}\left(\operatorname{chomp}_{1}^{\varphi}\left(\varphi, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, \iota, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]}=\operatorname{true}\right\} \\
= & \left\{h \mid h \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}[y \mapsto h]} \wedge \llbracket \varphi \rrbracket \mathcal{E}[y \mapsto h]=\operatorname{true}\right\} \\
= & \llbracket\left\{z: \tau_{1} \mid \varphi\right\} \rrbracket_{\mathcal{E}[y \mapsto h]}
\end{aligned}
$$

Case $\tau=\tau_{1}\left[z \mapsto \tau_{2}\right]:$

$$
\begin{aligned}
& \llbracket \text { heapRef }_{1}\left(\operatorname{chompRef}_{1}\left(\tau_{1}\left[z \mapsto \tau_{2}\right], y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} \\
= & \llbracket \text { heapRef }_{1}\left(\operatorname{chompRef}_{1}\left(\tau_{1}, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right)[z \mapsto \\
& \text { heapRef } \left._{1}\left(\operatorname{chompRef}_{1}\left(\tau_{2}, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, t, 1\right)\right] \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} \\
= & \left\{h \mid h_{2} \in \llbracket \operatorname{heapRef}_{1}\left(\operatorname{chompRef}_{1}\left(\tau_{2}, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, \iota, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]} \wedge\right. \\
& h \in \llbracket \operatorname{heapRef}_{1}\left(\operatorname{chompRef}_{1}\left(\tau_{1}, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, \iota, 1 \rrbracket \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}, z \mapsto h_{2}\right]}\right\} \\
= & \left.\left\{h \mid h_{2} \in \llbracket \tau_{2}\right]_{\mathcal{E}[y \mapsto h]} \wedge h \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}\left[y \mapsto h, z \mapsto h_{2}\right]}\right\} \\
= & \llbracket \tau_{1}\left[z \mapsto \tau_{2}\right] \rrbracket_{\mathcal{E}[y \mapsto h]}
\end{aligned}
$$

Lemma A. 39 (Semantic Chomp ${ }_{1}$ ). For all heap types $\tau$, environments $\mathcal{E}$ and $\mathcal{E}^{\prime}$, and variables $x$, if $\mathcal{E}^{\prime}=\mathcal{E}[x \mapsto(\langle \rangle,\langle \rangle,[\iota \mapsto v])]$, and if $x \in \operatorname{dom}(\mathcal{E}), v=\mathcal{E}(x)(\iota) @ h\left(p k t_{i n}\right)[\mathrm{o}$ : 1] and $\mathcal{E}(x)\left(p k t_{\text {in }}\right)=\mathcal{E}(x)\left(p k t_{\text {out }}\right)=\langle \rangle$, otherwise $v=h\left(p k t_{\text {in }}\right)[0: 1]$, then $\forall h \in$ $\llbracket \tau \rrbracket \mathcal{E} \cdot\left|h\left(p k t_{i n}\right)\right| \geq 1 \Longrightarrow \exists h^{\prime} \in \llbracket \operatorname{chompRec}(\tau, 1, l, x) \rrbracket \mathbb{E}^{\prime} \cdot h^{\prime}=\operatorname{chomp}^{\Downarrow}(h, 1)$

Proof.

```
    \(\forall h \in \llbracket \tau \rrbracket_{\mathcal{E}} \cdot\left|h\left(p k t_{i n}\right)\right| \geq 1 \Longrightarrow \exists h^{\prime} \in \llbracket \operatorname{chompRec}(\tau, 1, l, x) \rrbracket_{\mathcal{E}^{\prime}} \cdot h^{\prime}=\operatorname{chomp}^{\Downarrow}(h, 1)\)
\(\Leftrightarrow\) (By definition of chompRec)
\(\forall h \in \llbracket \tau \rrbracket \mathcal{E} \cdot\left|h\left(p k t_{i n}\right)\right| \geq 1 \Longrightarrow\)
    \(\exists h^{\prime} \in \llbracket \operatorname{heapRef}_{1}\left(\operatorname{chomp}_{1}\left(\tau, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket \mathcal{E}^{\prime} \cdot h^{\prime}=\operatorname{chomp}^{\Downarrow}(h, 1)\)
```

Proof by induction on $\tau$.
Case $\tau=\varnothing: \llbracket \varnothing \rrbracket_{\mathcal{E}}=\{ \}$. As there are no heaps in the semantics, the case holds.
Case $\tau=\mathrm{T}$ : Let $h$ be some heap from $\llbracket T \rrbracket \mathcal{E}=H$. Let $h^{\prime}=h$ except that $h^{\prime}\left(p k t_{\text {in }}\right)=$ $h\left(p k t_{i n}\right)[1:]$.
By definition of chomp ${ }_{1}$ and heapRef ${ }_{1}$, heapRef $_{1}\left(\operatorname{chomp}_{1}\left(T, b_{0}\right), b_{0}, x, l, 1\right)=T$ and $\llbracket T \rrbracket_{\mathcal{E}^{\prime}}=H$.
We can conclude that $h^{\prime} \in \llbracket$ heapRef ${ }_{1}\left(\operatorname{chomp}_{1}\left(\mathrm{~T}, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, t, 1\right) \rrbracket \mathcal{E}^{\prime}$ and $h^{\prime}=$ chomp ${ }^{\Downarrow}(h, 1)$ follows by construction of $h^{\prime}\left(p k t_{i n}\right)=h\left(p k t_{i n}\right)[1:]$.

Case $\tau=\Sigma y: \tau_{1} \cdot \tau_{2}$ : Let $h$ be some heap from $\llbracket \Sigma y: \tau_{1} \cdot \tau_{2} \rrbracket \mathcal{E}$. We know $h=h_{1}++h_{2}$, for some $h_{1} \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}}$ and some $h_{2} \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}\left[y \mapsto h_{1}\right]}$.
We have to show that there exists some

$$
h^{\prime} \in \llbracket \operatorname{heapRef}_{1}\left(\operatorname{chomp}_{1}\left(\Sigma y: \tau_{1} \cdot \tau_{2}, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, \iota, 1\right) \rrbracket_{\mathcal{E}^{\prime}}
$$

such that $h^{\prime}=$ chomp $^{\Downarrow}(h, 1)$.
We deconstruct $\llbracket \operatorname{heapRef}_{1}\left(\operatorname{chomp}_{1}\left(\tau, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}}$ :

$$
\begin{aligned}
& \text { 【heapRef }{ }_{1}\left(\operatorname{chomp}_{1}\left(\Sigma y: \tau_{1} . \tau_{2}, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}} \\
& =\llbracket \operatorname{heapRef}_{1}\left(\Sigma y: \operatorname{chomp}_{1}\left(\tau_{1}, \mathrm{~b}_{0}\right) \cdot \operatorname{chompRef}_{1}\left(\tau_{2}, y, \mathrm{~b}_{0}\right)+\right. \\
& \left.\Sigma y:\left\{z: \tau_{1}| | z \cdot p k t_{i n} \mid=0\right\} . \text { chomp }_{1}\left(\tau_{2}, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}} \\
& =\llbracket \Sigma y \text { : heapRef }{ }_{1}\left(\operatorname{chomp}_{1}\left(\tau_{1}, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, \iota, 1\right) \text {. } \\
& \text { heapRef }_{1}\left(\operatorname{chompRef}_{1}\left(\tau_{2}, x, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right)+ \\
& \Sigma y: \operatorname{heapRef}_{1}\left(\left\{z: \tau_{1}| | z \cdot p k t_{i n} \mid=0\right\}, \mathrm{b}_{0}, x, l, 1\right) \text {. } \\
& \text { heapRef }_{1}\left(\operatorname{chomp}_{1}\left(\tau_{2}, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}} \\
& =\llbracket \Sigma y \text { : heapRef }{ }_{1}\left(\operatorname{chomp}_{1}\left(\tau_{1}, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, \iota, 1\right) \text {. } \\
& \text { heapRef }_{1}\left(\operatorname{chompRef}_{1}\left(\tau_{2}, x, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}} \cup \\
& \llbracket \Sigma y:\left\{z: \tau_{1}| | z \cdot p k t_{i n} \mid=\mathrm{o}\right\} . \text { heapRef }_{1}\left(\operatorname{chomp}_{1}\left(\tau_{2}, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, t, 1\right) \rrbracket_{\mathcal{E}^{\prime}} \\
& =\left\{h_{1}^{\prime}++h_{2}^{\prime} \mid h_{1}^{\prime} \in \llbracket \operatorname{heapRef}_{1}\left(\operatorname{chomp}_{1}\left(\tau_{1}, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}} \wedge\right. \\
& h_{2}^{\prime} \in\left[\operatorname{heapRef}_{1}\left(\operatorname{chompRef}_{1}\left(\tau_{2}, x, \mathrm{~b}_{0}\right)\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h_{1}^{\prime}\right]}\right\} \cup \\
& \left\{h_{1}^{\prime}++h_{2}^{\prime} \mid h_{1}^{\prime} \in \llbracket\left\{z: \tau_{1}| | z \cdot p k t_{i n} \mid=0\right\} \rrbracket_{\mathcal{E}^{\prime}} \wedge\right. \\
& \left.h_{2}^{\prime} \in \llbracket \operatorname{heapRef}_{1}\left(\operatorname{chomp}_{1}\left(\tau_{2}, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, t, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h_{1}^{\prime}\right]}\right\}
\end{aligned}
$$

By case distinction on the length of $p k t_{i n}$ in $h_{1}$.

Subcase $\left|h_{1}\left(p k t_{i n}\right)\right|=0$ :
By definition of $\mathcal{E}^{\prime}, \llbracket \tau_{1} \rrbracket_{\mathcal{E}}=\llbracket \tau_{1} \rrbracket_{\mathcal{E}^{\prime}}$, because, $\tau_{1}$ can't contain a reference to the newly added bit in $\mathcal{E}^{\prime}(x)(\iota)$, from which follows that $h_{1} \in \llbracket \tau_{1} \rrbracket \mathcal{E}^{\prime}$.
By semantics of heap types, $h_{1} \in \llbracket\left\{z: \tau_{1}| | z \cdot p k t_{i n} \mid=0\right\} \rrbracket_{\mathcal{E}^{\prime}}$.
By IH there exists

$$
h_{2}^{\prime} \in \llbracket \text { heapRef }_{1}\left(\operatorname{chomp}_{1}\left(\tau_{2}, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h_{1}\right]}
$$

such that $h_{2}^{\prime}=$ chomp $^{\Downarrow}\left(h_{2}, 1\right)$.
Let $h^{\prime}=h_{1}++h_{2}^{\prime}$.
We conclude that

$$
\begin{aligned}
h^{\prime} \in\left\{h_{1}^{\prime}++h_{2}^{\prime} \mid h_{1}^{\prime}\right. & \in \llbracket\left\{z: \tau_{1}| | z \cdot p k t_{i n} \mid=0\right\} \rrbracket_{\mathcal{E}^{\prime}} \wedge \\
h_{2}^{\prime} & \left.\in \llbracket \operatorname{heapRef}_{1}\left(\operatorname{chomp}_{1}\left(\tau_{2}, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h_{1}^{\prime}\right]}\right\}
\end{aligned}
$$

and thus $h^{\prime} \in \llbracket \operatorname{heapRef}_{1}\left(\operatorname{chomp}_{1}\left(\Sigma y: \tau_{1} . \tau_{2}, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}}$.
By assumption $\left|h_{1}\left(p k t_{i n}\right)\right|=0$, we can conclude that

$$
h_{1}++\operatorname{chomp}^{\Downarrow}\left(h_{2}, 1\right)=\text { chomp }^{\Downarrow}\left(h_{1}++h_{2}, 1\right)
$$

thus

$$
\begin{aligned}
h^{\prime} & =h_{1}++h_{2}^{\prime} \\
& =h_{1}++\operatorname{chomp}^{\Downarrow}\left(h_{2}, 1\right) \\
& =\operatorname{chomp}^{\Downarrow}\left(h_{1}++h_{2}, 1\right) \\
& =\operatorname{chomp}^{\Downarrow}(h, 1)
\end{aligned}
$$

Subcase $\left|h_{1}\left(p k t_{i n}\right)\right| \neq 0$ :
By IH there exists $h_{1}^{\prime} \in \llbracket \operatorname{heapRef}_{1}\left(\operatorname{chomp}_{1}\left(\tau_{1}, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}}$, such that $h_{1}^{\prime}=$ chomp $^{\Downarrow}\left(h_{1}, 1\right)$.
By Lemma A. 38 follows that

$$
h_{2} \in \llbracket \operatorname{heapRef}_{1}\left(\operatorname{chompRef}_{1}\left(\tau_{2}, x, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, t, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h_{1}^{\prime}\right]}
$$

Let $h^{\prime}=h_{1}^{\prime}++h_{2}$.
We conclude that

$$
\begin{aligned}
h^{\prime} \in\left\{h_{1}^{\prime}++h_{2} \mid h_{1}^{\prime}\right. & \in \llbracket \operatorname{heapRef}_{1}\left(\operatorname{chomp}_{1}\left(\tau_{1}, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}} \wedge \\
h_{2} & \left.\in \llbracket \operatorname{heapRef}_{1}\left(\operatorname{chompRef}_{1}\left(\tau_{2}, x, \mathrm{~b}_{0}\right)\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h_{1}^{\prime}\right]}\right\}
\end{aligned}
$$

and thus $h^{\prime} \in \llbracket \operatorname{heapRef}_{1}\left(\operatorname{chomp}_{1}\left(\Sigma y: \tau_{1} . \tau_{2}, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}}$.
With

$$
\begin{aligned}
h^{\prime} & =h_{1}^{\prime}++h_{2} \\
& =\text { chomp }^{\Downarrow}\left(h_{1}, 1\right)++h_{2} \\
& =\text { chomp }^{\Downarrow}\left(h_{1}++h_{2}, 1\right) \\
& =\text { chomp }^{\Downarrow}(h, 1)
\end{aligned}
$$

we can conclude this case.

## Case $\tau=\tau_{1}+\tau_{2}$ :

Let $h$ be some heap from $\llbracket \tau_{1}+\tau_{2} \rrbracket \mathcal{E}$. By the semantics of heap types, we know $h \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}}$ or $h \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}}$.

Subcase $h \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}}$ :
By IH we know that there exists

$$
h^{\prime} \in \llbracket \text { heapRef }_{1}\left(\operatorname{chomp}_{1}\left(\tau_{1}, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}}
$$

such that $h^{\prime}=\operatorname{chomp}^{\Downarrow}(h, 1)$.
By set theory and

$$
\begin{aligned}
& \llbracket \text { heapRef }_{1}\left(\operatorname{chomp}_{1}\left(\tau_{1}+\tau_{2}, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket \mathcal{E}^{\prime} \\
= & \llbracket \text { heapRef }_{1}\left(\operatorname{chomp}_{1}\left(\tau_{1}, \mathrm{~b}_{0}\right)+\operatorname{chomp}_{1}\left(\tau_{2}, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}} \\
= & \llbracket \text { heapRef }_{1}\left(\operatorname{chomp}_{1}\left(\tau_{1}, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right)+ \\
& \text { heapRef }_{1}\left(\operatorname{chomp}_{1}\left(\tau_{2}, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket \mathcal{E}^{\prime} \\
= & \llbracket \text { heapRef }_{1}\left(\operatorname{chomp}_{1}\left(\tau_{1}, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}} \cup \\
& \llbracket \text { heapRef }_{1}\left(\operatorname{chomp}_{1}\left(\tau_{2}, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}}
\end{aligned}
$$

we conclude $h^{\prime} \in \llbracket$ heapRef $_{1}\left(\operatorname{chomp}_{1}\left(\tau_{1}+\tau_{2}, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, t, 1\right) \rrbracket_{\mathcal{E}^{\prime}}$.
Subcase $h \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}}:$ Symmetric to previous subcase.
Case $\tau=\left\{y: \tau_{1} \mid \varphi\right\}$ :
Let $h$ be some heap from $\llbracket\left\{y: \tau_{1} \mid \varphi\right\} \rrbracket \rrbracket_{\mathcal{E}}$.
By the semantics of heap types, we know that $h \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}}$ and $\llbracket \varphi \rrbracket_{\mathcal{E}[y \rightarrow h]}=$ true. By induction hypothesis there exists $h^{\prime} \in \llbracket$ heapRef $_{1}\left(\operatorname{chomp}_{1}\left(\tau_{1}, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}}$ such that $h^{\prime}=$ chomp $^{\Downarrow}(h, 1)$.
By Lemma A. 37 we know that

$$
\llbracket \varphi \rrbracket_{\mathcal{E}[y \mapsto h]}=\llbracket \operatorname{heapRef}_{1}\left(\operatorname{chomp}_{1}^{\varphi}\left(\varphi, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]}
$$

To apply Lemma A.37, we must show that $x$ is not free in $\varphi$, if $x \notin \operatorname{dom}(\mathcal{E})$. If this does not hold, there is no $h \in \llbracket \tau \rrbracket_{\mathcal{E}}$, which violates our initial assumption. With

$$
\begin{aligned}
& \llbracket \text { heapRef } \\
1 & \left(\operatorname{chomp}_{1}\left(\left\{y: \tau_{1} \mid \varphi\right\}, \mathrm{b}_{0}\right), \mathrm{b}_{0}, x, t, 1\right) \rrbracket_{\mathcal{E}^{\prime}} \\
= & \llbracket \operatorname{heapRef}_{1}\left(\left\{y: \operatorname{chomp}_{1}\left(\tau_{1}, \mathrm{~b}_{0}\right) \mid \operatorname{chomp}_{1}^{\varphi}\left(\varphi, y, \mathrm{~b}_{0}\right)\right\}, \mathrm{b}_{0}, x, l, 1\right) \rrbracket \mathcal{E}^{\prime} \\
= & \llbracket\left\{y: \operatorname{heapRef}_{1}\left(\operatorname{chomp}_{1}\left(\tau_{1}, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \mid\right. \\
& \left.\operatorname{heapRef}_{1}\left(\operatorname{chomp}_{1}^{\varphi}\left(\varphi, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right)\right\} \rrbracket_{\mathcal{E}^{\prime}} \\
= & \left\{h \mid h \in \llbracket \operatorname{heapRef}_{1}\left(\operatorname{chomp}_{1}\left(\tau_{1}, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime} \wedge} \wedge\right. \\
& \left.\llbracket \operatorname{heapRef}_{1}\left(\operatorname{chomp}_{1}^{\varphi}\left(\varphi, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}[y \mapsto h]}\right\}
\end{aligned}
$$

and with our assumptions from Lemma A. 37 and the induction hypothesis, we conclude that $h^{\prime} \in \llbracket$ heapRef ${ }_{1}\left(\operatorname{chomp}_{1}\left(\left\{y: \tau_{1} \mid \varphi\right\}, \mathrm{b}_{0}\right), \mathrm{b}_{0}, x, t, 1\right) \rrbracket_{\mathcal{E}^{\prime}}$ such that $h^{\prime}=$ chomp $^{\Downarrow}(h, 1)$.

Case $\tau=\tau_{1}\left[y \mapsto \tau_{2}\right]$ :
Let $h$ be some heap from $\llbracket \tau_{1}\left[y \mapsto \tau_{2}\right] \rrbracket_{\mathcal{E}}$. We know that $h \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}\left[y \mapsto h_{2}\right]}$ for some $h_{2} \in \llbracket \tau_{2} \rrbracket \mathcal{E}$.

By IH there exists some $h^{\prime} \in \llbracket$ heapRef ${ }_{1}\left(\operatorname{chomp}_{1}\left(\tau_{1}, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h_{2}\right]}$ such that $h^{\prime}=\operatorname{chomp}^{\Downarrow}(h, 1)$.

To conclude this case, we must show that

$$
h^{\prime} \in \llbracket \text { heapRef }_{1}\left(\operatorname{chomp}_{1}\left(\tau_{1}\left[y \mapsto \tau_{2}\right], \mathrm{b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}}
$$

From heapRef ${ }_{1}$, chomp $_{1}$ and the semantics of heap types, we get:

$$
\begin{aligned}
& \llbracket \text { heapRef }_{1}\left(\operatorname{chomp}_{1}\left(\tau_{1}\left[y \mapsto \tau_{2}\right], \mathrm{b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket \mathcal{E}^{\prime} \\
= & \llbracket \text { heapRRef }_{1}\left(\operatorname{chomp}_{1}\left(\tau_{1}, \mathrm{~b}_{0}\right)\left[y \mapsto \tau_{2}\right], \mathrm{b}_{0}, x, l, 1\right) \rrbracket \mathcal{E}^{\prime} \\
= & \llbracket \text { heapRef }_{1}\left(\operatorname{chomp}_{1}\left(\tau_{1}, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right)\left[y \mapsto \tau_{2}\right] \rrbracket_{\mathcal{E}^{\prime}} \\
= & \left\{h_{11} \mid h_{22} \in \llbracket \tau_{2} \rrbracket \mathcal{E}^{\prime} \wedge h_{11} \in \llbracket \operatorname{heapRef}_{1}\left(\operatorname{chomp}_{1}\left(\tau_{1}, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket \mathcal{E}^{\prime}\left[y \mapsto h_{22}\right]\right\}
\end{aligned}
$$

With $h_{11}=h^{\prime}$ and $h_{22}=h_{2}$, we can conclude that

$$
h^{\prime} \in \llbracket \operatorname{heapRef}_{1}\left(\operatorname{chomp}_{1}\left(\tau_{1}\left[y \mapsto \tau_{2}\right], \mathrm{b}_{0}\right), \mathrm{b}_{0}, x, t, 1\right) \rrbracket_{\mathcal{E}^{\prime}}
$$

To argue that $h_{2} \in \llbracket \tau_{2} \rrbracket \mathcal{E}^{\prime}$, we make a case distinction on $x \in \operatorname{dom}(\mathcal{E})$.
Subcase $x \notin \operatorname{dom}(\mathcal{E}): x$ cannot appear free in $\tau$ and thereby also not in $\tau_{2}$, otherwise there would be no $h \in \llbracket \tau \rrbracket_{\mathcal{E}}$, thus for all $h_{2} \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}}, h_{2} \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}^{\prime}}$.
Subcase $x \in \operatorname{dom}(\mathcal{E})$ : By assumption, there is some $h_{2} \in \llbracket \tau_{2} \rrbracket \mathcal{E}$. By semantics of heap types, all formulae in $\tau_{2}$ referencing $x$, evaluate to true, i.e., they only refer to information contained in $\mathcal{E}(x)$. By assumption, $\mathcal{E}(x)\left(p k t_{\text {in }}\right)=$ $\mathcal{E}^{\prime}(x)\left(p k t_{\text {in }}\right), \mathcal{E}(x)\left(p k t_{\text {out }}\right)=\mathcal{E}^{\prime}(x)\left(p k t_{\text {out }}\right)$, and $\mathcal{E}(x)(\iota) @ h\left(p k t_{\text {in }}\right)[\mathrm{o}$ : $1]=\mathcal{E}^{\prime}(x)(\iota)$. Since all information of $\mathcal{E}(x)$ is preserved in $\mathcal{E}^{\prime}, h_{2} \in \llbracket \tau_{2} \rrbracket \mathcal{E}^{\prime}$.

Lemma A. 40 (ChompRec Unroll). For all instances $\iota$, if $m+1 \leq \operatorname{sizeof}(t)$, then

$$
\operatorname{chompRec}(\operatorname{chompRec}(\tau, m, x, \iota), 1, x, \iota)=\operatorname{chompRec}(\tau, m+1, x, \iota)
$$

Proof. By induction on $n$.
Case $n=0$ :

$$
\begin{aligned}
& \operatorname{chompRec}(\operatorname{chompRec}(\tau, 0, x, \iota), 1, x, \iota) \\
= & \operatorname{chompRec}(\tau, 1, x, \iota)
\end{aligned}
$$

Case $n=1$ :

```
    chompRec(chompRec}(\tau,1,x,l),1,x,l
    = chompRec(chompRec(heapRef}\mp@subsup{}{1}{(chomp
    = chompRec(heapRef
    = chompRec(\tau,2,x,l)
```


## Case $n=m$ :

We assume that the lemma holds for $n=m$. We now have to show that the lemma also holds for $n=m+1$.

```
    chompRec(chompRec(\tau,m+1,x,l),1,x,l)
= chompRec(chompRec(heapRef
    1,x,l)
= chompRec(heapRef}\mp@subsup{}{1}{}(\mp@subsup{\operatorname{chomp}}{1}{}(\tau,\mp@subsup{b}{0}{}),\mp@subsup{\textrm{b}}{0}{},\iota,m+1),m+1,x,\iota
= chompRec( }\tau,m+2,x,l
```

Lemma A. 41 (Semantic Chomp Unroll). For all heaps $h$ and all $n \in \mathbb{N}$, if $\left|h\left(p k t_{i n}\right)\right| \geq$ $n+1$, then chomp $^{\Downarrow}\left(\operatorname{chomp}^{\Downarrow}(h, n), 1\right)=$ chomp $^{\Downarrow}(h, n+1)$

Proof. By definition of chomp $\downarrow$,

$$
\begin{aligned}
& \operatorname{chomp}^{\Downarrow}\left(\text { chomp }^{\Downarrow}(h, n), 1\right) \\
= & h\left[p k t_{i n} \mapsto h\left(p k t_{i n}\right)[n:]\right]\left[p k t_{i n} \mapsto h\left[p k t_{i n} \mapsto h\left(p k t_{i n}\right)[n:]\right]\left(p k t_{i n}\right)[1:]\right] \\
= & h\left[p k t_{i n} \mapsto h\left[p k t_{\text {in }} \mapsto h\left(p k t_{\text {in }}\right)[n:]\right]\left(p k t_{i n}\right)[1:]\right] \\
= & h\left[p k t_{i n} \mapsto h\left(p k t_{\text {in }}\right)[n+1:]\right] \\
= & \operatorname{chomp}^{\Downarrow}(h, n+1)
\end{aligned}
$$

Lemma A. 42 (Chomp Slice). For all heaps $h$ and all $n \in \mathbb{N}, i f\left|h\left(p k t_{i n}\right)\right| \geq n+1$, then

$$
\operatorname{chomp}^{\Downarrow}(h, n)\left(p k t_{i n}\right)[\mathrm{o}: 1]=h\left(p k t_{i n}\right)[n: n+1]
$$

Proof. By definition of chomp $\downarrow$,

$$
\operatorname{chomp}^{\Downarrow}(h, n)\left(p k t_{i n}\right)[0: 1]=h\left[p k t_{i n} \mapsto h\left(p k t_{i n}\right)[n:]\right]\left(p k t_{i n}\right)[0: 1]
$$

Let $b v=h\left(p k t_{i n}\right)=\left\langle b_{0}, \ldots, b_{n}, \ldots, b_{m}\right\rangle$.
Let $b v^{\prime}$ be the bit vector we obtain after removing the first $n$ bits from $b v^{\prime}, b v^{\prime}=$ $\left\langle b_{n}, \ldots, b_{m}\right\rangle$.

Accessing the first bit of $b v^{\prime}$ gives us bit $b_{n}$, which is also the $n$-th bit in $b v$, i.e., $b v[n: n+1]$.

Lemma A. 43 (Semantic ChompRec). For all heap types $\tau$, environments $\mathcal{E}$ and $\mathcal{E}^{\prime}$, variables $x$ and $n \in \mathbb{N}$, if $x$ does not appear free in $\tau$, and $\mathcal{E}^{\prime}=\mathcal{E}[x \mapsto(\langle \rangle,\langle \rangle,[\iota \mapsto$ $\left.\left.\left.h\left(p k t_{\text {in }}\right)[\mathrm{o}: n]\right]\right)\right]$, then

$$
\forall h \in \llbracket \tau \rrbracket \mathcal{E} \cdot\left|h\left(p k t_{i n}\right)\right| \geq n \Longrightarrow \exists h^{\prime} \in \llbracket \operatorname{chompRec}(\tau, n, l, x) \rrbracket_{\mathcal{E}^{\prime}} \cdot h^{\prime}=\operatorname{chomp}^{\Downarrow}(h, n) .
$$

Proof. Proof by induction on $n$.

Case $n=0$ :

$$
\begin{aligned}
\forall h \in \llbracket \tau \rrbracket \mathcal{E} \cdot\left|h\left(p k t_{i n}\right)\right| \geq 0 \Longrightarrow \\
\quad \exists h^{\prime} \in \llbracket \operatorname{chompRec}(\tau, \mathrm{o}, \iota, x) \rrbracket_{\mathcal{E}^{\prime}} \cdot h^{\prime}=\text { chomp }^{\Downarrow}(h, \mathrm{o}) \\
\Leftrightarrow \forall h \in \llbracket \tau \rrbracket_{\mathcal{E}} \cdot \exists h^{\prime} \in \llbracket \tau \rrbracket_{\mathcal{E}^{\prime}} \cdot h^{\prime}=h
\end{aligned}
$$

Let $h^{\prime}=h$, i.e., we have to show that $h \in \llbracket \tau \rrbracket_{\mathcal{E}^{\prime}}$. By assumption, $x$ is not free in $\tau$, i.e., the binding of $x$ in $\mathcal{E}^{\prime}$ has no effect on the semantics of $\tau$. Since $\mathcal{E}$ and $\mathcal{E}^{\prime}$ are otherwise identical, $\tau$ evaluated in both environments is described by the same set of heaps, from which we can conclude that $h \in \llbracket \tau \rrbracket_{\mathcal{E}^{\prime}}$.

Case $n=1$ :

$$
\begin{aligned}
& \forall h \in \llbracket \tau \rrbracket \mathcal{E} \cdot\left|h\left(p k t_{i n}\right)\right| \geq 1 \Longrightarrow \\
& \quad \exists h^{\prime} \in \llbracket \operatorname{chompRec}(\tau, 1, \iota, x) \rrbracket_{\mathcal{E}^{\prime}} \cdot h^{\prime}=\operatorname{chomp}^{\Downarrow}(h, 1)
\end{aligned}
$$

The result directly follows by Lemma A.39.

## Case $n=m+1$ :

We assume that the lemma holds for $n=m$, i.e.,

$$
\begin{aligned}
& \forall h \in \llbracket \tau \rrbracket_{\mathcal{E}}\left|h\left(p k t_{\text {in }}\right)\right| \geq m \Longrightarrow \\
& \quad \exists h^{\prime} \in \llbracket \operatorname{chompRec}(\tau, m, \iota, x) \rrbracket_{\mathcal{E}^{\prime}} \cdot h^{\prime}=\operatorname{chomp}^{\Downarrow}(h, m)
\end{aligned}
$$

We have to show that the lemma also holds for $n=m+1$, i.e.,

$$
\begin{aligned}
& \forall h \in \llbracket \tau \rrbracket \mathcal{E} \cdot\left|h\left(p k t_{i n}\right)\right| \geq m+1 \Longrightarrow \\
& \quad \exists h^{\prime} \in \llbracket \operatorname{chompRec}(\tau, m+1, \iota, x) \rrbracket_{\mathcal{E}^{\prime}} \cdot h^{\prime}=\operatorname{chomp}^{\Downarrow}(h, m+1)
\end{aligned}
$$

where $\mathcal{E}^{\prime}=\mathcal{E}\left[x \mapsto\left[\iota \mapsto h\left(p k t_{\text {in }}\right)[\mathrm{o}: m+1], p k t_{\text {in }} \mapsto\langle \rangle, p k t_{\text {out }} \mapsto\langle \rangle\right]\right]$.
Let $h$ be some heap $h \in \llbracket \tau \rrbracket_{\mathcal{E}_{0}}$.
By induction hypothesis, there exists some $h^{\prime} \in \llbracket \operatorname{chompRec}(\tau, m, l, x) \rrbracket_{\mathcal{E}_{0}^{\prime}}$, such that $h^{\prime}=$ chomp $^{\Downarrow}(h, m)$.
Let $\mathcal{E}_{1}=\mathcal{E}_{0}^{\prime}=\mathcal{E}_{0}\left[x \mapsto\left(\langle \rangle,\langle \rangle,\left[\iota \mapsto h\left(p k t_{i n}\right)[\mathrm{o}: m]\right]\right)\right]$. We use (A) to refer to this assumption.
By Lemma A.39, for all $h_{1} \in \llbracket \operatorname{chompRec}(\tau, m, \iota, x) \rrbracket \mathcal{E}_{1}$, there exists some heap $h_{1}^{\prime}$ such that $h_{1}^{\prime} \in \llbracket \operatorname{chompRec}(\operatorname{chompRec}(\tau, m, l, x), 1, l, x) \rrbracket_{\mathcal{E}_{1}^{\prime}}$ and $\mathcal{E}_{1}^{\prime}=\mathcal{E}_{1}[x \mapsto$ $\left(\left\rangle,\langle \rangle,\left[\iota \mapsto \mathcal{E}_{1}(x)(\iota) @ h_{1}\left(p k t_{i n}\right)[\mathrm{o}: 1]\right]\right)\right]$, and $h_{1}^{\prime}=\operatorname{chomp}^{\Downarrow}\left(h_{1}, 1\right)$.
Since, by assumption, $h^{\prime} \in \llbracket \operatorname{chompRec}(\tau, m, l, x) \rrbracket_{\mathcal{E}_{o}^{\prime}}$ and also $\mathcal{E}_{1}=\mathcal{E}_{0}^{\prime}$, we can define $h_{1}$ to be equal to $h^{\prime}$, i.e., $h_{1}^{\prime}=$ chomp $^{\Downarrow}\left(h^{\prime}, 1\right)$.
From $h^{\prime}=$ chomp $^{\Downarrow}(h, m)$ follows $h_{1}^{\prime}=$ chomp $^{\Downarrow}\left(\right.$ chomp $\left.^{\Downarrow}(h, m), 1\right)$.
From Lemma A. 41 also follows that $h_{1}^{\prime}=$ chomp $^{\Downarrow}(h, m+1)$.
We must show that $h_{1}^{\prime} \in \llbracket \operatorname{chompRec}(\tau, m+1, \iota, x) \rrbracket_{\mathcal{E}^{\prime}}$.
We know that

$$
h_{1}^{\prime} \in \llbracket \operatorname{chompRec}(\operatorname{chompRec}(\tau, m, l, x), 1, l, x) \rrbracket_{\mathcal{E}_{1}^{\prime}}
$$

and by Lemma A.40, $h_{1}^{\prime} \in \llbracket \operatorname{chompRec}(\tau, m+1, l, x) \rrbracket_{\mathcal{E}_{1}^{\prime}}$, so we must show that $\mathcal{E}^{\prime}=\mathcal{E}_{1}^{\prime}$.
By assumption, $\mathcal{E}_{1}^{\prime}=\mathcal{E}_{1}\left[x \mapsto\left(\langle \rangle,\langle \rangle,\left[\iota \mapsto \mathcal{E}_{1}(x)(\iota) @ h_{1}\left(p k t_{i n}\right)[\mathrm{o}: 1]\right]\right)\right]$, where $h_{1}=$ chomp $^{\Downarrow}(h, m)($ by IH).
Also by assumption, $\mathcal{E}(x)(\iota)=h\left(p k t_{i n}\right)[\mathrm{o}: m]$, i.e., $\mathcal{E}_{1}^{\prime}=\mathcal{E}[x \mapsto(\langle \rangle,\langle \rangle,[\iota \mapsto$ $\left.\left.\left.h\left(p k t_{i n}\right)[\mathrm{o}: m] @ h_{1}\left(p k t_{\text {in }}\right)[\mathrm{o}: 1]\right]\right)\right]$.
Again, substituting $h_{1}$ with $h^{\prime}$, and by $h^{\prime}=$ chomp $^{\Downarrow}(h, m)$, we obtain $\mathcal{E}_{1}^{\prime}=\mathcal{E}[x \mapsto$ $\left(\left\rangle,\langle \rangle,\left[\iota \mapsto h\left(p k t_{i n}\right)[\mathrm{o}: m] @ \operatorname{chomp}^{\Downarrow}(h, m)\left(p k t_{i n}\right)[0: 1]\right]\right)\right]$.
By Lemma A. 42 and by definition of bit vector concatenation, $\mathcal{E}_{1}^{\prime}=\mathcal{E}[x \mapsto$ $\left(\left\rangle,\langle \rangle,\left[\iota \mapsto h\left(p k t_{i n}\right)[\mathrm{o}: m+1]\right]\right)\right]=\mathcal{E}^{\prime}$.

Lemma A. 44 (Semantic Chomp). If $x$ does not appear free in $\tau$, then forall heaps $h \in \llbracket \tau \rrbracket_{\mathcal{E}}$ where $\left|h\left(p k t_{\text {in }}\right)\right| \geq \operatorname{sizeof}(\iota)$, there exists $h^{\prime} \in \llbracket \operatorname{chomp}(\tau, \iota, x) \rrbracket_{\mathcal{E}^{\prime}}$ such that $h^{\prime}=\operatorname{chomp}^{\Downarrow}(h, \operatorname{sizeof}(\iota))$ where $\mathcal{E}^{\prime}=\mathcal{E}\left[x \mapsto\left(\langle \rangle,\langle \rangle,\left[\iota \mapsto h\left(p k t_{i n}\right)[0: \operatorname{sizeof}(\iota)]\right]\right)\right]$.

Proof. By definition of chomp, we know that

$$
\operatorname{chomp}(\tau, \iota, x)=\operatorname{chompRec}(\tau, \operatorname{sizeof}(\iota), x, \iota)
$$

The result follows from Lemma A. 43 .
Lemma 4.45 (Semantic Chomp ${ }_{1}$ Inverse). For all $x, v, \tau, \mathcal{E}^{\prime}$ and $h^{\prime}$ such that $\mathcal{E}^{\prime}(x)=$ $\left(\rangle,\langle \rangle,[\iota \mapsto v])\right.$ and $h^{\prime} \in \llbracket \operatorname{chompRec}(\tau, 1, x, \iota) \rrbracket_{\mathcal{E}^{\prime}}$ and $x$ not free in $\tau$ and $\operatorname{sizeof}(v) \geq 1$, there exists $h$ and $\mathcal{E}$ such that,
(1) $h \in \llbracket \tau \rrbracket \mathcal{E}$ and
(2) $h^{\prime}=\operatorname{chomp}^{\Downarrow}(h, 1)$ and
and
(3) $\mathcal{E}=\mathcal{E}^{\prime} \backslash x$ and
(4) $v=h\left(p k t_{i n}\right)[0: 1]$
or
(3) $x \in \operatorname{dom}(\mathcal{E})$ and
(4) $v=\mathcal{E}(x)(\iota) @ h\left(p k t_{\text {in }}\right)[0: 1]$ and
(5) $\mathcal{E}=\mathcal{E}^{\prime}[x \mapsto(\langle \rangle,\langle \rangle,[\iota \mapsto v[0: \operatorname{sizeof}(v)-1]])]$

Proof. We refer to the general assumptions as follows:
(A) $\mathcal{E}^{\prime}(x)=(\langle \rangle,\langle \rangle,[\iota \mapsto v])$
(B) $h^{\prime} \in \llbracket \operatorname{chompRec}(\tau, 1, x, \iota) \rrbracket_{\mathcal{E}^{\prime}}$ and
(C) $x$ not free in $\tau$ and
(D) $\operatorname{sizeof}(v) \geq 1$

Proof by induction on $\tau$. By definition of chompRec follows that

$$
\llbracket \operatorname{chompRec}(\tau, 1, x, \iota) \rrbracket_{\mathcal{E}}=\llbracket \operatorname{heapRef}_{1}\left(\operatorname{chomp}_{1}\left(\tau, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, \iota, 1\right) \rrbracket_{\mathcal{E}}
$$

Case $\tau=\varnothing$ :

$$
\begin{aligned}
& \llbracket \text { heapRef }_{1}\left(\operatorname{chomp}_{1}\left(\varnothing, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}} \\
= & \llbracket \text { heapRef }_{1}\left(\varnothing, \mathrm{~b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}} \\
= & \llbracket \varnothing \rrbracket_{\mathcal{E}^{\prime}} \\
= & \}
\end{aligned}
$$

As there is no $h^{\prime} \in \llbracket \operatorname{chompRec}(\tau, 1, l, x) \rrbracket \mathcal{E}^{\prime}$, this case is immediate.
Case $\tau=\mathrm{T}$ :

$$
\begin{aligned}
& \llbracket \text { heapRef }_{1}\left(\operatorname{chomp}_{1}\left(\mathrm{~T}, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket \mathbb{E}^{\prime} \\
= & \llbracket \text { heapRef }_{1}\left(\mathrm{~T}, \mathrm{~b}_{0}, x, l, 1\right) \rrbracket \mathcal{E}^{\prime} \\
= & \llbracket \tau \rrbracket_{\mathcal{E}^{\prime}} \\
= & \mathcal{H}
\end{aligned}
$$

Let $\mathcal{E}^{\prime}$ where $\mathcal{E}^{\prime}(x)=(\langle \rangle,\langle \rangle,[\iota \mapsto v])$ be arbitrary.
Let $h^{\prime} \in \llbracket \operatorname{chompRec}(\mathrm{T}, 1, l, x) \rrbracket \mathcal{E}^{\prime}=\mathcal{H}$ be arbitrary. We have to distinguish two cases.

Subcase $\operatorname{sizeof}(v)=1$ :
Let $\mathcal{E}=\mathcal{E}^{\prime} \backslash x$ and let $h=h^{\prime}\left[p k t_{i n} \mapsto v @ h^{\prime}\left(p k t_{i n}\right)\right]$, i.e., $h\left(p k t_{i n}\right)[\mathrm{o}: 1]=v$. (1) follows by the semantics of heap types. (2) follows by the definition of chomp $\downarrow$. (3) and (4) immediately follow from the definition of $h$ and $\mathcal{E}$.
Subcase sizeof $(v)>1$ :
Let $\mathcal{E}=\mathcal{E}^{\prime}[x \mapsto(\langle \rangle,\langle \rangle,[\iota \mapsto v[0: \operatorname{sizeof}(v)-1]])]$ and let $h=h^{\prime}\left[p k t_{i n} \mapsto\right.$ $\left.v[\operatorname{sizeof}(v)-1: \operatorname{sizeof}(v)] @ h^{\prime}\left(p k t_{i n}\right)\right]$. (1) follows by the semantics of heap types. (2) follows by the definition of chomp ${ }^{\Downarrow}$ (3) and (4) immediately follow from the definition of $h$ and $\mathcal{E}$

Case $\tau=\tau_{1}+\tau_{2}$ : Let $\mathcal{E}^{\prime}$ where $\mathcal{E}^{\prime}(x)=(\langle \rangle,\langle \rangle,[\iota \mapsto v])$ be arbitrary.
Let $h^{\prime} \in \llbracket \operatorname{chompRec}\left(\tau_{1}+\tau_{2}, 1, \iota, x\right) \rrbracket \mathcal{E}^{\prime}$ be arbitrary.

$$
\begin{aligned}
& \llbracket \operatorname{chompRec}\left(\tau_{1}+\tau_{2}, 1, l, x\right) \rrbracket_{\mathcal{E}^{\prime}} \\
&= \llbracket \operatorname{heapRef}_{1}\left(\operatorname{chomp}_{1}\left(\tau_{1}+\tau_{2}, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}} \\
&= \llbracket \operatorname{heapRef}_{1}\left(\operatorname{chomp}_{1}\left(\tau_{1}, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, \iota, 1\right) \mathbb{\mathcal { E }}^{\prime} \cup \\
& \llbracket \operatorname{heapRef} \\
& 1
\end{aligned}\left(\operatorname{chomp}{ }_{1}\left(\tau_{1}, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, \iota, 1\right) \rrbracket_{\mathcal{E}^{\prime}} .
$$

We have to distinguish two cases,

1. $h^{\prime} \in \llbracket \operatorname{chompRec}\left(\tau_{1}, 1, x, \iota\right) \rrbracket_{\mathcal{E}^{\prime}}$ and
2. $h^{\prime} \in \llbracket \operatorname{chompRec}\left(\tau_{2}, 1, x \iota\right) \rrbracket_{\mathcal{E}^{\prime}}$.

Subcase $h^{\prime} \in \llbracket \operatorname{chompRec}\left(\tau_{1}, 1, x \iota\right) \rrbracket_{\mathcal{E}^{\prime}}$ : We further distinguish between the size of $v$.
Subcase $\operatorname{sizeof}(v)=1$ :
By IH, there exists $h_{1}, \mathcal{E}_{1}$, such that
(A1) $\left.h_{1} \in \llbracket \tau_{1}\right]_{\mathcal{E}_{1}}$
(A2) $h^{\prime}=$ chomp $^{\Downarrow}\left(h_{1}, 1\right)$
(A3) $\mathcal{E}_{1}=\mathcal{E}^{\prime} \backslash x$
(A4) $v=h_{1}\left(p k t_{i n}\right)[\mathrm{o}: 1]$
Let $h=h_{1}$ and $\mathcal{E}=\mathcal{E}_{1}, h \in \llbracket \tau_{1}+\tau_{2} \rrbracket_{\mathcal{E}}$ follows from the semantics of heap types and by ( A 1 ). The rest is immediate.
Subcase $\operatorname{sizeof}(v)>1$ : By IH, there exists $h_{1}, \mathcal{E}_{1}$, such that
(A1) $h_{1} \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}_{1}}$
(A2) $h^{\prime}=$ chomp $^{\Downarrow}\left(h_{1}, 1\right)$
(A3) $\mathcal{E}_{1}=\mathcal{E}^{\prime}[x \mapsto(\langle \rangle,\langle \rangle,[\iota \mapsto v[\mathrm{o}: \operatorname{sizeof}(v)-1]])]$
(A4) $v=\mathcal{E}(x)(\iota) @ h_{1}\left(p k t_{i n}\right)[\mathrm{o}: 1]$
Let $h=h_{1}$ and $\mathcal{E}=\mathcal{E}_{1} . h \in \llbracket \tau_{1}+\tau_{2} \rrbracket \mathcal{E}$ follows from the semantics of heap types and by ( A 1 ). The rest is immediate.
Subcase $h^{\prime} \in \llbracket \operatorname{chompRec}\left(\tau_{2}, 1, x, l\right) \rrbracket_{\mathcal{E}^{\prime}}:$
Symmetric to previous subcase.
Case $\tau=\Sigma y$ : $\tau_{1} . \tau_{2}$ : Let $\mathcal{E}^{\prime}$ where $\mathcal{E}^{\prime}(x)=(\langle \rangle,\langle \rangle,[\iota \mapsto v])$ be arbitrary. Let $h^{\prime} \in$
$\llbracket c h o m p R e c\left(\Sigma y: \tau_{1} \cdot \tau_{2}, 1, l, x\right) \rrbracket_{\mathcal{E}^{\prime}}$ be arbitrary.

$$
\begin{aligned}
& \llbracket \operatorname{chompRec}\left(\Sigma y: \tau_{1} \cdot \tau_{2}, 1, l, x\right) \rrbracket_{\mathcal{E}^{\prime}} \\
= & \left\{h_{1}^{\prime}++h_{2}^{\prime} \mid h_{1}^{\prime} \in \llbracket \operatorname{chompRec}\left(\tau_{1}, 1, l, x\right) \rrbracket_{\mathcal{E}^{\prime}} \wedge\right. \\
h_{2}^{\prime} \in \llbracket \operatorname{heapRef} & \left.1\left(\operatorname{chompRef}_{1}\left(\tau_{2}, x, \mathrm{~b}_{0}\right)\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h_{1}^{\prime}\right]}\right\} \cup \\
& \left\{h_{1}^{\prime}++h_{2}^{\prime} \mid h_{1}^{\prime} \in \llbracket\left\{z: \tau_{1}| | z \cdot p k t_{i n} \mid=0\right\} \rrbracket_{\mathcal{E}^{\prime}} \wedge\right. \\
& \left.h_{2}^{\prime} \in \llbracket \operatorname{chompRec}\left(\tau_{2}, 1, l, x\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h_{1}^{\prime}\right]}\right\}
\end{aligned}
$$

Case distinction on the membership of $h^{\prime}$.
Subcase $h^{\prime}$ contained in the first subset:
(A1) $h^{\prime}=h_{1}^{\prime}++h_{2}^{\prime}$ and
(A2) $h_{1}^{\prime} \in \llbracket \operatorname{chompRec}\left(\tau_{1}, 1, l, x\right) \rrbracket_{\mathcal{E}^{\prime}}$ and
(A3) $h_{2}^{\prime} \in \llbracket \operatorname{heapRef}_{1}\left(\operatorname{chompRef}_{1}\left(\tau_{2}, x, \mathrm{~b}_{0}\right)\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h_{1}^{\prime}\right]}$.
We distinguish two additional cases.
Subcase $\operatorname{sizeof}(v)=1:$ By IH, there exists $h_{1}, \mathcal{E}_{1}$, such that
(A4) $h_{1} \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}_{1}}$
(A5) $h_{1}^{\prime}=$ chomp $^{\Downarrow}\left(h_{1}, 1\right)$
(A6) $\mathcal{E}_{1}=\mathcal{E}^{\prime} \backslash x$
(A7) $v=h_{1}\left(p k t_{i n}\right)[\mathrm{o}: 1]$
By Lemma A. $38, h_{2}^{\prime} \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h_{1}\right]}$. Since $h_{1}^{\prime}++h_{2}^{\prime}$ is defined, i.e., they have disjoint sets of headers and chomp ${ }^{\Downarrow}$ does not affect the validity of headers, $h_{1}++h_{2}^{\prime}$ is defined.

Let $h=h_{1}+h_{2}^{\prime}$ and $\mathcal{E}=\mathcal{E}_{1}$.
(1) follows by (A4) and $h_{2}^{\prime} \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}_{1}\left[y \mapsto h_{1}\right]}$. The latter holds, because $x$ is not free in $\tau_{2}$ by assumption.
To show (2), we must show that $h_{1}^{\prime}++h_{2}^{\prime}=$ chomp $^{\Downarrow}\left(h_{1}++h_{2}^{\prime}, 1\right) \Leftrightarrow$ chomp ${ }^{\Downarrow}\left(h_{1}, 1\right)++h_{2}^{\prime}=$ chomp $^{\Downarrow}\left(h_{1}++h_{2}^{\prime}, 1\right)$. This equality holds, because chomping of one bit from the input packet of $h_{1}$ and then concatenating $h_{2}^{\prime}$ yields the same heap as concatenating both heaps and then removing the first bit of the input packet.
(3) follows by (A6) and (4) follows by (A7).

Subcase $\operatorname{sizeof}(v)>1$ : By IH, there exists $h_{1}, \mathcal{E}_{1}$, such that
(A8) $h_{1} \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}_{1}}$
(A9) $h_{1}^{\prime}=$ chomp $^{\Downarrow}\left(h_{1}, 1\right)$
(A10) $\mathcal{E}_{1}=\mathcal{E}^{\prime}[x \mapsto(\langle \rangle,\langle \rangle,[\iota \mapsto v[0: \operatorname{sizeof}(v)-1]])]$
(A11) $v=\mathcal{E}_{1}(x)(\iota) @ h_{1}\left(p k t_{i n}\right)[0: 1]$
By Lemma A. $38, h_{2}^{\prime} \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}}\left[y \mapsto h_{1}\right]$. Since $h_{1}^{\prime}++h_{2}^{\prime}$ is defined, i.e., they have disjoint sets of headers and chomp ${ }^{\Downarrow}$ does not affect the validity of headers, $h_{1}++h_{2}^{\prime}$ is defined.
Let $h=h_{1}+h_{2}^{\prime}$ and $\mathcal{E}=\mathcal{E}_{1}$.
(1) follows by (A8) and $h_{2}^{\prime} \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}_{1}\left[y \mapsto h_{1}\right]}$. The latter holds, because $x$ is not free in $\tau_{2}$ by assumption.
To show (2), we must show that $h_{1}^{\prime}++h_{2}^{\prime}=$ chomp $^{\Downarrow}\left(h_{1}++h_{2}^{\prime}, 1\right) \Leftrightarrow$ chomp ${ }^{\Downarrow}\left(h_{1}, 1\right)++h_{2}^{\prime}=$ chomp $^{\Downarrow}\left(h_{1}++h_{2}^{\prime}, 1\right)$. This equality holds, because chomping of one bit from the input packet of $h_{1}$ and then concatenating $h_{2}^{\prime}$ yields the same heap as concatenating both heaps and then removing the first bit of the input packet.
(3) follows by (A10) and (4) follows by (A11) and (5) follows by (A10).

Subcase $h^{\prime}$ contained in the second subset:
(A1) $h^{\prime}=h_{1}^{\prime}++h_{2}^{\prime}$ and
(A2) $h_{1}^{\prime} \in \llbracket\left\{z: \tau_{1}| | z . p k t_{i n} \mid=0\right\} \rrbracket_{\mathcal{E}^{\prime}}$ and
(A3) $h_{2}^{\prime} \in \llbracket \operatorname{chompRec}\left(\tau_{2}, 1, l, x\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h_{1}^{\prime}\right]}$.
We distinguish two cases.
Subcase $\operatorname{sizeof}(v)=1:$ By IH , for every $h_{1}^{\prime} \in \llbracket\left\{z: \tau| | z \cdot p k t_{i n} \mid=0\right\} \rrbracket_{\mathcal{E}^{\prime}}$, there exists $h_{2}, \mathcal{E}_{2}$, such that
(A4) $h_{2} \in \llbracket \tau_{2} \rrbracket \mathcal{E}_{2}$
(A5) $h_{2}^{\prime}=$ chomp $^{\Downarrow}\left(h_{2}, 1\right)$
(A6) $\mathcal{E}_{2}=\mathcal{E}^{\prime}\left[y \mapsto h_{1}^{\prime}\right] \backslash x$
(A7) $v=h_{2}\left(p k t_{i n}\right)[\mathrm{o}: 1]$
Let $h=h_{1}^{\prime}+h_{2}$ and $\mathcal{E}=\mathcal{E}_{2} \backslash y$. We have to show that $h \in \llbracket \Sigma x: \tau_{1} \cdot \tau_{2} \rrbracket \mathcal{E}$. By assumption, $x$ is not free in $\tau_{1}$ and $\tau_{2}$. By (A2) and by the fact that $\mathcal{E}=\mathcal{E}^{\prime} \backslash x, h_{1}^{\prime} \in \llbracket\left\{z: \tau_{1}| | z . p k t_{i n} \mid=0\right\} \rrbracket \mathcal{E}$ and by subtyping, $h_{1}^{\prime} \in \llbracket \tau_{1} \rrbracket \mathcal{E}$.
(1) follows together with (A4).

To show (2), we must show that $h_{1}^{\prime}++h_{2}^{\prime}=$ chomp $^{\Downarrow}\left(h_{1}^{\prime}++h_{2}, 1\right) \Leftrightarrow$ $h_{1}^{\prime}++$ chomp $^{\Downarrow}\left(h_{2}, 1\right)=$ chomp $^{\Downarrow}\left(h_{1}^{\prime}++h_{2}, 1\right)$. Since by (A2), the input packet of $h_{1}^{\prime}$ is empty, the input packet of both heaps are equal.
(3) follows by (A6) and (4) follows by (A7).

Subcase $\operatorname{sizeof}(v)>1:$ By IH, for every $h_{1}^{\prime} \in \llbracket\left\{z: \tau| | z \cdot p k t_{i n} \mid=0\right\} \rrbracket_{\mathcal{E}^{\prime}}$, there exists $h_{2}, \mathcal{E}_{2}$, such that
(A8) $h_{2} \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}_{2}}$
(A9) $h_{2}^{\prime}=$ chomp $^{\Downarrow}\left(h_{2}, 1\right)$
(A10) $\mathcal{E}_{2}=\mathcal{E}^{\prime}\left[y \mapsto h_{1}^{\prime}, x \mapsto(\langle \rangle,\langle \rangle,[\iota \mapsto v[0: \operatorname{sizeof}(v)-1]])\right]$
(A11) $v=\mathcal{E}_{2}(x)(l) @ h_{2}\left(p k t_{i n}\right)[0: 1]$
Let $h=h_{1}^{\prime}++h_{2}$ and $\mathcal{E}=\mathcal{E}_{2} \backslash y$. To show that $h \in \llbracket \Sigma x: \tau_{1} . \tau_{2} \rrbracket \mathcal{E}$. By assumption that $x$ is not free in $\tau_{1}$ and by ( $\mathrm{A}_{2}$ ), $h_{1}^{\prime} \in \llbracket\left\{z: \tau_{1}| | z . p k t_{i n} \mid=\right.$ o\} $\rrbracket_{\mathcal{E}}$ and by subtyping, $h_{1}^{\prime} \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}}$.
(1) follows together with (A8).

To show (2), we must show that $h_{1}^{\prime}++h_{2}^{\prime}=$ chomp $^{\Downarrow}\left(h_{1}^{\prime}++h_{2}, 1\right) \Leftrightarrow$ $h_{1}^{\prime}++$ chomp $^{\Downarrow}\left(h_{2}, 1\right)=$ chomp $^{\Downarrow}\left(h_{1}^{\prime}++h_{2}, 1\right)$. Since by $\left(\mathrm{A}_{2}\right)$, the input packet of $h_{1}^{\prime}$ is empty, the input packet of both heaps are equal.
(3) follows by (A10) and (4) follows by (A11) and (5) follows by (A10).

Case $\tau=\left\{y: \tau_{1} \mid \varphi\right\}$ :
Let $\mathcal{E}^{\prime}$ where $\mathcal{E}^{\prime}(x)=(\langle \rangle,\langle \rangle,[\iota \mapsto v])$ be arbitrary.
Let $h^{\prime} \in \llbracket \operatorname{chompRec}\left(\left\{y: \tau_{1} \mid \varphi\right\}, 1, \iota, x\right) \rrbracket \mathcal{E}^{\prime}$ be arbitrary.

$$
\begin{aligned}
& \llbracket \operatorname{chompRec}\left(\left\{y: \tau_{1} \mid \varphi\right\}, 1, \iota, x\right) \rrbracket_{\mathcal{E}^{\prime}} \\
= & \llbracket \text { heapRef }_{1}\left(\operatorname{chomp}_{1}\left(\left\{y: \tau_{1} \mid \varphi\right\}, \mathrm{b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}} \\
= & \left\{h^{\prime} \mid h^{\prime} \in \llbracket \operatorname{chompRec}\left(\tau_{1}, 1, \iota, x\right) \rrbracket_{\mathcal{E}^{\prime}} \wedge\right. \\
& \left.\llbracket \operatorname{heapRef}_{1}\left(\operatorname{chomp}_{1}^{\varphi}\left(\varphi, y, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h^{\prime}\right]}\right\}
\end{aligned}
$$

We distinguish two cases.
Subcase $\operatorname{sizeof}(v)=1$ :
By IH, there exists $h_{1}, \mathcal{E}_{1}$, such that
(A1) $h_{1} \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}_{1}}$
(A2) $h^{\prime}=$ chomp $^{\Downarrow}\left(h_{1}, 1\right)$
(A3) $\mathcal{E}_{1}=\mathcal{E}^{\prime} \backslash x$
(A4) $v=h_{1}\left(p k t_{i n}\right)[0: 1]$
Let $\mathcal{E}=\mathcal{E}_{1}$ and $h=h_{1}$. To show (1), we must show that $h_{1} \in \llbracket\left\{y: \tau_{1} \mid \varphi\right\} \rrbracket_{\mathcal{E}_{1}}$, which follows by ( $\mathrm{A}_{1}$ ) and Lemma A.37.
(2) follows by (A2), (3) follows by ( $\mathrm{A}_{3}$ ) and (4) follows by ( $\mathrm{A}_{4}$ ).

Subcase $\operatorname{sizeof}(v)>1$ :
By IH, there exists $h_{1}, \mathcal{E}_{1}$, such that
(A1) $h_{1} \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}_{1}}$
(A2) $h^{\prime}=$ chomp $^{\Downarrow}\left(h_{1}, 1\right)$
(A3) $\mathcal{E}_{1}=\mathcal{E}^{\prime}[x \mapsto(\langle \rangle,\langle \rangle,[\iota \mapsto v[0: \operatorname{sizeof}(v)-1]])]$
(A4) $v=\mathcal{E}_{1}(x)(\iota) @ h_{1}\left(p k t_{i n}\right)[\mathrm{o}: 1]$
Let $\mathcal{E}=\mathcal{E}_{1}$ and $h=h_{1}$. To show (1), we must show $h_{1} \in \llbracket\left\{y: \tau_{1} \mid \varphi\right\} \rrbracket_{\mathcal{E}_{1}}$, which follows by ( $\mathrm{A}_{1}$ ) and Lemma A. 37 .
(2) follows by (A2), (3) follows by (A3), (4) follows by (A4) and (5) follows by ( $\mathrm{A}_{3}$ ).

Case $\tau=\tau_{1}\left[y \mapsto \tau_{2}\right]:$
Let $\mathcal{E}^{\prime}$ where $\mathcal{E}^{\prime}(x)=(\langle \rangle,\langle \rangle,[\iota \mapsto v])$ be arbitrary.
Let $h^{\prime} \in \llbracket \operatorname{chompRec}\left(\tau_{1}\left[y \mapsto \tau_{2}\right], 1, \iota, x\right) \rrbracket \mathcal{E}^{\prime}$ be arbitrary.

$$
\begin{aligned}
& \llbracket \operatorname{chompRec}\left(\tau_{1}\left[y \mapsto \tau_{2}\right], 1, l, x\right) \rrbracket \mathcal{E}^{\prime} \\
= & \llbracket \text { heapRef }_{1}\left(\operatorname{chomp}_{1}\left(\tau_{1}\left[y \mapsto \tau_{2}\right], \mathrm{b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right) \rrbracket \mathcal{E}^{\prime} \\
= & \llbracket \operatorname{heapRef}_{1}\left(\operatorname{chomp}_{1}\left(\tau_{1}, \mathrm{~b}_{0}\right)\left[y \mapsto \tau_{2}\right], \mathrm{b}_{0}, x, l, 1\right) \rrbracket \mathcal{E}^{\prime} \\
= & \llbracket \operatorname{heapRef}_{1}\left(\operatorname{chomp}_{1}\left(\tau_{1}, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right)\left[y \mapsto \text { heapRef }_{1}\left(\tau_{2}, \mathrm{~b}_{0}, x, l, 1\right)\right] \rrbracket \mathcal{E}^{\prime} \\
= & \llbracket \operatorname{heapRef}_{1}\left(\operatorname{chomp}_{1}\left(\tau_{1}, \mathrm{~b}_{0}\right), \mathrm{b}_{0}, x, l, 1\right)\left[y \mapsto \tau_{2}\right] \rrbracket_{\mathcal{E}^{\prime}} \\
= & \left\{h_{1}^{\prime} \mid h_{2} \in \llbracket \tau_{2} \rrbracket \mathcal{E}^{\prime} \wedge h_{1}^{\prime} \in \llbracket \operatorname{chompRec}\left(\tau_{1}, 1, x, \iota\right) \rrbracket_{\mathcal{E}^{\prime}}\left[y \mapsto h_{2}\right]\right\}
\end{aligned}
$$

We distinguish two cases.
Subcase sizeof $(v)=1$ :
By IH, for every $h_{2} \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}^{\prime}}$, there exists $h_{1}, \mathcal{E}_{1}$, such that
(A1) $\left.h_{1} \in \llbracket \tau_{1} \rrbracket\right]_{\mathcal{E}_{1}\left[y \mapsto h_{2}\right]}$
(A2) $h_{1}^{\prime}=$ chomp $^{\Downarrow}\left(h_{1}, 1\right)$
(A3) $\mathcal{E}_{1}=\mathcal{E}^{\prime}\left[y \mapsto h_{2}\right] \backslash x$
(A4) $v=h_{1}\left(p k t_{i n}\right)[0: 1]$
Let $\mathcal{E}=\mathcal{E}_{1}$ and $h=h_{1}$. Since $x$ and $y$ not free in $\tau_{2}$, for every heap $h_{2} \in$ $\llbracket \tau_{2} \rrbracket_{\mathcal{E}^{\prime}}$ also holds that $h_{2} \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}_{1}}$, from which we can conclude (1), i.e., $h \in \llbracket \tau_{1}\left[y \mapsto \tau_{2}\right] \rrbracket_{\mathcal{E}}$.
To show (2), we must show that $h^{\prime}=$ chomp $^{\Downarrow}(h, 1)=$ chomp $^{\Downarrow}\left(h_{1}, 1\right)$, which follows from (A2) and the fact that $h^{\prime}=h_{1}^{\prime}$.
(3) follows by choice of $\mathcal{E}$ and (A3). (4) follows by (A4).

Subcase sizeof $(v)>1$ :
By IH, for every $h_{2} \in \llbracket \tau_{2} \rrbracket \mathcal{E}^{\prime}$, there exists $h_{1}, \mathcal{E}_{1}$, such that
(A1) $h_{1} \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}_{1}\left[y \rightarrow h_{2}\right]}$
(A2) $h_{1}^{\prime}=$ chomp $^{\Downarrow}\left(h_{1}, 1\right)$
(A3) $\mathcal{E}_{1}=\mathcal{E}^{\prime}\left[y \mapsto h_{2}, x \mapsto(\langle \rangle,\langle \rangle,[\iota \mapsto v[0: \operatorname{sizeof}(v)-1]])\right]$
(A4) $v=\mathcal{E}_{1}(x)(\iota) @ h_{1}\left(p k t_{i n}\right)[\mathrm{o}: 1]$
Let $\mathcal{E}=\mathcal{E}_{1}$ and $h=h_{1}$.
To show (1), we must show that $\left.h \in \llbracket \tau_{1}\left[y \mapsto \tau_{2}\right]\right]_{\mathcal{E}_{1}}$. Since $x$ and $y$ not free in $\tau_{2}$, for every heap $h_{2} \in \llbracket \tau_{2} \rrbracket \mathcal{E}^{\prime}$ also holds that $h_{2} \in \llbracket \tau_{2} \rrbracket \mathcal{E}_{1}$. The result follows by the semantics of heap types.
To show (2), we must show that $h^{\prime}=$ chomp $^{\Downarrow}\left(h_{1}, 1\right)$, which follows from (A2) and $h^{\prime}=h_{1}^{\prime}$.
(3) follows by ( $\mathrm{A}_{3}$ ), (4) follows by ( $\mathrm{A}_{4}$ ) and (5) follows also by ( $\mathrm{A}_{3}$ ).

Lemma A. 46 (Semantic ChompRec Inverse). For all variables $x$, values $v, n \in \mathbb{N}$, heap types $\tau$, environments $\mathcal{E}^{\prime}$ and heaps $h^{\prime}$ such that $\mathcal{E}^{\prime}(x)=(\langle \rangle,\langle \rangle,[\iota \mapsto v])$ and $h^{\prime} \in \llbracket \operatorname{chompRec}(\tau, n, l, x) \rrbracket_{\mathcal{E}^{\prime}}$ and $x$ not free in $\tau$ and $\operatorname{sizeof}(v)=n$, there exists $h$ and $\mathcal{E}$ such that $h \in \llbracket \tau \rrbracket_{\mathcal{E}}$ and $\mathcal{E}=\mathcal{E}^{\prime} \backslash x$ and chomp ${ }^{\Downarrow}(h, n)=h^{\prime}$.

Proof. Proof by induction on $n$.
Case $n=0$ :
By definition of chompRec, $\operatorname{chompRec}(\tau, 0, \iota, x)=\tau$. Together with assumption $h^{\prime} \in \llbracket \operatorname{chompRec}(\tau, n, l, x) \rrbracket_{\mathcal{E}^{\prime}}$, we know that $h^{\prime} \in \llbracket \tau \rrbracket_{\mathcal{E}^{\prime}}$. Let $h=h^{\prime}$. By assumption, $x$ is not free in $\tau$, thus the binding of $x$ in $\mathcal{E}^{\prime}$ does not affect the semantics of $\tau$. We can therefore remove the binding altogether, so $\tau$ describes the same set of heaps both in $\mathcal{E}^{\prime}$ and in $\mathcal{E}$. We can conclude that $h \in \llbracket \tau \rrbracket_{\mathcal{E}}$. chomp ${ }^{\Downarrow}(h, o)=h^{\prime}$ follows from the definition of semantic chomp.

## Case $n=1$ :

The result directly follows by Lemma A. 45 .

## Case $n=m+1$ :

We assume that the lemma holds for $n=m$. We have to show that the lemma also holds for $n=m+1$. Let $h_{\mathrm{o}}^{\prime}$ be some heap such that $h_{\mathrm{o}}^{\prime} \in \llbracket \operatorname{chompRec}(\tau, m+$ $1, l, x) \rrbracket_{\mathcal{E}_{o}^{\prime}}$. Together with Lemma A.40, we can conclude that

$$
h_{\mathrm{o}}^{\prime} \in \llbracket \operatorname{chompRec}(\operatorname{chompRec}(\tau, m, \iota, x), 1, \iota, x) \rrbracket \mathcal{E}_{\mathrm{o}}^{\prime}
$$

By Lemma A. 45 , there is some $h_{1}^{\prime}, \mathcal{E}_{1}^{\prime}$ such that $h_{1}^{\prime} \in \llbracket \operatorname{chompRec}(\tau, m, \iota, x) \rrbracket_{\mathcal{E}_{1}^{\prime}}$ and $h_{\circ}^{\prime}=\operatorname{chomp}^{\Downarrow}\left(h_{1}^{\prime}, 1\right)$ where $\mathcal{E}_{1}^{\prime}=\mathcal{E}_{0}^{\prime}[x \mapsto(\langle \rangle,\langle \rangle,[\iota \mapsto v[0: m]])]$. By IH , there exists a $h_{1}$ and $\mathcal{E}_{1}$, such that $h_{1} \in \llbracket \tau \rrbracket_{\mathcal{E}_{1}}$ where $\mathcal{E}_{1}=\mathcal{E}_{1}^{\prime} \backslash x$ and $h_{1}^{\prime}=$ chomp $^{\Downarrow}\left(h_{1}, m\right)$. From $h_{\mathrm{o}}^{\prime}=$ chomp $^{\Downarrow}\left(\operatorname{chomp}^{\Downarrow}\left(h_{1}, m\right), 1\right)$ and Lemma A.41, follows $h_{\mathrm{o}}^{\prime}=$ chomp $^{\Downarrow}\left(h_{1}, m+1\right)$.

Lemma A. 47 (Semantic Chomp Inverse). For all variables $x$, values $v$, instances $t$, heap types $\tau$, environments $\mathcal{E}^{\prime}$ and heaps $h^{\prime}$ such that $\mathcal{E}^{\prime}(x)=(\langle \rangle,\langle \rangle,[\iota \mapsto v])$ and $h^{\prime} \in \llbracket \operatorname{chomp}(\tau, \iota, x) \rrbracket_{\mathcal{E}^{\prime}}$ and $x$ not free in $\tau$, there exists $h$ and $\mathcal{E}$ such that $h \in \llbracket \tau \rrbracket_{\mathcal{E}}$ and $\mathcal{E}=\mathcal{E}^{\prime} \backslash x$ and $\operatorname{chomp}^{\Downarrow}(h, \operatorname{sizeof}(\iota))=h^{\prime}$.

Proof. By definition of chomp, we know that

$$
\operatorname{chomp}(\tau, \iota, x)=\operatorname{chompRec}(\tau, \operatorname{sizeof}(\iota), x, \iota)
$$

The result follows by Lemma A. 46 .
Lemma A. 48 (Weakening). If $\Gamma \vdash c:\left(x: \tau_{1}\right) \rightarrow \tau_{2}$ and variable $z$ does not appear free in $\tau_{1}$ or $\tau_{2}$, then $\Gamma, z: \tau \vdash c:\left(x: \tau_{1}\right) \rightarrow \tau_{2}$ for any heap type $\tau$.

Proof. By induction on the typing derivation.
Lemma A. 49 (Input Type Strengthening). If $\Gamma \vdash c:\left(x: \tau_{1}\right) \rightarrow \tau_{2}$ and $\llbracket \tau_{1}^{\prime} \rrbracket \mathcal{E}^{\prime} \subseteq \llbracket \tau_{1} \rrbracket_{\mathcal{E}}$ and $\mathcal{E} \vDash \Gamma$ and $\mathcal{E}^{\prime} \vDash \Gamma^{\prime}$ and $\Gamma \subseteq \Gamma^{\prime}$ and $\mathcal{E} \subseteq \mathcal{E}^{\prime}$, then $\exists \tau_{2}^{\prime} \cdot \Gamma^{\prime} \vdash c:\left(x: \tau_{1}^{\prime}\right) \rightarrow \tau_{2}^{\prime}$ and $\forall h^{\prime} \in \llbracket \tau_{1}^{\prime} \rrbracket_{\mathcal{E}} \cdot \llbracket \tau_{2}^{\prime} \rrbracket_{\mathcal{E}^{\prime}\left[x \mapsto h^{\prime}\right]} \subseteq \llbracket \tau_{2} \rrbracket_{\mathcal{E}\left[x \mapsto h^{\prime}\right]}$

Proof. By induction on a derivation of $\Gamma \vdash c:\left(x: \tau_{1}\right) \rightarrow \tau_{2}$ with case analysis on the last rule used. We refer to the proof goals as follows:
(1) $\exists \tau_{2}^{\prime} \cdot \Gamma^{\prime} \vdash c:\left(x: \tau_{1}^{\prime}\right) \rightarrow \tau_{2}^{\prime}$
(2) $\forall h^{\prime} \in \llbracket \tau_{1}^{\prime} \rrbracket_{\mathcal{E}} \cdot \llbracket \tau_{2}^{\prime} \rrbracket_{\mathcal{E}^{\prime}\left[x \mapsto h^{\prime}\right]} \subseteq \llbracket \tau_{2} \rrbracket_{\mathcal{E}\left[x \mapsto h^{\prime}\right]}$

We refer to the assumptions as follows:
(A) $\llbracket \tau_{1}^{\prime} \rrbracket_{\mathcal{E}^{\prime}} \subseteq \llbracket \tau_{1} \rrbracket_{\mathcal{E}}$
(B) $\mathcal{E} \vDash \Gamma$
(C) $\mathcal{E}^{\prime} \vDash \Gamma^{\prime}$
(D) $\Gamma \subseteq \Gamma^{\prime}$
(E) $\mathcal{E} \subseteq \mathcal{E}^{\prime}$

## Case T-Add:

By inversion of rule T-Add, we get
(A1) Excludes $\Gamma \tau \iota$ and
(A2) init $_{\mathcal{H} \mathcal{T}(\iota)}=v$
(A3) $\tau_{2}=\Sigma y:\left\{z: \tau_{1} \mid z \equiv x\right\} \cdot\left\{z: \iota \mid z \cdot p k t_{\text {out }}=z \cdot p k t_{\text {in }}=\langle \rangle \wedge z \cdot \iota=v\right\}$.
Let $\tau_{2}^{\prime}=\Sigma y:\left\{z: \tau_{1}^{\prime} \mid z \equiv x\right\} \cdot\left\{z: \iota \mid z \cdot p k t_{\text {out }}=z \cdot p k t_{\text {in }}=\langle \rangle \wedge z . \iota=v\right\}$. By assumptions (A1), (A) and (C) we can conclude that Excludes $\tau_{1}^{\prime} \iota \Gamma^{\prime}$ must also hold. (1) follows by T-Add.
Let $\left.h^{\prime} \in \llbracket \tau_{1}^{\prime}\right]_{\mathcal{E}^{\prime}}$ be arbitrary.

$$
\begin{aligned}
& \llbracket \Sigma y:\left\{z: \tau_{1}^{\prime} \mid z \equiv x\right\} \cdot\left\{z: \iota \mid z \cdot p k t_{\text {out }}=z . p k t_{\text {in }}=\langle \rangle \wedge z . l=v\right\} \rrbracket_{\mathcal{E}^{\prime}\left[x \mapsto h^{\prime}\right]} \subseteq \\
& \llbracket \Sigma y:\left\{z: \tau_{1} \mid z \equiv x\right\} \cdot\left\{z: \iota \mid z \cdot p k t_{\text {out }}=z . p k t_{\text {in }}=\langle \rangle \wedge z . l=v\right\} \rrbracket_{\mathcal{E}\left[x \mapsto h^{\prime}\right]} \\
& \Leftrightarrow\left\{h_{1}++h_{2} \mid h_{1} \in \llbracket\left\{z: \tau_{1}^{\prime} \mid z \equiv x\right\} \rrbracket_{\mathcal{E}^{\prime}\left[x \mapsto h^{\prime}\right]} \wedge\right. \\
&\left.h_{2} \in \llbracket\left\{z: \iota \mid z \cdot p k t_{\text {out }}=z \cdot p k t_{\text {in }}=\langle \rangle \wedge z . l=v\right\} \rrbracket_{\mathcal{E}^{\prime}\left[x \mapsto h^{\prime}, y \mapsto h_{1}\right]}\right\} \subseteq \\
&\left\{h_{1}++h_{2} \mid h_{1} \in \llbracket\left\{z: \tau_{1} \mid z \equiv x\right\} \rrbracket_{\mathcal{E}\left[x \mapsto h^{\prime}\right]} \wedge\right. \\
&\left.h_{2} \in \llbracket\left\{z: \iota \mid z . p k t_{\text {out }}=z . p k t_{\text {in }}=\langle \rangle \wedge z . l=v\right\} \rrbracket_{\mathcal{E}\left[x \mapsto h^{\prime}, y \mapsto h_{1}\right]}\right\} \\
& \Leftrightarrow\left\{h^{\prime}++h_{2} \mid h_{2} \in \llbracket\left\{z: \iota \mid z . p k t_{\text {out }}=z . p k t_{\text {in }}=\langle \rangle \wedge z . l=v\right\} \rrbracket_{\mathcal{E}^{\prime}\left[x \mapsto h^{\prime}, y \mapsto h^{\prime}\right]}\right\} \subseteq \\
&\left\{h^{\prime}++h_{2} \mid h_{2} \in \llbracket\left\{z: \iota \mid z . p k t_{\text {out }}=z . p k t_{\text {in }}=\langle \rangle \wedge z . l=v\right\} \rrbracket_{\mathcal{E}\left[x \mapsto h^{\prime}, y \mapsto h^{\prime}\right]}\right\}
\end{aligned}
$$

The type $\left\{z: \iota \mid z . p k t_{\text {out }}=z . p k t_{\text {in }}=\langle \rangle \wedge z . \iota=v\right\}$ does not contain any free variables, so the semantics does not depend on the environment. In fact, the sets of heaps described by $\tau_{2}^{\prime}$ and $\tau_{2}$ is actually equivalent, which shows (2).

## Case T-Ascribe:

By inversion of rule T-Ascribe, we get
(A1) $c=c_{a}$ as $\left(x: \tau_{1}\right) \rightarrow \tau_{2}$ and
(A2) $\Gamma \vdash c_{a}:\left(x: \tau_{1}\right) \rightarrow \tau_{2}$
From assumptions (A2) and (D) together with Lemma A. 48 follows that
(A3) $\Gamma^{\prime} \vdash c_{a}:\left(x: \tau_{1}\right) \rightarrow \tau_{2}$
Since $\mathcal{E}^{\prime}$ differs from $\mathcal{E}$ only in that it potentially contains additional bindings, we can conclude that
(A4) $\llbracket \tau_{1} \rrbracket_{\mathcal{E}}=\llbracket \tau_{1} \rrbracket_{\mathcal{E}^{\prime}}$ and together with assumption (B)
(A5) $\Gamma^{\prime} \vdash \tau_{1}^{\prime}<: \tau_{1}$.

By assumption ( $\mathrm{A}_{3}$ ) and T-Ascribe we get
(A6) $\Gamma^{\prime} \vdash c_{a}$ as $\left(x: \tau_{1}\right) \rightarrow \tau_{2}:\left(x: \tau_{1}\right) \rightarrow \tau_{2}$

Let $\tau_{2}^{\prime}=\tau_{2}$. (1) follows by T-Sub.

For (2), we have to show that $\forall h^{\prime} \in \llbracket \tau_{1}^{\prime} \rrbracket \mathcal{E}^{\prime} \cdot \llbracket \tau_{2} \rrbracket_{\mathcal{E}^{\prime}\left[x \mapsto h^{\prime}\right]} \subseteq \llbracket \tau_{2} \rrbracket_{\mathcal{E}\left[x \mapsto h^{\prime}\right]}$. In fact $\tau_{2}$ describes the same set of heaps, both in $\mathcal{E}$ and $\mathcal{E}^{\prime}$. Variable $x$ binds to the same heap and both environments provide the same bindings for any other free variable in $\tau_{2}$.

## Case T-Extract:

By inversion of rule T-Extract, we get
(A1) $\Gamma \vdash \operatorname{sizeof}_{p k t_{i n}}\left(\tau_{1}\right) \geq \operatorname{sizeof}(\iota)$
(A2) $\tau_{2}=\Sigma y:\left\{z: \iota \mid z \cdot p k t_{\text {in }}=z \cdot p k t_{\text {out }}=\langle \rangle\right\} \cdot\left\{z: \operatorname{chomp}\left(\tau_{1}, \iota, y\right) \mid\right.$ $\left.y . \iota @ z . p k t_{\text {in }}=x . p k t_{\text {in }} \wedge z . p k t_{\text {out }}=x . p k t_{\text {out }} \wedge z \equiv_{l} x\right\}$

Let $\tau_{2}^{\prime}=\Sigma y:\left\{z: \iota \mid z \cdot p k t_{\text {in }}=z . p k t_{\text {out }}=\langle \rangle\right\} .\left\{z: \operatorname{chomp}\left(\tau_{1}^{\prime}, \iota, y\right) \mid\right.$ $\left.y . \iota @ z . p k t_{\text {in }}=x . p k t_{\text {in }} \wedge z . p k t_{\text {out }}=x . p k t_{\text {out }} \wedge z \equiv, x\right\}$. By assumptions $(\mathrm{A})$ and (A1) follows that sizeof ${ }_{p k t_{i n}}\left(\tau_{1}^{\prime}\right) \geq \operatorname{sizeof}(\iota)(1)$ follows by T-Extract.

Let $h^{\prime} \in \llbracket \tau_{1}^{\prime} \rrbracket_{\mathcal{E}^{\prime}}$ be arbitrary.

$$
\begin{aligned}
& \llbracket \Sigma y:\left\{z: \iota \mid z \cdot p k t_{\text {in }}=z \cdot p k t_{\text {out }}=\langle \rangle\right\} . \\
& \left\{z: \operatorname{chomp}\left(\tau_{1}^{\prime}, l, y\right) \mid y . l @ z . p k t_{i n}=x . p k t_{i n} \wedge\right. \\
& \left.z . p k t_{\text {out }}=x . p k t_{\text {out }} \wedge z \equiv x\right\} \rrbracket_{\mathcal{E}^{\prime}\left[x \mapsto h^{\prime}\right]} \subseteq \\
& \llbracket \Sigma y:\left\{z: \iota \mid z \cdot p k t_{\text {in }}=z \cdot p k t_{\text {out }}=\langle \rangle\right\} . \\
& \left\{z: \operatorname{chomp}\left(\tau_{1}, \iota, y\right) \mid y . \iota @ z . p k t_{i n}=x . p k t_{i n} \wedge\right. \\
& \left.z . p k t_{\text {out }}=x . p k t_{\text {out }} \wedge z \equiv x\right\} \rrbracket_{\mathcal{E}\left[x \mapsto h^{\prime}\right]} \\
& \Leftrightarrow\left\{h_{1}+h_{2} \mid h_{1} \in \llbracket\left\{z: \iota \mid z . p k t_{\text {in }}=z . p k t_{\text {out }}=\langle \rangle\right\} \rrbracket_{\mathcal{E}\left[x \mapsto h^{\prime}\right]} \wedge\right. \\
& h_{2} \in \llbracket\left\{z: \operatorname{chomp}\left(\tau_{1}^{\prime}, l, y\right) \mid y . \iota @ z . p k t_{i n}=x . p k t_{i n} \wedge\right. \\
& z . p k t_{\text {out }}=x . p k t_{\text {out }} \wedge \\
& \left.\left.z \equiv_{\iota} x\right\} \rrbracket_{\mathcal{E}\left[x \mapsto h^{\prime}, y \mapsto h_{1}\right]}\right\} \subseteq \\
& \left\{h_{1}++h_{2} \mid h_{1} \in \llbracket\left\{z: \iota \mid z \cdot p k t_{\text {in }}=z . p k t_{\text {out }}=\langle \rangle\right\} \rrbracket_{\mathcal{E}\left[x \mapsto h^{\prime}\right]} \wedge\right. \\
& h_{2} \in \llbracket\left\{z: \operatorname{chomp}\left(\tau_{1}, l, y\right) \mid y . \iota @ z . p k t_{i n}=x \cdot p k t_{i n} \wedge\right. \\
& z . p k t_{\text {out }}=x \cdot p k t_{\text {out }} \wedge \\
& \left.\left.z \equiv_{\iota} x\right\} \rrbracket_{\mathcal{E}\left[x \mapsto h^{\prime}, y \mapsto h_{1}\right]}\right\} \\
& \Leftrightarrow\left\{h_{1}++h_{2} \mid h_{1} \in \llbracket\left\{z: \imath \mid z \cdot p k t_{\text {in }}=z \cdot p k t_{\text {out }}=\langle \rangle\right\} \rrbracket_{\mathcal{E}\left[x \mapsto h^{\prime}\right]} \wedge\right. \\
& h_{2} \in\left\{h_{22} \mid h_{22} \in \llbracket \operatorname{chomp}\left(\tau_{1}^{\prime}, l, y\right) \rrbracket_{\mathcal{E}\left[x \mapsto h^{\prime}, y \mapsto h_{1}\right]} \wedge\right. \\
& \llbracket y . \iota @ z . p k t_{\text {in }}=x . p k t_{\text {in }} \wedge z . p k t_{\text {out }}=x . p k t_{\text {out }} \wedge \\
& \left.z \equiv, x \rrbracket \mathcal{E}\left[x \mapsto h^{\prime}, y \mapsto h_{1}, z \mapsto h_{22}\right]\right\} \subseteq \subseteq \\
& \left\{h_{1}++h_{2} \mid h_{1} \in \llbracket\left\{z: \iota \mid z . p k t_{\text {in }}=z . p k t_{\text {out }}=\langle \rangle\right\} \rrbracket_{\mathcal{E}\left[x \mapsto h^{\prime}\right]} \wedge\right. \\
& h_{2} \in\left\{h_{22} \mid h_{22} \in \llbracket \operatorname{chomp}\left(\tau_{1}, l, y\right) \rrbracket_{\mathcal{E}\left[x \mapsto h^{\prime}, y \mapsto h_{1}\right]} \wedge\right. \\
& \llbracket y . \iota @ z . p k t_{\text {in }}=x . p k t_{\text {in }} \wedge z . p k t_{\text {out }}=x . p k t_{\text {out }} \wedge \\
& \left.\left.z \equiv \iota x \rrbracket_{\mathcal{E}\left[x \mapsto h^{\prime}, y \mapsto h_{1}, z \mapsto h_{22}\right]}\right\}\right\}
\end{aligned}
$$

By Lemma A.44, we obtain all heaps contained in $\llbracket \operatorname{chomp}\left(\tau_{1}^{\prime}, l, y\right) \rrbracket_{\mathcal{E}\left[x \mapsto h^{\prime}, y \mapsto h_{1}\right]}$ by taking all heaps from $\llbracket \tau_{1}^{\prime} \rrbracket_{\mathcal{E}^{\prime}}$ and removing the first sizeof $(\iota)$ bits from the input packet. From assumption (A) we know that all heaps described by $\tau_{1}^{\prime}$ are also contained in the set of heaps described by $\tau_{1}$ and when we remove the first sizeof $(t)$ bits from the input packet, the relation still holds. Since the rest of the types are identical this also holds for the concatenated heaps. This shows (2) and concludes the case.

## Case T-If:

By inversion of rule T-If, we get
(A1) $\Gamma ; \tau_{1} \vdash e: \mathbb{B}$
$\left(\mathrm{A}_{2}\right) \Gamma \vdash c_{1}:\left(x:\left\{y: \tau_{1} \mid \varphi[y /\right.\right.$ heap $\left.\left.]\right\}\right) \rightarrow \tau_{12}$
$\left(\mathrm{A}_{3}\right) \Gamma \vdash c_{2}:\left(x:\left\{y: \tau_{1} \mid \neg \varphi[y /\right.\right.$ heap $\left.\left.]\right\}\right) \rightarrow \tau_{22}$
(A4) $\tau_{2}=\left\{y: \tau_{12} \mid \varphi[x /\right.$ heap $\left.]\right\}+\left\{y: \tau_{22} \mid \varphi[x /\right.$ heap $\left.]\right\}$
(A5) $c=i f(\varphi) c_{1}$ else $c_{2}$
To be able to conclude (1) by T-If, we must show that
(1.1) $\Gamma^{\prime} ; \tau_{1}^{\prime} \vdash \varphi: \mathbb{B}$
(1.2) $\Gamma^{\prime} \vdash c_{1}:\left(x:\left\{y: \tau_{1}^{\prime} \mid \varphi[y /\right.\right.$ heap $\left.\left.]\right\}\right) \rightarrow \tau_{12}^{\prime}$
(1.3) $\Gamma^{\prime} \vdash c_{2}:\left(x:\left\{y: \tau_{1}^{\prime} \mid \neg \varphi[y /\right.\right.$ heap $\left.\left.]\right\}\right) \rightarrow \tau_{22}^{\prime}$

To apply the IH to $c_{1}$, we need some $\tau_{I H_{1}}^{\prime}$ such that

$$
\llbracket \tau_{I H_{1}}^{\prime} \rrbracket_{\mathcal{E}^{\prime}} \subseteq \llbracket\left\{y: \tau_{1} \mid \varphi[y / \text { heap }]\right\} \rrbracket_{\mathcal{E}}
$$

Let $\tau_{I H_{1}}^{\prime}=\left\{y: \tau_{1}^{\prime} \mid \varphi[y /\right.$ heap $\left.]\right\}$.
By IH, there exists $\tau_{12}^{\prime}$ such that
(A6) $\Gamma^{\prime} \vdash c_{1}:\left(x:\left\{y: \tau_{1}^{\prime} \mid \varphi[y /\right.\right.$ heap $\left.\left.]\right\}\right) \rightarrow \tau_{12}^{\prime}$
(A7) $\forall h_{1}^{\prime} \in \llbracket\left\{y: \tau_{1}^{\prime} \mid \varphi[y /\right.$ heap $\left.]\right\} \rrbracket_{\mathcal{E}^{\prime}} \cdot \llbracket \tau_{12}^{\prime} \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h_{1}^{\prime}\right]} \subseteq \llbracket \tau_{12} \rrbracket_{\mathcal{E}\left[y \mapsto h_{1}^{\prime}\right]}$
With a similar argument as before, also by IH, there exists $\tau_{22}^{\prime}$ such that
(A8) $\Gamma^{\prime} \vdash c_{2}:\left(x:\left\{y: \tau_{1}^{\prime} \mid \neg \varphi[y /\right.\right.$ heap $\left.\left.]\right\}\right) \rightarrow \tau_{22}^{\prime}$
(A9) $\forall h_{2}^{\prime} \in \llbracket\left\{y: \tau_{1}^{\prime} \mid \neg \varphi[y /\right.$ heap $\left.]\right\} \rrbracket_{\mathcal{E}^{\prime}} \cdot \llbracket \tau_{22}^{\prime} \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h_{2}^{\prime}\right]}=\llbracket \tau_{22} \rrbracket_{\mathcal{E}\left[y \mapsto h_{2}^{\prime}\right]}$
$\Gamma^{\prime} ; \tau_{1}^{\prime} \vdash \varphi: \mathbb{B}$ also holds, because the subtyping relation between $\tau_{1}^{\prime}$ and $\tau_{1}$ ensures that heaps described by $\tau_{1}^{\prime}$ have the same shape (i.e., the same instances are valid) and thus we can typecheck formula $e$ in the context of type $\tau_{1}^{\prime}$.
(1) follows by T-IF.

Let $\tau_{2}^{\prime}=\left\{y: \tau_{12}^{\prime} \mid \varphi[x /\right.$ heap $\left.]\right\}+\left\{y: \tau_{22}^{\prime} \mid \neg \varphi[x /\right.$ heap $\left.]\right\}$
To show $\forall h_{1}^{\prime} \in \llbracket \tau_{1}^{\prime} \rrbracket_{\mathcal{E}^{\prime}} . \llbracket\left\{y: \tau_{12}^{\prime} \mid \varphi[x /\right.$ heap $\left.]\right\}+\left\{y: \tau_{22} \mid \neg \varphi[x /\right.$ heap $\left.]\right\} \rrbracket_{\mathcal{E}^{\prime}\left[x \mapsto h_{1}^{\prime}\right]} \subseteq$ $\llbracket\left\{y: \tau_{12} \mid \varphi[x /\right.$ heap $\left.]\right\}+\left\{y: \tau_{22} \mid \neg \varphi[x /\right.$ heap $\left.]\right\} \rrbracket_{\mathcal{E}\left[x \mapsto h_{1}^{\prime}\right]}$
Let $h_{1}^{\prime} \in \llbracket \tau_{1}^{\prime} \rrbracket \mathcal{E}^{\prime}$ be arbitrary. Case distinction on wheter the formula $\varphi$ in $h_{1}^{\prime}$ evaluates to true or false.

## Subcase e evaluates to true:

$$
\begin{aligned}
& \llbracket\left\{y: \tau_{12}^{\prime} \mid \varphi[x / \text { heap }]\right\}+\left\{y: \tau_{22} \mid \neg \varphi[x / \text { heap }]\right\} \rrbracket_{\mathcal{E}^{\prime}\left[x \mapsto h_{1}^{\prime}\right]} \subseteq \\
& \llbracket\left\{y: \tau_{12} \mid \varphi[x / \text { heap }]\right\}+\left\{y: \tau_{22} \mid \neg \varphi[x / \text { heap }]\right\} \rrbracket_{\mathcal{E}\left[x \mapsto h_{1}^{\prime}\right]} \\
\Leftrightarrow & \llbracket\left\{y: \tau_{12}^{\prime} \mid \text { true }\right\}+\left\{y: \tau_{22} \mid \text { false }\right\} \rrbracket_{\mathcal{E}}\left[\left\{x \mapsto h_{1}^{\prime}\right] \subseteq\right. \\
& \llbracket\left\{y: \tau_{12} \mid \text { true }\right\}+\left\{y: \tau_{22} \mid \text { false }\right\} \rrbracket_{\mathcal{E}\left[x \mapsto h_{1}^{\prime}\right]} \\
\Leftrightarrow & \llbracket \tau_{12}^{\prime} \rrbracket_{\mathcal{E}^{\prime}\left[x \mapsto h_{1}^{\prime}\right]} \subseteq \llbracket \tau_{12} \rrbracket_{\mathcal{E}\left[x \mapsto h_{1}^{\prime}\right]}
\end{aligned}
$$

The result follows by (A7).
Subcase e evaluates to false: Symmetric to previous subcase. The result follows by (A9).

## Case T-Mod:

(A1) $\tau_{2}=\left\{y: \top \mid \varphi_{p k t} \wedge \varphi_{t} \wedge \varphi_{f} \wedge y . l . f=e[x /\right.$ heap $\left.]\right\}$
(A2) Includes $\Gamma \tau_{1} \downarrow$
( $\left.\mathrm{A}_{3}\right) \Gamma ; \tau \vdash t: \mathrm{BV}$
(A4) $\mathcal{F}(t, f)=\mathrm{BV}$
(A5) $c=t . f:=e$

To show: There exists $\tau_{2}^{\prime}$ such that
(1) $\Gamma^{\prime} \vdash \iota . f:=e:\left(x: \tau_{1}^{\prime}\right) \rightarrow \tau_{2}^{\prime}$ and
(2) $\forall h_{1}^{\prime} \in \llbracket \tau_{1}^{\prime} \rrbracket_{\mathcal{E}^{\prime}} \cdot \llbracket \tau_{2}^{\prime} \rrbracket_{\mathcal{E}\left[x \mapsto h_{1}^{\prime}\right]} \subseteq \llbracket \tau_{2} \rrbracket_{\mathcal{E}\left[x \mapsto h_{1}^{\prime}\right]}$

Let $\tau_{2}^{\prime}=\tau_{2}$.
Includes $\Gamma \tau_{1}^{\prime} \iota$ follows by assumptions ( $\mathrm{A}_{2}$ ) and (A) and set theory. By assumption (A), we know that $\tau_{1}^{\prime}$ has the same shape (contains the same instances) as $\tau_{1}$, so we can typecheck expression $e$ in context $\tau_{1}^{\prime}$ with a bit vector type, from which follows that $\Gamma^{\prime} ; \tau_{1}^{\prime} \vdash e: \mathrm{BV}$. (1) follows by T-Mod.
Let $h_{1}^{\prime} \in \llbracket \tau_{1}^{\prime} \rrbracket_{\mathcal{E}^{\prime}}$ be arbitrary. To show (2), we must show that

$$
\begin{aligned}
& \llbracket \tau_{2} \rrbracket_{\mathcal{E}^{\prime}\left[x \mapsto h_{1}^{\prime}\right]} \subseteq \llbracket \tau_{2} \rrbracket_{\mathcal{E}\left[x \mapsto h_{1}^{\prime}\right]} \\
\Leftrightarrow & \llbracket\left\{y: T \mid \varphi_{p k t} \wedge \varphi_{\iota} \wedge \varphi_{f} \wedge y . \iota . f=t[x / \text { heap }]\right\} \rrbracket_{\mathcal{E}^{\prime}\left[x \mapsto h_{1}^{\prime}\right]} \subseteq \\
& \llbracket\left\{y: T \mid \varphi_{p k t} \wedge \varphi_{\iota} \wedge \varphi_{f} \wedge y . \iota . f=t[x / \text { heap }]\right\} \rrbracket_{\mathcal{E}\left[x \mapsto h_{1}^{\prime}\right]}
\end{aligned}
$$

Since the only free variable is $x$

$$
\begin{aligned}
\Leftrightarrow & \llbracket\left\{y: \top \mid \varphi_{p k t} \wedge \varphi_{\iota} \wedge \varphi_{f} \wedge y . \iota . f=t[x / \text { heap }]\right\} \rrbracket_{\left[x \mapsto h_{1}^{\prime}\right]} \subseteq \\
& \llbracket\left\{y: \top \mid \varphi_{p k t} \wedge \varphi_{\iota} \wedge \varphi_{f} \wedge y . \iota . f=t[x / \text { heap }]\right\} \rrbracket_{\left[x \mapsto h_{1}^{\prime}\right]}
\end{aligned}
$$

The result is immediate.

## Case T-Remit:

By inversion of rule T-Remit, we get
(A1) Includes $\Gamma \tau$
(A2) $\tau_{2}=\Sigma y:\left\{z: \tau_{1} \mid z \equiv x\right\} \cdot\left\{z: \epsilon \mid z \cdot p k t_{\text {in }}=\langle \rangle \wedge z \cdot p k t_{\text {out }}=x . \iota\right\}$
Let $\tau_{2}^{\prime}=\Sigma y:\left\{z: \tau_{1}^{\prime} \mid z \equiv x\right\} .\left\{z: \epsilon \mid z \cdot p k t_{\text {in }}=\langle \rangle \wedge z . p k t_{\text {out }}=x . l\right\}$.
Includes $\Gamma^{\prime} \tau_{1}^{\prime} \iota$ follows by assumptions (A1) and (A) and set theory. (1) follows by T-Remit. Let $h \in \llbracket \tau_{1}^{\prime} \rrbracket_{\mathcal{E}}$ be arbitrary.

$$
\begin{aligned}
& \llbracket \Sigma y:\left\{z: \tau_{1} \mid z \equiv x\right\} \cdot\left\{z: \epsilon \mid z \cdot p k t_{\text {in }}=\langle \rangle \wedge z . p k t_{\text {out }}=x . l\right\} \rrbracket_{\mathcal{E}[x \mapsto h]} \\
= & \left\{h_{1}+h_{2} \mid h_{1} \in \llbracket\left\{z: \tau_{1} \mid z \equiv x\right\} \rrbracket_{\mathcal{E}[x \mapsto h]} \wedge\right. \\
& \left.h_{2} \in \llbracket\left\{z: \epsilon \mid z \cdot p k t_{\text {in }}=\langle \rangle \wedge z \cdot p k t_{\text {out }}=x . l\right\} \rrbracket_{\mathcal{E}\left[x \mapsto h, y \mapsto h_{1}\right]}\right\} \\
= & \left\{h++h_{2} \mid h_{2} \in \llbracket\left\{z: \epsilon \mid z \cdot p k t_{\text {in }}=\langle \rangle \wedge z . p k t_{\text {out }}=x . l\right\} \rrbracket_{\mathcal{E}[x \mapsto h, y \mapsto h]}\right\}
\end{aligned}
$$

$x$ is the only free variable in $\left\{z: \epsilon \mid z \cdot p k t_{\text {in }}=\langle \rangle \wedge z . p k t_{\text {out }}=x . \iota\right\}$, which maps to the same heap $h$ in both environments $\mathcal{E}[x \mapsto h, y \mapsto h]$ and $\mathcal{E}^{\prime}[x \mapsto h, y \mapsto h]$.

$$
\begin{aligned}
& =\left\{h++h_{2} \mid h_{2} \in \llbracket\left\{z: \epsilon \mid z \cdot p k t_{\text {in }}=\langle \rangle \wedge z . p k t_{\text {out }}=x . l\right\} \rrbracket_{\mathcal{E}^{\prime}[x \mapsto h, y \mapsto h]}\right\} \\
& =\left\{h_{1}++h_{2} \mid h_{1} \in \llbracket\left\{z: \tau_{1}^{\prime} \mid z \equiv x\right\} \rrbracket_{\mathcal{E}^{\prime}[x \mapsto h] \wedge}\right. \\
& \left.\quad h_{2} \in \llbracket\left\{z: \epsilon \mid z . p k t_{\text {in }}=\langle \rangle \wedge z \cdot p k t_{\text {out }}=x . l\right\} \rrbracket_{\mathcal{E}^{\prime}\left[x \mapsto h, y \mapsto h_{1}\right]}\right\} \\
& =\llbracket \Sigma y:\left\{z: \tau_{1}^{\prime} \mid z \equiv x\right\} .\left\{z: \epsilon \mid z . p k t_{\text {in }}=\langle \rangle \wedge z . p k t_{\text {out }}=x . l\right\} \rrbracket_{\mathcal{E}^{\prime}[x \mapsto h]}
\end{aligned}
$$

This concludes the case by showing (2).

## Case T-Remove:

By inversion of rule T-Remove, we get
(A1) Includes $\Gamma \tau_{1}{ }_{l}$
(A2) $\tau_{2}=\left\{y: \top \mid \varphi_{\imath} \wedge \varphi_{p k t} \wedge \neg y . .\right.$. valid $\}$
Let $\tau_{2}^{\prime}=\tau_{2}$.
Includes $\Gamma \tau_{1}^{\prime} \iota$ follows by assumptions ( $\mathrm{A}_{1}$ ) and (A) and set theory. (1) then follows by T-Remove.
$\llbracket\left\{y: \top \mid \varphi_{\iota} \wedge \varphi_{p k t} \wedge \neg y . \iota . v a l i d\right\} \rrbracket_{\mathcal{E}^{\prime}\left[x \mapsto h^{\prime}\right]} \subseteq \llbracket\left\{y: \top \mid \varphi_{\iota} \wedge \varphi_{p k t} \wedge \neg y . \iota\right.$. valid $\} \rrbracket_{\mathcal{E}\left[x \mapsto h^{\prime}\right]}$ also holds, because the only free variable is $x$, which maps to the same heap in both environments. This shows (2) and concludes the case.

## Case T-Reset:

By inversion of rule T-Reset, we get
(A1) $c=$ reset
(A2) $\tau_{2}=\Sigma y:\left\{z: \epsilon \mid z \cdot p k t_{\text {out }}=\langle \rangle \wedge z \cdot p k t_{\text {in }}=x \cdot p k t_{\text {out }}\right\} \cdot\left\{z: \epsilon \mid z \cdot p k t_{\text {out }}=\right.$ $\left\rangle \wedge z \cdot p k t_{i n}=x \cdot p k t_{i n}\right\}$
Let $\tau_{2}^{\prime}=\tau_{2}$. (1) follows by T-Reset. Let $h_{1}^{\prime} \in \llbracket \tau_{1}^{\prime} \rrbracket \mathcal{E}^{\prime}$ be arbitrary.

$$
\begin{aligned}
& \llbracket \Sigma y:\left\{z: \epsilon \mid z \cdot p k t_{\text {out }}=\langle \rangle \wedge z . p k t_{\text {in }}=x . p k t_{\text {out }}\right\} . \\
& \\
& \left\{z: \epsilon \mid z \cdot p k t_{\text {out }}=\langle \rangle \wedge z \cdot p k t_{\text {in }}=x \cdot p k t_{\text {in }}\right\} \rrbracket_{\mathcal{E}^{\prime}\left[x \mapsto h_{1}^{\prime}\right]} \subseteq \\
& \llbracket \Sigma y:\left\{z: \epsilon \mid z \cdot p k t_{\text {out }}=\langle \rangle \wedge z . p k t_{\text {in }}=x \cdot p k t_{\text {out }}\right\} . \\
& \\
& \left\{z: \epsilon \mid z \cdot p k t_{\text {out }}=\langle \rangle \wedge z \cdot p k t_{\text {in }}=x . p k t_{\text {in }}\right\} \rrbracket_{\mathcal{E}\left[x \mapsto h_{1}^{\prime}\right]}
\end{aligned}
$$

Both sets are actually equal, because $x$ is the only free variable in $\tau_{2}$ and $\tau_{2}^{\prime}$ respectively. Thus, all other bindings in the environments $\mathcal{E}$ and $\mathcal{E}^{\prime}$ have no effect on the semantics of $\tau_{2}$ and $\tau_{2}^{\prime}$ respectively. This shows (2) and concludes the case.

## Case T-SEQ:

By inversion of rule T-SEQ, we get
(A1) $c=c_{1} ; c_{2}$
(A2) $\Gamma \vdash c_{1}:\left(x: \tau_{1}\right) \rightarrow \tau_{12}$
(A3) $\Gamma, x: \tau_{1} \vdash c_{2}:\left(y: \tau_{12}\right) \rightarrow \tau_{22}$
(A4) $\tau_{2}=\tau_{22}\left[y \mapsto \tau_{12}\right]$
By IH with (A2), (A), (B) and (C), there exists some $\tau_{12}^{\prime}$ such that
(A5) $\Gamma^{\prime} \vdash c_{1}:\left(x: \tau_{1}^{\prime}\right) \rightarrow \tau_{12}^{\prime}$
(A6) $\forall h_{1}^{\prime} \in \llbracket \tau_{1}^{\prime} \rrbracket_{\mathcal{E}} \cdot \llbracket \tau_{12}^{\prime} \rrbracket_{\mathcal{E}^{\prime}\left[x \mapsto h_{1}^{\prime}\right]} \subseteq \llbracket \tau_{12} \rrbracket_{\mathcal{E}\left[x \mapsto h_{1}^{\prime}\right]}$
Apply the IH again to $c_{2}$ :
Let $h_{1}^{\prime} \in \llbracket \tau_{1}^{\prime} \rrbracket_{\mathcal{E}}$, be arbitrary. By ( $\mathrm{A}_{3}$ ) $\Gamma, x: \tau_{1} \vdash c_{2}:\left(y: \tau_{12}\right) \rightarrow \tau_{22}$. By (A6), $\llbracket \tau_{12}^{\prime} \rrbracket_{\mathcal{E}}\left[x \mapsto h_{1}^{\prime}\right] \subseteq \llbracket \tau_{12} \rrbracket_{\mathcal{E}\left[x \mapsto h_{1}^{\prime}\right]} \cdot \mathcal{E}\left[x \mapsto h_{1}^{\prime}\right] \vDash \Gamma, x: \tau_{1}$ because by assumption $\mathcal{E} \vDash \Gamma$
the entailment holds for all $x_{i} \neq x$. For $x$ there exists a binding to heap $h_{1}^{\prime} \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}}$ (with assumption (A)) and the entailment between $h$ and $\tau_{1}$ trivially holds.
To show $\mathcal{E}^{\prime}\left[x \mapsto h_{1}^{\prime}\right] \vDash \Gamma^{\prime}, x: \tau_{1}^{\prime}$, we must show that $\forall x_{i}, \tau_{i} . \Gamma^{\prime}\left(x_{i}\right)=\tau_{i} \Rightarrow \mathcal{E}^{\prime}[x \mapsto$ $\left.h_{1}^{\prime}\right]\left(x_{i}\right)=h_{i} \wedge h_{i} \vDash_{\mathcal{E}\left[x \mapsto h_{1}^{\prime}\right]} \tau_{i}$. Case $x_{i} \neq x$ : this holds by assumption (C). Case $x_{i}=x . \mathcal{E}^{\prime}\left[x \mapsto h_{1}^{\prime}\right](x)=h_{1}^{\prime}$. To show that $\left.h_{1}^{\prime} \vDash \mathcal{E}^{\prime}\left[x \mapsto h_{1}^{\prime}\right] \tau_{1}^{\prime} \Leftrightarrow h_{1}^{\prime} \in \llbracket \tau_{1}^{\prime}\right]_{\mathcal{E}^{\prime}\left[x \mapsto h_{1}^{\prime}\right]}$. By assumption, $x$ is not free in $\tau_{1}^{\prime}$, so we can equivalently show that $h_{1}^{\prime} \in \llbracket \tau_{1}^{\prime} \rrbracket \mathcal{E}^{\prime}$, which holds by assumption.
Again by IH, there exists some $\tau_{22}^{\prime}$ such that
(A7) $\Gamma^{\prime}, x: \tau_{1}^{\prime} \vdash c_{2}\left(y: \tau_{12}^{\prime}\right) \rightarrow \tau_{22}^{\prime}$
(A8) $\forall h_{12}^{\prime} \in \llbracket \tau_{12}^{\prime} \rrbracket_{\mathcal{E}^{\prime}\left[x \mapsto h_{1}^{\prime}\right]} \cdot \llbracket \tau_{22}^{\prime} \rrbracket_{\mathcal{E}^{\prime}\left[x \mapsto h_{1}, y \mapsto h_{12}^{\prime}\right]} \subseteq \llbracket \tau_{22} \rrbracket_{\mathcal{E}\left[x \mapsto h_{1}, y \mapsto h_{12}^{\prime}\right]}$
Let $\tau_{2}^{\prime}=\tau_{22}^{\prime}\left[y \mapsto \tau_{12}^{\prime}\right]$. (1) follows by T-Seq.
For (2), we must show that $\left.\forall h_{1}^{\prime} \in \llbracket \tau_{1}^{\prime}\right]_{\mathcal{E}^{\prime}} \cdot \llbracket \tau_{2}^{\prime} \rrbracket_{\mathcal{E}^{\prime}\left[x \mapsto h_{1}^{\prime}\right]} \subseteq \llbracket \tau_{2} \rrbracket_{\mathcal{E}\left[x \mapsto h_{1}^{\prime}\right]}$
Let $h_{1}^{\prime} \in \llbracket \tau_{1}^{\prime} \rrbracket \mathcal{E}^{\prime}$ be arbitrary.

$$
\begin{aligned}
& \llbracket \tau_{22}^{\prime}\left[y \mapsto \tau_{12}^{\prime}\right] \rrbracket_{\mathcal{E}^{\prime}\left[x \mapsto h_{1}^{\prime}\right]} \subseteq \llbracket \tau_{22}\left[y \mapsto \tau_{12}\right] \rrbracket_{\mathcal{E}\left[x \mapsto h_{1}^{\prime}\right]} \\
& \Leftrightarrow \bigcup_{\left.h_{12}^{\prime} \in \llbracket \tau_{12}^{\prime}\right]_{\mathcal{E}^{\prime}\left[x \mapsto h_{1}^{\prime}\right]}} \llbracket \tau_{22}^{\prime} \rrbracket_{\mathcal{E}}\left[x \mapsto h_{1}^{\prime}, y \mapsto h_{12}^{\prime}\right] \\
& \bigcup_{\left.h_{12} \in \llbracket \tau_{12}\right]_{\mathcal{E}\left[x \mapsto h_{1}^{\prime}\right]}} \llbracket \tau_{22} \rrbracket_{\mathcal{E}\left[x \mapsto h_{1}^{\prime}, y \mapsto h_{12}\right]}
\end{aligned}
$$

The result follows by (A6), (A8) and set theory.

## Case T-Skip:

By inversion of rule T-Skip, we get
(A1) $c=s k i p$
(A2) $\tau_{2}=\left\{y: \tau_{1} \mid y \equiv x\right\}$
Let $\tau_{2}^{\prime}=\left\{y: \tau_{1}^{\prime} \mid y \equiv x\right\}$. (1) follows by T-Skip.
To show (2), let $h^{\prime} \in \llbracket \tau_{1}^{\prime} \rrbracket \mathcal{E}^{\prime}$ be an arbitrary heap.

$$
\begin{align*}
& \llbracket\left\{y: \tau_{1}^{\prime} \mid y \equiv x\right\} \rrbracket_{\mathcal{E}^{\prime}\left[x \mapsto h^{\prime}\right]} \subseteq \llbracket\left\{y: \tau_{1} \mid y \equiv x\right\} \rrbracket_{\mathcal{E}\left[x \mapsto h^{\prime}\right]} \\
\Leftrightarrow & \left\{h^{\prime}\right\} \subseteq \llbracket\left\{y: \tau_{1} \mid y \equiv x\right\} \rrbracket_{\mathcal{E}\left[x \mapsto h^{\prime}\right]} \\
\Leftrightarrow & \left\{h^{\prime}\right\} \subseteq\left\{h^{\prime}\right\} \tag{A}
\end{align*}
$$

## Case T-Sub:

By inversion of rule T-Sub, we get
$\left(\mathrm{A}_{1}\right) \Gamma \vdash c:\left(x: \tau_{3}\right) \rightarrow \tau_{4}$
(A2) $\Gamma \vdash \tau_{1}<: \tau_{3}$
(А3) $\Gamma, x: \tau_{1} \vdash \tau_{4}<: \tau_{2}$
By assumption $\llbracket \tau_{1}^{\prime} \rrbracket_{\mathcal{E}} \subseteq \llbracket \tau_{1} \rrbracket_{\mathcal{E}}$ and from (A2) follows that $\llbracket \tau_{1} \rrbracket_{\mathcal{E}} \subseteq \llbracket \tau_{3} \rrbracket_{\mathcal{E}}$ and thus $\llbracket \tau_{1}^{\prime} \rrbracket_{\mathcal{E}^{\prime}} \subseteq \llbracket \tau_{3} \rrbracket_{\mathcal{E}}$. By IH, there exists $\tau_{4}^{\prime}$ such that
(A4) $\Gamma^{\prime} \vdash c:\left(x: \tau_{1}^{\prime}\right) \rightarrow \tau_{4}^{\prime}$
(A5) $\forall h^{\prime} \in \llbracket \tau_{1}^{\prime} \rrbracket_{\mathcal{E}^{\prime}} \cdot \llbracket \tau_{4}^{\prime} \rrbracket_{\mathcal{E}^{\prime}\left[x \leftrightarrow h^{\prime}\right]} \subseteq \llbracket \tau_{4} \rrbracket_{\mathcal{E}\left[x \mapsto h^{\prime}\right]}$

Let $\tau_{2}^{\prime}=\tau_{4}^{\prime}$. (1) follows by (A4).
For (1), we have to show that $\forall h^{\prime} \in \llbracket \tau_{1}^{\prime} \rrbracket_{\mathcal{E}^{\prime}} \cdot \llbracket \tau_{4}^{\prime} \rrbracket_{\mathcal{E}^{\prime}}\left[x \mapsto h^{\prime}\right] \subseteq \llbracket \tau_{2} \rrbracket_{\mathcal{E}\left[x \mapsto h^{\prime}\right]}$, which follows by ( $\mathrm{A}_{3}$ ) and ( $\mathrm{A}_{5}$ ) and by set theory.

Lemma A. 50 (Formulae Preservation). If $\Gamma ; \tau \vdash \varphi: \mathbb{B}$ and $\mathcal{E} \vDash \Gamma$ and $(I, O, H) \vDash_{\mathcal{E}} \tau$ and $\langle I, O, H, \varphi\rangle \rightarrow \varphi^{\prime}$ then $\Gamma ; \tau \vdash \varphi^{\prime}: \mathbb{B}$.

Proof. By induction on a derivation of $\Gamma ; \tau \vdash \varphi: \mathbb{B}$.
Lemma A. 51 (Semantic Formulae Preservation). If $\Gamma ; \tau \vdash \varphi: \mathbb{B}$ and $\mathcal{E} \vDash \Gamma$ and $(I, O, H) \vDash_{\mathcal{E}} \tau$ and $\langle I, O, H, \varphi\rangle \rightarrow \varphi^{\prime}$ then

$$
\llbracket \varphi[x / \text { heap }] \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]}=\llbracket \varphi^{\prime}[x / \text { heap }] \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]}
$$

Proof. By induction on a derivation of $\Gamma ; \tau \vdash \varphi: \mathbb{B}$.
Lemma A. 52 (Expression Preservation). If $\Gamma ; \tau \vdash e: \sigma$ and $\mathcal{E} \vDash \Gamma$ and $(I, O, H) \vDash \mathcal{E} \tau$ and $\langle I, O, H, e\rangle \rightarrow e^{\prime}$ then $\Gamma ; \tau \vdash e^{\prime}: \sigma$.

Proof. By induction on a typing derivation of $\Gamma ; \tau \vdash e: \sigma$.
Lemma A. 53 (Semantic Expression Preservation). If $\Gamma ; \tau \vdash e: \sigma$ and $\mathcal{E} \vDash \Gamma$ and $(I, O, H) \vDash_{\mathcal{E}} \tau$ and $\langle I, O, H, e\rangle \rightarrow e^{\prime}$ then

$$
\llbracket e[x / \text { cmdVar }] \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]}=\llbracket e^{\prime}[x / \text { heap }] \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]}
$$

Proof. By induction on a typing derivation of $\Gamma ; \tau \vdash e: \sigma$.
Theorem A. 54 (Preservation). If $\Gamma \vdash c:\left(x: \tau_{1}\right) \rightarrow \tau_{2},\langle I, O, H, c\rangle \rightarrow\left\langle I^{\prime}, O^{\prime}, H^{\prime}, c^{\prime}\right\rangle$, $\mathcal{E} \vDash \Gamma$ and $(I, O, H) \vDash \mathcal{E} \tau_{1}$, then $\exists \Gamma^{\prime}, \mathcal{E}^{\prime}, x^{\prime}, \tau_{1}^{\prime}, \tau_{2}^{\prime} . \Gamma^{\prime} \vdash c^{\prime}:\left(x^{\prime}: \tau_{1}^{\prime}\right) \rightarrow \tau_{2}^{\prime}$ and $\mathcal{E}^{\prime} \vDash \Gamma^{\prime}$ and $\Gamma \subseteq \Gamma^{\prime}$ and $\mathcal{E} \subseteq \mathcal{E}^{\prime}$ and $\left(I^{\prime}, O^{\prime}, H^{\prime}\right) \vDash \mathcal{E}^{\prime} \quad \tau_{1}^{\prime}$ and $\llbracket \tau_{2}^{\prime} \rrbracket \mathcal{E}^{\prime}\left[x^{\prime} \mapsto\left(I^{\prime}, O^{\prime}, H^{\prime}\right)\right] \subseteq$ $\llbracket \tau_{2} \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]}$

Proof. By induction on a derivation of $\Gamma \vdash c:\left(x: \tau_{1}\right) \rightarrow \tau_{2}$ with case analysis on the last rule used. We refer to the proof goals as follows:
(1) $\Gamma^{\prime} \vdash c^{\prime}:\left(x^{\prime}: \tau_{1}^{\prime}\right) \rightarrow \tau_{2}^{\prime}$
(2) $\mathcal{E}^{\prime} \vDash \Gamma^{\prime}$
(3) $\Gamma \subseteq \Gamma^{\prime}$
(4) $\mathcal{E} \subseteq \mathcal{E}^{\prime}$
(5) $\left(I^{\prime}, O^{\prime}, H^{\prime}\right){\vDash \mathcal{E}^{\prime}}^{\tau_{1}^{\prime}}$
(6) $\llbracket \tau_{2}^{\prime} \rrbracket_{\mathcal{E}^{\prime}\left[x^{\prime} \mapsto\left(I^{\prime}, O^{\prime}, H^{\prime}\right)\right]} \subseteq \llbracket \tau_{2} \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]}$

General assumptions:
(A) $\Gamma \vdash c:\left(x: \tau_{1}\right) \rightarrow \tau_{2}$
(B) $\langle I, O, H, c\rangle \rightarrow\left\langle I^{\prime}, O^{\prime}, H^{\prime}, c^{\prime}\right\rangle$
(C) $\mathcal{E} \vDash \Gamma$
(D) $(I, O, H) \vDash_{\mathcal{E}} \tau_{1}$

## Case T-Add:

By inversion of rule T-ADD, we get
(A1) $c=\operatorname{add}(\iota)$
(A2) Excludes $\Gamma \tau_{1} \iota$
(A3) $\tau_{2}=\Sigma y:\left\{z: \tau_{1} \mid z \equiv x\right\} .\left\{z: \iota \mid z \cdot p k t_{\text {in }}=z . p k t_{\text {out }}=\langle \rangle \wedge z . \iota=v\right\}$
Only evaluation rule E-ADD applies to $c$ :
(A4) $\iota \notin \operatorname{dom}(H)$
(A5) $\mathcal{H} \mathcal{T}(\iota)=\eta$
(A6) init $=v$.
(A7) $I^{\prime}=I$ and $O^{\prime}=O$ and $H^{\prime}=H[\iota \mapsto v]$ and $c^{\prime}=s k i p$
Let $\Gamma^{\prime}=\Gamma, x: \tau_{1}$ and $\mathcal{E}^{\prime}=\mathcal{E}[x \mapsto(I, O, H)]$.
Let $\tau_{1}^{\prime}=\Sigma y:\left\{z: \tau_{1} \mid z \equiv x\right\} .\left\{z: \iota \mid z . p k t_{\text {in }}=z . p k t_{\text {out }}=\langle \rangle \wedge z . \iota=v\right\}$ and $\tau_{2}^{\prime}=\left\{w: \Sigma y:\left\{z: \tau_{1} \mid z \equiv x\right\} .\left\{z: \iota \mid z . p k t_{i n}=z . p k t_{\text {out }}=\langle \rangle \wedge z . \iota=v\right\} \mid w \equiv x^{\prime}\right\}$.
(1) follows by T-Skip and (2) follows by assumptions (C) and (D) and Lemma A.30. (3) and (4) are immediate.

To show (5), we must show that $(I, O, H[\iota \mapsto v]) \vDash_{\mathcal{E}[x \mapsto(I, O, H)]} \Sigma y:\left\{z: \tau_{1} \mid z \equiv\right.$ $x\} .\left\{z: \iota \mid z . p k t_{\text {in }}=z . p k t_{\text {out }}=\langle \rangle \wedge z . l=v\right\}$. By Ent-Sigma, we must show that
(5.1) $(I, O, H) \vDash_{\mathcal{E}[x \mapsto(I, O, H)]}\left\{z: \tau_{1} \mid z \equiv x\right\}$ and
(5.2) $\left(\rangle,\langle \rangle,[\iota \mapsto v]) \vDash_{\mathcal{E}[x \mapsto(I, O, H) y \mapsto(I, O, H)]}\left\{z: \iota \mid z . p k t_{\text {in }}=\langle \rangle \wedge z . p k t_{\text {out }}=\right.\right.$ $\rangle \wedge z . \iota=v\}$.
(5.1) follows by Ent-Refine and (D). To show (5.2), by Ent-Refine, we must show that
(5.2.1) $\left(\rangle,\langle \rangle,[\iota \mapsto v]) \vDash_{\mathcal{E}[y \mapsto(I, O, H)]}\left\{z: T \mid z . \iota . v a l i d \wedge \wedge_{\kappa \in \operatorname{dom}(\mathcal{H} \mathcal{T})} \neg \mathcal{K}\right.\right.$. valid $\}$ and
(5.2.2) $\llbracket z . p k t_{\text {in }}=\langle \rangle \wedge z . p k t_{\text {out }}=\langle \rangle \wedge z . \iota=v \rrbracket_{\mathcal{E}[y \mapsto(I, O, H), z \mapsto(\langle \rangle,\langle \rangle,[\mapsto \sim])]}=$ true
(5.2.1) follows by Ent-Refine, Ent-Top and the semantics of formulae. (5.2.2) follows from the semantics of formulae.
(6) follows by

$$
\begin{aligned}
& \llbracket \Sigma y:\left\{z: \tau_{1} \mid z \equiv x\right\} \cdot\left\{z: \iota \mid z \cdot p k t_{\text {in }}=z \cdot p k t_{\text {out }}=\langle \rangle \wedge z . \iota=v\right\} \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]} \\
&=\{(I, O, H[\iota \mapsto v])\} \\
&=\left\{h \mid h \in \llbracket \Sigma y:\left\{z: \tau_{1} \mid z \equiv x\right\} \cdot\left\{z: \iota \mid z \cdot p k t_{\text {in }}=z \cdot p k t_{\text {out }}=\langle \rangle \wedge\right.\right. \\
&z \cdot l=v\} \rrbracket_{\mathcal{E}\left[x \mapsto(I, O, H), x^{\prime} \mapsto(I, O, H[\iota \mapsto v])\right]} \wedge \\
&= \llbracket w \equiv x^{\prime} \rrbracket_{\left.\mathcal{E}\left[x \mapsto(I, O, H), x^{\prime} \mapsto(I, O, H[\iota \mapsto v]), w \mapsto h\right]\right\}} \\
&= \llbracket w: \Sigma y:\left\{z: \tau_{1} \mid z \equiv x\right\} \cdot\left\{z: \iota \mid z \cdot p k t_{\text {in }}=z \cdot p k t_{\text {out }}=\langle \rangle \wedge z . \iota=v\right\} \mid \\
&\left.w \equiv x^{\prime}\right\} \rrbracket_{\mathcal{E}\left[x \mapsto(I, O, H), x^{\prime} \mapsto(I, O, H[\mapsto \nu])\right]}
\end{aligned}
$$

## Case T-Ascribe:

By inversion of rule T-Ascribe, we get
(A1) $c=c_{a}$ as $\left(x: \tau_{1}\right) \rightarrow \tau_{2}$
(A2) $\Gamma \vdash c_{a}:\left(x: \tau_{1}\right) \rightarrow \tau_{2}$
There is one evaluation rule that applies to $c$, E-Ascribe, so $I^{\prime}=I$ and $O^{\prime}=O$ and $H^{\prime}=H$ and $c^{\prime}=c_{a}$. Let $\Gamma^{\prime}=\Gamma, \mathcal{E}^{\prime}=\mathcal{E}, \tau_{1}^{\prime}=\tau_{1}$ and $\tau_{2}^{\prime}=\tau_{2}$. (1) follows by (A2), (2) follows by assumption (C). (3) and (4) are immediate. (5) follows by assumption (D) and (6) follows from the equality of $\tau_{2}$ and $\tau_{2}^{\prime}$, which itself follows by reflexivity.

## Case T-Extract:

By inversion of rule T-Extract, we get
(A1) $c=\operatorname{extract}(\iota)$
(A2) $\Gamma \vdash \operatorname{sizeof}_{p k t_{i n}}\left(\tau_{1}\right)=\operatorname{sizeof}(\iota)$
(A3) $\tau_{2}=\Sigma y:\left\{z: \iota \mid z \cdot p k t_{\text {in }}=z \cdot p k t_{\text {out }}=\langle \rangle\right\} \cdot\left\{z: \operatorname{chomp}\left(\tau_{1}, \iota, y\right) \mid\right.$

$$
\text { y.ı@z.pkt } \left.{ }_{\text {in }}=x . p k t_{\text {in }} \wedge z . p k t_{\text {out }}=x . p k t_{\text {out }} \wedge z \equiv, x\right\}
$$

Only evaluation rule E-Extract applies to $c$ :
(A4) deserialize $\mathcal{H \mathcal { T }}_{(ı)}(I)=\left(v, I^{\prime}\right)$
(A5) $O^{\prime}=O$
(A6) $H^{\prime}=H[\iota \mapsto v]$
(A7) $c^{\prime}=s k i p$
Let $\Gamma^{\prime}=\Gamma, x: \tau_{1}, \mathcal{E}^{\prime}=\mathcal{E}[x \mapsto(I, O, H)]$.
Let $\tau_{1}^{\prime}=\Sigma y:\left\{z: \iota \mid z . p k t_{i n}=z \cdot p k t_{\text {out }}=\langle \rangle\right\} \cdot\left\{z: \operatorname{chomp}\left(\tau_{1}, l, y\right) \mid\right.$ $\left.y . \iota @ z . p k t_{\text {in }}=x . p k t_{\text {in }} \wedge z . p k t_{\text {out }}=x . p k t_{\text {out }} \wedge z \equiv, x\right\}$.
Let $\tau_{2}^{\prime}=\left\{v: \Sigma y:\left\{z: \iota \mid z . p k t_{\text {in }}=z . p k t_{\text {out }}=\langle \rangle\right\} .\left\{z: \operatorname{chomp}\left(\tau_{1}, \iota, y\right) \mid\right.\right.$ $\left.\left.y . \iota @ z . p k t_{\text {in }}=x . p k t_{\text {in }} \wedge z . p k t_{\text {out }}=x . p k t_{\text {out }} \wedge z \equiv, x\right\} \mid v \equiv x\right\}$.
(1) follows by T-SKiP and (2) follows by assumptions (C) and (D) and Lemma A.30. (3) and (4) are immediate.

To show (5), we must show that $\left(I^{\prime}, O, H[\iota \mapsto v]\right) \vDash_{\mathcal{E}[x \mapsto(I, O, H)]} \Sigma y:\{z$ : $\left.\iota \mid z \cdot p k t_{\text {in }}=z \cdot p k t_{\text {out }}=\langle \rangle\right\} .\left\{z: \operatorname{chomp}\left(\tau_{1}, \iota, y\right) \mid y . \iota @ z . p k t_{\text {in }}=x \cdot p k t_{\text {in }} \wedge\right.$ $\left.z . p k t_{\text {out }}=x \cdot p k t_{\text {out }} \wedge z \equiv, x\right\}$.
By Ent-Sigma, we must show that
(5.1) $\left(\rangle,\langle \rangle,[\iota \mapsto v]) \vDash_{\mathcal{E}[x \mapsto(I, O, H)]}\left\{z: \iota \mid z \cdot p k t_{\text {in }}=z \cdot p k t_{\text {out }}=\langle \rangle\right\}\right.$, which follows by Ent-Refine and the semantics of types and
(5.2) $\left(I^{\prime}, O, H\right) \vDash_{\mathcal{E}[x \mapsto(I, O, H), y \mapsto(( \rangle,\langle \rangle,[\iota \mapsto v])]}\left\{z: \operatorname{chomp}\left(\tau_{1}, l, y\right) \mid y . l @ z . p k t_{i n}=\right.$ $\left.x \cdot p k t_{\text {in }} \wedge z \cdot p k t_{\text {out }}=x \cdot p k t_{\text {out }} \wedge z \equiv, x\right\}$
By Ent-Refine, we must show that
(5.2.1) $\left(I^{\prime}, O, H\right) \vDash_{\mathcal{E}[x \mapsto(I, O, H), y \mapsto(\langle \rangle,\langle \rangle,[\iota \mapsto v])]} \operatorname{chomp}\left(\tau_{1}, \iota, y\right)$ By Lemma A.28, it is sufficient to show that

$$
\left(I^{\prime}, O, H\right) \in \llbracket \operatorname{chomp}\left(\tau_{1}, l, y\right) \rrbracket_{\mathcal{E}[x \mapsto(I, O, H), y \mapsto(\langle \rangle,\langle \rangle,[\mapsto \nu])]}
$$

By assumption (D) and by Lemma A. 28 follows $(I, O, H) \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}}$. By Lemma A.44, there exists some heap

$$
h \in \llbracket \operatorname{chomp}\left(\tau_{1}, \iota, y\right) \rrbracket_{\mathcal{E}[x \mapsto(I, O, H), y \mapsto(\langle \rangle,\langle \rangle,[\lfloor\mapsto \nu])]}
$$

such that $h=\operatorname{chomp}^{\Downarrow}((I, O, H)$, sizeof $(\iota))$.
From the definition of chomp ${ }^{\Downarrow}$ follows that

$$
\operatorname{chomp}^{\Downarrow}((I, O, H), \operatorname{sizeof}(\iota))=\left(I^{\prime \prime}, O, H\right)
$$

where $I^{\prime \prime}=I[\operatorname{sizeof}(\iota):]=I^{\prime}$.
(5.2.2) $\llbracket y . \iota @ z . p k t_{i n}=x . p k t_{\text {in }} \wedge z . p k t_{\text {out }}=x . p k t_{\text {out }} \wedge$ $z \equiv \wedge]_{\mathcal{E}\left[x \mapsto(I, O, H), y \mapsto(\langle \rangle,\langle \rangle,[\mapsto \nu]), z \mapsto\left(I^{\prime}, O, H\right)\right]}$, which follows by the definition of deserialize and the semantics of formulae and expressions.

Finally, we must show that $\llbracket \tau_{2} \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]} \subseteq \llbracket \tau_{2}^{\prime} \rrbracket_{\mathcal{E}\left[x \mapsto\left(I^{\prime}, O, H[\mapsto \nu]\right)\right]}$.

$$
\begin{gathered}
\llbracket \tau_{2} \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]}=\llbracket \Sigma y:\left\{z: \iota \mid z \cdot p k t_{\text {in }}=z \cdot p k t_{\text {out }}=\langle \rangle\right\} . \\
\left\{z: \operatorname{chomp}\left(\tau_{1}, \iota, y\right) \mid y . \iota @ z \cdot p k t_{\text {in }}=x \cdot p k t_{\text {in }} \wedge z \cdot p k t_{\text {out }}=x \cdot p k t_{\text {out }} \wedge\right. \\
z \equiv, x\} \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]} \\
=\left\{h_{1}++h_{2} \mid h_{1} \in \llbracket\left\{z: \iota \mid z . p k t_{\text {in }}=z \cdot p k t_{\text {out }}=\langle \rangle\right\} \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]} \wedge\right. \\
h_{2} \in \llbracket\left\{z: \operatorname{chomp}\left(\tau_{1}, \iota, y\right) \mid y \cdot \iota @ z \cdot p k t_{\text {in }}=x \cdot p k t_{\text {in }} \wedge\right. \\
\left.\left.z . p k t_{\text {out }}=x \cdot p k t_{\text {out }} \wedge z \equiv, x\right\} \rrbracket_{\mathcal{E}\left[x \mapsto(I, O, H), y \mapsto h_{1}\right]}\right\} \\
=\left\{h_{1}++h_{2} \mid h_{1} \in \llbracket\left\{z: \iota \mid z . p k t_{\text {in }}=z \cdot p k t_{\text {out }}=\langle \rangle\right\} \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]} \wedge\right. \\
h_{2} \in\left\{h_{2}^{\prime} \in \llbracket \operatorname{chomp}\left(\tau_{1}, \iota, y\right) \rrbracket_{\mathcal{E}\left[x \mapsto(I, O, H), y \mapsto h_{1}\right]} \wedge\right. \\
\llbracket y . \iota @ z \cdot p k t_{\text {in }}=x \cdot p k t_{\text {in }} \wedge z \cdot p k t_{\text {out }}=x . p k t_{\text {out }} \wedge \\
\left.\left.z \equiv, x \rrbracket_{\mathcal{E}\left[x \mapsto(I, O, H), y \mapsto h_{1}, z \mapsto h_{2}^{\prime}\right]}\right\}\right\}
\end{gathered}
$$

By Lemma A. 44

$$
\begin{gathered}
=\left\{h_{1}++h_{2} \mid h_{1} \in \llbracket\left\{z: \iota \mid z \cdot p k t_{\text {in }}=z \cdot p k t_{\text {out }}=\langle \rangle\right\} \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]} \wedge\right. \\
h_{2} \in\left\{h \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]} \wedge h_{2}^{\prime}=\operatorname{chomp}^{\Downarrow}(h, \operatorname{sizeof}(\iota)) \wedge\right. \\
\llbracket y \cdot \iota @ z \cdot p k t_{\text {in }}=x \cdot p k t_{\text {in }} \wedge z \cdot p k t_{\text {out }}=x \cdot p k t_{\text {out }} \wedge \\
\left.\left.z \equiv, x \rrbracket_{\mathcal{E}\left[x \mapsto(I, O, H), y \mapsto h_{1}, z \mapsto h_{2}^{\prime}\right]}\right\}\right\}
\end{gathered}
$$

By definition of chomp ${ }^{\Downarrow}$ and semantics of types

$$
\begin{aligned}
& =\left\{\left(I^{\prime}, O, H[\iota \mapsto v]\right)\right\} \\
& =\left\{h \mid h \in \llbracket \Sigma y:\left\{z: \iota \mid z \cdot p k t_{\text {in }}=z . p k t_{\text {out }}=\langle \rangle\right\} .\right. \\
& \left\{z: \operatorname{chomp}\left(\tau_{1}, \iota, y\right) \mid y . \iota @ z . p k t_{i n}=x . p k t_{i n} \wedge\right. \\
& \left.z . p k t_{\text {out }}=x . p k t_{\text {out }} \wedge z \equiv x\right\} \rrbracket_{\mathcal{E}\left[x \mapsto(I, O, H), x^{\prime} \mapsto\left(I^{\prime}, O, H[\lfloor\mapsto])\right]\right.} \wedge \\
& \llbracket v \equiv x^{\prime} \rrbracket_{\mathcal{E}}\left[x \mapsto(I, O, H), x^{\prime} \mapsto\left(I^{\prime}, O, H[\lfloor\nu]), v \mapsto h\right]\right\} \\
& =\llbracket\left\{v: \Sigma y:\left\{z: \iota \mid z \cdot p k t_{\text {in }}=z \cdot p k t_{\text {out }}=\langle \rangle\right\} .\right. \\
& \left\{z: \operatorname{chomp}\left(\tau_{1}, \iota, y\right) \mid y . \iota @ z . p k t_{i n}=x . p k t_{i n} \wedge\right. \\
& \left.z \cdot p k t_{\text {out }}=x . p k t_{\text {out }} \wedge z \equiv, x\right\} \mid \\
& v \equiv x\} \rrbracket_{\mathcal{E}\left[x \mapsto(I, O, H), x^{\prime} \mapsto\left(I^{\prime}, O, H[\mapsto v]\right)\right]}
\end{aligned}
$$

This shows (6) and concludes the case.

## Case T-If:

By inversion of rule T-IF, we get
(A1) $\Gamma ; \tau_{1} \vdash \varphi: \mathbb{B}$
$\left(\mathrm{A}_{2}\right) \Gamma \vdash c_{1}:\left(x:\left\{y: \tau_{1} \mid \varphi[y /\right.\right.$ heap $\left.\left.]\right\}\right) \rightarrow \tau_{12}$
(A3) $\Gamma \vdash c_{2}:\left(x:\left\{y: \tau_{1} \mid \neg \varphi[y /\right.\right.$ heap $\left.\left.]\right\}\right) \rightarrow \tau_{22}$
(A4) $\tau_{2}=\left\{y: \tau_{12} \mid \varphi[x /\right.$ heap $\left.]\right\}+\left\{y: \tau_{22} \mid \neg \varphi[x /\right.$ heap $\left.]\right\}$
T-If There are three evaluation rules that apply to $c$.

## Subcase E-IF:

(A5) $c^{\prime}=i f\left(\varphi^{\prime}\right) c_{1}$ else $c_{2}$
(A6) $I^{\prime}=I, O^{\prime}=O, H^{\prime}=H$
Let $\Gamma^{\prime}=\Gamma, \mathcal{E}^{\prime}=\mathcal{E}, x^{\prime}=x$,
$\tau_{1}^{\prime}=\tau_{1}$ and
$\tau_{2}^{\prime}=\left\{y: \tau_{12} \mid \varphi^{\prime}[x /\right.$ heap $\left.]\right\}+\left\{y: \tau_{22} \mid \neg \varphi^{\prime}[x /\right.$ heap $\left.]\right\}$.
By Lemma A.50, $\Gamma ; \tau_{1} \vdash \varphi^{\prime}: \mathbb{B}$.
By Lemma A.51, $\llbracket \varphi\left[x /\right.$ heap $\rrbracket \rrbracket_{\mathcal{E}[x \rightarrow(I, O, H)]}=\llbracket \varphi^{\prime}[x /$ heap $] \rrbracket_{\mathcal{E}[x \rightarrow(I, O, H)]}$.
From (A2) follows that $\Gamma \vdash c_{1}:\left(x:\left\{y: \tau_{1} \mid \varphi^{\prime}[y /\right.\right.$ heap $\left.\left.]\right\}\right) \rightarrow \tau_{12}$ and from (A3) follows that $\Gamma \vdash c_{2}:\left(x:\left\{y: \tau_{1} \mid \neg \varphi^{\prime}[y /\right.\right.$ heap $\left.\left.]\right\}\right) \rightarrow \tau_{22}$ (1) follows by T-If.
(2) follows by assumption (C), (3) and (4) are immediate.
(5) follows by assumption (D).
(6) follows together with the assumption

$$
\llbracket \varphi[x / \text { heap }] \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]}=\llbracket \varphi^{\prime}[x / \text { heap }] \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]}
$$

from the equality $\llbracket \tau_{2}^{\prime} \rrbracket_{\mathcal{E}}\left[x \rightarrow\left(I^{\prime}, O^{\prime}, H^{\prime}\right)\right]=\llbracket \tau_{2} \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]}$.

$$
\begin{aligned}
& \llbracket \tau_{2} \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]} \\
= & \llbracket\left\{y: \tau_{12} \mid \varphi[x / \text { heap }]\right\}+\left\{y: \tau_{22} \mid \neg \varphi[x / \text { heap }]\right\} \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]} \\
= & \llbracket\left\{y: \tau_{12} \mid \varphi^{\prime}[x / \text { heap }]\right\}+\left\{y: \tau_{22} \mid \neg \varphi^{\prime}[x / \text { heap }]\right\} \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]} \\
= & \llbracket \tau_{2}^{\prime} \rrbracket_{\mathcal{E}}\left[x \mapsto\left(I^{\prime}, O^{\prime}, H^{\prime}\right)\right]
\end{aligned}
$$

## Subcase E-If True:

$$
\begin{aligned}
& \text { (A7) } c^{\prime}=c_{1} \\
& \text { (A8) } I^{\prime}=I, O^{\prime}=O, H^{\prime}=H \\
& \text { (A9) } \varphi=\text { true } \\
& \text { Let } \Gamma^{\prime}=\Gamma, \mathcal{E}^{\prime}=\mathcal{E}, x^{\prime}=x \\
& \tau_{1}^{\prime}=\left\{y: \tau_{1} \mid \varphi[y / \text { heap }]\right\}=\left\{y: \tau_{1} \mid \text { true }\right\}=\tau_{1} \text { and } \\
& \tau_{2}^{\prime}=\tau_{12} .
\end{aligned}
$$

(1) holds by assumption (A2), (2) holds by assumption (C). (3) and (4) are immediate and (5) $\left(I^{\prime}, O^{\prime}, H^{\prime}\right) \vDash \mathcal{E}^{\prime} \tau_{1}^{\prime} \Leftrightarrow(I, O, H) \vDash \mathcal{E} \quad \tau_{1}$ holds by assumption (D).
(6) follows from

$$
\begin{aligned}
& \llbracket\left\{y: \tau_{12} \mid \varphi[x / \text { heap }]\right\}+\left\{y: \tau_{22} \mid \neg \varphi[x / \text { heap }]\right\} \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]} \\
= & \llbracket\left\{y: \tau_{12} \mid \text { true }\right\}+\left\{y: \tau_{22} \mid \text { false }\right\} \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]} \\
= & \llbracket \tau_{12} \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]} \\
= & \llbracket \tau_{2}^{\prime} \rrbracket_{\mathcal{E}^{\prime}\left[x \mapsto\left(I^{\prime}, O^{\prime}, H^{\prime}\right)\right]}
\end{aligned}
$$

## Subcase E-If False:

Symmetric to previous subcase.

## Case T-Mod:

By inversion of rule T-Mod, we get
(A1) $c=\iota . f:=e$
(A2) Includes $\Gamma \tau_{1} \iota$
(A3) $\mathcal{F}(t, f)=\mathrm{BV}$
(A4) $\Gamma ; \tau_{1} \vdash e: B V$
(A5) $\tau_{2}=\left\{y: \top \mid \varphi_{p k t} \wedge \varphi_{\imath} \wedge \varphi_{f} \wedge y . \iota . f=e[x /\right.$ heap $\left.]\right\}$
(A6) $\varphi_{p k t} \triangleq y \cdot p k t_{i n}=x \cdot p k t_{\text {in }} \wedge y \cdot p k t_{\text {out }}=x \cdot p k t_{\text {out }}$
(A7) $\varphi_{\imath} \triangleq \forall \kappa \in \operatorname{dom}(\mathcal{H} \mathcal{T}) . \iota \neq \kappa \rightarrow y . \kappa=x . \kappa$
(A8) $\varphi_{f} \triangleq \forall g \in \operatorname{dom}(\mathcal{H} \mathcal{T}(\iota)) \cdot f \neq g \rightarrow y . \iota . g=x . \iota . g$
There are two evaluation rule that apply to $c$.

## Subcase E-Mod:

(A9) $H(\iota)=r$
(A10) $r^{\prime} \triangleq\{\mathrm{r}$ with $f=v\}$
(A11) $\left(I^{\prime}, O^{\prime}, H^{\prime}\right)=\left(I, O, H\left[\iota \mapsto r^{\prime}\right]\right)$
(A12) $c^{\prime}=s k i p$
(A13) $e=v$
(A14) $\tau_{2}=\left\{y: T \mid \varphi_{p k t} \wedge \varphi_{\iota} \wedge \varphi_{f} \wedge y . l . f=v[x /\right.$ heap $\left.]\right\}$

Let $\Gamma^{\prime}=\Gamma, x: \tau_{1}$ and $\mathcal{E}^{\prime}=\mathcal{E}[x \mapsto(I, O, H)]$ and $\tau_{1}^{\prime}=\tau_{2}$ and $\tau_{2}^{\prime}=\left\{z: \tau_{2} \mid\right.$ $z \equiv x\}$.
(1) follows by T-Skip and T-Sub. (2) follows by assumptions (C) and (D) and Lemma A.30. (3) and (4) are immediate.
To show (5), we must show that $\left(I, O, H\left[\iota \mapsto r^{\prime}\right]\right) \vDash_{\mathcal{E}[x \mapsto(I, O, H)]} \top$ and $\llbracket \varphi_{p k t} \wedge \varphi_{\iota} \wedge \varphi_{f} \wedge y . \iota . f=v[x /$ heap $] \rrbracket_{\mathcal{E}\left[x \mapsto(I, O, H), y \mapsto\left(I, O, H\left[\iota \leftrightarrow r^{\prime}\right]\right)\right]}=$ true. Since $\mathcal{E}(x)(I)=\mathcal{E}(y)(I)$ and $\mathcal{E}(x)(O)=\mathcal{E}(y)(O), \varphi_{p k t}$ holds. Similarly, $\mathcal{E}(x)(H)=\mathcal{E}(y)(H)$ in every aspect, except for field $f$ of instance $\iota$, so $\varphi_{f}$ and $\varphi_{\text {inst }}$ also hold. $y$.ı. $f=v[x /$ heap $]$ also holds, because

$$
\begin{aligned}
& \llbracket y . \iota . f \rrbracket_{\mathcal{E}\left[x \mapsto(I, O, H), y \mapsto\left(I, O, H\left[\mapsto r^{\prime}\right]\right)\right]}= \\
& \llbracket v[x / \text { heap }] \rrbracket_{\mathcal{E}\left[x \mapsto(I, O, H), y \mapsto\left(I, O, H\left[\mapsto r^{\prime}\right]\right)\right]} \\
\Leftrightarrow & v=\llbracket v[x / \text { heap }] \rrbracket_{[x \mapsto(I, O, H)]} \\
\Leftrightarrow & v=v
\end{aligned}
$$

To show

$$
\begin{aligned}
& \llbracket \tau_{2}^{\prime} \rrbracket \mathcal{E}^{\prime}\left[x^{\prime} \mapsto\left(I^{\prime}, O^{\prime}, H^{\prime}\right)\right] \subseteq \llbracket \tau_{2} \rrbracket \mathcal{E}[x \mapsto(I, O, H)] \\
\Leftrightarrow & \llbracket\left\{z: \tau_{2} \mid z \equiv x\right\} \rrbracket \mathcal{E}\left[x \mapsto(I, O, H), x^{\prime} \mapsto\left(I^{\prime}, O^{\prime}, H^{\prime}\right)\right] \subseteq \llbracket \tau_{2} \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]} \\
\Leftrightarrow & \left\{\left(I^{\prime}, O^{\prime}, H^{\prime}\right)\right\} \subseteq \llbracket \tau_{2} \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]} \\
\Leftrightarrow & \left\{\left(I, O, H\left[\iota \mapsto r^{\prime}\right]\right)\right\} \subseteq \llbracket \tau_{2} \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]}
\end{aligned}
$$

To show (6), let $h=\left(I, O, H\left[\iota \mapsto r^{\prime}\right]\right)$. Therefore,
(A15) $h\left(p k t_{i n}\right)=I$ and
(A16) $h\left(p k t_{\text {out }}\right)=O$ and
(A17) for all $\kappa \in \operatorname{dom}(\mathcal{H} \mathcal{T})$ such that $\kappa \neq \iota, h(\kappa)=H(\kappa)$
(A18) for all $g \in \operatorname{dom}(\mathcal{H} \mathcal{T}(\iota))$ such that $g \neq f, h(\iota)(g)=H(\imath)(g)$
(A19) $h(\iota)(f)=v[x /$ heap $]=v$
From the semantics of types follows that $h \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]}$.

## Subcase E-Modi:

(A20) $\langle I, O, H, t\rangle \rightarrow e^{\prime}$
(A21) $c^{\prime}=\iota . f:=e^{\prime}$
(A22) $I^{\prime}=I, O^{\prime}=O$ and $H^{\prime}=H$
Let $\Gamma^{\prime}=\Gamma$ and $\mathcal{E}^{\prime}=\mathcal{E}$ and $x^{\prime}=x$.
Let $\tau_{1}^{\prime}=\tau_{1}$ and
$\tau_{2}^{\prime}=\left\{y: T \mid \varphi_{p k t} \wedge \varphi_{\iota} \wedge \varphi_{f} \wedge y . \iota . f=e^{\prime}[x /\right.$ heap $\left.]\right\}$.
By semantic expression preservation (Lemma A.53), we know that $\llbracket e[x /$ heap $] \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]}=\llbracket e^{\prime}[x /$ heap $] \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]}$.
If $\Gamma ; \tau_{1} \vdash e: \mathrm{BV}$ and $\llbracket e[x /$ heap $] \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]}=\llbracket e^{\prime}[x /$ heap $] \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]}$ holds, then it must hold that $\Gamma ; \tau_{1} \vdash e^{\prime}: \mathrm{BV}$.
(1) follows by T-Mod, (2) follows from assumption (C). (3) and (4) are immediate. (5) follows from assumption (D).
(6) follows from

$$
\begin{aligned}
& \llbracket \tau_{2}^{\prime} \rrbracket_{\mathcal{E}}\left(\left[x^{\prime} \mapsto\left(I^{\prime}, O^{\prime}, H^{\prime}\right)\right]\right. \\
\Leftrightarrow & \llbracket_{2} \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]} \\
\Leftrightarrow & \llbracket \tau_{2}^{\prime} \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]} \subseteq \tau_{2} \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]} \\
\Leftrightarrow & \llbracket\left\{y: \top \mid \varphi_{p k t} \wedge \varphi_{\iota} \wedge \varphi_{f} \wedge y . \iota . f=e^{\prime}[x / \text { heap }]\right\} \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]} \subseteq \\
& \llbracket\left\{y: \top \mid \varphi_{p k t} \wedge \varphi_{\iota} \wedge \varphi_{f} \wedge y \text { y.ı. }=e[x / \text { heap }]\right\} \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]}
\end{aligned}
$$

together with assumption

$$
\llbracket e[x / \text { heap }] \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]}=\llbracket e^{\prime}[x / \text { heap }] \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]}
$$

and the semantics of types.

## Case T-Remit:

By inversion of rule T-Remit, we get
(A1) $c=$ remit
(A2) $\tau_{2}=\Sigma y:\left\{z: \tau_{1} \mid z \equiv x\right\} .\left\{z: \epsilon \mid z \cdot p k t_{i n}=\langle \rangle \wedge z . p k t_{o u t}=x . l\right\}$
( $\mathrm{A}_{3}$ ) Includes $\Gamma \tau_{1}$ !
There is only evaluation rule E-Remit that applies to $c$.
(A4) $l \in \operatorname{dom}(H)$
(A5) $\mathcal{H} \mathcal{T}(\iota)=\eta$
(A6) $\operatorname{serialize}_{\eta}(H(\iota))=b v$
(A7) $I^{\prime}=I, O^{\prime}=O @ b v, H^{\prime}=H, c^{\prime}=s k i p$
Let $\Gamma^{\prime}=\Gamma, x: \tau_{1}$ and $\mathcal{E}^{\prime}=\mathcal{E}[x \mapsto(I, O, H)]$.
Let $\tau_{1}^{\prime}=\Sigma y:\left\{z: \tau_{1} \mid z \equiv x\right\} .\left\{z: \epsilon \mid z . p k t_{\text {in }}=\langle \rangle \wedge z . p k t_{\text {out }}=x . l\right\}$ and $\tau_{2}^{\prime}=\left\{v: \Sigma y:\left\{z: \tau_{1} \mid z \equiv x\right\} .\left\{z: \epsilon \mid z \cdot p k t_{i n}=\langle \rangle \wedge z . p k t_{\text {out }}=x . \iota\right\} \mid v \equiv x^{\prime}\right\}$.
(1) follows by T-SKIP and (2) follows by assumptions (C) and (D) and Lemma A.30.
(3) and (4) are immediate.

For (5) we have to show that $(I, O @ b v, H) \vDash_{\mathcal{E}[x \mapsto(I, O, H)]} \Sigma y:\left\{z: \tau_{1} \mid z \equiv x\right\} .\{z$ : $\left.\epsilon \mid z . p k t_{\text {in }}=\langle \rangle \wedge z . p k t_{\text {out }}=x . l\right\}$. By Ent-Sigma, we must show that
(5.1) $(I, O, H) \vDash_{\mathcal{E}[x \mapsto(I, O, H)]}\left\{z: \tau_{1} \mid z \equiv x\right\}$ and
(5.2) $\left(\rangle, b v,[]) \vDash_{\mathcal{E}[x \mapsto(I, O, H), y \mapsto(I, O, H)]}\left\{z: \epsilon \mid z \cdot p k t_{\text {in }}=\langle \rangle \wedge z . p k t_{\text {out }}=x . \iota\right\}\right.$

For (5.1) we must show by Ent-Refine that $(I, O, H) \vDash_{\mathcal{E}[x \mapsto(I, O, H)]} \tau_{1}$, which follows by assumption (D) and the fact that x is not free in $\tau_{1}$. We must also show that $\llbracket z \equiv x \rrbracket_{\mathcal{E}[x \mapsto(I, O, H), z \mapsto(I, O, H)]}$, which follows by the semantics of formulae and by reflexivity.

To show (5.2), by Ent-Refine, we must show that

$$
(\rangle, b v,[]) \vDash \mathcal{E}[x \mapsto(I, O, H), y \mapsto(I, O, H)] \epsilon
$$

and that

$$
\llbracket z . p k t_{\text {in }}=\langle \rangle \wedge z . p k t_{\text {out }}=x . \downarrow \rrbracket_{\mathcal{E}[x \mapsto(I, O, H), y \mapsto(I, O, H), z \mapsto(( \rangle, b v,[])]}=\text { true }
$$

The first follows after unfolding the definition of $\epsilon$ by Ent-Top, Ent-Refine and the semantics of formulae. The second follows by the semantics of formulae and (A6).
(6) follows by

$$
\begin{aligned}
& \llbracket \tau_{2} \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]} \\
&=\llbracket \Sigma y:\left\{z: \tau_{1} \mid z \equiv x\right\} \cdot\left\{z: \epsilon \mid z \cdot p k t_{\text {in }}=\langle \rangle \wedge z \cdot p k t_{\text {out }}=x . l\right\} \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]} \\
&=\left\{h_{1}++h_{2} \mid h_{1} \in \llbracket\left\{z: \tau_{1} \mid z \equiv x\right\} \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)] \wedge}\right. \\
&\left.\quad h_{2} \in \llbracket\left\{z: \epsilon \mid z \cdot p k t_{\text {in }}=\langle \rangle \wedge z \cdot p k t_{\text {out }}=x . l\right\} \rrbracket_{\mathcal{E}\left[x \mapsto(I, O, H), y \mapsto h_{1}\right]}\right\} \\
&=\{(I, O, H)++(\langle \rangle, b v,[])\} \\
&=\{(I, O @ b v, H)\} \\
&=\left\{h \mid h \in \llbracket \Sigma y:\left\{z: \tau_{1} \mid z \equiv x\right\} .\right. \\
&\left\{z: \epsilon \mid z \cdot p k t_{\text {in }}=\langle \rangle \wedge\right. \\
&\left.\quad z \cdot p k t_{\text {out }}=x . l\right\} \rrbracket_{\mathcal{E}\left[x \mapsto(I, O, H), x^{\prime} \mapsto(I, O @ b v, H)\right]} \wedge \\
&= \llbracket\left\{v: \Sigma y:\left\{z: \tau_{1} \mid z \equiv x\right\} .\right. \\
&\left\{z: \epsilon \mid z \cdot p k t_{\text {in }}=\langle \rangle \wedge z \cdot p k t_{\text {out }}=x . l\right\} \mid \\
&\left.\quad v \equiv x^{\prime}\right\} \rrbracket_{\mathcal{E}\left[x \mapsto(I, O, H), x^{\prime} \mapsto(I, O @ b v, H)\right]} \\
&= \llbracket \tau_{2}^{\prime} \rrbracket_{\mathcal{E}^{\prime}\left[x^{\prime} \mapsto\left(I^{\prime}, O^{\prime}, H^{\prime}\right)\right]}
\end{aligned}
$$

## Case T-Remove:

By inversion of rule T-Remove, we get
(A1) $c=\operatorname{remove}(\iota)$
(A2) $\tau_{2}=\left\{y: \top \mid \varphi_{1} \wedge \varphi_{p k t} \wedge \neg y . \iota\right.$. valid $\}$
Only evaluation rule E-Remove applies to c.
$\left(\mathrm{A}_{3}\right) c^{\prime}=s k i p$
(A4) $I^{\prime}=I$ and $O^{\prime}=O$ and $H^{\prime}=H \backslash \iota$
Let $\Gamma^{\prime}=\Gamma, x: \tau_{1}$ and $\mathcal{E}^{\prime}=\mathcal{E}[x \mapsto(I, O, H)]$.
Let $\tau_{1}^{\prime}=\tau_{2}$ and let $\tau_{2}^{\prime}=\left\{z: \tau_{2} \mid z \equiv x^{\prime}\right\}$.
(1) follows by T-Skip. (2) follows by assumptions (C) and (D) and Lemma A.30.
(3) and (4) are immediate.

To show (5), we must show that $(I, O, H \backslash \iota) \vDash_{\mathcal{E}[x \mapsto(I, O, H)]}\left\{y: T \mid \varphi_{t} \wedge \varphi_{p k t} \wedge\right.$ $\neg y . \iota . v a l i d\}$, which follows by Ent-Refine, Ent-Top and the semantics of formulae and expressions.
(6) follows by

$$
\begin{aligned}
& \llbracket \tau_{2}^{\prime} \rrbracket_{\mathcal{E}} \mathcal{E}^{\prime}\left[x^{\prime} \mapsto\left(I^{\prime}, O^{\prime}, H^{\prime}\right)\right] \\
= & \llbracket\left\{z:\left\{y: \uparrow \mid \varphi_{\iota} \wedge \varphi_{p k t} \wedge \neg y . \iota . \text { valid }\right\} \mid z \equiv x^{\prime}\right\} \rrbracket_{\mathcal{E}\left[x \mapsto(I, O, H), x^{\prime} \mapsto(I, O, H \backslash)\right]} \\
= & \left\{h \mid h \in \llbracket\left\{y: \top \mid \varphi_{\iota} \wedge \varphi_{p k t} \wedge \neg y . . . v a l i d\right\}\right\} \rrbracket_{\mathcal{E}\left[x \mapsto(I, O, H), x^{\prime} \mapsto(I, O, H \backslash \iota)\right]} \wedge \\
& \left.\llbracket z \equiv x^{\prime} \rrbracket \mathcal{E}\left[x \mapsto(I, O, H), x^{\prime} \mapsto(I, O, H \backslash \iota), z \mapsto h\right]\right\} \\
= & \{(I, O, H \backslash \iota)\} \\
= & \llbracket\left\{y: \top \mid \varphi_{\iota} \wedge \varphi_{p k t} \wedge \neg y . \iota . \text { valid }\right\} \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]}
\end{aligned}
$$

## Case T-Reset:

By inversion of rule T-Reset, we get
(A1) $c=$ reset
(A2) $\tau_{2}=\Sigma y:\left\{z: \epsilon \mid z \cdot p k t_{\text {out }}=\langle \rangle \wedge z \cdot p k t_{\text {in }}=x \cdot p k t_{\text {out }}\right\} \cdot\left\{z: \epsilon \mid z \cdot p k t_{\text {out }}=\right.$ $\left\rangle \wedge z \cdot p k t_{i n}=x \cdot p k t_{i n}\right\}$

There is only one evaluation rule that applies to $c$, E-Reset.
$\left(\mathrm{A}_{3}\right) c^{\prime}=$ skip
(A4) $I^{\prime}=O @ I, O^{\prime}=\langle \rangle$ and $H^{\prime}=[]$
Let $\mathcal{E}^{\prime}=\mathcal{E}[x \mapsto(I, O, H)]$ and $\Gamma^{\prime}=\Gamma, x: \tau_{1}$.
Let $\tau_{1}^{\prime}=\Sigma y:\left\{z: \epsilon \mid z \cdot p k t_{\text {out }}=\langle \rangle \wedge z \cdot p k t_{\text {in }}=x \cdot p k t_{\text {out }}\right\} \cdot\left\{z: \epsilon \mid z \cdot p k t_{\text {out }}=\right.$ $\left\rangle \wedge z . p k t_{i n}=x . p k t_{i n}\right\}$
Let $\tau_{2}^{\prime}=\left\{v: \Sigma y:\left\{z: \epsilon \mid z \cdot p k t_{\text {out }}=\langle \rangle \wedge z \cdot p k t_{\text {in }}=x \cdot p k t_{\text {out }}\right\} \cdot\left\{z: \epsilon \mid z \cdot p k t_{\text {out }}=\right.\right.$ $\left.\left\rangle \wedge z . p k t_{i n}=x . p k t_{i n}\right\} \mid v \equiv x^{\prime}\right\}$
(1) follows by T-Skip and (2) follows by assumptions (C) and (D) and Lemma A. 30 .
(3) and (4) are immediate.

To show (5), we must show that

$$
\begin{aligned}
(O @ I,\langle \rangle,[]) \vDash_{\mathcal{E}[x \mapsto(I, O, H)]} \Sigma y:\left\{z: \epsilon \mid z \cdot p k t_{\text {out }}\right. & \left.=\langle \rangle \wedge z \cdot p k t_{\text {in }}=x \cdot p k t_{\text {out }}\right\} . \\
\left\{z: \epsilon \mid z \cdot p k t_{\text {out }}\right. & \left.=\langle \rangle \wedge z \cdot p k t_{\text {in }}=x \cdot p k t_{\text {in }}\right\}
\end{aligned}
$$

By Ent-Sigma, we must show that

$$
(O,\langle \rangle,[]) \vDash_{\mathcal{E}[x \mapsto(I, O, H)]}\left\{z: \epsilon \mid z . p k t_{\text {out }}=\langle \rangle \wedge z . p k t_{\text {in }}=x . p k t_{\text {out }}\right\}
$$

and

$$
(I,\langle \rangle,[]) \vDash_{\mathcal{E}[x \mapsto(I, O, H), y \mapsto(O,( \rangle,[])]}\left\{z: \epsilon \mid z . p k t_{\text {out }}=\langle \rangle \wedge z . p k t_{\text {in }}=x . p k t_{\text {in }}\right\}
$$

Both follow after unfolding the definition of $\epsilon$ by Ent-Refine, Ent-Top and the semantics of formulae.
(6) follows by

$$
\begin{aligned}
& \llbracket \Sigma y: \\
&\left\{z: \epsilon \mid z \cdot p k t_{\text {out }}=\langle \rangle \wedge z \cdot p k t_{\text {in }}=x \cdot p k t_{\text {out }}\right\} . \\
&=\left\{z: \epsilon \mid z \cdot p k t_{\text {out }}=\langle \rangle \wedge z \cdot p k t_{\text {in }}=x \cdot p k t_{\text {in }}\right\} \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]} \mid \\
& h_{1} \in \llbracket\left\{z: \epsilon \mid z \cdot p k t_{\text {out }}=\langle \rangle \wedge z \cdot p k t_{\text {in }}=x \cdot p k t_{\text {out }}\right\}_{\mathcal{E}[x \mapsto(I, O, H)]} \rrbracket \wedge \\
&\left.h_{2} \in \llbracket\left\{z: \epsilon \mid z \cdot p k t_{\text {out }}=\langle \rangle \wedge z \cdot p k t_{\text {in }}=x \cdot p k t_{\text {in }}\right\} \rrbracket_{\mathcal{E}\left[x \mapsto(I, O, H), y \mapsto h_{1}\right]}\right\} \\
&=\{(O @ I,\langle \rangle)+[+(I,\langle \rangle,[])\} \\
&= \llbracket\left\{v: \Sigma y:\left\{z: \epsilon \mid z \cdot p k t_{\text {out }}=\langle \rangle \wedge z \cdot p k t_{\text {in }}=x \cdot p k t_{\text {out }}\right\} .\right. \\
&\left\{z: \epsilon \mid z \cdot p k t_{\text {out }}=\langle \rangle \wedge z \cdot p k t_{\text {in }}=x \cdot p k t_{\text {in }}\right\} \mid \\
&\left.v \equiv x^{\prime}\right\} \rrbracket_{\mathcal{E}\left[x \mapsto(O @ I,\langle \rangle,[]), x^{\prime} \mapsto(O @ I,\langle \rangle,[])\right]}
\end{aligned}
$$

## Case T-SEQ:

By inversion of rule T-SEQ, we get
(A1) $c=c 1 ; c 2$
(A2) $\Gamma \vdash c_{1}:\left(x: \tau_{1}\right) \rightarrow \tau_{12}$
(A3) $\Gamma, x: \tau_{1} \vdash c_{2}:\left(y: \tau_{12}\right) \rightarrow \tau_{22}$
(A4) $\tau_{2}=\tau_{22}\left[y \mapsto \tau_{12}\right]$

## Subcase E-SEQ:

(A5) $c_{1}=$ skip
(A6) $c^{\prime}=c_{2}$
By E-Seq, $I^{\prime}=I, O^{\prime}=O, H^{\prime}=H$.
Let $\Gamma^{\prime}=\Gamma, x: \tau_{1}, \mathcal{E}^{\prime}=\mathcal{E}[x \mapsto(I, O, H)]$ and $x^{\prime}=y$.
Let $\tau_{1}^{\prime}=\tau_{12}$ and $\tau_{2}^{\prime}=\tau_{22}$.
(1) follows by ( $\mathrm{A}_{3}$ ), (2) follows by assumptions (C) and (D) and Lemma A.30. (3) and (4) are immediate.

To show (5), we must show that $(I, O, H) \vDash_{\mathcal{E}[x \mapsto(I, O, H)]}\left\{z: \tau_{1} \mid z \equiv x\right\}$ holds. By Ent-Refine, we must show that $(I, O, H) \vDash_{\mathcal{E}[x \rightarrow(I, O, H)]} \tau_{1}$ and $\llbracket z \equiv x \rrbracket_{\mathcal{E}[x \mapsto(I, O, H), z \mapsto(I, O, H)]}$, which follows by assumption (D) and the semantics of formulae.
For (6), we must show that

$$
\begin{aligned}
& \llbracket \tau_{2}^{\prime} \rrbracket_{\mathcal{E}}\left[\left[x \mapsto\left(I^{\prime}, O^{\prime}, H^{\prime}\right)\right]\right. \\
\Leftrightarrow & \llbracket \tau_{22} \rrbracket_{\mathcal{E}[x \mapsto(I, O, H), y \mapsto(I, O, H)]} \subseteq \tau_{2} \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]} \\
& \llbracket \tau_{22}\left[y \mapsto\left\{z: \tau_{22} \mid z \equiv x\right\}\right] \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]} \\
= & \left.\left\{h_{22} \mid h_{12}\right] \rrbracket_{\mathcal{E}[x \mapsto(I, O . H)]} \in \llbracket\left\{z: \tau_{1} \mid z \equiv x\right\} \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]} \wedge h_{22} \in \llbracket \tau_{22} \rrbracket_{\mathcal{E}\left[x \mapsto(I, O, H), y \mapsto h_{12}\right]}\right\} \\
= & \left\{h_{22} \mid h_{12}=(I, O, H) \wedge h_{22} \in \llbracket \tau_{22} \rrbracket_{\mathcal{E}\left[x \mapsto(I, O, H), y \mapsto h_{12}\right]}\right\} \\
= & \left\{h_{22} \mid h_{22} \in \llbracket \tau_{22} \rrbracket_{\mathcal{E}[x \mapsto(I, O, H), y \mapsto(I, O, H)]}\right\} \\
= & \llbracket \tau_{22} \rrbracket_{\mathcal{E}[x \mapsto(I, O, H), y \mapsto(I, O, H)]}
\end{aligned}
$$

This shows (6) and concludes this subcase.

## Subcase E-SEQ1:

(A7) $c^{\prime}=c_{1}^{\prime} ; c_{2}$
(A8) $\left\langle I, O, H, c_{1}\right\rangle \rightarrow\left\langle I^{\prime}, O^{\prime}, H^{\prime}, c_{1}^{\prime}\right\rangle$
By IH with (A2), (A7), (C) and (D), there exists $\Gamma^{\prime}, \mathcal{E}^{\prime}, \tau_{1}^{\prime}, \tau_{12}^{\prime}, x^{\prime}$, such that,
(A9) $\Gamma^{\prime} \vdash c_{1}^{\prime}:\left(x^{\prime}: \tau_{1}^{\prime}\right) \rightarrow \tau_{12}^{\prime}$ where
(A10) $\mathcal{E}^{\prime} \vDash \Gamma^{\prime}$
(A11) $\Gamma \subseteq \Gamma^{\prime}$
(A12) $\mathcal{E} \subseteq \mathcal{E}^{\prime}$
(A13) $\left(I^{\prime}, O^{\prime}, H^{\prime}\right) \vDash \mathcal{E}^{\prime} \tau_{1}^{\prime}$
(A14) $\llbracket \tau_{12}^{\prime} \rrbracket_{\mathcal{E} '\left[x^{\prime} \mapsto\left(I^{\prime}, O^{\prime}, H^{\prime}\right)\right]} \subseteq \llbracket \tau_{12} \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]}$
(1) follows by T-SEQ, if we can show that there exists some $\tau_{22}^{\prime}$, such that $\Gamma^{\prime}, x^{\prime}: \tau_{1}^{\prime} \vdash c_{2}:\left(y: \tau_{12}^{\prime}\right) \rightarrow \tau_{22}^{\prime}$ where $\tau_{2}^{\prime}=\tau_{22}^{\prime}\left[y \mapsto \tau_{12}^{\prime}\right]$ :

$$
\frac{\begin{array}{l}
\text { T-SEQ } \\
\Gamma^{\prime} \vdash c_{1}^{\prime}:\left(x: \tau_{1}^{\prime}\right) \rightarrow \tau_{12}^{\prime}
\end{array} \quad \Gamma^{\prime}, x: \tau_{1}^{\prime} \vdash c_{2}:\left(y: \tau_{12}^{\prime}\right) \rightarrow \tau_{22}^{\prime}}{\Gamma^{\prime} \vdash c_{1}^{\prime} ; c_{2}:\left(x: \tau_{1}^{\prime}\right) \rightarrow \tau_{22}^{\prime}\left[y \mapsto \tau_{12}^{\prime}\right]}
$$

By Lemma A. 49 with ( $\mathrm{A}_{3}$ ) and (A13), there exists some $\tau_{22}^{\prime}$ such that
(A15) $\Gamma^{\prime}, x^{\prime}: \tau_{1}^{\prime} \vdash c_{2}:\left(y: \tau_{1}^{\prime}\right) \rightarrow \tau_{22}^{\prime}$
(A16) $\left.\forall h^{\prime} \in \llbracket \tau_{12}^{\prime} \rrbracket_{\mathcal{E}}\left[x^{\prime} \mapsto\left(I^{\prime}, O^{\prime}, H^{\prime}\right)\right] \cdot \llbracket \tau_{22}^{\prime} \rrbracket\right]_{\mathcal{E}^{\prime}\left[x^{\prime} \mapsto\left(I^{\prime}, O^{\prime}, H^{\prime}\right), y \mapsto h^{\prime}\right]} \subseteq$

$$
\llbracket \tau_{22} \rrbracket_{\mathcal{E}\left[x \mapsto(I, O, H), y \mapsto h^{\prime}\right]}
$$

(2) follows by (A10) and (3) follows by (A11), (4) follows by (A12) and (5) follows by (A13).
(6) follows by

$$
\begin{aligned}
& \llbracket \tau_{2}^{\prime} \rrbracket_{\mathcal{E}^{\prime}\left[x^{\prime} \mapsto\left(I^{\prime}, O^{\prime}, H^{\prime}\right)\right]} \subseteq \llbracket \tau_{2} \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]} \\
& \left.\Leftrightarrow \llbracket \tau_{22}^{\prime}\left[y \mapsto \tau_{12}^{\prime}\right]\right]_{\mathcal{E}^{\prime}\left[x \mapsto\left(I^{\prime}, O^{\prime}, H^{\prime}\right)\right]} \subseteq \llbracket \tau_{22}\left[y \mapsto \tau_{12}\right] \rrbracket_{\mathcal{E}[x \mapsto(I, O, H)]} \\
& \Leftrightarrow \bigcup_{h^{\prime} \in \llbracket \tau_{\tau_{12}^{\prime}} \mathbb{E}_{\mathcal{E}^{\prime}\left[x \leftrightarrow\left(I^{\prime}, O^{\prime}, H^{\prime}\right)\right]} \llbracket \tau_{22}^{\prime} \rrbracket_{\mathcal{E}^{\prime}\left[x \mapsto\left(I^{\prime}, O^{\prime}, H^{\prime}\right), y \mapsto h^{\prime}\right]} \subseteq} \\
& \bigcup_{h \in \llbracket \tau_{12} \rrbracket \mathcal{E}[x \mapsto(I, O, H)]} \llbracket \tau_{22} \rrbracket_{\mathcal{E}[x \mapsto(I, O, H), y \mapsto h]}
\end{aligned}
$$

and by (A14), (A16) and the semantics of heap types.

## Case T-Skip:

Immediately holds as there is no $c^{\prime}$ such that $\langle I, O, H, c\rangle \rightarrow\left\langle I^{\prime}, O^{\prime}, H^{\prime}, c^{\prime}\right\rangle$.

## Case T-Sub:

(A1) $\Gamma \vdash c:\left(x: \tau_{3}\right) \rightarrow \tau_{4}$
(A2) $\Gamma \vdash \tau_{1}<: \tau_{3}$
(A3) $\Gamma, x: \tau_{1} \vdash \tau_{4}<: \tau_{2}$
By Lemma A. 29 with assumptions (C), (D) and (A2),
(A4) $(I, O, H) \vDash \mathcal{E} \tau_{3}$.
By IH with ( $\mathrm{A}_{1}$ ), (B), (C) and ( $\mathrm{A}_{4}$ ), there exists $\Gamma^{\prime}, \mathcal{E}^{\prime}, \tau_{3}^{\prime}, \tau_{4}^{\prime}, x^{\prime}$ such that
(A5) $\Gamma^{\prime} \vdash c:\left(x^{\prime}: \tau_{3}^{\prime}\right) \rightarrow \tau_{4}^{\prime}$
(A6) $\mathcal{E}^{\prime} \vDash \Gamma^{\prime}$
(A7) $\Gamma \subseteq \Gamma^{\prime}$
(A8) $\mathcal{E} \subseteq \mathcal{E}^{\prime}$
(A9) $\left(I^{\prime}, O^{\prime}, H^{\prime}\right) \vDash_{\mathcal{E}^{\prime}} \tau_{3}^{\prime}$
(A10) $\left.\llbracket \tau_{4}^{\prime} \rrbracket_{\mathcal{E}^{\prime}\left[x^{\prime} \mapsto\left(I^{\prime}, O^{\prime}, H^{\prime}\right)\right]} \subseteq \llbracket \tau_{4}\right]_{\mathcal{E}[x \mapsto(I, O, H)]}$
Let $\tau_{1}^{\prime}=\tau_{3}^{\prime}$ and $\tau_{2}^{\prime}=\tau_{4}^{\prime}$. (1) follows by (A5), (2) follows by assumption (A6), (3) follows by assumption (A7), (4) follows by assumption (A8), (5) follows by assumption (A9) and (6) follows by ( $\mathrm{A}_{3}$ ) and (A10).

## A.2.2 Algorithmic Typing Correctness

Lemma A. 55 (Subtype Reflexivity). For all subtyping contexts $\Gamma$ and heap types $\tau, \Gamma \vdash$ $\tau<: \tau$.

Proof. Immediate.
Lemma A. 56 (Subtype Transitivity). If $\Gamma \vdash \tau_{1}<: \tau_{2}$, and $\Gamma \vdash \tau_{2}<: \tau_{3}$, then $\Gamma \vdash \tau_{1}<: \tau_{3}$.
Proof. Assume $\Gamma \vdash \tau_{1}<: \tau_{2}$ and also assume $\Gamma \vdash \tau_{2}<: \tau_{3}$. Let $\mathcal{E} \vDash \Gamma$ and $h \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}}$ be arbitrary. By the first assumption $h \in \llbracket \tau_{2} \rrbracket \mathcal{E}$. By the second assumption $h \in \llbracket \tau_{3} \rrbracket \mathcal{E}$.

Lemma A. 57 (Environment Entails Subtype). If $\Gamma \vdash \tau_{1} \ll \tau_{3}$, and $\mathcal{E} \vDash \Gamma, x: \tau_{3}$, then $\mathcal{E} \vDash \Gamma, x: \tau_{1}$.

Proof. Let $\Gamma \vdash \tau_{1}<: \tau_{3}$ and $\mathcal{E} \vDash \Gamma,\left(x: \tau_{3}\right)$. Let $\mathcal{E} \vDash \Gamma,\left(x: \tau_{1}\right)$. We can write $\mathcal{E}=\mathcal{E}^{\prime}\left[x \mapsto h_{1}\right]$, such that $\mathcal{E}^{\prime} \vDash \Gamma$, and $\left.h_{1} \in \llbracket \tau_{1}\right] \mathcal{E}^{\prime}$. The definition of subtyping gives $h_{1} \in \llbracket \tau_{3} \rrbracket \mathcal{E}^{\prime}$. The result follows by definition of entailment.

Lemma A. 58 (Context Strengthening). If $\Gamma \vdash \tau_{1}<: \tau_{3}$ and $\Gamma, x: \tau_{3} \vdash \tau_{2}<: \tau_{4}$ then $\Gamma, x: \tau_{1} \vdash \tau_{2}<: \tau_{4}$.

Proof. Assume $\Gamma \vdash \tau_{1}<: \tau_{3}$ and further assume $\Gamma, x: \tau_{3} \vdash \tau_{2}<: \tau_{4}$. Let $\mathcal{E} \vDash \Gamma, x: \tau_{1}$. By Lemma A. 57 and the first assumption, $\mathcal{E} \vDash \Gamma,\left(x: \tau_{3}\right)$. Let $h \in \llbracket \tau_{2} \rrbracket \mathcal{E}$, the second assumption gives that $h \in \llbracket \tau_{2} \rrbracket \mathcal{E}$, and we're done.

Lemma A. 59 (Packet Bound Subtype). $\Gamma \vdash \operatorname{sizeof}_{p k t_{i n}}(\tau) \geq N$ iff $\Gamma \vdash \tau<:\{x: \top \mid$ $\left.\left|x . p k t_{i n}\right| \geq \operatorname{sizeof}(\iota)\right\}$.

Proof. We show each direction separately.
$(\Rightarrow)$ Assume $\Gamma \vdash \operatorname{sizeof}_{p k t_{i n}}(\tau) \geq N$. Let $\mathcal{E} \vDash \Gamma$ and $h \in \llbracket \tau \rrbracket \mathcal{E}$ be arbitrary. By definition, $\left|h\left(p k t_{i n}\right)\right| \geq N$. By the definition of subtyping, it suffices to show $h \in \llbracket\left\{x: \top| | x . p k t_{i n} \mid \geq N\right\} \rrbracket_{\mathcal{E}}$. By definition, $\llbracket\left\{x: \top| | x . p k t_{i n} \mid \geq N\right\} \rrbracket_{\mathcal{E}}=\{h \mid$ $\left.h \in \mathcal{H} \wedge h\left(p k t_{i n}\right) \geq N\right\}$, which concludes this case.
$(\Leftarrow)$ Assume $\Gamma \vdash \tau<:\left\{x: \top| | x . p k t_{i n} \mid \geq N\right\}$. We have to show that $\Gamma \vdash$ sizeof $_{p k t_{i n}}(\tau) \geq N$. Let $\mathcal{E} \vDash \Gamma$ and $h \in \llbracket \tau \rrbracket \mathcal{E}$ be arbitrary. By the definition of subtyping, $h \in \llbracket\left\{x: T| | x . p k t_{i n} \mid \geq N\right\} \rrbracket \mathcal{E}$. By definition of the semantics, we can conclude $h\left(p k t_{i n}\right) \geq N$.

Lemma A. 60 (Chomp Subtype). If $x$ not free in $\tau$ and $\tau^{\prime}$, and $\Gamma \vdash \operatorname{sizeof}_{p k t_{i n}}(\tau) \geq$ $\operatorname{sizeof}(\iota)$ and $\Gamma \vdash \tau<: \tau^{\prime}$, then $\Gamma, x:\left\{y: \iota \mid y \cdot p k t_{\text {in }}=y . p k t_{\text {out }}=\langle \rangle\right\} \vdash$ $\operatorname{chomp}(\tau, \iota, x)<: \operatorname{chomp}\left(\tau^{\prime}, \iota, x\right)$.

Proof. Given some heap $h^{\prime} \in \llbracket \operatorname{chomp}\left(\tau^{\prime}, l, x\right) \rrbracket_{\mathcal{E}^{\prime}}$. By Lemma A. 47 there exists some $\mathcal{E}$ and $h^{\prime \prime} \in \llbracket \tau^{\prime} \rrbracket \mathcal{E}$, such that $\mathcal{E}^{\prime}=\mathcal{E}\left[x \mapsto\left(\langle \rangle,\langle \rangle,\left[\iota \mapsto h^{\prime \prime}\left(p k t_{\text {in }}\right)[\right.\right.\right.$ o:sizeof $\left.\left.\left.(\iota)]\right]\right)\right]$ and $h^{\prime}=$ chomp $^{\Downarrow}\left(h^{\prime \prime}, \operatorname{sizeof}(\iota)\right)$. By assumption $\Gamma \vdash \tau<: \tau^{\prime}$, we also know that $h^{\prime \prime} \in \llbracket \tau \rrbracket \mathcal{E}$.

By Lemma A.44, we know that there exists $h \in \llbracket \operatorname{chomp}(\tau, t, x) \rrbracket_{\mathcal{E}^{\prime \prime}}$ such that $h=$ chomp ${ }^{\Downarrow}\left(h^{\prime \prime}, \operatorname{sizeof}(\iota)\right)$ and $\mathcal{E}^{\prime \prime}=\mathcal{E}\left[x \mapsto\left(\langle \rangle,\langle \rangle,\left[\iota \mapsto h^{\prime \prime}\left(p k t_{i n}\right)[\mathrm{o}: \operatorname{sizeof}(\iota)]\right]\right)\right]$. From $h^{\prime}=\operatorname{chomp}^{\Downarrow}\left(h^{\prime \prime}, \operatorname{sizeof}(\iota)\right)$ and $h=\operatorname{chomp}^{\Downarrow}\left(h^{\prime \prime}, \operatorname{sizeof}(\iota)\right)$ follows by the transitivity of equality that $h^{\prime}=h$. By the fact that $\mathcal{E}^{\prime}=\mathcal{E}^{\prime \prime}$ follows that for every heap $h^{\prime} \in \llbracket \operatorname{chomp}\left(\tau^{\prime}, l, x\right) \rrbracket_{\mathcal{E}^{\prime}}$ also holds that $h^{\prime} \in \llbracket \operatorname{chomp}(\tau, \iota, y) \rrbracket_{\mathcal{E}^{\prime \prime}}$.

Lemma A. 61 (Refinement Subtype). If $\Gamma \vdash \tau^{\prime}<: \tau$, then $\Gamma \vdash\left\{x: \tau^{\prime} \mid \varphi\right\}<:\{x: \tau \mid \varphi\}$.
Proof. Let $\mathcal{E} \vdash \Gamma$ and $h \in \llbracket\left\{x: \tau^{\prime} \mid \varphi\right\} \rrbracket_{\mathcal{E}}$. Then $h \in \llbracket \tau^{\prime} \rrbracket_{\mathcal{E}}$, and $\llbracket \varphi \rrbracket_{\mathcal{E}[x \mapsto h]}=$ true. By assumption $\Gamma \vdash \tau^{\prime}<: \tau, h \in \llbracket \tau \rrbracket_{\mathcal{E}}$. Conclude $h \in \llbracket\{x: \tau \mid \varphi\} \rrbracket_{\mathcal{E}}$ by definition. The result follows.

Lemma A. 62 (Sigma Left-Subtype). If $\Gamma \vdash \tau_{1}^{\prime}<: \tau_{1}$ then $\Gamma \vdash \Sigma x: \tau_{1}^{\prime} . \tau_{2}<: \Sigma x: \tau_{1} . \tau_{2}$.
Proof. Let $\mathcal{E} \vDash \Gamma$, and $h \in \llbracket \Sigma x: \tau_{1}^{\prime} \cdot \tau_{2} \rrbracket_{\mathcal{E}}$. By definition of the semantics, $h=h_{1}+$ $+h_{2}$, where $h_{1} \in \llbracket \tau_{1}^{\prime} \rrbracket_{\mathcal{E}}$ and $h_{2} \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}[x \mapsto h]}$. By assumption $\Gamma \vdash \tau_{1}^{\prime}<: \tau_{1}$ follows, $h_{1} \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}}$. By the definition of the semantics, $h \in \llbracket \Sigma x: \tau_{1} \cdot \tau_{2} \rrbracket \mathcal{E}$. The result follows.

Lemma A. 63 (Sigma Right-Subtype). If $\Gamma, x: \tau_{1} \vdash \tau_{2}^{\prime}<: \tau_{2}$ then $\Gamma \vdash \Sigma x: \tau_{1} . \tau_{2}^{\prime}<: \Sigma x$ : $\tau_{1} . \tau_{2}$.

Proof. Let $\mathcal{E} \vDash \Gamma$ and $h \in \llbracket \Sigma x: \tau_{1} \cdot \tau_{2}^{\prime} \rrbracket_{\mathcal{E}}$. By definition, $h=h_{1}++h_{2}$ such that $h_{1} \in \llbracket \tau_{1} \rrbracket \mathcal{E}$, and $h_{2} \in \llbracket \tau_{2}^{\prime} \rrbracket_{\mathcal{E}\left[x \mapsto h_{1}\right]}$. Notice that $\mathcal{E}\left[x \mapsto h_{1}\right] \vDash \Gamma,\left(x: \tau_{1}\right)$, so by assumption $\Gamma, x: \tau_{1} \vdash$ $\tau_{2}^{\prime}<: \tau_{2}$ follows $h_{2} \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}\left[x \mapsto h_{1}\right]}$. By the definition of the semantics $h \in \llbracket \Sigma x: \tau_{1} \cdot \tau_{2} \rrbracket_{\mathcal{E}}$. The result follows.

Lemma A. 64 (Substitution Subtype). If $\Gamma \vdash \tau_{1}^{\prime}<: \tau_{1}$ and $\Gamma,\left(x: \tau_{1}^{\prime}\right) \vdash \tau_{2}^{\prime}<: \tau_{2}$, then $\Gamma \vdash \tau_{2}^{\prime}\left[x \mapsto \tau_{1}^{\prime}\right]<: \tau_{2}\left[x \mapsto \tau_{1}\right]$.

Proof. Let $\mathcal{E} \vDash \Gamma$ and $\left.h_{2} \in \llbracket \tau_{2}^{\prime}\left[x \mapsto \tau_{1}^{\prime}\right]\right]_{\mathcal{E}}$. Then we know $h_{1} \in \llbracket \tau_{1}^{\prime} \rrbracket_{\mathcal{E}}$ and $h_{2} \in$ $\left.\llbracket \tau_{2}^{\prime}\right]_{\mathcal{E}\left[x \mapsto h_{1}\right]}$. Assumption $\Gamma \vdash \tau_{1}^{\prime}<: \tau_{1}$ tells us that $h_{1} \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}}$. Notice that $\mathcal{E}[x \mapsto$ $\left.h_{1}\right] \vDash \Gamma,\left(x: \tau_{1}^{\prime}\right)$. Assumption $\Gamma,\left(x: \tau_{1}\right) \vdash \tau_{2}^{\prime}<: \tau_{2}$ gives $h_{2} \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}\left[x \mapsto h_{1}\right]}$. By the definition of the semantics of heap types, $h_{2} \in \llbracket \tau_{2}\left[x \mapsto \tau_{1}\right] \rrbracket_{\mathcal{E}}$.

Lemma A. 65 (Choice Subtype). If $\Gamma \vdash \tau_{1}^{\prime}<: \tau_{1}$, and $\Gamma \vdash \tau_{2}^{\prime}<: \tau_{2}$, then $\Gamma \vdash \tau_{1}^{\prime}+\tau_{2}^{\prime}<$ : $\tau_{1}+\tau_{2}$

Proof. Let $\mathcal{E} \vDash \Gamma$. Let $h \in \llbracket \tau_{1}^{\prime}+\tau_{2}^{\prime} \rrbracket_{\mathcal{E}}$. By semantics of heap types, either $h \in \llbracket \tau_{1}^{\prime} \rrbracket_{\mathcal{E}}$ or $h \in \llbracket \tau_{2}^{\prime} \rrbracket_{\mathcal{E}}$.

Subcase $h \in \llbracket \tau_{1}^{\prime} \rrbracket_{\mathcal{E}}:$ By assumption $\Gamma \vdash \tau_{1}^{\prime} \ll \tau_{1}$ it also holds that $h \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}}$ and we can conclude that $h \in \llbracket \tau_{1}+\tau_{2} \rrbracket_{\mathcal{E}}$.

Subcase $h \in \llbracket \tau_{2}^{\prime} \rrbracket \mathcal{E}:$ By assumption $\Gamma \vdash \tau_{2}^{\prime}<: \tau_{2}$ it also holds that $h \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}}$ and we can conclude that $h \in \llbracket \tau_{1}+\tau_{2} \rrbracket \mathcal{E}$.

Lemma A. 66 (Context-Bound Refinement Subtype). If heap is the only free binder in $\varphi$, and $\Gamma, x:\left\{y: \tau_{1} \mid \varphi[y /\right.$ heap $\left.]\right\} \vdash \tau_{2}^{\prime}<: \tau_{2}$ then $\Gamma, x: \tau_{1} \vdash\left\{y: \tau_{2}^{\prime} \mid \varphi[x /\right.$ heap $\left.]\right\}<:\{y:$ $\tau_{2} \mid \varphi[x /$ heap $\left.]\right\}$.

Proof. Let $\mathcal{E} \vDash \Gamma, x: \tau_{1}$. We can write this as $\mathcal{E}=\mathcal{E}^{\prime}\left[x \mapsto h_{1}\right]$, where $h_{1} \in \llbracket \tau_{1} \rrbracket \mathcal{E}$. Let $h_{2} \in \llbracket\left\{y: \tau_{2}^{\prime} \mid \varphi[x /\right.$ heap $\rrbracket\} \rrbracket_{\mathcal{E}}$. Then $h_{2} \in \llbracket \tau_{2}^{\prime} \rrbracket \mathcal{E}$ and $\llbracket \varphi[x /$ heap $] \rrbracket_{\mathcal{E}\left[y \mapsto h_{2}\right]}=$ true . Compute as follows, recalling that heap is the only free binder in $\varphi$ :

$$
\begin{aligned}
& \llbracket \varphi[x / \text { heap }] \rrbracket_{\mathcal{E}\left[y \mapsto h_{2}\right]} \\
= & \llbracket \varphi[x / \text { heap }] \rrbracket \rrbracket_{\mathcal{E}} \\
= & \llbracket \varphi[x / \text { heap }] \rrbracket_{\mathcal{E}^{\prime}\left[x \mapsto h_{1}\right]} \\
= & \llbracket \varphi[y / \text { heap }] \rrbracket_{\mathcal{E}^{\prime}}\left[y \mapsto h_{1}\right]
\end{aligned}
$$

Together with assumption $h_{1} \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}}$, we get $h_{1} \in \llbracket\left\{y: \tau_{1} \mid \varphi[y /\right.$ heap $\rrbracket\} \rrbracket_{\mathcal{E}^{\prime}}$, and thus $\mathcal{E}^{\prime}\left[x \mapsto h_{1}\right] \vDash \Gamma, x:\left\{y: \tau_{1} \mid \varphi[y /\right.$ heap $\left.]\right\}$.

With assumption $\Gamma, x:\left\{y: \tau_{1} \mid \varphi[y /\right.$ heap $\left.]\right\} \vdash \tau_{2}^{\prime}<: \tau_{2}$, we can conclude that $h_{2} \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}}$. Since we already have that $\llbracket \varphi[x /$ heap $] \rrbracket_{\mathcal{E}\left[y \mapsto h_{2}\right]}=$ true, it follows that $h_{2} \in \llbracket\left\{y: \tau_{2} \mid \varphi[x /\right.$ heap $\rrbracket\} \rrbracket_{\mathcal{E}}$, which is what we wanted to show.

Lemma A. 67 (If Choice Subtype). If heap is the only free binder in $\varphi$ and $\Gamma, x:\{y$ : $\tau_{1}^{\prime} \mid \varphi[y /$ heap $\left.]\right\} \vdash \tau_{12}^{\prime}<: \tau_{12}$, and $\Gamma, x:\left\{y: \tau_{1}^{\prime} \mid \neg \varphi[y /\right.$ heap $\left.]\right\} \vdash \tau_{22}^{\prime}<: \tau_{22}$, then $\Gamma, x: \tau_{1}^{\prime} \vdash\left\{y: \tau_{12}^{\prime} \mid \varphi[x /\right.$ heap $\left.]\right\}+\left\{y: \tau_{22}^{\prime} \mid \varphi[x /\right.$ heap $\left.]\right\}<:\left\{y: \tau_{12} \mid \varphi[x /\right.$ heap $\left.]\right\}+\{y:$ $\tau_{22} \mid \neg \varphi[x /$ heap $\left.]\right\}$

Proof. By Lemmas A. 66 and A. 65 .
Lemma A. 68 (Algorithmic Weakening). If $\Gamma \vdash c:\left(x: \tau_{1}\right) \leadsto \tau_{2}$ and variable $y$ does not appear free in $\tau_{1}$ or $\tau_{2}$, then $\Gamma, y: \tau \vdash c:\left(x: \tau_{1}\right) \leadsto \tau_{2}$ for any heap type $\tau$.

Proof. By induction on the typing derivation.
Lemma A. 69 (Typing Context Subtype). If $\Gamma, x: \tau_{1} \vdash c:\left(y: \tau_{12}\right) \leadsto \tau_{22}$ and $\Gamma \vdash \tau_{1}^{\prime}<$ : $\tau_{1}$, then $\Gamma, x: \tau_{1}^{\prime} \vdash c:\left(y: \tau_{12}\right) \leadsto \tau_{22}$.

Proof. If $x$ is not free in $\tau_{12}$ or $\tau_{22}$, the result follows from Lemma A.68. Otherwise we proceed by induction on the typing derivation. We refer to the general assumptions as follows:
(A) $\Gamma, x: \tau_{1} \vdash c:\left(y: \tau_{12}\right) \leadsto \tau_{22}$
(B) $\Gamma \vdash \tau_{1}^{\prime} \ll \tau_{1}$

Case T-Extract-Algo: By inversion of T-Extract-Algo, we know

$$
\text { (A1) } \Gamma, x: \tau_{1} \vdash \operatorname{extract}(\iota):\left(y: \tau_{12}\right) \leadsto \tau_{22}
$$

(A2) $\tau_{22}=\Sigma z:\left\{v: \iota \mid \varphi_{1}\right\} \cdot\left\{v: \operatorname{chomp}\left(\tau_{12}, \iota, z\right) \mid \varphi_{2}\right\}$
(A3) $\Gamma, x: \tau_{1} \vdash \tau_{12}<:\left\{z: \top| | z \cdot p k t_{i n} \mid \geq \operatorname{sizeof}(\iota)\right\}$
By Lemma A. 58 applied to ( $\mathrm{A}_{3}$ ) and (B) follows
(A4) $\Gamma, x: \tau_{1}^{\prime} \vdash \tau_{12}<:\left\{z: \top| | z \cdot p k t_{i n} \mid \geq \operatorname{sizeof}(\iota)\right\}$
The result follows by T-Extract-Algo.
Case T-Seq-Algo: By inversion of T-Seq-Algo, we know
(A1) $\Gamma, x: \tau_{1} \vdash c_{1} ; c_{2}:\left(y: \tau_{12}\right) \leadsto \tau_{22}$
(A2) $\Gamma, x: \tau_{1} \vdash c_{1}:\left(y: \tau_{12}\right) \leadsto \tau_{12}^{\prime}$
(A3) $\Gamma, x: \tau_{1}, y: \tau_{12} \vdash c_{2}:\left(z: \tau_{12}^{\prime}\right) \leadsto \tau_{22}^{\prime}$
(A4) $\tau_{22}=\tau_{22}^{\prime}\left[z \mapsto \tau_{12}^{\prime}\right]$
By IH applied to (A2) and (B) follows
(A5) $\Gamma, x: \tau_{1}^{\prime} \vdash c_{1}:\left(y: \tau_{12}\right) \leadsto \tau_{12}^{\prime}$
By IH applied to ( $\mathrm{A}_{3}$ ) and (B) follows
(A6) $\Gamma, x: \tau_{1}^{\prime}, y: \tau_{12} \vdash c_{2}:\left(z: \tau_{12}^{\prime}\right) \leadsto \tau_{22}^{\prime}$
The result follows by T-Seq-Algo with (A5) and (A6).
Case T-Skip-Algo: The result immediately follows by T-Skip-Algo.
Case T-Remit-Algo: By inversion of T-Remit-Algo, we know
(A1) $\Gamma, x: \tau_{1} \vdash \tau_{12}<: \iota_{\approx}$
By Lemma A. 58 applied to (A1) and (B) follows
(A2) $\Gamma, x: \tau_{1}^{\prime} \vdash \tau_{12}<: \iota_{\approx}$
The result follows by T-Remit-Algo.
Case T-Reset-Algo: The result immediately follows by T-Reset-Algo.
Case T-Ascribe-Algo: By inversion of T-Ascribe-Algo, we know
(A1) $\Gamma, x: \tau_{1} \vdash c_{0}$ as $\left(y: \hat{\tau}_{12}\right) \rightarrow \tau_{22}:\left(y: \tau_{12}\right) \leadsto \tau_{22}$
(A2) $\Gamma, x: \tau_{1} \vdash c_{0}:\left(y: \hat{\tau}_{12}\right) \leadsto \tau_{22}^{\prime}$
(A3) $\Gamma, x: \tau_{1} \vdash \tau_{12}<: \hat{\tau}_{12}$
(A4) $\Gamma, x: \tau_{1}, y: \hat{\tau}_{12} \vdash \tau_{22}^{\prime}<: \tau_{22}$
By IH applied to ( $\mathrm{A}_{2}$ ) and (B) follows
(A5) $\Gamma, x: \tau_{1}^{\prime} \vdash c_{0}:\left(y: \hat{\tau}_{12}\right) \leadsto \tau_{22}^{\prime}$
By Lemma A. 58 applied to ( $\mathrm{A}_{3}$ ) and (B) follows
(A6) $\Gamma, x: \tau_{1}^{\prime} \vdash \tau_{12}<: \hat{\tau}_{12}$
By Lemma A. 58 applied to ( $\mathrm{A}_{4}$ ) and (B) follows
(A7) $\Gamma, x: \tau_{1}^{\prime}, y: \hat{\tau}_{12} \vdash \tau_{22}^{\prime}<: \tau_{22}$
The result follows by T-Ascribe-Algo with (A5), (A6) and (A7).

## Case T-If-Algo:

By inversion of T-If-Algo, we know
(A1) $\Gamma, x: \tau_{1} \vdash c_{1}:\left(y:\left\{z: \tau_{12} \mid \varphi[z /\right.\right.$ heap $\left.\left.]\right\}\right) \leadsto \tau_{12}^{\prime}$
(A2) $\Gamma, x: \tau_{1} \vdash c_{1}:\left(y:\left\{z: \tau_{12} \mid \neg \varphi[z /\right.\right.$ heap $\left.\left.]\right\}\right) \leadsto \tau_{12}^{\prime \prime}$
(A3) $\tau_{22}=\left\{z: \tau_{12}^{\prime} \mid \varphi[x /\right.$ heap $\left.]\right\}+\left\{z: \tau_{12}^{\prime \prime} \mid \neg \varphi[x /\right.$ heap $\left.]\right\}$
(A4) $\Gamma, x: \tau_{1} ; \tau_{12} \vdash \varphi: \mathbb{B}$
By IH applied to ( A 1 ) and (B) follows
(A5) $\Gamma, x: \tau_{1}^{\prime} \vdash c_{1}:\left(y:\left\{z: \tau_{12} \mid \varphi[z /\right.\right.$ heap $\left.\left.]\right\}\right) \leadsto \tau_{12}^{\prime}$
By IH applied to ( $\mathrm{A}_{2}$ ) and (B) follows
(A6) $\Gamma, x: \tau_{1}^{\prime} \vdash c_{1}:\left(y:\left\{z: \tau_{12} \mid \neg \varphi[z /\right.\right.$ heap $\left.\left.]\right\}\right) \leadsto \tau_{12}^{\prime \prime}$
Since $\Gamma, x: \tau_{1} ; \tau_{12} \vdash \varphi: \mathbb{B}$, it also holds that $\Gamma, x: \tau_{1}^{\prime} ; \tau_{12} \vdash \varphi: \mathbb{B}$. The result then follows by T-If-Algo.

Case T-Mod-Algo: By inversion of T-Mod-Algo, we know
(A1) $\Gamma, x: \tau_{1} \vdash \tau_{12}<: \iota_{\approx}$
(A2) $\Gamma, x: \tau_{1} ; \tau_{12} \vdash e: \mathrm{BV}$
By Lemma A. 58 applied to (A1) and (B) follows
(A3) $\Gamma, x: \tau_{1}^{\prime} \vdash \tau_{12}<: \iota_{\approx}$
Since $\Gamma, x: \tau_{1} ; \tau_{12} \vdash e: \mathrm{BV}$, it also holds that $\Gamma, x: \tau_{1}^{\prime} ; \tau_{12} \vdash e: \mathrm{BV}$. The result follows by T-Mod-Algo.

Case T-Add-Algo: By inversion of T-Add-Algo, we know
(A1) $\Gamma, x: \tau_{1} \vdash \tau_{12}<:\{x: \top \mid \neg x . . \iota$. valid $\}$
By Lemma A. 58 applied to (A1) and (B) follows
(A2) $\Gamma, x: \tau_{1}^{\prime} \vdash \tau_{12}<:\{x: \top \mid \neg x . ı . v a l i d\}$
The result follows by T-Add-Algo.
Case T-Remove-Algo: By inversion of T-Remove-Algo, we know
(A1) $\Gamma, x: \tau_{1} \vdash \tau_{12}<: \iota_{\approx}$
By Lemma A. 58 applied to (A1) and (B) follows
(A2) $\Gamma, x: \tau_{1}^{\prime} \vdash \tau_{12}<: \iota_{\approx}$
The result follows by T-Remove-Algo.

Lemma A. 7 o (Algorithmic Input Subtype). If $\Gamma \vdash \tau_{1}^{\prime}<: \tau_{1}$ and $\Gamma \vdash c:\left(x: \tau_{1}\right) \leadsto \tau_{2}$ such that $x$ is not free in $\tau_{1}$ or $\tau_{1}^{\prime}$, then there exists $\tau_{2}^{\prime}$ such that $\Gamma \vdash c:\left(x: \tau_{1}^{\prime}\right) \leadsto \tau_{2}^{\prime}$ and $\Gamma, x: \tau_{1}^{\prime} \vdash \tau_{2}^{\prime}<: \tau_{2}$.

Proof. By induction on the typing derivation. We refer to the general assumptions as follows:
(A) $\Gamma \vdash \tau_{1}^{\prime} \ll \tau_{1}$
(B) $\Gamma \vdash c:\left(x: \tau_{1}\right) \leadsto \tau_{2}$
(C) $x$ not free in $\tau_{1}$ or $\tau_{1}^{\prime}$

We refer to the proof goals as follows:
(1) $\Gamma \vdash c:\left(x: \tau_{1}^{\prime}\right) \leadsto \tau_{2}^{\prime}$
(2) $\Gamma, x: \tau_{1}^{\prime} \vdash \tau_{2}^{\prime}<: \tau_{2}$

Case T-Extract-Algo: By inversion of T-Extract-Algo, we know
(A1) $c=\operatorname{extract}(\imath)$
(A2) $\Gamma \vdash \tau_{1}<:\left\{x: T| | x \cdot p k t_{i n} \mid \geq \operatorname{sizeof}(\iota)\right\}$
(A3) $\varphi_{1} \triangleq z \cdot p k t_{\text {in }}=z \cdot p k t_{\text {out }}=\langle \rangle$
(A4) $\varphi_{2} \triangleq y . \iota @ z . p k t_{\text {in }} \wedge z . p k t_{\text {out }}=x . p k t_{\text {out }} \wedge z \equiv$, $x$
(A5) $\tau_{21} \triangleq\left\{z: \iota \mid \varphi_{1}\right\}$
(A6) $\tau_{22} \triangleq\left\{z: \operatorname{chomp}\left(\tau_{1}, l, y\right) \mid \varphi_{2}\right\}$
(A7) $\tau_{2}=\Sigma y: \tau_{21} \cdot \tau_{22}$
By Lemma A. 56 with (A) and (A2),
(A8) $\Gamma \vdash \tau_{1}^{\prime}<:\left\{x: \top| | x . p k t_{i n} \mid \geq \operatorname{sizeof}(\iota)\right\}$
By Lemma A. 60 with (A2) and (A),
(A9) Г, $y: \tau_{21} \vdash \operatorname{chomp}\left(\tau_{1}^{\prime}, l, y\right)<: \operatorname{chomp}\left(\tau_{1}, l, y\right)$
(1) follows by T-Extract-Algo. (2) follows by Lemmas A. 61 and A. 63 and (A9).

Case T-Seq-Algo: By inversion of T-Seq-Algo, we know
(A1) $c=c_{1} ; c_{2}$
(A2) $\Gamma \vdash c_{1}:\left(x: \tau_{1}\right) \leadsto \tau_{12}$
(A3) $\Gamma, x: \tau_{1} \vdash c_{2}:\left(y: \tau_{12}\right) \leadsto \tau_{22}$
(A4) $\tau_{2}=\tau_{22}\left[y \mapsto \tau_{12}\right]$
By IH applied to (A2) and (A), there is some $\tau_{12}^{\prime}$ such that
(A5) $\Gamma \vdash c_{1}:\left(x: \tau_{1}^{\prime}\right) \leadsto \tau_{12}^{\prime}$, and
(A6) $\Gamma, x: \tau_{1}^{\prime} \vdash \tau_{12}^{\prime}<: \tau_{12}$
By Lemma A. 69 with (A) and (A3),
(A7) $\Gamma, x: \tau_{1}^{\prime} \vdash c_{2}:\left(y: \tau_{12}\right) \leadsto \tau_{22}$
By IH applied to (A7) and (A6), there is some $\tau_{22}^{\prime}$ such that
(A8) $\Gamma, x: \tau_{1}^{\prime} \vdash c_{2}:\left(y: \tau_{12}^{\prime}\right) \leadsto \tau_{22}^{\prime}$, and
(A9) $\Gamma, x: \tau_{1}^{\prime}, y: \tau_{12}^{\prime} \vdash \tau_{22}^{\prime}<: \tau_{22}$
(1) follows by T-Seq-AlGo with (A5) and (A8). To show (2), we just need to show that $\Gamma, x: \tau_{1}^{\prime} \vdash \tau_{22}^{\prime}\left[y \mapsto \tau_{12}^{\prime}\right]<: \tau_{22}\left[y \mapsto \tau_{12}\right]$. This follows by Lemma A. 64 applied to (A6) and (A9).

Case T-Skip-Algo: Immediate by T-Skip-Algo.
Case T-Remit-Algo: By inversion of T-Remit-Algo, we know
(A1) $\Gamma \vdash \tau_{1}<: \iota_{\approx}$
(A2) $c=\operatorname{remit}(\iota)$
(A3) $\tau_{2}=\Sigma y:\left\{z: \tau_{1} \mid z \equiv x\right\} .\left\{z: \epsilon \mid z \cdot p k t_{\text {in }}=\langle \rangle \wedge z . p k t_{\text {out }}=x . \iota\right\}$
By Lemma A. 56 with (A1) and (A) follows
(A4) $\Gamma \vdash \tau_{1}^{\prime}<: \iota_{\approx}$
Let $\tau_{2}^{\prime}=\Sigma y:\left\{z: \tau_{1}^{\prime} \mid z \equiv x\right\} .\left\{z: \epsilon \mid z \cdot p k t_{\text {in }}=\langle \rangle \wedge z . p k t_{\text {out }}=x . \iota\right\}$. (1) follows by T-Remit-Algo. By Lemma A.58, and since $x$ does not occur free in $\tau_{1}$ or $\tau_{1}^{\prime}$, we know
(A5) $\Gamma, x: \tau_{1}^{\prime} \vdash \tau_{1}^{\prime}<: \tau_{1}$
By Lemma A. 61 we know
(A6) $\Gamma, x: \tau_{1}^{\prime} \vdash\left\{z: \tau_{1}^{\prime} \mid z \equiv x\right\}<:\left\{z: \tau_{1} \mid z \equiv x\right\}$
By Lemma A. 62 with $\tau_{22}=\left\{z: \epsilon \mid z . p k t_{\text {in }}=\langle \rangle \wedge z . p k t_{\text {out }}=x . l\right\}$ follows
(A7) $\Gamma, x: \tau_{1}^{\prime} \vdash \Sigma y:\left\{z: \tau_{1}^{\prime} \mid z \equiv x\right\} . \tau_{22}<: \Sigma y:\left\{z: \tau_{1} \mid z \equiv x\right\} . \tau_{22}$
This shows (2) and concludes this case.
Case T-Reset-Algo: By inversion of T-Reset-Algo, we know
(A1) $\tau_{2}=\Sigma y:\left\{z: \epsilon \mid z \cdot p k t_{\text {out }}=\langle \rangle \wedge z \cdot p k t_{\text {in }}=x \cdot p k t_{\text {out }}\right\} \cdot\left\{z: \epsilon \mid z \cdot p k t_{\text {out }}=\right.$ $\left\rangle \wedge z . p k t_{i n}=z . p k t_{i n}\right\}$

Let $\tau_{2}^{\prime}=\tau_{2}$. (1) follows by T-Reset-Algo and (2) follows by Lemma A. 55 .
Case T-Ascribe-Algo: By inversion of T-Ascribe-Algo, we know
(A1) $c=c_{0}$ as $\left(x: \hat{\tau}_{1}\right) \rightarrow \tau_{2}$
(A2) $\Gamma \vdash \mathcal{c}_{0}:\left(x: \hat{\tau}_{1}\right) \leadsto \tau_{c}$
(A3) $\Gamma \vdash \tau_{1}<: \hat{\tau}_{1}$
(A4) $\Gamma, x: \hat{\tau}_{1} \vdash \tau_{c}<: \tau_{2}$
By Lemma A. 56 applied to (A) and (A3) follows that
(A5) $\Gamma \vdash \tau_{1}^{\prime}<: \hat{\tau}_{1}$

Let $\tau_{2}^{\prime}=\tau_{2}$. (1) follows by T-Ascribe-Algo with (A2), (A4) and (A5). (2) follows by Lemma A. 55 .

Case T-If-Algo: By inversion of T-If-Algo, we know
(A1) $c=i f(\varphi) c_{1}$ else $c_{2}$
(A2) $\Gamma ; \tau_{1} \vdash \varphi: \mathbb{B}$
(A3) $\Gamma \vdash c_{1}:\left(x:\left\{y: \tau_{1} \mid \varphi[y /\right.\right.$ heap $\left.\left.]\right\}\right) \leadsto \tau_{12}$
(A4) $\Gamma \vdash c_{2}:\left(x:\left\{y: \tau_{1} \mid \neg \varphi[y /\right.\right.$ heap $\left.\left.]\right\}\right) \leadsto \tau_{22}$
(A5) $\tau_{2}=\left\{y: \tau_{12} \mid \varphi[x /\right.$ heap $\left.]\right\}+\left\{y: \tau_{22} \mid \neg \varphi[x /\right.$ heap $\left.]\right\}$
By (A) and Lemma A.61, we know
(A6) $\Gamma \vdash\left\{y: \tau_{1}^{\prime} \mid \varphi[y /\right.$ heap $\left.]\right\}<:\left\{y: \tau_{1} \mid \varphi[y /\right.$ heap $\left.]\right\}$, and
(A7) $\Gamma \vdash\left\{y: \tau_{1}^{\prime} \mid \neg \varphi[y /\right.$ heap $\left.]\right\}<:\left\{y: \tau_{1} \mid \neg \varphi[y /\right.$ heap $\left.]\right\}$
By applying the IH to (A6) and ( $\mathrm{A}_{3}$ ) we get $\tau_{12}^{\prime}$ such that
(A8) $\Gamma \vdash c_{1}:\left(x:\left\{y: \tau_{1}^{\prime} \mid \varphi[y /\right.\right.$ heap $\left.\left.]\right\}\right) \leadsto \tau_{12}^{\prime}$, and
(A9) $\Gamma, x:\left\{y: \tau_{1}^{\prime} \mid \varphi[y /\right.$ heap $\left.]\right\} \vdash \tau_{12}^{\prime}<: \tau_{12}$
By applying the IH to ( $\mathrm{A}_{7}$ ) and ( $\mathrm{A}_{4}$ ) we get $\tau_{22}^{\prime}$ such that
(A1o) $\Gamma \vdash c_{2}:\left(x:\left\{y: \tau_{1}^{\prime} \mid \neg \varphi[y /\right.\right.$ heap $\left.\left.]\right\}\right) \leadsto \tau_{22}^{\prime}$, and
(A11) $\Gamma, x:\left\{y: \tau_{1}^{\prime} \mid \neg \varphi[y /\right.$ heap $\left.]\right\} \vdash \tau_{22}^{\prime}<: \tau_{22}$
From (A2) and (A), we can conclude that
(A12) $\Gamma ; \tau_{1}^{\prime} \vdash \varphi: \mathbb{B}$
Let $\tau_{2}^{\prime}=\left\{y: \tau_{12}^{\prime} \mid \varphi[x /\right.$ heap $\left.]\right\}+\left\{y: \tau_{22}^{\prime} \mid \neg \varphi[x /\right.$ heap $\left.]\right\}$. (2) follows by Lemma A.67. (1) follows by T-If-Algo with (A8), (A10) and (A12).

Case T-Mod-Algo: By inversion of T-Mod-Algo, we know
(A1) $\Gamma \vdash \tau_{1}<: \iota_{\approx}$
(A2) $\Gamma ; \tau_{1} \vdash e: B V$
(A3) $\tau_{2}=\left\{y: T \mid \varphi \wedge \varphi_{\imath} \wedge \varphi_{f} \wedge y . t . f=e[x /\right.$ heap $\left.]\right\}$
By Lemma A. 56 with (A) and (A1) follows
(A4) $\Gamma \vdash \tau_{1}^{\prime}<: \iota_{\approx}$
From (A2) and (A), we can conclude
(A5) $\Gamma ; \tau_{1}^{\prime} \vdash e: B V$
Let $\tau_{2}^{\prime}=\tau_{2}$. (1) follows by T-Mod-Algo with ( $\mathrm{A}_{4}$ ) and ( $\mathrm{A}_{5}$ ). (2) follows by Lemma A. 55 .

Case T-Add-Algo: By inversion of T-Add-Algo, we know
(A1) $\Gamma \vdash \tau_{1}<:\{x: \top \mid \neg x . \iota$. valid $\}$
(A2) $\tau_{2}=\Sigma y:\left\{z: \tau_{1} \mid z \equiv x\right\} .\left\{z: \iota \mid z . p k t_{\text {in }}=z . p k t_{\text {out }}=\langle \rangle \wedge z . \iota=v\right\}$

By Lemma A. 56 with (A1) and (A) follows
(A3) $\Gamma \vdash \tau_{1}^{\prime}<:\{x: \top \mid \neg x . ı$. valid $\}$
Let $\tau_{2}^{\prime}=\Sigma y:\left\{z: \tau_{1}^{\prime} \mid z \equiv x\right\} .\left\{z: \iota \mid z \cdot p k t_{\text {in }}=z . p k t_{\text {out }}=\langle \rangle \wedge z . \iota=v\right\}$. (1) follows by T-Add-Algo. By Lemma A. 58 with (A) and (C),
(A4) $\Gamma, x: \tau_{1}^{\prime} \vdash \tau_{1}^{\prime}<: \tau_{1}$
By Lemma A. 61 follows
(A5) $\Gamma, x: \tau_{1}^{\prime} \vdash\left\{z: \tau_{1}^{\prime} \mid z \equiv x\right\}<:\left\{z: \tau_{1} \mid z \equiv x\right\}$
(2) follows by Lemma A. 62 with ( $\mathrm{A}_{5}$ ).

Case T-Remove-Algo: By inversion of T-Remove-Algo, we know
(A1) $\Gamma \vdash \tau_{1}<: \iota_{\approx}$
By Lemma A. 56 with ( $\mathrm{A}_{1}$ ) and (A),
(A2) $\Gamma \vdash \tau_{1}^{\prime}<: \iota_{\approx}$
Let $\tau_{2}^{\prime}=\tau_{2}$. (1) follows by T-Remove-Algo with (A2). (2) follows by Lemma A. 55 .

Lemma A. 71 (Includes Subtype). Includes $\Gamma \tau \iota \Longleftrightarrow \Gamma \vdash \tau<\iota_{\approx}$
Proof. Prove each direction separately
$(\Rightarrow)$ Assume Includes $\Gamma \tau \iota$. Let $\mathcal{E} \vDash \Gamma$ and $h \in \llbracket \tau \rrbracket_{\mathcal{E}}$. By definition of the inclusion relation, $\iota \in \operatorname{dom}(h)$. By the definition of subtyping, it suffices to show $h \in \llbracket \iota_{\sim} \rrbracket \mathcal{E}$. By definition, $\llbracket \iota_{\approx} \rrbracket_{\mathcal{E}}=\{h \mid \iota \in \operatorname{dom}(h)\}$, and we're done.
$(\Leftarrow)$ Assume $\Gamma \vdash \tau<: \iota_{\approx}$. Show Includes $\Gamma \tau \iota$. To that end, let $\mathcal{E} \vDash \Gamma$ and $h \in \llbracket \tau \rrbracket \mathcal{E}$ be arbitrary. By the definition of subtyping, $h \in \llbracket \iota_{\approx} \rrbracket \mathcal{E}$. By definition of the semantics, conclude $t \in \operatorname{dom}(h)$.

Lemma A. 72 (Excludes Subtype). Excludes $\Gamma \tau \iota$ iff $\Gamma \vdash \tau<:\{x: \top \mid \neg x . \iota$.valid $\}$
Proof. Prove each direction separately
$(\Rightarrow)$ Assume Excludes $\Gamma \tau \iota$. Let $\mathcal{E} \vDash \Gamma$ and $h \in \llbracket \tau \rrbracket \mathcal{E}$. By definition of the exclusion relation, $\iota \notin \operatorname{dom}(h)$. By the definition of subtyping, it suffices to show $h \in \llbracket\{x$ : $\top \mid \neg x . \iota . v a l i d\} \rrbracket_{\mathcal{E}}$. By definition, $\llbracket\{x: \top \mid \neg x . \iota . v a l i d ~\} \rrbracket_{\mathcal{E}}=\{h \mid \iota \notin \operatorname{dom}(h)\}$, and we're done.
$(\Leftarrow)$ Assume $\Gamma \vdash \tau<:\{x: \top \mid \neg x . ı . v a l i d\}$. Show Excludes $\Gamma \tau \iota$. To that end, let $\mathcal{E} \vDash \Gamma$ and $h \in \llbracket \tau \rrbracket_{\mathcal{E}}$ be arbitrary. By the definition of subtyping, $h \in \llbracket\{x: \top \mid$ $\neg x . . \iota . v a l i d\}\rangle \rrbracket_{\mathcal{E}}$. By definition of the semantics, conclude $\iota \notin \operatorname{dom}(h)$.

Theorem A. 73 (Algorithmic Typing Correctness). For all subtyping contexts $\Gamma$, commands $c$, variables $x$, heap types $\tau_{1}$, and $\tau_{2}$, where $x$ is not free in $\tau_{1}, \Gamma \vdash c:\left(x: \tau_{1}\right) \rightarrow \tau_{2}$ if and only if there is some $\tau_{2}^{\prime}$ such that $\Gamma \vdash c:\left(x: \tau_{1}\right) \leadsto \tau_{2}^{\prime}$, and $\Gamma,\left(x: \tau_{1}\right) \vdash \tau_{2}^{\prime} \ll \tau_{2}$.

Proof. $\quad(\Rightarrow)$ Assume $\Gamma \vdash c:\left(x: \tau_{1}\right) \rightarrow \tau_{2}$. Proceed by induction on the typing derivation, leaving $\Gamma$ general. We refer to the proof goals as follows:
(1) $\Gamma \vdash c:\left(x: \tau_{1}\right) \leadsto \tau_{2}^{\prime}$
(2) $\Gamma,\left(x: \tau_{1}\right) \vdash \tau_{2}^{\prime}<: \tau_{2}$

## Case T-Extract:

(A1) $c=\operatorname{extract}(\iota)$
(A2) $\Gamma \vdash \operatorname{sizeof}_{p k t_{i n}}\left(\tau_{1}\right) \geq \operatorname{sizeof}(i)$
(A3) $\varphi_{1} \triangleq z \cdot p k t_{\text {in }}=z \cdot p k t_{o u t}=\langle \rangle$
(A4) $\varphi_{2} \triangleq y . \iota @ z . p k t_{i n}=x . p k t_{\text {in }} \wedge z . p k t_{\text {out }}=x . p k t_{\text {out }} \wedge z \equiv, x$
(A5) $\tau_{2}=\Sigma y:\left\{z: \iota \mid \varphi_{1}\right\} .\left\{z: \operatorname{chomp}\left(\tau_{1}, \iota, y\right) \mid \varphi_{2}\right\}$
The only algorithmic rule that applies to extract $(\iota)$ is T-Extract-Algo. Since $\Gamma \vdash \tau_{1}<:\left\{x: \top| | x . p k t_{i n} \mid \geq \operatorname{sizeof}(\iota)\right\}$ by (A2) and Lemma A.59, T-Extract-Algo produces $\tau_{2}^{\prime}$ such that
(A6) $\varphi_{1}^{\prime} \triangleq z \cdot p k t_{\text {in }}=z \cdot p k t_{\text {out }}=\langle \rangle$
(A7) $\varphi_{2}^{\prime} \triangleq y . \iota @ z . p k t_{\text {in }}=x . p k t_{\text {in }} \wedge z . p k t_{\text {out }}=x \cdot p k t_{\text {out }} \wedge z \equiv x$
(A8) $\tau_{2}^{\prime}=\Sigma y:\left\{z: \iota \mid \varphi_{1}^{\prime}\right\} \cdot\left\{z: \operatorname{chomp}\left(\tau_{1}, \iota, y\right) \mid \varphi_{2}^{\prime}\right\}$
which shows (1). (2) follows by Lemma A. 55 .

## Case T-SEQ:

(A1) $c=c_{1} ; c_{2}$
(A2) $\Gamma \vdash c_{1}:\left(x: \tau_{1}\right) \rightarrow \tau_{12}$
(A3) $\Gamma,\left(x: \tau_{1}\right) \vdash c_{2}:\left(y: \tau_{12}\right) \rightarrow \tau_{22}$
(A4) $\tau_{2}=\tau_{22}\left[x \mapsto \tau_{12}\right]$
By applying the IH to ( A 2 ), we get $\tau_{12}^{\prime}$ such that
(A5) $\Gamma \vdash c_{1}:\left(x: \tau_{1}\right) \leadsto \tau_{12}^{\prime}$, and
(A6) $\Gamma, x: \tau_{1} \vdash \tau_{12}^{\prime}<: \tau_{12}$
By applying the IH to ( $\mathrm{A}_{3}$ ), we get $\tau_{22}^{\prime}$ such that
(A7) $\Gamma, x: \tau_{1} \vdash c_{2}:\left(y: \tau_{12}\right) \leadsto \tau_{22}^{\prime}$, and
(A8) $\Gamma, x: \tau_{1}, y: \tau_{12} \vdash \tau_{22}^{\prime}<: \tau_{22}$
By Lemma A.7o with (A6) and (A7) there exists $\tau_{22}^{\prime \prime}$ such that
(A9) $\Gamma, x: \tau_{1} \vdash c_{2}:\left(y: \tau_{12}^{\prime}\right) \leadsto \tau_{22}^{\prime \prime}$, and
(A10) $\Gamma, x: \tau_{1}, y: \tau_{12}^{\prime} \vdash \tau_{22}^{\prime \prime}<: \tau_{22}^{\prime}$
By T-Seq-Algo with (A5) and (A9) follows
(A11) $\Gamma \vdash c_{1} ; c_{2}:\left(x: \tau_{1}\right) \leadsto \tau_{22}^{\prime \prime}\left[y \mapsto \tau_{12}^{\prime}\right]$
which shows (1). By Lemma A. 58 with (A6) and (A8) follows
(A12) $\Gamma, x: \tau_{1}, y: \tau_{12}^{\prime} \vdash \tau_{22}^{\prime}<: \tau_{22}$
By Lemma A. 56 with (A10) and (A12) follows
(A13) $\Gamma, x: \tau_{1}, y: \tau_{12}^{\prime} \vdash \tau_{22}^{\prime \prime}<: \tau_{22}$

By Lemma A. 64 follows
(A14) $\Gamma, x: \tau_{1} \vdash \tau_{22}^{\prime \prime}\left[y \mapsto \tau_{12}^{\prime}\right]<: \tau_{22}\left[y \mapsto \tau_{12}\right]$ which shows (2) and concludes this case.
Case T-Skip: (1) follows by T-Skip-Algo and (2) follows by Lemma A. 55.
Case T-Remit: By inversion of T-Remit, we know
(A1) $c=\operatorname{remit}(\iota)$
(A2) Includes $\Gamma \tau_{1} \iota$
(A3) $\varphi \triangleq z . p k t_{\text {in }}=\langle \rangle \wedge z . p k t_{\text {out }}=x .1$
(A4) $\tau_{2}=\Sigma y:\left\{z: \tau_{1} \mid z \equiv x\right\} .\{z: \epsilon \mid \varphi\}$
By Lemma A. 71 and (A2), T-Remit-Algo computes $\tau_{2}^{\prime}$ such that
(A5) $\Gamma \vdash \operatorname{remit}\left(\iota_{i}\right):\left(\left(x: \tau_{1}\right)\right) \leadsto \tau_{2}^{\prime}$, and
(A6) $\tau_{2}^{\prime}=\Sigma y:\left\{z: \tau_{1} \mid z \equiv x\right\} .\left\{z: \epsilon \mid z . p k t_{i n}=\langle \rangle \wedge z . p k t_{i n}=x . l\right\}$
which shows (1). Since $\tau_{2}^{\prime}=\tau_{2}$, (2) follows by Lemma A. 55 .
Case T-Reset: (1) follows by T-Reset-Algo and (2) follows by Lemma A. 55.
Case T-Ascribe: By inversion of T-Ascribe, we know
(A1) $c=c_{0}$ as $\left(x: \tau_{1}\right) \rightarrow \tau_{2}$
(A2) $\Gamma \vdash \mathcal{c}_{0}:\left(x: \tau_{1}\right) \rightarrow \tau_{2}$
By IH applied to (A2), there exists $\hat{\tau}_{2}$ such that
(A3) $\Gamma \vdash c_{0}:\left(x: \tau_{1}\right) \leadsto \hat{\tau}_{2}$
(A4) $\Gamma, x: \tau_{1} \vdash \hat{\tau}_{2}<: \tau_{2}$
By T-Ascribe-Algo with (A3), (A4) and Lemma A.55,
(A5) $\Gamma \vdash c_{0}$ as $\left(x: \tau_{1}\right) \rightarrow \tau_{2}:\left(x: \tau_{1}\right) \leadsto \tau_{2}$
showing (1). (2) follows by Lemma A. 55 .
Case T-IF: By inversion of T-IF, we know
(A1) $c=i f(\varphi) c_{1}$ else $c_{2}$
(A2) $\Gamma ; \tau_{1} \vdash \varphi: \mathbb{B}$
$\left(\mathrm{A}_{3}\right) \Gamma \vdash c_{1}:\left(x:\left\{y: \tau_{1} \mid \varphi[y /\right.\right.$ heap $\left.\left.]\right\}\right) \rightarrow \tau_{12}$
(A4) $\Gamma \vdash c_{2}:\left(x:\left\{y: \tau_{1} \mid \neg \varphi[y /\right.\right.$ heap $\left.\left.]\right\}\right) \rightarrow \tau_{22}$
(A5) $\tau_{2}=\left\{y: \tau_{12} \mid \varphi[x /\right.$ heap $\left.]\right\}+\left\{y: \tau_{22} \mid \neg \varphi[x /\right.$ heap $\left.]\right\}$
By the IH applied to (A3) there exists $\tau_{12}^{\prime}$ such that
(A6) $\Gamma \vdash c_{1}:\left(x:\left\{y: \tau_{1} \mid \varphi[y /\right.\right.$ heap $\left.\left.]\right\}\right) \leadsto \tau_{12}^{\prime}$
(A7) $\Gamma, x:\left\{y: \tau_{1} \mid \varphi[y /\right.$ heap $\left.]\right\} \vdash \tau_{12}^{\prime}<: \tau_{12}$
By the IH applied to (A4) there exists $\tau_{22}^{\prime}$ such that
(A8) $\Gamma \vdash c_{2}:\left(x:\left\{y: \tau_{1} \mid \neg \varphi[y /\right.\right.$ heap $\left.\left.]\right\}\right) \leadsto \tau_{22}^{\prime}$
(A9) $\Gamma, x:\left\{y: \tau_{1} \mid \neg \varphi[y /\right.$ heap $\left.]\right\} \vdash \tau_{22}^{\prime}<: \tau_{22}$
(1) follows by T-IF-Algo with (A2), (A6) and (A8).
(2) follows by Lemma A. 67 with (A7) and (A9).

Case T-Mod: By inversion of T-Mod, we know
(A1) $c=\iota . f:=e$
(A2) Includes $\Gamma \tau_{1} \iota$
(A3) $\mathcal{F}(\iota, f)=\mathrm{BV}$
(A4) $\Gamma ; \tau_{1} \vdash e: \mathrm{BV}$
(A5) $\tau_{2}=\left\{y: \mathrm{T} \mid \varphi_{p k t} \wedge \varphi_{t} \wedge \varphi_{f} \wedge y . t . f=e[x /\right.$ heap $\left.]\right\}$
By Lemma A. 71 and (A2),
(A6) $\Gamma \vdash \tau_{1}<: \iota_{i}$
(1) follows by T-Mod-Algo with (A3),(A4), (A5), and (A6). (2) follows by Lemma A. 55 .
Case T-Add: By inversion of T-Add, we know
(A1) Excludes $\Gamma \tau_{1} \iota$
(A2) init $_{\mathcal{H} \mathcal{T}(\iota)}=v$
(A3) $\tau_{2}=\Sigma y:\left\{z: \tau_{1} \mid z \equiv x\right\} \cdot\left\{z: \iota \mid z \cdot p k t_{i n}=z . p k t_{\text {out }}=\langle \rangle \wedge z . \iota=v\right\}$
By Lemma A. 72 and (A1),
(A4) $\Gamma \vdash \tau<:\{x: \top \mid \neg x . ı$. valid $\}$
(1) follows by T-Add-Algo with (A1) and (A2). (2) follows by Lemma A. 55 .

Case T-Sub: By inversion of T-Sub, there exists some $\tau_{3}$ and $\tau_{4}$ such that
(A1) $\Gamma \vdash \tau_{1}<: \tau_{3}$
(A2) $\Gamma, x: \tau_{1} \vdash \tau_{4}<: \tau_{2}$
(A3) $\Gamma \vdash c:\left(x: \tau_{3}\right) \rightarrow \tau_{4}$
By applying the IH to $\left(\mathrm{A}_{3}\right)$ there is some $\tau_{4}^{\prime}$ such that
(A4) $\Gamma \vdash c:\left(x: \tau_{3}\right) \leadsto \tau_{4}^{\prime}$, and
(A5) $\Gamma, x: \tau_{3} \vdash \tau_{4}^{\prime}<: \tau_{4}$
By Lemma A. 58 together with (A1) and (A5), follows
(A6) $\Gamma, x: \tau_{1} \vdash \tau_{4}^{\prime}<: \tau_{4}$
By applying Lemma A. 70 to (A1) and (A4) we get $\tau_{4}^{\prime \prime}$ such that
(A7) $\Gamma \vdash c:\left(x: \tau_{1}\right) \leadsto \tau_{4}^{\prime \prime}$, and
(A8) $\Gamma, x: \tau_{1} \vdash \tau_{4}^{\prime \prime}<: \tau_{4}^{\prime}$
(1) follows by (A7). (2) follows by repeated application of Lemma A. 56 with (A2), (A6) and (A8).
Case T-Remove: By inversion of T-Remove, we know
(A1) Includes $\Gamma \tau_{1} /$
(A2) $\tau_{2}=\left\{y: \top \mid \varphi_{\iota} \wedge \varphi_{p k t} \wedge \neg y . \iota\right.$. valid $\}$
By Lemma A. 71 with (A1), we get
(A3) $\Gamma \vdash \tau_{1}<: \iota_{\approx}$
(1) follows by T-Remove-Algo. (2) follows by Lemma A. 55 .
$(\Leftarrow)$ Proceed by induction on the typing derivation. We refer to the general assumptions as follows:
(A) $\Gamma \vdash c:\left(x: \tau_{1}\right) \leadsto \tau_{2}^{\prime}$
(B) $\Gamma,\left(x: \tau_{1}\right) \vdash \tau_{2}^{\prime} \ll \tau_{2}$

Case T-Extract-Algo: By inversion of T-Extract-Algo, we know
(A1) $\Gamma \vdash \tau_{1}<:\left\{x: \top| | x . p k t_{i n} \mid \geq \operatorname{sizeof}(\iota)\right\}$
(A2) $\tau_{2}^{\prime}=\Sigma y:\left\{z: \iota \mid \varphi_{1}\right\} \cdot\left\{z: \operatorname{chomp}\left(\tau_{1}, \iota, y\right) \mid \varphi_{2}\right\}$
By Lemma A. 59 follows
(A3) $\Gamma \vdash \operatorname{sizeof}_{p k t_{i n}}\left(\tau_{1}\right) \geq \operatorname{sizeof}(\iota)$
The result follows by T-Extract with ( $\mathrm{A}_{3}$ ) and Lemma A. 55 .
Case T-Seq-Algo: By inversion of T-Seq-Algo, we know
(A1) $c=c_{1} ; c_{2}$
(A2) $\Gamma \vdash c_{1}:\left(\left(x: \tau_{1}\right)\right) \leadsto \tau_{12}$
(A3) $\Gamma,\left(x: \tau_{1}\right) \vdash c_{2}:\left(\left(y: \tau_{12}\right)\right) \leadsto \tau_{22}$
(A4) $\tau_{2}^{\prime}=\tau_{22}\left[y \mapsto \tau_{12}\right]$
(A5) $\Gamma, x: \tau_{1} \vdash \tau_{2}^{\prime}<: \tau_{2}$
With Lemma A. 55 follows
(A6) $\Gamma, x: \tau_{1} \vdash \tau_{12}<: \tau_{12}$
By IH with (A2) and (A6) follows
(A7) $\Gamma \vdash c_{1}:\left(x: \tau_{1}\right) \rightarrow \tau_{12}$
Similarly, applying the IH to (A3) gives
(A8) $\Gamma,\left(x: \tau_{1}\right) \vdash c_{2}:\left(x: \tau_{12}\right) \rightarrow \tau_{22}$
By T-Seq with (A7) and (A8) follows
(A9) $\Gamma \vdash c_{1} ; c_{2}:\left(x: \tau_{1}\right) \rightarrow \tau_{22}\left[y \mapsto \tau_{12}\right]$
By (A4), (A5), Lemma A. 55 and T-Sub follows
(A10) $\Gamma \vdash c_{1} ; c_{2}:\left(x: \tau_{1}\right) \rightarrow \tau_{2}$
which concludes this case.

## Case T-Skip-Algo:

The result follows by T-Skip, Lemma A.55, and T-Sub.
Case T-Remit-Algo: The result follows by T-Remit, Lemma A.71, Lemma A.55, and T-Sub.

Case T-Reset-Algo: The result follows by T-Reset, Lemma A.55, and T-Sub.
Case T-Ascribe-Algo: By inversion of T-Ascribe-Algo, we know
(A1) $c=c_{0}$ as $\left(x: \hat{\tau}_{1}\right) \rightarrow \tau_{2}^{\prime}$
(A2) $\Gamma \vdash c_{0}:\left(x: \hat{\tau}_{1}\right) \leadsto \tau_{2}^{\prime \prime}$
(A3) $\Gamma \vdash \tau_{1}<: \hat{\tau}_{1}$
(A4) $\Gamma, x: \hat{\tau}_{1} \vdash \tau_{2}^{\prime \prime}<: \tau_{2}^{\prime}$
By IH applied to (A2) and (A4), we get
(A5) $\Gamma \vdash c_{0}:\left(x: \hat{\tau}_{1}\right) \rightarrow \tau_{2}^{\prime}$
By T-Ascribe follows from (A5) that
(A6) $\Gamma \vdash c_{0}$ as $\left(x: \hat{\tau}_{1}\right) \rightarrow \tau_{2}^{\prime}:\left(x: \hat{\tau}_{1}\right) \rightarrow \tau_{2}^{\prime}$
The result follows by T-SUB with assumptions (B), (A3) and (A6).
Case T-If-Algo: By inversion of T-If-Algo, we know
$\left(\mathrm{A}_{1}\right) c=i f(\varphi) c_{1}$ else $c_{2}$
(A2) $\Gamma ; \tau_{1} \vdash \varphi: \mathbb{B}$
(A3) $\Gamma \vdash c_{1}:\left(x:\left\{y: \tau_{1} \mid \varphi[y /\right.\right.$ heap $\left.\left.]\right\}\right) \leadsto \tau_{12}$
(A4) $\Gamma \vdash c_{2}:\left(x:\left\{y: \tau_{1} \mid \neg \varphi[y /\right.\right.$ heap $\left.\left.]\right\}\right) \leadsto \tau_{22}$
(A5) $\tau_{2}^{\prime}=\left\{y: \tau_{12} \mid \varphi[x /\right.$ heap $\left.]\right\}+\left\{y: \tau_{22} \mid \neg \varphi[x /\right.$ heap $\left.]\right\}$
By Lemma A.55, we can conclude
(A6) $\Gamma, x:\left\{y: \tau_{1} \mid \varphi[y /\right.$ heap $\left.]\right\} \vdash \tau_{12}<: \tau_{12}$
By IH with ( $\mathrm{A}_{3}$ ) and (A6) follows
(A7) $\Gamma \vdash c_{1}:\left(x:\left\{y: \tau_{1} \mid \varphi[y /\right.\right.$ heap $\left.\left.]\right\}\right) \rightarrow \tau_{12}$
We can reason similarly as before to conclude
(A8) $\Gamma \vdash c_{2}:\left(x:\left\{y: \tau_{1} \mid \neg \varphi[y /\right.\right.$ heap $\left.\left.]\right\}\right) \rightarrow \tau_{22}$
The result follows by T-If with (A2), (A7), (A8) and (A5) and by T-SuB with assumption (B).
Case T-Mod-Algo: The result follows by T-Mod with Lemmas A. 55 and A. 71 and T-Sub.

Case T-Add-Algo: The result follows by T-Add with Lemmas A. 55 and A. 72 and T-Sub.

Case T-Remove-Algo: The result follows by T-Remove with Lemmas A. 55 and A. 71 and T-Sub.

## A.2.3 Decidability of Typechecking

Lemma A. 74 (Refinement Bound). For every $\Gamma, x, \tau, \varphi, N$, such that $\Gamma \vdash \tau \leq N$, $\Gamma \vdash\{x: \tau \mid \varphi\} \leq N$.

Proof. Let $\Gamma, x, \tau, \varphi, N$, be given such that $\Gamma \vdash \tau \leq N$. Let $\mathcal{E} \vDash \Gamma$. Further, let $h \in \llbracket\{x: \tau \mid \varphi\} \rrbracket \mathcal{E}$. By the semantics of heap types, we also know that $h \in \llbracket \tau \rrbracket \mathcal{E}$. Assumption $\Gamma \vdash \tau \leq N$ gives us that $\left|h\left(p k t_{\text {in }}\right)\right|+\left|h\left(p k t_{\text {out }}\right)\right| \leq N$, which is what we want to show.

Lemma A. 75 (Bound Constraints). For every $\Gamma, x, y, \tau_{1}, \tau_{2}$, and $\varphi$, such that heap is the only free variable in $\varphi, \Gamma, x: \tau_{1} \vdash\left\{y: \tau_{2} \mid \varphi[x /\right.$ heap $\left.]\right\} \leq N$, if and only if $\Gamma, x:\left\{y: \tau_{1} \mid \varphi[y /\right.$ heap $\left.]\right\} \vdash \tau_{2} \leq N$.

Proof. Let $\Gamma, x, y, \tau_{1}, \tau_{2}$, and $\varphi$ be given. Prove each direction separately.
$(\Rightarrow)$ Assume $\Gamma, x: \tau_{1} \vdash\left\{y: \tau_{2} \mid \varphi[x /\right.$ heap $\left.]\right\} \leq N$.
Let $\mathcal{E} \vDash \Gamma, x:\left\{y: \tau_{1} \mid \varphi[y /\right.$ heap $\left.]\right\}$. We can write $\mathcal{E}=\mathcal{E}^{\prime}[x \mapsto h]$ for some $\mathcal{E}^{\prime} \vDash \Gamma$, and some $h_{1} \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}^{\prime}}$, such that $\llbracket \varphi[y /$ heap $] \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h_{1}\right]}=\llbracket \varphi \rrbracket_{\mathcal{E}^{\prime}\left[\text { heap } \mapsto h_{1}\right]}=$ true . Since $y$ does not occur in $\varphi$, then we also have $\llbracket \varphi]_{\mathcal{E}\left[\text { heap } \leftrightarrow h, y \mapsto h_{2}\right]}=$ true.
Now, consider $h_{2} \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}^{\prime}\left[x \mapsto h_{1}\right]}$. To show that $\left|h_{2}\left(p k t_{\text {in }}\right)\right|+\left|h_{2}\left(p k t_{\text {out }}\right)\right| \leq N$. Since $\llbracket \varphi \rrbracket_{\mathcal{E}\left[\text { heap } \mapsto h, y \mapsto h_{2}\right]}=\operatorname{true}=\llbracket \varphi[x /$ heap $] \rrbracket_{\mathcal{E}\left[x \mapsto h, y \mapsto h_{2}\right]}$, we can conclude that $h_{2} \in \llbracket\left\{y: \tau_{2} \mid \varphi[x /\right.$ heap $\left.]\right\} \rrbracket_{\mathcal{E}^{\prime}[x \mapsto h]}$. Now, since $\mathcal{E}^{\prime}[x \mapsto h] \vDash \Gamma,\left(x: \tau_{1}\right)$, the result follows by our initial assumption assumption $\Gamma,\left(x: \tau_{1}\right) \vdash\left\{y: \tau_{2} \mid \varphi[x /\right.$ heap $\left.]\right\} \leq$ $N$.
$(\Leftarrow)$ Assume $\Gamma,\left(x:\left\{y: \tau_{1} \mid \varphi[y /\right.\right.$ heap $\left.\left.]\right\}\right) \vdash \tau_{2} \leq N$.
Let $\mathcal{E} \vDash \Gamma,\left(x: \tau_{1}\right)$. We can write $\mathcal{E}=\mathcal{E}^{\prime}\left[x \mapsto h_{1}\right]$ where $h_{1} \in \llbracket \tau_{1} \rrbracket \mathcal{E}^{\prime}$ and $\mathcal{E}^{\prime} \vDash \Gamma$.
Now consider $h_{2} \in \llbracket\left\{y: \tau_{2} \mid \varphi[x /\right.$ heap $\left.]\right\} \rrbracket_{\mathcal{E}^{\prime}\left[x \mapsto h_{1}\right]}$. To show $\left|h_{2}\left(p k t_{i n}\right)\right|+$ $\left|h_{2}\left(p k t_{\text {out }}\right)\right|<N$.
By the semantics of heap types, we have $\llbracket \varphi \rrbracket_{\mathcal{E}^{\prime}\left[x \mapsto h_{1}, y \mapsto h_{2}\right]}=\llbracket \varphi \rrbracket_{\mathcal{E}^{\prime}}\left[\right.$ heap $\left.\leftrightarrow h_{1}, y \mapsto h_{2}\right]=$ true. Since $y$ is not free in $\varphi$, we also have $\llbracket \varphi \rrbracket_{\mathcal{E}^{\prime}\left[\text { heap } \mapsto h_{1}\right]}=\llbracket \varphi[y /$ heap $] \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h_{1}\right]}=$ true, so we can conclude that $h_{1} \in \llbracket\left\{y: \tau_{1} \mid e[y /\right.$ heap $\left.]\right\} \rrbracket_{\mathcal{E}^{\prime}}$. By our initial assumption, every heap in $\tau_{2}$ is bounded and as such also heap $h_{2} \in \llbracket\left\{y: \tau_{2} \mid\right.$ $\varphi[x /$ heap $]\} \rrbracket_{\mathcal{E}^{\prime}\left[x \mapsto h_{1}\right]}$.

Lemma A. 76 (Bound Choice). If $\Gamma \vdash \tau_{1} \leq N$ and $\Gamma \vdash \tau_{2} \leq M$, then $\Gamma \vdash \tau_{1}+\tau_{2} \leq$ $\max (M, N)$.

Proof. Let $\mathcal{E} \vDash \Gamma$ and $h \in \llbracket \tau_{1}+\tau_{2} \rrbracket \mathcal{E}$. We have to show that $\left|h\left(p k t_{i n}\right)\right|+\left|h\left(p k t_{\text {out }}\right)\right| \leq$ $\max (M, N)$.

Case $M=N$ :
Assume $M=N$, so $\max (M, N)=M=N$. By the semantics of heap types $h \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}}$ or $h \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}}$.

Subcase $h \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}}$ : The result immediately follows by assumption $\Gamma \vdash \tau_{1} \leq N$.
Subcase $h \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}}$ : The result immediately follows by assumption $\Gamma \vdash \tau_{2} \leq M$.
Case $M>N$ : Without loss of generality, we assume that $M>N$, so $\max (M, N)=M$. By the semantics of heap types $h \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}}$ or $h \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}}$.

Subcase $h \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}}$ : By assumption $\Gamma \vdash \tau_{1} \leq N$ and since by assumption $N<M$, it follows $\Gamma \vdash \tau_{1} \leq M$.
Subcase $h \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}}$ : The result immediately follows by assumption $\Gamma \vdash \tau_{2} \leq M$.

Lemma A. 77 (Bound Substitution). For all $\Gamma, y, \tau_{1}, \tau_{2}, N, \Gamma \vdash \tau_{2}\left[y \mapsto \tau_{1}\right] \leq N$ if and only if $\Gamma,\left(y: \tau_{1}\right) \vdash \tau_{2} \leq N$.

Proof. Let $\Gamma, y, \tau_{1}, \tau_{2}$, and $N$ be given. Prove each direction separately:
$(\Rightarrow)$ Assume $\Gamma \vdash \tau_{2}\left[y \mapsto \tau_{2}\right] \leq N$. Let $\mathcal{E} \vDash \Gamma, y: \tau_{1}$ such that $h_{2} \in \llbracket \tau_{2} \rrbracket \mathcal{E}$. This means there is some $\mathcal{E}^{\prime} \vDash \Gamma$ and $h_{1} \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}^{\prime}}$ such that $h_{2} \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}^{\prime}\left[y \mapsto h_{1}\right]}$. By the semantics of heap types follows that $\left.h_{2} \in \llbracket \tau_{2}\left[y \mapsto \tau_{1}\right]\right]_{\mathcal{E}^{\prime}}$. With the initial assumption $\Gamma \vdash \tau_{2}\left[y \mapsto \tau_{1}\right] \leq N$, we can conclude that $\left|h_{2}\left(p k t_{\text {in }}\right)\right|+\left|h_{2}\left(p k t_{\text {out }}\right)\right| \leq N$.
$(\Leftarrow)$ Assume $\Gamma, y: \tau_{1} \vDash \tau_{2} \leq N$. Let $\mathcal{E} \vDash \Gamma$ and let $h_{2} \in \llbracket \tau_{2}\left[y \mapsto \tau_{1}\right] \rrbracket_{\mathcal{E}}$. By the semantics of heap types, there is some $h_{1} \in \llbracket \tau_{1} \rrbracket \mathcal{E}$ such that $h_{2} \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}\left[y \rightarrow h_{1}\right]}$. Notice that $\mathcal{E}\left[y \mapsto h_{1}\right] \vDash \Gamma, y: \tau_{1}$. The initial assumption proves that $\left|h\left(p k t_{i n}\right)\right|+$ $\left|h\left(p k t_{\text {out }}\right)\right| \leq N$.

Theorem A. 78 (Forwards MTU Bound). For every $\Gamma, c, x, \tau_{1}, \tau_{2}$, and $N \in \mathbb{N}$, if $\Gamma \vdash$ $\tau_{1} \leq N$ and $\Gamma \vdash c:\left(x: \tau_{1}\right) \leadsto \tau_{2}$ and every ascribed type in $c$ is also bounded by $N$, then $\Gamma, x: \tau_{1} \vdash \tau_{2} \leq N+\operatorname{emit}(c)$

Proof. Proceed by induction on $c$, leaving $\Gamma$ and $N$ general. We refer to the general assumptions as follows:
(A) $\Gamma \vdash \tau_{1} \leq N$
(B) $\Gamma \vdash c:\left(x: \tau_{1}\right) \leadsto \tau_{2}$

Case extract ( $\iota$ ):
The only algorithmic typing rule that applies to extract ( $\iota$ ) is T-Extract-Algo. By inversion, we know
(A1) $\tau_{2}=\Sigma y:\left\{z: \iota \mid \varphi_{1}\right\} \cdot\left\{z: \operatorname{chomp}\left(\tau_{1}, \iota, y\right) \mid \varphi_{2}\right\}$
(A2) $\varphi_{1}=z . p k t_{\text {in }}=z \cdot p k t_{\text {out }}=\langle \rangle$
(A3) $\varphi_{2}=y \cdot \iota @ z . p k t_{\text {in }}=x \cdot p k t_{\text {in }} \wedge z \cdot p k t_{\text {out }}=x \cdot p k t_{\text {out }} \wedge z \equiv, x$.
Since emit $(\operatorname{extract}(\iota))=0$, it suffices to show $\Gamma, x: \tau_{1} \vdash \tau_{2} \leq N$.
Let $\mathcal{E} \vDash \Gamma, x: \tau_{1}$ and let $h_{2} \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}}$. We can write $\mathcal{E}$ as $\mathcal{E}^{\prime}\left[x \mapsto h_{1}\right]$ where $h_{1} \in \llbracket \tau_{1} \rrbracket \mathbb{\mathcal { E }}^{\prime}$.
By definition of the semantics of heap types we know there are some $h_{21}$ and $h_{22}$ such that
(A4) $h_{2}=h_{21}++h_{22}$,
(A5) $h_{21}\left(p k t_{\text {in }}\right)=h_{21}\left(p k t_{\text {out }}\right)=\langle \rangle$
(A6) $h_{1}\left(p k t_{i n}\right)=h_{21}(\iota) @ h_{22}\left(p k t_{i n}\right)$
(A7) $h_{1}\left(p k t_{\text {out }}\right)=h_{22}\left(p k t_{\text {out }}\right)$
We can further conclude that
(A8) $h_{2}\left(p k t_{\text {out }}\right)=h_{1}\left(p k t_{\text {out }}\right)$,
(A9) $h_{2}\left(p k t_{i n}\right)=h_{22}\left(p k t_{i n}\right)=h_{1}\left(p k t_{\text {in }}\right)[|\iota|:]$
From assumption (A) follows
(A1o) $\left|h_{1}\left(p k t_{\text {in }}\right)\right|+\left|h_{1}\left(p k t_{\text {out }}\right)\right| \leq N$
Together with (A8) and (A9), we can conclude that $\left|h_{2}\left(p k t_{\text {in }}\right)\right|+\left|h_{2}\left(p k t_{\text {out }}\right)\right|<$ $\left|h_{1}\left(p k t_{\text {in }}\right)\right|+\left|h_{1}\left(p k t_{\text {out }}\right)\right| \leq N$.
Case add ( $\iota$ ):
The only algorithmic typing rule that applies to $\operatorname{add}(\iota)$ is T-Add-Algo. By inversion, we know
(A1) $\tau_{2}=\Sigma y:\left\{z: \tau_{1} \mid z \equiv x\right\}\left\{z: \iota \mid z \cdot p k t_{\text {in }}=z \cdot p k t_{\text {out }}=\langle \rangle \wedge z . \iota=v\right\}$.
Since emit $(\operatorname{add}(\iota))=0$, it suffices to show that $\Gamma, x: \tau_{1} \vdash \tau_{2} \leq N$.
Let $\mathcal{E} \vDash \Gamma, x: \tau_{1}$, and $h_{2} \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}}$. We can write $\mathcal{E}$ as $\mathcal{E}^{\prime}\left[x \mapsto h_{1}\right]$ where $h_{1} \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}^{\prime}}$. By definition of the semantics of heap types we know there are some $h_{21}$ and $h_{22}$ such that
(A2) $h_{2}=h_{21}++h_{22}$
(A3) $h_{21}=h_{1}$
(A4) $h_{22}\left(p k t_{\text {in }}\right)=h_{22}\left(p k t_{\text {out }}\right)=\langle \rangle$
From these three equations we can conclude that
(A5) $h_{2}\left(p k t_{i n}\right)=h_{1}\left(p k t_{i n}\right)$
(A6) $h_{2}\left(p k t_{\text {out }}\right)=h_{1}\left(p k t_{\text {out }}\right)$
The result follows by assumption (A).
Case remove( $\iota$ ):
The only algorithmic typing rule that applies to remove( $(1)$ is T-Remove-Algo. By inversion, we know
(A1) $\tau_{2}=\left\{y: \top \mid \varphi_{\imath} \wedge \varphi_{p k t} \wedge \neg y . \iota\right.$. valid $\}$
(A2) $\varphi_{p k t}=y \cdot p k t_{\text {in }}=x \cdot p k t_{\text {in }} \wedge y \cdot p k t_{\text {out }}=x \cdot p k t_{\text {out }}$
Since emit $(\operatorname{remove}(\iota))=0$, it suffices to show that $\Gamma, x: \tau_{1} \vdash \tau_{2} \leq N$.
Let $\mathcal{E} \vDash \Gamma, x: \tau_{1}$ and $h_{2} \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}}$. We can write $\mathcal{E}$ as $\mathcal{E}^{\prime}\left[x \mapsto h_{1}\right]$ where $h_{1} \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}^{\prime}}$. From assumption (A1) and by the semantics of heap types follows
(A3) $h_{2}\left(p k t_{i n}\right)=h_{1}\left(p k t_{i n}\right)$
(A4) $h_{2}\left(p k t_{\text {out }}\right)=h_{1}\left(p k t_{\text {out }}\right)$
The result follows by assumption (A).
Case $1 . f:=e$ :
The only algorithmic typing rule that applies to $\tau . f:=e$ is T-Mod-Algo. By inversion, we know
(A1) $\tau_{2}=\left\{y: T \mid \varphi_{p k t} \wedge \varphi_{1} \wedge \varphi_{f} \wedge y . l . f=e[x /\right.$ heap $\left.]\right\}$
(A2) $\varphi_{p k t}=y \cdot p k t_{\text {in }}=x \cdot p k t_{\text {in }} \wedge y \cdot p k t_{o u t}=x \cdot p k t_{\text {out }}$
Since emit $(\iota . f:=e)=0$, it suffices to show that $\Gamma, x: \tau_{1} \vdash \tau_{2} \leq N$.
Let $\mathcal{E} \vDash \Gamma, x: \tau_{1}$ and $h_{2} \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}}$. We can write $\mathcal{E}$ as $\mathcal{E}^{\prime}\left[x \mapsto h_{1}\right]$ where $h_{1} \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}^{\prime}}$. From assumption ( $\mathrm{A}_{1}$ ) and by the semantics of heap types follows
$\left(\mathrm{A}_{3}\right) h_{2}\left(p k t_{i n}\right)=h_{1}\left(p k t_{i n}\right)$
(A4) $h_{2}\left(p k t_{\text {out }}\right)=h_{1}\left(p k t_{\text {out }}\right)$
The result follows by assumption (A).
Case remit(ı):
The only algorithmic typing rule that applies to remit $(\iota)$ is T-Remit-Algo. By inversion, we know
(A1) $\tau_{2}=\Sigma y:\left\{z: \tau_{1} \mid z \equiv x\right\} .\{z: \epsilon \mid \varphi\}$
(A2) $\varphi=z \cdot p k t_{\text {in }}=\langle \rangle \wedge z . p k t_{\text {out }}=x . \iota$

Since emit $(\operatorname{remit}(\iota))=\operatorname{sizeof}(\iota)$, we have to show that $\Gamma, x: \tau_{1} \vdash \tau_{2} \leq N+$ sizeof $(\iota)$.
Let $\mathcal{E} \vDash \Gamma, x: \tau_{1}$ and $h_{2} \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}}$. We can write $\mathcal{E}$ as $\mathcal{E}^{\prime}\left[x \mapsto h_{1}\right]$ where $h_{1} \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}^{\prime}}$. By the semantics of heap types, there exists $h_{21}$ and $h_{22}$ such that
(A3) $h_{2}=h_{21}++h_{22}$
(A4) $h_{21}=h_{1}$
(A5) $h_{22}\left(p k t_{i n}\right)=\langle \rangle$
(A6) $h_{22}\left(p k t_{\text {out }}\right)=h_{1}(\iota)$
From (A4) and (A5), we can conclude that
(A7) $h_{2}\left(p k t_{i n}\right)=h_{1}\left(p k t_{i n}\right)$
From (A4) and (A6), we can further conclude that
(A8) $h_{2}\left(p k t_{\text {out }}\right)=h_{1}\left(p k t_{\text {out }}\right) @ h_{1}(\iota)$
From (A7) and (A8) then follows
(A9) $\left|h_{2}\left(p k t_{\text {in }}\right)\right|+\left|h_{2}\left(p k t_{\text {out }}\right)\right|=\left|h_{1}\left(p k t_{\text {in }}\right)\right|+\left|h_{1}\left(p k t_{\text {out }}\right)\right|+\operatorname{sizeof}(\iota)$
Together with assumption (A), we can conclude that $\left|h_{2}\left(p k t_{\text {in }}\right)\right|+\left|h_{2}\left(p k t_{\text {out }}\right)\right| \leq$ $N+\operatorname{sizeof}(\iota)$.

Case reset:
The only algorithmic typing rule that applies to reset is T-Reset-Algo. By inversion, we know
(A1) $\tau_{2}=\Sigma y:\left\{z: \epsilon \mid \varphi_{1}\right\} \cdot\left\{z: \epsilon \mid \varphi_{2}\right\}$
(A2) $\varphi_{1}=z \cdot p k t_{\text {out }}=\langle \rangle \wedge z \cdot p k t_{\text {in }}=x \cdot p k t_{\text {out }}$
(A3) $\varphi_{2}=z \cdot p k t_{\text {out }}=\langle \rangle \wedge z \cdot p k t_{\text {in }}=x \cdot p k t_{\text {in }}$
Since emit $($ reset $)=0$, we have to show that $\Gamma, x: \tau_{1} \vdash \tau_{2} \leq N$.
Let $\mathcal{E} \vDash \Gamma, x: \tau_{1}$ and $h_{2} \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}}$. We can write $\mathcal{E}$ as $\mathcal{E}^{\prime}\left[x \mapsto h_{1}\right]$ where $h_{1} \in \llbracket \tau_{1} \rrbracket_{\mathcal{E}}$. By the semantics of heap types, there exists $h_{21}$ and $h_{22}$ such that
(A4) $h_{2}=h_{21}++h_{22}$
(A5) $h_{21}\left(p k t_{\text {out }}\right)=\langle \rangle$
(A6) $h_{21}\left(p k t_{\text {in }}\right)=h_{1}\left(p k t_{\text {out }}\right)$
(A7) $h_{22}\left(p k t_{\text {out }}\right)=\langle \rangle$
(A8) $h_{22}\left(p k t_{i n}\right)=h_{1}\left(p k t_{i n}\right)$
By (A4), (A5) and (A7) follows
$(\mathrm{A} 9) h_{2}\left(p k t_{\text {out }}\right)=\langle \rangle$
and by (A4), (A6) and (A8) follows
$\left(\mathrm{A}_{10}\right) h_{2}\left(p k t_{\text {in }}\right)=h_{1}\left(p k t_{\text {out }}\right) @ h_{1}\left(p k t_{\text {in }}\right)$

Since by assumption (A), $\left|h_{1}\left(p k t_{\text {in }}\right)\right|+\left|h_{1}\left(p k t_{\text {out }}\right)\right| \leq N$, by (A9) and (A1o), $\left|h_{2}\left(p k t_{\text {in }}\right)\right|+\left|h_{2}\left(p k t_{\text {out }}\right)\right| \leq N$.

Case if $(\varphi) c_{1}$ else $c_{2}$ : The only algorithmic typing rule that applies is T-If-Algo. By inversion, we know
(A1) $\tau_{2}=\left\{y: \tau_{12} \mid \varphi[x /\right.$ heap $\left.]\right\}+\left\{y: \tau_{22} \mid \neg \varphi[x /\right.$ heap $\left.]\right\}$
(A2) $\Gamma \vdash c_{1}:\left(x:\left\{y: \tau_{1} \mid \varphi[y /\right.\right.$ heap $\left.\left.]\right\}\right) \leadsto \tau_{12}$
(A3) $\Gamma \vdash c_{1}:\left(x:\left\{y: \tau_{1} \mid \neg \varphi[y /\right.\right.$ heap $\left.\left.]\right\}\right) \leadsto \tau_{22}$
Since emit $\left(i f(\varphi) c_{1}\right.$ else $\left.c_{2}\right)=\max \left(\operatorname{emit}\left(c_{1}\right)\right.$, emit $\left.\left(c_{2}\right)\right)$, we have to show that $\Gamma, x: \tau_{1} \vdash \tau_{2} \leq N+\max \left(\operatorname{emit}\left(c_{1}\right)\right.$, emit $\left.\left(c_{2}\right)\right)$.
By Lemma A. 74 and assumption (A) follows
(A4) $\Gamma \vdash\left\{y: \tau_{1} \mid \varphi[y /\right.$ heap $\left.]\right\} \leq N$
(A5) $\Gamma \vdash\left\{y: \tau_{1} \mid \neg \varphi[y /\right.$ heap $\left.]\right\} \leq N$
Applying the IH to $\left(\mathrm{A}_{2}\right)$ and $\left(\mathrm{A}_{2}\right)$ with $\left(\mathrm{A}_{4}\right)$ and $\left(\mathrm{A}_{5}\right)$ respectively, gives
(A6) $\Gamma, x:\left\{y: \tau_{1} \mid \varphi[y /\right.$ heap $\left.]\right\} \vdash \tau_{12} \leq N+\operatorname{emit}\left(c_{1}\right)$
(A7) $\Gamma, x:\left\{y: \tau_{1} \mid \neg \varphi[y /\right.$ heap $\left.]\right\} \vdash \tau_{22} \leq N+\operatorname{emit}\left(c_{2}\right)$
By Lemma A. 75 with (A6) and (A7) respectively follows
(A8) $\Gamma, x: \tau_{1} \vdash\left\{y: \tau_{12} \mid \varphi[x /\right.$ heap $\left.]\right\} \leq N+\operatorname{emit}\left(c_{1}\right)$
(A9) $\Gamma, x: \tau_{1} \vdash\left\{y: \tau_{22} \mid \neg \varphi[x /\right.$ heap $\left.]\right\} \leq N+\operatorname{emit}\left(c_{2}\right)$
By Lemma A. 76 with (A8) and (A9) follows
(A1o) $\Gamma, x: \tau_{1} \vdash\left\{y: \tau_{12} \mid \varphi[x /\right.$ heap $\left.]\right\}+\left\{y: \tau_{22} \mid \neg \varphi[x /\right.$ heap $\left.]\right\} \leq \max (N+$ $\left.\operatorname{emit}\left(c_{1}\right), N+\operatorname{emit}\left(c_{2}\right)\right)$

The result follows together with the fact that $\max (A+B, A+C)=A+\max (B, C)$.

## Case $c_{1} ; c_{2}$ :

The only algorithmic typing rule that applies to $c_{1} ; c_{2}$ is T-SEQ-Algo. By inversion, we know
(A1) $\Gamma \vdash c_{1}:\left(x: \tau_{1}\right) \leadsto \tau_{12}$
(A2) $\Gamma, x: \tau_{1} \vdash c_{2}:\left(y: \tau_{12}\right) \leadsto \tau_{22}$
(A3) $\tau_{2}=\tau_{22}\left[y \mapsto \tau_{12}\right]$
Since emit $\left(c_{1} ; c_{2}\right)=\operatorname{emit}\left(c_{1}\right)+\operatorname{emit}\left(c_{2}\right)$, we have to show that $\Gamma, x: \tau_{1} \vdash \tau_{2} \leq$ $N+\operatorname{emit}\left(c_{1}\right)+\operatorname{emit}\left(c_{2}\right)$. Let $\mathcal{E} \vDash \Gamma, x: \tau_{1}$ and let $h_{2} \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}}$.
By applying the IH to ( $\mathrm{A}_{1}$ ), we get
(A4) $\Gamma, x: \tau_{1} \vdash \tau_{12} \leq N+\operatorname{emit}\left(c_{1}\right)$
Since we left $\Gamma$ and $N$ general, we can apply the IH again to ( $\mathrm{A}_{2}$ ) and get
(A5) $\Gamma, x: \tau_{1}, y: \tau_{12} \vdash \tau_{22} \leq N+\operatorname{emit}\left(c_{1}\right)+\operatorname{emit}\left(c_{2}\right)$
The result follows by Lemma A. 77 with (A5).

## Case skip:

The only algorithmic typing rule that applies to skip is T-Skip-Algo. By inversion, we know
(A1) $\tau_{2}=\left\{y: \tau_{1} \mid y \equiv x\right\}$
Since emit $($ skip $)=0$, we have to show that $\Gamma, x: \tau_{1} \vdash \tau_{2} \leq N$.
Let $\mathcal{E} \vDash \Gamma, x: \tau_{1}$ and $h_{2} \in \llbracket \tau_{2} \rrbracket_{\mathcal{E}}$. We can write $\mathcal{E}$ as $\mathcal{E}^{\prime}\left[x \mapsto h_{1}\right]$ where $h_{1} \in \llbracket \tau_{1} \rrbracket \mathcal{E}^{\prime}$. By the semantics of heap types, follows
(A2) $h_{2}\left(p k t_{i n}\right)=h_{1}\left(p k t_{i n}\right)$
(A3) $h_{2}\left(p k t_{\text {out }}\right)=h_{1}\left(p k t_{\text {out }}\right)$
The result follows by assumption (A).
Case $c_{0}$ as $\left(x: \hat{\tau}_{1}\right) \rightarrow \tau_{2}$ :
The only algorithmic typing rule that applies is T-Ascribe-Algo. By inversion, we know
(A1) $\Gamma \vdash c_{0}:\left(x: \hat{\tau}_{1}\right) \leadsto \tau_{c}$
(A2) $\Gamma \vdash \tau_{1}<: \hat{\tau}_{1}$
(А3) $\Gamma, x: \hat{\tau}_{1} \vdash \tau_{c}<: \tau_{2}$
Since emit $\left(c_{0}\right.$ as $\left.\sigma\right)=\operatorname{emit}\left(c_{\mathrm{o}}\right)$, we have to show that $\Gamma, x: \tau_{1} \vdash \tau_{2} \leq N+$ emit $\left(c_{0}\right)$. By our initial assumption, every ascribed type is also bounded by $N$. We therefore have $\Gamma, x: \tau_{1} \vdash \tau_{2} \leq N$ from which the result immediately follows.

Theorem A. 79 (Decidability). If $\Gamma, \tau_{1}, \tau_{2}$ and every ascribed type in $c$ are bounded by the MTU N, then $\Gamma \vdash c:\left(x: \tau_{1}\right) \rightarrow \tau_{2}$ is decidable.

Proof. By Theorem A. 73 (Algorithmic Typing Correctness), we can equivalently show that $\Gamma \vdash c:\left(x: \tau_{1}\right) \leadsto \tau_{2}^{\prime}$ and $\Gamma, x: \tau_{1} \vdash \tau_{2}^{\prime}<: \tau_{2}$ are decidable. By Theorem A. $78, \tau_{2}^{\prime}$ is bounded. $\Gamma, x: \tau_{1} \vdash \tau_{2}^{\prime}<: \tau_{2}$ is therefore decidable by finite enumeration.

To show that $\Gamma \vdash c:\left(x: \tau_{1}\right) \leadsto \tau_{2}^{\prime}$ is decidable, we proceed by induction on the algorithmic typing derivation.

## Case T-Skip-Algo:

Immediate, because T-Skip-Algo does not perform any subtyping checks.

## Case T-Reset-Algo:

Also immediate, because T-Reset-Algo does not perform any subtyping checks.

## Case T-Seq-Algo:

By inversion of T-Seq-Algo,
(A1) $\Gamma \vdash c_{1}:\left(x: \tau_{1}\right) \leadsto \tau_{12}$
(A2) $\Gamma, x: \tau_{1} \vdash c_{2}:\left(y: \tau_{12}\right) \leadsto \tau_{22}$
By ( $\mathrm{A}_{1}$ ) and Theorem A.78,
(A3) $\Gamma, x: \tau_{1} \vdash \tau_{12} \leq N+\operatorname{emit}\left(c_{1}\right)$
Applying the IH to ( $\mathrm{A}_{1}$ ) with assumption $\Gamma \vdash \tau_{1} \leq N$ and ( $\mathrm{A}_{3}$ ) gives us that $\Gamma \vdash c_{1}:\left(x: \tau_{1}\right) \leadsto \tau_{12}$ is decidable.
Again, by Theorem A. 78 with ( $\mathrm{A}_{2}$ ) and ( $\mathrm{A}_{3}$ ), follows
(A4) $\Gamma, x: \tau_{1}, y: \tau_{12} \vdash \tau_{22} \leq N+\operatorname{emit}\left(c_{1}\right)+\operatorname{emit}\left(c_{2}\right)$
By IH follows that $\Gamma, x: \tau_{1} \vdash c_{2}:\left(y: \tau_{12}\right) \leadsto \tau_{22}$ is decidable and thus type checking the sequence of both commands is decidable.

## Case T-Add-Algo:

By inversion, we know that T-Add-Algo performs the subtyping check $\Gamma \vdash \tau_{1}<$ : $\{x: \top \mid \neg x . . ı$ valid $\}$. To show that type checking is decidable in this case, we must show that $\Gamma \vdash \tau_{1}<:\{x: \top \mid \neg x . \iota . v a l i d\}$ is decidable. This is the case because we can finitely enumerate the heaps $h$ described by $\tau_{1}$ and check wether every $h$ is a member of $\{x: \top \mid \neg y . ı$ valid $\}$.

## Case T-Extract-Algo:

By inversion, we know that T-Extract-Algo performs the subtyping check $\Gamma \vdash \tau_{1}<:\left\{x: \top| | x . p k t_{i n} \mid \geq \operatorname{sizeof}(\iota)\right\}$. To show that type checking is decidable in this case, we must show that $\Gamma \vdash \tau_{1}<:\left\{x: \top| | x . p k t_{i n} \mid \geq \operatorname{sizeof}(\iota)\right\}$ is decidable. This is the case because we can finitely enumerate the heaps $h$ described by $\tau_{1}$ and check wether every $h$ is a member of $\left\{x: T| | x . p k t_{i n} \mid \geq \operatorname{sizeof}(\iota)\right\}$.

## Case T-Remove-Algo:

By inversion, we know that T-Remove-Algo performs the subtyping check $\Gamma \vdash \tau_{1}<:\{x: \mathrm{T} \mid$ x.ı.valid $\}$. To show that type checking is decidable in this case, we must show that $\Gamma \vdash \tau_{1}<:\{x: \top \mid x . . \iota . v a l i d\}$ is decidable. This is the case because we can finitely enumerate the heaps $h$ described by $\tau_{1}$ and check wether every $h$ is a member of $\{x: \top \mid x$ x.ı.valid $\}$.

## Case T-Remit-Algo:

Identical to the previous subcase.

## Case T-Mod-Algo:

Identical to the previous subcase.

## Case T-If-Algo:

Since $\tau_{1}$ is bounded by assumption and refining the input type does not increase the size, $\left\{y: \tau_{1} \mid \varphi[y /\right.$ heap $\left.]\right\}$ and $\left\{y: \tau_{1} \mid \neg \varphi[y /\right.$ heap $\left.]\right\}$ are still bounded. By Theorem A. 78 then follows that the output types of $c_{1}$ and $c_{2}$ are also bounded. By IH applied to $c_{1}$ and $c_{2}$, we get that the algorithmic type checking applied to $c_{1}$ and $c_{2}$ respectively is decidable and thus checking the conditional is decidable.

## Case T-Ascribe-Algo:

By assumption, $\Gamma, \tau_{1}$ and $\hat{\tau}_{1}$ are bounded, so $\Gamma \vdash \tau_{1}<: \hat{\tau}_{1}$ is decidable by finite enumeration. Since by assumption $\Gamma \vdash \tau_{1} \leq N$, by Theorem A. 78 follows that $\Gamma, x: \hat{\tau}_{1} \vdash \tau_{c} \leq N+\operatorname{emit}\left(c_{o}\right)$. By IH then follows that $\Gamma \vdash c:\left(x: \tau_{1}\right) \leadsto \tau_{2}^{\prime}$ is decidable. Since $\tau_{c}$ is bounded and by assumption also $\tau_{2}^{\prime}$ is bounded, we can finitely enumerate, so $\Gamma, x: \hat{\tau}_{1} \vdash \tau_{c}<: \tau_{2}^{\prime}$ is also decidable and thus type checking an ascribed command is decidable.

## A.2. 4 Type Equivalences

Lemma A.8o (Rewrite Sigma Extract). For all $\Gamma, x, \tau$ and $\iota$, if $\Gamma \vdash \operatorname{sizeof}_{p k t_{i n}}(\tau) \geq$ sizeof $(\iota)$ and $x$ does not occur free in $\tau$, then

$$
\begin{aligned}
& \Gamma, x: \tau \vdash \\
& \Sigma y:\left\{\begin{array}{l|l}
z: \iota & \begin{array}{l}
z . p k t_{\text {in }}=\langle \rangle \wedge \\
z \cdot p k t_{\text {out }}=\langle \rangle
\end{array}
\end{array}\right\} \cdot\left\{\begin{array}{ll}
z: \operatorname{chomp}(\tau, \iota, y) & \begin{array}{l}
y . \iota @ z \cdot p k t_{\text {in }}=x \cdot p k t_{\text {in }} \wedge \\
z \cdot p k t_{\text {out }}=x \cdot p k t_{\text {out }} \wedge \\
z \equiv, x
\end{array}
\end{array}\right\}
\end{aligned}
$$

Proof. Proof each direction separately.
$(\Rightarrow)$ Let $\mathcal{E} \vDash \Gamma, x: \tau$. We know $\mathcal{E}=\mathcal{E}^{\prime}[x \mapsto h]$ such that $h \in \llbracket \tau \rrbracket_{\mathcal{E}^{\prime}}$. Let $h_{\Sigma} \in$ $\llbracket \Sigma y:\left\{z: \iota \mid z . p k t_{\text {in }}=z . p k t_{\text {out }}=\langle \rangle\right\} .\left\{z: \operatorname{chomp}(\tau, \iota, y) \mid y . \iota @ z . p k t_{i n}=\right.$ $\left.x . p k t_{\text {in }} \wedge z . p k t_{\text {out }}=x . p k t_{\text {out }} \wedge z \equiv, x\right\} \rrbracket_{\mathcal{E}^{\prime}[x \mapsto h]}$ be arbitrary. By the semantics of heap types follows
(A1) $h_{\Sigma}=h_{1}++h_{2}$
(A2) $h_{1} \in \llbracket\left\{z: \iota \mid z . p k t_{\text {in }}=z . p k t_{\text {out }}=\langle \rangle\right\} \rrbracket_{\mathcal{E}^{\prime}[x \mapsto h]}$
(A3) $h_{2} \in \llbracket\left\{z: \operatorname{chomp}(\tau, \iota, y) \mid y . \iota @ z . p k t_{\text {in }}=x . p k t_{\text {in }} \wedge z . p k t_{\text {out }}=x . p k t_{\text {out }} \wedge\right.$ $z \equiv x\} \rrbracket_{\mathcal{E}^{\prime}\left[x \mapsto h, y \mapsto h_{1}\right]}$
(A4) $h_{2} \in \llbracket \operatorname{chomp}(\tau, \iota, y) \rrbracket_{\mathcal{E}^{\prime}\left[x \mapsto h, y \mapsto h_{1}\right]}$
(A5) 【y.ı@z.pkt $t_{\text {in }}=x \cdot p k t_{\text {in }} \wedge z \cdot p k t_{\text {out }}=x \cdot p k t_{\text {out }} \wedge z \equiv, x \rrbracket_{\mathcal{E}^{\prime}\left[x \mapsto h, y \mapsto h_{1}, z \mapsto h_{2}\right]}=$ true
By Lemma A.47, there exists $\hat{h}_{2} \in \llbracket \tau \rrbracket_{\mathcal{E}^{\prime}[x \mapsto h]}$ such that
(A6) $h_{2}=$ chomp $^{\Downarrow}\left(\hat{h}_{2}\right.$, sizeof $\left.(t)\right)$
Together with (A5), we can conclude that
(A7) $\hat{h}_{2}\left(p k t_{\text {out }}\right)=h\left(p k t_{\text {out }}\right)$
(A8) $\hat{h}_{2}\left(p k t_{i n}\right)=h\left(p k t_{i n}\right)[\operatorname{sizeof}(\iota):]$
(A9) $\forall \kappa \neq \iota . \hat{h}_{2}(\kappa)=h(\kappa)$
(A10) $\iota \notin \operatorname{dom}\left(\hat{h}_{2}\right)$
$h_{\Sigma} \in \llbracket\left\{y: \top \mid y . \iota . v a l i d \wedge \wedge_{\kappa \in \operatorname{dom}(\mathcal{H} \mathcal{T}) \wedge \kappa \neq \iota} y . \kappa=x . \kappa \wedge y . \iota @ y . p k t_{i n}=x . p k t_{i n} \wedge\right.$ $\left.y . p k t_{\text {out }}=x . p k t_{\text {out }}\right\} \rrbracket_{\mathcal{E}}$ follows by the semantics of heap types with (A1), (A6), (A7), (A8), (A9) and (A10).
$(\Leftarrow)$ Let $\mathcal{E} \vDash \Gamma, x: \tau$. We know $\mathcal{E}=\mathcal{E}^{\prime}[x \mapsto h]$ such that $h \in \llbracket \tau \rrbracket \mathcal{E}^{\prime}$. Let $\hat{h} \in \llbracket\{y$ : $\top \mid y . \iota . v a l i d \wedge \wedge_{\kappa \in \operatorname{dom}(\mathcal{H T}) \wedge \kappa \neq \iota} y . \kappa=x . \kappa \wedge y . \iota @ y . p k t_{i n}=x . p k t_{i n} \wedge y . p k t_{\text {out }}=$ $\left.x . p k t_{\text {out }}\right\} \rrbracket_{\mathcal{E}^{\prime}[x \mapsto h]}$
By the semantics of heap types,
$\left(\mathrm{A}_{1}\right) \hat{h}(\iota)=h\left(p k t_{\text {in }}\right)[\mathrm{o}: \operatorname{sizeof}(\iota)]$
(A2) $\forall \kappa \neq \iota . \hat{h}(\kappa)=h(\kappa)$
(A3) $\hat{h}\left(p k t_{o u t}\right)=h\left(p k t_{\text {out }}\right)$
(A4) $\hat{h}(\iota) @ \hat{h}\left(p k t_{\text {in }}\right)=h\left(p k t_{i n}\right) \Leftrightarrow \hat{h}\left(p k t_{i n}\right)=h\left(p k t_{\text {in }}\right)[\operatorname{sizeof}(\iota):]$
To show that $\hat{h} \in \llbracket \Sigma y:\left\{z: \iota \mid z \cdot p k t_{\text {in }}=z \cdot p k t_{\text {out }}=\langle \rangle\right\} .\{z: \operatorname{chomp}(\tau, \iota, y) \mid$ y. $\left.@ z . p k t_{\text {in }}=x . p k t_{\text {in }} \wedge z . p k t_{\text {out }}=x . p k t_{\text {out }} \wedge z \equiv_{\iota} x\right\} \rrbracket_{\mathcal{E}^{\prime}[x \mapsto h]}$, we have to show that there exists $h_{1}$ and $h_{2}$ such that $\hat{h}=h_{1}++h_{2}$ and $h_{1} \in \mathbb{\{ z : ~} \mid$ $\left.z . p k t_{i_{n}}=z . p k t_{\text {out }}=\langle \rangle\right\} \rrbracket_{\mathcal{E}^{\prime}[x \mapsto h]}$ and $h_{2} \in \llbracket\left\{z: \operatorname{chomp}(\tau, \iota, y) \mid y . \iota @ z . p k t_{\text {in }}=\right.$ $\left.x . p k t_{\text {in }} \wedge z \cdot p k t_{\text {out }}=x \cdot p k t_{\text {out }} \wedge z \equiv, x\right\} \rrbracket_{\mathcal{E}^{\prime}\left[x \mapsto h, y \mapsto h_{1}\right]}$.
Let $h_{1}\left(p k t_{\text {in }}\right)=h_{1}\left(p k t_{\text {out }}\right)=\langle \rangle$ and $h_{1}(\iota)=h\left(p k t_{\text {in }}\right)[0: \operatorname{sizeof}(\iota)]$ and no other instances be valid in heap $h_{1} . h_{1} \in \llbracket\left\{z: \iota \mid z . p k t_{\text {in }}=z . p k t_{\text {out }}=\langle \rangle\right\} \rrbracket_{\mathcal{E}^{\prime}[x \leftrightarrow h]}$ then follows by the semantics of heap types. By Lemma A.44, there exists $h_{2} \epsilon$ $\llbracket \operatorname{chomp}(\tau, \iota, y) \rrbracket_{\mathcal{E}^{\prime}\left[y \rightarrow h_{1}\right]}$ such that $h_{2}=\operatorname{chomp}^{\Downarrow}(h$, sizeof $(t))$. Since $x$ not free in $\tau$, it also holds that $h_{2} \in \llbracket \operatorname{chomp}(\tau, l, y) \rrbracket_{\mathcal{E}^{\prime}\left[x \mapsto h, y \mapsto h_{1}\right]}$.
Since $h_{1}(\iota) @ h_{2}\left(p k t_{\text {in }}\right)=h\left(p k t_{\text {in }}\right), h_{2}\left(p k t_{\text {out }}\right)=h\left(p k t_{\text {out }}\right)$ and since chomp does not change already valid header instances also for all $\kappa \neq \iota, h_{2}(\kappa)=h(\kappa)$, we can conclude that $\llbracket y . \iota @ z . p k t_{\text {in }}=x . p k t_{\text {in }} \wedge z . p k t_{\text {out }}=x . p k t_{\text {out }} \wedge z \equiv_{\iota}$ $x \rrbracket_{\mathcal{E}^{\prime}\left[x \leftrightarrow h, y \mapsto h_{1}, z \mapsto h_{2}\right]}=$ true and thus $h_{2} \in \llbracket\left\{z: \operatorname{chomp}(\tau, \iota, y) \mid y . l @ z . p k t_{\text {in }}=\right.$ $\left.x . p k t_{\text {in }} \wedge z . p k t_{\text {out }}=x \cdot p k t_{\text {out }} \wedge z \equiv \iota x\right\} \rrbracket_{\mathcal{E}^{\prime}\left[x \leftrightarrow h, y \mapsto h_{1}\right]}$.
By the semantics of heap types, we can further conclude that
(A5) $\left(h_{1}++h_{2}\right)\left(p k t_{i n}\right)=h_{2}\left(p k t_{\text {in }}\right)=h^{\prime}\left(p k t_{\text {in }}\right)$
(A6) $\left(h_{1}++h_{2}\right)\left(p k t_{\text {out }}\right)=h\left(p k t_{\text {out }}\right)=h^{\prime}\left(p k t_{\text {out }}\right)$
(A7) $\left(h_{1}++h_{2}\right)(\iota)=h\left(p k t_{\text {in }}\right)[\mathrm{o}: \operatorname{sizeof}(\iota)]=h^{\prime}(\iota)$
(A8) $\forall \kappa \neq \iota .\left(h_{1}++h_{2}\right)(\kappa)=h_{2}(\kappa)=h(\kappa)=h^{\prime}(\kappa)$
This shows that actually $h^{\prime}=h_{1}++h_{2}$ and concludes this case.


[^0]:    ${ }^{1}$ There are other actions in the real implementation, but these three actions demonstrate the core of the problem.

[^1]:    ${ }^{1}$ In this work we focus on exact and ternary matches as well as matches on the validity bit.

[^2]:    ${ }^{1} \mathrm{P}_{4}$ allows loops within parsers, but because programs are restricted to finite state, the language specification allows implementations to unroll loops.

[^3]:    ${ }^{1}$ This is a design decision we made; instead we could have used the binder of the $\Sigma$-type for the left projection and only introduce a fresh binder for the right projection.

[^4]:    ${ }^{2}$ We consider the optimized algorithmic typing rules presented in Figure 6.10.

[^5]:    ${ }^{3}$ Assuming an MTU of 1500 bytes

[^6]:    ${ }^{1}$ Strictly speaking, IPv 4 requires a special ICMP message to be returned to the sender to indicate the error, but here we will simply drop the packet.

[^7]:    ${ }^{1}$ https://opennetworking.org/ng-sdn/
    ${ }^{2}$ https://opennetworking.org/onos/

[^8]:    ${ }^{3}$ As explained in Section 2.2.2, this is the default behavior in $\mathrm{P}_{4}$

