

Securing P4 Programs by Information Flow Control

Anoud Alshnakat*, Amir M. Ahmadian*, Musard Balliu*, Roberto Guanciale* and Mads Dam*

*KTH Royal Institute of Technology

Abstract—Software-Defined Networking (SDN) has transformed network architectures by decoupling the control and data-planes, enabling fine-grained control over packet processing and forwarding. P4, a language designed for programming data-plane devices, allows developers to define custom packet processing behaviors directly on programmable network devices. This provides greater control over packet forwarding, inspection, and modification. However, the increased flexibility provided by P4 also brings significant security challenges, particularly in managing sensitive data and preventing information leakage within the data-plane.

This paper presents a novel security type system for analyzing information flow in P4 programs that combines security types with interval analysis. The proposed type system allows the specification of security policies in terms of input and output packet bit fields rather than program variables. We formalize this type system and prove it sound, guaranteeing that well-typed programs satisfy noninterference. Our prototype implementation, TAP4S, is evaluated on several use cases, demonstrating its effectiveness in detecting security violations and information leakages.

I. INTRODUCTION

Software-Defined Networking (SDN) [1] is a software-driven approach to networking that enables programmatic control of network configuration and packet processing rules. SDN achieves this by decoupling the routing process, performed in the control-plane, from the forwarding process performed in the data-plane. The control-plane is often implemented by a logically-centralized SDN controller that is responsible for network configuration and the setting of forwarding rules. The data-plane consists of network devices, such as programmable switches, that process and forward packets based on instructions received from the control-plane. Before SDN, hardware providers had complete control over the supported functionalities of the devices, leading to lengthy development cycles and delays in deploying new features. SDN has shifted this paradigm, allowing application developers and network engineers to implement specific network behaviors, such as deep packet inspection, load balancing, and VPNs, and execute them directly on networking devices.

Network Functions Virtualization (NFV) further expands upon this concept, enabling the deployment of multiple virtual data-planes over a single physical infrastructure [2]. SDN and NFV together offer increased agility and optimization, making them cornerstones of future network architectures. Complementing this evolution, the Programming Protocol-independent Packet Processors (P4) [3] domain-specific language has emerged as a leading standard for programming the data-plane’s programmable devices, such as FPGAs and switches. Additionally, P4 serves as a specification language

to define the behavior of the switches as it provides a suitable level of abstraction, yet is detailed enough to accurately capture the behavior of the switch. It maintains a level of simplicity and formalism that allows for effective automated analysis [4].

NFVs and SDNs introduce new security challenges that extend beyond the famous and costly outages caused by network misconfigurations [5]. Many data-plane applications process sensitive data, such as cryptographic keys and internal network topologies. The complexity of these applications, the separation of ownership of platform and data-plane in virtualized environments, and the integration of third-party code facilitate undetected information leakages. Misconfiguration may deliver unencrypted packets to a public network, bugs may leak sensitive packet metadata or routing configurations that expose internal network topology, and malicious code may build covert channels to exfiltrate data via legitimate packet fields such as TCP sequence numbers and TTL fields [6].

In this domain, the core challenge lies in the data dependency of what is observable, what is secret, and the packet forwarding behavior. An attacker may be able to access only packets belonging to a specific subnetwork, only packets for a specific network protocol may be secret, and switches may drop packets based on the matching of their fields with routing configurations. These data dependencies make information leakage a complex problem to address in SDN-driven networks.

Existing work in the area of SDN has focused on security of routing configurations by analyzing network flows that are characterized by port numbers and endpoints. However, these works ignore indirect flows that may leak information via other packet fields. In the programming languages area, current approaches (including P4BID [7]) substantially ignore data dependencies and lead to overapproximations unsuitable for SDN applications. For example, the sensitivity of a field in a packet might depend on the packet’s destination.

We develop a new approach to analyze information flow in P4 programs. A key idea is to augment a security type system (which is a language-based approach to check how information can flow in a program) with interval analysis, which in the domain of SDNs can be used to abstract over the network’s parameters such as subnetwork segments, port ranges, and non-expired TTLs. Therefore, in our approach, in addition to a security label, the security type also keeps track of an interval.

The analysis begins with an input policy, expressed as an assignment of types to fields of the input packet. For instance, a packet might be considered sensitive only if its source IP

belongs to the internal network. The analysis conservatively propagates labels and intervals throughout the P4 program in a manner reminiscent of dynamic information flow control [8] and symbolic execution, cf. [9]. This process is not dependent on a prior assignment of security labels to internal program variables, thus eliminating the need for the network engineer to engage with P4 program internals. The proposed analysis produces multiple final output packet typings, corresponding to different execution paths. These types are statically compared with the output security policy, which allows to relate observability of the output to intervals of fields of the resulting packets and their metadata.

The integration of security types and intervals is challenging. On one hand, the analysis should be path-sensitive and be driven by values in the packet fields to avoid rejecting secure programs due to overapproximation. On the other hand the analysis must be sound and not miss indirect information flows. Another challenge is that the behaviors of P4 programs depend on tables and external functions, but these components are not defined in P4. We address this by using user-defined contracts that overapproximate their behavior.

Summary of contributions.

- We propose a security type system which combines security labels and abstract domains to provide noninterference guarantees on P4 programs.
- Our approach allows defining data-dependent policies without the burden of annotating P4 programs.
- We implement the proposed type system in a prototype tool TAP4S [10] and evaluate the tool on a test suite and 5 use cases.

II. P4 LANGUAGE AND SECURITY CHALLENGES

This section provides a brief introduction to the P4 language and its key features, while motivating the need for novel security analysis that strikes a balance between expressiveness of security policies and automation of the verification process.

P4 manipulates and forwards packets via a pipeline consisting of three stages: parser, match-action, and deparser. The parser stage dissects incoming packets, converting the byte stream into structured header formats. In the match-action stage, these headers are matched against rules to determine the appropriate actions, such as modifying, dropping, or forwarding the packet to specific ports. Finally, the deparser stage reconstructs the processed packet back into a byte stream, ready for transmission over the network.

We use Program 1 as a running example throughout the paper. The program implements a switch that manages congestion in the network of Fig. 1. In an IPv4 packet, the Explicit Congestion Notification (ECN) field provides the status of congestion experienced by switches while transmitting the packet along the path from source to destination. ECN value 0 indicates that at least one of the traversed switches does not support the ECN capability, values 1 and 2 indicate that all traversed switches support ECN and the packet can be marked if congestion occurs, and 3 indicates that the packet has experienced congestion in at least one of the switches.

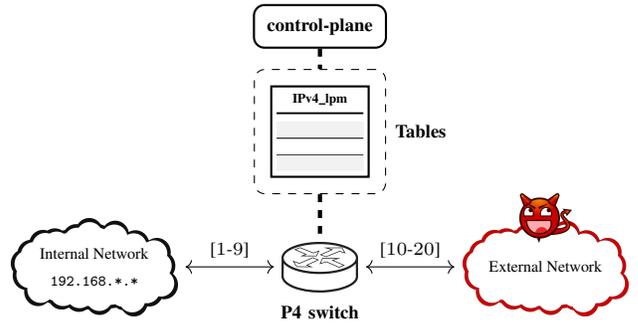


Fig. 1: Congestion notifier network layout

We assume for this example that the switch is the *only* ingress and egress point for traffic entering and exiting the internal network, connecting it to external networks as shown in Fig. 1. To illustrate our approach, we assume that the switch is designed to prevent any information leakage about internal network congestion to the external network. In addition to standard packet forwarding, the switch sets ECN to 3 if its queue length exceeds a predefined threshold. This holds only if the packet’s destination is within the internal network. Conversely, if the packet is destined to the external network, the switch sets the ECN field to 0. This indicates that ECN is not supported for outbound traffic and ensures that congestion signals experienced within the internal network are not exposed externally.¹

P4 structs and headers. Structs are records used to define the format of P4 packets. Headers are special structs with an additional implicit boolean indicating the header’s validity, which is set when the header is extracted. Special function `isValid` (line 53) is used to check the validity of a header.

For example, the struct `headers` on line 3 has two headers of type `ethernet_t` and `ipv4_t`, as depicted in Fig. 2. The fields of the `ethernet_t` specify the source and destination MAC addresses and the Ethernet type. The header `ipv4_t` represents a standard IPv4 header with fields such as ECN, time-to-live (TTL), and source and destination IP addresses.

Parser. The parser dissects incoming raw packets (packet on line 12), extracts the raw bits, and groups them into headers. The parser’s execution begins with the `start` state and terminates either in `reject` state or `accept` state accepting the packet and moving to the next stage of the pipeline.

For example, `MyParser` consists of three states. The parsing begins at the `start` state (line 17) and transitions to `parse_ethernet` extracting the Ethernet header from the input packet (line 22), which automatically sets the header’s validity boolean to `true`. Next, depending on the value of `hdr.eth.etherType`, which indicates the packet’s protocol, the parser transitions to either state `parse_ipv4` or state

¹In scenarios with multiple ingress and egress points, where external traffic may fully traverse the internal network, the identification of outbound packets cannot rely solely on IP addresses. Instead, classification would need to be based on the forwarding port to allow the use of the ECN field while the packet is inside the internal network.

accept. If the value is 0x0800, indicating an IPv4 packet, the parser transitions to state `parse_ipv4` and extracts the IPv4 header (line 30). Finally, it transitions to the state `accept` (line 31), accepts the packet, and moves to the match-action stage.

Match-Action. This stage processes packets as instructed by control-plane-configured tables. A table consists of key-action rows and each row determines the action to be performed based on the key value. Key-action rows are updated by the control-plane, externally to P4. By applying a table, the P4 program matches the key value against table entries and executes the corresponding action. An action is a programmable function performing operations on a packet, such as forwarding, modifying headers, or dropping the packet.

The match-action block `MyCtrl` of Program 1 starts at line 34. If the IPv4 header is not valid (line 53) the packet is dropped. Otherwise, if the packet’s destination (line 54-56) is the internal network, the program checks for congestion. The standard metadata’s `enq_qdepth` field indicates the length of the queue that stores packets waiting to be processed. A predefined `THRESHOLD` is used to determine the congestion status and store it in the `ecn` field (line 57 and 59). Finally, the packet is forwarded by applying the `ipv4_lpm` table (line 61). This table, defined at line 46, matches based on longest prefix (*lpm*) of the IPv4 destination address (`hdr.ipv4.dstAddr`), and has two actions (shown on line 48): `ipv4_forward` which forwards the packet and `drop` which drops the packet. If no match exists, the default action on line 49 is invoked.

Calling conventions. P4 is a heapless language, implementing a unique copy-in/copy-out calling convention that allows static allocation of resources. P4 function parameters are optionally annotated with a direction (`in`, `inout` or `out`). The direction indicates how arguments are handled during function invocation and termination, offering fine-grained control over data visibility and potential side effects.

For example, `inout` indicates that the invoked function can both read from and write to a local copy of the caller’s argument. Once the function terminates, the caller receives the updated value of that argument. For instance, assume `hdr.ipv4.ttl` value is 10 in line 43. The invocation of `decrease_copies-in` the value 10 to parameter `x`, and the assignment on line 9 modifies `x` to value 9. Upon termination, the function copies-out the value 9 back to the caller’s parameter, changing the value of `hdr.ipv4.ttl` to 9 in line 44.

Externs. Externs are functionalities that are implemented outside the P4 program and their behavior is defined by the underlying hardware or software platform. Externs are typically used for operations that are either too complex or not directly expressible in P4’s standard constructs. This includes operations like hashing, checksum computations, and cryptographic functions. Externs can directly affect the global architectural state that is external to the P4 state, but their effects to the P4 state are controlled by the copy-in/copy-out calling convention.

For example, the extern function `mark_to_drop` (line 38) signals to the forwarding pipeline that a packet should be

discarded. Generally, the packet is sent to the port identified by the standard metadata’s `egress_spec` field, and dropping a packet is achieved by setting this field to the drop port of the switch. The drop port’s value depends on the target switch; we assume the value is 0.

A. Problem statement

The power and flexibility of P4 to programmatically process and forward packets across different networks provides opportunities for security vulnerabilities such as information leakage and covert channels. For instance, in Program 1, additionally to the ECN, the standard metadata’s `enq_qdepth` field, which indicates the length of the queue that stores packets waiting to be processed, indirectly reveals the congestion status of the current switch.

Programming errors and misconfigurations can cause information leakage. Consider the application of the `ipv4_lpm` table (line 61) which forwards the packet to a table-specified port. A bug in the branch condition on line 54, which checks the least significant bits of the `dstAddr` (e.g. by mistakenly checking `hdr.ipv4.dstAddr[7:0] == 192` instead), would result in setting the `ecn` field on the packets leaving the internal network, thus causing the packets forwarded to an external network to leak information about the internal network’s congestion state. Covert channels can also result from buggy or malicious programs. For example, by encoding the `ecn` field into the `ttl` field, an adversary can simply inspect `ttl` to deduce the congestion status.

To detect these vulnerabilities, we set out to study the security of P4 programs by means of information flow control (IFC). IFC tracks the flow of information within a program, preventing leakage from sensitive sources to public sinks. Information flow security policies are typically expressed by assigning security labels to the sources and sinks and the flow relations between security labels describe the allowed (and disallowed) information flows. In our setting, the sensitivity of sources (sinks) depends on predicates on the input (output) packets and standard metadata. Therefore, we specify the security labels of sources (i.e. input packet and switch state) by an *input policy*, while the security labels of the sinks (i.e. output packet and switch state) are specified by an *output policy*.

The input policy of Program 1, describing the security label of its sources is defined as:

If the switch’s input packet has the protocol IPv4 (i.e. `hdr.eth.etherType` is 0x0800) and its IPv4 source address `hdr.ipv4.srcAddr` belongs to the internal network subnet 192.168..*, then the `ecn` field is secret, otherwise it is public. All the other fields of the input packet are always public, while the switch’s `enq_qdepth` is always secret.* (1)

Program 1 should not leak sensitive information to external networks. An output policy defines public sinks by the ports associated with the external network and labels the fields of the corresponding packets as public.

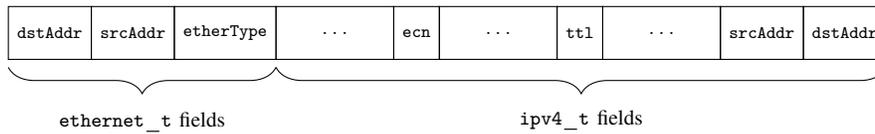


Fig. 2: Packet header

```

1  const bit<19> THRESHOLD = 10;
2
3  struct headers {
4    ethernet_t eth;
5    ipv4_t     ipv4;
6  }
7
8  void decrease (inout bit<8> x) {
9    x = x - 1;
10 }
11
12 parser MyParser(packet_in packet, out headers hdr,
13                inout metadata meta,
14                inout standard_metadata_t standard_metadata) {
15
16   state start {
17     transition parse_ethernet;
18   }
19
20   state parse_ethernet {
21     packet.extract(hdr.eth);
22     transition select(hdr.eth.etherType) {
23       0x0800: parse_ipv4;
24       default: accept;
25     }
26   }
27
28   state parse_ipv4 {
29     packet.extract(hdr.ipv4);
30     transition accept;
31   }
32 }
33
34 control MyCtrl(inout headers hdr,
35               inout metadata meta,
36               inout standard_metadata_t standard_metadata) {
37   action drop() {
38     mark_to_drop(standard_metadata);
39   }
40   action ipv4_forward(bit<48> dstAddr, bit<9> port) {
41     standard_metadata.egress_spec = port;
42     hdr.eth.srcAddr = hdr.eth.dstAddr;
43     hdr.eth.dstAddr = dstAddr;
44     decrease(hdr.ipv4.ttl);
45   }
46   table ipv4_lpm {
47     key = { hdr.ipv4.dstAddr: lpm; }
48     actions = { ipv4_forward; drop; }
49     default_action = drop();
50   }
51
52   apply {
53     if (hdr.ipv4.isValid()) {
54       if (hdr.ipv4.dstAddr[31:24] == 192 &&
55           hdr.ipv4.dstAddr[23:16] == 168){
56         if (standard_metadata.enq_qdepth >= THRESHOLD)
57           hdr.ipv4.ecn = 3;
58       } else {
59         hdr.ipv4.ecn = 0;
60       }
61       ipv4_lpm.apply(); //forward all valid packets
62     } else {
63       drop();
64     }
65   }
66 }

```

Program 1: Congestion notifier

Packets leaving the switch through ports 10 to 20 are forwarded to the external network and are observable by attackers. Therefore, all fields of such packets should be public. All the other packets are not observable by attackers. (2)

Our goal is to design a static security analysis that strikes a balance between expressiveness and automation of the verification process. We identify three main challenges that a security analysis of P4 programs should address:

- 1) Security policies are data-dependent. For instance, the `ecn` field is sensitive only if the packet is IPv4 and its IP source address is in the range `192.168.*.*`.
- 2) The analysis should be value- and path-sensitive, reflecting the different values of header fields. For example, the value of the field `etherType` determines the packet's protocol and its shape. This information influences the reachability of program paths; for instance if the packet is IPv4 the program will not go through the parser states dedicated to processing IPv6 packets.

- 3) Externs and tables behavior are not defined in P4. Tables are statically-unknown components and configured at runtime. For example, a misconfiguration of the `ipv4_lpm` table may insecurely forward packets with sensitive fields to an external network.

Note that P4 lacks many features that could negatively affect analysis precision, including heap, memory aliasing, recursion, and loops.

Threat model. Our threat model considers a network attacker that knows the code of the P4 program and observes data on public sinks, as specified by a policy. We also assume that the keys and the actions of the tables are public and observable, but tables can pass secret data as the arguments of the actions. Because of the batch-job execution model, security policies can be specified as data-dependent security types over the initial and final program states. We aim at protecting against storage channels pertaining to explicit and implicit flows, while deferring other side channels, e.g. timing, to future work.

III. SOLUTION OVERVIEW

We develop a novel combination of security type systems and interval abstractions to check information flow policies. We argue that our lightweight analysis of P4 programs provides a sweet spot balancing expressiveness, precision and automation.

Data-dependent policies are expressed by security types augmented with intervals, and the typing rules ensure that the program has no information flows from secret (H) sources to public (L) sinks. Specifically, a security type is a pair (I, ℓ) of an interval I indicating a range of possible values and a security label $\ell \in \{L, H\}$. For simplicity, we use the standard two-element security lattice $\{L, H\}$ ordered by \sqsubseteq with $\text{lub } \sqcup$. For example, the type $(\langle 1, 5 \rangle, L)$ of the `ttl` field of the `ipv4` header specifies that the `ttl` field contains public data ranging between 1 and 5.

The security types allow to precisely express data-dependent policies such as (1). The input and output policies in our approach specify the shape of the input and output packets. Since packets can have many different shapes (e.g. IPv4 or IPv6), these policies may result in multiple distinct policy cases. For example, input policy (1) results in two cases:

In the *first input policy case*, the packet’s `hdr.eth.etherType` is `0x0800`, its IPv4 source address is in the internal network of interval $\langle 192.168.0.0, 192.168.255.255 \rangle$, `hdr.ipv4.ecn` and standard metadata’s `enq_qdepth` can contain any value (represented as $\langle * \rangle$) but are classified as H , while all other header fields are $(\langle * \rangle, L)$ (omitted here). We express this policy using our security types as follows:

```

hdr.eth.etherType : ( $\langle 0x0800, 0x0800 \rangle, L$ )
hdr.ipv4.srcAddr : ( $\langle 192.168.0.0, 192.168.255.255 \rangle, L$ )
hdr.ipv4.ecn : ( $\langle * \rangle, H$ )
standard_metadata.enq_qdepth : ( $\langle * \rangle, H$ )

```

The intervals and labels in these security types describe the values and labels of the initial state of the program under this specific input policy case.

The *second input policy case* describes all the packets where `hdr.eth.etherType` is not `0x0800` or IPv4 source address is outside the range $\langle 192.168.0.0, 192.168.255.255 \rangle$, all of the packet header fields are $(\langle * \rangle, L)$, while the standard metadata’s `enq_qdepth` is still $(\langle * \rangle, H)$.

Similarly, the output policy (2) can be expressed with the output policy case: “if the standard metadata’s `egress_spec` is $(\langle 10, 20 \rangle, L)$, then all of the packet’s header fields are $(\langle * \rangle, L)$.”

It turns out that this specific output policy case is the only interesting one, even though output policy (2) can result in two distinct policy cases. In the alternative case, the fact that the attacker is unable to observe the output packet can be represented by assigning $(\langle * \rangle, H)$ to all the fields of the packet. The flow relation among security labels, as determined by the ordering of the security labels, only characterizes flows from H sources to L sinks as insecure. This implies that any policy cases where the source is L or the sink is H cannot result in

insecure flows. Thus, the alternative case is irrelevant and can be safely ignored.

Driven by the data-dependent types, we develop a new security type system that uses the intervals to provide a finer-grained assignment of security labels. Our interval analysis allows the type system to statically eliminate execution paths that are irrelevant to the security policy under consideration, thus addressing the second challenge of precise analysis. For example, our interval analysis can distinguish between states where `hdr.eth.etherType` is `0x0800` and states where it is not, essentially providing a path-sensitive analysis. This enables the analysis to avoid paths where `hdr.eth.etherType` is *not* `0x0800` when checking the policy of IPv4 packets. As a result, we exclude paths visited by non-IPv4 packets when applying the `ipv4_lpm` table in line 61. This reduces the complexity of the analysis as we avoid exploring irrelevant program paths, and helps reduce false positives in the results.

Finally, to address the challenge of tables and externs, we rely on user-defined contracts which capture a bounded model of the component’s behavior. Upon analyzing these components, the type system uses the contracts to drive the analysis. For Program 1, the contract for a correctly-configured table `ipv4_lpm` ensures that if the packet’s `hdr.ipv4.dstAddr` belongs to the internal network, then the action `ipv4_forward` (line 40) forwards the packet to ports and MAC addresses connected to the internal network.

Even if the `ipv4_lpm` table is correctly configured and its contract reflects that, bugs in the program can still cause unintended information leakage. For example, on line 54, the branch condition might have been incorrect and instead of checking the 8 most significant bits (i.e. [31:24]) of the `hdr.ipv4.dstAddr`, it checks the least significant bits (i.e. [7:0]). This bug causes the `hdr.ipv4.ecn` field in some packets destined for the external network to include congestion information, leading to unintended information leaks. Such errors are often overlooked but can be detected by our type system.

We ensure that the program does not leak sensitive information by checking the final types produced by the type system against an output policy. If these checks succeed, the program is deemed secure. The details of this process and the role of the interval information in the verification process are explained in Section VI.

IV. SEMANTICS

In this section, we briefly summarize a big-step semantics of P4. The language’s program statements, denoted by s , include standard constructs such as assignments, conditionals, and sequential composition. Additionally, P4 supports transition statements, function calls, table invocations, and extern invocations as shown in Fig. 3.

Values, represented by v , are either big-endian bitvectors \bar{b} (raw packets) or structs $\{f_1 = v_1, \dots, f_n = v_n\}$ (representing headers).

P4 states m are mappings from variables x to values v . In this slightly simplified semantics, variables are either global

$$\begin{aligned}
v &::= \bar{b} \mid \{f_1 = v_1, \dots, f_n = v_n\} \\
e &::= v \mid x \mid \ominus e \mid e \oplus e' \mid e.f \mid e[\bar{b} : bitv'] \mid \{f_1 = e_1, \dots, f_n = e_n\} \\
lval &::= x \mid lval.f \mid lval[\bar{b} : \bar{b}'] \\
s &::= \text{skip} \mid lval := e \mid s_1; s_2 \mid \text{if } e \text{ then } s_1 \text{ else } s_2 \mid \text{apply } tbl \mid \\
&\quad f(e_1, \dots, e_n) \mid \text{transition select } e \{v_1 : st_1, \dots, v_n : st_n\} st
\end{aligned}$$

Fig. 3: Syntax

or local. States can thus be represented as disjoint unions (m_g, m_l) , where m_g (m_l) maps global (local) variables only.

While externs in P4 can modify the architectural state, they cannot change the P4 state itself. To simplify our model, we integrate the architectural state into P4's global state, treating it as a part of the global state. Therefore, in our model the externs are allowed to modify the global state of P4. To maintain isolation between the program's global variables and the architectural state, we assume that the variable names used to represent the global state are distinct from those used for the architectural state.

Expressions e use a standard selection of operators including binary \oplus , unary \ominus , comparison \otimes , and struct field access, as well as bitvector slicing $e[b : a]$ extracting the slice from index a to index b of e , and $m(e)$ is the evaluation of e in state m . An lvalue $lval$ is an assignable expression, either a variable, a field of a struct, or a bitvector slice. The semantics of expressions is standard and consists of operations over bitvectors and record access.

The semantics of statements uses a mapping E from function names f to pairs $(s, (x, d))$, where (x, d) is the signature of f , a sequence of pairs (x_i, d_i) of function parameters with their directions $d_i \in \{\text{in}, \text{out}, \text{inout}\}$. Additionally, E maps parser state names st to their bodies. Furthermore, since P4 programs may depend on external components, E also maps externs f and tables t to their respective implementations.

The semantic rules presented in Fig. 4 rely on judgments of the form $E : m_1 \xrightarrow{s} m_2$ to represent the execution of statement s under mapping E which starts from state m_1 and terminates in m_2 .

Many of the rules in Fig. 4 are standard and are therefore not explained here. Rule S-CALL fetches the invoked function's body s and signature, and copies in the arguments into m'_l , which serves as the local state for the called function and is used to execute the function's body. Note that the function's body can modify the global state, but cannot change the caller's local state due to P4's calling conventions. After executing the function's body, the variables in final local state m''_l must be copied-out according to the directions specified in the function's signature. Given a direction d_i , the auxiliary function `isOut` returns true if the direction is out or inout. We rely on this function to copy-out the values from m''_l back to the callee only for parameters with out and inout direction.

For example, in Program 1 let $m_l = \{\text{hdr.ipv4.ttl} \mapsto 2\}$ when invoking `decrease` at line 44. The local state of

$$\begin{array}{c}
\text{S-SEQ} \\
\frac{E : m \xrightarrow{s_1} m' \quad E : m' \xrightarrow{s_2} m''}{E : m \xrightarrow{s_1; s_2} m''} \\
\text{S-SKIP} \quad \frac{}{E : m \xrightarrow{\text{skip}} m} \quad \text{S-ASSIGN} \quad \frac{m' = m[lval \mapsto m(e)]}{E : m \xrightarrow{lval := e} m'} \\
\text{S-COND-T} \quad \frac{m(e) = \text{true} \quad E : m \xrightarrow{s_1} m'}{E : m \xrightarrow{\text{if } e \text{ then } s_1 \text{ else } s_2} m'} \quad \text{S-COND-F} \quad \frac{m(e) = \text{false} \quad E : m \xrightarrow{s_2} m'}{E : m \xrightarrow{\text{if } e \text{ then } s_1 \text{ else } s_2} m'} \\
\text{S-CALL} \quad \frac{(s, (x, d)) = E(f) \quad m'_l = \{x_i \mapsto (m_g, m_l)(e_i)\} \quad E : (m_g, m'_l) \xrightarrow{s} (m'_g, m''_l)}{E : (m_g, m_l) \xrightarrow{f(e_1, \dots, e_n)} (m'_g, m_l)[e_i \mapsto m''_l(x_i) \mid \text{isOut}(d_i)]} \\
\text{S-EXTERN} \quad \frac{(sem_f, (x, d)) = E(f) \quad m'_l = \{x_i \mapsto (m_g, m_l)(e_i)\} \quad (m'_g, m''_l) = sem_f(m_g, m'_l)}{E : m \xrightarrow{f(e_1, \dots, e_n)} (m'_g, m_l)[e_i \mapsto m''_l(x_i) \mid \text{isOut}(d_i)]} \\
\text{S-TRANS} \quad \frac{st' = \begin{cases} st_i & \text{if } m(e) = v_i \\ st & \text{otherwise} \end{cases} \quad E : m \xrightarrow{E(st')} m'}{E : m \xrightarrow{\text{transition select } e \{v_1 : st_1, \dots, v_n : st_n\} st} m'} \\
\text{S-TABLE} \quad \frac{(\bar{e}, sem_{tbl}) = E(tbl) \quad sem_{tbl}((m_g, m_l)(e_1), \dots, (m_g, m_l)(e_n)) = (a, \bar{v}) \quad (s, (x_1, \text{none}), \dots, (x_n, \text{none})) = E(a) \quad m'_l = \{x_i \mapsto v_i\} \quad E : (m_g, m'_l) \xrightarrow{s} (m'_g, m''_l)}{E : (m_g, m_l) \xrightarrow{\text{apply } tbl} (m'_g, m_l)}
\end{array}$$

Fig. 4: Semantic rules

`decrease` (i.e. the copied-in state) becomes $m'_l = \{x \mapsto 2\}$. After executing the function's body (line 8), the final local state will be $m''_l = \{x \mapsto 1\}$ while the global state m_g remains unchanged. Finally, the copying out operation updates the caller's state to $m'' = (m_g, \{\text{ttl} \mapsto 1\})$ by updating its local state.

The S-EXTERN rule is similar to S-CALL. The key difference is that instead of keeping a body in E , we keep the extern's behavior defined through sem_f . This function takes a state containing the global m_g and copied-in state m'_l and returns (possibly) modified global and local states, represented as $sem_f(m_g, m'_l) = (m'_g, m''_l)$. Finally, the extern rule performs a copy-out procedure similar to the function call.

The S-TRANS rule defines how the program transitions between parser states based on the evaluation of expression e . It includes a default state name st for unmatched cases. If in program state m , expression e evaluates to value v_i , the program transitions to state name st_i according to the defined value-state pattern. However, if the evaluation result does not match any of the v_i values, the program instead

transitions to the default state st . For example, assume that $m(\text{hdr.eth.etherType}) = 0x0800$ on line 23 of Program 1. The select expression within the transition statement will transition to the state `parse_ipv4`, and executes its body.

Rule S-TABLE fetches from E the table’s implementation sem_{tbl} and a list of expressions \bar{e} representing table’s keys. It then proceeds to evaluate each of these expressions in the current state (m_g, m_l) , passing the evaluated values as key values to sem_{tbl} . The table’s implementation sem_{tbl} then returns an action a and its arguments \bar{v} . We rely on E again to fetch the body and signature of action a , however, since in P4 action parameters are directionless we use none in the signature to indicate there is no direction. Finally, similar to S-CALL we copy-in the arguments into m'_l , which serves as the local state for the invoked action and is used to execute the action’s body. For example, let $m(\text{hdr.ipv4.dstAddr}) = 192.168.2.2$ at line 61, and the semantics of table `ipv4_lpm` contains:

```
192.168.2.2  $\mapsto$  ipv4_forward (4A:5B:6C:7D:8E:9F, 5)
```

then the table invokes action `ipv4_forward` with arguments 4A:5B:6C:7D:8E:9F and 5.

V. TYPES AND SECURITY CONDITION

In our approach types are used to represent and track both bitvector abstractions (i.e. intervals) and security labels, and we use the same types to represent input and output policies.

In P4, bitvector values represent packet fragments, where parsing a bitvector involves slicing it into sub-bitvectors (i.e. slices), each with different semantics such as payload data or header fields like IP addresses and ports. These header fields are typically evaluated against various subnetwork segments or port ranges. Since header fields or their slices are still bitvectors, they can be conveniently represented as integers, enabling us to express the range of their possible values as $I = \langle a, b \rangle$, the interval of integers between a and b .

We say a bitvector v is typed by type τ , denoted as $v : \tau$, if τ induces a slicing of v that associates each slice with a suitable interval I and security label $\ell \in \{L, H\}$. We use the shorthand I_i^ℓ to represent a slice of length i , with interval $I \subseteq \langle 0, 2^i - 1 \rangle$ and security label ℓ . The bitvector type can therefore be presented as $\tau = I_{n_{i_n}}^{\ell_n} \dots I_{1_{i_1}}^{\ell_1}$, representing a bitvector of length $\sum_{j=1}^n i_j$ with n slices, where each slice i has interval I_i and security label ℓ_i . Singleton intervals are abbreviated $\langle a \rangle$, $\langle \rangle$ is the empty interval, and $\langle * \rangle$ is the complete interval, that is, the range $\langle 0, 2^i - 1 \rangle$ for a slice of length i . Function $\text{lbl}(\tau)$ indicates the least upper bound of the labels of slices in τ .

To illustrate this, let τ_1 be $\langle * \rangle_2^H \cdot \langle 0, 1 \rangle_3^L$ which types a bitvector of length 5 consisting of two slices. The first slice has a length of 3, with values drawn from the interval $\langle 0, 1 \rangle$ and security label L . The second slice, with a length of 2, has a security label H , and its values drawn from the complete interval $\langle 0, 3 \rangle$ (indicated by $*$). Accordingly, $\text{lbl}(\tau_1)$ evaluates to $H \sqcup L = H$.

Type τ is also used to denote a record type, where record $\{f_1 = v_1, \dots, f_n = v_n\}$ is typed as $\langle f_1 : \tau_1; \dots; f_n : \tau_n \rangle$ if each value v_i is typed with type τ_i .

In this setting, the types are not unique, as it is evident from the fact that a bitvector can be sliced in many ways and a single value can be represented by various intervals. For example, bitvector $\boxed{100}$ can be typed as $\langle 4 \rangle_3^L$, or $\langle * \rangle_1^L \cdot \langle 0, 1 \rangle_2^L$, or $\langle 2 \rangle_2^L \cdot \langle * \rangle_1^L$.

State types. A type environment, or *state type*, $\gamma = (\gamma_g, \gamma_l)$ is a pair of partial functions from variable names x to types τ . Here, γ_g and γ_l represent global and local state types, respectively, analogous to the global (m_g) and local (m_l) states in the semantics. We say that γ can type state m , written as $\gamma \vdash m$, if for every $lval$ in the domain of m , the value $m(lval)$ belongs to the interval specified by $\gamma(lval)$; formally, $\forall lval \in \text{domain}(m), m(lval) : \gamma(lval)$. Note that the typing judgment $\gamma \vdash m$ is based on the interval inclusion and it is independent of any security labels. For example, if $m(x) = 257$ then 257 is considered well-typed wrt. γ if and only if $\gamma(x)$ is an interval that contains 257 (e.g., $\langle 0, 257 \rangle$, $\langle 257, 257 \rangle$, or $\langle 100, 300 \rangle$).

A state type might include a type with an empty interval; we call this state type *empty* and denote it as \bullet .

Let $\text{lblOf}(lval, \gamma)$ be the least upper bound of the security labels of all the slices of $lval$ in state type γ . The states m_1 and m_2 are considered *low equivalent with respect to* γ , denoted as $m_1 \sim_\gamma m_2$, if for all $lval$ such that $\text{lblOf}(lval, \gamma) = L$, then $m_1(lval) = m_2(lval)$ holds.

Example 1. Assume a state type $\gamma = \{x \mapsto \langle * \rangle_1^H \cdot \langle 0, 1 \rangle_2^L\}$. The following states $m_1 = \{x \mapsto \boxed{000}\}$ and $m_2 = \{x \mapsto \boxed{100}\}$ are low equivalent wrt. γ . However, states $m_1 = \{x \mapsto \boxed{000}\}$ and $m_3 = \{x \mapsto \boxed{101}\}$ are not low equivalent even though both can be typed by γ .

Contracts. A table consists of key-action rows, and in our threat model, we assume the keys and actions of the tables are always public (i.e. L), but the arguments of the actions can be secret (i.e. H). Given that tables are populated by the control-plane, the behavior of a table is unknown at the time of typing. We rely on user-specified contracts to capture a bounded model of the behavior of the tables. In our model, a table’s contract has the form $(\bar{e}, \text{Cont}_{tbl})$, where \bar{e} is a list of expressions indicating the keys of the table, and Cont_{tbl} is a set of tuples $(\phi, (a, \bar{\tau}))$, where ϕ is a boolean expression defined on \bar{e} , and a denotes an action to be invoked with argument types $\bar{\tau}$ when ϕ is satisfied.

For instance, the `ipv4_lpm` table of Program 1 uses `hdr.ipv4.dstAddr` as its key, and can invoke two possible actions: `drop` and `ipv4_forward`. An example of a contract for this table is depicted in Fig. 5. This contract models a table that forwards the packets with `hdr.ipv4.dstAddr = 192.*.*` to ports 1-9, the ones with `hdr.ipv4.dstAddr = 10.*.*` to ports 10-20, and drops all the other packets. Notice that in the first case, the first argument resulting from the table look up is secret.

The table contracts are essentially the security policies of the tables, where ϕ determines a subset of table rows that invoke the same action (a) with the same argument types ($\bar{\tau}$).

```

(hdr.ipv4.dstAddr),
{ (dstAddr[31 : 24] = 192, (ipv4_forward, [(*)H48, (1, 9)L9]))
  (dstAddr[31 : 24] = 10, (ipv4_forward, [(*)L48, (10, 20)L9]))
  (dstAddr[31 : 24] ≠ 192 ∧ dstAddr[31 : 24] ≠ 10, (drop, [])) }

```

Fig. 5: The contract of ipv4_lpm table

Using the labels in $\bar{\tau}$, and given action arguments \bar{v}_1 and \bar{v}_2 , we define $\bar{v}_1 \sim_{\bar{\tau}} \bar{v}_2$ as $|\bar{v}_1| = |\bar{v}_2| = |\bar{\tau}|$ and for all i , $v_{1_i} : \tau_i$ and $v_{2_i} : \tau_i$, and if $\text{lbl}(\tau_i) = L$ then $v_{1_i} = v_{2_i}$. Note that $\text{lbl}(\tau_i)$ returns the least upper bound of the labels of all τ_i 's slices, hence if there is even one H slice in τ_i , $\text{lbl}(\tau_i)$ would be H . We use mapping T to associate table names tbl with their contracts.

We say that two mappings E_1 and E_2 are considered *indistinguishable* wrt. T , denoted as $E_1 \sim_T E_2$, if for all tables tbl such that $(\bar{e}_1, \text{sem}_{1_{tbl}}) = E_1(tbl)$, $(\bar{e}_2, \text{sem}_{2_{tbl}}) = E_2(tbl)$, $(\bar{e}, \text{Cont}_{tbl}) = T(tbl)$ then $\bar{e}_1 = \bar{e}_2 = \bar{e}$, and for all $(\phi, (a, \bar{\tau})) \in \text{Cont}_{tbl}$, and for all arbitrary states m_1 and m_2 , such that $m_1(\bar{e}) = m_2(\bar{e}) = \bar{v}$ and $m_1(\phi) \Leftrightarrow m_2(\phi)$, if $m_1(\phi)$ then \bar{v}_1, \bar{v}_2 exist such that $\text{sem}_{1_{tbl}}(\bar{v}) = (a, \bar{v}_1)$, $\text{sem}_{2_{tbl}}(\bar{v}) = (a, \bar{v}_2)$, and $\bar{v}_1 \sim_{\bar{\tau}} \bar{v}_2$. In other words, T -indistinguishability of E_1 and E_2 guarantees that given equal key values, E_1 and E_2 return the same actions with $\bar{\tau}$ -indistinguishable arguments \bar{v}_1 and \bar{v}_2 such that these arguments are in bound wrt. their type $\bar{\tau}$.

Security condition. As explained in Section III the input and output policy cases are expressed by assigning types to program variables. State types, specifying security types of program variables, are used to formally express input and output policy cases. Hereafter, we use γ_i and γ_o to denote input and output policy cases, respectively. Using this notation, the input policy, denoted by Γ_i , is represented as a set of input policy cases γ_i . Similarly, the output policy is expressed as a set of output policy cases γ_o and denoted by Γ_o .

Given this intuition, we say two states m_1 and m_2 are *indistinguishable* wrt. a policy case γ if $\gamma \vdash m_1$, $\gamma \vdash m_2$, and $m_1 \sim_{\gamma} m_2$. Relying on this, we present our definition of noninterference as follows:

Definition 1 (Noninterference). A program s is *noninterfering* wrt. the input and output policy cases γ_i and γ_o , and table contract mapping T , if for all mappings E_1, E_2 and states m_1, m_2, m'_1 such that:

- $E_1 \sim_T E_2$,
- $\gamma_i \vdash m_1, \gamma_i \vdash m_2$, and $m_1 \sim_{\gamma_i} m_2$,

- $E_1 : m_1 \xrightarrow{s} m'_1$

there exists a state m'_2 such that:

- $E_2 : m_2 \xrightarrow{s} m'_2$,
- if $\gamma_o \vdash m'_1$, then $\gamma_o \vdash m'_2$ and $m'_1 \sim_{\gamma_o} m'_2$.

The existential quantifier over the state m'_2 does not mean that the language is non-deterministic, in fact if such state

exists it is going to be unique. This existential quantifier guarantees that our security condition is termination sensitive, meaning that it only accepts the cases where the program terminates for both initial states m_1 and m_2 .

Intuitively, Definition 1 relies on two different equivalence relations: one induced by the input policy case and one by the output policy case. The former induces a partial equivalence relation (PER) [11], $P_{\gamma_i}(m_1, m_2) = \gamma_i \vdash m_1 \wedge \gamma_i \vdash m_2 \wedge m_1 \sim_{\gamma_i} m_2$, such that the domain contains only states that satisfy the intervals of γ_i . Similarly, the latter induces $Q_{\gamma_o}(m_1, m_2) = (\gamma_o \vdash m_1 \wedge \gamma_o \vdash m_2 \wedge m_1 \sim_{\gamma_o} m_2) \vee (\gamma_o \not\vdash m_1 \wedge \gamma_o \not\vdash m_2)$, which is an equivalence relation (ER). A program is then noninterfering wrt. γ_i and γ_o if every class of the PER P_{γ_i} is mapped to a class of the ER Q_{γ_o} .

This condition implies the following intuitive assumptions: (1) the policy cases are public knowledge, (2) entailment of a state on the intervals of a policy case is public knowledge, (3) the states that do not entail the intervals of an input policy (i.e., those outside the domain of the PER) are considered entirely public, and their corresponding execution is unconstrained, (4) the attacker can observe whether the final state entails the intervals of the output policy, and (5) the attacker cannot observe any additional information about states that do not entail the intervals of the output policy. We use the following examples to further discuss our security condition.

Example 2. Consider the input and output policy cases:

$$\begin{aligned} \gamma_i &= \{a \mapsto \langle 0, 256 \rangle^H, b \mapsto \langle * \rangle^L_9\} \\ \gamma_o &= \{a \mapsto \langle * \rangle^H_9, b \mapsto \langle 1025, * \rangle^L_9\} \end{aligned}$$

- The inputs $a_1 = 257$ and $a_2 = 258$ are distinguishable by the attacker, since they do not fall in the H interval $\langle 0, 256 \rangle$ under γ_i .
- The inputs $a_1 = 0$ and $a_2 = 256$ are indistinguishable, since they belong to the H interval $\langle 0, 256 \rangle$ under γ_i .
- Under γ_o , any value $b \geq 1025$ is distinguishable by the attacker, otherwise it is indistinguishable since it falls outside the L interval $\langle 1025, * \rangle$.

We consider pairs of states m_1, m_2 such that $\gamma_i \vdash m_1$, $\gamma_i \vdash m_2$, and $m_1 \sim_{\gamma_i} m_2$. For example,

- 1) $m_1(a) = a_1$ and $m_2(a) = a_2$, where $a_1, a_2 \in \langle 0, 256 \rangle$.
- 2) $m_1(b) = m_2(b) = b_0$.

We use the above policies and states to discuss the security condition of the following one-line programs:

- **b=a**
This program yields $m'_1(b) = a_1$ and $m'_2(b) = a_2$. Since $m'_1(b) \notin \langle 1025, * \rangle$, then $\gamma_o \not\vdash m'_1$, hence noninterference is trivially satisfied. Intuitively, despite the variable a being H , the output on the variable b is not observable by the attacker.
- **if (a<=1024) then b=a else skip**
This program yields $m'_1(b) = a_1$ and $m'_2(b) = a_2$. Since both a_1 and a_2 are in $\langle 0, 256 \rangle$, the program executes **b=a**

in the true branch. Thus, $\gamma_o \not\vdash m'_1$ and noninterference is trivially satisfied.

- $b=a+1000$

(i) Let $a_1 = 25$ and $a_2 = 26$. Then $m'_1(b) = 1025$ and $m'_2(b) = 1026$ indicating $\gamma_o \vdash m'_1$ and $\gamma_o \vdash m'_2$, however $m'_1 \not\sim_{\gamma_o} m'_2$, hence the program is interfering.

(ii) Let $a_1 = 25$ and $a_2 = 0$. Then $m'_1(b) = 1025$ and $m'_2(b) = 1000$ indicating $\gamma_o \vdash m'_1$ and $\gamma_o \not\vdash m'_2$, hence the program is interfering.

Example 3. Assume program `if y==1 then x=1 else x=x+1`, input policy case $\gamma_i = \{x \mapsto \langle * \rangle_2^H, y \mapsto \langle 1 \rangle_3^L\}$, and initial states $m_1 = \{x \mapsto \boxed{10}, y \mapsto \boxed{001}\}$ and $m_2 = \{x \mapsto \boxed{01}, y \mapsto \boxed{001}\}$. We can see that $\gamma_i \vdash m_1$, $\gamma_i \vdash m_2$, and $m_1 \sim_{\gamma_i} m_2$. In a scenario where the only initial states are m_1 and m_2 , executing this program would result in final states $m'_1 = \{x \mapsto \boxed{01}, y \mapsto \boxed{001}\}$ and $m'_2 = \{x \mapsto \boxed{01}, y \mapsto \boxed{001}\}$, respectively. Given output policy case $\gamma_o = [x \mapsto \langle * \rangle_2^L, y \mapsto \langle 1 \rangle_3^L]$, we say that this program is noninterfering wrt. γ_o because $\gamma_o \vdash m'_1$, $\gamma_o \vdash m'_2$, and $m'_1 \sim_{\gamma_o} m'_2$.

We extend the definition of noninterference to input policies Γ_i and output policies Γ_o , requiring the program to be noninterfering for *every pair* of input and output policy cases. In our setting, the output policy, which indicates the shape of the output packets, describes what the attacker observes. As such, it is typically independent of the shape of the input packet and the associated input policy. Thus, our approach does not directly pair input and output policy cases. Instead, it ensures that the program is noninterfering for all combinations of input and output policy cases.

VI. SECURITY TYPE SYSTEM

We introduce a security type system that combines security types and interval abstractions. Our approach begins with an input policy case and conservatively propagates labels and intervals of P4 variables. In the following, we assume that the P4 program is well-typed.

A. Typing of expressions

The typing judgment for expressions is $\gamma \vdash e : \tau$. Rules for values, variables, and records are standard and omitted here.

P4 programs use bitvectors to represent either raw packets (e.g. `packet_in` packet of line 12) or finite integers (e.g. `x` of line 8). While there is no distinction between these two cases at the language level, it is not meaningful to add or multiply two packets, as it is not extracting a specific byte from an integer representing a time-to-live value. For this reason, we expect that variables used to marshal records have multiple slices but are not used in arithmetic operations, while variables used for integers have one single slice and are not used for sub-bitvector operations. This allows us to provide a relatively simple semantics of the slice domain, which is sufficient for many P4 applications.

T-SINGLESLICEBS

$$\frac{\gamma \vdash e_1 : \langle I_1 \rangle_i^{\ell_1} \quad \gamma \vdash e_2 : \langle I_2 \rangle_i^{\ell_2}}{\gamma \vdash e_1 \oplus e_2 : \langle I_1 \oplus I_2 \rangle_i^{\ell_1 \sqcup \ell_2}}$$

$$\gamma \vdash \ominus e_1 : \langle \ominus I_1 \rangle_i^{\ell_1}$$

$$\gamma \vdash e_1 \otimes e_2 : \langle I_1 \otimes I_2 \rangle_i^{\ell_1 \sqcup \ell_2}$$

T-SINGLESLICEBS rule allows the reuse of standard interval analysis for binary, unary, and comparison operations over bitvectors that have only *one* single slice. The resulting label is the least upper bound of labels associated with the input types.

T-ALIGNEDSLICE

$$\frac{\gamma \vdash e : \langle I_n \rangle_{i_n}^{\ell_n} \cdots \langle I_1 \rangle_{i_1}^{\ell_1}}{\gamma \vdash e[\sum_{j=1}^b i_j : \sum_{j=1}^a i_j] : \langle I_b \rangle_{i_b}^{\ell_b} \cdots \langle I_a \rangle_{i_a}^{\ell_a}}$$

T-NONALIGNEDSLICE

$$\frac{\gamma \vdash e : \langle I_n \rangle_{i_n}^{\ell_n} \cdots \langle I_1 \rangle_{i_1}^{\ell_1}}{\gamma \vdash e[b : a] : \langle * \rangle_{b-a}^{\ell_i}}$$

In the slicing rules, sub-bitvector (i.e. $e[b : a]$) preserves precision only if the slices of the input are aligned with sub-bitvector's indexes, otherwise sub-bitvector results in $\langle * \rangle$, representing all possible values. The following lemmas show that interval and labeling analysis of expressions is sound:

Lemma 1. *Given expression e , state m , and state type γ such that $\gamma \vdash m$, if the expression is well-typed $\gamma \vdash e : \tau$, and evaluates to a value $m(e) = v$, then:*

- v is well-typed wrt. to the interval of type τ (i.e. $v : \tau$).
- for every state m' such that $m \sim_{\gamma} m'$, if $\text{lbl}(\tau) = L$, then $m'(e) = v$.

B. Typing of statements

To present the typing rules for statements, we rely on some auxiliary notations and operations to manipulate state types, which are introduced informally here due to space constraints. The properties guaranteed by these operations are reported in Appendix C.

$\gamma[lval \mapsto \tau]$ indicates updating the type of `lval`, which can be a part of a variable, in state type γ . $\gamma \uparrow \uparrow \gamma'$ updates γ such that for every variable in the domain of γ' , the type of that variable in γ is updated to match γ' . $\text{refine}(\gamma, e)$ returns an overapproximation of γ that satisfy the abstraction of γ and the predicate e . $\text{join}(\gamma_1, \gamma_2)$ returns an overapproximation of γ_1 , whose labels are at least as restrictive as γ_1 and γ_2 . These operations tend to overapproximate, potentially causing a loss of precision in either the interval or the security label, as illustrated in the following example:

Example 4. Let x be mapped to an interval between 2 and 8, or in binary, bitvectors between $\boxed{0010}$ and $\boxed{1000}$, in γ . That is, $\gamma = \{x \mapsto \langle 2, 8 \rangle_4^L\}$. The following update $\gamma[x[3 : 3] \mapsto \langle 0 \rangle_1^H]$ modifies the slice $x[3 : 3]$ and results in the state type $\{x \mapsto \langle 0 \rangle_1^H \cdot \langle * \rangle_3^L\}$. Here, `lvalue` $x[2 : 0]$

loses precision because after updating $x[3 : 3]$, the binary representation of the interval of lvalue $x[2 : 0]$ would be between $\boxed{010}$ and $\boxed{000}$, that is every 3-bit value except $\boxed{001}$. Such value set cannot be represented by a single continuous interval, hence we overapproximate to the complete interval $\langle * \rangle$.

Similarly, the operation $\text{refine}(\gamma, x[3 : 3] < 1)$ updates the interval of lvalue $x[3 : 3]$ which results in $\{x \mapsto \langle 0 \rangle_1 \cdot \langle * \rangle_3^L\}$ where lvalue $x[2 : 0]$ again loses precision.

On the other hand, an operation such as $\text{join}(\gamma, \{x \mapsto \langle * \rangle_1^H \cdot \langle * \rangle_3^L\})$ does not modify the intervals of γ , but since the $\langle * \rangle_1^H$ slice overlaps with a slice of x in γ its label should be raised, which results in $\gamma' = \{x \mapsto \langle 2, 8 \rangle_4^H\}$.

The security typing of statement s uses judgments of the form $T, pc, \gamma \vdash s : \Gamma$, where pc is the security label of the current program context, T is a static mapping, and γ is a state type. We use T to map a parser state name (st) or function name (f) to their bodies. For functions, T also returns their signatures. Moreover, as described in Section V, we also use T to map externs and tables to their contracts. The typing judgment concludes with Γ , which is a set of state types. In our type system, the security typing is not an on-the-fly check that immediately rejects a program when encountering an untypeable statement. Instead, we proceed with typing the program and produce a state type for each path and accumulate all of those in a final set Γ . This is done in order to increase precision, by minimizing the need to unify, and hence overapproximate, intermediate typings during type derivation. This is indeed one of the key technical innovations of our type system, as explained in more detail below. Once the final set Γ is obtained, the state types within Γ are then verified against the output security policy Γ_o , ensuring that they meet all the output policy cases γ_o in Γ .

In the following rules, we use $\text{raise}(\tau, \ell)$ to return a type where each label ℓ' within τ has been updated to $\ell' \sqcup \ell$.

$$\text{T-ASSIGN} \quad \frac{\gamma \vdash e : \tau \quad \tau' = \text{raise}(\tau, pc) \quad \gamma' = \gamma[lval \mapsto \tau']}{T, pc, \gamma \vdash lval := e : \{\gamma'\}}$$

T-ASSIGN rule follows the standard IFC convention. It updates the type of the left-hand-side of the assignment (i.e. $lval$) with the type of expression e while raising its security label to the current security context pc in order to capture indirect information flows.

$$\text{T-SEQ} \quad \frac{T, pc, \gamma \vdash s_1 : \Gamma_1 \quad \forall \gamma_1 \in \Gamma_1. T, pc, \gamma_1 \vdash s_2 : \Gamma_2^{\gamma_1} \quad \Gamma' = \bigcup_{\gamma_1 \in \Gamma_1} \Gamma_2^{\gamma_1}}{T, pc, \gamma \vdash s_1; s_2 : \Gamma'}$$

T-SEQ types the sequential composition of two statements. This rule type checks the first statement s_1 , gathering all possible resulting state types into an intermediate set Γ_1 . Then,

for each state type in this intermediate set, the rule type checks the second statement s_2 , and accumulates all resulting state types into the final state type set Γ' .

$$\text{T-COND} \quad \frac{\gamma \vdash e : \tau \quad \ell = \text{lbl}(\tau) \quad pc' = pc \sqcup \ell \quad T, pc', (\text{refine}(\gamma, e)) \vdash s_1 : \Gamma_1 \quad T, pc', (\text{refine}(\gamma, \neg e)) \vdash s_2 : \Gamma_2}{T, pc, \gamma \vdash \text{if } e \text{ then } s_1 \text{ else } s_2 : \text{joinOnHigh}(\Gamma_1 \cup \Gamma_2, \ell)}$$

T-COND rule types the two branches using state types refined with the branch condition and its negation, which results in the state type sets Γ_1 and Γ_2 , respectively. The final state type set is a simple union of Γ_1 and Γ_2 .

However, in order to prevent implicit information leaks, if the branch condition is H , the security labels of Γ_1 and Γ_2 should be joined. We do this by the auxiliary function joinOnHigh , defined as follows:

$$\text{joinOnHigh}(\Gamma, \ell) = \begin{cases} \text{join}(\Gamma) & \text{if } \ell = H \\ \Gamma & \text{otherwise} \end{cases}$$

where the join operator has been lifted to Γ and defined as $\text{join}(\Gamma) = \{\text{join}(\gamma, \Gamma) \mid \gamma \in \Gamma\}$, $\text{join}(\gamma, \{\gamma'\} \cup \Gamma) = \text{join}(\text{join}(\gamma, \gamma'), \Gamma)$ and $\text{join}(\gamma, \emptyset) = \gamma$.

Example 5. Consider the conditional statement on line 56 of Program 1, where initially $\gamma = \{\text{enq_qdepth} \mapsto \langle * \rangle_{19}^H, \text{hdr.ipv4.ecn} \mapsto \langle * \rangle_2^L, \dots\}$. Since the label of enq_qdepth is H , after the assignment on line 57, hdr.ipv4.ecn becomes H in Γ_1 . However, since there is no `else` branch, s_2 is trivially `skip`, meaning that hdr.ipv4.ecn remains L in Γ_2 . Typically, in IFC, the absence of an update for hdr.ipv4.ecn in the `else` branch leaks that the if statement's condition does not hold. To prevent this, we join the security labels of all state types if the branch condition is H . Therefore, in the final state set Γ' , hdr.ipv4.ecn is labeled H .

Even on joining the security labels, **T-COND** does not merge the final state types in order to maintain abstraction precision. To illustrate this consider program `if b then x[0:0]=0 else x[0:0]=1`, where the pc and the label of b are both L , and an initial state type $\gamma = \{x \mapsto \langle * \rangle_3^H; b \mapsto \langle * \rangle_1^L\}$. After typing both branches, the two typing state sets are $\Gamma_1 = \{\{x \mapsto \langle * \rangle_2^H \cdot \langle 0 \rangle_1^L\}\}$ and $\Gamma_2 = \{\{x \mapsto \langle * \rangle_2^H \cdot \langle 1 \rangle_1^L\}\}$. Performing a union after the conditional preserves the labeling and abstraction precision of $x[0:0]$, whereas merging them would result in a loss of precision.

T-TRANS rule types parser transitions. Similar to **T-COND**, it individually types each state's body and then joins or unions the final state types based on the label of pc .

T-EMPTYTYPE Refining a state type might lead to an empty abstraction for some variables. We call these states empty and denote them by \bullet . An empty state indicates that there is no state m such that $\bullet \vdash m$. The rule states that from an empty state type, any statement can result in any final state type,

$$\begin{array}{c}
\text{T-TRANS} \\
\gamma \vdash e : \tau \quad \ell = \text{lbl}(\tau) \quad pc' = pc \sqcup \ell \\
\gamma'_i = \text{refine}(\gamma, e = v_i \wedge \bigwedge_{j < i} e \neq v_j) \quad T, pc', \gamma'_i \vdash T(st_i) : \Gamma_i \\
\gamma'_d = \text{refine}(\gamma, \bigwedge_{j < i} e \neq v_j) \quad T, pc', \gamma'_d \vdash T(st) : \Gamma_d \\
\Gamma' = \Gamma_d \cup \left(\bigcup_i \Gamma_i \right) \quad \Gamma'' = \text{joinOnHigh}(\Gamma', \ell) \\
\hline
T, pc, \gamma \vdash \text{transition select } e \{v_1 : st_1, \dots, v_n : st_n\} st : \Gamma''
\end{array}$$

$$\begin{array}{c}
\text{T-EMPTYTYPE} \\
\hline
T, L, \bullet \vdash s : \Gamma
\end{array}$$

since there is no concrete state that matches the initial state type. Notice that Γ can simply be *empty* and allow the analysis to prune unsatisfiable paths. This rule applies *only* when pc is L . For cases where pc is H , simply pruning the empty states is *unsound*, as illustrate by the following example:

Example 6. Assume the state type $\gamma = \{\text{enq_qdepth} \mapsto \langle 5 \rangle_{19}^H, \text{hdr.ipv4.ecn} \mapsto \langle * \rangle_2^L, \dots\}$, upon reaching the conditional statement on line 56 of Program 1. The refinement of the *then* branch under the condition $\text{enq_qdepth} \geq \text{THRESHOLD}$ (where THRESHOLD is a constant value 10) results in the empty state $\bullet = \{\text{enq_qdepth} \mapsto \langle \rangle_{19}^H, \dots\}$, where $\langle \rangle_{19}^H$ denotes an empty interval. If we prune this empty state type, the final state type set Γ' contains only the state types obtained from the *else* branch (which is *skip*). This is *unsound* because a L -observer would be able see that the value of hdr.ipv4.ecn has remained unchanged and infer that the H field enq_qdepth was less than 10.

There is a similar problem of implicit flows in dynamic information flow control, where simply upgrading a L variable to H in only one of the branches when pc is H might result in partial information leakage. This is because the variable contains H data in one execution while it might remain L on an alternative execution. To overcome this problem, many dynamic IFC methods employ the so-called no-sensitive-upgrade (NSU) check [12], which terminates the program's execution whenever a L variable is updated in a H context. Here, to overcome this problem, we type all the statements in all branches whenever the pc is H , even when the state type is empty [13], [14]. For instance, in Example 6, we type-check the *then* branch under an empty state type, and by rule T-COND the security labels of the final state types of both branches are joined, resulting in hdr.ipv4.ecn 's label being H in all the final state types.

$$\begin{array}{c}
\text{T-CALL} \\
\gamma \vdash \bar{e} : \bar{\tau} \quad \text{tCall}(T, f, pc, \bar{\tau}, \gamma, \Gamma) \\
\hline
T, pc, \gamma \vdash f(\bar{e}) : \Gamma
\end{array}$$

T-CALL rule types function calls. It individually types the function arguments e_i to obtain their types τ_i , and passes them

to auxiliary function tCall , defined as:

$$\begin{array}{c}
(s, \overline{(x, d)}) = T(f) \quad \gamma_f = \{x_i \mapsto \tau_i\} \quad T, pc, (\gamma_g, \gamma_f) \vdash s : \Gamma' \\
\Gamma = \{(\gamma'_g, \gamma_l)[e_i \mapsto \gamma'_f(x_i) \mid \text{isOut}(d_i)] \mid (\gamma'_g, \gamma'_f) \in \Gamma'\}
\end{array}$$

which retrieves the function's body s and its signature $\overline{(x, d)}$ from the mapping T . Creates a new local state type γ_f by assigning each argument to its corresponding type (i.e. copy-in), and then types the function's body to obtain the resulting state type set Γ' . Finally, tCall produces Γ by copying out the out and inout parameters (identified by the isOut function), which means updating the passed lvalues (i.e. e_i) with the final types of their corresponding parameters (i.e. $\gamma'_f(x_i)$).

Example 7. Assume that at line 44 of Program 1, the ttl in the state type is mapped to $\langle 1, 10 \rangle_8^L$. Calling decrease entails creating a new local state type and copying in the arguments, which yields $\gamma_{\text{decrease}} = \{x \mapsto \langle 1, 10 \rangle_8^L\}$. Typing the function's body ($x = x - 1$) results in the state type $\gamma'_{\text{decrease}} = \{x \mapsto \langle 0, 9 \rangle_8^L\}$. The final Γ'' is produced by copying out arguments back to the initial state type which would map ttl to $\langle 0, 9 \rangle_8^L$.

In contrast to standard type systems, we directly type the body of the function, instead of typing functions separately and in isolation. The main reason is that the intervals and labels of the types of actual arguments can be different for each invocation of the function. Notice that the nested analysis of the invoked function does not hinder termination of our analysis since P4 does not support recursion, eliminating the need to find a fix point for the types [15].

$$\begin{array}{c}
\text{T-TABLE} \\
(\bar{e}, \text{Cont}_{tbl}) = T(tbl) \quad \gamma \vdash e_i : \tau_i \\
\ell = \bigsqcup_i \text{lbl}(\tau_i) \quad pc' = pc \sqcup \ell \quad \forall (\phi_j, (a_j, \bar{\tau}_j)) \in \text{Cont}_{tbl}. \\
\gamma_j = \text{refine}(\gamma, \phi_j) \quad \text{tCall}(T, a_j, pc', \bar{\tau}_j, (\gamma_g, \gamma_l), \Gamma_j) \\
\hline
T, pc, \gamma \vdash \text{apply } tbl : \text{joinOnHigh}(\cup_j \Gamma_j, \ell)
\end{array}$$

T-TABLE rule is similar to T-COND and T-CALL. It relies on user-specified contracts to type the tables. A contract, as introduced in Section V, has the form $(\bar{e}, \text{Cont}_{tbl})$, where Cont_{tbl} consists of a set of triples $(\phi, (a, \bar{\tau}))$. Each triple specifies a condition ϕ , under which an action a is executed with arguments of specific types $\bar{\tau}$. A new context pc' is produced by the initial pc with the least upper bound of the labels of the keys.

For each triple $(\phi_j, (a_j, \bar{\tau}_j))$, T-TABLE relies on tCall to type the action a_j 's body under pc' , similar to T-CALL, and accumulates the resulting state types into a set (i.e. $\cup_j \Gamma_j$). Finally, T-TABLE uses $\text{joinOnHigh}(\cup_j \Gamma_j, \ell)$ to join their labels if ℓ was H .

Example 8. Given the table contract depicted in Fig. 5, assume a state type γ where pc is L and hdr.ipv4.dstAddr is typed as $\langle 192 \rangle_8^L \cdot \langle 168 \rangle_8^L \cdot \langle * \rangle_{16}^L$. According to T-TABLE, refining γ produces three state types, out of which only one is not empty: $\text{refine}(\gamma, \text{dstAddr}[31 : 24] = 192)$. This

refined state is used to type the action `ipv4_forward` with arguments $[\langle * \rangle_{48}^H, \langle 1, 9 \rangle_9^L]$. The two empty states should be used to type the actions `ipv4_forward` (with arguments $[\langle * \rangle_{48}^L, \langle 10, 20 \rangle_9^L]$) and `drop`. However, these states can be pruned by T-EMPTYTYPE, since pc' is L .

$$\begin{array}{c}
\text{T-EXTERN} \\
(\gamma_g, \gamma_t) \vdash e_i : \tau_i \quad (\text{Cont}_E, (x_1, d_1), \dots, (x_n, d_n)) = T(f) \\
\gamma_f = \{x_i \mapsto \tau_i\} \quad \forall (\gamma_i, \phi, \gamma_t) \in \text{Cont}_E. (\gamma_g, \gamma_f) \sqsubseteq \gamma_i \\
\Gamma' = \{\gamma' \uparrow \text{raise}(\gamma_t, pc) \mid (\gamma_i, \phi, \gamma_t) \in \text{Cont}_E \\
\quad \wedge \text{refine}((\gamma_g, \gamma_f), \phi) = \gamma' \neq \bullet\} \\
\Gamma'' = \{(\gamma'_g, \gamma_t)[e_i \mapsto \gamma'_f(x_i) \mid \text{isOut}(d_i)] \mid (\gamma'_g, \gamma'_f) \in \Gamma'\} \\
\hline
T, pc, (\gamma_g, \gamma_t) \vdash f(e_1, \dots, e_n) : \Gamma''
\end{array}$$

T-EXTERN types the invocation of external functions. It is similar to T-CALL with the main difference that the semantics of external functions are not defined in P4, therefore, we rely on user-specified contracts to approximate their behavior. An extern contract is a set of tuples $(\gamma_i, \phi, \gamma_t)$, where γ_i is the input state type, ϕ is a boolean expression defined on the parameters of the extern, and γ_t indicates the state type components updated by the extern function (i.e., its side effects).

γ_i denotes a contract-defined state type that must be satisfied prior to the invocation of the extern, and the rule T-EXTERN ensure that the initial state (γ_g, γ_f) is at most as restrictive as γ_i . This approach is standard in type systems where functions are type-checked in isolation using predefined pre- and post-typing environments. For each $(\gamma_i, \phi, \gamma_t)$ tuple in the contract, T-EXTERN refines the initial state type (γ_g, γ_t) by ϕ yielding γ' , and filters out all γ' s that do not satisfy ϕ (i.e., the refinement $\text{refine}((\gamma_g, \gamma_f), \phi)$ is \bullet). This is sound because we assume for all the variables appeared in ϕ , the least upper bound of their labels within γ_i is less restrictive than the lower bound of γ_t . We raise the label of all elements in the γ_t to pc to capture indirect flows arising from updating the state type γ' in a H context, and then use \uparrow operation to update γ' with the types in γ_t . The final state type set Γ' is produced by copying out the out and inout parameters from γ' .

Example 9. In Program 1, let the contract for `mark_to_drop` at line 38 be defined as:

$$\{ \langle \text{egress_spec} \mapsto \langle * \rangle_9^L \rangle, \text{true}, \{ \langle \text{egress_spec} \mapsto \langle 0 \rangle_9^L \rangle \}$$

which indicates that given an input state type $\langle \text{egress_spec} \mapsto \langle * \rangle_9^L \rangle$ the extern always sets the value of `egress_spec` to zero. Assuming an initial state type $\gamma = \langle \text{egress_spec} \mapsto \langle 7 \rangle_9^L \rangle$. Since the condition of the contract is true the refinement in this state type does not modify γ . This state type will be updated with the contract's γ_t to become $\langle \text{egress_spec} \mapsto \langle 0 \rangle_9^L \rangle$ if pc is L , otherwise $\langle \text{egress_spec} \mapsto \langle 0 \rangle_9^H \rangle$.

To guarantee the abstraction soundness of externs, for any input state m to the externs semantics $m' = \text{sem}_f(m)$ and the

contracts set $(\gamma_i, \phi, \gamma_t)$ must satisfy the following properties:

- 1) Every input state m must satisfy some condition in the contract set ϕ , i.e., $\exists \phi. \phi(m)$
- 2) All modified variables in output state m' must be in the domain of γ_t , and their abstraction types in γ_t must hold, i.e. $\{x. m(x) \neq m'(x)\} \subseteq \text{domain}(\gamma_t)$, and for all $x \in \text{domain}(\gamma_t)$ holds $m'(x) : \gamma_t(x)$.

Additionally, to guarantee the labeling soundness of externs, the contracts must satisfy the following properties:

- 1) Conditions must preserve secrecy with respect to the output state type. For all variable names $\{x_1, \dots, x_n\}$ appearing in the contract's condition ϕ , holds $\text{lbl}(\gamma_i(x_1)) \sqcup \dots \sqcup \text{lbl}(\gamma_i(x_n)) \sqsubseteq \text{lb}(\gamma_t)$.
- 2) Extern semantics must preserve low-equivalence. Given any states m_1 and m_2 , If $\phi(m_1)$ and $m_1 \sim_{\gamma_i} m_2$, then $m'_1 = \text{sem}_f(m_1)$, $m'_2 = \text{sem}_f(m_2)$, then the difference between the two output states must be also low equivalent $(m'_1 \setminus m_1) \sim_{\gamma_t} (m'_2 \setminus m_2)$.

C. Soundness

Given initial state types γ_1 and γ_2 , and initial states m_1 and m_2 , we write $m_1 \stackrel{\gamma_2}{\sim}_{\gamma_1} m_2$ to indicate that $\gamma_1 \vdash m_1$, $\gamma_2 \vdash m_2$, and $m_1 \sim_{\gamma_1 \sqcup \gamma_2} m_2$.

The type system guarantees that a well-typed program terminates, and the final result is well-typed wrt. at least one of the resulting state types.

Lemma 2 (Soundness of abstraction and labeling). *Given initial state types γ_1 and γ_2 , and initial states m_1 and m_2 , such that $T, pc, \gamma_1 \vdash s : \Gamma_1$ and $T, pc, \gamma_2 \vdash s : \Gamma_2$, and $E_1 \sim_T E_2$, and $m_1 \stackrel{\gamma_2}{\sim}_{\gamma_1} m_2$ then there exists m'_1 and m'_2 such that $E_1 : m_1 \xrightarrow{s} m'_1$, $E_2 : m_2 \xrightarrow{s} m'_2$, $\gamma'_1 \in \Gamma_1$, $\gamma'_2 \in \Gamma_2$, and $m'_1 \stackrel{\gamma'_2}{\sim}_{\gamma'_1} m'_2$.*

Lemma 2 states that starting from two indistinguishable states wrt. $\gamma_1 \sqcup \gamma_2$, a well-typed program results in two indistinguishable states wrt. *some* final state types in Γ_1 and Γ_2 that can also type the resulting states m'_1 and m'_2 .

We rely on Theorem 1 to establish noninterference, that is, if every two states m_1 and m_2 that are indistinguishable wrt. *any* two final state types are also indistinguishable by the output policy, then the program is noninterfering:

Theorem 1 (Noninterference). *Given input policy case γ_i and output policy Γ_o , if $T, pc, \gamma_i \vdash s : \Gamma$ and for every $\gamma_a, \gamma_b \in \Gamma$, such that $m_1 \stackrel{\gamma_b}{\sim}_{\gamma_a} m_2$ it holds also that $m_1 \stackrel{\gamma_o}{\sim}_{\gamma_a} m_2$ for all $\gamma_o \in \Gamma_o$, then s is noninterfering wrt. the input policy case γ_i and the output policy Γ_o .*

Theorem 1 is required to be proved for *every* possible pair of states. To make the verification process feasible, we rely on the following lemma to show that this condition can be verified by simply verifying a relation between the final state types (Γ) and the output policy (Γ_o):

Lemma 3 (Sufficient Condition). *Assume for every $\gamma_1, \gamma_2 \in \Gamma$ and every $\gamma_o \in \Gamma_o$ such that $\gamma_1 \cap \gamma_o \neq \bullet$ that*

- (1) $\gamma_2 \cap \gamma_o \neq \bullet$ implies $\gamma_1 \sqcup \gamma_2 \sqsubseteq \gamma_o$, and
- (2) for every $lval$ either $\gamma_2(lval) \subseteq \gamma_o(lval)$ or $\gamma_1 \sqcup \gamma_2(lval) = L$.

Then for every $\gamma_1, \gamma_2 \in \Gamma$ such that $m_1 \stackrel{\gamma_2}{\sim} m_2$, and every $\gamma_o \in \Gamma_o$ such that $\gamma_o \vdash m_1$ also $\gamma_o \vdash m_2$ and moreover $m_1 \stackrel{\gamma_o}{\sim} m_2$.

In the statement of Lemma 3 we use $\gamma_2(lval) \subseteq \gamma_o(lval)$ to indicate that the interval of $lval$ in γ_2 is included in the interval specified in γ_o .

Intuitively, Lemma 3 formalizes that the least upper bound of any pair in the set of final state types (Γ) should not be more restrictive than the output policy (e.g. if H information has flown to a variable, that variable should also be H in the output policy cases) and the abstractions specified in the output policy cases (i.e. the intervals) are either always satisfied or do not depend on H variables.

D. Revisiting the basic congestion program

We revisit Program 1 to illustrate some key aspects of our typing rules. Here, we only consider the first policy case of the input policy (1) of Section II, where the input packet is IPv4 and it is coming from the internal network.

For the initial state derived from this input policy case, since (1) the parser’s transitions depend on L variables, (2) the type system does not merge state types, and (3) the type system prunes the unreachable transition to accept from `parse_ethernet`, then the parser terminates in a single state type where both `hdr.eth` and `hdr.ipv4` are valid, and their respective headers include the slices, intervals, and labels defined by the initial state type.

After the parsing stage is finished, the program’s control flow reaches the `MyCtrl` control block. Since `hdr.ipv4` is valid and `pc` is L , pruning empty states allows us to ignore the `else` branch on line 62. Afterwards, the nested `if` statement at line 54 entails two possible scenarios. First scenario, when the destination address is in range `192.168.*.*`, as described in Example 5, the two state types resulting from the `if` at line 56 have `hdr.ipv4.ecn` set to H . As in Example 8, these state types satisfy only the first condition of the table’s contract, which results in assigning the type $\langle 1, 9 \rangle_9^L$ to `egress_spec` and producing the state types γ_{int}^1 and γ_{int}^2 .

Second scenario, when the destination address on line 54 does not match `192.168.*.*`, the state type is refined for the `else` branch, producing one state type under condition `ipv4.dstAddr \geq 192.169.0.0` and one under `ipv4.dstAddr $<$ 192.168.0.0`. For both of these state types, `hdr.ipv4.ecn` is set to $\langle 0 \rangle_2^L$ by assignment on line 59. Since in this case all branch conditions were L , there is no H field left in the headers. The first of these two refined state types only satisfies the first condition of the table contract, resulting in one single (after pruning empty states) final state type, γ_{int}^3 , where the packet has been forwarded to the internal

network and `egress_spec` is set to $\langle 1, 9 \rangle_9^L$. The second refined state type however satisfies all the conditions of the table contract, resulting in three final state types γ_{int}^4 , γ_{ext}^1 , γ_{drop}^1 with `egress_spec` being set to $\langle 1, 9 \rangle_9^L$, $\langle 10, 20 \rangle_9^L$, and $\langle 0 \rangle_9^L$, respectively. Notice that among these states, only γ_{int}^3 and γ_{int}^4 contains a H fields (i.e. `ipv4.dstAddr`) due to the first argument returned by the table being $\langle * \rangle_{48}^H$.

We finally check the sufficient condition for the output policy (2) and its only output policy case γ_o , which states that when `egress_spec` is $\langle 10, 20 \rangle_9^L$ (i.e. the packet leaves the internal network) all header fields are L . Only state type γ_{ext}^1 matches the output policy case (i.e. $\cap \gamma_o \neq \bullet$), and this state type satisfies $\gamma_{ext}^1 \sqcup \gamma_{ext}^1 \sqsubseteq \gamma_o$ since all header fields and `egress_spec` are L in γ_{ext}^1 . All other state types (i.e. γ_{int}^1 , γ_{int}^2 , γ_{int}^3 , γ_{int}^4 , and γ_{drop}^1) do not match the output policy condition (i.e. $\cap \gamma_o = \bullet$), since they do not correspond to packets sent to the external network (i.e. their `egress_spec` is not in range $\langle 10, 20 \rangle$). Therefore, we conclude that for this specific input policy case, Program 1 is non-interfering wrt. the output policy case γ_o .

Our analysis can also detect bugs. Assume a bug on line 54 of Program 1. To illustrate this, assume that the program is buggy and instead of checking the 8 most significant bits (i.e. [31:24]) of the `hdr.ipv4.dstAddr`, it checks the least significant bits (i.e. [7:0]). This means that IPv4 packets with destination address is in range `*.168.*.192` would satisfy the condition of the `if` statement on line 54. Similar to the non-buggy program, the `if` at line 56 would produce two state types with `hdr.ipv4.ecn` set to H . These state types satisfy all the conditions of the table contract. For presentation purposes, let us focus on only one of these state types. Applying the table on line 61 would produce three final state types γ_{int}^1 , γ_{ext}^1 , γ_{drop}^1 with `egress_spec` being set to $\langle 1, 9 \rangle_9^L$, $\langle 10, 20 \rangle_9^L$, and $\langle 0 \rangle_9^L$, respectively. Note that in all these final state types, `hdr.ipv4.ecn` is H . When checking the sufficient condition, state type γ_{ext}^1 matches the output policy case (i.e. $\cap \gamma_o \neq \bullet$) but it does not satisfy $\gamma_{ext}^1 \sqcup \gamma_{ext}^1 \sqsubseteq \gamma_o$, because the `hdr.ipv4.ecn` field H in γ_{ext}^1 and L in γ_o . Hence, this buggy program will be marked as interfering, highlighting the fact that some of the packets destined for the external network contain congestion information and unintentionally leak sensitive information.

The benefit of value- and path-sensitivity of our approach can also be demonstrated here. For all other input policy cases that describe non-IPv4 packets, the `parse_ipv4` state is not going to be visited. A path-insensitive analysis, which merges the results of the parser transitions, would lose the information about the validity of the `hdr.ip4` header. This would then lead to the rejection of the program as insecure because an execution where the `parse_ipv4` state has not been visited, yet the `if` branch on line 53 has been taken, will be considered feasible.

Our analysis, on the other hand, identifies that any execution that has not visited `parse_ipv4` results in an invalid `hdr.ip4` header. Consequently, for all such executions, it produces a final state type where the packet is dropped, and `egress_spec`

is set to $(0)_9^L$. This state type satisfies the sufficient condition, since egress_spec does not intersect $\gamma_o(\text{egress_spec})$ and is L .

VII. IMPLEMENTATION AND EVALUATION

To evaluate our approach we developed TAP4S [10], a prototype tool which implements the security type system of Section VI. TAP4S is developed in Python and uses the lark parser library [16] to parse P4 programs.

TAP4S takes as input a P4 program, an input policy, and an output policy. Initially, it parses the P4 program, generates an AST, and relies on this AST and the input policy γ_i to determine the initial type of input packet fields and the standard metadata. Because the input policy is data-dependent, the result of this step can generate multiple state types $(\gamma_1, \dots, \gamma_m)$, one state type for each input interval. TAP4S uses each of these state types as input for implementing the type inference on the program. During this process TAP4S occasionally interacts with a user-defined contract file to retrieve the contracts of the tables and externs. Finally, TAP4S yields a set of final state types $(\gamma'_1, \dots, \gamma'_m)$ which are checked against an output policy, following the condition in Lemma 3. If this check is successful the program is deemed secure wrt. the output policy, otherwise the program is rejected as insecure.

Test suite. To validate our implementation we rely on a functional test suite of 25 programs. These programs are P4 code snippets designed to validate the support for specific functionalities of our implementation, such as extern calls, refinement, and table application.

Use cases. We evaluate TAP4S on 5 use cases, representing different real-world scenarios. The results of this evaluation are summarized in Table I. Due to space constraints, detailed descriptions of these use cases are provided in Appendix A. We also implement and evaluate the use cases from P4BID [7]. These use cases are described in Appendix B, and their corresponding evaluation results are included in Table I. They serve as a baseline for comparing the feasibility of TAP4S with P4BID. On average, P4BID takes 30 ms to analyze these programs, whereas TAP4S takes 246 ms. Despite the increased time, this demonstrates that TAP4S performs the analysis with an acceptable overhead. On the other hand, due to the data-dependent nature of our use cases and their reliance on P4-specific features such as slicing and externs, P4BID cannot reliably check these scenarios, leading to their outright rejection in all cases.

VIII. RELATED WORK

IFC for P4. Our work draws inspiration from P4BID [7], which adapts and implements a security type system [17] for P4, ensuring that well-typed programs satisfy noninterference. By contrast, we show that security policies are inherently data-dependent, thus motivating the need for combining security types with interval-based abstractions. This is essential enforcing IFC in real-world P4 programs without code modifications, as demonstrated by our 5 use cases. Moreover, our analysis

TABLE I: Evaluation results

	Time (ms)			
	Total	Typing	Security Check	Number of Final γ s
Basic Congestion	5930	966	4794	97
Basic Tunneling	610	157	290	15
Multicast	199	16	23	6
Firewall	4560	1015	3378	44
MRI	7646	523	6957	23
Data-plane Routing	274	109	8	12
In-Network Caching	261	91	14	6
Resource Allocation	256	87	10	9
Network Isolation - Alice	243	27	62	3
Network Isolation - Top	242	23	63	3
Topology	202	40	4	3

handles P4 features such as slicing and externs, while supporting the different stages of the P4 pipeline, beyond a single control block of the match-action stage.

IFC policy enforcement. Initial attempts at enforcing data-dependent policies [18]–[22] used dynamic information flow control. The programmer declaratively specifies data-dependent policies and delegates the enforcement to a security-enhanced runtime, thus separating the policy specification from the code implementation.

Our approach shares similarities with static enforcement of data-dependent IFC policies such as *dependent information flow types* and *refinement information flow types*. Dependent information flow types [23] rely on dependent type theory and propose a dependent security type system, in which the security level of a type may depend on its runtime value. Eichholz et al. [24] introduced a dependent type system for the P4 language, called $\Pi 4$, which ensures properties such as preventing the forwarding of expired packets and invalid header accesses. Value-dependent security labels [25] partition the security levels by indexing their labels with values, resulting in partitions that classify data at a specific level, depending on the value. Dependent information flow types provide a natural way to express data-centric policies where the security level of a data structure’s field may depend on values stored in other fields.

Later approaches have focussed on trade-offs between automation and decidability of the analysis. Liquid types [26], [27] are an expressive yet decidable refinement type system [28] to statically express and enforce data-dependent information flow polices. LIFTY [29] provides tool support for specifying data-dependent policies and uses Haskell’s liquid type-checker [27] to verify and repair the program against these policies. STORM [30] is a web framework that relies on liquid types to build MVC web applications that can statically enforce data-dependent policies on databases using liquid types.

Our interval-based security types can be seen as instantiations of refinement types and dependent types. Our simple interval analysis appears to precisely capture the key ingredients of P4 programs, while avoiding challenges with

more expressive analysis. By contrast, the compositionality of analysis based on refinement and dependent types can result in precision loss and is too restrictive for our intended purposes (as shown in Section VI-D), due to merging types of different execution paths. We solve this challenges by proposing a global path-sensitive analysis that avoids merging abstract state in conditionals. We show that our simple yet tractable abstraction is sufficient to enforce the data-dependent policies while precisely modeling P4-specific constructs such as slicing, extract, and emit.

Other works use abstract interpretation in combination with IFC. De Francesco and Martini [31] implement information-flow analysis for stack-based languages like Java. They analyze the instructions an intermediate language by using abstract interpretation to abstractly execute a program on a domain of security levels. Their method is flow-sensitive but not path-sensitive. Cortesi and Halder [32] study information leakage in databases interacting with Hibernate Query Language (HQL). Their method uses a symbolic domain of positive propositional formulae that encodes the variable dependencies of database attributes to check information leaks. Amtoft and Banerjee formulate termination-insensitive information-flow analysis by combining abstract interpretation and Hoare logic [33]. They also show how this logic can be extended to form a security type system that is used to encode noninterference. This work was later extended to handle object-oriented languages in [34].

Analysis and verification of network properties. Existing works on network analysis and verification do not focus on information flow properties. Symbolic execution is widely used for P4 program debugging, enabling tools to explore execution paths, find bugs, and generate test cases. Vera [35] uses symbolic execution to explore all possible execution paths in a P4 program, using symbolic input packets and table entries. Vera catches bugs such as accesses fields of invalid headers and checking that the `egress_spec` is zero for dropped packets. Additionally, it allows users to specify policies, such as; ensuring that the NAT table translates packets before reaching the output ports, and the NAT drops all packets if its entries are empty. Recently, Scaver [36] uses symbolic execution to verify forwarding properties of P4 programs. To address the path explosion problem, they propose multiple pruning strategies to reduce the number of explored paths. ASSERT-P4 [37] combines symbolic execution with assertion checking to find bugs in P4 programs, for example, that the packets with TTL value of zero are not dropped and catching invalid fields accesses. Tools like P4Testgen [38] and p4pktgen [39] use symbolic execution to automatically generate test packets. This approach supports test-driven development and guarantees the correct handling of packets by synthesizing table entries for thorough testing of P4 programs.

Abstract interpretation has also been used to verify functional properties such as packet reachability and isolation. While these properties ensure that packets reach their intended destinations, they do not address the flow of information within

the network. Alpernas et al. [40] introduce an abstract interpretation algorithm for networks with stateful middleboxes (such as firewalls and load balancers). Their method abstracts the order and cardinality of packet on channels, and the correlation between middleboxes states, allowing for efficient and sound analysis. Beckett et al. [41] develop ShapeShifter, which uses abstract interpretation to abstract routing algebras to verify reachability in distributed network control-planes, including objects such as path vectors and IP addresses and methods such as path lengths, regular expressions, intervals, and ternary abstractions.

IX. CONCLUSION

This paper introduced a novel type system that combines security types with interval analysis to ensure noninterference in P4 programs. Our approach effectively prevents information leakages and security violations by statically analyzing data-dependent flows in the data-plane. The type system is both expressive and precise, minimizing overapproximation while simplifying policy specification for developers. Additionally, our type system successfully abstracts complex elements like match-action blocks, tables, and external functions, providing a robust framework for practical security verification in programmable networks. Our implementation, TAP4S, demonstrated the applicability of the security type system on real-world P4 use cases without losing precision due to overapproximations. Future research includes adding support for declassification, advanced functionalities such as cryptographic constructs, and extending the type system to account for side channels.

ACKNOWLEDGMENT

Thanks are due to the anonymous reviewers for their insightful comments and feedback. This work was partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation, the KTH Digital Futures research program, and the Swedish Research Council (VR).

REFERENCES

- [1] P. Goransson, C. Black, and T. Culver, *Software defined networks: a comprehensive approach*. Morgan Kaufmann, 2016.
- [2] J. Matias, J. Garay, N. Toledo, J. Unzilla, and E. Jacob, "Toward an SDN-enabled NFV architecture," *IEEE Communications Magazine*, vol. 53, no. 4, pp. 187–193, 2015.
- [3] P. Bosshart, D. Daly, G. Gibb, M. Izzard, N. McKeown, J. Rexford, C. Schlesinger, D. Talayco, A. Vahdat, G. Varghese *et al.*, "P4: Programming protocol-independent packet processors," *ACM SIGCOMM Computer Communication Review*, vol. 44, no. 3, pp. 87–95, 2014.
- [4] K. D. Albab, J. DiLorenzo, S. Heule, A. Kheradmand, S. Smolka, K. Weitz, M. Timarzi, J. Gao, and M. Yu, "SwitchV: automated SDN switch validation with P4 models," in *Proceedings of the ACM SIGCOMM 2022 Conference*, 2022, pp. 365–379.
- [5] J. C. C. Chica, J. C. Imbachi, and J. F. B. Vega, "Security in SDN: A comprehensive survey," *Journal of Network and Computer Applications*, vol. 159, p. 102595, 2020.
- [6] S. Zander, G. Armitage, and P. Branch, "Covert channels in the IP time to live field," in *Australian Telecommunication Networks and Application Conference (ATNAC) 2006*, 2006.

- [7] K. Grewal, L. D’Antoni, and J. Hsu, “P4BID: information flow control in P4,” in *Proceedings of the 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation*, 2022, pp. 46–60.
- [8] G. Le Guernic, A. Banerjee, T. Jensen, and D. A. Schmidt, “Automata-based confidentiality monitoring,” in *Annual Asian Computing Science Conference*, 2006, pp. 75–89.
- [9] M. Balliu, M. Dam, and G. Le Guernic, “Encover: Symbolic exploration for information flow security,” in *2012 IEEE 25th Computer Security Foundations Symposium*, 2012, pp. 30–44.
- [10] A. Alshnakat, A. M. Ahmadian, M. Balliu, R. Guanciale, and M. Dam, “TAP4S,” 2025, software release. [Online]. Available: <https://github.com/KTH-LangSec/TAP4S>
- [11] A. Sabelfeld and D. Sands, “A per model of secure information flow in sequential programs,” *High. Order Symb. Comput.*, vol. 14, no. 1, pp. 59–91, 2001.
- [12] T. H. Austin and C. Flanagan, “Efficient purely-dynamic information flow analysis,” in *Proceedings of the ACM SIGPLAN Fourth Workshop on Programming Languages and Analysis for Security*, 2009, pp. 113–124.
- [13] N. Vachharajani, M. Bridges, J. Chang, R. Rangan, G. Ottoni, J. Blome, G. Reis, M. Vachharajani, and D. August, “Rifle: An architectural framework for user-centric information-flow security,” in *37th International Symposium on Microarchitecture (MICRO-37’04)*, 2004, pp. 243–254.
- [14] M. Balliu, D. Schoepe, and A. Sabelfeld, “We are family: Relating information-flow trackers,” in *Computer Security - ESORICS 2017 - 22nd European Symposium on Research in Computer Security*, Oslo, Norway, September 11-15, 2017, *Proceedings, Part I*, ser. Lecture Notes in Computer Science, vol. 10492, 2017, pp. 124–145.
- [15] S. Hunt and D. Sands, “On flow-sensitive security types,” in *Proceedings of the 33rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2006, Charleston, South Carolina, USA, January 11-13, 2006*, 2006, pp. 79–90.
- [16] “Lark parser.” [Online]. Available: <https://github.com/lark-parser/lark>
- [17] D. Volpano, C. Irvine, and G. Smith, “A sound type system for secure flow analysis,” *J. Comput. Secur.*, vol. 4, no. 2–3, pp. 167–187, Jan. 1996.
- [18] D. B. Giffin, A. Levy, D. Stefan, D. Terei, D. Mazieres, J. C. Mitchell, and A. Russo, “Hails: Protecting data privacy in untrusted web applications,” in *10th USENIX Symposium on Operating Systems Design and Implementation (OSDI 12)*, 2012, pp. 47–60.
- [19] D. Stefan, D. Mazières, J. C. Mitchell, and A. Russo, “Flexible dynamic information flow control in the presence of exceptions,” *Journal of Functional Programming*, vol. 27, p. e5, 2017.
- [20] J. Yang, K. Yessenov, and A. Solar-Lezama, “A language for automatically enforcing privacy policies,” *ACM SIGPLAN Notices*, vol. 47, no. 1, pp. 85–96, 2012.
- [21] M. Guarnieri, M. Balliu, D. Schoepe, D. A. Basin, and A. Sabelfeld, “Information-flow control for database-backed applications,” in *IEEE European Symposium on Security and Privacy, EuroS&P 2019*. IEEE, 2019, pp. 79–94.
- [22] J. Parker, N. Vazou, and M. Hicks, “LWeb: Information flow security for multi-tier web applications,” *Proceedings of the ACM on Programming Languages*, vol. 3, no. POPL, pp. 1–30, 2019.
- [23] L. Lourenço and L. Caires, “Dependent information flow types,” in *Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, 2015, pp. 317–328.
- [24] M. Eichholz, E. H. Campbell, M. Krebs, N. Foster, and M. Mezini, “Dependently-typed data plane programming,” *Proceedings of the ACM on Programming Languages*, vol. 6, no. POPL, pp. 1–28, 2022.
- [25] L. Lourenço and L. Caires, “Information flow analysis for valued-indexed data security compartments,” in *International Symposium on Trustworthy Global Computing*, 2013, pp. 180–198.
- [26] N. Vazou, P. M. Rondon, and R. Jhala, “Abstract refinement types,” in *European Symposium on Programming*, 2013, pp. 209–228.
- [27] N. Vazou, E. L. Seidel, R. Jhala, D. Vytiniotis, and S. Peyton-Jones, “Refinement types for Haskell,” in *Proceedings of the 19th ACM SIGPLAN international conference on Functional programming*, 2014, pp. 269–282.
- [28] R. Jhala, N. Vazou *et al.*, “Refinement types: A tutorial,” *Foundations and Trends® in Programming Languages*, vol. 6, no. 3–4, pp. 159–317, 2021.
- [29] N. Polikarpova, D. Stefan, J. Yang, S. Itzhaky, T. Hance, and A. Solar-Lezama, “Liquid information flow control,” *Proceedings of the ACM on Programming Languages*, vol. 4, no. ICFP, pp. 1–30, 2020.
- [30] N. Lehmann, R. Kunkel, J. Brown, J. Yang, N. Vazou, N. Polikarpova, D. Stefan, and R. Jhala, “STORM: Refinement types for secure web applications,” in *15th USENIX Symposium on Operating Systems Design and Implementation OSDI 21*, 2021, pp. 441–459.
- [31] N. De Francesco and L. Martini, “Instruction-level security analysis for information flow in stack-based assembly languages,” *Information and Computation*, vol. 205, no. 9, pp. 1334–1370, 2007.
- [32] A. Cortesi and R. Halder, “Information-flow analysis of hibernate query language,” in *Future Data and Security Engineering: First International Conference, FDSE 2014, Ho Chi Minh City, Vietnam, November 19-21, 2014, Proceedings*, 2014, pp. 262–274.
- [33] T. Amtoft and A. Banerjee, “Information flow analysis in logical form,” in *International Static Analysis Symposium*, 2004, pp. 100–115.
- [34] T. A. S. B. A. Banerjee, “A logic for information flow in object-oriented programs.”
- [35] R. Stoescu, D. Dumitrescu, M. Popovici, L. Negreanu, and C. Raiciu, “Debugging P4 programs with vera,” in *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication*, 2018, pp. 518–532.
- [36] Y. Yao, Z. Cui, L. Tian, M. Li, F. Pan, and Y. Hu, “Scaver: A scalable verification system for programmable network,” in *Proceedings of the 2024 SIGCOMM Workshop on Formal Methods Aided Network Operation*, 2024, pp. 14–19.
- [37] L. Freire, M. Neves, L. Leal, K. Levchenko, A. Schaeffer-Filho, and M. Barcellos, “Uncovering bugs in P4 programs with assertion-based verification,” in *Proceedings of the Symposium on SDN Research*, 2018, pp. 1–7.
- [38] F. Ruffly, J. Liu, P. Kotikalapudi, V. Havel, H. Tavante, R. Sherwood, V. Dubina, V. Peschanenko, A. Sivaraman, and N. Foster, “P4Testgen: An extensible test oracle for P4-16,” in *Proceedings of the ACM SIGCOMM 2023 Conference*, 2023, pp. 136–151.
- [39] A. Nötzli, J. Khan, A. Fingerhut, C. Barrett, and P. Athanas, “P4pktgen: Automated test case generation for P4 programs,” in *Proceedings of the Symposium on SDN Research*, 2018, pp. 1–7.
- [40] K. Alpernas, R. Manevich, A. Panda, M. Sagiv, S. Shenker, S. Shoham, and Y. Velner, “Abstract interpretation of stateful networks,” in *Static Analysis: 25th International Symposium, SAS 2018, Freiburg, Germany, August 29–31, 2018, Proceedings 25*, 2018, pp. 86–106.
- [41] R. Beckett, A. Gupta, R. Mahajan, and D. Walker, “Abstract interpretation of distributed network control planes,” *Proceedings of the ACM on Programming Languages*, vol. 4, no. POPL, pp. 1–27, 2019.
- [42] K. Subramanian, A. Abhashkumar, L. D’Antoni, and A. Akella, “D2r: Policy-compliant fast reroute,” in *Proceedings of the ACM SIGCOMM Symposium on SDN Research (SOSR)*, 2021, pp. 148–161.

APPENDIX A USE CASES

A. Basic Tunneling

Our first use case, shown in Program 2, outlines procedures for handling standard IPv4 packets and encapsulated tunneling packets. The parser `MyParser` starts by extracting the Ethernet header on line 6, and for `etherType 0x1212` (tunneled packet), it transitions to `parse_myTunnel` state (line 14), extracts the tunnel header, checks the `proto_id` field, and transitions to `parse_ipv4` state (line 21) if an IPv4 packet is indicated. For `etherType 0x0800` (IPv4 packet), it directly transitions to `parse_ipv4` and extracts the IPv4 header. Once the headers are parsed, the pipeline proceeds to the `MyCtrl` control block, starting from the `apply` block on line 42 which contains two if statements: If only the IPv4 header is valid (line 43), the `ipv4_lpm` table is applied which forward or drop the packet based on a longest prefix match (lpm) on the destination IPv4 address. If the tunnel header is valid (line 46), the `myTunnel_exact` table forwards the packet based on

an exact match of the `myTunnel` header's `dst_id`, using the `myTunnel_forward` action.

Since the header fields of the tunneled packets are not modified while they are forwarded (Lines 34-36), to keep the source MAC address of the packets within the internal networks private, the program should not forward tunneled packets to an external network. The input policy in this use case indicates that if the input packet the packet is tunneled (i.e., its `etherType` is `0x1212`) then the packet's `srcAddr` is `H`. The output policy relies on the output port the packet is sent to, and ensures that "if the `egress_spec` is between 10-511 then the packet has left the internal network, therefore all field of the packet's headers should be `L`."

The general behavior of table `ipv4_lpm` is reflected in its contract as it makes sure the packets with `ipv4` destination address `198.*.*.*` are forwarded to the ports connected to the internal network, while all the other destination addresses are forwarded to ports connected to the external network.

The contract of table `myTunnel_exact` plays a crucial role in the security of Program 2. A correct behavior for this table only forwards the tunneled packets to ports connected to the internal network. The evaluation reported in Table I is performed under a contract that reflected this behavior, which results in TAP4S accepting the program as secure. If this table is somehow misconfigured and forwards the tunneled packet to any port connected to the external network, TAP4S can capture this and flag the program as insecure.

B. Multicast

Our next use case is Program 3 which is capable of multicasting packets to a group of ports. Upon receiving a packet, the switch looks up its destination MAC address `dstAddr`, if it is destined to any of the hosts connected to the switch, the packet is forwarded to its destination (line 11), otherwise the switch broadcasts the packet on ports belonging to a multicast group by setting the `standard_metadata.mcast_grp` to 1 (line 9). Fig. 6 illustrates the network schema of this scenario.

To implement this functionality, the program utilizes the table `mac_lookup` which is populated by the control-plane, and contains the mac addresses and the port information needed to forward non-multicast packets.

While the broadcast packets are sent to all of the multicast ports, it is desirable to ensure the packets that are not supposed to be broadcast are indeed not broadcasted. Our input security policy in this scenario sets the packets destined to any of the hosts connected to the switch as `H`, while labeling the broadcast packets `L`. The contract of the table `mac_lookup`'s needs to capture the essence of this use case, that is, packets with the `dstAddr` of any of the hosts need to be forwarded by invoking the `mac_forward` action (line 11), and all the other packets need to be broadcast by invoking the `multicast` action (line 8). To ensure the program behaves desirably, the output policy checks that all the packets send to the multicast ports (which have their `mcast_grp` set to 1 according to line 9) are `L`.

As illustrated in Table I, under these policies and contracts, Program 3 is secure. It results in 6 final state types (γ), and takes approximately 220 milliseconds to verify the program is policy compliance.

C. Firewall

This use case models a scenario where the switch is running the firewall Program 4 which allows it to monitor the connections between an internal and an external network. The network schema of this scenario is presented in Fig. 7.

After parsing an input packet, the switch applies the `ipv4_lpm` table (line 26), which based on the packet's IPv4 destination address forwards or drops the packet. Next, it applies the `check_ports` table, which based on the input port number (`ingress_port`) identifies whether the packet is coming from the external or the internal network. As depicted in Fig. 7 port 4 is connected to the external network and ports 1–3 are connected to the hosts of the internal network, therefor if the standard metadata's `ingress_port` was 4, the `check_ports` table sets the direction to 1 which indicates the packet is coming from the external network.

The policy of the firewall is that the hosts in the internal network are allowed to communicate with the outside networks, but the hosts in the external network are only allowed to `ssh` to the internal hosts. To this end, for all the packets with direction 1, the program will drop all the packets whose `tcp` port (`hdr.tcp.srcPort`) is not 22 (line 31).

To enforce such policy, we rely on integrity labels (instead of confidentiality) to designate which packets are allowed (trusted), and which packets are not. The input security policy in this scenario sets the packets coming the internal network, identified by their `ingress_port` as trusted (`L`), while any packet coming from the external network (with `ingress_port` 4) is untrusted (`H`) except when its TCP source port `srcPort` is 22.

The contract of the `ipv4_lpm` captures the behavior of this table by making sure the packets with `ipv4` destination address `198.*.*.*` are forwarded to the ports connected to the internal network (ports 1 to 3), and all the other destination addresses are forwarded to port 4. The contract of table `check_ports` updates the direction by checking the `ingress_port` of the incoming packet, setting the direction to 1 if the `ingress_port` was 4.

The output policy checks that all the packets leaving the switch are trusted and `L`. Table I depicts the results of TAP4S for this use case. Under these policies and contracts Program 3 is deemed secure. TAP4S produces 44 final state types, and takes approximately 6 seconds to verify the security of the program.

D. Multi-Hop Route Inspection

Program 5 implements a simplified version of In-Band Network Telemetry, called Multi-Hop Route Inspection (MRI). The purpose of MRI is to let the users to track the path and the length of queues that every packet travels through. To do this, the P4 program adds an ID and queue length to the header

```

1 struct headers {
2   ethernet_t eth; myTunnel_t myTunnel; ipv4_t ipv4;
3 }
4 parser MyParser(/* omitted */) {
5   state start { transition parse_ethernet; }
6   state parse_ethernet {
7     packet.extract(hdr.eth);
8     transition select(hdr.eth.etherType) {
9       0x1212: parse_myTunnel;
10      0x0800: parse_ipv4;
11      default: accept;
12    }
13  }
14  state parse_myTunnel {
15    packet.extract(hdr.myTunnel);
16    transition select(hdr.myTunnel.proto_id) {
17      0x0800: parse_ipv4;
18      default: accept;
19    }
20  }
21  state parse_ipv4 {
22    packet.extract(hdr.ipv4);
23    transition accept;
24  }
25 }

```

```

26 control MyCtrl(/* omitted */) {
27   action drop()
28     { /* drops the packet */ }
29   action ipv4_forward(bit<48> dstAddr, bit<9> port)
30     { /* basic forward */ }
31   table ipv4_lpm
32     { /* omitted */ }
33
34   action myTunnel_forward(bit<19> port) {
35     standard_metadata.egress_spec = port;
36   }
37   table myTunnel_exact {
38     key = {hdr.myTunnel.dst_id: exact;}
39     actions = { myTunnel_forward; drop;}
40     default_action = drop();
41   }
42   apply {
43     if (hdr.ipv4.isValid() && !hdr.myTunnel.isValid()) {
44       ipv4_lpm.apply(); // Process non-tunneled packets
45     }
46     if (hdr.myTunnel.isValid()) {
47       myTunnel_exact.apply(); // Process tunneled packets
48     }
49   }
50 }

```

Program 2: Basic tunneling

```

1 struct headers {
2   ethernet_t ethernet;
3 }
4 control MyCtrl(/* omitted */) {
5   action drop() {
6     mark_to_drop(standard_metadata);
7   }
8   action multicast() {
9     standard_metadata.mcast_grp = 1;
10  }
11  action mac_forward(bit<9> port) {
12    standard_metadata.egress_spec = port;
13  }
14  table mac_lookup {
15    key = { hdr.ethernet.dstAddr : exact; }
16    actions = { multicast; mac_forward; drop; }
17  }
18  apply {
19    if (hdr.ethernet.isValid())
20      mac_lookup.apply();
21  }
22 }

```

Program 3: Multicast

stack of every packet (line 10). Upon reaching the destination, the sequence of switch IDs shows the path the packet took, and each ID is followed by the queue length at that switch.

After parsing the packet, the program applies the `ipv4_lpm` table to forward the packet based on its IPv4 destination address. Afterwards, in the `MyEgress` control block, the `swtrace` table (line 27), based on the port information specified in the `egress_spec`, decides whether to add the queue length data to `swtraces` header or not.

While it makes sense to add this information for packets that are traveling within a local network, similar to the basic

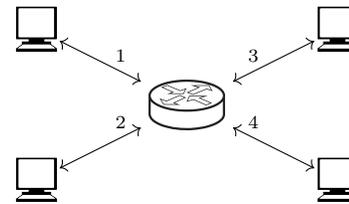


Fig. 6: Multicast schema

congestion example (Program 1) the id of the switches and their queue length can give an external adversary information about the state of the local network. Therefore it is desirable to protect the local network by making sure that Program 5 only adds this data to the packets being forwarded within the local network.

The input policy of this scenario labels the input packet as `L` and only marks the `deq_qdepth` of the standard metadata as `H`. The contract of the `ipv4_lpm` table forwards the packets with `ipv4` destination address `198.*.*.*` to the ports connected to the internal network, while all the other destination addresses are forwarded to ports connected to the external network. If the `egress_spec` indicates ports connected to the internal network, the contract of the `swtrace` table invokes the `add_swtrace` action, adding the queue length data to `swtraces` header, otherwise `NoAction` takes place.

The output policy ensures that in all of the packets going to the external network, identified by their `hdr.ipv4.dstAddr` being anything other than `198.*.*.*`, have `L` `switch_t` header.

Table I depicts the results of type checking this program with TAP4S. It generates 23 final state types and takes approximately 4 seconds to verify the security of the program.

```

1 header tcp_t{
2   bit<16> srcPort;
3   // omitted
4 }
5 struct headers {
6   ethernet_t ethernet; ipv4_t ipv4; tcp_t tcp;
7 }
8 control MyCtrl(/* omitted */) {
9   action drop()
10    { /* drops the packet */ }
11   action ipv4_forward(bit<48> dstAddr, bit<9> port)
12    { /* basic forward */ }
13   table ipv4_lpm
14    { /* omitted */}
15
16   action set_direction(bit<1> dir) {
17     direction = dir;
18   }
19   table check_ports {
20     key = { standard_metadata.ingress_port: exact; }
21     actions = { set_direction; NoAction; }
22   }
23
24   apply {
25     if (hdr.ipv4.isValid()) {
26       ipv4_lpm.apply();
27     }
28     if (hdr.tcp.isValid()) {
29       check_ports.apply();
30       // only allow ssh connections from outside
31       if (direction == 1) { // Packet is from outside
32         if (hdr.tcp.srcPort != 22) { drop(); }
33       }
34     }
35   }
36 }

```

Program 4: Firewall

```

1 header switch_t {
2   switchID_t swid;
3   qdepth_t qdepth;
4 }
5 struct headers {
6   ethernet_t ethernet;
7   ipv4_t ipv4;
8   ipv4_option_t ipv4_option;
9   mri_t mri;
10  switch_t[9] swtraces;
11 }
12 control MyIngress(/* omitted */) {
13   action drop()
14    { /* drops the packet */ }
15   action ipv4_forward(bit<48> dstAddr, bit<9> port)
16    { /* basic forward */ }
17   table ipv4_lpm
18    { /* omitted */}
19   apply {
20     if (hdr.ipv4.isValid()) { ipv4_lpm.apply(); }
21   }
22 }
23 control MyEgress(/* omitted */) {
24   action add_swtrace(switchID_t swid) {
25     // updates the swtraces header
26   }
27   table swtrace {
28     key = { standard_metadata.egress_spec: exact; }
29     actions = { add_swtrace; NoAction; }
30   }
31   apply {
32     if (hdr.mri.isValid()) {
33       swtrace.apply();
34     }
35   }
36 }

```

Program 5: Multi-Hop Route Inspection

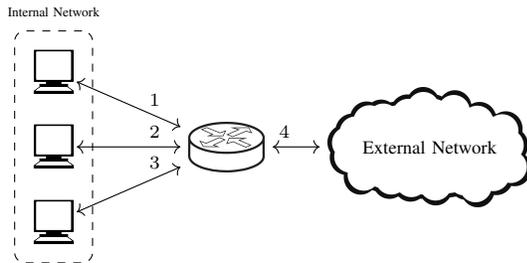


Fig. 7: Firewall schema

swtrace table is crucial for the security of this Program and if it is misconfigured and calls the `add_swtrace` action on outgoing packets, the program will be rejected by TAP4S.

APPENDIX B P4BID USE CASES

We implemented the use cases of P4BID [7] in TAP4S to ensure that it can correctly evaluate all of their use cases and that its verdict is inline with the results reported in [7]. These use cases and their corresponding policies are simpler than our own use cases because P4BID does not support data-dependent policies and hence only labels program variables

without taking the value of the packet header fields into account. The results of this evaluation is depicted in Table I.

Dataplane Routing Routing is the process of determining how to send a packet from its source to its destination. In traditional networks the control-plane is responsible for routing, but recently, Subramanian et al. [42] proposed an approach to implement the routing in the data plane. Their approach uses pre-loaded information about the network topology and link failures to perform a breadth-first search (BFS) and find a path to the destination.

In this scenario we do not care about the details of this BFS search algorithm, but we want to make sure that the sensitive information about the private network (such as the number of hops in the network) do not leak to an external network. Similar to P4BID [7] labeling the number of hops as H will result in the program being rejected by TAP4S because the forwarding action uses this information to update the packet's priority field, which results in an indirect leakage of sensitive information.

In-Network Caching In order to enable the fast retrieval of popular items, switches keep track of the frequently requested items in a cache and only query the controller when an item cannot be found in the cache. Similar to any cache system, the

result of a query is the same regardless of where the item is stored. However, from a security perspective, an observer can potentially detect variations in item retrieval time. This timing side-channel can potentially allow an adversary to learn about the state of the system.

To model the cache in this scenario, we mark the request query as sensitive, because whether this query is a *hit* or a *miss* leaks information about the internal state of the switch. The variable marking the state of the result in the cache (`response.hit`) is not sensitive because it is considered observable by the adversary. Similar to P4BID [7] this labeling will result in the program being rejected by TAP4S because a sensitive query can indirectly affect the value of `response.hit`, resulting in the leakage of sensitive information.

Resource Allocation This use case models a simple resource allocation program, where the switch increases the priority of the packets belonging to latency-sensitive applications. The application ID in the packet’s header will indicate which application the packet belongs to. A table will match on this application ID and sets the packet’s priority by modifying the priority field of the `ipv4` header.

The problem is that a malicious client can manipulate the application ID to increase the priority of their packets. We rely on integrity labels (instead of confidentiality) to address this issue, that is, the application ID will be labeled as untrusted (H) and the `ipv4` priority field will be labeled as trusted (L). Since the program sets the priority field based on the value of the application ID, the priority field will also be labeled untrusted by the type system, which results in the rejection of the program.

Network Isolation This use case models a private network used by two clients, Alice and Bob. Each client runs its own P4 program, but the packets sent between these two clients have a shared header with separate fields for Alice and Bob. In this scenario we want to make sure that Alice does not touch Bob’s fields, and vice versa.

The isolation property in this example can be modeled by a four-point lattice with labels $\{A, B, \top, \perp\}$, where A is the label of Alice’s data, B is for Bob’s data, \top is the top element confidential to both Alice and Bob, and \perp is public. By IFC, data from level ℓ can flow to ℓ' if and only if $\ell \sqsubseteq \ell'$.

In this use case, we consider Alice’s program in which she updates the fields belonging to herself. Additionally, we use label \top to label the telemetry data which can be updated by Alice’s program, but she cannot leak information from \top -labeled data into her own fields.

Since TAP4S only support simple lattice with two levels, we type check this program twice, with two policies. First where Alice is H and everything else is L , and a second time where \top is H and everything else is L . The same process can be repeated for Bob’s program as well. This program is accepted by TAP4S, and the results for both cases are reported in Table I.

Topology This use case is a P4 program which processes packets as they enters a local network. The incoming packets

refers to a virtual address which needs to be translated to a physical address as the packet is routed in the local network.

Our security policy dictates that the routing details of this local network should not leak into fields that are visible when the packet leaves the network. As such, the program relies on a separate header to store the local information, and as long as the packet is inside the local network, the switches do not modify the `ipv4` and Ethernet headers, instead, they parse, use, and update this local header with the routing information.

As explained in P4BID [7], this program has a bug where it incorrectly stores the local `ttl` in the `ipv4` header instead of the local header. Marking the local fields as H , TAP4S flags this program as insecure and facilitates the process of catching and fixing these types of errors.

APPENDIX C STATE TYPE OPERATIONS

A type τ' is considered an overapproximation of type τ (written as $\tau \leq \tau'$) iff for every value $v : \tau$ it holds that $v : \tau'$ and $\text{lbl}(\tau) \sqsubseteq \text{lbl}(\tau')$. We denote two non-overlapping lvalues by $\text{lval} \ast \text{lval}'$, which means if lval is a record lval' is not one of its fields, and if lval is a bitvector lval' is not one of its sub-slices.

We present the properties that operators over the state types must guarantee:

- $\gamma[\text{lval} \mapsto \tau] = \gamma'$ indicates updating the type of lval , which can be a part of a variable, in state type γ . This operator guarantees that $\gamma' \vdash \text{lval} : \tau$ and for every lvalue $\text{lval}' \ast \text{lval}$ such that $\gamma \vdash \text{lval}' : \tau'$ and $\gamma' \vdash \text{lval}' : \tau''$ then $\tau' \leq \tau''$.
- $\gamma \uparrow \gamma'$ updates γ such that for every variable in the domain of γ' , the type of that variable in γ is updated to match γ' .
- $\text{refine}(\gamma, e) = \gamma'$ returns an overapproximation of states that satisfy the abstraction of γ and the predicate e . It guarantees that if $\gamma \vdash m$ and e evaluates to *true* in m then $\gamma' \vdash m$ and for every lval , if $\gamma \vdash \text{lval} : \tau$ and $\gamma' \vdash \text{lval} : \tau'$ then $\text{lbl}(\tau) \sqsubseteq \text{lbl}(\tau')$
- $\text{join}(\gamma_1, \gamma_2) = \gamma_3$ returns an overapproximation of γ_1 , whose labels are at least as restrictive as γ_1 and γ_2 . This operator guarantees that if $\gamma_1 \vdash m$ then $\gamma_3 \vdash m$ and for every lval , then $\text{lbl}(\gamma_1(\text{lval})) \sqcup \text{lbl}(\gamma_2(\text{lval})) \sqsubseteq \text{lbl}(\gamma_3(\text{lval}))$.

APPENDIX D PROOFS AND GUARANTEES

We use $T \vdash E$ to represent the abstraction and labeling soundness guarantees for externs and tables, and in the following we assume that this condition holds.

A. Sufficient condition proof

Lemma 3 (Sufficient Condition). *Assume for every $\gamma_1, \gamma_2 \in \Gamma$ and every $\gamma_o \in \Gamma_o$ such that $\gamma_1 \cap \gamma_o \neq \bullet$ that*

- (1) $\gamma_2 \cap \gamma_o \neq \bullet$ implies $\gamma_1 \sqcup \gamma_2 \sqsubseteq \gamma_o$, and
- (2) for every lval either $\gamma_2(\text{lval}) \subseteq \gamma_o(\text{lval})$ or $\gamma_1 \sqcup \gamma_2(\text{lval}) = L$.

Then for every $\gamma_1, \gamma_2 \in \Gamma$ such that $m_1 \stackrel{\gamma_2}{\sim} m_2$, and every $\gamma_o \in \Gamma_o$ such that $\gamma_o \vdash m_1$ also $\gamma_o \vdash m_2$ and moreover $m_1 \stackrel{\gamma_o}{\sim} m_2$.

Proof. First, we prove $\gamma_o \vdash m_2$. By definition, it is sufficient to prove that for every $lval$, $m_2(lval) : \gamma_o(lval)$.

From the definition of $m_1 \stackrel{\gamma_1}{\sim} m_2$ we know that $\gamma_1 \vdash m_1$ and $\gamma_2 \vdash m_2$ hold. Since $\gamma_o \vdash m_1$ and $\gamma_1 \vdash m_1$, then trivially $\gamma_1 \cap \gamma_o \neq \bullet$ holds. From the second hypothesis of the sufficient condition, two cases are possible.

- 1) $\gamma_2(lval) \subseteq \gamma_o(lval)$. Since $\gamma_2 \vdash m_2$ hold, then by definition $m_2(lval) : \gamma_2(lval)$ holds, therefore trivially $m_2(lval) : \gamma_o(lval)$ holds.
- 2) $\gamma_1 \sqcup \gamma_2(lval) = L$. Since $m_1 \stackrel{\gamma_1}{\sim} m_2$ then, indeed $m_1(lval) = m_2(lval)$ holds. By definition of $\gamma_o \vdash m_1$, we know that $m_1(lval) : \gamma_o(lval)$ also holds. Therefore, we can trivially show that $m_2(lval) : \gamma_o(lval)$.

Second, we prove $m_1 \stackrel{\gamma_o}{\sim} m_2$. Previously, we showed that $\gamma_o \vdash m_2$ holds, and since $\gamma_2 \vdash m_2$, then trivially $\gamma_2 \cap \gamma_o \neq \bullet$ holds. From the first hypothesis of the sufficient condition, we can show that $\gamma_1 \sqcup \gamma_2 \subseteq \gamma_o$. By definition of $m_1 \stackrel{\gamma_1}{\sim} m_2$, we know that $m_1 \stackrel{\gamma_1 \sqcup \gamma_2}{\sim} m_2$ holds. Therefore, trivially $m_1 \stackrel{\gamma_o}{\sim} m_2$. \square

B. Hypothesis for refine

Hyp 1 (Interval typedness - boolean expressions' refinement).

$$\begin{aligned} \gamma \vdash e : \tau \wedge \gamma \vdash m &\implies \\ (m(e) = true &\implies \\ (\text{refine}(\gamma, e) \vdash m \wedge m(e) = false &\implies \\ (\text{refine}(\gamma, \neg e) \vdash m) & \end{aligned}$$

Hyp 2 (Interval typedness - select expressions' refinement).

$$\begin{aligned} \gamma \vdash e : \tau \wedge m(e) = v \wedge \gamma \vdash m \\ \wedge i = \min\{i. v = v_i \vee i = n + 1\} \implies \\ \gamma_1, \dots, \gamma_n, \gamma_{n+1} = (\text{refine}(\gamma, e = v_i)) \implies \\ \gamma_i \vdash m \end{aligned}$$

Hyp 3 (Interval typedness - externs and tables refinement).

$$\gamma \vdash m \wedge \phi(m) \implies \text{refine}(\gamma, \phi) \vdash m$$

Hyp 4 (Label typedness - boolean expressions' refinement).

$$\begin{aligned} \gamma \vdash e : \tau_1 \wedge \gamma \vdash e : \tau_2 \\ \wedge \text{lbl}(\tau_1) = L \wedge \text{lbl}(\tau_2) = L \\ \wedge m_1 \stackrel{\gamma_1 \sqcup \gamma_2}{\sim} m_2 \implies \\ \left((\text{refine}(\gamma_1, e) = \gamma'_1 \wedge \text{refine}(\gamma_2, e) = \gamma'_2 \right. \\ \wedge m_1(e) = true \wedge m_2(e) = true \implies m_1 \stackrel{\gamma_1' \sqcup \gamma_2'}{\sim} m_2) \\ \wedge (\text{refine}(\gamma_1, \neg e) = \gamma'_1 \wedge \text{refine}(\gamma_2, \neg e) = \gamma'_2 \\ \wedge m_1(\neg e) = true \wedge m_2(\neg e) = true \implies \\ \left. m_1 \stackrel{\gamma_1' \sqcup \gamma_2'}{\sim} m_2) \right) \end{aligned}$$

Hyp 5 (Label typedness - select expressions' refinement).

$$\begin{aligned} \gamma \vdash e : \tau \wedge \text{lbl}(\tau) = L \wedge m_1 \stackrel{\gamma}{\sim} m_2 \implies \\ \left(\gamma_1, \dots, \gamma_n, \gamma_{n+1} = \text{refine}(\gamma, e = v_i) \wedge m_1(e) = v \implies \right. \\ \left. \forall j \leq n + 1. m_1 \stackrel{\gamma_j}{\sim} m_2 \right) \end{aligned}$$

C. Lemmas

Lemma 4 (Expression reduction preserves the type).

$$\gamma \vdash m \wedge \gamma \vdash e : \tau \wedge m(e) = v \implies v : \tau$$

Lemma 5 (lvalue updates preserves the type).

$$\gamma \vdash m \wedge v : \tau \implies \gamma[lval \mapsto \tau] \vdash m[lval \mapsto v]$$

Lemma 6 (Expressions have types same as their values).

$$\gamma \vdash m \wedge \gamma \vdash e : \tau \implies \exists v. m(e) = v \wedge v : \tau$$

Lemma 7 (Join does not modify the intervals).

$$\begin{aligned} \gamma \vdash m \wedge \gamma \in \Gamma_1 \implies \\ \forall \Gamma_2 \dots \Gamma_n. \exists \gamma' \in \text{join}(\Gamma_1 \cup \Gamma_2 \cup \dots \cup \Gamma_n). \\ \gamma' \vdash m \wedge \gamma \sqsubseteq \gamma' \end{aligned}$$

Lemma 8 (Expression evaluation of consistent states).

$$\begin{aligned} m_1 \stackrel{\gamma_1 \sqcup \gamma_2}{\sim} m_2 \wedge \gamma_1 \vdash m_1 \wedge \gamma_2 \vdash m_2 \wedge \\ \gamma_1 \vdash e : \tau_1 \wedge \gamma_2 \vdash e : \tau_2 \wedge \\ \text{lbl}(\tau_1) = L \wedge \text{lbl}(\tau_2) = L \implies \\ (m_1(e) = v \wedge m_2(e) = v) \end{aligned}$$

Lemma 9 (State equivalence preservation).

$$\gamma' \sqsubseteq \gamma \wedge m_1 \stackrel{\gamma'}{\sim} m_2 \implies m_1 \stackrel{\gamma}{\sim} m_2$$

Lemma 10 (Branch on high - state preservation).

$$\begin{aligned} E : m \xrightarrow{s} m' \wedge T, H, \gamma \vdash s : \Gamma \wedge \gamma' \in \Gamma \implies \\ (\gamma'(lval) = \tau \wedge \text{lbl}(\tau) = L) \implies \\ m(lval) = m'(lval) \end{aligned}$$

Lemma 11 (Join's low label implication).

$$\begin{aligned} \Gamma &= \text{join}(\Gamma_1 \cup \dots \cup \Gamma_n) \wedge \gamma \in \Gamma \wedge \\ \gamma(\text{lval}) &= \tau \wedge \text{lbl}(\tau) = \mathbf{L} \implies \\ \forall \gamma' \in \Gamma_1, \dots, \Gamma_n. \gamma'(\text{lval}) &= \tau' \wedge \text{lbl}(\tau') = \mathbf{L} \end{aligned}$$

Lemma 12 (High program's final types).

$$T, \mathbf{H}, \gamma \vdash s : \Gamma \implies \forall \gamma' \in \Gamma. \gamma \sqsubseteq \gamma'$$

Lemma 13 (Branch on high - state lemma).

$$\begin{aligned} T, \mathbf{H}, \gamma \vdash s : \Gamma \wedge \gamma' \in \Gamma \wedge \\ \gamma'(\text{lval}) = \tau' \wedge \text{lbl}(\tau') = \mathbf{L} \implies \\ \forall \gamma'' \gamma'''. T, pc, \gamma'' \vdash s : \Gamma' \wedge \gamma''' \in \Gamma' \implies \\ \left(\gamma''(\text{lval}) = \tau'' \wedge \gamma'''(\text{lval}) = \tau''' \right. \\ \left. \implies \text{lbl}(\tau'') \sqsubseteq \text{lbl}(\tau''') \right) \end{aligned}$$

Lemma 14 (Branch on high is never empty).

$$T, \mathbf{H}, \gamma \vdash s : \Gamma \implies \Gamma \neq \emptyset$$

Lemma 15 (Low equivalence distribution).

$$\begin{aligned} (m_1, m'_1)_{(\gamma_a, \gamma'_a) \sqcup (\gamma_b, \gamma'_b)} \sim (m_2, m'_2) \Leftrightarrow \\ (m_1 \sim_{\gamma_a \sqcup \gamma_b} m_2 \wedge m'_1 \sim_{\gamma'_a \sqcup \gamma'_b} m'_2) \end{aligned}$$

Lemma 16 (Low equivalence update).

$$\begin{aligned} m_1 \sim_{\gamma_a \sqcup \gamma_b} m_2 \wedge m_3 \sim_{\gamma_c \sqcup \gamma_d} m_4 \implies \\ m_1[\text{lval} \mapsto m_3(e)] \sim_{\gamma_a[\text{lval} \mapsto \gamma_c(e)] \sqcup \gamma_b[\text{lval} \mapsto \gamma_d(e)]} m_2[\text{lval} \mapsto m_4(e)] \end{aligned}$$

D. Soundness of abstraction

Theorem 2.

$$\begin{aligned} \forall s T m \gamma pc \Gamma. T, pc, \gamma \vdash s : \Gamma \implies \\ T \vdash E \wedge \gamma \vdash m \implies \\ \exists m'. E : m \xrightarrow{s} m' \wedge \exists \gamma' \in \Gamma. \gamma' \vdash m' \end{aligned}$$

Proof. In this proof we assume that the program is well typed and does not get stuck according to HOL4P4 type system. By induction on the typing tree of the program $stmt$, i.e. s . Note that, initially $\gamma = (\gamma_g, \gamma_m)$ and $m = (m_g, m_l)$. In the following proof, we know that $\gamma \vdash m$ holds in all subcases.

◇ **Case assignment:** Here $stmt$ is $\text{lval} := e$. From assignment typing rule we know:

- 1) $\gamma \vdash e : \tau$
- 2) $\tau' = \text{raise}(\tau, pc)$
- 3) $\gamma' = \gamma[\text{lval} \mapsto \tau']$
- 4) $\Gamma = \{\gamma'\}$

We need to prove $\exists m'. E : m \xrightarrow{\text{lval} := e} m'$ and $\exists \gamma'' \in \Gamma. \gamma'' \vdash m'$. From the assignment reduction definition, we can rewrite the goal's conjunctions to:

- 1) $m(e) = v$ (from assignment reduction)
- 2) $\exists m'. m' = m[\text{lval} \mapsto v]$ (from assignment reduction)
- 3) $\exists \gamma'' \in \{\gamma'\}. \gamma'' \vdash m'$

Goal 1: we know that the initial variable map is typed using the state type, i.e. $\gamma \vdash m$ from assumptions. Using $\gamma \vdash m$ and 1 from typing rule, then we can use Lemma 6 to directly infer that the expression's reduction indeed keeps the type, i.e. $\exists v. m(e) = v \wedge v : \tau$ and this resolves goal 1.

Goal 2: trivial, as we can instantiate m' to be $m[\text{lval} \mapsto v]$.

Goal 3: we know that assignment typing rule produces a singleton set from 4 of typing rule, thus we can instantiate γ'' to be $\gamma[\text{lval} \mapsto \tau']$, thus allows us to rewrite the goal to $\gamma[\text{lval} \mapsto \tau'] \vdash m[\text{lval} \mapsto v]$.

Given $\gamma \vdash m$, this entails (by definition) that γ and m they contain the same variable name in the domain, and also the values are well-typed, i.e. $\text{domain}(\gamma) = \text{domain}(m) \wedge \forall x \in \text{domain}(m). m(x) : \gamma(x)$.

Using Lemma 4, we know that the assigned value v can be typed with τ , i.e. $v : \tau$. The assignment typing rule raises the labels using the function raise , however by definition we know that the abstraction is unaffected, so the interval of τ' and τ are the same, thus we can infer that $v : \tau'$.

Using Lemma 5, we can see that the update preserves the type, thus goal proved.

◇ **Case condition:** Here $stmt$ is if e then s_1 else s_2 . From conditional typing rule we know:

- 1) $\gamma \vdash e : \tau$
- 2) $\ell = \text{lbl}(\tau)$
- 3) $pc' = pc \sqcup \ell$
- 4) $T, pc', (\text{refine}(\gamma, e)) \vdash s_1 : \Gamma_1$
- 5) $T, pc', (\text{refine}(\gamma, \neg e)) \vdash s_2 : \Gamma_2$
- 6) $\Gamma' = \Gamma_1 \cup \Gamma_2$
- 7) $\Gamma'' = \begin{cases} \text{join}(\Gamma') & \text{if } \ell = \mathbf{H} \\ \Gamma' & \text{otherwise} \end{cases}$

We need to prove $\exists m'. E : m \xrightarrow{\text{if } e \text{ then } s_1 \text{ else } s_2} m'$ and $\exists \gamma'' \in \Gamma''. \gamma'' \vdash m'$.

In the goal, the boolean guard e evaluates to *true* or *false*. So we will get two goal cases with similar proofs. This only solves for e , while case $\neg e$ follows the same proof strategy. We can rewrite the goal to:

- 1) $\exists m'. E : m \xrightarrow{s_1} m'$
- 2) $\exists \gamma'' \in \Gamma''. \gamma'' \vdash m'$

In this proof, we will get two induction hypotheses for statement: for s_1 call it IH1 and for s_2 call it IH2.

IH1 is the following (note that IH2 is the same, but instantiated for s_2):

$$\begin{aligned} \forall T m \gamma pc \Gamma. T, pc, \gamma \vdash s_1 : \Gamma \implies \\ T \vdash E \wedge \gamma \vdash m \implies \\ \exists m'. E : m \xrightarrow{s_1} m' \wedge \exists \gamma' \in \Gamma. \gamma' \vdash m' \end{aligned}$$

We first prove that the set of refined state types using e can still type the staring concrete memory m . This we can show, because we know initially $\gamma \vdash m$, and we know that e is typed as a boolean (assumed to be well-typed), and we know that e reduces to *true* from the reduction rule, these allow us to infer $(\text{refine}(\gamma, e)) \vdash m$ using Hyp 1.

Now we instantiate the induction hypothesis IH1 using the following $(T, m, \text{refine}(\gamma, e), pc', \Gamma_1)$, to show that exists m' such that $E : m \xrightarrow{s_1} m'$, i.e. there indeed exists a transition to a final configuration in the semantics to m' (which resolves goal 1). Additionally, we can show from IH1 that exists $\gamma' \in \Gamma_1$ such that $\gamma' \vdash m'$.

Now we implement cases on the expression's label ℓ being H or L :

case $\text{lbl}(\tau) = H$: we need to prove $\exists \gamma'' \in \text{join} \{\Gamma_1 \cup \Gamma_2\}. \gamma'' \vdash m'$. Then it is easy to deduct that the goal holds; because we showed that there is a state type $\gamma' \in \Gamma_1$ such that it is a sound abstraction of the final state $\gamma' \vdash m'$, and we know that join operation does not change the abstraction, it just modifies the security label. Hence, indeed there exists a state type γ'' in $\text{join} \{\Gamma_1 \cup \Gamma_2\}$ such that it is also a sound abstraction of final state $\gamma'' \vdash m'$.

case $\text{lbl}(\tau) = L$: we need to prove $\exists \gamma'' \in \{\Gamma_1 \cup \Gamma_2\}. \gamma'' \vdash m'$, which is trivially true.

For the negation case, use IH2 and follow the same steps.

◇ **Case sequence**: Here stmt is $s_1; s_2$. From sequence typing rule we know:

- 1) $T, pc, \gamma \vdash s_1 : \Gamma_1$
- 2) $\forall \gamma_1 \in \Gamma_1. T, pc, \gamma_1 \vdash s_2 : \Gamma_2^{\gamma_1}$
- 3) $\Gamma' = \bigcup_{\gamma_1 \in \Gamma_1} \Gamma_2^{\gamma_1}$

In this proof, we will get two induction hypotheses for statement: for s_1 call it IH1 and for s_2 call it IH2.

IH1 is (note that IH2 is the same, but instantiated for s_2):

$$\begin{aligned} \forall T m \gamma pc \Gamma. T, pc, \gamma \vdash s_1 : \Gamma &\implies \\ T \vdash E \wedge \gamma \vdash m &\implies \\ \exists m'. E : m \xrightarrow{s_1} m' \wedge \exists \gamma' \in \Gamma. \gamma' \vdash m' & \end{aligned}$$

We need to prove $\exists m''. E : m \xrightarrow{s_1; s_2} m''$ and $\exists \gamma'' \in \Gamma'. \gamma'' \vdash m''$.

We can rewrite the goal using the definition of sequence case to:

- 1) $E : m \xrightarrow{s_1} m'$
- 2) $\exists m''. E : m' \xrightarrow{s_2} m''$
- 3) $\exists \gamma'' \in \Gamma'. \gamma'' \vdash m''$

Goal 1: We can instantiate the IH1 with $(T, m, \gamma, pc, \Gamma_1)$ to infer that exists m'_1 such that $E : m \xrightarrow{s_1} m'_1$, and also exists $\gamma' \in \Gamma_1$ such that $\gamma' \vdash m'_1$. Since the semantics are deterministic, m'_1 and m' are equivalent, thus it holds $E : m \xrightarrow{s_1} m'$ and $\gamma' \vdash m'$.

Goal 2 and 3: We showed from IH1 that $\gamma' \vdash m'$, now we can instantiate 2 from sequence typing rule with γ' . Now we can to instantiate IH2 with $(T, m', \gamma', pc, \Gamma_2^{\gamma'})$, and infer that exists m'_2 such that $E : m' \xrightarrow{s_2} m'_2$, and also exists $\gamma''' \in \Gamma_2^{\gamma'}$ such that $\gamma''' \vdash m'_2$. Since the semantics are deterministic, m'_2 and m'' are equivalent, thus it holds $E : m \xrightarrow{s_2} m''$ and $\gamma''' \vdash m''$.

From 3 in sequence typing rule, we know that Γ' is the union of all resulted state type sets, such that it can type s_2 ,

since we know that $\gamma''' \in \Gamma_2^{\gamma'}$ will be in Γ' , then we prove the goal $\exists \gamma'' \in \Gamma'. \gamma'' \vdash m''$.

◇ **Case function call**: Here stmt is $f(e_1, \dots, e_n)$. From call typing rule we know (note that here we explicitly write the global and local state type):

- 1) $(\gamma_g, \gamma_l) \vdash e_i : \tau_i$
- 2) $(s, (x_1, d_1), \dots, (x_n, d_n)) = (C, F)(f)$
- 3) $\gamma_f = \{x_i \mapsto \tau_i\}$
- 4) $(C, F), pc, (\gamma_g, \gamma_f) \vdash s : \Gamma'$
- 5) $\Gamma'' = \{(\gamma'_g, \gamma'_l)[e_i \mapsto \gamma'_f(x_i) \mid \text{isOut}(d_i)] \mid (\gamma'_g, \gamma'_f) \in \Gamma'\}$

In the following proof, we know that $(\gamma_g, \gamma_l) \vdash (m_g, m_l)$ and $(C, F) \vdash (X, F)$ hold. Additionally, we get an induction hypothesis IH for the body of the function s .

$$\begin{aligned} \forall (C, F) m \gamma pc \Gamma. (C, F), pc, \gamma \vdash s : \Gamma &\implies \\ (C, F) \vdash (X, F) \wedge \gamma \vdash m &\implies \\ \exists m'. (X, F) : m \xrightarrow{s} m' \wedge \exists \gamma' \in \Gamma. \gamma' \vdash m' & \end{aligned}$$

We need to prove $\exists m'. (X, F) : (m_g, m_l) \xrightarrow{f(e_1, \dots, e_n)} m' \wedge \exists \gamma' \in \Gamma''. \gamma' \vdash m'$

From call reduction rule we can rewrite the goal to:

- 1) $\exists s (x_1, d_1), \dots, (x_n, d_n). (s, (x_1, d_1), \dots, (x_n, d_n)) = (X, F)(f)$
- 2) $\exists m_f. m_f = \{x_i \mapsto (m_g, m_l)(e_i)\}$
- 3) $\exists (m'_g, m'_f). (X, F) : (m_g, m_f) \xrightarrow{s} (m'_g, m'_f)$
- 4) $\exists m''. m'' = (m'_g, m_l)[e_i \mapsto m'_f(x_i) \mid \text{isOut}(d_i)]$
- 5) $\exists \gamma' \in \Gamma''. \gamma' \vdash m''$

Goal 1: From $(C, F) \vdash (X, F)$ we know indeed the function's body and signature found in the semantics is the same found in the typing rule.

Goal 2: Trivial, the existence can be instantiated with $\{x_i \mapsto (m_g, m_l)(e_i)\}$.

Goal 3: To prove that there exists a state where the body of the function reduces to, we need to use the IH. Thus, we first need to show that the resulted copy-in map is also well-typed i.e. $\gamma_f \vdash m_f$, more specifically $\forall i. \{x_i \mapsto \tau_i\} \vdash \{x_i \mapsto (m_g, m_l)(e_i)\}$. Given that initially the state type can type the state $(\gamma_g, \gamma_l) \vdash (m_g, m_l)$ from assumptions and given that the expressions e_i have a type τ_i from typing rule 1, now we can use Lemma 4 to infer that $\forall v_i : \tau_i$ such that v_i is the evaluation of $(m_g, m_l)(e_i)$. This leads us to trivially infer that $\forall i. \{x_i \mapsto \tau_i\} \vdash \{x_i \mapsto v_i\}$ holds.

Now we can use the IH by instantiating it to $((C, F), (m_g, m_f), (\gamma_g, \gamma_f), pc, \Gamma')$ in order to infer that exists (m'_g, m'_f) such that $(X, F) : (m_g, m_f) \xrightarrow{s} (m'_g, m'_f)$ and exists $(\gamma''_g, \gamma''_f) \in \Gamma'$ such that $(\gamma''_g, \gamma''_f) \vdash (m'_g, m'_f)$.

Goal 4: Trivial, we can instantiate the existence by $(m'_g, m_l)[e_i \mapsto m'_f(x_i) \mid \text{isOut}(d_i)]$.

Goal 5: The goal is to prove copy-out operation to be well-typed, thus we can rewrite the goal to $\exists \gamma' \in \{(\gamma'_g, \gamma'_l)[e_i \mapsto \gamma'_f(x_i) \mid \text{isOut}(d_i)] \mid (\gamma'_g, \gamma'_f) \in \Gamma'\}$ such that $\gamma' \vdash (m'_g, m_l)[e_i \mapsto m'_f(x_i) \mid \text{isOut}(d_i)]$.

We can choose γ' to be $(\gamma''_g, \gamma'_l)[e_i \mapsto \gamma''_f(x_i) \mid \text{isOut}(d_i)]$.

Since we are able to choose a state type that types the final state, we can rewrite the goal to prove again to be

$(\gamma_g'', \gamma_l)[e_i \mapsto \gamma_f''(x_i) \mid \text{isOut}(d_i)] \vdash (m_g', m_l)[e_i \mapsto m_f'(x_i) \mid \text{isOut}(d_i)]$.

First, we know that $\gamma_l \vdash m_l$ holds from the assumptions. Also, we showed in (Goal 3) that $(\gamma_g'', \gamma_f'') \vdash (m_g', m_f')$, thus trivially $\gamma_g'' \vdash m_g'$ and $\gamma_f'' \vdash m_f'$ hold. Thus, we can deduct that $(\gamma_g'', \gamma_l) \vdash (m_g'', m_l)$, and $m_f'(x_i) : \gamma_f''(x_i)$, and then we can use Lemma 5 to deduct that the update preserves the well-typedness, i.e. $(\gamma_g'', \gamma_l)[e_i \mapsto \gamma_f''(x_i)] \vdash (m_g', m_l)[e_i \mapsto m_f'(x_i)]$, which proves the goal.

◇ **Case extern:** Here *stmt* is $f(e_1, \dots, e_n)$. From extern typing rule we know (note that here we explicitly write the global and local state type):

- 1) $(\gamma_g, \gamma_l) \vdash e_i : \tau_i$
- 2) $(\text{Cont}_E, (x_1, d_1), \dots, (x_n, d_n)) = (C, F)(f)$
- 3) $\gamma_f = \{x_i \mapsto \tau_i\}$
- 4) $\forall (\gamma_i, \phi, \gamma_t) \in \text{Cont}_E. (\gamma_g, \gamma_l) \sqsubseteq \gamma_i$
- 5) $\Gamma' = \{\gamma' \# \text{raise}(\gamma_t, pc) \mid (\gamma_i, \phi, \gamma_t) \in \text{Cont}_E \wedge \text{refine}((\gamma_g, \gamma_f), \phi) = \gamma' \neq \bullet\}$
- 6) $\Gamma'' = \{(\gamma_g', \gamma_l)[e_i \mapsto \gamma_f'(x_i) \mid \text{isOut}(d_i)] \mid (\gamma_g', \gamma_f') \in \Gamma'\}$

In the following proof, we know that $(\gamma_g, \gamma_l) \vdash (m_g, m_l)$ and $(C, F) \vdash (X, F)$ hold.

We need to prove $\exists m'. (X, F) : (m_g, m_l) \xrightarrow{f(e_1, \dots, e_n)} m' \wedge \exists \gamma'' \in \Gamma''. \gamma'' \vdash m'$

From extern reduction rule we can rewrite the goal to:

- 1) $\exists (\text{sem}_f, (x_1, d_1), \dots, (x_n, d_n)). (\text{sem}_f, (x_1, d_1), \dots, (x_n, d_n)) = (X, F)(f)$
- 2) $\exists m_f. m_f = \{x_i \mapsto (m_g, m_l)(e_i)\}$
- 3) $\exists (m_g', m_f'). (m_g', m_f') = \text{sem}_f(m_g, m_f)$
- 4) $\exists m''. m'' = (m_g', m_l)[e_i \mapsto m_f'(x_i) \mid \text{isOut}(d_i)]$
- 5) $\exists \gamma'' \in \Gamma''. \gamma'' \vdash m''$

Goal 1: From the environment's well-typedness $(C, F) \vdash (X, F)$, we know $\text{domain}(C) \cap \text{domain}(F) = \emptyset$ and $\text{extWT } C \ X$ holds. This mean that indeed the extern is defined only in C . Additionally, from well-typedness $(C, F) \vdash (X, F)$, that if $C(f) = (\text{sem}_f, (x, d))$ then $X(f) = (\text{Cont}_E, (x, d))$, thus indeed exist Cont_E and signature $(x_1, d_1), \dots, (x_n, d_n)$.

Goal 2: Trivial, by instantiating m_f to be $\{x_i \mapsto (m_g, m_l)(e_i)\}$.

We can here also prove that the resulted copy-in map is also well-typed i.e. $\gamma_f \vdash m_f$, more specifically $\forall i. \{x_i \mapsto \tau_i\} \vdash \{x_i \mapsto (m_g, m_l)(e_i)\}$. Given that initially the typing state types the state $(\gamma_g, \gamma_l) \vdash (m_g, m_l)$ from assumptions and given that the expressions e_i have a type τ_i from typing rule 1, now we can use Lemma 4 to infer that $\forall v_i : \tau_i$ such that v_i is the evaluation of $(m_g, m_l)(e_i)$. This leads us to trivially infer that $\forall i. \{x_i \mapsto \tau_i\} \vdash \{x_i \mapsto v_i\}$ holds.

Goal 3: From $(C, F) \vdash (X, F)$, we know that $\text{extWT } C \ X$ holds, and from its definition, we know that indeed exists $(\gamma_i, \phi, \gamma_t)$ such that $\phi(m_g, m_f)$, this means that indeed there exists a contract's predicate satisfied by the values in the initial concrete input state, i.e. $\phi(m_g, m_f)$. This implies, from the definition of $\text{extWT } C \ X$, that indeed exists (m_g', m_f') such that $(m_g', m_f') = \text{sem}_f(m_g, m_f)$.

Goal 4: Trivial, by instantiating m'' to $(m_g', m_l)[e_i \mapsto m_f'(x_i) \mid \text{isOut}(d_i)]$.

Goal 5: We can rewrite the goal to exists $\gamma'' \in \{(\gamma_g', \gamma_l)[e_i \mapsto \gamma_f'(x_i) \mid \text{isOut}(d_i)] \mid (\gamma_g', \gamma_f') \in \Gamma'\}$ such that $\gamma'' \vdash (m_g', m_l)[e_i \mapsto m_f'(x_i) \mid \text{isOut}(d_i)]$.

We can prove this goal by first find a $\gamma_A \in \Gamma'$ such that it types the final states of the extern's semantic (m_g', m_f') . Second, we find the $\gamma'' \in \Gamma''$ such that it can type the final state m'' after copying out the extern.

We previously established $\gamma_f \vdash m_f$ (from goal 2), also given that $\gamma_g \vdash m_g$ from assumptions we can trivially $(\gamma_g, \gamma_f) \vdash (m_g, m_f)$. Since we previously showed that $\phi(m_g, m_f)$ (from goal 3), now we can use Hyp 3 in order to infer that the refined state $\text{refine}((\gamma_g, \gamma_f), \phi)$ can also type (m_g, m_f) also we infer that it is not empty, i.e. $\text{refine}((\gamma_g, \gamma_f), \phi) \vdash (m_g, m_f)$ and $\text{refine}((\gamma_g, \gamma_f), \phi) \neq \bullet$.

The definition of $\text{extWT } C \ X$ states that $\gamma_t \vdash (m_g', m_f')$ holds for the set of variables that the extern's semantics has changed i.e. $\{x. (m_g, m_f)(x) \neq (m_g', m_f')(x)\} \subseteq \text{domain}(\gamma_t)$.

Consequently, we can further prove that indeed exists a γ_A in Γ' (from 4 in typing rule of extern) that can type the output of the extern's semantics (m_g', m_f') including the unchanged variables. Thus, we can make cases on extern's semantics input and output as following:

case $(m_g, m_f)(x) = (m_g', m_f')(x)$: This means that variable x not in the domain of γ_t , thus it is unchanged, therefore it is typed by the refined state $(m_g', m_f')(x) : (\text{refine}((\gamma_g, \gamma_f), \phi))(x)$. Trivially, we can also infer that $(m_g', m_f')(x) : (\text{refine}((\gamma_g, \gamma_f), \phi) \# \text{raise}(\gamma_t, pc))(x)$.

case $(m_g, m_f)(x) \neq (m_g', m_f')(x)$: This means that variable x is in the domain of γ_t , thus it is changed, and the new type of it is in γ_t . Therefore, we can trivially conclude that $(m_g, m_f)(x) : \gamma_t(x)$, Consequently, since we know that raise does not change the abstraction, and just change labels, we can therefore conclude that $(m_g', m_f')(x) : (\text{refine}((\gamma_g, \gamma_f), \phi) \# \text{raise}(\gamma_t, pc))(x)$

These cases show that we can select γ_A such that $\gamma_A \in \Gamma'$ to be $(\text{refine}((\gamma_g, \gamma_f), \phi) \# \text{raise}(\gamma_t, pc))$, because it can indeed type (m_g', m_f') . i.e. $(\text{refine}((\gamma_g, \gamma_f), \phi) \# \text{raise}(\gamma_t, pc)) \vdash (m_g', m_f')$. For simplicity in the rest of the proof, let us rewrite $\gamma_A = (\gamma_{A_g}, \gamma_{A_f})$, where γ_{A_g} is the global part of the pair $(\text{refine}((\gamma_g, \gamma_f), \phi) \# \text{raise}(\gamma_t, pc))$ and γ_{A_f} is the local part of the pair. Thus, we can say that $\gamma_{A_g} \vdash m_g'$ and $\gamma_{A_f} \vdash m_f'$.

Now we need to find $\gamma'' \in \Gamma''$ such that it types $(m_g', m_l)[e_i \mapsto m_f'(x_i) \mid \text{isOut}(d_i)]$ order to prove Goal 5.

Given from assumptions $\gamma_l \vdash m_l$, and we showed that $\gamma_{A_g} \vdash m_g'$ and $\gamma_{A_f} \vdash m_f'$. In 5 of the typing rules, we pick (γ_g', γ_f') such that it is in Γ' to be $(\gamma_{A_g}, \gamma_{A_f})$, Now we conduct cases on the direction of the parameter $\text{isOut}(d_i)$ being out or not.

case $\neg \text{isOut}(d_i)$: Then (γ_{A_g}, γ_l) are unchanged, similarly in the semantics (m_g', m_l) are also unchanged. Thus they can still be typed as $(\gamma_{A_g}, \gamma_l) \vdash (m_g', m_l)$

case $\text{isOut}(d_i)$: Then (γ_{A_g}, γ_l) are updated with $e_i \mapsto \gamma_{A_f}(x)$, similarly the semantics state (m_g', m_l) is updated with

$e_i \mapsto m'_f(x)$. Previously we showed that $\gamma_{A_f} \vdash m'_f$, this entails by the definition of state typedness $m'_f(x) : \gamma_{A_f}(x)$. This leads to the point that the modifications of e_i 's type in state type (γ_{A_g}, γ_l) and value in state (m'_g, m_l) keeps them well typed. Therefore, $(\gamma_{A_g}, \gamma_l) \vdash (m'_g, m_l)$ holds.

We can finally conclude that goal 5 can be resolved by picking the γ'' to be (γ_{A_g}, γ_l) , the goal is now proven.

◇ **Case table application:** Here *stmt* is apply *tbl*. From table typing rule we know:

- 1) $(\bar{e}, \text{Cont}_{\text{tbl}}) = (C, F)(\text{tbl})$
- 2) $\gamma \vdash e_i : \tau_i$
- 3) $\ell = \bigsqcup_i \text{lbl}(\tau_i)$
- 4) $pc' = pc \sqcup \ell$
- 5) $\forall (\phi_j, (a_j, \bar{\tau}_j)) \in \text{Cont}_{\text{tbl}}. (\gamma_{g_j}, \gamma_{l_j}) = \text{refine}(\gamma, \phi_j) \wedge (s_j, (x_{j_1}, \text{none}), \dots, (x_{j_n}, \text{none})) = (C, F)(a_j) \wedge \gamma_{a_j} = \{x_{j_i} \mapsto \tau_{j_i}\} \wedge T, pc', (\gamma_{g_j}, \gamma_{a_j}) \vdash s_j : \Gamma_j$
- 6) $\Gamma' = \cup_j \{(\gamma'_{g_j}, \gamma'_{l_j}) \mid (\gamma'_{g_j}, \gamma'_{a_j}) \in \Gamma_j\}$
- 7) $\Gamma'' = \begin{cases} \text{join}(\Gamma') & \text{if } \ell = H \\ \Gamma' & \text{otherwise} \end{cases}$

In the following proof, we know that $(\gamma_g, \gamma_l) \vdash (m_g, m_l)$ and $(C, F) \vdash (X, F)$ hold.

We need to prove $\exists m'. E : m \xrightarrow{\text{apply tbl}} m'$ and $\exists \gamma'' \in \Gamma''. \gamma'' \vdash m'$.

And from table reduction rule we can rewrite the goal to:

- 1) $\exists \bar{e} \text{ sem}_{\text{tbl}}. (\bar{e}, \text{sem}_{\text{tbl}}) = (X, F)(\text{tbl})$
- 2) $\exists (a, \bar{v}). \text{sem}_{\text{tbl}}((m_g, m_l)(e_1), \dots, (m_g, m_l)(e_n)) = (a, \bar{v})$
- 3) $\exists s (x_1, \dots, x_n). (s, (x_1, \text{none}), \dots, (x_n, \text{none})) = E(a)$
- 4) $\exists m_a. m_a = \{x_i \mapsto v_i\}$
- 5) $\exists (m_{g'}, m_{a'}). E : (m_g, m_a) \xrightarrow{s} (m_{g'}, m_{a'})$
- 6) $\exists \gamma'' \in \Gamma''. \gamma'' \vdash (m_{g'}, m_l)$

In this proof, we will get an induction hypothesis for action call $a(\bar{v})$ (formalized as a function call), call it IH.

$$\begin{aligned} \forall (C, F) m \gamma pc \Gamma. (C, F), pc, \gamma \vdash a(\bar{v}) : \Gamma &\implies \\ (C, F) \vdash (X, F) \wedge \gamma \vdash m &\implies \\ \exists m'. (X, F) : m \xrightarrow{a(\bar{v})} m' \wedge \exists \gamma' \in \Gamma. \gamma' \vdash m' & \end{aligned}$$

Goal 1: Trivial, by the well-typedness condition $(C, F) \vdash (X, F)$, we know that if the table has a contract (from 1 in typing rule), then indeed there is semantics for it sem_{tbl} and a key list \bar{e} that matches the one in the table typing rule.

Goal 2: From $(C, F) \vdash (X, F)$, we can deduce from condition $\text{tblWT } C \ X$ that indeed exists an action and value list pair (a, \bar{v}) in the contract Cont_{tbl} correlated to a $\phi_j(m)$ that holds.

Goal 3: From $(C, F) \vdash (X, F)$ we know indeed the actions's a body and signature found in the semantics is the same found in the typing rule.

Goal 4: Trivial, by setting m_a to be $\{x_i \mapsto v_i\}$.

Goal 5: To prove this goal, we need to use IH. And in order to use IH, we must first show that $(\gamma_{g_j}, \gamma_{a_j}) \vdash (m_g, m_a)$ by proving $\gamma_{g_j} \vdash m_g$ where $(\gamma_{g_j}, \gamma_{l_j}) = \text{refine}(\gamma, \phi_j)$ and the copied in is well typed $\gamma_{a_j} \vdash m_a$.

Prove $\gamma_{g_j} \vdash m_g$: Since initially given that $\gamma \vdash m$ i.e. $(\gamma_g, \gamma_l) \vdash (m_g, m_l)$, also we know from $\text{tblWT } C \ X$ that there is j such that the predicate ϕ_j is satisfied in (m_g, m_l) i.e. $\phi_j(m_g, m_l)$, thus we can use Hyp 3 to deduce that $\text{refine}(\gamma, \phi_j) \vdash (m_g, m_l)$, i.e. we can infer that the refined state type is able to type the initial state. Given 5 in the typing rule, we know that $(\gamma_{g_j}, \gamma_{l_j}) = \text{refine}(\gamma, \phi_j)$ thus $(\gamma_{g_j}, \gamma_{l_j}) \vdash (m_g, m_l)$, therefore $\gamma_{g_j} \vdash m_g$ holds.

Prove $\gamma_{a_j} \vdash m_a$: This goal can be rewritten as $\{x_i \mapsto \tau_i\} \vdash \{x_i \mapsto v_i\}$. The proof is trivial by WF definition of tables we know that $v_i : \tau_i$, thus the variable x_i is well typed.

Now, we can instantiate IH to using $((C, F), (m_g, m_a), (\gamma_{g_j}, \gamma_{a_j}), pc', \Gamma_j)$, so we can infer that exists m' such that $(X, F) : m \xrightarrow{s} m'$ (thus goal 5 is resolved).

Goal 6: We can also infer from IH that exists $\gamma' \in \Gamma_j$ such that $\gamma' \vdash m'$. Let $(\gamma'_{g_j}, \gamma'_{a_j}) = \gamma'$ and $(m'_{g'}, m'_{a'}) = m'$, thus indeed $(\gamma'_{g_j}, \gamma'_{a_j}) \vdash (m'_{g'}, m'_{a'})$ holds trivially.

Line 6 of the typing rule iterates over each final state type set and collects the modified global state and the refined local state, thus $(\gamma'_{g_j}, \gamma_{l_j})$ is indeed in Γ' . Since we proved that $\gamma'_{g_j} \vdash m'_{g'}$ holds in the previous step, and also proved $\gamma_{l_j} \vdash m_l$ in goal 5, therefore, $(\gamma'_{g_j}, \gamma_{l_j}) \vdash (m'_{g'}, m_l)$.

Line 7 of the typing rule changes the labels but not the abstraction, thus the abstraction of the state type $(\gamma'_{g_j}, \gamma_{l_j})$ in Γ' indeed exists in Γ'' with labels changed so goal 6 holds. \square

E. Soundness of labeling

Theorem 3.

$$\forall s \ T \ pc \ m_1 \ m_2 \ m'_1 \ m'_2 \ \gamma_1 \ \gamma_2 \ \Gamma_1 \ \Gamma_2 \ E_1 \ E_2.$$

$$T, pc, \gamma_1 \vdash s : \Gamma_1 \wedge T, pc, \gamma_2 \vdash s : \Gamma_2 \implies$$

$$T \vdash E_1 \wedge T \vdash E_2 \wedge E_1 \underset{T}{\sim} E_2 \wedge$$

$$\gamma_1 \vdash m_1 \wedge \gamma_2 \vdash m_2 \wedge m_1 \underset{\gamma_1 \sqcup \gamma_2}{\sim} m_2 \wedge$$

$$E_1 : m_1 \xrightarrow{s} m'_1 \wedge E_2 : m_2 \xrightarrow{s} m'_2 \implies$$

$$(\exists \gamma'_1 \in \Gamma_1 \wedge \gamma'_2 \in \Gamma_2. \gamma'_1 \vdash m'_1$$

$$\wedge \gamma'_2 \vdash m'_2 \wedge m'_1 \underset{\gamma'_1 \sqcup \gamma'_2}{\sim} m'_2)$$

Proof. In this proof we assume that the program is well typed and does not get stuck according to HOL4P4 type system. by induction on the typing tree of the program *stmt* i.e. s . Note that, initially $\gamma = (\gamma_g, \gamma_m)$ and $m = (m_g, m_l)$. In the following proof, we know that $m_1 \underset{\gamma_1 \sqcup \gamma_2}{\sim} m_2$, we also know that $\gamma_1 \vdash m_1$ and $\gamma_2 \vdash m_2$.

◇ **Case assignment:** Here *stmt* is $\text{lval} := e$. From assignment typing rule we know:

- 1) $\gamma_1 \vdash e : \tau_1$
- 2) $\tau'_1 = \text{raise}(\tau_1, pc)$
- 3) $\gamma'_1 = \gamma_1[\text{lval} \mapsto \tau'_1]$
- 4) $\Gamma_1 = \{\gamma'_1\}$
- 5) $\gamma_2 \vdash e : \tau_2$
- 6) $\tau'_2 = \text{raise}(\tau_2, pc)$
- 7) $\gamma'_2 = \gamma_2[\text{lval} \mapsto \tau'_2]$

8) $\Gamma_2 = \{\gamma'_2\}$

And from assignment reduction rule we know:

- 1) $m_1(e) = v_1$
- 2) $m'_1 = m_1[lval \mapsto v_1]$
- 3) $m_2(e) = v_2$
- 4) $m'_2 = m_2[lval \mapsto v_2]$

Prove $\exists \gamma'_1 \in \Gamma_1 \wedge \gamma'_2 \in \Gamma_2. \gamma'_1 \vdash m'_1 \wedge \gamma'_2 \vdash m'_2 \wedge m'_1 \sim_{\gamma'_1 \sqcup \gamma'_2} m'_2$.

Since that assignment typing rule produces one state type (from 3,4,7,and 8), then from SOUNDNESS OF ABSTRACTION, we can infer $\gamma'_1 \vdash m'_1$ and $\gamma'_2 \vdash m'_2$, therefore the first two conjunctions of the goal holds.

Now, the final remaining goal to prove is that $m'_1 \sim_{\gamma'_1 \sqcup \gamma'_2} m'_2$.

This entails proving that all $lval'$ in both m'_1 and m'_2 are low equivalent with respect to the least upper bound of the final state types that types the final states.

Now we do cases on the label of $lval'$ being H or L in $\gamma'_1 \sqcup \gamma'_2$.

case label of $lval'$ is H : If the label is H , i.e. $\gamma'_1 \sqcup \gamma'_2 \vdash lval' : \tau \wedge \text{lbl}(\tau) = H$, then the property of labeling soundness holds after the assignment trivially. That's because soundness property checks the equality of the state type's low ranges only.

case label of $lval'$ is L : If the label is L , i.e., $\gamma'_1 \sqcup \gamma'_2 \vdash lval' : \tau \wedge \text{lbl}(\tau) = L$, this implies that in each state type γ'_1 and γ'_2 individually, the typing label of $lval'$ is L .

In this section, we conduct a case analysis on possible sub-cases relations between of $lval$ being assigned and $lval'$, thus we will have the following subcases: $lval' \subsetneq lval$, $lval \subsetneq lval'$, $lval' = lval$, $lval' \subseteq lval$, and $lval \subseteq lval'$.

case $lval' \subsetneq lval$ and $lval \subsetneq lval'$:

Now for all $lval'$ that are not equal to the $lval$ we assign to, or not sub- $lval$ of it: our goal is to show $m'_1(lval') = m'_2(lval')$ by demonstrating that initially $m_1(lval') = m_2(lval')$ also holds. We know that the assignment doesn't alter those parts of the states. Thus, the semantic update should keep the values of $lval'$ the same, so $m_1(lval') = m'_1(lval')$ and $m_2(lval') = m'_2(lval')$. Likewise, the typing update should keep the type of $lval'$ unchanged, so $\gamma'_1(lval') = \gamma_1(lval')$ and $\gamma'_2(lval') = \gamma_2(lval')$. This indeed mean that the labels of $lval'$ were also L in both γ_1 and γ_2 , and given the assumption that $m_1 \sim_{\gamma_1 \sqcup \gamma_2} m_2$, thus indeed $m_1(lval') = m_2(lval')$.

case $lval' = lval$:

For $lval' = lval$, we know that in $\gamma'_1 \sqcup \gamma'_2$ we type the $lval$ as L , i.e. $\gamma'_1 \sqcup \gamma'_2(lval) = \tau \wedge \text{lbl}(\tau) = L$ thus the same property holds in individual state types $\gamma'_1(lval) = \tau' \wedge \text{lbl}(\tau') = L$ and also $\gamma'_2(lval) = \tau'' \wedge \text{lbl}(\tau'') = L$. We need to prove $m'_1(lval) = m'_2(lval)$. For $lval$ to be L after the update function in either γ'_1 or γ'_2 , it is necessary for the types of the expression e to be L in both initial state types γ_1 in 1 and γ_2 in 5 in typing rules, i.e., $(\text{lbl}(\tau_1) = L \text{ and } \text{lbl}(\tau_2) = L)$.

This condition holds because otherwise, if the typing labels of e were H , then the goal would be trivially true (as the update would make $lval$ H in γ'_1 and γ'_2 , contradicting the assumptions). Since the typing label of e is L in 1 and 5 of typing rule, when reduced to a value in 1 and 3 of the semantics rule, this indicates that they reduce to the same value $v_1 = v_2$ (using Lemma 8). Since we update $lval$ in m_1 and m_2 with the same value such that we produce m'_1 and m'_2 respectively, it follows that $m'_1(lval) = m'_2(lval)$.

case $lval' \subseteq lval$:

For $lval' \subseteq lval$, we know that $lval'$ can be a shorter variation of the $lval$. The proof is the same as the previous case.

case $lval \subseteq lval'$:

For $lval \subseteq lval'$, we know that $lval$ can be a shorter variation of the $lval'$, thus the update of $lval$ affects part of $lval'$ type while the rest of it stays unchanged. Therefore, the proof is straightforward by conducting the same steps of the first two cases.

Given the last four subcases, we can now show that $m'_1 \sim_{\gamma'} m'_2$.

◇ **Case condition:** Here $stmt$ is if e then s_1 else s_2 . From conditional typing rule we know:

- 1) $\gamma \vdash e : \tau_1$
- 2) $\ell_1 = \text{lbl}(\tau_1)$
- 3) $pc_1 = pc \sqcup \ell_1$
- 4) $T, pc_1, (\text{refine}(\gamma_1, e)) \vdash s_1 : \Gamma_1$
- 5) $T, pc_1, (\text{refine}(\gamma_1, \neg e)) \vdash s_2 : \Gamma_2$
- 6) $\Gamma_3 = \begin{cases} \text{join}(\Gamma_1 \cup \Gamma_2) & \text{if } \ell_1 = H \\ \Gamma_1 \cup \Gamma_2 & \text{otherwise} \end{cases}$
- 7) $\gamma \vdash e : \tau_2$
- 8) $\ell_2 = \text{lbl}(\tau_2)$
- 9) $pc_2 = pc \sqcup \ell_2$
- 10) $T, pc_2, (\text{refine}(\gamma_2, e)) \vdash s_1 : \Gamma_4$
- 11) $T, pc_2, (\text{refine}(\gamma_2, \neg e)) \vdash s_2 : \Gamma_5$
- 12) $\Gamma_6 = \begin{cases} \text{join}(\Gamma_4 \cup \Gamma_5) & \text{if } \ell_2 = H \\ \Gamma_4 \cup \Gamma_5 & \text{otherwise} \end{cases}$

We also know that both initial states are $\gamma_1 \vdash m_1$ and $\gamma_2 \vdash m_2$ and also $m_1 \sim_{\gamma_1 \sqcup \gamma_2} m_2$. And we know that the conditional statement is executed with m_1 and m_2 resulting m'_1 and m'_2 consequently.

In addition to that, we get induction hypothesis for s_1 IH1 and s_2 IH2 (we only show IH1):

$$\begin{aligned} & \forall T \ pc \ m_a \ m_b \ m'_a \ m'_b \ \gamma_a \ \gamma_b \ \Gamma_a \ \Gamma_b \ E_a \ E_b. \\ & T, pc, \gamma_a \vdash s_1 : \Gamma_a \ \wedge \ T, pc, \gamma_b \vdash s_1 : \Gamma_b \implies \\ & T \vdash E_a \ \wedge \ T \vdash E_b \ \wedge \ E_a \sim_T E_b \ \wedge \\ & \gamma_a \vdash m_a \ \wedge \ \gamma_b \vdash m_b \ \wedge \ m_a \sim_{\gamma_a \sqcup \gamma_b} m_b \ \wedge \\ & E_a : m_a \xrightarrow{s_1} m'_a \ \wedge \ E_b : m_b \xrightarrow{s_1} m'_b \\ & \implies \end{aligned}$$

$$(\exists \gamma'_a \in \Gamma_a \ \wedge \ \gamma'_b \in \Gamma_b. \ \gamma'_a \vdash m'_a \ \wedge \ \gamma'_b \vdash m'_b \ \wedge \ m'_a \sim_{\gamma'_a \sqcup \gamma'_b} m'_b)$$

We start by cases on labels ℓ_1 and ℓ_2 of e .

case $\ell_1 = \ell_2 = L$: We need to prove that $\exists \gamma'_1 \in \Gamma_3$ and $\exists \gamma'_2 \in \Gamma_6$ it holds $m'_1 \underset{\gamma'_1 \sqcup \gamma'_2}{\sim} m'_2$. We can directly use Lemma 8 to

infer that e is evaluation is indistinguishable in the states, and since e is assumed to be typed as boolean, thus we get two subcases where: $m_1(e) = true$ and $m_2(e) = true$, or $m_1(e) = false$ and $m_2(e) = false$.

case $m_1(e) = true$ and $m_2(e) = true$: when looking into the reduction rule of the both if statements in the assumption, we only reduce the first branch of each. Hence, we have $E_1 : m_1 \xrightarrow{s_1} m'_1$ and $E_2 : m_2 \xrightarrow{s_1} m'_2$. From Hyp 4, we can show that $m_1 \underset{(\text{refine}(\gamma_1, e)) \sqcup (\text{refine}(\gamma_2, e))}{\sim} m_2$. Additionally, we can infer that $(\text{refine}(\gamma_1, e)) \vdash m_1$ and $(\text{refine}(\gamma_2, e)) \vdash m_2$ using Hyp 1.

Now we can directly instantiate and apply IH1 using the following $(T, pc \sqcup L, m_1, m_2, m'_1, m'_2, (\text{refine}(\gamma_1, e)), (\text{refine}(\gamma_2, e)), \Gamma_1, \Gamma_4, E_1, E_2)$ to infer that the states after executing s_1 are low equivalent, i.e. exists $\gamma''_1 \in \Gamma_1$ and exists $\gamma''_2 \in \Gamma_4$ such that $\gamma''_1 \vdash m'_1$ and $\gamma''_2 \vdash m'_2$ and $m'_1 \underset{\gamma''_1 \sqcup \gamma''_2}{\sim} m'_2$. Since $\Gamma_1 \subseteq \Gamma_3$ and $\Gamma_4 \subseteq \Gamma_6$, thus the goal holds.

case $m_1(e) = false$ and $m_2(e) = false$ same proof as the previous case.

case $\ell_2 = H$: We initiate the proof by fixing ℓ_2 to be H , and the value of e to be reduced to $false$, thus it executes s_2 (starting from configuration m_2 , and yields m'_2 , note that if e reduces to $true$ the proof is identical as this case). In this proof, we refer to these as the second configuration. Now, consider the following scenario where we start from m_1 in the semantics rule and γ_1 in typing rule, we generalize the proof for any boolean expression e_i such that i ranges over $true$ and $false$, where e_{true} is e , e_{false} is $\neg e$, s_{true} is the first branch s_1 , and s_{false} is the second branch s_2 . In this sub-case of the proof, ℓ_1 denotes the label associated e_i 's typing label, and let the refinement of the initial typing scope γ_1 to be represented as $\text{refine}(\gamma_1, e_i)$. Suppose the executed branch is s_i , yielding a final set of state types denoted as Γ_i . In this proof, we refer to these as the first configuration.

Given the previous generalizations, we can rewrite the assumptions to:

- (a) $T, \ell_1, (\text{refine}(\gamma_1, e_i)) \vdash s_i : \Gamma_i$
- (b) $E_1 : m_1 \xrightarrow{s_i} m'_1$
- (c) $T, H, (\text{refine}(\gamma_2, \neg e)) \vdash s_2 : \Gamma_5$
- (d) $T, H, (\text{refine}(\gamma_2, e)) \vdash s_1 : \Gamma_4$
- (e) $E_2 : m_2 \xrightarrow{s_2} m'_2$

What we aim to prove is the existence of $\gamma'_1 \in \Gamma_3$ such that $\gamma'_1 \vdash m'_1$. Additionally, we need to establish the existence of $\gamma'_2 \in \Gamma_6$ where $\Gamma_6 = \text{join}(\Gamma_4 \cup \Gamma_5)$, such that $\gamma'_2 \vdash m'_2$. Furthermore, we must prove that $m'_1 \underset{\gamma'_1 \sqcup \gamma'_2}{\sim} m'_2$.

From the SOUNDNESS OF ABSTRACTION, we know that for the second configuration indeed exists $\gamma''_2 \in \Gamma_5$ such that it types m'_2 (i.e. $\gamma''_2 \vdash m'_2$).

Given that $\gamma''_2 \in \Gamma_5$ and $\Gamma_6 = \text{join}(\Gamma_4 \cup \Gamma_5)$, we can deduce (by Lemma 7) the existence of $\overline{\gamma''_2} \in \text{join}(\Gamma_4 \cup \Gamma_5)$ such that it is more restrictive than γ''_2 , denoted as $\gamma''_2 \sqsubseteq \overline{\gamma''_2}$. Using the same lemma, we conclude that $\overline{\gamma''_2} \vdash m'_2$. Now on, we choose $\overline{\gamma''_2}$ to be used in the proof and resolve the second conjunction of the goal.

From the SOUNDNESS OF ABSTRACTION, we know that for the first configuration indeed exists $\gamma''_i \in \Gamma_i$ such that it types m'_1 (i.e. $\gamma''_i \vdash m'_1$).

In the first configuration, we generalized the proof according to the evaluation of e_i . Consequently, the final state type set Γ_3 can be either a union (if $\ell_1 = L$) or a join (if $\ell_1 = H$) of all final state type sets $\forall i \leq 1$. Γ_i resulting from typing their corresponding s_i . In either case (union or join), we can establish the existence of $\overline{\gamma''_i} \in \Gamma_3$ such that $\gamma''_i \sqsubseteq \overline{\gamma''_i}$ and indeed $\overline{\gamma''_i} \vdash m'_1$. Note that if Γ_3 resulted from a join, we infer this using Lemma 7; otherwise, if it resulted from a union, it is trivially true. In fact, we can directly choose $\gamma''_i = \overline{\gamma''_i}$ when union the final state types sets.

Next, we proceed to implement cases based on whether an $lval$'s type label is H or L .

case $(\overline{\gamma''_i} \sqcup \overline{\gamma''_2}(lval) = \tau) \wedge \text{lbl}(\tau) = H$: the goal holds trivially.

case $(\overline{\gamma''_i} \sqcup \overline{\gamma''_2}(lval) = \tau) \wedge \text{lbl}(\tau) = L$: this case entails that each state type individually holds $\overline{\gamma''_i}(lval) = \tau'_1 \wedge \text{lbl}(\tau'_1) = L$ and also $\overline{\gamma''_2}(lval) = \tau''_2 \wedge \text{lbl}(\tau''_2) = L$.

Given that the $lval$'s type is L in $\overline{\gamma''_2}$, and considering $\overline{\gamma''_2} \in \text{join}(\Gamma_4 \cup \Gamma_5)$, it follows that the $lval$ is also L in any state type within $\text{join}(\Gamma_4 \cup \Gamma_5)$. Consequently, the $lval$ is L in the state types of both Γ_5 and Γ_4 (if not empty) individually before the join operation. Hence, since $\gamma''_2 \in \Gamma_5$, it implies that $lval$ is also L in γ''_2 expressed as $\gamma''_2(lval) = \tau'''_2 \wedge \text{lbl}(\tau'''_2) = L$ (using Lemma 11). We also know that typing a statement in a H context in (d) entails that the final Γ_4 is not empty (using Lemma 14), thus $lval$'s type label is also L in all the state types in Γ_4 .

In the second configuration, we type the statement s_2 with a H context in (c), and s_2 reduces to m'_2 in (e). Furthermore, from the previous step, we inferred that the $lval$'s type is L in γ''_2 . Hence, we can use Lemma 10 to infer that the initial and final states remain unchanged for L lvalues, which means $m_2(lval) = m'_2(lval)$. Then we can use Lemma 12 to infer that $\text{refine}(\gamma_2, \neg e) \sqsubseteq \gamma''_2$, this entails that the $lval$'s type label is indeed L in the refined state $\text{refine}(\gamma_2, \neg e)$. It is easy to see that $\gamma_2 \sqsubseteq \text{refine}(\gamma_2, \neg e)$, thus $lval$'s type is also L in the initial state type γ_2 , i.e. $\gamma_2(lval) = \tau'_2 \wedge \text{lbl}(\tau'_2) = L$.

For the first configuration, in this sub-case, we have $\overline{\gamma''_i}(lval) = \tau'_1 \wedge \text{lbl}(\tau'_1) = L$, and $\overline{\gamma''_i} \in \Gamma_3$. We previously showed $\gamma''_i \in \Gamma_i$ and $\gamma''_i \sqsubseteq \overline{\gamma''_i}$, where Γ_i such that is the final state type of typing s_i according

to the evaluation of e_i . Since $lval$'s typing label is L in $\overline{\gamma_i''}$ in Γ_3 and we know that the state types in $\overline{\gamma_i''}$ are more restrictive than the state types in γ_i'' , we can conclude that $\gamma_i'' \in \Gamma_i$ also types $lval$ as L $\gamma_i''(lval) = \tau_1'' \wedge \text{lbl}(\tau_1'') = L$.

Previously, we demonstrated that the $lval$'s typing label is L in all final state types in Γ_4 and Γ_5 , and we showed that Γ_4 is not empty. Consequently, neither s_1 nor s_2 can modify $lval$. Considering the assumptions (a) and (b) (related to the first configuration), where s_i can be either s_1 or s_2 , we conclude that the $lval$ remains unchanged there as well. Here, we can apply Lemma 13 to deduce that the $lval$'s typing label in $\gamma_i'' \in \Gamma_i$ are more restrictive than the one we find in the refined state $\text{refine}(\gamma_1, e_i)$, thus $\text{refine}(\gamma_1, e_i)(lval) = \tau_i' \wedge \text{lbl}(\tau_i') = L$. Now we can apply Lemma 10 for any s_i to show that $m_1(lval) = m_1'(lval)$.

Finally, since the typing label of $lval$ in $\text{refine}(\gamma_1, e_i)$ is L , then trivially we know that the typing label of $lval$ in γ_1 is also L because $\gamma_1 \sqsubseteq \text{refine}(\gamma_1, e_i)$. We previously showed that $lval$'s typing label is L in γ_2 , thus we now can show that $\gamma_1 \sqcup \gamma_2(lval) = \tau' \wedge \text{lbl}(\tau') = L$. Now, we can deduce that $m_1(lval) = m_2(lval)$ from the definition of the assumption $m_1 \underset{\gamma_1 \sqcup \gamma_2}{\sim} m_2$. Thus, the goal holds.

case $\ell_1 = H$: when fixing the first configuration, we implement same proof as previous case.

◇ **Case sequence**: Here $stmt$ is $s_1; s_2$. From sequence typing rule we know:

- 1) $T, pc, \gamma_1 \vdash s_1 : \Gamma_1$
- 2) $\forall \gamma_1' \in \Gamma_1. T, pc, \gamma_1' \vdash s_2 : \Gamma_2^{\gamma_1'}$
- 3) $\Gamma' = \bigcup_{\gamma_1' \in \Gamma_1} \Gamma_2^{\gamma_1'}$
- 4) $T, pc, \gamma_2 \vdash s_1 : \Gamma_3$
- 5) $\forall \gamma_2' \in \Gamma_3. T, pc, \gamma_2' \vdash s_2 : \Gamma_4^{\gamma_2'}$
- 6) $\Gamma'' = \bigcup_{\gamma_2' \in \Gamma_3} \Gamma_4^{\gamma_2'}$

And from sequence reduction rule we know:

- 1) $E_1 : m_1 \xrightarrow{s_1} m_1'$
- 2) $E_1 : m_1' \xrightarrow{s_2} m_1''$
- 3) $E_2 : m_2 \xrightarrow{s_1} m_2'$
- 4) $E_2 : m_2' \xrightarrow{s_2} m_2''$

We initially know that $\gamma_1 \vdash m_1, \gamma_2 \vdash m_2$, and $m_1 \underset{\gamma_1 \sqcup \gamma_2}{\sim} m_2$.

In addition to that, we get induction hypothesis for s_1 IH1 and s_2 IH2 (we only show IH1):

$$\begin{aligned} & \forall T \ pc \ m_a \ m_b \ m'_a \ m'_b \ \gamma_a \ \gamma_b \ \Gamma_a \ \Gamma_b \ E_a \ E_b. \\ & T, pc, \gamma_a \vdash s_1 : \Gamma_a \ \wedge \ T, pc, \gamma_b \vdash s_1 : \Gamma_b \ \implies \\ & \quad T \vdash E_a \ \wedge \ T \vdash E_b \ \wedge \ E_a \underset{T}{\sim} E_b \ \wedge \\ & \quad \gamma_a \vdash m_a \ \wedge \ \gamma_b \vdash m_b \ \wedge \ m_a \underset{\gamma_a \sqcup \gamma_b}{\sim} m_b \ \wedge \\ & \quad E_a : m_a \xrightarrow{s_1} m'_a \ \wedge \ E_b : m_b \xrightarrow{s_1} m'_b \\ & \quad \implies \end{aligned}$$

$$(\exists \gamma'_a \in \Gamma_a \ \wedge \ \gamma'_b \in \Gamma_b. \ \gamma'_a \vdash m'_a \ \wedge \ \gamma'_b \vdash m'_b \ \wedge \ m'_a \underset{\gamma'_a \sqcup \gamma'_b}{\sim} m'_b)$$

We need to prove there are two state types $\gamma_1'' \in \Gamma'$ and $\gamma_2'' \in \Gamma''$ such they type the final states $\gamma_1'' \vdash m_1''$ and $\gamma_2'' \vdash m_2''$ and indeed $m_1'' \underset{\gamma_1'' \sqcup \gamma_2''}{\sim} m_2''$ holds.

We start by using IH1, and instantiating it with $(T, pc, m_1, m_2, m'_1, m'_2, \gamma_1, \gamma_2, \Gamma_1, \Gamma_3, E_1, E_2)$ to infer that there exists $\gamma_1' \in \Gamma_1$ and $\gamma_2' \in \Gamma_3$ such that $\gamma_1' \vdash m_1' \ \wedge \ \gamma_2' \vdash m_2'$ and also $m_1' \underset{\gamma_1' \sqcup \gamma_2'}{\sim} m_2'$.

Then, in the typing rule, we instantiate 2 with γ_1' and 5 with γ_2' . Now we can use IH2, and instantiating it with $(T, pc, m'_1, m'_2, m''_1, m''_2, \gamma_1', \gamma_2', \Gamma_2^{\gamma_1'}, \Gamma_4^{\gamma_2'}, E_1, E_2)$. From that we can infer that indeed there exists state types $\gamma_1'' \in \Gamma_2^{\gamma_1'}$ and $\gamma_2'' \in \Gamma_4^{\gamma_2'}$ such that they type the final states $\gamma_1'' \vdash m_1''$ and $\gamma_2'' \vdash m_2''$, where they keep the states low equivalent as $m_1'' \underset{\gamma_1'' \sqcup \gamma_2''}{\sim} m_2''$.

We know that the final set of state type of interest is simply the union of all state types that can type the second statement in 3 and 6. It is easy to see that since $\gamma_1'' \in \Gamma_2^{\gamma_1'}$ then $\gamma_1'' \in \Gamma'$. Similarly, $\gamma_2'' \in \Gamma_4^{\gamma_2'}$ then $\gamma_2'' \in \Gamma''$. Thus, the goal is proven.

◇ **Case function call**: Here $stmt$ is $f(e_1, \dots, e_n)$. From call typing rule we know (note that here we explicitly write the global and local state type):

- 1) $(\gamma_{g1}, \gamma_{l1}) \vdash e_i : \tau_{i1}$
- 2) $(s, (x_1, d_1), \dots, (x_n, d_n)) = (C, F)(f)$
- 3) $\gamma_{f1} = \{x_i \mapsto \tau_{i1}\}$
- 4) $(C, F), pc, (\gamma_{g1}, \gamma_{f1}) \vdash s : \Gamma'_1$
- 5) $\Gamma''_1 = \{(\gamma'_{g1}, \gamma_{l1})[e_i \mapsto \gamma'_{f1}(x_i) \mid \text{isOut}(d_i)] \mid (\gamma'_{g1}, \gamma'_{f1}) \in \Gamma'_1\}$
- 6) $(\gamma_{g2}, \gamma_{l2}) \vdash e_i : \tau_{i2}$
- 7) $\gamma_{f2} = \{x_i \mapsto \tau_{i2}\}$
- 8) $(C, F), pc, (\gamma_{g2}, \gamma_{f2}) \vdash s : \Gamma'_2$
- 9) $\Gamma''_2 = \{(\gamma'_{g2}, \gamma_{l2})[e_i \mapsto \gamma'_{f2}(x_i) \mid \text{isOut}(d_i)] \mid (\gamma'_{g2}, \gamma'_{f2}) \in \Gamma'_2\}$

And from call reduction rule we know:

- 1) $(s, (x_1, d_1), \dots, (x_n, d_n)) = (X_1, F)(f)$
- 2) $m_{f1} = \{x_i \mapsto (m_{g1}, m_{l1})(e_i)\}$
- 3) $(X_1, F) : (m_{g1}, m_{f1}) \xrightarrow{s} (m'_{g1}, m'_{f1})$
- 4) $m''_1 = (m'_{g1}, m_{l1})[e_i \mapsto m'_{f1}(x_i) \mid \text{isOut}(d_i)]$
- 5) $(s, (x_1, d_1), \dots, (x_n, d_n)) = (X_2, F)(f)$
- 6) $m_{f2} = \{x_i \mapsto (m_{g2}, m_{l2})(e_i)\}$
- 7) $(X_2, F) : (m_{g2}, m_{f2}) \xrightarrow{s} (m'_{g2}, m'_{f2})$
- 8) $m''_2 = (m'_{g2}, m_{l2})[e_i \mapsto m'_{f2}(x_i) \mid \text{isOut}(d_i)]$

Let $m_1 = (m_{g1}, m_{l1})$, $m_2 = (m_{g2}, m_{l2})$, let $\gamma_1 = (\gamma_{g1}, \gamma_{l1})$, and $\gamma_2 = (\gamma_{g2}, \gamma_{l2})$. We initially know that: $\gamma_1 \vdash m_1, \gamma_2 \vdash m_2$, and $m_1 \underset{\gamma_1 \sqcup \gamma_2}{\sim} m_2$.

In addition to that, we get induction hypothesis for s IH:

$$\begin{aligned} & \forall T \text{ pc } m_a \ m_b \ m'_a \ m'_b \ \gamma_a \ \gamma_b \ \Gamma_a \ \Gamma_b \ E_a \ E_b. \\ & T, \text{pc}, \gamma_a \vdash s : \Gamma_a \ \wedge \ T, \text{pc}, \gamma_b \vdash s : \Gamma_b \implies \\ & T \vdash E_a \ \wedge \ T \vdash E_b \ \wedge \ E_a \sim_T E_b \ \wedge \end{aligned}$$

$$\gamma_a \vdash m_a \ \wedge \ \gamma_b \vdash m_b \ \wedge \ m_a \sim_{\gamma_a \sqcup \gamma_b} m_b \ \wedge$$

$$E_a : m_a \xrightarrow{s} m'_a \ \wedge \ E_b : m_b \xrightarrow{s} m'_b$$

$$\implies$$

$$(\exists \gamma'_a \in \Gamma_a \ \wedge \ \gamma'_b \in \Gamma_b. \ \gamma'_a \vdash m'_a \ \wedge \ \gamma'_b \vdash m'_b \ \wedge \ m'_a \sim_{\gamma'_a \sqcup \gamma'_b} m'_b)$$

We need to prove there are two state types $\gamma''_1 \in \Gamma''_1$ and $\gamma''_2 \in \Gamma''_2$ such they type the final states $\gamma''_1 \vdash m''_1$ and $\gamma''_2 \vdash m''_2$ and indeed $m''_1 \sim_{\gamma''_1 \sqcup \gamma''_2} m''_2$ holds.

First we need to prove that the resulted copy-in map is also well-typed i.e. $\gamma_{f1} \vdash m_{f1}$. Note that the same proof applies to prove $\gamma_{f2} \vdash m_{f2}$:

Given that initially the typing state types the state $(\gamma_{g1}, \gamma_{l1}) \vdash (m_{g1}, m_{l1})$ from assumptions and given that the expressions e_i have a type τ_{i1} from typing rule 1, now we can use Lemma 4 to infer that for all $v_i : \tau_i$ such that v_i is the evaluation of $(m_{g1}, m_{l1})(e_i)$. This leads us to trivially infer that $\forall i. \{x_i \mapsto \tau_i\} \vdash \{x_i \mapsto v_i\}$ holds, thus $\gamma_{f1} \vdash m_{f1}$.

Now we want to show that $(m_{g1}, m_{f1}) \sim_{(\gamma_{g1}, \gamma_{f1}) \sqcup (\gamma_{g2}, \gamma_{f2})} (m_{g2}, m_{f2})$. This can be broken down and rewritten into two goals according to Lemma 15:

Goal 1. $m_{g1} \sim_{\gamma_{g1} \sqcup \gamma_{g2}} m_{g2}$:

We know that the domains of m_{g1} and m_{l1} are distinct and do not intersect (similarly for m_{g2} and m_{l2}), and given that they are initially low equivalent with respect to $(\gamma_{g1}, \gamma_{l1}) \sqcup (\gamma_{g2}, \gamma_{l2})$ as $(m_{g1}, m_{l1}) \sim_{(\gamma_{g1}, \gamma_{l1}) \sqcup (\gamma_{g2}, \gamma_{l2})} (m_{g2}, m_{l2})$. Then we can use Lemma 15 to show that the goal holds.

Goal 2. $m_{f1} \sim_{\gamma_{f1} \sqcup \gamma_{f2}} m_{f2}$:

First we start by rewriting the goal as following:

$$\begin{aligned} & \{x_i \mapsto (m_{g1}, m_{l1})(e_i)\} \sim_{\{x_i \mapsto (\gamma_{g1}, \gamma_{l1})(e_i)\} \sqcup \{x_i \mapsto (\gamma_{g2}, \gamma_{l2})(e_i)\}} \\ & \{x_i \mapsto (m_{g2}, m_{l2})(e_i)\}. \end{aligned}$$

It is easy to see that the empty map is low equivalent as: $\{\} \sim_{\{\} \sqcup \{\}} \{\}$. Now we can use

Lemma 16 to show that the goal holds.

Now we can use IH, and instantiate it with: $((C, F), \text{pc}, (m_{g1}, m_{f1}), (m_{g2}, m_{f2}), (m'_{g1}, m'_{f1}), (m'_{g2}, m'_{f2}), (\gamma_{g1}, \gamma_{f1}), (\gamma_{g2}, \gamma_{f2}), \Gamma'_1, \Gamma'_2, (X_1, F), (X_2, F))$, so that we can deduct that exists $\gamma'_1 \in \Gamma'_1$ such that $\gamma'_1 \vdash (m'_{g1}, m'_{f1})$, also exists $\gamma'_2 \in \Gamma'_2$ such that $\gamma'_2 \vdash (m'_{g2}, m'_{f2})$. We can also infer that $(m'_{g1}, m'_{f1}) \sim_{\gamma'_1 \sqcup \gamma'_2} (m'_{g2}, m'_{f2})$.

Let $\overline{\gamma'_1} = (\overline{\gamma'_{g1}}, \overline{\gamma'_{f1}})$ and $\overline{\gamma'_2} = (\overline{\gamma'_{g2}}, \overline{\gamma'_{f2}})$ then we can also conclude from the previous:

- (a) $\overline{\gamma'_{g1}} \vdash m'_{g1}$
- (b) $\overline{\gamma'_{g2}} \vdash m'_{g2}$
- (c) $\overline{\gamma'_{f1}} \vdash m'_{f1}$

- (d) $\overline{\gamma'_{f2}} \vdash m'_{f2}$
- (e) $m'_{g1} \sim_{\overline{\gamma'_{g1}} \sqcup \overline{\gamma'_{g2}}} m'_{g2}$
- (f) $m'_{f1} \sim_{\overline{\gamma'_{f1}} \sqcup \overline{\gamma'_{f2}}} m'_{f2}$

Since the final goal is to prove there are two state types $\gamma''_1 \in \Gamma''_1$ and $\gamma''_2 \in \Gamma''_2$ such they type the final states $\gamma''_1 \vdash m''_1$ and $\gamma''_2 \vdash m''_2$ and indeed $m''_1 \sim_{\gamma''_1 \sqcup \gamma''_2} m''_2$ holds, then we can select

γ''_1 to be $(\overline{\gamma'_{g1}}, \overline{\gamma'_{f1}})$, i.e. the copied-out map is $(\overline{\gamma'_{g1}}, \gamma_{l1})[e_i \mapsto \overline{\gamma'_{f1}}(x_i) \mid \text{isOut}(d_i)]$.

Similarly, we can select γ''_2 to be $(\overline{\gamma'_{g2}}, \overline{\gamma'_{f2}})$, i.e. the copied-out map is $(\overline{\gamma'_{g2}}, \gamma_{l2})[e_i \mapsto \overline{\gamma'_{f2}}(x_i) \mid \text{isOut}(d_i)]$.

Goal 1:

$$(\overline{\gamma'_{g1}}, \gamma_{l1})[e_i \mapsto \overline{\gamma'_{f1}}(x_i)] \vdash (m'_{g1}, m_{l1})[e_i \mapsto m'_{f1}(x_i)]$$

Since we know $\overline{\gamma'_{g1}} \vdash m'_{g1}$ ((a) from previous step) and $\gamma_{l1} \vdash m_{l1}$ from assumptions rewrites, this entails that $(\overline{\gamma'_{g1}}, \gamma_{l1}) \vdash (m'_{g1}, m_{l1})$. We also proved that $\overline{\gamma'_{f1}} \vdash m'_{f1}$ ((c) from previous step), thus for all $x \in \text{domain}(\overline{\gamma'_{f1}})$ holds $m'_{f1}(x) : \overline{\gamma'_{f1}}(x)$. Now we can use Lemma 5 to show that the goal holds.

Goal 2:

$$(\overline{\gamma'_{g2}}, \gamma_{l2})[e_i \mapsto \overline{\gamma'_{f2}}(x_i)] \vdash (m'_{g2}, m_{l2})[e_i \mapsto m'_{f2}(x_i)]$$

Same proof as the previous sub goal, using (b) and (d) the previous step, and the initial assumption $\gamma_{l2} \vdash m_{l2}$.

Goal 3:

$$\begin{aligned} & (m'_{g1}, m_{l1})[e_i \mapsto m'_{f1}(x_i)] \\ & \sim_{(\overline{\gamma'_{g1}}, \gamma_{l1})[e_i \mapsto \overline{\gamma'_{f1}}(x_i)] \sqcup (\overline{\gamma'_{g2}}, \gamma_{l2})[e_i \mapsto \overline{\gamma'_{f2}}(x_i)]} \\ & (m'_{g2}, m_{l2})[e_i \mapsto m'_{f2}(x_i)] \end{aligned}$$

We know that $m'_{g1} \sim_{\gamma'_{g1} \sqcup \gamma'_{g2}} m'_{g2}$ (from (e) of previous step), we also know that $m'_{l1} \sim_{\gamma_{l1} \sqcup \gamma_{l2}} m_{l2}$ (from assumptions), now using Lemma 15 we can combine them to infer an equivalence before copying out or (no copy-out because there are not out directed parameters), i.e. :

$$(m'_{g1}, m_{l1}) \sim_{(\overline{\gamma'_{g1}}, \gamma_{l1}) \sqcup (\overline{\gamma'_{g2}}, \gamma_{l2})} (m'_{g2}, m_{l2}).$$

We previously proved that $m_{f1} \sim_{\gamma_{f1} \sqcup \gamma_{f2}} m_{f2}$, now we can prove this goal directly using Lemma 16.

◇ **Case extern:** Here stmt is $f(e_1, \dots, e_n)$. From call typing rule we know (Note that here we explicitly write the global and local state type):

- 1) $(\gamma_{g1}, \gamma_{l1}) \vdash e_i : \tau_{i1}$
- 2) $(\text{Cont}_{\text{E}}, (x_1, d_1), \dots, (x_n, d_n)) = (C, F)(f)$
- 3) $\gamma_{f1} = \{x_i \mapsto \tau_{i1}\}$
- 4) $\forall (\gamma_i, \phi, \gamma_t) \in \text{Cont}_{\text{E}}. (\gamma_{g1}, \gamma_{l1}) \sqsubseteq \gamma_i$
- 5) $\Gamma'_1 = \{\gamma'_1 \# \text{raise}(\gamma_t, \text{pc}) \mid (\gamma_i, \phi, \gamma_t) \in \text{Cont}_{\text{E}} \ \wedge \ \text{refine}((\gamma_{g1}, \gamma_{f1}), \phi) = \gamma'_1 \neq \bullet\}$
- 6) $\Gamma''_1 = \{(\gamma'_g, \gamma_{l1})[e_i \mapsto \gamma'_f(x_i) \mid \text{isOut}(d_i)] \mid (\gamma'_g, \gamma'_f) \in \Gamma'_1\}$

- 7) $(\gamma_{g2}, \gamma_{l2}) \vdash e_i : \tau_{i2}$
- 8) $\gamma_{f2} = \{x_i \mapsto \tau_{i2}\}$
- 9) $\forall (\gamma_i, \phi, \gamma_t) \in \text{Cont}_E. (\gamma_{g2}, \gamma_{l2}) \sqsubseteq \gamma_i$
- 10) $\Gamma'_2 = \{\gamma'_2 \# \text{raise}(\gamma_t, pc) \mid (\gamma_i, \phi, \gamma_t) \in \text{Cont}_E \wedge \text{refine}((\gamma_{g2}, \gamma_{f2}), \phi) = \gamma'_2 \neq \bullet\}$
- 11) $\Gamma''_2 = \{(\gamma'_g, \gamma_{l2})[e_i \mapsto \gamma'_f(x_i) \mid \text{isOut}(d_i)] \mid (\gamma'_g, \gamma'_f) \in \Gamma'_2\}$

And from extern reduction rule we know:

- 1) $(\text{sem}_f, (x_1, d_1), \dots, (x_n, d_n)) = (X_1, F)(f)$
- 2) $m_{f1} = \{x_i \mapsto (m_{g1}, m_{l1})(e_i)\}$
- 3) $(m'_{g1}, m'_{f1}) = \text{sem}_f(m_{g1}, m_{f1})$
- 4) $m''_1 = (m'_{g1}, m_{l1})[e_i \mapsto m'_{f1}(x_i) \mid \text{isOut}(d_i)]$
- 5) $(\text{sem}_f, (x_1, d_1), \dots, (x_n, d_n)) = (X_2, F)(f)$
- 6) $m_{f2} = \{x_i \mapsto (m_{g2}, m_{l2})(e_i)\}$
- 7) $(m'_{g2}, m'_{f2}) = \text{sem}_f(m_{g2}, m_{f2})$
- 8) $m''_2 = (m'_{g2}, m_{l2})[e_i \mapsto m'_{f2}(x_i) \mid \text{isOut}(d_i)]$

Let $m_1 = (m_{g1}, m_{l1})$, $m_2 = (m_{g2}, m_{l2})$, let $\gamma_1 = (\gamma_{g1}, \gamma_{l1})$, and $\gamma_2 = (\gamma_{g2}, \gamma_{l2})$. We initially know that : $\gamma_1 \vdash m_1$, $\gamma_2 \vdash m_2$, and $m_1 \underset{\gamma_1 \sqcup \gamma_2}{\sim} m_2$.

We need to prove there are two state types $\gamma''_1 \in \Gamma''_1$ and $\gamma''_2 \in \Gamma''_2$ such they type the final states $\gamma''_1 \vdash m''_1$ and $\gamma''_2 \vdash m''_2$ and indeed $m''_1 \underset{\gamma''_1 \sqcup \gamma''_2}{\sim} m''_2$ holds.

First we need to prove that the resulted copy-in map is also well-typed i.e. $\gamma_{f1} \vdash m_{f1}$ and $\gamma_{f2} \vdash m_{f2}$: same proof as the function call case.

Now we want to show that $(m_{g1}, m_{f1}) \underset{(\gamma_{g1}, \gamma_{f1}) \sqcup (\gamma_{g2}, \gamma_{f2})}{\sim} (m_{g2}, m_{f2})$: same proof as the function call case.

We know from extWT $C X_1$ and extWT $C X_2$ relation in the well-typedness $(X_1, F) \vdash (C, F)$ and $(X_2, F) \vdash (C, F)$ that indeed there exists an input state type, condition and output state type i.e. $(\gamma_i, \phi, \gamma_t)$ in the contract of the extern such that the condition satisfies the input state $\phi(m_{g1}, m_{f1})$.

From SOUNDNESS OF ABSTRACTION proof, we know that $(\gamma_{g1}, \gamma_{f1}) \vdash (m_{g1}, m_{f1})$. Using 4 of extern's typing rule we know that $(\gamma_{g1}, \gamma_{f1})$ is no more restrictive than γ_i , i.e. $(\gamma_{g1}, \gamma_{f1}) \sqsubseteq \gamma_i$ holds. Using the same strategy and 9 of extern's typing rule, we can also prove $(\gamma_{g2}, \gamma_{f2}) \sqsubseteq \gamma_i$. We can also deduct that $(m_{g1}, m_{f1}) \underset{\gamma_i}{\sim} (m_{g2}, m_{f2})$ using Lemma 9.

Let the variables $\{x_1, \dots, x_n\}$ be the ones used in the condition ϕ . Using the definition of extern's well-typedness again, we know that the least upper bound of typing label of $\{x_1, \dots, x_n\}$ in γ_i is no more restrictive that the lower bound of the output state type γ_t . The entails that since ϕ holds on m_1 (i.e. $\phi(m_1)$) then it indeed holds for m_2 (i.e. $\phi(m_2)$).

Now we split the proof into two cases:

A. For the changed variables by extern in the state: given 3 and 7 of the extern reduction rule, and using $\phi(m_1)$ we can use the definition of extern's well-typedness again to infer that the variables that are changed by the semantics are a subset of the domain of γ_t and low equivalent with respect to γ_t :

$$((m'_{g1}, m'_{f1}) \setminus (m_{g1}, m_{f1})) \underset{\gamma_t}{\sim} ((m'_{g2}, m'_{f2}) \setminus (m_{g2}, m_{f2}))$$

B. Now for the unchanged variables by extern in the state: it is easy to see that the refined state $\text{refine}((\gamma_{g1}, \gamma_{f1}), \phi)$ is more restrictive than $(\gamma_{g1}, \gamma_{f1})$, thus we can say $(\gamma_{g1}, \gamma_{f1}) \sqsubseteq \text{refine}((\gamma_{g1}, \gamma_{f1}), \phi)$. Now we can use Lemma 9 to show:

$$(m_{g1}, m_{f1}) \underset{\text{refine}((\gamma_{g1}, \gamma_{f1}), \phi)}{\sim} (m_{g2}, m_{f2}) \underset{\text{refine}((\gamma_{g1}, \gamma_{f1}), \phi)}{\sim} (m'_{g1}, m'_{f1})$$

This property also holds on the variable names x that are unchanged by the behavior of the extern, i.e.,:

$$(m'_{g1}, m'_{f1}) \underset{\text{refine}((\gamma_{g1}, \gamma_{f1}), \phi)}{\sim} (m'_{g2}, m'_{f2}) \underset{\text{refine}((\gamma_{g1}, \gamma_{f1}), \phi)}{\sim} (m_{g1}, m_{f1})$$

We can rename $\text{refine}((\gamma_{g1}, \gamma_{f1}), \phi) \# \gamma_t$ to $\overline{\gamma_3}$ and $\text{refine}((\gamma_{g2}, \gamma_{f2}), \phi) \# \gamma_t$ to $\overline{\gamma_4}$, and we can easily infer from A and B that : $(m'_{g1}, m'_{f1}) \underset{\overline{\gamma_3} \sqcup \overline{\gamma_4}}{\sim} (m'_{g2}, m'_{f2})$.

The rest of the proof is the same as the function call case.

◇ **Case table application:** Here *stmt* is apply *tbl*. From table rule we know:

- 1) $(\overline{e}, \text{Cont}_{tbl}) = (C, F)(tbl)$
- 2) $\gamma_1 \vdash e_i : \tau_{i1}$
- 3) $\ell_1 = \bigsqcup_i \text{lbl}(\tau_{i1})$
- 4) $pc'_1 = pc \sqcup \ell_1$
- 5) $\forall (\phi_j, (a_j, \overline{\tau}_j)) \in \text{Cont}_{tbl}. (\gamma_{g_j}, \gamma_{l_j}) = \text{refine}(\gamma_1, \phi_j) \wedge (s_j, (x_{j1}, \text{none}), \dots, (x_{jn}, \text{none})) = (C, F)(a_j) \wedge \gamma_{a_j} = \{x_{j_i} \mapsto \tau_{j_i}\} \wedge (C, F), pc'_1, (\gamma_{g_j}, \gamma_{a_j}) \vdash s_j : \Gamma_j$
- 6) $\Gamma'_1 = \cup_j \{(\gamma'_{g_j}, \gamma_{l_j}) \mid (\gamma'_{g_j}, \gamma'_{a_j}) \in \Gamma_j\}$
- 7) $\Gamma''_1 = \begin{cases} \text{join}(\Gamma'_1) & \text{if } \ell_1 = H \\ \Gamma'_1 & \text{otherwise} \end{cases}$
- 8) $\gamma_2 \vdash e_i : \tau_{i2}$
- 9) $\ell_2 = \bigsqcup_i \text{lbl}(\tau_{i2})$
- 10) $pc'_2 = pc \sqcup \ell_2$
- 11) $\forall (\phi_k, (a_k, \overline{\tau}_k)) \in \text{Cont}_{tbl}. (\gamma_{g_k}, \gamma_{l_k}) = \text{refine}(\gamma_2, \phi_k) \wedge (s_k, (x_{k1}, \text{none}), \dots, (x_{kn}, \text{none})) = (C, F)(a_k) \wedge \gamma_{a_k} = \{x_{k_i} \mapsto \tau_{k_i}\} \wedge (C, F), pc'_2, (\gamma_{g_k}, \gamma_{a_k}) \vdash s_k : \Gamma_k$
- 12) $\Gamma'_2 = \cup_k \{(\gamma'_{g_k}, \gamma_{l_k}) \mid (\gamma'_{g_k}, \gamma'_{a_k}) \in \Gamma_k\}$
- 13) $\Gamma''_2 = \begin{cases} \text{join}(\Gamma'_2) & \text{if } \ell_2 = H \\ \Gamma'_2 & \text{otherwise} \end{cases}$

And from table reduction rule we know:

- 1) $(\overline{e}, \text{sem}_{tbl1}) = (X_1, F)(tbl)$
- 2) $\text{sem}_{tbl1}((m_{g1}, m_{l1})(e_1), \dots, (m_{g1}, m_{l1})(e_n)) = (a_1, \overline{v})$
- 3) $(s_1, (x_1, \text{none}), \dots, (x_n, \text{none})) = (X_1, F)(a_1)$
- 4) $m_{a1} = \{x_i \mapsto v_i\}$
- 5) $(X_1, F) : (m_{g1}, m_{a1}) \xrightarrow{s_1} (m'_{g1}, m'_{a1})$
- 6) $m_{final1} = (m'_{g1}, m_{l1})$
- 7) $(\overline{e'}, \text{sem}_{tbl2}) = (X_2, F)(tbl)$
- 8) $\text{sem}_{tbl2}((m_{g2}, m_{l2})(e'_1), \dots, (m_{g2}, m_{l2})(e'_n)) = (a_2, \overline{v'})$
- 9) $(s_2, (x'_2, \text{none}), \dots, (x'_n, \text{none})) = (X_2, F)(a_2)$
- 10) $m_{a2} = \{x'_i \mapsto v'_i\}$
- 11) $(X_2, F) : (m_{g2}, m_{a2}) \xrightarrow{s_2} (m'_{g2}, m'_{a2})$
- 12) $m_{final2} = (m'_{g2}, m_{l2})$

In addition to that, we get induction hypothesis for s_1 IH1 (similarly for s_2 IH2):

$$\begin{aligned} \forall T \text{ pc } m_a m_b m'_a m'_b \gamma_a \gamma_b \Gamma_a \Gamma_b E_a E_b. \\ T, \text{pc}, \gamma_a \vdash s : \Gamma_a \wedge T, \text{pc}, \gamma_b \vdash s : \Gamma_b \implies \\ T \vdash E_a \wedge T \vdash E_b \wedge E_a \underset{T}{\sim} E_b \wedge \\ \gamma_a \vdash m_a \wedge \gamma_b \vdash m_b \wedge m_a \underset{\gamma_a \sqcup \gamma_b}{\sim} m_b \wedge \\ E_a : m_a \xrightarrow{s_1} m'_a \wedge E_b : m_b \xrightarrow{s_1} m'_b \\ \implies \end{aligned}$$

$$(\exists \gamma'_a \in \Gamma_a \wedge \gamma'_b \in \Gamma_b. \gamma'_a \vdash m'_a \wedge \gamma'_b \vdash m'_b \wedge m'_a \underset{\gamma'_a \sqcup \gamma'_b}{\sim} m'_b)$$

Let $m_1 = (m_{g1}, m_{l1})$, $m_2 = (m_{g2}, m_{l2})$, let $\gamma_1 = (\gamma_{g1}, \gamma_{l1})$, and $\gamma_2 = (\gamma_{g2}, \gamma_{l2})$. We initially know that : $\gamma_1 \vdash m_1$, $\gamma_2 \vdash m_2$, and $m_1 \underset{\gamma_1 \sqcup \gamma_2}{\sim} m_2$.

The final goal is to prove there are two state types $\gamma''_1 \in \Gamma''_1$ and $\gamma''_2 \in \Gamma''_2$ such they type the final states $\gamma''_1 \vdash (m'_{g1}, m_{l1})$ and $\gamma''_2 \vdash (m'_{g2}, m_{l2})$ and indeed $(m'_{g1}, m_{l1}) \underset{\gamma''_1 \sqcup \gamma''_2}{\sim} (m'_{g2}, m_{l2})$ holds.

Initially, we prove expression (table keys) found in 1 in reduction rule, with 1 and 7 of typing rule are the same. We use $(X_1, F) \underset{(C,F)}{\sim} (X_2, F)$ to show that 1 and 7 in the typing rule have the same expression (i.e. $\bar{e} = \bar{e}'$). Then, using tblWT $C X_1$ and tblWT $C X_2$, we confirm that the expression in rule 1 of the reduction rule matches those in rules 1 and 7 of the typing rule. Therefore, all relevant expressions are equivalent.

First, we conduct a case analysis on ℓ_1 and ℓ_2 being equivalent:

case $\ell_1 = \ell_2 = L$: Given that ℓ_1 and ℓ_2 are L in 3 and 9 of the typing rule, respectively, we can conclude that the evaluation of the key expressions \bar{e} in states m_1 and m_2 are identical. This follows directly from Lemma 8, which establishes that $m_1(e_i) = m_2(e_i)$ for all e_i .

By the definition of $(X_1, F) \underset{(C,F)}{\sim} (X_2, F)$, we know that

for any memory states m_1 and m_2 , if $m_1(e_i) = m_2(e_i)$ for all table's keys e_i , then $m_1(\phi) \Leftrightarrow m_2(\phi)$. This implies that the condition for both tables to match is identical. Applying this definition again, we deduce that the actions a_1 and a_2 in 2 and 8 of the reduction rule are identical i.e. $a_1 = a_2 = a$, therefore their corresponding action bodies and signatures must also be the same in (X_1, F) and (X_2, F) , thus $s_1 = s_2 = s$ and $\bar{x} = \bar{x}'$. Applying this definition $(X_1, F) \underset{(C,F)}{\sim} (X_2, F)$, yet again, we can also deduce that the action's values \bar{v} and \bar{v}' are low equivalent wrt. $\bar{\tau}$ i.e. $\bar{v} \underset{\bar{\tau}}{\sim} \bar{v}'$.

Now we can rewrite the table reduction rule to:

- $(\bar{e}, \text{sem}_{tbl1}) = (X_1, F)(tbl)$
- $\text{sem}_{tbl1}((m_{g1}, m_{l1})(e_1), \dots, (m_{g1}, m_{l1})(e_n)) = (a, \bar{v})$
- $(s, (x_1, \text{none}), \dots, (x_n, \text{none})) = (X_1, F)(a)$
- $m_{a1} = \{x_i \mapsto v_i\}$
- $(X_1, F) : (m_{g1}, m_{a1}) \xrightarrow{s} (m'_{g1}, m'_{a1})$

- $m_{final1} = (m'_{g1}, m_{l1})$
- $(\bar{e}, \text{sem}_{tbl2}) = (X_2, F)(tbl)$
- $\text{sem}_{tbl2}((m_{g2}, m_{l2})(e_1), \dots, (m_{g2}, m_{l2})(e_n)) = (a, \bar{v}')$
- $(s, (x_2, \text{none}), \dots, (x_n, \text{none})) = (X_2, F)(a)$
- $m_{a2} = \{x_i \mapsto v'_i\}$
- $(X_2, F) : (m_{g2}, m_{a2}) \xrightarrow{s} (m'_{g2}, m'_{a2})$
- $m_{final2} = (m'_{g2}, m_{l2})$

Given $(C, F) \vdash (X_1, F)$ that indeed exists condition ϕ_j in Cont_{tbl} that satisfies the table input state (m_{g1}, m_{l1}) , and also exists a list of types $\bar{\tau}$ such that it can type that values of the table's semantics (in 2 of the typing rule) i.e. $v_i : \tau_i$. Similarly, from $(C, F) \vdash (X_2, F)$, we know that exists ϕ_k that satisfies the table input state (m_{g2}, m_{l2}) , and exists $\bar{\tau}'$ such that $v'_i : \tau'_i$.

We can instantiate 5 from the table rule with $(\phi_j, (a, \bar{\tau}))$, and instantiate 11 with $(\phi_k, (a, \bar{\tau}'))$.

We need to prove $(m_{g1}, m_{a1}) \underset{(\gamma_{g_j}, \gamma_{a_j}) \sqcup (\gamma_{g_k}, \gamma_{a_k})}{\sim} (m_{g2}, m_{a2})$ using the following sub-goals:

Goal 1. prove $m_{g1} \underset{\gamma_{g_j} \sqcup \gamma_{g_k}}{\sim} m_{g2}$ and $m_{l1} \underset{\gamma_{l_j} \sqcup \gamma_{l_k}}{\sim} m_{l2}$:

It is easy to see that $\gamma_1 \sqsubseteq \text{refine}(\gamma_1, \phi_j)$ and $\gamma_2 \sqsubseteq \text{refine}(\gamma_2, \phi_k)$ trivially hold, and can be rewritten to $\gamma_1 \sqsubseteq (\gamma_{g_j}, \gamma_{l_j})$ and $\gamma_2 \sqsubseteq (\gamma_{g_k}, \gamma_{l_k})$. Since initially we know that $(m_{g1}, m_{l1}) \underset{\gamma_1 \sqcup \gamma_2}{\sim} (m_{g2}, m_{l2})$, therefore using Lemma 9 $(m_{g1}, m_{l1}) \underset{(\gamma_{g_j}, \gamma_{l_j}) \sqcup (\gamma_{g_k}, \gamma_{l_k})}{\sim} (m_{g2}, m_{l2})$.

This entails using Lemma 15 that $m_{g1} \underset{\gamma_{g_j} \sqcup \gamma_{g_k}}{\sim} m_{g2}$ and also $m_{l1} \underset{\gamma_{l_j} \sqcup \gamma_{l_k}}{\sim} m_{l2}$ hold.

Additionally, note that Hyp 3 states also that $(\gamma_{g_j}, \gamma_{l_j})$ can still type the state (m_{g1}, m_{l1}) , i.e. $(\gamma_{g_j}, \gamma_{l_j}) \vdash (m_{g1}, m_{l1})$. Similarly, $(\gamma_{g_k}, \gamma_{l_k}) \vdash (m_{g2}, m_{l2})$.

Goal 2. prove $m_{a1} \underset{\gamma_{a_j} \sqcup \gamma_{a_k}}{\sim} m_{a2}$: From $(X_1, F) \underset{(C,F)}{\sim} (X_2, F)$

we know it holds $\bar{v} \underset{\bar{\tau}}{\sim} \bar{v}'$, note that these are the table's semantic output (i.e. will be action's arguments). Thus, whenever τ_i is L , then the values of arguments are equivalent $v_i = v'_i$. This entails that constructing states $m_{a1} = \{x_i \mapsto v_i\}$ and $m_{a2} = \{x_i \mapsto v'_i\}$ must be low equivalent with respect to $\gamma_{a_j} = \{x_i \mapsto \tau_i\}$, i.e. $m_{a1} \underset{\gamma_{a_j}}{\sim} m_{a2}$. Similarly, from $(X_1, F) \underset{(C,F)}{\sim} (X_2, F)$ and whenever τ'_i is L , we can also deduce that $\gamma_{a_k} = \{x_i \mapsto \tau'_i\}$, i.e. $m_{a1} \underset{\gamma_{a_k}}{\sim} m_{a2}$. Therefore, this entails (using Lemma 15) that $m_{a1} \underset{\gamma_{a_j} \sqcup \gamma_{a_k}}{\sim} m_{a2}$.

From the last two sub-goals, we can use Lemma 15 to deduce $(m_{g1}, m_{a1}) \underset{(\gamma_{g_j}, \gamma_{a_j}) \sqcup (\gamma_{g_k}, \gamma_{a_k})}{\sim} (m_{g2}, m_{a2})$.

Now we can use IH and instantiate it with $((C, F), \text{pc} \sqcup \ell, (m_{g1}, m_{a1}), (m_{g2}, m_{a2}), (m'_{g1}, m'_{a1}), (m'_{g2}, m'_{a2}), (\gamma_{g_j}, \gamma_{a_j}), (\gamma_{g_k}, \gamma_{a_k}), \Gamma_j, \Gamma_k, (X_1, F), (X_2, F))$ to deduce that indeed exist $\bar{\gamma}_1 \in \Gamma_j$ and $\bar{\gamma}_2 \in \Gamma_k$ such that $\bar{\gamma}_1 \vdash (m'_{g1}, m'_{a1})$ and $\bar{\gamma}_2 \vdash (m'_{g2}, m'_{a2})$ and indeed $(m'_{g1}, m'_{a1}) \underset{\bar{\gamma}_1 \sqcup \bar{\gamma}_2}{\sim} (m'_{g2}, m'_{a2})$. In the following, let $\bar{\gamma}_1 = (\bar{\gamma}_{1g}, \bar{\gamma}_{1a})$ and $\bar{\gamma}_2 = (\bar{\gamma}_{2g}, \bar{\gamma}_{2a})$.

We can rewrite the IH results as following: exist $(\overline{\gamma_{1g}}, \overline{\gamma_{1l}}) \in \Gamma_j$ and $(\overline{\gamma_{2g}}, \overline{\gamma_{2l}}) \in \Gamma_k$ such that $(\overline{\gamma_{1g}}, \overline{\gamma_{1l}}) \vdash (m'_{g1}, m'_{a1})$ and $(\overline{\gamma_{2g}}, \overline{\gamma_{2l}}) \vdash (m'_{g2}, m'_{a2})$ and indeed $(m'_{g1}, m'_{a1}) \sim_{(\overline{\gamma_{1g}}, \overline{\gamma_{1l}}) \sqcup (\overline{\gamma_{2g}}, \overline{\gamma_{2l}})} (m'_{g2}, m'_{a2})$.

Clearly, from 6 in the typing rule, and we know that Γ'_1 is the union of all the changed global state types by the action's body, with the refined starting local state type $(\gamma_{gj}, \gamma_{lj}) = \text{refine}(\gamma_1, \phi_j)$. Thus, we know that indeed $(\overline{\gamma_{1g}}, \overline{\gamma_{1l}}) \in \Gamma'_1$. Similarly, from 12 in the typing rule, we know that $(\overline{\gamma_{2g}}, \overline{\gamma_{2l}}) \in \Gamma'_2$.

We choose $(\overline{\gamma_{1g}}, \overline{\gamma_{1l}})$ and $(\overline{\gamma_{2g}}, \overline{\gamma_{2l}})$ to finish the proof of the goal in this subcase.

Since $\ell_1 = \ell_2$ and is L in this sub-case, then trivially $\Gamma''_1 = \Gamma'_1$ in 7 of the typing rule, and $\Gamma''_2 = \Gamma'_2$ in 13 of the typing rule. Since we proved that $\gamma_{lj} \vdash m_{l1}$ and $\gamma_{lk} \vdash m_{l2}$, therefore $(\overline{\gamma_{1g}}, \overline{\gamma_{1l}}) \vdash (m'_{g1}, m_{l1})$ and $(\overline{\gamma_{2g}}, \overline{\gamma_{2l}}) \vdash (m'_{g2}, m_{l2})$ hold. Additionally, using Lemma 15 $(m'_{g1}, m_{l1}) \sim_{(\overline{\gamma_{1g}}, \overline{\gamma_{1l}}) \sqcup (\overline{\gamma_{2g}}, \overline{\gamma_{2l}})} (m'_{g2}, m_{l2})$.

case $\ell_1 \neq \ell_2$: This proof is similar to the conditional case. We initiate the proof by fixing ℓ_2 to be H while ℓ_1 can be either H or L , thus the evaluation of \bar{e} in the states (m_{g1}, m_{l1}) and (m_{g2}, m_{l2}) differs. Consequently, we (possibly) end up with two different actions and their corresponding arguments (a_1, \bar{v}) and (a_2, \bar{v}') .

From $(C, F) \vdash (X_1, F)$ we know that indeed exists condition ϕ_j in Cont_{tbl} that satisfies the table input state (m_{g1}, m_{l1}) , and also exists a list of types $\bar{\tau}$ such that it can type that values of the table's semantics (in 2 of the typing rule) i.e. $v_i : \tau_i$. Similarly, from $(C, F) \vdash (X_2, F)$, we know that exists ϕ_k that satisfies the table input state (m_{g2}, m_{l2}) , and exists $\bar{\tau}'$ such that $v'_i : \tau'_i$.

We can instantiate 5 (we refer to these as the first configuration) from the table rule with $(\phi_j, (a_1, \bar{\tau}))$, and instantiate 11 (we refer to these as the second configuration) with $(\phi_k, (a_2, \bar{\tau}'))$. Since the actions are different, then we let s_1 be the body of action a_1 , and s_2 be the body of action a_2 (we refer to these as the first configuration). Using the same steps in the previous sub-case, we know that $m_{g1} \sim_{\gamma_{gj} \sqcup \gamma_{gk}} m_{g2}$ and $m_{l1} \sim_{\gamma_{lj} \sqcup \gamma_{lk}} m_{l2}$.

The final goal is to prove there are two state types $\gamma''_1 \in \Gamma''_1$ and $\gamma''_2 \in \Gamma''_2$ such they type the final states $\gamma''_1 \vdash (m'_{g1}, m_{l1})$ and $\gamma''_2 \vdash (m'_{g2}, m_{l2})$ and indeed $(m'_{g1}, m_{l1}) \sim_{\gamma''_1 \sqcup \gamma''_2} (m'_{g2}, m_{l2})$ holds.

From the SOUNDNESS OF ABSTRACTION, we know that for the second configuration indeed exists $\gamma'_2 \in \Gamma_k$ such that it types the action's resulted state (m'_{g2}, m'_{a2}) (i.e. $\gamma'_2 \vdash (m'_{g2}, m'_{a2})$). In the following, let $(\gamma'_{g2}, \gamma'_{l2}) = \gamma'_2$, thus indeed $(\gamma'_{g2}, \gamma'_{l2}) \vdash (m'_{g2}, m'_{a2})$. Given that Γ'_2 in 12 of the typing rule is a union of all Γ_i ; thus indeed it includes Γ_k that has the resulted global state type γ'_{g2} , and the refined caller's local state type γ_{lk} while removing the callee's resulted local state type γ_{ak} . In short, we know that indeed for the second configuration will find $(\gamma'_{g2}, \gamma_{lk}) \in \Gamma'_2$. Note that in the SOUNDNESS

OF ABSTRACTION previously we proved that $\gamma_{lk} \vdash m_{l2}$, therefore $(\gamma'_{g2}, \gamma_{lk}) \vdash (m'_{g2}, m_{l2})$. Since ℓ_2 is H , then $\Gamma''_2 = \text{join}(\Gamma'_2)$, so we can deduce (by Lemma 7) the existence of $\gamma'_2 \in \text{join}(\Gamma'_2)$ such that it is more restrictive than γ'_2 , denoted as $\gamma'_2 \sqsubseteq \overline{\gamma'_2}$. Using the same lemma, we conclude that $\overline{\gamma'_2} \vdash (m'_{g2}, m_{l2})$. Now on, we choose $\overline{\gamma'_2}$ to be used in the proof and resolve the second conjunction of the goal.

From the SOUNDNESS OF ABSTRACTION, we know that for the first configuration indeed exists $\gamma'_1 \in \Gamma_j$ such that it types action's resulted state (m'_{g1}, m_{a1}) (i.e. $\gamma'_1 \vdash (m'_{g1}, m_{a1})$). In the first configuration, the final state type set Γ''_1 can be either a union (if $\ell_1 = L$) or a join (if $\ell_1 = H$) of all final state type sets and including Γ'_1 . In either case (union or join), we can establish the existence of $\overline{\gamma'_1} \in \Gamma''_1$ such that $\gamma'_1 \sqsubseteq \overline{\gamma'_1}$ and indeed $\overline{\gamma'_1} \vdash (m'_{g1}, m_{l1})$. Note that if Γ''_1 resulted from a join, we follow the same steps of the (second configuration) in the previous step; otherwise, if it resulted from a union, it is trivially true. Thus, the first conjunction of the final goal is proved.

The final goal left to prove is $(m'_{g1}, m_{l1}) \sim_{\overline{\gamma'_1} \sqcup \overline{\gamma'_2}} (m'_{g2}, m_{l2})$ holds. In the following, let $(\overline{\gamma'_{g1}}, \overline{\gamma'_{l1}}) = \overline{\gamma'_1}$ and $(\overline{\gamma'_{g2}}, \overline{\gamma'_{l2}}) = \overline{\gamma'_2}$.

Next, we proceed to implement cases based on whether an $lval$'s type label is H or L .

case $(\overline{\gamma'_1} \sqcup \overline{\gamma'_2}(lval) = \tau) \wedge \text{lbl}(\tau) = H$: holds trivially.
case $(\overline{\gamma'_1} \sqcup \overline{\gamma'_2}(lval) = \tau) \wedge \text{lbl}(\tau) = L$: this case entails that each state type individually holds $\overline{\gamma'_1}(lval) = \tau'_1 \wedge \text{lbl}(\tau'_1) = L$ and also $\overline{\gamma'_2}(lval) = \tau'_2 \wedge \text{lbl}(\tau'_2) = L$.

Given that the $lval$'s type is L in $\overline{\gamma'_2}$, and considering $\overline{\gamma'_2} \in \text{join}(\Gamma'_2)$, it follows that the $lval$ is also L in any state type within (Γ'_2) . Consequently, the $lval$ is L in Γ_k , thus L in $(\gamma'_{g2}, \gamma_{lk})$.

In the second configuration, we type the action's body s_2 with a H context, where s_2 reduces to (m'_{g2}, m'_{a2}) . And since we previously showed that $lval$ is L in $(\gamma'_{g2}, \gamma_{lk})$ such that $(\gamma'_{g2}, \gamma_{lk}) \vdash (m'_{g2}, m_{l2})$, then $lval$ is in m'_{g2} or m_{l2} , however not in m'_{a2} .

Since $lval$'s type is L in γ'_{g2} , we can use Lemma 10 to infer that the global initial and final states remain unchanged for L lvalues, which means $m_{g2}(lval) = m'_{g2}(lval)$. Then we can use Lemma 12 to infer that $\text{refine}(\gamma_2, \phi_k) \sqsubseteq (\gamma'_{g2}, \gamma'_{l2})$, this entails that the $lval$'s type label is indeed L in the global of refined state γ_{gk} . It is easy to see that $\gamma_2 \sqsubseteq \text{refine}(\gamma_2, \phi_k)$, thus $(\gamma_{g2}, \gamma_{l2}) \sqsubseteq (\gamma_{gk}, \gamma_{lk})$, therefore $lval$'s type is also L in the global initial state type γ_{g2} , i.e. $\gamma_{g2}(lval) = \tau'_2 \wedge \text{lbl}(\tau'_2) = L$.

For the first configuration, in this sub-case, we have $\overline{\gamma'_1}(lval) = \tau'_1 \wedge \text{lbl}(\tau'_1) = L$, and $\overline{\gamma'_1} \in \Gamma''_1$. We previously showed $\gamma'_1 \sqsubseteq \overline{\gamma'_1}$. Since $lval$'s typing label is L in $\overline{\gamma'_1}$ and we know that the state types in $\overline{\gamma'_1}$ are more restrictive than the state types in γ'_1 , we can conclude that $\gamma'_1 \in \Gamma_j$ also types $lval$ as L .

Now, we will prove that $m'_{g1} \xrightarrow[\gamma'_{g1} \sqcup \gamma'_{g2}]{} m'_{g2}$ by conducting

cases on ℓ_1 :

case If ℓ_1 is **H**:

We can replicate the same exact steps done for the second configuration to deduct $m_{g1}(lval) = m'_{g1}(lval)$ and $(\gamma_{g1}, \gamma_{l1}) \sqsubseteq (\gamma_{g_j}, \gamma_{l_j})$ and the $lval$'s type is **L** in the global initial state type γ_{g1} , i.e. $\gamma_{g1}(lval) = \tau'_2 \wedge \text{lbl}(\tau'_2) = \mathbf{L}$.

case If ℓ_1 is **L**:

Previously, we demonstrated that the $lval$'s typing label is **L** in all final state types in all Γ_i in Γ_2 . Consequently, none of the actions' bodies can modify $lval$, this is true because all actions in the first and second semantics and the contracts are identical, therefore $m_{g1}(lval) = m'_{g1}(lval)$. Thus, we know that indeed s_1 is typed under a high pc in the second configuration, we conclude that the $lval$ remains unchanged there as well. Consequently, we can now apply Lemma 13 to deduce that the $lval$'s typing label in the first configuration $\gamma'_1 \in \Gamma_j$ are more restrictive than the one we find in the refined state $\text{refine}(\gamma_1, \phi_j)$, and because $(\gamma_{g_j}, \gamma_{l_j}) = \text{refine}(\gamma_1, \phi_j)$ then $(\gamma_{g_j}, \gamma_{l_j})(lval) = \tau'_1 \wedge \text{lbl}(\tau'_1) = \mathbf{L}$. Therefore, $lval$'s type is also **L** in the global initial state type γ_{g2} , i.e. $\gamma_{g2}(lval) = \tau'_2 \wedge \text{lbl}(\tau'_2) = \mathbf{L}$.

Then, we need to prove $m_{l1} \xrightarrow[\gamma'_{l1} \sqcup \gamma'_{l2}]{} m_{l2}$. First, we prove

that $\overline{\gamma'_{l1}} = \gamma_{l_j}$ directly from the definition of join as $\text{join}(\gamma_{l_j}, \gamma_{l_j}) = \gamma_{l_j}$ and we know $\text{join}(\gamma_{l_j}, \gamma_{l_j}) = \overline{\gamma'_{l1}}$ thus $\overline{\gamma'_{l1}} = \gamma_{l_j}$ holds. Similarly, we know that $\overline{\gamma'_{l2}} = \gamma_{l_k}$ holds. Since $lval$ is **L** in $\overline{\gamma'_{l2}}$ then it is also **L** in γ_{l_k} . And since $lval$ is **L** in $\overline{\gamma'_{l2}}$, then it is also **L** in γ_{l_j} . From the previous subgoal, we proved $m_{l1} \xrightarrow[\gamma_{l_j} \sqcup \gamma_{l_k}]{} m_{l2}$. This is property hold.

Finally, since we proved $m'_{g1} \xrightarrow[\gamma'_{g1} \sqcup \gamma'_{g2}]{} m'_{g2}$ and

$m_{l1} \xrightarrow[\gamma'_{l1} \sqcup \gamma'_{l2}]{} m_{l2}$ now we can use Lemma 15 to deduct

that $(m'_{g1}, m_{l1}) \xrightarrow[\gamma'_1 \sqcup \gamma'_2]{} (m'_{g2}, m_{l2})$ holds.

◇ **Case transition:** similar to the conditional statement proof.

□