

LOFT — Verified Migration of Linux Firewalls to SDN

Julius Michaelis and Cornelius Diekmann

May 20, 2026

Abstract

We present LOFT — *Linux firewall OpenFlow Translator*, a system that transforms the main routing table and `FORWARD` chain of iptables of a Linux-based firewall into a set of static OpenFlow rules. Our implementation is verified against a model of a simplified Linux-based router and we can directly show how much of the original functionality is preserved.

Please note that this document is organized in two distinct parts. The first part contains the necessary definitions, helper lemmas and proofs in all their technicality as made in the theory code. The second part reiterates the most important definitions and proofs in a manner that is more suitable for human readers and enriches them with detailed explanations in natural language. Any interested reader should start from there.

Many of the considerations that have led to the definitions made here have been explained in [8].

Contents

I	Code	2
II	Documentation	25
1	Configuration Translation	25
1.1	Linux Firewall Model	25
1.1.1	Routing Table	26
1.1.2	iptables Firewall	27
1.2	OpenFlow Switch Model	27
1.2.1	Matching Flow Table entries	28
1.2.2	Evaluating a Flow Table	28
1.3	Translation Implementation	30
1.3.1	Chaining Firewalls	30
1.3.2	Translation Implementation	31
1.3.3	Comparison to Exodus	33
2	Evaluation	33
2.1	Mininet Examples	33
2.2	Performance Evaluation	35
3	Conclusion and Future Work	36

Part I

Code

```
theory OpenFlow-Matches
imports IP-Addresses.Prefix-Match
         Simple-Firewall.Simple-Packet
         HOL-Library.Monad-Syntax

         HOL-Library.List-Lexorder
         HOL-Library.Char-ord
begin

datatype of-match-field =
  IngressPort string
| EtherSrc 48 word
| EtherDst 48 word
| EtherType 16 word
| VlanId 16 word
| VlanPriority 16 word

| IPv4Src 32 prefix-match
| IPv4Dst 32 prefix-match
| IPv4Proto 8 word

| L4Src 16 word 16 word
| L4Dst 16 word 16 word

schematic-goal of-match-field-typeset: (field-match :: of-match-field) ∈ {
  IngressPort (?s::string),
  EtherSrc (?as::48 word), EtherDst (?ad::48 word),
  EtherType (?t::16 word),
  VlanId (?i::16 word), VlanPriority (?p::16 word),
  IPv4Src (?pms::32 prefix-match),
  IPv4Dst (?pmd::32 prefix-match),
  IPv4Proto (?ipp :: 8 word),
  L4Src (?ps :: 16 word) (?ms :: 16 word),
  L4Dst (?pd :: 16 word) (?md :: 16 word)
}
⟨proof⟩

function prerequisites :: of-match-field ⇒ of-match-field set ⇒ bool where
prerequisites (IngressPort -) - = True |

prerequisites (EtherDst -) - = True |

prerequisites (EtherSrc -) - = True |

prerequisites (EtherType -) - = True |
```

```

prerequisites (VlanId -) = True |
prerequisites (VlanPriority -) m = (∃ id. let v = VlanId id in v ∈ m ∧ prerequisites v m) |
prerequisites (IPv4Proto -) m = (let v = EtherType 0x0800 in v ∈ m ∧ prerequisites v m) |
prerequisites (IPv4Src -) m = (let v = EtherType 0x0800 in v ∈ m ∧ prerequisites v m) |
prerequisites (IPv4Dst -) m = (let v = EtherType 0x0800 in v ∈ m ∧ prerequisites v m) |

prerequisites (L4Src - -) m = (∃ proto ∈ {TCP,UDP,L4-Protocol.SCTP}. let v = IPv4Proto proto in v ∈ m ∧ prerequisites v m) |
prerequisites (L4Dst - -) m = prerequisites (L4Src undefined undefined) m
⟨proof⟩

```

fun match-sorter :: of-match-field ⇒ nat **where**

```

match-sorter (IngressPort -) = 1 |
match-sorter (VlanId -) = 2 |
match-sorter (VlanPriority -) = 3 |
match-sorter (EtherType -) = 4 |
match-sorter (EtherSrc -) = 5 |
match-sorter (EtherDst -) = 6 |
match-sorter (IPv4Proto -) = 7 |
match-sorter (IPv4Src -) = 8 |
match-sorter (IPv4Dst -) = 9 |
match-sorter (L4Src - -) = 10 |
match-sorter (L4Dst - -) = 11

```

termination prerequisites ⟨proof⟩

definition less-eq-of-match-field1 :: of-match-field ⇒ of-match-field ⇒ bool

where less-eq-of-match-field1 (a::of-match-field) (b::of-match-field) ←→ (case (a, b) of
(IngressPort a, IngressPort b) ⇒ a ≤ b |
(VlanId a, VlanId b) ⇒ a ≤ b |
(EtherDst a, EtherDst b) ⇒ a ≤ b |
(EtherSrc a, EtherSrc b) ⇒ a ≤ b |
(EtherType a, EtherType b) ⇒ a ≤ b |
(VlanPriority a, VlanPriority b) ⇒ a ≤ b |
(IPv4Proto a, IPv4Proto b) ⇒ a ≤ b |
(IPv4Src a, IPv4Src b) ⇒ a ≤ b |
(IPv4Dst a, IPv4Dst b) ⇒ a ≤ b |
(L4Src a1 a2, L4Src b1 b2) ⇒ if a2 = b2 then a1 ≤ b1 else a2 ≤ b2 |
(L4Dst a1 a2, L4Dst b1 b2) ⇒ if a2 = b2 then a1 ≤ b1 else a2 ≤ b2 |
(a, b) ⇒ match-sorter a < match-sorter b)

instantiation of-match-field :: linorder

begin

definition

less-eq-of-match-field (a::of-match-field) (b::of-match-field) ←→ less-eq-of-match-field1 a b

definition

less-of-match-field (*a::of-match-field*) (*b::of-match-field*) $\longleftrightarrow a \neq b \wedge \text{less-eq-of-match-field1 } a \ b$

instance

<proof>

end

fun *match-no-prereq* :: *of-match-field* \Rightarrow (*32*, 'a) *simple-packet-ext-scheme* \Rightarrow *bool* **where**

match-no-prereq (*IngressPort* *i*) *p* = (*p-iiiface* *p* = *i*) |
match-no-prereq (*EtherDst* *i*) *p* = (*p-l2src* *p* = *i*) |
match-no-prereq (*EtherSrc* *i*) *p* = (*p-l2dst* *p* = *i*) |
match-no-prereq (*EtherType* *i*) *p* = (*p-l2type* *p* = *i*) |
match-no-prereq (*VlanId* *i*) *p* = (*p-vlanid* *p* = *i*) |
match-no-prereq (*VlanPriority* *i*) *p* = (*p-vlanprio* *p* = *i*) |
match-no-prereq (*IPv4Proto* *i*) *p* = (*p-proto* *p* = *i*) |
match-no-prereq (*IPv4Src* *i*) *p* = (*prefix-match-semantic* *i* (*p-src* *p*)) |
match-no-prereq (*IPv4Dst* *i*) *p* = (*prefix-match-semantic* *i* (*p-dst* *p*)) |
match-no-prereq (*L4Src* *i* *m*) *p* = (*p-sport* *p* && *m* = *i*) |
match-no-prereq (*L4Dst* *i* *m*) *p* = (*p-dport* *p* && *m* = *i*)

definition *match-prereq* :: *of-match-field* \Rightarrow *of-match-field set* \Rightarrow (*32*, 'a) *simple-packet-ext-scheme* \Rightarrow *bool option* **where**
match-prereq *i* *s* *p* = (if *prerequisites* *i* *s* then *Some* (*match-no-prereq* *i* *p*) else *None*)

definition *set-seq* *s* \equiv if ($\forall x \in s. x \neq \text{None}$) then *Some* (*the* ' *s*) else *None*

definition *all-true* *s* $\equiv \forall x \in s. x$

term *map-option*

definition *OF-match-fields* :: *of-match-field set* \Rightarrow (*32*, 'a) *simple-packet-ext-scheme* \Rightarrow *bool option* **where** *OF-match-fields*
m *p* = *map-option* *all-true* (*set-seq* (($\lambda f. \text{match-prereq } f \ m \ p$) ' *m*))

definition *OF-match-fields-unsafe* :: *of-match-field set* \Rightarrow (*32*, 'a) *simple-packet-ext-scheme* \Rightarrow *bool* **where**
OF-match-fields-unsafe *m* *p* = ($\forall f \in m. \text{match-no-prereq } f \ p$)

definition *OF-match-fields-safe* *m* \equiv *the* \circ *OF-match-fields* *m*

definition *all-prerequisites* *m* $\equiv \forall f \in m. \text{prerequisites } f \ m$

lemma

all-prerequisites *p* \implies

L4Src *x* *y* $\in p \implies$

IPv4Proto ' {*TCP*, *UDP*, *L4-Protocol.SCTP*} $\cap p \neq \{\}$

<proof>

lemma *of-safe-unsafe-match-eq*: *all-prerequisites* *m* \implies *OF-match-fields* *m* *p* = *Some* (*OF-match-fields-unsafe* *m* *p*)
<proof>

lemma *of-match-fields-safe-eq*: **assumes** *all-prerequisites* *m* **shows** *OF-match-fields-safe* *m* = *OF-match-fields-unsafe* *m*
<proof>

lemma *OF-match-fields-alt*: *OF-match-fields* *m* *p* =

(if $\exists f \in m. \neg \text{prerequisites } f \ m$ then *None* else

if $\forall f \in m. \text{match-no-prereq } f \ p$ then *Some True* else *Some False*)

<proof>

lemma *of-match-fields-safe-eq2*: **assumes** *all-prerequisites* *m* **shows** *OF-match-fields-safe* *m* *p* \longleftrightarrow *OF-match-fields* *m* *p* = *Some True*

<proof>

end
theory *OpenFlow-Action*
imports
 OpenFlow-Matches
begin

datatype *of-action* = *Forward* (*oiface-sel*: *string*) | *ModifyField-l2dst* 48 *word*

fun *of-action-semantic* **where**
of-action-semantic *p* [] = {} |
of-action-semantic *p* (*a#as*) = (*case a of*
 Forward i ⇒ *insert* (*i,p*) (*of-action-semantic* *p as*) |
 ModifyField-l2dst a ⇒ *of-action-semantic* (*p*(*p-l2dst := a*)) *as*)

value *of-action-semantic* *p* []
value *of-action-semantic* *p* [*ModifyField-l2dst* 66, *Forward* "*oif*"]

end
theory *Semantics-OpenFlow*
imports *List-Group Sort-Descending*
 IP-Addresses.IPv4
 OpenFlow-Helpers
begin

datatype *'a flowtable-behavior* = *Action* *'a* | *NoAction* | *Undefined*

definition *option-to-ftb* *b* ≡ *case b of Some a* ⇒ *Action a* | *None* ⇒ *NoAction*
definition *ftb-to-option* *b* ≡ *case b of Action a* ⇒ *Some a* | *NoAction* ⇒ *None*

datatype (*'m, 'a*) *flow-entry-match* = *OFEntry* (*ofe-prio*: 16 *word*) (*ofe-fields*: *'m set*) (*ofe-action*: *'a*)

find-consts ((*'a* × *'b*) ⇒ *'c*) ⇒ *'a* ⇒ *'b* ⇒ *'c*

find-consts ('a ⇒ 'b ⇒ 'c) ⇒ ('a × 'b) ⇒ 'c

definition *split3* f p ≡ case p of (a,b,c) ⇒ f a b c

find-consts ('a ⇒ 'b ⇒ 'c ⇒ 'd) ⇒ ('a × 'b × 'c) ⇒ 'd

type-synonym ('m, 'a) flowtable = (('m, 'a) flow-entry-match) list

type-synonym ('m, 'p) field-matcher = ('m set ⇒ 'p ⇒ bool)

definition *OF-same-priority-match2* :: ('m, 'p) field-matcher ⇒ ('m, 'a) flowtable ⇒ 'p ⇒ 'a flowtable-behavior **where**

OF-same-priority-match2 γ flow-entries packet ≡ let s =

{ofe-action f | f. f ∈ set flow-entries ∧ γ (ofe-fields f) packet ∧

(∀ fo ∈ set flow-entries. ofe-prio fo > ofe-prio f → ¬γ (ofe-fields fo) packet)} in

case card s of 0 ⇒ NoAction

| (Suc 0) ⇒ Action (the-elem s)

| - ⇒ Undefined

definition *check-no-overlap* γ ft = (∀ a ∈ set ft. ∀ b ∈ set ft. ∀ p ∈ UNIV. (ofe-prio a = ofe-prio b ∧ γ (ofe-fields a) p ∧ a ≠ b) → ¬γ (ofe-fields b) p)

definition *check-no-overlap2* γ ft = (∀ a ∈ set ft. ∀ b ∈ set ft. (a ≠ b ∧ ofe-prio a = ofe-prio b) → ¬(∃ p ∈ UNIV. γ (ofe-fields a) p ∧ γ (ofe-fields b) p))

lemma *check-no-overlap-alt*: *check-no-overlap* γ ft = *check-no-overlap2* γ ft

⟨proof⟩

lemma *no-overlap-not-undefined*: *check-no-overlap* γ ft ⇒ *OF-same-priority-match2* γ ft p ≠ Undefined

⟨proof⟩

fun *OF-match-linear* :: ('m, 'p) field-matcher ⇒ ('m, 'a) flowtable ⇒ 'p ⇒ 'a flowtable-behavior **where**

OF-match-linear - [] - = NoAction |

OF-match-linear γ (a#as) p = (if γ (ofe-fields a) p then Action (ofe-action a) else *OF-match-linear* γ as p)

lemma *OF-match-linear-ne-Undefined*: *OF-match-linear* γ ft p ≠ Undefined

⟨proof⟩

lemma *OF-match-linear-append*: *OF-match-linear* γ (a @ b) p = (case *OF-match-linear* γ a p of NoAction ⇒ *OF-match-linear* γ b p | x ⇒ x)

⟨proof⟩

lemma *OF-match-linear-match-allsameaction*: [gr ∈ set oms; γ gr p = True]

⇒ *OF-match-linear* γ (map (λx. split3 OFEntry (pri, x, act)) oms) p = Action act

⟨proof⟩

lemma *OF-lm-noa-none-iff*: *OF-match-linear* γ ft p = NoAction ↔ (∀ e ∈ set ft. ¬ γ (ofe-fields e) p)

⟨proof⟩

lemma *set-eq-rule*: (∧x. x ∈ a ⇒ x ∈ b) ⇒ (∧x. x ∈ b ⇒ x ∈ a) ⇒ a = b ⟨proof⟩

lemma *unmatching-insert-agnostic*: ¬ γ (ofe-fields a) p ⇒ *OF-same-priority-match2* γ (a # ft) p = *OF-same-priority-match2* γ ft p

⟨proof⟩

lemma *OF-match-eq: sorted-descending* (map ofe-prio ft) \implies check-no-overlap γ ft \implies
OF-same-priority-match2 γ ft p = *OF-match-linear* γ ft p
 ⟨proof⟩

lemma *overlap-sort-invar[simp]*: check-no-overlap γ (sort-descending-key k ft) = check-no-overlap γ ft
 ⟨proof⟩

lemma *OF-match-eq2*:
assumes check-no-overlap γ ft
shows *OF-same-priority-match2* γ ft p = *OF-match-linear* γ (sort-descending-key ofe-prio ft) p
 ⟨proof⟩

lemma *prio-match-matcher-alt*: {f. f \in set flow-entries \wedge γ (ofe-fields f) packet \wedge
 (\forall fo \in set flow-entries. ofe-prio fo > ofe-prio f \implies \neg γ (ofe-fields fo) packet)}
 = (
 let matching = {f. f \in set flow-entries \wedge γ (ofe-fields f) packet}
 in {f. f \in matching \wedge (\forall fo \in matching. ofe-prio fo \leq ofe-prio f)}
)
 ⟨proof⟩

lemma *prio-match-matcher-alt2*: (
 let matching = {f. f \in set flow-entries \wedge γ (ofe-fields f) packet}
 in {f. f \in matching \wedge (\forall fo \in matching. ofe-prio fo \leq ofe-prio f)}
) = set (
 let matching = filter (λ f. γ (ofe-fields f) packet) flow-entries
 in filter (λ f. \forall fo \in set matching. ofe-prio fo \leq ofe-prio f) matching
)
 ⟨proof⟩

definition *OF-priority-match where*

OF-priority-match γ flow-entries packet \equiv
 let m = filter (λ f. γ (ofe-fields f) packet) flow-entries;
 m' = filter (λ f. \forall fo \in set m. ofe-prio fo \leq ofe-prio f) m in
 case m' of [] \Rightarrow NoAction
 | [s] \Rightarrow Action (ofe-action s)
 | - \Rightarrow Undefined

definition *OF-priority-match-ana where*

OF-priority-match-ana γ flow-entries packet \equiv
 let m = filter (λ f. γ (ofe-fields f) packet) flow-entries;
 m' = filter (λ f. \forall fo \in set m. ofe-prio fo \leq ofe-prio f) m in
 case m' of [] \Rightarrow NoAction
 | [s] \Rightarrow Action s
 | - \Rightarrow Undefined

lemma *filter-singleton*: [x \leftarrow s. f x] = [y] \implies f y \wedge y \in set s ⟨proof⟩

lemma *OF-spm3-get-fe*: *OF-priority-match* γ ft p = Action a \implies \exists fe. ofe-action fe = a \wedge fe \in set ft \wedge *OF-priority-match-ana*
 γ ft p = Action fe
 ⟨proof⟩

fun *no-overlaps where*

no-overlaps - [] = True |
no-overlaps γ (a#as) = (no-overlaps γ as \wedge (

$\forall b \in \text{set as. ofe-prio } a = \text{ofe-prio } b \longrightarrow \neg(\exists p \in \text{UNIV. } \gamma (\text{ofe-fields } a) p \wedge \gamma (\text{ofe-fields } b) p))$

lemma *no-overlap-ConsI*: $\text{check-no-overlap2 } \gamma (x\#xs) \Longrightarrow \text{check-no-overlap2 } \gamma xs$
(*proof*)

lemma *no-overlapsI*: $\text{check-no-overlap } \gamma t \Longrightarrow \text{distinct } t \Longrightarrow \text{no-overlaps } \gamma t$
(*proof*)

lemma *check-no-overlapI*: $\text{no-overlaps } \gamma t \Longrightarrow \text{check-no-overlap } \gamma t$
(*proof*)

lemma $(\bigwedge e p. e \in \text{set } t \Longrightarrow \neg\gamma (\text{ofe-fields } e) p) \Longrightarrow \text{no-overlaps } \gamma t$
(*proof*)

lemma *no-overlaps-append*: $\text{no-overlaps } \gamma (x @ y) \Longrightarrow \text{no-overlaps } \gamma y$
(*proof*)

lemma *no-overlaps-ne1*: $\text{no-overlaps } \gamma (x @ a \# y @ b \# z) \Longrightarrow ((\exists p. \gamma (\text{ofe-fields } a) p) \vee (\exists p. \gamma (\text{ofe-fields } b) p)) \Longrightarrow a \neq b$
(*proof*)

lemma *no-overlaps-defeq*: $\text{no-overlaps } \gamma fe \Longrightarrow \text{OF-same-priority-match2 } \gamma fe p = \text{OF-priority-match } \gamma fe p$
(*proof*)

lemma *distinct fe*: $\text{distinct } fe \Longrightarrow \text{check-no-overlap } \gamma fe \Longrightarrow \text{OF-same-priority-match2 } \gamma fe p = \text{OF-priority-match } \gamma fe p$
(*proof*)

theorem *OF-eq*:

assumes *no*: $\text{no-overlaps } \gamma f$

and *so*: *sorted-descending* (*map ofe-prio f*)

shows $\text{OF-match-linear } \gamma f p = \text{OF-priority-match } \gamma f p$
(*proof*)

corollary *OF-eq-sort*:

assumes *no*: $\text{no-overlaps } \gamma f$

shows $\text{OF-priority-match } \gamma f p = \text{OF-match-linear } \gamma (\text{sort-descending-key ofe-prio } f) p$

(*proof*)

lemma *OF-lm-noa-none*: $\text{OF-match-linear } \gamma ft p = \text{NoAction} \Longrightarrow \forall e \in \text{set } ft. \neg \gamma (\text{ofe-fields } e) p$
(*proof*)

lemma *OF-spm3-noa-none*:

assumes *no*: $\text{no-overlaps } \gamma ft$

shows $\text{OF-priority-match } \gamma ft p = \text{NoAction} \Longrightarrow \forall e \in \text{set } ft. \neg \gamma (\text{ofe-fields } e) p$

(*proof*)

lemma *no-overlaps-not-undefined*: $\text{no-overlaps } \gamma ft \Longrightarrow \text{OF-priority-match } \gamma ft p \neq \text{Undefined}$
(*proof*)

end

theory *OpenFlow-Serialize*

imports *OpenFlow-Matches*

OpenFlow-Action

Semantics-OpenFlow

Simple-Firewall.Primitives-toString
IP-Addresses.Lib-Word-toString

begin

definition *serialization-test-entry* \equiv *OFEntry* 7 {*EtherDst* 0x1, *IPv4Dst* (*PrefixMatch* 0xA00201 32), *IngressPort* "s1-lan", *L4Dst* 0x50 0, *L4Src* 0x400 0x3FF, *IPv4Proto* 6, *EtherType* 0x800} [*ModifyField-l2dst* 0xA641F185E862, *Forward* "s1-wan"]

value (*map* ((<<) (1::48 word) \circ (*) 8) \circ *rev*) [0..<6]

definition *serialize-mac* (*m*::48 word) \equiv (*intersperse* (*CHR* "'") \circ *map* (*hex-string-of-word* 1 \circ ($\lambda h. (m \gg h * 8)$) && 0xff)) \circ *rev*) [0..<6]

lemma *serialize-mac* 0xdeadbeefcafe = "de:ad:be:ef:ca:fe" <proof>

definition *serialize-action pids a* \equiv (case a of
 Forward oif \Rightarrow "output:" @ pids oif |
 ModifyField-l2dst na \Rightarrow "mod-dl-dst:" @ *serialize-mac* na)

definition *serialize-actions pids a* \equiv if length a = 0 then "drop" else (*intersperse* (*CHR* "'") \circ *map* (*serialize-action pids*)) a

lemma *serialize-actions* (λ oif. "42") (*ofe-action* *serialization-test-entry*) =
 "mod-dl-dst:a6:41:f1:85:e8:62,output:42" <proof>

lemma *serialize-actions* anything [] = "drop"
 <proof>

definition *prefix-to-string pfx* \equiv *ipv4-cidr-toString* (*pfxm-prefix* pfx, *pfxm-length* pfx)

primrec *serialize-of-match* **where**

serialize-of-match pids (*IngressPort* p) = "in-port=" @ pids p |
 serialize-of-match - (*VlanId* i) = "dl-vlan=" @ *dec-string-of-word0* i |
 serialize-of-match - (*VlanPriority* -) = undefined |
 serialize-of-match - (*EtherType* i) = "dl-type=0x" @ *hex-string-of-word0* i |
 serialize-of-match - (*EtherSrc* m) = "dl-src=" @ *serialize-mac* m |
 serialize-of-match - (*EtherDst* m) = "dl-dst=" @ *serialize-mac* m |
 serialize-of-match - (*IPv4Proto* i) = "nw-proto=" @ *dec-string-of-word0* i |
 serialize-of-match - (*IPv4Src* p) = "nw-src=" @ *prefix-to-string* p |
 serialize-of-match - (*IPv4Dst* p) = "nw-dst=" @ *prefix-to-string* p |
 serialize-of-match - (*L4Src* i m) = "tp-src=" @ *dec-string-of-word0* i @ (if m = - 1 then [] else "/0x" @ *hex-string-of-word* 3 m) |
 serialize-of-match - (*L4Dst* i m) = "tp-dst=" @ *dec-string-of-word0* i @ (if m = - 1 then [] else "/0x" @ *hex-string-of-word* 3 m)

definition *serialize-of-matches* :: (*string* \Rightarrow *string*) \Rightarrow *of-match-field set* \Rightarrow *string*

where

serialize-of-matches pids \equiv (@) "hard-timeout=0,idle-timeout=0," \circ *intersperse* (*CHR* "'") \circ *map* (*serialize-of-match pids*)
 \circ *sorted-list-of-set*

lemma *serialize-of-matches pids of-matches* =
 (*List.append* "hard-timeout=0,idle-timeout=0,"
 (*intersperse* (*CHR* "'") (*map* (*serialize-of-match pids*) (*sorted-list-of-set of-matches*))))
 <proof>

export-code *serialize-of-matches* **checking** *SML*

lemma *serialize-of-matches* ($\lambda oif. "42"$) (*ofe-fields serialization-test-entry*) =

"hard-timeout=0,idle-timeout=0,in-port=42,dl-type=0x800,dl-dst=00:00:00:00:01,nw-prot=6,nw-dst=10.0.2.1/32,tp-src=1024
<proof>

definition *serialize-of-entry* *pids e* \equiv (*case e of* (*OFEntry p f a*) \Rightarrow "priority=" @ *dec-string-of-word0 p* @ ", " @ *serialize-of-matches pids f* @ ", " @ "action=" @ *serialize-actions pids a*)

lemma *serialize-of-entry* (*the* \circ *map-of* [(*s1-lan*, "42"),(*s1-wan*, "1337")]) *serialization-test-entry* =

"priority=7,hard-timeout=0,idle-timeout=0,in-port=42,dl-type=0x800,dl-dst=00:00:00:00:01,nw-prot=6,nw-dst=10.0.2.1/32,tp-src=1024
<proof>

end

theory *Featherweight-OpenFlow-Comparison*

imports *Semantics-OpenFlow*

begin

inductive *guha-table-semantics* :: ('*m*, '*p*) *field-matcher* \Rightarrow ('*m*, '*a*) *flowtable* \Rightarrow '*p* \Rightarrow '*a* *option* \Rightarrow **bool** **where**

guha-matched: γ (*ofe-fields fe*) *p* = **True** \Longrightarrow

$\forall fe' \in \text{set } (ft1 @ ft2). \text{ofe-prio } fe' > \text{ofe-prio } fe \longrightarrow \gamma$ (*ofe-fields fe'*) *p* = **False** \Longrightarrow

guha-table-semantics γ (*ft1* @ *fe* # *ft2*) *p* (*Some* (*ofe-action fe*)) |

guha-unmatched: $\forall fe \in \text{set } ft. \gamma$ (*ofe-fields fe*) *p* = **False** \Longrightarrow

guha-table-semantics γ *ft* *p* **None**

lemma *guha-table-semantics-ex2res*:

assumes *ta*: $CARD('a) \geq 2$

assumes *ms*: $\exists ff. \gamma$ *ff* *p*

shows $\exists ft (a1 :: 'a) (a2 :: 'a). a1 \neq a2 \wedge \text{guha-table-semantics } \gamma$ *ft* *p* (*Some* *a1*) \wedge *guha-table-semantics* γ *ft* *p* (*Some* *a2*)
<proof>

lemma *guha-umstaendlich*:

assumes *ae*: *a* = *ofe-action fe*

assumes *ele*: *fe* \in *set ft*

assumes *rest*: γ (*ofe-fields fe*) *p*

$\forall fe' \in \text{set } ft. \text{ofe-prio } fe' > \text{ofe-prio } fe \longrightarrow \neg \gamma$ (*ofe-fields fe'*) *p*

shows *guha-table-semantics* γ *ft* *p* (*Some* *a*)

<proof>

lemma *guha-matched-rule-inversion*:

assumes *guha-table-semantics* γ *ft* *p* (*Some* *a*)

shows $\exists fe \in \text{set } ft. a = \text{ofe-action } fe \wedge \gamma$ (*ofe-fields fe*) *p* \wedge ($\forall fe' \in \text{set } ft. \text{ofe-prio } fe' > \text{ofe-prio } fe \longrightarrow \neg \gamma$ (*ofe-fields fe'*)
p)

<proof>

lemma *guha-equal-Action*:

assumes *no*: *no-overlaps* γ *ft*

assumes *spm*: *OF-priority-match* γ *ft* *p* = *Action a*

shows *guha-table-semantics* γ *ft* *p* (*Some* *a*)

<proof>

lemma *guha-equal-NoAction*:

assumes *no*: *no-overlaps* γ *ft*
assumes *spm*: *OF-priority-match* γ *ft* *p* = *NoAction*
shows *guha-table-semantic*s γ *ft* *p* *None*
<proof>

lemma *guha-equal-hlp*:

assumes *no*: *no-overlaps* γ *ft*
shows *guha-table-semantic*s γ *ft* *p* (*ftb-to-option* (*OF-priority-match* γ *ft* *p*))
<proof>

lemma *guha-deterministic1*: *guha-table-semantic*s γ *ft* *p* (*Some* *x1*) \implies \neg *guha-table-semantic*s γ *ft* *p* *None*
<proof>

lemma *guha-deterministic2*: \llbracket *no-overlaps* γ *ft*; *guha-table-semantic*s γ *ft* *p* (*Some* *x1*); *guha-table-semantic*s γ *ft* *p* (*Some* *a*) \rrbracket
 \implies *x1* = *a*
<proof>

lemma *guha-equal*:

assumes *no*: *no-overlaps* γ *ft*
shows *OF-priority-match* γ *ft* *p* = *option-to-ftb* *d* \iff *guha-table-semantic*s γ *ft* *p* *d*
<proof>

lemma *guha-nondeterministicD*:

assumes \neg *check-no-overlap* γ *ft*
shows \exists *fe1* *fe2* *p*. *fe1* \in *set* *ft* \wedge *fe2* \in *set* *ft*
 \wedge *guha-table-semantic*s γ *ft* *p* (*Some* (*ofe-action* *fe1*))
 \wedge *guha-table-semantic*s γ *ft* *p* (*Some* (*ofe-action* *fe2*))
<proof>

The above lemma does indeed not hold, the reason for this are (possibly partially) shadowed overlaps. This is exemplified below: If there are at least three different possible actions (necessary assumption) and a match expression that matches all packets (convenience assumption), it is possible to construct a flow table that is admonished by *check-no-overlap* but still will never run into undefined behavior.

lemma

assumes $CARD('action) \geq 3$
assumes $\forall p. \gamma x p$
shows $\exists ft::('a, 'action) flow-entry-match list. \neg check-no-overlap \gamma ft \wedge$
 $\neg(\exists fe1 fe2 p. fe1 \in set ft \wedge fe2 \in set ft \wedge fe1 \neq fe2 \wedge ofe-prio fe1 = ofe-prio fe2$
 $\wedge guha-table-semantic \gamma ft p (Some (ofe-action fe1))$
 $\wedge guha-table-semantic \gamma ft p (Some (ofe-action fe2)))$
<proof>

end

theory *LinuxRouter-OpenFlow-Translation*

imports *IP-Addresses.CIDR-Split*

Automatic-Refinement.Misc

Simple-Firewall.Generic-SimpleFw

Semantics-OpenFlow

OpenFlow-Matches

OpenFlow-Action

Routing.Linux-Router

Pure-ex.Guess

begin

hide-const *Misc.uncurry*

hide-fact *Misc.uncurry-def*

definition *route2match* $r =$

$(iiface = ifaceAny, oiface = ifaceAny,$

$src = (0,0), dst=(pfxm-prefix (routing-match r),pfxm-length (routing-match r)),$

$proto=ProtoAny, sports=(0,- 1), ports=(0,- 1))$

definition *toprefixmatch* **where**

toprefixmatch $m \equiv (let\ pm = PrefixMatch\ (fst\ m)\ (snd\ m)\ in\ if\ pm = PrefixMatch\ 0\ 0\ then\ None\ else\ Some\ pm)$

lemma *prefix-match-semantic-simple-match:*

assumes *some:* *toprefixmatch* $m = Some\ pm$

assumes *vld:* *valid-prefix* pm

shows *prefix-match-semantic* $pm = simple-match-ip\ m$

$\langle proof \rangle$

definition *simple-match-to-of-match-single* ::

$(32, 'a)\ simple-match-scheme$

$\Rightarrow char\ list\ option \Rightarrow protocol \Rightarrow (16\ word \times 16\ word)\ option \Rightarrow (16\ word \times 16\ word)\ option \Rightarrow of-match-field\ set$

where

simple-match-to-of-match-single $m\ iif\ prot\ sport\ dport \equiv$

$uncurry\ L4Src\ 'option2set\ sport \cup\ uncurry\ L4Dst\ 'option2set\ dport$

$\cup\ IPv4Proto\ ' (case\ prot\ of\ ProtoAny \Rightarrow \{\} \mid Proto\ p \Rightarrow \{p\})$ — protocol is an 8 word option anyway...

$\cup\ IngressPort\ 'option2set\ iif$

$\cup\ IPv4Src\ 'option2set\ (toprefixmatch\ (src\ m)) \cup\ IPv4Dst\ 'option2set\ (toprefixmatch\ (dst\ m))$

$\cup\ \{EtherType\ 0x0800\}$

definition *simple-match-to-of-match* :: $32\ simple-match \Rightarrow string\ list \Rightarrow of-match-field\ set\ list$ **where**

simple-match-to-of-match $m\ ifs \equiv (let$

$npm = (\lambda p. fst\ p = 0 \wedge snd\ p = - 1);$

$sb = (\lambda p. (if\ npm\ p\ then\ [None]\ else\ if\ fst\ p \leq\ snd\ p$

$then\ map\ (Some \circ (\lambda pfx. (pfxm-prefix\ pfx, Bit-Operations.not\ (pfxm-mask\ pfx))))\ (wordinterval-CIDR-split-prefixmatch$
 $(WordInterval\ (fst\ p)\ (snd\ p)))\ else\ []))$

$in\ [simple-match-to-of-match-single\ m\ iif\ (proto\ m)\ sport\ dport.$

$iif \leftarrow (if\ iiface\ m = ifaceAny\ then\ [None]\ else\ [Some\ i. i \leftarrow ifs, match-iface\ (iiface\ m)\ i]),$

$sport \leftarrow sb\ (sports\ m),$

$dport \leftarrow sb\ (dports\ m)]$

)

lemma *smtoms-eq-hlp:* *simple-match-to-of-match-single* $r\ a\ b\ c\ d = simple-match-to-of-match-single\ r\ f\ g\ h\ i \iff (a = f \wedge b = g \wedge c = h \wedge d = i)$

$\langle proof \rangle$

lemma *simple-match-to-of-match-generates-prereqs:* *simple-match-valid* $m \implies r \in set\ (simple-match-to-of-match\ m\ ifs) \implies all-prerequisites\ r$

$\langle proof \rangle$

lemma *and-assoc*: $a \wedge b \wedge c \longleftrightarrow (a \wedge b) \wedge c$ *<proof>*

lemmas *custom-simpset* = *Let-def set-concat set-map map-map comp-def UN-iff fun-app-def Set.image-iff image-image*

abbreviation *simple-fw-prefix-to-wordinterval* \equiv *prefix-to-wordinterval* \circ *uncurry PrefixMatch*

lemma *simple-match-port-alt*: *simple-match-port* m $p \longleftrightarrow p \in$ *wordinterval-to-set* (*uncurry WordInterval* m) *<proof>*

lemma *simple-match-src-alt*: *simple-match-valid* $r \implies$
simple-match-ip (*src* r) $p \longleftrightarrow$ *prefix-match-semantic* (*PrefixMatch* (*fst* (*src* r)) (*snd* (*src* r))) p
<proof>

lemma *simple-match-dst-alt*: *simple-match-valid* $r \implies$
simple-match-ip (*dst* r) $p \longleftrightarrow$ *prefix-match-semantic* (*PrefixMatch* (*fst* (*dst* r)) (*snd* (*dst* r))) p
<proof>

lemma $x \in$ *set* (*wordinterval-CIDR-split-prefixmatch* w) \implies *valid-prefix* x
<proof>

lemma *simple-match-to-of-matchI*:
assumes *mv*: *simple-match-valid* r
assumes *mm*: *simple-matches* r p
assumes *ii*: *p-iiface* $p \in$ *set ifs*
assumes *ippkt*: *p-l2type* $p = 0x800$
shows *eq*: $\exists gr \in$ *set* (*simple-match-to-of-match* r *ifs*). *OF-match-fields* gr $p =$ *Some True*
<proof>

lemma *prefix-match-00*[*simp,intro!*]: *prefix-match-semantic* (*PrefixMatch* 0 0) p
<proof>

lemma *simple-match-to-of-matchD*:
assumes *eg*: $gr \in$ *set* (*simple-match-to-of-match* r *ifs*)
assumes *mo*: *OF-match-fields* gr $p =$ *Some True*
assumes *me*: *match-iface* (*oiface* r) (*p-oiface* p)
assumes *mv*: *simple-match-valid* r
shows *simple-matches* r p
<proof>

primrec *annotate-rlen* **where**

annotate-rlen $[] = []$ |

annotate-rlen ($a\#as$) = (*length* as , a) $\#$ *annotate-rlen* as

lemma *annotate-rlen "asdf"* = [(3, *CHR* "a"), (2, *CHR* "s"), (1, *CHR* "d"), (0, *CHR* "f")] *<proof>*

lemma *fst-annotate-rlen-le*: $(k, a) \in$ *set* (*annotate-rlen* l) $\implies k <$ *length* l
<proof>

lemma *distinct-fst-annotate-rlen*: *distinct* (*map* *fst* (*annotate-rlen* l))
<proof>

lemma *distinct-annotate-rlen*: *distinct* (*annotate-rlen* l)
<proof>

lemma *in-annotate-rlen*: $(a,x) \in$ *set* (*annotate-rlen* l) $\implies x \in$ *set* l
<proof>

lemma *map-snd-annotate-rlen*: *map* *snd* (*annotate-rlen* l) = l
<proof>

lemma *sorted-descending* (*map* *fst* (*annotate-rlen* l))

<proof>
lemma *annotate-rlen* $l = \text{zip } (\text{rev } [0..<\text{length } l]) \ l$
<proof>

primrec *annotate-rlen-code* **where**
annotate-rlen-code $[] = (0, []) \mid$
annotate-rlen-code $(a\#as) = (\text{case } \text{annotate-rlen-code } as \text{ of } (r, aas) \Rightarrow (\text{Suc } r, (r, a) \# aas))$
lemma *annotate-rlen-len*: $\text{fst } (\text{annotate-rlen-code } r) = \text{length } r$
<proof>
lemma *annotate-rlen-code[code]*: $\text{annotate-rlen } s = \text{snd } (\text{annotate-rlen-code } s)$
<proof>

lemma *suc2plus-inj-on*: $\text{inj-on } (\text{of-nat} :: \text{nat} \Rightarrow ('l :: \text{len}) \text{ word}) \ \{0.. \text{unat } (\text{max-word} :: 'l \text{ word})\}$
<proof>

lemma *distinct-of-nat-list*:
 $\text{distinct } l \Longrightarrow \forall e \in \text{set } l. e \leq \text{unat } (\text{max-word} :: ('l :: \text{len}) \text{ word}) \Longrightarrow \text{distinct } (\text{map } (\text{of-nat} :: \text{nat} \Rightarrow 'l \text{ word}) \ l)$
<proof>

lemma *annotate-first-le-hlp*:
 $\text{length } l < \text{unat } (\text{max-word} :: ('l :: \text{len}) \text{ word}) \Longrightarrow \forall e \in \text{set } (\text{map } \text{fst } (\text{annotate-rlen } l)). e \leq \text{unat } (\text{max-word} :: 'l \text{ word})$
<proof>
lemmas *distinct-of-prio-hlp* = *distinct-of-nat-list*[OF *distinct-fst-annotate-rlen* *annotate-first-le-hlp*]

lemma *fst-annotate-rlen*: $\text{map } \text{fst } (\text{annotate-rlen } l) = \text{rev } [0..<\text{length } l]$
<proof>

lemma *sorted-word-upt*:
defines[*simp*]: $\text{won} \equiv (\text{of-nat} :: \text{nat} \Rightarrow ('l :: \text{len}) \text{ word})$
assumes $\text{length } l \leq \text{unat } (\text{max-word} :: 'l \text{ word})$
shows *sorted-descending* $(\text{map } \text{won } (\text{rev } [0..<\text{Suc } (\text{length } l)]))$
<proof>

lemma *sorted-annotated*:
assumes $\text{length } l \leq \text{unat } (\text{max-word} :: ('l :: \text{len}) \text{ word})$
shows *sorted-descending* $(\text{map } \text{fst } (\text{map } (\text{apfst } (\text{of-nat} :: \text{nat} \Rightarrow 'l \text{ word})) (\text{annotate-rlen } l)))$
<proof>

l3 device to l2 forwarding

definition *lr-of-tran-s3* $\text{ifs } \text{ard} = ($
 $[(p, b, \text{case } a \text{ of } \text{simple-action.Accept} \Rightarrow [\text{Forward } c] \mid \text{simple-action.Drop} \Rightarrow []).$
 $(p, r, (c, a)) \leftarrow \text{ard}, b \leftarrow \text{simple-match-to-of-match } r \ \text{ifs}]$

definition *oif-ne-iif-p1* $\text{ifs} \equiv [(\text{simple-match-any}(\text{oiface} := \text{Iface } \text{oif}, \text{iiface} := \text{Iface } \text{iif}), \text{simple-action.Accept}). \text{oif} \leftarrow \text{ifs},$
 $\text{iif} \leftarrow \text{ifs}, \text{oif} \neq \text{iif}]$

definition *oif-ne-iif-p2* $\text{ifs} = [(\text{simple-match-any}(\text{oiface} := \text{Iface } i, \text{iiface} := \text{Iface } i), \text{simple-action.Drop}). i \leftarrow \text{ifs}]$

definition *oif-ne-iif* $\text{ifs} = \text{oif-ne-iif-p2 } \text{ifs} \ @ \ \text{oif-ne-iif-p1 } \text{ifs}$

definition *lr-of-tran-s4* $\text{ard } \text{ifs} \equiv \text{generalized-fw-join } \text{ard } (\text{oif-ne-iif } \text{ifs})$

definition *lr-of-tran-s1* $\text{rt} = [(\text{route2match } r, \text{output-iface } (\text{routing-action } r)). r \leftarrow \text{rt}]$

definition *lr-of-tran-fbs* *rt fw ifs* \equiv *let*
gfw = *map simple-rule-dtor fw*; — generalized simple fw, hopefully for FORWARD
ftr = *lr-of-tran-s1 rt*; — *rt* as fw
prd = *generalized-fw-join ftr gfw*
in prd

definition *pack-OF-entries* *ifs ard* \equiv (*map (split3 OFEntry) (lr-of-tran-s3 ifs ard)*)

definition *no-oif-match* \equiv *list-all* ($\lambda m. \text{oiface (match-sel m) = ifaceAny}$)

definition *lr-of-tran* *rt fw ifs* \equiv
if \neg (*no-oif-match fw* \wedge *has-default-policy fw* \wedge *simple-fw-valid fw* \wedge *valid-prefixes rt* \wedge *has-default-route rt* \wedge *distinct ifs*)
then Inl "Error in creating OpenFlow table: prerequisites not satisfied"
else (
let nrd = *lr-of-tran-fbs rt fw ifs*;
ard = *map (apfst of-nat) (annotate-rlen nrd)* — give them a priority
in
if length nrd < unat (- 1 :: 16 word)
then Inr (pack-OF-entries ifs ard)
else Inl "Error in creating OpenFlow table: priority number space exhausted")

definition *is-iface-name* *i* $\equiv i \neq [] \wedge \neg \text{Iface.iface-name-is-wildcard } i$

definition *is-iface-list* *ifs* \equiv *distinct ifs* \wedge *list-all is-iface-name ifs*

lemma *max-16-word-max[simp]*: (*a* :: 16 word) $\leq 0xffff$

<proof>

lemma *replicate-FT-hlp*: $x \leq 16 \wedge y \leq 16 \implies \text{replicate (16 - x) False} @ \text{replicate x True} = \text{replicate (16 - y) False} @ \text{replicate y True} \implies x = y$

<proof>

lemma *mask-inj-hlp1*: *inj-on* (*mask* :: nat \Rightarrow 16 word) {0..16}

<proof>

lemma *distinct-simple-match-to-of-match-portlist-hlp*:

fixes *ps* :: (16 word \times 16 word)

shows *distinct ifs* \implies

distinct

(*if fst ps* = 0 \wedge *snd ps* = *max-word* *then* [None]

else if fst ps \leq *snd ps*

then map (Some \circ ($\lambda pfx. (\text{pfxm-prefix } pfx, \sim\sim (\text{pfxm-mask } pfx)))$)

(*wordinterval-CIDR-split-prefixmatch (WordInterval (fst ps) (snd ps))*)

else [])

<proof>

lemma *distinct-simple-match-to-of-match*: *distinct ifs* \implies *distinct (simple-match-to-of-match m ifs)*

<proof>

lemma *inj-inj-on*: *inj F* \implies *inj-on F A* *<proof>*

lemma *no-overlaps-lroft-hlp2*: *distinct (map fst amr)* \implies ($\bigwedge r. \text{distinct (fm } r)$) \implies
distinct (concat (map ($\lambda(p, r, c, a). \text{map } (\lambda b. (p, b, fs a c))$) (fm r)) amr))

<proof>

lemma *distinct-lroft-s3*: $\llbracket \text{distinct } (\text{map fst } \text{amr}); \text{distinct ifs} \rrbracket \implies \text{distinct } (\text{lr-of-tran-s3 ifs } \text{amr})$
<proof>

lemma *no-overlaps-lroft-hlp3*: $\text{distinct } (\text{map fst } \text{amr}) \implies$
 $(\text{aa}, \text{ab}, \text{ac}) \in \text{set } (\text{lr-of-tran-s3 ifs } \text{amr}) \implies (\text{ba}, \text{bb}, \text{bc}) \in \text{set } (\text{lr-of-tran-s3 ifs } \text{amr}) \implies$
 $\text{ac} \neq \text{bc} \implies \text{aa} \neq \text{ba}$
<proof>

lemma *no-overlaps-lroft-s3-hlp-hlp*:
 $\llbracket \text{distinct } (\text{map fst } \text{amr}); \text{OF-match-fields-unsafe } \text{ab } \text{p}; \text{ab} \neq \text{ad} \vee \text{ba} \neq \text{bb}; \text{OF-match-fields-unsafe } \text{ad } \text{p};$
 $(\text{ac}, \text{ab}, \text{ba}) \in \text{set } (\text{lr-of-tran-s3 ifs } \text{amr}); (\text{ac}, \text{ad}, \text{bb}) \in \text{set } (\text{lr-of-tran-s3 ifs } \text{amr}) \rrbracket$
 $\implies \text{False}$
<proof>

lemma *no-overlaps-lroft-s3-hlp*: $\text{distinct } (\text{map fst } \text{amr}) \implies \text{distinct ifs} \implies$
 $\text{no-overlaps } \text{OF-match-fields-unsafe } (\text{map } (\text{split3 } \text{OFEntry}) (\text{lr-of-tran-s3 ifs } \text{amr}))$
<proof>

lemma *lr-of-tran-no-overlaps*: **assumes** *distinct ifs* **shows** $\text{Inr } t = (\text{lr-of-tran } \text{rt } \text{fw } \text{ifs}) \implies \text{no-overlaps}$
 $\text{OF-match-fields-unsafe } t$
<proof>

lemma *sorted-lr-of-tran-s3-hlp*: $\forall x \in \text{set } f. \text{fst } x \leq a \implies b \in \text{set } (\text{lr-of-tran-s3 } s \text{ } f) \implies \text{fst } b \leq a$
<proof>

lemma *lr-of-tran-s3-Cons*: $\text{lr-of-tran-s3 ifs } (a \# \text{ard}) = ($
 $\llbracket (p, b, \text{case } a \text{ of simple-action.Accept} \Rightarrow [\text{Forward } c] \mid \text{simple-action.Drop} \Rightarrow [])$
 $(p, r, (c, a)) \leftarrow [a], b \leftarrow \text{simple-match-to-of-match } r \text{ ifs}] \rrbracket @ \text{lr-of-tran-s3 ifs } \text{ard}$
<proof>

lemma *sorted-lr-of-tran-s3*: $\text{sorted-descending } (\text{map fst } f) \implies \text{sorted-descending } (\text{map fst } (\text{lr-of-tran-s3 } s \text{ } f))$
<proof>

lemma *sorted-lr-of-tran-hlp*: $(\text{ofe-prio} \circ \text{split3 } \text{OFEntry}) = \text{fst}$ *<proof>*

lemma *lr-of-tran-sorted-descending*: $\text{Inr } r = \text{lr-of-tran } \text{rt } \text{fw } \text{ifs} \implies \text{sorted-descending } (\text{map } \text{ofe-prio } r)$
<proof>

lemma *lr-of-tran-s1-split*: $\text{lr-of-tran-s1 } (a \# \text{rt}) = (\text{route2match } a, \text{output-iface } (\text{routing-action } a)) \# \text{lr-of-tran-s1 } \text{rt}$
<proof>

lemma *route2match-correct*: $\text{valid-prefix } (\text{routing-match } a) \implies \text{prefix-match-antics } (\text{routing-match } a) (p\text{-dst } p) \longleftrightarrow \text{simple-matches } (\text{route2match } a) (p)$
<proof>

lemma *s1-correct*: $\text{valid-prefixes } \text{rt} \implies \text{has-default-route } (\text{rt}::('i::\text{len}) \text{ prefix-routing}) \implies$
 $\exists \text{rm } \text{ra}. \text{generalized-sfw } (\text{lr-of-tran-s1 } \text{rt}) \text{ } p = \text{Some } (\text{rm}, \text{ra}) \wedge \text{ra} = \text{output-iface } (\text{routing-table-antics } \text{rt } (p\text{-dst } p))$
<proof>

definition *to-OF-action* $a \equiv (\text{case } a \text{ of } (p, d) \Rightarrow (\text{case } d \text{ of simple-action.Accept} \Rightarrow [\text{Forward } p] \mid \text{simple-action.Drop} \Rightarrow []))$

definition *from-OF-action* $a = (\text{case } a \text{ of } [] \Rightarrow (''', \text{simple-action.Drop}) \mid [\text{Forward } p] \Rightarrow (p, \text{simple-action.Accept}))$

lemma *OF-match-linear-not-noD*: $OF\text{-match-linear } \gamma \text{ oms } p \neq \text{NoAction} \implies \exists \text{ome. } \text{ome} \in \text{set oms} \wedge \gamma (\text{ofe-fields ome}) p$
 ⟨proof⟩

lemma *s3-noaction-hlp*: $\llbracket \text{simple-match-valid } ac; \neg \text{simple-matches } ac \text{ } p; \text{match-iface } (\text{oiface } ac) (p\text{-oiface } p) \rrbracket \implies$
 $OF\text{-match-linear } OF\text{-match-fields-safe } (\text{map } (\lambda x. \text{split3 } OF\text{Entry } (x1, x, \text{case } ba \text{ of } \text{simple-action.Accept} \Rightarrow [\text{Forward } ad] \mid$
 $\text{simple-action.Drop} \Rightarrow [])) (\text{simple-match-to-of-match } ac \text{ ifs})) p = \text{NoAction}$
 ⟨proof⟩

lemma *aux*:
 ⟨ $v = \text{Some } x \implies \text{the } v = x$ ⟩
 ⟨proof⟩

lemma *s3-correct*:
assumes *vsfwm*: $\text{list-all simple-match-valid } (\text{map } (\text{fst} \circ \text{snd}) \text{ ard})$
assumes *ippkt*: $p\text{-l2type } p = 0x800$
assumes *iiifs*: $p\text{-iiface } p \in \text{set ifs}$
assumes *oiifs*: $\text{list-all } (\lambda m. \text{oiface } (\text{fst } (\text{snd } m)) = \text{ifaceAny}) \text{ ard}$
shows $OF\text{-match-linear } OF\text{-match-fields-safe } (\text{pack-OF-entries } ifs \text{ ard}) p = \text{Action } ao \iff (\exists r \text{ af. } \text{generalized-sfw } (\text{map } \text{snd}$
 $\text{ard}) p = (\text{Some } (r, \text{af})) \wedge (\text{if } \text{snd } \text{af} = \text{simple-action.Drop} \text{ then } ao = [] \text{ else } ao = [\text{Forward } (\text{fst } \text{af})]))$
 ⟨proof⟩

context
notes *valid-prefix-00*[*simp, intro!*]

begin
lemma *lr-of-tran-s1-valid*: $\text{valid-prefixes } rt \implies \text{gsfw-valid } (\text{lr-of-tran-s1 } rt)$
 ⟨proof⟩

end

lemma *simple-match-valid-fbs-rlen*: $\llbracket \text{valid-prefixes } rt; \text{simple-fw-valid } fw; (a, aa, ab, b) \in \text{set } (\text{annotate-rlen } (\text{lr-of-tran-fbs } rt$
 $\text{fw } ifs)) \rrbracket \implies \text{simple-match-valid } aa$
 ⟨proof⟩

lemma *simple-match-valid-fbs*: $\llbracket \text{valid-prefixes } rt; \text{simple-fw-valid } fw \rrbracket \implies \text{list-all simple-match-valid } (\text{map } \text{fst } (\text{lr-of-tran-fbs}$
 $rt \text{ fw } ifs))$
 ⟨proof⟩

lemma *lr-of-tran-prereqs*: $\text{valid-prefixes } rt \implies \text{simple-fw-valid } fw \implies \text{lr-of-tran } rt \text{ fw } ifs = \text{Inr } oft \implies$
 $\text{list-all } (\text{all-prerequisites} \circ \text{ofe-fields}) oft$
 ⟨proof⟩

lemma *OF-unsafe-safe-match3-eq*:
 $\text{list-all } (\text{all-prerequisites} \circ \text{ofe-fields}) oft \implies$
 $OF\text{-priority-match } OF\text{-match-fields-unsafe } oft = OF\text{-priority-match } OF\text{-match-fields-safe } oft$
 ⟨proof⟩

lemma *OF-unsafe-safe-match-linear-eq*:
 $\text{list-all } (\text{all-prerequisites} \circ \text{ofe-fields}) oft \implies$
 $OF\text{-match-linear } OF\text{-match-fields-unsafe } oft = OF\text{-match-linear } OF\text{-match-fields-safe } oft$
 ⟨proof⟩

lemma *simple-action-ne*[*simp*]:
 $b \neq \text{simple-action.Accept} \iff b = \text{simple-action.Drop}$

$b \neq \text{simple-action.Drop} \iff b = \text{simple-action.Accept}$
(proof)

lemma *map-snd-apfst*: $\text{map snd (map (apfst x) l)} = \text{map snd l}$
(proof)

lemma *match-ifaceAny-eq*: $\text{iface } m = \text{ifaceAny} \implies \text{simple-matches } m \ p = \text{simple-matches } m \ (p(\!|p\text{-iface} := \text{any}))$
(proof)

lemma *no-oif-matchD*: $\text{no-oif-match } fw \implies \text{simple-fw } fw \ p = \text{simple-fw } fw \ (p(\!|p\text{-iface} := \text{any}))$
(proof)

lemma *lr-of-tran-fbs-acceptD*:
assumes *s1*: *valid-prefixes rt has-default-route rt*
assumes *s2*: *no-oif-match fw*
shows *generalized-sfw (lr-of-tran-fbs rt fw ifs) p = Some (r, oif, simple-action.Accept) \implies*
simple-linux-router-nol12 rt fw p = Some (p(\!|p-oiface := oif))
(proof)

lemma *lr-of-tran-fbs-acceptI*:
assumes *s1*: *valid-prefixes rt has-default-route rt*
assumes *s2*: *no-oif-match fw has-default-policy fw*
shows *simple-linux-router-nol12 rt fw p = Some (p(\!|p-oiface := oif)) \implies*
 $\exists r. \text{generalized-sfw (lr-of-tran-fbs rt fw ifs) } p = \text{Some (r, oif, simple-action.Accept)}$
(proof)

lemma *lr-of-tran-fbs-dropD*:
assumes *s1*: *valid-prefixes rt has-default-route rt*
assumes *s2*: *no-oif-match fw*
shows *generalized-sfw (lr-of-tran-fbs rt fw ifs) p = Some (r, oif, simple-action.Drop) \implies*
simple-linux-router-nol12 rt fw p = None
(proof)

lemma *lr-of-tran-fbs-dropI*:
assumes *s1*: *valid-prefixes rt has-default-route rt*
assumes *s2*: *no-oif-match fw has-default-policy fw*
shows *simple-linux-router-nol12 rt fw p = None \implies*
 $\exists r \text{ oif. } \text{generalized-sfw (lr-of-tran-fbs rt fw ifs) } p = \text{Some (r, oif, simple-action.Drop)}$
(proof)

lemma *no-oif-match-fbs*:
 $\text{no-oif-match } fw \implies \text{list-all } (\lambda m. \text{iface (fst (snd m))} = \text{ifaceAny}) \ (\text{map (apfst of-nat) (annotate-rlen (lr-of-tran-fbs rt fw ifs))})$
(proof)

lemma *lr-of-tran-correct*:
fixes *p* :: (32, 'a) *simple-packet-ext-scheme*
assumes *nerr*: *lr-of-tran rt fw ifs = Inr oft*
and *ippkt*: *p-l2type p = 0x800*
and *ifvld*: *p-iface p \in set ifs*
shows *OF-priority-match OF-match-fields-safe oft p = Action [Forward oif] \iff simple-linux-router-nol12 rt fw p = (Some (p(\!|p-oiface := oif)))*
OF-priority-match OF-match-fields-safe oft p = Action [] \iff simple-linux-router-nol12 rt fw p = None

$OF\text{-priority-match } OF\text{-match-fields-safe } oft\ p \neq NoAction\ OF\text{-priority-match } OF\text{-match-fields-safe } oft\ p \neq Undefined$
 $OF\text{-priority-match } OF\text{-match-fields-safe } oft\ p = Action\ ls \longrightarrow length\ ls \leq 1$
 $\exists ls. length\ ls \leq 1 \wedge OF\text{-priority-match } OF\text{-match-fields-safe } oft\ p = Action\ ls$
 <proof>

end

theory *OF-conv-test*

imports

Iptables-Semantics.Parser

Simple-Firewall.SimpleFw-toString

Routing.IpRoute-Parser

../LinuxRouter-OpenFlow-Translation

../OpenFlow-Serialize

begin

parse-iptables-save *SQRL-fw=iptables-save*

term *SQRL-fw*

thm *SQRL-fw-def*

thm *SQRL-fw-FORWARD-default-policy-def*

value[code] *map* ($\lambda(c,rs). (c, \text{map } (\text{quote-rewrite} \circ \text{common-primitive-rule-toString})\ rs)$) *SQRL-fw*

definition *unfolded* = *unfold-ruleset-FORWARD SQRL-fw-FORWARD-default-policy (map-of-string-ipv4 SQRL-fw)*

lemma *map* (*quote-rewrite* \circ *common-primitive-rule-toString*) *unfolded* =

["-p icmp -j ACCEPT",

"-i s1-lan -p tcp -m tcp --spts [1024:65535] -m tcp --dpts [80] -j ACCEPT",

"-i s1-wan -p tcp -m tcp --spts [80] -m tcp --dpts [1024:65535] -j ACCEPT",

"-j DROP"] <proof>

lemma *length unfolded* = 4 <proof>

value[code] *map* (*quote-rewrite* \circ *common-primitive-rule-toString*) (*upper-closure unfolded*)

lemma *length* (*upper-closure unfolded*) = 4 <proof>

value[code] *upper-closure* (*packet-assume-new unfolded*)

lemma *length* (*lower-closure unfolded*) = 4 <proof>

lemma *check-simple-fw-preconditions* (*upper-closure unfolded*) = *True* <proof>

lemma $\forall m \in \text{get-match'set } (\text{upper-closure } (\text{packet-assume-new } \text{unfolded}))$. *normalized-nnf-match* *m* <proof>

lemma $\forall m \in \text{get-match'set } (\text{optimize-matches } \text{abstract-for-simple-firewall } (\text{upper-closure } (\text{packet-assume-new } \text{unfolded})))$. *normalized-nnf-match* *m* <proof>

lemma *check-simple-fw-preconditions* (*upper-closure* (*optimize-matches* *abstract-for-simple-firewall* (*upper-closure* (*packet-assume-new unfolded*)))) <proof>

lemma *length* (*to-simple-firewall* (*upper-closure* (*packet-assume-new unfolded*))) = 4 <proof>

lemma (*lower-closure* (*optimize-matches* *abstract-for-simple-firewall* (*lower-closure* (*packet-assume-new unfolded*)))) = *lower-closure unfolded*

lower-closure unfolded = upper-closure unfolded
(upper-closure (optimize-matches abstract-for-simple-firewall (upper-closure (packet-assume-new unfolded)))) = upper-closure unfolded *<proof>*

value_[code] (getParts (to-simple-firewall (lower-closure (optimize-matches abstract-for-simple-firewall (lower-closure (packet-assume-new unfolded))))))

definition *SQRL-fw-simple* \equiv *remdups-rev (to-simple-firewall (upper-closure (optimize-matches abstract-for-simple-firewall (upper-closure (packet-assume-new unfolded))))))*

value_[code] *SQRL-fw-simple*

lemma *simple-fw-valid SQRL-fw-simple* *<proof>*

parse-ip-route *SQRL-rtbl-main* = *ip-route*

value *SQRL-rtbl-main*

lemma *SQRL-rtbl-main* = [(*routing-match* = *PrefixMatch 0xA000100 24*, *metric* = 0, *routing-action* = (*output-iface* = "s1-lan", *next-hop* = None)),
(*routing-match* = *PrefixMatch 0xA000200 24*, *metric* = 0, *routing-action* = (*output-iface* = "s1-wan", *next-hop* = None)),
(*routing-match* = *PrefixMatch 0 0*, *metric* = 0, *routing-action* = (*output-iface* = "s1-wan", *next-hop* = Some 0xA000201))] *<proof>*

value *dotdecimal-of-ipv4addr 0xA0D2500*

lemma *SQRL-rtbl-main* = [

rr-ctor (10,0,1,0) 24 "s1-lan" None 0,
rr-ctor (10,0,2,0) 24 "s1-wan" None 0,
rr-ctor (0,0,0,0) 0 "s1-wan" (Some (10,0,2,1)) 0
]

<proof>

definition *SQRL-rtbl-main-sorted* \equiv *rev (sort-key (λr . pfxm-length (routing-match r)) SQRL-rtbl-main)*

value *SQRL-rtbl-main-sorted*

definition *SQRL-ifs* \equiv [

(*iface-name* = "s1-lan", *iface-mac* = 0x10001),
(*iface-name* = "s1-wan", *iface-mac* = 0x10002)
]

value *SQRL-ifs*

definition *SQRL-macs* \equiv [

(("s1-lan", (ipv4addr-of-dotdecimal (10,0,1,2), 0x1)),
("s1-lan", (ipv4addr-of-dotdecimal (10,0,1,3), 0x2)),
("s1-wan", (ipv4addr-of-dotdecimal (10,0,2,1), 0x3))
(("s1-wan", (ipv4addr-of-dotdecimal (10,0,2,4), 0x4)),
(("s1-wan", (ipv4addr-of-dotdecimal (10,0,2,5), 0x5))
)

definition *SQRL-ports* \equiv [

("s1-lan", "1"),
("s1-wan", "2")
]

lemma *let fw = SQRL-fw-simple in no-oif-match fw \wedge has-default-policy fw \wedge simple-fw-valid fw* *<proof>*

lemma *let rt = SQRL-rtbl-main-sorted in valid-prefixes rt \wedge has-default-route rt* *<proof>*

lemma *let ifs = (map iface-name SQRL-ifs) in distinct ifs* *<proof>*

definition *ofi* ≡

case (*lr-of-tran* *SQRL-rtbl-main-sorted* *SQRL-fw-simple* (map *iface-name* *SQRL-ifs*)
of (*Inr* *openflow-rules*) ⇒ map (*serialize-of-entry* (the ◦ *map-of* *SQRL-ports*)) *openflow-rules*)

lemma *ofi* =

```
["priority=11,hard-timeout=0,idle-timeout=0,dl-type=0x800,nw-proto=1,nw-dst=10.0.2.0/24,action=output:2",  
"priority=10,hard-timeout=0,idle-timeout=0,in-port=1,dl-type=0x800,nw-proto=6,nw-dst=10.0.2.0/24,tp-src=1024/0xfc00,tp-dst=80,action=output:2",  
"priority=10,hard-timeout=0,idle-timeout=0,in-port=1,dl-type=0x800,nw-proto=6,nw-dst=10.0.2.0/24,tp-src=2048/0xf800,tp-dst=80,action=output:2",  
"priority=10,hard-timeout=0,idle-timeout=0,in-port=1,dl-type=0x800,nw-proto=6,nw-dst=10.0.2.0/24,tp-src=4096/0xf000,tp-dst=80,action=output:2",  
"priority=10,hard-timeout=0,idle-timeout=0,in-port=1,dl-type=0x800,nw-proto=6,nw-dst=10.0.2.0/24,tp-src=8192/0xe000,tp-dst=80,action=output:2",  
"priority=10,hard-timeout=0,idle-timeout=0,in-port=1,dl-type=0x800,nw-proto=6,nw-dst=10.0.2.0/24,tp-src=16384/0xc000,tp-dst=80,action=output:2",  
"priority=10,hard-timeout=0,idle-timeout=0,in-port=1,dl-type=0x800,nw-proto=6,nw-dst=10.0.2.0/24,tp-src=32768/0x8000,tp-dst=80,action=output:2",  
"priority=9,hard-timeout=0,idle-timeout=0,in-port=2,dl-type=0x800,nw-proto=6,nw-dst=10.0.2.0/24,tp-src=80,tp-dst=1024/0xfc00,action=output:2",  
"priority=9,hard-timeout=0,idle-timeout=0,in-port=2,dl-type=0x800,nw-proto=6,nw-dst=10.0.2.0/24,tp-src=80,tp-dst=2048/0xf800,action=output:2",  
"priority=9,hard-timeout=0,idle-timeout=0,in-port=2,dl-type=0x800,nw-proto=6,nw-dst=10.0.2.0/24,tp-src=80,tp-dst=4096/0xf000,action=output:2",  
"priority=9,hard-timeout=0,idle-timeout=0,in-port=2,dl-type=0x800,nw-proto=6,nw-dst=10.0.2.0/24,tp-src=80,tp-dst=8192/0xe000,action=output:2",  
"priority=9,hard-timeout=0,idle-timeout=0,in-port=2,dl-type=0x800,nw-proto=6,nw-dst=10.0.2.0/24,tp-src=80,tp-dst=16384/0xc000,action=output:2",  
"priority=9,hard-timeout=0,idle-timeout=0,in-port=2,dl-type=0x800,nw-proto=6,nw-dst=10.0.2.0/24,tp-src=80,tp-dst=32768/0x8000,action=output:2",  
"priority=8,hard-timeout=0,idle-timeout=0,dl-type=0x800,nw-dst=10.0.2.0/24,action=drop",  
"priority=7,hard-timeout=0,idle-timeout=0,dl-type=0x800,nw-proto=1,nw-dst=10.0.1.0/24,action=output:1",  
"priority=6,hard-timeout=0,idle-timeout=0,in-port=1,dl-type=0x800,nw-proto=6,nw-dst=10.0.1.0/24,tp-src=1024/0xfc00,tp-dst=80,action=output:2",  
"priority=6,hard-timeout=0,idle-timeout=0,in-port=1,dl-type=0x800,nw-proto=6,nw-dst=10.0.1.0/24,tp-src=2048/0xf800,tp-dst=80,action=output:2",  
"priority=6,hard-timeout=0,idle-timeout=0,in-port=1,dl-type=0x800,nw-proto=6,nw-dst=10.0.1.0/24,tp-src=4096/0xf000,tp-dst=80,action=output:2",  
"priority=6,hard-timeout=0,idle-timeout=0,in-port=1,dl-type=0x800,nw-proto=6,nw-dst=10.0.1.0/24,tp-src=8192/0xe000,tp-dst=80,action=output:2",  
"priority=6,hard-timeout=0,idle-timeout=0,in-port=1,dl-type=0x800,nw-proto=6,nw-dst=10.0.1.0/24,tp-src=16384/0xc000,tp-dst=80,action=output:2",  
"priority=6,hard-timeout=0,idle-timeout=0,in-port=1,dl-type=0x800,nw-proto=6,nw-dst=10.0.1.0/24,tp-src=32768/0x8000,tp-dst=80,action=output:2",  
"priority=5,hard-timeout=0,idle-timeout=0,in-port=2,dl-type=0x800,nw-proto=6,nw-dst=10.0.1.0/24,tp-src=80,tp-dst=1024/0xfc00,action=output:2",  
"priority=5,hard-timeout=0,idle-timeout=0,in-port=2,dl-type=0x800,nw-proto=6,nw-dst=10.0.1.0/24,tp-src=80,tp-dst=2048/0xf800,action=output:2",  
"priority=5,hard-timeout=0,idle-timeout=0,in-port=2,dl-type=0x800,nw-proto=6,nw-dst=10.0.1.0/24,tp-src=80,tp-dst=4096/0xf000,action=output:2",  
"priority=5,hard-timeout=0,idle-timeout=0,in-port=2,dl-type=0x800,nw-proto=6,nw-dst=10.0.1.0/24,tp-src=80,tp-dst=8192/0xe000,action=output:2",  
"priority=5,hard-timeout=0,idle-timeout=0,in-port=2,dl-type=0x800,nw-proto=6,nw-dst=10.0.1.0/24,tp-src=80,tp-dst=16384/0xc000,action=output:2",  
"priority=5,hard-timeout=0,idle-timeout=0,in-port=2,dl-type=0x800,nw-proto=6,nw-dst=10.0.1.0/24,tp-src=80,tp-dst=32768/0x8000,action=output:2",  
"priority=4,hard-timeout=0,idle-timeout=0,dl-type=0x800,nw-dst=10.0.1.0/24,action=drop",  
"priority=3,hard-timeout=0,idle-timeout=0,dl-type=0x800,nw-proto=1,action=output:2",  
"priority=2,hard-timeout=0,idle-timeout=0,in-port=1,dl-type=0x800,nw-proto=6,tp-src=1024/0xfc00,tp-dst=80,action=output:2",  
"priority=2,hard-timeout=0,idle-timeout=0,in-port=1,dl-type=0x800,nw-proto=6,tp-src=2048/0xf800,tp-dst=80,action=output:2",  
"priority=2,hard-timeout=0,idle-timeout=0,in-port=1,dl-type=0x800,nw-proto=6,tp-src=4096/0xf000,tp-dst=80,action=output:2",  
"priority=2,hard-timeout=0,idle-timeout=0,in-port=1,dl-type=0x800,nw-proto=6,tp-src=8192/0xe000,tp-dst=80,action=output:2",  
"priority=2,hard-timeout=0,idle-timeout=0,in-port=1,dl-type=0x800,nw-proto=6,tp-src=16384/0xc000,tp-dst=80,action=output:2",  
"priority=2,hard-timeout=0,idle-timeout=0,in-port=1,dl-type=0x800,nw-proto=6,tp-src=32768/0x8000,tp-dst=80,action=output:2",  
"priority=1,hard-timeout=0,idle-timeout=0,in-port=2,dl-type=0x800,nw-proto=6,tp-src=80,tp-dst=1024/0xfc00,action=output:2",  
"priority=1,hard-timeout=0,idle-timeout=0,in-port=2,dl-type=0x800,nw-proto=6,tp-src=80,tp-dst=2048/0xf800,action=output:2",  
"priority=1,hard-timeout=0,idle-timeout=0,in-port=2,dl-type=0x800,nw-proto=6,tp-src=80,tp-dst=4096/0xf000,action=output:2",  
"priority=1,hard-timeout=0,idle-timeout=0,in-port=2,dl-type=0x800,nw-proto=6,tp-src=80,tp-dst=8192/0xe000,action=output:2",  
"priority=1,hard-timeout=0,idle-timeout=0,in-port=2,dl-type=0x800,nw-proto=6,tp-src=80,tp-dst=16384/0xc000,action=output:2",  
"priority=1,hard-timeout=0,idle-timeout=0,in-port=2,dl-type=0x800,nw-proto=6,tp-src=80,tp-dst=32768/0x8000,action=output:2",  
"priority=0,hard-timeout=0,idle-timeout=0,dl-type=0x800,action=drop"] (proof)
```

value[*code*] *ofi*

end

theory *RFC2544*

```

imports
  Iptables-Semantics.Parser
  Routing.IpRoute-Parser
  ../../LinuxRouter-OpenFlow-Translation
  ../../OpenFlow-Serialize
begin

parse-iptables-save SQRL-fw=iptables-save

term SQRL-fw
thm SQRL-fw-def
thm SQRL-fw-FORWARD-default-policy-def

value[code] map ( $\lambda(c,rs). (c, \text{map} (\text{quote-rewrite} \circ \text{common-primitive-rule-toString}) rs)$ ) SQRL-fw
definition unfolded = unfold-ruleset-FORWARD SQRL-fw-FORWARD-default-policy (map-of-string-ipv4 SQRL-fw)

lemma length unfolded = 26 <proof>

value[code] unfolded
value[code] (upper-closure unfolded)
value[code] map ( $\text{quote-rewrite} \circ \text{common-primitive-rule-toString}$ ) (upper-closure unfolded)
lemma length (upper-closure unfolded) = 26 <proof>

value[code] upper-closure (packet-assume-new unfolded)

lemma length (lower-closure unfolded) = 26 <proof>

lemma check-simple-fw-preconditions (upper-closure unfolded) <proof>
lemma  $\forall m \in \text{get-match'set} (\text{upper-closure} (\text{packet-assume-new unfolded})). \text{normalized-nnf-match } m$  <proof>
lemma  $\forall m \in \text{get-match'set} (\text{optimize-matches abstract-for-simple-firewall} (\text{upper-closure} (\text{packet-assume-new unfolded}))). \text{normalized-nnf-match } m$  <proof>
lemma check-simple-fw-preconditions (upper-closure (optimize-matches abstract-for-simple-firewall (upper-closure (packet-assume-new unfolded)))) <proof>
lemma length (to-simple-firewall (upper-closure (packet-assume-new unfolded))) = 26 <proof>

lemma (lower-closure (optimize-matches abstract-for-simple-firewall (lower-closure (packet-assume-new unfolded)))) = lower-closure unfolded
  lower-closure unfolded = upper-closure unfolded
  (upper-closure (optimize-matches abstract-for-simple-firewall (upper-closure (packet-assume-new unfolded)))) = upper-closure unfolded <proof>

value[code] (getParts (to-simple-firewall (lower-closure (optimize-matches abstract-for-simple-firewall (lower-closure (packet-assume-new unfolded)))))))

definition SQRL-fw-simple  $\equiv \text{remdups-rev} (\text{to-simple-firewall} (\text{upper-closure} (\text{optimize-matches abstract-for-simple-firewall} (\text{upper-closure} (\text{packet-assume-new unfolded}))))))$ 
value[code] SQRL-fw-simple
lemma simple-fw-valid SQRL-fw-simple <proof>

parse-ip-route SQRL-rtbl-main = ip-route

```

value *SQRL-rtbl-main*

lemma *SQRL-rtbl-main* = [(*routing-match* = *PrefixMatch 0xC6120100 24*, *metric* = 0, *routing-action* = (*output-iface* = "ip1", *next-hop* = None)),
(*routing-match* = *PrefixMatch 0xC6130100 24*, *metric* = 0, *routing-action* = (*output-iface* = "op1", *next-hop* = None)),
(*routing-match* = *PrefixMatch 0 0*, *metric* = 0, *routing-action* = (*output-iface* = "op1", *next-hop* = Some 0xC6130102)))]
{*proof*}

lemma *SQRL-rtbl-main* = [
 rr-ctor (198,18,1,0) 24 "ip1" None 0,
 rr-ctor (198,19,1,0) 24 "op1" None 0,
 rr-ctor (0,0,0,0) 0 "op1" (Some (198,19,1,2)) 0
]
{*proof*}

definition *SQRL-ports* ≡ [
 ("ip1", "1"),
 ("op1", "2")
]

definition *ofi* ≡

case (*lr-of-tran SQRL-rtbl-main SQRL-fw-simple* (map *fst SQRL-ports*)
 of (*Inr openflow-rules*) ⇒ map (*serialize-of-entry* (the ∘ *map-of SQRL-ports*)) *openflow-rules*)

lemma *ofi* =

[*"priority=27,hard-timeout=0,idle-timeout=0,dl-type=0x800,nw-dst=192.18.1.0/24,action=drop"*,
"priority=26,hard-timeout=0,idle-timeout=0,dl-type=0x800,nw-dst=192.19.1.0/24,action=drop",
"priority=25,hard-timeout=0,idle-timeout=0,dl-type=0x800,nw-src=192.18.1.1/32,nw-dst=192.18.101.1/32,action=drop",
"priority=24,hard-timeout=0,idle-timeout=0,dl-type=0x800,nw-src=192.18.2.2/32,nw-dst=192.18.102.2/32,action=drop",
"priority=23,hard-timeout=0,idle-timeout=0,dl-type=0x800,nw-src=192.18.3.3/32,nw-dst=192.18.103.3/32,action=drop",
"priority=22,hard-timeout=0,idle-timeout=0,dl-type=0x800,nw-src=192.18.4.4/32,nw-dst=192.18.104.4/32,action=drop",
"priority=21,hard-timeout=0,idle-timeout=0,dl-type=0x800,nw-src=192.18.5.5/32,nw-dst=192.18.105.5/32,action=drop",
"priority=20,hard-timeout=0,idle-timeout=0,dl-type=0x800,nw-src=192.18.6.6/32,nw-dst=192.18.106.6/32,action=drop",
"priority=19,hard-timeout=0,idle-timeout=0,dl-type=0x800,nw-src=192.18.7.7/32,nw-dst=192.18.107.7/32,action=drop",
"priority=18,hard-timeout=0,idle-timeout=0,dl-type=0x800,nw-src=192.18.8.8/32,nw-dst=192.18.108.8/32,action=drop",
"priority=17,hard-timeout=0,idle-timeout=0,dl-type=0x800,nw-src=192.18.9.9/32,nw-dst=192.18.109.9/32,action=drop",
"priority=16,hard-timeout=0,idle-timeout=0,dl-type=0x800,nw-src=192.18.10.10/32,nw-dst=192.18.110.10/32,action=drop",
"priority=15,hard-timeout=0,idle-timeout=0,dl-type=0x800,nw-src=192.18.11.11/32,nw-dst=192.18.111.11/32,action=drop",
"priority=14,hard-timeout=0,idle-timeout=0,dl-type=0x800,nw-src=192.18.12.12/32,nw-dst=192.18.112.12/32,action=drop",
"priority=13,hard-timeout=0,idle-timeout=0,dl-type=0x800,nw-src=192.19.1.2/32,nw-dst=192.19.65.1/32,action=output:2",
"priority=12,hard-timeout=0,idle-timeout=0,dl-type=0x800,nw-src=192.18.13.13/32,nw-dst=192.18.113.13/32,action=drop",
"priority=11,hard-timeout=0,idle-timeout=0,dl-type=0x800,nw-src=192.18.14.14/32,nw-dst=192.18.114.14/32,action=drop",
"priority=10,hard-timeout=0,idle-timeout=0,dl-type=0x800,nw-src=192.18.15.15/32,nw-dst=192.18.115.15/32,action=drop",
"priority=9,hard-timeout=0,idle-timeout=0,dl-type=0x800,nw-src=192.18.16.16/32,nw-dst=192.18.116.16/32,action=drop",
"priority=8,hard-timeout=0,idle-timeout=0,dl-type=0x800,nw-src=192.18.17.17/32,nw-dst=192.18.117.17/32,action=drop",
"priority=7,hard-timeout=0,idle-timeout=0,dl-type=0x800,nw-src=192.18.18.18/32,nw-dst=192.18.118.18/32,action=drop",
"priority=6,hard-timeout=0,idle-timeout=0,dl-type=0x800,nw-src=192.18.19.19/32,nw-dst=192.18.119.19/32,action=drop",
"priority=5,hard-timeout=0,idle-timeout=0,dl-type=0x800,nw-src=192.18.20.20/32,nw-dst=192.18.120.20/32,action=drop",
"priority=4,hard-timeout=0,idle-timeout=0,dl-type=0x800,nw-src=192.18.21.21/32,nw-dst=192.18.121.21/32,action=drop",
"priority=3,hard-timeout=0,idle-timeout=0,dl-type=0x800,nw-src=192.18.22.22/32,nw-dst=192.18.122.22/32,action=drop",
"priority=2,hard-timeout=0,idle-timeout=0,dl-type=0x800,nw-src=192.18.23.23/32,nw-dst=192.18.123.23/32,action=drop",
"priority=1,hard-timeout=0,idle-timeout=0,dl-type=0x800,nw-src=192.18.24.24/32,nw-dst=192.18.124.24/32,action=drop",
"priority=0,hard-timeout=0,idle-timeout=0,dl-type=0x800,action=drop"] {*proof*}

value[*code*] *length ofi*

end

Part II

Documentation

1 Configuration Translation

All the results we present in this section are formalized and verified in Isabelle/HOL [11]. This means that their formal correctness can be trusted a level close to absolute certainty. The definitions and lemmas stated here are merely a repetition of lemmas stated in other theory files. This means that they have been directly set to this document from Isabelle and no typos or hidden assumptions are possible. Additionally, it allows us to omit various helper lemmas that do not help the understanding. However, it causes some notation inaccuracy, as type and function definitions are stated as lemmas or schematic goals.

theory *OpenFlow-Documentation*

1.1 Linux Firewall Model

We want to write a program that translates the configuration of a linux firewall to that of an OpenFlow switch. We furthermore want to verify that translation. For this purpose, we need a clear definition of the behavior of the two device types – we need their models and semantics. In case of a linux firewall, this is problematic because a linux firewall is a highly complex device that is ultimately capable of general purpose computation. Creating a comprehensive semantics that encompasses all possible configuration types of a linux firewall is thus highly non-trivial and not useful for the purpose of analysis. We decided to approach the problem from the other side: we created a model that includes only the most basic features. (This implies neglecting IPv6.) Fortunately, many of the highly complex features are rarely essential and even our basic model is still of some use.

We first divided the firewall into subsystems. Given a routing table rt , the firewall rules fw , the routing decision for a packet p can be obtained by *routing-table-antics* rt (p -dst p), the firewall decision by *simple-fw* fw p . We draft the first description of our linux router model:

1. The destination MAC address of an arriving packet is checked: Does it match the MAC

address of the ingress port? If it does, we continue, otherwise, the packet is discarded.

2. The routing decision $rd \equiv$ *routing-table-antics* rt p is obtained.
3. The packet's output interface is updated based on rd ¹.
4. The firewall is queried for a decision: *simple-fw* fw p . If the decision is to *Drop*, the packet is discarded.
5. The next hop is computed: If rd provides a next hop, that is used. Otherwise, the destination address of the packet is used.
6. The MAC address of the next hop is looked up; the packet is updated with it and sent.

We decided that this description is best formalized as an abortable program in the option monad:

```
lemma simple-linux-router  $rt$   $fw$   $mlf$   $ifl$   $p \equiv do$  {  
  -  $\leftarrow$  iface-packet-check  $ifl$   $p$ ;  
   $let$   $rd$  — (routing decision) = routing-table-antics  $rt$   
    ( $p$ -dst  $p$ );  
   $let$   $p = p$ ( $p$ -oiface := output-iface  $rd$ );  
   $let$   $fd$  — (firewall decision) = simple-fw  $fw$   $p$ ;  
  -  $\leftarrow$  (case  $fd$  of Decision FinalAllow  $\Rightarrow$  Some () | Decision  
    FinalDeny  $\Rightarrow$  None);  
   $let$   $nh =$  (case next-hop  $rd$  of None  $\Rightarrow$   $p$ -dst  $p$  | Some  $a \Rightarrow$   
     $a$ );  
   $ma \leftarrow$   $mlf$   $nh$ ;  
  Some ( $p$ ( $p$ -l2dst :=  $ma$ ))  
}
```

(proof)

where mlf is a function that looks up the MAC address for an IP address.

There are already a few important aspects that have not been modelled, but they are not core essential for the functionality of a firewall. Namely, there is no local traffic from/to the firewall. This is problematic since this model can not generate ARP replies — thus, an equivalent OpenFlow device will not do so, either. Furthermore, this model is problematic because it requires access to a function that looks up a MAC address, something that may not be known at the time of time running a translation to an OpenFlow configuration.

¹Note that we assume a packet model with input and output interfaces. The origin of this is explained in Section 1.1.2

It is possible to circumvent these problems by inserting static ARP table entries in the directly connected devices and looking up their MAC addresses *a priori*. A test-wise implementation of the translation based on this model showed acceptable results. However, we deemed the *a priori* lookup of the MAC addresses to be rather inelegant and built a second model.

definition *simple-linux-router-altered* $rt\ fw\ ifl\ p \equiv do \{$
 $let\ rd = routing-table-semantic\ rt\ (p-dst\ p);$
 $let\ p = p(p-oiface := output-iface\ rd);$
 $- \leftarrow if\ p-oiface\ p = p-iiface\ p\ then\ None\ else\ Some\ ();$
 $let\ fd = simple-fw\ fw\ p;$
 $- \leftarrow (case\ fd\ of\ Decision\ FinalAllow \Rightarrow Some\ () \mid Decision\ FinalDeny \Rightarrow None);$
 $Some\ p$
 $\}$

In this model, all access to the MAC layer has been eliminated. This is done by the approximation that the firewall will be asked to route a packet (i.e. be addressed on the MAC layer) iff the destination IP address of the packet causes it to be routed out on a different interface. Because this model does not insert destination MAC addresses, the destination MAC address has to be already correct when the packet is sent. This can only be achieved by changing the subnet of all connected device, moving them into one common subnet².

While a test-wise implementation based on this model also showed acceptable results, the model is still problematic. The check $p-oiface\ p = p-iiface\ p$ and the firewall require access to the output interface. The details of why this cannot be provided are elaborated in Section 1.3. The intuitive explanation is that an OpenFlow match can not have a field for the output interface. We thus simplified the model even further:

lemma *simple-linux-router-nol12* $rt\ fw\ p \equiv do \{$
 $let\ rd = routing-table-semantic\ rt\ (p-dst\ p);$
 $let\ p = p(p-oiface := output-iface\ rd);$
 $let\ fd = simple-fw\ fw\ p;$
 $- \leftarrow (case\ fd\ of\ Decision\ FinalAllow \Rightarrow Some\ () \mid Decision\ FinalDeny \Rightarrow None);$
 $Some\ p$
 $\}$ *<proof>*

We continue with this definition as a basis for our translation. Even this strongly altered version and the original linux firewall still behave the same in a substantial amount of cases:

²There are cases where this is not possible — A limitation of our system.

theorem

$\llbracket iface-packet-check\ ifl\ pii \neq None;$
 $mlf\ (case\ next-hop\ (routing-table-semantic\ rt\ (p-dst\ pii))$
 $of\ None \Rightarrow p-dst\ pii \mid Some\ a \Rightarrow a) \neq None \rrbracket \implies$
 $\exists x. map-option\ (\lambda p. p(p-l2dst := x))$
 $(simple-linux-router-nol12\ rt\ fw\ pii) = simple-linux-router$
 $rt\ fw\ mlf\ ifl\ pii$
<proof>

The conditions are to be read as “The check whether a received packet has the correct destination MAC never returns *False*” and “The next hop MAC address for all packets can be looked up”. Obviously, these conditions do not hold for all packets. We will show an example where this makes a difference in Section 2.1.

1.1.1 Routing Table

The routing system in linux features multiple tables and a system that can use the iptables firewall and an additional match language to select a routing table. Based on our directive, we only focused on the single most used **main** routing table.

We define a routing table entry to be a record (named tuple) of a prefix match, a metric and the routing action, which in turn is a record of an output interface and an optional next-hop address.

schematic-goal $(?rtbl-entry :: ('a::len)\ routing-rule) = (\$
 $routing-match = PrefixMatch\ pfx\ len, metric = met, routing-action = (\$
 $output-iface = oif-string, next-hop = (h :: 'a\ word\ option) \}) \}$ *<proof>*

A routing table is then a list of these entries:

lemma $(rtbl :: ('a :: len)\ prefix-routing) = (rtbl :: 'a\ routing-rule\ list)$ *<proof>*

Not all members of the type *prefix-routing* are sane routing tables. There are three different validity criteria that we require so that our definitions are adequate.

- The prefixes have to be 0 in bits exceeding their length.
- There has to be a default rule, i.e. one with prefix length 0. With the condition above, that implies that all its prefix bits are zero and it thus matches any address.
- The entries have to be sorted by prefix length and metric.

The first two are set into code in the following way:

lemma *valid-prefix* (*PrefixMatch pfx len*) \equiv *pfx* && (2 ^ (32 - len) - 1) = (0 :: 32 word)
 <proof>

lemma *has-default-route* *rt* \longleftrightarrow ($\exists r \in$ *set rt*. *pfm-length* (*routing-match r*) = 0)
 <proof>

The third is not needed in any of the further proofs, so we omit it.

The semantics of a routing table is to simply traverse the list until a matching entry is found.

schematic-goal *routing-table-semantics* (*rt-entry # rt*) *dst-addr* = (if *prefix-match-semantics* (*routing-match rt-entry*) *dst-addr* then *routing-action* *rt-entry* else *routing-table-semantics* *rt* *dst-addr*) <proof>

If no matching entry is found, the behavior is undefined.

1.1.2 iptables Firewall

The firewall subsystem in a linux router is not any less complex than any of the other systems. Fortunately, this complexity has been dealt with in [6, 5] already and we can directly use the result.

In short, one of the results is that a complex *iptables* configuration can be simplified to be represented by a single list of matches that only support the following match conditions:

- (String) prefix matches on the input and output interfaces.
- A *prefix-match* on the source and destination IP address.
- An exact match on the layer 4 protocol.
- Interval matches on the source or destination port, e.g. $p_d \in \{1..1023\}$

The model/type of the packet is adjusted to fit that: it is a record of the fields matched on. This also means that input and output interface are coded to the packet. Given that this information is usually stored alongside the packet content, this can be deemed a reasonable model. In case the output interface is not needed (e.g., when evaluating an OpenFlow table), it can simply be left blank.

Obviously, a simplification into the above match type cannot always produce an equivalent firewall, and the

set of accepted packets has to be over- or underapproximated. The reader interested in the details of this is strongly referred to [6]; we are simply going to continue with the result: *simple-fw*.

One property of the simplification is worth noting here: The simplified firewall does not know state and the simplification approximates stateful matches by stateless ones. Thus, the overapproximation of a stateful firewall ruleset that begins with accepting packets of established connections usually begins with a rule that accepts all packets. Dealing with this by writing a meaningful simplification of stateful firewalls is future work.

1.2 OpenFlow Switch Model

In this section, we present our model of an OpenFlow switch. The requirements for this model are derived from the fact that it models devices that are the target of a configuration translation. This has two implications:

- All configurations that are representable in our model should produce the correct behavior wrt. their semantics. The problem is that correct here means that the behavior is the same that any real device would produce. Since we cannot possibly account for all device types, we instead focus on those that conform to the OpenFlow specifications. To account for the multiple different versions of the specification (e.g. [2, 3]), we tried making our model a subset of both the oldest stable version 1.0 [2] and the newest available specification version 1.5.1 [3].
- Conversely, our model does not need to represent all possible behavior of an OpenFlow switch, just the behavior that can be invoked by the result of our translation. This is especially useful regarding controller interaction, but also for MPLS or VLANs, which we did not model in Section 1.1.

More concretely, we set the following rough outline for our model.

- A switch consists of a single flow table.
- A flow table entry consists of a priority, a match condition and an action list.
- The only possible action (we require) is to forward the packet on a port.

- We do not model controller interaction.

Additionally, we decided that we wanted to be able to ensure the validity of the flow table in all qualities, i.e. we want to model the conditions ‘no overlapping flow entries appear’, ‘all match conditions have their necessary preconditions’. The details of this are explained in the following sections.

1.2.1 Matching Flow Table entries

Table 3 of Section 3.1 of [2] gives a list of required packet fields that can be used to match packets. This directly translates into the type for a match expression on a single field:

```
schematic-goal (field-match :: of-match-field) ∈ {
  IngressPort (?s::string),
  EtherSrc (?as::48 word), EtherDst (?ad::48 word),
  EtherType (?t::16 word),
  VlanId (?i::16 word), VlanPriority (?p::16 word),
  IPv4Src (?pms::32 prefix-match),
  IPv4Dst (?pmd::32 prefix-match),
  IPv4Proto (?ipp :: 8 word),
  L4Src (?ps :: 16 word) (?ms :: 16 word),
  L4Dst (?pd :: 16 word) (?md :: 16 word)
} <proof>
```

Two things are worth additional mention: L3 and L4 “addressess”. The *IPv4Src* and *IPv4Dst* matches are specified as “can be subnet masked” in [2], whereas [3] states clearly that arbitrary bitmasks can be used. We took the conservative approach here. Our alteration of *L4Src* and *L4Dst* is more grave. While [2] does not state anything about layer 4 ports and masks, [3] specifically forbids using masks on them. Nevertheless, OpenVSwitch [1] and some other implementations support them. We will explain in detail why we must include bitmasks on layer 4 ports to obtain a meaningful translation in Section 1.3.

One *of-match-field* is not enough to classify a packet. To match packets, we thus use entire sets of match fields. As Guha *et al.* [7] noted³, executing a set of given *of-match-fields* on a packet requires careful consideration. For example, it is not meaningful to use *IPv4Dst* if the given packet is not actually an IP packet, i.e. *IPv4Dst* has the prerequisite of *EtherType 2048* being among the match fields. Guha *et al.* decided to use the fact that the preconditions can be arranged on

³See also: [8, §2.3]

a directed acyclic graph (or rather: an acyclic forest). They evaluated match conditions in a manner following that graph: first, all field matches without preconditions are evaluated. Upon evaluating a field match (e.g., *EtherType 2048*), the matches that had their precondition fulfilled by it (e.g., *IPv4Src* and *IPv4Dst* in this example) are evaluated. This mirrors the faulty behavior of some implementations (see [7]). Adopting that behavior into our model would mean that any packet matches against the field match set $\{IPv4Dst (PrefixMatch 134744072 32)\}$ instead of just those destined for 8.8.8.8 or causing an error. We found this to be unsatisfactory.

To solve this problem, we made three definitions. The first, *match-no-prereq* matches an *of-match-field* against a packet without considering prerequisites. The second, *prerequisites*, checks for a given *of-match-field* whether its prerequisites are in a set of given match fields. Especially:

lemma

```
prerequisites (VlanPriority pri) m = (∃ id. let v = VlanId
id in v ∈ m ∧ prerequisites v m)
prerequisites (IPv4Proto pr) m = (let v = EtherType
0x0800 in v ∈ m ∧ prerequisites v m)
prerequisites (IPv4Src a) m = (let v = EtherType 0x0800
in v ∈ m ∧ prerequisites v m)
prerequisites (IPv4Dst a) m = (let v = EtherType 0x0800
in v ∈ m ∧ prerequisites v m)
prerequisites (L4Src p msk) m = (∃ proto ∈
{TCP,UDP,L4-Protocol.SCTP}. let v = IPv4Proto proto
in v ∈ m ∧ prerequisites v m)
prerequisites (L4Dst p msk) m = prerequisites (L4Src
undefined) m
<proof>
```

Then, to actually match a set of *of-match-field* against a packet, we use the option type:

lemma *OF-match-fields* *m* *p* =

```
(if ∃ f ∈ m. ¬prerequisites f m then None else
if ∀ f ∈ m. match-no-prereq f p then Some True else
Some False)
<proof>
```

1.2.2 Evaluating a Flow Table

In the previous section, we explained how we match the set of match fields belonging to a single flow entry against a packet. This section explains how the correct flow entry from a table can be selected. To prevent too much entanglement with the previous section, we assume an arbitrary match function γ . This function γ

takes the match condition m from a flow entry $OFEntry$ priority m action and decides whether a packet matches those.

The flow table is simply a list of flow table entries *flow-entry-match*. Deciding the right flow entry to use for a given packet is explained in the OpenFlow specification [2], Section 3.4:

Packets are matched against flow entries based on prioritization. An entry that specifies an exact match (i.e., has no wildcards) is always the highest priority⁴. All wildcard entries have a priority associated with them. Higher priority entries must match before lower priority ones. If multiple entries have the same priority, the switch is free to choose any ordering.

We use the term “overlapping” for the flow entries that can cause a packet to match multiple flow entries with the same priority. Guha *et al.* [7] have dealt with overlapping. However, the semantics for a flow table they presented [7, Figure 5] is slightly different from what they actually used in their theory files. We have tried to reproduce the original inductive definition (while keeping our abstraction γ), in Isabelle/HOL⁵:

lemma γ (ofe-fields fe) p = True \implies
 $\forall fe' \in set (ft1 @ ft2). ofe-prio fe' > ofe-prio fe \implies \gamma$
 (ofe-fields fe') p = False \implies
 guha-table-semantics γ (ft1 @ fe # ft2) p (Some (ofe-action fe))
 $\forall fe \in set ft. \gamma$ (ofe-fields fe) p = False \implies
 guha-table-semantics γ ft p None *<proof>*

Guha *et al.* have deliberately made their semantics non-deterministic, to match the fact that the switch “may choose any ordering”. This can lead to undesired results:

lemma $CARD('action) \geq 2 \implies \exists ff. \gamma ff p \implies \exists ft (a1 :: 'action) (a2 :: 'action). a1 \neq a2 \wedge guha-table-semantics \gamma$
 $ft p (Some a1) \wedge guha-table-semantics \gamma ft p (Some a2)$
<proof>

This means that, given at least two distinct actions exist and our matcher γ is not false for all possible match conditions, we can say that a flow table and two actions exist such that both actions are executed. This

⁴This behavior has been deprecated.

⁵The original is written in Coq [4] and we can not use it directly.

can be misleading, as the switch might choose an ordering on some flow table and never execute some of the (overlapped) actions.

Instead, we decided to follow Section 5.3 of the specification [3], which states:

If there are multiple matching flow entries, the selected flow entry is explicitly undefined.

This still leaves some room for interpretation, but it clearly states that overlapping flow entries are undefined behavior, and undefined behavior should not be invoked. Thus, we came up with a semantics that clearly indicates when undefined behavior has been invoked:

lemma

OF-priority-match γ flow-entries packet = (
 let $m = filter (\lambda f. \gamma$ (ofe-fields f) packet) flow-entries;
 $m' = filter (\lambda f. \forall fo \in set m. ofe-prio fo \leq ofe-prio f)$
 m in
 case m' of [] $\implies NoAction$
 | [s] $\implies Action$ (ofe-action s)
 | - $\implies Undefined$)

<proof>

The definition works the following way⁶:

1. The flow table is filtered for those entries that match, the result is called m .
2. m is filtered again, leaving only those entries for which no entries with lower priority could be found, i.e. the matching flow table entries with minimal priority. The result is called m' .
3. A case distinction on m' is made. If only one matching entry was found, its action is returned for execution. If m is empty, the flow table semantics returns *NoAction* to indicate that the flow table does not decide an action for the packet. If, not zero or one entry is found, but more, the special value *Undefined* for indicating undefined behavior is returned.

The use of *Undefined* immediately raises the question in which condition it cannot occur. We give the following definition:

⁶Note that the order of the flow table entries is irrelevant. We could have made this definition on sets but chose not to for consistency.

lemma *check-no-overlap* γ *ft* = $(\forall a \in \text{set } ft. \forall b \in \text{set } ft. (a \neq b \wedge \text{ofe-prio } a = \text{ofe-prio } b) \longrightarrow \neg(\exists p. \gamma (\text{ofe-fields } a) p \wedge \gamma (\text{ofe-fields } b) p))$ *<proof>*

Together with distinctness of the flow table, this provides the absence of *Undefined*⁷:

lemma $\llbracket \text{check-no-overlap } \gamma \text{ } ft; \text{distinct } ft \rrbracket \implies \text{OF-priority-match } \gamma \text{ } ft \text{ } p \neq \text{Undefined}$ *<proof>*

Given the absence of overlapping or duplicate flow entries, we can show two interesting equivalences. the first is the equality to the semantics defined by Guha *et al.*:

lemma $\llbracket \text{check-no-overlap } \gamma \text{ } ft; \text{distinct } ft \rrbracket \implies \text{OF-priority-match } \gamma \text{ } ft \text{ } p = \text{option-to-ftb } d \longleftrightarrow \text{guha-table-semantics } \gamma \text{ } ft \text{ } p \text{ } d$ *<proof>*

where *option-to-ftb* maps between the return type of *OF-priority-match* and an option type as one would expect.

The second equality for *OF-priority-match* is one that helps reasoning about flow tables. We define a simple recursive traversal for flow tables:

lemma
 $\text{OF-match-linear } \gamma \text{ } [] \text{ } p = \text{NoAction}$
 $\text{OF-match-linear } \gamma \text{ } (a \# as) \text{ } p = (\text{if } \gamma (\text{ofe-fields } a) \text{ } p \text{ then } \text{Action } (\text{ofe-action } a) \text{ else } \text{OF-match-linear } \gamma \text{ } as \text{ } p)$ *<proof>*

For this definition to be equivalent, we need the flow table to be sorted:

lemma
 $\llbracket \text{no-overlaps } \gamma \text{ } f \text{ ;sorted-descending } (\text{map } \text{ofe-prio } f) \rrbracket \implies \text{OF-match-linear } \gamma \text{ } f \text{ } p = \text{OF-priority-match } \gamma \text{ } f \text{ } p$ *<proof>*

As the last step, we implemented a serialization function for flow entries; it has to remain unverified. The serialization function deals with one little inaccuracy: We have modelled the *IngressPort* match to use the interface name, but OpenFlow requires numerical interface IDs instead. We deemed that pulling this translation step into the main translation would only make the correctness lemma of the translation more complicated while not increasing the confidence in the correctness significantly. We thus made replacing interface names by their ID part of the serialization.

⁷It is slightly stronger than necessary, overlapping rules might be shadowed and thus never influence the behavior.

Having collected all important definitions and models, we can move on to the conversion.

1.3 Translation Implementation

This section explains how the functions that are executed sequentially in a linux firewall can be compressed into a single OpenFlow table. Creating this flow table in a single step would be immensely complicated. We thus divided the task into several steps using the following key insights:

- All steps that are executed in the linux router can be formulated as a firewall, more specifically, a generalization of *simple-fw* that allows arbitrary actions instead of just accept and drop.
- A function that computes the conjunction of two *simple-fw* matches is already present. Extending this to a function that computes the join of two firewalls is relatively simple. This is explained in Section 1.3.1

1.3.1 Chaining Firewalls

This section explains how to compute the join of two firewalls.

The basis of this is a generalization of *simple-fw*. Instead of only allowing *Accept* or *Drop* as actions, it allows arbitrary actions. The type of the function that evaluates this generalized simple firewall is *generalized-sfw*. The definition is straightforward:

lemma
 $\text{generalized-sfw } [] \text{ } p = \text{None}$
 $\text{generalized-sfw } (a \# as) \text{ } p = (\text{if } (\text{case } a \text{ of } (m, -) \Rightarrow \text{simple-matches } m \text{ } p) \text{ then } \text{Some } a \text{ else } \text{generalized-sfw } as \text{ } p)$ *<proof>*

Based on that, we asked: if fw_1 makes the decision a (where a is the second element of the result tuple from *generalized-sfw*) and fw_2 makes the decision b , how can we compute the firewall that makes the decision (a, b) ⁸. One possible answer is given by the following definition:

lemma $\text{generalized-fw-join } l1 \text{ } l2 \equiv [(u, a, b). (m1, a) \leftarrow l1, (m2, b) \leftarrow l2, u \leftarrow (\text{case } \text{simple-match-and } m1 \text{ } m2 \text{ of } \text{None} \Rightarrow [] \mid \text{Some } s \Rightarrow [s])]$

⁸Note that tuples are right-associative in Isabelle/HOL, i.e., (a, b, c) is a pair of a and the pair (b, c)

<proof>

This definition validates the following lemma:

lemma *generalized-sfw (generalized-fw-join fw₁ fw₂) p = Some (u, d₁, d₂) \longleftrightarrow (\exists r₁ r₂. generalized-sfw fw₁ p = Some (r₁, d₁) \wedge generalized-sfw fw₂ p = Some (r₂, d₂) \wedge Some u = simple-match-and r₁ r₂)*
<proof>

Thus, *generalized-fw-join* has a number of applications. For example, it could be used to compute a firewall ruleset that represents two firewalls that are executed in sequence.

definition *simple-action-conj a b \equiv (if a = simple-action.Accept \wedge b = simple-action.Accept then simple-action.Accept else simple-action.Drop)*

definition *simple-rule-conj \equiv (uncurry SimpleRule \circ apsnd (uncurry simple-action-conj))*

theorem *simple-fw rs₁ p = Decision FinalAllow \wedge simple-fw rs₂ p = Decision FinalAllow \longleftrightarrow simple-fw (map simple-rule-conj (generalized-fw-join (map simple-rule-dtor rs₁) (map simple-rule-dtor rs₂))) p = Decision FinalAllow*
<proof>

Using the join, it should be possible to compute any *n*-ary logical operation on firewalls. We will use it for something somewhat different in the next section.

1.3.2 Translation Implementation

This section shows the actual definition of the translation function, in Figure 1. Before beginning the translation, the definition checks whether the necessary preconditions are valid. This first two steps are to convert *fw* and *rt* to lists that can be evaluated by *generalized-sfw*. For *fw*, this is done by *map simple-rule-dtor*, which just deconstructs *simple-rules* into tuples of match and action. For *rt*, we made a firewall ruleset with rules that use prefix matches on the destination IP address. The next step is to join the two rulesets. The result of the join is a ruleset with rules *r* that only match if both, the corresponding firewall rule *fwr* and the corresponding routing rule *rr* matches. The data accompanying *r* is the port from *rr* and the firewall decision from *fwr*. Next, descending priorities are added to the rules using *map (apfst word-of-nat) \circ annotate-rln*. If the number of rules is too large to fit into the 2¹⁶ priority classes, an error is returned. Otherwise, the function *pack-OF-entries* is used to convert the (16 word \times 32 simple-match \times

char list \times simple-action) list to an OpenFlow table. While converting the *char list \times simple-action* tuple is straightforward, converting the *simple-match* to an equivalent list of *of-match-field set* is non-trivial. This is done by the function *simple-match-to-of-match*.

The main difficulties for *simple-match-to-of-match* lie in making sure that the prerequisites are satisfied and in the fact that a *simple-match* operates on slightly stronger match expressions.

- A *simple-match* allows a (string) prefix match on the input and output interfaces. Given a list of existing interfaces on the router *ifs*, the function has to insert flow entries for each interface matching the prefix.
- A *simple-match* can match ports by an interval. Now it becomes obvious why Section 1.2.1 added bitmasks to *L4Src* and *L4Dst*. Using the algorithm to split word intervals into intervals that can be represented by prefix matches from [6], we can efficiently represent the original interval by a few (32 in the worst case) prefix matches and insert flow entries for each of them.⁹

The following lemma characterizes *simple-match-to-of-match*:

lemma *simple-match-to-of-match*:

assumes

simple-match-valid r
p-iiface p \in set ifs
match-iface (oiface r) (p-oiface p)
p-l2type p = 0x800

shows

simple-matches r p \longleftrightarrow (\exists gr \in set (simple-match-to-of-match r ifs). OF-match-fields gr p = Some True)

<proof>

The assumptions are to be read as follows:

- The match *r* has to be valid, i.e. it has to use *valid-prefix* matches, and it cannot use anything other than 0-65535 for the port matches unless its protocol match ensures *TCP*, *UDP* or *L4-Protocol.SCTP*.

⁹It might be possible to represent the interval match more efficiently than a split into prefixes. However, that would produce overlapping matches (which is not a problem if we assign separate priorities) and we did not have a verified implementation of an algorithm that does so.

```

lemma lr-of-tran rt fw ifs ≡
if ¬ (no-oif-match fw ∧ has-default-policy fw ∧ simple-fw-valid fw ∧ valid-prefixes rt ∧ has-default-route rt ∧
distinct ifs)
  then Inl "Error in creating OpenFlow table: prerequisites not satisfied"
  else (
let
  nfw = map simple-rule-dtor fw;
  firt = map (λr. (route2match r, output-iface (routing-action r))) rt;
  nrd = generalized-fw-join firt nfw;
  ard = (map (apfst of-nat) ∘ annotate-rlen) nrd
in
if length nrd < unat (- 1 :: 16 word)
  then Inr (pack-OF-entries ifs ard)
  else Inl "Error in creating OpenFlow table: priority number space exhausted"
)
⟨proof⟩

```

Figure 1: Function for translating a '*i* simple-rule list, a '*i* routing-rule list, and a list of interfaces to a flow table.

- *simple-match-to-of-match* cannot produce rules for packets that have input interfaces that are not named in the interface list.
- The output interface of p has to match the output interface match of r . This is a weakened formulation of *oiface* $r = ifaceAny$, since

match-iface ifaceAny i

. We require this because OpenFlow field matches cannot be used to match on the output port — they are supposed to match a packet and decide an output port.

- The *simple-match* type was designed for IP(v4) packets, we limit ourselves to them.

The conclusion then states that the *simple-match* r matches iff an element of the result of *simple-match-to-of-match* matches. The third assumption is part of the explanation why we did not use *simple-linux-router-altered: simple-match-to-of-match* cannot deal with output interface matches. Thus, before passing a generalized simple firewall to *pack-OF-entries*, we would have to set the output ports to *ifaceAny*. A system replace output interface matches with destination IP addresses has already been formalized and will be published in a

future version of [5]. For now, we limit ourselves to firewalls that do not do output port matching, i.e., we require *no-oif-match fw*.

Given discussed properties, we present the central theorem for our translation in Figure 2. The first two assumptions are limitations on the traffic we make a statement about. Obviously, we will never see any packets with an input interface that is not in the interface list. Furthermore, we do not state anything about non-IPv4 traffic. (The traffic will remain unmatched in by the flow table, but we have not verified that.) The last assumption is that the translation does not return a run-time error. The translation will return a run-time error if the rules can not be assigned priorities from a 16 bit integer, or when one of the following conditions on the input data is not satisfied:

lemma

```

¬ no-oif-match fw ∨
¬ has-default-policy fw ∨
¬ simple-fw-valid fw ∨
¬ valid-prefixes rt ∨
¬ has-default-route rt ∨
¬ distinct ifs ⇒
∃ err. lr-of-tran rt fw ifs = Inl err ⟨proof⟩

```

```

theorem
fixes
  p :: (32, 'a) simple-packet-ext-scheme
assumes
  p-iiface p ∈ set ifs and p-l2type p = 0x800
  lr-of-tran rt fw ifs = Inr oft
shows
  OF-priority-match OF-match-fields-safe oft p = Action [Forward oif]  $\longleftrightarrow$  simple-linux-router-nol12 rt fw p =
  (Some (p(p-oiface := oif)))
  OF-priority-match OF-match-fields-safe oft p = Action []  $\longleftrightarrow$  simple-linux-router-nol12 rt fw p = None
  OF-priority-match OF-match-fields-safe oft p  $\neq$  NoAction  $\longleftrightarrow$  OF-priority-match OF-match-fields-safe oft p  $\neq$ 
  Undefined
  OF-priority-match OF-match-fields-safe oft p = Action ls  $\longrightarrow$  length ls  $\leq$  1
   $\exists$  ls. length ls  $\leq$  1  $\wedge$  OF-priority-match OF-match-fields-safe oft p = Action ls
⟨proof⟩

```

Figure 2: Central theorem on *lr-of-tran*

1.3.3 Comparison to Exodus

We are not the first researchers to attempt automated static migration to SDN. The (only) other attempt we are aware of is *Exodus* by Nelson *et al.* [10].

There are some fundamental differences between Exodus and our work:

- Exodus focuses on Cisco IOS instead of linux.
- Exodus does not produce OpenFlow rulesets, but FlowLog [9] controller programs.
- Exodus is not limited to using a single flow table.
- Exodus requires continuous controller interaction for some of its functions.
- Exodus attempts to support as much functionality as possible and has implemented support for dynamic routing, VLANs and NAT.
- Nelson *et al.* reject the idea that the translation could or should be proven correct.

2 Evaluation

In Section 1, we have made lots of definitions and created lots of models. How far these models are in accordance with the real world has been up to the vigilance of

the reader. This section attempts to alleviate this burden by providing some examples.

2.1 Mininet Examples

The first example is designed to be minimal while still showing the most important properties of our conversion. For this purpose, we used a linux firewall F, that we want to convert. We gave it two interfaces, and connected one client each. Its original configuration and the ruleset resulting from the translation is shown in Figure 3. (The list of interfaces can be extracted from the routing table; `s1-lan` received port number 1.) While the configuration does not fulfil any special function (especially, no traffic from the interface `s1-wan` is permitted), it is small enough to let us have a detailed look. More specifically, we can see how the only firewall rule (Line 2) got combined with the first rule of the routing table to form Line 1 of the OpenFlow rules. This also shows why the bitmasks on the layer 4 ports are necessary. If we only allowed exact matches, we would have 2^{15} rules instead of just one. Line 2 of the OpenFlow ruleset has been formed by combining the default drop policy with Line 1 of the routing table. In a similar fashion, Line 2 of the routing rules has also been combined with the two firewall rules. However, as 10.0.2.0/24 from the firewall and 10.0.1.0/24 from the routing table have no common elements, no rule results

```

1 :FORWARD DROP [0:0]
2 -A FORWARD -d 10.0.2.0/24 -i s1-lan -p tcp -
  m tcp --sport 32768:65535 --dport 80 -j
  ACCEPT

```

(a) FORWARD chain

```

1 10.0.2.0/24 dev s1-wan proto kernel scope
  link src 10.0.2.4
2 10.0.1.0/24 dev s1-lan proto kernel scope
  link src 10.0.1.1
3 default via 10.0.2.1 dev s1-wan

```

(b) Routing table (sorted)

```

1 priority=4,hard_timeout=0,idle_timeout=0,in_port=1,dl_type=0x800,nw_proto=6,nw_dst=10.0.2.0/24,
  tp_src=32768/0x8000,tp_dst=80,action=output:2
2 priority=3,hard_timeout=0,idle_timeout=0,dl_type=0x800,nw_dst=10.0.2.0/24,action=drop
3 priority=2,hard_timeout=0,idle_timeout=0,dl_type=0x800,nw_dst=10.0.1.0/24,action=drop
4 priority=1,hard_timeout=0,idle_timeout=0,in_port=1,dl_type=0x800,nw_proto=6,nw_dst=10.0.2.0/24,
  tp_src=32768/0x8000,tp_dst=80,action=output:2
5 priority=0,hard_timeout=0,idle_timeout=0,dl_type=0x800,action=drop

```

(c) Resulting OpenFlow rules

Figure 3: Example Network 1 – Configuration

from combining Line 2 and Line 2. In a similar fashion, the rest of the OpenFlow ruleset can be explained.

We feel that it is also worth noting again that it is necessary to change the IP configuration of the two devices attached to F. Assuming they are currently configured with, e.g., 10.0.1.100/24 and 10.0.2.1/24, the subnet would have to be changed from 24 to 22 or lower to ensure that a common subnet is formed and the MAC layer can function properly.

Next, we show a somewhat more evolved example. Its topology is depicted in Figure 4a. As before, we called the device to be replaced F. It is supposed to implement the simple policy that the clients H1 and H2 are allowed to communicate with the outside world via HTTP, ICMP is generally allowed, any other traffic is to be dropped (we neglected DNS for this example). We used the iptables configuration that is shown in Figure 4b. The routing table is the same as in the first example network.


The topology has been chosen for a number of reasons: we wanted one device which is not inside a common subnet with F and thus requires no reconfiguration for the translation. Moreover, we wanted two devices in a network that can communicate with each other while being overheard by F. For this purpose, we added two clients H1 and H2 instead of just one. We connected them with a broadcasting device.¹⁰

¹⁰For the lack of a hub in mininet, we emulated one with an OpenFlow switch.

Executing our conversion function results in 36 rules¹¹, we decided not to include them here. Comparing to the first example network, the size of the ruleset seems relatively high. This can be explained by the port matches: 1024-65535 has to be expressed by 6 different matches, `tp_src=1024/0xfc00`, `tp_src=2048/0xf800`, ..., `tp_src=32768/0x8000` (or `tp_dst` respectively). When installing these rules, we also have to move all of H1, H2 and S1 into a common subnet. We chose 10.0.0.0/16 and updated the IP configuration of the three hosts accordingly. As discussed, the configuration of S2 did not have to be updated, as it does not share any subnet with F. We then tested reachability for TCP 22 and 80 and ICMP. The connectivity between all pairs of hosts (H1,H2,S1 and S2) remained the same compared to before the conversion. This shows that the concept can be made to work.

However, the example also reveals a flaw: When substituting the more complete model of a linux firewall with the simple one in Section 1.1, we assumed that the check whether the correct MAC address is set and the packets are destined for the modelled device would never fail — we assumed that all traffic arriving at a device is actually destined for it. Obviously, this network violates this assumption. We can trigger this in many ways, for example by sending an ICMP ping from H1 to H2. This will cause the generated rule `priority=7, icmp, nw_dst=10.0.1.0/24 actions=output:1` (where port

¹¹If we had implemented some spoofing protection by adding `! -s 10.0.1.0/24` to the respective rule, the number of rules would have been increased to 312. This is because a cross product of two prefix splits would occur.



(a) Topology

```

1 :FORWARD DROP [0:0]
2 -A FORWARD -p icmp -j ACCEPT
3 -A FORWARD -i s1-lan -p tcp -m tcp --sport
  1024:65535 --dport 80 -j ACCEPT
4 -A FORWARD -d 10.0.1.0/24 -i s1-wan -p tcp -m tcp
  --sport 80 --dport 1024:65535 -j ACCEPT

```

(b) FORWARD chain

Figure 4: Example Network 2

1 is the port facing H1 and H2) to be activated twice. This is obviously not desired behavior. Dealing with this is, as mentioned, future work.

2.2 Performance Evaluation

Unfortunately, we do not have any real-world data that does not use output port matches as required in Section 1.3. There is thus no way to run the translation on the real-world firewall rulesets we have available and obtain a meaningful result. Nevertheless, we can use a real-world ruleset to evaluate the performance of our translation. For this purpose, we picked the largest firewall from the firewall collection from [6]. A significant amount of time is necessary to convert its `FORWARD` chain including 4946 rules¹² to the required simplified firewall form. Additionally to the simplified firewall, we acquired the routing table (26 entries) from the same machine. We then evaluated the time necessary to complete the translation and the size of the resulting rule-set when using only the first n simple firewall rules and the full routing table. The result is shown in Figure 5. Given the time necessary to complete the conversion of the iptables firewall to a simple firewall, it is reasonable to say that the translation function is efficient enough.

¹²In the pre-parsed and already normalized version we used for this benchmark, it took 45s. The full required time lies closer to 11min as stated in [6].




Figure 5: Benchmark

At first glance, size of the resulting ruleset seems high. This can be explained by two facts:

- The firewall contains a large number of rules with port matches that allow the ports 1-65535, which requires 16 OpenFlow rules.
- Some combinations of matches from the firewall and the routing table cannot be ruled out, since the firewall match might only contain an output port and the rule can thus only apply for the packets matching a few routing table entries. However, the translation is not aware of that and can thus not remove the combination of the firewall rule and other routing table entries.

In some rules, the conditions above coincide, resulting in 416 ($= 16 \cdot 26$) rules. To avoid the high number of rules resulting from the port matches, rules that forbid packets with source or destination port 0 could be added to the start of the firewall and the 1-65535 could be removed; dealing with the firewall / routing table problem is part of the future work on output interfaces.

3 Conclusion and Future Work

We believe that we have shown that it is possible to translate at least basic configurations of a linux firewall into OpenFlow rulesets while preserving the most important aspects of the behavior. We recognize that our system has limited practical applicability. One possible example would be a router or firewall inside a company network whose state tables have been polluted by special attack traffic. Our translation could provide an OpenFlow based stateless replacement. However, given the current prerequisites the implementation has on the configuration, this application is relatively unlikely.

For the configuration translation, we have contributed formal models of a linux firewall and of an OpenFlow switch. Furthermore, the function that joins two firewalls and the function that translates a simplified match from [6] to a list of equivalent OpenFlow field match sets are contributions that we think are likely to be of further use.

We want to explicitly formulate the following two goals for our future work:

- We want to deal with output interface matches. The idea is to formulate and verify a destination

interface / destination IP address rewriting that can exchange output interfaces and destination IP addressed in a firewall, based on the information from the routing table.¹³

- We want to develop a system that can provide a stricter approximation of stateful matches so our translation will be applicable in more cases.

References

- [1] Open vSwitch. <http://openvswitch.org/>.
- [2] OpenFlow Switch Specification v1.0.0, December 2009.
- [3] OpenFlow Switch Specification v1.5.1, March 2015.
- [4] B. Barras, S. Boutin, C. Cornes, J. Courant, J.-C. Filliatre, E. Gimenez, H. Herbelin, G. Huet, C. Munoz, C. Murthy, et al. The coq proof assistant reference manual: Version 6.1. 1997.
- [5] C. Diekmann and L. Hupel. Iptables Semantics. *Archive of Formal Proofs*, Sept. 2016. http://isa-afp.org/entries/Iptables_Semantics.shtml, Formal proof development.
- [6] C. Diekmann, J. Michaelis, M. Haslbeck, and G. Carle. Verified iptables Firewall Analysis. In *Proceedings of IFIP Networking 2016 (NET-WORKING 16)*, May 2016.
- [7] A. Guha, M. Reitblatt, and N. Foster. Machine-verified Network Controllers. In *Proceedings of the 34th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '13*, pages 483–494, New York, NY, USA, 2013. ACM.
- [8] J. Michaelis and C. Diekmann. Middlebox models in network verification research. In *Proceedings of the Seminars Future Internet (FI) and Innovative Internet Technologies and Mobile Communications (IITM), Winter Semester 2015/2016*, volume 17, 2016.
- [9] T. Nelson, A. D. Ferguson, M. J. Scheer, and S. Krishnamurthi. Tierless programming and reasoning for software-defined networks. In *11th USENIX*

¹³As of now this has already been implemented, but is not yet fully ready.

Symposium on Networked Systems Design and Implementation (NSDI 14), pages 519–531, 2014.

- [10] T. Nelson, A. D. Ferguson, D. Yu, R. Fonseca, and S. Krishnamurthi. Exodus: toward automatic migration of enterprise network configurations to SDNs. In *Proceedings of the 1st ACM SIGCOMM*

Symposium on Software Defined Networking Research, page 13. ACM, 2015.

- [11] T. Nipkow, L. C. Paulson, and M. Wenzel. *Isabelle/HOL — A Proof Assistant for Higher-Order Logic*, volume 2283 of *LNCS*. Springer, 2015.