

CS 244

Network Verification

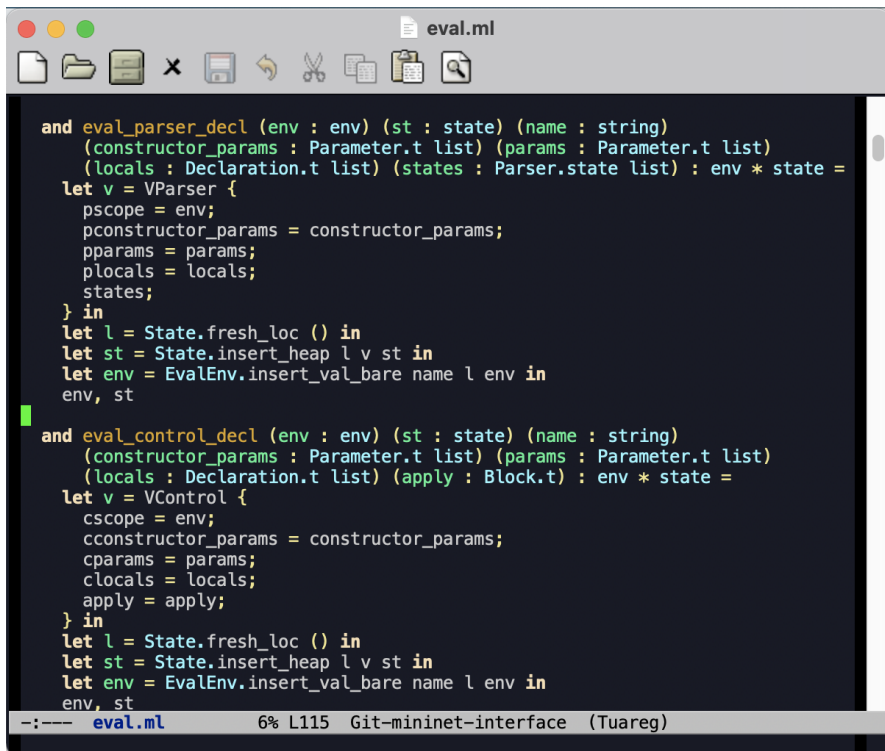
Nate Foster
Cornell University



Opening Survey

<https://pollev.com/jnfoster>

When I need to troubleshoot buggy code...

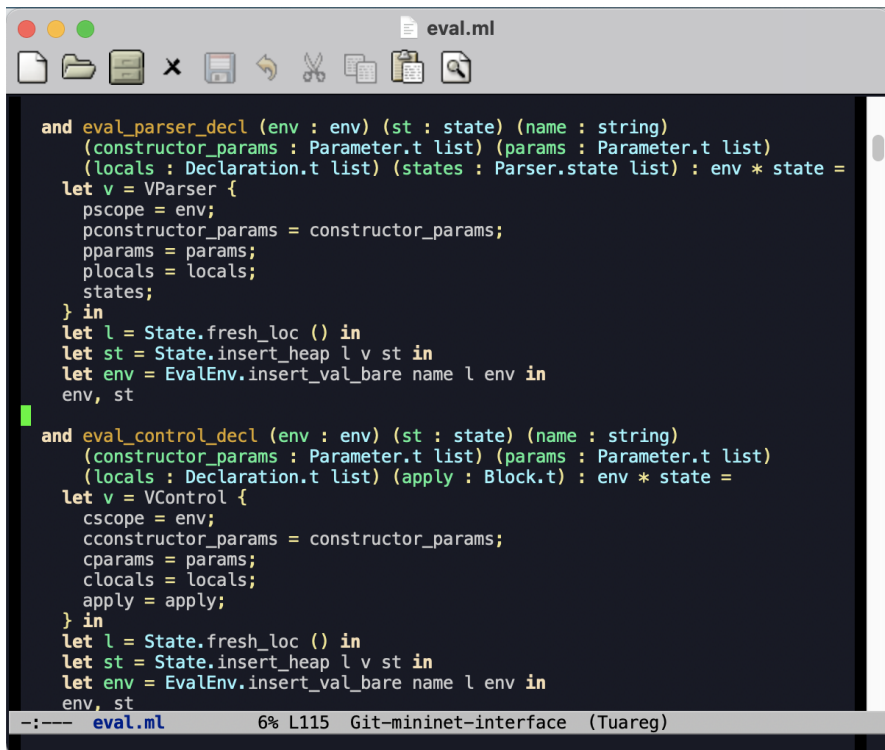


```
and eval_parser_decl (env : env) (st : state) (name : string)
  (constructor_params : Parameter.t list) (params : Parameter.t list)
  (locals : Declaration.t list) (states : Parser.state list) : env * state =
  let v = VParser {
    pscope = env;
    pconstructor_params = constructor_params;
    pparams = params;
    plocals = locals;
    states;
  } in
  let l = State.fresh_loc () in
  let st = State.insert_heap l v st in
  let env = EvalEnv.insert_val_bare name l env in
  env, st

and eval_control_decl (env : env) (st : state) (name : string)
  (constructor_params : Parameter.t list) (params : Parameter.t list)
  (locals : Declaration.t list) (apply : Block.t) : env * state =
  let v = VControl {
    cscope = env;
    cconstructor_params = constructor_params;
    cparams = params;
    clocals = locals;
    apply = apply;
  } in
  let l = State.fresh_loc () in
  let st = State.insert_heap l v st in
  let env = EvalEnv.insert_val_bare name l env in
  env, st
```

--:---- eval.ml 6% L115 Git-mininet-interface (Tuareg)

When I need to troubleshoot buggy code...



```
eval.ml

and eval_parser_decl (env : env) (st : state) (name : string)
  (constructor_params : Parameter.t list) (params : Parameter.t list)
  (locals : Declaration.t list) (states : Parser.state list) : env * state =
  let v = VParser {
    pscope = env;
    pconstructor_params = constructor_params;
    pparams = params;
    plocals = locals;
    states;
  } in
  let l = State.fresh_loc () in
  let st = State.insert_heap l v st in
  let env = EvalEnv.insert_val_bare name l env in
  env, st

and eval_control_decl (env : env) (st : state) (name : string)
  (constructor_params : Parameter.t list) (params : Parameter.t list)
  (locals : Declaration.t list) (apply : Block.t) : env * state =
  let v = VControl {
    cscope = env;
    cconstructor_params = constructor_params;
    cparams = params;
    clocals = locals;
    apply = apply;
  } in
  let l = State.fresh_loc () in
  let st = State.insert_heap l v st in
  let env = EvalEnv.insert_val_bare name l env in
  env, st

--:-- eval.ml 6% L115 Git-mininet-interface (Tuareg)
```



Credit: Wikipedia

Software Validation: A Spectrum

Social

- Code reviews
- Pair programming

Methodological


- Test-driven development
- Version control
- Bug tracking

Technological

- Static "linters"
- Fuzzers

Mathematical

- Type systems
- Formal verification



Less formal: techniques are easy to use but may miss some problems in programs

Best practice: all of these techniques should be used!



More formal: eliminate *with certainty* as many problems as possible, but may be hard to use

What about computer networks?



```
nate — ping 8.8.8.8 — 70x24
Last login: Tue Apr 27 05:20:52 on ttys000
~
05:25 $ ping 8.8.8.8
PING 8.8.8.8 (8.8.8.8): 56 data bytes
64 bytes from 8.8.8.8: icmp_seq=0 ttl=117 time=29.864 ms
64 bytes from 8.8.8.8: icmp_seq=1 ttl=117 time=30.463 ms
64 bytes from 8.8.8.8: icmp_seq=2 ttl=117 time=29.036 ms
64 bytes from 8.8.8.8: icmp_seq=3 ttl=117 time=30.427 ms
64 bytes from 8.8.8.8: icmp_seq=4 ttl=117 time=36.375 ms
64 bytes from 8.8.8.8: icmp_seq=5 ttl=117 time=32.785 ms
64 bytes from 8.8.8.8: icmp_seq=6 ttl=117 time=34.360 ms
```

"You're On Your Own Mate" —Nick McKeown

Network Verification Pre-History

- Proposed a static analysis to determine network reachability
- Based on an underlying model that captures IP forwarding
- Extension supports richer behaviors like NAT and middleboxes

On Static Reachability Analysis of IP Networks

Geoffrey G. Xie^{*}
Albert Greenberg[†]

Jibin Zhan[†]
Gisli Hjalmtysson[‡]

David A. Maltz[†]
Jennifer Rexford[‡]

Hui Zhang[†]
Jennifer Rexford[‡]

ABSTRACT

The primary purpose of a network is to provide reachability between applications running on end hosts. In this paper, we describe how to compute the reachability a network provides from a snapshot of the configuration state from each of the routers. Our primary contribution is the precise definition of the *potential reachability* of a network and a substantial simplification of the problem through a unified modeling of packet filters and routing protocols. In the end, we reduce a complex, important practical problem to computing the transitive closure to set union and intersection operations on reachability set representations. We then extend our algorithm to model the influence of packet transformations (e.g., by NATs or ToS remapping) along the path. Our technique for static analysis of network reachability is valuable for verifying the intent of the network designer, troubleshooting reachability problems, and performing “what-if” analysis of failure scenarios.

Index Terms—Routing, Static Configuration Analysis.

I. INTRODUCTION

While the ultimate goal of networking is to enable communication between hosts that are not directly connected, a wide variety of mechanisms are being used to *limit* the set of destinations the hosts can reach. For example, backbone networks may provide Virtual Private Network services to connect only remote offices belonging to the same enterprise, and enterprise networks themselves are often segmented into departments or offices whose hosts must

Determining what kinds of packets can be exchanged between two hosts connected to a network is a difficult and critical problem facing network designers and operators. To our knowledge, the problem is largely unexamined in the networking research literature. Solving the problem requires knowing far more than the network’s topology or the routing protocols it uses. For example, despite having a route to a remote end-point, a sender’s packets may be discarded by a packet filter on one of the links in the path. The network’s packet filters, routing policies, and packet transformations all must be taken into account to even ask the simple and very important question of “can these two hosts communicate?”

This paper crystallizes the problem of calculating the *reachability* provided by a network. By mapping packet filters, routing information, and packet transformations to a single unified model of reachability we have determined how to transform this seemingly intractable problem into a classical graph problem that can be solved with polynomial time algorithms such as transitive closure. This is the primary contribution of this paper.

A. Advantages of Automated Static Analysis

Currently, the common practice to determine if packets can reach from one point in a network to another is to use tools such as `ping` and `traceroute` to send probe traffic that experimentally test whether reachability exists. In contrast, we have developed a *static-analysis* approach that can be applied even if only a description of the network is available. Static analysis has many advantages

Network Verification in SDN

Pre-SDN

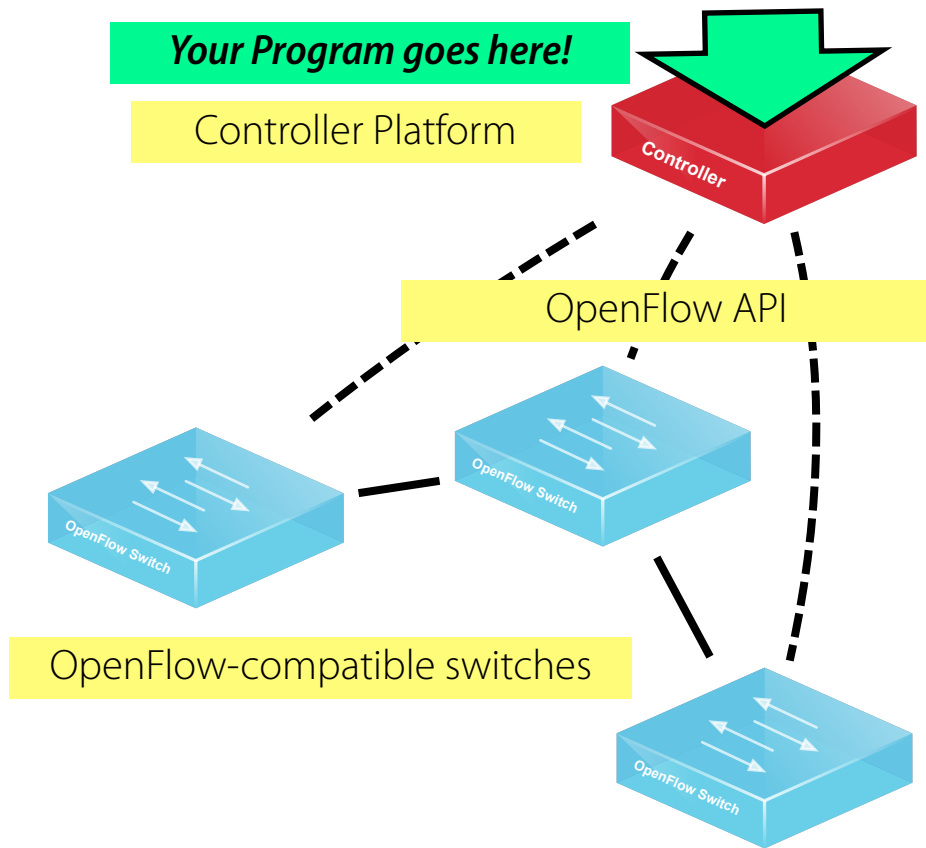
- Distributed control plane
- Complex data plane
 - Dozens of protocols
 - Tricky, undocumented semantics

Post-SDN

- Centralized control plane
- Streamlined data plane
 - OpenFlow 1.0: only 12 protocols!
 - Clear semantics

Key Insight

- Can instrument control plane to build a model of the data plane behavior
- Can reason statically about network-wide forwarding properties



Plan for Today

Header Space Analysis

NetKAT


Key Questions:

- How to encode networks and properties?
- How to automate reasoning?
- ~~How scalable and how fast?~~


HSA [NSDI '12]

Packet Model

- Fix the set of headers (Ethernet, IPv4, TCP, UDP, etc.) used by devices
- If L is the total number of bits used to encode all headers...
- Then a packet can be seen as a point in an L -dimensional space
- Can also assign each port a unique identifier and add a "pseudo header" to track packet's location
- Formally: $\{0, 1\}^L \times \{1, \dots, P\}$

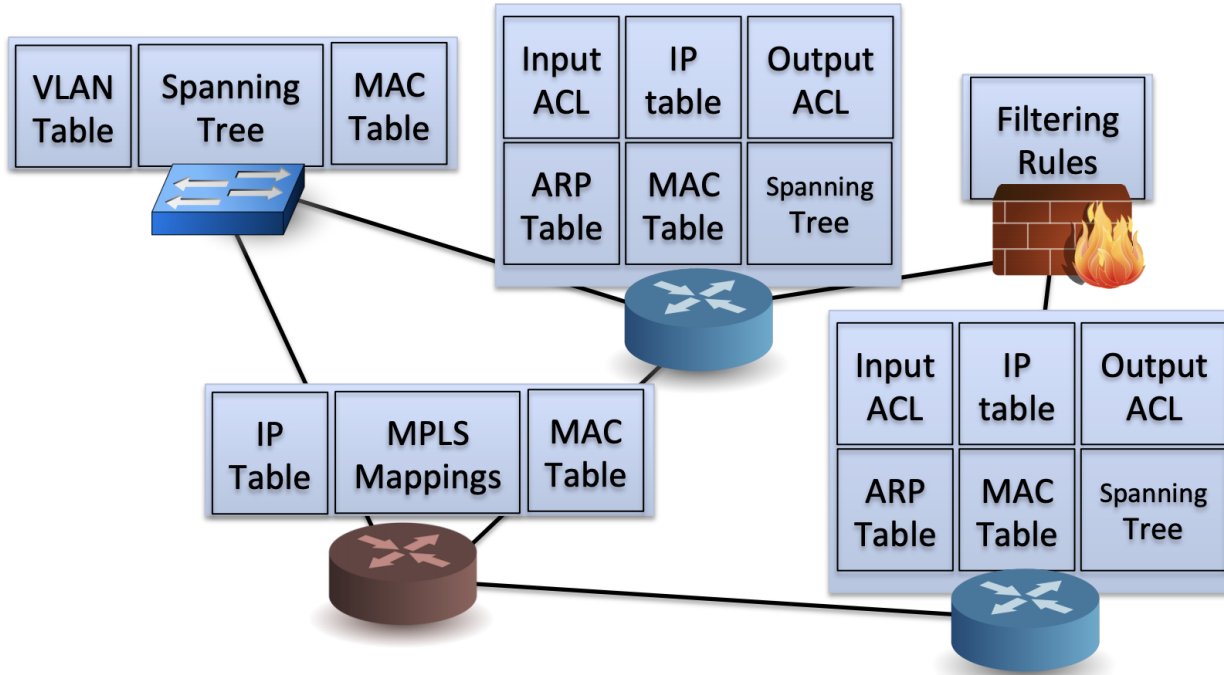


Forwarding Model



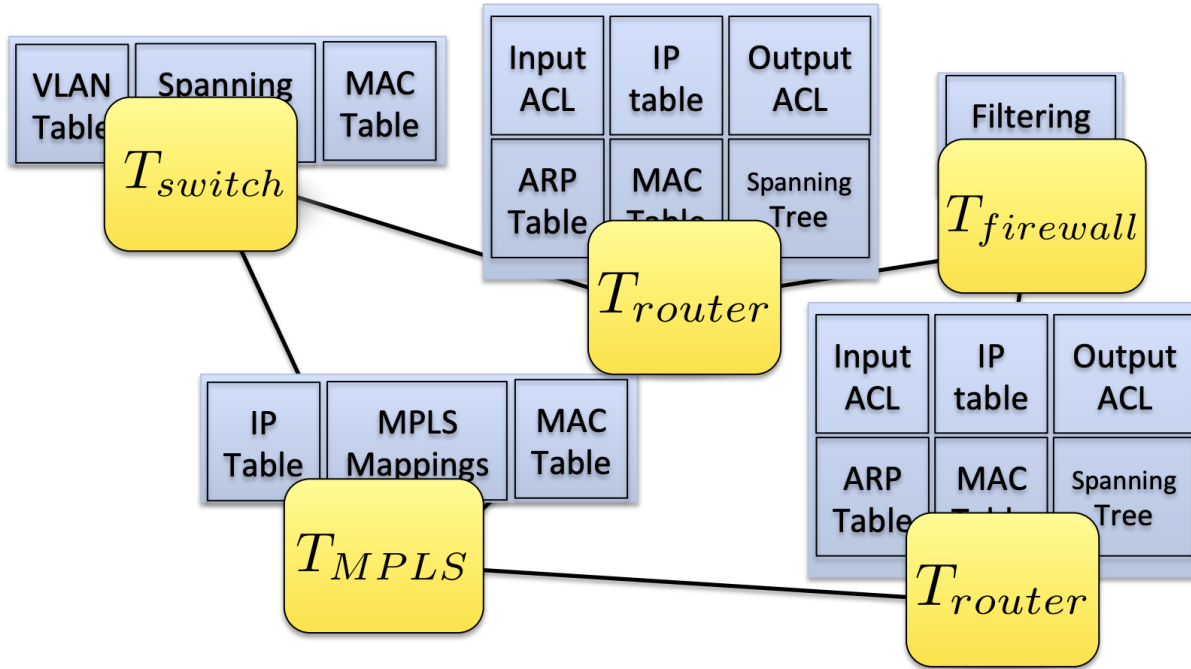
Packet forwarding can be viewed as a transformation on *header space*

Device Complexity



Network devices seem complex, with many different features and protocols...

Transfer Functions



... but ultimately they too can be modeled as *transfer functions* on packet space


Formalizing Transfer Functions

$T \in \text{Header} \times \text{Port} \rightarrow (\text{Header} \times \text{Port}) \text{ Set}$

Formalizing Transfer Functions

$T \in \text{Header} \times \text{Port} \rightarrow (\text{Header} \times \text{Port}) \text{ Set}$

- 172.24.74.0 255.255.255.0 Port1
- 172.24.128.0 255.255.255.0 Port2
- 171.67.0.0 255.255.0.0 Port3




$$T(h, p) = \begin{cases} (h, 1) & \text{if } \text{dst_ip}(h) = 172.24.74.x \\ (h, 2) & \text{if } \text{dst_ip}(h) = 172.24.128.x \\ (h, 3) & \text{if } \text{dst_ip}(h) = 171.67.x.x \end{cases}$$

Symbolic Representation


Scaling Challenges

- We now have a foundational model of forwarding behavior
- But the space of packets is huge...
- ...and the space of functions on that space is larger still...
- Unclear how to realize this model in an implementation that scales well!



Insight: Networks are (Mostly) Uniform

- In practice, most transfer functions transform the packet space in a (mostly) uniform way
- Sets of packets can be viewed as regions (i.e., hypercubes) in the same N-dimensional space
- So we can build symbolic representations that manipulate regions of header space rather than individual points in packet space



Header Space Algebra

Elements

Every region of header space can be represented as a union of ternary "wildcard expressions" in $\{0, 1, x\}^*$

Operations

- Equivalence & inclusion
- Union
- Intersection
- Difference
- Complement

Intersection

Single-bit intersection

| $b_i \backslash b'_i$ | 0 | 1 | x |
|-----------------------|---|---|---|
| 0 | 0 | z | 0 |
| 1 | z | 1 | 1 |
| x | 0 | 1 | x |

Definition

Intersect bit-wise, yield \emptyset if any bit is "z"

Example

$$11000xxx \cap xx00010x = 1100010x$$

Union

Definition

Simply take union of wildcard expressions, simplify if possible

Example

$1111xxxx \cup 0000xxxx$

Optimization Example

$1100xxxx \cup 1000xxxx = 1x00xxxx$

Complement


Definition

- Flip each non-wildcard bit, wildcard every other bits
- Result is union of all such expressions

Example


$$(010x)^c = 1xxx \cup x0xx \cup xx1x$$

Composing Transfer Functions



We can model network-wide behavior as the *composition* of transfer functions

Composing Transfer Functions




We can model network-wide behavior as the *composition* of transfer functions

Question: how do you reconcile the different input and output types?


Answer: "lift" the second function to $(\text{Header} \times \text{Port}) \text{ Set} \rightarrow (\text{Header} \times \text{Port}) \text{ Set}$

Domain and Range



We can compute the *domain* and *range* of an transfer function symbolically in terms of header spaces represented as wildcard expressions

Inverse Transfer Function



Can also compute the header space produced by the *inverse* of a transfer function, yielding a model of the *inputs* that map to a given set of outputs...

HSA Applications

Reachability

Goal

Want to know whether packet originating at A can get to B

Approach

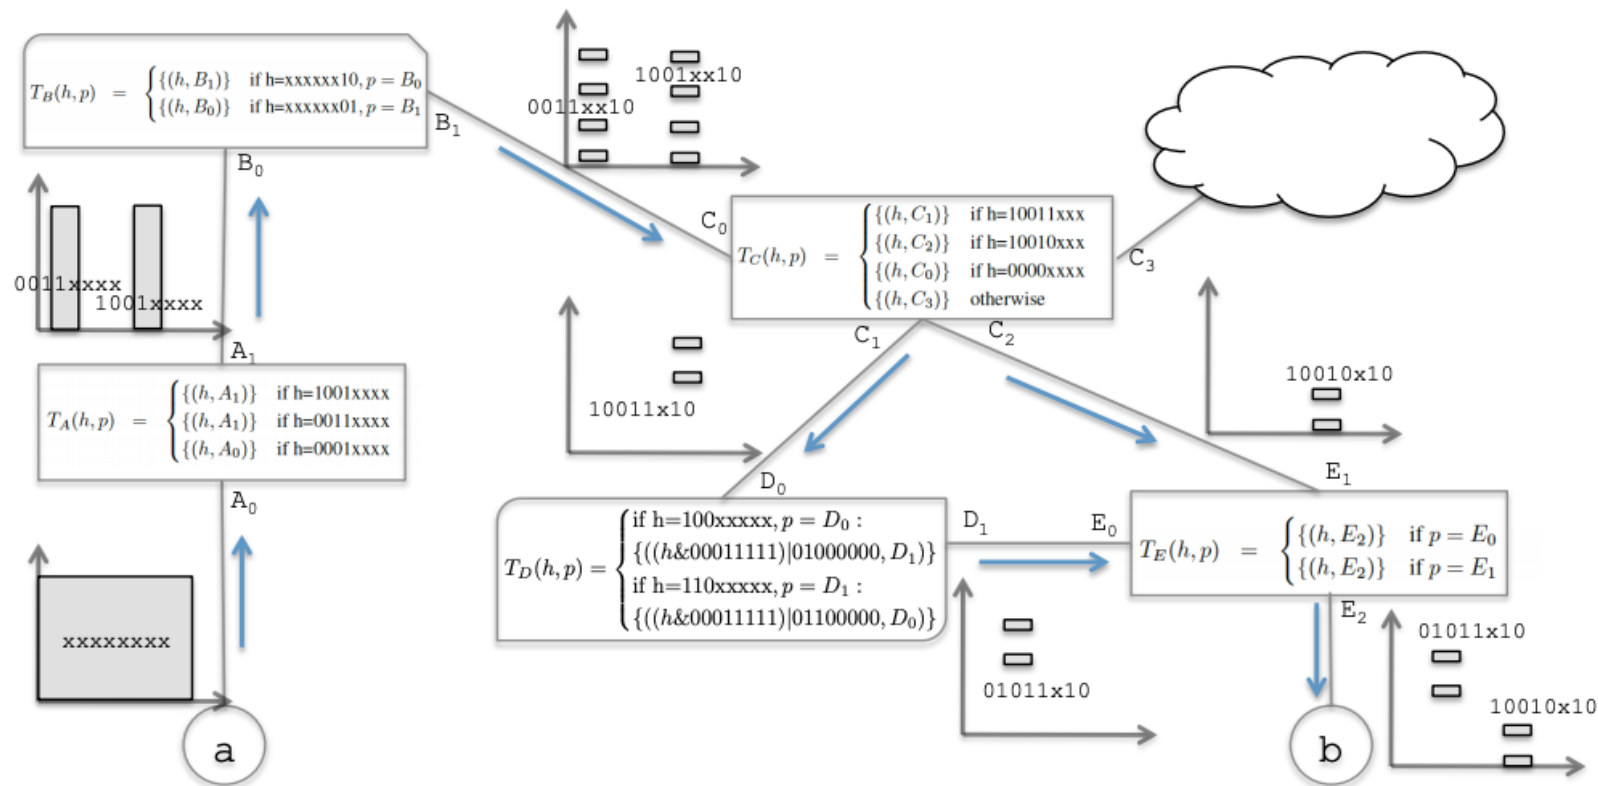
- Symbolically execute $R_{a \rightarrow b}$ on an "all wildcard" packet "xxxx..."
- Compose transfer functions for the devices path to get $R_{a \rightarrow b}$
- The result models all packets that reach B from A

Extensions

- Waypointing, Blackholing, etc.

Reachability Example

[NSDI '12, Fig 2]



Loop Freedom

Goal

Want to know whether packets can loop infinitely...



Distinction

- **Generic Loop:** a packet loops back to the same switch
- **Infinite Loop:** an *identical packet* loops back to the same switch


Approach

- Use reachability to identify generic loops
- Then analyze header spaces to identify infinite loops


Loop Freedom Example



Finite Loop



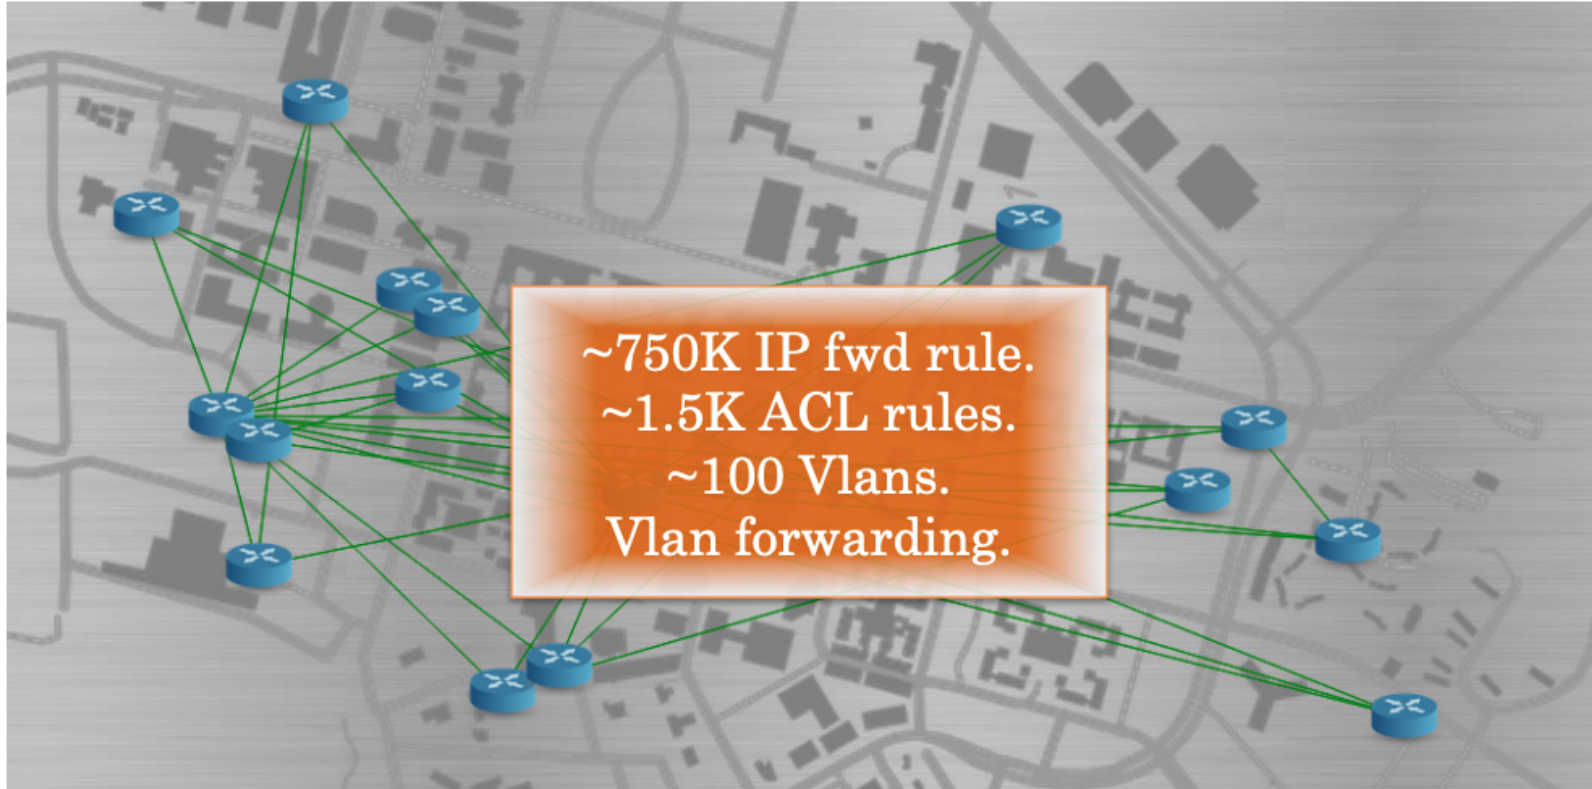
Infinite Loop



?

Performance

Stanford Campus Network (ca. 2012)



HSA Performance

On a single machine with 4 cores and 4GB Ram


| | |
|--------------------------------|-----------|
| Generating TF Rules | ~150 sec |
| Loop Detection Test (30 ports) | ~560 sec |
| Average Per Port | ~18 sec |
| Min Per Port | ~ 8 sec |
| Max Per Port | ~ 135 sec |
| Reachability Test (Avg) | ~13 sec |

NetKAT [POPL '14]

Why I ❤️ NetKAT



Theory




Applications

Why I ❤️ NetKAT



Theory



Applications


- Compilation
- Verification
- New Features

Why I ❤️ NetKAT




Theory

$\llbracket p \rrbracket$
denotational
semantics




automata
theory

$\vdash p \equiv q$
sound & complete
axiomatization



symbolic
representation



Applications

- Compilation
- Verification
- New Features

NetKAT Roadmap

Language Design &
Modeling

Reasoning &
Verification

Programming &
Compilation

NetKAT Roadmap




Language Design &
Modeling

Reasoning &
Verification


Programming &
Compilation

Essential Features

Essential Features




Essential Features




Regular Expressions

$+$, i , $*$



Essential Features




Regular Expressions

$+$, $.$, $*$




Boolean Algebra


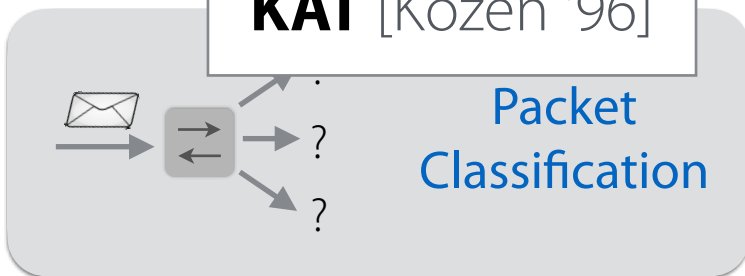
true, false, $f=n$,
 $a \& b$, $a | b$, $\neg a$



Essential Features



KAT [Kozen '96]




Regular Expressions

$+$, $|$, $*$


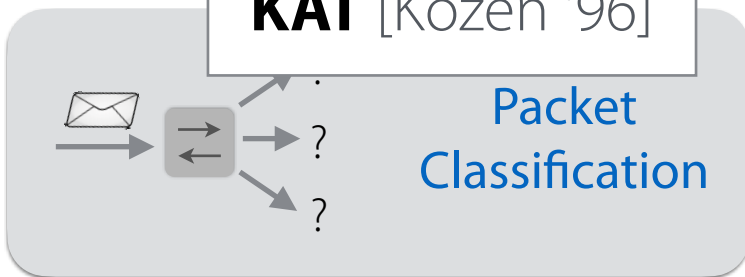
Boolean Algebra

true, false, $f=n$,
 $a \& b$, $a|b$, $\neg a$

Essential Features



KAT [Kozen '96]



Regular Expressions

$+$, $.$, $*$


Boolean Algebra

true, false, $f=n$,
 $a \& b$, $a|b$, $\neg a$

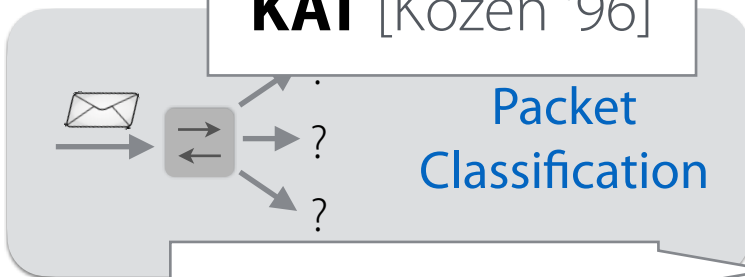
Network Primitives

$f:=n$, $A \rightarrow B$

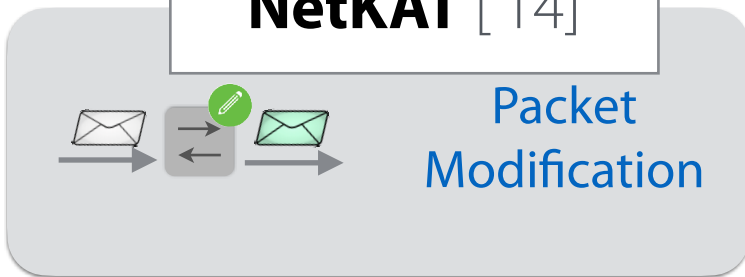
Essential Features



KAT [Kozen '96]



NetKAT ['14]



Regular Expressions

$+$, $:$, $*$

Boolean Algebra

true, false, $f=n$,
 $a \& b$, $a|b$, $\neg a$

Network Primitives


$f:=n$, $A \rightarrow B$

Example

```
port = 88; switch = 6;  
dest := 10.0.0.1;  
(port := 50 + port := 51)
```


“For all packets incoming on port 88 of switch 6,
set the destination IP address to 10.0.0.1 and
multicast the packet out of ports 50 and 51.”

Design Goal: Modular Composition



program fragments can be composed to form larger programs

Sequential Composition



“First filter out untrusted traffic, then forward.”

Sequential Composition


```
if dstport=22 then false  
else true
```

•
/

```
if dest=10.0.0.1 then port:=1  
elif dest=10.0.0.2 then port:=2  
elif dest=10.0.0.3 then port:=3  
else false
```

“First filter out untrusted traffic, then forward.”

Parallel Composition



“Execute both Monitor and Forward on all incoming packets”

Multicast: `port:=1 + port:=2`

Language Design &
Modeling

Reasoning &
Verification

Programming &
Compilation

Language Design &
Modeling

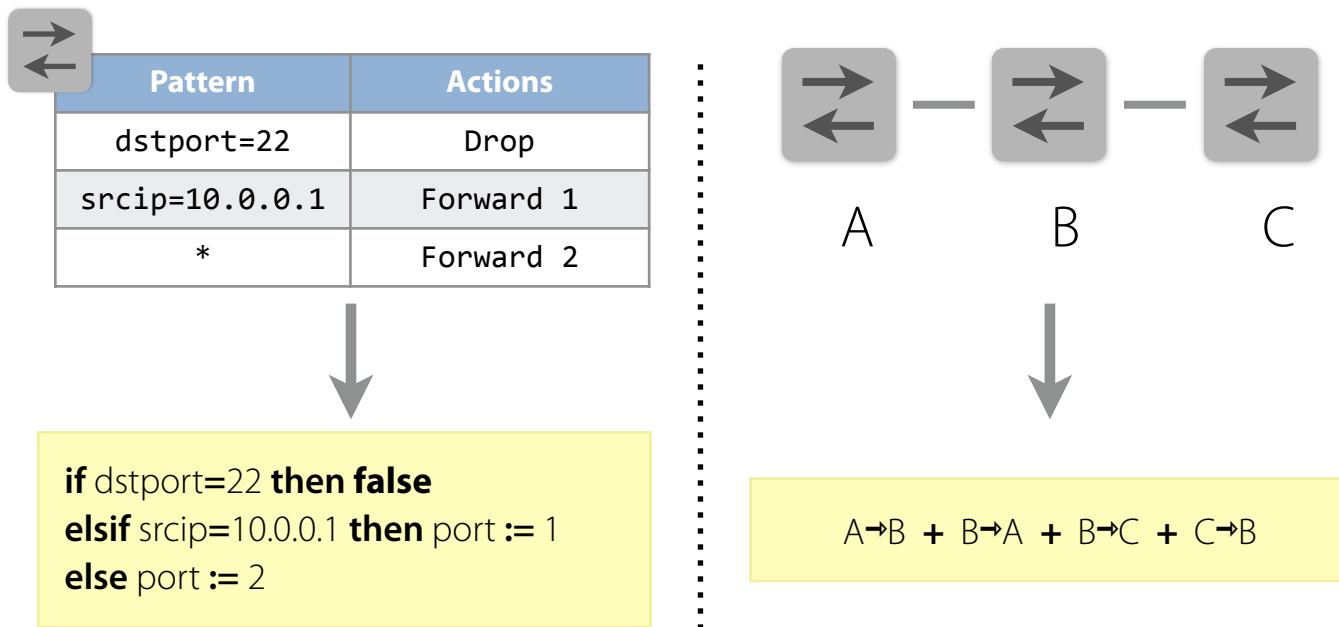


Reasoning &
Verification

Programming &
Compilation


Encoding Networks

Forwarding tables and topologies can be represented in NetKAT using straightforward encodings



Encoding Networks


A network can be encoded in NetKAT by interleaving steps of processing by switches and topology



$(\text{topology}; \text{switch})^*$

Encoding Networks


A network can be encoded in NetKAT by interleaving steps of processing by switches and topology



$(\text{topology}; \text{switch})^*$

Encoding Networks


A network can be encoded in NetKAT by interleaving steps of processing by switches and topology



$(\text{topology}; \text{switch})^*$

Encoding Networks


A network can be encoded in NetKAT by interleaving steps of processing by switches and topology



$(\text{topology}; \text{switch})^*$

Encoding Networks


A network can be encoded in NetKAT by interleaving steps of processing by switches and topology



$(\text{topology}; \text{switch})^*$


Encoding Networks

A network can be encoded in NetKAT by interleaving steps of processing by switches and topology



$(\text{topology}; \text{switch})^*$

Checking Reachability



Given a network encoded this way, we'd like to be able to automatically answer questions like:

“Does the network forward from ingress to egress?”

Can reduce this question (and others) to program equivalence

`in; (topology; switch)*; out ≠ false`


Denotational Semantics

```
pol ::=  
  | false  
  | true  
  | field = val  
  | field := val  
  | pol1 + pol2  
  | pol1 ; pol2  
  | ¬pol  
  | pol*  
  | B → A
```

Denotational Semantics

```
pol ::=  
| false  
| true  
| field = val  
| field := val  
| pol1 + pol2  
| pol1 ; pol2  
| ¬pol  
| pol*  
| B → A
```


Local: input-output behavior of switches




Denotational Semantics

```
pol ::=  
| false  
| true  
| field = val  
| field := val  
| pol1 + pol2  
| pol1 ; pol2  
| ¬pol  
| pol*  
| B → A
```

Local: input-output behavior of switches



Global: network-wide paths



NetKAT Axioms

Kleene Algebra Axioms

$$p + (q + r) \equiv (p + q) + r$$

$$p + q \equiv q + p$$

$$p + \mathbf{false} \equiv p$$

$$p + p \equiv p$$

$$p ; (q ; r) \equiv (p ; q) ; r$$

$$p ; (q + r) \equiv p \bullet q + p ; r$$

$$(p + q) ; r \equiv p ; r + q ; r$$

$$\mathbf{true} ; p \equiv p$$

$$p \equiv p ; \mathbf{true}$$

$$\mathbf{false} ; p \equiv \mathbf{false}$$

$$p ; \mathbf{false} \equiv \mathbf{false}$$

$$\mathbf{true} + p ; p^* \equiv p^*$$

$$\mathbf{true} + p^* ; p \equiv p^*$$

$$p + q ; r + r \equiv r \Rightarrow p^* ; q + r \equiv r$$

$$p + q ; r + q \equiv q \Rightarrow p ; r^* + q \equiv q$$

Boolean Algebra Axioms

$$a + (b ; c) \equiv (a + b) ; (a + c)$$

$$a + \mathbf{true} \equiv \mathbf{true}$$

$$a + \neg a \equiv \mathbf{true}$$

$$a ; b \equiv b ; a$$

$$a ; \neg a \equiv \mathbf{false}$$

$$a ; a \equiv a$$

Packet Axioms

$$f := n ; f' := n' \equiv f' := n' ; f := n \quad \text{if } f \neq f'$$

$$f := n ; f' = n' \equiv f' = n' ; f := n \quad \text{if } f \neq f'$$

$$f := n ; f = n \equiv f := n$$

$$f = n ; f := n \equiv f = n$$

$$f := n ; f := n' \equiv f := n'$$

$$f = n ; f = n' \equiv \mathbf{false} \quad \text{if } n \neq n'$$

$$\mathbf{A} \rightarrow \mathbf{B} ; f = n \equiv f = n ; \mathbf{A} \rightarrow \mathbf{B} \quad \text{if } f \notin \{\text{switch, port}\}$$

$$\sum_i f = n_i \equiv \mathbf{true}$$

NetKAT Axioms

Kleene Algebra Axioms

$$p + (q + r) \equiv (p + q) + r$$

$$p + q \equiv q + p$$

$$p + \text{false} \equiv p$$

$$p + p \equiv p$$

$$p ; (q ; r) \equiv$$

$$p ; (q + r) \equiv$$

$$(p + q) ; r \equiv$$

$$\text{true} ; p \equiv p$$

$$p \equiv p ; \text{true}$$

$$\text{false} ; p \equiv \text{false}$$

$$p ; \text{false} \equiv \text{false}$$

$$\text{true} + p ; p^* \equiv p^*$$

$$\text{true} + p^* ; p \equiv p^*$$

$$p + q ; r + r \equiv r \Rightarrow p^* ; q + r \equiv r$$

$$p + q ; r + q \equiv q \Rightarrow p ; r^* + q \equiv q$$

Boolean Algebra Axioms

$$a + (b ; c) \equiv (a + b) ; (a + c)$$

$$a + \text{true} \equiv \text{true}$$

$$a + \neg a \equiv \text{true}$$

$$a \cdot b = b \cdot a$$

Soundness: If $\vdash p \equiv q$, then $\llbracket p \rrbracket = \llbracket q \rrbracket$

Completeness: If $\llbracket p \rrbracket = \llbracket q \rrbracket$, then $\vdash p \equiv q$

$$f := n ; f' = n' \equiv f' = n' ; f := n \quad \text{if } f \neq f'$$

$$f := n ; f = n \equiv f := n$$

$$f = n ; f := n \equiv f = n$$

$$f = n ; f := n \equiv f = n$$

$$f := n ; f := n' \equiv f := n'$$

$$f = n ; f = n' \equiv \text{false} \quad \text{if } n \neq n'$$

$$\mathbf{A} \rightarrow \mathbf{B} ; f = n \equiv f = n ; \mathbf{A} \rightarrow \mathbf{B} \quad \text{if } f \notin \{\text{switch, port}\}$$

$$\sum_i f = n_i \equiv \text{true}$$

Decision Procedure



Decides program equivalence fully automatically!

Theoretical Insight: NetKAT programs \leftrightarrow NetKAT automata

Decision Procedure

Decides program equivalence fully automatically!


Theoretical Insight: NetKAT programs \leftrightarrow NetKAT automata



Decision Procedure

Decides program equivalence fully automatically!

Theoretical Insight: NetKAT programs \leftrightarrow NetKAT automata



Algorithm checks bisimilarity of automata

Language Design &
Modeling

Reasoning &
Verification

Programming &
Compilation

Language Design &
Modeling




Reasoning &
Verification




Programming &
Compilation


Example




Local Program



Local Program




Local Program



```
port=1; tag:=1; port:=3 +  
port=2; tag:=2; port:=3
```

???


Local Program



port=1; tag:=1; port:=3
+
port=2; tag:=2; port:=3

tag=1; port:=5
+
tag=2; port:=6

Local Program




```
port=1; tag:=1; port:=3  
+  
port=2; tag:=2; port:=3
```


```
tag=1; port:=5  
+  
tag=2; port:=6
```

Tedious for programmers... difficult to get right!

Global Program




Global Program



port=1; **A→B**; port:=5
+
port=2; **A→B**; port:=6


Global Program




```
port=1; A→B; port:=5  
+  
port=2; A→B; port:=6
```

Simple and elegant!

Virtual Program




Virtual Program



virtual "big switch"


Virtual Program



```
port=1; port:=5  
+  
port=2; port:=6
```

Even simpler!


Virtual Program



```
port=1; port:=5  
+  
port=2; port:=6
```

Even simpler!

Virtual Program



firewall

;

```
port=1; port:=5  
+  
port=2; port:=6
```

Even simpler!

Virtual Program



Can implement **multiple** arbitrary **virtual networks**
on top of **single physical network**

firewall

;

port=1; port:=5
+
port=2; port:=6

Even simpler!

Compilation [ICEP '15]

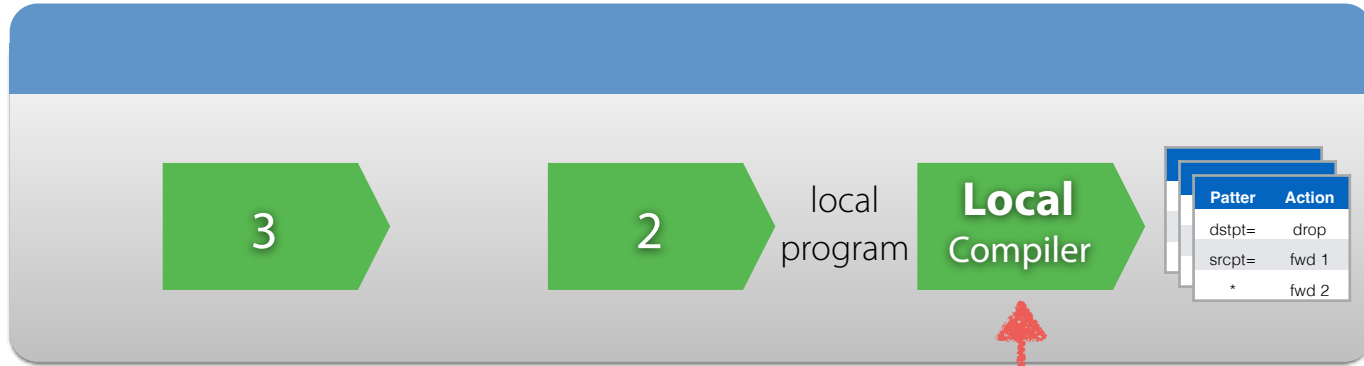


3

2

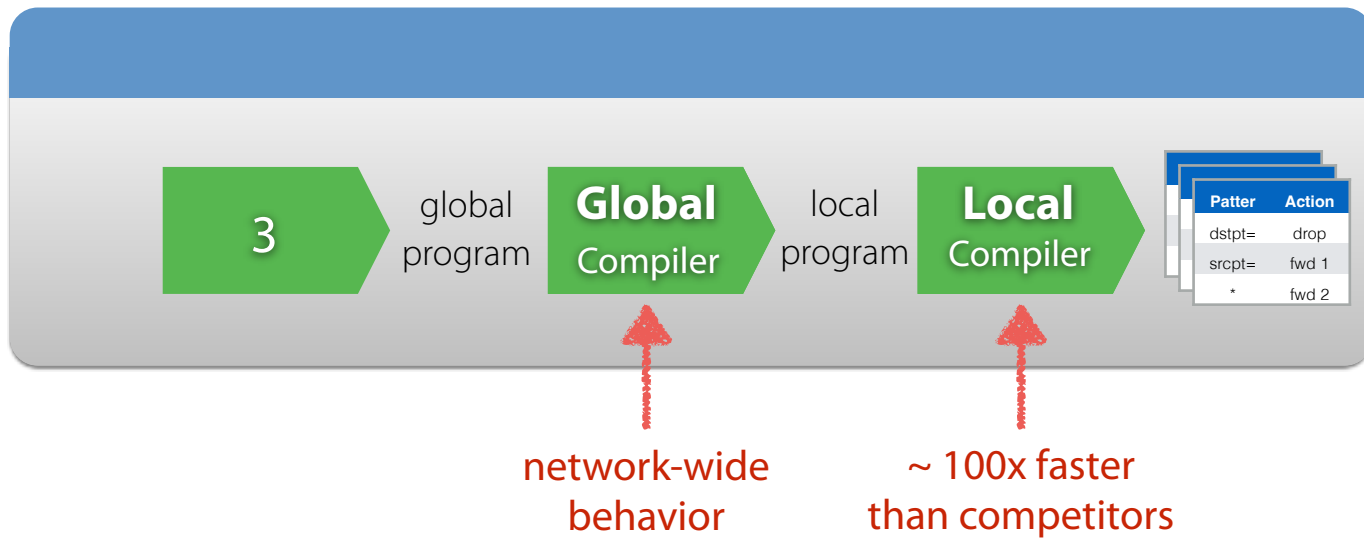
1

Compilation [ICEP '15]

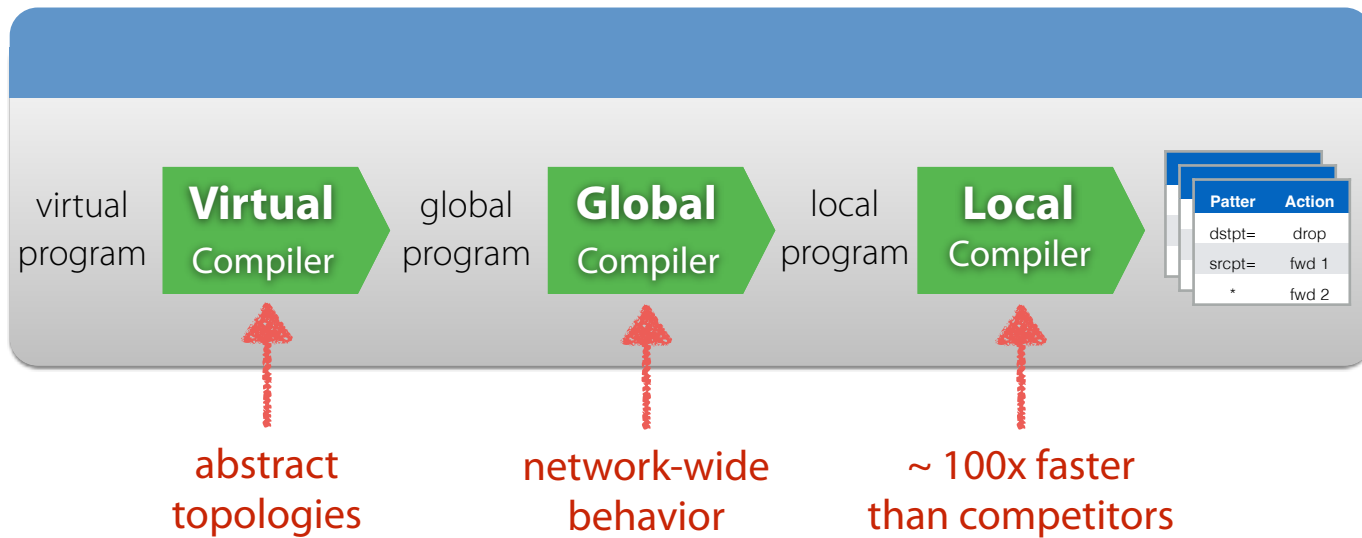


~ 100x faster
than competitors

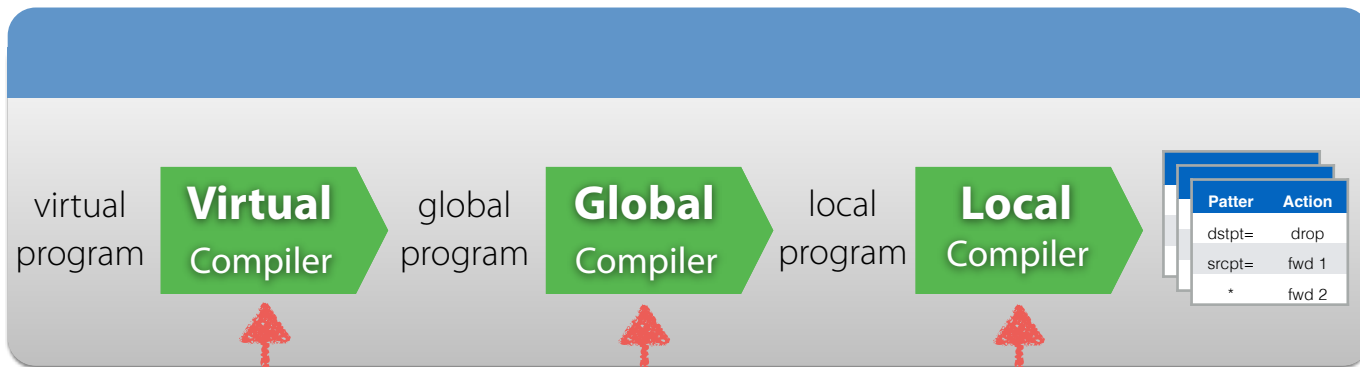
Compilation [ICEP '15]



Compilation [ICEP '15]



Compilation [ICEP '15]




abstract
topologies

network-wide
behavior

~ 100x faster
than competitors




Global Compilation




Global Compilation

1. Adding Extra State
"Tagging"




Global Compilation


1. Adding Extra State
"Tagging"



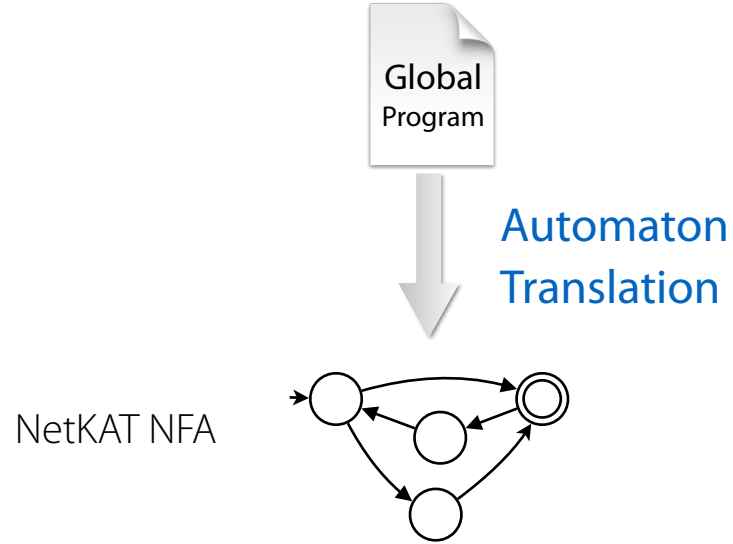
2. Avoiding Duplication
(naive tagging is unsound!)



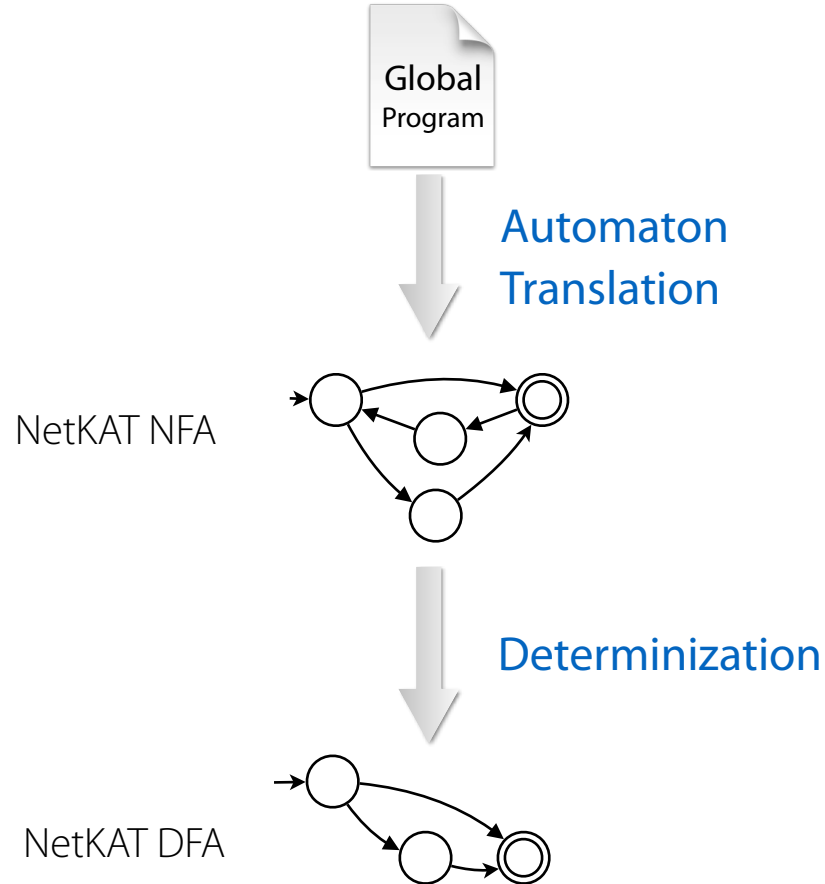
Global Compilation



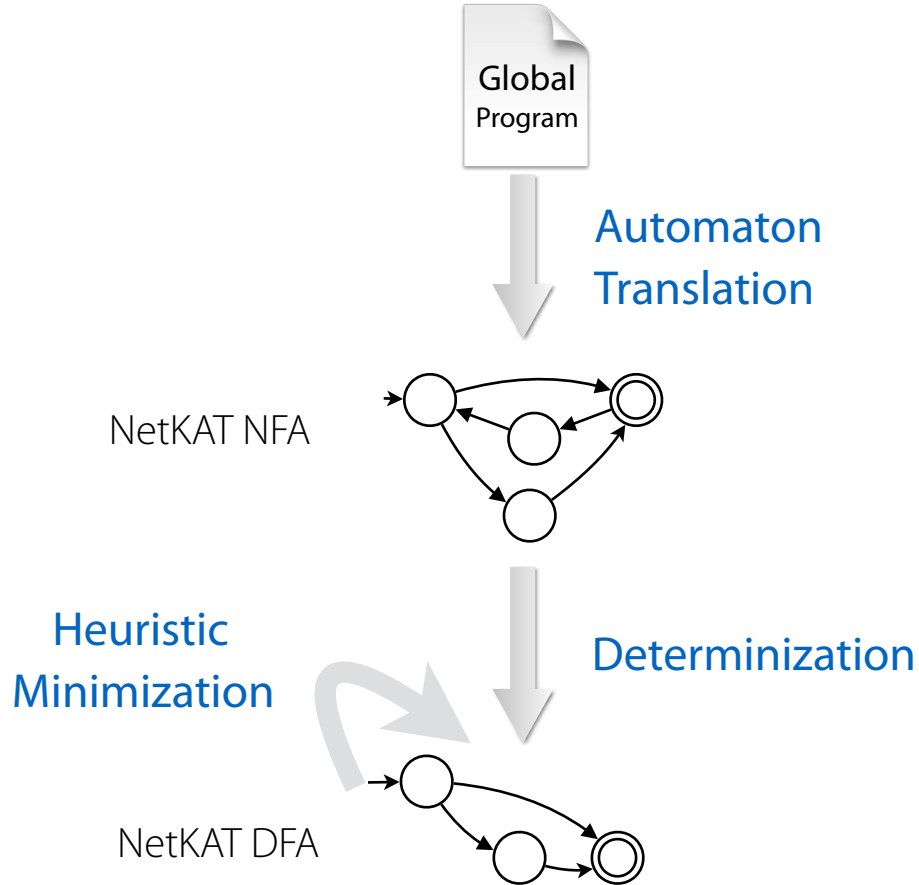
Global Compilation



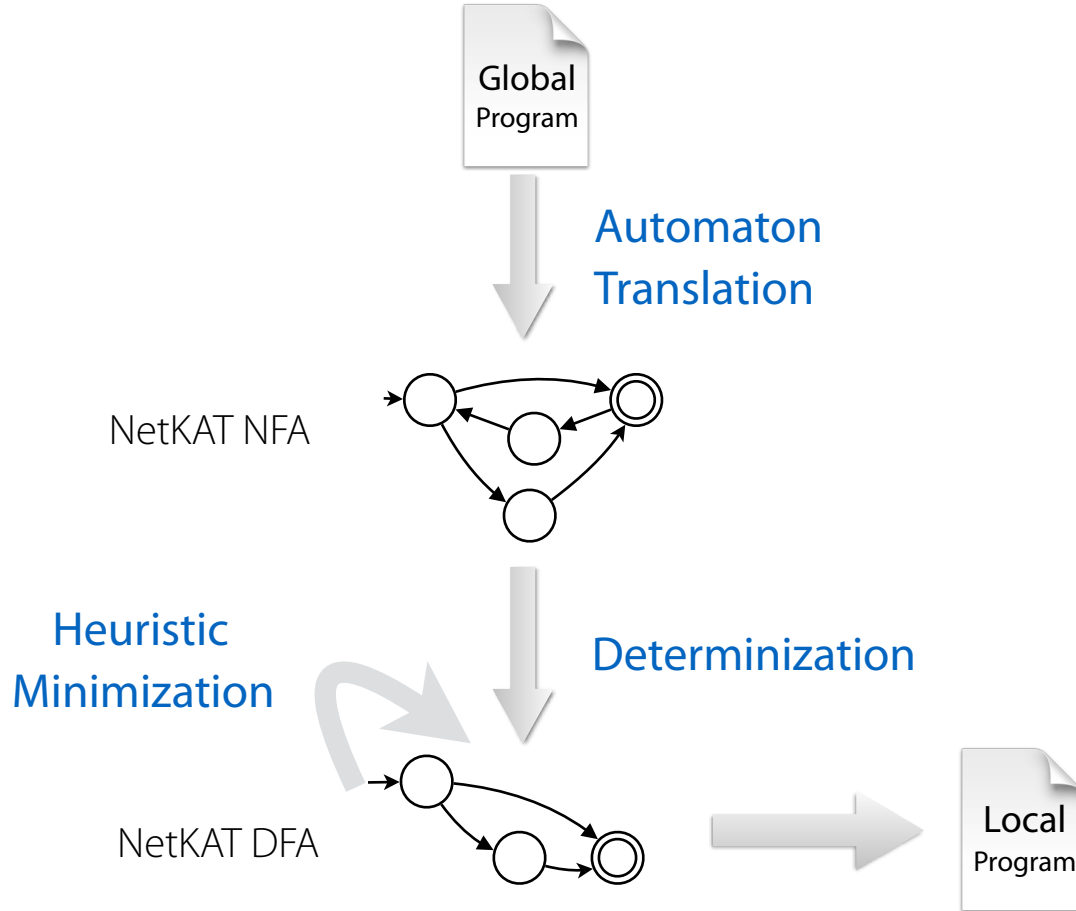
Global Compilation



Global Compilation



Global Compilation



Questions?

$p ::=$ **false**

| **true**

| $f = n$

| $f ::= n$

| $p_1 + p_2$

| $p_1 ; p_2$

| $\neg p$

| p^*

| $A \rightarrow B$





false drops its input


```
p ::= false
| true
| f = n
| f := n
| p1 + p2
| p1 ; p2
| ¬p
| p*
| A → B
```





true copies its input

```
p ::= false
| true
| f = n
| f := n
| p1 + p2
| p1 ; p2
| ¬p
| p*
| A → B
```




$f = n$ copies its input if $pk.f = n$ and otherwise drops it

```
p ::= false
| true
| f = n
| f := n
| p1 + p2
| p1 ; p2
| ¬p
| p*
| A → B
```




$f = n$ copies its input if $pk.f = n$ and otherwise drops it

$p ::=$ **false**
| **true**
| $f = n$
| $f := n$
| $p_1 + p_2$
| $p_1 ; p_2$
| $\neg p$
| p^*
| $A \rightarrow B$




$f := n$ sets the input's f component to n

```
p ::= false
| true
| f = n
| f := n
| p1 + p2
| p1 ; p2
| ¬p
| p*
| A → B
```




$p_1 + p_2$ duplicates the input, sends one copy to each sub-policy, and takes the *union* of their outputs

```
p ::= false
| true
| f = n
| f := n
| p1 + p2
| p1 ; p2
| ¬p
| p*
| A → B
```




$p_1 ; p_2$ runs the input through pol_1 and then runs every output produced by p_1 through p_2

```
p ::= false
| true
| f = n
| f := n
| p1 + p2
| p1 ; p2
| ¬p
| p*
| A → B
```





$\neg p$ drops the input if p produces any output and copies it otherwise

```
p ::= false  
| true  
| f = n  
| f := n  
| p1 + p2  
| p1 ; p2  
| ¬p  
| p*  
| A → B
```






p^* repeatedly runs packets through p to a fixpoint


```
p ::= false
| true
| f = n
| f := n
| p1 + p2
| p1 ; p2
| ¬p
| p*
| A → B
```



$A \rightarrow B$ duplicates the packet and moves it across the link

```
p ::= false
| true
| f = n
| f := n
| p1 + p2
| p1 ; p2
| ¬p
| p*
| A → B
```



$A \rightarrow B$ duplicates the packet and moves it across the link

NetKAT Semantics

NetKAT Semantics

$\llbracket p \rrbracket \in \text{History} \rightarrow \text{History Set}$

NetKAT Semantics

$\llbracket p \rrbracket \in \mathbf{History} \rightarrow \mathbf{History Set}$

$\llbracket \mathbf{true} \rrbracket h = \{ h \}$

NetKAT Semantics

$\llbracket p \rrbracket \in \mathbf{History} \rightarrow \mathbf{History Set}$

$\llbracket \mathbf{true} \rrbracket h = \{ h \}$

$\llbracket \mathbf{false} \rrbracket h = \{ \}$

NetKAT Semantics

$\llbracket p \rrbracket \in \mathbf{History} \rightarrow \mathbf{History Set}$

$\llbracket \mathbf{true} \rrbracket h = \{ h \}$

$\llbracket \mathbf{false} \rrbracket h = \{ \}$

$\llbracket f = n \rrbracket pk :: h = \begin{cases} \{ pk :: h \} & \text{if } pk.f = n \\ \{ \} & \text{otherwise} \end{cases}$

NetKAT Semantics

$\llbracket p \rrbracket \in \mathbf{History} \rightarrow \mathbf{History Set}$

$\llbracket \mathbf{true} \rrbracket h = \{ h \}$

$\llbracket \mathbf{false} \rrbracket h = \{ \}$

$\llbracket f = n \rrbracket pk :: h = \begin{cases} \{ pk :: h \} & \text{if } pk.f = n \\ \{ \} & \text{otherwise} \end{cases}$

$\llbracket \neg p \rrbracket h = \{ h \} \setminus \llbracket p \rrbracket h$

NetKAT Semantics

$\llbracket p \rrbracket \in \mathbf{History} \rightarrow \mathbf{History Set}$

$\llbracket \mathbf{true} \rrbracket h = \{ h \}$

$\llbracket \mathbf{false} \rrbracket h = \{ \}$

$\llbracket f = n \rrbracket pk :: h = \begin{cases} \{ pk :: h \} & \text{if } pk.f = n \\ \{ \} & \text{otherwise} \end{cases}$

$\llbracket \neg p \rrbracket h = \{ h \} \setminus \llbracket p \rrbracket h$

$\llbracket f := n \rrbracket pk :: h = \{ pk[f:=n] :: h \}$

NetKAT Semantics

$\llbracket p \rrbracket \in \mathbf{History} \rightarrow \mathbf{History\ Set}$

$\llbracket \mathbf{true} \rrbracket h = \{ h \}$

$\llbracket \mathbf{false} \rrbracket h = \{ \}$

$\llbracket f = n \rrbracket pk :: h = \begin{cases} \{ pk :: h \} & \text{if } pk.f = n \\ \{ \} & \text{otherwise} \end{cases}$

$\llbracket \neg p \rrbracket h = \{ h \} \setminus \llbracket p \rrbracket h$

$\llbracket f := n \rrbracket pk :: h = \{ pk[f:=n] :: h \}$

$\llbracket p_1 + p_2 \rrbracket h = \llbracket p_1 \rrbracket h \cup \llbracket p_2 \rrbracket h$

NetKAT Semantics

$\llbracket p \rrbracket \in \mathbf{History} \rightarrow \mathbf{History Set}$

$\llbracket \mathbf{true} \rrbracket h = \{ h \}$

$\llbracket \mathbf{false} \rrbracket h = \{ \}$

$\llbracket f = n \rrbracket pk :: h = \begin{cases} \{ pk :: h \} & \text{if } pk.f = n \\ \{ \} & \text{otherwise} \end{cases}$

$\llbracket \neg p \rrbracket h = \{ h \} \setminus \llbracket p \rrbracket h$

$\llbracket f := n \rrbracket pk :: h = \{ pk[f:=n] :: h \}$

$\llbracket p_1 + p_2 \rrbracket h = \llbracket p_1 \rrbracket h \cup \llbracket p_2 \rrbracket h$

$\llbracket p_1 ; p_2 \rrbracket h = (\llbracket p_1 \rrbracket \bullet \llbracket p_2 \rrbracket) h$

$f, g \in \mathbf{History} \rightarrow \mathbf{History Set}$

$(f \bullet g) h = \bigcup \{ g h' \mid h' \in f h \}$

NetKAT Semantics

$\llbracket p \rrbracket \in \mathbf{History} \rightarrow \mathbf{History Set}$

$\llbracket \mathbf{true} \rrbracket h = \{ h \}$

$\llbracket \mathbf{false} \rrbracket h = \{ \}$

$\llbracket f = n \rrbracket pk :: h = \begin{cases} \{ pk :: h \} & \text{if } pk.f = n \\ \{ \} & \text{otherwise} \end{cases}$

$\llbracket \neg p \rrbracket h = \{ h \} \setminus \llbracket p \rrbracket h$

$\llbracket f := n \rrbracket pk :: h = \{ pk[f:=n] :: h \}$

$\llbracket p_1 + p_2 \rrbracket h = \llbracket p_1 \rrbracket h \cup \llbracket p_2 \rrbracket h$

$\llbracket p_1 ; p_2 \rrbracket h = (\llbracket p_1 \rrbracket \bullet \llbracket p_2 \rrbracket) h$

$\llbracket p^* \rrbracket h = (\bigcup_i \llbracket p \rrbracket^i h)$

$f, g \in \mathbf{History} \rightarrow \mathbf{History Set}$

$(f \bullet g) h = \bigcup \{ g h' \mid h' \in f h \}$

NetKAT Semantics

$\llbracket p \rrbracket \in \mathbf{History} \rightarrow \mathbf{History\ Set}$

$\llbracket \mathbf{true} \rrbracket h = \{ h \}$

$\llbracket \mathbf{false} \rrbracket h = \{ \}$

$\llbracket f = n \rrbracket pk :: h = \begin{cases} \{ pk :: h \} & \text{if } pk.f = n \\ \{ \} & \text{otherwise} \end{cases}$

$\llbracket \neg p \rrbracket h = \{ h \} \setminus \llbracket p \rrbracket h$

$\llbracket f := n \rrbracket pk :: h = \{ pk[f:=n] :: h \}$

$\llbracket p_1 + p_2 \rrbracket h = \llbracket p_1 \rrbracket h \cup \llbracket p_2 \rrbracket h$

$\llbracket p_1 ; p_2 \rrbracket h = (\llbracket p_1 \rrbracket \bullet \llbracket p_2 \rrbracket) h$

$\llbracket p^* \rrbracket h = (\bigcup_i \llbracket p \rrbracket^i h)$

$\llbracket A \rightarrow B \rrbracket pk :: h = \begin{cases} \{ pk[sw:=B] :: pk :: h \} & \text{if } pk.sw = A \\ \{ \} & \text{otherwise} \end{cases}$

$f, g \in \mathbf{History} \rightarrow \mathbf{History\ Set}$

$(f \bullet g) h = \bigcup \{ g h' \mid h' \in f h \}$