Efficient Programming Abstractions for SDN



Steffen Smolka

Networks are becoming programmable



Networks are becoming programmable



Language Design

Language Design







Rich Packet Classification Network-wide Abstractions

Modular Composition NetKAT

Model

Packets are records of values. Programs are functions on packets.



{ **switch** = A, **port** = 3, ethSrc = 8:8::::8:8, ethDst = 2:2:::2:2, vLan = 8, ipSrc = 192.168.2.1, ipDst = 127.0.0.1,

pol ::= false true field = val field **:=** val $pol_1 + pol_2$ $pol_1; pol_2$!pol pol* | S**→**S'



Boolean Algebra



Boolean Algebra + Kleene Algebra "Regular Expressions"

pol ::= false true field = val field := val $pol_1 + pol_2$ $pol_1; pol_2$!pol pol* S→S'

Boolean Algebra + Kleene Algebra "Regular Expressions" + Packet Primitives

pol ::= | false

Boolean Algebra

if p then q else r \triangleq p;q + !p;r

while $p do q \triangleq p;q*;!p$

| pol₁; pol₂ | !pol | pol* | S→S'

Packet Primitives

NetKAT Semantics



NetKAT Semantics



Local NetKAT: input-output behavior of switches



[pol] ∈ Packet → Packet Set

NetKAT Semantics



Local NetKAT: input-output behavior of switches



[pol] ∈ Packet → Packet Set

Global NetKAT: network-wide behavior



[pol] ∈ History → History Set

Example











Tedious for programmers... difficult to get right!

Global Program



Global Program



Simple and elegant!





virtual "big switch"



virtual "big switch"

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



virtual "big switch"

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





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



Compilation



Compilation



NetKAT Compiler



Local Compilation



Input: local program

Output: collection of flow tables, one per switch

Challenges: efficiency and size of generated tables

let route =

if ipDst = 10.0.0.1 then
 port := 1
else if ipDst = 10.0.0.2 then
 port := 2
else
 port := learn



let monitor =

if (tcpSrc = 22 + tcpDst = 22) then
 port:=console
else
 false

let route =

```
if ipDst = 10.0.0.1 then
    port := 1
else if ipDst = 10.0.0.2 then
    port := 2
else
    port := learn
```



let monitor =

if (tcpSrc = 22 + tcpDst = 22) then
 port:=console
else
 false

Pattern	Actions
src=10.0.0.1	Fwd 1
src=10.0.0.2	Fwd 2
*	Controller

Pattern	Actions
tcpSrc=22	Controller
tcpDst=22	Controller
*	Drop

let route =

```
if ipDst = 10.0.0.1 then
   port := 1
else if ipDst = 10.0.0.2 then
   port := 2
else
   port := learn
```



let monitor =

if (tcpSrc = 22 + tcpDst = 22) then
 port:=console
else
 false





Pattern	Actions
tcpSrc=22	Controller
tcpDst=22	Controller
*	Drop



Inefficient!

Tables are a hardware abstraction, not an efficient data structure!!



if ipDst = 10.0.0.1 then
 port := 1
else if ipDst = 10.0.0.2 then
 port := 2
else
 port := learn

let monitor =

if (tcpSrc = 22 + tcpDst = 22) then
 port:=console
else

false

÷

let route =

if ipDst = 10.0.0.1 then
 port := 1
else if ipDst = 10.0.0.2 then
 port := 2
else
 port := learn

let monitor =

if (tcpSrc = 22 + tcpDst = 22) then
 port:=console
else

false









let route =

if ipDst = 10.0.0.1 then
 port := 1
else if ipDst = 10.0.0.2 then
 port := 2
else

let monitor =

```
if (tcpSrc = 22 + tcpDst = 22) then
    port:=console
else
    false
```

Key Data Structure:

42

Forwarding Decision Diagram

→ now widely adopted



Global Compilation



Input: NetKAT program (with links)

Output: equivalent local program (without links)

Main Challenges



1. Adding Extra State "Tagging"

2. Avoiding Duplication (naive tagging is unsound!)



Global Program

Adding Extra State = Translation to Automaton





Global Program

Adding Extra State = Translation to Automaton

NetKAT NFA



Avoiding Duplication = Determinization



NetKAT DFA

Global Program

Adding Extra State = Translation to Automaton



NetKAT NFA

Automaton Minimization = Tag Elimination Avoiding Duplication = Determinization

NetKAT DFA

Global Program

Adding Extra State = Translation to Automaton



Automaton Minimization = Tag Elimination

Avoiding Duplication = Determinization



NetKAT DFA

NetKAT NFA

NetKAT Automata [Foster et al, POPL '15]

Transition relation $\delta : Q \rightarrow Packet \rightarrow P(Q \times Packet)$

"Alphabet size": Packet x Packet



NetKAT Automata [Foster et al, POPL '15]

Transition relation $\delta : Q \rightarrow Packet \rightarrow P(Q \times Packet)$

"Alphabet size": Packet x Packet



Can represent δ symbolically using FDDs!

Automata construction: Antimirov partial derivatives & Position Automata

Virtual Compilation



Input: program written against virtual topology

Output: global program that simulates virtual behavior









Observation: can formulate execution of a virtual program as a two-player game

Compiler: synthesizes physical program p that encodes a winning strategy to all instances of that game

Evaluation

Local Compiler vs State of the Art



about 100x speedup













Conclusion



Fast, Flexible, and Fully implemented in OCaml: <u>http://github.com/frenetic-lang/frenetic/</u>

Go ahead and use it! (others are using it already)





