Adriaan van Wijngaarden meets Scott Domain-Theoretic Foundations for Probabilistic Network Programming

Alexandra Silva (UCL)





The gang











Networks



Network configuration largely a black art



built and programmed the same way

- Difficult to implement end-to-end routing policies and optimisations that require a global perspective
- Difficult to extend with new functionality
- Effectively impossible to reason precisely about behaviour

interfaces

Software-Defined Networks



Openflow

[McKeown & al., SIGCOMM 08]

- Specifies capabilities and behaviour of switch hardware
- A language for manipulating network configurations
- Very low-level: easy for hardware to implement, difficult for humans to write and reason about

But...

- ✓ is platform independent
- ✓ provides an open standard that any vendor can implement

Verification of networks

Trend in PL&Verification after Software-Defined Networks

- Design high-level languages that model essential network features
- Develop *semantics* that enables reasoning precisely about behaviour
- Build tools to synthesise low-level implementations automatically

- Frenetic [Foster & al., ICFP 11]
- Pyretic [Monsanto & al., NSDI 13]
- Maple [Voellmy & al., SIGCOMM 13]
- FlowLog [Nelson & al., NSDI 14]
- Header Space Analysis [Kazemian & al., NSDI 12]
- VeriFlow [Khurshid & al., NSDI 13]
- NetKAT [Anderson & al., POPL 14]
- and many others . . .

NetKAT

Kleene algebra with tests (KAT) + additional specialized constructs particular to

network topology and packet switching



Stephen Cole Kleene (1909–1994)

$(0 + 1(01^*0)^*1)^*$ {multiples of 3 in binary}



$$(ab)^*a = a(ba)^*$$

 $\{a, aba, ababa, \ldots\}$
 $\rightarrow O_{b}^a$

$$(a+b)^* = a^*(ba^*)^*$$

{all strings over $\{a, b\}$ } $\rightarrow \bigcirc a + b$

 $(K, B, +, \cdot, *, -, 0, 1), \quad B \subseteq K$

- $(K, +, \cdot, *, 0, 1)$ is a Kleene algebra
- $(B, +, \cdot, -, 0, 1)$ is a Boolean algebra
- $(B, +, \cdot, KAT = simple imperative language$
- **p**, **q**, **r**, . **If** b **then** p **else** q = b;p + !b;q

While b do $p = (bp)^*!b$

- a packet π is an assignment of constant values *n* to fields *x*
- a packet history is a nonempty sequence of packets $\pi_1 :: \pi_2 :: \cdots :: \pi_k$
- the head packet is π_1

NetKAT

- ► assignments x ← n assign constant value n to field x in the head packet
- tests x = n if value of field x in the head packet is n, then pass, else drop
- dup duplicate the head packet

Networks in NetKAT

sw=6;pt=8;dst := 10.0.1.5;pt:=5

For all packets located at port 8 of switch 6, set the destination address to 10.0.1.5 and forward it out on port 5.

Networks in NetKAT

The behaviour of an entire network can be encoded in NetKAT by interleaving steps of processions by switches and topology



Semantics



$$\llbracket e \rrbracket : H \to 2^{H}$$
$$\llbracket x \leftarrow n \rrbracket (\pi_{1} :: \sigma) \triangleq \{\pi_{1}[n/x] :: \sigma\}$$
$$\llbracket x = n \rrbracket (\pi_{1} :: \sigma) \triangleq \begin{cases} \{\pi_{1} :: \sigma\} & \text{if } \pi_{1}(x) = n \\ \varnothing & \text{if } \pi_{1}(x) \neq n \end{cases}$$
$$\llbracket dup \rrbracket (\pi_{1} :: \sigma) \triangleq \{\pi_{1} :: \pi_{1} :: \sigma\}$$

Verification using NetKAT

Reachability

Can host A communicate with host B? Can every host communicate with every other host?

Security

Does all untrusted traffic pass through the intrusion detection system located at C?

Loop detection

Is it possible for a packet to be forwarded around a cycle in the network?

Verification using NetKAT

Soundness and Completeness [Anderson et al. 14]

▶ $\vdash p = q$ if and only if $\llbracket p \rrbracket = \llbracket q \rrbracket$

Decision Procedure [Foster et al. 15]

- NetKAT coalgebra
- efficient bisimulation-based decision procedure
- implementation in OCaml
- deployed in the Frenetic suite of network management tools

Limitations

$\llbracket e \rrbracket \colon H \to 2^H$

*Packet-processing **function**

*Applicability limited to simple connectivity or routing behavior

Probabilities are needed

* expected congestion
* reliability
* randomized routing

ProbNetKAT

$p\oplus_r q$



$\mathsf{dst} = \mathsf{h}_3; \mathsf{pt} \leftarrow 2 \oplus_{.5} \mathsf{pt} \leftarrow 4$

ProbNetKAT by example



$$net \triangleq in; (p;t)^*; p; out$$

$$\begin{array}{l} l_{1} & \triangleq (\mathsf{sw} = S_1; \mathsf{pt} = 2; \mathsf{dup}; \mathsf{sw} \leftarrow S_2; \mathsf{pt} \leftarrow 1; \mathsf{dup}) \\ \texttt{Ingress} = \mathsf{gress}; \mathsf{pt} \leftarrow 1; \\ \texttt{dst} = h_2^2; \mathsf{pt} \leftarrow 1; \mathsf{dup}; \mathsf{sw} \leftarrow S_1; \mathsf{pt} \leftarrow 2; \mathsf{dup}) \\ & \& (\mathsf{dst} = h_2^2; \mathsf{pt} \leftarrow 1; \mathsf{dup}; \mathsf{sw} \leftarrow S_1; \mathsf{pt} \leftarrow 2; \mathsf{dup}) \\ & i \& (\mathsf{dst} = h_2^2; \mathsf{pt} \leftarrow 1; \mathsf{dup}; \mathsf{sw} \leftarrow S_1; \mathsf{pt} \leftarrow 2; \mathsf{dup}) \\ & i \& (\mathsf{dst} = h_2^2; \mathsf{pt} \leftarrow 1; \mathsf{pt} \leftarrow 1; \mathsf{pt} \leftarrow 2; \mathsf{pt} = 2) \& \dots \\ & i \& (\mathsf{dst} = h_2^2; \mathsf{pt} \leftarrow 1; \mathsf{pt} \leftarrow 1; \mathsf{pt} \leftarrow 2; \mathsf{pt} = 2) \& \dots \\ & o u \& (\mathsf{dst} = \mathsf{st} \cdot \mathsf{swh}_4; \mathsf{lpt} \cdot \mathsf{pt} \leftrightarrow 1) \& (\mathsf{sw} = 2; \mathsf{pt} = 2) \& \dots \end{array}$$

Forwarding policy

$$p \triangleq (\mathsf{sw} = S_1; p_1) \& (\mathsf{sw} = S_2; p_2) \& (\mathsf{sw} = S_3; p_3) \& (\mathsf{sw} = S_4; p_4)$$

Topology

 $t \triangleq l_{1,2} \& l_{2,3} \& l_{3,4} \& l_{1,4}$

Semantics

 $\llbracket p \rrbracket \in 2^{\mathsf{H}} \to \{\mu : \mathcal{B} \to [0,1] \mid \mu \text{ is a probability measure} \}$

 ${\cal B}$ Borel sets of $2^{\rm H}$ using Cantor topology

$$\begin{bmatrix} x \leftarrow n \end{bmatrix}(a) = \delta_{\{\pi[n/x]:\sigma \mid \pi:\sigma \in a\}} \\ \begin{bmatrix} x = n \end{bmatrix}(a) = \delta_{\{\pi:\sigma \mid \pi:\sigma \in a, \pi(x) = n\}} \\ \begin{bmatrix} dup \end{bmatrix}(a) = \delta_{\{\pi:\pi:\sigma \mid \pi:\sigma \in a\}} \\ \\ \begin{bmatrix} skip \end{bmatrix}(a) = \delta_a \\ \\ \begin{bmatrix} drop \end{bmatrix}(a) = \delta_{\varnothing} \end{bmatrix}$$

 $\llbracket p \& q \rrbracket(a) = \llbracket p \rrbracket(a) \& \llbracket q \rrbracket(a)$ $(\mu \& \nu)(A) \triangleq (\mu \times \nu)(\{(a, b) \mid a \cup b \in A\}).$

$$[\![p+_r q]\!](a) = r[\![p]\!](a) + (1-r)[\![p]\!](a)$$

Semantics

 $\llbracket p \rrbracket \in 2^{\mathsf{H}} \to \{ \mu : \mathcal{B} \to [0,1] \mid \mu \text{ is a probability measure} \}$

$$[\![p^*]\!] = ?$$

Ideally: $[\![p^*]\!] = [\![1\&pp^*]\!]$

least fix point? which order?

Ad-hoc attempt: infinite stochastic process



Congestion Query: Random Variable $\mathbf{Q}: 2^{H} \rightarrow [0,\infty]$

$$Q(a) \triangleq \sum_{h \in a} \ \#_l(h)$$

Expected Congestion: $E_v[Q]$

$$\mathbf{E}_{\nu}[Q] = \int Q \, d\nu$$

Issues with previous semantics



The importance of continuity



Perform computation f on a f continuous

$$f(a_1), f(a_2), \cdots$$
 limit $f(a)$

The importance of continuity for network analysis



 $\mathbf{E}_{\mu}[f]$ — expected value of a continuous map is continuous

monotonically improving sequence of approximations for performance metrics such as latency and congestion

New semantics

 $\llbracket p \rrbracket \in 2^{\mathsf{H}} \to \{ \mu : \mathcal{B} \to [0, 1] \mid \mu \text{ is a probability measure} \}$ $\mathcal{B} \text{ Borel sets of}$ using Scott topology

$\llbracket p^* \rrbracket = \mathsf{lfp} \ X \mapsto 1 \ \& \ \llbracket p \rrbracket; X$

Finite approximations -----> Practical implementation

Implementation and Case studies

Interpreter in OCaml

Approximates the answer monotonically

Several case studies



Internet2's Abilene backbone network

Conclusions

First language-based framework for specifying and verifying **probabilistic network behavior**.



Questions?





roperties



(c) Max congestion

(d) Throughput

Values converge monotonically



roperties





Routing

Equal Cost Multipath Routing (ECMP)

k-Shortest Paths (KSP)

Multipath Routing (Multi)

Oblivious Routing (Raecke)