Probabilistic Network Programming

Nate Foster Cornell



Plan

I: Introduction

- Semantics Primer
- Software-Defined Networking

II: NetKAT

- Language Design
- Formal Semantics

III: Probabilistic NetKAT

- Formal Semantics
- Applications

Quiz

What does this program do?

n := 0
while !(x = 0)
do
n := n + 1;
(x := 0 ⊕ x := 1);

Quiz

What does this program do?

n := 0 while !(x = 0)do n := n + 1; $(x := 0 \oplus x := 1);$ Random Choice

Randomized Routing

Topology



Assume each node can send and receive traffic at rate ${\bf r}$

Demand Matrix



The D_{ij} represents the demand from S_i to S_j

Routing Scheme



Assigns the demand to paths while respecting link capacities

Demand Matrix



Restrict attention to *feasible* demand matrices

- Send: **Σ**_i **D**_{ji} < **r**
- Receive: Σ_j D_{ij} < r

Demand Matrix



Restrict attention to *feasible* demand matrices

- Send: **Σ**_i **D**_{ji} < **r**
- Receive: Σ_j D_{ij} < r

Theorem [Keslassy '05]. Mesh with capacities **2r/N** is optimal design that routes all feasible demand matrices













- 1. To a *random* intermediary
- 2. To destination



- 1. To a *random* intermediary
- 2. To destination



- 1. To a *random* intermediary
- 2. To destination



Theorem [Valiant '82]. Two-stage routing can handle all feasible demand matrices.

Intuition: the randomization in the first stage balances the offered load across all nodes (with high probability)

- 1. To a *random* intermediary
- 2. To destination

Oblivious Routing

•	Robust and efficient traffic eng ×													
\leftarrow		ecure https	://researc	h.fb.com/robust-a	nd-efficient-traffic-	engineeri	ng-with	Q 🕁	1	AP				
								Search		Q				
	Research Areas	Publications	People	Academic Programs	Downloads & Projects	Careers	Blog							

Robust and efficient traffic engineering with oblivious routing

By: Chiun Lin Lim, Petr Lapukhov



Related

Publication

The Effect of Computer-Generated Descriptions on Photo-Sharing Experiences of People with Visual Impairments

Yuhang Zhao, Shaomei Wu, Lindsay Reynolds, Shiri Azenkot

CSCW 2018

November 3, 2018

Publication

The effects of natural scene statistics on text readability in additive displaye

Path Selection



The choice of paths can have a major impact on the congestion induced in wide-area networks

Oblivious Routing

Can generalize VLB to arbitrary topologies by randomly routing traffic using a set of well-chosen tree-structured overlays



Path Budget



Latency



Probablistic NetKAT

A language for modeling & reasoning about networks probabilistically.

Prob + NetKAT

probabilistic primitive p⊕_rq network primitives

f:=n, dup

A language for modeling & reasoning about networks probabilistically.

ProbNetKA	Г				
•		NetKAT			2013
		•		KAT	1996
				1956	1847
Prob	+	Net	+	KA	+ T
probabilistic primitives		network primitives		regular expressions	boolean 5 tests
$p \oplus_r q$		f:=n, dup		+, ·, *	f=n

A language for modeling & reasoning about networks probabilistically.

 $\llbracket p \rrbracket \in 2^{H} \rightarrow \text{Dist}(2^{H})$



Probabilistic Semantics





ProbNetKAT model **p**, input distribution **µ**



ProbNetKAT model **p**, input distribution **µ**

→ output distribution $\mathbf{v} = \llbracket p \rrbracket(\mu) \in \text{Dist}(2^{H})$



ProbNetKAT model **p**, input distribution **µ**

→ output distribution $\mathbf{v} = \llbracket p \rrbracket(\mu) \in \text{Dist}(2^{H})$



ProbNetKAT model **p**, input distribution **µ**

→ output distribution $\mathbf{v} = \llbracket p \rrbracket(\mu) \in \text{Dist}(2^{H})$

utilization query: $\mathbf{Q}: 2^{H} \rightarrow [0,\infty]$

expected utilization: $E_{v}[Q]$

How to implement this?



How to implement this?



How to implement this?


$((f:=0 \oplus f:=1) \bullet dup)^*$

 $((f:=0 \oplus f:=1) \bullet dup)^*$



execution \cong infinite path \cong random output $\in 2^{H}$

 $((f:=0 \oplus f:=1) \bullet dup)^*$



How many paths are there? \rightarrow one for every r \in [0,1]

 $((f:=0 \oplus f:=1) \bullet dup)^*$



How many paths are there? \rightarrow one for every r \in [0,1]

What's the probability of any particular path? $\rightarrow 0$

Key Idea

limits + continuity \rightarrow approximation

Key Idea

limits + continuity → approximation



Key Idea

limits + **continuity** → approximation



1) Iteration-free programs generate only finite distributions

- 1) Iteration-free programs generate only finite distributions
- 2) Iteration may introduce continuous distributions but can be approximated by bounded iteration

- 1) Iteration-free programs generate only finite distributions
- 2) Iteration may introduce continuous distributions but can be approximated by bounded iteration

3) All programs can be approximated

- 1) Iteration-free programs generate only finite distributions
- Iteration may introduce continuous distributions ...
 but can be approximated by bounded iteration



3) All programs can be approximated

- 1) Iteration-free programs generate only finite distributions
- 2) Iteration may introduce continuous distributions but can be approximated by bounded iteration



3) All programs can be approximated

continuity of ∫, E[–]

4) Queries can be approximated











Quantitative Verification

Decidable Semantics

Restricted to the local fragment, a NetKAT program can be seen as a **random** function on packets:



Decidable Semantics

Restricted to the local fragment, a NetKAT program can be seen as a **random** function on packets:



Such functions can be modeled as Markov chains:

- states are given by sets of packet: $S = 2^{Pk}$
- Chain is given by transition matrix $\mathbf{B} \in [0,1]^{S \times S}$
- B(a, b): probability of producing output b on input a
- state space large, but finite!

Example: Sequential Composition

ProbNetKAT primitives can be mapped to matrices and program operators map to matrix operations.



Similarly for $p \& q, p \oplus_{\mathbf{r}} q$.

Problem: Kleene Star

Question

Let $\pi_1!$ be the program that generates packet π_1 Consider $p = (true \oplus \pi_1!)$. What is **B**[[p*]]({ π_0 }, { π_0 })? **Answer**

 $\mathbf{B}[\![p^{(n)}]\!](\{\pi_0\}, \{\pi_0\}) = (1/_2)^n$

Thus,

Β[[p^{*}]]({
$$\pi_0$$
}, { π_0 })
m Β[[p⁽ⁿ⁾]]({ π_0 } { π_0 }

ע**דו** (כ) א

$$= \lim_{n \to \infty} \mathbf{B}[[p^{(n)}]](\{\pi_0\}, \{\pi_0\})$$

$$= \lim_{n \to \infty} (1/_2)^n = 0$$

But how to compute these limits in general?

Small Step Semantics

1 step in $S[p] \approx 1$ iteration of p^*

In one iteration, p*:

- executes p to get
 new set of packets
- emits previous set of packets

 $B[[p^*]] := \lim S[[p]]^n$



Small Step Semantics

1 step in $S[p] \approx 1$ iteration of p^*



Absorbing Markov Chains

S[[p]] can be "massaged" into an **absorbing Markov chain**

Absorbing state: 1



Absorbing chain: any state can reach absorbing state

Crucial property: for #steps $\rightarrow \infty$, will reach absorbing state with probability 1 (no matter the start state)

$$T = \begin{bmatrix} I & 0 \\ R & Q \end{bmatrix}$$
$$\lim_{n \to \infty} T^n = \begin{bmatrix} I & 0 \\ (I - Q)^{-1} R & 0 \end{bmatrix}$$

Absorbing Markov Chains

S[[p]] can be "massaged" into an **absorbing Markov chain**

Absorbing stateAbsorbingSo $B[[p^*]] = \lim_{n \to \infty} S[[p]]^n can$
be computed explicitly!Crucial probe computed explicitly!state with proc

$$\int = \begin{bmatrix} I & 0 \\ R & Q \end{bmatrix}$$
$$\lim_{n \to \infty} T^n = \begin{bmatrix} I & 0 \\ (I - Q)^{-1} R & 0 \end{bmatrix}$$

Evaluation

Scalability

Because semantics is formulated in terms of (sparse) matrices, can engineer an efficient implementation



Case Study: Topology

An ABFatTree is much like a regular FatTree



But it provides shorter detours around failures

Case Study: Topology

An ABFatTree is much like a regular FatTree



But it provides shorter detours around failures

Case Study: Topology

An ABFatTree is much like a regular FatTree



But it provides shorter detours around failures











Case Study: k-Resilience

We verified k-resilience using ProbNetKAT

Case Study: k-Resilience

We verified k-resilience using ProbNetKAT

Sophistication of Routing Scheme			
k	F10 ₀	F10 ₃	F10 _{3,5}
0	\checkmark	\checkmark	X
1	×	\checkmark	×
2	×	\checkmark	×
3	×	×	×
4	×	×	×
∞	X	X	X

k = number of failures $\checkmark = 100\%$ packet delivery

Case Study: k-Resilience

We verified k-resilience using ProbNetKAT



k = number of failures $\checkmark = 100\%$ packet delivery
Case Study: k-Resilience

We verified k-resilience using ProbNetKAT



k = number of failures $\checkmark = 100\%$ packet delivery

Case Study: k-resilience

After fixing the bug...



k = number of failures $\checkmark = 100\%$ packet delivery

Case Study: Results

Resilience

Latency



Wrapping Up...

Conclusion

- Randomized algorithms are a powerful tool for solving problems in many domains, including networking
- Programming languages based on probabilistic semantics are needed to express and reason about these algorithms
- Can build practical tools that analyze quantitative network properties such as peak congestion, resilience to failures, and latency automatically

Reading

 Nate Foster, Dexter Kozen, Konstantinos Mamouras, Mark Reitblatt, and Alexandra Silva. **Probabilistic NetKAT**. In *European Symposium on Programming (ESOP),* April 2016.

Steffen Smolka, Praveen Kumar, Nate Foster, Dexter Kozen, and Alexandra Silva. **Cantor Meets Scott: Semantic Foundations for Probabilistic Networks**. In ACM SIGPLAN—SIGACT Symposium on Principles of Programming Languages (POPL), Paris, France, January 2017.

Acknowlegments

- Dexter Kozen (Cornell)
- David Kahn (CMU)
- Konstantinos Mamouras (Rice)
- Mark Reitblatt (Facebook)
- Alexandra Silva (UCL)
- Steffen Smolka (Cornell)
- Justin Hsu (Wisconsin)