


# Probabilistic NetKAT

Nate Foster, Dexter Kozen (Cornell U), Konstantinos Mamouras (Penn), Mark Reitblatt (Facebook), **Alexandra Silva (UCL)**



# Context

Formal specification and verification of networks have recently become a reality

- ❖ 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 . . .

# Context

Formal specification and verification of networks have recently become a reality

## Trend in PL&Verification after Software-Defined Networks

- Design *high-level languages* that model essential network features
  - Develop *semantics* that enables reasoning precisely about behavior
  - Build *tools* to synthesize low-level implementations automatically
- ❖ NetKAT [Anderson & al., POPL 14]
  - ❖ and many others . . .

# Probabilistic NetKAT in a nutshell

[Foster & al., POPL15][Smolka & al., ICFP15][Anderson & al., POPL14]

- \* A probabilistic extension of NetKAT, a programming language/ logic for specification/verification/programming of packet switching networks
- \* Programs denote functions that give **probability distributions** on sets of packet histories
- \* Enables reasoning about **probabilistic routing protocols** or behavior of deterministic protocols on **random inputs**
- \* Can handle scenarios involving **congestion**, **failure**, and **randomized routing**

# ProbNetKAT language

```
pol ::= drop
      | skip
      | field = val
      | pol1 & pol2
      | pol1 ; pol2
      | !pol
      | pol +r pol
      | pol*
      | field := val
      | dup
```

# ProbNetKAT language

## Boolean Algebra

```
pol ::= drop
      | skip
      | field = val
      | pol1 & pol2
      | pol1 ; pol2
      | !pol
      | pol +r pol
      | pol*
      | field := val
      | dup
```

# ProbNetKAT language

pol ::= drop

| skip

| field = val

| pol1 & pol2

| pol1 ; pol2

| !pol

| pol +<sub>r</sub> pol

| pol\*

| field := val

| dup

**Boolean Algebra**

+

**Kleene Algebra**

# ProbNetKAT language

```
pol ::= drop
      | skip
      | field = val
      | pol1 & pol2
      | pol1 ; pol2
      | !pol
      | pol +r pol
      | pol*
      | field := val
      | dup
```

**Boolean Algebra**

+

**Kleene Algebra**

+

**Packet Primitives**



# ProbNetKAT language

pol ::= drop

| skip

| field = val

| pol1 & pol2

| pol1 ; pol2

| !pol

| pol +<sub>r</sub> pol

| pol\*

| field := val

| dup

**Boolean Algebra**

+

**Kleene Algebra**

+

**Packet Primitives**

+

**Probabilistic choice**

# ProbNetKAT language

```
pol ::= drop
      | skip
      | field = val
      | pol1 & pol2
      | pol1 ; pol2
      | !pol
      | pol +r pol
      | pol*
      | field := val
      | dup
```

**Boolean Algebra**

+

**Kleene Algebra**

+

**Packet Primitives**

+

**Probabilistic choice**

} KAT (Kozen'96)

# ProbNetKAT language

```
pol ::= drop
      | skip
      | field = val
      | pol1 & pol2
      | pol1 ; pol2
      | !pol
      | pol +r pol
      | pol*
      | field := val
      | dup
```

**Boolean Algebra**

+

**Kleene Algebra**

+

**Packet Primitives**

+

**Probabilistic choice**

# ProbNetKAT language

```
pol ::= drop
      | skip
      | field = val
      | pol1 & pol2
      | pol1 ; pol2
      | !pol
      | pol +r pol
      | pol*
      | field := val
      | dup
```

**Boolean Algebra**

+

**Kleene Algebra**

+

**Packet Primitives**

+

**Probabilistic choice**

NetKAT  
(Anderson et al'14)

# ProbNetKAT language

pol ::= drop

| skip

| field = val

| pol1 & pol2

| pol1 ; pol2

| !pol

| pol +<sub>r</sub> pol

| pol\*

| field := val

| dup

**Boolean Algebra**

+

KAT = simple imperative language

**If** b **then** p **else** q = b;p + !b;q

**While** b **do** p = (bp)\*!b

**Probabilistic choice**

NetKAT  
Anderson et al'14)


# Networks in NetKAT

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

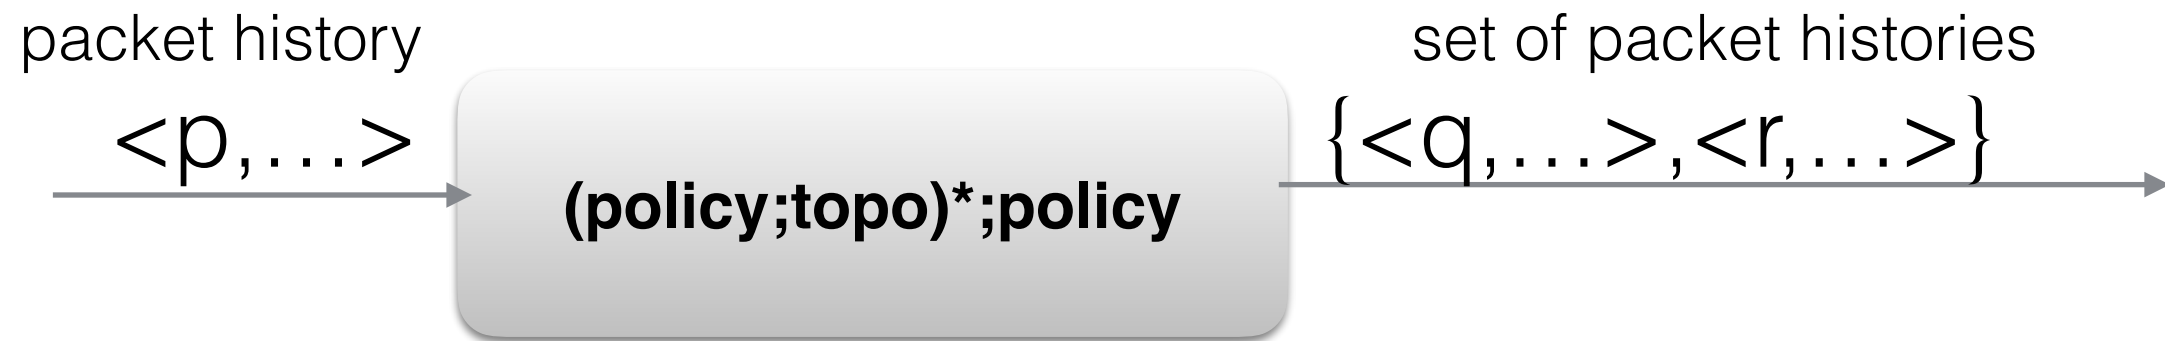
*For all packets located at port 8 of switch 6, set the destination address to 10.0.1.5 and forward it out on port 5.*

# Networks in NetKAT

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


$$\begin{aligned} & \text{policy} \\ & + \\ & (\text{policy}; \text{topo}; \text{policy}) \\ & + \\ & (\text{policy}; \text{topo}; \text{policy}; \text{topo}); \text{policy} \\ & \vdots \\ & (\text{policy}; \text{topo})^*; \text{policy} \end{aligned}$$

# Semantics




$$\llbracket e \rrbracket : H \rightarrow 2^H$$

\*Packet-processing **function**

\*Applicability limited to simple connectivity or routing behavior




# Example



Configure the switches to forward traffic on the two left-to-right paths from  $I_1$  to  $E_1$  and  $I_2$  to  $E_2$  .

$p = (\text{sw} = I_1 ; \text{dup} ; \text{sw} := C ; \text{dup} ; \text{sw} := D ; \text{dup} ; \text{sw} := E_1) \ \& \ (\text{sw} = I_2 ; \text{dup} ; \text{sw} := C ; \text{dup} ; \text{sw} := D ; \text{dup} ; \text{sw} := E_2)$


# Example



Calculate not just **where** traffic is routed but also **how much** traffic is sent across each link.

# Example

In each time period, the number of packets originating at  $I_1/I_2$  is either 0, 1 or 2, with equal probability.





Calculate not just **where** traffic is routed but also **how much** traffic is sent across each link.



# Example



In each time period, the number of packets originating at  $I_1/I_2$  is either 0, 1 or 2, with equal probability.



Calculate not just **where** traffic is routed but also **how much** traffic is sent across each link.

# Example

In each time period, the number of packets originating at  $I_1/I_2$  is either 0, 1 or 2, with equal probability.




Calculate not just **where** traffic is routed but also **how much** traffic is sent across each link.

$$d1 = \text{drop} + \frac{1}{3} \text{ (purple square)@I1} + \frac{1}{3} ( \text{ (purple square)@I1} \& \text{ (blue square)@I1} )$$

# Example

In each time period, the number of packets originating at  $I_1/I_2$  is either 0, 1 or 2, with equal probability.




Calculate not just **where** traffic is routed but also **how much** traffic is sent across each link.

$$d1 = \text{drop} + \frac{1}{3} \text{ (purple square)@I1} + \frac{1}{3} ( \text{purple square@I1} \& \text{blue square@I1} )$$



# Example

In each time period, the number of packets originating at  $I_1/I_2$  is either 0, 1 or 2, with equal probability.



Calculate not just **where** traffic is routed but also **how much** traffic is sent across each link.

$$d1 = \text{drop} + \frac{1}{3} \text{ (purple @ } I_1) + \frac{1}{3} ( \text{ (purple @ } I_1 \& \text{ (blue @ } I_1) )$$

$$d2 = \text{drop} + \frac{1}{3} \text{ (red @ } I_2) + \frac{1}{3} ( \text{ (red @ } I_2 \& \text{ (green @ } I_2) )$$



# Probabilistic semantics

$$d1 = \text{drop} + \frac{1}{3} \text{ (purple square @I1)} + \frac{1}{3} \left( \text{purple square @I1} \& \text{blue square @I1} \right)$$

$$d2 = \text{drop} + \frac{1}{3} \text{ (red square @I2)} + \frac{1}{3} \left( \text{red square @I2} \& \text{green square @I2} \right)$$

## Distributions on set of histories

Full input distribution to the network : **d1 & d2**



# Probabilistic semantics

$$d1 = \text{drop} + \frac{1}{3} \text{ (purple square)}_{@I1} + \frac{1}{3} ( \text{ (purple square)}_{@I1} \& \text{ (blue square)}_{@I1} )$$


$$d2 = \text{drop} + \frac{1}{3} \text{ (red square)}_{@I2} + \frac{1}{3} ( \text{ (red square)}_{@I2} \& \text{ (green square)}_{@I2} )$$

## Distributions on set of histories

Full input distribution to the network : **d1 & d2**

Compositionality

# Probabilistic semantics



$$d1 = \text{drop} + \frac{1}{3} \quad \text{purple square} @I1 + \frac{1}{3} ( \text{purple square} @I1 \& \text{blue square} @I1 )$$

$$d2 = \text{drop} + \frac{1}{3} \quad \text{red square} @I2 + \frac{1}{3} ( \text{red square} @I2 \& \text{green square} @I2 )$$

$$p = (\text{sw} = I1 ; \text{dup} ; \text{sw} := C ; \text{dup} ; \text{sw} := D ; \text{dup} ; \text{sw} := E1) \& \\ (\text{sw} = I2 ; \text{dup} ; \text{sw} := C ; \text{dup} ; \text{sw} := D ; \text{dup} ; \text{sw} := E2)$$

Amount of congestion on links in the network?

$$\llbracket d; p \rrbracket =$$

$\frac{1}{9} . \{ \} +$  **no packet**  
 $\frac{1}{9} . \{ E1;1:C2;1:C1;1:I1;1 \} +$  **one packet from I1-E1**  
 $\frac{1}{9} . \{ E1;1:C2;1:C1;1:I1;1; E1;2:C2;2:C1;2:I1;2 \} +$  **two packets from I1-E1**  
 ....

# Probabilities are needed

- \* **expected congestion**: the network operator wishes to calculate the expected congestion on each link, given a model of incoming traffic
- \* **reliability**: the network operator wishes to calculate the probability of successful packet delivery given probability of failure of some network components
- \* **randomized routing**: the network operator wishes to use randomized routing schemes such as equal-cost multi-path routing (ECMP) or Valiant load balancing (VLB) to balance load across multiple paths

# Yet another (Net)KAT extension?

**The obvious extension does not work...**

$$\llbracket e \rrbracket : H \rightarrow 2^H$$

$$\llbracket e \rrbracket : H \rightarrow \mathcal{D}(2^H)$$

$$\llbracket e \rrbracket(h) = \delta(\llbracket e \rrbracket(h)) \} \text{ for the deterministic fragment}$$

$$\llbracket p +_r q \rrbracket(h) = r \llbracket p \rrbracket(h) + (1 - r) \llbracket q \rrbracket(h)$$

# Yet another (Net)KAT extension?

**The obvious extension does not work...**

$$\llbracket e \rrbracket : H \rightarrow 2^H$$

$$\llbracket e \rrbracket : H \rightarrow \mathcal{D}(2^H)$$

$$\llbracket e \rrbracket(h) = \delta(\llbracket e \rrbracket(h)) \} \text{ for the deterministic fragment}$$

$$\llbracket p +_r q \rrbracket(h) = r \llbracket p \rrbracket(h) + (1 - r) \llbracket q \rrbracket(h)$$

The meaning of a program on an input set of packet histories is **not** uniquely determined by its action on individual histories

# Yet another (Net)KAT extension?

**The obvious extension does not work...**

$$\llbracket e \rrbracket : H \rightarrow 2^H$$

$$\llbracket e \rrbracket : H \rightarrow \mathcal{D}(2^H)$$

$$\llbracket e \rrbracket(h) = \delta(\llbracket e \rrbracket(h)) \} \text{ for the deterministic fragment}$$

$$\llbracket p +_r q \rrbracket(h) = r \llbracket p \rrbracket(h) + (1 - r) \llbracket q \rrbracket(h)$$

The meaning of a program ~~on a set of histories~~ is **not** uniquely determined by its meaning on a set of histories

**Not compositional!!**

# Example bad semantics

$$\llbracket \pi_0! + .5 \pi_1! \rrbracket(\pi_1) = \llbracket (\pi_0! \& \pi_1!) + .5 \textit{drop} \rrbracket(\pi_1) = 0.5$$

**Problem 1** Different from desired meaning

$$\llbracket \pi_0! + .5 \pi_1! \rrbracket = 0.5\delta_{\{\pi_0\}} + 0.5\delta_{\{\pi_1\}}$$

$$\llbracket (\pi_0! \& \pi_1!) + .5 \textit{drop} \rrbracket = 0.5\delta_{\{\pi_0, \pi_1\}} + 0.5\delta_{\emptyset}$$

# Example bad semantics

$$\llbracket \pi_0! + .5 \pi_1! \rrbracket(\pi_1) = \llbracket (\pi_0! \& \pi_1!) + .5 \text{ drop} \rrbracket(\pi_1) = 0.5$$

**Problem 1** Different from desired meaning

$$\llbracket \pi_0! + .5 \pi_1! \rrbracket = 0.5\delta_{\{\pi_0\}} + 0.5\delta_{\{\pi_1\}}$$

$$\llbracket (\pi_0! \& \pi_1!) + .5 \text{ drop} \rrbracket = 0.5\delta_{\{\pi_0, \pi_1\}} + 0.5\delta_{\emptyset}$$

**Problem 2**

$$\llbracket (\pi_0! + .5 \pi_1!); \pi_0! \rrbracket = \delta_{\{\pi_0!\}}$$

$$\llbracket ((\pi_0! \& \pi_1!) + .5 \text{ drop}); \pi_0! \rrbracket = 0.5\delta_{\{\pi_0!\}} + 0.5\delta_{\emptyset}$$



# Example bad semantics

$$\llbracket \pi_0! + .5 \pi_1! \rrbracket(\pi_1) = \llbracket (\pi_0! \& \pi_1!) + .5 \text{ drop} \rrbracket(\pi_1) = 0.5$$

**Problem 1** Different from desired meaning

$$\llbracket \pi_0! + .5 \pi_1! \rrbracket = 0.5\delta_{\{\pi_0\}} + 0.5\delta_{\{\pi_1\}}$$

$$\llbracket (\pi_0! \& \pi_1!) + .5 \text{ drop} \rrbracket = 0.5\delta_{\{\pi_0, \pi_1\}} + 0.5\delta_{\emptyset}$$

**Problem 2**

**Not compositional!!**

$$\llbracket (\pi_0! + .5 \pi_1!); \pi_0! \rrbracket = \delta_{\{\pi_0!\}}$$

$$\llbracket ((\pi_0! \& \pi_1!) + .5 \text{ drop}); \pi_0! \rrbracket = 0.5\delta_{\{\pi_0!\}} + 0.5\delta_{\emptyset}$$

# It gets worse....

Discrete measures are not enough

$$p = \pi_0! + .5 \pi_1!$$

$$p; (dup; p)^*$$

# It gets worse....

Discrete measures are not enough

$$p = \pi_0! + .5 \pi_1!$$

$$p; (dup; p)^*$$

*Sets the input packet to either 0 or 1  
with equal probability,  
then repeat:  
(i) output the current packet,  
(ii) duplicate the current packet, and  
(iii) set the new current packet to 0  
or 1 with equal probability.*

# It gets worse....

Discrete measures are not enough

$$p = \pi_0! + .5 \pi_1!$$

$$p; (dup; p)^*$$

$\llbracket p; (dup; p)^* \rrbracket(\pi_0)$  is a continuous measure

# Our solution

## Markov Kernels

$$[[e]]: 2^H \times \mathcal{B} \rightarrow \mathbb{R}$$

measurable on the first argument  
probability measure on the second argument

$$\mathcal{B} \subseteq 2^{2^H}$$

smallest sigma-algebra  
containing

$$\mathcal{B}_\tau = \{a \in 2^H \mid \tau \in a\}$$

# Our solution

## Markov Kernels

$$\llbracket e \rrbracket : 2^H \times \mathcal{B} \rightarrow \mathbb{R}$$

$$\mathcal{B} \subseteq 2^{2^H}$$

smallest sigma-algebra  
containing

$$\mathcal{B}_\tau = \{a \in 2^H \mid \tau \in a\}$$

$$\llbracket x \leftarrow n \rrbracket(a) = \delta_{\{\pi[n/x] : \sigma \mid \pi : \sigma \in a\}}$$

$$\llbracket x = n \rrbracket(a) = \delta_{\{\pi : \sigma \mid \pi : \sigma \in a, \pi(x)=n\}}$$

$$\llbracket \text{dup} \rrbracket(a) = \delta_{\{\pi : \pi : \sigma \mid \pi : \sigma \in a\}}$$

$$\llbracket \text{skip} \rrbracket(a) = \delta_a$$

$$\llbracket \text{drop} \rrbracket(a) = \delta_\emptyset$$

# Our solution

## Markov Kernels

$$\llbracket e \rrbracket : 2^H \times \mathcal{B} \rightarrow \mathbb{R}$$

$$\mathcal{B} \subseteq 2^{2^H}$$

smallest sigma-algebra  
containing

$$\mathcal{B}_\tau = \{a \in 2^H \mid \tau \in a\}$$

$$\llbracket x \leftarrow n \rrbracket(a) = \delta_{\{\pi[n/x] : \sigma \mid \pi : \sigma \in a\}}$$

$$\llbracket x = n \rrbracket(a) = \delta_{\{\pi : \sigma \mid \pi : \sigma \in a, \pi(x)=n\}}$$

$$\llbracket \text{dup} \rrbracket(a) = \delta_{\{\pi : \pi : \sigma \mid \pi : \sigma \in a\}}$$

$$\llbracket \text{skip} \rrbracket(a) = \delta_a$$

$$\llbracket \text{drop} \rrbracket(a) = \delta_\emptyset$$

$$\llbracket p \ \& \ q \rrbracket(a) = \llbracket p \rrbracket(a) \ \& \ \llbracket q \rrbracket(a)$$

$$(\mu \ \& \ \nu)(A) \triangleq (\mu \times \nu)(\{(a, b) \mid a \cup b \in A\}).$$

# Our solution

## Markov Kernels

$$\llbracket e \rrbracket : 2^H \times \mathcal{B} \rightarrow \mathbb{R}$$

$$\mathcal{B} \subseteq 2^{2^H}$$

smallest sigma-algebra  
containing

$$\mathcal{B}_\tau = \{a \in 2^H \mid \tau \in a\}$$

$$\llbracket x \leftarrow n \rrbracket(a) = \delta_{\{\pi[n/x] : \sigma \mid \pi : \sigma \in a\}}$$

$$\llbracket x = n \rrbracket(a) = \delta_{\{\pi : \sigma \mid \pi : \sigma \in a, \pi(x) = n\}}$$

$$\llbracket \text{dup} \rrbracket(a) = \delta_{\{\pi : \pi : \sigma \mid \pi : \sigma \in a\}}$$

$$\llbracket \text{skip} \rrbracket(a) = \delta_a$$

$$\llbracket \text{drop} \rrbracket(a) = \delta_\emptyset$$

$$\llbracket p \ \& \ q \rrbracket(a) = \llbracket p \rrbracket(a) \ \& \ \llbracket q \rrbracket(a)$$

$$(\mu \ \& \ \nu)(A) \triangleq (\mu \times \nu)(\{(a, b) \mid a \cup b \in A\}).$$

$$\llbracket p +_r q \rrbracket(a) = r \llbracket p \rrbracket(a) + (1 - r) \llbracket p \rrbracket(a)$$


**set of histories!**



# Properties

## Conservative Extension

For deterministic programs, ProbNetKAT semantics and NetKAT semantics agree



The NetKAT axioms are sound and complete for deterministic ProbNetKAT programs.

# Some more properties

$$[[p \ \& \ \text{drop}]] = [[\text{drop} \ \& \ p]] = [[p]]$$

$$[[p \ +_r \ p]] = [[p]]$$


$$[[p \ +_r \ q]] = [[q \ +_{1-r} \ p]]$$

$$[[ (p \ \& \ q) \ \& \ s ]] = [[ p \ \& \ (q \ \& \ s) ]]$$

...


# Applications

## Fault Tolerance




Probability packets are delivered at  $S_4$  if  $S_1 \rightarrow S_2$  fails 10%

## Load balancing



Maximum number of packets traversing a link


## Gossiping protocols



Expected number of infected nodes after  $n$  rounds


# Applications

## Fault Tolerance




Probability packets are delivered at S4 if S1→S2 fails 10%

## Load balancing




Maximum number of packets traversing a link

## Gossiping protocols




Expected number of infected nodes after n rounds

# Fault Tolerance




# Fault Tolerance

What is the probability that a packet that originates at  $S_1$  will be successfully delivered to  $S_4$  ?



# Fault Tolerance

What is the probability that a packet that originates at S1 will be successfully delivered to S4 ?



```
t = sw=S1 ; pt =2 ;((sw:=2 ; pt:=1) +.9 drop)
    & sw=S1 ; pt =2 ;sw:=3 ; pt:=1)
    & sw=S2 ; pt =4 ;sw:=4 ; pt:=2)
    & sw=S3 ; pt =4 ;sw:=4 ; pt:=3)
```

# Fault Tolerance

## Switch behavior

**p** = (sw=1; pt:=2) & (sw =2 ; pt:=4)


all traffic via S2

**q** = (sw=1; (pt:=2 +.5 pt:=3))  
& (sw =2 ; pt:=4)& (sw=3 ; pt:=4)

traffic split between S2 and S4

**e** = sw=4

egress predicate





# Fault Tolerance

## Switch behavior

$\mathbf{p} = (\text{sw}=1; \text{pt}:=2) \ \& \ (\text{sw} = 2 \ ; \ \text{pt}:=4)$


all traffic via S2

What is the probability that a packet that originates at S1 will be successfully delivered to S4 ?

traffic split between S2 and S4

$\mathbf{e} = \text{sw}=4$

egress predicate



# Fault Tolerance

## Switch behavior

$\mathbf{p} = (\text{sw}=1; \text{pt}:=2) \ \& \ (\text{sw} = 2 \ ; \ \text{pt}:=4)$


all traffic via S2

What is the probability that a packet that originates at S1 will be successfully delivered to S4 ?

traffic split between S2 and S4

$\mathbf{e} = \text{sw}=4$

egress predicate



$(\mathbf{p};\mathbf{t})^*;\mathbf{e}$

90%

$[[ - ]]$

$(\mathbf{q};\mathbf{t})^*;\mathbf{e}$

95%

# Conclusions

- ◆ First language-based framework for specifying and verifying **probabilistic network behavior**.
- ◆ Formal semantics for ProbNetKAT based on Markov kernels (conservative over NetKAT).
- ◆ Notion of approximation — every ProbNetKAT program is arbitrarily closely approximated by loop-free programs.
- ◆ Several case studies — fault tolerance, load balancing, and a probabilistic gossip protocol.

# Future work

Axiomatizations

Decision procedure

Simulation

Certified  
Compiler

# Future work

Axiomatizations

Decision procedure

Simulation

Certified  
Compiler

**Questions?**