


# Automatically Correcting Networks with *NEAt*

Wenxuan Zhou, Jason Croft, Bingzhe Liu,  
Elaine Ang, Matthew Caesar

University of Illinois at Urbana-Champaign




# A simple idea about complexity...



Networks are so complex it's hard to make sure they're doing the right thing.

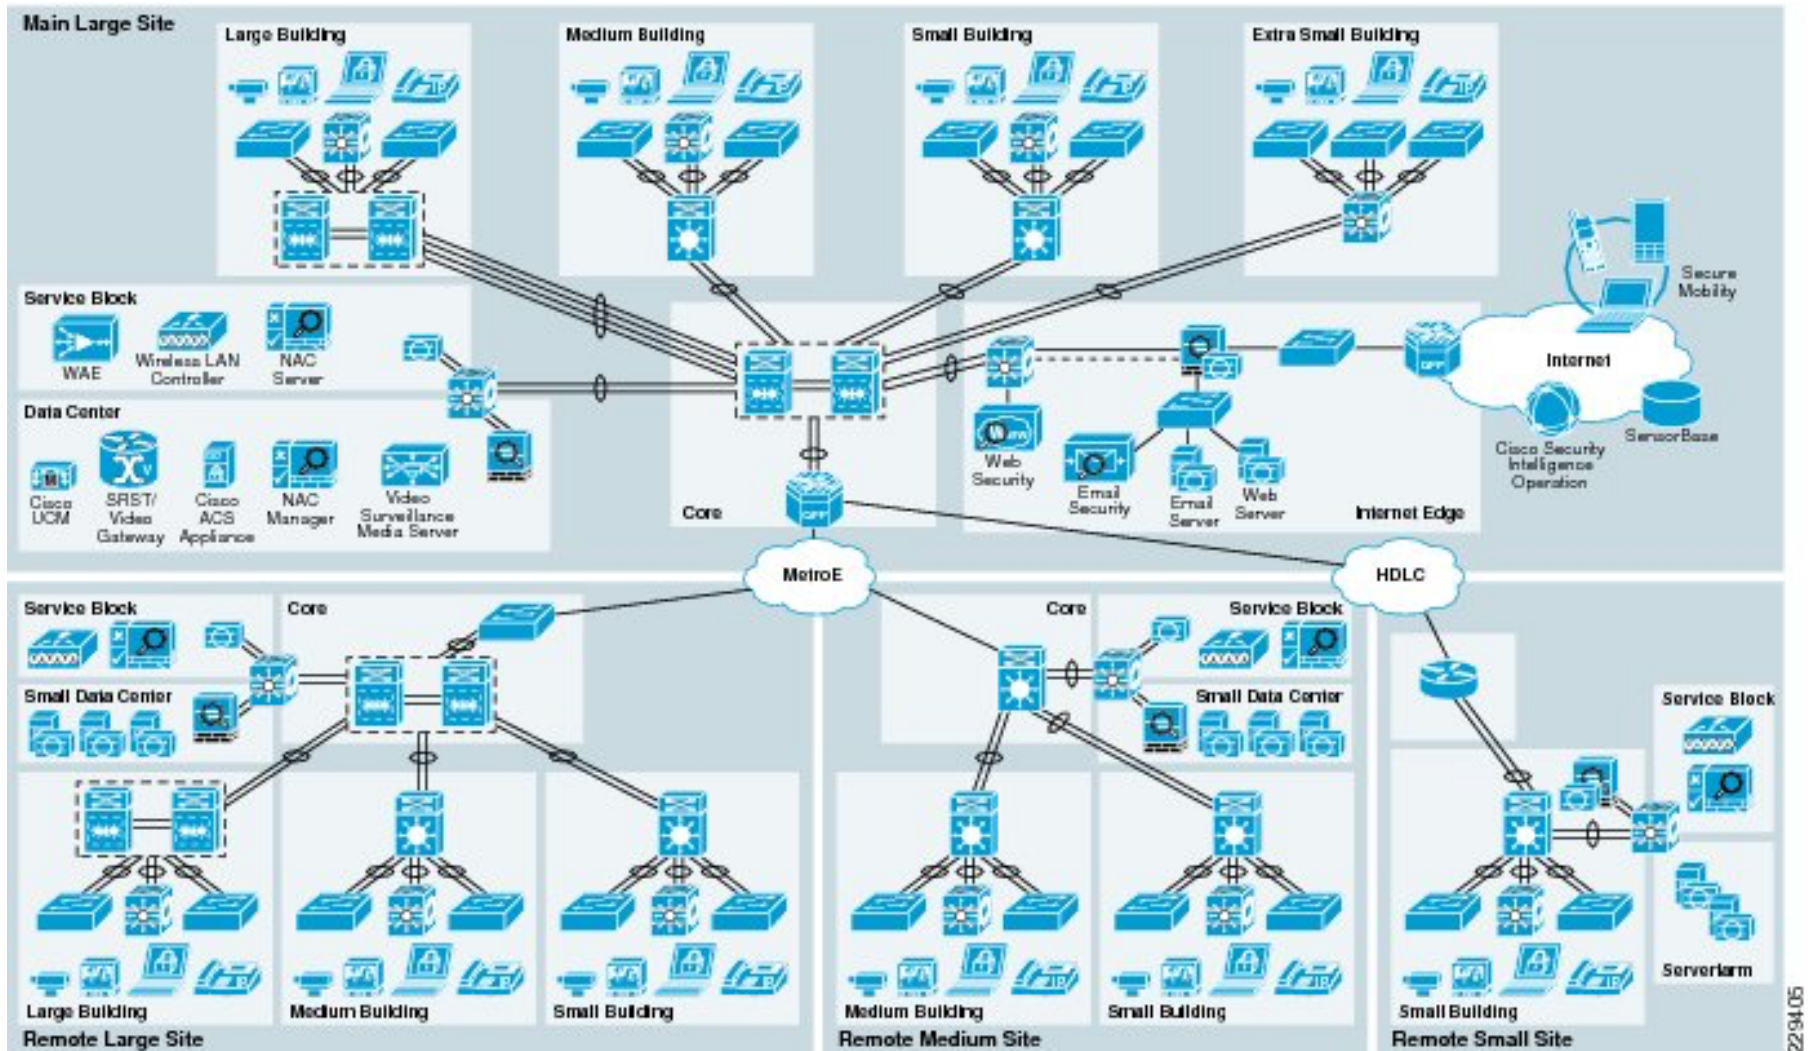
# What is a network supposed to do?



They instill policies

- no untrusted traffic entering a secure zone
- the preference of one path over another
- loop & black hole avoidance
- ...

# A Typical Enterprise Network



# Network errors are common



Lots of problems arise today



**89%** of operators are not certain their configuration changes are safe. *[Kinetic NSDI'15]*

Networks are so complex it's hard to make sure they're doing the right thing.


Let's automate.

What to automate?


# Many have tried.




## Automatically identifying errors in networks (Verification)



## Automatically synthesize correct networks



# What if we can automatically correct networks on the fly?







# What if we can automatically correct networks on the fly?



## Auto-correct is not a new idea.



c++ unordered\_set implementation

All Images Videos News Shopping More Settings Tools

About 127,000 results (0.52 seconds)

Showing results for **c++ unordered set implementation**

[ 16%] Building CXX object CMakeFiles/neat\_cpp.dir/src/clus

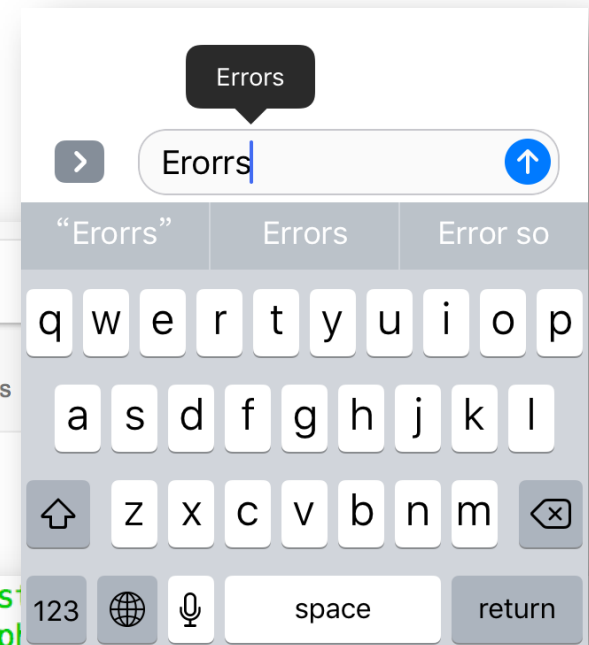
[ 33%] Building CXX object CMakeFiles/neat\_cpp.dir/src/Graph

[ 50%] Building CXX object CMakeFiles/neat\_cpp.dir/src/BloomFilter.cpp.o

[/Users/violet/Documents/projects/gcc/corrector/neat\\_cpp/src/clustering.cpp:18:16: error: member](#)

reference type 'Vertex \*' is a pointer; maybe you meant to use '->'


VID cur\_id = v.id;




# Network Error Auto-Correct



What if we can automatically correct networks on the fly?



## Network Abstraction



Operating on the data plane simplifies our work

- *Diagnose problems as close as possible to actual network behavior*
- *Data plane is a “narrower waist” than configuration*



Goal: Improve upon a manual effort with transparency in both performance and architecture.

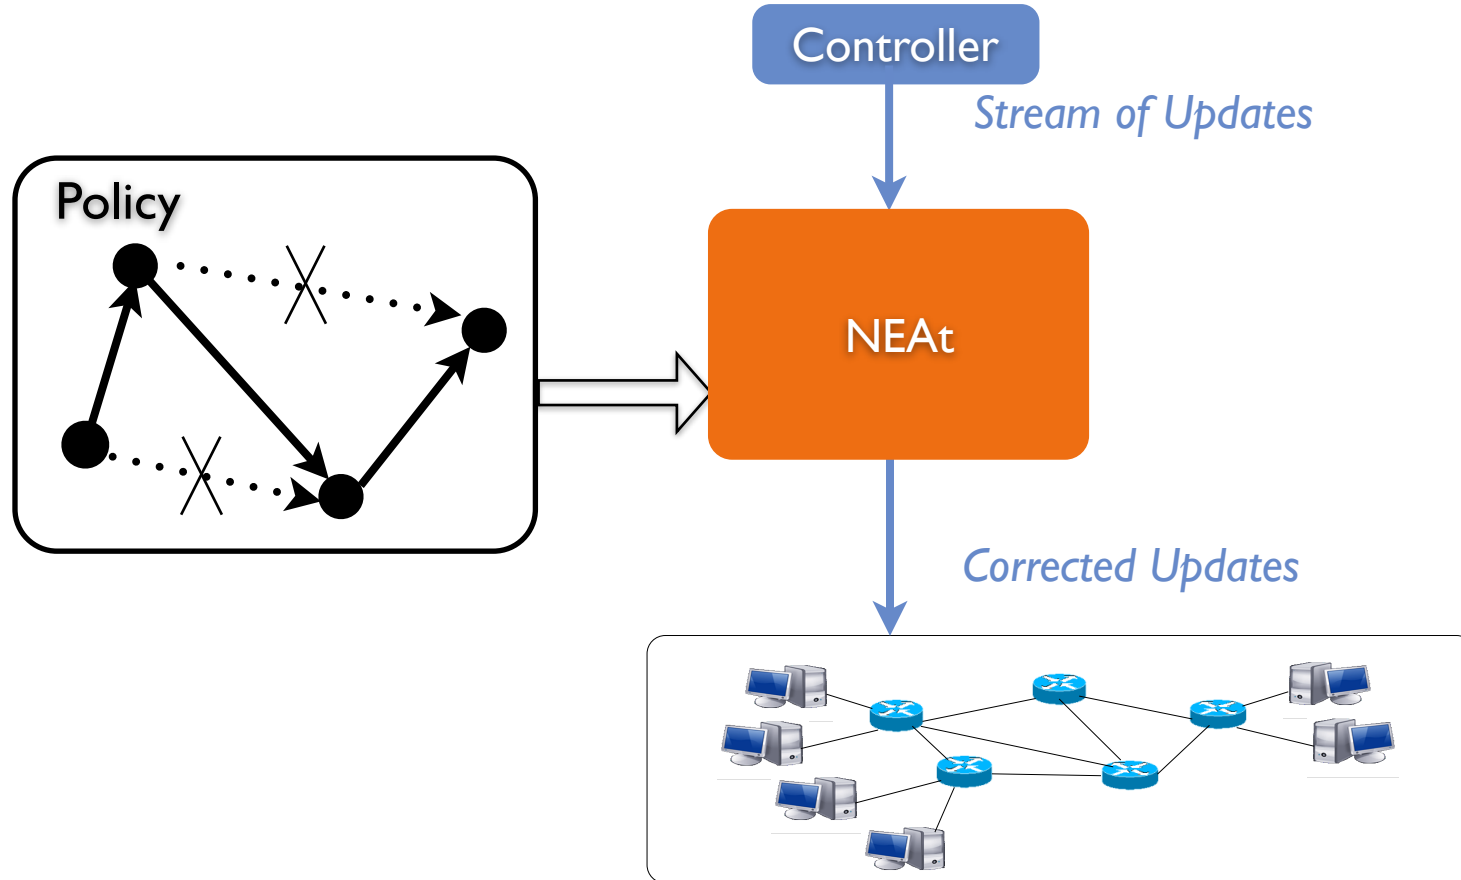
## Challenge 1: Repair speed

- Based on real-time verification technique
- Derive fixes via linear optimization, with min. changes
- Topology limitation & graph compression

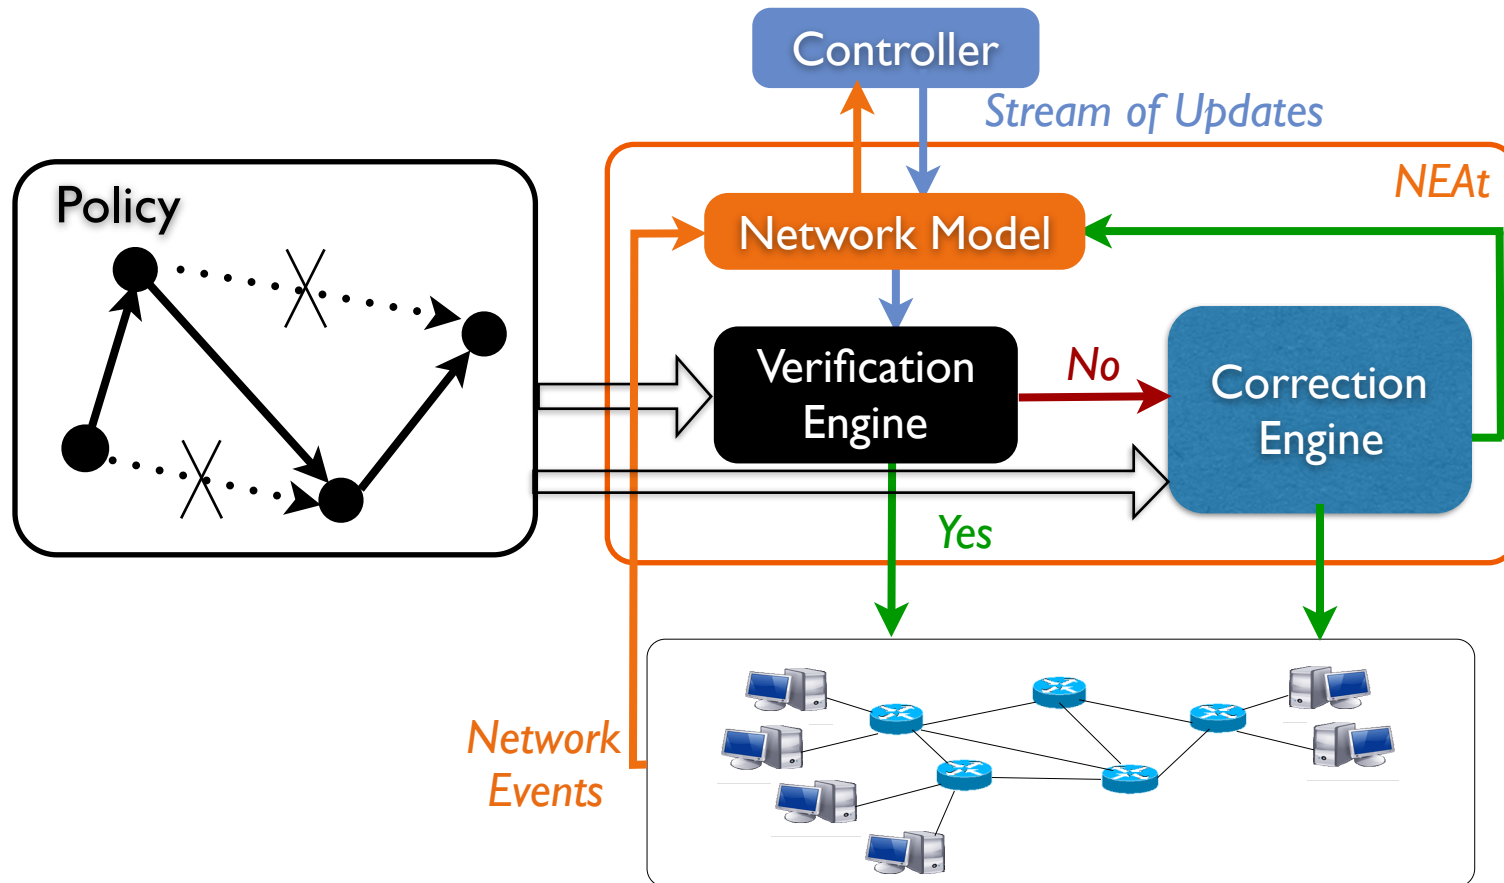
## Challenge 2: Zero/minimal architecture/application changes

- Minimal changes
- Pass-through mode
- Interactive mode


# Design of NEAt



# Design of NEAt




# Application Mode



# Policy as Graphs




Graphs are **neat**



- Network state synthesis → viewing the network as a whole.
- Graphs → richer set of policies.

A *policy graph* is defined on a packet header pattern

- ip dst 10.0.1.0/24, port 443.




Reachability	$m = 1$
Bounded path length (shortest path)	$m = 1$ $n = \text{path\_length}$
Multipath/Resillience	$m = k \ (k > 1)$
Isolation	$m = 0$



# Policy as Graphs (Cont'd)



## Service Chaining



## Load balancing



Use *policy graphs* to express both qualitative and quantitative reachability constraints



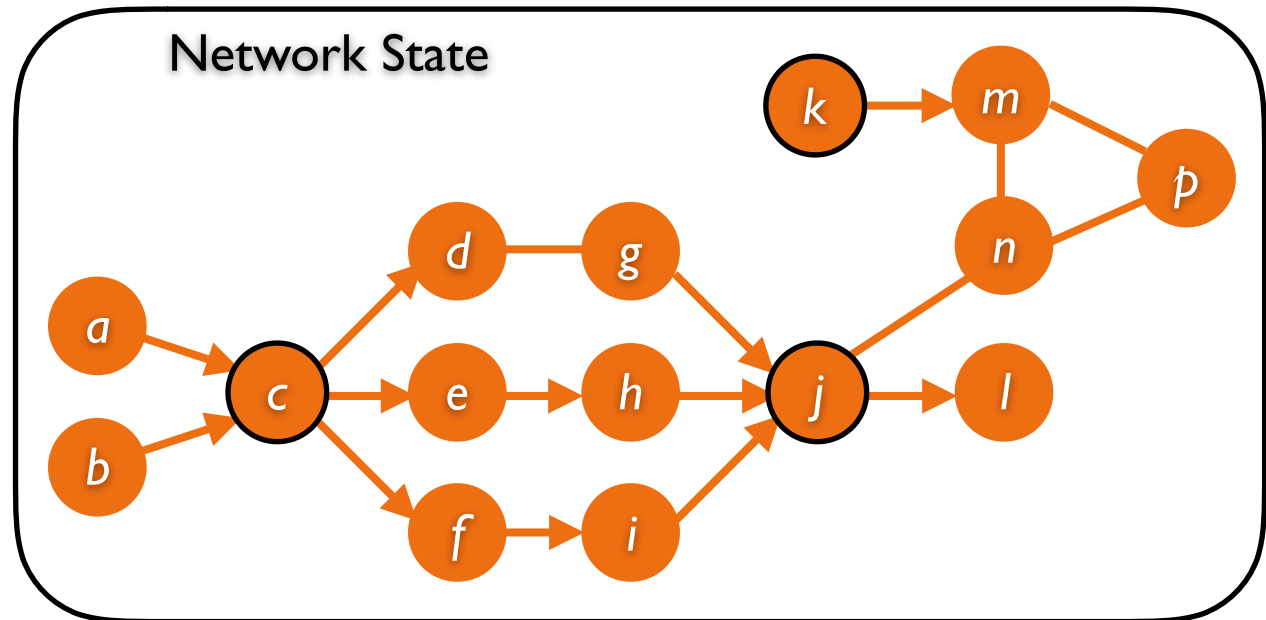

# Repair Algorithm



Cast the problem as an optimization problem:

- Map forwarding graph to policy graph
- Minimize # of changes


Correction  
Engine



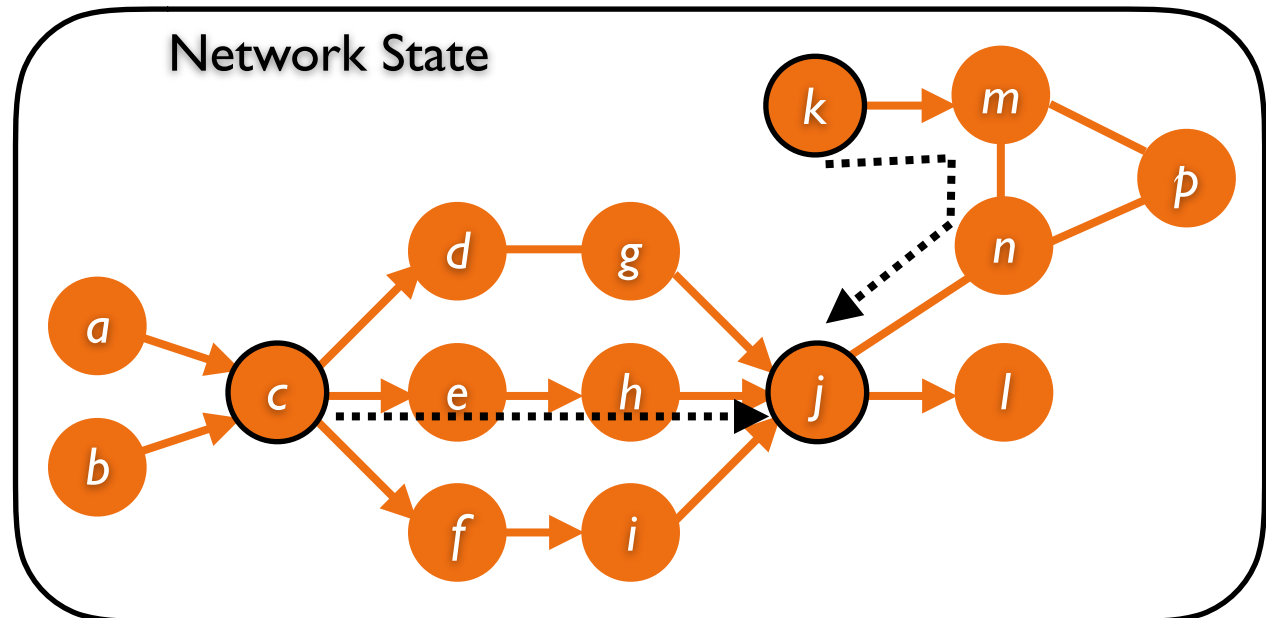

# Repair - Basic Reachability



Cast the problem as an optimization problem:

- Map forwarding graph to policy graph
- Minimize # of changes
- boolean variable  $x_{i,j,p,q}$ :
  - *topology edge*  $(i,j)$   *policy edge*  $(p,q)$
  - s.t., policy level reachability  $(p,q)$

Correction  
Engine



# Repair - Generalized Reachability



## Basic Reachability

$$\forall(i, j) \quad x_{i,j} \geq \sum_{(p,q) \in E_{\emptyset c}} \frac{x_{i,j,p,q}}{N(E_{\emptyset c})}$$

$$\forall(j, i) \quad x_{j,i} \geq \sum_{(p,q) \in E_{\emptyset c}} \frac{x_{j,i,p,q}}{N(E_{\emptyset c})}$$

$$\forall(j, i) \quad x_{i,j} + x_{j,i} \leq 1$$

$$\forall(p, q), \forall i \in T:$$

$$\begin{cases} \sum_{j \in NB_T(i)} x_{i,j,p,q} = 1 & \text{if } i = p \\ \sum_{j \in NB_T(i)} x_{j,i,p,q} = 0 & \end{cases}$$

$$\begin{cases} \sum_{j \in NB_T(i)} x_{i,j,p,q} = 0 & \text{if } i = q \\ \sum_{j \in NB_T(i)} x_{j,i,p,q} = 1 & \end{cases}$$

$$\left\{ \sum_{j \in NB_T(i)} (x_{i,j,p,q} - x_{j,i,p,q}) = 0 \quad \text{otherwise} \right.$$

## • Isolation

$$\begin{cases} \sum_{j \in NB_T(i)} x_{i,j,p,q} = 0 & \text{if } i = DROP \\ \sum_{j \in NB_T(i)} x_{j,i,p,q} = 1 & \text{if } i = q \\ \sum_{j \in NB_T(i)} x_{i,j,p,q} = 0 & \text{if } i = q \end{cases}$$

## • Service Chaining

$$\begin{cases} \sum_{j \in NB_T(i)} x_{i,j,p,q} = 0 & \text{if } i = p \text{ and } \exists i, q \in \mathcal{P}_q(i) \\ \sum_{j \in NB_T(i)} x_{j,i,p,q} = 0 & \end{cases}$$

## • Bounded or Equal Path Length

$$\sum_{(i,j) \in E_T} (x_{i,j,p,q} + x_{j,i,p,q}) \leq n$$

## • MultiPath (Resilience)

$$\begin{cases} \sum_{j \in NB_T(i)} x_{i,j,p,q} \geq m & \text{if } i = p \\ \sum_{j \in NB_T(i)} x_{j,i,p,q} = 0 & \\ \sum_{j \in NB_T(i)} x_{i,j,p,q} = 0 & \text{if } i = q \\ \sum_{j \in NB_T(i)} x_{j,i,p,q} \geq m & \end{cases}$$

## • Load Balancing

$$\prod_{(i,j) \in E_T} x_{i,j,p,q} = m$$



The preceding algorithm operates on a *loop-free* graph.

First check for and remove loops before repairing other type violations.

Objective: Minimize changes

- Remove the minimal # of rules.
- Affect few packets as possible.
  - E.g. remove a rule matching 10.0.0.1/32 over one for 10.0.0.0/8.

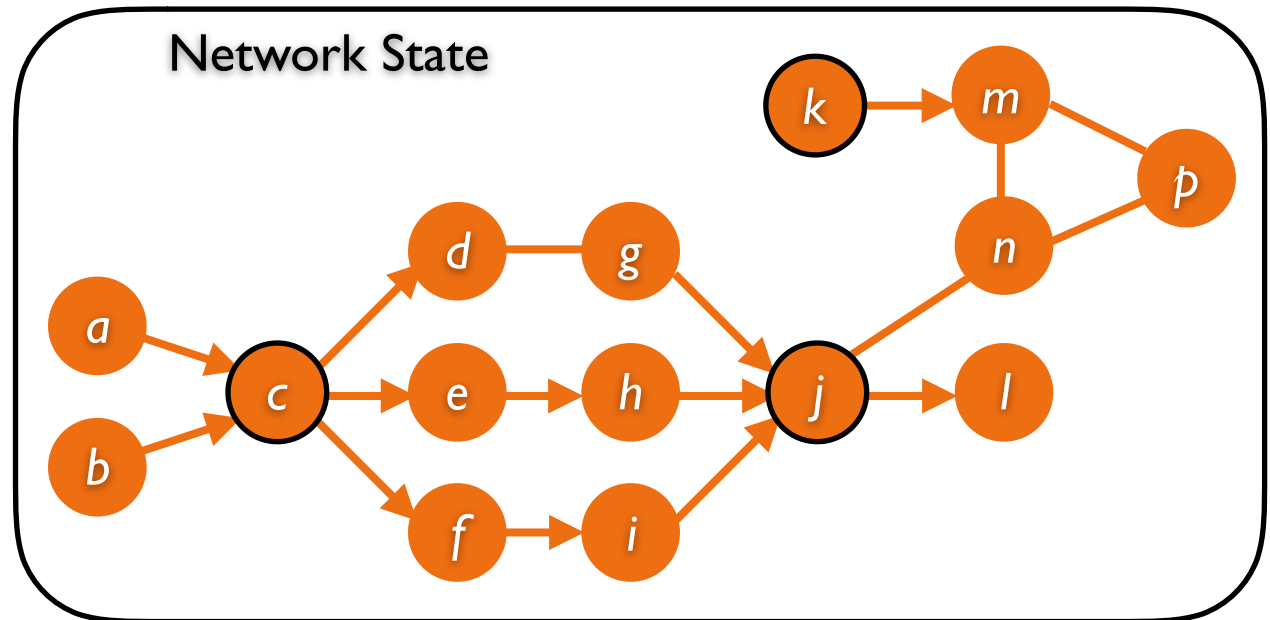

# Scalability Challenge and Solution



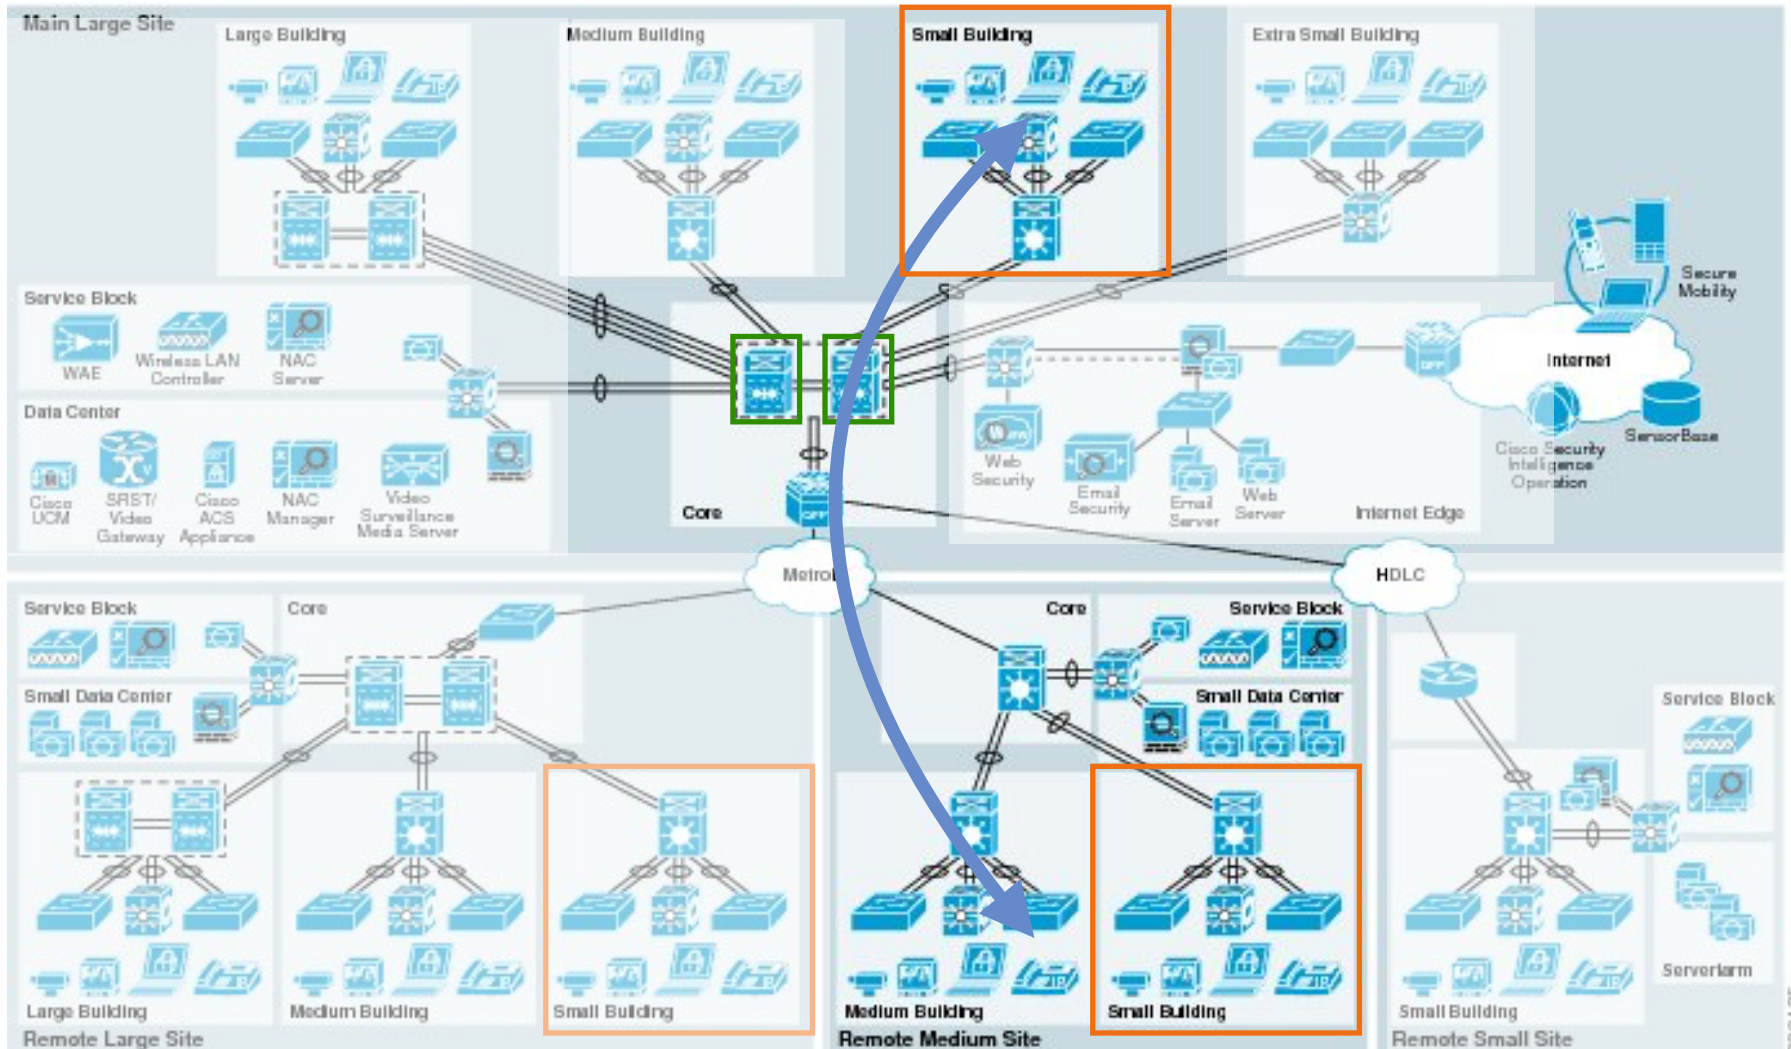
## Scalability challenge

- # of variables  $\approx |E(G_{\text{topo}})| \times |E(G_{\text{policy}})|$
- Easily exceeds 100k

Solution: ?



# A Typical Enterprise Network

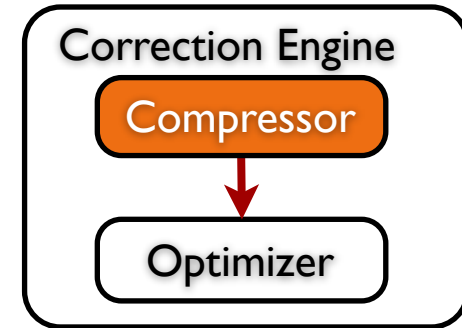


# Scalability Challenge and Solution



## Scalability challenge

- # of variables  $\approx |E(G_{\text{topo}})| \times |E(G_{\text{policy}})|$
- Easily exceeds 100k



## Solution

- Topology Limitation
- Graph Compression *w.r.t policy*
  - Key: Compressed graph == original graph
  - Bisimulation Based Graph Compression



## Prototype implementation in Python

- Use Gurobi within optimization engine
- Pass-through mode: proxy
- Interactive mode: XML-RPC API

## Datasets:

- Synthetic fat-tree configurations
- SDN applications
- 244-router enterprise network trace




# Application End-to-End Delay



## Pox + Mininet

- Learning switch app (pass-through)
- Load balancing app (interactive)



# Enterprise Network Trace Study



244 routers, one million forwarding rules

Policy: loop freedom & reachability

Issues found and repaired:

- Loops caused by default route
- Load balancing shouldn't be turned on

# Repair vs. Synthesis



## Synthesizer (NetGen) as repair tool

#TopoLinks	NEAt	NetGen	NetGen-C
96	5.9ms	743.2ms	513.2ms
324	7.2ms	4404.0ms	1160.8ms
768	9.0ms	16337.7m	2056.3ms

## NEAt as synthesizer

#TopoLinks	NEAt	NetGen
96	921.7ms	7.1min
324	16.3s	381.7min
768	2.9min	173.2hrs

# Isn't that NEAt?



*NEAt*, a system analogous to a smartphone's autocorrect.

- Casting the problem as an optimization problem
- Millisecond to second repair speed
- Generic policy support

Future work:

- Evolving & richer policies
- Different optimization goals
- Repair relevance study





# Graph Compression



*w.r.t policy*  
Key: Compressed graph == original graph

Major building block:

- Bisimulation Based Graph Compression\*




\*Query preserving graph compression. *SIGMOD 2012. W. Fan et al*

# Graph Compression (Cont'd)



- Customized policy-preserving compression

Topology	Compression Ratio
Fattree (6750 hosts, 1125 switches)	99.38%
Enterprise (236 routers)	88.98%



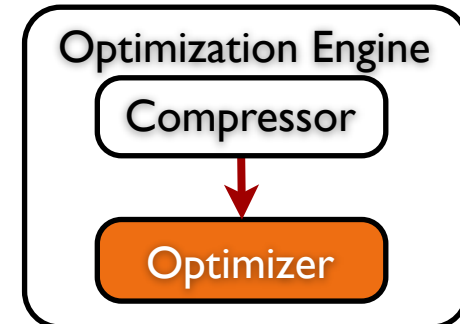
- Incremental Compression

- Repair compressed graphs

$$\begin{cases} \sum_{j \in NB_{cp}(i)} (x_{j,p,q} * weight_{ij}) \geq m \\ \sum_{j \in NB_{tcp}(i)} x_{j,i,p,q} = 0 \end{cases}$$

- Mapping back

- Proved Policy Perseverance




# Backup: Network Model




Correction  
Engine

Model packet space as a set of Equivalence Classes

*Equivalence class (EC)*: Packets experiencing the same forwarding actions throughout the network.



Model forwarding behavior of each EC as a directed graph






## Preventing errors at run-time

- Allow arbitrary SDN applications to run on top
- Not restricted to any programming language
- Influence updates, rather than synthesize from scratch

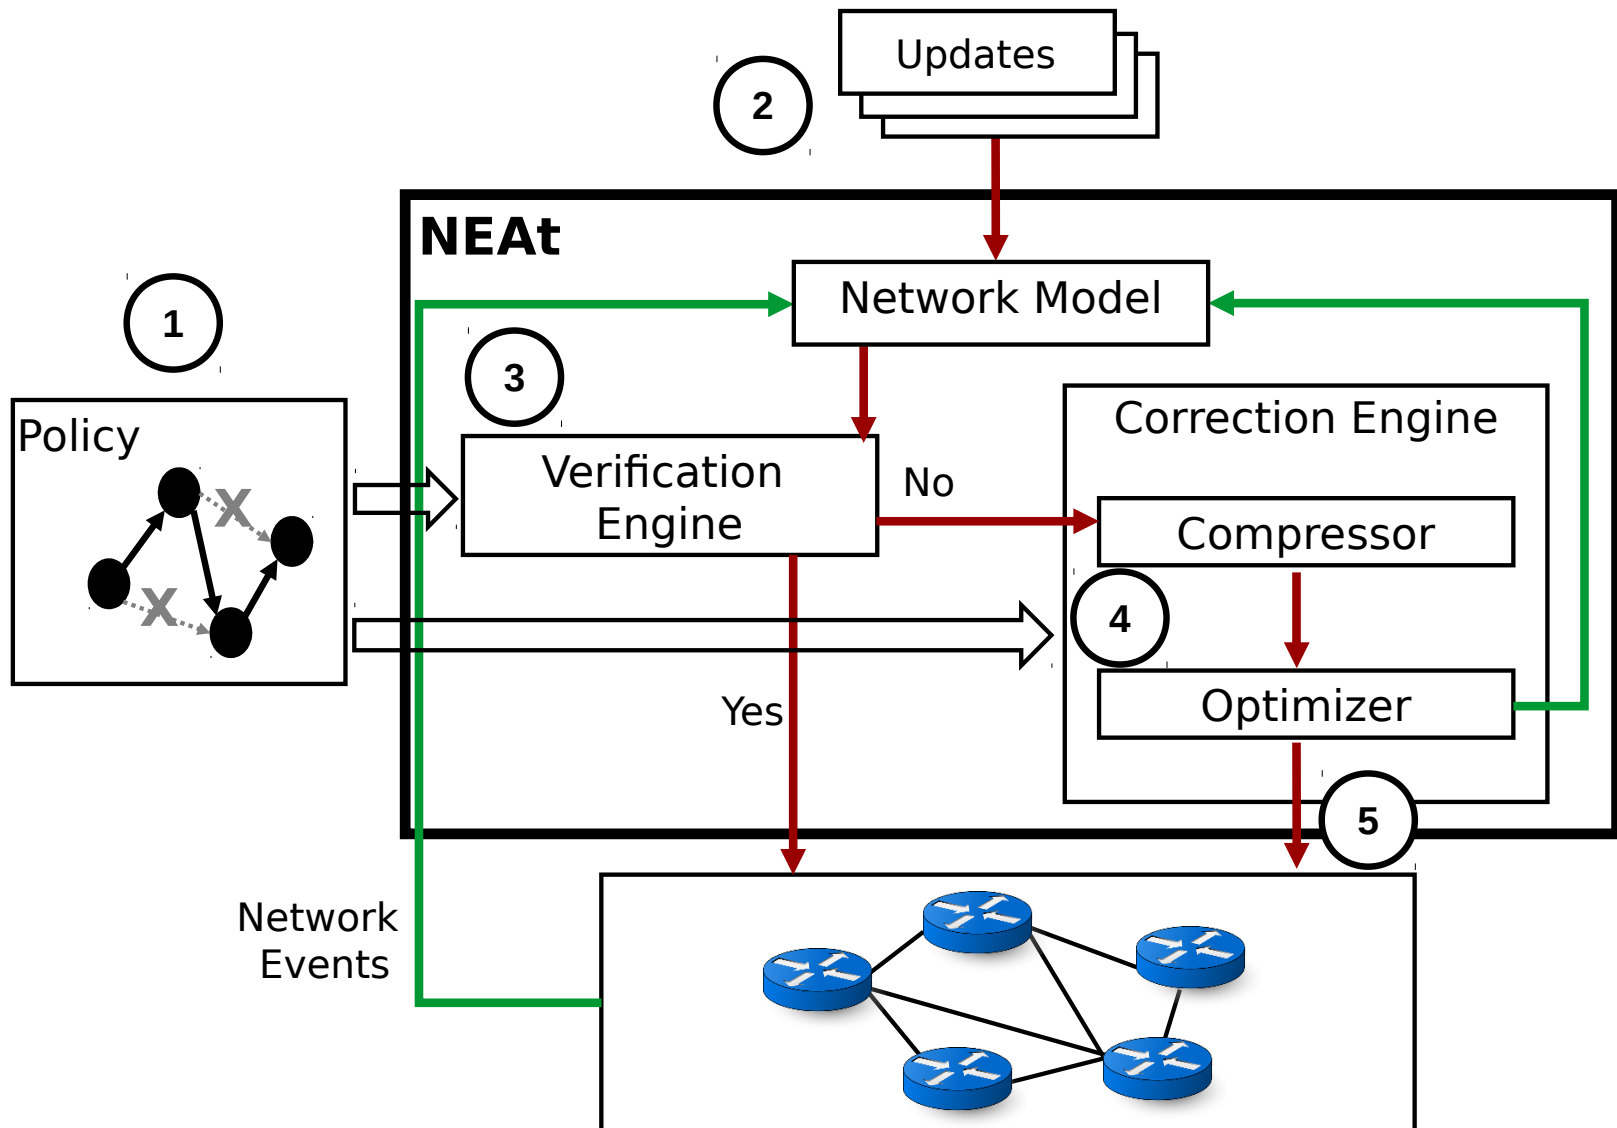
## Graphs are **neat**

- Networks are graphs
- Model network forwarding behaviors as directed graphs
- Represent operator intents as a policy graph



## Discovering repairs

- Equivalent to modifying network state graph so that there exists a mapping between it and the policy graph



# Approach: Data plane verification & repair



Configuration

Control plane

Data plane

Network  
behavior

Configuration	Data-plane
Prediction is difficult: <ul style="list-style-type: none"><li>• Various configuration languages</li><li>• Dynamic distributed protocols</li></ul>	Closer to actual network behavior
Misses control-plane bugs	Unified analysis for multiple control-plane protocols
Test prior to deployment	Can catch control-plane bugs
	Only detects bugs that are present in the data plane

Operating on the data plane simplifies our work

- *Diagnose problems as close as possible to actual network behavior*
- *Data plane is a “narrower waist” than configuration*