

# **PoiRoot:** Investigating the Root Cause of Interdomain Path Changes

Umar Javed\*, Italo Cunha<sup>^</sup>, **David Choffnes**<sup>\*#</sup>, Ethan Katz-Bassett<sup>\$</sup>, Tom Anderson\*, Arvind Krishnamurthy\*

\*U. Washington



<sup>^</sup>UFMG



#Northeastern



<sup>\$</sup>USC



# The case of the missing trigger

---

- ▶ **BGP: the glue that holds the Internet together**
  
- ▶ **When paths change, bad things can happen**
  - ▶ 100ms or more additional user latency caused by interdomain path changes 40% of the time (Google, Zhu et al. '12)
  - ▶ 50% of unintelligible VoIP samples caused by a BGP update (Kushman et al. '07)
  
- ▶ **When bad things happen, who do you blame?**

# Root cause of interdomain path changes

---

- ▶ **Goal:** Find the network triggering a BGP path change
- ▶ **Challenge:** BGP is an information-*hiding* protocol
- ▶ **Our contributions**
  - ▶ New model for path change propagation
  - ▶ Algorithm for identifying the root cause of an arbitrary change
  - ▶ Evaluation using controlled experiments on real Internet routes

# Outline


---

- ▶ Intro
- ▶ **Motivation**
- ▶ **Root Cause Isolation**
- ▶ **System & Evaluation**

# Assumptions in this talk



---

- ▶ Consider a routing toward a single prefix, owned by AS D
  - ▶ Path changes for D occur one at a time
- ▶ Model ASN as a single node
  - ▶ Approaches apply to finer granularity
- ▶ Vantage point: Location at which we measure a path to D
  - ▶ BGP feed
  - ▶ Traceroute (forward and reverse)



# Approach from previous work



- ▶ Feldmann et al. '04, Caesar et al. '03
  - ▶ Key assumption: root cause is on either the **new** or **old** path (NOOP)
  - ▶ Intersection-based approach



Vantage points: {Q, V}

# Approach from previous work



- ▶ Feldmann et al. '04, Caesar et al. '03
  - ▶ Key assumption: root cause is on either the **new** or **old** path (NOOP)
  - ▶ Intersection-based approach



Vantage points: {Q, V}

# Approach from previous work

- ▶ Feldmann et al. '04, Caesar et al. '03
  - ▶ Key assumption: root cause is on either the **new or old path** (NOOP)
  - ▶ Intersection-based approach





Vantage points: {Q, V}



# Approach from previous work



- ▶ Feldmann et al. '04, Caesar et al. '03
  - ▶ Key assumption: root cause is on either the **new or old path** (NOOP)
  - ▶ Intersection-based approach



Vantage points: {Q, V}

# Approach from previous work



## ▶ NOOP algorithm



Vantage points: {Q, V}

# Approach from previous work

- ▶ NOOP algorithm
  - ▶ Union of all old and new paths: {Q,V,X,Y,Z}




Vantage points: {Q,V}

# Approach from previous work

- ▶ **NOOP algorithm**



- ▶ Union of all old and new paths:  $\{Q, V, X, Y, Z\}$
- ▶ Intersection of all old and new paths:  $\{Q, V, X, Z\}$



# Approach from previous work

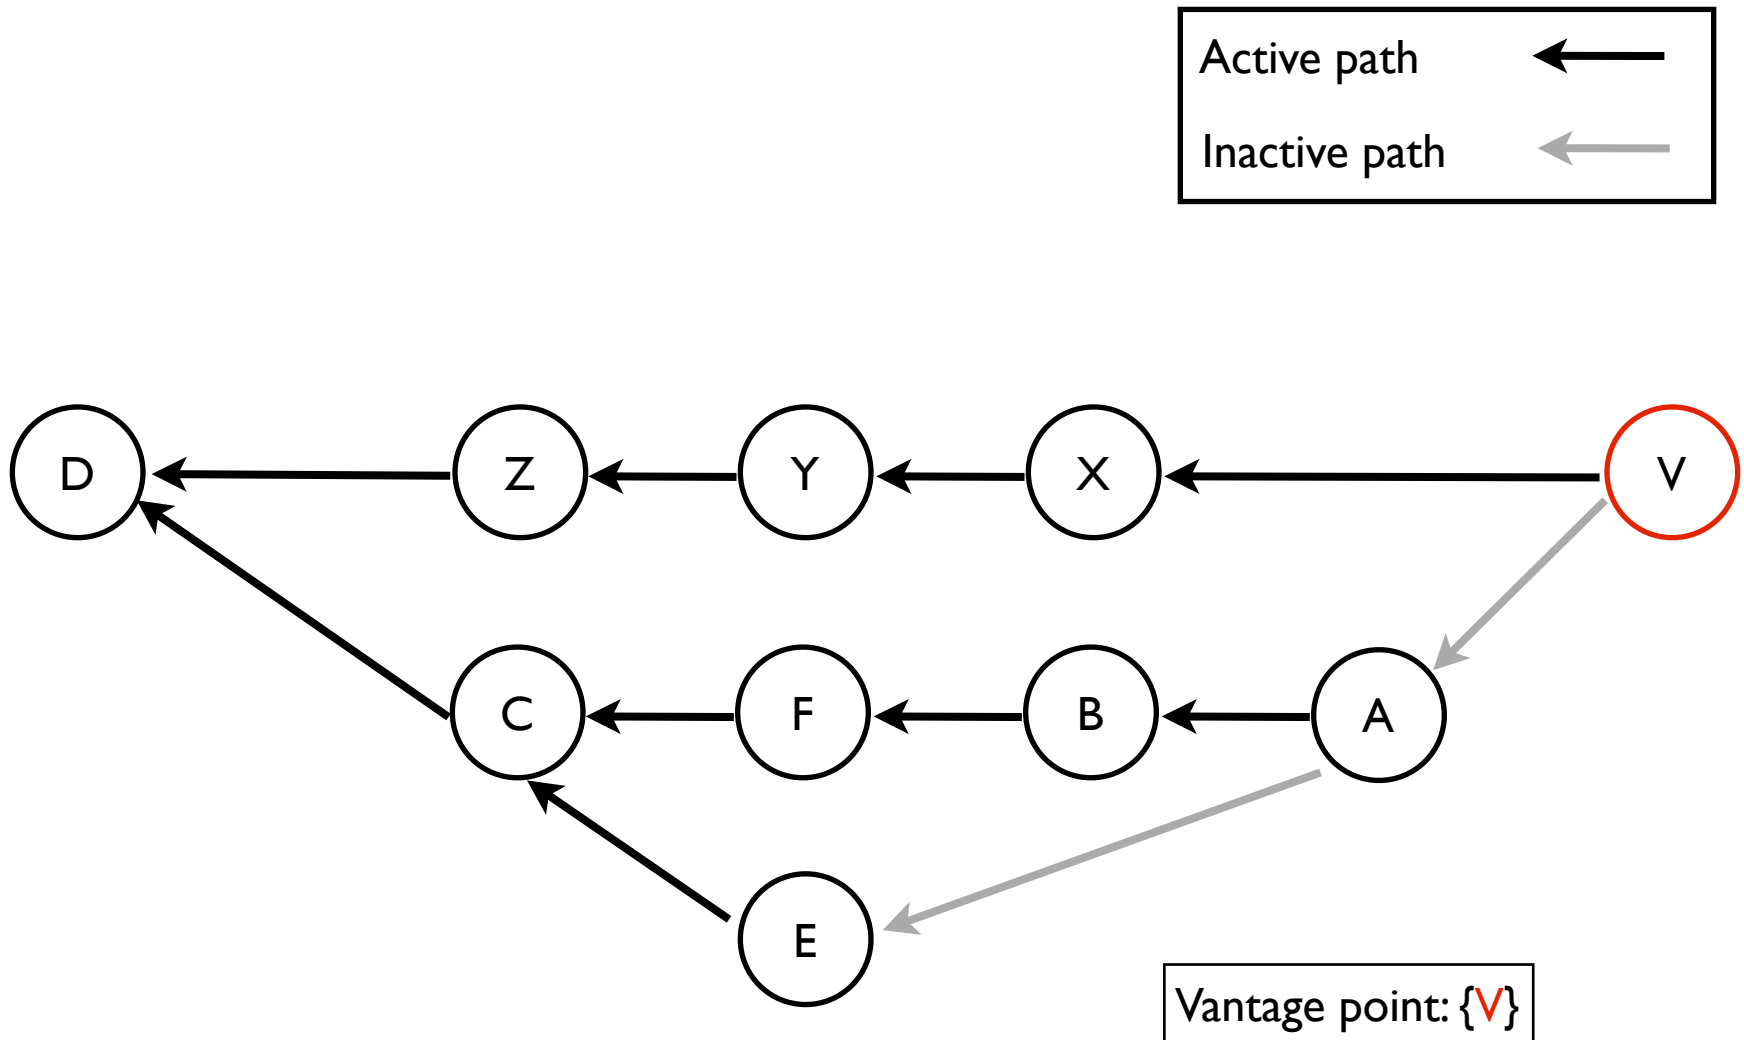
## ▶ NOOP algorithm

- ▶ Union of all old and new paths:  $\{Q, V, X, Y, Z\}$
- ▶ Intersection of all old and new paths:  $\{Q, V, X, Z\}$
- ▶ Root cause is  $U - \cap$ :  $\{Y\}$

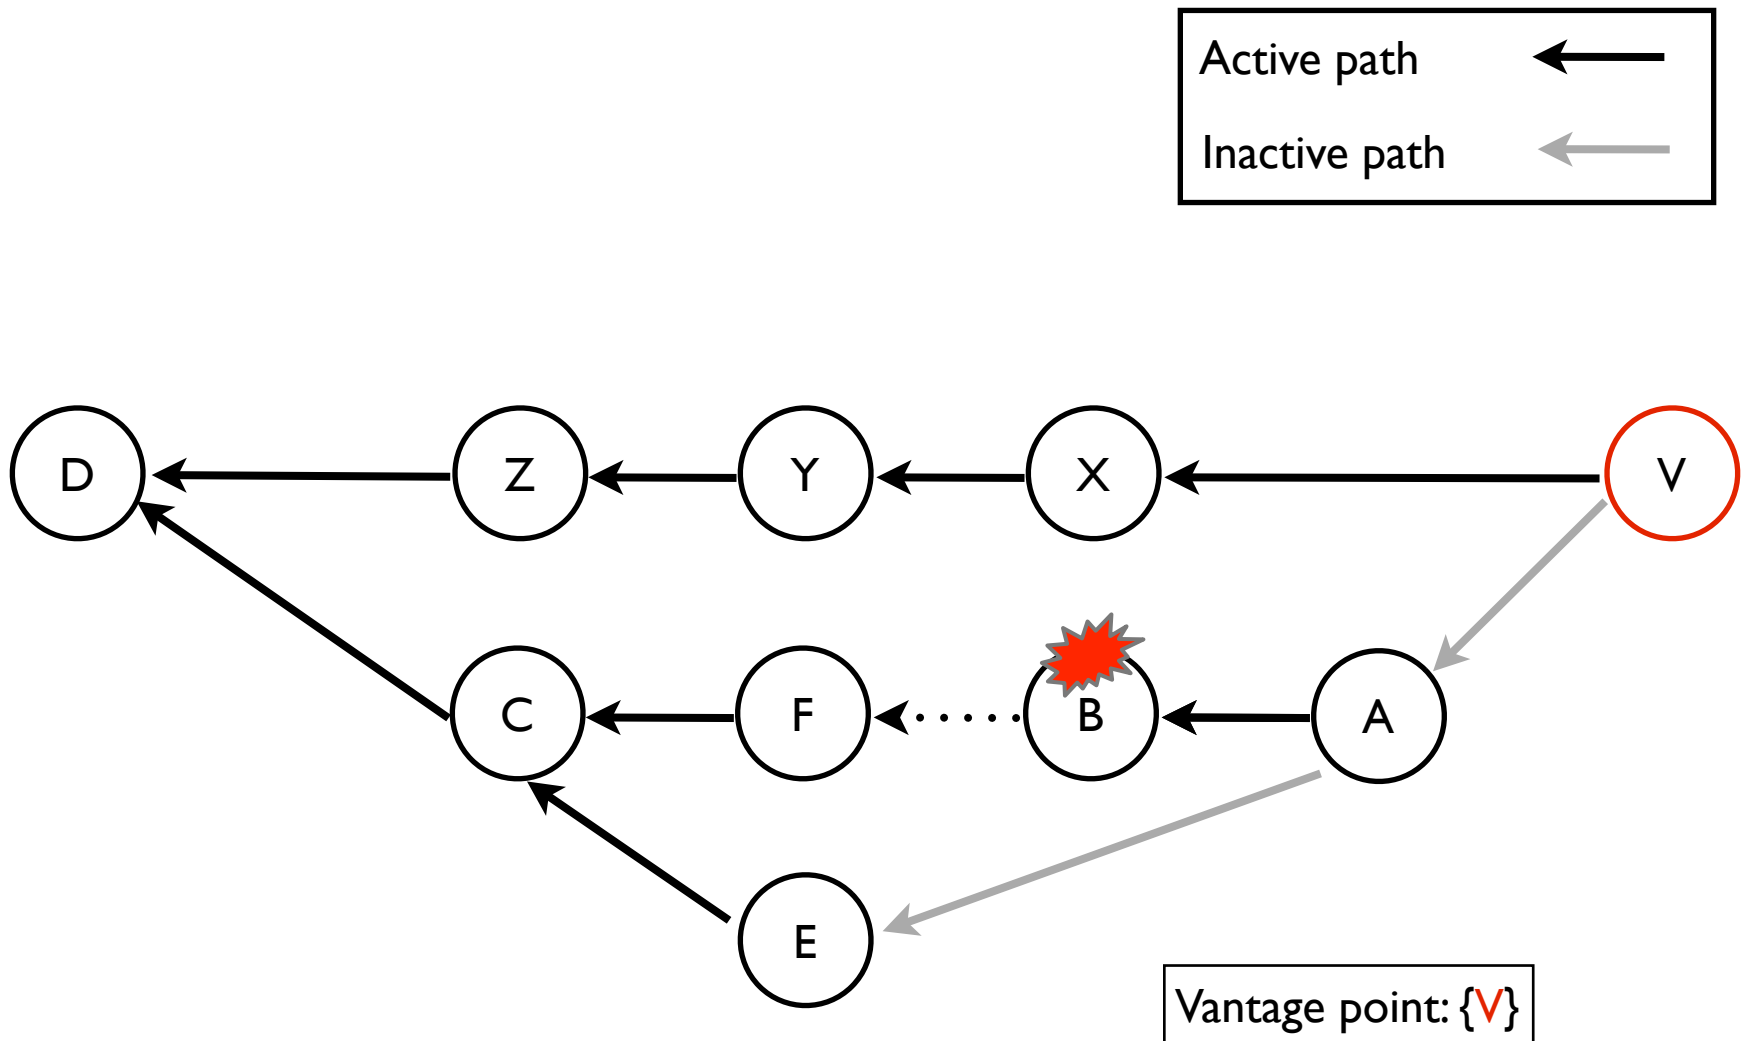


Vantage points:  $\{Q, V\}$

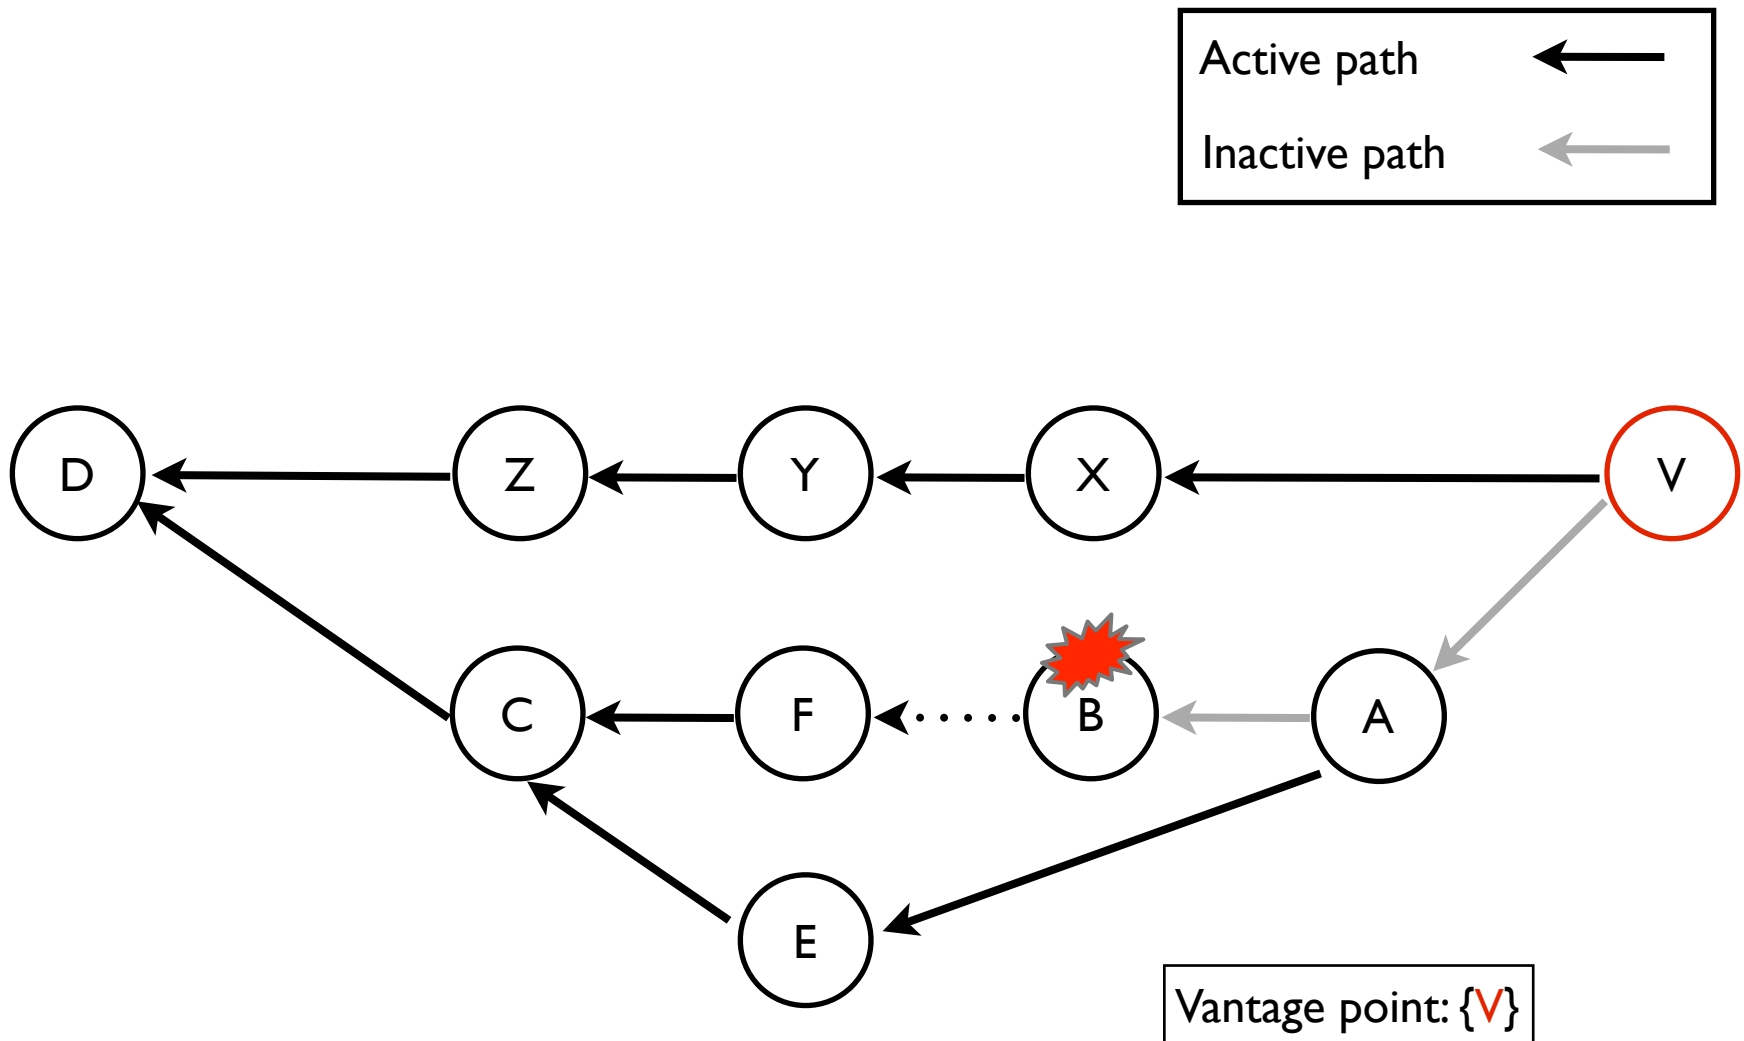
# The case of the induced path change



# The case of the induced path change

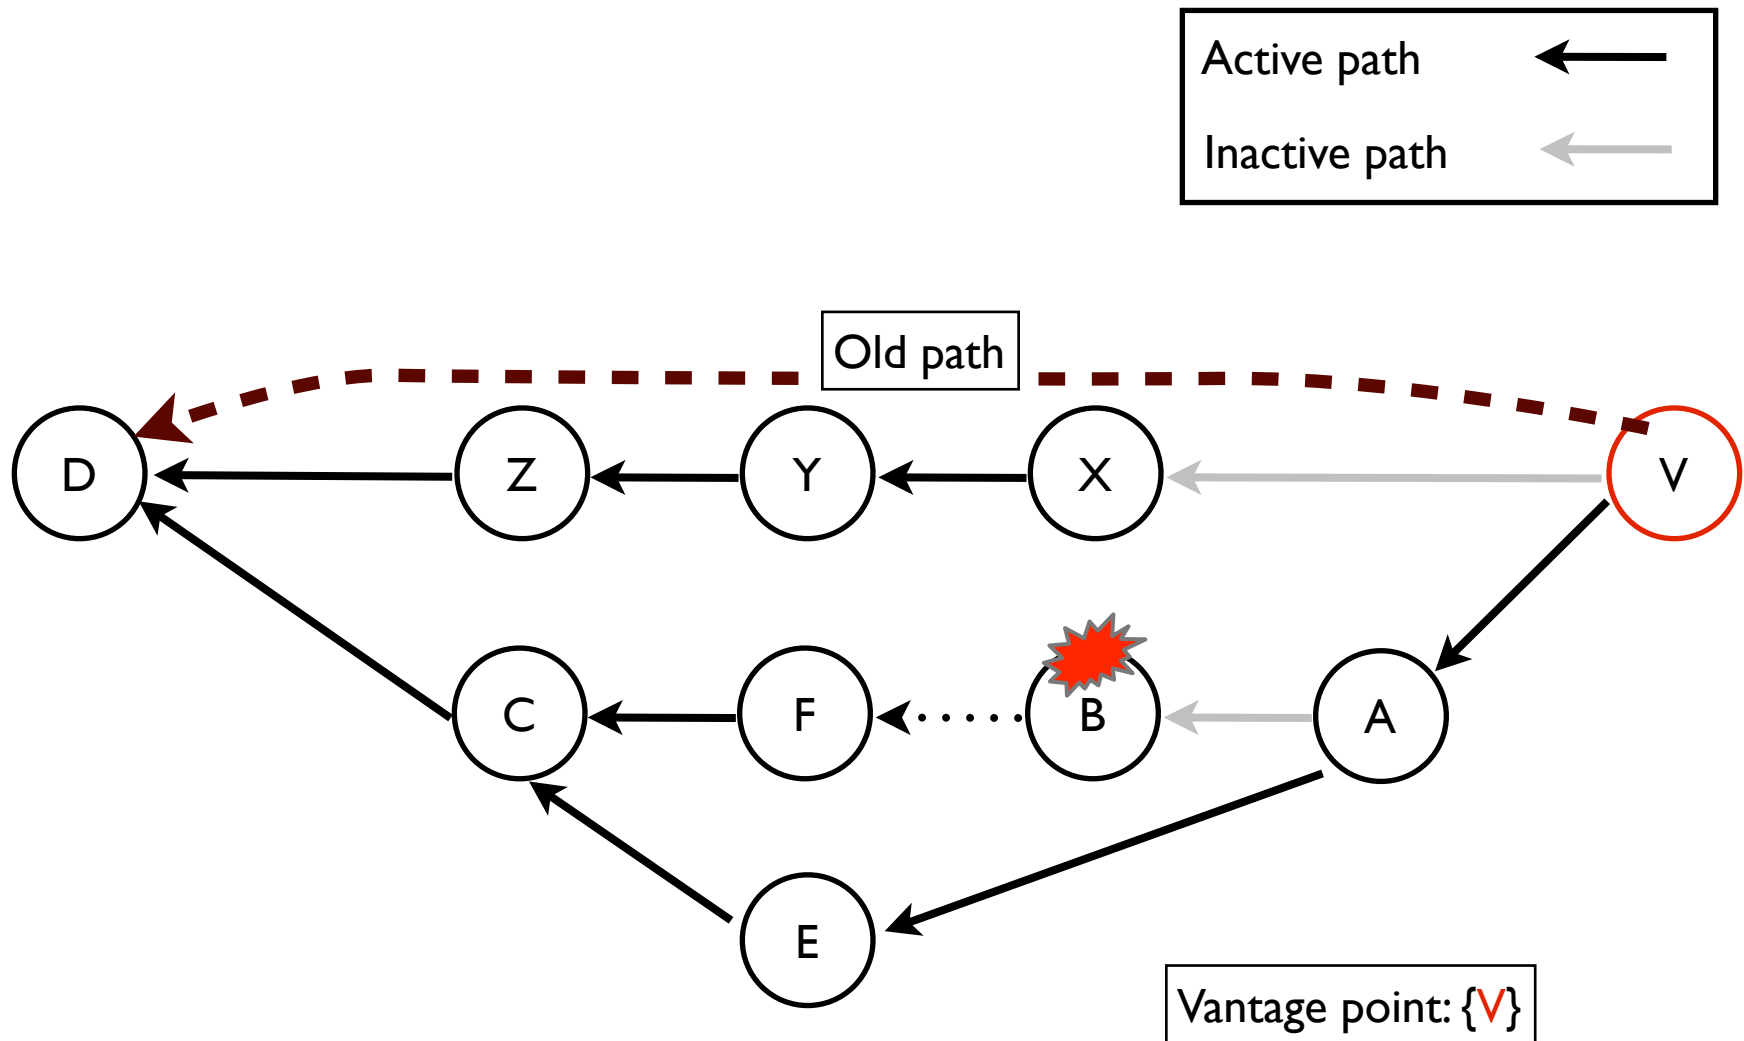


# The case of the induced path change






# The case of the induced path change




# The case of the induced path change



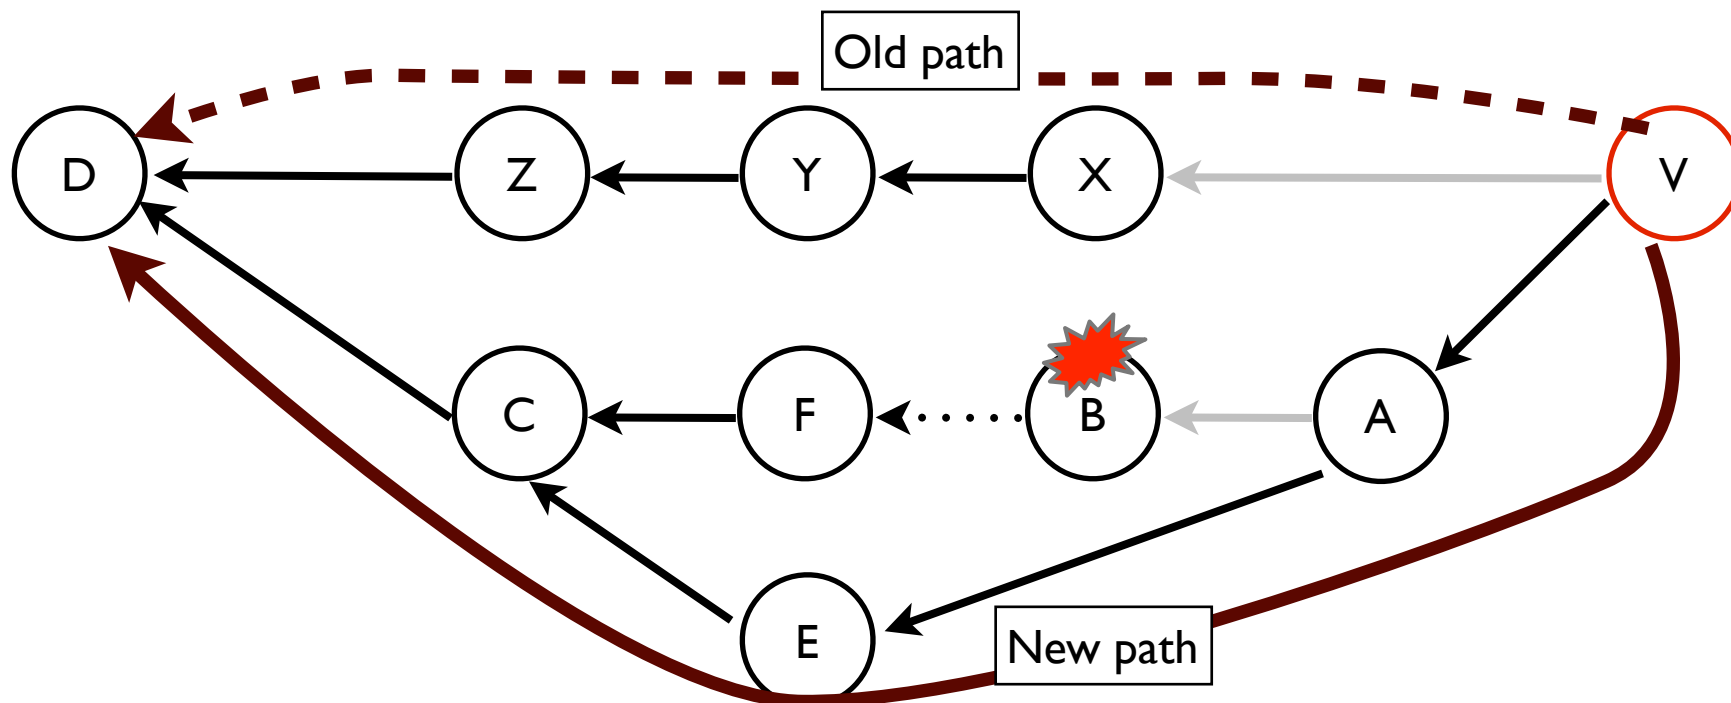
# The case of the induced path change

- ▶ Root cause (**B**) neither on old or new path!



# The case of the induced path change

- ▶ Root cause (**B**) neither on old or new path!
  - ▶ Need to revisit assumptions and model of path changes



# Outline

---

- ▶ Intro
- ▶ Motivation
- ▶ **Root Cause Isolation**
- ▶ **System & Evaluation**

# PoiRoot: Revisiting assumptions

---

- ▶ Theory of candidate ASes responsible for a change
  - ▶ Policy (i.e., localPref) trumps path length
  - ▶ Policy depends only on next-hop AS (Gill et al. @ NANOG)

# PoiRoot: Revisiting assumptions

---

- ▶ Theory of candidate ASes responsible for a change
  - ▶ Policy (i.e., localPref) trumps path length
  - ▶ Policy depends only on next-hop AS (Gill et al. @ NANOG)
- ▶ **NOOP** says ASes on old or new path from VP
  - ▶ We know this is wrong. Do we have to consider all ASes?
  - ▶ No, we prove the root cause can also be:
    - ▶ Any AS on old paths from ASes in new path from VP
    - ▶ Any AS on new paths from ASes in old path from VP

# PoiRoot: Revisiting assumptions

---

- ▶ Theory of candidate ASes responsible for a change
  - ▶ Policy (i.e., localPref) trumps path length
  - ▶ Policy depends only on next-hop AS (Gill et al. @ NANOG)
- ▶ **NOOP** says ASes on old or new path from VP
  - ▶ We know this is wrong. Do we have to consider all ASes?
  - ▶ No, we prove the root cause can also be:
    - ▶ Any AS on old paths from ASes in new path from VP
    - ▶ Any AS on new paths from ASes in old path from VP



# PoiRoot: Revisiting assumptions

---

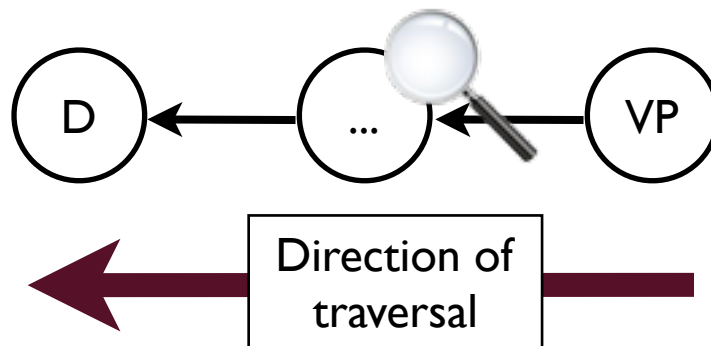
- ▶ Theory of candidate ASes responsible for a change
  - ▶ Policy (i.e., localPref) trumps path length
  - ▶ Policy depends only on next-hop AS (Gill et al. @ NANOG)
- ▶ NOOP says ASes on old or new path from VP
  - ▶ We know this is wrong. Do we have to consider all ASes?
  - ▶ No, we prove the root cause can also be:
    - ▶ Any AS on old paths from ASes in new path from VP
    - ▶ Any AS on new paths from ASes in old path from VP

*For now, let's assume I can obtain this.*

# PoiRoot: Theory in action

---


- ▶ Model tells us which networks *may* be root cause in general
- ▶ Algorithm to identify cause of *specific* change
  - ▶ Visit all ASes in candidate set
  - ▶ Move from VP toward D (DFS traversal)
  - ▶ Eliminate ASes that cannot the root cause



*Let's walk through an example*

# PoiRoot: Recursion


- ▶ Start at V
  - ▶ Check ASes on old path  $O(V)$
  - ▶ After reaching D, check ASes on new path  $N(V)$



# PoiRoot: Recursion

## Axiom I:

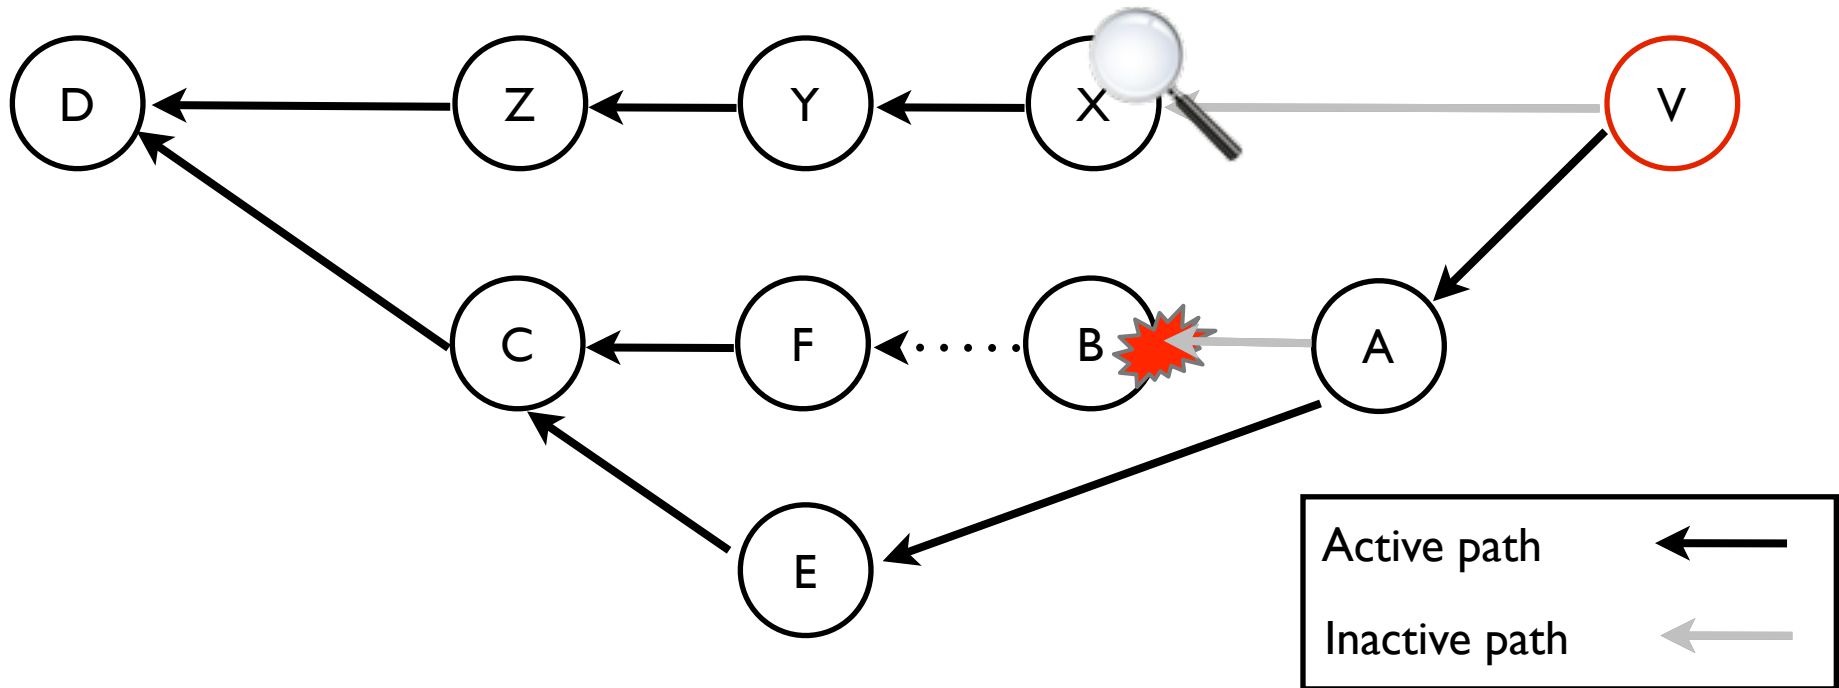
- ▶ If  $O(V) \neq N(V)$ ,  $V$  **may be** the root cause  
**but ASes upstream of  $V$  are not** the root cause



# PoiRoot: Recursion

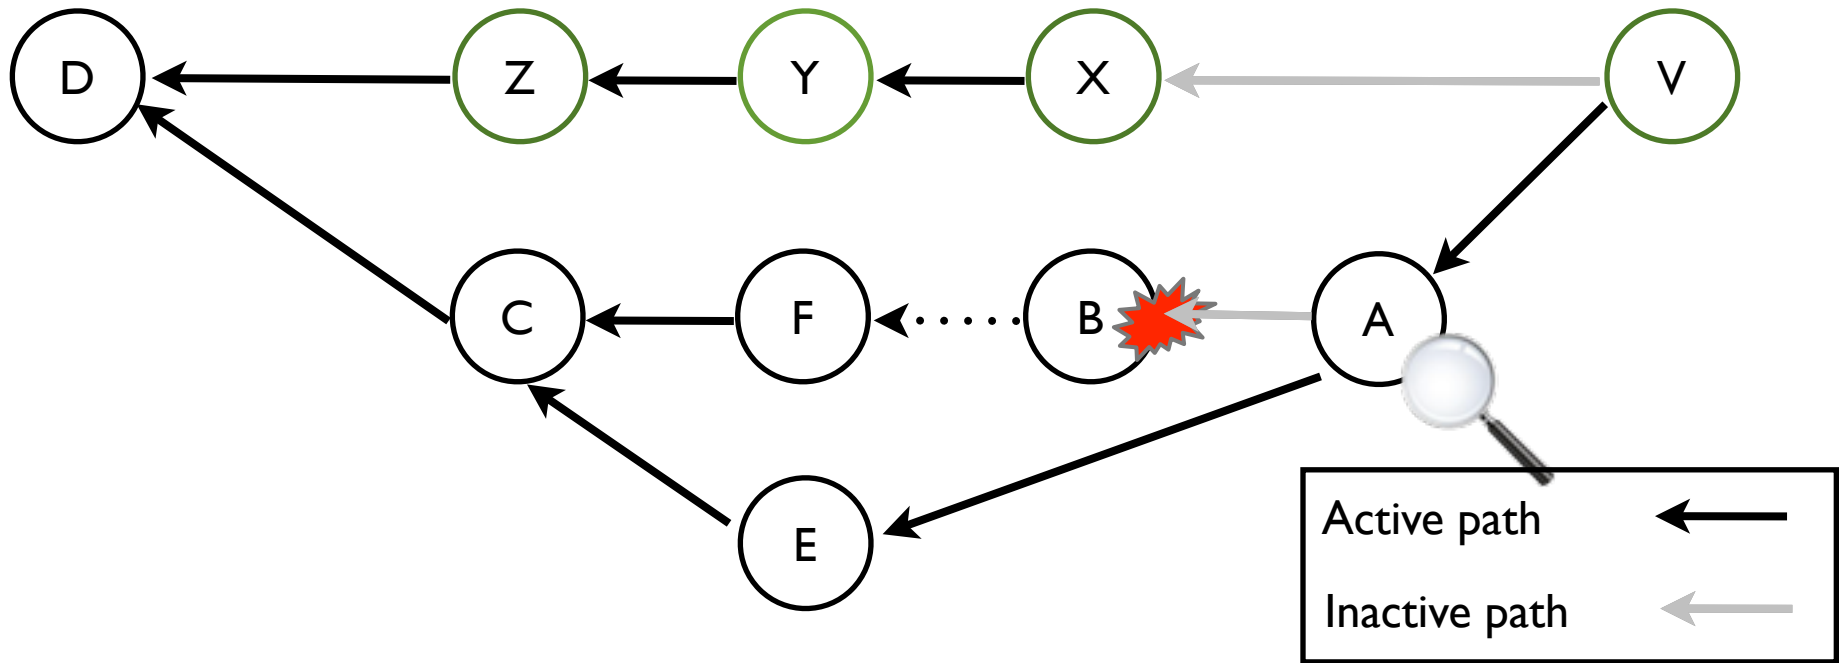
## Axiom 2:

- ▶ If  $O(X) == N(X)$ ,  $X$  is **not** the root cause




# PoiRoot: Recursion

- ▶  $O(A) \neq N(A)$ :  
A might be root cause, but V is not



# PoiRoot: Recursion

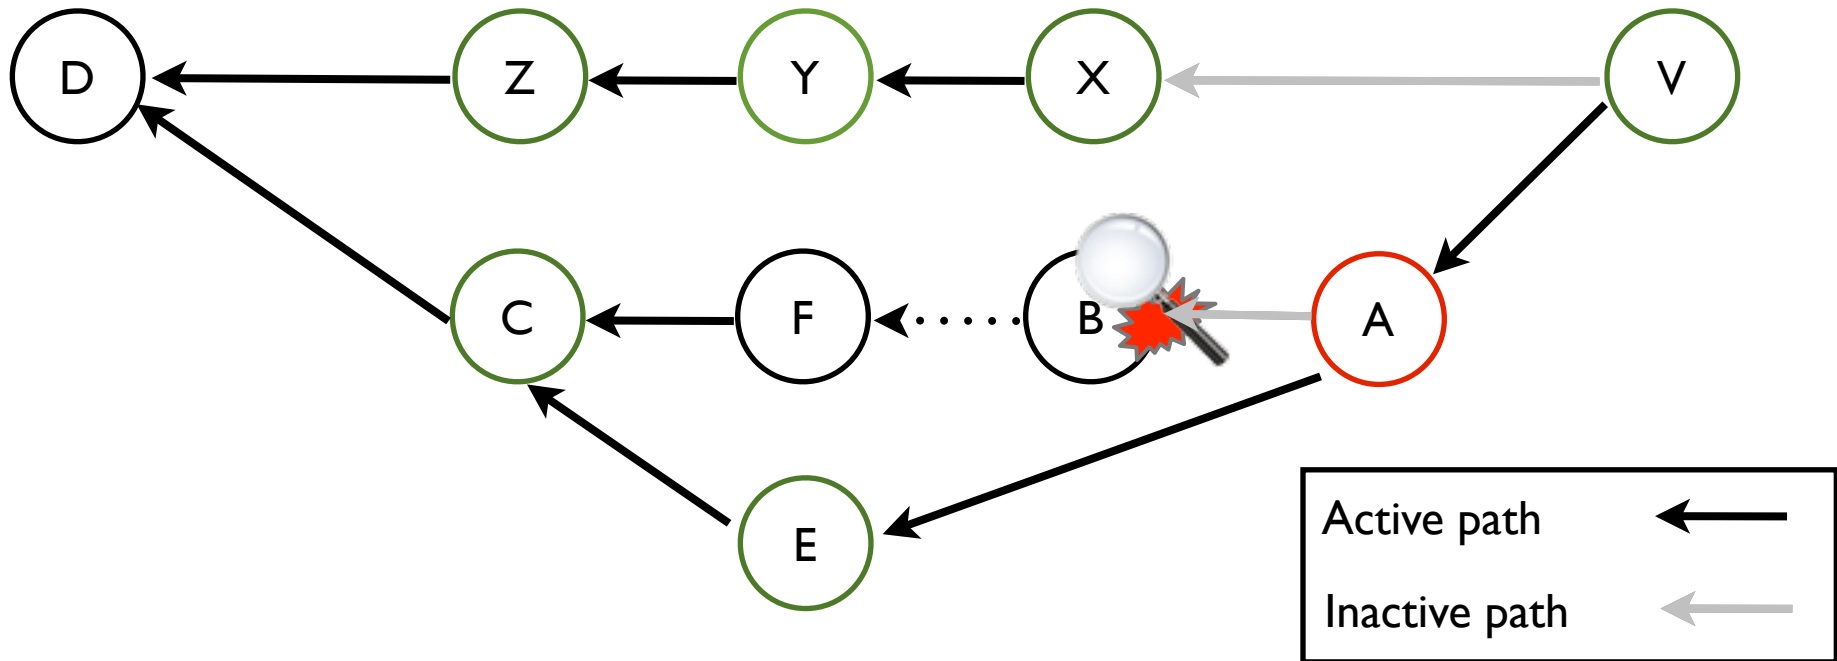
- ▶  $O(E) == N(E)$ :  
E is not the root cause



# PoiRoot: Recursion

Now check ASes on **old paths** from ASes in  $N(V)$

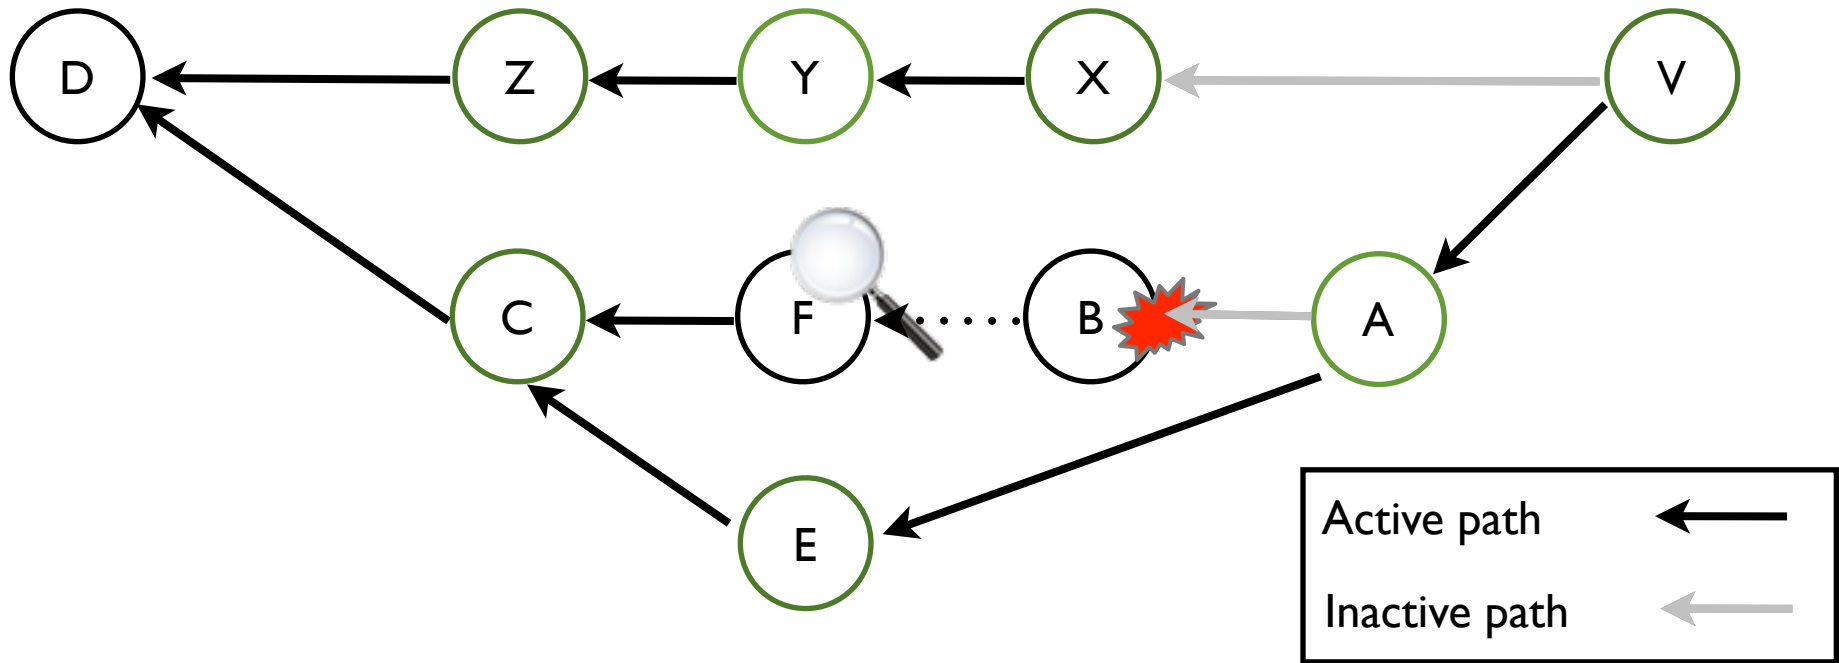
- ▶  $O(B) \neq N(B)$ :  
B may be the root cause





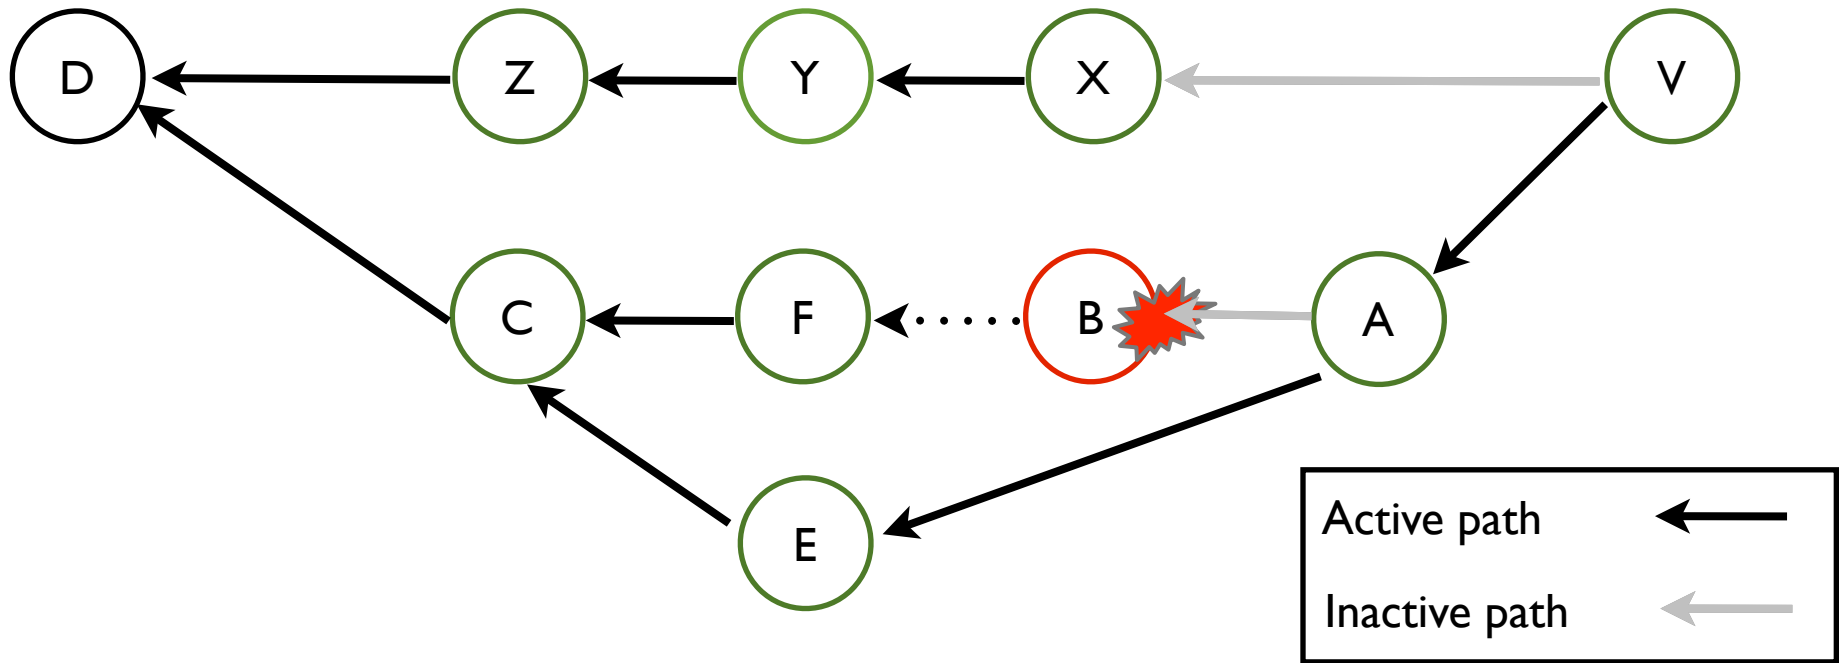
# PoiRoot: Recursion

- ▶  $O(F) == N(F)$ :  
F is not the root cause



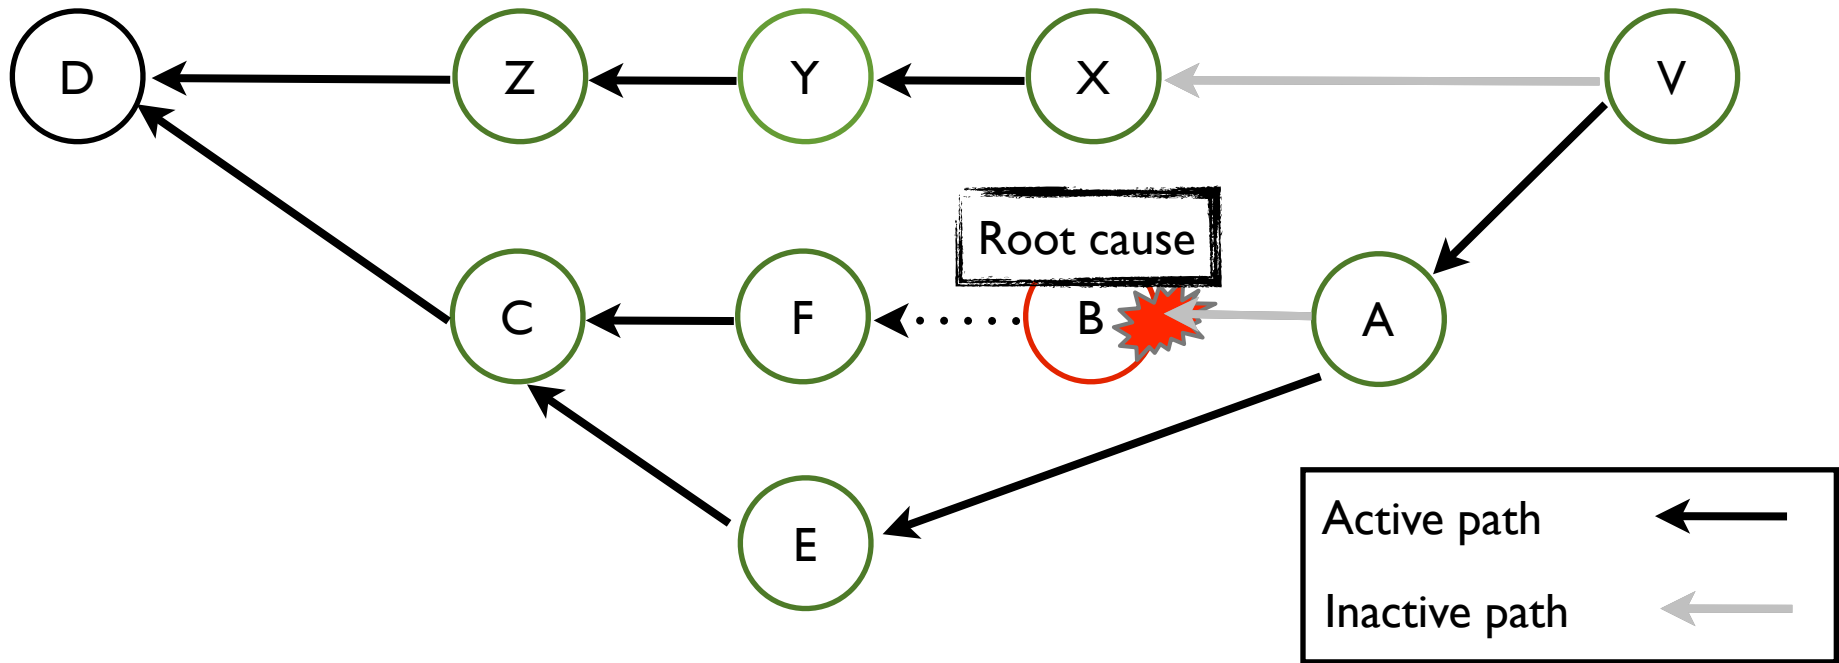
# PoiRoot: Recursion

- ▶ Only candidate AS left is B, the root cause



# PoiRoot: Recursion

- ▶ Only candidate AS left is B, the root cause





# Other considerations

---

- ▶ Requires measuring paths that *might* be used in the future
  - ▶ Use *poisoning* to explore less preferred paths *before* a change
  - ▶ Measure those paths (e.g., via reverse traceroutes)



# Poisoning on the BGP Links

- ▶ Normal case (no poisoning)





# Poisoning on the BGP Links

- ▶ Normal case (no poisoning)





# Poisoning on the BGP Links

- ▶ Normal case (no poisoning)



# Poisoning on the BGP Links



- ▶ Normal case (no poisoning)







# Poisoning on the BGP Links

- ▶ Poisoning Y to force X to explore less preferred paths





# Poisoning on the BGP Links

- ▶ Poisoning Y to force X to explore less preferred paths





# Poisoning on the BGP Links

- ▶ Poisoning Y to force X to explore less preferred paths





# Poisoning on the BGP Links

- ▶ Poisoning Y to force X to explore less preferred paths



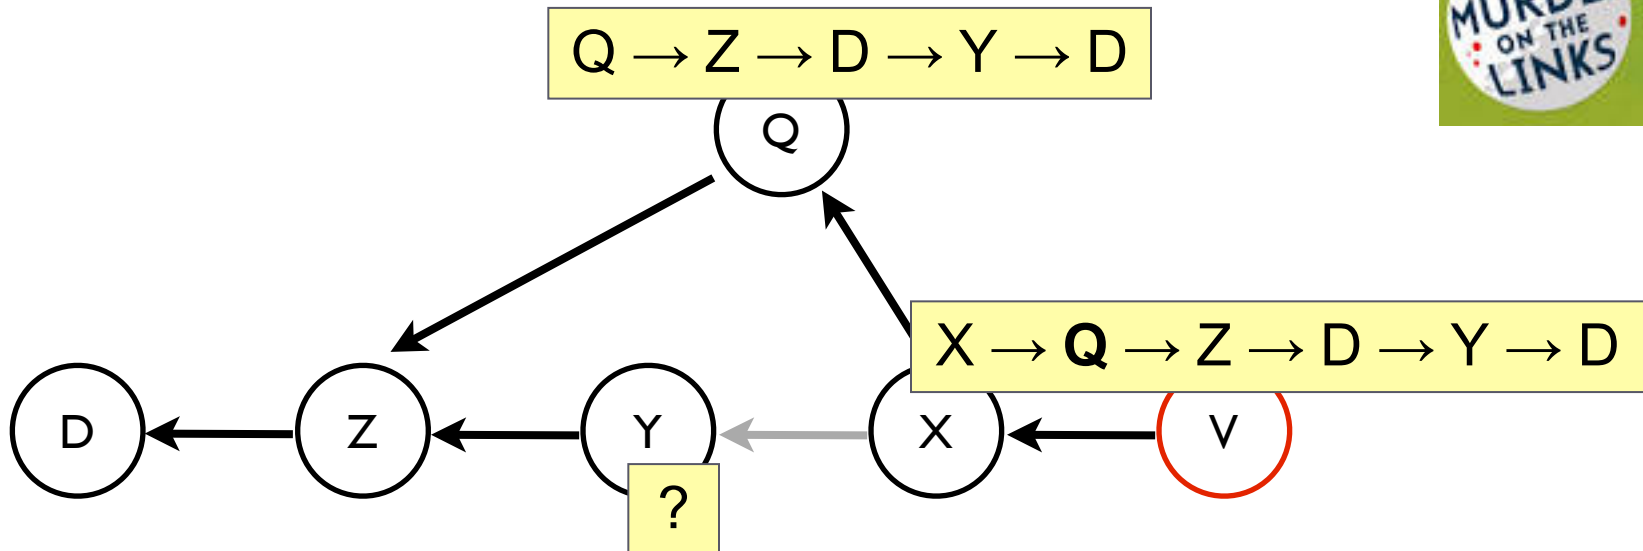

# Poisoning on the BGP Links

- ▶ Poisoning Y to force X to explore less preferred paths



# Poisoning on the BGP Links

- ▶ Poisoning Y to force X to explore less preferred paths



# Other considerations

---

- ▶ Requires measuring paths that *might* be used in the future
  - ▶ Use *poisoning* to explore less preferred paths *before* a change
  - ▶ Measure those paths (e.g., via traceroutes)
- ▶ Potential set of ASes to monitor can grow quite large
  - ▶ Consider next-less-preferred and all more preferred paths
- ▶ May not have all measurements
  - ▶ Use correlation to reduce candidate set

# Outline

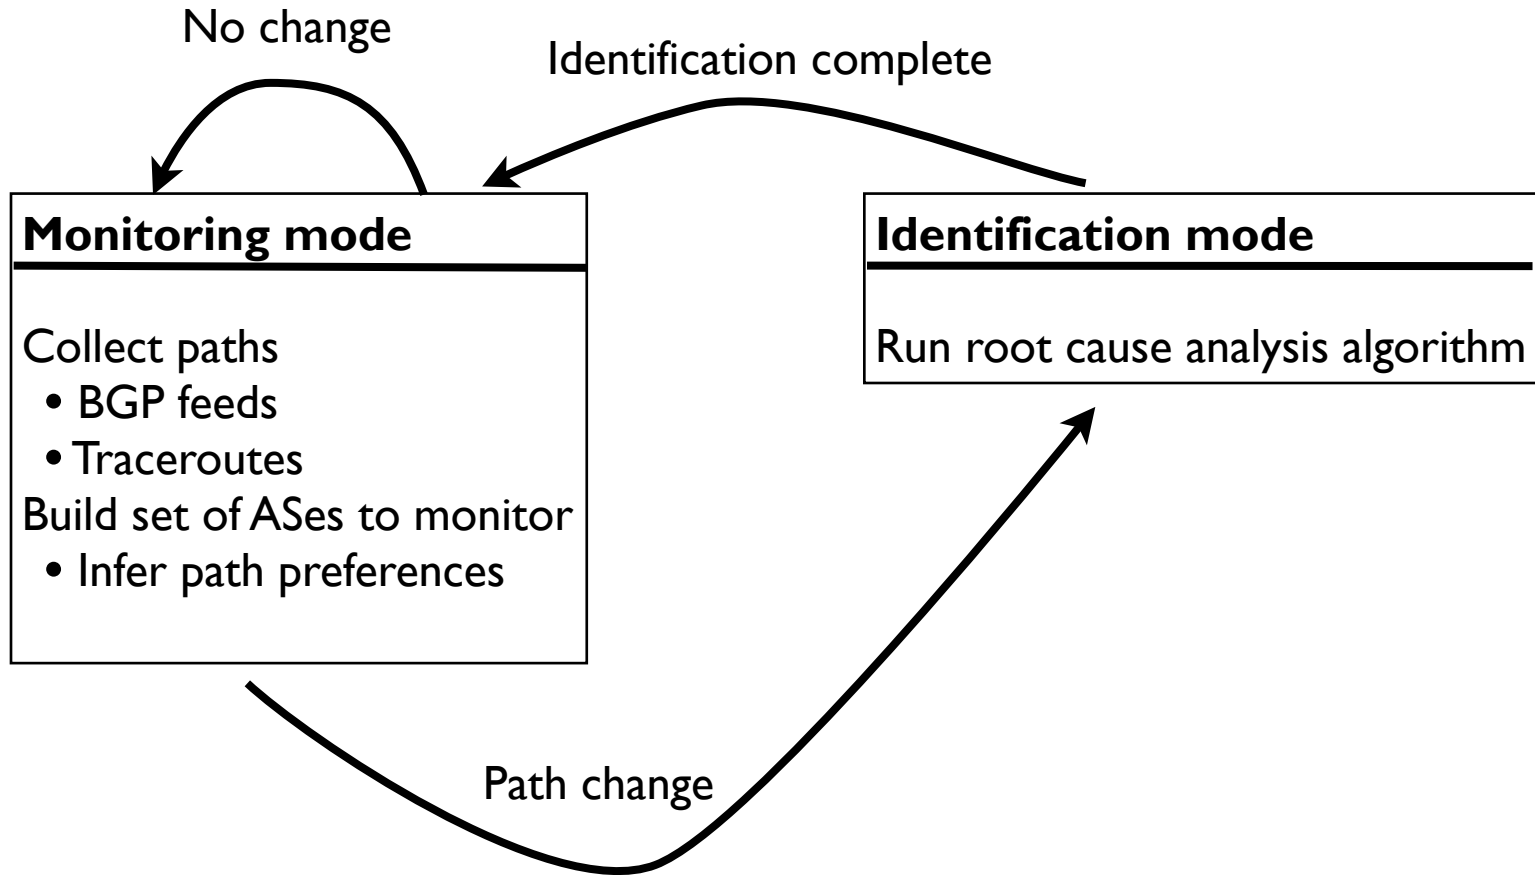
---

- ▶ Intro
- ▶ Motivation
- ▶ Root Cause Isolation
- ▶ **System & Evaluation**



# System overview

---



# Evaluation

---

- ▶ What we want to know
  - ▶ **Accuracy:** How often do I identify the root cause?
  - ▶ **Precision:** How large is the candidate set?
- ▶ Key challenge: **Ground truth!**
  - ▶ Requires an AS to experiment with, ...
  - ▶ ... known triggers for path changes, ...
  - ▶ and running code to identify the root cause during the change

# Evaluation

---

- ▶ What we want to know
  - ▶ **Accuracy:** How often do I identify the root cause?
  - ▶ **Precision:** How large is the candidate set?
- ▶ Key challenge: **Ground truth!**
  - ▶ Requires: an AS to experiment with, ...
    - ▶ Transit Portal: uses 5 universities as providers
  - ▶ ... known triggers for path changes, ...
  - ▶ and running code to identify the root cause during the change

# Evaluation

---

- ▶ What we want to know
  - ▶ **Accuracy:** How often do I identify the root cause?
  - ▶ **Precision:** How large is the candidate set?
- ▶ Key challenge: **Ground truth!**
  - ▶ Requires: an AS to experiment with, ...
    - ▶ Transit Portal: uses 5 universities as providers
  - ▶ ... known triggers for path changes, ...
    - ▶ Our good friend BGP path poisoning
  - ▶ and running code to identify the root cause during the change

# Evaluation

---

- ▶ What we want to know
  - ▶ **Accuracy:** How often do I identify the root cause?
  - ▶ **Precision:** How large is the candidate set?
- ▶ Key challenge: **Ground truth!**
  - ▶ Requires: an AS to experiment with, ...
    - ▶ Transit Portal: uses 5 universities as providers
  - ▶ ... known triggers for path changes, ...
    - ▶ Our good friend BGP path poisoning
  - ▶ and running code to identify the root cause during the change
    - ▶ Continuous path measurements from a large set of VPs


# Accuracy and precision

---

- ▶ **Accuracy: 100%**
  - ▶ PoiRoot never excludes an AS that is the root cause
  - ▶ NOOP does exclude when doing correlation (38% incorrect)


# Accuracy and precision

- ▶ **Accuracy: 100%**
  - ▶ PoiRoot never excludes an AS that is the root cause
  - ▶ NOOP does exclude when doing correlation (38% incorrect)
- ▶ **Precision: Mean suspect set size is 1.66**



# Accuracy and precision

- ▶ Accuracy: 100%
  - ▶ PoiRoot never excludes an AS that is the root cause
  - ▶ NOOP does exclude when doing correlation (38% incorrect)
- ▶ Precision: Mean suspect set size is 1.66





# Conclusion

---

- ▶ **Identify network triggering an interdomain path change**
  - ▶ Both accurate and precise
  - ▶ Handles arbitrary changes
- ▶ **New model to reason about path changes**
  - ▶ Relies on few, generally valid, assumptions about routing
- ▶ **Evaluation using controlled experiments on Internet paths**

# Questions?

---

