




Improving AS Relationship Inference Using PoPs

Lior Neudorfer Yuval Shavitt Noa Zilberman
School of Electrical Engineering

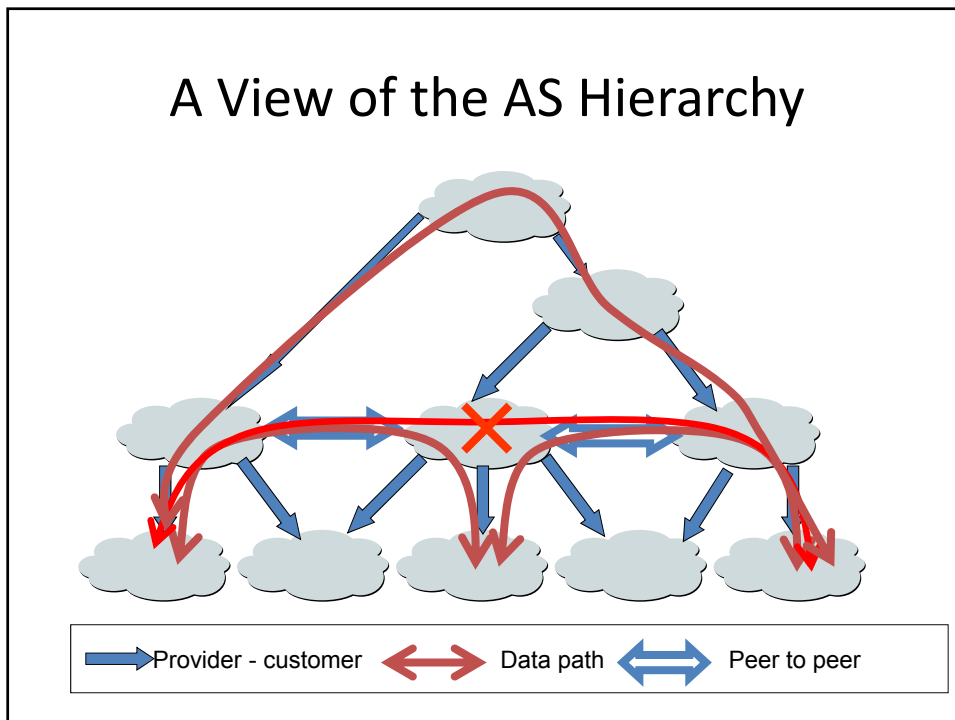
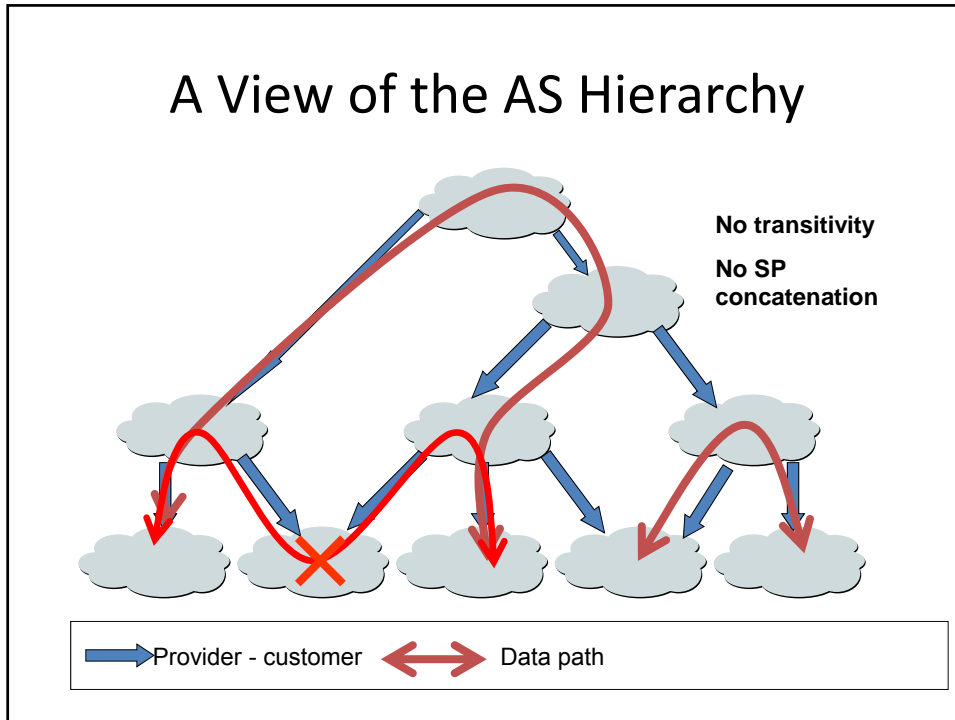


shavitt@eng.tau.ac.il
<http://www.eng.tau.ac.il/~shavitt>

A View of the AS Hierarchy



➔ Provider - customer



AS ToR Inference

[Gao 00] Introduced the problem

[Subramanian 02] formalized it as an optimization problem

[Dimitropoulos 06] a solution that maximizes the number of valley-free paths is not necessarily correct

[Shavitt 09] near deterministic solution

ToRs are Complex

[Gao 01] complex AS relationships exist

[Subramanian 02] introduced AS path anomalies classification

[Dimitropoulos 06] a survey with several large ISPs, revealed backup links and hybrid c2p/p2p relationships.

Our Contribution




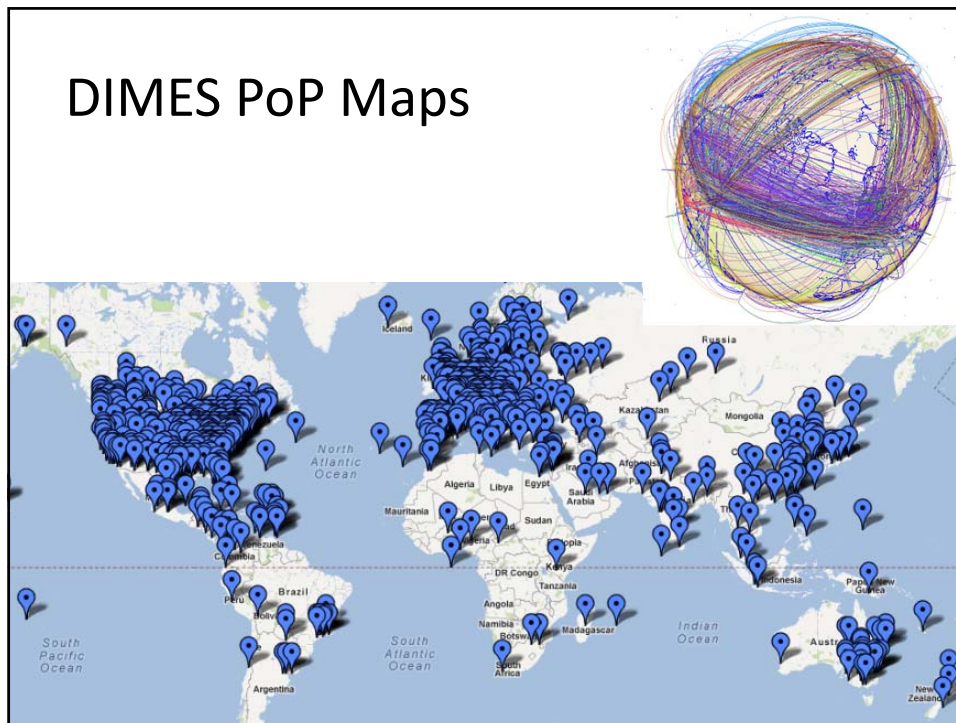
- Revealing complex ToRs
- Cleaning errors in ToR inference
- How?
 - Using PoP level analysis
- Why?
 - AS level cannot detect complex ToRs
 - AS level too gross to detect various errors
 - Router level: too detailed, hard to work with

Point of Presence (PoP)



- A concentration of an ISP's networking equipment in a single location
- DIMES generates periodical PoP maps






Process in a Nutshell


- Convert IP traceroutes to PoP level traceroute
- Add *known* AS ToRs to links
 - Augment with new inference
- Find all paths violating valley-freedom
 - Mark suspected violating PoP-PoP links
- Deduce: errors, complex ToRs, anomalous ToRs

IP to PoP Traceroute

- IXP links are short-jumped




Augmenting ToRs




Augmenting ToRs

deterministic augmentation




Augmenting ToRs




Augmenting ToRs

Non-deterministic augmentation:
Can be either P-P, P-C, or C-P




Finding PoP Level Anomalies

- Consider only paths with a single **violating link**
- **Anomalous candidates:**
 - PoP level links that can fix the single violation
- Every single **violating link** path can have 1 or 2 **anomalous candidates**



Finding PoP Level Anomalies

- Consider only paths with a single **violating link**
- **Anomalous candidates:**
 - PoP level links that can fix the single violation
- Every single **violating link** path can have 1 or 2 **anomalous candidates**



A PoP-PoP Link Anomaly

A PoP-PoP Link is anomalous if

- It appears in a few valid paths
- Changing its ToR will correct many paths
- Its measurements are significant
 - Diverse routes from many nodes

Data Used in Study



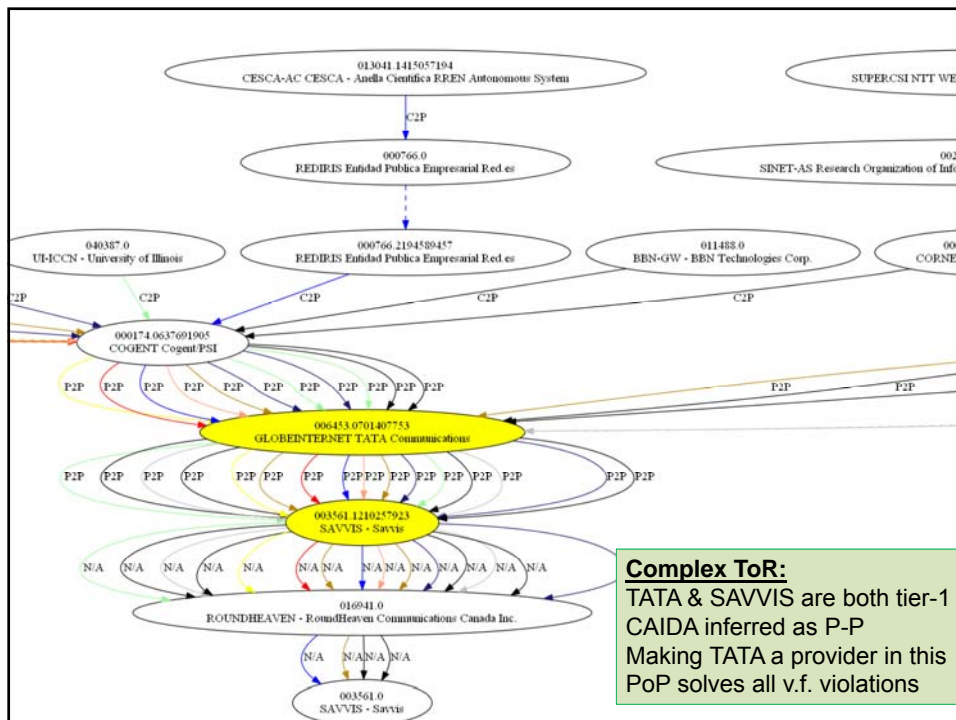
- DIMES dataset:
 - 29M traceroute measurements to 2.3M unique IP destinations
 - 5347 detected PoPs from 2636 different ASes
 - 1.6M unique AS/PoP paths after preprocessing
- CAIDA AS Relationships dataset
 - constructed from BGP routes and WHOIS queries
 - 119,924 known ToRs

ToR Augmentation

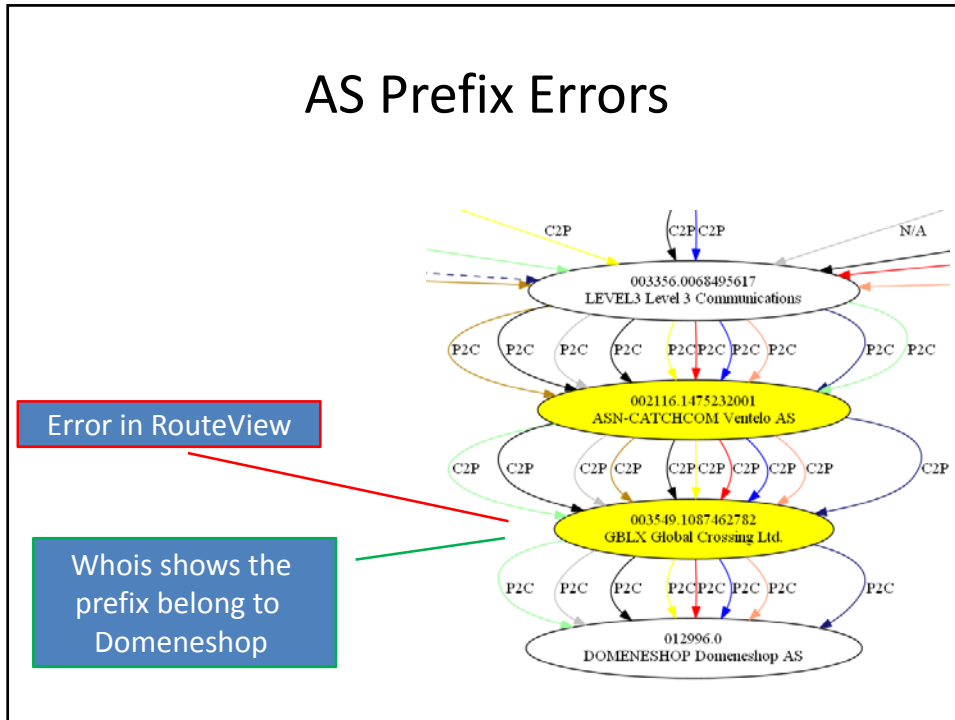
- 70,714 AS-AS links in DIMES measurements
 - out of them, CAIDA inferred ToRs for only 45,202 (64%)
- We augment ToRs assuming valley-free routing
- We define p as the minimum number of paths required to infer a ToR
 - With low confidence ($p = 1$), we infer 13601 additional AS ToRs (overall 83%)
 - With higher condence ($p = 5$), we infer 6699 additional AS ToRs (overall 73%)

Anomaly Detection

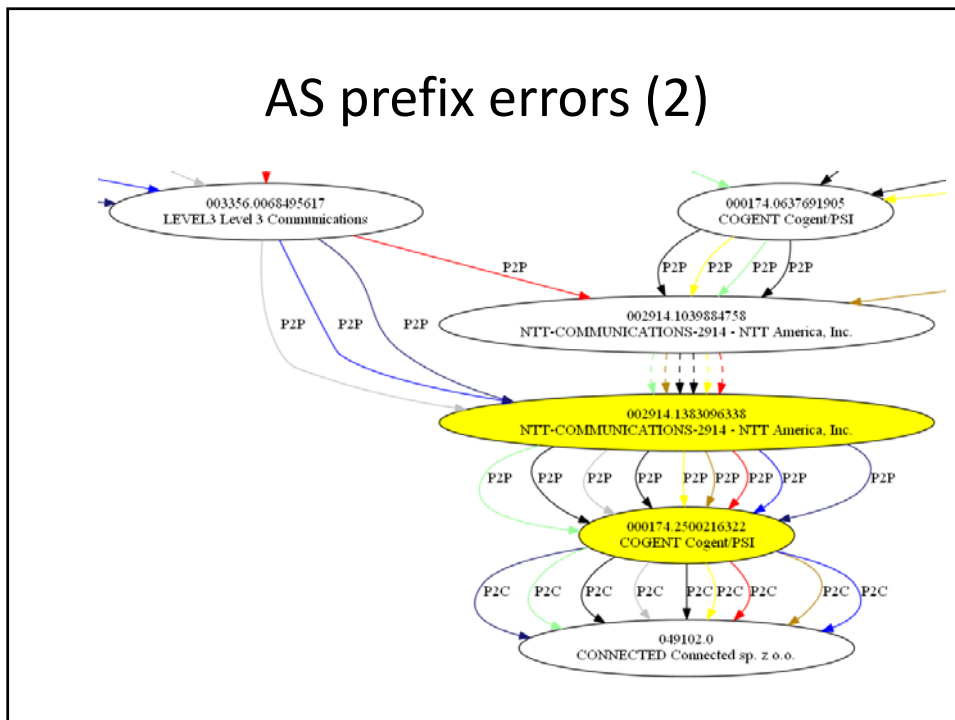
- The number of the anomalies was small enough to be manually examined
- Types of anomalies we've detected:
 - complex AS relationships (c2p/p2p)
 - wrong AS prefix resolutions
 - Internet eXchange Points (IXPs)
 - previously unknown sibling relationships
 - CAIDA ToR mistakes
 - exotic academic connectivity

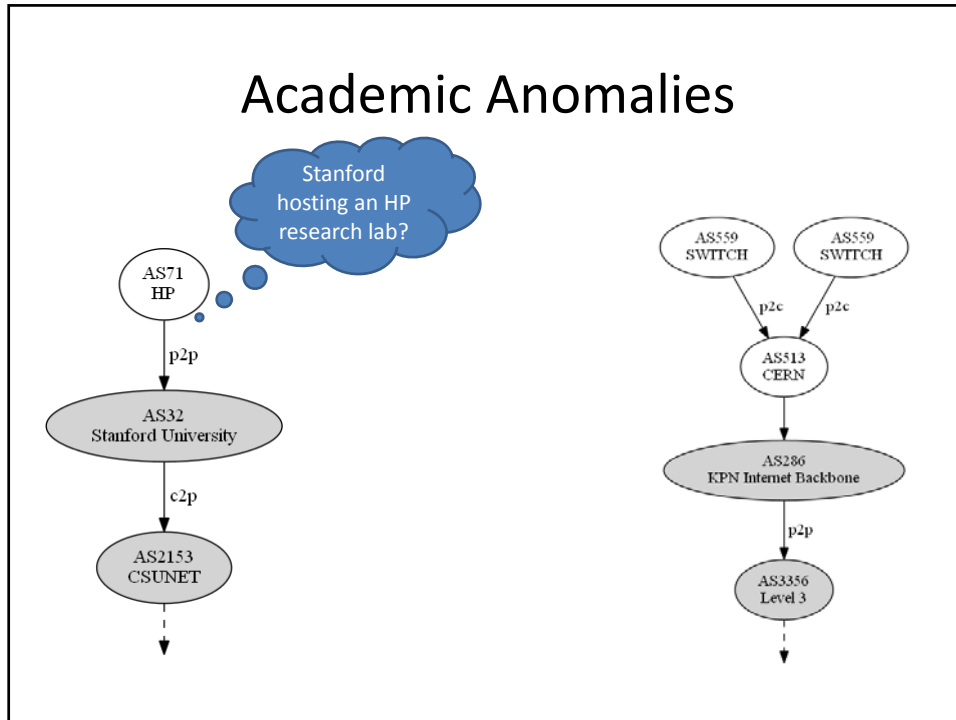


AS Prefix Errors



AS prefix errors (2)





Summary

- Using PoP level analysis we can identify a small number of suspect links
 - Can be manually checked for root cause
 - Can detect complex ToRs
 - Can find errors
 - ToR classifications
 - AP assignments
 - Traceroute errors
 - Can augment ToR classification
- Can be used with any initial ToR inference