Evaluating Potential Routing Diversity for Internet Failure Recovery

*Chengchen Hu, +Kai Chen, +Yan Chen, *Bin Liu

*Tsinghua University,

+Northwestern University

Internet Failures

- Failure is part of everyday life in IP networks
  - e.g., 675,000 excavation accidents in 2004 [Common Ground Alliance]
  - Network cable cuts every few days …
- Real-world emergencies or disasters can lead to substantial Internet disruption
  - Earthquakes
  - Storms
  - Terrorist incident
  - …

Example: Taiwan earthquake incident

- Large earthquakes hit south of Taiwan on 26 December 2006
- Only two of nine cross-sea cables **not** affected
- There were still abundant physical level connectivity there, but it took too long for ISPs to find them and use them.

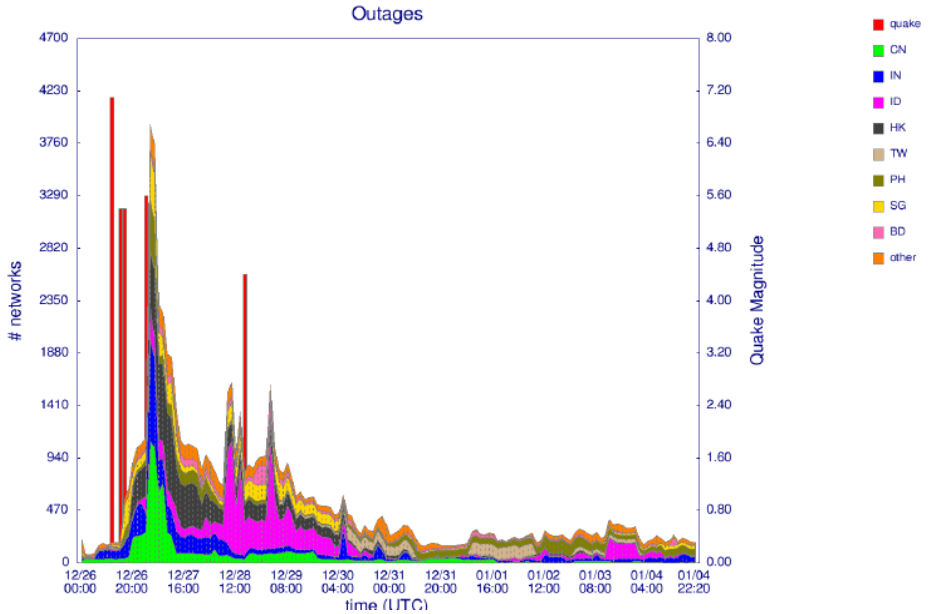

figures cited from "Aftershocks from the Taiwan Earthquakes: Shaking up Internet transit in Asia, NANOG42"

How reliable the Internet is?

- Internet is not as reliable as people expected! [Wu, CoNEXT'07]
  - 32% ASes are vulnerable to a single critical customer-provider link cut
  - 93.7% Tier-1 ISP's single-homed customers are lost from the peered ISP due to Tier-1 depeering

- Our question: can we find more resources to increase the Internet reliability especially when Internet emergency happens?

Roadmap

- Background
- Where are the potential resources?

Basic Idea

- Two places where we can find more routing diversities:
  - Internet eXchange Points (IXPs)
    - Co-location where multiple ASes exchange their traffic
    - Participant ASes in an IXP may not be connected via BGP
  - Internet valley-free routing policy
    - AS relationships: customer-provider, peering, sibling
    - Peering relaxation (PR): allow one AS to carry traffic from the other to its provider
    - Mentioned in [Wu, CoNEXT'07], but no evaluation
- Our main focus:
  - How much can we gain from these two potential resources, i.e., IXP and PR?

Roadmap

- Background
- Where are the potential resources?
- How much potential resources are there?

Dataset for Evaluation

- Most complete AS topology graph
  - BGP data
    - Route Views, RIPE/RIS, Abilene, CERNET BGP View
  - P2P traceroute
    - Traceroute data from 992, 000 IPs in over 3, 700 ASes
  - In total, 120K AS links with AS relationships
  - http://aqualab.cs.northwestern.edu/projects/SidewalkEnds.html [Chen et al, CoNEXT'09]
- IXP data
  - PCH + Peeringdb + Euro-IX (~200 IXPs)
  - 3468 participant ASes

Failure Models

- peering link teardown
  - Tier-1 depeering (Cogent and Level3 depeering)
- provider-customer link teardown
  - Several breakdowns of Tier-1 provider-customer links
- Mixed types of link breakdown
  - Large victims in Taiwan earthquakes

More results please check our report
"Evaluating Potential Routing Diversity for Internet Failure Recovery"
@ http://s-router.cs.tsinghua.edu.cn/pub/IER_report.pdf

Evaluation Metrics

- Recovery Ratio
  - # of recovered <src-dst> AS pairs versus total # of affected <src-dst> AS pairs
- Path Diversity
  - # of increased link-disjoint AS paths between affected <src-dst> AS pairs
- Shifted Path
  - # of link-disjoint AS paths shifted onto a normal link after we use IXP or PR resources

Results: Tier-1 Depeering

- 36 experiments for 9 Tier-1 ASes
- Recovery ratio: most of the lost AS pairs can be recovered

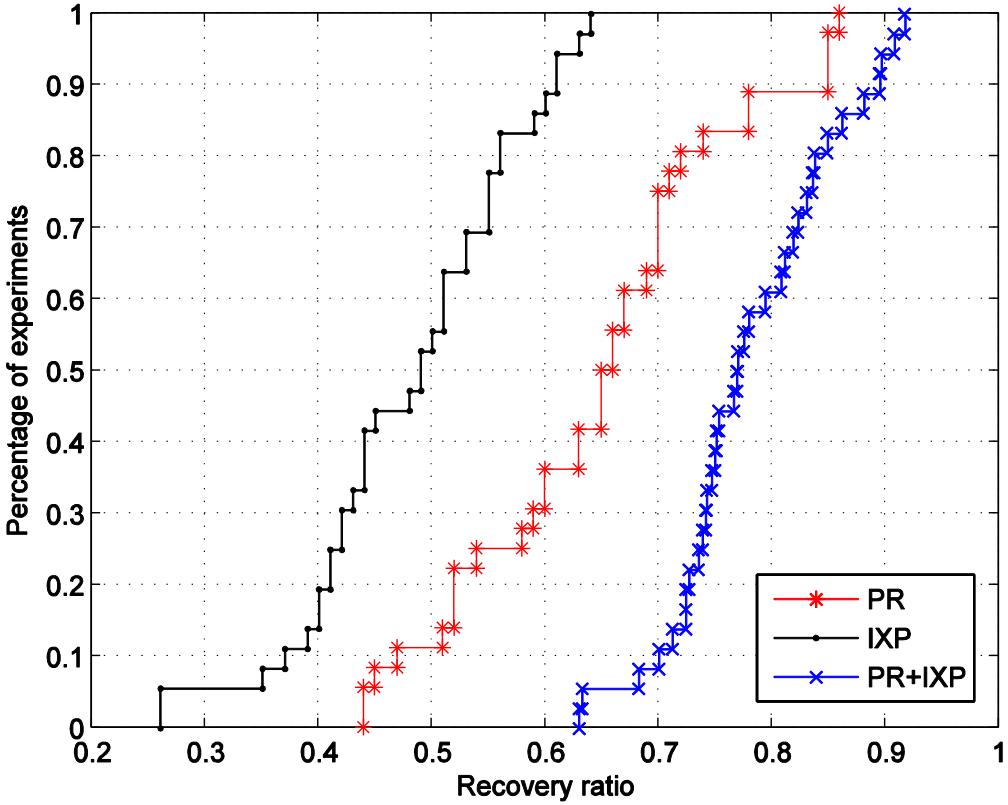

Results: Tier-1 Depeering

- Path diversity: multiple AS paths between lost AS pairs

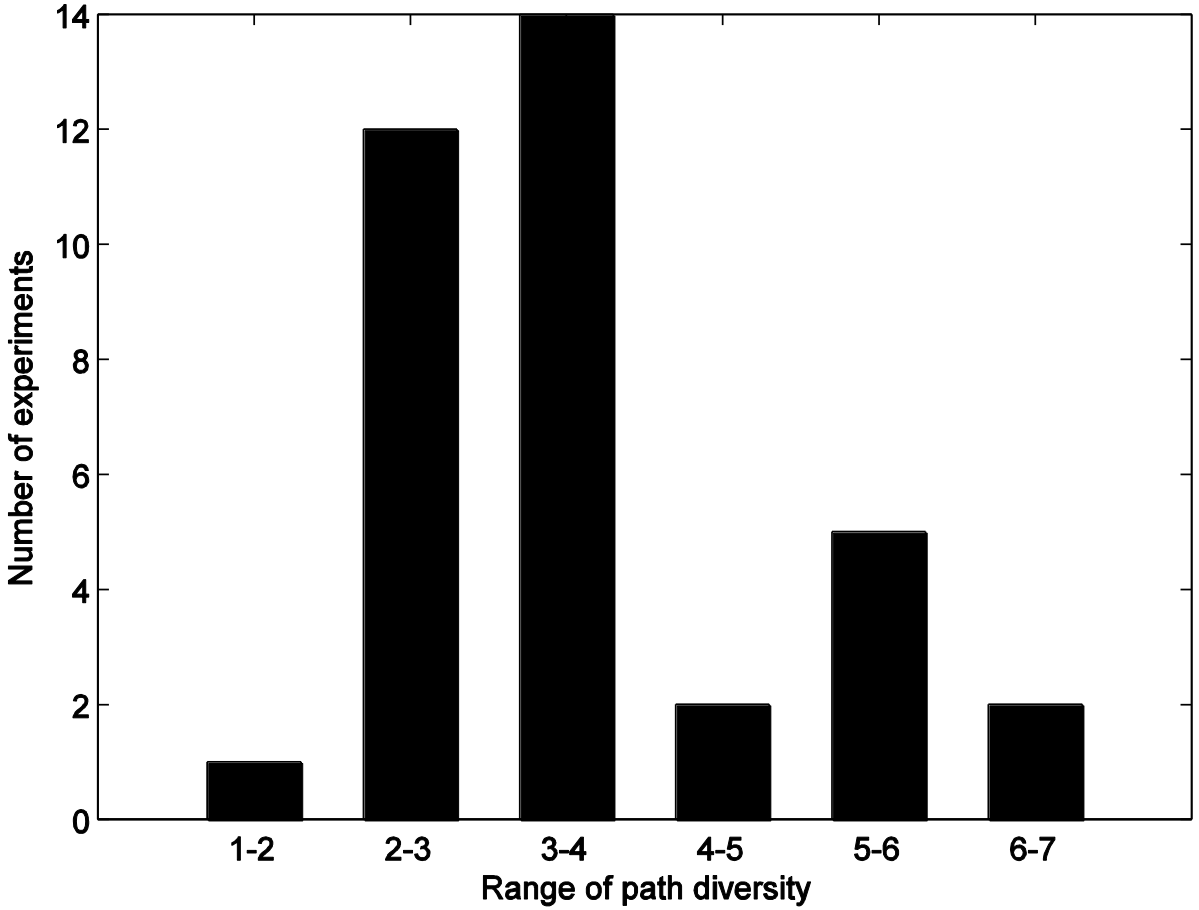

Results: Tier-1 Depeering

- Shifted path
  - On average, 3.75 ~ 17.2 for all 36 experiments
  - Moderate traffic load shifted onto the unaffected links

Roadmap

- Background
- Where are the potential resources?
- How much potential resources are there?
- How to use the potential resources?

Economic model

- B pays to A for recovery

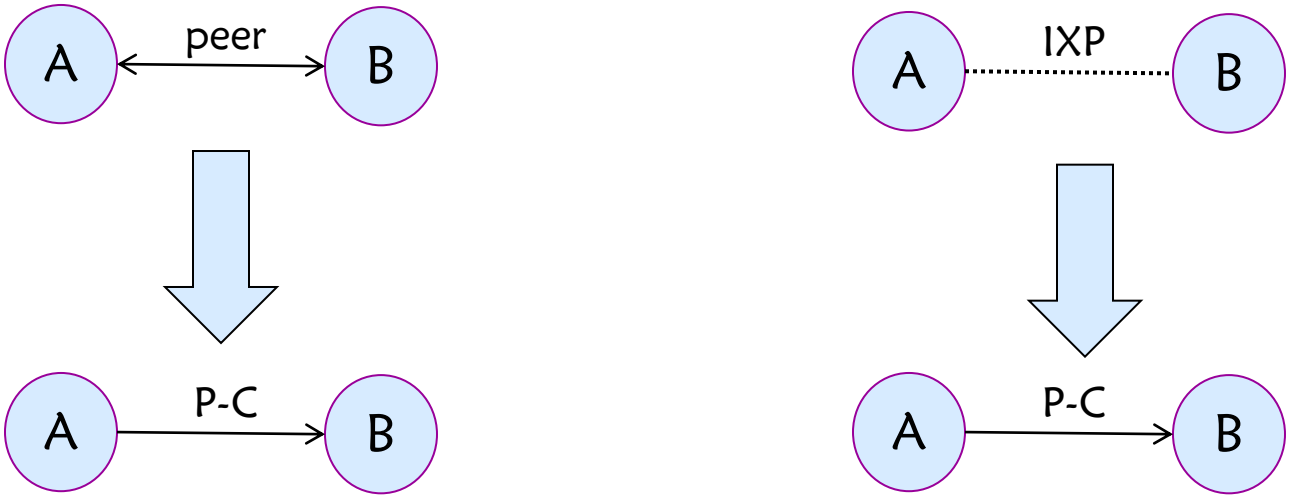

- Risk alliance (like airlines): price is determined beforehand
- pay on bandwidth & duration or bits (95 percentile)

Communication channel

- For peers
  - Have direct connections to peers
- For co-located ASes in the same IXP
  - ASes are connected by switches in modern IXPs
  - Messages are broadcasted via switches
  - Message confidentiality through public key crypto

Automatic communications: query phase

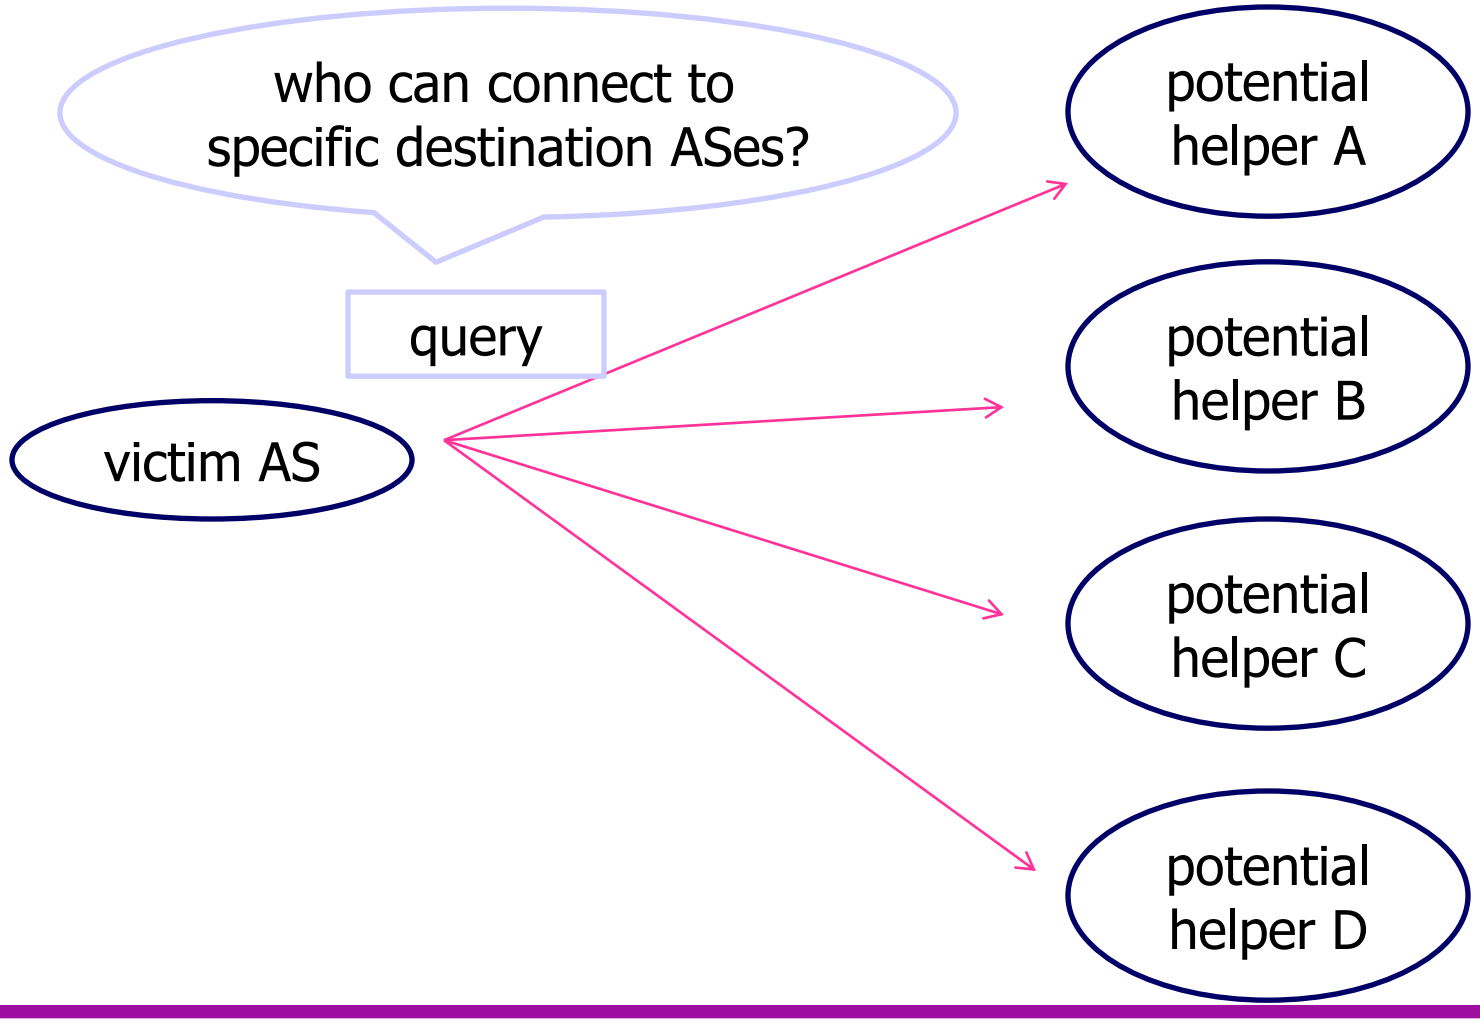

Automatic communications: Check availability

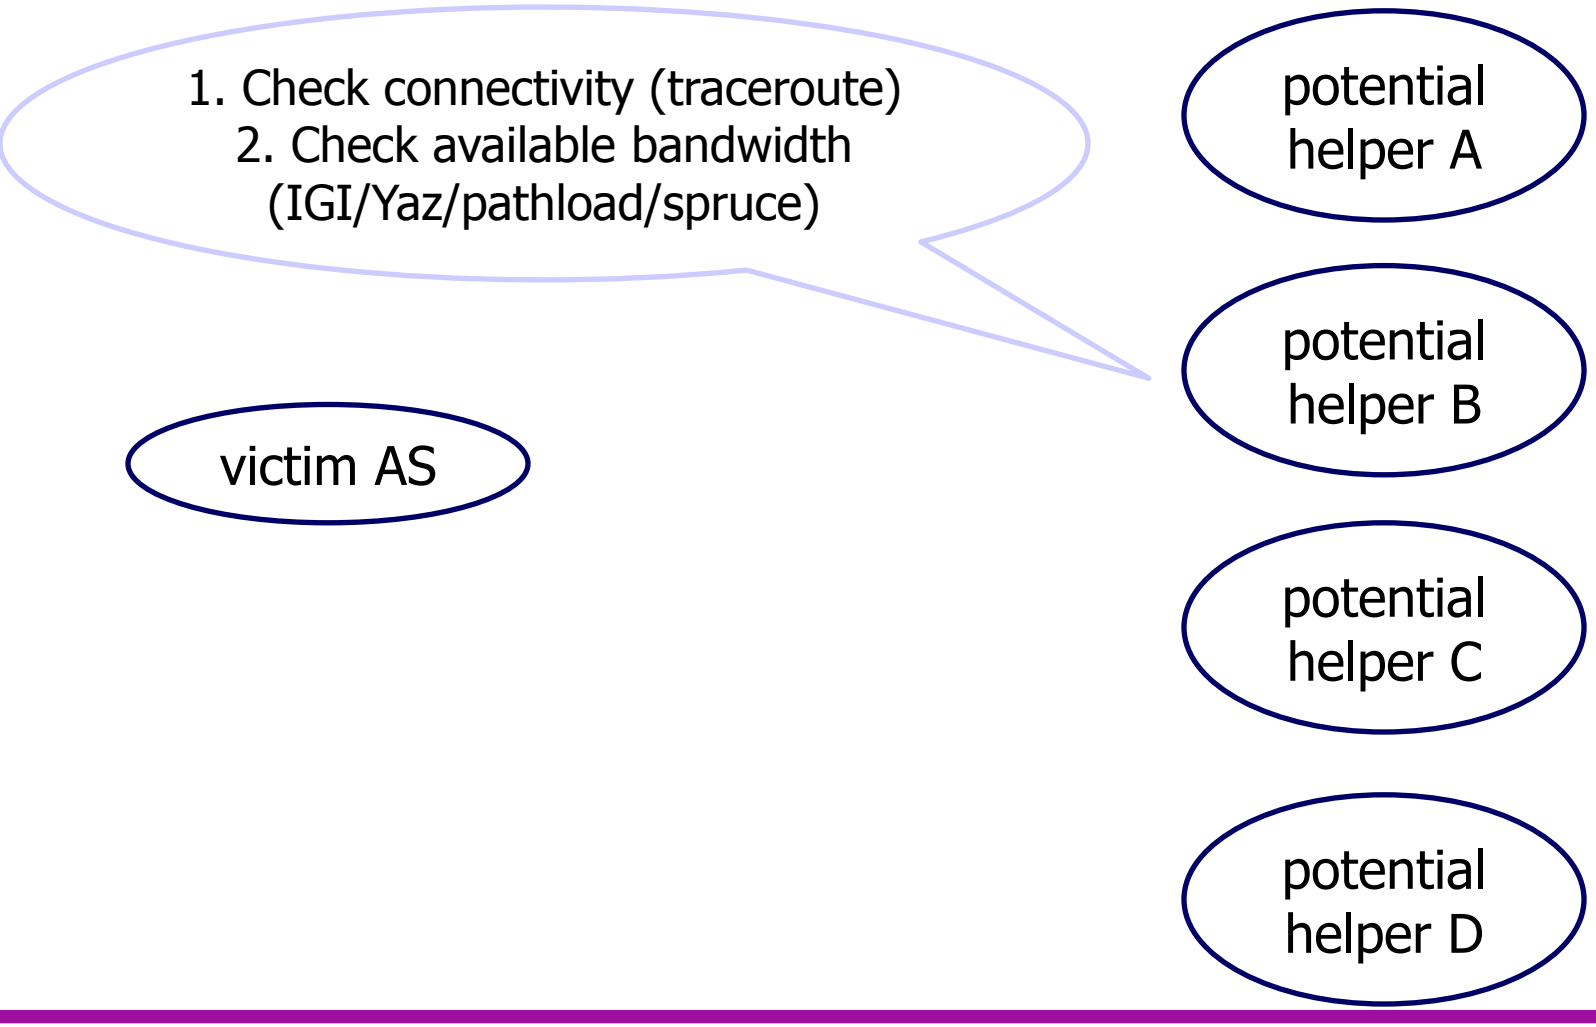

Automatic communications: reply phase

Message confidentiality with public key crypto

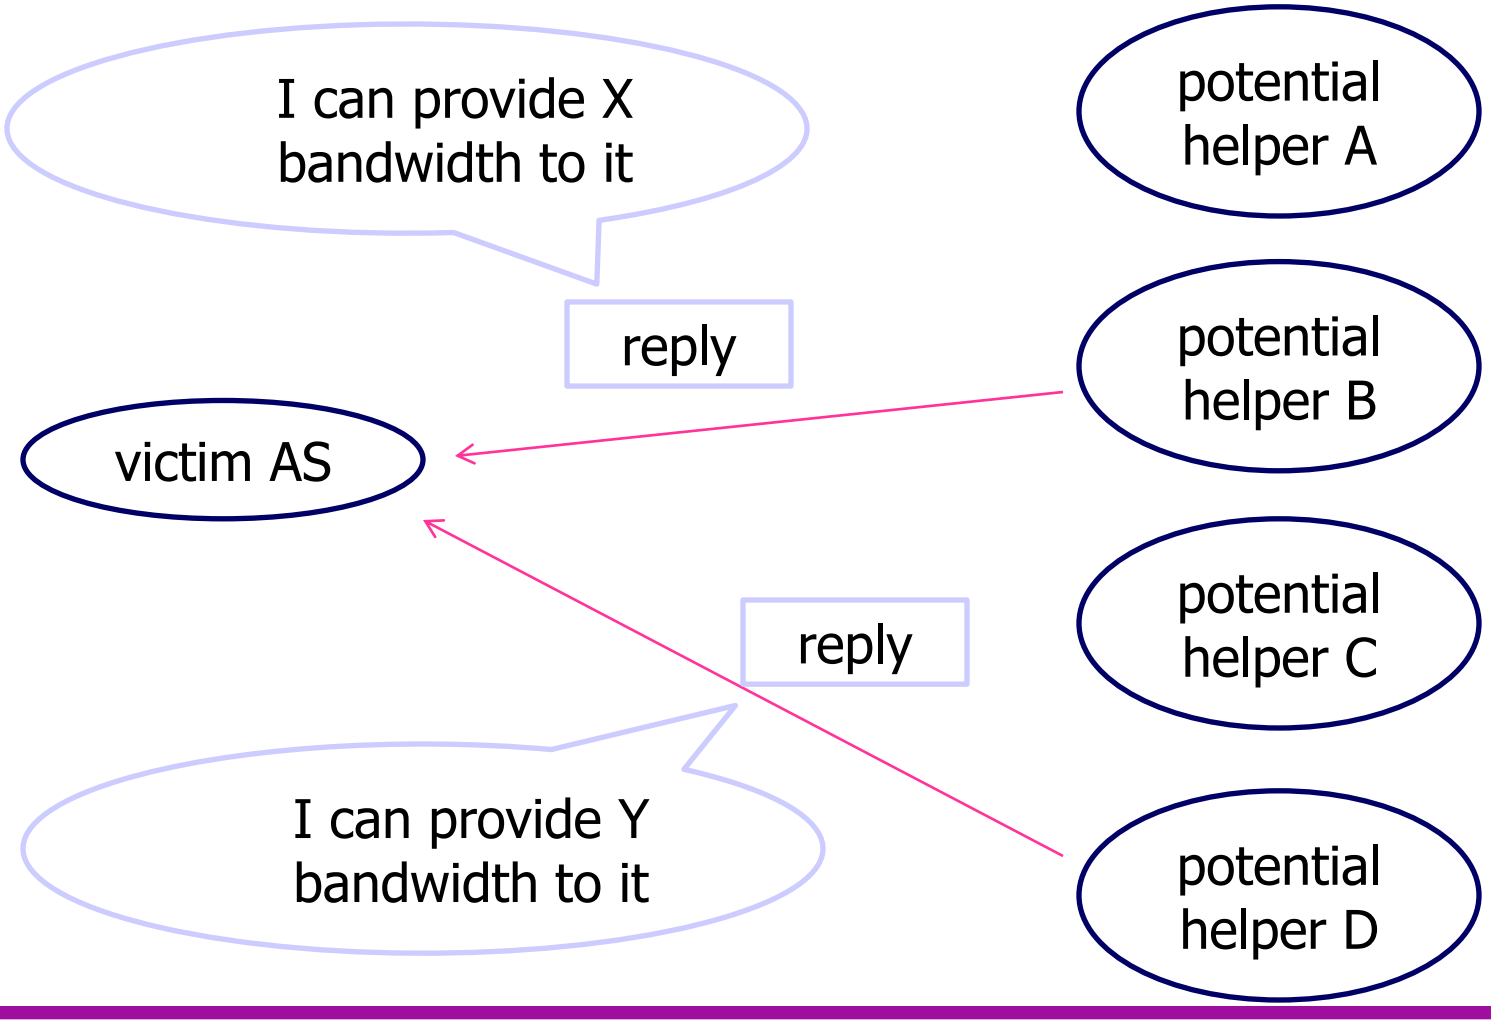

Automatic communications: ACK phase

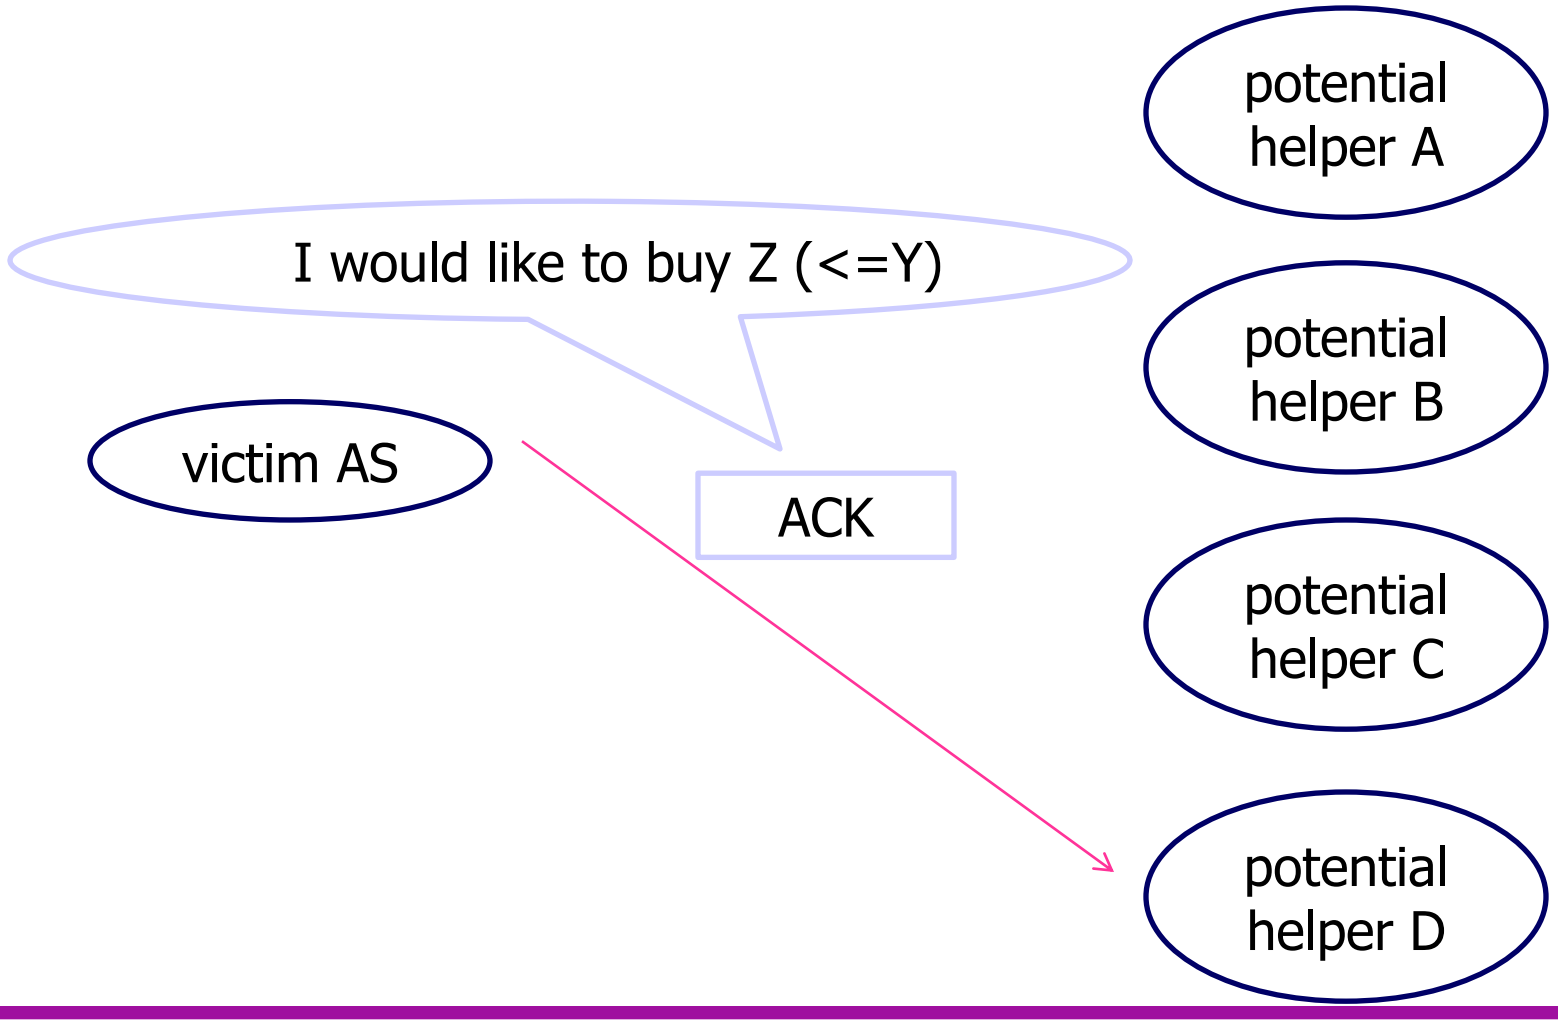

Automatic communications: new BGP session

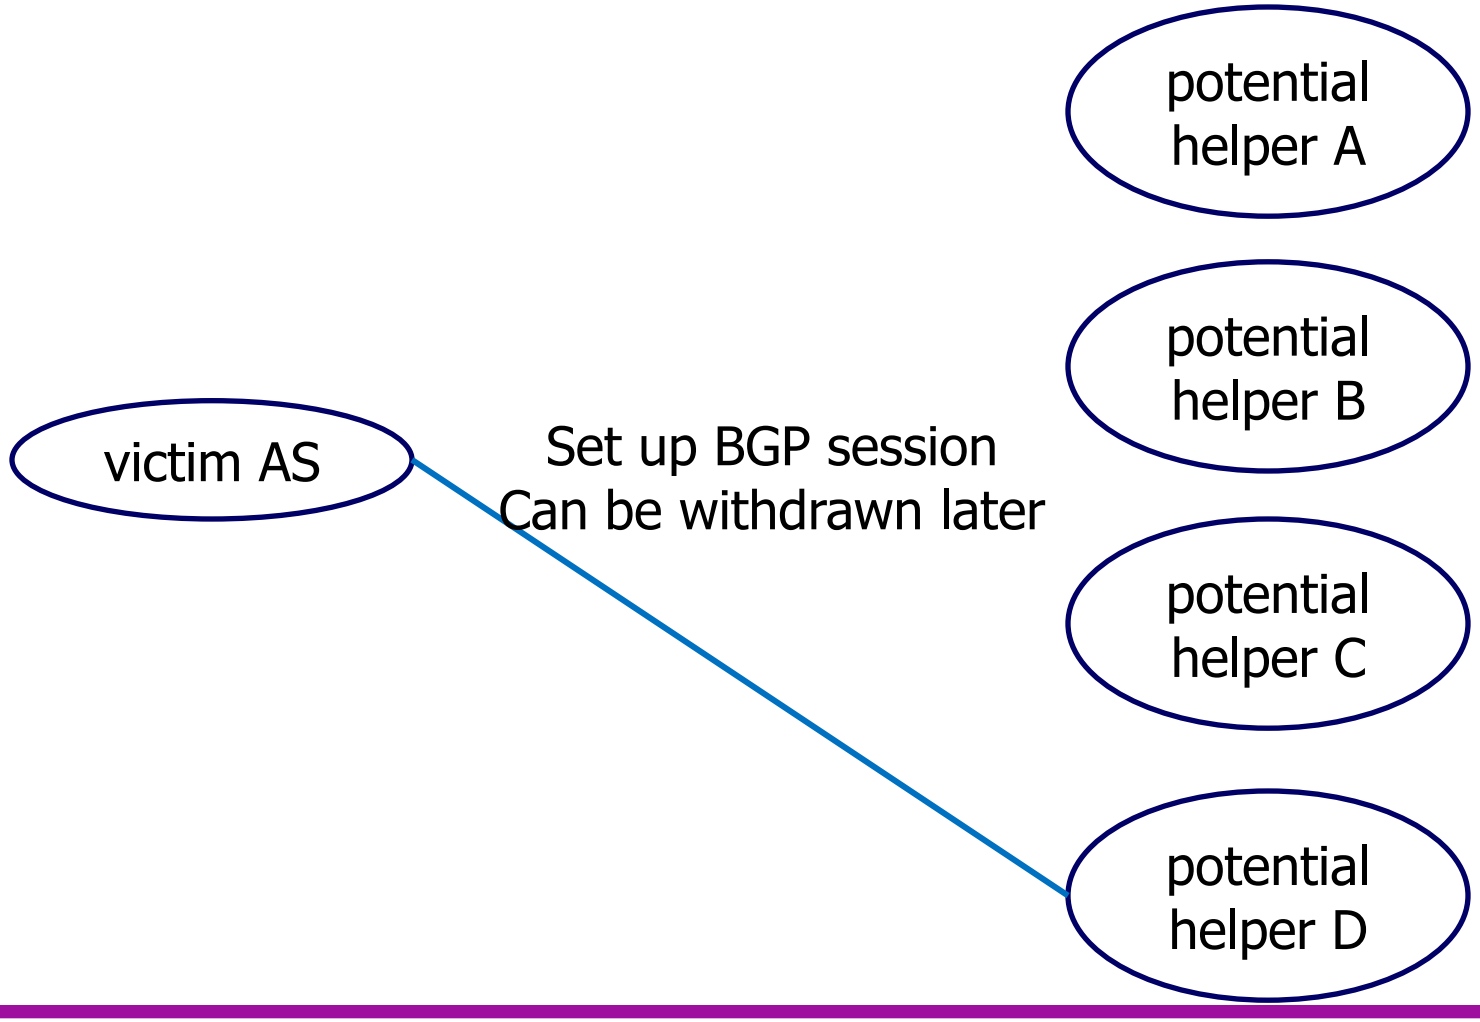

Optimal selection of helper ISPs

- From a single victim ISP perspective
  - Buy transit from a minimal number of ASes
  - Recover all the (prioritized) traffic
  - Least cost

Summary

- Point out a new venue for Internet failure recovery.
- Evaluate the potential routing diversity via IXP and PR with the most complete AS topology graph.
- 40%-80% of affected <Src, Dst> AS pairs can be recovered via IXP and PR with multiple paths and moderate shifted paths.
- Possible and practical mechanisms to utilize potential routing diversity.

- Look forward to feedback and collaborations from IXP/ISPs!