Internet Emergency Response through Reconnection in Internet Exchange Points

Chengchen Hu
Tsinghua University

Kai Chen, Rahul Potharaju, Yan Chen
Northwestern University

Bin Liu
Tsinghua University

Yang R. Yang
Yale University

Outline

1. Motivation & idea
2. How many potential helpers in an IXP?
3. How to discover the available helpers inside an IXP in emergency?
4. Business considerations
5. How to reconnect?
6. Experimental evaluations
7. Summary

1. Motivation & idea: Internet emergency

- 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 …
- Internet emergency can lead to substantial disruption
  - Dec. 26, 2006, Taiwan quake damaged 7 (of 9) cables causing a disruption from Asia to America, which lasted for days.

Example: Taiwan quake incident

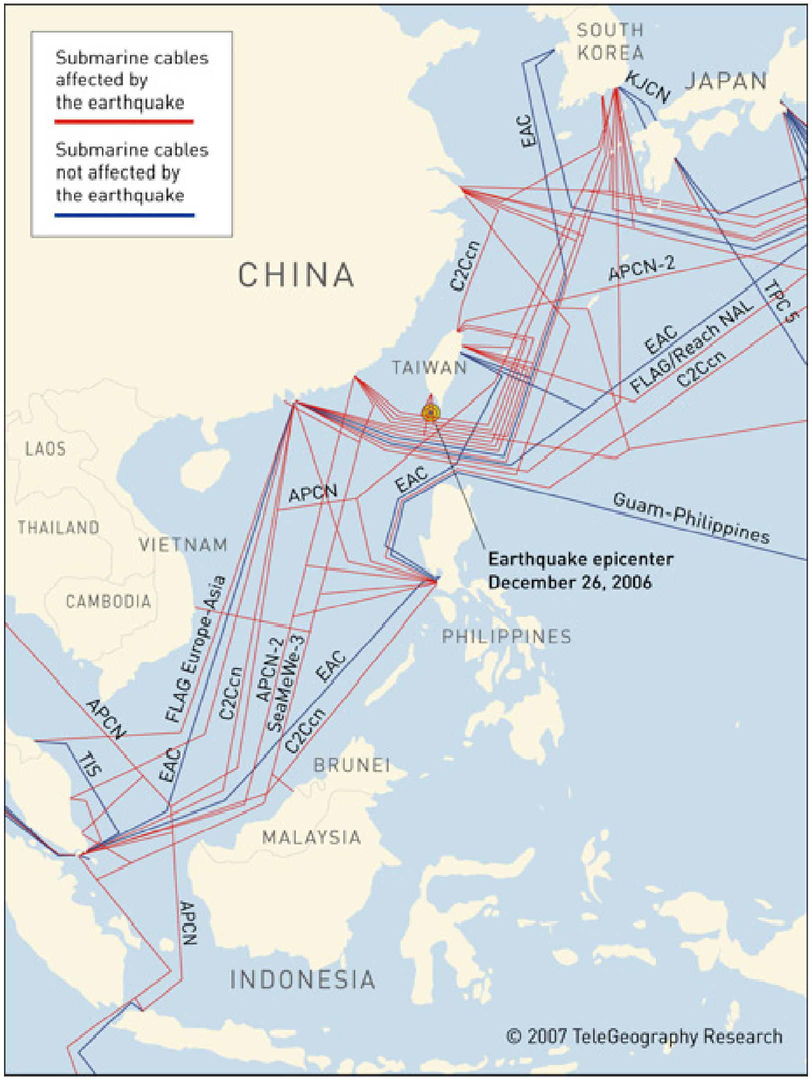

- Large earthquakes hit Luzon Strait, south of Taiwan on 26 December, 2006
- Only two of nine cables **not** impacted
- All cables reported repaired as of February 14, 2007

Taiwan quake incident (cont’d)

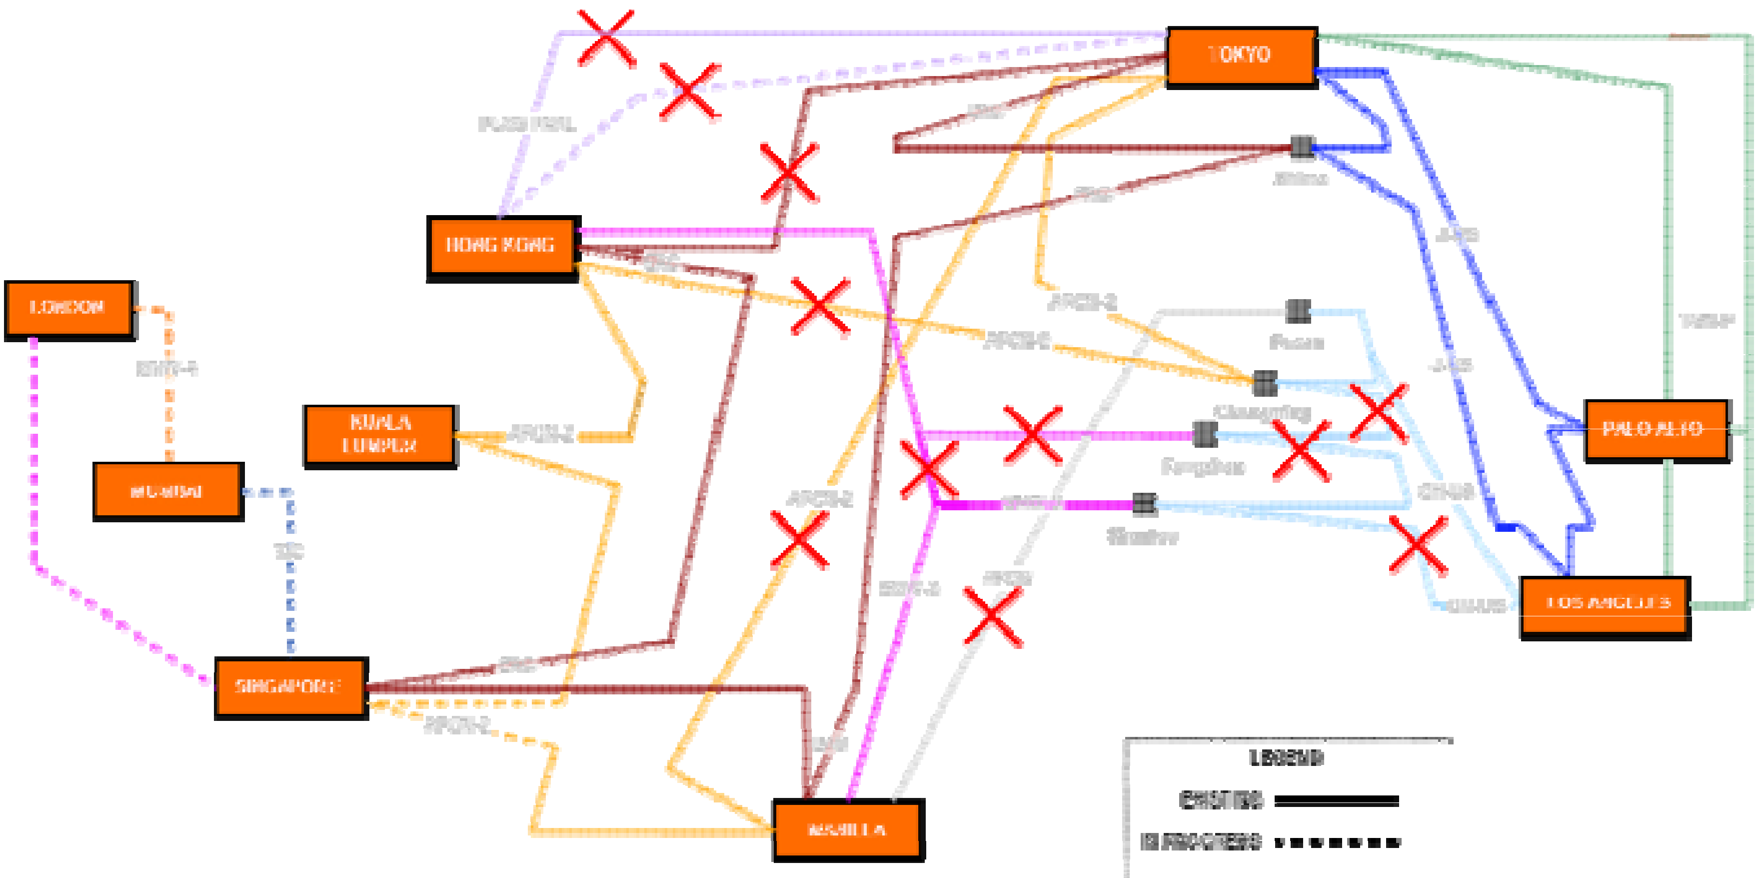

Outage by origin AS (Asia part)

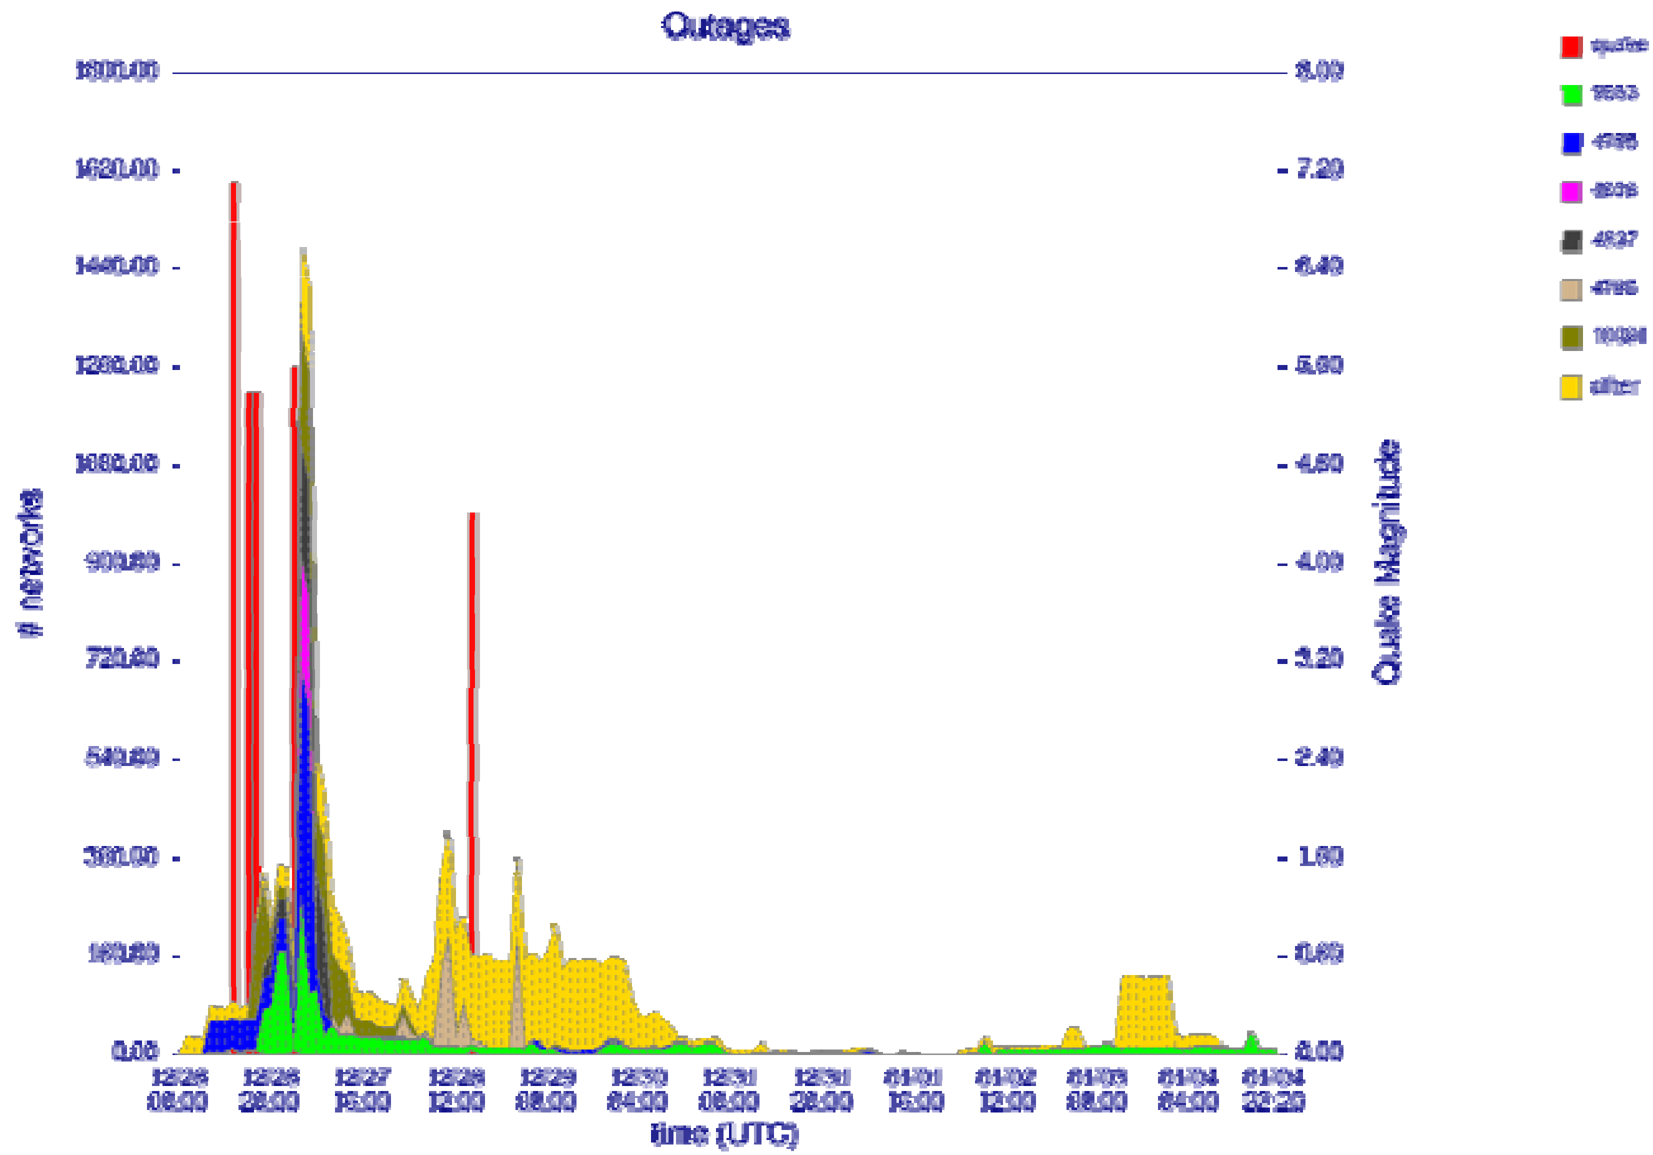

1. Motivation & idea: Internet emergency

- Internet Emergency Response
  - The cable repairing is slow
  - Alternative & indirection solution?
    - Quick response (Automatic/semi-automatic)
    - Recover routes
    - Involve multiple ASes (Inter-AS)

But not easy job…

Problems and challenges :

- How to design an effective system architecture for Internet emergency response?
- Once a disaster is detected, to where and who we will lunch our rescue inquiries?
- How to find the available helpers?
- If no direct connection beforehand, how can we determine alternative routes?
- How many resources are available? If not enough, how to get a global balancing?
- how to establish business relationships with actual helpers and achieve resource allocations?

Breaking the fence of legal issue, business restriction and making the mechanism incentive-compatible

Relatively easier to get help from neighbors

Much easier to regain connection from neighbors

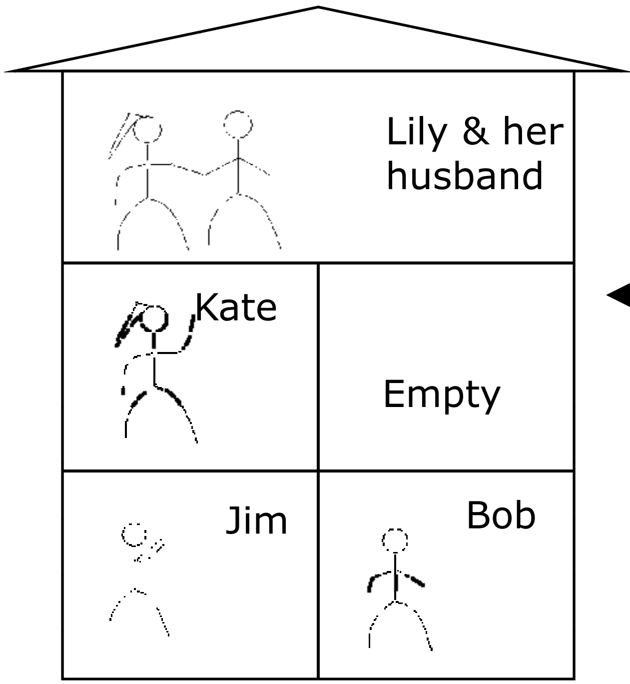

Hotel

IXP is a hotel for ASes & routers

Don't know who is your neighbor for privacy reasons

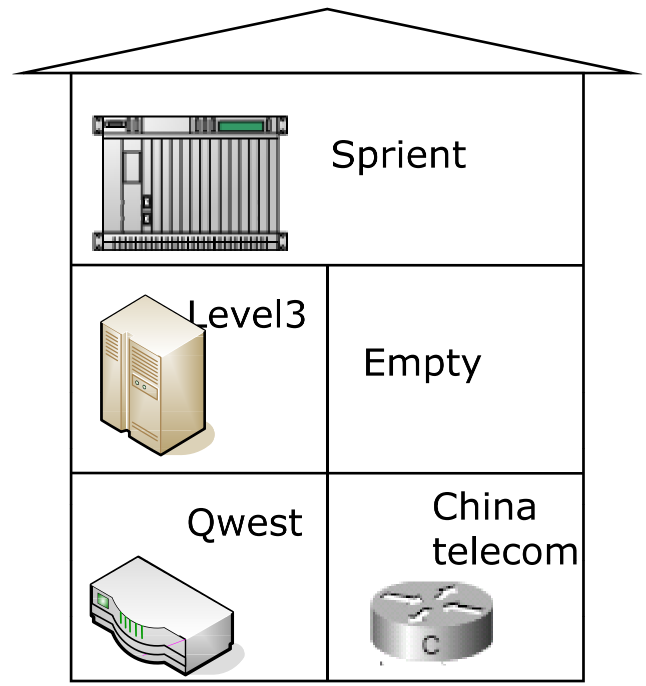

IXP- Internet eXchange point

High level framework

- In emergency, A suffered network (buyer)
  - Discover the collocated AS/routers (Potential helpers)
  - Check the availability of potential helpers (available helpers)
  - Select helpers and get permission to use the resource (actual helpers)
  - Reconnect to the actual helpers
- Note
  - An AS may locate in several IXPs, and it tries to discover and utilize the helpers in all IXPs it located.

We try to solve…

- How many potential helpers in an IXP?
- How to discover potential/available helpers inside an IXP in emergency?
- What is the business considerations on actual helpers selection and allocation?
- How to reconnect?

2. How many potential helpers in an IXP?

AS Relationships

- Provider-to-customer
  - one pays money to another network for Internet access
- Peer-to-peer
  - two networks exchange traffic between each other's *customers*
  - The traffic from peers will not delivered to its providers

What can be helpful

- Upgrade peering connection
- Add new provider-customer connections (non-existing but easy connected in an IXP)

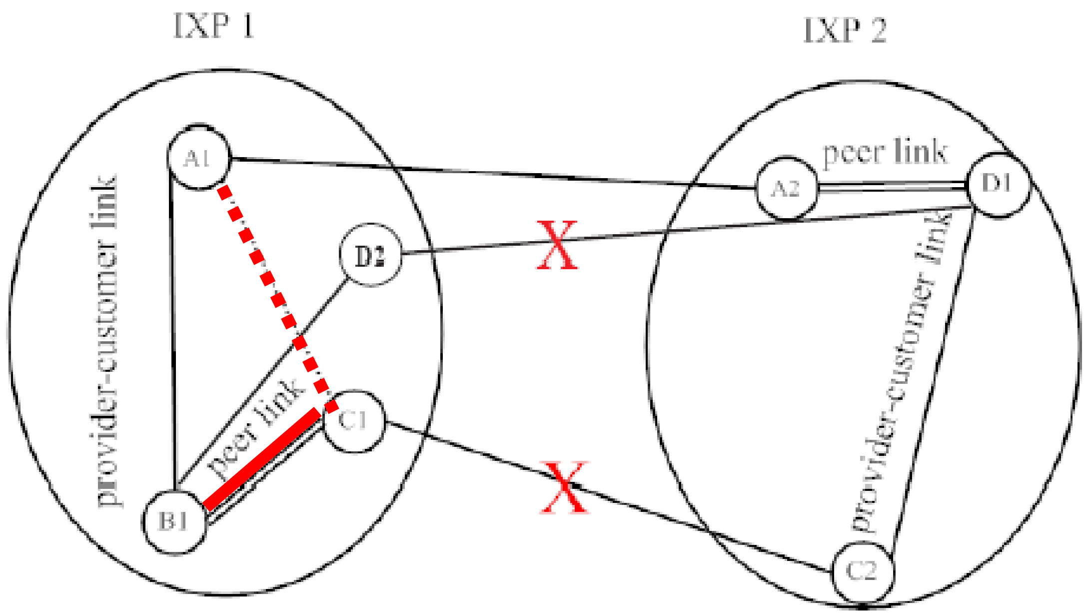

Evaluation on peers

- We use the method/data presented in CAIDA AS Ranking project to infer the relationships among ASes
- The data are from RouteView collected on Oct. 8, 2007
- The # of ASes is ~20,000
- The average peer links for each AS is about 0.77
- The distribution of the peer links are quit uneven.

Only peers are not enough

1. The number of peering links is small

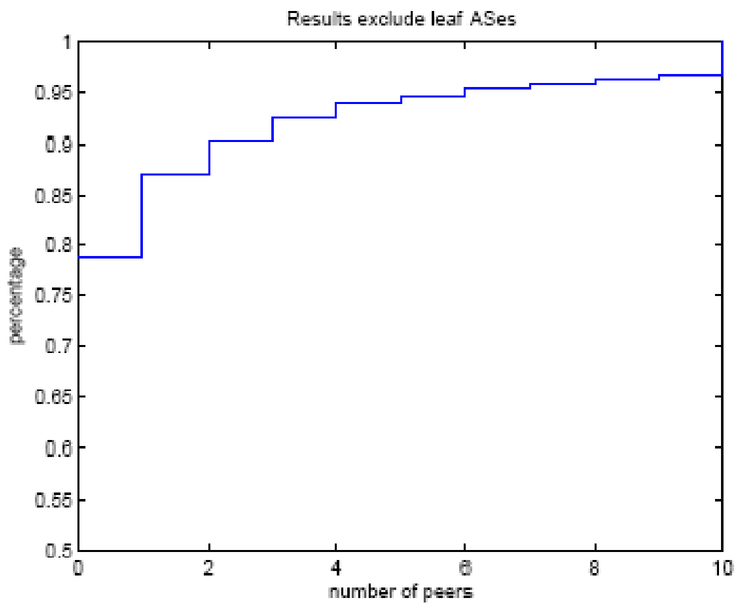

Figure 4. Cumulative distribution of the number of the peers links (exclude leaf ASes).

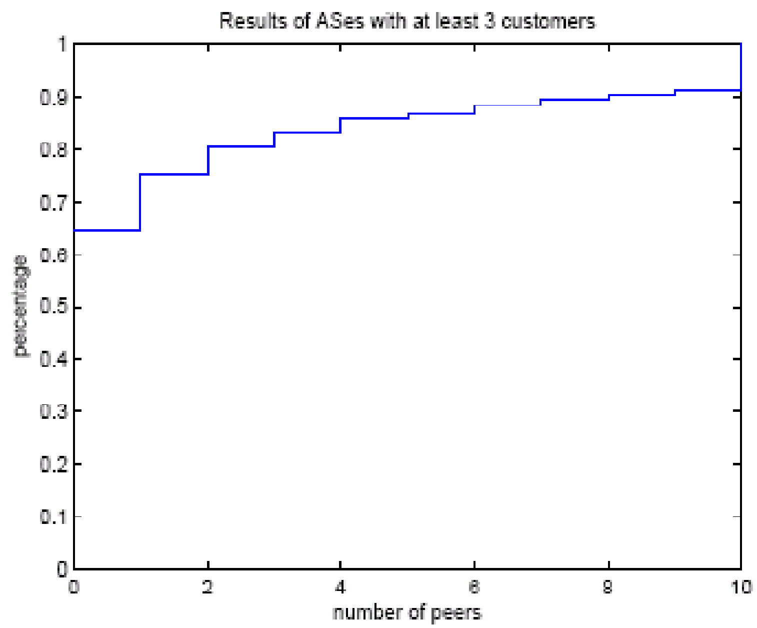

Figure 5. Cumulative distribution of the number of the peers links in ASes who have at least 3 customers.

2. Furthermore, the bandwidth of peer link is also small

Evaluation on non-existing links

- No data or method is available for discovery of the router/As in an IXP
- We develop an measurement experiment to get the result
- Data set:
  - trace route data from iPlane on Nov. 17, 2007
  - 201 vantage points, 110k prefixes
  - Cover 90% edge prefixes

IXP discovery

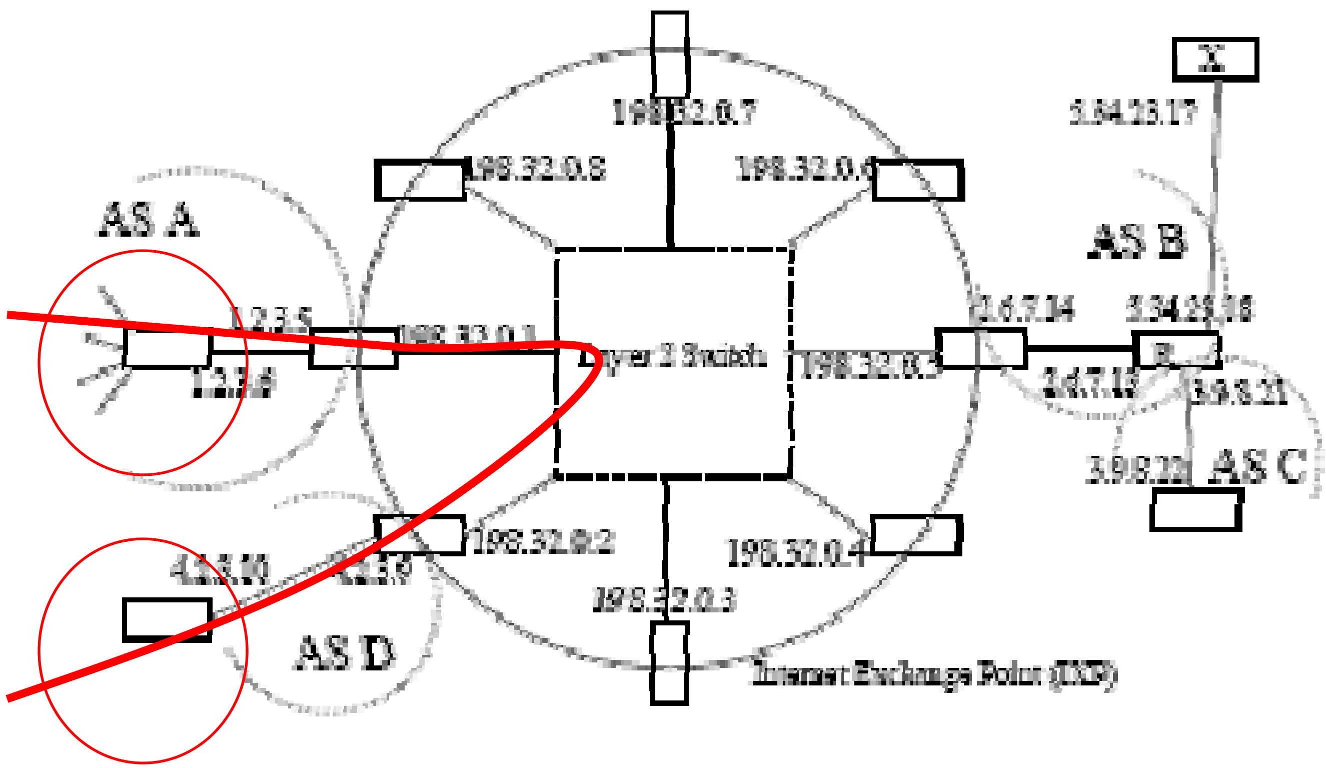

Y. He, G. Siganos, M. Faloutsos, and S. V. Krishnamurthy, “A systematic framework for unearthing the missing links: Measurements and impact,” in *USENIX/SIGCOMM NSDI*, 2007.

Results

Table 2: The number of potential helpers to important destinations. PH1 are the helpers to MSN, PH2 are the helpers to Google, PH3 are the helpers to DNS, PH4 are the helpers to Yahoo, and PH5 are the helpers that can access all the above.

Name of IXP	# of Participants	# of PH1	# of PH2	# of PH3	# of PH4	# of PH5
HKIX	100	97	96	96	96	96
JPIX	112	110	110	110	110	110
AMSIX	289	275	276	276	276	275
DECIX	220	217	217	217	217	217
LINX	262	256	257	256	256	256
PARIX	44	41	42	41	41	41
AIX	57	54	54	54	54	53
NYIIX	101	100	100	100	100	100
SIX	105	101	101	101	101	101

Table 3: The number of "low bandwidth path"(LBP) and "high bandwidth path"(HBP).LBP1 and HBP1 are the paths to MSN, LBP2 and HBP2 are the paths to Google, LBP3 and HBP3 are the paths to DNS, LBP4 and HBP4 are the paths to Yahoo.

Name of IXP	HBP1	LBP1	HBP2	LBP2	HBP3	LBP3	HBP4	LBP4
HKIX	47	50	67	29	64	32	63	33
JPIX	17	93	86	24	42	68	84	26
AMSIX	75	200	194	82	119	157	186	89
DECIX	70	147	146	71	101	116	141	76
LINX	83	173	192	65	136	120	181	75
PARIX	15	26	33	9	23	18	31	10
AIX	16	38	40	14	25	29	35	18
NYIIX	33	67	83	17	49	51	80	20
SIX	23	78	71	30	40	61	61	40
AVERAGE	42	97	101	38	67	72	96	43

Results(Cont.)

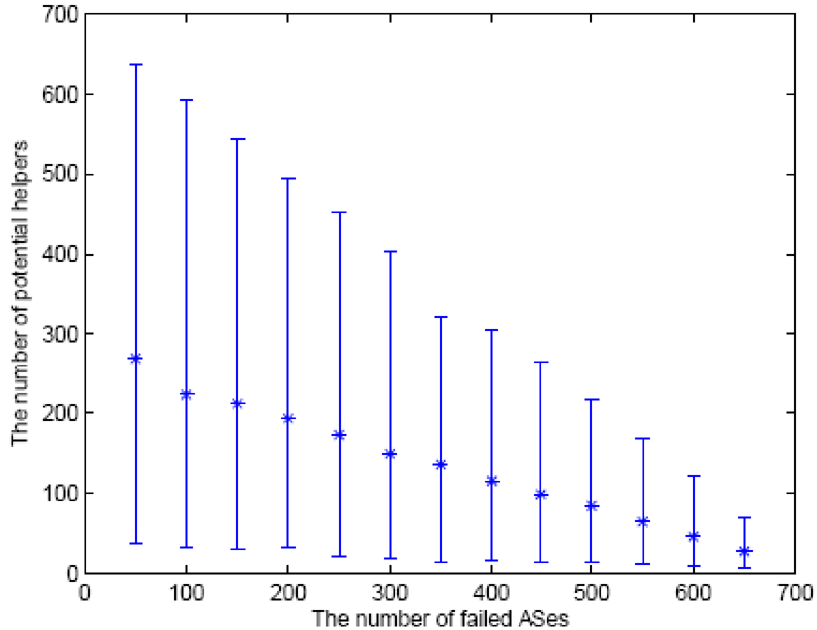

Figure 8. The number of potential helpers in simulated disaster. In this simulation disaster, we involve all the unique 751 ASes in HKIX, JPIX, LNIX, AMSIX, PARIX, DECIX. We just randomly select a few ASes to be failed in a disaster. Note that, if an AS has located in more than one IXP, it can find potential helpers in all the IXPs it locates. The star in the figure is the average value, and two ends are the minimal and maximal number of potential helpers.

3. How to discover the available helpers inside an IXP in emergency?

Available helpers

- Potential helpers, who can provide helper in the same IXP
- Still reachable to specific networks

Regulator-based solution may not apply

- Regulator-based solution
  - There exists a regulator who knows all the routers in the IXP
  - Everyone who wants to know the information, just query the regulator.
- Not incentive-compatible
  - ASes do not want to disclose the locations of their routers
  - Legal problems

Possible solution

- Use methodology in the last section could get potential helpers; after that, check their availability using looking glass
- Advantage: infrastructure-compatible
- Limitation:
  - # of looking glass sites
  - Inference method may have error

Patch current infrastructure

- Build controlling paths among routers in an IXP
- Add a signaling protocol among routers in an IXP
- Add a script on border routers for queries that are similar to looking glass

Communication Channel

- Modify classical spanning tree algorithm
- Limit broadcast inside the IXP
- Run in MAC layer

Spanning tree

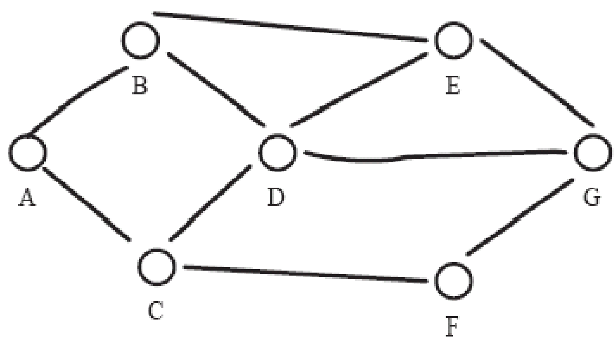

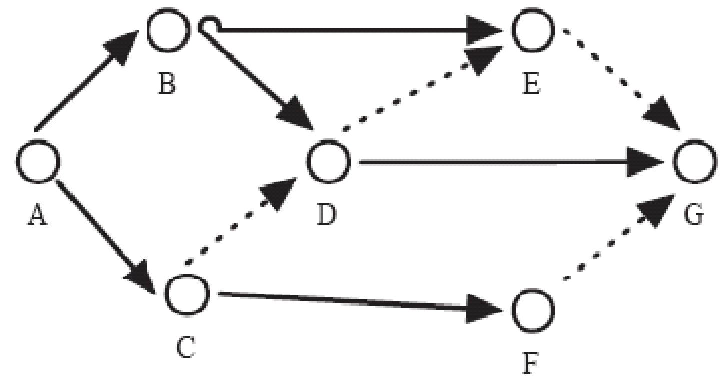

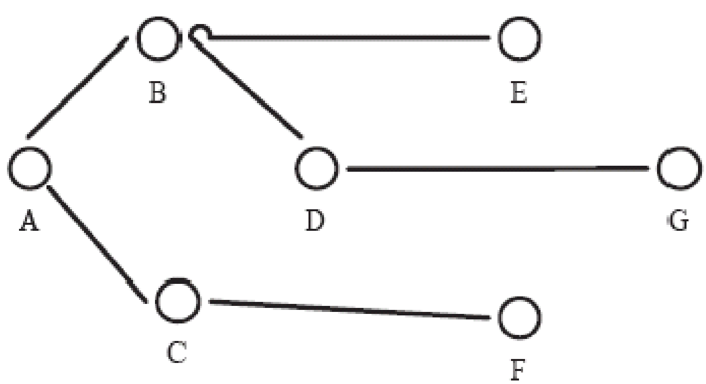

Only flood to neighbors who:

1) belongs to different ASes and are in the same IXP
2) are in same AS with buyer but are known in the same IXP

Signaling process

- Query
  - A buyer broadcasts queries through controlling path.
- Availability check
  - Potential helpers check the reachability to specific networks in the queries.
- Reply
  - Only available helpers send a reply to the buyer

Looking glass script

- Provide any one interface from the following
  - Show ip bgp
  - Ping
  - traceroute

4. How Business considerations

- Like BGP, we should give selection freedom to different networks
- It is not main points of our work, and we just give 3 possible usage models in order to provide some insights
  - Free bid model
  - Broker-based model
  - Double auction model

Free bid

- Keep the list of available helpers in each individual network
- Select one with its own preference
- Pre-agreement may apply for the relationship & price
- Like BGP and like mechanisms between airlines

Broker-based model

- Currently, bandwidth sells at coarse granularity, e.g., 1G, 2.5G, 10G
- A broker agent buys large qualities of bandwidth from helpers and resell it in a more flexible way.
  - The broker own money from trading
  - the buyer can save money by always choosing a larger bandwidth
  - the provider save cost to build agreement with large number of buyers

Double auction model

- Like a stock market…
- Helpers and buyers bid in different IXP with different prices.
- The bandwidth in different IXP = different stock

5. How to reconnect?

Two methods

- Using existing link
- Adding a new link

Using existing link

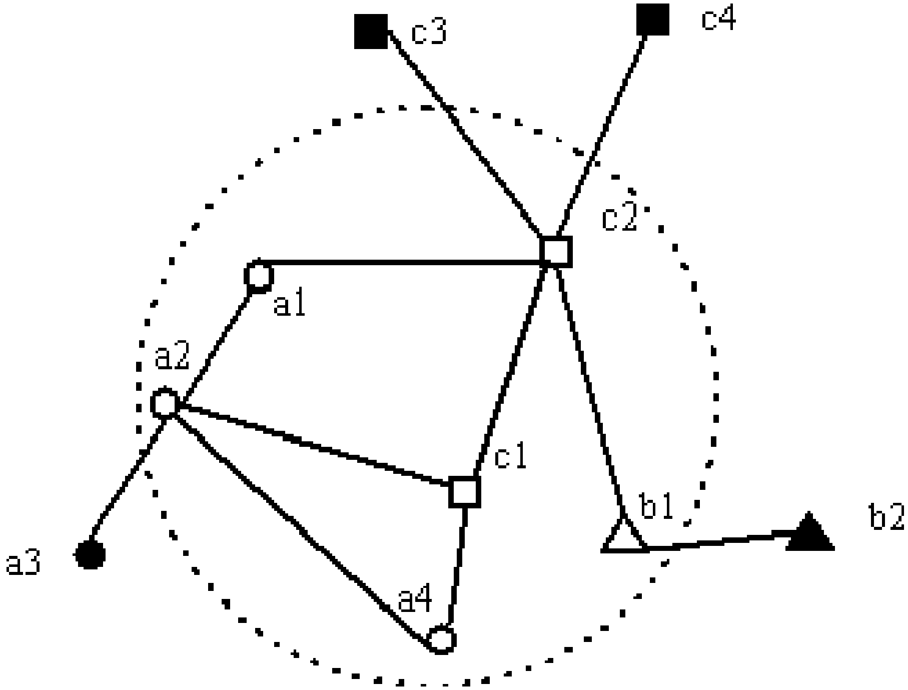

- b1 wanna connect to its helper a1 through c2.
- c2 should be b1's customer or peer originally
- Actually, b1 are using c2 as a direct helper instead of a1
- Just modify the relationship of b1c2 temporarily.
- c2 exports some route to specific network from a1 to b1.
- Current link bandwidth may not be sufficient
- Chain effect when using helpers hops away

Adding a new link

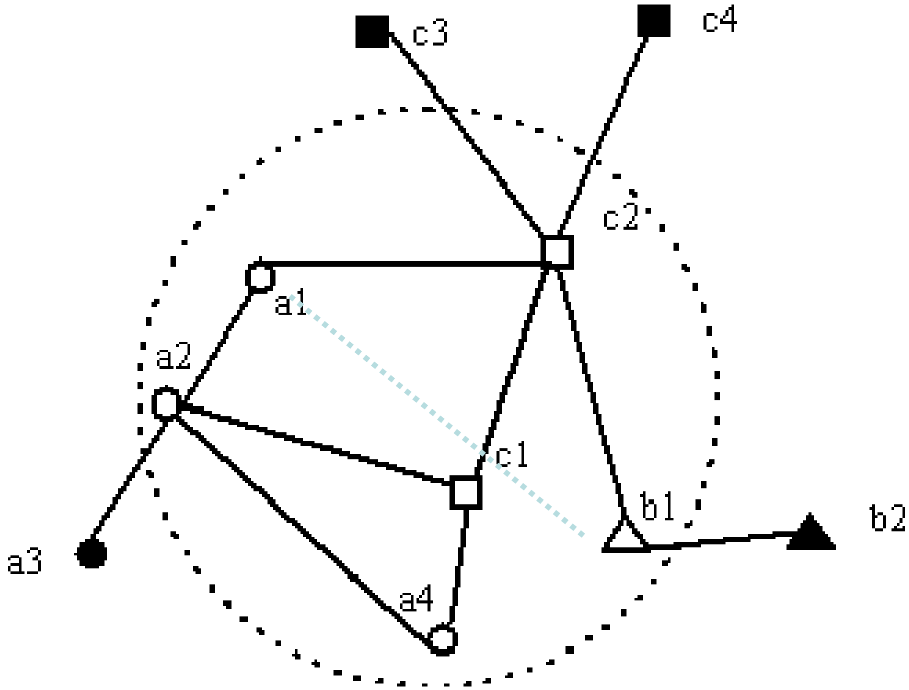

- By adding a1b1, we connect b1 to its helper a1.
- Need manual configuration for direct interconnection model – slow
- Introduce new links and bandwidth
- A full provider-customer link may affect existing traffic a bit.
- setting a temporary partial provider-customer link where only routes to specific network is exported ?

Reconnection

- Both two methods may affect existing traffic
- Consider response speed as major selection criterion
- Direct Interconnection is recommended to use the existing link if possible
- Exchange-based interconnection is recommended to add new links by switch configuration.

6. Partial result on Taiwan earthquake

Taiwan earthquake recovery
on AS7473, AS4143, AS24077

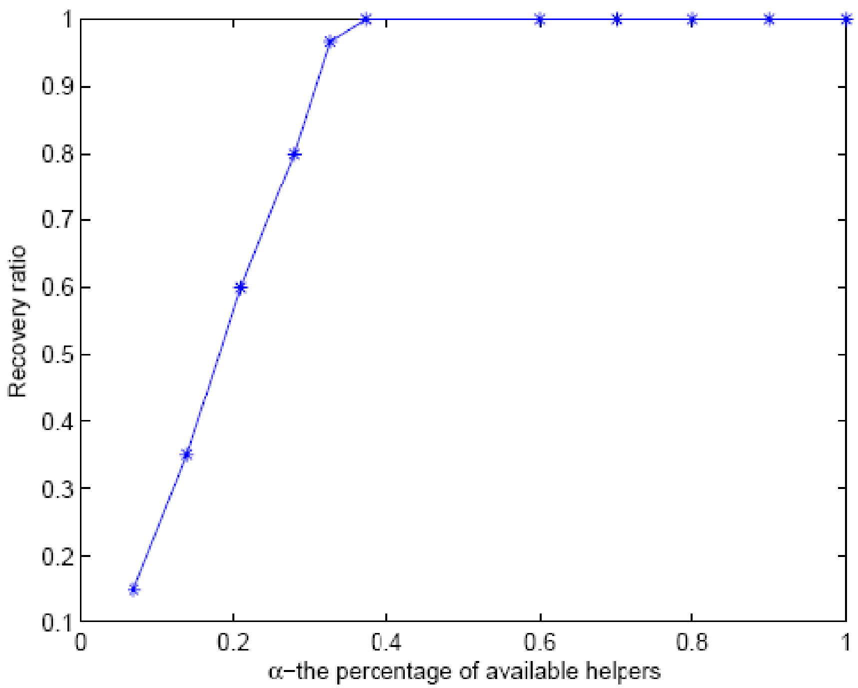

- Note that less than 40% available helpers can recover all the traffic of these three ASes.

7. Summary

- Internet disaster response is an important and practical issue
- Existing recovery process is manual, slow and inefficient
- We propose a systemized solution including the recovery architecture, communication protocol, reconnection-building strategy as well as the resource allocation mechanisms
- We simulate an evaluation during the after-math scenario of the recent Taiwan earthquake to demonstrate the effectiveness of our design
- The common issues and several guidelines can direct the future development of Internet disaster recovery