A characterization of routing dynamics between neighbor ASes

Renata Teixeira[†], Steve Uhlig[§], Augustin Soule[‡], Christophe Diot[‡]

[†] LIP6, UPMC, France
[§] CSE Dept., UCL, Belgium
[‡] Thomson, France

Agenda

- Interactions between neighbor ASes
- Case study: GEANT and Abilene
- Methodology
- Results
- Implications
- Conclusions

Interactions between neighbor ASes

Interactions between neighbor ASes

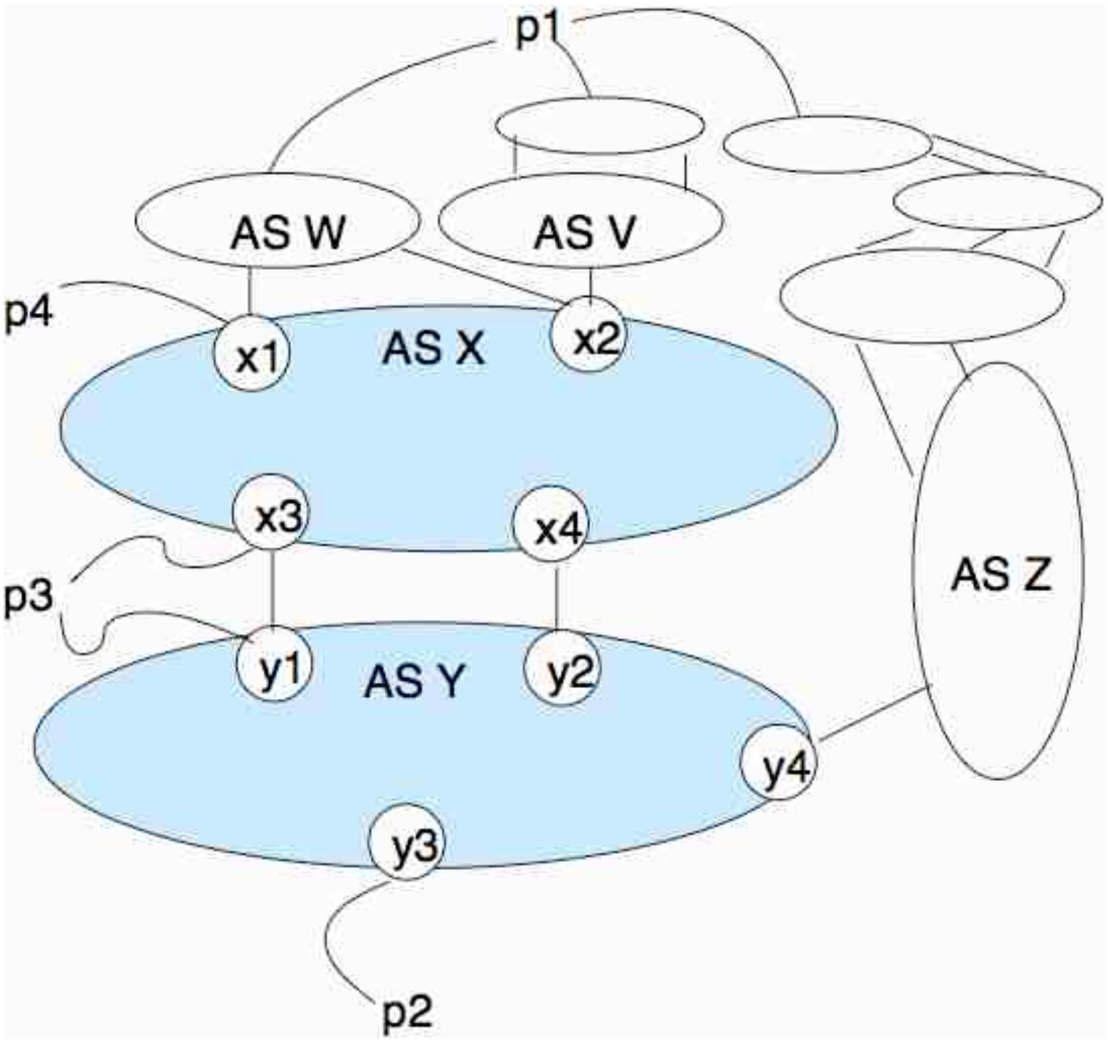

- Egress router for *p*: router that receives a best route
- Egress-set for *p*: set of egress routers for *p*
- $P_{X \to Y}$: *set of prefixes advertised by X to Y*
- $P_{X \to Y} = \{p1,p3,p4\}$
- $P_{Y \to X} = \{p2,p3\}$

Case study: GEANT and Abilene

GEANT and Abilene: interconnectivity

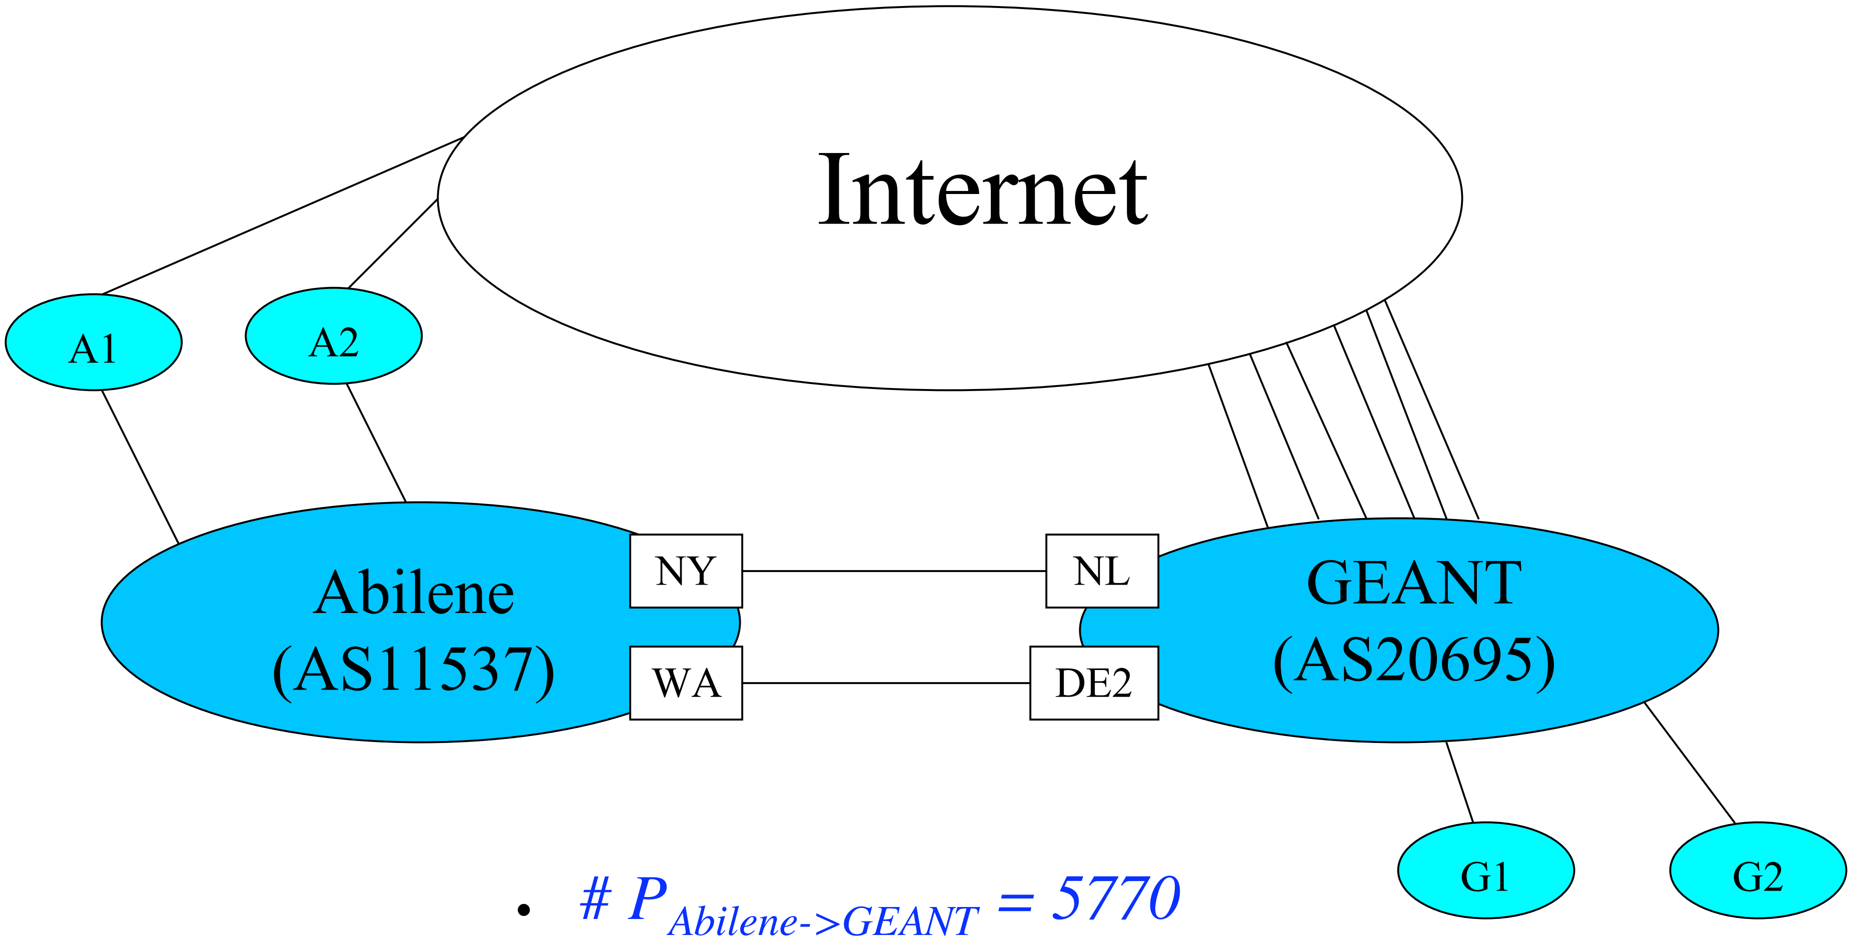

- $\# P_{Abilene\text{->}GEANT} = 5770$
- $\# P_{GEANT\text{->}Abilene} = 2200$

GEANT and Abilene: measurement infrastructure

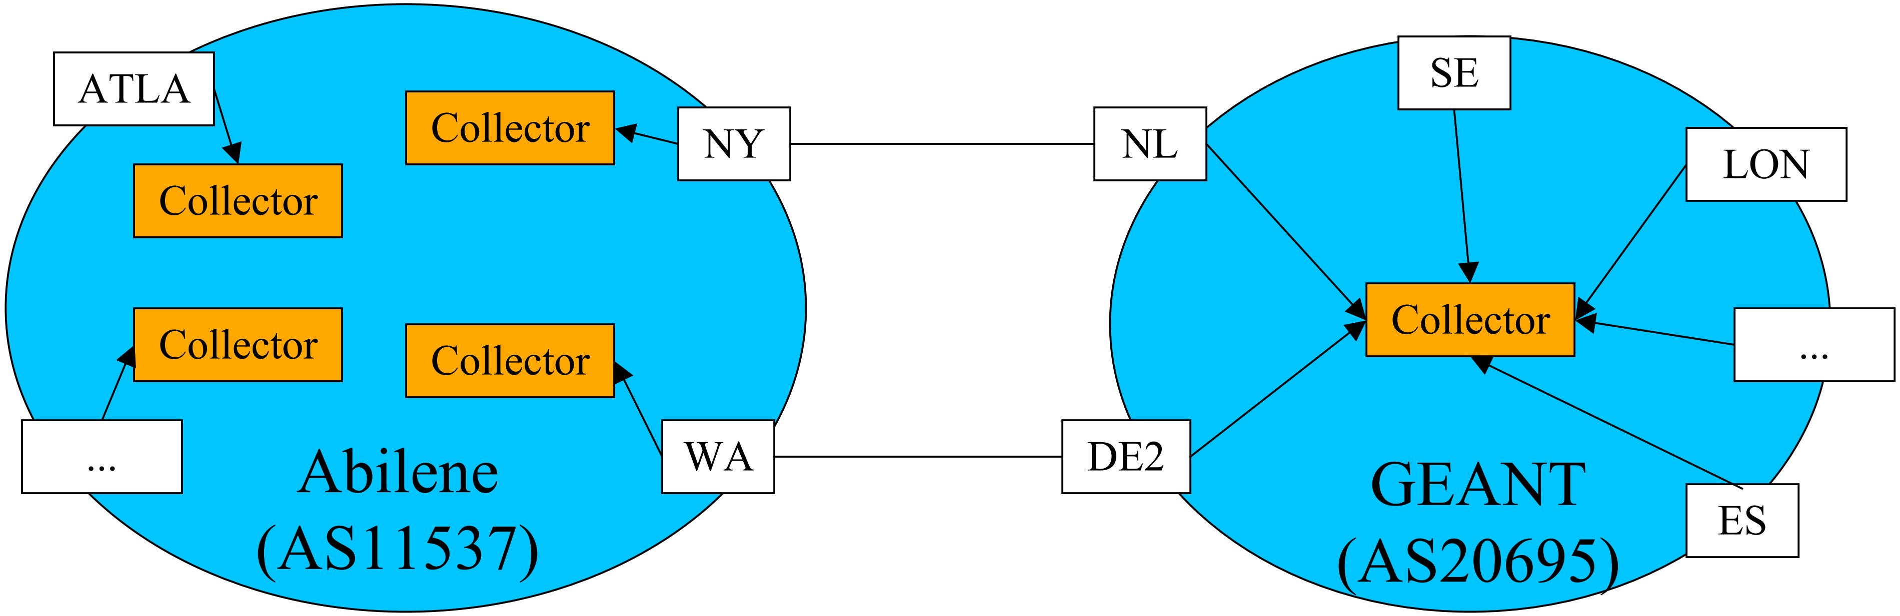

- Routing: 1 collector per POP (BGP + ISIS)
- Netflow: 1/100 sampling

- Routing: 1 collector in the network (BGP, ISIS)
- Netflow: 1/1000 sampling

Methodology

Methodology: egress-set changes

- *Egress-set change* ≡ BGP event that changes the egress-set for *p*
- BGP path exploration => many transient states to egress-set
- Filtering of BGP changes: group transient egress-set changes close in time (75% filtered)

Methodology: taxonomy

- Prefix down for X: $\Delta\ \#P_{X \to Y} \downarrow$
- Prefix up for X: $\Delta\ \#P_{X \to Y} \uparrow$
- Egress-set change for X: $\Delta\ \#P_{X \to Y} = 0$

Methodology: correlating egress-set changes

- AS X is said *source* of egress-set change if routing change observed first at AS X and then at AS Y (*destination*)
- Correlating algorithm:
  1. Selection of relevant egress-set change
  2. Identification of the time window T
  3. Matching related egress-set changes

Methodology: selection of relevant egress-set changes

- Considered prefixes:
  - $P_{G \rightarrow A} \cup P_{A \rightarrow G}$
  - Select prefixes whose BGP messages contain the other network as next hop
- If $P_{G \rightarrow A} \cap P_{A \rightarrow G} \neq \emptyset$ then causal relationship is unclear

Methodology: identification of time window T

- Propagation of BGP messages takes time:
  - route flap damping: avoid propagating messages for flapping prefix
  - out-delay (Juniper): delay messages before sending them
  - router's load: largely depends on number of prefixes and peers
- $T_{A,G} > 90$ min (lots of prefixes, out-delay and damping)
- $T_{G,A} = 0$ (no out-delay, lightly loaded)

Methodology: matching related egress-changes

- Order egress-changes in *destination* with time
- Look for events at *source* that might have triggered change at *destination*
- Compatible change: for each change of type *c* at *destination* at time *t*, an event of type *c'* at source is *compatible* if $t \leq t' \leq t+T(p)$ and if *c'* is compatible with *c*

Results: mutual impact between GEANT and Abilene

Analysis of mutual impact

Egress-set change			Impact		
Type	**Abilene**	**GEANT**	**Type**	**A→G**	**G→A**
prefix down	19,109	4,318	prefix down egress worse egress equiv.	5,496 126 3,510	1,506 0 94
prefix up	22,262	6,214	prefix up egress better egress equiv.	7,467 113 4,690	2,558 0 316
egress to worse	499	114	egress better egress equiv. egress worse	0 80 2	0 0 0
egress to better	598	442		0	0
egress to equiv.	5,828	3,035		0	0
total	48,296	14,123	total impact	21,484	4,474

Propagation time

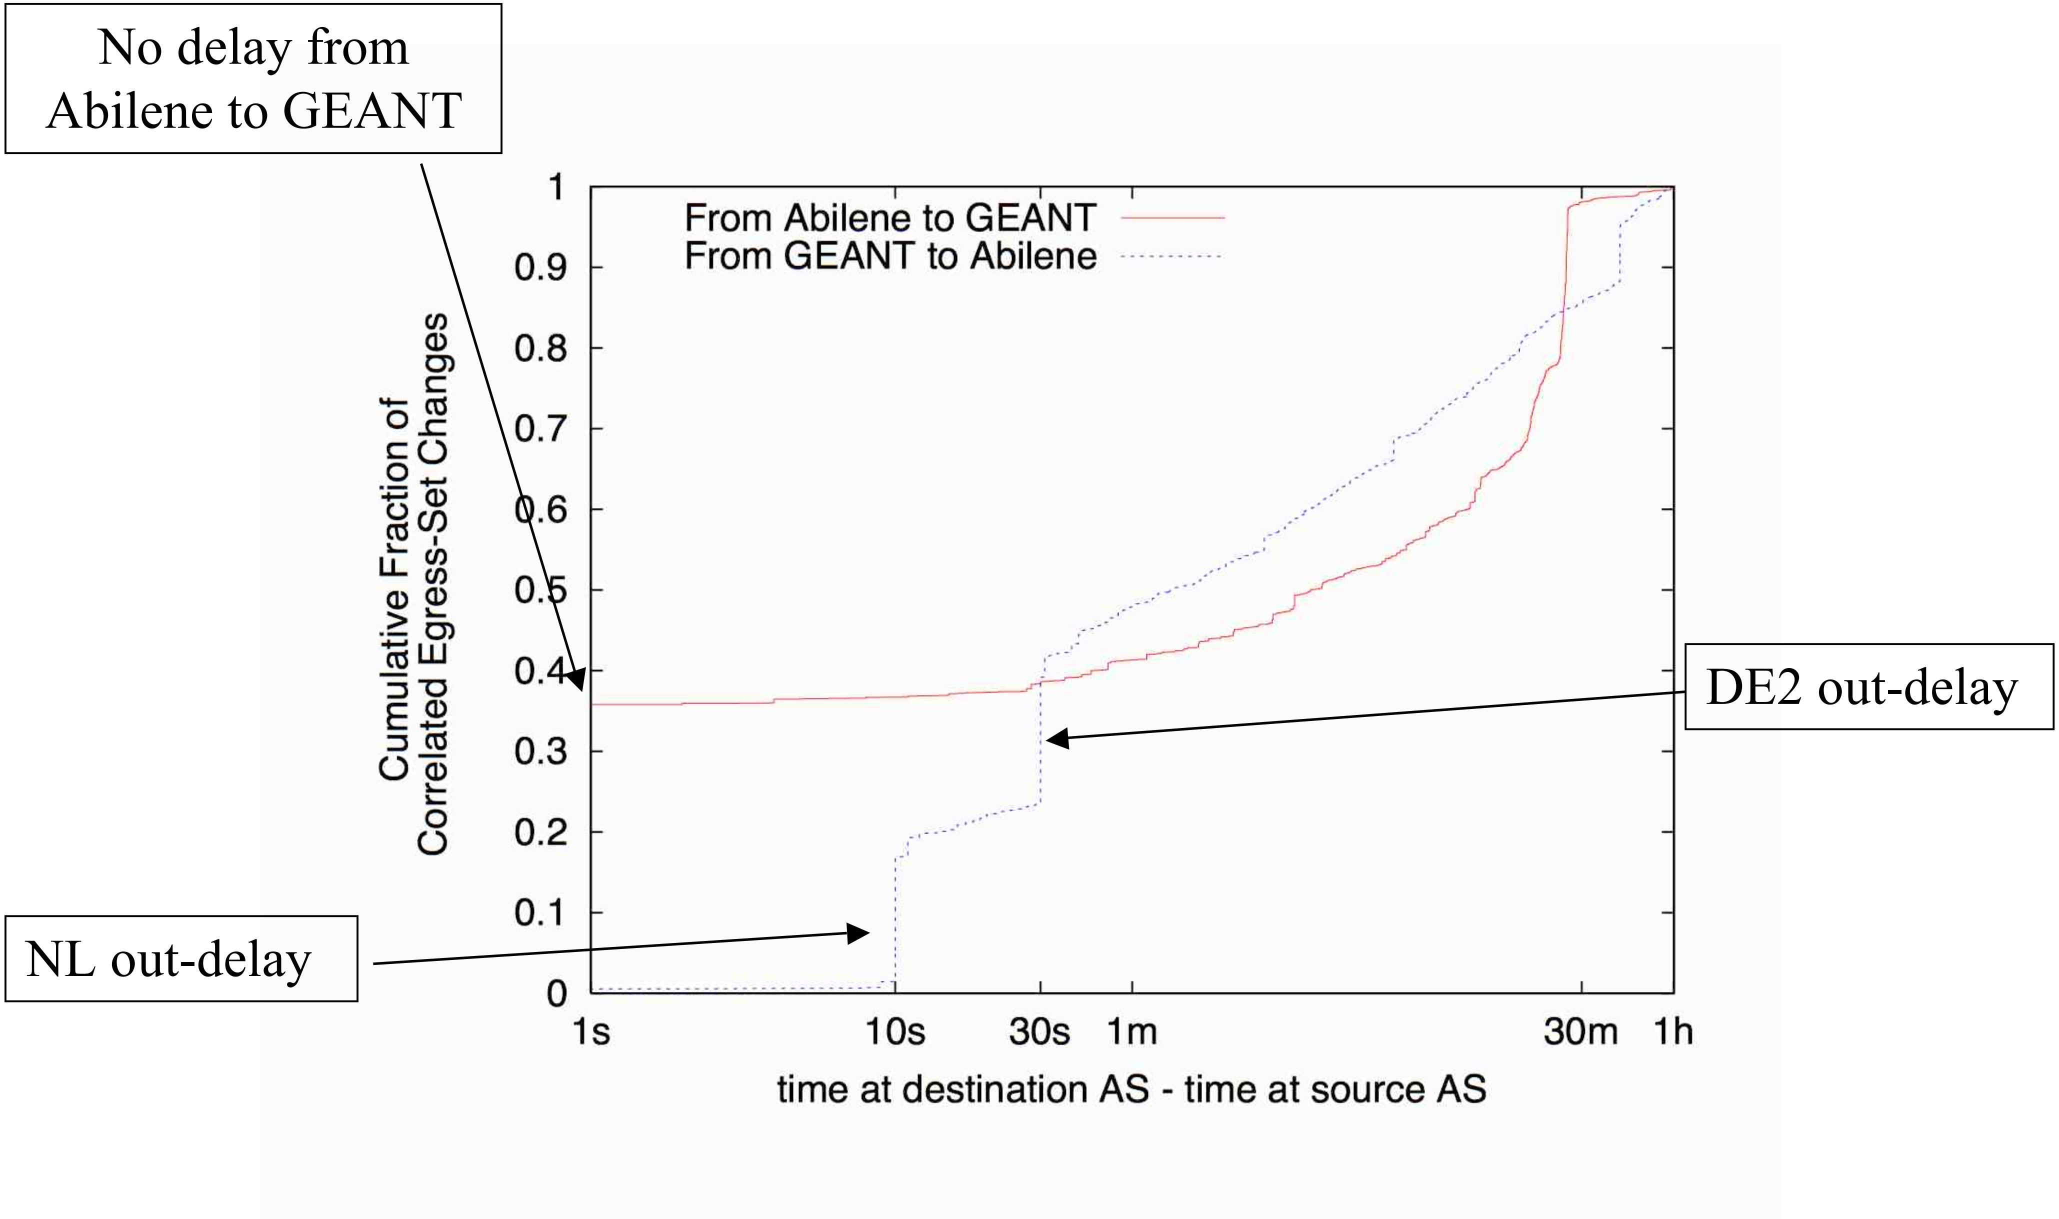

Implications

Implications on inter-AS diagnosis

- Diagnosing routing changes between neighbor ASes using egress-change sharing:
  - type of change
  - list of concerned prefixes
  - time of observation
  - traffic ? (rank in top hitters)
- Interest:
  - Based on this information, operators might decide what is a routing anomaly that propagated across their network
  - Bounds on events propagation: proposed methodology can be used recursively from observation AS to originating AS

Conclusions

Conclusions

- Most interactions between GEANT and Abilene are reachability events
- Key aspects: network engineering, configuration and peers
- Impossible to understand impact from one AS only
- Better understanding of AS interactions requires finer matching of routing data from multiple vantage points