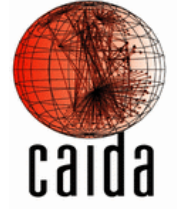

IPv6 AS Relationships, Cliques, and Congruence

Vasileios Giotsas, Matthew Luckie, Bradley Huffaker, kc claffy

vgiotsas@caida.org

UCSD/CAIDA

Passive and Active Measurement (PAM) Conference 2015, New York City, NY, March 19-20, 2015

Business relationships determine the economics of routing

- customer-to-provider (c2p):
  - An AS buys ***transit*** access to the routing table of a better-connected AS
- peer-to-peer (p2p):
  - peers provide mutual access to subset of each others routing table
- Tier-1 clique:
  - A set of of ASes that access the entire routing table through non-provider links
  - Tier-1 ASes need to peer with each other to achieve global reachability

IPv6 AS relationship inference (more) challenging

- Low deployment and traffic mean different economics than IPv4
- IPv6 business policies are less rigorously enforced, leading to more policy violations
- IPv6 graph is not fully connected due to peering disputes between large transit-free providers

IPv6 AS relationship inference (more) challenging

- Low deployment and traffic mean different economics than IPv4
- IPv6 business policies are less rigorously enforced, leading to more policy violations
- IPv6 graph is not fully connected due to peering disputes between large transit-free providers

Limited academic or commercial efforts for IPv6 relationship inference

Adapting an IPv4 algorithm to IPv6

- Increasing similarity between IPv4 & IPv6 topologies and paths[1]:
  - The fraction of IPv6 Enterprise Customer ASes converges to IPv4
  - Dual-stack paths only 5% identical in 2007 → 50% in 20012

- IPv6 traffic is maturing[2]
  - 13-fold increase in IPv6 traffic between 2010 and 2013
  - IPv6 traffic mix more similar to IPv4 than in the past

1. Dhamdhere, A., et al.: Measuring the deployment of IPv6: Topology, routing and performance. IMC 2012
2. Czyz, J., et al.: Measuring IPv6 adoption. SIGCOMM 2014

Abbreviated IPv4 relationship inference

1. Sanitize input BGP paths (remove loops, reserved ASNs, IXPs)
2. Rank ASes by *transit degree*
3. Infer clique at the top of the AS topology
4. Remove path poisoning
5. Infer c2p relationships
   - neighbour passes route to a provider
   - neighbour is in clique and passes route to another clique AS
   - Infer c2p relationships between stub and clique Ases
6. Infer all other links as p2p

Abbreviated IPv4 relationship inference

1. Sanitize input BGP paths (remove loops, reserved ASNs, IXPs)
2. **Rank ASes by *transit degree***
3. **Infer clique at the top of the AS topology**
4. Remove path poisoning
5. Infer c2p relationships
   - neighbour passes route to a provider
   - neighbour is in clique and passes route to another clique AS
   - Infer c2p relationships between stub and clique Ases
6. Infer all other links as p2p

IPv4 tier-1 clique inference

- Sort ASes by transit degree (TD):

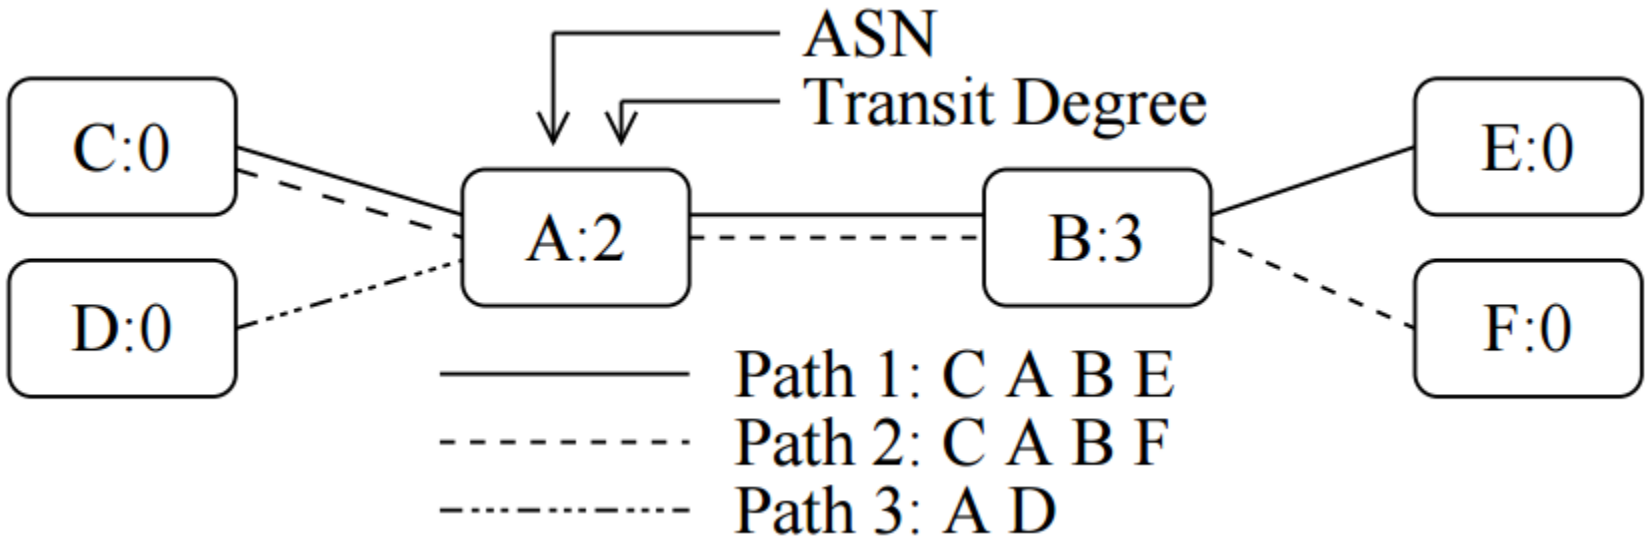

- Apply Bron-Kerbosch algorithm to find maximal clique starting from the 10 largest ASes (SEED) in terms of transit degree

Problems with IPv6 clique inference

- IPv4 ASes with the largest TD have *restrictive* peering policies

  ➔ form cliques *only* with other large ASes

BUT

- ✗ In IPv6 ASes with the largest TD have *open* peering policies, form large peering meshes ASes of varying sizes, often not transit-free
- ✗ Seed ASes may have partial IPv6 reachability due to peering disputes
- ✗ IPv6 topology is more dynamic, making transit degree a volatile metric

IPv6-specific seed requirements

- **Seed ASes should not have open peering policy:**
  - Avoid ASes that aggressively establish peering meshes
- **Seed ASes should have ≥ 90% reachability degree:**
  - Fraction of the BGP-visible IPv6 address space that an AS announces
  - Avoid partitioned ASes
- **Number of seed ASes should reflect the topology size:**
  - Size of IPv6 topology significantly smaller than IPv4
  - Reduce seed size to be proportional to the topology size

IPv4 relationship inference algorithm

1. Sanitize input BGP paths (remove loops, reserved ASNs, IXPs)
2. **Rank ASes by *transit degree***
3. **Infer clique at the top of the AS topology**
4. Remove path poisoning
5. Infer c2p relationships
   - neighbour passes route to a provider
   - neighbour is in clique and passes route to another clique AS
   - ~~**Infer c2p relationships between stub and clique ASes**~~
6. Infer all other links as p2p

IPv6 relationship inference algorithm

1. Sanitize input BGP paths (remove loops, reserved ASNs, IXPs)
2. **Rank ASes by *transit degree***
3. **Infer clique at the top of the AS topology**
4. Remove path poisoning
5. Infer c2p relationships
   - neighbour passes route to a provider
   - neighbour is in clique and passes route to another clique AS
   - ~~**Infer c2p relationships between stub and clique ASes**~~
6. Infer all other links as p2p

IPv6 clique members can have open peering policy, peer with stub ASes

Three validation datasets

- **BGP Communities** that denote relationship type
- **RPSL** objects that are used to store routing policies in IRRs
- **Local Preference** values for Hurricane Electric
- Validation data cover ~25% of the visible IPv6 AS links for July 2014

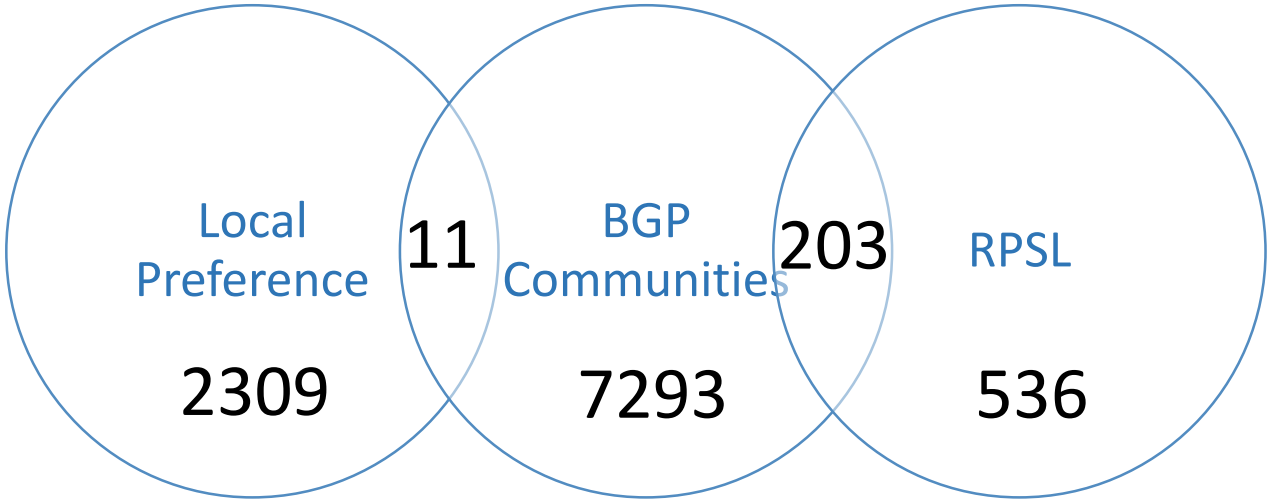

Validation results

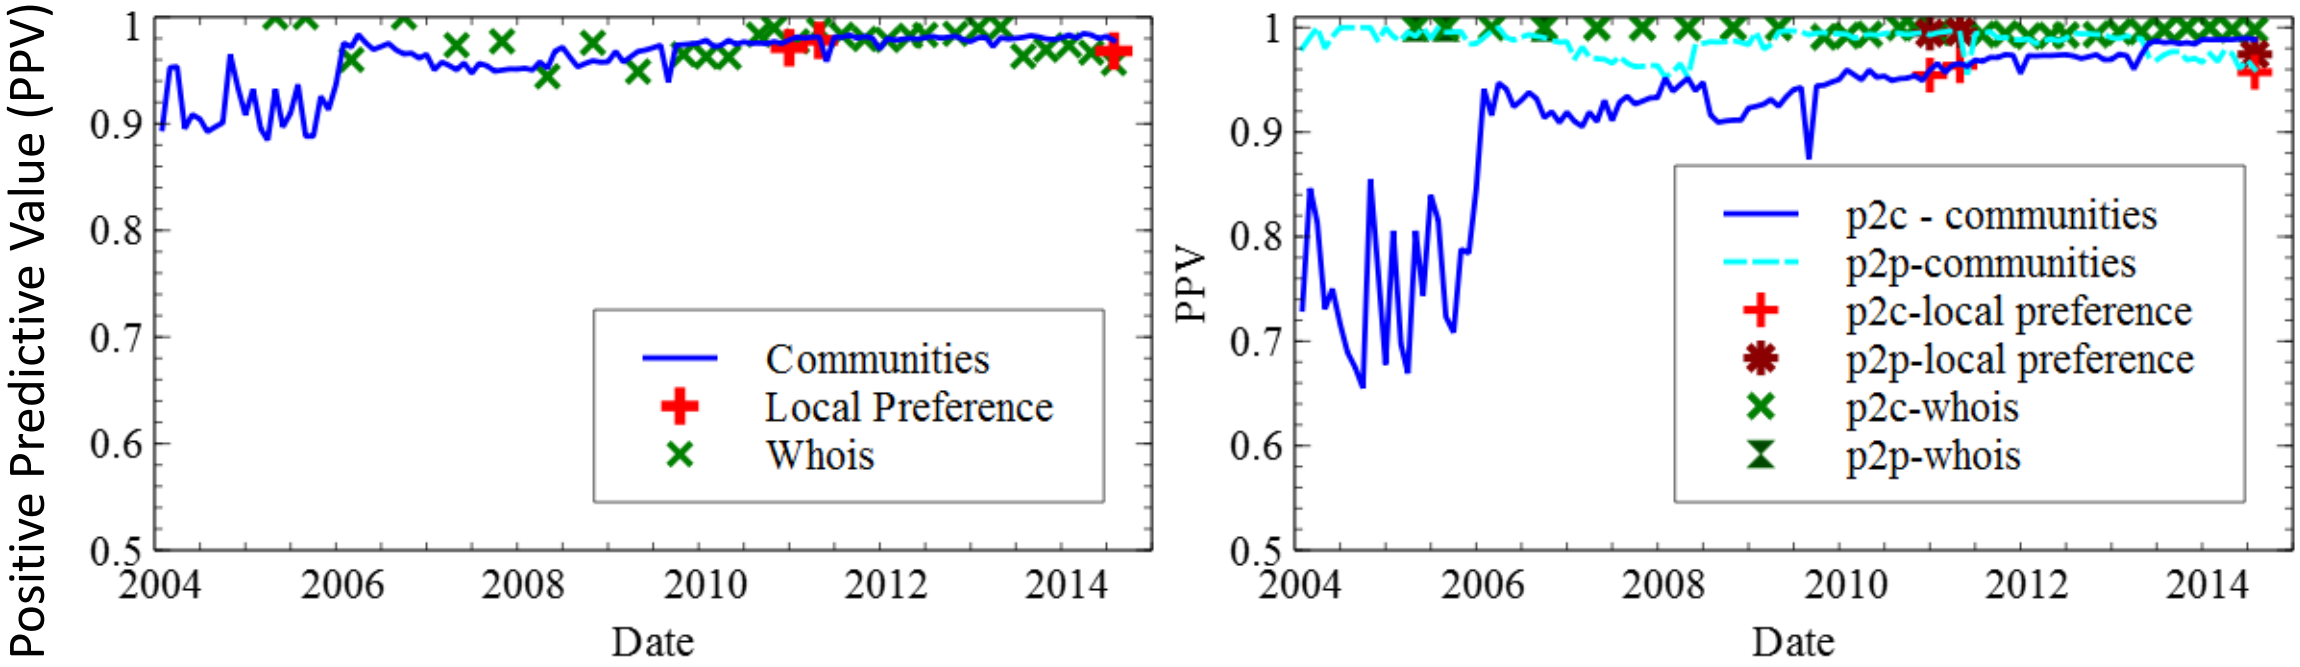

- PPV ≥ 90% for all three datasets across all the snapshots
- PPV increases along time

Validation results

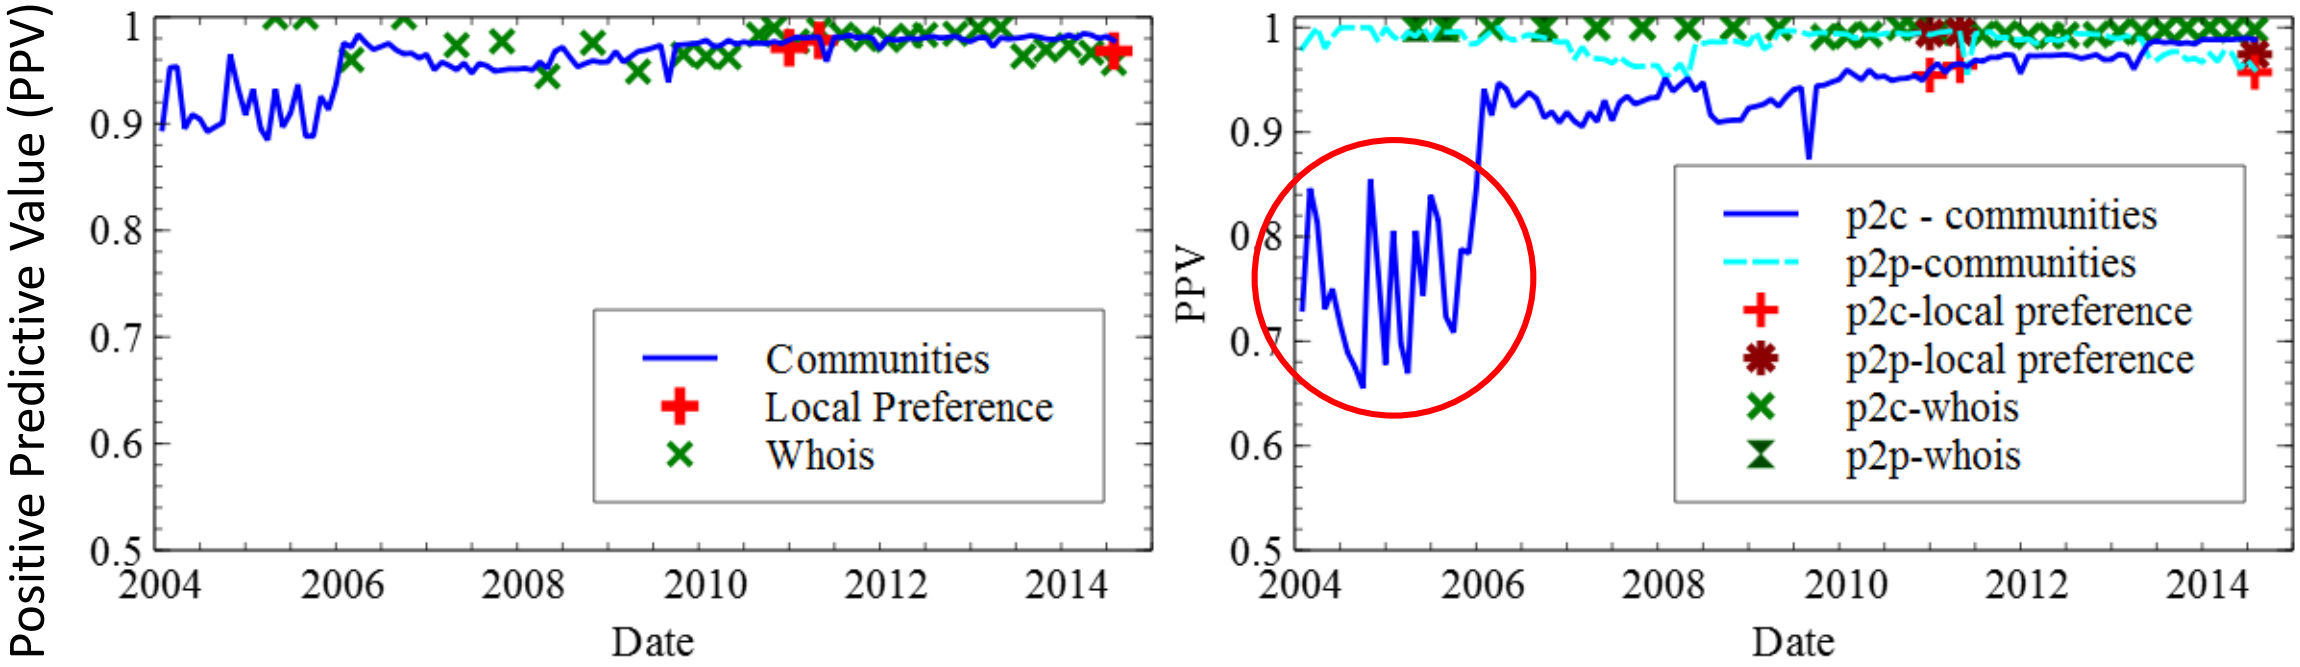

- For small topology sizes mis-inferences for a single AS can affect the overall PPV, but smoothes out as topology grows

Convergence of IPv4/IPv6 Tier-1 cliques

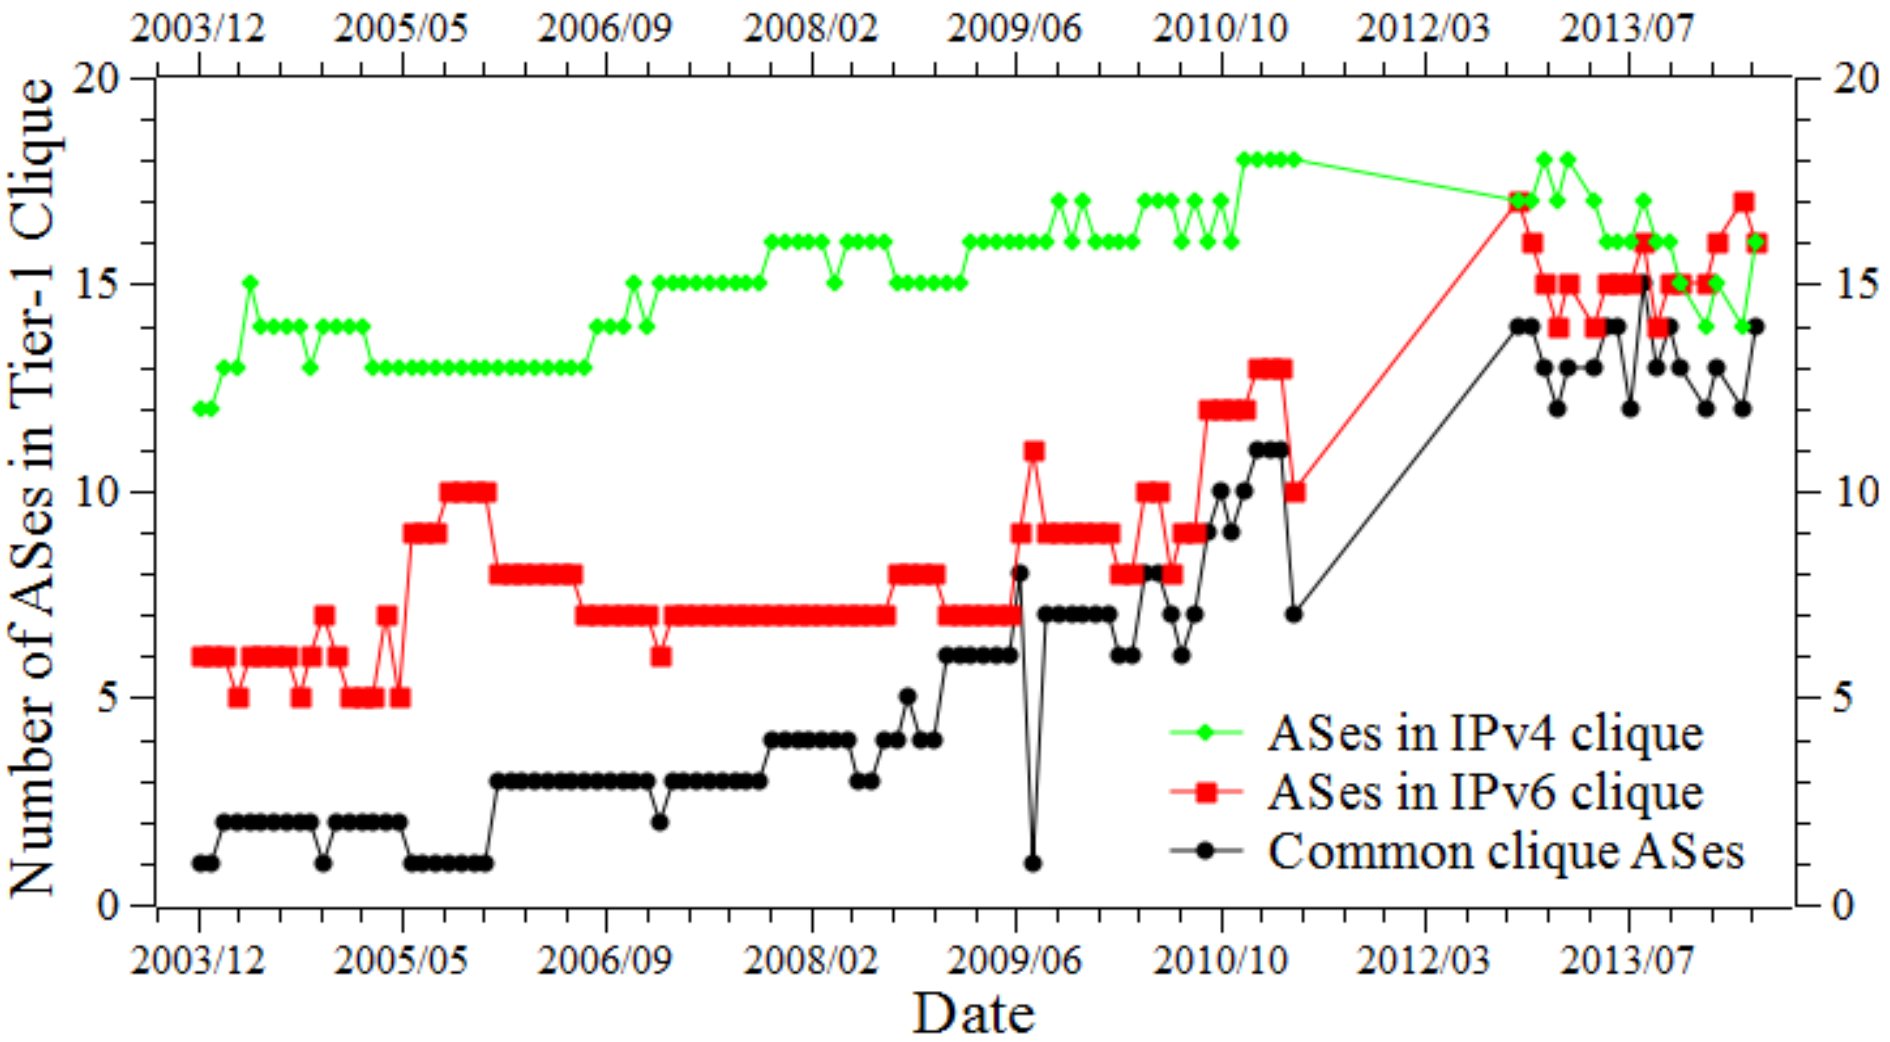

Decreasing relationship disparity

- AS links that appear in both IPv4 and IPv6 topologies are called dual-stack links
- If the relationship type of dual-stack links is the same in both IP version we call the ***congruent***, otherwise ***disparate***

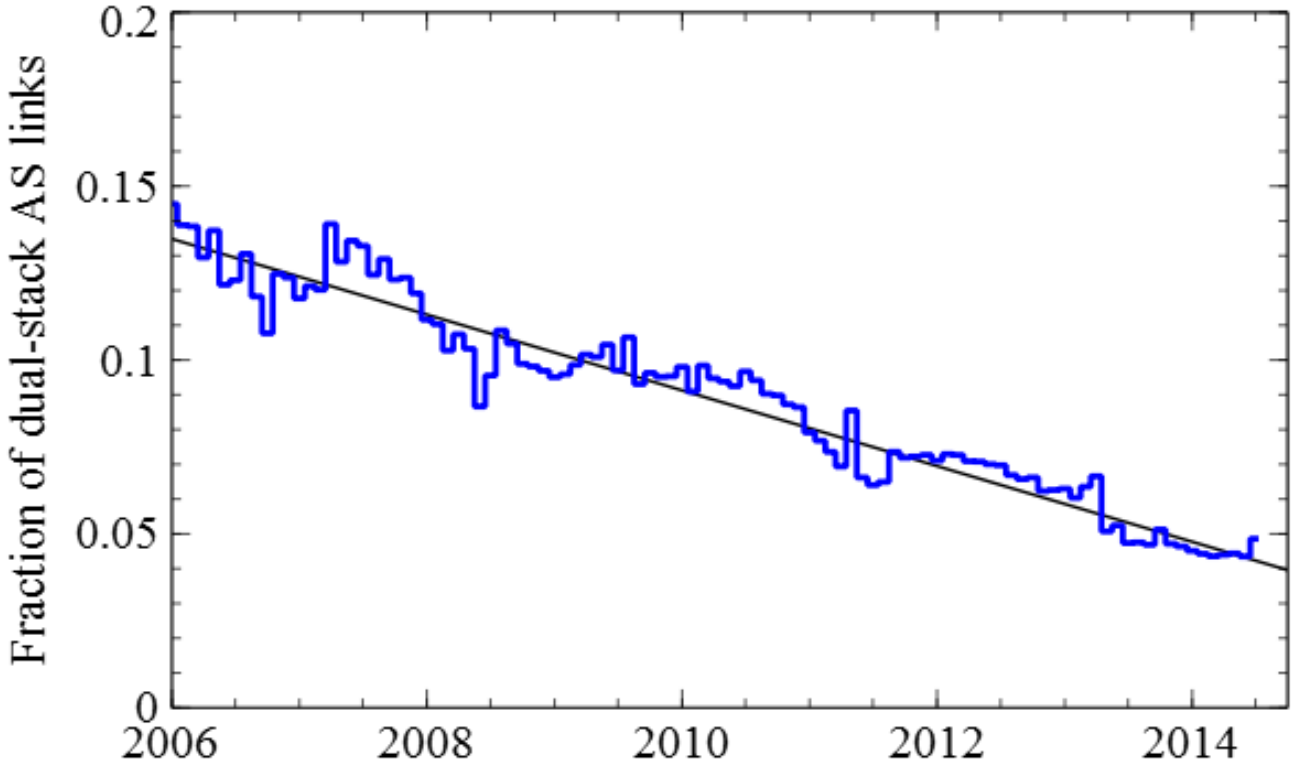

Use customer cones to asses the influence of Tier-1 ASes

- Customer cone is the set of ASes that are reached from a given AS following only customer links in the BGP paths we observe
- *A*'s customer cone contains *A*, plus *A*'s customers, plus its customers' customers, and so on.

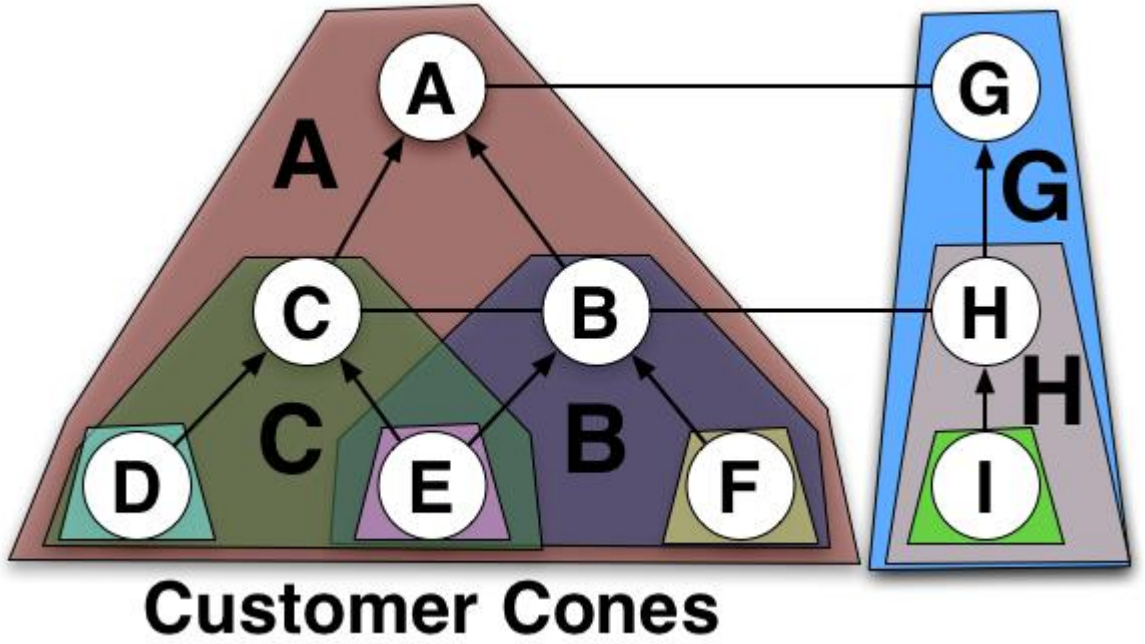

Tier-1 clique ASes less prominent in IPv6

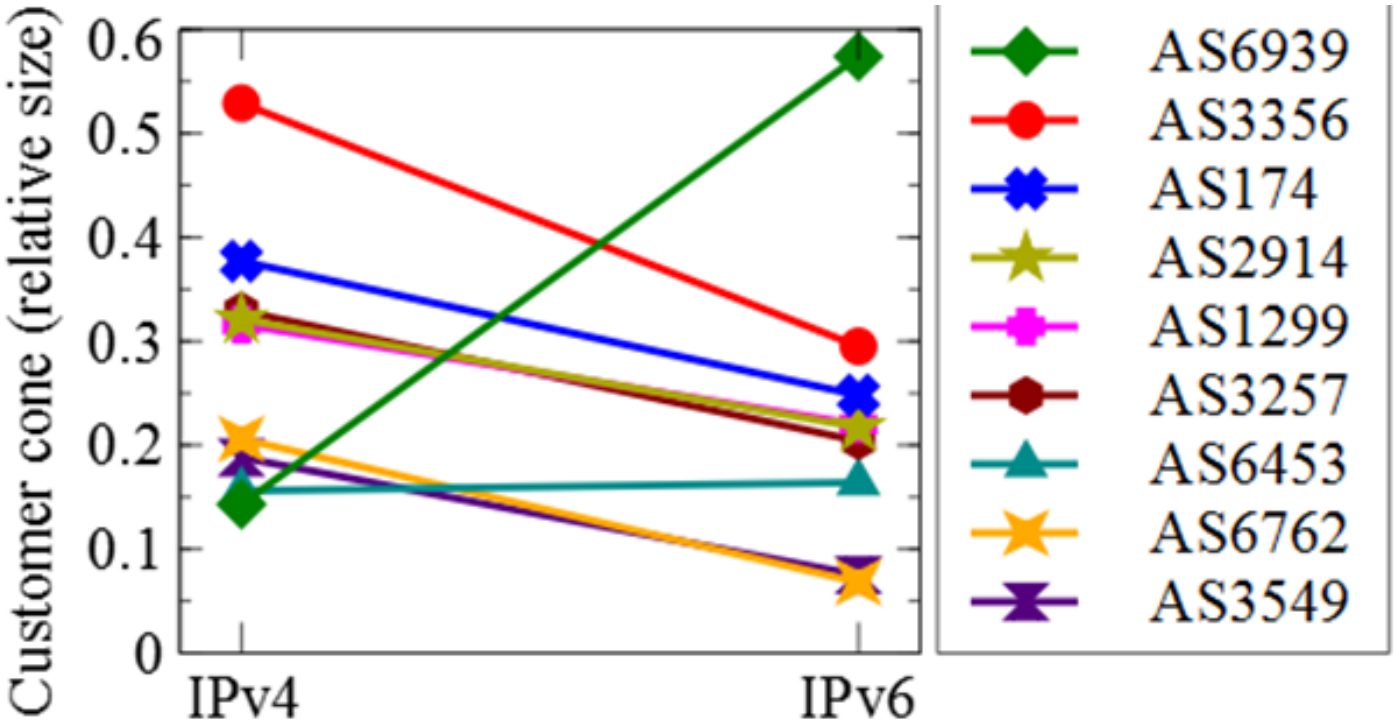

- Only AS6939 (Hurricane Electric) increased its customer cone size in IPv6
- AS6939 contributes > 50% of disparate dual-stack relationships

Clique ASes increase their market-share in IPv6

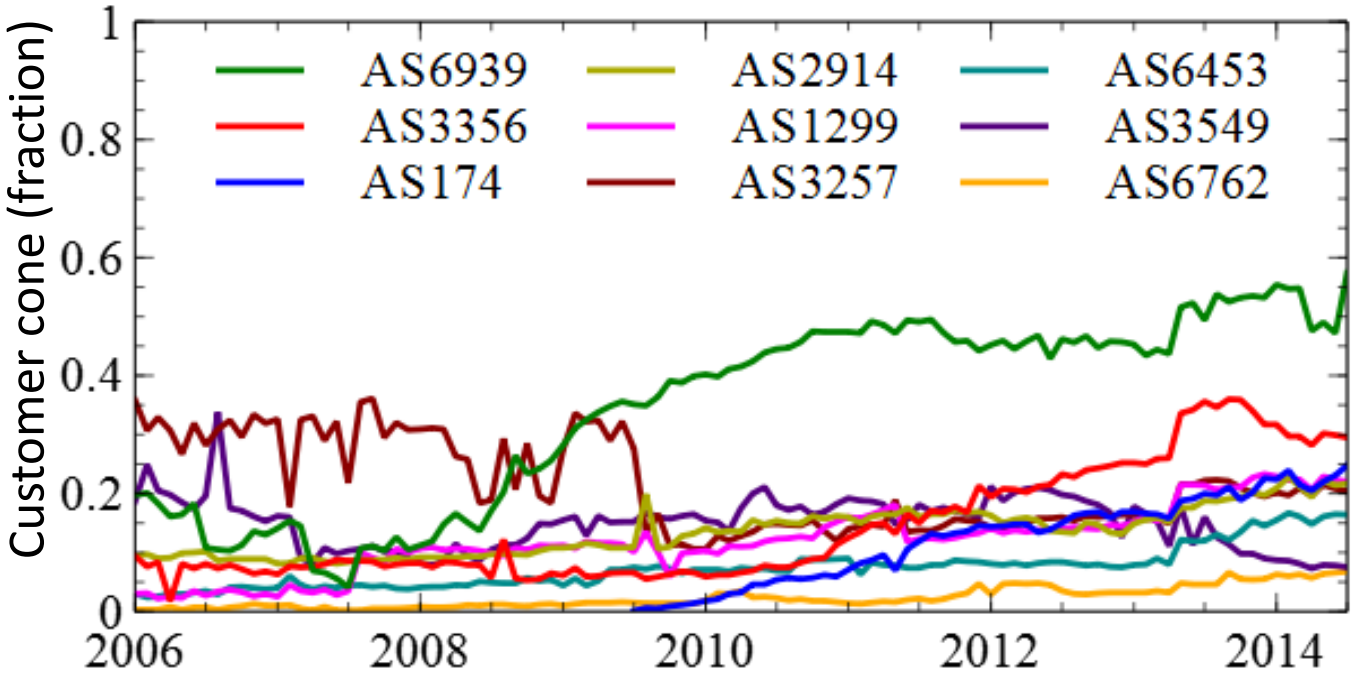

IPv6 Clique

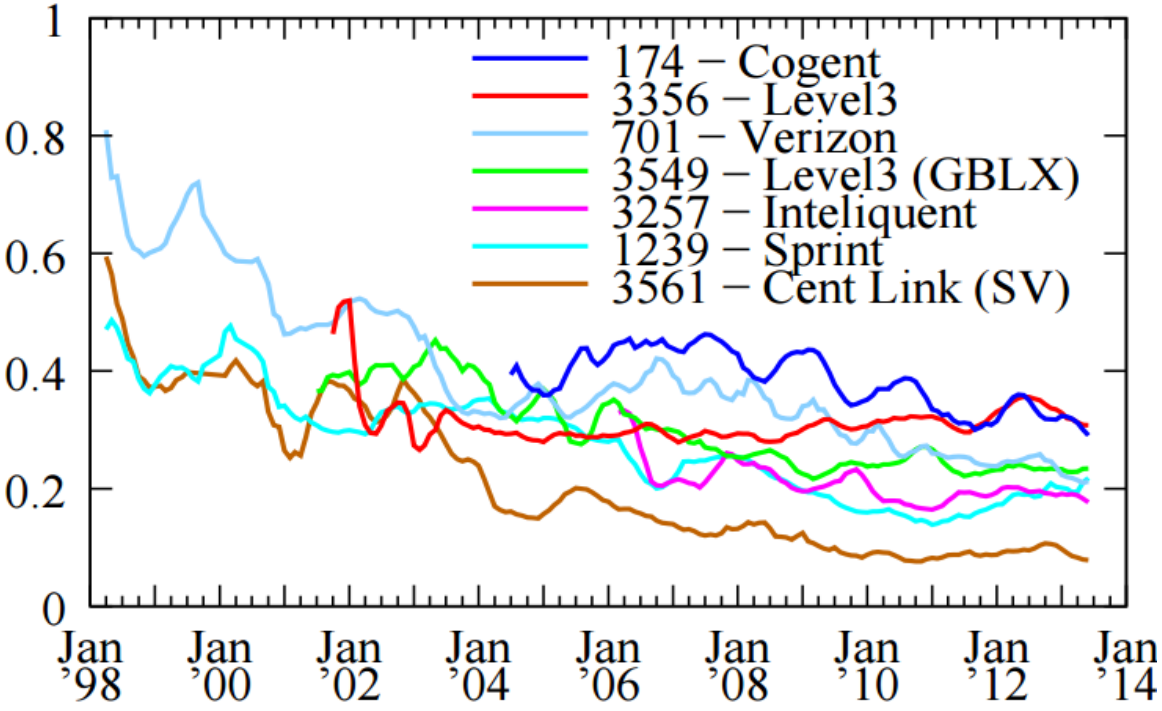

IPv4 Clique[3]

3. Luckie, M., et al.: AS Relationships, Customer Cones, and Validation. IMC 2013

Conclusions

- Validated 25% of visible links against three datasets
  - Positive Predictive Value consistently above 90% over a decade
- Dual-stack relationships increasingly congruent
  - 15% disparity in 2006, to 5% in 2014
- HE is the largest contributor of disparate relationships
  - Largest customer cone in IPv6 topology
- IPv6 transit market small but with trend of growth in contrast to IPv4

http://data.caida.org/datasets/2015-asrank6-data-supplement/data/

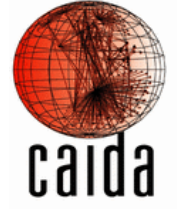

Thank you!

Backup slide 1: Comparison of clique inference techniques

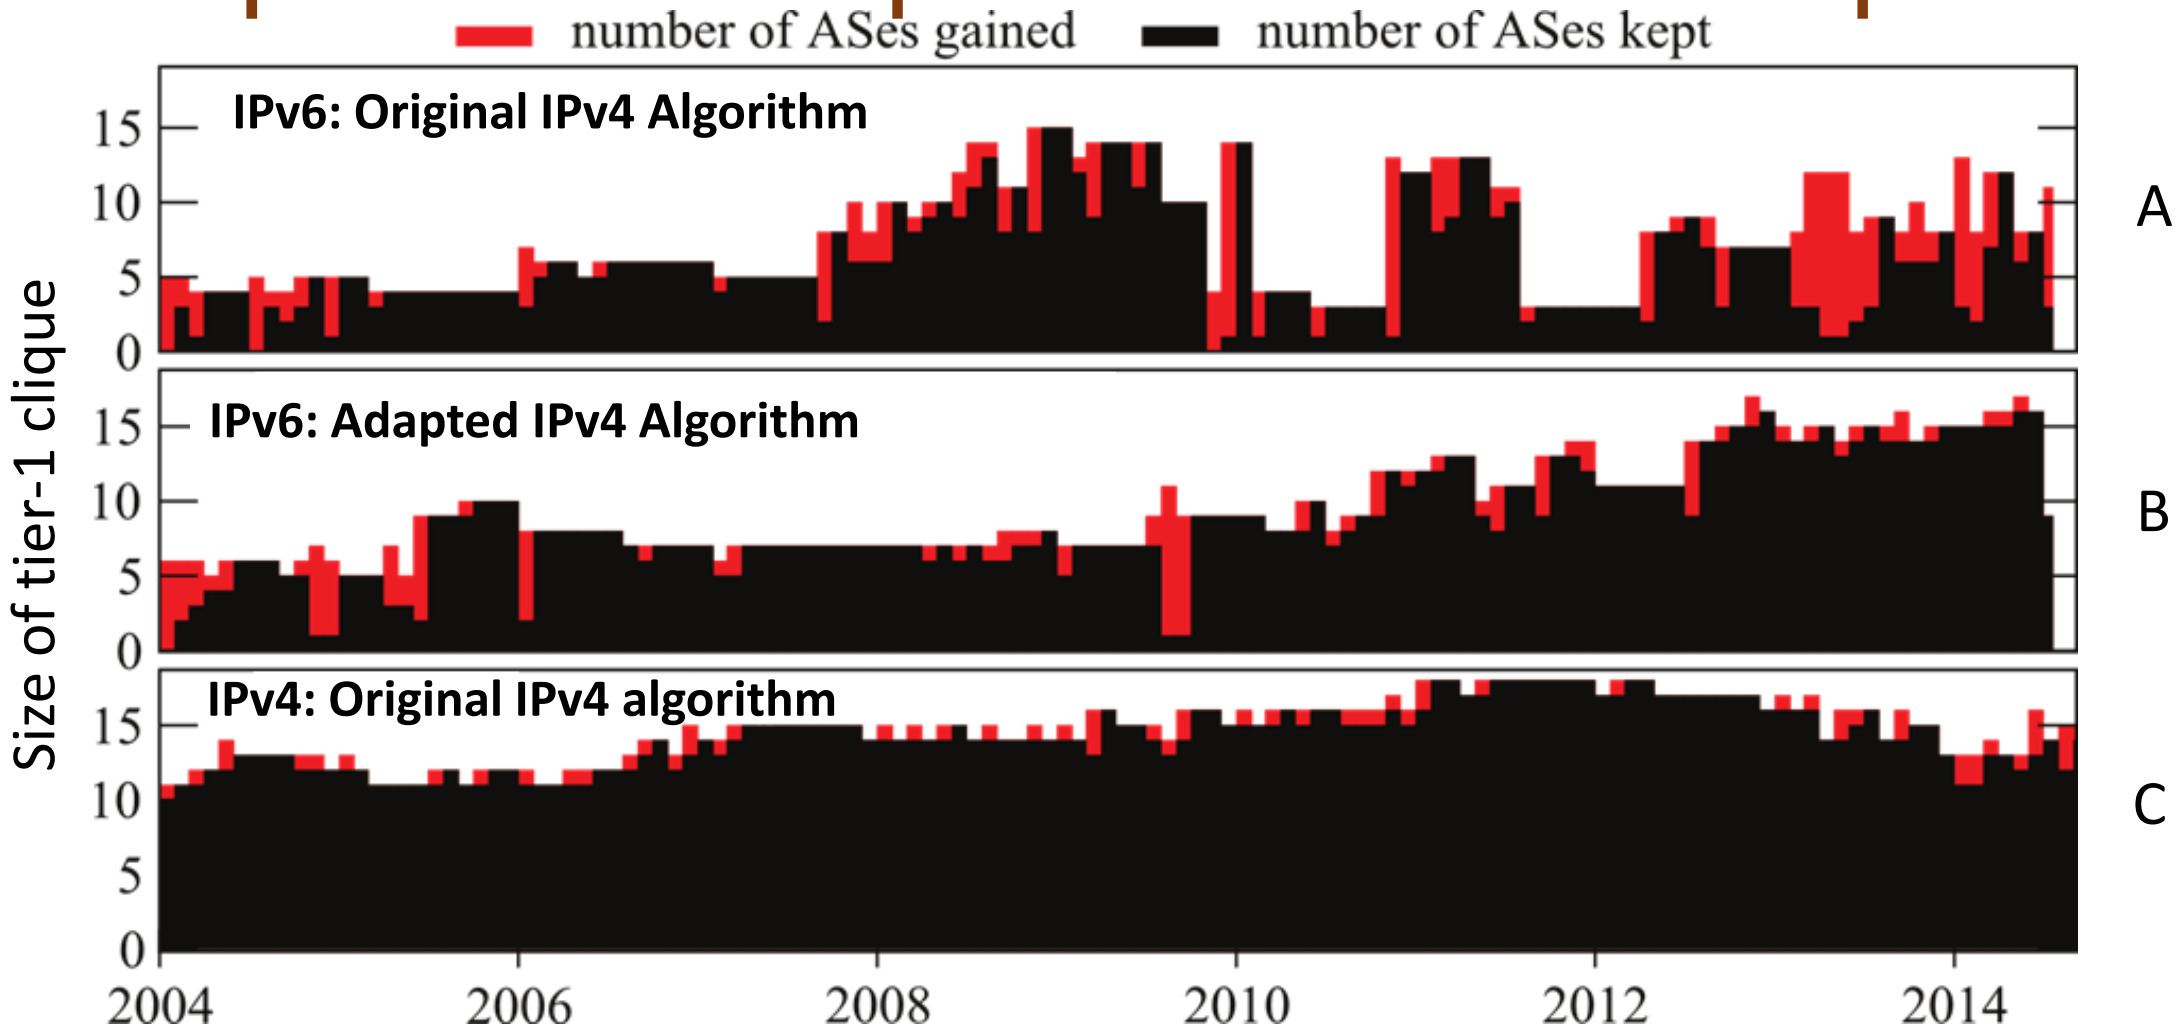

Backup slide 2: AS6939 Local Preference Values

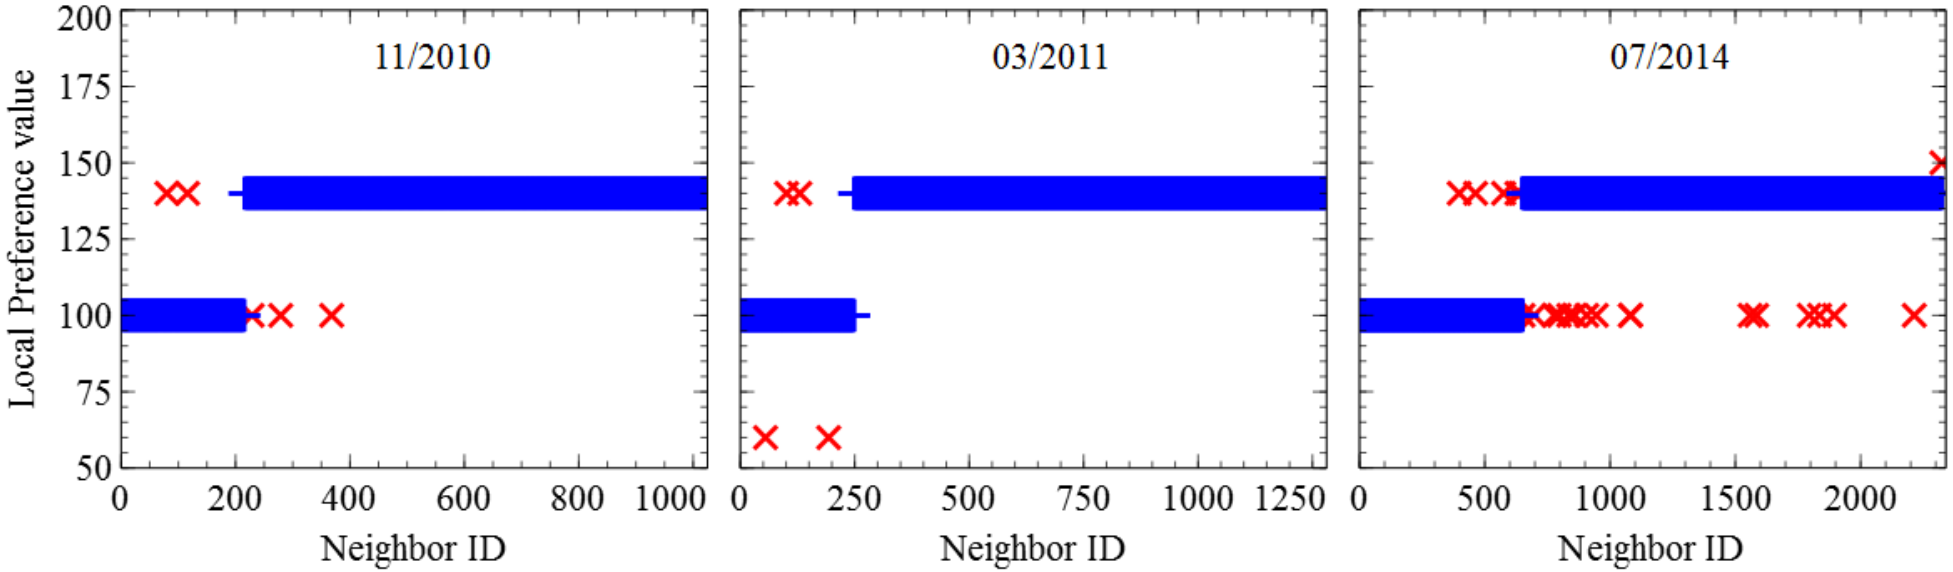

Backup slide 3: Evolution of the IPv6 Clique