BGP and Inferring Autonomous System Relationships.

- by Vaibhav Mhaske

- Guide – Prof. G. P. Saraph

Outline

- BGP routing parameters
- Inference of Autonomous systems relationships.

BGP – Border Gateway Protocol.

- Interdomain routing protocol.
- EBGP and IBGP.
- The route advertised to the neighbour is the optimal path to the destination.

BGP – Border Gateway Protocol.

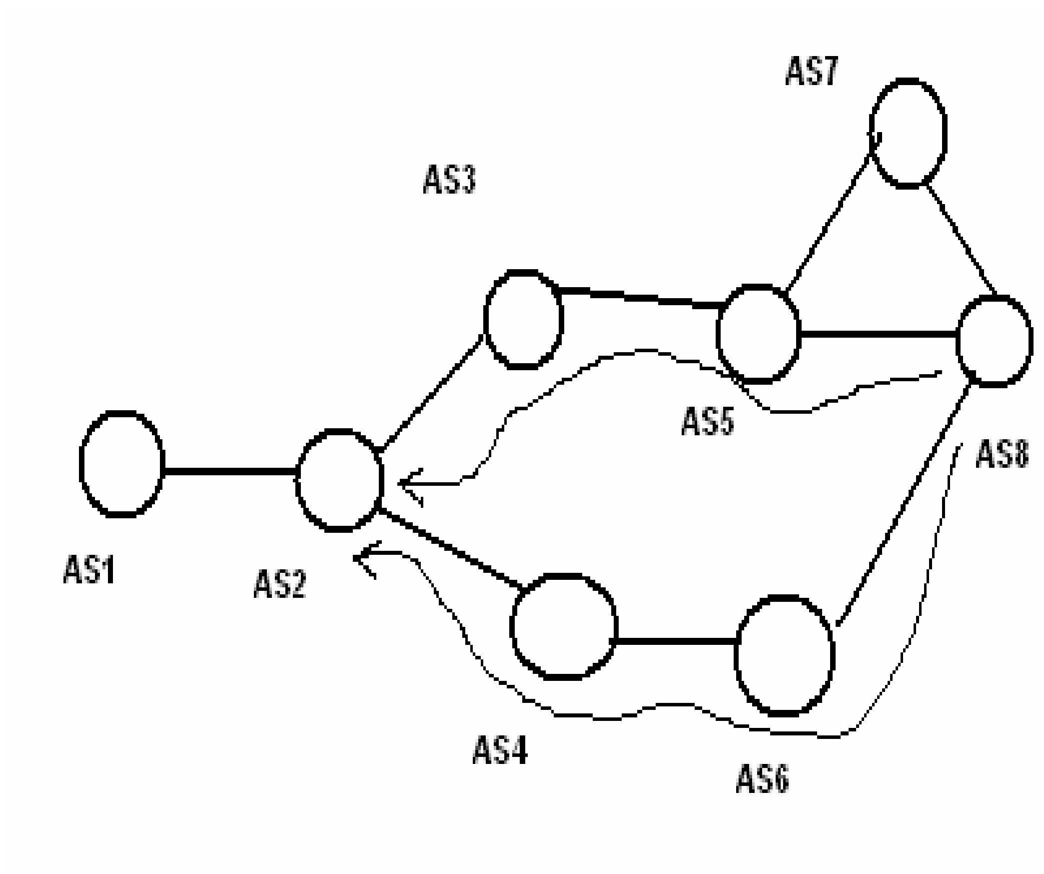

Routing Attributes

- BGP routing tables have more than 1 lakh entries.
- Scalability – routing parameters.
- Determine best route.
- Route selection influenced by routing parameters.

Weight Attribute

- Not advertised to neighbour routers.

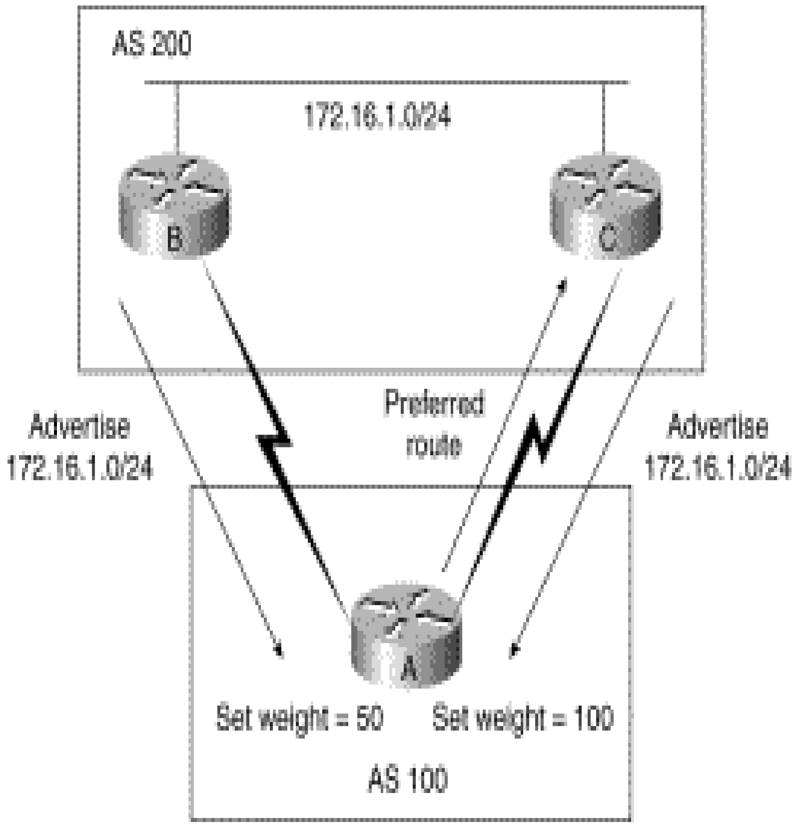

http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/bgp.htm#1020595

Local Preference

- Is used to prefer exit point from an AS.
- It is propagated throughout the local AS.

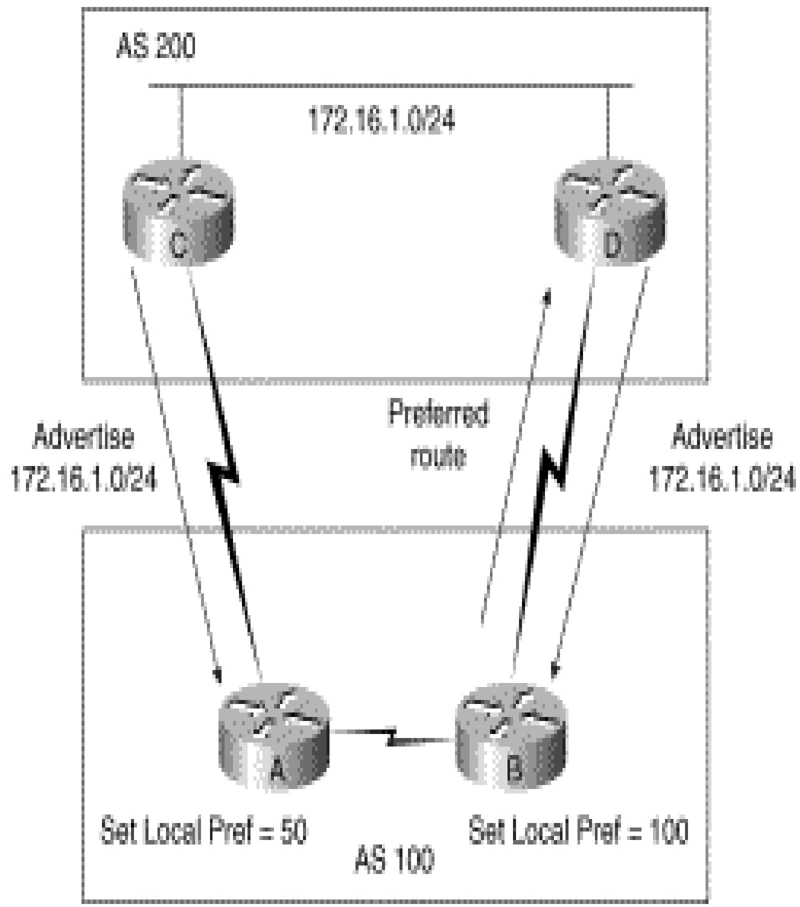

http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/bgp.htm#1020595

Multi Exit Discriminator (MED)

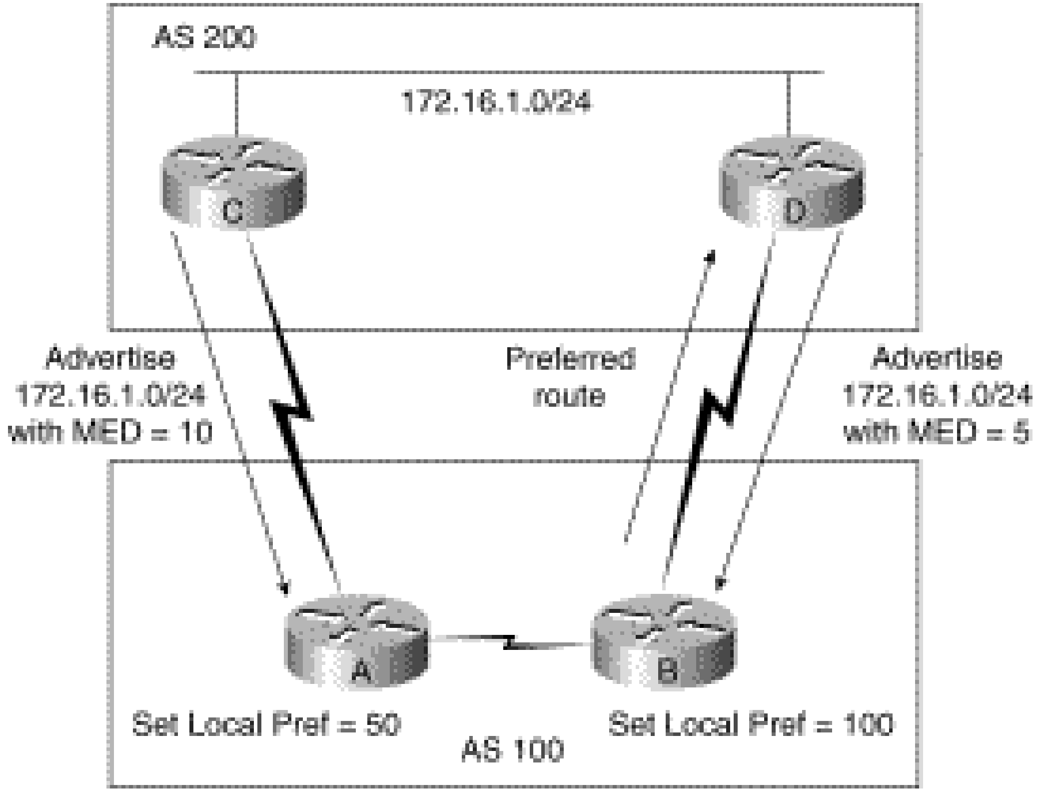

http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/bgp.htm#1020595

Community Attributes

- No export.
- Doesnot advertise it to any other AS.

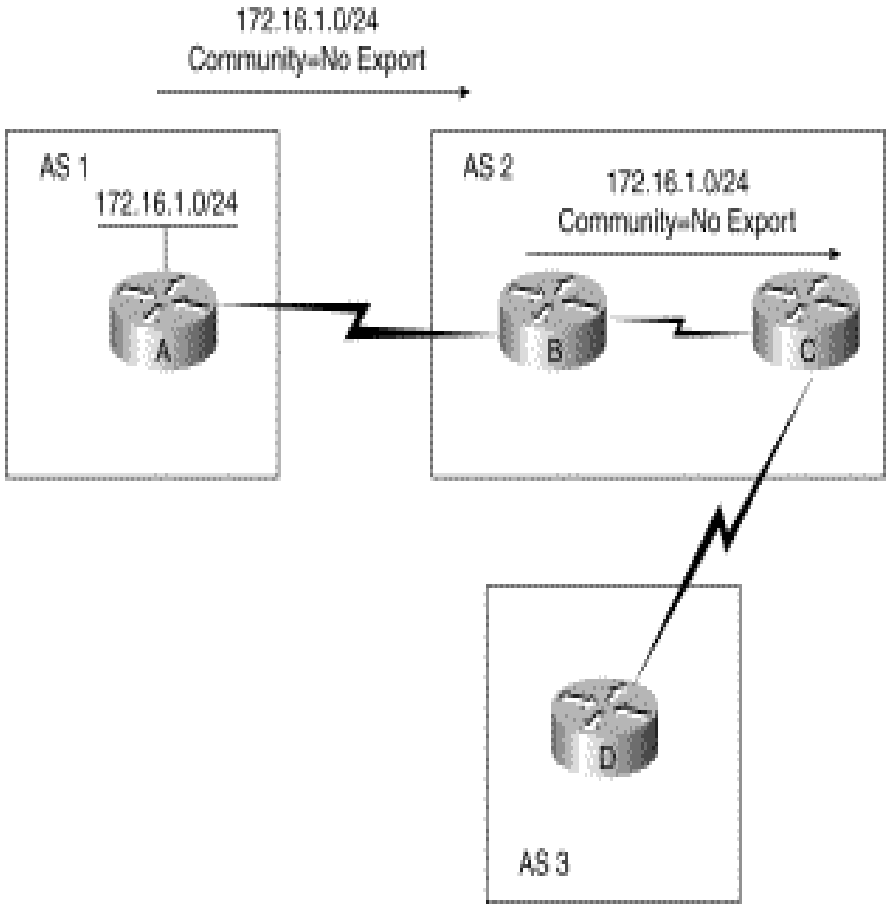

http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/bgp.htm#1020595

Community Attributes

- No advertise.
- Doesnot propagate the route to any other router.

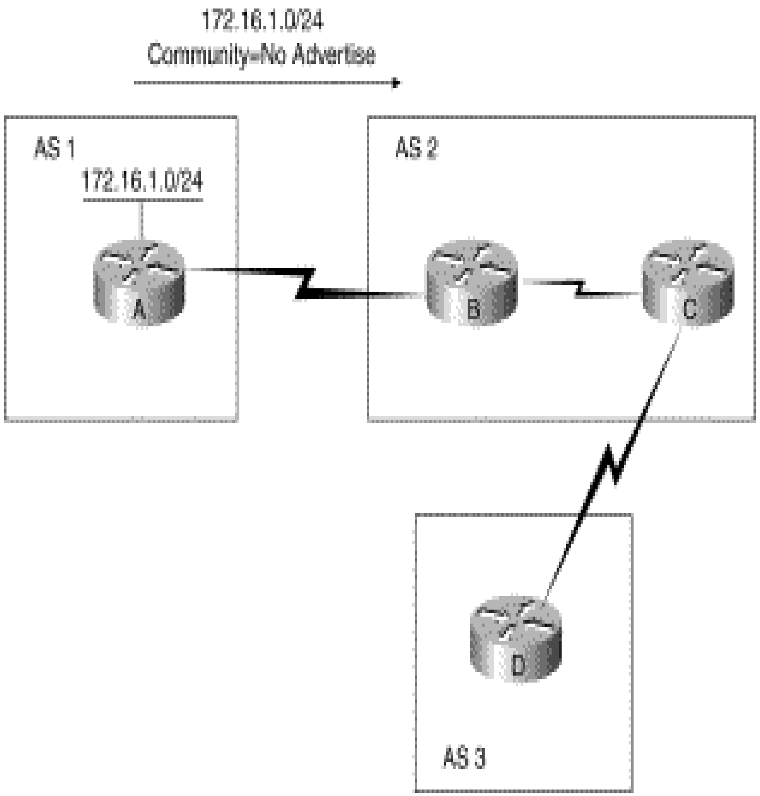

http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/bgp.htm#1020595

Community Attributes

- Internet.
- No limitations for advertisements.

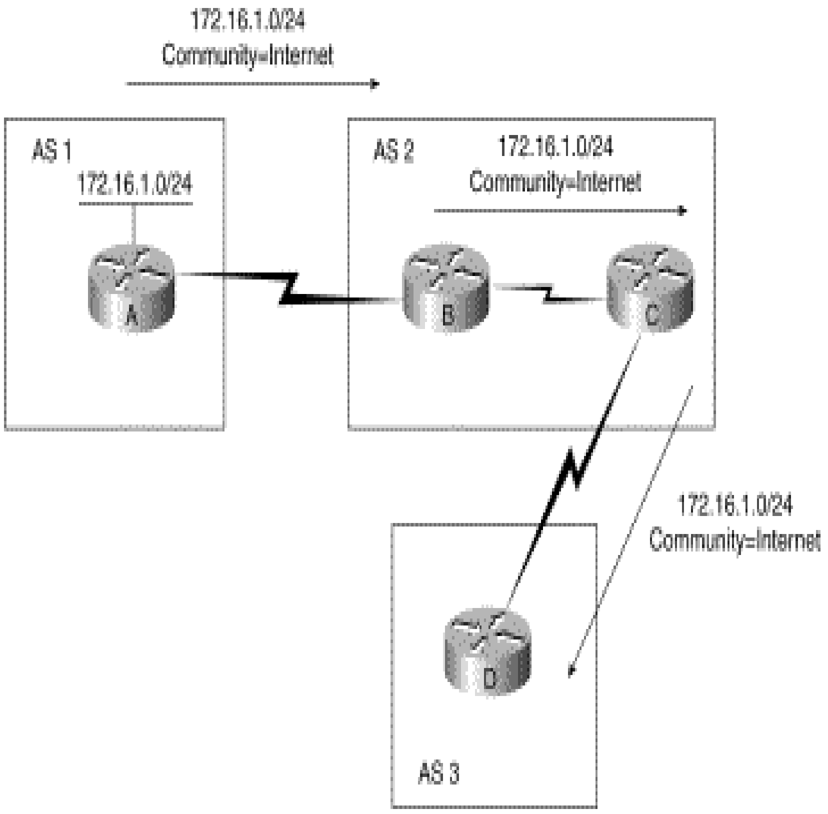

http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/bgp.htm#1020595

BGP path selection

- Largest Weight attribute.
- Largest Local preference.
- Shortest AS path.
- Lowest MED attribute.
- EBGP learned routes are preferred to IBGP learned routes.

Routing policies and BGP routing tables

- Loop avoidance rule : reject routes having its own AS number.
- This avoids cycle in the AS path.

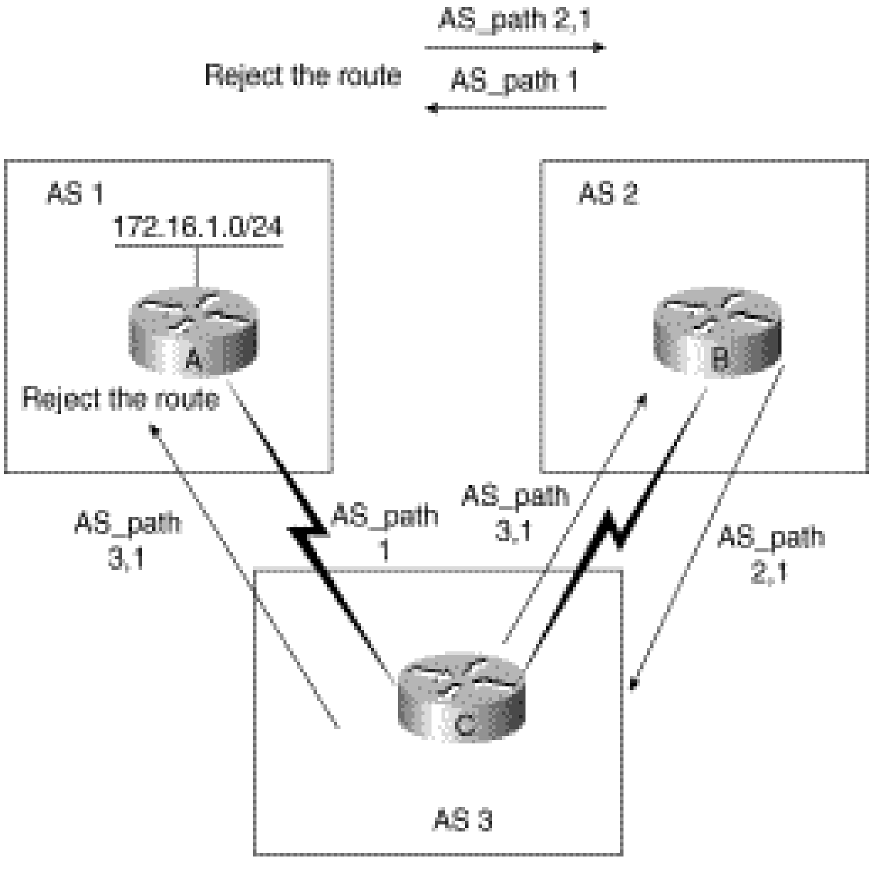

L. Gao. On inferring autonomous system relationships in the Internet. www.caida.org, www.caida.org, June 2002.

AS graph.

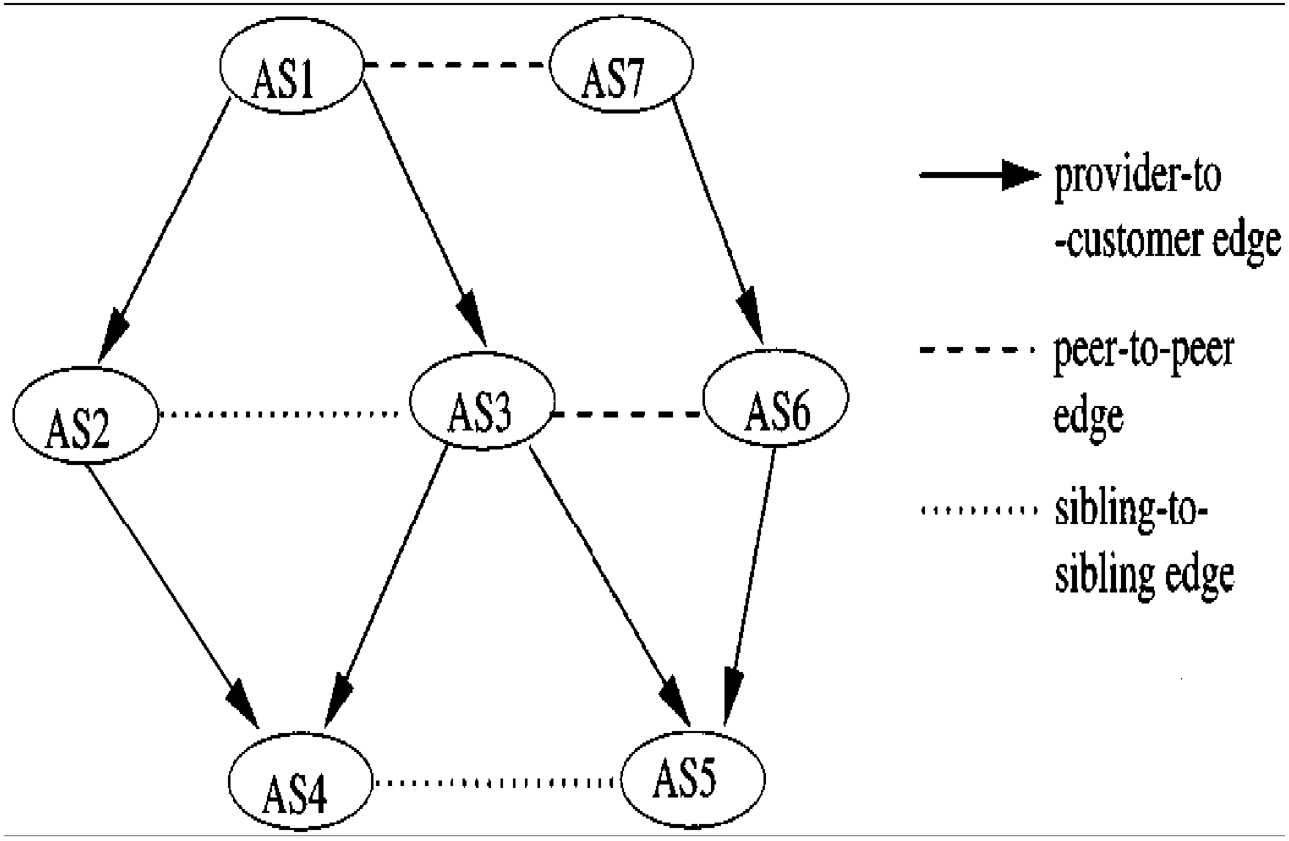

L. Gao. On inferring autonomous system relationships in the Internet. www.caida.org, www.caida.org, June 2002.

AS graph

Routes

- Each AS system sets up its export policies according to its relationship with its neighbouring ASes.
- Routes : customer, provider or peer routes.
- r.AS path : $(u_1,u_2,u_3,\ldots,u_n)$, r is a route.
- (u_i,u_{i+1}) is sibling – sibling edge for all $i<j$ and (u_J,u_{J+1}) is provider – customer edge (customer-provider or peer-peer), then r is a customer route(provider or peer).

Selective Export Rule

- Exporting information to a provider,customer,peer and sibling.
- r is provider or peer route, export(v,u)[{r}] = {}
- r is a provider or a peer route, export(v,u)[{r}] =/={}.
- u transits traffic for v only if it's a provider or peer route.

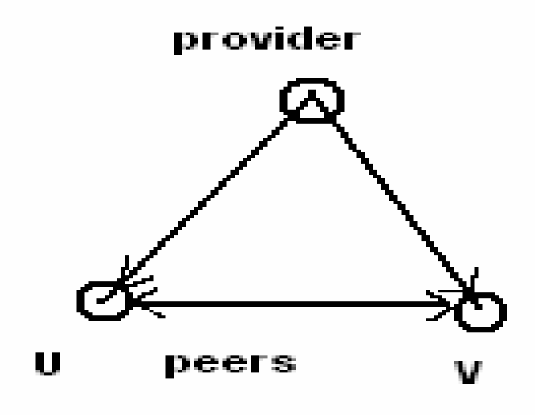

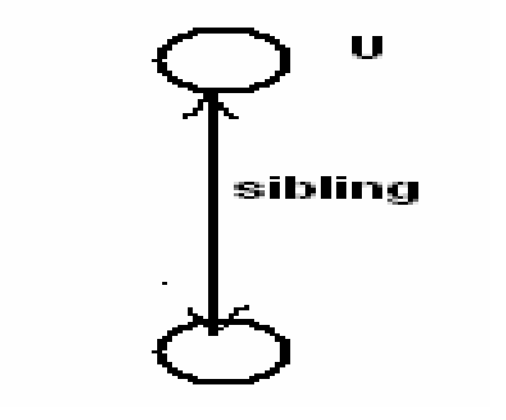

Selective Export Rule

- u_0's BGP routing table contains entry e for prefix d such that e.aspath = $(u_1,u_2,u_3,\ldots,u_n)$.
- u_i selects $(u_{i+1},\ldots,u_n)$ as the best path and exports it to u_{i-1} as the best route.
- Selective export rule and the above property ensures that the as path of the BGP routing table should be valley free.
- Provider to customer edge followed by only provider to customer edge and peer to peer edge followed by provider to customer edge.

Valley free path

- (1,2,6,3) : valley free
- (1,4,5,3) : not valley free

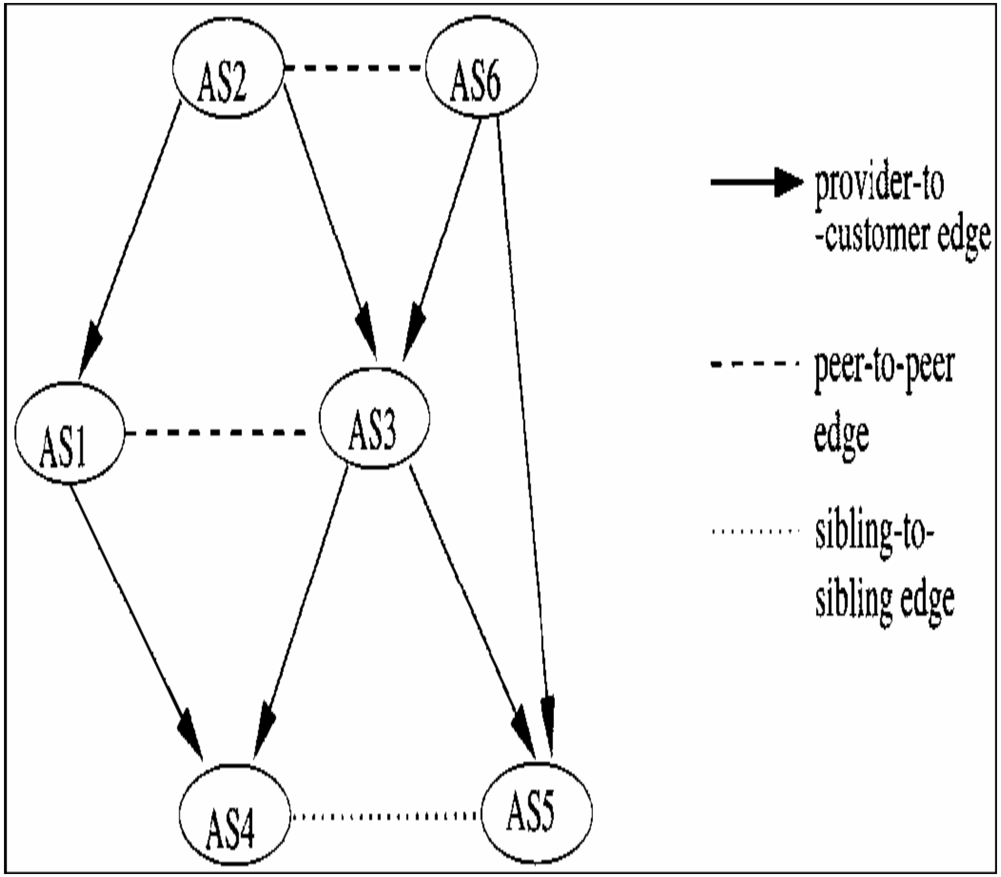

L. Gao. On inferring autonomous system relationships in the Internet. www.caida.org,
www.caida.org, June 2002.

AS path pattern in BGP routing table

- 1) an uphill path.
- 2) a downhill path.
- 3) an uphill path followed by a downhill path.
- 4) an uphill path followed by a peer-to-peer edge.
- 5) a peer-to-peer edge followed by a downhill path.
- 6) an uphill path followed by a peer-to-peer edge, which is followed by a downhill path.

Heuristic algorithms

- Inferring provider customer and sibling relationships.
- Inferring peering relationships.
- Input is BGP routing table and Output is AS graph.

Algorithms

- Basic algorithm.
- Refined algorithm.
- as.path (u,w,v).
- Parameter L.

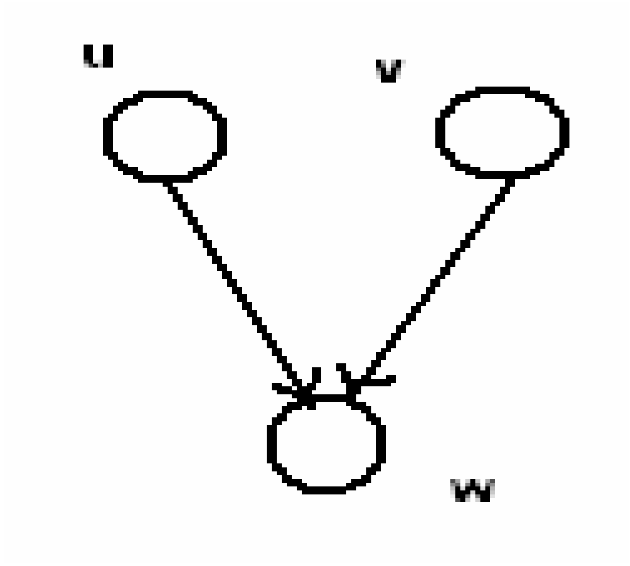

Final algorithm

- Algorithm for inferring peering relationships.
- u,v are peers if they donot transit traffic for each other.
- Top provider can have peering relationship with atmost one of its neighbors.
- Parameter R.

Final algorithm

Final Algorithm:
Input: BGP routing tables
Output: Annotated AS graph G

Phase 1: Use either Basic or Refined algorithm to coarsely classify AS pairs into provider-customer or sibling relationships

Phase 2: Identify AS pairs that can not have a peering relationship

For each AS path $(u_1, u_2, \ldots, u_n)$,
  find the AS u_j such that degree[u_j]=$\max_{1 \le i \le n}$ degree[u_i]
  for $i = 1, \ldots, j-2$,
    notpeering[u_i, u_{i+1}]=1
  for $i = j+1, \ldots, n-1$,
    notpeering[u_i, u_{i+1}]=1
  if edge[u_{j-1}, u_j]$\neq$sibling-to-sibling and edge[u_j, u_{j+1}]$\neq$ sibling-to-sibling
    if degree[u_{j-1}]>degree[u_{j+1}]
      notpeering[u_j, u_{j+1}]= 1
    else
      notpeering[u_{j-1}, u_j]= 1

Phase 3: Assign peering relationships to AS pairs

For each AS path $(u_1, u_2, \ldots, u_n)$,
  for j=1, ..., n-1,
    if notpeering[u_j, u_{j+1}]$\neq$ 1 and notpeering[u_{j+1}, u_j]$\neq$ 1 and
      degree[u_j]/degree[u_{j+1}]$< R$ and degree[u_j]/degree[u_{j+1}]$> 1/R$
      edge[u_j, u_{j+1}] = peer-to-peer

L. Gao. On inferring autonomous system relationships in the Internet. www.caida.org,
www.caida.org, June 2002.

Experimental results

- We run the algorithms for the BGP routing tables from September 27, 1999, January 2, 2000, and March 9, 2000 collected from the RouteViews server.
- From the BGP routing table on September 27,1999, the Basic and Final algorithms infer that among 11288 AS graph edges, there are 10745 provider-to-customer edges, 149 sibling-to-sibling edges, and 884 peer-to-peer edges.

Inference results.

	Total Routing Entries	Total edges	Sibling-to-sibling edges inferred by Basic (Percentage)	Sibling-to-sibling edges inferred by Refined[$L = 1$] (Ignored Entries)	Peer-to-peer edges inferred by Final[$R = \infty$] (Percentage)	Peer-to-peer edges inferred by Final[$R = 60$] (Percentage)
1999/9/27	968674	11288	149 (1.3%)	124 (25)	884 (7.8%)	733 (6.5%)
2000/1/2	936058	12571	186 (1.47%)	135 (51)	838 (6.7%)	668 (5.3%)
2000/3/9	1227596	13800	203 (1.47%)	157 (46)	857 (6.2%)	713 (5.7%)

L. Gao. On inferring autonomous
system relationships in the Internet.
www.caida.org,
www.caida.org, June 2002.

Conclusions and future work.

- Consistency : more than 90.5 % provider – customer edges, less than 1.5 % sibling-sibling and less than 8% peer-peer edges.
- Small percent of peer to peer edges since route view routers peers with only tier 1 providers.
- Increase in sibling sibling edges due to increasing number of complex AS relationships and ISP mergers.
- Improve accuracy for AS relationships.
- ISPs can reduce misconfiguration and debug router configuration files.
- An ISP can scan its BGP routing tables periodically to identify the erroneous routes and inform the originating AS.

References

- L. Gao. On inferring autonomous system relationships in the Internet. www.caida.org, www.caida.org, June 2002.
- http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/bgp.htm.
- Routeviews. http://www.routeviews.org.

- Thank you.