Inter-domain Routing

– The structure of Internet

Qinsi Wang

Computer Science Department, Carnegie Mellon

September 15, 2010

Outline

- Lecture 4: Interdomain Routing;
- L. Gao, On inferring autonomous system relationships in the Internet, IEEE/ACM Trans. Networking, vol. 9, no. 6, December 2001;
- L. Subramanian, S. Agarwal, J. Rexford, and R. H. Katz, Characterizing the Internet hierarchy from multiple vantage points, in Proc. IEEE INFOCOM, June 2002, pp. 618¨C627;

Outline

- the structure of the Internet;
- Inferring the hierarchy of Internet;
- Inferring the relationships of autonomous systems;
- Problem aroused.

Constructive Parts of Internet

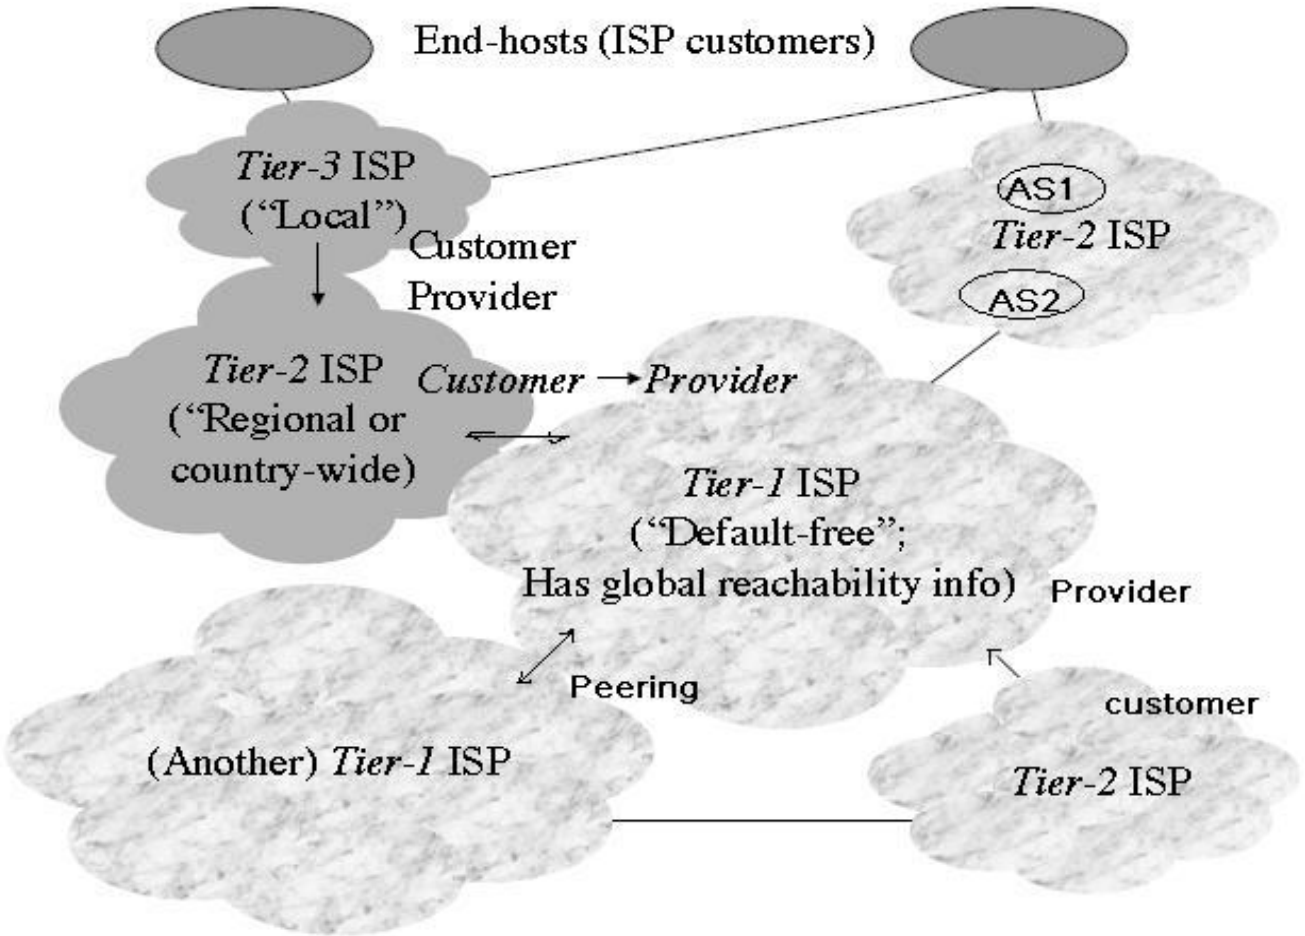

Figure 1: The structure of the Internet

Internet Hierarchy

- Tier 1 ISPs: those that have access to the global Internet Routing Table and don't buy network capacity from other providers.
  UUNet, Cable & Wireless, Sprint, Qwest, Genuity, and AT&T.
- Tier 2 ISPs: generally has a smaller national presence than a Tier 1 ISP and may lease part or all of its network from a Tier 1.
  America Online, Broadwing, and @home.
- Tier 3 ISPs: typically a regional provider with no national backbone.
  RCN, Verizon, Log On America, and so on.

Internet Hierarchy

- Tier 1 ISPs: those that have access to the global Internet Routing Table and don't buy network capacity from other providers.
  UUNet, Cable & Wireless, Sprint, Qwest, Genuity, and AT&T.
- Tier 2 ISPs: generally has a smaller national presence than a Tier 1 ISP and may lease part or all of its network from a Tier 1.
  America Online, Broadwing, and @home.
- Tier 3 ISPs: typically a regional provider with no national backbone.
  RCN, Verizon, Log On America, and so on.

Autonomous system: a collection of IP networks and routers under the control of one entity.

Types of AS Relationships

- Provider-to-Customer (Customer-to-Provider): the customer pays the provider to gain the access to all (or most) destinations in provider's routing table.
- Peering: two ASes provide mutual access to a subset of each other's routing tables without any financial settlement.
- Sibling: two ASes, belonging to the same commercial organization, sharing all their routing tables.
- Backup: to provide connectivity in the event of a failure.
- Maybe some new comings...

Inferring the AS Relationships

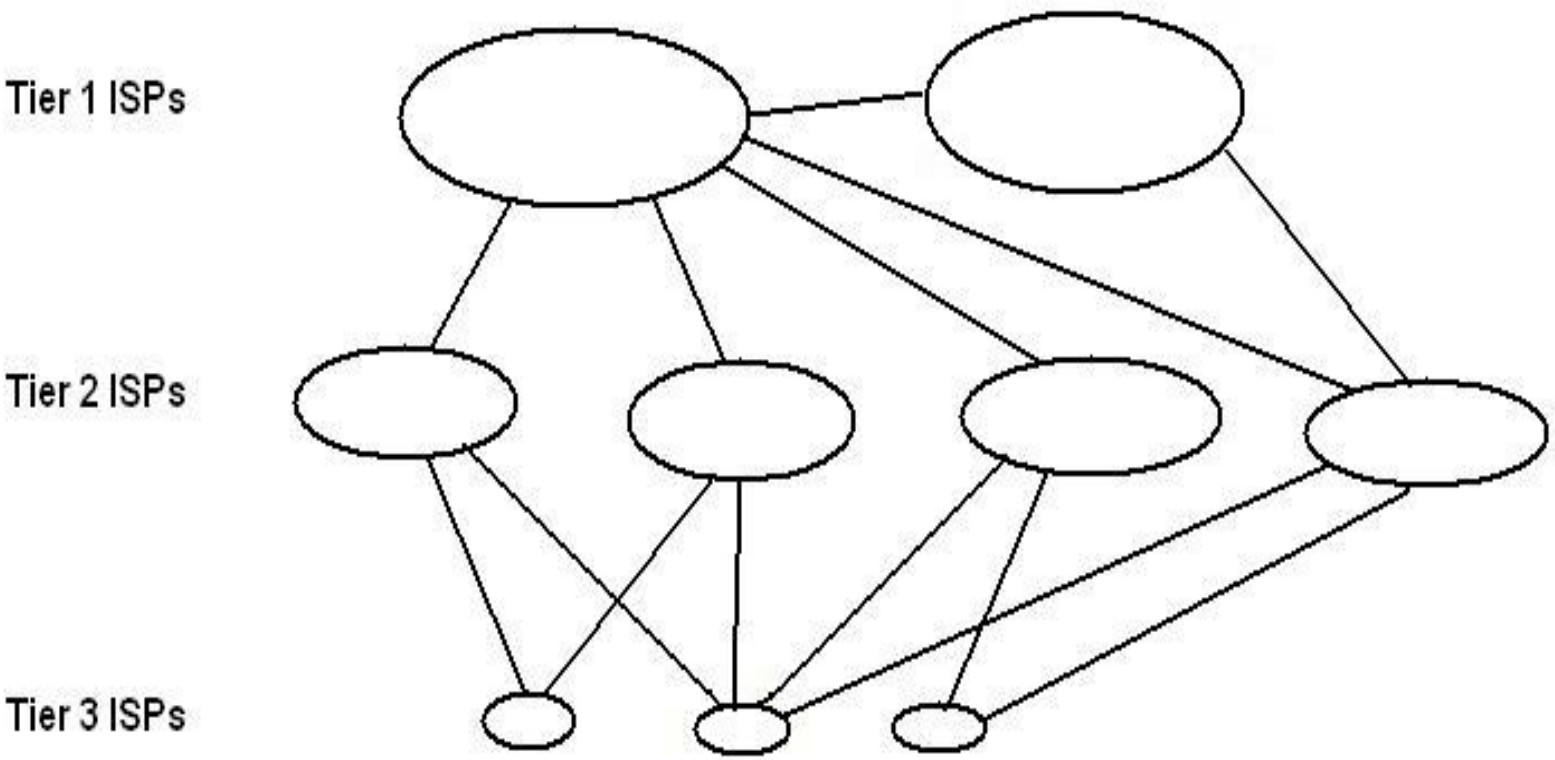

Figure 2: A undirected graph representing the connectivity view of part of the Internet

Inferring the AS Relationships

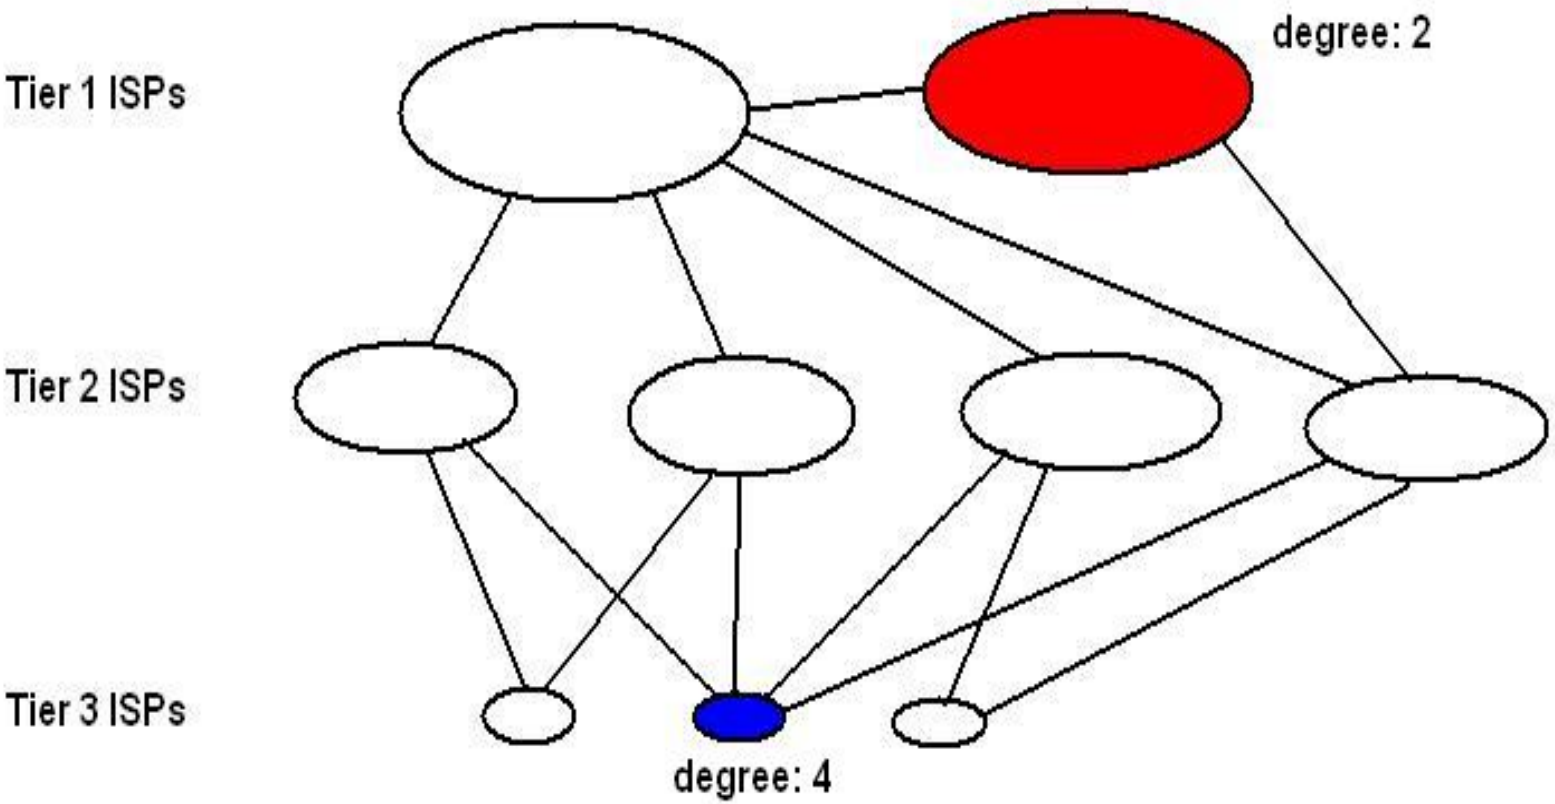

Figure 2: A undirected graph representing the connectivity view of part of the Internet

Inferring the AS Relationships

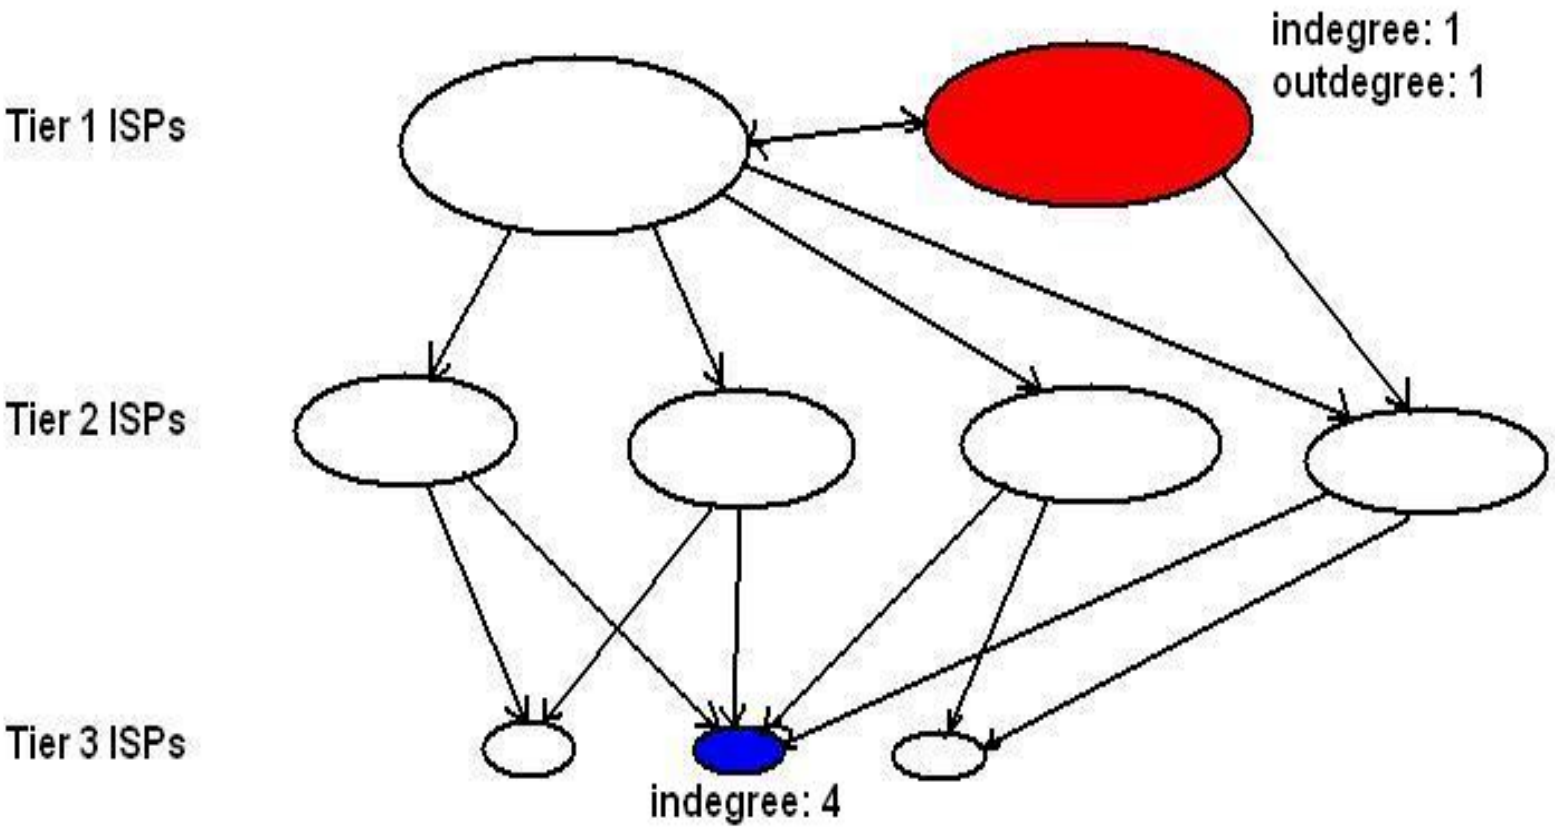

Figure 3: A directed graph representing the connectivity view of part of the Internet

Inferring the Positions of ASes in Internet Hierarchy

The relationships translate directly into policies for exporting route advertisements via BGP sessions with neighboring ASes:

- Exporting to a provider: In exchanging routing information with a provider, an AS can export its routes and routes of its customers, but cannot export routes learned from other providers or peers.
- Exporting to a customer: An AS can export its routes, routes of its customers, and routes learned from other providers and peers to its customer.
- Exporting to a peer: In exchanging routing information with a peer, an AS can export its routes and the routes of its customers, but cannot export routes learned from other providers or peers.
- Exporting to a sibling: An AS can export its routes, routes of its customers, and routes learned from other providers and peers to its customer.

If every AS obeys the export policies, then every advertised path has the following form:

$$((customer-to-provider)^*, (sibing)^*)^*, (peering)^X, ((provider-to-customer)^*, (sibing)^*)^*,$$

Where X equals to 0 or 1.

Gao's Algorithm

- Source data: A collection of AS paths gained from a single AS.
- Algorithm:
  - Step One: to choose the first node with the highest degree to be the boundary between the uphill and downhill portions of the path;
  - Step Two: To assign the relationship of any continuous nodes in the uphill part to be customer-to-provider; To assign the relationship of any continuous nodes in the downhill part to be provider-to-customer;
  - Step Three: To apply step 1 and 2 to other paths in the collection until all the paths have been issued;
  - Step Four: To go through all the paths, and then assign the relationship of two continuous nodes to be sibling if they have both customer-to-provider and provider-to-customer relationships;
  - Step Six: To compare the degree of the boundary node with that of its next node. If the difference is within some threshold, the relationship between these two nodes could be peering.
- Results: High accuracy for provider-to-customer and peering relationships, but comparatively low accuracy for sibling relationships.

Subramanian's Algorithm – The Main Difference

Using different source data: A collection of AS paths gained from multiple vantage points (ASes).

Because of multiple vantage points,

- we could have a more complete view of the whole Internet;
- we could examine the relationship between two ASes from different angles;
- we would have enough conflict paths which can be used to figure out the configuration mistakes;
- ...

However, we should also notice that improper choice of vantage points may introduce "REDUNDANCY".

Problems Raised

The Privacy Issue!

Acknowledgement

Thanks very much for your time!