CSE588 Notes
April 16, 2002
Sam Park

Q: Suppose you are an ISP, or campus or corporation? What intra-domain routing protocol do you use?
Suggestions: OSPF, Centralized arbiter

Q: What considerations do we have?
A: Scalability (communication cost, computation cost)
Single point of failure (for centralized arbiter)

Aside: Its important to distinguish forwarding (the actual delivery of packets) and routing (how to determine forwarding)

Distributed Route Computation
Distance Vector -> exchange tables, keep best alternative

In Class: Demonstration of Distance Vector routing by sending a message between two students in opposite corners of the room. The message propagates through adjacent neighbors.

Illustration of Distance Vector (in ascii…sorry folks)

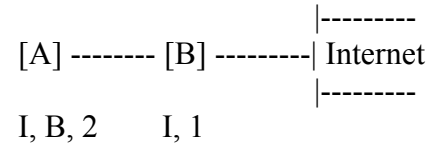

Node A has a route entry to the internet through B, at a cost of 2
Node B has a route entry to the internet at a cost of 1

Possible problem with Distance Vector:
Now if Bs link the internet fails, B might find that A has a route to the internet at a cost of 2, then thus it will update its route table with I, A, 3. Node A then may update itself because B advertises a route to the internet at cost 3 with the entry I, B, 4.

Some more stuff about Distance Vector: Why polls regularly?
Determine if neighbors are alive (hello)
Determine if topology has changed
Don't trust others to have correct info (soft state)
Dynamic routing possibilities (congestion control)

Link State
Send everyone info about your neighbors
Periodic update

Give everyone topology info

Q: When there is a change, should router 1) Send and process, or 2) process and send?
Process and Send. If a neighbor updates before I could (faster processor?), you might receive packets before you are ready to route them.
Consequence is that routing information propagates slower.

Q: How do we do QOS?
Reservation system
Circuit based
A separate control network (like phone system)
Comment on control network: Internet is 1 network, lots simpler.

Inter-domain Routing (BGP)

Scenario: High bandwidth connections between various cities:

[Seattle] ------ [SF] ------ [Denver] ------ [Chicago]
OC48 OC48 OC48

Now suppose a T1 line is added between Seattle and Chicago.
If routing decision was based on number of AS domains traveled, the T1 line would be used, although it's the lowest bandwidth connection.
The costs per link could be calculated using a 1/capacity. This also has its limitations.

Miscellaneous comments: A lot was discussed about ATT and Sprint being the backbone providers of the internet. They have a near parallel network in major areas, and they can exchange information free of charge. Most of the rest of the internet is various organizations that must pay for connection to the backbone which is charged by volume. Several have tried to build their own backbone (360, Qwest, Level3, Global Xing) but have not yet succeeded because they don't have enough of a customer (network) base.

BGP Policy "Knobs"
Link selection (possibly managing who I am buying from)
Advertisement (which prefixes do I advertise)
Preferences (Which links to use for outgoing traffic. Ex: INET 2)
Meds
Community
Peers

Additional Comments: BGP policies are decided at an organizational level.
For instance, if an AS is connected to ATT and Sprint, they can choose to prefer one or the other to try to minimize billing. In the case of INET2, INET2 is intended only for university to university communications, so a BGP policy can decide based on destination to use INET2, and all others go the some other backbone.