

Routing

Brighten Godfrey
cs598pbg Sept 28 2010

slides ©2010 by Brighten Godfrey unless otherwise noted



Dealing with difficult readings



- **Readings are difficult to understand**
 - Next time you read papers on this topic, they will make more sense
 - Ask questions!
- **Readings are difficult to criticize in the reviews**
 - Goal is to **think critically** about the paper, not to write the definitive judgement of the work
 - This is part of the process of understanding!
- **But, we'll introduce upcoming topics to provide some context before you read**



Choosing paths along which messages will travel from source to destination.

- Often defined as the job of Layer 3 (IP). But...
 - Ethernet spanning tree protocol (Layer 2)
 - Distributed hash tables, content delivery overlays, ... (Layer 4+)

Problems for intradomain routing



- Distributed path finding
- Optimize link utilization
- React to dynamics
- High reliability even with failures
- Scale

Problems for interdomain routing



- All of intradomain's problems
- Bigger scale
- Multiple parties
 - No central control
 - Conflicting interests
 - Attacks
- Harder to change architecture
 - Intradomain evolution: RIP, ISIS, OSPF, MPLS
 - Interdomain: BGP.

The two classic approaches



- Distance vector & Link State
- These are far from the only two approaches! We'll see more later..

Distance vector routing



- Original ARPANET: distance vector routing
- Remember vector of distances to each destination and exchange this vector with neighbors
 - Initially: distance 0 from myself
 - Upon receipt of vector: my distance to each destination = min of all my neighbors' distances + 1
- Send packet to neighbor with lowest dist.
- **Slow convergence** and **looping** problems
 - E.g., consider case of disconnection from destination
 - Fix for loops in BGP: store path instead of distance

Link-state routing



- New ARPANET algorithm: link state routing (“shortest path first (SPF)”)
 - McQuillan, Richer, and Rosen 1980; Perlman 1983; led to OSPF
- Broadcast the entire topology to everyone
- Locally run shortest path algorithm
- Send packet to neighbor along computed shortest path
- How can this result in forwarding inconsistencies?

Interdomain routing




- BGP: Border Gateway Protocol
- Remember path instead of distance (“path vector” instead of “distance vector”)
 - Avoid loops; but more importantly...
 - **Support policies:** can pick any path offered by neighbors, not necessarily the shortest

Common policies




- Route selection: prefer customer over peer over provider



Common policies



- Route export: “valley-free”: to/from customer only



What's to come



- Today: interdomain routing basics, and a radically different design
- Thursday: **reliability**
- Next Tuesday: **scalability**
- Next Thursday: **selfishness & conflicting interests**