Justin Morgan
cse588 – 3rd week, 1st half

Intradomain routing protocols

- Centralized arbiter
  scalability – function of communication cost and computation
  fault tolerance (single point failure)
  simpler
  propagation of updates
- Distributed routing computation
  Distance vector routing – each node talks only to neighbors, exchanges tables, and keeps best alternative
  - fundamental problem – propagation not deterministic, info gets lost, can lead to routing loops or flaps
  - poison reverse – remember immediate source of info
  - path vector (BGP) – remember complete path

  Link-state routing – (e.g. OSPF) send info about all neighbors to everyone, compute routes (shortest path)
  - problem – temporary routing loops during update propagation (no synchronization)
  - Is it better to send update message and then compute update or compute the update and then send update message? No clear answer.
    a) SEND FIRST – propagate news quickly
    b) UPDATE FIRST – prevents neighbor from updating first and creating loop, but takes longer to stabilize. (OSPF)

Periodic updates – why?
1) hello, is link up?
2) congestion control
3) routing updates (changes)
4) soft state recovery

BGP (interdomain routing) only sends deltas rather than entire table with each update.

Difficult to guarantee end-to-end quality of service if route changes, requiring circuit tear down and re-establishment.

Exercise – what happens if a router misbehaves, sending incorrect updates?

Metrics?
1) Hop count – poor metric, adding paths can degrade performance.
2) Transmission + delay + queue time

Route selection today -
1) Heuristics – do the best you can
2) Circuits and policy
3) Optimization algorithms

Adaptive routing at local level for congestion control can lead to route oscillation.