ECE 435 – Network Engineering
Lecture 14

Vince Weaver

http://web.eece.maine.edu/~vweaver

vincent.weaver@maine.edu

24 October 2016

Announcements

- RFC791 post-it note (having trouble finding original source):

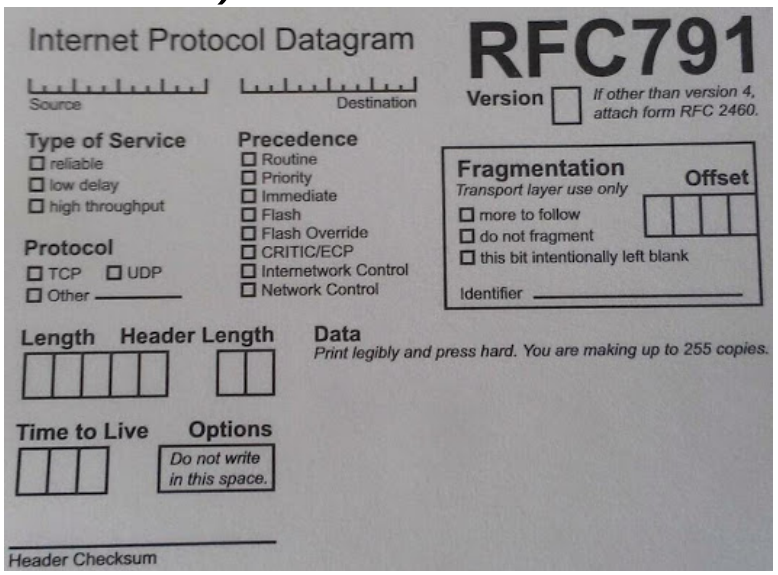

- DDoS of internet on Friday

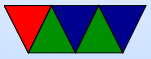

Hierarchical Routing

- Would you want to have all routers in network on flat network? Crazy routing table
- Split into a hierarchy
- Autonomous Systems (ASes)
- Systems under same command (same ISP) use intra-domain routing protocol, or interior gateway protocol (IGP)
- Border routers connect to border routers of others
- Inter-domain routing, EGP (exterior gateway protocol)

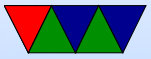

- types
  - Stub AS – like ISP with customers, one gateway to internet
  - Multihomed AS – multiple gateways (why?) redundancy. traffic generally doesn't flow through
  - Transit AS – traffic can flow through network

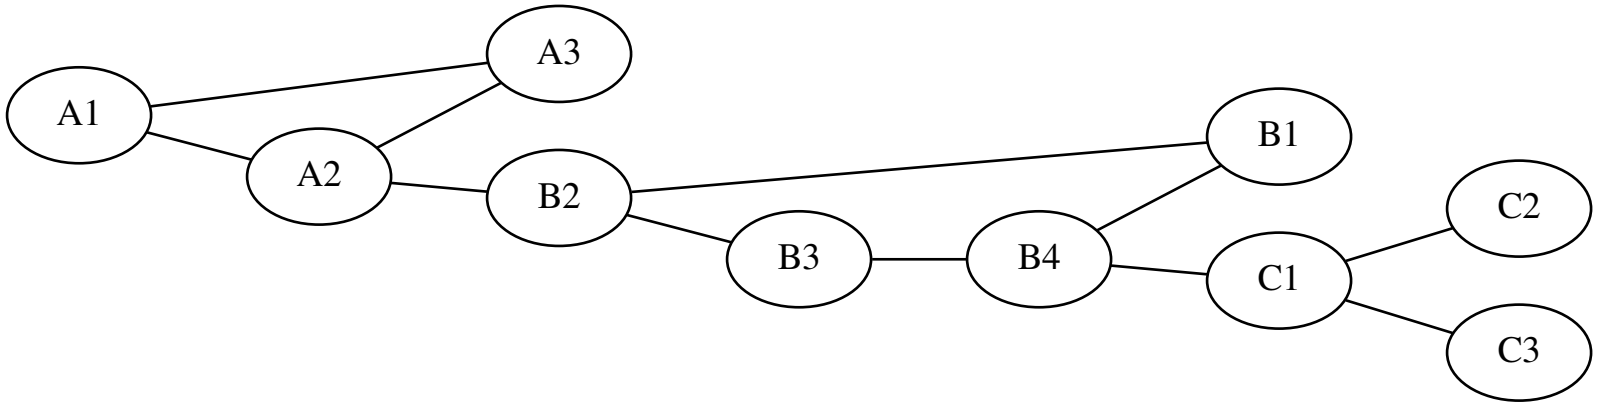

- Packet A1 - A3 internal A1 - B2 goes to border router

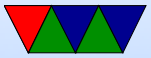

and across, then local A1 - C2 goes to border router to B network, across local to B/C border, then finally to C

- If flat network, need to know 10 machines in routing table
- In hierarchical only need to communicate to 2-3 other routers, find way to border router

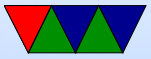

Intra-Domain Routing

- Interior Gateway Protocols
- RIP (Routing Information Protocol)
  by Xerox, included in BSD, routed RFC 2543
  distance vector, with hop count, max 15 hops
  RIP advertisements over UDP port 52
  Send advertisement when changes, or 30s
  25 vectors: dest, next hop, distance
  timeout on not functioning route, garbage collection
  split horizon, poison reverse, and hold down timers to

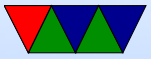

avoid loops

- OSPF (Open Shortest Path First)
  successor to RIP. RFC2328 (5340 for IPv6)
  link-state-routing protocol
  link-state flooded to all routers in domain, each router uses Dijkstra to find least cost for self, builds table
  load balancing – supports equal-cost multipath routing (can equally use equal cost routes)
  supports CIDR routing
  support available for multicast
  8-byte password for authentication

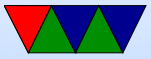

supports hierarchical
example? complex

- Implementations: zebra? userspace multi-protocol router daemon

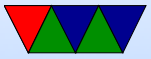

Inter-Domain Routing

- Can be complicated. Say company with network, and two connections to outside X,Y. Don't want to send packets out and back even if it looks like lower cost. Also don't want to transit packets between X and Y for outsiders. Policy.

- BGP (border gateway protocol)
  BGP4 RFC 4271
  interior and exterior BGP
  iBGP makes sure that the setups for multiple gateway

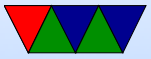

routers are kept synchronized

eBGP used to talk between other exterior routers at peers. Uses TCP (reliable) port 179.

Due to size of internet, uses distance vector over link state.

Uses path vector rather than distance vector: exchanges info with neighbor, but includes complete path info to avoid looping. Each AS has unique number, so if it sees itself in the path knows there is a loop.

Four types of messages

Whole table not passed around (Due to size) but updates

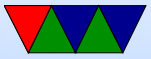

Keeps track of all feasible paths, but only advertises the "best" one

Example. Full BGP of internet backbone router might have more than 300,000 entries (2010) now over 600,000. `http://bgp.potaroo.net/`

Some routers had limit of 512k so on August 12 2014 part of internet went down when crossed the border. (Ipv6 currently only around 20k)

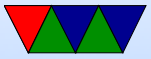

Peering

- How companies agree to connect their networks together. There's not really a master connection, but instead companies agree to have routers talk to each other via BGP.

- Transit – pay money to pass through network.

- Peering – In many cases no money changes hands. Why? Well if you have a lot of users, but no content, people won't stay with you. Same if you have content but no

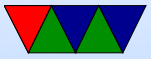

access to users. Averages out and is mutually beneficial.

  - Increased redundancy
  - Increased capacity
  - Increased routing control
  - Improved performance
  - Fame (high-tier network)
  - Ease of requesting aid (?)

- Customer – you buy an internet connection

- Peering locations, often in large data centers.
  At one point there were 4 major ones MAE-East

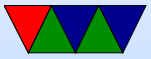

(Virginia) [in basement of parking garage, at one point half of internet went through here], Chicago, NY, SF. All defunct now

- Depeering – if you think you aren't getting a good deal, break up. Some situations there is a fight, a hope that the customers lose enough performance will have to repeer.
- Related – net neutrality
- Tiers – Tier 1 network is one that can reach rest of internet without paying for transit; Tier 2 peers with

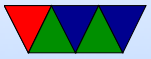

some but purchases for other; Tier 3 only purchases

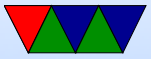

Routing Issues

- Problems – routing black hole, use BGP to send addresses intentionally to 0.0.0.0 and get dropped. BGP will propagate
- router update mistakes can accidentally blackhole parts of the internet
- In 2008 Pakistan was trying to blackhole Youtube and accidentally announced to world via BGP and took it down world wide

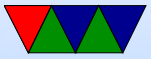

Broadcast Routing

- Send to each destination? Waste bandwidth, but also need to know all possible destinations
- Flooding? Also too much bandwidth
- Multi-destination routing
- Spanning tree
- Reverse-path forwarding

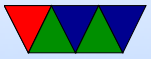

Multicast IP

- For IP, just join a class D network
- To both sender and receiver it's like sending/receiving a unicast packet
- all the hard work done by routers
- How do you join a multicast group?
- Router two tasks: group membership management, packet delivery.

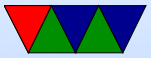

- IGMP (Internet Group Management Protocol)
  IGMPv3 RFC 3376
  query, report, leave
  querier and noquerier
  router with lowest IP is querier
  no real controls on who can join or send

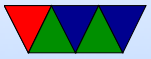

Mobile Routing

- Steiner tree – NP complete
- Heuristics, but none generate entire tree as need centralized and global knowledge
- DVMRP (Distance-Vector Routing Protocol)
  original protocol, MBONE (tell story)
  reverse path broadcast, then prune vs reverse path multicast
- Protocol Independent Multicast (PIM)

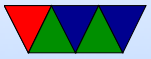

DVRMP not scalable for multicast groups with sparse members

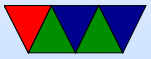

Bridging vs Routing

- Bridge is link-layer, Routing is network layer
  Bridge uses MAC address, router uses IP
- Bridge uses self-learning, router uses routing algorithms
- Bridge needs to use spanning tree to avoid loops
- Bridge separates a collision domain, router a broadcast domain
- Scalability. Bridged networks are all on same broadcast

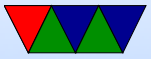

domain, so if millions of devices all broadcasting can have problems

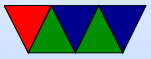

Other types of Routing

- Mobile – what do you do when machines can come and go?
  have a "home" location. Packets go there. When you get on network, update with actual location. Network gets packets at home location, encapsulates and sends to actual location

- Ad Hoc Routing
  Bunch of machines in an area, routers and devices can come or go more or less randomly.

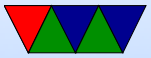

route discovery

- Peer to Peer File Sharing
  - Centralized server? Napster? Easy to take down.
  - Want Distributed, no central control.
  - Flooding: connect to one other connected node. Floods requests (sort of like broadcast) until it finds who has file, then direct connect to transfer.

  - distributed hash table

- Secret routing

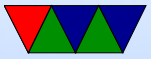

TOR / The onion Router

Packet encrypted multiple times, in layers. Randomly sent to next machine which decrypts that layer, passed on

At end comes out random “exit node” and drops onto regular internet