

ECSE-6600: Internet Protocols

Informal Quiz #04

Shivkumar Kalyanaraman:
GOOGLE: "Shiv RPI"
shivkuma@ecse.rpi.edu



Routing (Slide set #5): Informal Quiz

Routing: Basic Ideas

T F

- Forwarding works in the control plane whereas routing works in the data plane
- Forwarding is an example of resolving an indirection (I.e. binding), and the routing table is the aid used.
- Routing tables are always necessary for the filtering and forwarding decision
- A routing protocol summarizes global information to setup a local next-hop entry in the forwarding table
- Internet routing guarantees that packets forwarded by using the routing tables at routers will ultimately reach their destination
- The internet routing infrastructure continuously works to increase the probability that packets forwarded using the routing tables will ultimately reach their destination.
- Internet routing guarantees shortest paths from source to destinations
- When the flag field in a host routing table is U, it means the corresponding interface (and peer) is likely to be up and working
- Layer 2 (bridge) forwarding tables tend to be inefficient when topology becomes larger and highly dynamic

Routing: Basic Ideas.

- □ A L3 routing protocol is always needed to make changes in a forwarding table.
- □ Bridges resort to flooding of data packets when they do not find a match for the packet's destination address in the forwarding table
- □ Routers resort to flooding of data packets when they do not find a match for the packet's destination address in the forwarding table
- □ It is possible that a L3 lookup in a forwarding table will not lead to any matches.
- □ It is possible that a L3 lookup in a forwarding table will lead to multiple matches due to the practice of supernetting (CIDR): this situation is resolved using the longest-prefix-match rule.
- □ Bridges coordinate with other bridges only to make sure that the flooding of packets does not go out of control (using a tree-topology to guide and filter flooding)
- □ ARP requests, multicast traffic and poorly populated L2 bridging tables can lead to “broadcast storms” on extended LANs: a factor limiting scalability of L2 bridged networks

Routing: Basic Ideas..

- □ Routing algorithms always build a comprehensive view (a map or graph) of the network and figure out next-hops to all destinations from these views.
- □ There is usually a tradeoff between consistency and completeness in routing design (assuming limited scale networks).
- □ If the topology is simple, the networks are small, and the network management goals are simple, then a simple routing protocol (eg: RIP) is preferred to complex routing protocols (eg: OSPF, IS-IS).
- □ Both the goals: consistency and completeness, put a limit on scalability of the routing protocol. This is one of the motivations for hierarchical routing (eg: OSPF and IS-IS areas, PNNI peer-groups, and the split between inter- and intra-domain routing).
- □ Global state of large, dynamic networks can be quickly and easily collected, and processed to create routing tables.
- □ The internet uses a signaled approach to routing state establishment
- □ When a routing protocol is in its transient phase (I.e. it has not converged), we can safely say that its routing tables are consistent

Routing: Basic Ideas...

- □ Inconsistent routing tables lead to wandering or looping packets.
- □ Delay-sensitive, fragile applications (eg: internet telephony, interactive multimedia) can function quite well over networks that have a high probability of being inconsistent in terms of routing table states.
- □ A simple consistency criterion is usually necessary in a fully-distributed routing approach to quickly converge to consistent routing tables.
- □ With a complete view of the network topology and dynamic state, more sophisticated routing and traffic engineering decisions (eg: multi-path or QoS routing) can be performed by the routing algorithms.
- □ Telephony uses an explicit, source-based routing approach.
- □ Centralized routing is not even used in relatively small networks because it is prone to failures and traffic surges
- □ The term “hop-by-hop” routing is used often instead of the term “fully distributed routing”
- □ The “loose source routing” option is quite frequently used in practice in the Internet.

Routing: Basic Ideas....

- □ Routed is an example of a static routing protocol
- □ The basic routing architecture in telephony is much simpler than the dynamic routing infrastructure of the Internet
- □ The “dumb” network analogy does not apply to the control-plane architecture (I.e. routing protocols) of the Internet
- □ Like the telephone networks (and its derivatives: frame-relay and ATM networks), the Internet also attempts to provide end-to-end QoS (I.e. quality-of-service) routing
- □ The core topology of the telephone networks is extremely simple (I.e. full mesh) and reliable, which facilitates a fairly simple routing approach
- □ Reliable and robust hardware and software components lead to huge surges of routing control traffic for conservatively sized and engineered networks.
- □ The telephone network builds in an alternate path routing mechanism to reduce the call-blocking probability.
- □ Telephone networks are organized as multiple autonomous systems and have multiple AS numbers within a single organization

Routing: Basic Ideas.....

- □ All intra-domain routing protocols provide support for sophisticated traffic engineering features
- □ One of the key goals of an inter-domain routing protocol is to allow policy-based path selection.
- □ Inter-domain routing protocols guarantee stability and reachability.
- □ Border routers in an AS may participate in both intra-domain and inter-domain routing
- □ Inter-domain routing may be understood as a routing protocol operating at a higher level of abstraction, I.e. providing routes between “virtual nodes” that are in fact “autonomous systems”
- □ Source-based (or loose-source) routing methods typically encode the route in the packet header and/or pre-setup the route using a signaling protocol
- □ Distance vector protocol collects dynamic state about every link in the network
- □ A path-vector protocol collects dynamic state about every possible path in the network

Routing: Basic Ideas....

- □ The maps collected by link state algorithms (link state databases) are guaranteed to be consistent at all times.
- □ Link state algorithms exchange more state information, but they tend to converge to consistent routing tables faster than simple distance-vector algorithms, especially for larger dynamic networks
- □ Distance vector algorithms collect per-node distance information in order to compute routes
- □ The final result of a routing algorithm is the enumeration of all possible routes to every destination in the network.
- □ The distance-vector protocol involves checking neighbors' distance vectors and updating its own distance vector.
- □ Both the distance-vector and link-state approaches could lead to transient routing loops because the information maintained could be incomplete.
- □ The simple consistency criterion used both by link state and distance vector algorithms in the Internet is that: “a subset of a shortest path is also a shortest path between the intermediate nodes”

Routing: Basic Ideas.....

- □ After m-iterations of both the distance vector or link state approach, the algorithm is dealing with a m-hop view of the network from the source node
- □ Even though both LS and DV algorithms deal with the same effective network view after m-iterations, the LS algorithm has more information, and the DV algorithm performs each iteration only after further information exchange with neighbors
- □ Bellman-Ford is an example of a link-state routing algorithm
- □ The count-to-infinity problem is an example of inconsistency issues in distance vector methods.
- □ The poisoned reverse modification of DV algorithm eliminates all forms of inconsistency problems in DV protocols
- □ The poisoned reverse modification of DV algorithm is less effective in cases where the cost of a remote link (not the first or second) in a path increases (because bad-news travels slowly).
- □ The link state method does not face the count-to-infinity problem because it has complete global information (a map in terms of link-states).

Routing: Basic Ideas.....

- □ The Dijkstra algorithm greedily relaxes edges of all nodes at the beginning (I.e. initialization) of the algorithm
- □ The Dijkstra algorithm computes the shortest-path-tree (SPT) that is a spanning tree to all nodes rooted at the source node
- □ Hierarchical addressing, and proper address assignment allows entire subnets to be viewed by interior routers as “virtual nodes”, leading to routing scalability
- □ Hierarchical addressing leads to optimal (I.e. shortest) routes between all pairs of nodes in the network