

CS 168

Introduction to the Internet: Architecture and Protocols

Spring 2024

Instructor: Sylvia Ratnasamy & Rob Shakir

Lecture: Tu/Th 3:30pm-4:59pm, [Dwinelle 145](#)

NOTE: This website is under construction.

Glossary

This page will be updated as the class progresses. Click on the 'Term' and 'Lecture' headers to sort.

| Term | Definition |
|-------------------|--|
| ACK | TCP packet with ACK flag set, that indicates data has been received. |
| ARP | Address Resolution Protocol. The protocol that allows devices to map MAC addresses to IP addresses. A device will send out (broadcast) an ARP Request message, to find out the MAC address corresponding to the IP Address. The device that is being queried will respond (unicast) with an ARP Response message. Mappings between MAC addresses and IP addresses are stored in the ARP table, which serves as a cache. Entries in the ARP table will time out (soft state). |
| Autonomous System | A network or set of networks that are all managed and supervised by a single entity or organization. A single ISP is often a single AS; however, some ISPs partition their network into multiple ASes. Each AS is assigned a number, which is used in BGP to identify paths. |
| AXE | A proposed alternative to STP, where loops are prevented without the need for a spanning tree (by using duplicate suppression instead). Source of some of the finest poetry the world has ever seen. |

| | |
|-------------------------|---|
| can happen to packets | Lost, corrupted, reordered, delayed, duplicated. |
| Bandwidth-Delay Product | This is the quantity (bandwidth)*(propagation delay) which represents the number of bits needed to “fill the pipe” (i.e., the number of bits that would have been sent but not yet received if the sender was sending at the bandwidth of the link). |
| Bellman Ford Equation | Equation that says your shortest distance to a destination is the minimum of your cost to a neighbour plus that neighbour’s distance to the link, for all neighbours. Or, more concretely, Node u’s cost to a given destination v is then: $d(u,v) = \min(\text{nbrs } w) [c(u,w) + d(w,v)]$ |
| Best Effort | On-demand delivery where the system provides no performance guarantees other than the system will try its best. |
| Border Router | Routers that are connected to routers in another network. |
| Checksum | This is used to detect corruption, and is a number computed over some portion of the packet (which depends on which protocol). |
| CIDR | Short for Classless Interdomain Routing. Couples IP address with a network mask to determine which bits are the network bits. Far more flexible than the original IP addressing scheme, or classful addressing. |
| Circuit Switching | Method of data transfer in which an end system reserves BW along a path - sets up a circuit - for communication. No need for packets. |
| Classful Addressing | A scheme for determining the network and host bits of an IP address. There are three classes (that we deal with): Class A addresses start with 0, and use the first 8 bits to identify the network. The last 24 bits identify the host. Class B addresses start with 10, and use the first 16 bits to identify the network. The last 16 bits identify the host. Class C addresses start with 110, and use the first 24 bits to identify the network. The last 8 bits identify the host.” |

| | |
|---------------------------|--|
| Control Plane | This refers to the network mechanisms used to compute routing tables and other forwarding information. |
| Convergence | We say that an algorithm has converged once all parties have up-to-date information and, barring a change in the topology of the network, all subsequent “updates” sent and received don’t affect the routing state. |
| Core/Backbone Router | Routers that are connected to other internal routers. |
| Cost Table | Data structure on routers that contains the set of costs to all neighbors. |
| Count-to-Infinity Problem | The name for a certain type of routing loop that can occur due to the asynchronous nature of information propagation using Distance Vector. Usually caused by a link going down, a router, originally using a broken path to a certain destination, believes their neighbour has a valid path to said destination and adopts this path without knowing that it contains a subset of its original broken path. Both neighbours repeatedly receive updates from each other, adopting this broken path. |
| Cumulative ACK | ACK means “I have received all packets (or bytes) up until this one”. |
| Data Plane | This refers to the network mechanisms used to forward data. |
| Datacenters | Massive collections of machines. |
| David Clark | The unsung hero of the Internet. He was the chief architect and authored the end-to-end principle. |
| Dead End | When a packet arrives at a router or switch but the forwarding decision does not yield an outgoing port, forcing the packet to be dropped. |
| Destination-Based Routing | Routing that only depends on the destination. Paths from two different sources to same destination must coincide once they overlap. |
| DHCP | Dynamic Host Configuration Protocol. The protocol that provides a host with its IP address upon connecting to a network. When a host connects to a new |

| | |
|-------------------------|--|
| | <p>that it needs an IP address. The server sends an Offer message, containing an offered IP address, a subnet mask, the IP address of the first-hop router, and a lease time. The host will send a Request, corresponding to the offer it would like to accept. The server responds with an Acknowledgement/Acceptance message. All DHCP messages are broadcasted.</p> |
| Distance Vector Routing | <p>Distance vector routing is a scalable and distributed routing algorithm in which each router keeps a “vector” of distances as well the next-hop router to each destination. Each node floods its vector shortest distances to its neighbors and upon receiving a vector each router uses Bellman-Ford to update its own vector.</p> |
| DNS | <p>Domain Name Service, a system which associates names with addresses and is commonly used to look up the address of a host given a name.</p> |
| Dotted-quad notation | <p>A notation that writes down IPv4 addresses as 4 numbers, one number per byte. For example, 12.34.158.5</p> |
| Duplicate ACKs | <p>A stream of cumulative ACKs that acknowledge the same received data multiple times: a sign of an isolated packet loss, because these additional ACKs show that data is still being received.</p> |
| Edge Router | <p>Routers to which end hosts are attached.</p> |
| End to End Principle | <p>Helps determine whether or not something should be implemented in the network, or only in the end hosts. This class presented three interpretations: Only-if-necessary: If a function can be implemented by the hosts, don’t implement it in the network. Only-if-sufficient: Only implement a function at this level if it can be completely implemented at this level and you can relieve burden from the hosts. Only-if-useful: Implement a function in the network if it can improve the performance of the function without burdening applications that don’t require it.</p> |
| Enterprises | <p>Companies and universities.</p> |

| | |
|----------------------|--|
| Fate Sharing | Store state in the entities that rely on that state, such that that entity will not be affected by other failures. |
| First-hop router | The router that a host sends a packet to when it wants to send that packet to a destination outside its L2 network. |
| Flooding | In this class, we use flooding to refer to the act of sending a packet out all ports (excepting the incoming port) in a single switch. |
| Flow | A stream of packets between two processes. |
| Forwarding | Sending a packet towards its destination. This is done by reading the address from the packet's header, searching the routing state for the correct output port, and sending the packet out that port. This is a local process within a router, done in the data plane, and it must be done quickly. |
| Forwarding Entry | An entry in the forwarding table that maps an address or set of addresses to an outgoing port. |
| Forwarding Table | A table the router computes for itself to guide its forwarding decisions. The forwarding table is computed using the information in the peer and cost tables. |
| Fragmentation | Dividing a packet into smaller packets to fit the maximum transmission unit (MTU) of a link. |
| Full-information ACK | An ACK that describes all data received so far, and can take the form "I have received all packets up until this one, plus these additional ones". |
| Hard State | Systems in "hard state" do not time out their information – they assume once they have been given some knowledge, it remains true and valid until explicitly told otherwise. |
| Host bits | The part of the IP address that identifies the host inside its network. |

| | |
|---|---|
| Host/End System | End-points of a network. These entities are responsible for the generation of data packets that are then routed across the network. |
| Individual ACK | An ACK that means “I received this single, specific packet”. |
| Internet | The core network infrastructure that links all connected computing devices. |
| IP Address | The addressing scheme used in layer 3. |
| IPv4 | Version 4 of the IP protocol. |
| ISP (Internet Service Provider)/ISP Network | A network of packet switches and communication links providing network access to end systems. |
| LAN | Local area network, an L2 network that spans a small geographical area, for example a house. |
| Layering | Generally, layering is splitting a complex system into separate levels that build upon/depend on each other. In the Internet context, this refers to a specific set of layers (physical = L1, datalink = L2, internetworking = L3, transport = L4) that only interact with the layers directly above or below. |
| Layers | Application (Network support for apps). 4: Transport (Reliable/Unreliable end-to-end delivery). 3: Network (Global best-effort delivery). 2: Datalink (local best-effort delivery). 1: Physical (bits being transmitted over some medium). |
| Learning Switches | Typically used at L2 in combination with the spanning tree protocol. Learning switches maintain a forwarding table mapping destination to output link. They learn from the “source” field of a packet. When a packet comes in, the switch checks if the destination is in it’s table. If it is, it forwards the packet down that link. If it isn’t, it floods the packet. |
| Linecard | Piece of hardware (in a router) that receives/sends packets. They update various fields (checksum, TTL, etc.) and select outgoing port. |

| | |
|---------------------------------|---|
| Link | The physical pieces or infrastructure that connects routers. |
| Link-State Routing | In Link-State routing, each router sends (using a protocol-specific broadcast mechanism) its link state to all other routers in the network. This way, every router learns the entire network graph. Then, every router computes the least-cost paths from themselves to all other nodes using any valid algorithm (for example, Dijkstra's). |
| Loop | When a packet cycles around the same set of nodes forever. |
| LPM | Longest-prefix-match: When an IP Address matches multiple prefixes, select the longest match (Think of traversing the prefix tree until the address 'falls off'). |
| MAC Address | Used for L2 routing, a MAC Address is a 48 bit number burned into the network interface of hosts and routers. The MAC address is encoded in the physical hardware stored in Read-Only memory, making it a permanent identifier. |
| Maximum Transmission Unit (MTU) | The largest number of bits a link can transfer as a single unit, the largest packet size that can be sent across a link. |
| Modularity | Decomposing a problem into tasks or abstractions. Leads to the design principles of layering. |
| Multihoming | Connecting one host to multiple, disparate networks, so that if one parent network goes offline, the host is still accessible. Prevents aggregation. |
| NACK | "Non-acknowledgement" message – "I did not receive this data [that I was expecting to]". |
| Network | When used informally, this refers to a system composing of end systems, routers/switches, and links that is able to transfer data between hosts (e.g. Berkeley's campus network). When used formally, it refers to a set of network elements that share the same network address in IPv4, and is often used synonymously with subnet. |

| | |
|---------------------|---|
| Network Address | The component of an IP address that refers to the network (or subnet), rather than the host. |
| Network bits | The part of the IP address that identifies the network the host is on. |
| Network mask | An IP-address-like string of bits used to identify the network portion of an IP Address. Made up of some set number of 1s (one per network address bit), followed by all 0s. |
| Network Name | The name of a host (Something human-friendly). |
| Network Stack | The networking software on the host, it replicates some functionality found at the routers and also adds additional functionality (e.g. Sockets, TCP header, etc). |
| Packet | Bags of bits. Consists of: Header with meaningful information for network and network stack to make decisions. Body containing a payload. Ex. A file, imagine, an application header, etc. |
| Packet Switching | Method of data transfer in which data is segmented into packets and routers/switches service each packet they receive independently by inspecting its header. |
| Path Vector Routing | Similar to distance vector routing, but when advertising to neighbors, instead of sending them your shortest distance, you send them your paths to destinations. |
| Payload | Data carried in packet. |
| Peer Table | Data structure on routers that contains copies of the information each of the router's "peers" or "neighbours" sent them. |
| Poison Reverse | Method attempting to mitigate the count to infinity problem by not advertising the ability to reach a destination (i.e. advertising a distance of infinity) to a neighbour you use on the path to said destination. For example, router A creates a temporary copy of its vector to send to router C that |

| | |
|------------------------------|--|
| | AC. |
| Port (Logical) | A number that an OS assigns to a socket that is used to identify the socket. |
| Port (Router) | The physical port which connects a router to another router through a link. |
| Prefix Aggregation | Combining routing table entries into one entry by using a common prefix (i.e. combining 101 and 100 to 10*). |
| Prefix Tree | Binary tree that represents matching bits in IP address lookup (how the lookup table is traversed). |
| Reliability (see Robustness) | Two interpretations: 1) The network recovers from failures quickly, allowing two, non-partitioned endpoints to communicate. 2) Network failures do not interfere with endpoint semantics. |
| Reliable Delivery | Building a reliable transport service on top of best-effort delivery. |
| Reliable Transport | A transport mechanism is “reliable” if and only if (a) it resends all dropped or corrupted packets, and (b) it attempts to make progress. |
| Resource Accountability | The ability to know who is using what resources (bandwidth) so that you can hold them accountable to it. A failure in Internet architecture. |
| Robustness (see Reliability) | As long as the network is not partitioned, two hosts should be able to communicate eventually, AND failures should never interfere with application semantics. |
| Route Aggregation | Instead of having one forwarding entry per host, have one entry per set of hosts with the same prefix that all go out the same port. |
| Route Poisoning | Procedure to mitigate network inconsistencies that says when a link goes down between A and B, router B should advertise to all its neighbours that it no longer has a link to router A (i.e. B advertises a distance of infinity), in order to signify that it can no longer reach A. |

| | |
|------------------------|---|
| Routing | Guiding packets from source to destination (can be done in many ways - see link-state, distance vector, spanning tree, etc). This is inherently a global process, so it must scale. This is done in the control plane, and may be done slowly. |
| Routing Table | This is similar to the Forwarding Table, but can refer to all the information a router (including from other peers) rather than just the best forwarding entries. |
| Slash notation | Notation for talking about a subnet. Looks like 1.2.0.0/10 where the first 10 bits of 1.2.0.0 are the subnet prefix. |
| Sliding window | A finite number of un-acked packets allowed to be in flight (for efficiency purposes) before we stop sending more. |
| Socket | An OS mechanism used to connect a process to the networking stack. |
| Soft State | The concept of allowing your stored knowledge to “time out”, under the assumption that it may have changed/no longer be valid/etc. Systems that operate under soft state will periodically “forget” what they know and need to “re-learn” it – by requesting the information again, waiting for new messages and information, etc. DHCP offers having a ‘lease time’, cached ARP entries timing out, and the periodic messages in Distance-Vector Routing are all examples of soft-state. |
| Spanning Tree Protocol | A distributed protocol in which switches send messages of the format (Y, d, X) from node X proposing Y as the root and advertising a distance of d to Y. This protocol identifies the node with the lowest ID and builds a spanning tree with that node as the root. |
| Split Horizon | Split horizon provides the same functionality as poison reverse, but expresses the information differently. Split Horizon is used in the context of full updates, and the router does not advertise any route to destination X to the neighbor it uses to reach destination X. |

| | |
|--------------------------|---|
| Statistical Multiplexing | Summing the maximum rate of flows is greater than combining the flows and finding the maximum. |
| Subnet | In this class we use this term to refer to a portion of a network that is connected by L2 and shares the same network address. |
| Time to Live (TTL) | In IP, this refers to the number of hops a packet can travel before being dropped, which is useful in preventing loops. More generally, TTL refers to the time until something expires (such as a cached entry). |
| Valid Routing State | A routing state is valid if and only if there are no loops and no deadends (assuming no packet replication). If there is packet replication, then this changes to saying at least one replica does not hit a deadend. |
| WAN | This can refer to any L3 network (i.e., not just a local area network), or to networks that span large geographic distances (i.e., not a datacenter). |

Last built: 2024-05-10 20:31 UTC