LECTURE 1
Connecting Computers with Packet Switching

[1]

This lecture discusses different ways of interconnecting links (of the same kind) to build a simple computer network. To achieve this task, we will use a device called a *switch*, and discuss several different ways of switching to move data across a network. We focus on *packet switching*, discussing its main ideas and principles.

This lecture assumes that the reader is familiar with standard ways of communicating digital information (bits and link-layer frames) over a single link. Most networking texts cover this material in depth; we provide a short summary of the essential ideas in L0.

■ 1.1 Interconnections

The rather limited scope—in terms of physical distance, number of connected hosts, and amount of sustainable traffic—of single-link networks leads us to examine ways of interconnecting single-link communication media together to form larger networks of computers. We start by discussing a few different interconnection techniques. The fundamental problem is that the most obvious way to build a computer network—by connecting each pair of computers with a dedicated link—is both prohibitively expensive (because of the sheer number of links required, a number that grows quadratically with the network size) and technically challenging (because signals attenuate with distance, requiring ways to regenerate information across large distances). The solution to these problems is to develop ways to *share* links between different communicating nodes, and to regenerate the information being communicated as it travels across the network.

The key component used for such interconnections is a *switch*, which is a specialized computing device that receives data frames (of bits) that arrive over links, processes them, and forwards them over one (or more) other links. Links are physically connected to switches at *attachment points* or *switch ports*.

The fundamental functions performed by switches are to multiplex and demultiplex data frames belonging to different computer-to-computer information transfer sessions (or "conversations"), and to determine the link(s) along which to forward any given data frame. This task is essential because a given physical link will usually be shared by several concurrent sessions between different computers.

Over time, two radically different techniques have developed for doing this. The first, used by networks like the telephone network, is called *circuit switching*. The second, used by networks like the Internet, is called *packet switching*. The key difference between the two is that, in circuit-switched networks, the frames do not need to carry any special information that tells the switches how to forward information, while in packet-switched networks, they do.

The transmission of information in circuit-switched networks usually occurs in two phases: first, a setup phase in which some state is configured at each switch along a path from source to destination, and second, the information transfer phase when the frames are actually sent. Of course, because the frames themselves contain no information about where they should go, the setup phase needs to instantiate the correct state in the switches to enable correct forwarding.

A common (but not the only) way to implement circuit switching is using *time-division multiplexing (TDM)*, also known as *isochronous transmission*. Here, the physical capacity of a link connected to a switch, C (in bits/s), is conceptually broken into some number N of virtual "channels," such that the ratio C/N bits/s is sufficient for each information transfer session (such as a telephone call between two parties). Call this ratio, R, the *rate* of each independent transfer session. Now, if we constrain each frame to be of some fixed size, s bits, then the switch can perform time multiplexing by allocating the link's capacity in time-slots of length s/C units each, and by associating the ith time-slice to the ith transfer (modulo N). It is easy to see that this approach provides each session with the required rate of R bits/s, because each session gets to send s bits over a time period of Ns/C seconds, and the ratio of the two is equal to $C/N = R$ bits/s.

Each data frame is therefore forwarded by simply using the time slot in which it arrives at the switch to decide which port it should be sent on. Thus, the state set up during the first phase has to associate one of these channels with the corresponding soon-to-follow data transfer by allocating the ith time-slice to the ith transfer. The end computers transmitting data send frames only at the specific time-slots that they have been told to do so by the setup phase.

Other ways of doing circuit switching include *wavelength division multiplexing (WDM)*, *frequency division multiplexing (FDM)*, and *code division multiplexing (CDM)*; the latter two (as well as TDM) are used in some cellular wireless networks. Various networking textbooks (e.g., Tanenbaum, Peterson and Davie, etc.) describe these schemes in some detail.

Circuit switching makes sense for a network where the workload is relatively uniform, with all information transfers using the same capacity, and where each transfer uses a *constant bit rate (CBR)* (or near-constant bit rate). The most compelling example of such a workload is telephony, and this is indeed how most telephone networks today are architected. (The other reason for this design choice is historical; circuit switching was invented long before packet switching.)

However, circuit-switching tends to waste link capacity if the workload has a *variable*

bit-rate, or if the frames arrive in bursts at a switch. Because a large number of computer applications induce burst data patterns, we should consider at other link sharing strategies for computer networks. It turns out that the once-radical packet-switching technique is a general way of getting better performance for such workloads, and is the fundamental multiplexing approach used in most networks today.

1.2 Packet switching

The best way to overcome the above inefficiencies is to allow for any sender to transmit data at any time, but yet allow the link to be shared. Packet switching is a way to accomplish this, and uses a tantalizingly simple idea: add to each frame of data a little bit of information that tells the switch how to forward it. This information is added to what is usually called a *header*. There are several different forms of packet switching, which differ in the details of what information is present in the header and what information the switch needs to perform forwarding. The combination of a data frame, with a header that tells switches something about the data's destination or path, is called a *packet*.[2]

The "purest" form of packet switching uses a *datagram* as the unit of framing, with the header containing the *address* of the destination. This address uniquely identifies the destination of data, which each switch uses to forward the datagram. The second form of packet switching is *source routing*, where the header contains a complete sequence of switches, or complete *route* that the datagram can take to reach the destination. Each switch now has a simple forwarding decision, provided the source of the datagram provides correct information. The third form of packet switching is actually a hybrid between circuit and packet switching, and uses an idea called *virtual circuits*. Because it uses a header, we classify it as a packet switching technique, although its use of a setup phase resembles circuit switching. We now look at each of these techniques in more detail.

Packet switches usually require *queues* to buffer packets that arrive in bursts. We will spend a fair amount of time discussing approaches to managing congestion and queues when we discuss network resource management schemes.

1.2.1 Datagram routing

In datagram routing, the sender transmits datagrams that include the address of the destination in the header; datagrams also usually include the sender's address to help the receiver send messages back to the sender. The job of the switch is to use the destination address as a key and perform a lookup on a data structure called a *routing table* (or *forwarding table*; the distinction between the two is sometimes important and will be apparent later in the course). This lookup returns an output port to forward the packet on towards the intended destination.

While forwarding is a relatively simple lookup in a data structure, the harder question is determining how the entries in the routing table are obtained. This occurs in a background process using a *routing protocol*, which is typically implemented in a distributed manner by the switches. There are several types of routing protocols possible (both in theory and practice, although some only in theory!), and we will study several in later lectures. For

[2]This term is due to Donald Davies from the early 1960s.

now, it is enough to understand that the result of running a routing protocol is to obtain routes (paths) in the network to every destination.

Switches in datagram networks that implement the functions described in this section are often called *routers*. Forwarding and routing of packets using the Internet Protocol (IP) in the Internet is an example of datagram routing.

■ 1.2.2 Source routing

Whereas switches implemented routing protocols to populate their routing tables in the "pure" datagram networks of the previous section, the equivalent function of determining which paths to use could also be performed by each sender. When this is done, the network can implement *source routing*, where the sender attaches an entire (and complete) sequence of switches (or more helpfully, per-switch next-hop ports) to each packet. Now, the task of each switch is rather simple; no table lookups are needed. However, it does require each sender to participate in a routing protocol to learn the topology of the network.

People tend not to build networks solely using source routing because of the above reason, but many networks (e.g., the Internet) allow source routing to co-exist with datagram routing.[3]

■ 1.2.3 Virtual circuits

Virtual circuit (VC) switching is an interesting hybrid between circuit and packet switching—it combines the setup phase of circuit switching with the explicit header of packet switching. The setup phase begins with the source sending a special *signaling* message addressed to a destination, which traverses a sequence of switches on its way to the destination. Each switch associates a local *tag* (or *label*), on a per-port basis, with this signaling message and sets this tag on the message before forwarding it to the next switch. When a switch receives a signaling message on one of its input ports, it first determines what output port will take the packet to its destination. It then associates the combination of input port and incoming tag to an entry in a local table, which maps this combination to an output port and outgoing tag (unique per-output).

Data transfer does not use the destination address in the packet header, but uses these tags instead. The forwarding task at each switch now consists of a tag lookup step, which yields and output port and replacement tag, and a tag swapping step which replaces the tag in the packet header.

The reason for the replacement tag is simply to avoid confusion; if global tags were used, then each source would have to be sure that any tag it chooses is not currently being used in the network.

There are many examples of network technologies that employ virtual circuit switching, including Frame Relay and Asynchronous Transfer Mode (ATM). These networks differ in the details of the tag formats and semantics (and these tags are known by different names; e.g., in ATM, a tag is a combination of a VPI or Virtual Path Identifier and VCI or Virtual Circuit Identifier, which can be thought of as a single tag whose structure is hierarchical), and in the details of how these tags are computed. "Multi-Protocol Label Switching" (MPLS) is a link technology-independent approach that network switches can

[3]It turns out that source routing isn't deployed widely in the Internet today for security reasons.

use to implement tag switching. The general principle in all these systems is as explained in this section.

Proponents of virtual circuit switching argue that it is advantageous over datagram routing because of several reasons, including:

- "It allows routes between source and destination to be "pinned" to a fixed route, which allows network operators to provision and engineer their network for various traffic patterns."
- "Tag lookups are more efficient than more-complex lookups based on various fields in the datagram header."
- "Because there is an explicit setup phase, applications that (think they) need resource reservation (e.g., link bandwidth, switch buffer space) can use this signaling to reserve resources."

The above claims are quoted, because they are all quite controversial, except perhaps the first one. Virtual circuit switching is arguably more complex than datagram routing, requiring a separate signaling protocol to set up switch state, and does not handle link or route failures as naturally. (This signaling is in addition to another protocol that allows the switches to discover the network topology, as in datagram networks.) Furthermore, the more complex forwarding table lookups required in IP datagrams are now efficient to implement even at high speeds, and the speed advantages of tag switching appear non-existent. The rationale and mechanism for resource reservation has been hotly debated in the community and will continue to be for some more years! (We will discuss these issues in later lectures.)

Virtual circuit technologies are common in the Internet infrastructure, and are often used to connect two IP routers in a so-called *transport network*.[4] Transport networks are the "link-layer" over which IP packets between two routers are communicated; examples include ATM-based link technologies, switched Ethernet-based local-area networks, 802.11-based wireless networks, etc.

Let us now look at a concrete example of a switched network in more detail, using the widely deployed LAN (local-area network) switching technology as a case study.

■ 1.3 Case study: LAN switching (aka "bridging")

[5]

A single shared medium segment, like a single Ethernet, is limited by the number of stations it can support and by the amount of traffic that can be shared on it. To extend the reach of a single LAN segment requires some way of interconnecting many of them together. Perhaps the simplest way of extending LANs is via "LAN switching," a technique historically known as "bridging". Bridges (or LAN switches; we will use the two terms interchangably) are datagram switches that extend the reach of a single shared physical

[4]Not to be confused with transport protocols, which are used by end points to communicate with each other.

[5]In 2005, this material will be discussed during a recitation.

medium. They work by looking at data frames arriving on one segment, capturing them, and transmitting them on one or more other segments.

Another reason to study bridges is because they are a great example of self-configuring ("plug-and-play") networks.

Bridges are switches that are characterized by:

1. *Promiscuous receive-and-forward* behavior with two or more ports. Each port has a unique identifier, or address, as does each network interface on computers attached to the LAN. "Promiscuous receive-and-forward" means that any packet that arrives on a particular bridge port is replicated by the bridge and forwarded on all other ports to which LANs are currently attached.

   Each bridge port has at most one LAN attached to it (i.e., there could be bridges with no attached LANs), with each LAN in turn possibly containing other bridges. In such LANs implementing promiscuous receive-and-forward behavior with no other mechanisms in place, the aggregate capacity does not exceed the capacity of the weakest segment, since any packet transmitted on any LAN appears on all others, including the slowest one.

2. *Learning*, wherein they learn which stations are on which LAN segments to forward packets more efficiently.

3. *Spanning tree construction*, by which bridge topologies with loops can form a loop-free topology and avoid packet duplication and implosion.

Bridges are *transparent* entities—they completely preserve the integrity of the packets they handle without changing them in any way. Of course, they may add a little to the delays experienced by packets and may occasionally (or frequently, under severe congestion) lose packets because they are, by nature, receive-and-forward devices, but they are transparent to end-points in terms of the functionality they provide.

■ 1.3.1 Learning bridges

The basic idea of a learning bridge is that the bridge "learns", by building a *forwarding table*, which stations are downstream of which port. Then, when a packet destined to a given destination (MAC address) arrives, it knows which port to forward it on. How does it build this table? It learns by looking at the *source* address of packets it sees. When it associates the source address of a packet with a LAN segment, it adds that source-segment pair to the forwarding table.

If the table doesn't have an entry for some destination, the bridge simply floods the packet on all ports except the one the packet just arrived on. Thus, the forwarding table state maintained by a learning bridge is akin to a cache, and is used only as an optimization (albeit a very useful one). The consistency of this cache is not essential for correctness, because the absence of any information causes the packet to be flooded.

This strategy works well, except when there are *loops* in the network topology. In fact, in the absence of loops, not only does the above strategy handle routing in the static case, it also properly handles nodes that move in the topology from time to time (mobile nodes). If a node moves to another location in the switched network, the first time it sends data,

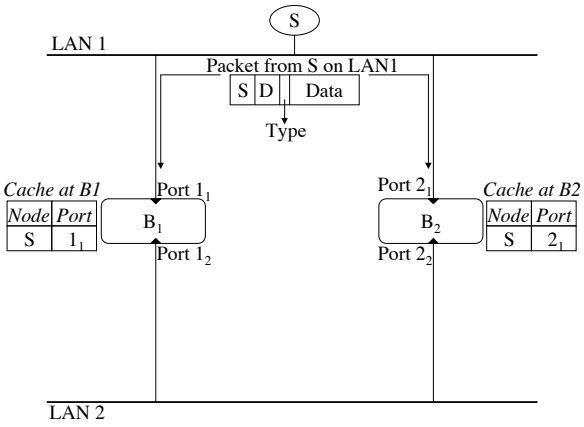

Figure 1-1: Learning bridges with loops. Such bridges suffer from packet duplication and proliferation.

the bridges along the (new) path to its destination cache the new point of attachment of the node that moved. Indeed, a variant of this method with a few optimizations is used in various wireless LAN access points to implement link-layer mobility. 802.11 ("WiFi") access points use this idea.

When there are loops in a switched LAN, significant problems arise. Consider the example shown in Figure 1.3.1, where bridges B_1 and B_2 have been configured to connect LAN1 and LAN2 together. (One reason for configuring such a network might be to add redundancy to the topology for reliability.) Consider a packet from source S transmitted on LAN1 for destination D. Since both B_1 and B_2 see this packet, they both pick it up and learn that node S is on LAN1, and add the appropriate information to their forwarding tables (caches) as shown in the picture. Then, if the bridges don't have any information in their tables for destination D, both bridges enqueue the packet for forwarding it on to LAN2. They compete for the channel (LAN2) according to the CSMA/CD protocol and one of them, say B_1, wins the contention race. When B_1 forwards the packet on to LAN2, the packet is seen by B_2.

B_2 has now seen a *duplicate* packet, but in addition, B_2 will believe that the source node S of this packet is on LAN2, when in fact it is on LAN1! The forwarding table has now been corrupted. However, B_2 would also enqueue the packet for transmission onto LAN1, because:

1. B_2 does not have an entry for D in its table, and

2. B_2 really has no way of telling that the packet is a duplicate short of carefully looking through every bit of each enqueued packet, a very inefficient process.

These problems are a direct consequence of one of the very reasons learning bridges are so attractive—their transparent behavior. Thus, the duplicate packet would continue to loop around forever (because bridges are transparent, there are no hop limit or time-to-live fields in the header).

But this looping isn't the worst part—packets in fact can *reproduce* over and over in

some cases! To see how, add another bridge B_3 between the two LANs. Now, each time a bridge sends one packet on to a LAN, the two other bridges enqueue one packet *each* for the other LAN. It's not hard to see that this state of affairs will go on for ever and make the system unusable when bridges form loops.

There are several possible solutions to this problem, including: 1) avoiding loops by contruction, 2) detecting loops automatically and informing the network administrator to fix the problem when loops are found, or 3) making packet forwarding somehow work in the presence of loops. Clearly the last alternative is the most preferred one, if we can figure out how to do this.

The trick is to find a loop-free subset of the network topology. LAN switches use a *distributed spanning tree algorithm* for this task.

■ 1.3.2 The Solution: Spanning Trees

There are many distributed spanning tree algorithms, and in this course we'll encounter different ones in the context of unicast routing, wireless routing, multicast routing, and overlay networks. Bridges use a rooted spanning tree algorithm that generates the same tree as Dijkstra's shortest path trees. The idea is quite simple: bridges elect a root and form shortest paths to the root. The spanning tree induced by the union of these shortest paths is the final tree.

More specifically, the problem is as follows. For each bridge in the network, determine which of its ports should participate in forwarding data, and which should remain inactive, such that in the end each LAN has exactly one bridge directly connected to it on a path from the LAN to the root.

Viewing a network of LAN segments and LAN switches as a graph over which a spanning tree should be constructed is a little tricky. It turns out that the way to view this in order to satisfy the problem statement of the previous paragraph is to construct a graph by associating a node with each LAN segment and with each LAN switch. Edges in this graph emanate from each LAN switch, and connect to the LAN segment nodes they are connected to in the network topology, or to other LAN switches they are connected to. The goal is to find the subset of edges that form a tree, which span all the LAN segment nodes in the graph (notice that there may be LAN switch nodes that may be eliminated by this method; that is fine, because those LAN switches are in fact redundant).

The first challenge is to achieve this goal using a distributed, asynchronous algorithm, rather than using a centralized controller. The goal is for each bridge to independently discover which of its ports belong to the spanning tree, since it must forward packets along them alone. The result of the algorithm is a loop-free forwarding topology. The second challenging part is handling bridges that fail (e.g., due to manual removal or bugs), and arrive (e.g., new bridges and LAN segments that are dynamically attached to the network), without bringing the entire network to a halt.

Each bridge has a unique ID assigned by the network administrator (in practice, bridges have a vendor-specified unique ID, but administrators can set their own IDs to arrange for suitable trees to be built). Each port (network interface) on each bridge, as well as each network interface on end point computers has a unique vendor-specified ID.

Each bridge periodically, and asynchronous with the other bridges, sends *configuration messages* to all other bridges on the LAN. This message includes the following information:

```
[bridge_unique_ID] [bridge's_idea_of_root] [distance_to_root]
```

By consensus, the bridge with the smallest unique ID is picked as the root of the spanning tree. Each bridge sends in its configuration message the ID of the bridge that it thinks is the root. These messages are not propagated to the entire network, but only on each LAN segment; the destination address usually corresponds to a well-known link-layer address corresponding to "ALL-BRIDGES" and is received and processed by only the bridges on the LAN segment. Initially, each bridge advertises itself as the root, since that's the smallest ID it would have heard about thus far. The root's estimate of the distance to itself is 0.

At any point in time, a bridge hears of and builds up information corresponding to the smallest ID it has heard about and the distance to it. The distance usually corresponds to the number of other bridge ports that need to be traversed to reach the root. It stores the port on which this message arrived, the "root port" and advertises the new root on all its other ports with a metric equal to one plus the metric it has stored. Finally, it does not advertise any messages if it hears someone else advertising a better metric on the same LAN. This last step is necessary to ensure that each LAN segment has exactly one bridge that configures itself to forward traffic for it. This bridge, called the *designated bridge* for the LAN, is the one that is closest to the root of the spanning tree. In the case of ties, the bridge with smallest ID performs this task.

It is not hard to see that this procedure converges to a rooted spanning tree if nothing changes for "a while." Each bridge only forwards packets on ports chosen as part of the spanning tree. To do this, it needs to know its root port and also which of its other ports are being used by other bridges as their preferred (or *designated*) root ports. Obtaining the former is trivial. Obtaining the latter is not hard either, since being a designated bridge for a LAN corresponds to this. When two bridges are directly connected to each other, each bridge views the other as a LAN segment.

The result of this procedure is a loop-free topology that is the same as a shortest-paths spanning tree rooted at the bridge with smallest ID.

Notice that the periodic announcements of configuration messages handle new bridges and LAN segments being attached to the network. The spanning tree topology reconfigures without bringing the entire network to a complete standstill in most cases.

This basic algorithm doesn't work when bridges fail. Failures are handled by timing information out in the absence of periodic updates. In this sense, bridges treat configuration announcements and their forwarding table entries as *soft state*. The absence of a configuration message from a designated bridge or the root triggers a spanning tree recalculation. Furthermore, an entry in the table that hasn't been used for a while may be timed out, resulting in packet flooding for that destination every once in a while. This approach trades of some bandwidth efficiency for robustness, and turns out to work well in practice in many networks.

The notion of "soft state" is an important idea that we will repeatedly see in this course, and is important to robust operation in a scalable manner. Together, periodic announcements to refresh soft state information inside the network enable *eventual consistency* to a loop-free spanning tree topology using the algorithm described above.

1.3.3 Virtual LANs

As described thus far, switched LANs do not scale well to large networks. The reasons for this include the linear scaling behavior of the spanning tree algorithm, and the fact that all broadcast packets in a switched LAN must reach all nodes on all connected LAN segments. *Virtual LANs* improve the scaling properties of switched LANs by allowing a single switched LAN to be partitioned into several separate virtual ones. Each virtual LAN is assigned a "color," and packets are forwarded by a LAN switch on to another only if the color matches. Thus, the algorithms described in the previous section are implemented over each color of the LAN separately, and each port of a LAN switch is configured to some subset of all the colors in the entire system.

1.4 Summary

Switches are specialized computers used to interconnect single-link communication media to form bigger networks. There are two main forms of switching—circuit switching and packet switching. Packet switching comes in various flavors, as do switched networks themselves. We studied some features of one of them, switched LANs.

This lecture illustrated two key ideas that we will encounter time and again:

- Soft state maintained by network elements.
- Distributed, asynchronous algorithms (spanning tree constuction in the case of LAN switches).

While LAN switches work well, they aren't enough to build a global network infrastructure. The primary reason for this is poor scalability. This approach may scale to a network with thousands of nodes (perhaps), but not to larger networks. The first scaling problem us caused by each LAN switch having to maintain per-destination information in its forwarding table. The second problem is due to the occasional flooding that's required.

A second reason for the approach described in this lecture being unattractive for a global network is that the approach may not work well over heterogeneous link technologies, many of which don't resemble Ethernet.A method for interconnecting heterogeneous link technologies is needed. This interface is what IP, the Internet Protocol, provides.

IP solves the "internetworking problem" of connecting different networks together. In the coming lectures, we will investigate its design.