ECE 333: Introduction to Communication Networks
Fall 2002

Lecture 21: Switching and Multiplexing

- Static Multiplexing and Circuit Switching
- Statistical Multiplexing and Packet Switching

Switching and Multiplexing

In the last several lectures our focus has been on medium access control techniques for broadcast networks. These networks are commonly used in a LAN environment. In this lecture, we begin to examine point-to-point networks, that is networks that consist of point-to-point links connecting routers and hosts in a mesh topology. This type of network is commonly used in a WAN environment.

First, let us consider why there is this difference between the WAN and LAN environement. One reason, as we have seen is that with most MAC techniques, performance goes down as the geographical size and the number of users increases. One solution to this is to use bridges to interconnect LANs. This still has limited use as size of network grows. Usually at most on the order of 10 LANs are interconnected this way. Another reason is differences in economics and traffic characteristics between WANs and LANs. These considerations favor using a point-to-point technology in a WAN.

Recall, from Lecture 1, that the topology of a WAN can be divided into ***hosts*** or ***users*** and the ***subnet.*** The subnet consists of ***nodes***, which are also called ***switches*** or ***routers.***

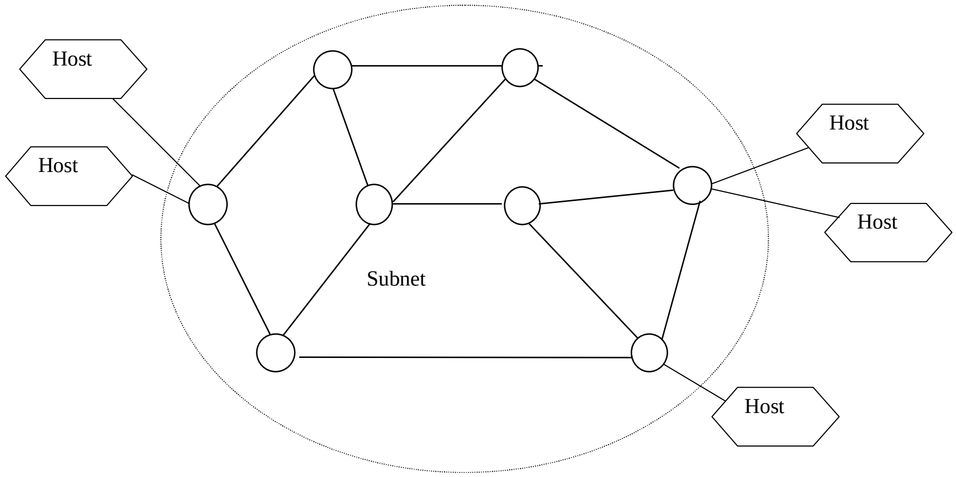

Routers/switches

Routers or switches are multi-input/multi-output devices. Each input/output is called a ***port*** and is connected to a transmission link. Switches receive data on input link and transmit the data on an output link. The details of how this is accomplished is referred to as ***switching.***

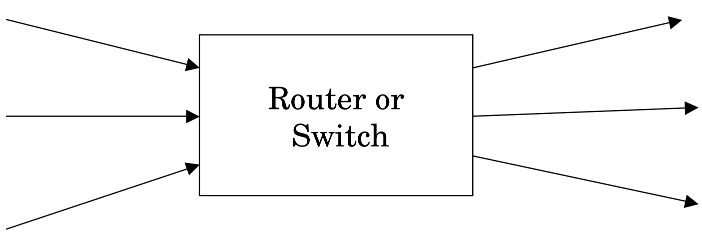

The transmission links in subnet usually carry traffic from several users. How this traffic is combined onto a link is referred to as ***Multiplexing.*** The reason for this is again primarily economic. In particular, transmission lines exhibit an ***economy of scale***, what this means is that the cost per unit bandwidth for installing a transmission line is decreasing with the bandwidth.

In this lecture we discuss switching and multiplexing in greater detail.

Switching and Multiplexing

Conceptually multiplexing is very similar to medium access control. The main difference is with medium access control, the users accessing the channel are spatially separate, so coordination and contention arise. With multiplexing the channel sharing is done at one location, namely in a switch, so contention and coordination are not as difficult (but, as we will see, these issues still can arise.) As with medium access control, multiplexing techniques can be divided into static and dynamic approaches (usually called statistical multiplexing). With static approaches each user is allocated a sub-channel, i.e. a fixed portion of the channel, while with dynamic approaches the allocation varies according to demand.

Recall, in Lecture 1 we identified 2 types of switching - circuit switching and packet switching. Circuit switching is used with static multiplexing techniques; this is the approach used in the public telephone network. Packet switching is used primarily with statistical multiplexing, this is the approach used in most data networks, such as the Internet. First we examine circuit switching and static multiplexing. We then turn to packet switching and statistical multiplexing.

Static Multiplexing Techniques

Static approaches to multiplexing include:

- Frequency division multiplexing (FDM)
- Wavelength division multiplexing (WDM)
- Time division multiplexing (TDM)

FDM: FDM is analogous to FDMA (see lecture 13). The frequency spectrum of a link is divided up into disjoint bands and a separate band is assigned to each session. FDM is used in the Cable-TV network. It is also used on legacy analog telephone connections.

Example: Analog telephone system - In the 1930's FDM began to be used in the U.S. telephone network. In the telephone network, each analog voice channel is limited to ~3kHz by filters in the phone system. Historically, long-distance carrier systems multiplexed many analog channels using FDM over high capacity links such as microwave or coax. Each voice channel is assigned to a 4kHz band to provide adequate guard bands. In the U.S., AT&T designated a ***hierarchy*** of FDM schemes, as shown in the table on the next page. A similar, but not identical system, was standardized internationally by the ITU-T.

Legacy FDM telephone hierarchy:

Number of voice channels	Spectrum	AT&T	ITU-T
12	60-108kHz	Group	Group
60	312-552 kHz	Supergroup	Supergroup
300	812-2044 kHz		Mastergroup
600	564-3084 kHz	Mastergroup	

For example, 12 voice channels can be multiplexed into ***group***; 5 groups can be multiplexed into a ***Supergroup***, etc. Various larger combinations of ***Mastergroups*** have also been defined.

WDM: When FDM is used in optical communication systems it is called WDM. In WDM systems, different sessions modulate light sources (LED's or lasers) at different wavelengths (frequencies) over the same fiber. These signals can then be separated using optical components such as diffraction gratings. Using WDM current systems can multiplex 40-128 channels onto single optical fiber at rates of 2.5 - 10 Gbps each. Low rate optical streams can be multiplexed together "all-optically" without any electrical conversion. This may also be done using passive optical components. Each WDM wavelength can be treated independently and used for analog or digital data. In many cases, WDM is combined with TDM.

TDM: TDM is analogous to TDMA. Time is divided up into slots and each session is assigned a given set of slots in a cyclical order. Each slot could correspond to the time to transmit either a single bit or byte from a user or a fixed size packet. The sequence of time-slots, 1 per user is called a ***frame***.

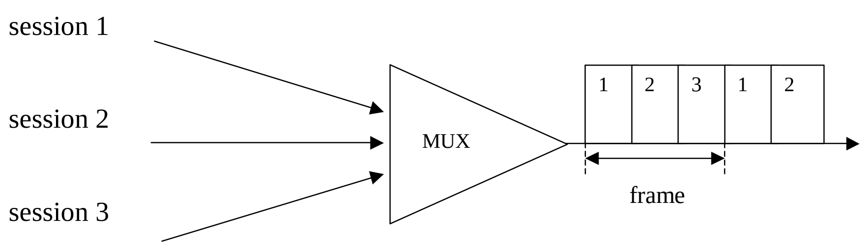

TDM is used with digital data. TDM was introduced into the public telephone network in the 1960's and has largely replaced FDM systems in the U.S.

Note with TDM (or FDM) slots are assigned to users when a call is set-up, thus no header information is required in each slot to identify a session. However one does need to ensure that the receiver is synchronized to the start of the frame. TDM systems usually provide some method for the receiver to gain frame synchronization if it is lost.

One method used for synchronization in TDM is known as ***added bit framing*** - an extra control bit is added to each TDM frame. On this "control channel" a predetermined sequence is transmitted, e.g. 10101010...The receiver searches for this pattern to synchronize, and monitors this bit to detect framing errors.

TDM is used to combine signals from constant rate incoming lines onto higher rate trunks. The synchronization of the low-rate lines is important to the operation of TDM multiplexers. For example suppose N low-rate lines are to be multiplexed at the bit level. Each line is supposed to have a rate of R bps and is to be combined onto a NR bps trunk. Thus every $1/R$ seconds the multiplexer expects to have a bit to send from a given line. If one line is not synchronized and generates bits at a rate greater than R bps, then eventually an arriving bit for this line will need to be dropped. Alternatively, if one line generates bits at a rate slower than R bps, then eventually there will not be an input bit for this line in a given frame. One solution to this is called ***bit stuffed synchronization*** - with this technique the multiplexer has slightly higher outgoing rate than NR and will periodically either add or remove extra bits to adjust the rate. The locations of these bit adjustments are signaled to the receiver so they can be accounted for.

We noted that the public switched telephone network was one example of where TDM is widely used. Besides telephone service, lines are often leased from the telephone network for connecting together nodes in WAN's. We take a closer look at the telephone network next

Public switched telephone network

The public switched telephone network basically has a hierarchical tree topology, with extra trunks added to improve efficiency (these are shown in green).

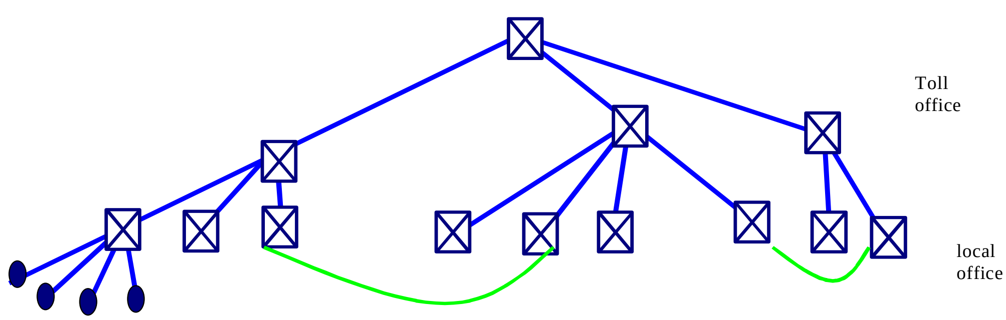

Most (residential) customers connect to a ***local central office*** through the ***local loop***. The local loop is generally copper wire and uses analog signaling. The first 3 digits of a telephone number (after the area code) identify the central office. Local central offices are connected to nearby ***toll offices*** by ***trunks***, which can be coaxial cable, microwave links, or optical fibers. The hierarchy extends up with primary secondary and regional offices, also connected via trunks. Trunks use digital signaling and TDM (in some places a combination of WDM and TDM is used.) A call between two users served by the same local office will be

switched in the local office; other calls are switched at a higher level in the hierarchy.
In U.S. there are now several parallel hierarchies. Lower branches (e.g. end offices and local loops) are divided into regional areas called ***Local Access and Transport Areas (LATA's)***; LATA's are operated by ***Local Exchange Carriers (LEC's)*** such as Ameritech, Bell Atlantic, etc. Traffic between LATA's is handled by ***Inter-eXchange Carriers (IXC's)***, such as AT&T, MCI WorldCom, etc. Each IXC has ***point of presence*** in the toll offices of a LATA that connects to the IXC's own upper branches. This is organization is changing over time due to effects such as deregulation.

At the local office, an incoming analog signal from a user is converted into a digital signal by a ***Codec*** (coder-decoder). The codec samples the analog signal at 8000 samples per seconds and quantizes each sample using 8 bits. This results in a 64 kbps stream. At the central office on the receiving end of a call, a Codec turns the 64 kbps stream back into analog signal. This sampling of the signal used in the telephone network is called: ***PCM - Pulse Code Modulation.***

TDM hierarchy in the telephone network

Digital signal are combined onto trunks using TDM. As with FDM, a TDM hierarchy has been standardized in the U.S. This hierarchy is based on the 64kbps signal generated by PCM. At the lowest level, 24 voice channels are multiplexed together at the byte level (note this correspond to one sample of a analog voice signal. As shown below, one additional bit is added to each frame for added bit framing, as discussed above. This results in a 1.544 Mbps signal; a trunk that carries this signal is called a "T1- line." Additional levels in this hierarchy are shown in the table on the next page.

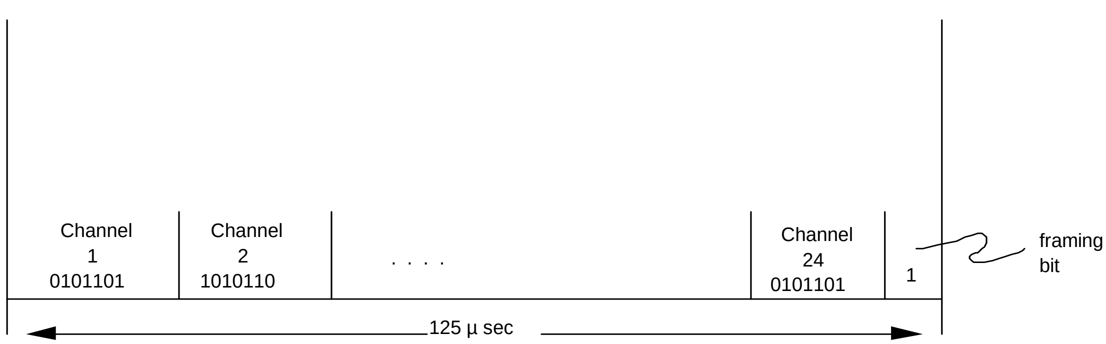

North American TDM hierarchy

# of voice circuits	Bit rate	Name
24	1.544Mbps	T1
96	6.312Mbps	T2
672	44.736Mbps	T3

A T2 and T3 frame contain extra bits used for bit-stuffed synchronization. To be precise a there are actually two set of standards the "T" standards which specifies the electrical specifications and the "DS" standards which specify the multiplexing; thus some will refer to a T1 line as a DS1 line. Most use these terms interchangeably. Outside North America and Japan a similar system is used and designated with an "E", e.g. an "E1" runs at 2.048 Mbps and contains 32 voice channels.

Organizations can lease a T1 line from the phone company for digital data service. This provides 23 channels worth of data plus a special synch byte. Each channel allows 7 bits per frame for data and 1 bit for control.

SONET/SDH

In the late 1980's, a new digital TDM hierarchy called ***SONET (Synchronous Optical NETwork)*** was proposed by Bellcore and eventually standardized by ANSI. A related standard called ***SDH (Synchronous digital hierarchy)*** was adopted by ITU-T. SONET was designed to extend the TDM hierarchy to higher rates that could be supported over optical fibers. SONET/SDH also provided an approach that unified previously incompatible hierarchies in U.S., Europe and Japan. Today, almost all long distance telephone traffic in U.S. runs over SONET links.

SONET provides a much more efficient approach to multiplexing than most previous systems and allows access to low speed channels without having to demultiplex an entire stream. This is accomplished in part by providing a master clock to all network elements and working hard to keep them synchronized. Bits are transmitted according to this master clock. SONET also provides control channels for communication between Network Elements. A major use of these channels is to provide information for ***protection switching***- the switching of circuits to alternate paths when a failure occurs.

SONET Rates

SONET specifies a TDM hierarchy starting with a basic rate of 51.84 Mbit/s, this is called an STS-1 or OC-1 signal (STS refers to the electrical standard, OC refers to the optical standard, both are often used interchangeably). In SDH the signals are called STM-n (synchronous transport mode) where n = 1,2, … Some other rates in the hierarchy are given below. The difference between the payload rate and the data rate is the bits required for control information.

ANSI SONET name		ITU SDH name	DATA Rate (Mbit/s)	Payload Rate (Mbit/s)
logical	optical			
STS-1	OC-1	--	51.84	50.112
STS-3	OC-3	STM-1	155.52	150.336
STS-12	OC-12	STM-4	622.08	601.344
STS-48	OC-48	STM-16	2488.32	2405.376
STS-192	OC-192	STM-64	9953.28	9621.504

A single fiber may use WDM to carry several wavelengths and then use SONET to time-division multiplex several lower rate streams onto each wavelength.

Circuit switching

A switch transfers signals from one input port to an appropriate output. A basic problem is then how to transfer traffic to the correct output port. In a circuit switch, the connection of input ports to output ports is determined at call set-up. This can be thought of as closing a circuit between the input and the corresponding output. In the early telephone network, operators closed circuits manually. In modern circuit switches this is done electronically in digital switches. If no circuit is available when a call is made, it will be blocked (rejected). When a call is finished a connection teardown is required to make the circuit available for another user. Switches must have some intelligence in order to perform the call setup/teardown. Note because a call set-up is required, circuit-switching networks are connection-oriented.

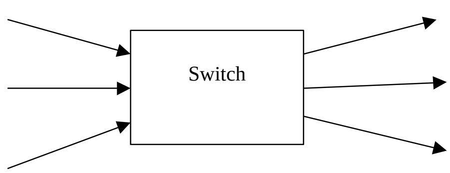

In the following we consider the design and operation of circuit switches in more detail.

Crossbar Switch - A basic example of a circuit switch is a ***crossbar switch,*** as shown below. A crossbar switch with N input lines and N output lines contains an N x N array of *cross points* that connect each input line to one output line. In modern switches, each cross point is a semiconductor gate.

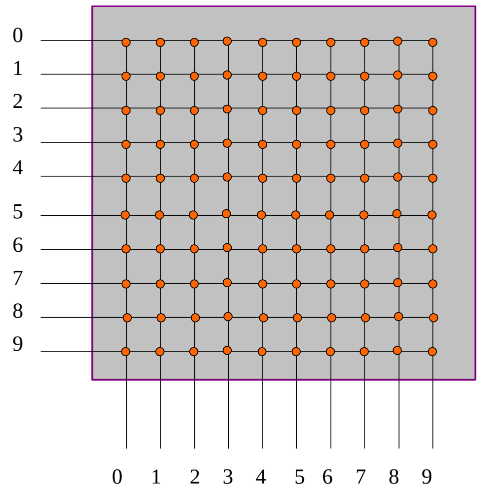

Space Division switches Crossbars are somewhat inefficient and require large number of cross-points. (The number of crosspoints is a measure of the complexity of the switch.) A more efficient design is to use several smaller crossbar switches and interconnect them. The resulting switch is called a **(multi-stage) space division switch.** An example is shown next.

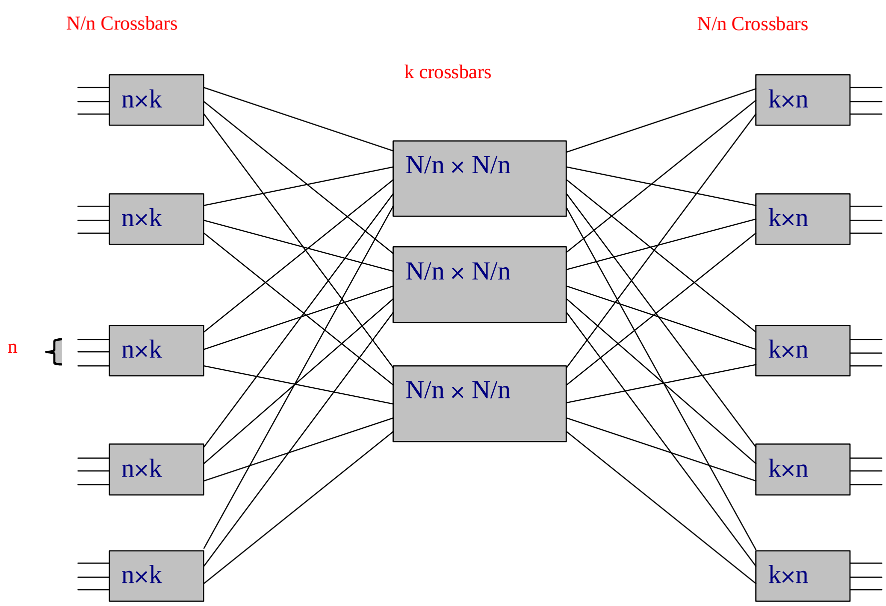

Suppose there are N incoming and outgoing ports. These are divided into N/n groups of n ports. Each group of n incoming ports is attached to a nxk first stage crossbar switch. The k output ports of each of these switches are connected to one of k second stage N/n x N/n switches. Finally there is a last stage of N/n kxn switches. An incoming call is switched through all 3 stages.

In this type of switch, the number of crosspoints is

$$2kN + k(N/n)^2$$

Notice there are $(kN)/n$ paths leaving the center stage. Thus if $(kN)/n<N$, a new connection may be blocked even if the final output line for that connection is available. In this case the switch is said to be ***blocking***. Increasing k will decrease the blocking probability, but result in a more expensive switch. It can be shown that if $k \geq 2n-1$ then the switch will be ***non-blocking***, i.e. if an input and output line are both available, a connection can be made.

When the lines coming into the switch use TDM, it is desirable to be able to switch data on any input sub-channel to any output sub-channel. One approach might be to reconfigure the switch during every TDM slot. However, this would allow a call on a given input slot to be switched only to the corresponding slot in the outgoing frame. Another approach is to demultiplex the input lines and feed each sub-channel into a port of the switch and then re-multiplex the output lines. The more common approach is to use a ***time-slot interchanger,*** to re-order the slots, before the switch. This works by reading all the slots in a frame into memory and then writing them out in another order. This requires memory access times to be fast enough to keep up with the line rate.

Multistage switches that utilize both time-slot interchangers and small crossbar switches can be made - these are called ***time-space-time*** switches.

Packet switched networks

Next we take a closer look at packet switching and statistical multiplexing. The Internet is one example of a packet switched network. A number of other packet switched network architectures have been standardized (many by the telephone companies) and used to provide leased services to customers. These include

- **X.25** - a standard based on layer 3 of the OSI reference model.
- **Switched Multimegabit Data Service (SMDS)**
- **Frame relay**
- **Asynchronous Transfer Mode (ATM)**

We will not get into the details of these networks, but will occasionally mention one as an example.

Statistical (time-division) multiplexing (SM)

With statistical multiplexing, packets are multiplexed together based on demand. One approach to this is to have all packets placed in a common queue and served FCFS. As we have seen (see Lecture 13 and Problem set 4) with bursty traffic such an approach can result in smaller delay and smaller buffer sizes than with a static multiplexing approach.

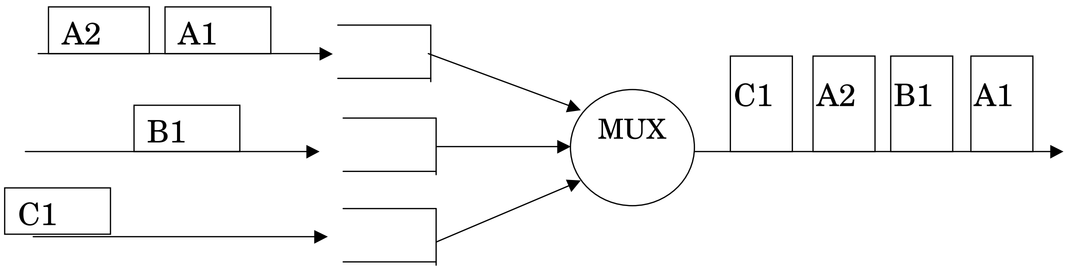

Statistical multiplexing is usually implemented at the packet level. Packets may have variable sizes or may all be a fixed size depending on the network. For example the Internet allows for variable packet sizes, while in ATM networks all packets have a fixed size of 53 bytes. The outgoing link may be synchronous and require each packet to be in a fixed slot, or it may be asynchronous allowing a packet to be sent at any time.

Statistical multiplexing

Statistical multiplexing is not without its cost. One cost is that it is much more difficult (than with static multiplexing) to guarantee service characteristics such as delay to users. Note with a static multiplexing approach the maximum delay the first packet from a user will have to wait in a multiplexer can be determined. Also for fixed rate users, with static multiplexers it can be guaranteed that no packets will be dropped, again this is more difficult with statistical multiplexing. Finally, with statistical multiplexing, users are not as isolated from each other as with static multiplexing.

Some of these issues can be addressed by implementing various service disciplines other than FCFS, such as priority service, round robin, etc. (see Lecture 11). Much of the current research in networking seeks to address how to provide users with an acceptable Quality of Service while still taking advantage of the benefits of statistical multiplexing. The service discipline used at the multiplexers is but one technique used to address this problem.

Packet Switching

In packet switched networks, a packet received at a switch is stored into memory and then forward on the correct output link when it is available. Because of this packet switching is also called ***store-and-forward*** switching. There are some variations of the above description, for example with ***cut-through switching*** a packet may be forwarded to the output port after only the header has been placed into memory. In circuit switched networks, the forwarding function of the switch is established during call set-up and based on the sub-channel data arrives in. In packet switched networks the forwarding decision must be based on information in the packet header. This is implemented in two different ways, called ***datagram switching*** and ***virtual circuit switching***.

Datagram switching- Datagram switching is a connectionless approach, requiring no call set-up. Each packet or datagram contains the destination address. The packet is forwarded according to a routing table stored in the router. This table may be static or periodically change. In datagram networks, the forwarding decision for each datagram is made independently; different packets for the same destination may take different routes and arrive out-of-order. Datagram switching is used in IP and SMDS networks.

Virtual Circuit - Virtual circuit switching is a connection-oriented approach. As in a circuit switched networks, a call set-up is required to select a route before data is transmitted, however with virtual circuit networks, no fixed-rate circuit is reserved. Data is routed along the fixed route using ***virtual circuit identifier*** (***VCI***); VCI's are usually much shorter than destination addresses. The switch may change the VCI in a packet at each hop. At call set-up, each node establishes a forwarding table entry that contains the incoming port and VCI for a connection and the corresponding outgoing port and VCI. Each session on a given link needs a unique VCI, but the same VCI may be re-used elsewhere in the network.

In some virtual circuit networks, permanent virtual circuits can be established between two locations. In this case no call set-up is required for a connection between these locations. Since virtual circuit networks require a call set-up, they can also block certain requests, for example to ensure that the delay of existing users is not too large.

Virtual circuits are used in X.25, Frame Relay and ATM networks. A virtual circuit approach, called Multi-Protocol Label Switching (MPLS) is also being standardized for the Internet.

Example of virtual circuit switching:

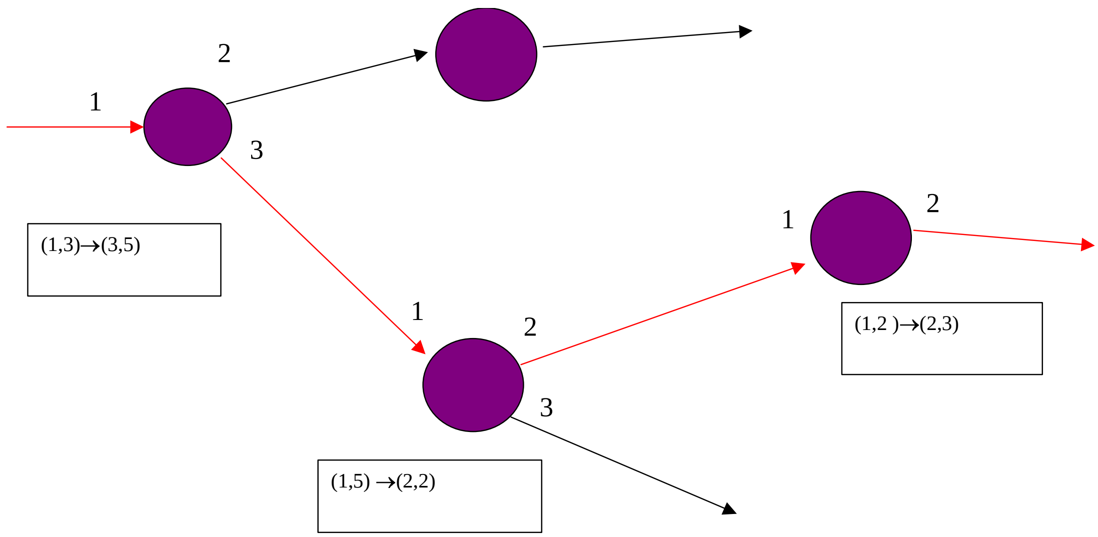

Each box shows the forwarding table entry for a connection along the path shown in red. The notation used is (incoming port, VCI)→(outgoing port, VCI). Thus at the first switch, packets arriving with VCI 1, will be forwarded to port 3 and their VCI will be set to 5.

Comparison of circuit switching & datagram switching:

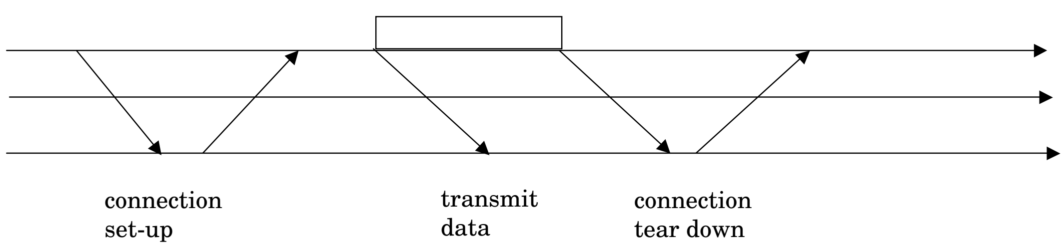

Circuit switched path

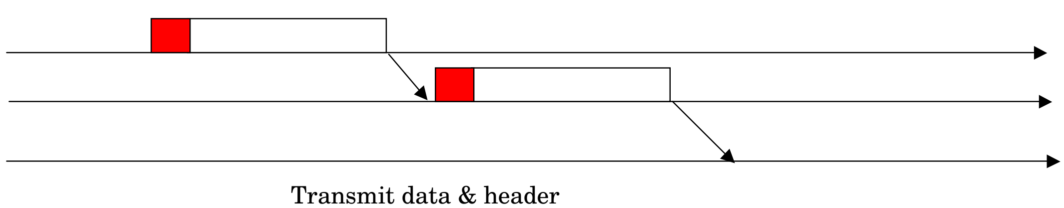

Datagram switched path

These timing diagrams show a single packet sent on a circuit switched network and a datagram switched network. Here we assume that in both cases the entire transmission rate is used for the packet. In the circuit switched case, if FDM or TDM is used, then the transmission time would be longer. For the circuit switched approach, extra delay is incurred for call set-up. While for the packet switched approach, extra delay is incurred for forwarding at intermediate nodes.

The delay incurred at intermediate nodes in store-and-forward switching can be reduced through *pipelining:*

Large packets:

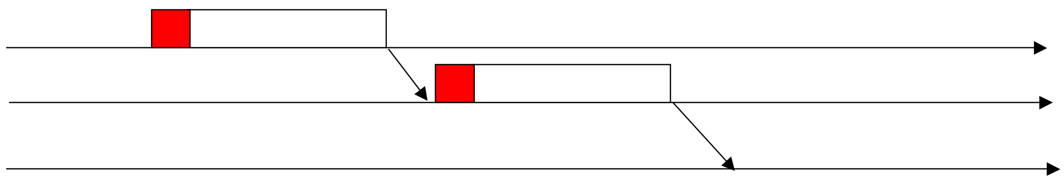

Shorter packets:

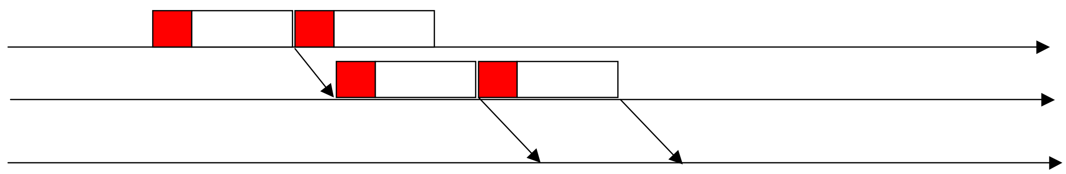

Trade-offs in packet size for pipelining:

Suppose you want to send an ***M*** bit message over ***L*** switches from a source to destination in a store-and-forward network.

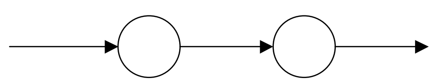

Ex. with L = 2.

Assume the message is sent as ***K*** packet*s* and each packet requires an ***H*** bit header. Also, assume that each link has a transmission rate of ***C*** bps. Thus, each packet will contain ***M/K*** bits of the message plus ***H*** header bits. (Assume ***M/K*** is an integer, if not then the last packet will contain fewer bits than the rest.) In this case the transmitted message will contain $M + KH$ bits, this will take $\frac{M+KH}{C}$ seconds for the source to transmit.

After transmission (ignoring propagation delays) it will take an additional $\frac{M/K+H}{C}$ seconds of transmission delay at each of the L switches. Thus the total delay is given by:

$$\frac{M+KH}{C} + L\left(\frac{M/K+H}{C}\right)$$

In this case, it can be shown that $K_{opt} \approx \sqrt{LM/H}$.

Comparisons of circuit switching and packet switching

- For bursty traffic, datagram switching usually achieves higher utilization than circuit switching.
- For short sessions, datagram switching can achieve less delay.
- With datagrams, end-to-end delay is variable, no guarantees. Also requires buffering and results in possible losses within network.
- With datagram switching, re-routing is easier than circuit switching, for example due to a failure.
- Datagram switching usually provides connectionless service, circuit switching is always connection oriented.
- Virtual circuit switching is a compromise. Less efficient than datagram switching, but better efficiency than circuit switching. Can provide some guarantees if admission control used.

Examples of packet switches

A basic design for a packet switch (router) is simply a workstation with mulitple interfaces, as shown below. (here NIC= Network Interface Card)

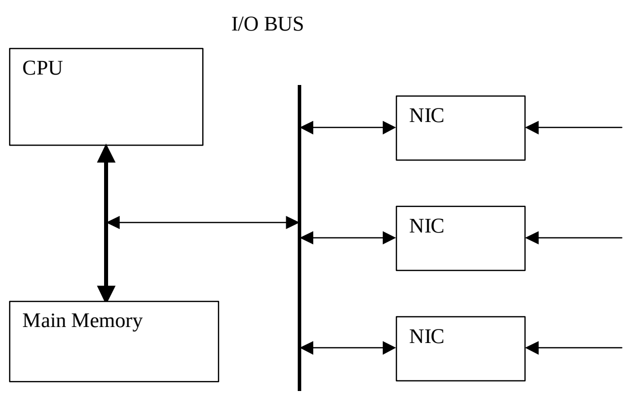

In this switch, when a packet arives, it gets transferred into memory. The CPU executes a routing table lookup for output port. The packet then gets transfered to the output NIC and forwarded. In this buffering is needed in the interfaces for each input, in case the I/O bus is busy. This design can be limited by the I/O bus speed, the memory bandwidth, or the CPU speed.

The maximum throughput of the switch is the maximum number of packets/sec that can be transferred through the switch. Suppose the I/O bus BW is limited to P packets/sec. Each packet must go over this bus at least twice - once going from the NIC to memory and once the other way. Thus the maximum throughput of the switch is limited to $P/2$ packets/sec. Suppose the switch has N input and output ports and each port is connected to a line with a transmission rate of R packets per second. Then for the bus to not be a bottleneck at high loads, we would need $P > 2NR$. Similar arguments can be made for the memory and CPU. In many cases, the speed of a switch is limited by the processing per packet, not per bit.

Several things can be done to improve rates in packet switches including to do more processing in lines cards, to use techniques for faster table look-ups and to replace the bus with a ***switching fabric***; the switching fabirc may be a cross-bar or space-division switch as discussed in the previous lecture. An example of this is shown below.

High-speed Packet Switch