Circuit Switching

This lecture describes circuit switching, a switching method that is efficient for constant bit rate (CBR) data such as uncompressed speech.
After this lecture you should be able to: draw a diagram of a network showing how data flows from one endpoint to another, list two reasons why circuit switching is more efficient than packet switching for CBR data sources, compute traffic intensity given the average call arrival rate and mean call duration, and compute blocking probability for a blocked-calls cleared model with known traffic intensity and number of trunks.

Switching

The purpose of a network is to allow communication between endpoints, also known as subscribers or terminals, depending on the context. Since it is not practical to provide a direct connections between each subscriber, a network connects these nodes to each other indirectly. Subscribers connect to a type of network node called a switch. The switches are themselves connected to each other, possibly by other switches.

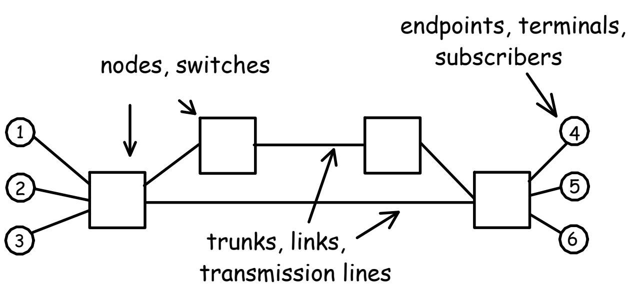

Exercise 1: What links and switches would have to be used for node 1 to communicate with node 6?

There are many types of networks. The endpoints may be telephones and other PSTN-connected devices such as fax machines, interactive data terminals (e.g. credit card terminals, smart phones, ...), or servers (public web sites, government databases, ...).

Exercise 2: Find examples of communication networks other than the PSTN or the Internet. Identify the endpoints, switching nodes and communication links.

Circuit vs Packet Switching

Switching is the mechanism that transfers data between network nodes. There are two broad classes of switching. "Circuit switching" connects one node to another for duration of a "call" while "packet switching" transfers a packet of information from one node to another until it reaches its destination.

Circuit switching was originally used for telephone networks and older data networks. Each transmission link, or time slot if a TDM link, is allocated to one call for the duration of that call.

Circuit switched networks have the following advantages compared to packet-switched networks:

- Signalling, the transfer of data required to start (set up) or end (terminate, tear down) a call, is only necessary at the start and end of a call. A packet-switched network has signalling overhead for every packet.
- There is less delay because it is not necessary to receive the whole packet before forwarding it on to the next network node.

Circuit-switched networks are not efficient for transmission of bursty data because transmission channels (or time slots) will be unused much of the time. For these uses packet switching works better.

Packet switching is now the predominant switching method and will be covered in detail in later lectures. The principles of circuit switching are still relevant because some packet switches networks use hybrids of packet and circuit switching. Examples include cell switching as is used in ATM networks and packet switched networks that reserve resources to avoid packet loss.

Telephony Circuit Switches

The best-known example of a circuit-switched network is the PSTN.

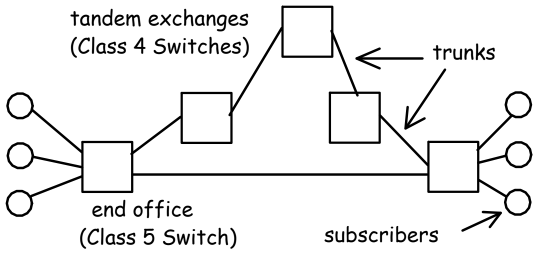

Subscribers connect over a local loop to a central office, also known as an end office. The switch at an end office is known as a "Class 5" switch that supports signalling required by subscribers (battery, ringing generators, etc.).

End offices can connect to each other directly or through trunks between "tandem exchanges." These use "Class 4" switches that are not designed to connect directly to subscribers but can handle much higher call volumes on each trunk.

Exercise 3: Assuming each CO has 10,000 directly-connected subscribers, how many CO's are required to service 3 million subscribers? What is the aggregate data rate at each CO if 10% of users are active with a data rate of 128kb/s each? What SONET rate (OC-n) would handle this data rate?

Circuit Switch Architecture

Early switches used electrical crosspoint switches that created an electrical path between two loops. The following diagram shows such a switch where the horizontal lines correspond to the links being interconnected and each 'X' indicates a potential connection.

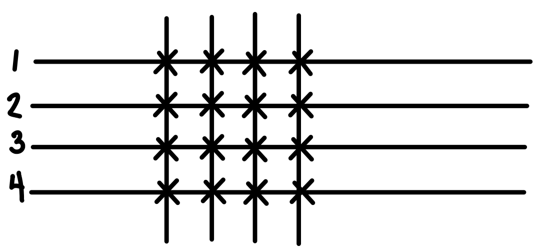

Exercise 4: What connections could be made to connect line 1 to line 4?

It is not common to design end offices that can simultaneously switch any number of subscribers because it would be very unusual for all of them to active at the same time. Instead multi-level switches are used that concentrate calls into a smaller number of lines which are then put into a crosspoint switch as shown below[1]

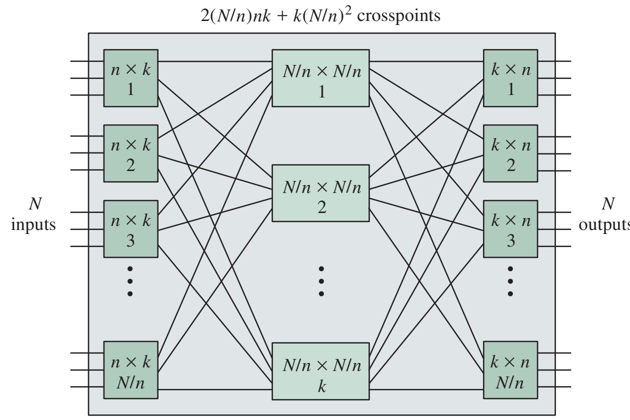

However this concentration can lead to calls being "blocked." For example, in the diagram below[2] connections have been made from input ports 5 and 6 to output ports 1 and 2. It is not possible to set up a new call to ports 3–5 because there is no free input to the third-stage (output) switch.

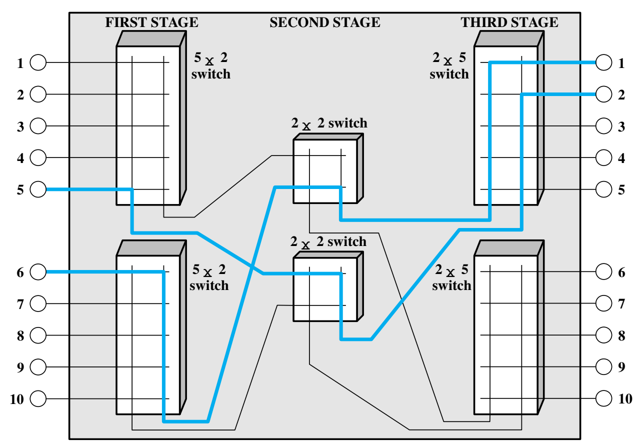

Exercise 5: What are the values of N, n and k in the diagram above?

Modern circuit switches operate on digitized signals. The switching hardware reads 8-bit PCM samples from one line card or TDM trunk time slot. It then writes the sample to a second line card or buffers it until it can be inserted into the proper time slot in

[1] From "Communication Networks" by A. Leon-Garcia et al, McGraw-Hill, 2001.

[2] From "Data and Computer Communications," 8th Edition by William Stallings, Prentice-Hall, 2007.

an outgoing TDM trunk. This is called Time Slot Interchange (TSI) switching. The switch capacity is thus limited by the speed with which it can read and write to the buffer memory.

Exercise 6: You want to design a 1000×1000 digital switch. How fast must you write and read samples from memory assuming a sample rate of 8 kHz? How much buffer memory is required?

Blocking Probability

As described above, under certain conditions the number of calls will exceed the capacity of the switch. This could be due to blocking in the switch "fabric" or because all of the trunks to other switches are busy. It is useful to be able to calculate the approximate probability that this will happen.

When a call is blocked it is possible that the call will be abandoned (cleared) or that it will be delayed (queued). The simplest model, and a reasonable one in many cases, is that blocked calls are cleared. We will assume this.

A common model for call statistics results from assuming that the probability of a subscriber deciding to make a call at any given instant of time is constant. If a call is already in progress the probability that it will end at any instant is also a (different) constant.

This model leads to relatively simple probability distributions for call arrival rate ("Poisson") and call duration (negative exponential) that can be described with a single statistic, their means (averages).

The mean call arrival rate, λ, is measured in calls per second and the mean call duration or holding time, H, is measured in seconds. The product of the two is a measure of the average number of call attempts. It is often called the offered load or call intensity, a, and is measured in "Erlangs:"

$$a = \lambda H$$

The Grade of Service (GoS) is the probability that a call will be blocked. It is usually specified for the average offered load during the busiest hour of the week rather than the weekly average (lower) or instantaneous peak (higher).

For example, the average residential busy hour offered load per subscriber is 0.05 to 0.10 Erlangs [3] with an average holding time of about 3–4 minutes.

[3] John C Bellamy, "Digital Telephony", Wiley 2001.

Exercise 7: What is the average call arrival rate?

Assuming the above model, the blocking probability is given by the Erlang-B formula:

$$P_b = \frac{\frac{a^c}{c!}}{\sum_{k=0}^{c} \frac{a^k}{k!}}$$

where c is the number of trunks or other limiting resource and a is the offered load.

The following graphs[4] show the blocking probability using the Erlang-B formula for various numbers of channels c and offered load a.

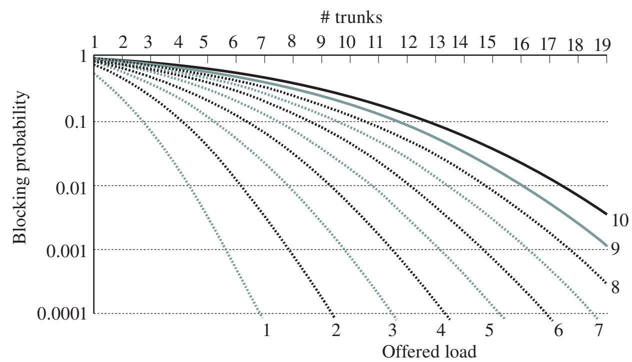

Exercise 8: A switch has 20 outgoing trunks available for 1200 users. Each users attempts to make one call every two hours and the average call time is 1.2 minutes. What is the probability that a call will be blocked? What if there were 2 trunks for 100 users?

Note that for the same blocking probability a system can handle a larger offered load per trunk as the number of trunks increases. For example, from the figure above a blocking probability of 1% with two Erlangs of traffic requires about 7 trunks (3.5 trunks/Erlang) while the same blocking probability for 10 Erlangs of traffic only requires about 18 trunks (1.8 trunks/Erlang). Thus the "trunking efficiency" of the system increases with the number of trunks.

[4] From Leon-Garcia, op cit and Rappaport, Theodore S. *Wireless Communications: Principles and Practice*. Second Edition. Prentice-Hall, 2002.

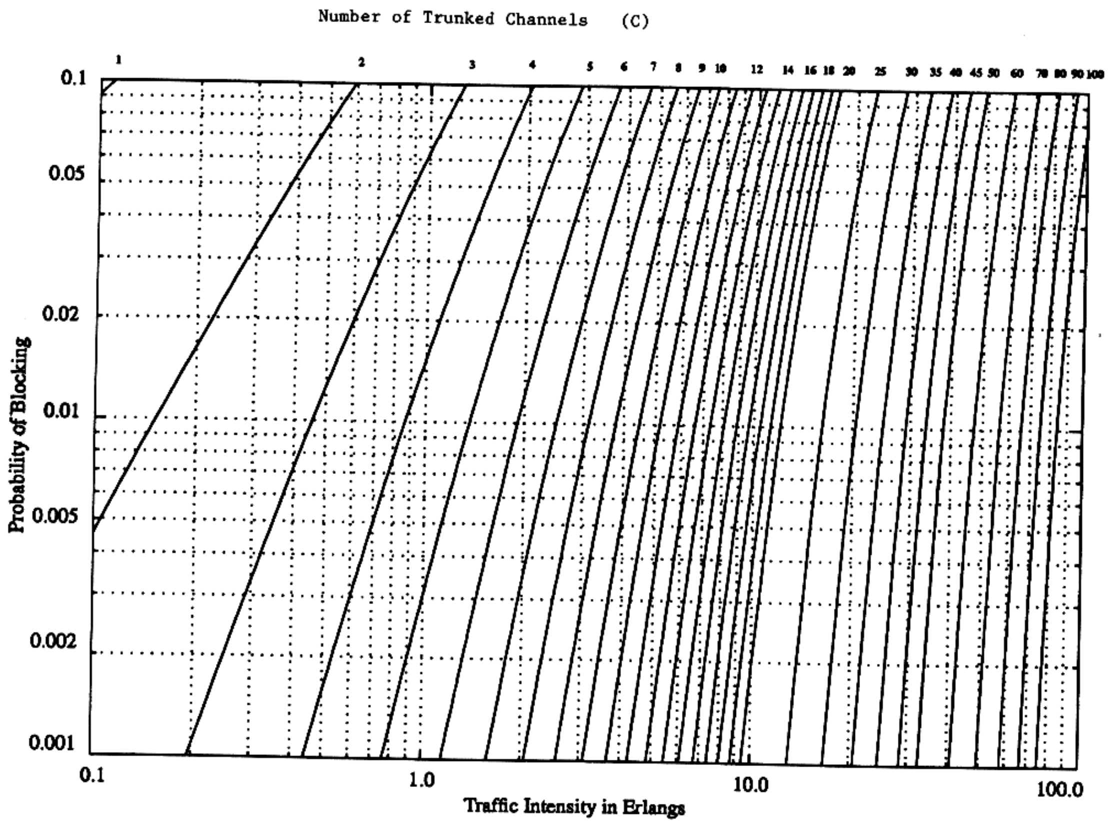

Figure 3.6 The Erlang B chart showing the probability of blocking as functions of the number of channels and traffic intensity in Erlangs.