Networking

Alex Sverdlov
alex@theparticle.com

1 Introduction

Networking is the glue that sits between all IT systems. It is used to share resources such as printers, get access to applications, get access to data, perform parallel processing, and enable communication via email or voice, etc.

2 Overview

When setting up communication between N end-points, there are a few decisions to be made, such as topology, security, etc.

Most of it boils down to sending a sequence of packets from source to destination, and expecting those packes to get there—in a way that destination can reconstruct the re-ordered sequence, and request lost packets if needed.

Below are some of the key ideas that drive networking decisions.

3 Circuit Switched vs Packed Switched

There are essentially two ways of connecting two end-points: via a circuit or via a sequence of packets. The old telephone network was circuit switched—meaning as we were dialing the phone number, the switches along the route would link up source with destination: the entire phone conversation would be conducted over that established circuit.

Lately, everyone converted to packet switched networks. Meaning that there's no circuit, and at every hop, the router determines where the packets will go—this may create a situation where packets get lost or arrive at a destination out of order, because they followed a different path.

Packet switched networks are generally fast enough that for phone conversations, any delays in packets is essentially not noticeble.

Some big corporations tweak packet switched networks to have a pre-defined (programmed) path—kinda like a circuit. This might be used by big tech companies to mirror data between east coast and west coast data centers.

4 Synchronous vs Asynchronous

Synchronous communication means that sender and reciever share a clock—or some other synthronization information. The sender cannot be sending faster than the recipient. Most communication that happens inside the computer is synchronous.

Asynchronous means there is no shared clock between a sender and reciever. It also creates a lot of other problems: such as how to notify the sender that a reciever cannot process data at a given rate.

4.1 Flow Control

Flow control is the idea that the reciever tells the sender when it is ready to recieve more data. The implmenetation is esentially to have a buffer 'window' of a certain size, where the sender may have N packets in-flight before waiting for an acknowledgement (ACK).

When the reciever gets a packet they respond with the ACK of the next expected packet—the imaginary window moves forward to enable N packets in-flight again. If a NAK (negative acknowledgement) arrives, the sender will re-transmit lost or damaged packets.

This mechanism must be paired with ARQ (automatic retransmission), meaning that if a sender never recieves an ACK or NAK, it will assume that the original data (or response) got lost, and will retransmit the packets automatically. If ARQ kicks in several times, the sender will assume the connection is gone.

5 Analog vs Digital

There are no "digital" transmissions—the signal on the wire is always analog. What differentiates analog from digital, is that a digital signal may be read, turned into a sequence of bits, and retransmitted perfectly. Analog signals may be amplified, etc., and may lose signal quality over time.

Just about all communication these days is digital in that sense.

There are still a few analog signals out there, such as AM/FM radio, and old style TV signals around the world.

6 Speed of transmission

There are several meanings of 'speed' when it comes to communication. One is capacity, and the other is actual propagation delay.

In copper and fiber, the propagation delay is 2/3rd the speed of light, or 2×10^8 meters per second. For radio and microwave, the propagation speed is almost the speed of light, or 3×10^8 meters per second.

This may seem fast, but networks and computer hardware is so fast these days that this delay is very noticeble and may be manipulated by business rivals (firms trading in both

New York and Chicago have an information gap—someone slightly faster may be able to know 'future' prices and lock in profits).

The other sense of speed of transmission is how many bits per second can be pushed through the channel. This has to do with bandwidth, and noise. Bandwidth is the block of frequencies used for transmission. More bandwidth, more bits per second.

The general formula (by Shannon) is: $C = B\log(1 + SNR)$, where B is bandwidth (in Hz), and SNR is signal to noise ratio, and capacity is how many bits per second we can get given that SNR. This formula is not very useful, since this gives us the limit, but no way to achieve it.

What is important for management to know is that there *is* a limit to how many bits can be pushed given a certain noise level.

7 Media

Media is what the data flies through. Often is it copper wire (the typical CAT-5 or CAT-7 wire), fiber, or radio.

The primary difference between CAT-5 and CAT-6 (and CAT-7) is the amount of shielding—going back to reducing the SNR from above.

There are several different fiber media. Single mode, has a thin and dense inner core with thik cladding. Multi-mode step index has a thick-ish glass inner core, with lower density cladding—the lower density cladding creates a mirror-like surface for the beam to bounce off. Multi-mode graded index has an inner core that is denser on the inside than outside—creating steps that push the signal towards the center of the inner core (avoiding mirror surface bounce as in multi-mode step index).

Radio is all the same 'media'—it is all 'light' (radio is light). The choice is essentially which frequencies and protocols are used.

8 Topology

The topology determines how things are connected. It may be a physical topology or a logical one. Below are some typical topologies:

8.1 Bus Network

A sub network is essentially a wire (the bus) with several workstations passively connected to it. Passively means they listen, but do not retransmit.

Because the signal will eventually lose power, the bus has a maximum length. Also, because most buses use a contention based protocol where they must activly listen for collissions while transmitting data, the length of the bus network is often quite short.

Needless to say, if the bus is cut, then no communication can take place.

8.2 Ring

A ring network, or token-ring, is a circle of workstations, organized as a ring, where the packets hop from workstation to workstation around the ring until they reach their destination.

A special packet, called a token, grants a workstation the ability to put new data on the ring. Otherwise, every workstations just reads and retransmits what it recieved—this is in contrast to the passive interface of Bus networks.

If the ring is cut, then data cannot make it past that point. Because ring nodes are actively reading and retransmitting data, there is a failure mode of a route workstation that reads and corrupts the data.

8.3 Star

A star topology has a main central hub or switch, to which everyone connects to. This is often the architecture of home routes, where all devices are connected to the main access point.

The failure mode is when the main hub or switch fails, then no communication can take place. If any individual workstation fails then generally the rest of the workstations can still continue to communicate.

8.4 Mesh

Mesh networks are not rings, busses, nor stars. They're often thought of as grids, but physically they may be distributed without any grid layout.

9 Persistence strategy

In a lot of networks, several workstations are sharing the same media—and need to share it. One form of sharing is not transmitting when another workstation is transmitting. Because of distances involved, that is not always an easy thing to determine.

For example, A starts transmitting at the same time as B. Each is monitoring the transmission for a collission signal—if detected, they both sound a collission detected alert. What happens next depends on the persistence strategy: A will back away, wait a random amount of time, and then listen for the media before begining to transmit again. If this happens several times, then a workstation may declare that it not capable of transmitting the data.

10 Types of networks

There are different kinds of networks. The most famous one is the Internet.

There is also the local LAN that many folks have at home. Then there are external networks—those we can connect to while outside. Phone netwroks are similar.

Then there are value-added networks: these networks connect specific users with specific systems—such as the SWIFT network, the wire transfer networks, etc.

11 Protocols

The primary protcol these days is TCP/IP. It is made up of 5-sub-protocols.

1. IP: Internet Protocol: defines the basic packet structure. IP packets are routed from IP address to another IP address.

2. ARP: address resolution protocol, used to convert an IP address into a MAC address.

3. ICMP, control protocol. It is a bare-bones IP packet used for diagnostic information. An example is a 'ping' packet.

4. UDP: user datagram packet. This is essentially an IP packet with a port address, which lets applications to send IP packets to other applications (again, IP packets go from computer to computer, and UDP goes from application to application, where application is determined by the port address).

5. TCP: a reliable connection oriented application protocol: lets applications write a stream of bytes to other applications, without the worry of packets. (under the hood, everything is just packets. TCP maintains a window, with ACKs and NAKs, ARQs, etc.)

Protcols that sit on top of TCP/IP are FTP, email (SMTP), HTTP, etc.