Ethernet*

Don Johnson

This work is produced by The Connexions Project and licensed under the
Creative Commons Attribution License †

Abstract

The ethernet is a network supporting the transfer of packets of information between computers.

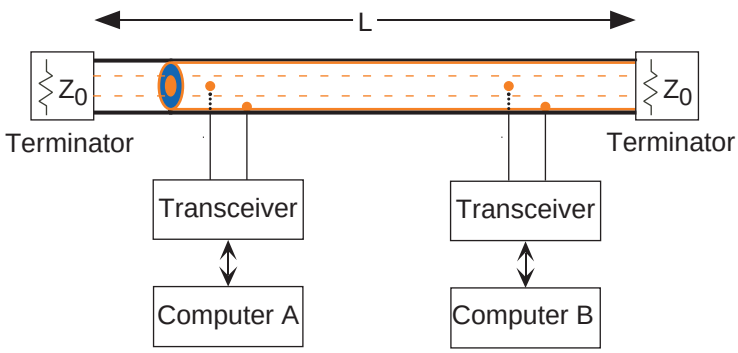

Figure 1: The Ethernet architecture consists of a single coaxial cable terminated at either end by a resistor having a value equal to the cable's characteristic impedance. Computers attach to the Ethernet through an interface known as a **transceiver** because it sends as well as receives bit streams represented as analog voltages.

Ethernet uses as its communication medium a single length of coaxial cable (Figure 1). This cable serves as the "ether", through which all digital data travel. Electrically, computers interface to the coaxial cable (Figure 1) through a device known as a **transceiver**. This device is capable of monitoring the voltage appearing between the core conductor and the shield as well as applying a voltage to it. Conceptually it consists of two op-amps, one applying a voltage corresponding to a bit stream (transmitting data) and another serving as an amplifier of Ethernet voltage signals (receiving data). The signal set for Ethernet resembles that shown in BPSK Signal Sets, with one signal the negative of the other. Computers are attached in parallel, resulting in the circuit model for Ethernet shown in Figure 2.

*Version 2.13: Jun 11, 2009 9:29 am -0500
†http://creativecommons.org/licenses/by/1.0

Exercise 1 ***(Solution on p. 4.)***

From the viewpoint of a transceiver's sending op-amp, what is the load it sees and what is the transfer function between this output voltage and some other transceiver's receiving circuit? Why should the output resistor R_{out} be large?

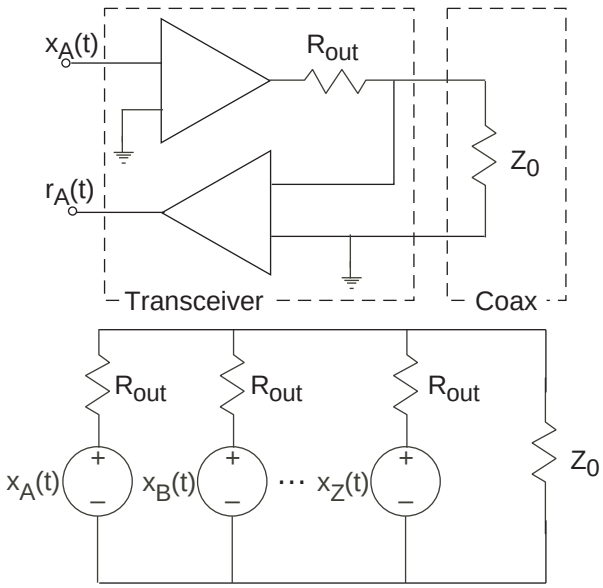

Figure 2: The top circuit expresses a simplified circuit model for a transceiver. The output resistance R_{out} must be much larger than Z_0 so that the sum of the various transmitter voltages add to create the Ethernet conductor-to-shield voltage that serves as the received signal $r(t)$ for all transceivers. In this case, the equivalent circuit shown in the bottom circuit applies.

No one computer has more authority than any other to control when and how messages are sent. Without scheduling authority, you might well wonder how one computer sends to another without the (large) interference that the other computers would produce if they transmitted at the same time. The innovation of Ethernet is that computers schedule themselves by a **random-access** method. This method relies on the fact that **all** packets transmitted over the coaxial cable can be received by **all** transceivers, regardless of which computer might actually be the intended recipient. In communications terminology, Ethernet directly supports broadcast. Each computer goes through the following steps to send a packet.

1. The computer senses the voltage across the cable to determine if some other computer is transmitting.
2. If another computer is transmitting, wait until the transmissions finish and go back to the first step. If the cable has no transmissions, begin transmitting the packet.
3. If the receiver portion of the transceiver determines that no other computer is also sending a packet, continue transmitting the packet until completion.
4. On the other hand, if the receiver senses interference from another computer's transmissions, immediately cease transmission, waiting a random amount of time to attempt the transmission again (go to step 1) until only one computer transmits and the others defer. The condition wherein two (or more) computers' transmissions interfere with others is known as a **collision**.

The reason two computers waiting to transmit may not sense the other's transmission immediately arises because of the finite propagation speed of voltage signals through the coaxial cable. The longest time any computer must wait to determine if its transmissions do not encounter interference is $\frac{2L}{c}$, where L is the coaxial cable's length. The maximum-length-specification for Ethernet is 1 km. Assuming a propagation speed of 2/3 the speed of light, this time interval is more than 10 μs. As analyzed in Problem 22[1], the number of these time intervals required to resolve the collision is, on the average, less than two!

Exercise 2 ***(Solution on p. 4.)***

Why does the factor of two enter into this equation? (Consider the worst-case situation of two transmitting computers located at the Ethernet's ends.)

Thus, despite not having separate communication paths among the computers to coordinate their transmissions, the Ethernet random access protocol allows computers to communicate without only a slight degradation in efficiency, as measured by the time taken to resolve collisions relative to the time the Ethernet is used to transmit information.

A subtle consideration in Ethernet is the minimum packet size P_{min}. The time required to transmit such packets equals $\frac{P_{\text{min}}}{C}$, where C is the Ethernet's capacity in bps. Ethernet now comes in two different types, each with individual specifications, the most distinguishing of which is capacity: 10 Mbps and 100 Mbps. If the minimum transmission time is such that the beginning of the packet has not propagated the full length of the Ethernet before the end-of-transmission, it is possible that two computers will begin transmission at the same time and, by the time their transmissions cease, the other's packet will not have propagated to the other. In this case, computers in-between the two will sense a collision, which renders both computer's transmissions senseless to them, without the two transmitting computers knowing a collision has occurred at all! For Ethernet to succeed, we must have the minimum packet transmission time exceed **twice** the voltage propagation time: $\frac{P_{\text{min}}}{C} > \frac{2L}{c}$ or

$$P_{\text{min}} > \frac{2LC}{c} \tag{1}$$

Thus, for the 10 Mbps Ethernet having a 1 km maximum length specification, the minimum packet size is 200 bits.

Exercise 3 ***(Solution on p. 4.)***

The 100 Mbps Ethernet was designed more recently than the 10 Mbps alternative. To maintain the same minimum packet size as the earlier, slower version, what should its length specification be? Why should the minimum packet size remain the same?

[1] "Information Communication Problems", Problem 31 <http://cnx.org/content/m10352/latest/#item6.25>

Solutions to Exercises in this Module

Solution to Exercise (p. 1)

The transmitting op-amp sees a load or $R_{\text{out}} + Z_0 \parallel \frac{R_{\text{out}}}{N}$, where N is the number of transceivers other than this one attached to the coaxial cable. The transfer function to some other transceiver's receiver circuit is R_{out} divided by this load.

Solution to Exercise (p. 3)

The worst-case situation occurs when one computer begins to transmit just before the other's packet arrives. Transmitters must sense a collision before packet transmission ends. The time taken for one computer's packet to travel the Ethernet's length **and** for the other computer's transmission to arrive equals the round-trip, not one-way, propagation time.

Solution to Exercise (p. 3)

The cable must be a factor of ten shorter: It cannot exceed 100 m. Different minimum packet sizes means different packet formats, making connecting old and new systems together more complex than need be.