Chapter 16
Analog and Digital Subscriber Loops

We have so far briefly touched on several aspects of the telephone system. Let us now review some of these, and extend our discussion to more of the system.

POTS

When we discussed POTS — the *Plain Old Telephone System* — also sometimes called PSTN for *Public Switched Telephone Network*, we introduced Fig. 16-1, a very simplified diagram of the overall system. At that time, we introduced the following terms:

The *transmitter* or microphone picks up your voice and converts it to an electrical signal. The *receiver* or earphone converts the incoming electrical signal back into sound.

The hybrid interfaces the two one-way or *simplex* connections (to the earphone, and from the microphone) to the two-way or *duplex* telephone line which leads to the telephone company's *central office* or *CO*. The CO contains the switching equipment ("the switch") which switches your calls in response to the number you dial.

The telephone line, called the *subscriber loop*, is a *twisted pair* — a balanced line which connects between your telephone and the central office. This line carries voice simultaneously in both directions.

At the CO, another hybrid splits the two-wire two-way subscriber loop back into two one-directional connections. The outgoing signal is converted from analog to digital, and stays digital all the way until it gets to the central office at the far end. At that point, a digital-to-analog converter converts the digital signal back to analog for transmission through the subscriber loop to the other party's telephone set.

We discussed more details of the analog-to-digital-to-analog conversion process in Chapter 14. We mentioned that the conversion is done at a sampling rate of 8000 times per second, and therefore anti-aliasing filters in the system cut off all audio above roughly 3500 Hz. We also discussed μ-law compression, which allows 8-bit samples to provide quality which otherwise would require more bits per sample.

Let us now discuss the system in more detail.

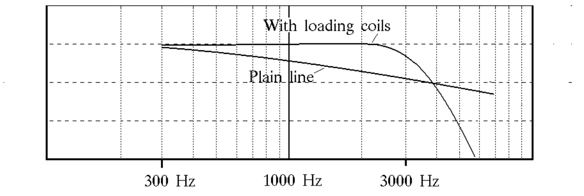

Fig. 16-2. Subscriber loop frequency response

The subscriber loop

The subscriber loop is an unshielded pair of (generally) 24- or 26-gauge wires. It is a balanced and twisted pair; the balanced connection helps to reduce the pickup of outside noise and hum, as well as crosstalk from other, adjacent wire pairs.

Because the wires are thin and close together, there is a sizable capacitance between them. This makes the circuit into a low-pass filter, which reduces the high-frequency response. Even with short cables, at just 3000 or 3500 Hz these high frequencies are attenuated and result in a noticeable lack of treble. The curve labelled "plain cable" in Fig. 16-2 shows the typical frequency response of a subscriber loop several thousand feet long.

In order to improve the high frequency response, the telephone company therefore often inserts *loading coils* into the subscriber loop. These are toroidal inductors, most often 44, 66, or 88 millihenries (see Fig. 16-3), which are connected in series with the line. These are not as common in large cities, where the distance from your phone to the nearest CO may be fairly small, but appear quite often in the suburbs or out in the country.

The loading coil is also sometimes called a *peaking coil*; it resonates with the line capacitance and produces a peak in the frequency response somewhere between 3000 and 4000 Hz; this increases the high-frequency re-

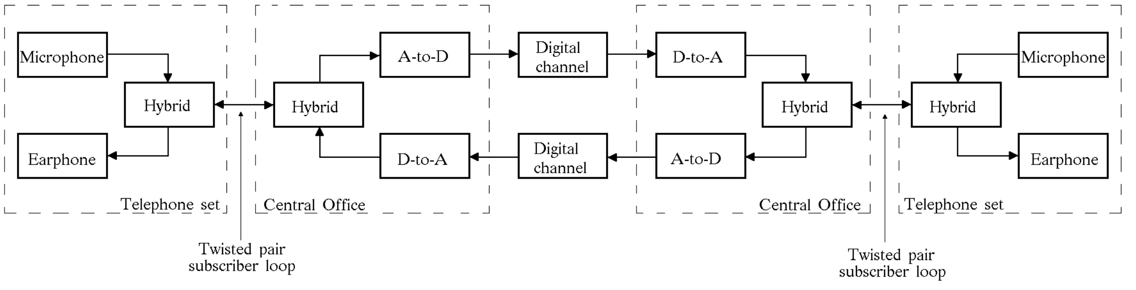

Fig. 16-1. Overall block diagram of the telephone system

Fig. 16-3. 88-mhy loading coil

sponse of the line, as shown in the “with loading coils” curve in Fig. 16-2. But you can see that, although the loading coil improves the frequency response in the high audio range up to about 3500 Hz or so, it actually makes things worse above that. (It also affects the phase of signals.) Signals above 4000 or 5000 Hz (as well as harmonics of any pulse or digital signals) are now almost totally attenuated. This explains why, when the telephone company decides to use an existing pair of wires for a digital circuit, it must remove any loading coils that were previously used for voice circuits.

Your telephone set

Also called a subscriber set (part of the CPE or *Customer Premises Equipment*), today your telephone is an analog instrument which converts between sound and electrical signals. In addition to the microphone and earphone, the telephone set also contains the dial (either a pulse dial or a tone dial which emits *DTMF — Dual Tone Multiple Frequencies*) and a bell (called a *ringer* in telephone parlance.) Fig. 16-4 shows a very simplified circuit of how the instrument connects to the rest of the network.

In the central office (shown at the right of Fig. 16-4) is a 48-volt battery. This is generally a series of wet cells (i.e., batteries with a wet electrolyte, much like the lead-acid battery in a car) connected in series to provide 48 volts. The advantage of using batteries is that they provide power even if the local electric utility fails. These batteries are constantly being charged, however, so that the actual voltage is a bit higher — typically 50 volts or so.

When you are not using the telephone, the hook switch (*switch hook* in telephone talk) is open, and only the ringer is connected to the line. A capacitor in series with the ringer prevents dc current flow through the ringer; hence the telephone set looks like an open dc circuit and there is no current through the loop. Even though there is some resistance in the subscriber loop (as well as in the circuitry in the CO), the absence of current means there is no voltage drop, and so the full 50 volts appears across your telephone set. The lack of current tells the CO that your phone is *on hook*.

When there is an incoming call, the telephone company rings your phone by sending an ac ringing voltage of 100 volts at 20 Hz. Since there is a capacitor in series with your ringer, the ac ringing signal is sent to the ringer. The ringer requires very little current, and thus rings.

When you pick up the handset to answer the phone, you close the hook switch under the handset. The name dates back to when an earphone or handset would hang on a hook at the side of the phone; the switch would close when you picked up the handset, and this was called *going off hook*). When the hook switch closes, dc current can now pass through the telephone. We show in Fig. 16-4 that the dc current passes through a subscriber relay; its closing then tells the telephone company that you have picked up the phone, at which point

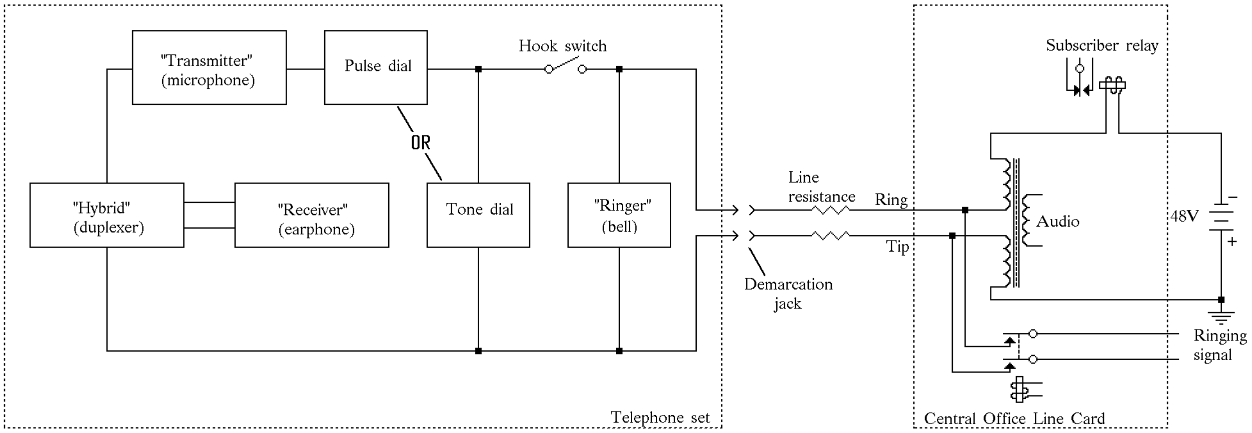

Fig. 16-4. Block diagram of a typical telephone set

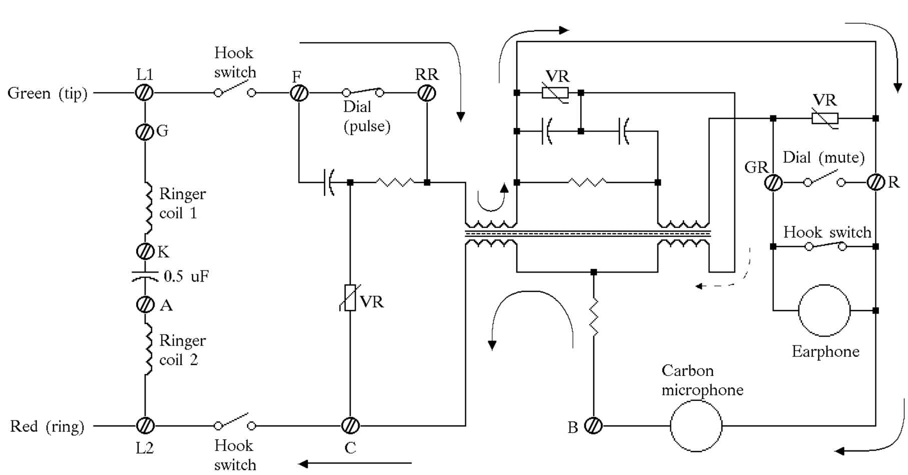

Fig. 16-5. Actual telephone schematic diagram

We have discussed the remaining major parts of the telephone instrument in earlier chapters — the dial in Chapter 11, and the hybrid in Chapter 12.

The Line Card

Located at the telephone company's central office, the switch does the switching of calls. Much of the circuitry in the switch is common to all users, but some circuitry must be duplicated for each line. As a result, all the incoming subscriber loops terminate in *line cards*, which handle just a small number of subscribers each.

Technicians often use the word BORSCHT to remind themselves of all that the line card does (Borscht is a kind of Russian red beet soup.) These letters stand for the following:

- **B**attery. It connects –48 or –50 volts dc to your line.
- **O**vervoltage protection, to protect the switch from lightning, short circuits to power lines, and similar problems.
- **R**inging. It connects the 100-volt 20 Hz ringing signal to ring your bell.
- **S**upervision. It monitors the dc current through your line to determine when you pick up your phone or hang up.
- **C**oding. It has the analog-to-digital and digital-to-analog converters, as well as the necessary anti-aliasing filters.
- **H**ybrid.
- **T**esting. The line card allows the switch to perform various testing on your line to make sure all is well.

Fig. 16-4 implied that the line card contained relays, transformers, and other large components. That is the way it used to be, but today's line cards are built with integrated circuits and other solid-state components. Fig. 16-7 shows a simplified diagram of the typical line card.

The incoming analog signal from the subscriber loop is connected to a specialized integrated circuit called a SLIC — a *Subscriber Line Interface Circuit.* Along with a few other components, the SLIC provides about half of the BORSCHT functions. It does not contain a ringing generator or battery, since these are external and

it stops ringing your bell. (In modern central offices, an integrated circuit is used instead of a relay to sense that dc current, as we will see shortly.)

Fig. 16-5 shows the schematic diagram of an actual (older!) telephone set. Keep in mind that there are many ways to wire a phone; this diagram shows an older 500-style rotary (pulse) dial set of the type in Fig. 16-6, and your phone may vary.

The ringer circuit consists of two coils separated by a 0.5 μF capacitor. The capacitor prevents dc current from flowing through the coils. This makes sure that the CO doesn't mistake the bell for an off-hook telephone.

The arrows in the diagram show the path of the dc current through the phone when it is off hook. Despite the fairly-large 50-volt battery in the CO, because of the various resistances in the circuit (in the CO, the wiring of the loop, and the telephone set itself), the current is typically limited to somewhere between 10 and 30 milliamperes (depending on the loop length and resistance.)

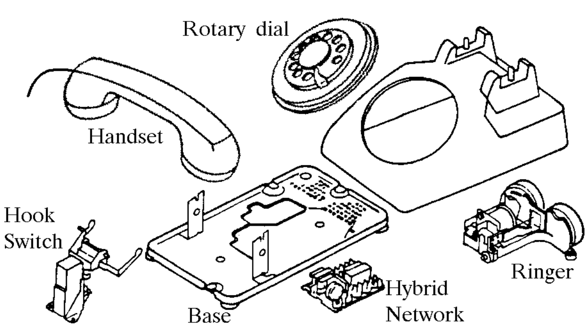

Fig. 16-6. A disassembled older telephone instrument

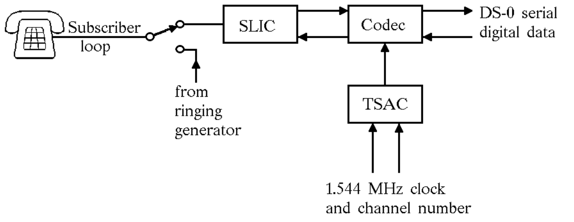

Fig. 16-7. Simplified block diagram of a line card

shared by all the line cards, but it controls their signals. The related components provide the overvoltage protection, and the SLIC handles supervision and testing, and contains the hybrid.

The SLIC's hybrid converts the two-wire full duplex subscriber loop into two one-way connections, and these go to the coder/decoder or codec. This is another special-purpose integrated circuit which handles the entire job of converting an analog signal into digital format, and also converting the incoming digital signal back into analog. It contains all the essential parts needed to do A-to-D and D-to-A conversion, including an anti-aliasing filter and a sample-and-hold circuit. It also does the μ-law compression and expansion. The codec does one more important job — it converts its serial data to and from serial, and even multiplexes it into a TDM signal.

As we have discussed earlier, the codec converts the telephone signal into 8-bit μ-coded binary numbers. It therefore outputs 8-bit numbers for the outgoing signal, and inputs 8-bit numbers for the returning signal, for a total of 16 bits. There is one codec for each telephone customer; if every codec had 16 wires carrying parallel data back and forth, the switch would be terribly complex. Thus a multiplexed serial connection is absolutely necessary to keep things within reason.

We will see how this is all done in the next chapter, but in the meantime here is a brief description. The codec outputs 64K bits per second: at a sampling rate of 8000 samples per second, it outputs an 8-bit number every $1/8000$ second; that is, once every 125 microseconds. But each of the eight bits lasts only about 0.65 microsecond, so that the eight bits come out in a short burst that only lasts about 5 microseconds. That leaves about 120 microseconds before the next burst; during this time, up to 23 other codecs can also squeeze in their 8-bit numbers. Thus 24 codecs can all share the same wire at the same time. (Actually, dozens of customers and their codecs can share the same TDM connection, since it is unlikely that they will all want to make a call at exactly the same time.)

The TSAC in Fig. 16-7 is a *Timing Slot Assignment Circuit*. Its job is to control the codec's timing, and tell it exactly when to output or input data via the serial connection. This will become clearer in the next chapter, when we discuss T carrier systems.

To recap: the normal audio telephone network sends analog audio from your telephone, through the subscriber loop, to the central office, where it is sampled and converted into a digital signal. It is carried digitally from then on, at a rate of 64K bps, until it is converted back to an analog audio signal in the central office at the other end, just before it is sent to the person you are speaking with. The sampling is done at an 8 kHz rate, with an 8-bit analog-to-digital converter, for an effective transmission rate of 64K bps within the network itself. But due to the analog subscriber loop, the maximum modem speeds available (as we discussed in Chapter 12) are about 52K bps in one direction and 33K bps in the other. Wouldn't it be nice if the entire 64K bps rate were available to the customer? Such is the case with ISDN.

ISDN — Integrated Services Digital Network

ISDN is a purely digital connection from the customer to the central office. It has been around quite a few years, is always on the verge of becoming popular ... and may already be a has-been.

ISDN changes the picture by moving the sampling and A-to-D / D-to-A conversion (if you use it for voice) from the central office back to your home or office. It is still done, but it is now the job of your equipment to do it. At the expense of some complexity, you get access to the full 64K bps data rate — and more.

There are two kinds of ISDN: Basic Rate Interface or BRI, and Primary Rate Interface or PRI. (There is a third kind called *Broadband ISDN*, but it is not strictly ISDN at all.)

Basic Rate Interface ISDN

BRI is intended for the home or small office. It provides two 64K data paths (called *Bearer Channels*) and one 16K bps *Data Channel* (plus some overhead bits), all over *one pair of wires* just like a normal subscriber loop. The system is often called 2B+D for that reason.

In terms of hardware, BRI is connected to the customer through the same pair of wires that used to handle an analog connection. At the central office switch, BRI ISDN requires a different line card plus some changes to the switch programming. (These changes are expensive, and so ISDN may not be available in areas where there are not enough customers to justify the extra cost.)

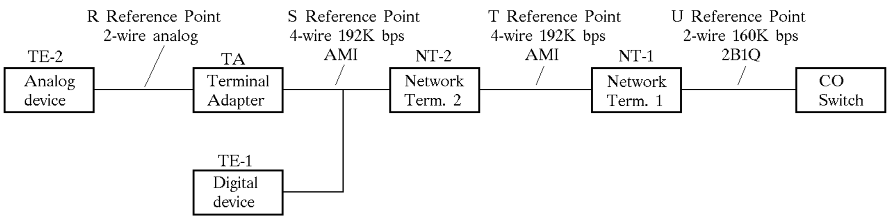

Fig. 16-8. BRI ISDN components

Of the 2B+D channels, the D or Data channel is used for signalling between your system and the central office, such as for dialing a number or ringing the phone, but it can also handle digital data. Each of the two B or Bearer channels can carry one 64K bps voice signal, which makes it attractive from the telephone company point of view — they can provide two voice circuits (plus a data circuit) over just one pair of wires.

From a user point of view, there are a number of advantages. Foremost is that, if you call from your ISDN line to someone else that also has an ISDN line, the 64K bps data you send out is delivered completely unchanged to the person you called. In other words, you need not use this channel just for voice — it can also be used for computer-to-computer or fax-to-fax communications, or for any other purpose that requires a direct data transfer. If 64K bps is not fast enough, you can *bind* the two B channels together; that is, you can use the two channels together to get 128K bps data transfer. (But note that this requires that you place, and pay for, two telephone calls, one for each channel. It also requires some additional equipment.)

NT-1 and the U Reference Point

Fig. 16-8 shows how a BRI ISDN line can be used. Shown on the right side is the connection between the central office switch and a *Network Terminator 1* called an NT-1. This is a 2-wire line — a normal subscriber loop, except that loading coils and any taps must be removed because it carries pure digital data (although short taps are sometimes permissible.)

This part of the circuit is called the *U Reference Point* or *U interface*. The line may be up to 18,000 feet (approximately 3½ miles) long. This may create some problems in outlying areas, but there are repeaters that can be inserted to lengthen the line.

The NT-1 terminator primarily terminates the line, converts the 2-wire U line to a 4-wire system, and protects your equipment against lightning and other faults. It does require power, though, and here is an interesting aspect of ISDN: normal analog POTS lines are powered from the central office (where the telephone company maintains sizable batteries and diesel generators), and generally work even when there is a power failure in your neighborhood. ISDN lines, on the other hand, do not provide power to your equipment (except in Europe). You must provide your own power to the NT-1 as well as all other equipment at your end. Fortunately, much ISDN equipment provides battery backup for use in emergencies, but this seldom provides much running time. Hence it is not a good idea to rely on ISDN service for your only telephone.

Sending 160 K bps digital data through a subscriber loop (especially one approaching 18,000 feet) is not easy, so the designers use a technique similar to that of modems — carrying two bits on each symbol. Each symbol has four possible values, and so the system is called 2B1Q — 2 bits on 1 quaternary symbol. In this case, however, the data is sent in pure digital form, with the four symbols being four different voltage levels:

1 0 = +3 volts
1 1 = +1 volt
0 1 = –1 volt
0 0 = –3 volts

Fig. 16-9 shows an idealized 2B1Q signal on the U interface; it is actually quite distorted by the time it travels down the line. Further, since the U connection is bidirectional, there are two of these signals traveling on the line at the same time.

Remember that the data on the U interface is 160K bps; the 2B1Q method reduces that to 80K baud; that works out to an effective 40 kHz, which is more manageable. You can see this in Fig. 16-9: the first four bits of data essentially make up one cycle of the signal; the waveform therefore has ¼ as many cycles as bits.

Worldwide, there is some confusion as to the U signal because it was never specified in the CCITT / ITU-T standard. In the US, the NT-1 is to be provided by the customer, and so the U interface is the demarcation be-

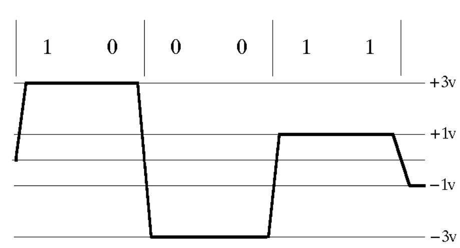

Fig. 16-9. 2B1Q signal

tween the telephone company and the customer. In other parts of the world, the NT-1 is provided by the telco, and the demarcation between telco and user comes after the NT-1; the U interface is considered proprietary by the telco. Because of this lack of a U standard, early US ISDN systems used incompatible U interface standards, and so some early NT-1 units do not work on modern systems.

NT-2 and the S and T Reference Points

The duplexer in the NT-1 splits the 2-wire U signal into a 4-wire system, with one pair in each direction, called the T (or S/T) Reference Point. At this point, the signal is a 192K bps Alternate Space Inversion (ASI) signal.

ASI is similar to AMI, the Alternate Mark Inversion shown in Fig. 13-9, but in this case a ONE is represented by 0 volts, while a ZERO is 1 volt; but the ZEROs alternate polarity. That is, one ZERO may be +1 volt, while the very next ZERO would be –1 volt, and the following again at +1 volt.. Thus alternate ZEROs (spaces) are inverted. Although this complicates the circuitry somewhat, this scheme has two advantages. First, as in 2B1Q, alternating polarities reduces the effective bandwidth of the signal. Second, if the receiver ever gets two ZEROs of the same polarity, it knows there has been an error; this is called a *bipolar violation.*

The NT-2 Network Terminator 2 is simply a splitter, which splits the T signal into a number of S signals for use by several devices. If the ISDN line is used for only one purpose, the NT-2 can be omitted, since the S and T signals are otherwise the same. In that case, the output of the NT-1 is called the S/T Reference Point.

TE-1 ISDN Devices

A number of actual devices can be connected to the S or S/T signals — voice phones, faxes, computer terminals, etc. These break down into two types: ISDN devices are called Terminal Equipment type 1 or TE-1, while non-ISDN devices are Terminal Equipment type 2 or TE-2.

TE-1 devices have a digital interface, and can directly talk to the S or S/T interface. For example, a TE-1 telephone would be a digital phone, which does its own a-to-d and d-to-a conversion.

ISDN fax machines can also be used; as opposed to ordinary analog (group 3) fax machines, the TE-1 machine is called a Group 4 fax, and it is purely digital. ISDN fax machines are still rare and expensive, but are five to six times faster than an ordinary analog fax machine.

ISDN computer terminals (or computers) require a bit more than just the plain RS-232 interface, because they must be able to command the NT-1 to dial out.

TE-2 Non-ISDN devices

TE-2 or non-ISDN devices are the plain, old analog kind, like a desk telephone or analog fax machine. Since these cannot talk to a digital S or S/T interface, they need external conversion to and from the digital world. The TA or Terminal Adapter in Fig. 16-8 does this job.

There is obviously no advantage in using an ISDN line plus a TA just to use the same analog devices that could be used on a plain POTS line. Still, the TA is useful during the transition from analog to all-digital. It is also potentially useful if the telephone company can only provide one wire pair, yet you need two circuits.

Multiple devices

The 2B + D ISDN line uses just one physical line to connect to the central office, but has three actual logical channels. Each of the two B channels can be used for a separate call (or they can be bonded together for higher speed), while the D channel, although normally used for controlling the B calls, can also be used as a separate data channel for packets. One ISDN line can therefore handle up to three simultaneous tasks.

But the S or S/T interface can handle up to eight different devices. Moreover, from the central office point of view, each of these can have a different telephone number. That is possible on a plain analog POTS line today (with distinctive ringing on one line and a corresponding switch box), but ISDN carries it further with more options.

The phrase “multiple devices” has another meaning as well. Although Fig. 16-8 shows the TE1, TE2, TA, NT1, and NT2 as separate boxes, a number of vendors are combining them in various combinations. For example, Motorola makes a device called a BitSurfr Pro, which combines an NT1, NT2, and two TA’s into one package. It hooks up to the ISDN line at one end, and provides an S interface plus two analog lines.

Primary Rate Interface (PRI)

Unlike BRI, which is two 64K bearer channels and one data channel and is designed for home or small office applications, PRI has twenty-four 64K channels, of which 23 are B channels, and one is a 64K D channel. As with BRI, the D channel is used for control purposes, such as sending dialing data or receiving a ringing signal. PRI is designed for users who need a large bandwidth, or who intend to parcel out the bandwidth to individual users. For example, a common application of PRI is to connect a PBX — a Private Branch eXchange, a local switchboard which services a number of telephones — to the central office.

PRI ISDN uses 24×64K bps, or 1536K bps, plus an additional 8K bps for some control information, for a total of 1544K bits per second. This is too fast for a subscriber loop, and is thus provided over a special connection called a T1 line (which we shall discuss in the next chapter.) As a result, PRI ISDN is generally a service for business, not individuals.

Both BRI and PRI have some interesting characteristics. Unlike a normal analog telephone subscriber loop, where dialing is done with either a rotary dial or a DTMF tone dial , an ISDN device dials by sending digital dialing info to the switch via the D channel. At either 16K bps (for the BRI system) or 64K bps (for the PRI system), this is much faster than analog dialing. Combined with the speed of today's digital switches, this means that the typical call can be established within a second or two. Moreover, when ISDN is used for digital data, all of the time delays normally associated with modems (where two modems may spend 10 or 20 seconds negotiating what speed and protocol they will use) are not required. Hence it only takes a second or two from the time you place a call until you're ready to send data. It thus becomes practical to hang up a connection, and dial again the next time you have something to send. You therefore need not pay for a call when nothing is being sent.

Likewise, when the switch rings your phone, it does not send the normal AC ringing signal to ring a bell. Instead, it sends a digital code via the D channel. This code not only tells your equipment to answer, but also identifies the type of device it wants to talk to. For example, on a normal analog line, when an incoming call comes in, your phone, your modem, and your fax machine all ring. On an ISDN line, only the device being addressed responds, and does so right away. Again, no delays.