

CHAPTER 7

Understanding Telephony

Utility is when you have one telephone, luxury is when you have two, opulence is when you have three—and paradise is when you have none.

—Doug Larson

We're now going to take a break from Asterisk for a chapter or two, because we want to spend some time discussing the technologies with which your Asterisk system will need to interface. In this chapter, we are going to talk about some of the technologies of the traditional telephone network—especially those that people most commonly want to connect to Asterisk.* While tomes could be written about the technologies in use in telecom networks, the material in this chapter was chosen based on our experiences in the community, which helped us define the specific items that might be most useful. Although this knowledge may not be strictly required in order to configure your Asterisk system, it will be of great benefit when interconnecting to systems (and talking with people) from the world of traditional telecommunications.

Analog Telephony

The purpose of the Public Switched Telephone Network (PSTN) is to establish and maintain audio connections between two endpoints.

Although humans can perceive sound vibrations in the range of 20–20,000 Hz,† most of the sounds we make when speaking tend to be in the range of 250–3,000 Hz. Since the purpose of the telephone network is to transmit the sounds of people speaking, it

* We'll discuss Voice over IP in the next chapter.

† If you want to play around with what different frequencies look like on an oscilloscope, grab a copy of Sound Frequency Analyzer, from Reliable Software. It's a really simple and fun way to visualize what sounds "look" like. The spectrograph gives a good picture of the complex harmonics our voices can generate, as well as an appreciation for the background sounds that always surround us. You should also try the delightfully annoying NCH Tone Generator, from NCH Swift Sound.


was designed with a bandwidth of somewhere in the range of 300–3,500 Hz. This limited bandwidth means that some sound quality will be lost (as anyone who’s had to listen to music on hold can attest to), especially in the higher frequencies.

Parts of an Analog Telephone

An analog phone is composed of five parts: the ringer, the dial pad, the hybrid (or network), and the hook switch and handset (both of which are considered parts of the hybrid). The ringer, the dial pad, and the hybrid can operate completely independently from one another.

Ringer

When the central office (CO) wants to signal an incoming call, it will connect an alternating current (AC) signal of roughly 90 volts to your circuit. This will cause the bell in your telephone to produce a ringing sound. (In electronic telephones, this ringer may be a small electronic warbler rather than a bell. Ultimately, a ringer can be anything that is capable of reacting to the ringing voltage—for example, strobe lights are often employed in noisy environments such as factories.)



Ringing voltage can be hazardous. Be very careful to take precautions when working with an in-service telephone line.

Many people confuse the AC voltage that triggers the ringer with the direct current (DC) voltage that powers the phone. Remember that the ringer will not respond to DC voltage, and you’ve got it.

In North America, the number of ringers you can connect to your line is dependent on the Ringer Equivalence Number (REN) of your various devices. (The REN must be listed on each device.) The total REN for all devices connected to your line cannot exceed 5.0. An REN of 1.0 is equivalent to an old-fashioned analog set with an electromechanical ringer. Some electronic phones have an REN of 0.3 or even less.

Dial pad

When you place a telephone call, you need some way of letting the network know the address of the party you wish to reach. The dial pad is the portion of the phone that provides this functionality. In the early days of the PSTN, dial pads were rotary devices that used pulses to indicate digits. This was a rather slow process, so the telephone companies eventually introduced touch-tone dialing. With touch-tone—also known as Dual-Tone Multi Frequency (DTMF)—dialing, the dial pad consists of 12 buttons. Each button has two frequencies assigned to it (see Table 7-1).

Table 7-1. DTMF digits

	1209 Hz	1336 Hz	1477 Hz	1633 Hz ^a
697 Hz	1	2	3	A
770 Hz	4	5	6	B
852 Hz	7	8	9	C
941 Hz	*	0	#	D

^a Notice that this column contains letters that are not typically present as keys on a telephone dial pad. They are part of the DTMF standard nonetheless, and any proper telephone contains the electronics required to create them, even if it doesn't contain the buttons themselves. (These buttons actually do exist on some telephones, which are mostly used in military and government applications.)

When you press a button on your dial pad, the two corresponding frequencies are transmitted down the line.

We assume that you've used a telephone, so we won't spend any more time on DTMF.

Hybrid (or network)

The hybrid is a type of transformer that handles the need to combine the signals transmitted and received across a single pair of wires in the PSTN and two pairs of wires in the handset. One of the functions the hybrid performs is regulating *sidetone*, which is the amount of your transmitted signal that is returned to your earpiece—its purpose is to provide a more natural-sounding conversation. Too much sidetone, and your voice will sound too loud; too little, and you'll think the line has gone dead.

Hook switch (or switch hook). This device signals the state of the telephone circuit to the CO. When you pick up your telephone, the hook switch closes the loop between you and the CO, which is seen as a request for a dial tone. When you hang up, the hook switch opens the circuit, which indicates that the call has ended.*

The hook switch can also be used for signaling purposes. Some electronic analog phones have a button labeled "Link" that causes an event called a *flash*. You can perform a flash manually by depressing the hook switch for a duration of between 200 and 1,200 milliseconds. If you leave it down for longer than that, the carrier will assume you've hung up. The purpose of the Link button is to handle this timing for you. If you've ever used call waiting or three-way calling on an analog line, you have performed a hook switch flash for the purpose of signaling the network.

* When referring to the state of an analog circuit, people often speak in terms of "off-hook" and "on-hook." When your line is "off-hook," your telephone is "on" a call. If your phone is "on-hook," the telephone is essentially "off."

Handset. The handset is composed of the transmitter and receiver. It performs the conversion between the sound energy humans use and the electrical energy the telephone network uses.

Tip and Ring

In an analog telephone circuit, there are two wires. In North America, these wires are referred to as Tip and Ring.* This terminology comes from the days when telephone calls were connected by live operators sitting at cord boards. The plugs they used had two contacts, one located at the tip of the plug and the other connected to the ring around the middle (Figure 7-1).




Figure 7-1. Tip and Ring

The Tip lead is the positive polarity wire. In North America, this wire is typically green and provides the return path. The Ring wire is the negative polarity wire. In North America, this wire is normally red. When your telephone is on-hook, this wire will have a potential of -48V DC with respect to Tip. Off-hook, this voltage drops to roughly -7V DC .

Digital Telephony

Analog telephony is almost dead.

In the PSTN, the famous Last Mile is the final remaining piece of the telephone network still using technology pioneered well over a hundred years ago.†

One of the primary challenges when transmitting analog signals is that all sorts of things can interfere with those signals, causing low volume, static, and all manner of other undesired effects. Instead of trying to preserve an analog waveform over distances that may span thousands of miles, why not simply measure the characteristics

* They may have other names elsewhere in the world (such as “A” and “B”).

† “The Last Mile” is a term that was originally used to describe the only portion of the PSTN that had not been converted to fiber optics: the connection between the central office and the customer. The Last Mile is more than that, however, as it also has significance as a valuable asset of the traditional phone companies—they own a connection into your home. The Last Mile is becoming more and more difficult to describe in technical terms, as there are now so many ways to connect the network to the customer. As a thing of strategic value to telecom, cable, and other utilities, its importance is obvious.

of the original sound and send that information to the far end? The original waveform wouldn't get there, but all the information needed to reconstruct it would.

This is the principle of all digital audio (including telephony): sample the characteristics of the source waveform, store the measured information, and send that data to the far end. Then, at the far end, use the transmitted information to generate a completely new audio signal that has the same characteristics as the original. The reproduction is so good that the human ear can't tell the difference.

The principle advantage of digital audio is that the sampled data can be mathematically checked for errors all along the route to its destination, ensuring that a perfect duplicate of the original arrives at the far end. Distance no longer affects quality, and interference can be detected and eliminated.

Pulse-Code Modulation

There are several ways to digitally encode audio, but the most common method (and the one used in telephony systems) is known as *Pulse-Code Modulation* (PCM). To illustrate how this works, let's go through a few examples.

Digitally encoding an analog waveform

The principle of PCM is that the amplitude of the analog waveform is sampled at specific intervals so that it can later be recreated. The amount of detail that is captured is dependent both on the bit-resolution of each sample and on how frequently the samples are taken. A higher bit-resolution and a higher sampling rate will provide greater accuracy, but more bandwidth will be required to transmit this more detailed information.

To get a better idea of how PCM works, consider the waveform displayed in Figure 7-2.

To digitally encode the wave, it must be sampled on a regular basis, and the amplitude of the wave at each moment in time must be measured. The process of slicing up a waveform into moments in time and measuring the energy at each moment is called *quantization*, or *sampling*.

The samples will need to be taken frequently enough and will need to capture enough information to ensure that the far end can recreate a sufficiently similar waveform. To achieve a more accurate sample, more bits will be required. To explain this concept, we will start with a very low resolution, using four bits to represent our amplitude. This will make it easier to visualize both the quantization process itself and the effect that resolution has on quality.

Figure 7-3 shows the information that will be captured when we sample our sine wave at four-bit resolution.




Figure 7-2. A simple sinusoidal (sine) wave




Figure 7-3. Sampling our sine wave using four bits

At each time interval, we measure the amplitude of the wave and record the corresponding intensity—in other words, we sample it. You will notice that the four-bit resolution limits our accuracy. The first sample has to be rounded to 0011, and the next quantization yields a sample of 0101. Then comes 0100, followed by 1001, 1011, and so forth. In total, we have 14 samples (in reality, several thousand samples must be taken per second). If we string together all the values, we can send them to the other side as:

0011 0101 0100 1001 1011 1011 1010 0001 0101 0101 0000 1100 1100 1010

On the wire, this code might look something like Figure 7-4.




Figure 7-4. PCM encoded waveform

When the far end's digital-to-analog (D/A) converter receives this signal, it can use the information to plot the samples, as shown in Figure 7-5.




Figure 7-5. Plotted PCM signal

From this information, the waveform can be reconstructed (see Figure 7-6).




Figure 7-6. Delineated signal

As you can see if you compare Figure 7-7 with Figure 7-8, this reconstruction of the waveform is not very accurate. This was done intentionally, to demonstrate an important point: the quality of the digitally encoded waveform is affected by the resolution and rate at which it is sampled. At too low a sampling rate, and with too low a sample resolution, the audio quality will not be acceptable.

Increasing the sampling resolution and rate

Let's take another look at our original waveform, this time using five bits to define our quantization intervals (Figure 7-7).




Figure 7-7. The same waveform, on a higher-resolution overlay



In reality, there is no such thing as five-bit PCM. In the telephone network, PCM samples are encoded using eight bits.*

We'll also double our sampling frequency. The points plotted this time are shown in Figure 7-8.

We now have twice the number of samples, at twice the resolution. Here they are:

```
00111 01000 01001 01001 01000 00101 10110 11000 11001 11001 11000 10111 10100 10001  
00010 00111 01001 01010 01001 00111 00000 11000 11010 11010 11001 11000 10110 10001
```

When received at the other end, that information can now be plotted as shown in Figure 7-9.

* Other digital audio methods may employ 16 bits or more.




Figure 7-8. The same waveform at double the resolution




Figure 7-9. Five-bit plotted PCM signal

From this information, the waveform shown in Figure 7-10 can then be generated. As you can see, the resultant waveform is a far more accurate representation of the original. However, you can also see that there is still room for improvement.



Note that 40 bits were required to encode the waveform at 4-bit resolution, while 156 bits were needed to send the same waveform using 5-bit resolution (and also doubling the sampling rate). The point is, there is a tradeoff: the higher the quality of audio you wish to encode, the more bits will be required to do it, and the more bits you wish to send (in real time, naturally), the more bandwidth you will need to consume.




Figure 7-10. Waveform delineated from five-bit PCM

Nyquist's Theorem

So how much sampling is enough? That very same question was considered in the 1920s by an electrical engineer (and AT&T/Bell employee) named Harry Nyquist. Nyquist's Theorem* states: "When sampling a signal, the *sampling frequency* must be greater than twice the bandwidth of the input signal in order to be able to reconstruct the original perfectly from the sampled version."

In essence, what this means is that to accurately encode an analog signal you have to sample it twice as often as the total bandwidth you wish to reproduce. Since the telephone network will not carry frequencies below 300 Hz and above 4,000 Hz, a sampling frequency of 8,000 samples per second will be sufficient to reproduce any frequency within the bandwidth of an analog telephone. Keep that 8,000 samples per second in mind; we're going to talk about it more later.

Logarithmic companding

So, we've gone over the basics of quantization, and we've discussed the fact that more quantization intervals (i.e., a higher sampling rate) give better quality but also require more bandwidth. Lastly, we've discussed the minimum sample rate needed to accurately measure the range of frequencies we wish to be able to transmit (in the case of the telephone, it's 8,000 Hz). This is all starting to add up to a fair bit of data being sent on the wire, so we're going to want to talk about companding.

* Nyquist published two papers, "Certain Factors Affecting Telegraph Speed" (1924) and "Certain Topics in Telegraph Transmission Theory" (1928), in which he postulated what became known as Nyquist's Theorem. Proven in 1949 by Claude Shannon ("Communication in the Presence of Noise"), it is also referred to as the Nyquist-Shannon sampling theorem.

Companding is a method of improving the dynamic range of a sampling method without losing important accuracy. It works by quantizing higher amplitudes in a much coarser fashion than lower amplitudes. In other words, if you yell into your phone, you will not be sampled as cleanly as you will be when speaking normally. Yelling is also not good for your blood pressure, so it's best to avoid it.

Two companding methods are commonly employed: μ -law* in North America, and A-law in the rest of the world. They operate on the same principles but are otherwise not compatible with each other.

Companding divides the waveform into *cords*, each of which has several *steps*. Quantization involves matching the measured amplitude to an appropriate step within a cord. The value of the band and cord numbers (as well as the sign—positive or negative) becomes the signal. The following diagrams will give you a visual idea of what companding does. They are not based on any standard, but rather were made up for the purpose of illustration (again, in the telephone network companding will be done at an eight-bit, not five-bit, resolution).

Figure 7-11 illustrates five-bit companding. As you can see, amplitudes near the zero-crossing point will be sampled far more accurately than higher amplitudes (either positive or negative). However, since the human ear, the transmitter, and the receiver will also tend to distort loud signals, this isn't really a problem.

A quantized sample might look like Figure 7-12. It yields the following bit stream:

```
00000 10011 10100 10101 01101 00001 00011 11010 00010 00001 01000 10011 10100 10100
00101 00100 00101 10101 10011 10001 00011 00001 00000 10100 10010 10101 01101 10100
00101 11010 00100 00000 01000
```

Aliasing

If you've ever watched the wheels on a wagon turn backward in an old Western movie, you've seen the effects of aliasing. The frame rate of the movie cannot keep up with the rotational frequency of the spokes, and a false rotation is perceived.

In a digital audio system (which the modern PSTN arguably is), aliasing always occurs if frequencies that are greater than one-half the sampling rate are presented to the analog-to-digital (A/D) converter. In PSTN, that is any audio frequencies above 4,000 Hz (half the sampling rate of 8,000 Hz). This problem is easily corrected by passing the audio through a low-pass filter† before presenting it to the A/D converter.

* μ -law is often referred to as "u-law" because, let's face it, how many of us have μ keys on our keyboards? μ is in fact the Greek letter Mu; thus, you will also see μ -law written (more correctly) as "Mu-law." When spoken, it is correct to confidently say "Mew-law," but if folks look at you strange, and you're feeling generous, you can help them out and tell them it's "u-law." Many people just don't appreciate trivia.

† A low-pass filter, as its name implies, allows through only frequencies that are lower than its cut-off frequency. Other types of filters are high-pass filters (which remove low frequencies) and band-pass filters (which filter out both high and low frequencies).




Figure 7-11. Five-bit companding

The Digital Circuit-Switched Telephone Network

For over a hundred years, telephone networks were exclusively circuit-switched. What this meant was that for every telephone call made, a dedicated connection was established between the two endpoints, with a fixed amount of bandwidth allocated to that circuit. Creating such a network was costly, and where distance was concerned, using that network was costly as well. Although we are all predicting the end of the circuit-switched network, many people still use it every day, and it really does work rather well.

The humble DS-0, the foundation of it all

Since the standard method of digitizing a telephone call is to record an 8-bit sample 8,000 times per second, we can see that a PCM-encoded telephone circuit will need a bandwidth of 64,000 bps. This 64-kbps channel is referred to as a DS-0 (that's Dee-Ess-Zero). The DS-0 is the fundamental building block of all digital telecommunications circuits.

Even the ubiquitous analog circuit is sampled into a DS-0 as soon as possible. Sometimes this happens where your circuit terminates at the central office, and sometimes well before.

T-carrier circuits

The venerable T-1 is one of the more recognized digital telephony terms. A T-1 is a digital circuit consisting of 24 DS-0s multiplexed together into a 1.544-Mbps bit stream.* This bit stream is properly defined as a DS-1. Voice is encoded on a T-1 using the μ -law companding algorithm.



The European version of the T-1 was developed by the European Conference of Postal and Telecommunications Administrations† (CEPT), and was first referred to as a CEPT-1. It is now called an E-1.

The E-1 is comprised of 32 DS-0s, but the method of PCM encoding is different—E-1s use A-law companding. This means that connecting between an E-1-based network and a T-1-based network will always require a transcoding step. Note that an E-1, although it has 32 channels, is also considered a DS-1.

The various other T-carriers (T-2, T-3, and T-4) are multiples of the T-1, each based on the humble DS-0. Table 7-2 illustrates the relationships between the different T-carrier circuits.

Table 7-2. T-carrier circuits

Carrier	Equivalent data bitrate	Number of DS-0s	Data bitrate
T-1	24 DS-0s	24	1.544 Mbps
T-2	4 T-1s	96	6.312 Mbps
T-3	7 T-2s	672	44.736 Mbps
T-4	6 T-3s	4032	274.176 Mbps

At densities above T-3, it is very uncommon to see a T-carrier circuit. For these speeds, optical carrier (OC) circuits may be used.

* The 24 DS-0s use 1.536 Mbps, and the remaining .008 Mbps is used by framing bits.

† Conférence Européenne des Administrations des Postes et des Télécommunications.

SONET and OC circuits

The Synchronous Optical Network (SONET) was developed out of a desire to take the T-carrier system to the next technological level: fiber optics. SONET is based on the bandwidth of a T-3 (44.736Mbps), with a slight overhead making it 51.84 Mbps. This is referred to as an OC-1 or STS-1. As Table 7-3 shows, all higher-speed OC circuits are multiples of this base rate.

Table 7-3. OC circuits

Carrier	Equivalent data bitrate	Number of DS-0s	Data bitrate
OC-1	1 DS-3 (plus overhead)	672	51.840 Mbps
OC-3	3 DS-3s	2016	155.520 Mbps
OC-12	12 DS-3s	8064	622.080 Mbps
OC-48	48 DS-3s	32256	2488.320 Mbps
OC-192	192 DS-3s	129024	9953.280 Mbps

SONET was created in an effort to standardize optical circuits, but due to its high cost, coupled with the value offered by many newer schemes, such as Dense Wave Division Multiplexing (DWDM), there is some controversy surrounding its future.

Digital Signaling Protocols

As with any circuit, it is not enough for the circuits used in the PSTN to just carry (voice) data between endpoints. Mechanisms must also be provided to pass information about the state of the channel between each endpoint. (Disconnect and answer supervision are two examples of basic signaling that might need to take place; Caller ID is an example of a more complex form of signaling.)

Channel Associated Signaling (CAS)

Also known as robbed-bit signaling, CAS is what you will use to transmit voice on a T-1 when ISDN is not available. Rather than taking advantage of the power of the digital circuit, CAS simulates analog channels. CAS signaling works by stealing bits from the audio stream for signaling purposes. Although the effect on audio quality is not really noticeable, the lack of a powerful signaling channel limits your flexibility.

When configuring a CAS T-1, the signaling options at each end must match. E&M (Ear & Mouth or recEive & transMit) signaling is generally preferred, as it offers the best supervision.

CAS is very rarely used on PSTN circuits anymore, due to the superiority of ISDN-PRI. One of the limitations of CAS is that it does not allow the dynamic assignment of channels to different functions. Also, Caller ID information (which may not even be supported) has to be sent as part of the audio stream. CAS is commonly used on the T-1 link in channel banks, although PRI is sometimes available (and preferable).

ISDN

The Integrated Services Digital Network (ISDN) has been around for over 20 years. Because it separates the channels that carry the traffic (the bearer channels, or B-channels) from the channel that carries the signaling information (the D-channel), ISDN allows for the delivery of a much richer set of features than CAS. In the beginning, ISDN promised to deliver much the same sort of functionality that the Internet has given us, including advanced capabilities for voice, video, and data transfer.

Unfortunately, rather than ratifying a standard and sticking to it, the respective telecommunications manufacturers all decided to add their own tweaks to the protocol, in the belief that their versions were superior and would eventually dominate the market. As a result, getting two ISDN-compliant systems to connect to each other was often a painful and expensive task. The carriers who had to implement and support this expensive technology in turn priced it so that it was not rapidly adopted. Currently, ISDN is rarely used for much more than basic trunking—in fact, the acronym ISDN has become a joke in the industry: “It Still Does Nothing.”

Having said that, ISDN has become quite popular for trunking, and it is now (mostly) standards-compliant. If you have a PBX with more than a dozen lines connected to the PSTN, there’s a very good chance that you’ll be running an ISDN-PRI circuit. Also, in places where DSL and cable access to the Internet are not available (or too expensive), an ISDN-BRI circuit might provide you with an affordable 128-kbps connection. In much of North America, the use of ISDN-BRI for Internet connectivity has been deprecated in favor of DSL and cable modems, but it’s still very popular in other parts of the world.

ISDN-BRI/BRA. Basic Rate Interface (or Basic Rate Access) is the flavor of ISDN designed to service small endpoints such as workstations.

The BRI flavor of the ISDN specification is often referred to simply as “ISDN,” but this can be a source of confusion, as ISDN is a protocol, not a type of circuit (not to mention that PRI circuits are also correctly referred to as ISDN!).

A Basic Rate ISDN circuit consists of two 64-kbps B-channels controlled by a 16-kbps D-channel, for a total of 144 kbps.

Basic Rate ISDN has been a source of much confusion during its life, due to problems with standards compliance, technical complexity, and poor documentation. Still, in European countries ISDN-BRI circuits remain quite a popular way of connecting to the PSTN.

ISDN-PRI/PRA. The Primary Rate Interface (or Primary Rate Access) flavor of ISDN is used to provide ISDN service over larger network connections. A Primary Rate ISDN circuit uses a single DS-0 channel as a signaling link (the D-channel); the remaining channels serve as B-channels.

In North America, Primary Rate ISDN is commonly carried on one or more T-1 circuits. Since a T-1 has 24 channels, a North American PRI circuit typically consists of 23 B-channels and 1 D-channel. For this reason, PRI is often referred to as 23B+D.*



In Europe, a 32-channel E-1 circuit is used, so a Primary Rate ISDN circuit is referred to as 30B+D (the final channel is used for synchronization).

Primary Rate ISDN is very popular, due to its technical benefits and generally competitive pricing. If you believe you will require more than a dozen or so PSTN lines, you should look into Primary Rate ISDN pricing.

From a technical perspective, ISDN-PRI is always preferable to CAS.

Signaling System 7

SS7 is the signaling system used by carriers. It is conceptually similar to ISDN, and it is instrumental in providing a mechanism for the carriers to transmit the additional information ISDN endpoints typically need to pass. However, the technology of SS7 is different from that of ISDN—one big difference is that SS7 runs on a completely separate network from the actual trunks that carry the calls.


SS7 support in Asterisk is on the horizon, as there is much interest in making Asterisk compatible with the carrier networks. An open source version of SS7 (<http://www.openss7.org>) exists, but work is still needed for full SS7 compliance, and as of this writing it is not known whether this will be integrated with Asterisk. Another promising source of SS7 support comes from Sangoma Technologies, who offer SS7 functionality in many of their products.

It should be noted that adding support for SS7 in Asterisk is not going to be as simple as writing a proper driver. Connecting equipment to an SS7 network will not be possible without that equipment having passed an extremely rigorous certification processes. Even then, it seems doubtful that any traditional carrier is going to be in a hurry to allow such a thing to happen.

Packet-Switched Networks

In the mid-1990s, network performance improved to the point where it became possible to send a stream of media information in real time across a network connection. Because the media stream is chopped up into segments, which are then

* PRI is actually quite a bit more flexible than that, as it is possible to span a single PRI circuit across multiple T-1 spans. This can give rise, for example, to a 47B+D circuit (where a single D-channel serves two T-1s) or a 46B+2D circuit (where primary and backup D-channels serve a pair of T-1s). You will sometimes see PRI described as nB+nD, because the number of B- and D-channels is, in fact, quite variable.



wrapped in an addressing envelope, such connections are referred to as *packet-based*. The challenge, of course, is to send thousands of these packets between two endpoints, ensuring that the packets arrive in the same order in which they were sent, in less than 300 milliseconds, with none lost. This is the essence of Voice over IP.

Conclusion

This chapter has explored the technologies currently in use in the PSTN. In the next chapter, we will discuss protocols for VoIP: the carrying of telephone connections across IP-based networks. These protocols define different mechanisms for carrying telephone conversations, but their significance is far greater than just that. Bringing the telephone network into the data network will finally erase the line between telephones and computers, which holds the promise of a revolutionary evolution in the way we communicate.

