SECTION 9

DSP APPLICATIONS

- High Performance Modems for Plain Old Telephone Service (POTS)
- Remote Access Server (RAS) Modems
- ADSL (Assymetric Digital Subscriber Line)
- Digital Cellular Telephones
- GSM Handset Using SoftFone™ Baseband Processor and Othello™ Radio
- Analog Cellular Basestations
- Digital Cellular Basestations
- Motor Control
- Codecs and DSPs in Voiceband and Audio Applications
- A Sigma-Delta ADC with Programmable Digital Filter

9.b

SECTION 9
DSP APPLICATIONS
Walt Kester

HIGH PERFORMANCE MODEMS FOR PLAIN OLD TELEPHONE SERVICE (POTS)

Modems (*Mo*dulator/*Dem*odulator) are widely used to transmit and receive digital data using analog modulation over the Plain Old Telephone Service (POTS) network as well as private lines. Although the data to be transmitted is digital, the telephone channel is designed to carry voice signals having a bandwidth of approximately 300 to 3300Hz. The telephone transmission channel suffers from delay distortion, noise, crosstalk, impedance mismatches, near-end and far-end echoes, and other imperfections. While certain levels of these signal degradations are perfectly acceptable for voice communication, they can cause high error rates in digital data transmission. The fundamental purpose of the transmitter portion of the modem is to prepare the digital data for transmission over the analog voice line. The purpose of the receiver portion of the modem is to receive the signal which contains the analog representation of the data , and reconstruct the original digital data at an acceptable error rate. High performance modems make use of digital techniques to perform such functions as modulation, demodulation, error detection and correction, equalization, and echo cancellation.

A block diagram of an ordinary telephone channel (often referred to as plain old telephone service – or POTS) is shown in Figure 9.1. Most voiceband telephone connections involve several connections through the telephone network. The 2-wire twisted pair subscriber line available at most sites is generally converted to a 4-wire signal at the telephone central office: two wires for transmit, and two wires for receive. The signal is converted back to a 2-wire signal at the far-end subscriber line. The 2- to 4-wire interface is implemented with a circuit called a *hybrid*. The hybrid intentionally inserts impedance mismatches to prevent oscillations on the 4-wire trunk line. The mismatch forces a portion of the transmitted signal to be reflected or echoed back to the transmitter. This echo can corrupt data the transmitter receives from the far-end modem.

Half-duplex modems are capable of passing signals in either direction on a 2-wire line, but not simultaneously. *Full-duplex* modems operate on a 2-wire line and can transmit and receive data simultaneously. Full-duplex operation requires the ability to separate a receive signal from the reflection (echo) of the transmitted signal. This is accomplished by assigning the signals in the two directions different frequency bands separated by filtering, or by echo cancelling in which a locally synthesized replica of the reflected transmitted signal is subtracted from the composite receive signal.

ANALOG MODEM USING PLAIN OLD TELEPHONE SERVICE (POTS) ANALOG CHANNEL

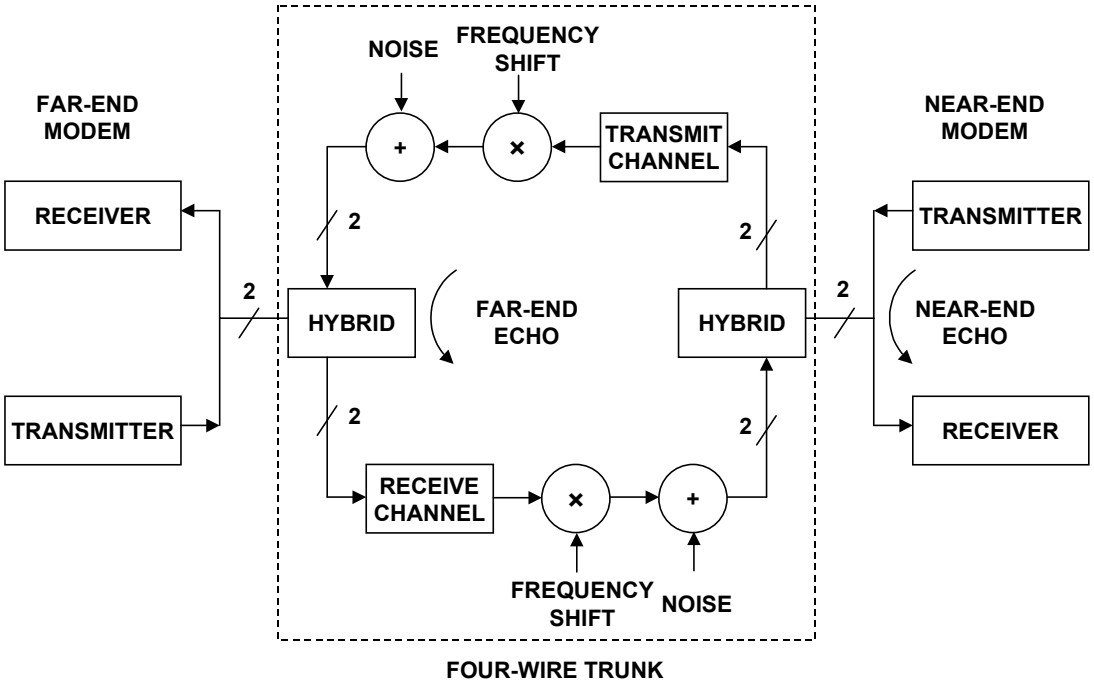

Figure 9.1

There are two types of echo in a typical voiceband telephone connection. The first echo is the reflection from the near-end hybrid, and the second echo is from the far-end hybrid. In long distance telephone transmissions, the transmitted signal is hetrodyned to and from a carrier frequency. Since local oscillators in the network are not exactly matched, the carrier frequency of the far-end echo may be offset from the frequency of the transmitted carrier signal. In modern applications this shift can affect the degree to which the echo signal can be canceled. It is therefore desirable for the echo canceller to compensate for this frequency offset.

For transmission over the telephone voice network, the digital signal is modulated onto an audio sinewave carrier, producing a modulated tone signal. The frequency of the carrier is chosen to be well within the telephone band. The transmitting modem modulates the audio carrier with the transmit data signal, and the receiving modem demodulates the tone to recover the receive data signal.

The baseband data signal may be used to modulate the amplitude, the frequency, or the phase of the audio carrier, depending on the data rate required. These three types of modulation are known as amplitude shift keying (ASK), frequency shift keying (FSK), and phase shift keying (PSK). In its simplest form the modulated carrier takes on one of two states - that is, one of two amplitudes, one of two frequencies, or one of two phases. The two states represent a logic 0 or a logic 1.

Low- to medium-speed data links usually use FSK up to 1,200 bits/s. Multiphase PSK are used for 2,400 bits/s and 4,800 bits/s links. PSK utilizes bandwidth more efficiently than FSK but is more costly to implement. ASK is least efficient and is

used only for very low speed links (less than 100 bits/s). For 9,600 bits/s up to 33,600 bits/s a combination of PSK and ASK is used, known as Quadrature Amplitude Modulation (QAM).

The International Telegraph and Telephone Consultative Committee (CCITT in France) has established standards and recommendations for modems which are given in Figure 9.2.

SOME MODEM STANDARDS

CCITT Rec.	Approximate Date	Speed (bits/s) maximum	Half Duplex/ Full Duplex/ Echo Cancel	Modulation Method
V.21	1964	300	FDX	FSK
V.22		1200	FDX	PSK
V.22bis		2400	FDX	16QAM
V.23		1200	HDX	FSK
V.26 bis		2400	HDX	PSK
V.26 ter		2400	FDX(EC)	PSK
V.27 ter		4800	HDX	8PSK
V.32		9600	FDX(EC)	32QAM
V.32 bis		14400	FDX(EC)	QAM
V.34		33600	FDX(EC)	QAM
V.90	1998	56000*	FDX(EC)	PCM
V.92	2001	56000**	FDX(EC)	PCM

* DOWNSTREAM ONLY, UPSTREAM IS V.34
** UPSTREAM AND DOWNSTREAM

Figure 9.2

The goal in designing high performance modems is to achieve the highest data transfer rate possible over the POTS network and avoid the expense of using dedicated conditioned private telephone lines. The V.90 recommendation describes a full-duplex (simultaneous transmission and reception) modem that operates on the POTS network. The V.90 modem communicates downstream from the central office to the subscriber modem at a rate of 56,000 bits/s using Pulse Code Modulation (PCM). Upstream communication from the subscriber to the central office is at the V.34 rate of up to 33,600 bits/s (QAM).

A simplified block diagram for a V.90 analog modem is shown in Figure 9.3. The diagram shows that the bulk of the signal processing is done digitally. Both the transmit and receive portions of the modem subject the digital signals to a number of DSP algorithms which can be efficiently run on modern processors.

V.90 ANALOG MODEM SIMPLIFIED BLOCK DIAGRAM

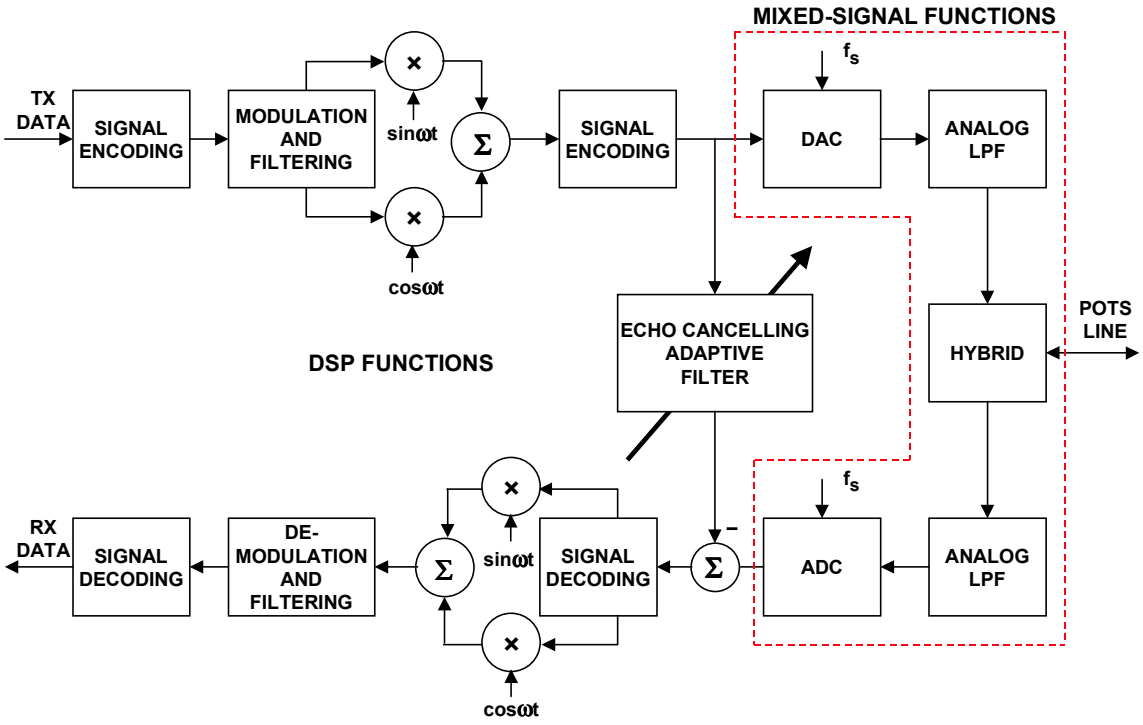

Figure 9.3

The TX input serial bit stream is first scrambled and encoded. Scrambling takes the input bit stream and produces a pseudo-random sequence. The purpose of the scrambler is to whiten the spectrum of the transmitted data. Without the scrambler, a long series of identical symbols could cause the receiver to lose carrier lock. Scrambling makes the transmitted spectrum resemble white noise, to utilize the bandwidth of the channel more efficiently, makes carrier recovery and timing synchronization easy, and makes adaptive equalization and echo cancellation possible.

The scrambled bit stream is divided into groups of bits, and the groups of bits are first differentially encoded and then convolutionally encoded.

The symbols are then mapped into the signal space using QAM as defined in the V.34 standard. The signal space mapping produces two coordinates, one for the real part of the QAM modulator and one for the imaginary part. As an example, a diagram of a 16-QAM signal constellation is shown in Figure 9.4. Larger constellations are used in V.90 modems, and the actual size of the constellation is adaptive and determined during the training, or "handshake" interval when the modems synchronize with each other for upstream or downstream signaling.

QUADRATURE AMPLITUDE MODULATED (QAM) SIGNAL TRANSMITS 4 BITS PER SYMBOL (16-QAM)

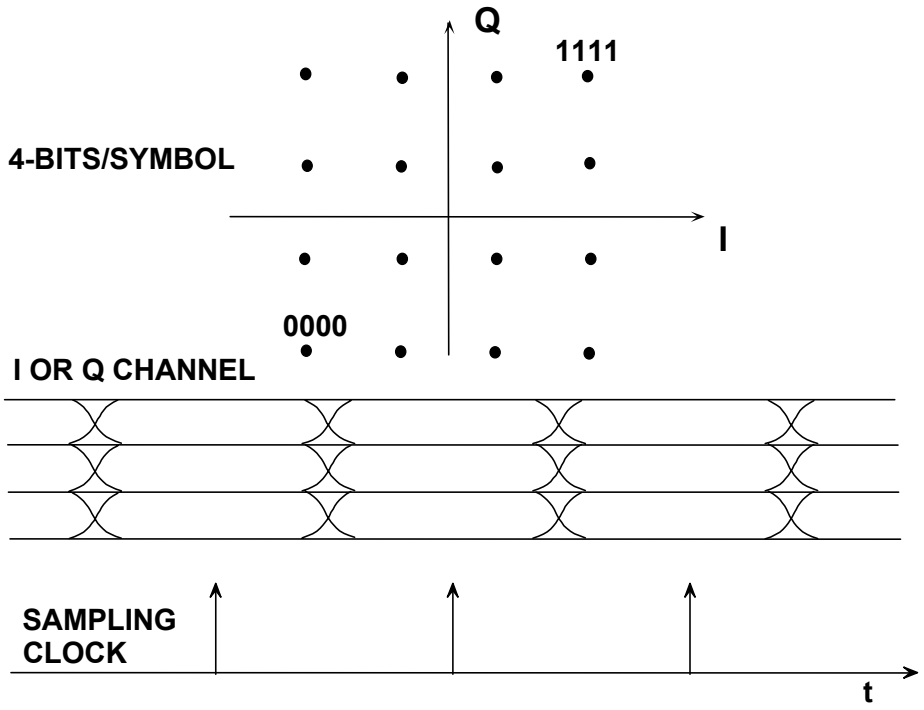

Figure 9.4

Used prior to modulation, digital pulse shaping filters attenuate frequencies above the Nyquist frequency that are generated in the signal mapping process. These filters are designed to have zero crossings at the appropriate frequencies to cancel intersymbol interference.

QAM is easily implemented in modern DSP processors. The process of modulation requires the access of a sine or cosine value, the access of an input symbol (x or y coordinate) and a multiplication. The parallel architecture of the ADSP-21xx-family permits all three operations to be performed in a single instruction cycle.

The output of the digital modulator drives a DAC. The output of the DAC is passed through an analog lowpass filter and to the 2-wire telephone line for transmission over the POTS network.

The receiver is made up of several functional blocks: the input antialiasing filter and ADC, a demodulator, an adaptive equalizer, a Viterbi decoder, an echo canceller, a differential decoder, and a descrambler. The receiver DSP algorithms are both memory-intensive and computation-intensive. The ADSP-218x-family addresses both needs, providing sufficient program memory RAM (for both code and data) on chip, data memory RAM on chip, and an instruction execution rate of up to 75MIPS.

The antialiasing filter and ADC in the receiver need to have a dynamic range from the largest echo signal to the smallest. The received signal can be as low as -40dBm,

while the near-end echo can be as high as -6dBm. In order to ensure that the analog front end of the receiver does not contribute any significant impairment to the channel under these conditions, an instantaneous dynamic range of 84dB and an SNR of 72dB is required.

In order to compensate for amplitude and phase distortion in the telephone channel, equalization is required to recover the transmitted data at an acceptably low bit error rate. In order to respond to rapidly changing conditions on the telephone line, adaptive equalization is required for the V.90 modem receiver. An adaptive equalizer can be implemented digitally in an FIR filter whose coefficients are continuously updated based on current line conditions.

Separation between the transmit and receive signal in the V.90 modem is accomplished using echo cancellation. Both near-end and far-end echo must be cancelled in order to yield reliable communication. Echo cancellation is achieved by subtracting an estimate of the echo return signal from the actual received signal. The predicted echo is determined by feeding the transmitted signal into an adaptive filter with a transfer function that approximates the telephone channel. The adaptive filter commonly used in echo cancellers is the FIR filter (chosen for its stability and linear phase response). The taps are determined using the least-mean-square (LMS) algorithm during a training sequence executed prior to full-duplex communications.

The most common technique for decoding the received data is Viterbi decoding. Named after its inventor, the Viterbi algorithm is a general-purpose technique for making an error-corrected decision. Viterbi decoding provides a certain degree of error correction by examining the received bit pattern over time to deduce the value that was the most likely to have been transmitted at a particular time. Viterbi decoding is computation-intensive. A history for each of the possible symbols sent at each symbol interval has to be maintained. At each symbol interval, the length of the path backward in time from each possible received symbol to a symbol sent some time ago is calculated. The symbol that has the shortest path back to the original signal is chosen to be the current decoded symbol. A complete description of Viterbi decoding and its implementation on the ADSP-21xx- family of processors is given in documentation available from Analog Devices (Reference 2).

Figure 9.5 shows a comparison between V.34 and V.90 modems. Note that in the case of V.34 (Figure 9.5A), the communication is between two *analog* modems. This requires and ADC/DAC in both the transmit and receive path as shown in the diagram. The V.90 system requires an all digital network and a V.90 digital modem as shown in Figure 9.5B. Note that the second ADC/DAC combination is eliminated, thereby allowing the faster downstream data rate of 56Kbits/s. Downstream communication to the V.90 analog modem uses 64Kbits/s PCM data, which is standard for all digital telephone networks. This serial data is converted into a pulse amplitude modulated (PAM) signal (8-bits, 8kSPS) using an 8-bit DAC. The signal from the DAC to the analog modem is therefore a 256K constellation with no imaginary component, i.e., the analog modem receiver must detect which one of the 256 levels is being sent during the symbol interval.

The V.90 standard allows downstream data rates of up to slightly less than 56Kbits/s, and upstream data rates of up to 33.6Kbits/s (V.34). The V.92 standard will allow 56Kbits/s transmission in both directions.

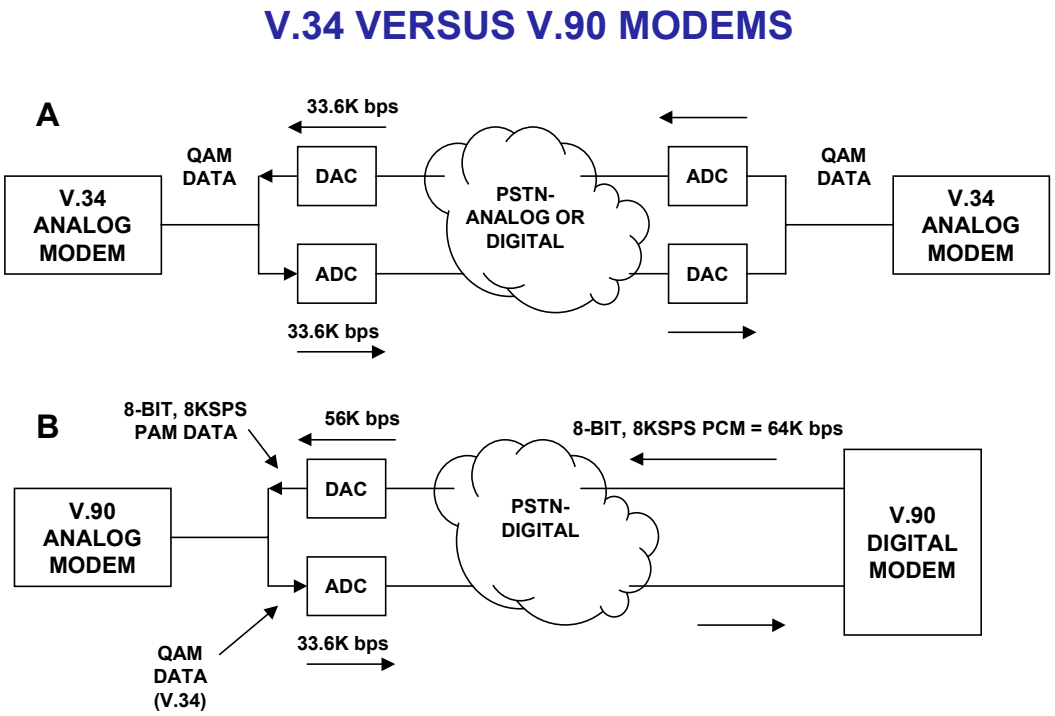

Figure 9.5

REMOTE ACCESS SERVER (RAS) MODEMS

Rapid growth and use of the Internet has created a problem in that there are more users trying to get on the Internet than there is equipment to accommodate all these users. Internet Service Providers (ISPs) like America On Line purchase modem equipment so their customers (referred to as subscribers) can remotely access a network (like the Internet from home). This application of accessing a network from a remote location is called Remote Network Access. The equipment used in this application is called a Remote Access Server (RAS) as shown in Figure 9.6. The Remote Access Server is made up of many modem ports; each modem port can connect to a different user. The RAS can use analog modems, which connects to a POTS line, or digital modems which are compatible with T1, E1, PRI, or BRI lines. Digital modems are used in most RAS systems since they are more efficient for 8 ports or more.

Network access equipment enables individuals, small offices, and traveling employees to connect to corporate networks (Intranets) and the Internet. Internet Service Providers use devices called *concentrators* to connect their telephone access lines to their networks. These concentrators are also referred to as Remote Access Servers. The rapid growth in the use of the Internet and Intranets has created a tremendous demand for modem equipment.

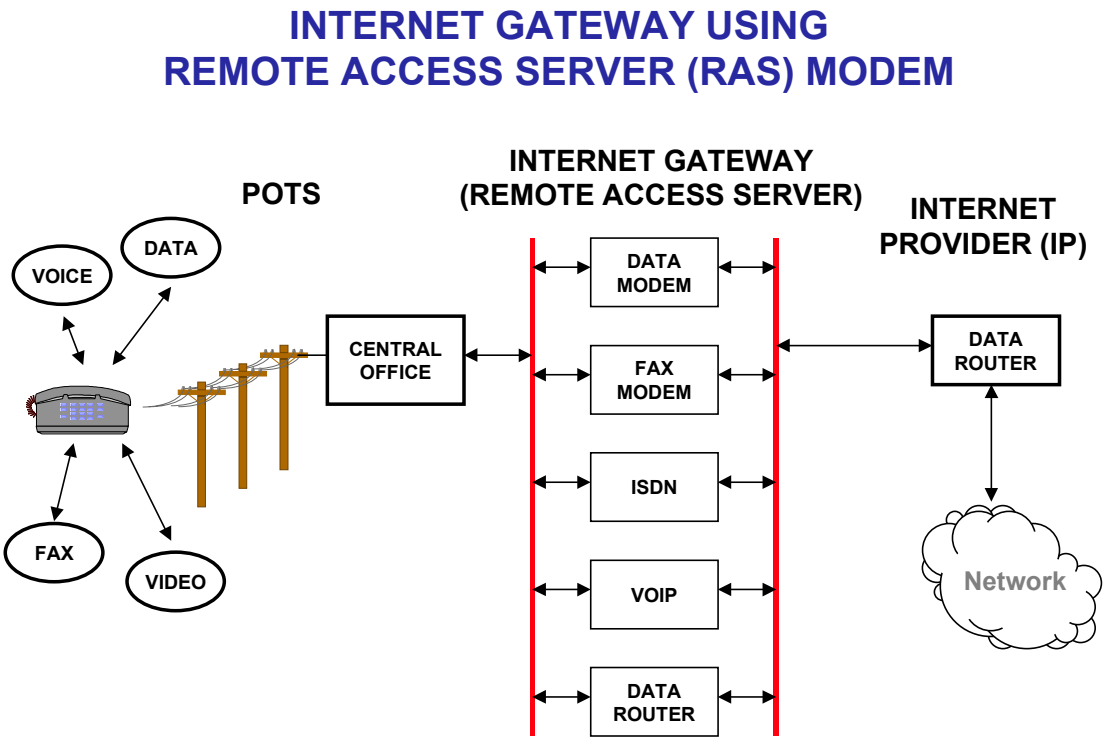

Figure 9.6

Remote access servers accommodate employees and small offices/home offices (SOHO) wishing to connect individual computers to LANs (Local Area Networks) or Intranets. If a remote access server is installed in a corporate LAN, remote users can access the network in such a way that their computers appear to be directly connected to the LAN. This allows them to work from a remote location as if they were sitting at their desk in their office.

The ADSP-21mod870 acts as a bridge between the voice-based continuous connection switched network and the data based IP network as shown in Figure 9.7. The high speed DMA interface and large on-chip RAM of the ADSP-21mod870 allows it to be configured to handle a large variety of tasks. The software with the ADSP-21mod870-100 can be configured for modem calls or HDLC (high bit-rate digital subscriber line) processing of digital ISDN (integrated services digital network) based calls. Since the ADSP-21mod870 is an open platform, other functions can be loaded by users. Examples of other functions include voice over internet and FAX over internet. In these applications, the ADSP-21mod870 is a gateway for voice network users to save toll charges by routing their call over the IP network. The ADSP-21mod870 DSP uses the ADSP-218x 16-bit fixed-point core and is code compatible with other members of the ADSP-21xx family.

As the number of remote network users has grown, capacity over the telephone central office switched network has often become strained. These bottlenecks often occur when thousands of calls from a metropolitan area are switched to a single point of presence (POP). To eliminate these bottle necks, RAS equipment can be pushed outward from the POP toward the edges of the switched network as shown in Figure 9.8. When RAS equipment is located in the switching center for local exchanges, data calls can be separated from voice calls, eliminating the strain on

the voice based switched network. RAS equipment integrated into switching equipment is referred to *on-switch* based RAS. This differs for standalone RAS systems in that on-switch RAS can separate data calls before they are connected to trunk lines exiting the switch.

DSP-BASED RAS MODEM USING THE ADSP-21modXXX - FAMILY

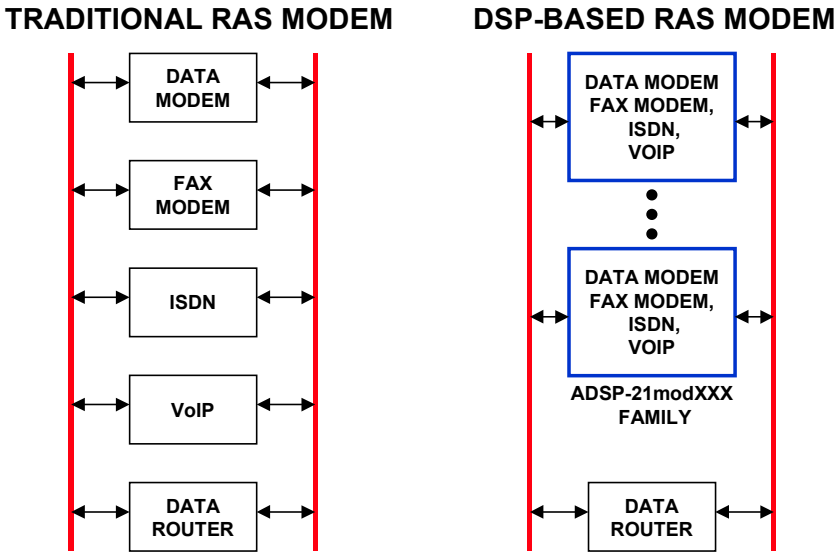

Figure 9.7

EXPANDING CENTRAL OFFICE CAPABILITY WITH ADSP-21modXXX FAMILY

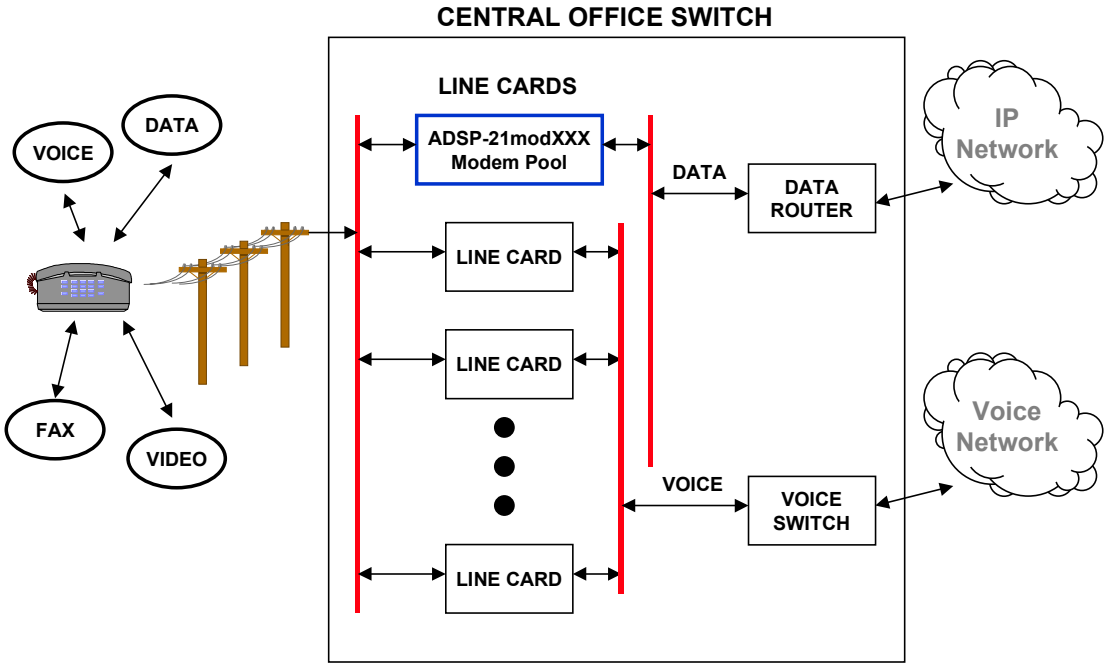

Figure 9.8

Several types of RAS equipment have evolved to satisfy the needs of the different remote network access users. RAS equipment can take several forms. RAS concentrators combine a modem pool with a router in a standalone chassis. NT server RAS uses off the shelf Windows NT PCs to perform routing functions with a modem pool connected to its PCI or ISA expansion bus. On-switch RAS integrates modem pools into the line card chassis in switching systems. These various forms of RAS equipment serve the different needs of end users. Local Exchange Companies (LECs) or the new Competing Local Exchange Companies (CLECs) can take advantage of On-switch RAS equipment to eliminate bottle necks in their switched networks. ISPs (Internet Service Providers) use RAS concentrators to collect calls from a large area for connection to the Internet. Large business users also use RAS concentrators to connect their employees to their LAN or Intranets. Small office and home office (SOHO) users can use low-cost NT Server based RAS to support all of their RAS, LAN, and other telecommunication needs.

The ADSP-21mod870 digital modem processor is the first complete digital RAS modem on a single chip. The device is V.34/56K and V.42/V.42bis compatible, has a 16-bit DMA port for software downloads, provides TDM serial port interfaces directly to T1/E1, has 160Kbytes on-chip SRAM, dissipates 140mW at +3.3V, and is packaged in a 16mm TQFP package. The ADSP-21mod870's small footprint and power efficiency will enable Internet Service Providers (ISPs) and central office access providers to quadruple the number of ports within the existing modem bank chassis. Moreover, the chip's unique capability to support any protocol on any port will improve ISP customer service and reduce business operating costs.

Part of a family of digital modem solutions from ADI, the ADSP-21mod870 includes silicon, software and service. ADI is singular in its ability to provide the complete solution. Also available are the ADSP-21mod970 (six modem channels, 31mm BGA) and the ADSP-21mod980 (eight modem channels, 35mm BGA).

MULTI-CHANNEL VOICE-OVER-INTERNET-PROVIDER (VOIP) SERVER

The ADSP-218x family of devices can efficiently implement multi-channel telephony applications (such as RAS/VOIP servers and gateways) due to their high performance and large on-chip memories. A typical system is shown in Figure 9.9. The programmable nature of the DSP architecture gives the system architect the flexibility to implement various speech coding algorithms in addition to baseline telephony functionality.

The ADSP-2188M is the most highly integrated member of the 218x family (with over 2MBit of on-chip SRAM). This high level of integration combined with 75MIPS performance can support up to 6 voice channels per DSP (depending on the voice coder chosen).

ADSP-218x IMPLEMENTATION OF A MULTI-CHANNEL VOIP SERVER

Multi-Channel VoIP Server	MIPS	Program Memory	Data Memory
G.723 annex A	18	8192	12288
G.728	28	7200	2000
G.729B	11.4	10240	5120
G.711	.4	512	512
DTMF Gen & Detect	1.9	300	400
Call Discrimination	1.3	1536	300
G.165	8	800	1024
OS Kernal	.2	512	
Total	N/A	29292	21644

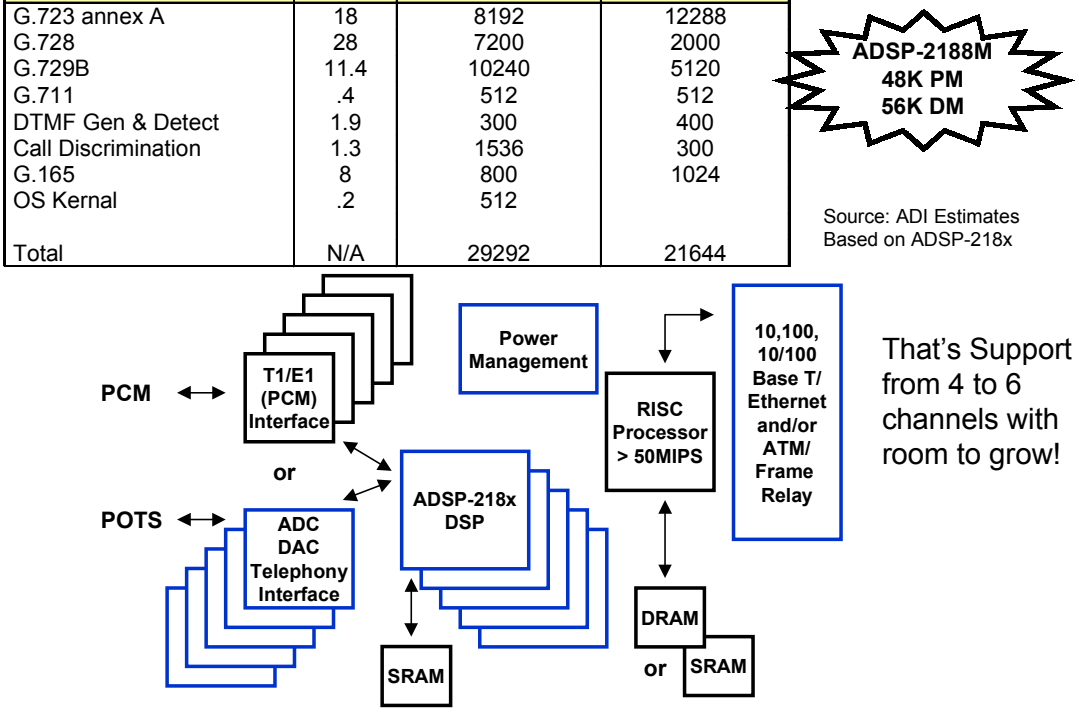

Figure 9.9

ADSL (ASYMMETRIC DIGITAL SUBSCRIBER LINE)

Thanks to the widespread popularity of the World Wide Web, Internet traffic is at an all-time high. A study recently conducted by the Wall Street Journal reported 58 million Internet users in the United States and Canada alone. Research firms all predict heavier Internet traffic as more and more people buy PCs and use the Internet for business, academia and recreational purposes.

Unless something is done to improve the way we access the Internet, user traffic will ultimately burden the public switched telephone network (PSTN) beyond its original design limits. Internet users are frustrated by the amount of time it takes to view simple text-based Web sites, especially from 8:00 am to 6:00 pm on business days when congestion and bottlenecks are most prevalent. The problem will only get worse as users attempt to view graphically complex sites, download new video and audio clips, and access other types of multimedia services becoming available over the Internet.

Today's analog modem and telephone switch technology is simply inadequate. Assuming little or no network delays, a 10 Megabyte data transmission - the equivalent of a four-minute audio/video clip - takes about 95 minutes to download when using a 14.4 Kbps analog modem, 45 minutes when using a 28.8 Kbps modem, and 25 minutes when using a 56Kbps modem. Lengthy on-line calls are tying up telephone systems originally designed to handle short (three-minute) voice calls and

switches built to accommodate approximately nine minutes per line during peak hours. How often do you browse the Internet for 10 minutes or less? It now appears that relief, in the form of ADSL (asymmetric digital subscriber line) technology, is on the way.

ADSL is a new high-speed digital switching/routing and signal processing technology. It promises to relieve network bottlenecks and provide enough user bandwidth for the Internet explosion. Originally conceived in 1994, ADSL delivers the huge amounts of bandwidth needed for interactive gaming, multimedia services and video-on-demand. These applications, along with videoconferencing, remote schooling and home shopping are still attractive applications. As more and more people use the Internet for electronic commerce in homes around the world, the need for high-speed network access becomes ever more paramount.

ADSL can transfer data over ordinary telephone lines nearly 200 times faster than today's contemporary modems and 90 times faster than ISDN. Early tests and trials have shown promising results all over the world. And while GTE and other large telephone companies begin to deploy ADSL systems in select regions of the US and abroad, others seek to deploy ADSL-based equipment in 1999 as standards-based systems and modems become commercially available.

Due to ADSL's technical complexity, only a few semiconductor manufacturers are presently developing ADSL silicon chips. Analog Devices is one of them, and is considered a pioneer in ADSL after having produced the first complete ADSL chipset back in 1997. Early adopters of ADSL rallied behind the AD20msp910 chipset after having successfully evaluated its high-speed, long reach capability. Soon thereafter, the world's most influential industry standards committees (ANSI, ETSI and ITU) approved an advanced discrete multi-tone (DMT) signal processing technique used by the AD20msp910. Today, Analog Devices enjoys the advantages of having the industry's first standards-based solution, the largest installed customer base, and the most design wins of any semiconductor manufacturer to date.

ADSL is attractive for the following reasons:

- ADSL is fast. The same 10 MB video clip that takes 90 minutes to download using a conventional modem would be downloaded in just 10 seconds using an ADSL modem. Superfast ADSL modems can transmit data at rates as high as 8 megabits per second.
- ADSL is easy to install. It uses existing copper twisted-pair telephone lines from a local exchange carrier (central office) to the subscriber's home or office. Little or no rewiring is necessary.
- ADSL is cost-effective. It requires no major upgrades in the existing telephone network infrastructure.
- ADSL is viable. Issues that have slowed the deployment of high-speed fiber networks to the home (e.g., prohibitive cost and installation) do not apply. ADSL works with existing POTS (plain-old telephone service). High-speed data transport can occur simultaneously with voice calls and fax transmissions.

Unlike other high-speed data transmission technologies, ADSL requires no rewiring over the "last mile" of the network. Although commonly referred to as the last mile, transmission length is typically 12,000 to 18,000 feet. That "last mile" (the leg from the central office to a user's home or office) operates over existing copper loops of twisted-pair telephone wires. But ADSL does require the installation of new equipment at local exchange central offices and major switch offices. However, the technology developed for central office equipment is the same as that used in PC modems and home splitter boxes, thus ensuring interoperability throughout the entire network. A simplified block diagram of an ADSL system is shown in Figure 9.10.

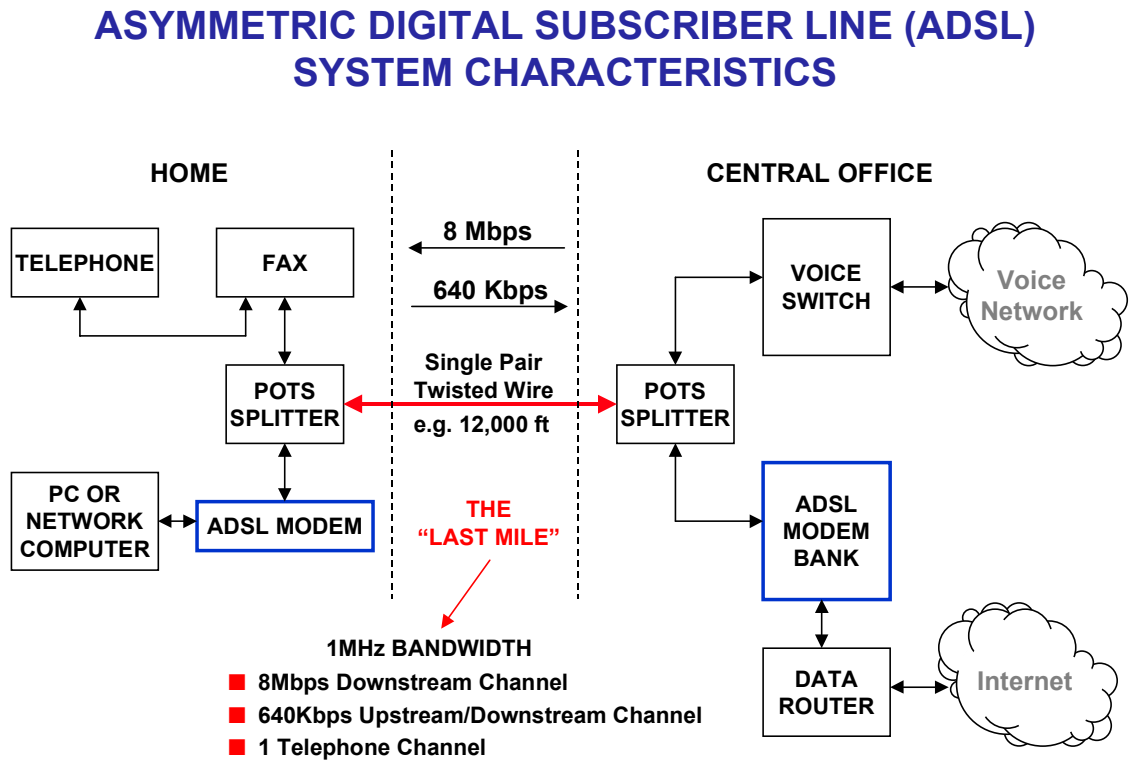

Figure 9.10

Using sophisticated digital signal processing techniques, ADSL modems push as much data as possible across copper wires by making the best use of available bandwidth. Many perceived the telephone network is as having only 4 kHz of bandwidth, but 4 kHz only represents the band used for the analog (voice) transmission. Using ADSL, the physical connection between the home and local exchange carrier (LEC) over conventional copper wire is capable of carrying 1 MHz. ADSL takes advantage of the portion of the bandwidth not used for voice calls. Essentially, it splits the 1 MHz bandwidth into three information channels: one high-speed downstream channel, one medium-speed duplex (upstream/downstream) channel, and one conventional voice channel. (Downstream refers to data transmitted from the telephone network to the customer's premises; upstream is data routed from the customer to the network.)

In addition to providing high bandwidth data transmission, ADSL preserves "lifeline" services (i.e., the voice network that people depend upon for day-to-day communications and emergency situations). This three-channel approach allows

subscribers to send an email, download a video for viewing, and talk on the telephone-all at the same time. Telecommuters can access their corporate local area network, and simultaneously videoconference with a customer. In fact, using ADSL's full capability provides enough bandwidth for running four channels of MPEG compressed video with no interruption to normal telephone service.

Most Internet applications require uneven amounts of upstream and downstream bandwidth. In other words, users tend to move greater amounts of data in one direction than they do in the other direction. Generally speaking, Internet users access more information and receive more data than they transmit. That's the nature of the Internet. They read more electronic mail than they send...they download more video information than they produce. A user's upstream capacity is typically limited to sending commands or transmitting small data files to a server. Much more information travels downstream. ADSL was designed to take advantage of this trend in uneven bandwidth. It provides data rates in excess of 8 Mbps from the network to the subscriber (downstream), and up to 640 Kbps from the subscriber to the network (upstream).

For telephone companies, ADSL may be the key to overcoming the congestion created by explosive Internet growth. If telephone companies hope to capture the Internet data carrier market, they can't use equipment designed for short voice traffic. The average voice call lasts about three minutes. According to a 1997 study by Bellcore, the research laboratory funded by US local telephone companies, (one of five submitted to the Federal Communications Commission) the average length of time for a typical Internet call is now in excess of 20 minutes. Telephony magazine reported as early as 1994 that nearly 20 percent of all on-line data service connections lasted more than an hour. Cable companies like MediaOne, based in the Northeast, charge one flat rate for Internet use. "Broadband" subscribers can stay on line 24-hours a day for roughly $40 a month.

ADSL was designed to take advantage of this natural unevenness in bandwidth requirements and make the most of available bandwidth, a limited resource. It provides for rates in excess of 8 Mbps from the network to the subscriber, and up to 640 Kbps from the subscriber to the network.

While some other telephony-based technologies offer relief for the PSTN, many carry high price tags and long-term implementation time frames. Replacing the copper network with fiber to the curb, for example, is a costly proposition. Industry experts estimate that creating a fiber-to-the-curb network will cost about $1,500 per telephone customer. Replacing the existing 560 million copper subscriber lines worldwide would cost more than three-quarters of a trillion dollars.

ADSL, by contrast, is easily installed. To complete an ADSL circuit, two modems are required, one at each end of a twisted-pair telephone line. One modem is located at the subscriber's premise; the other modem (usually a rack of modems with line cards) is located at the local telephone company's central office. Figure 9.10 shows a simplified connection.

At the customer premise, an ADSL modem is inside a PC or connected to the network computer, as well as the telephone and/or facsimile machine. A copper twisted pair telephone line links the POTS splitter to the central office. At the

central office, where the switch and rack of modems (line cards) are located, a POTS splitter separates voice from data. Voice calls are forwarded to the central office switch for relay over the public telephone network. Data transmissions are sent through an Ethernet switch and a router and eventually over a high-speed connection (for example, 155 Mbps OC-3) to the Internet service provider. The connection is made to the Internet backbone for access to the World Wide Web.

The key to ADSL's high-speed data transmission is advanced digital signal processing (DSP). Drawing on its analog innovation and advanced digital signal processing capabilities, Analog Devices developed the first generation AD20msp910 chipset (see Figure 9.11). The AD20msp910 offers three competitive advantages to modem manufacturers:

- It is the most functionally complete chipset solution on the market.
- It is fully compliant to ANSI, ETSI and ITU industry standards.
- It is compatible with nearly every DLC (digital loop carrier) and modem manufacturer.

The AD20msp910 solution includes both hardware and software. It integrates a DSP host processor, line driver, and control software, plus DMT technology. Other supplier's chips offer only portions of the modem. The AD20msp910 speeds and simplifies development of ADSL modems for high-speed Internet access and multimedia services and was released in 1997. The second generation AD20msp918 adds ATM functionality, improved performance and support for ADSL over ISDN for the European market. Both chipsets are fully compliant to all standards (ANSI T1.413 Issue 2, ETSI TR328, ITU G.dmt and ITU G.lite for splitterless) and both are complete solutions, including all microcontroller hardware and software functions.

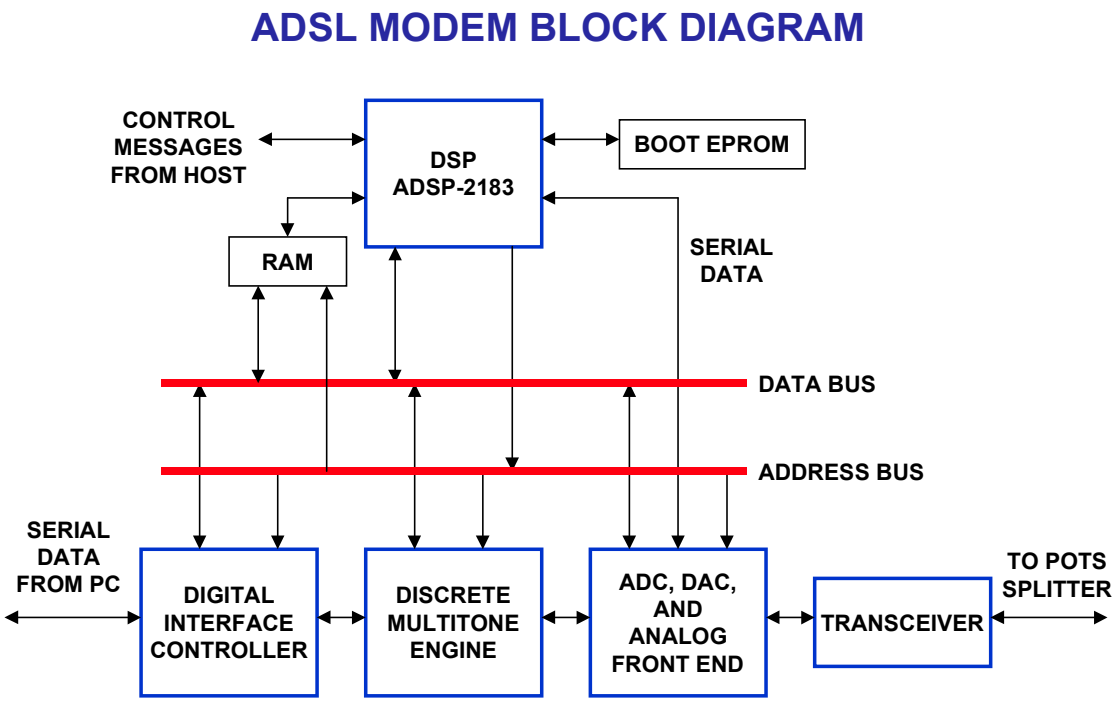

Figure 9.11

Now on our third generation chipset, the AD20msp930, Analog Devices enables modem manufacturers to reduce their development time and costs. To further speed and streamline design efforts, Analog Devices will also provide a PC board layout, schematics, and a reference design. As a result, modem manufacturers can focus their development efforts on adding value to the core ADSL modem technology.

DIGITAL CELLULAR TELEPHONES

In the early 1990s, the GSM (Global System for Mobile Communications aka Groupe Speciale Mobile) was introduced in Europe heralding the introduction of digital cellular radio. Out of necessity, other countries, including the US, have adopted various digital standards in order to relieve the congestion and other problems faced by analog cellular systems such as AMPS (Advanced Mobile Phone Service). The limitations of analog cellular radio are well known and include call blocking during busy hours, misconnects and disconnects due to rapidly fading signals, lack or privacy and security, and limited data transmission rates.

The basic cellular system is shown in Figure 9.12. A region is broken up into cells, with each cell having its own base station and its own group of assigned frequencies. Because the radius of each cell is small (10 miles, for example) low power transmitters and receivers can be used.

CELLULAR PHONE FREQUENCY REUSE

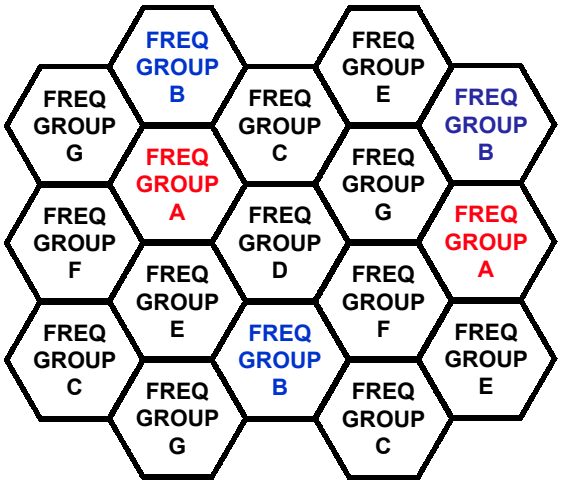

- Each Cell Covers a Radius of 5 to 10 miles
- Each Cell Requires its own Basestation for Reception and Re-Transmission
- Each Cell Must Handle Multiple Callers Simultaneously
- Callers May Be Handed Off from One Cell to Another

Figure 9.12

The cellular system lends itself to frequency reuse, since cells which are far enough apart can utilize the same band of frequencies without interference. The base stations must be linked together with an elaborate central control network so that a call may be handed-off to another cell when the signal strength from the mobile unit becomes too low for the current cell to handle.

The frequency spectrum allocation for analog cellular radio (AMPS) in the United States is approximately 825 to 850MHz and 870 to 895MHz. Conventional architectures (both analog and digital) are channelized. The total spectrum is divided up into a large number of relatively narrow channels, defined by a carrier frequency. The carrier frequency is frequency-modulated with the voice signal using analog techniques. Each full-duplex channel requires a pair of frequencies, each with a bandwidth of approximately 30kHz. A user is assigned both frequencies for the duration of the call. The forward and reverse channel are widely separated, to help the radio keep the transmit and receive functions separated. The signal bandwidth for the "A" or "B" carriers serving a particular geographical area is 12.5MHz (416 channels, each 30kHz wide). There can only be one caller at a time per channel.

Time Division Multiple Access (TDMA) allocates bandwidth on a time-slot basis. In the United States TDMA system, the entire 30kHz channel is assigned to a particular transmission, but only for a short period of time. A 3:1 multiplexing scheme means that three conversations can take place with TDMA using the same amount of bandwidth as one analog cellular conversation does. Each transmit/receive sequence occurs on time slots lasting 6.7ms. The TDMA system relies on an extensive amount of DSP technology to reduce the coded speech bit-rate as well as to prepare the digital data for transmission over the analog medium. The TDMA approach was chosen for the GSM system and will be discussed later in more detail.

The second digital approach used in the United States is called Code Division Multiple Access (CDMA). This technique has been used in secure military communications for a number of years under the name of *spread spectrum*. In spread spectrum, the transmitter transmits in a pseudo-random sequence of frequency hops over a relatively wide frequency range. The receiver has access to the same random sequence and can decode the transmission. The effect of adding additional users on the system is to decrease the overall signal to noise ratio for all the users. With this technique, the effect of allowing more calls than the normal capacity is to increase the bit-error rate for all users. New callers can keep coming in, interference levels will rise gradually, until at some point the process will become self-regulating: the quality of the voice link will become so bad that users will cut short or refrain from making additional calls. No one is ever blocked in the conventional sense, as they are in FDMA or TDMA systems when all channels or slots are full.

Both TDMA and CDMA digital systems make extensive use of DSP algorithms in both speech encoding and in preparing the signal for transmission. In the receiver, DSP techniques are used for demodulation and decoding the speech signal.

At the present time, both the analog and digital systems are in use in the US. In many cases, analog and digital systems must co-exist in the same geographical area

and within the same service area. This implies that the cellular basestations must handle both analog and digital formats, and therefore there is extensive use of digital techniques in basestation design to simplify the hardware.

The remainder of this section will concentrate on speech processing and channel coding as they relate to the GSM system. This will serve to illustrate the fundamental principles which are applicable to all digital mobile radio systems.

DIGITAL MOBILE RADIO STANDARDS

- **Frequency Division Multiple Access (FDMA) - User Allocation based on Frequency Slots (AMPS System)**
- **Time Division Multiple Access (TDMA) - User Allocation Based on Time Slots (At least 3X More Capacity than FDMA) GSM is an Example of TDMA**
- **Code Division Multiple Access (CDMA) - Based on Spread Spectrum Technology - More Users Cause Graceful Degradation in Bit-Error Rate**
- **Both TDMA and CDMA make Extensive Use of DSP in Speech Encoding and Channel Coding for Transmission and Reception**

Figure 9.13

The GSM System

Figure 9.14 shows a simplified block diagram of the GSM Digital Cellular Telephone System. The *speech encoder and decoder* and *discontinuous transmission* function will be described in detail. Up conversion and downconversion portions of the system contain a digital modem similar to the ones previously discussed. Similar functions are performed digitally such as equalization, convolutional coding, Viterbi decoding, modulation and demodulation.

The standard for encoding voice has been set in the T-Carrier digital transmission system. In this system, speech is logarithmically encoded to 8 bits at a sampling rate of 8kSPS. The logarithmic encoding and decoding to 8 bits is equivalent to linear encoding and decoding to 13 bits of resolution. This produces a bit-rate of 104kb/s. In most handsets, a 16-bit sigma-delta ADC is used, so the effective bit-rate is 128kb/s. The Speech Encoder portion of the GSM system compresses the speech signal to 13kb/s, and the decoder expands the compressed signal at the receiver. The speech encoder is based on an enhanced version of linear predictive coding (LPC). The LPC algorithm uses a model of the human vocal tract that represents the throat as a series of concentric cylinders of various diameters. An excitation (breath) is forced into the cylinders. This model can be mathematically represented by a series of simultaneous equations which describe the cylinders.

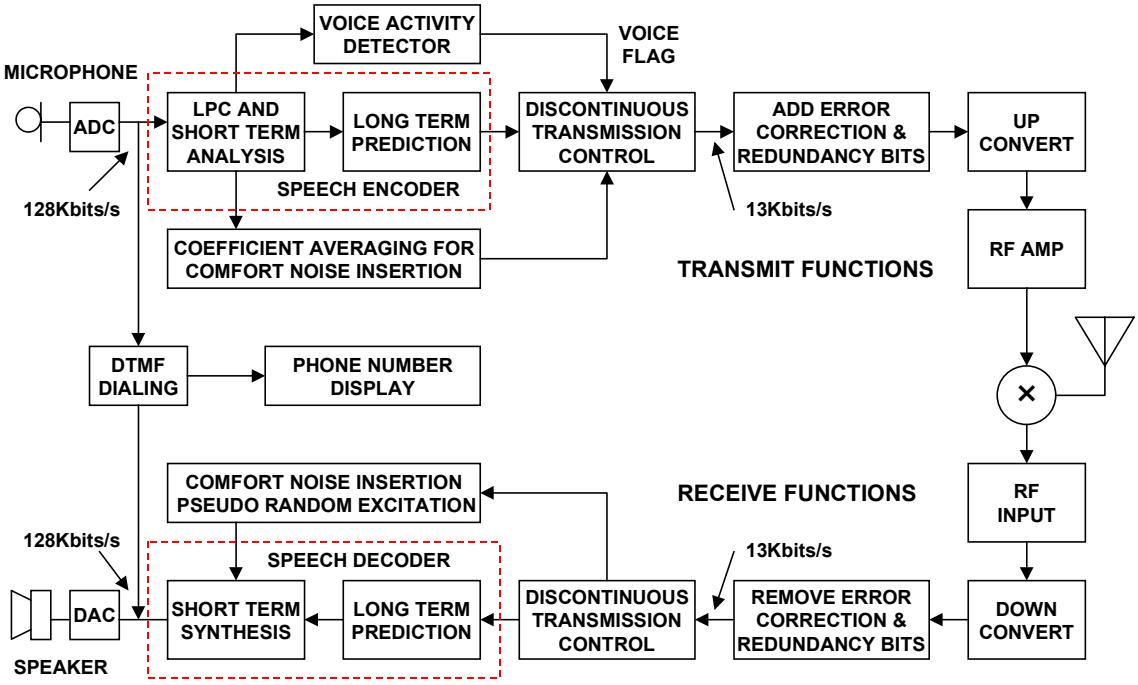

Figure 9.14

The excitation signal is passed through the cylinders, producing an output signal. In the human body, the excitation signal is air moving over the vocal cords or through a constriction in the vocal tract. In a digital system, the excitation signal is a series of pulses for vocal excitation, or noise for a constriction. The signal is input to a digital lattice filter. Each filter coefficient represents the size of a cylinder.

An LPC system is characterized by the number of cylinders it uses in the model. Eight cylinders are used in the GSM system, and eight reflection coefficients must be generated.

Early LPC systems worked well enough to understand the encoded speech, but often the quality was too poor to recognize the voice of the speaker. The GSM LPC system employs two advanced techniques that improve the quality of the encoded speech. These techniques are *regular pulse excitation* (RPE) and *long term prediction* (LTP). When these techniques are used, the resulting quality of encoded speech is nearly equal to that of logarithmic pulse code modulation (companded PCM as in the T-Carrier system).

The actual input to the speech encoder is a series of 16-bit samples of uniform PCM speech data. The sampling rate is 8kHz. The speech encoder operates on a 20ms window (160 samples) and reduces it to 76 coefficients (260 bits total), resulting in an encoded data rate of 13kb/s.

Discontinuous transmission (DTX) allows the system to shut off transmission during the pauses between words. This reduces transmitter power consumption and increases the overall GSM system's capacity.

Low power consumption prolongs battery life in the handset and is an important consideration for hand-held portable phones. Call capacity is increased by reducing the interference between channels, leading to better spectral efficiency. In a typical conversation each speaker talks for less than 40% of the time, and it has been estimated that DTX can approximately double the call capacity of the radio system.

The voice activity detector (VAD) is located at the transmitter. Its job is to distinguish between speech superimposed on the background noise and noise with no speech present. The input to the voice activity detector is a set of parameters computed by the speech encoder. The VAD uses this information to decide whether or not each 20ms frame of the encoder contains speech.

Comfort noise insertion (CNI) is performed at the receiver. The comfort noise is generated when the DTX has switched off the transmitter; it is similar in amplitude and spectrum to the background noise at the transmitter. The purpose of the CNI is to eliminate the unpleasant effect of switching between speech with noise, and silence. If you were listening to a transmission without CNI, you would hear rapid alternating between speech in a high-noise background (i.e. in a car), and silence. This effect greatly reduces the intelligibility of the conversation.

When DTX is in operation, each burst of speech is transmitted followed by a *silence descriptor* (SID) frame before the transmission is switched off. The SID serves as an end of speech marker for the receive side. It contains characteristic parameters of the background noise at the transmitter, such as spectrum information derived through the use of linear predictive coding.

The SID frame is used by the receiver's comfort noise generator to obtain a digital filter which, when excited by pseudo-random noise, will produce noise similar to the background noise at the transmitter. This comfort noise is inserted into the gaps between received speech bursts. The comfort noise characteristics are updated at regular intervals by the transmission of SID frames during speech pauses.

Redundant bits are then added by the processor for error detection and correction at the receiver, increasing the final encoded bit rate to 22.8kb/s. The bits within one window, and their redundant bits, are interleaved and spread across several windows for robustness.

GSM HANDSET USING SOFTFONE™ BASEBAND PROCESSOR AND OTHELLO™ RADIO

Analog Devices recently announced two chipsets which comprise the majority of a GSM handset. The SoftFone™ chipset performs the baseband and DSP functions, while the Othello™ radio chipset handles the RF functions.

The frequency bands originally allocated to GSM were 890MHz to 915MHz for mobile transmitting and 935MHz to 960MHz for mobile receiving. Another frequency allocation was made to further expand GSM capacity. This band, allocated to digital communications services (DCS), was 1710MHz to 1785MHz and 1805MHz to 1880MHz. All countries adopting GSM use one of these two pairs of frequencies. In the United States, these bands were already allocated by the FCC. In the mid-1990s, a set of bands were made available for GSM in the United States: 1850MHz to 1910MHz, and 1930MHz to 1990MHz.

Because of the frequency allocations in GSM countries (other than the U.S.), most GSM handsets must be dual-band: capable of handling both GSM and DCS frequencies. The SoftFone and Othello chipsets supply the main functions necessary for implementing dual- or triple-band radios for GSM cellular phones. The AD20msp430 SoftFone™ chipset comprises the baseband portion of the GSM handset. The AD20msp430 baseband processing chipset uses a combination of GSM system knowledge and advanced analog and digital signal processing technology to provide a new benchmark in GSM/GPRS terminal design.

The SoftFone architecture is entirely RAM-based. The software is loaded from FLASH memory and is executed from the on-chip RAM. This allows fast development cycles since no ROM-code turns are required. Furthermore, the handset software can be updated in the field to enable new features. Combined with the Analog Devices “Othello” RF chipset, a complete multiband handset design contains less than 200 components, fits in a 20 cm^2 single sided PCB layout, and has a total bill-of-materials cost 20-30% lower than previous solutions. A simplified block diagram of the handset is shown in Figure 9.15.

The AD20msp430 chipset is comprised of two chips, the AD6522 DSP-based baseband processor and the AD6521 voiceband/baseband mixed-signal codec. Together with the “Othello ” radio chipset, the AD20msp430 allows a significant reduction in the component count and bill-of-materials (BOM) cost of GSM voice handsets and data terminals. The software and hardware foundations of the AD20msp430 chipset enjoy a long history of successful integration into GSM handsets.

OTHELLO™ RADIO AND SOFTPHONE™ CHIPSETS MAKE COMPLETE GSM/DCS HANDSET

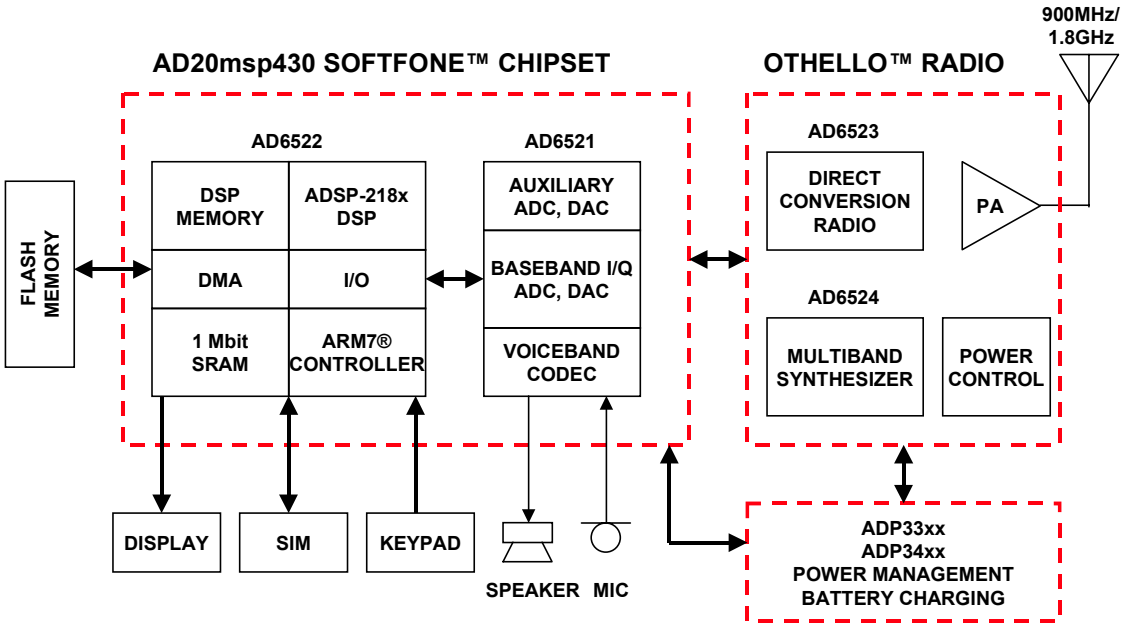

Figure 9.15

This is Analog Devices' 4th generation of GSM chipsets, each of which has passed numerous type- approvals and network operator approvals in OEM handsets. In each generation, additional features have been added, while cost and power have been reduced. Numerous power-saving features have been included in the AD20msp430 chipset to reduce the total power consumption. A programmable state machine allows events to be controlled with a resolution of one-quarter of a bit period. This reduces current in standby mode to 1mA. This allows a handset to operate in standby mode for up to 1000 hours on a battery charge. The AD20msp430 chipset uses the SoftFone™ architecture, where all software resides in RAM or FLASH memory. Since ROM is not used, development time is reduced and additional features can be field-installed easily. A basic dual-band GSM terminal typically requires a single 8Mb FLASH memory chip.

There are two processors in the AD20msp430 chipset. The DSP processor is the ADSP-218x core, proven in previous generations of GSM chipsets, and operated at 65MIPS in the AD20msp430. This DSP performs the voiceband and channel coding functions previously discussed. The AD6521 voiceband/baseband codec chip contains all analog and mixed-signal functions. These include the I/Q channel ADCs and DACs, high performance multichannel voiceband codec, and several auxiliary ADCs and DACs for AGC (automatic gain control), AFC (automatic frequency control),and power- amplifier ramp control. The microcontroller is an ARM7 TDMI, running at 39MIPS. The ARM7 handles the protocol stack and the man-machine interface functions. Both processors are field-proven in digital wireless applications.

The AD20msp430 chipset is fully supported by a suite of development tools and software. The development tools allow easy customization of the DSP and/or ARM® controller software to allow handset and terminal manufacturers to optimize the feature set and user interface of the end equipment. Software is available from Analog Devices' software partner TTPCom for all layers, including both voice and data applications, and is updated as new features become available. The system DMA and interrupt controllers are designed to allow easy upgrades to future generations of DSP and controller cores. The display interface can be used with either parallel or serial-interface displays. System development can be shortened by the use of the debugging features in the AD20msp430. Most critical signals can be routed under software control to the Universal System Connector. This allows system debugging to take place in the final form factor. In addition, the architecture includes high speed logger and address trace functions in the DSP and single-wire trace/debug in the ARM controller.

Analog Devices recently announced the revolutionary Othello™ direct-conversion radio for mobile applications. By eliminating intermediate-frequency (IF) stages, this chip set will permit the mobile electronics industry to reduce the size and cost of radio sections and enable flexible, multi-standard, multi-mode operation. The radio consists of two integrated circuits, the AD6523 Zero-IF Transceiver and the AD6524 Multi-Band Synthesizer.

The AD6523 contains the main functions necessary for both a direct-conversion receiver and a direct VCO transmitter, known as the Virtual-IF™ transmitter. It also includes the local-oscillator generation block and a complete on-chip regulator that supplies power to all active circuitry for the radio. The AD6524 is a fractional-N synthesizer that features extremely fast lock times to enable advanced data services over cellular telephones-such as high-speed circuit-switched data (HSCSD) and general packet radio services (GPRS).

Most digital cellular phones today include at least one "downconversion" in their signal chain. This frequency conversion shifts the desired signal from the allocated RF band for the standard (say, at 900 MHz) to some lower intermediate frequency (IF), where channel selection is performed with a narrow channel-select filter (usually a surface acoustic-wave (SAW) or a ceramic type). The now-filtered signal is then further down-converted to either a second IF or directly to baseband, where it is digitized and demodulated in a digital signal processor (DSP). Figure 9.16 shows the comparison between this *superheterodyne* architecture and the superhomodyne™ architecture of the Othello radio receiver.

The idea of using direct-conversion for receivers has long been of interest in RF design. The reason is obvious: in consumer equipment conversion stages add cost, bulk, and weight. Each conversion stage requires a local oscillator, (often including a frequency synthesizer to lock the LO onto a given frequency), a mixer, a filter, and (possibly) an amplifier. No wonder, then, that direct conversion receivers would be attractive. All intermediate stages are eliminated, reducing the cost, volume, and weight of the receiver.

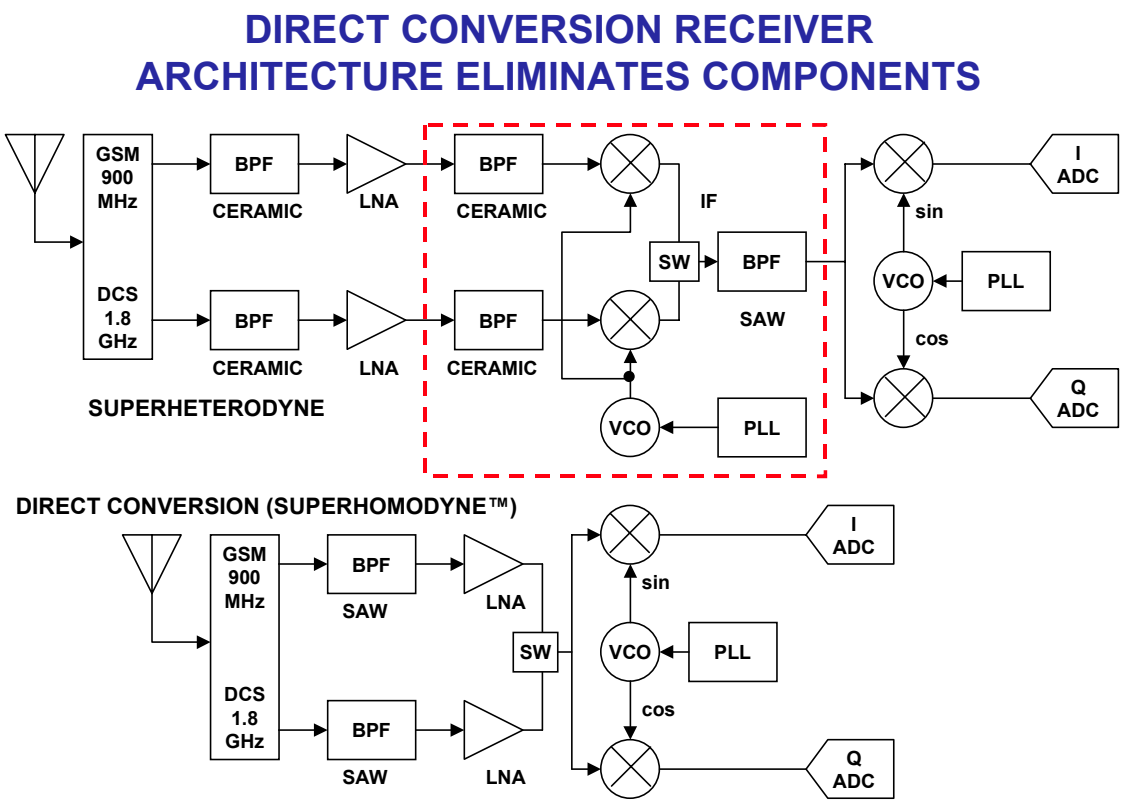

Figure 9.16

The Othello™ radio reduces the component count even more by integrating the front-end GSM low-noise amplifier (LNA). This eliminates an RF filter (the "image" filter) that is necessary to eliminate the image, or unwanted mixing product of a mixer and the off chip LNA. This stage, normally implemented with a discrete transistor, plus biasing and matching networks, accounts for a total of about 12 components. Integrating the LNA saves a total of about 15 to 17 components, depending on the amount of matching called for by the (now-eliminated) filter.

A functional block diagram of the Othello™ dual band GSM radio's architecture is shown in Figure 9.17. The receive section is at the top of the figure. From the antenna connector, the desired signal enters the transmit/receive switch and exits on the appropriate path, either 925-960 MHz for the GSM band or 1805-1880 MHz for DCS. The signal then passes through an RF band filter (a so-called "roofing filter") that serves to pass the entire desired frequency band while attenuating all other out-of-band frequencies (blockers-including frequencies in the transmission band) to prevent them from saturating the active components in the radio front end. The roofing filter is followed by the low-noise amplifier (LNA). This is the first gain element in the system, effectively reducing the contribution of all following stages to system noise. After the LNA, the direct-conversion mixer translates the desired signal from radio frequency (RF) all the way to baseband by multiplying the desired signal with a local oscillator (LO) output at the same frequency.

SUPERHOMODYNE™ DIRECT CONVERSION DUAL-BAND TRANSCEIVER USING THE AD6523/AD6524 CHIPSET

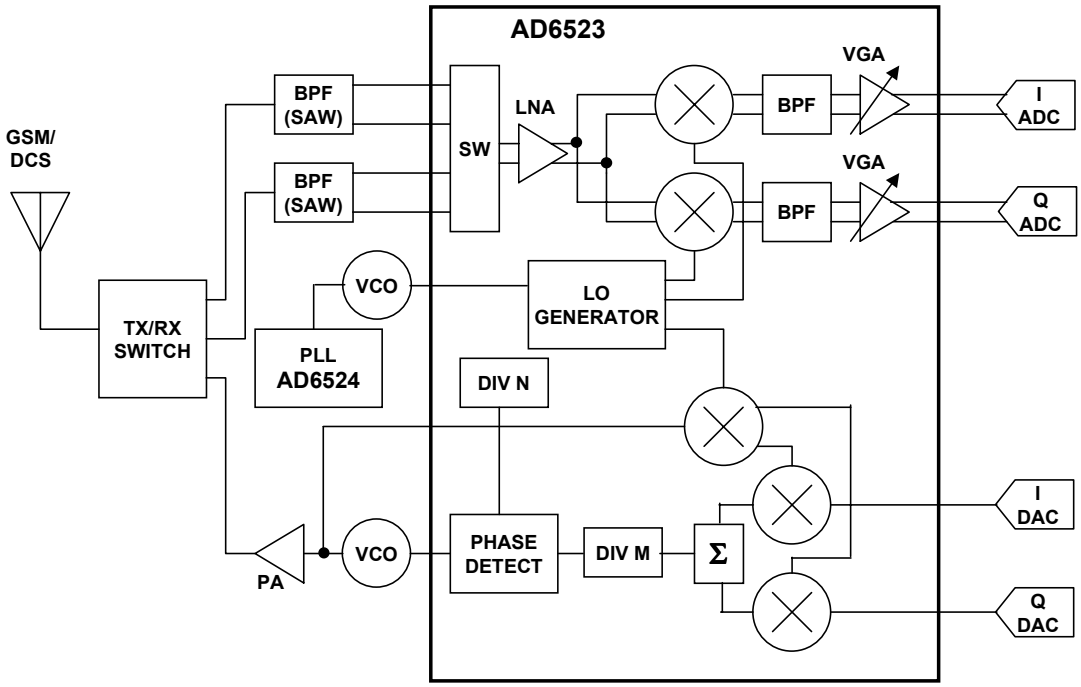

Figure 9.17

The output of the mixer stage is then sent in quadrature (I and Q channels) to the variable-gain baseband amplifier stage. The VGA also provides some filtering of adjacent channels, and attenuation of in-band blockers. These blocking signals are other GSM channels that are some distance from the desired channel, say 3 MHz and beyond. The baseband amplifiers filter these signals so that they will not saturate the Receive ADCs. After the amplifier stage, the desired signal is digitized by the Receive ADCs.

The Transmit section begins on the right, at the multiplexed I and Q inputs/outputs. Because the GSM system is a time division duplex (TDD) system, the transmitter and receiver are never on at the same time. The Othello™ radio architecture takes advantage of this fact to save four pins on the transceiver IC's package. The quadrature transmit signals enter the transmitter through the multiplexed I/Os. These I and Q signals are then modulated onto a carrier at an intermediate frequency greater than 100 MHz.

The output of the modulator goes to a phase-frequency detector (PFD), where it is compared to a reference frequency that is generated from the external channel selecting LO. The output of the PFD is a charge pump, operating at above 100 MHz, whose output is filtered by a fairly wide (1MHz) loop filter. The output of the loop filter drives the tuning port of a voltage-controlled oscillator (VCO), with frequency ranges that cover the GSM and DCS transmit bands.

The output of the transmit VCO is sent to two places. The main path is to the transmit power amplifier (PA), which amplifies the transmit signal from about +3dBm to +35dBm, sending it to the transmit/receive switch and low pass filter (which attenuates power-amplifier harmonics). The power amplifiers are dual band, with a simple CMOS control voltage for the band switch. The VCO output also goes to the transmit feedback mixer by means of a coupler, which is either a printed circuit, built with discrete inductors and capacitors, or a monolithic (normally ceramic) coupling device. The feedback mixer downconverts the transmit signal to the transmit IF, and uses it as the local oscillator signal for the transmit modulator. This type of modulator has several names, but the most descriptive is probably "translation loop". The translation loop modulator takes advantage of one key aspect of the GSM standard: the modulation scheme is Gaussian-filtered minimum-shift keying (GMSK). This type of modulation does not affect the envelope amplitude, which means that a power amplifier can be saturated and still not distort the GMSK signal sent through it.

GMSK can be generated in several different ways. In another European standard (for cordless telephones), GMSK is created by directly modulating a free running-VCO with the Gaussian filtered data stream. In GSM, the method of choice has been quadrature modulation. Quadrature modulation creates accurate phase GMSK, but imperfections in the modulator circuit (or up-conversion stages) can produce envelope fluctuations, which can in turn degrade the phase trajectory when amplified by a saturated power amplifier. To avoid such degradations, GSM phone makers have been forced to use amplifiers with somewhat higher linearity, at the cost of reduced efficiency and talk time per battery charge cycle.

The translation loop modulator combines the advantages of directly modulating the VCO and the inherently more accurate quadrature modulation. In effect, the scheme creates a phase locked loop (PLL), comprising the modulator, the LO signal, and the VCO output and feedback mixer. The result is a directly modulated VCO output with a perfectly constant envelope and almost perfect phase trajectory. Phase trajectory errors as low as 1.5 degrees have been measured in the AD6523 transceiver IC, using a signal generator as the LO signal to provide a reference for the loop.

Because Othello™ radios can be so compact, they enable GSM radio technology to be incorporated in many products from which it has been excluded, such as very compact phones or PCMCIA cards. However, the real power of direct conversion will be seen when versatile third-generation phones are designed to handle multiple standards. With direct-conversion, hardware channel-selection filters will be unnecessary, because channel selection is performed in the digital signal-processing section, which can be programmed to handle multiple standards. Contrast this with the superheterodyne architecture, where multiple radio circuits are required to handle the different standards (because each will require different channel-selection filters), and all the circuits will have to be crowded into a small space. With direct conversion, the same radio chain could in concept be used for several different standards, bandwidths, and modulation types. Thus, Web-browsing and voice services could, in concept, occur over the GSM network using the same radio in the handset.

ANALOG CELLULAR BASESTATIONS

Consider the analog superheterodyne receiver invented in 1917 by Major Edwin H. Armstrong (see Figure 9.18). This architecture represented a significant improvement over single-stage direct conversion (homodyne) analog receivers which had previously been constructed using tuned RF amplifiers, a single detector, and an audio gain stage. (Note that the homodyne technique has now gained favor in DSP-based receivers as explained above). A significant advantage of the superhetrodyne analog receiver is that it is much easier and more economical to have the gain and selectivity of a receiver at fixed intermediate frequencies (IF) than to have the gain and frequency-selective circuits "tune" over a band of frequencies.

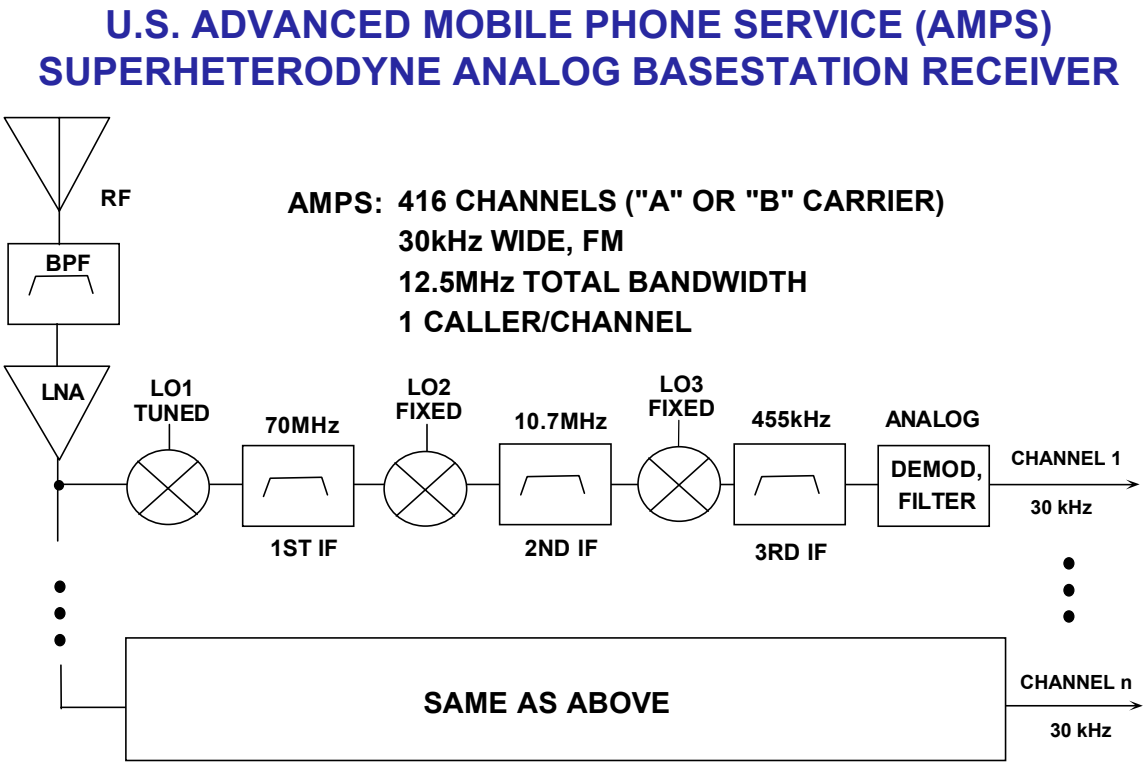

Figure 9.18

The frequencies shown in Figure 9.18 correspond to the AMPS (Advanced Mobile Phone Service) analog cellular phone system currently used in the U.S. The receiver is designed for AMPS signals at 900MHz RF. The signal bandwidth for the "A" or "B" carriers serving a particular geographical area is 12.5MHz (416 channels, each 30kHz wide). The receiver shown uses triple conversion, with a first IF frequency of 70MHz, a second IF of 10.7MHz, and a third of 455kHz. The image frequency at the receiver input is separated from the RF carrier frequency by an amount equal to twice the first IF frequency (illustrating the point that using relatively high first IF frequencies makes the design of the image rejection filter easier).

The output of the third IF stage is demodulated using analog techniques (discriminators, envelope detectors, synchronous detectors, etc.). In the case of AMPS the modulation is FM. An important point to notice about the above scheme

is that there is *one receiver required per channel*, and only the antenna, prefilter, and LNA can be shared.

It should be noted that in order to make the receiver diagrams more manageable, the interstage amplifiers are not shown. They are, however, an important part of the receiver, and the reader should be aware that they must be present.

It should be seen by now that analog receiver design is a complicated art, and there are many tradeoffs that can be made between IF frequencies, single-conversion vs. double-conversion or triple conversion, filter cost and complexity at each stage in the receiver, demodulation schemes, etc. There are many excellent references on the subject, and the purpose of this discussion is to form an historical frame of reference for the following discussions on the application of digital techniques in the design of advanced communications and basestation receivers.

DIGITAL CELLULAR BASESTATIONS

Cellular telephone basestations form the backbone of the modern wireless cellular infrastructure. They must receive multiple calls, process the calls, and re-transmit them. Handoffs to basestations in adjacent cells must be done seamlessly with respect to the mobile customer. In addition, the basestation must often handle multiple standards simultaneously. Many carriers in certain areas of the United States use multiple technologies in the same geographical area; AMPS and CDMA, for example.

Flexibility, performance, and low per-channel costs are the keys to modern basestations. Maximizing the use of DSP in the tranceiver allows multiple standards to be handled without the necessity for changing hardware. This has led to the widespread proliferation of so-called *software radios* which dominate the current basestation market.

As in the case of the cell phone handset, direct conversion techniques are used in basestations. The signal is digitizied by a high performance wideband ADC after only one stage of downconversion. Figure 9.19 shows two fundamental approaches to a digital receiver: *narrowband* and *wideband*.

By *narrowband*, we mean that sufficient prefiltering has been done such that all undesired signals have been eliminated and that only the channel of interest is presented to the ADC input. *Wideband* simply means that a number of channels are presented to the input of the ADC and further filtering, tuning, and processing is performed digitally. Usually, a wideband receiver is designed to receive an entire band; cellular or other similar wireless service such as PCS or CDMA. In fact, one wideband digital receiver can be used to receive all channels within the band simultaneously, allowing almost all of the analog hardware (including the ADC) to be shared among all the channels.

NARROWBAND AND WIDEBAND DIGITAL RECEIVERS FOR CELLULAR BASESTATIONS

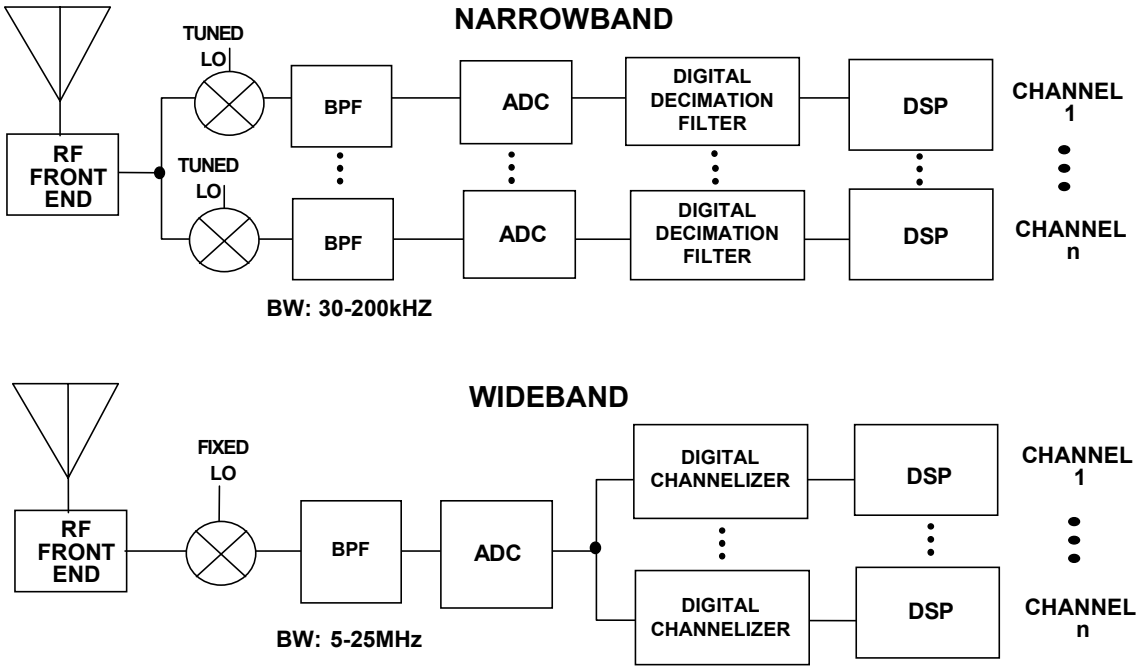

Figure 9.19

The wideband approach places severe constraints on the ADC and requires high spurious free dynamic range (SFDR) and signal-to-noise ratio (SNR), especially in cellular applications where signal strengths of channels can differ by more than 100dB. This requires ADCs with bandwidths of greater than 100MHz and sampling frequencies greater than 50MSPS (required to handle a multicarrier bandwidth of 25MHz, for example). On the other hand, the narrowband approach provides more processing gain because each channel is highly oversampled, but this approach also requires more ADCs to process the same number of channels.

ADI's SoftCell™ chipset addresses key issues for wireless operators including the cost of coverage, flexibility and size, as well as quality of service. Basestations containing the SoftCell chip set are easy to modify in terms of adding services, additional channels, and changing wireless standards incrementally. In effect, operators will have the ability to use and move between any air interface standard (e.g., GSM, PHS, D-AMPS), deploy a higher number of channels, and offer frequency plans with greater efficiency. The new architecture also eliminates redundant channel radios for both transmitters and receivers.

The SoftCell chip set is optimized for four RF carrier channels and is easily expandable. This solution enables equipment manufacturers to create highly scalable multicarrier, multimode base stations at a fraction of the cost of traditional multichannel base stations using analog techniques. A block diagram of a system using the SoftCell chipset is shown in Figure 9.20.

BASESTATION BLOCK DIAGRAM USING SoftCell™ MULTICARRIER WIDEBAND CHIPSET

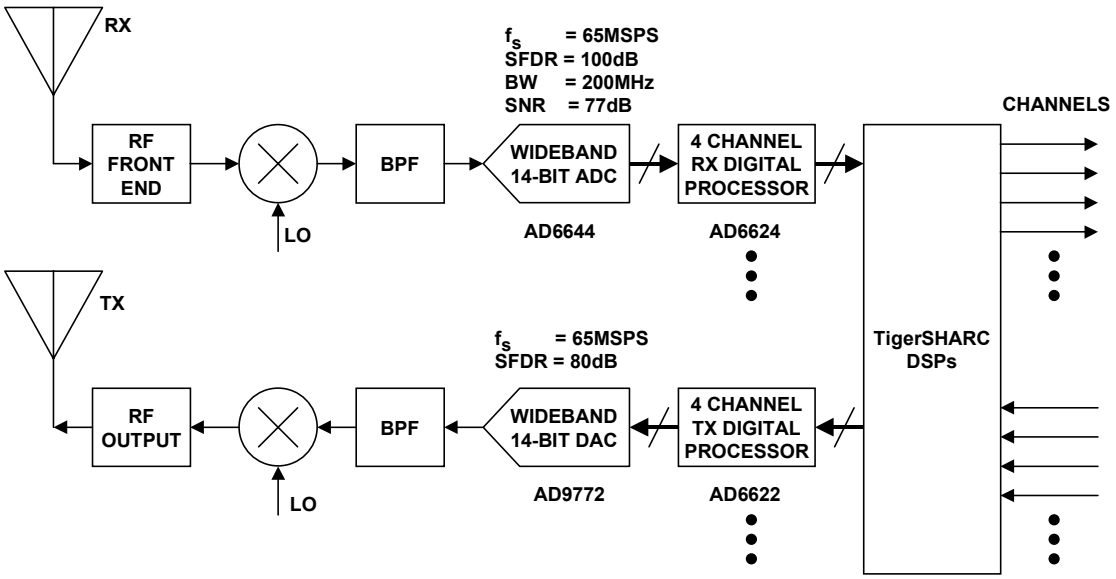

Figure 9.20

The decrease in size and cost savings associated with SoftCell means basestations can be deployed in higher numbers and in tighter locations. The result is better coverage, better quality, and fewer busy signals for users. The deployment capability also makes the SoftCell chip set well-suited for wireless office applications. In addition, the software radio technology implemented in the chipset enables new applications such as Smart Antenna or phased-array antenna services that enable power efficiencies and cost savings, and small *picocell* installations for additional coverage capacity and in-building wireless systems.

The SoftCell chip set consists of ADI's AD6644 14-bit ADC, AD6624 quad digital receive signal processor (RSP), AD9772 TxDAC® 14-bit DAC, and the AD6622 quad digital transmit signal processor (TSP). The solution leverages the benefits of digital signal processors to enable channel separation, equalization, error correction and decoding with greater flexibility and efficiencies. This new chip set is optimized to work with ADI's TigerSharc™ multiprocessor digital signal processor (DSP).

The TigerSHARC DSP is a ground breaking processor optimized for communications applications capable of performing approximately 1 billion 16-bit multiply/accumulates per second at 150 MHz. Additionally, the TigerSHARC is unique in its ability to support 8-, 16-, and 32-bit word sizes on a single chip. Modulation/demodulation, channel encode/decoding, and other radio functions can be multiplexed to support multiple carriers in a single DSP.

SoftCell complements ADI's recently introduced AD6600 diversity receiver ADC. The AD6600 addresses narrowband applications where a multicarrier architecture is still not feasible but allows direct IF sampling of signals at frequencies up to 250MHz. Configured with the appropriate digital receiver signal processor, the AD6600 can address a variety of air interface standards, including GSM Macrocell.

Classic basestation architectures require a complete transceiver for every RF carrier processed (from 4 to 80 channels for digital and analog systems respectively). These radios must be duplicated for diversity and sectorized antennas as well. It is easy to see why base station electronics consume so much space, power, and cost. The advantage of a multicarrier software radio is the elimination of redundant radios in favor of a single, high-performance radio per antenna, where each RF carrier is processed in the digital domain. Deployment of true software radios has been limited by the performance of analog-to-digital converters (ADC) which must digitize the enormous dynamic range demanded in a spectrum composed of multiple carriers, blockers, and adjacent channel interference.

Multicarrier transmitters have similar challenges to meet the performance demands of the newest air interface specifications. Digital-to-analog converters (DAC) and multicarrier-power-amplifiers (MCPAs) must preserve the spectrum of several digitally generated carriers without corruption or spurious signal generation in adjacent channels. The AD9772 is a 14-bit interpolating DAC optimized to accurately convert multiple carriers to a single IF frequency. The AD9772 is the latest in ADI's TxDAC® family of high-performance converters.

The heart of ADI's SoftCell chip set is the AD6644, a 14-bit, 65-MSPS ADC that provides up to 100dB of spurious-free-dynamic-range (SFDR) and 77dB signal-to-noise ratio (SNR). This provides the receiver performance needed to implement a multicarrier digitizing radio for many applications. Shifting the channel tuning, filtering, and demodulation to the digital domain allows the flexibility to support different air interface standards, number of channels, and frequency plans with a single radio design.

Following the ADC, a digital receive signal processor (RSP) performs the channel tuning, filtering, and decimation required to provide baseband I and Q signals to a digital signal processor (DSP). The AD6624 is an 65MSPS Quad Digital RSP developed to support GSM, IS136, and other narrowband standards. The AD6624's four channels are independently programmable to change air interface characteristics on demand. It is a simple matter to add AD6624s in parallel for additional channel capacity. The AD6624 can also be configured to support EDGE extensions of GSM and IS136.

The AD6622 is a four channel Digital Transmit Signal Processor that takes baseband I and Q inputs from a DSP and provides all the signal processing functions required to drive the AD9772 DAC. Each independent channel can be programmed to provide the desired channel filtering for most air interface standards. The AD6622 supports IS95 and WCDMA standards and can be daisy-chained to combine an arbitrary number of channels onto a single 18-bit digital output.

MOTOR CONTROL

Long known for its simplicity of construction, low-cost, high efficiency and long-term dependability, the AC induction motor has been limited by the inability to control its dynamic performance in all but the crudest fashion. This has severely restricted the application of AC induction motors where dynamic control of speed, torque and response to changing load is required. However, recent advances in digital signal processing (DSP) and mixed-signal integrated circuit technology are providing the AC induction motor with performance never before thought possible. Manufacturers anxious to harness the power and economy of Vector Control can reduce R&D costs and time to market for applications ranging from industrial drives to electric automobiles and locomotives by utilizing a standard chipset/development system.

It is unlikely that Nikola Tesla (1856-1943), the inventor of the induction motor, could have envisaged that this workhorse of industry could be rejuvenated into a new class of motor that is competitive in most industrial applications.

Before discussing the advantages of Vector Control it is necessary to have a basic understanding of the fundamental operation of the different types of electric motors in common use.

Until recently, motor applications requiring servo-control tasks such as tuned response to dynamic loads, constant torque and speed control over a wide range were almost exclusively the domain of DC brush and DC permanent magnet synchronous motors. The fundamental reason for this preference was the availability of well understood and proven control schemes. Although easily controlled, DC brush motors suffer from several disadvantages; brushes wear and must be replaced at regular intervals, commutators wear and can be permanently damaged by inadequate brush maintenance, brush/commutator assemblies are a source of particulate contaminants, and the arcing of mechanical commutation can be a serious fire hazard is some environments.

The availability of power inverters capable of controlling high-horsepower motors allowed practical implementation of alternate motor architectures such as the DC permanent magnet synchronous motor (PMSM) in servo control applications. Although eliminating many of the mechanical problems associated with DC brush motors, these motors required more complex control schemes and suffered from several drawbacks of their own. Aside from being costly, DC PMSMs in larger, high-horsepower configurations suffer from high rotor moment-of-inertia as well as limited use in high speed applications due to mechanical constraints of rotor construction and the need to implement field weakening to exceed baseplate speed.

In the 1960's, advances in control theory, in particular the development of *indirect field-oriented control*, provided the theoretical basis for dynamic control of AC induction motors. Indirect field-oriented control makes use of reference frame theory. Using these techniques, it is possible to transform the phase variable machine description of a motor to another reference frame. By judicious choice of the reference frame, it is possible to simplify considerably the complexity of the mathematical machine model. While these techniques were initially developed for the analysis and simulation of AC motors, they are now invaluable tools in the

digital control of such machines. As digital control techniques are extended to the control of the currents, torque, and flux of such machines, the need for compact, accurate motor models is obvious.

Fortunately, the theory of reference frames is equally applicable to the synchronous machines, such as the Permanent Magnet Synchronous Machine (PMSM). This motor is sometimes known as the sinusoidal brushless motor or the brushless AC machine, and is very popular as a high performance servo drive.

Because of the intensive mathematical computations required by indirect field-oriented control, now commonly referred to as *vector control* or *reference frame theory*, practical implementation was not possible for many years. Available hardware could not perform the high-speed precision sensing of rotor position and near real-time computation of dynamic flux vectors. The current availability of precision optical encoders, isolated gate bipolar transistors (IGBTs), high-speed resolver-to-digital converters and high-speed digital signal processors (DSPs) has pushed vector control to the forefront of motor development due to the advantages inherent in the AC induction motor.

A simplified block diagram of an AC induction motor control system is shown in Figure 9.21. The inputs to the controller are the motor currents (normally three-phase) and the motor rotor position and velocity. Hall-effect sensors are often used to monitor the currents, and a resolver and a resolver-to-digital converter (RDC) monitor the rotor position and velocity. The DSP is used to perform the real time vector-type calculations necessary to generate the control outputs to the inverter processors. The transformations required for reference frame transformations and vector control are also accomplished with the DSP.

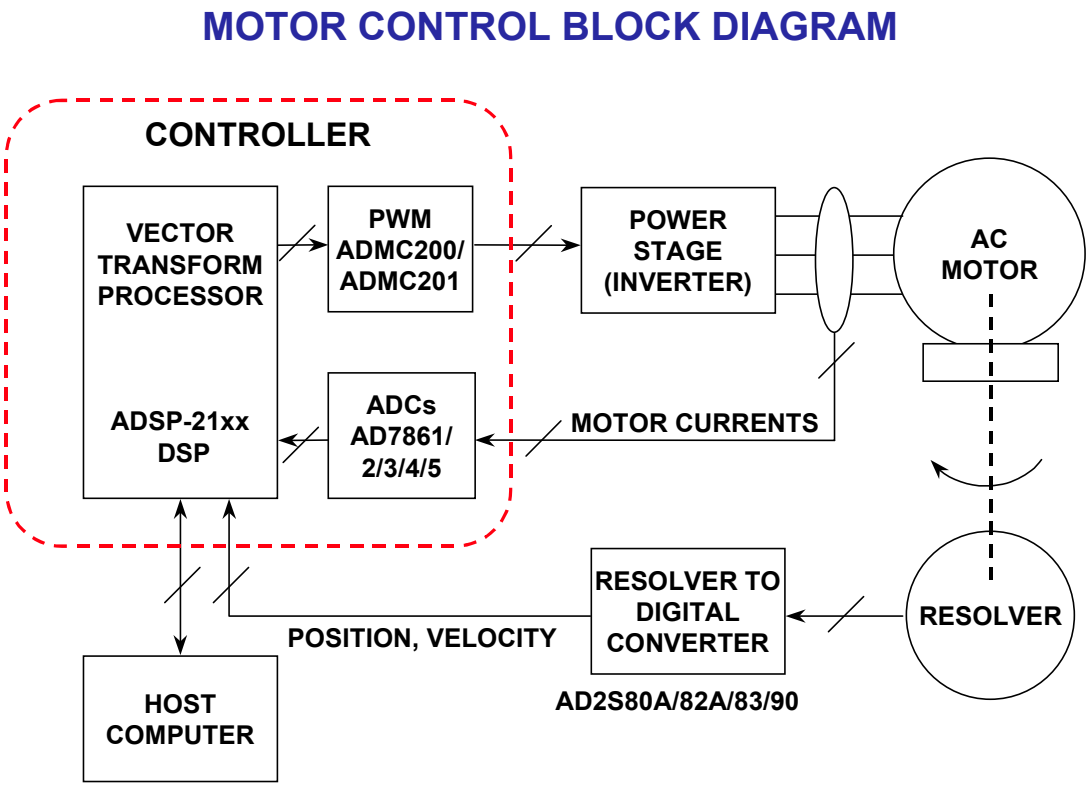

Figure 9.21

The functions in the controller block have been integrated into the Analog Devices' ADMC300, ADMC331, ADMC401, and the ROM-based ADMC326 and ADMC328 DSP motor controller chips. These devices include the peripheral circuitry such as ADCs, voltage references, PWM controllers, timers, etc., required to perform all the functions shown in Figure 9.21.

The latest members of the family are the ADMCF326 and ADMCF328 motor controllers, referred to as DashDSP™, meaning *digital* plus *analog* plus *flash* memory (see Figure 9.22). The use of flash memory allows the devices to be field programmable, thereby providing greater flexibility and shortening development time. These controllers include a 20MIPS 16-bit fixed-point core based on the ADSP-217x family architecture. The memory consists of a 512 x 24-bit program memory RAM, a 512 x 16-bit data memory RAM, 4K x 24-bit program memory ROM, and a 4K x 24-bit programmable flash memory. A completely integrated ADC analog subsystem allows the three-phase motor currents to be monitored. A 16-bit 3-phase PWM generates the control signals to the external inverter power stage. The parts are packaged in a 28-pin SOIC or PDIP package. A block diagram of the ADMCF328 is shown in Figure 9.23.

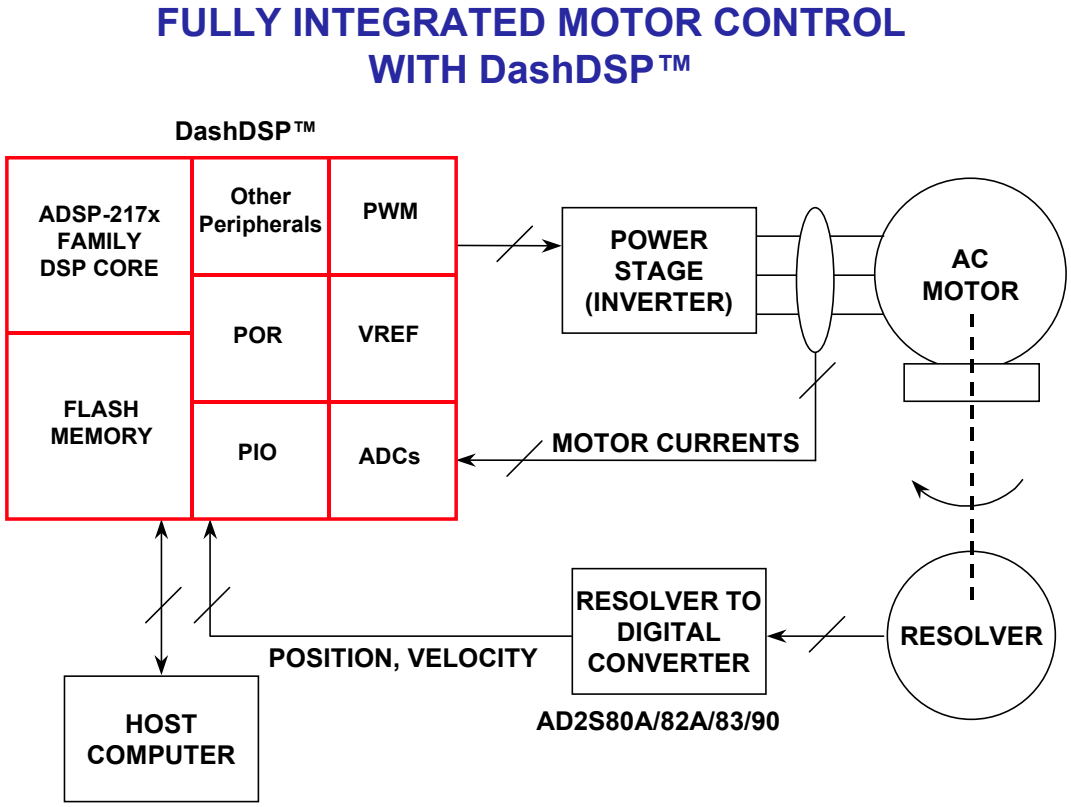

Figure 9.22

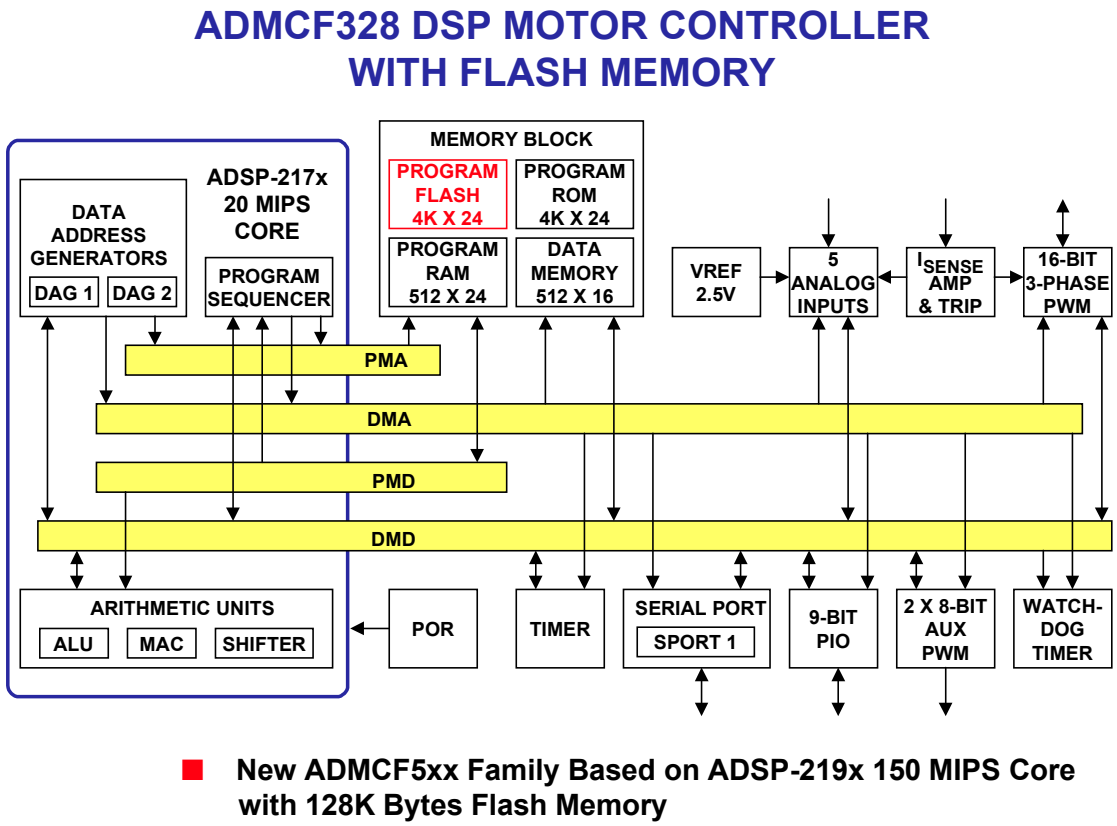

Figure 9.23

Third-party DSP software, reference designs, and evaluation systems are available to facilitate motor control system development using these chips.

Analog Devices recently announced a new ADMCF5xx DashDSP family based on the ADSP-219x, 150MIPS 16-bit fixed point core. At 0.4mA/MIP, this solution provides engineers a low power and wide ranging speed device. The integration of up to 128KBytes of embedded flash memory eases development and enables flexible system software modifications and system upgrades. This also allows users to boot directly from flash memory into RAM. With 10-, 12-, and 14-bit ADCs, customers can select a device that matches their system sensing specifications and cost requirements. The family is supported by ADI's VisualDSP tools, which incorporate the industry's first C++ compiler. The ADMCF5xx series is augmented by a wide range of high performance peripherals, such as 3 to 6 phase 16-bit PWMs with single or dual encoder interfaces for multi axis motor control. Motor current sampling can be achieved through the use of DC link or inverter shunt techniques, with user programmable converter sampling times. A precision voltage reference, a power-on reset circuit, and auxiliary PWMs capable of implementing power factor correction are also included. Additionally, various communication peripherals such as a full CAN bus, UARTs, SPORTs, and a JTAG interface have been integrated into the family.

CODECS AND DSPS IN VOICEBAND AND AUDIO APPLICATIONS

In voiceband and audio applications such as hands-free car telephone kits and modems, analog front ends (AFEs), also called codecs, make excellent system building blocks along with DSPs.

Analog Devices recently announced the ADSP-21ESP202-series specifically designed for embedded speech processing applications such as hands-free car kits. This part is based on the AD73322 dual AFE and the ADSP-218x 16-bit fixed point core. 40KBytes of on-chip RAM is configured as 8K x 24-bit program memory RAM, and 8K x 24-bit data memory RAM. 24KBytes of on-chip ROM is configured as 8K x 24-bit program memory ROM.

A full-duplex hands-free car kit can be implemented with the ADSP-21ESP202 as shown in Figure 9.24. This series represents a breakthrough in both the level and diversity of functional integration.

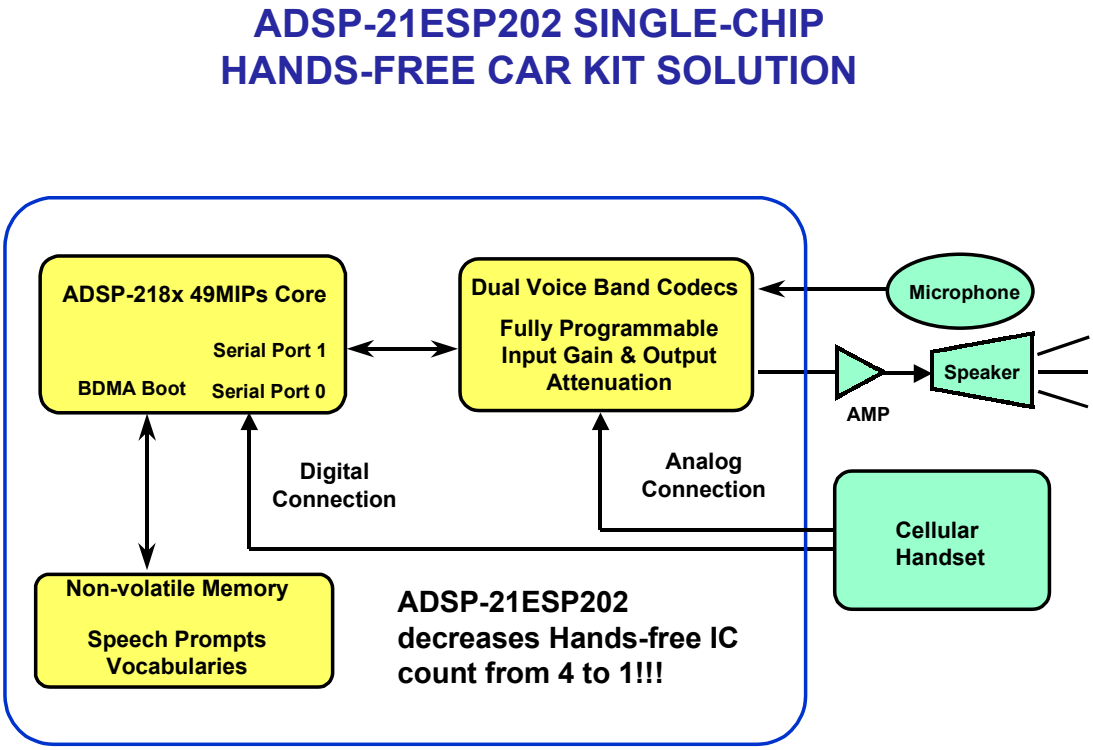

Figure 9.24

Specifically, the ADSP-21ESP202-series is the first general purpose ADSP-218x-based product to include analog functionality (see Figure 9.25). All members of the series contain two 16-bit sigma-delta voiceband codecs which are fully programmable in sampling rate (up to 64kSPS), input gain, and output gain. There are two analog comparators which can be used to implement voice detection algorithms and detect cable insertion/de-insertion events as well as generate processor interrupts. Dual current sources with switches are included which can be

configured to be driven by the PWM functionality of the enhanced timer. They have both trickle and rapid charge settings. The current switches can also be used in automatic gain control (AGC) algorithms and also for input clock frequency/phase adjustments.

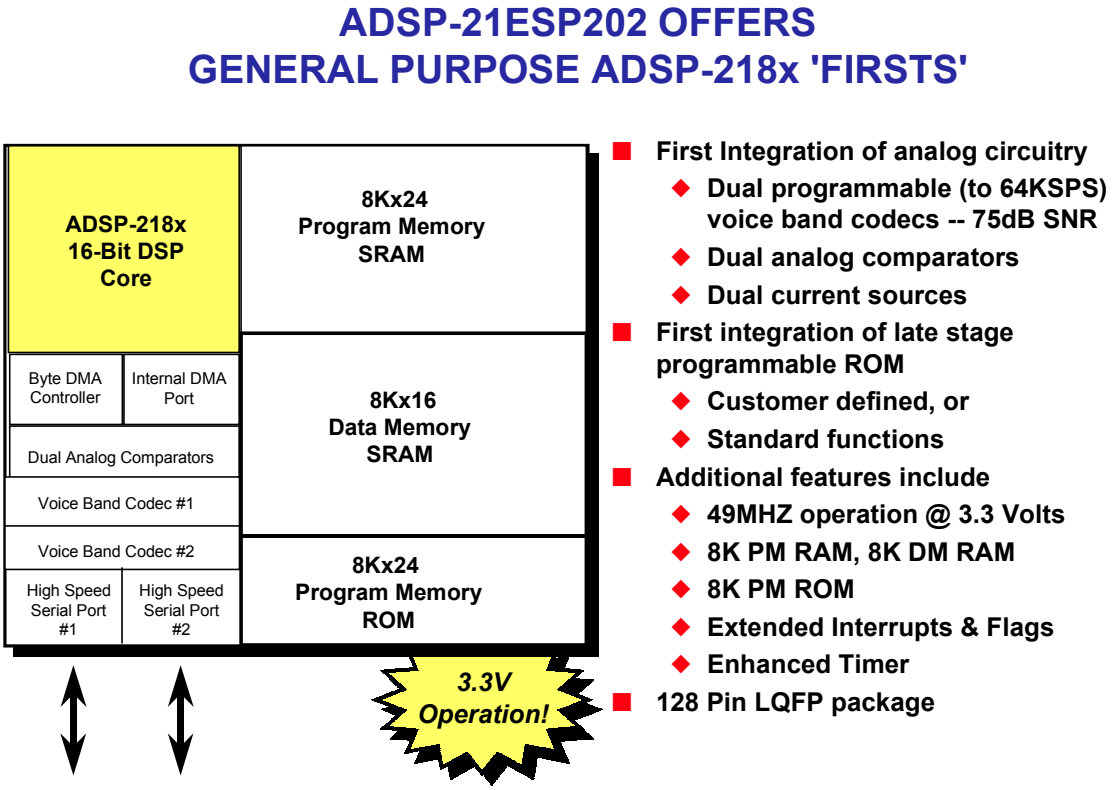

Figure 9.25

The ADSP-21ESP202 also contains a programmable block of 8K words program memory ROM. Analog Devices will offer multiple versions of the series with application specific algorithms embedded into this space. Analog Devices will also accept customer defined ROM codes.

The system shown in Figure 9.24 provides a DSP engine for noise/echo cancellation and speech recognition, codecs to interface to the speaker/microphone and cellular terminal, and non-volatile memory for program storage and speech recognition vocabularies and prompts. The ADSP-21ESP202 series integrates all these components in a single package providing a single-chip solution for full-duplex hands-free car kits. This represents a 75% reduction in IC count over the previous generation solutions.

PC applications such as voice processing and modems also require high performance codecs. Figure 9.26 shows a generalized audio/modem application based on the AD1819B SoundPort® codec.

This codec is fully compliant with the AC'97 specification (Audio Codec '97, Component Specification, Revision 1.03, © 1996, Intel Corporation). In addition, the AD1819 supports multiple codec configurations (up to three per AC Link), a DSP

serial mode, variable sample rates, modem sample rates and filtering, and built-in Phat™ Stereo 3D enhancement.

GENERALIZED PC AUDIO, MODEM APPLICATION

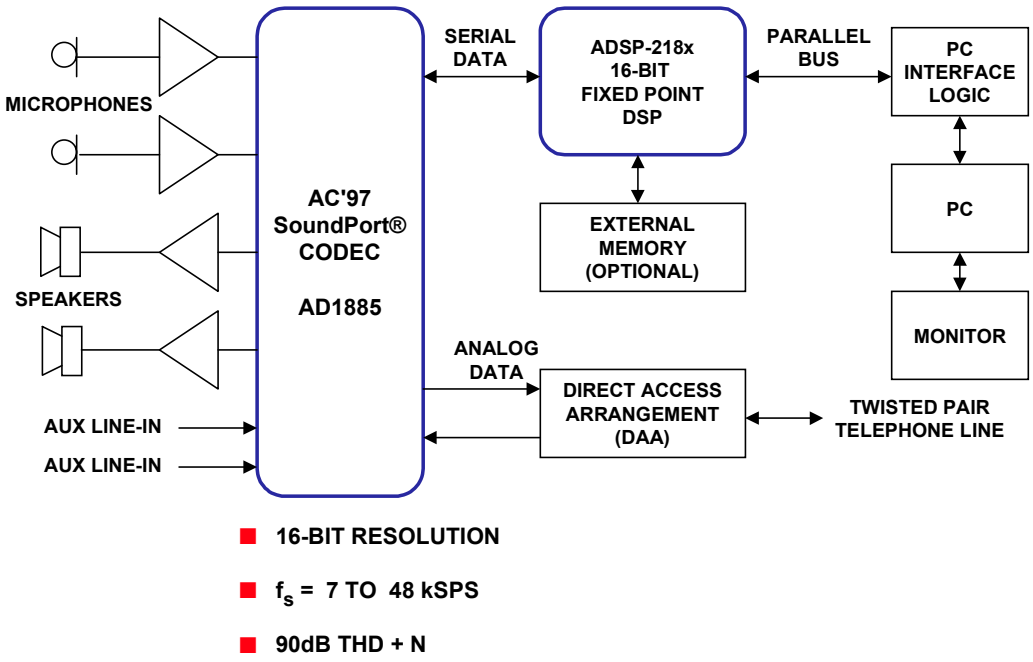

Figure 9.26

The AD1819B is an analog front end for high performance PC audio, modem, or DSP applications. The main architectural features of the AD1819B are the high quality analog mixer section, two channels of 16-bit sigma-delta A/D conversion, two channels of 16-bit sigma-delta D/A conversion, and a serial port. THD + N is 90dB, and the sampling frequency is variable from 7kSPS to 48kSPS.

The Analog Devices' 32-bit floating point SHARC DSP has demonstrated the highest Dolby Digital AC-3 audio quality. The DSP architecture of the reference design shown in Figure 9.27 uses the ADSP-21065L SHARC and an integrated mixed signal IC (AD1836) to provide low cost, high quality, multi-channel audio. Key applications include A/V receivers for home theater and high-end automotive audio. The AD1836 provides all the mixed signal functions including 4 ADC input channels and 6 DAC output channels. The AD1836 codec provides 97dB THD + N and 105dB SNR for high quality audio. While a fixed-function DSP could be used, a programmable DSP provides greater flexibility. This reference design can be programmed for MP3, Dolby Digital AC-3, THX, or DTS decoding. Other audio processing algorithms can easily be added through additional programs.

AUTOMOTIVE AND HOME THEATER REFERENCE DESIGN USING 32-BIT SHARC

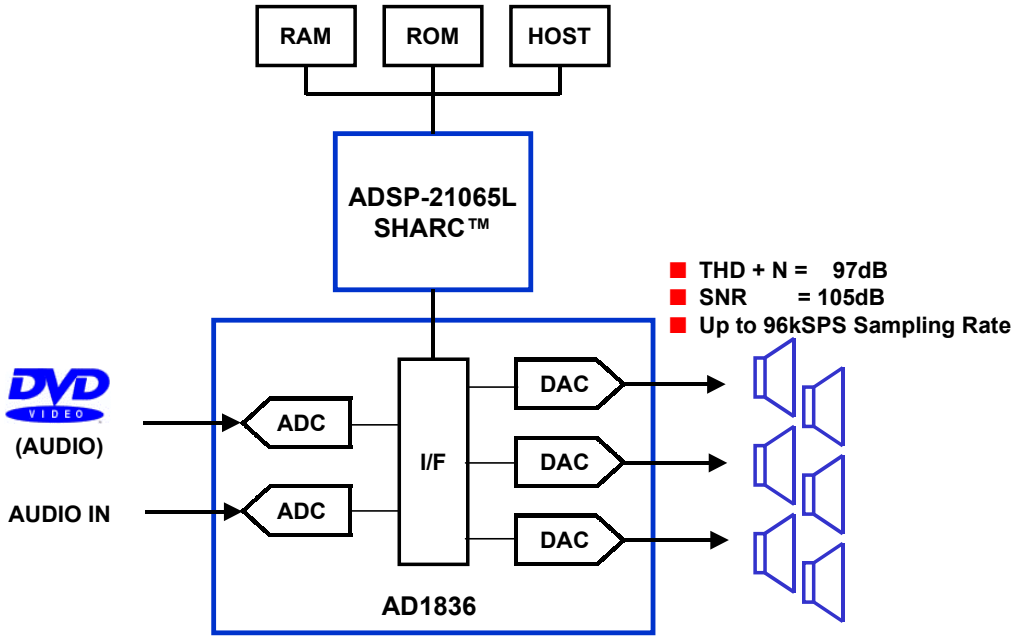

Figure 9.27

In large digital audio systems it is often necessary to distribute the signal processing tasks between multiple processors. Figure 9.28 shows a 16-channel mixer that uses two ADSP-21160s. The data comes from sixteen 24-bit ADCs and is passed to the FPGA. This FPGA converts the serial ADC data into parallel data and sends it to the two ADSP-21160's external ports. The external port on each DSP has hardware support to enable the FPGA to write to both DSPs at once. DMA engines inside the DSP receive this data and move it to desired locations in internal memory. The hardware support and DMA engines reduce the complexity of the FPGA design, because the FPGA is simply required to drive the data onto the bus. No arbitration logic or address generation is necessary in the FPGA.

The DSPs perform various algorithms, such as mixing, left/right panning, equalization, and additional effects processing, such as reverberation or compression/expansion. The audio output samples are then transmitted to a 24-bit stereo DAC. These algorithms could be mapped in such a way that one of the DSPs does the mixing and effects, while the other does the equalization. Another possibility is to give each DSP the task of processing half of the data channels. The optimal algorithmic mapping depends on what type of processing is required.

For this example, it is estimated that the two ADSP-21160s have enough computation power to perform various algorithms on 16 channels of data sampled at 48kSPS. In a 20µs period each DSP is capable of executing 2000 instructions in the core. If each DSP is responsible for half of the channels of data (8 channels), that means that the DSP can allocate 250 instructions to each channel.

With the ADSP-21160 SIMD architecture, there are roughly enough instructions for a 3-band (high, mid, and low) equalizer, mixing, delay effects, and compression for each channel. The movement of data into the memory does not enter into the processing loading calculation since it is a zero-overhead task.

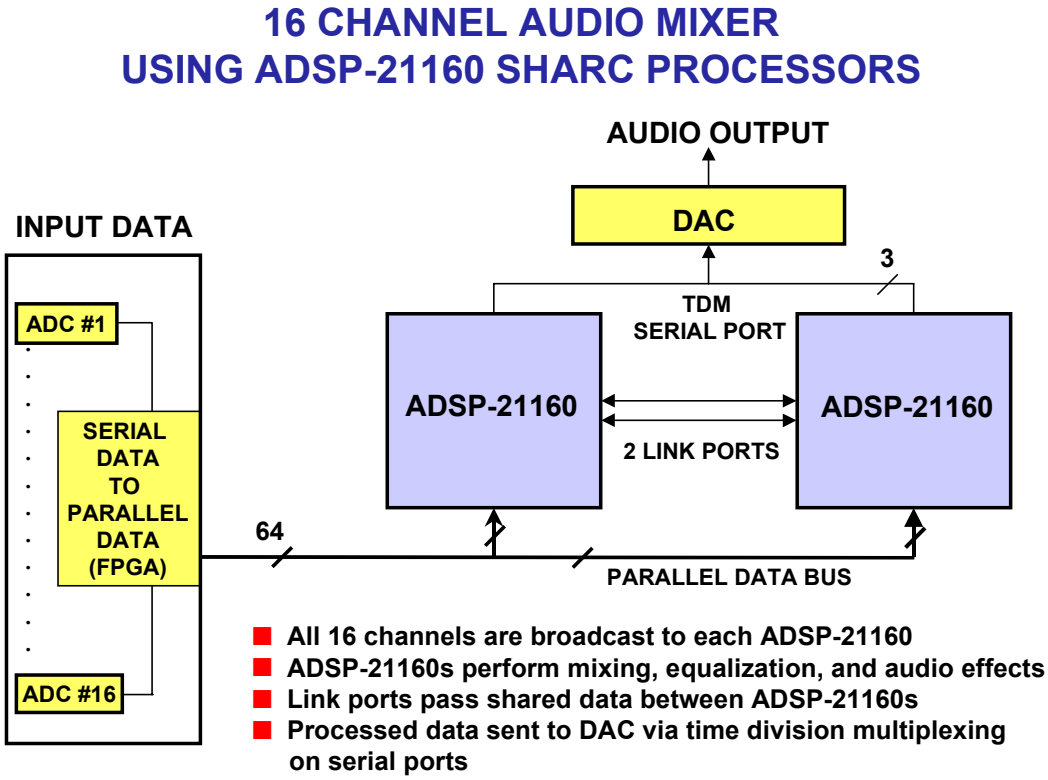

Figure 9.28

A SIGMA-DELTA ADC WITH PROGRAMMABLE DIGITAL FILTER

Most sigma-delta ADCs have a fixed internal digital filter. The filter's cutoff frequency (and the ADC output data rate) scales with the master clock frequency. The AD7725 is a 16-bit sigma-delta ADC with a programmable internal digital filter. The block diagram in Figure 9.29 shows that the modulator operates at a maximum oversampling rate of 19.2MSPS. The modulator is followed by a preset FIR filter which decimates the modulator output by a factor of 8, yielding an output data rate of 2.4MSPS. The response of the preset FIR filter is shown in Figure 9.29. The output of the preset filter drives a programmable digital filter. The diagram shows typical response for a 300kHz lowpass FIR filter.

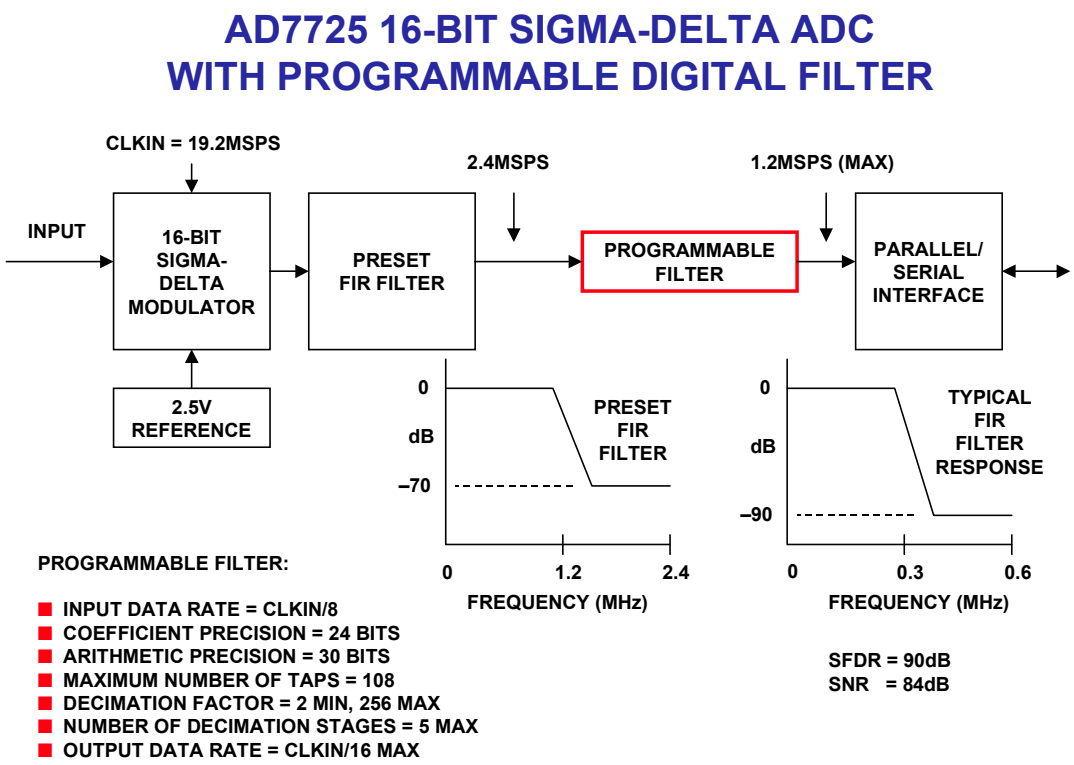

Figure 9.29

The programmable filter is flexible with respect to number of taps and decimation rate. The filter can have up to 108 taps, up to 5 decimation stages, and a decimation factor between 2 and 256. Coefficient precision is 24-bits, and arithmetic precision is 30-bits.

The AD7725 contains Systolix's PulseDSP™ (trademark of Systolix) post processor which permits the filter characteristics to be programmed through the parallel or serial microprocessor interface.

The post processor is a fully programmable core which provides processing power of up to 130 million multiply-accumulates (MAC) per second. To program the post processor, the user must produce a configuration file which contains the programming data for the filter function. This file is generated by the FilterWizard compiler which is available from Analog Devices. The AD7725 compiler accepts filter coefficient data as an input and automatically generates the required device programming data.

The coefficient file for the filter response can be generated using a digital filter design package such as Systolix FilterExpress™ (http://www.systolix.co.uk) or QEDesign™ from Momentum Data Systems (http://www.mds.com). The response of the filter can be plotted so the user knows the response before generating the filter coefficients. The data is available to the processor at a 2.4MSPS rate. When decimation is employed in a multistage filter, the first filter will be operated at 2.4MSPS, and the user can then decimate between stages. The number of taps which can be contained in the processor is 108. Therefore, a single filter with 108

taps can be generated, or a multistage filter can be designed whereby the total number of taps adds up to 108. The filter characteristic can be lowpass, highpass, bandstop, or bandpass and can be either FIR or IIR.

The AD7725 operates on a single +5V supply, has an on-chip 2.5V reference, and is packaged in a 44-pin PQFP. Power dissipation is approximately 350mW when operating at full power. A half-power mode is available with a master clock frequency of 10MSPS maximum. Power consumption in the standby mode is 200mW maximum.

SUMMARY

More examples of DSP applications are summarized in Figure 9.30. There are many other practical applications of DSP in today's rapidly expanding industrial, communication, medical, military, and consumer markets. A discussion of each would constitute a book in itself. It has been the purpose of this section to highlight just a few of the more common applications and give the reader a sense of how DSP affects practically every aspect of modern life.

OTHER DSP APPLICATIONS

- Hands-Free Car Telephone Kits
- Digital Answering Machines
- Voice Recognition Systems
- Cable Modems
- Computer Sound Cards
- Digital Audio: Professional and Consumer
- Digital Video Signal Processing
- High Definition Television (HDTV)
- Computer Generated Graphics
- Digital Special Effects
- Direct Broadcast Satellite (DBS)
- Global Positioning Systems (GPS)
- Medical: Ultrasound, MRI, CT Scanners
- Military: Radar, Missile Guidance

Figure 9.30

REFERENCES

1. John Bingham, **The Theory and Practice of Modem Design**, John-Wiley, 1988.

2. **Digital Signal Processing Applications Using the ADSP-2100 Family, Vol. 1 and Vol. 2**, Analog Devices, Free download at: http://www.analog.com

3. George Calhoun, **Digital Cellular Radio**, Artech House, Norwood MA, 1988.

4. Nick Morley, *New Cellular Scheme Muscles In*, **EDN News**, November 15, 1990, p. 1.

5. P. Vary, K. Hellwig, R. Hofmann, R. J. Sluyter, C. Galand, M. Rosso, *Speech Codec for the European Mobile Radio System,* **ICASSP 1988 Proceedings,** p. 227.

6. D. K. Freeman, G. Cosier, C. B. Southcott, I. Boyd, *The Voice Activity Detector for the Pan-European Digital Cellular System,* **ICASSP 1989 Proceedings,** p. 369.

7. Brad Brannon, *Using Wide Dynamic Range Converters for Wide Band Radios*, **RF Design**, May 1995, pp.50-65.

8. Joe Mitola, *The Software Radio Architecture*, **IEEE Communications Magazine**, Vol. 33, No.5, May 1995, pp. 26-38.

9. Jeffery Wepman, *Analog-to-Digital Converters and Their Applications in Radio Receivers*, **IEEE Communications Magazine**, Vol. 33, No.5, May 1995, pp. 39-45.

10. Rupert Baines, *The DSP Bottleneck*, **IEEE Communications Magazine**, Vol. 33, No.5, May 1995, pp. 46-54.

11. Brad Brannon, *Overcoming Converter Nonlinearities with Dither*, **Application Note AN-410**, Analog Devices, 1995.

12. P.J.M. Coussens, et al, *Three Phase Measurements with Vector Rotation Blocks in Mains and Motion Control*, **PCIM Conference, Intelligent Motion, April 1992 Proceedings**, Available from Analog Devices.

13. Dennis Fu, *Digital to Synchro and Resolver Conversion with the AC Vector Processor AD2S100*, Available from Analog Devices.

14. Dennis Fu, *Circuit Applications of the AD2S90 Resolver-to-Digital Converter*, AN-230, Analog Devices.

15. Aengus Murray and P. Kettle, *Towards a Single Chip DSP Based Motor Control Solution*, **Proceedings PCIM - Intelligent Motion**, May 1996, Nurnberg Germany, pp. 315-326. Also available at http://www.analog.com.

16. D. J. Lucey, P. J. Roche, M. B. Harrington, and J. R. Scannell, *Comparison of Various Space Vector Modulation Strategies*, **Proceedings Irish DSP and Control Colloquium**, July 1994, Dublin, Ireland, pp. 169-175.

17. Niall Lyne, *ADCs Lend Flexibility to Vector Motor Control Applications*, **Electronic Design**, May 1, 1998, pp. 93-100.

18. **Motor Control Products, Application Notes, and Tools**, CDROM, Analog Devices, Spring 2000.

19. **Application of Digital Signal Processing in Motion Control Seminar**, Analog Devices, Spring 2000.

20. Analog Devices' Motor Control Website: http://www.analog.com/motorcontrol