A Portable Multimedia Terminal

Successful personal communications terminals will depend upon the smooth integration of computation and communications facilities in a lightweight unit.

Samuel Sheng, Anantha Chandrakasan, and R.W. Brodersen

The personal communications industry has seen explosive growth in the past several years, especially in the number and types of services and technologies. In voiceband communications, systems such as mobile analog cellular telephony, radio pagers, and cordless telephones have become commonplace, despite their limited nature and sometimes poor quality of transmission. In computing, portable "notebook" computers are boasting capabilities far in excess of the desktop machines of five years ago, and multi-MIPS, RISC-based, portable workstations are available. Despite the myriad technologies to be had, however, little integration of these diverse services — the combination of computation and communications facilities in a portable unit — has occurred. Thus, our vision of a future personal communications system (PCS) centers on such integration of services to provide ubiquitous access to data and communications via a specialized, wireless multimedia terminal.

A schematic view of such a system is shown in Fig. 1. The wireless terminal is in full duplex communication with a networked base station, which serves as a gateway between the wired and wireless mediums. Via a base station, users access services over the high-speed communications backbone, including communicating with another person also linked into the network. This idea can also be extended to a user communicating not only with another person, but with network "servers." Because the data bandwidth of future fiber-optic networks is easily in excess of 10 Gb/s, these centralized servers can provide a wide variety of information services to users. A personal communications system will likely include the four key features that follow.

Access to large commercial databases that contain information such as international and domestic news, financial information, traffic data, transportation schedules, voice mail, telephone numbers, news, bulletin boards, and educational material is necessary. The continuous connectivity afforded by personal communications systems has several advantages. Many sources of information are of a transitory nature, such as stock pricing, local news, and so on, making distribution by other means such as CD-ROM impractical. Furthermore, given sufficiently large database servers, libraries of books, journal archives, and other currently "paper-intensive" media can be placed on-line; these databases would allow for instantaneous recovery of all types of information, without the need to be at a terminal physically attached to the wired network.

Second, a PCS would have access to digital video databases containing both entertainment and educational media, such as animated information sequences, taped lectures, movies, news clips, and other isochronous data. Unlike today's television broadcasts, video databases can be made available on-demand, giving users the freedom to access video information as needed. Video data will be necessarily stored in a compressed format for minimization of both storage space and transmission bandwidth, thus requiring that the wireless terminals at least support video decompression.

Simplified entry mechanisms such as voice-recognition and handwriting-recognition interfacing to access the above functions would also be available. The design of an effective user interface to access such a vast information storehouse is a critical issue. By using speech recognition and pen-based input, supported by large, speaker-independent recognizers placed on the network, such interfacing and information access can be tremendously simplified. Placing the recognition units on the network conserves power in portable units and enables much larger and more complex recognition algorithms to be employed. Recognizer servers can also make use of context-sensitive analysis, which can increase recognition accuracy by determining which words are most likely to be used in a given application [1].

Fourth, the system would provide support for a distributed computing environment, such as MIT's X-Window system. In distributed computing environments, computation need not take place on a local machine; instead, computation is performed by programs executing on one or more remote machines, which may have no computing capability except that required to act as an intelligent display device. Many such inexpensive "X-terminals" already exist. Unlike

ROBERT W. BRODERSEN is a member of the Electrical Engineering and Computer Science faculty of the University of California, Berkeley.

ANANTHA P. CHANDRAKASAN and SAMUEL SHENG are Ph.D. candidates in the Electrical Engineering and Computer Science department at the University of California, Berkeley.

0163-6804/92/$03.00 1992© IEEE

TTY terminals that can only communicate with a single host machine, X-terminals possess all of the necessary networking capability to communicate with as many remote servers as needed. The multimedia terminal will be based on this model: remote computation servers will be used to run applications like spreadsheets, word processors, etc., with the results being transmitted to the terminal. Likewise, supercomputer-class servers will perform intensive tasks requiring simulations, 3-D image rendering, and computer-aided design.

Clearly, the cornerstone of the entire system lies in the ability of multimedia terminals and wireless communications links to support all of the aforementioned services. Correspondingly, it is this desire for portability that translates directly into design constraints on the size and weight of the terminals, the power they consume, and the frequency bandwidth needed in the wireless links.

A diagram of a portable terminal is shown in Fig. 2. To minimize power, only those functions that are absolutely necessary are implemented: the analog RF transceiver; baseband processing for communications, such as equalization, coding, and packetization; the image decompression unit and a display driver; and the speech codec. Because size is also an issue, a pen input system is integrated directly onto a compact flat-panel display, eliminating the need for a large keyboard and providing greater visual feedback than possible with mouse- or trackball-based interfaces. More important, the system is asymmetric in nature: High-quality, full-motion video is only supported in the downlink from the base station to the portable. This must be accounted for the in the design, as the bandwidth requirements in the reverse link from the portable unit are thus considerably less than the link from the base station. For video teleconferencing, a low-rate, reduced-quality video uplink might also be supported; however, the asymmetric bandwidth requirements will still remain.

The diagram shows that no direct user computation is supported within the portable itself; instead, it is wholly dependent on the network servers to provide desired functionality. Although this has immediate benefits in terms of reducing power consumption, it provides another advantage: Data that is highly sensitive to corruption will not be transmitted over the wireless network.

Existing distributed computation environments are dependent on the fact that data transmitted over the network has high integrity — bit-error rates on wired Ethernet are typically on the order of 1 in 10^{12} bits, and further protection is gained by packet retransmission after errors. On wireless networks, however, this is not true. Even after extensive application of error-correction coding, it is still difficult to attain error rates even remotely as low as this. User "computation" data, such as spreadsheets or simulation results, simply cannot be allowed to sustain any corruption. For wireless systems, this translates into an inordinate amount of transmission overhead in terms of coding and data retransmission. On the other hand, user "multimedia" information, such as voice and image data, is relatively tolerant of bit errors — an error in a single video frame or an audio sample will not significantly change the meaning or usefulness of the data. The ability to coexist with an error-prone transmission environment has tremendous impact on the overall system design.

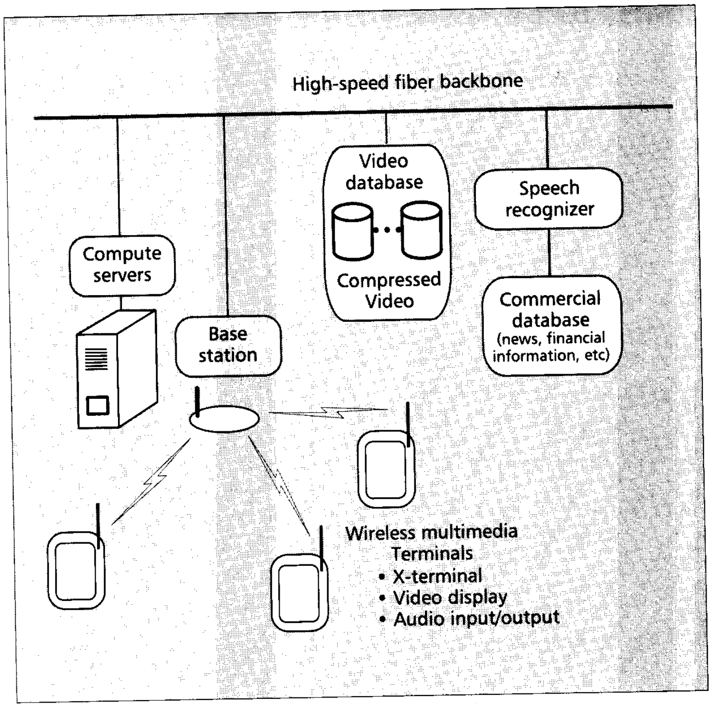

■ **Figure 1.** *Personal communications system overview.*

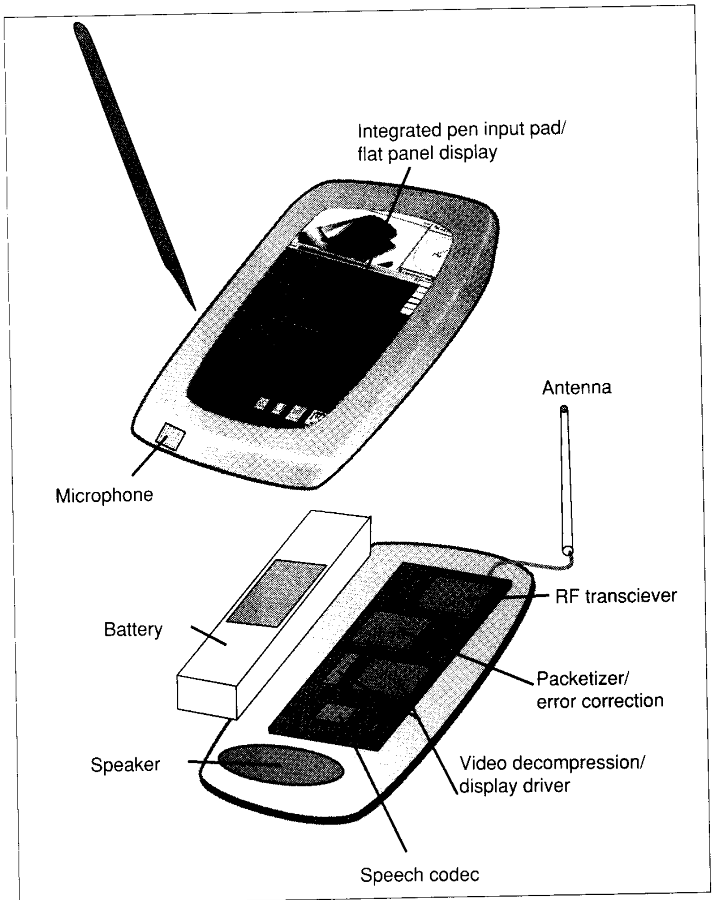

■ **Figure 2.** *Diagram of the portable wireless multimedia terminal.*

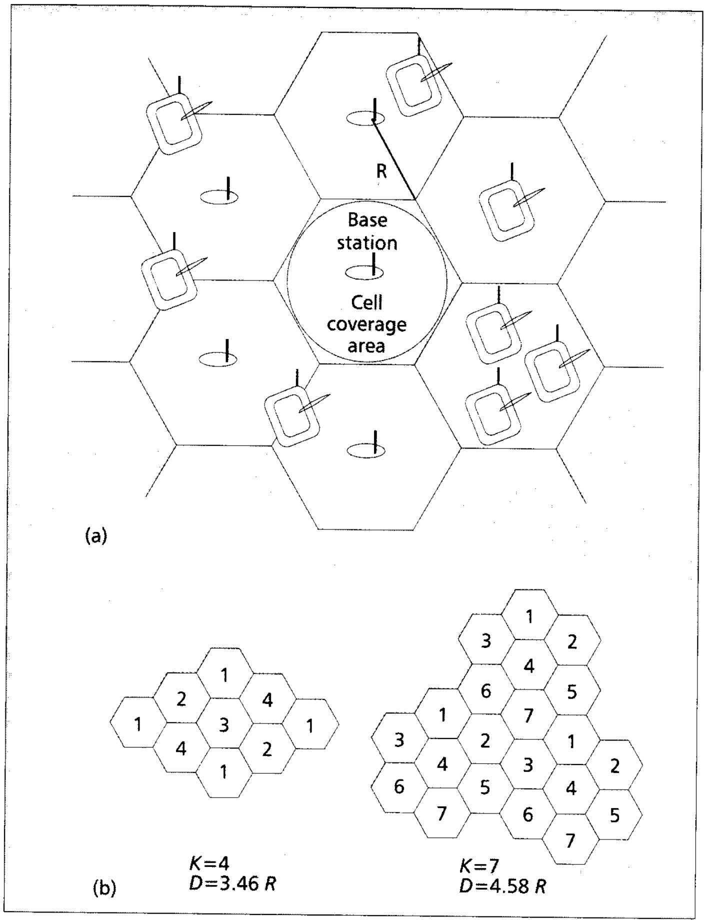

■ **Figure 3.** *(a) Cellular communications system; (b) typical frequency reuse patterns.*

Thus, the portable unit described above is truly a terminal dedicated to multimedia personal communications, and not simply a notebook computer with a wireless LAN/modem attached to it.

The remainder of this article will focus on several of the major design issues behind portable multimedia terminals: spectrally efficient picocellular networking, low-power digital design, video data compression, and integrated wireless RF transceivers. Optimizing performance in each of these areas is crucial in meeting the performance requirements of the overall system and providing a small, lightweight terminal for personal communications.

Picocell Networking and Spectrum Usage

Because full-motion digital video is to be supported, spectrum usage is of great concern — per-user data rates can easily exceed 2 Mb/s even with the best compression schemes reported to date. This data rate is not needed on a continuous basis; regular computation tasks such as word processing or use of a spreadsheet require only slight screen changes on a frame-by-frame basis over a small region, usually on the order of a single character or a few pixels. Hence, it is probable that the peak data rate required by users will be much larger than the overall time-average data rate. Minimizing overall system bandwidth consumption while supporting a large number of users accessing data simultaneously is paramount.

One method to achieve this goal is to physically reduce peak user data rates via data compression techniques (we will discuss this later). Another technique, applied at the system level, is to utilize cellular networking techniques to achieve spatial frequency reuse. Because such a personal communications system will first be used as a step beyond conventional wireless LANs, an indoor picocellular transmission environment will be of primary concern. (The techniques described in this article are applicable to both indoor and outdoor environments.)

The advantages in improved spectral efficiency afforded by cellular systems have been employed extensively in present-day analog mobile radiotelephony, where large-scale cells exploit these advantages to a limited extent. By scaling down cell sizes, tremendous increases in spectral efficiency can be achieved. A simple cellular scheme, as shown in Fig. 3a, consists of dividing the entire service area for the personal communication system into "cells" of radius R, with a single base station serving all mobile users within that cell. Each cell uses its own distinct set of frequencies. As users move from cell to cell, their transactions with the network are "handed off" from base station to base station, reconfiguring the network dynamically as the need arises.

The key benefit of cellular systems is that they allow the network to achieve spatial multiple access. If two cells are separated by sufficient distance, each can use the same frequency bands at the same time without resulting in disastrous cochannel interference. Thus, extensive frequency reuse becomes possible, as opposed to an umbrella scheme where every user must be assigned a different frequency slot. Fig. 3b shows several classical reuse patterns [2]; such patterns are typically characterized by a frequency reuse factor K, which represents the number of distinct frequency sets that need to be used to cover the entire service area. Instead of one user per frequency band, the network can now support N users per band, where N is the number of cells in the service area using that band. From the point of view of spectrum usage, each user effectively consumes only B/N Hz of bandwidth, where B is the physical bandwidth needed to support transmission, thus drastically increasing overall spectral efficiency.

Clearly, minimizing the physical distance D between cells using the same frequency, by reducing the cell size R, yields the greatest frequency reuse, and hence the greatest gains in efficiency.[1] Therefore, it is clear that the number of users supportable within the same overall system bandwidth increases quadratically as R decreases, because of the increased number of cellular subdivisions within the service area. Minimizing R (and hence D) is critical in achieving high levels of spectral efficiency.

With an indoor environment, it is no longer feasible to have only a single network transceiver station serving all of the terminals in the building. Due to the 5 to 15 dB attenuation through walls, the total microwave output power from all of the transmitters would have to be inordinately (and dangerously) high [3]. However, this attenuation can be taken advantage of by a cellular network — each room naturally becomes its own cell. Likewise, the cellular

---

[1] *The frequency reuse distance is geometrically related to K and R by* $D = R\sqrt{3K}$

scheme now moves into three dimensions because the floors also provide RF isolation. Even if walls are not present, the use of electrically-adjustable directional antennas (such as a small phased-array device) can provide the same effect. The cells are now extremely small, on the order of about five meters; R is usually dictated by the size of the room, and K can be as low as 3 to 4, depending on how much attenuation is provided by the walls. If K is increased to 6 or 7, the assumption that cochannel interference is negligible becomes reasonable for most indoor office environments.

In light of the above considerations, the total amount of spectrum that will be consumed to provide the outlined services can now be addressed. After examination of the user density in a typical office environment,[2] such as those found in modern buildings with open-area soft-partition cubicles, cells with a radius of five meters typically contain 12 to 16 active users. In the worst-case, a 2 Mb/s data rate for full-motion video using a linear DQPSK (differential quadrature phase-shift keying) modulation scheme and design parameters from existing systems [4-5], would require a transmission bandwidth of approximately 1.5 MHz per user using a 20 percent excess bandwidth raised-cosine pulse shape and a 25 percent loss for packetization, equalizer training, and other overhead. Assuming that half of the 16 users in the cell demand the complete 2 Mb/s data rate for full-motion digital video, and the remainder need 256 kb/s each[3] for lower data rate applications, a picocellular system with $K = 7$ would use approximately 100 MHz of bandwidth. Although 100 MHz is a considerable amount of spectrum, this is amortized over large numbers of people using this spectrum simultaneously within multiple buildings. Because the bandwidth of 100 MHz is designed to support full motion video and other multimedia network services for all users, this allocation of spectrum is not unreasonable, given the level of service provided by the system, especially when compared to the spectrum allocated for existing systems such as NTSC television.

There is another significant advantage to picocellular wireless systems: because transmit power is scaled down as cells move closer together to reduce interference, the power consumed in the portable's transmitter to drive the antenna is correspondingly reduced. Whereas existing cellular systems use 1 watt of transmit power for voiceband RF links in 5 mile cells, a picocellular system with 5 meter cells requires only milliwatts to maintain the link [6].

Implementing Portable Terminals

Picocellular networking ameliorates several important issues in providing portable multimedia-based communications systems. Many challenges remain, however, in building the required functionality into terminal hardware.

A signal flow diagram for a terminal is shown in Fig. 4. Incoming data can be one of three types: digital video, screen graphics, or sampled audio. Outgoing data can be either voice or pen input. The asymmetry in the data rates of the uplink and the downlink is clear; no high-rate signals are intended for transmission from the mobile to the base station. The hardware design must also reflect this asymmetry: within the analog RF block, the receiver design becomes critical, because it must demodulate a high-rate signal corrupted by noise and distortion, whereas the transmitter is relatively simple, with low data rate and output power requirements. Likewise, the algorithm chosen for the video decompression should be designed, if possible, to make decompression as simple as possible and with little consideration for compression complexity, because compression can be performed by one of the network servers.

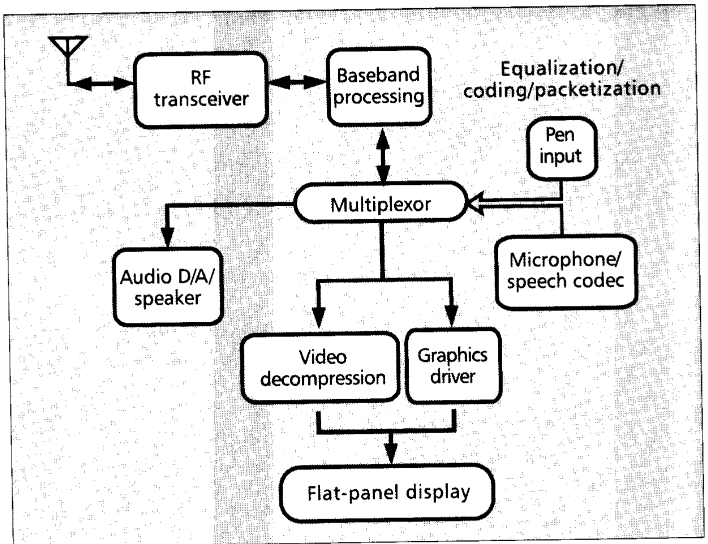

■ **Figure 4.** *Signal flow diagram for the proposed multimedia terminal.*

One key consideration is how long the portable can function between battery rechargings. Ideally, it should be able to operate for one work day, or eight to ten hours of battery life. Given that conventional batteries typically possess 20 watt-hours for each pound of battery weight, and a limit of one pound of batteries in the portable, the entire portable can consume no more than 2 watts of power. Furthermore, projections of progress in battery technologies show that only a 20 percent improvement in battery capacity will occur over the next ten years. Thus, power minimization becomes a serious concern.

The largest power consumer in current portables is the backlighting of flat-panel displays. As display technologies improve screen contrast, however, this requirement will be significantly relaxed, implying that low-power techniques for implementing the analog and digital core circuitry are needed.

With present-day technology — single-chip packaging using printed circuit boards for interconnection — one-third or more of the total power is consumed by the chip's input/output (I/O) because the capacitances at the chip boundaries are usually much larger than the capacitances inside the chip. Typical values range from a few 10s of femtofarads at internal chip nodes, up to 10s of picofarads at the chip interface attributed to pad capacitances and board traces. To reduce the power consumed in the I/O, low-capacitance, high-density interconnect methods will be employed, such as the emerging multichip module (MCM) technologies. MCM integrates many individual chip die into a single structure, reducing the size of interchip capacitances to the same order-of-mag-

---

[2] *For example, the EECS graduate student research facility at the University of California, Berkeley.*

[3] *256 kb/s is typical of the peak data rate afforded by most wireless LAN systems, and is reflective of data rates used in existing X-terminals; this value may be considerably lower depending on the level of user activity.*

DCT Algorithm	Multiplies	Additions
Brute force	4096	4096
Row-col DCT	1024	1024
Chen's algorithm	256	416
Lee's algorithm	192	464
Feig's algorithm	54	462

■ **Table 1.** *Complexity of various DCT algorithms (8 x 8 blocks).*

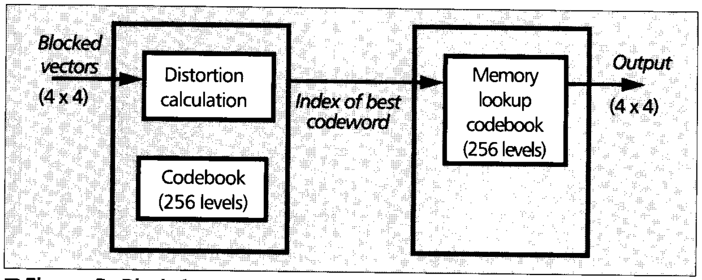

■ **Figure 5.** *Block diagram of a VQ coder.*

nitude as on-chip capacitances and minimizing the power consumed in the I/O drivers. Thus, with MCM, the majority of the power is consumed within the functional core of the chip itself, as opposed to the interface. Also, because the packing density has increased, and with the ever-decreasing size of CMOS circuitry (down to 0.2 μm line widths), over 10^{10} transistors can be placed within a single eight-by-eleven-inch MCM substrate. Area constraints imposed by available silicon are no longer of great issue, allowing for greater possibilities in power optimization, as we will discuss later.

For analog RF transceivers, however, there are other design considerations beyond low-power implementation. Due to size considerations, traditional discrete element design is not feasible for a small, portable unit such as the proposed multimedia terminal; single-chip integration techniques that exploit advances in silicon CMOS (as opposed to gallium arsenide, GaAs) must be explored, to address cost and manufacturability concerns. Likewise, the fact that digital circuitry is readily available on-chip also opens up new possibilities: analog performance requirements can be reduced at the expense of increased digital signal processing.

To examine these implementation details more fully, we will discuss three distinct design issues. The first is image compression, both as a means of spectrum reduction and as an example of how the choice of algorithms can take advantage of the aforementioned asymmetry to reduce power consumption. Second, we will show how low-power digital system design leads to large reductions in power consumption. Third, we will analyze the design of analog RF transceiver to exploit monolithic integration techniques and the underlying digital nature of the transmitted signal.

Image Compression

As stated above, frequency reuse is only one means of reducing spectrum consumption. Concomitant with frequency reuse is the idea of reducing the amount of physical bandwidth needed by each user, and hence reducing the bandwidth consumed by the overall system. Because high-resolution, full-motion video is to be supported, transmission of 640-by-480 pixel images, digitized at 24 bits/pixel, would require a bandwidth of 220 Mb/s in an uncompressed format at 30 frames per second. Thus, it is clear that video compression techniques are crucial in making wireless video transmission feasible.

The video module performs the decoding and display interface functions and converts a compressed data stream to an uncompressed video stream that is displayed on the LCD display. The decompression module can be implemented using a variety of algorithms, such as transform-based schemes, vector quantization and subband coding. The selection of the algorithm for the portable terminal depends not only on the traditional criteria of achievable compression ratio and the quality of reconstructed images, but also on computational complexity (and hence power) and robustness to higher bit error rates. The choice of an algorithm to implement the decompression function is the most important in meeting the power constraints. The basic complexity of the computation must be optimized and, as shown in the next section, the ability to parallelize an algorithm will be critical.

Most current compression standards (for example, JPEG and MPEG) are based upon the Discrete Cosine Transform (DCT). The basic idea in intraframe schemes such as JPEG is to apply a two-dimensional DCT on a blocked image (typically eight pixels by eight pixels) followed by quantization to remove correlations within a given frame. In the transform domain, most of the image energy is packed into only a few of the resulting coefficients, and compression is achieved by transmitting only a carefully chosen subset of the coefficients. While these standards specify the use of DCT as the transformation to be applied, they do not specify the algorithm to be used. Table 1 shows a comparison of a few algorithms that can be used to implement the DCT [7-8]. Minimizing the operation count is important in minimizing the switching events and the hence the power consumption.

A primary characteristic of the DCT is the symmetric nature of the computation; that is, the coder and decoder have equal computational complexity. However, an alternative compression scheme is that of vector quantization (VQ) coding, which is asymmetrical in nature and has been unpopular due to its complex coder requirements. The basic idea behind VQ coding is to group the image data into a vector and quantize it. Fig. 5 shows a block diagram of a VQ coder. On the coder side, the input is first blocked into a vector (for example, a four-by-four block of video is 16 bytes). This vector is compared to the entries in the codebook with the goal of minimizing the expected error or distortion between the input vector and its reproduction for a given bit rate. The codeword corresponding to the closest match (or the match that minimizes the distortion) is transmitted (in this case, the index is one byte long and 16:1 compression is achieved). On the decoder side, a simple memory look-up is used to reproduce the image data. The design of the codebook and distortion measure have been extensively discussed in literature. Clearly, the distortion computation is much more computationally intensive than the decoder, as the entire decoder

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts

Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research

With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips

Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.

WHAT WILL YOU BUILD? | **sales@docketalarm.com** | **1-866-77-FASTCASE**