Technology Directions for Portable Computers

ERIK P. HARRIS, MEMBER, IEEE, STEVEN W. DEPP, MEMBER, IEEE,
WILLIAM E. PENCE, MEMBER, IEEE, SCOTT KIRKPATRICK, SENIOR MEMBER, IEEE,
M. SRI-JAYANTHA, MEMBER, IEEE, AND RONALD R. TROUTMAN, FELLOW, IEEE

Invited Paper

This paper contains an evaluation of trends in the key system parameters (e.g., size, weight, function, performance, battery life) for battery-powered portable computers, together with a review of development trends in the technologies required for such systems. The discussion will focus on notebook-size portable computers. Those technologies which will have substantial impact on battery life and power budgets of future notebook computers will receive the primary emphasis in this paper: for example, liquid crystal displays, storage technology, wireless communication technology, and low power electronics. System power management will also be addressed. The basic theme of this paper is first to develop a view of what the key attributes of future notebook computers will be, and then to discuss how technologies must evolve to allow such systems to be advanced over the current state of the art in terms of portability and battery life.

I. INTRODUCTION

Information systems with reduced power consumption are assuming increased importance in today's environment. This is in part a reflection of an overall trend in the information technology industry toward providing the customer with greater function at smaller size, lower weight, reduced electrical and thermal loads, increased reliability, and lower cost, but the principal driver for the development of low-power technologies has been the advent of battery-powered portable information systems with significant computational and communication capability. It is these portable systems and the technologies required for them which will be the subject of this paper.

The growth of the market for battery-powered portables is now a dominant trend in the information technology business and is expected to remain so for the foreseeable future. Manufacturers of portable systems are working hard to develop smaller, lighter, cheaper portables with enhanced performance, function, and battery life. Increased battery life is seen as an important source of leverage in this intensely competitive marketplace. In this paper, we will present an evaluation of trends in key system parameters (e.g., size, weight, function, performance, power budget, battery life) for battery-powered portable computers, together with a review of development trends in the technologies required for such systems. Those technologies which will have substantial impact on battery life and power budgets of future portable systems will receive the primary emphasis in this paper. The largest power consumer in many portable systems is the display. This is particularly true for backlit liquid-crystal (LC) color displays, and such displays are discussed in detail. Low-power CMOS devices and circuits are discussed in detail elsewhere in this issue, and are considered here primarily to emphasize the enormous leverage which exists to reduce the power consumption of CMOS logic through scaling and voltage reduction. Wireless communication technology will provide important function in future portables and its power implications are discussed in detail. Storage technology is also treated, with particular emphasis on hard disk drives. Input technology is treated briefly, as it is expected to be a relatively small power consumer. There are system integration issues to be addressed when considering how to reduce system power consumption, for example: system power management and dc power supplies. System power management will be treated in this paper. dc power supplies—and, in particular, battery technology—will not be discussed in detail here. For a recent view of battery technologies for portable computers, the reader is referred to the paper by J. F. Freiman [1], and to other articles in the same conference proceedings.

The basic theme of this paper is first to develop a view of what the key attributes of future portable computers will be, and then to discuss how portable technologies must evolve to allow such systems to be significantly advanced over current systems from a portability and battery-life viewpoint. In the next section, we examine the power budgets and other system parameters of A4-format notebook computers and develop a view of how such systems and their power budgets are likely to evolve in coming years. In the following sections, we discuss

Manuscript received June 17, 1994; revised November 14, 1994.
The authors are with IBM Research Division, T. J. Watson Research Center, Yorktown Heights, NY 10598 USA.
IEEE Log Number 9408755.

development trends in key portable technologies with the objective of showing where the principal challenges will lie in improving system portability and battery life while at the same time enhancing performance and functionality. In the final section, we summarize our findings and present conclusions.

II. Portable System Environment and Trends

A. Growing Market for Portable Systems

Portable systems incorporating low-power electronics technology are not new: low-current, low-voltage CMOS technology has been exploited in watches for over two decades [2]. However, battery–powered systems of the types we will be concerned with here first began to appear in the late 1980's. The term "portable system" today covers a broad range of system types, including palmtop and notepad computers, personal digital assistants (PDA's), keyboard-input sub-notebook and notebook computers, and larger laptop and ac-powered computers. Of these, the largest market today is for notebook computers (battery-powered portables which have the letter-size A4 format—29.7 cm by 21.0 cm—and weigh about 3 kg or less). Notebook computers are available with a wide variety of features and capabilities, but many common notebooks from the year-end 1993 timeframe have a 25–33 MHz 486-generation microprocessor, about 4 MB RAM, 120 MB hard disk, 1.44 MB floppy disk, a 10–in monochrome super-twisted nematic (STN) LC display, weigh about 2.5–3 kg, and have a 0.6-kg NiCd or nickel–metal–hydride (NiMH) battery which provides about 3–4 h of battery life in continuous operation. With power–saving features enabled, battery life can be increased by about 25–50% in many applications.

Strong market growth for portable systems is expected to continue through this decade, driven by technology enhancements which will make it possible to capture increasing function and performance in small, highly portable, battery-powered systems with long battery life. Added capabilities such as speech recognition, handwriting recognition, and wireless communications will draw new classes of users to the portable marketplace and enable new applications which go far beyond those available with today's systems. Portability and battery life have been, and will continue to be, key considerations for portable customers.

A schematic view of the likely future directions of portable system evolution is shown in Fig. 1. One direction will be an evolution of systems with roughly today's level of performance (but adding some new capabilities such as wireless communications capability) toward much smaller size and lower power dissipation over time [3], for instance to meet the needs of mobile users of handheld pen–based systems. Already, many handheld systems with significant capability are appearing (see, for example, [4]). A second direction will be the evolution of the notebook computer in a way which preserves the popular A4 format but moves toward steadily increasing function, performance, and battery life over time. The A4 format will remain

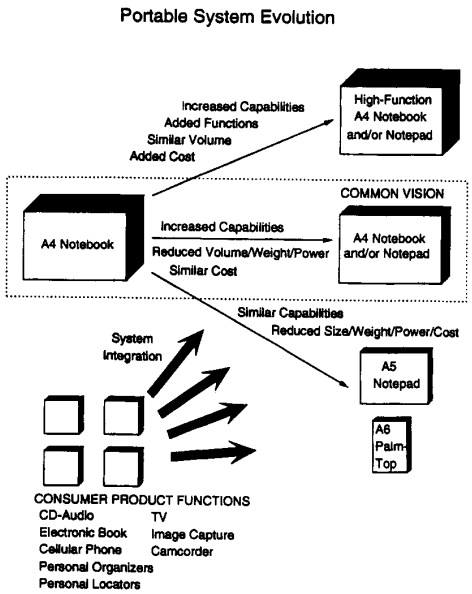

Fig. 1. A schematic view of portable system evolution.

a desirable size, but reduced weight and thickness will be advantageous. The future notebook we envision will have both keyboard and other modes of input (such as handwriting and speech recognition capability) and also will have substantial on-board communications capability. This evolutionary path is represented in Fig. 1 as the "common vision." We turn next to a more detailed look at trends in the key parameters for systems of this type: the A4-format notebook computer.

B. Key Parameters for Notebook Computers

Before discussing future trends, let us consider the power budgets and other key parameters of current notebook computers. Current notebook computers with a 25–33 MHz microprocessor, a hard disk drive, and a monochrome display typically have a 30 Wh NiCd or NiMH battery and dissipate about 8 W of power when running at full speed with the hard disk spinning (at idle) and the LC display on, thereby yielding about 3–4 h of battery life in continuous operation. A representative power budget is shown in Fig. 2. (These power budget numbers are similar to those presented elsewhere in the literature [5].) This figure shows the power consumption for the system and its constituent parts under various operating conditions. Here, "full-on" refers to operation at maximum speed with LC display and hard disk drive (HDD) on and power management disabled. Note that in this case about 4 W of power (50% of the total) is consumed in the CMOS logic and memory (including the video chipset). The "power-managed" case refers to operation on a particular user-input task with power management enabled, which allows about a 25% reduction in power, primarily in the CMOS microprocessor and other logic. This case is similar to what a user could expect in many applications, although the

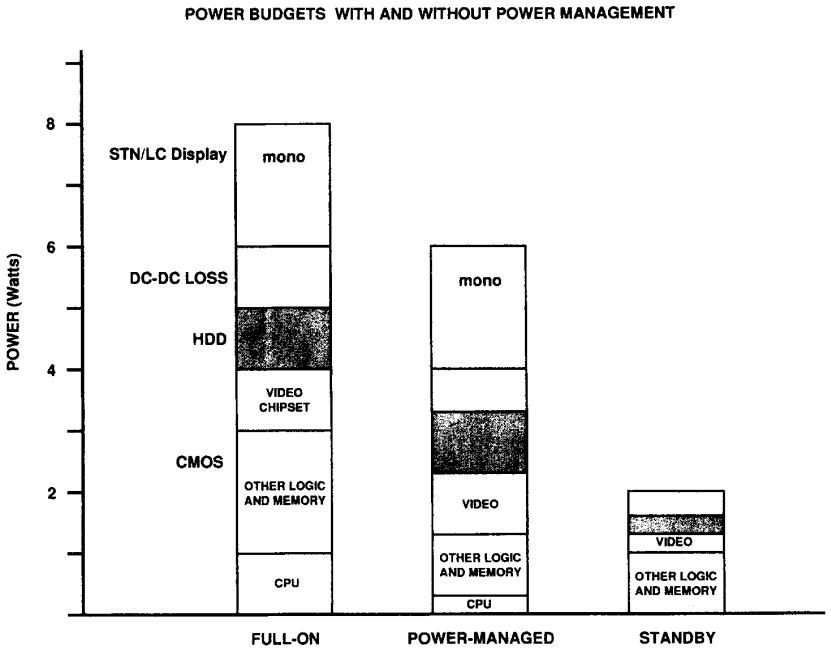

Fig. 2. Typical power budgets for a 25 MHz notebook computer in the 1993 timeframe with and without power management. "Full-On" corresponds to full-speed operation with HDD and LC display on and power management disabled. "Power-Managed" corresponds to normal operation on a particular task with power management enabled and HDD and LC display on. "Standby" corresponds to HDD spun down, LC display turned off, and the CPU chip in a sleep state.

actual power consumption will depend on the application. In the "standby" case, the LC display is off, the HDD is spun down, and the microprocessor is in a sleep state.

The power budget of Fig. 2 is not best-of-breed today. Today's most advanced monochrome notebooks dissipate substantially less power. The trend in notebook computer power dissipation is shown in Fig. 3, where approximate power budgets for systems available at various points in the 1992–1994 timeframe are displayed. In each case the power dissipation displayed corresponds to the "full-on" conditions described above. There are several points to be made about these power budgets. First, systems with active-matrix color displays dissipate significantly more power than their monochrome counterparts in the same timeframe, because the color displays require much more power for backlighting. Second, there have been dramatic reductions in power dissipation over the past few years for both monochrome and color systems. In the case of monochrome systems, this is due primarily to reduction in the power dissipation of CMOS logic and memory which has accompanied the transition from the previous *de facto* industry standard 5 V supply voltage to a 3.3 V supply voltage for many logic and memory parts. CMOS power dissipation has decreased even though processor clock frequency and performance have increased significantly in the same timeframe. For color systems, there have also been substantial reductions in the power consumed by backlighting. In fact, the power dissipation of today's 486-based active-matrix color notebooks are about the same as the 386-based monochrome notebooks of 2–3 years ago. Finally, it should be pointed out that the improvements in battery life which have occurred in this timeframe are attributable almost entirely to these reductions in power dissipation, since battery energy capacities have changed relatively little over this period.

C. Future Notebook Computers

Given the progress which has been made in notebook computer power dissipation over the past few years, it is interesting to consider how these trends may be extended into the future. Our projections of the key parameters for future notebook computers, corresponding to the "common vision" of Fig. 1 and projecting out to about 1997, are contained in Table 1 and also in Fig. 3. A basic assumption of our "common vision" is that to be competitive, future notebook computers will need to combine increased performance and function with thinner, lighter weight packages and longer battery life.

The first column of Table 1 is representative of monochrome notebook computers in the 1993 timeframe, with a power budget similar to that of Fig. 2. Shown in the second column of Table 1 are our projections of what might be possible for a notebook computer in 1997. We have assumed a substantial enhancement in CPU performance at much lower power—in large part due to lower voltage operation and scaling to smaller ground rules. The enhanced performance will be needed to support added capabilities such as speech recognition [6]. It is assumed that by mid-

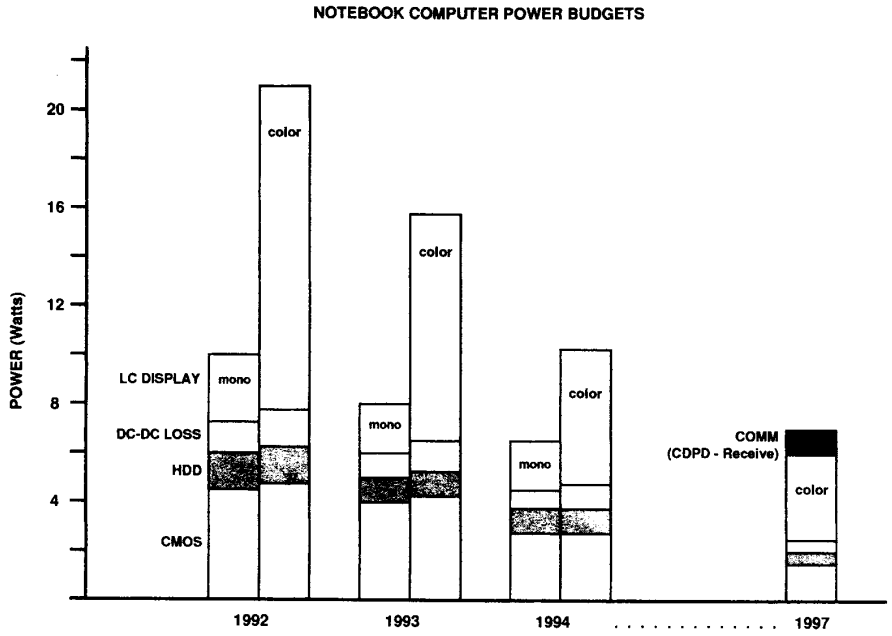

Fig. 3. Representative power budgets for monochrome and active-matrix color notebook computers in the 1992–1994 timeframe. Also shown is the projected power budget of a hypothetical 1997 notebook computer.

decade the display of choice will be the active-matrix color LC display. The A4 format is held constant but weight and thickness are assumed to drop by half over the time period indicated, and battery weight is scaled with system weight. Aggressive objectives for battery life improvement are indicated, going from 4 h in 1993 to 8 h in 1997. We have also assumed that by 1997, a nonnegligible portion of the power budget will be taken up by communications capability. Here, we have assumed that 1 W of power will be required for a CDPD (Cellular Digital Packet Data) wide-area communications adapter mounted in a PCMCIA (Personal Computer Memory Card International Association) card and operating in the receive mode. The 1 W power estimate is extrapolated from current values of about 2 W, and is about 1/2 of what would be required in the transmit mode. This will be discussed in more detail later in this paper. If notebooks with the characteristics shown are to meet these battery life objectives, substantial advances must be made in battery specific energy (energy per unit weight), as well as in the power-efficiency of many notebook subsystems. In the table we have indicated the power budgets and battery specific energies that will be required. The parameter values shown are consistent with the material presented on technology trends in subsequent sections of this paper. Li-Polymer batteries offer the prospect of much higher specific energy than today's batteries—about 200 Wh/kg has been projected [1]—but it is uncertain when such batteries will be commercially available. We have assumed here that 175-Wh/kg Li-Polymer batteries will be available by 1997. Note that even if this significant increase in battery specific energy is realized, total power budgets will be more stringent in the future than they are today if the indicated weight reduction and battery life goals are to be met. The projected advances outlined above for the 1997 notebook computer, while substantial, generally represent evolutionary extensions of today's technology. In this paper we choose to emphasize this evolutionary path, realizing that more revolutionary (albeit more speculative) developments—for example, the achievement of low-power reflective color LC displays with the image quality of backlit active-matrix displays—could alter the power budgets of future portable computers significantly.

D. Low-Power Technology Requirements

Outlined below, in brief, are some of the technology areas in which significant efforts will be required to reach the levels of function, performance, and power-efficiency envisioned in the year–1997 projection in Table 1. Several of these areas will be treated in more detail in subsequent sections of this paper.

1) Low-Power Color LC Displays: For future portables one would like to have a high-resolution low-cost color display readable in most lighting conditions. The best candidate known today is the active-matrix color LC display. Such displays dissipate significantly more power than monochrome STN displays because the backlighting must pass through color filters. Hence, they will consume a much larger fraction of the total power budget. The color display for 1997 in Table 1 is projected to require about 3.5 W. This will require improvement from today's levels for active-matrix color displays and is about the minimum that can be envisioned for a 10-in high-resolution display

Table 1 Projected Key Parameters for Notebook Computers. Projections Correspond to the "Common Vision" of Fig. 1. A4 Format is Assumed. Power and Battery Life Numbers Correspond to Full-Speed Operation with LC Display On and HDD Idle

NiMH	1993	1997
CPU		
Ground Rule	0.75-μm	0.35-μm
Clock frequency (MHz)	25	120
V_{DD} (Volts)	3.3	1.5
Performance (SPECint92)	14	100
Power (Watts)	1	0.4
Memory	4 MB	32 MB
HDD Storage		
Capacity	120 MB	700 MB
Form Factor	2.5-in.	1.8-in.
Display		
Resolution	VGA (640 × 480 pixels)	1 MPel
Type	STN LC	Active Matrix LC
Mono/color	Mono	Color
Size	10-in.	10-in.
Communications	Wireline Modem (9.6 kbaud)	Wireline Modem (19.6 kbaud) IR (4 Mb/s) RF Modem (CDPD)
Input	Keyboard	Keyboard Speech Handwriting
Thickness	4.0 cm	2.0 cm
Weight	2.5 kg	1.25 kg
dc Power		
Battery type	NiMH	Li-Polymer
Specific Energy	50 Wh/kg	175 Wh/kg
Battery weight	0.64 kg	0.32 kg
Battery energy	32 Wh	56 Wh
Battery life	4 h	8 h
Power	8 W	7 W
Power Budget		
Logic and Memory	4 W	1.5 W
Display	2 W	3.5 W
Communications	—	1 W
DC Conversion Loss	1 W	0.4 W
Storage/Other	1 W	0.6 W

with sufficient brightness for good readability. Getting much below about 3 W will require breakthroughs in light management.

2) Low-Power CMOS Electronics: Because the color display consumes so much of the power budget in the year–1997 portable in Table 1, very little power will be left for the electronics. Considering for now just the microprocessor, it is clear that an aggressive low-power design strategy will be needed to meet the requirements of the year–1997 example. Today's most power-efficient microprocessors [7], [8] provide approximately 30 SPECint92/W. That number will have to increase by 8x to meet the 1997 power budget. Active power in other logic and memory components will have to scale correspondingly. Improvements of this magnitude should be possible through voltage reduction, scaling to smaller groundrules, and other improvements, as will be discussed later in this paper.

3) Wireless Communications: A large fraction of future portable computers in all but the smallest sizes will require some form of wireless connectivity, either built-in or provided by optional communications adapters, in order to combine existing office-computer functions with new applications to be used away from an office by users in the field. There are different wireless I/O technologies which are likely to find significant application: RF and IR short-range links for cable replacements, RF and IR local-area networks (LAN's), and RF wide-area networks (WAN's). With reasonable assumptions about duty cycles, it should be possible to reduce average power to levels that will be tolerable in future notebook computers.

4) Dense, Low-Power Storage Devices: Portables require hard disk drives (HDD's) which are physically small, power-efficient, light-weight, and rugged. In Table 1, we have assumed that HDD power consumption can be reduced by about 0.6× by 1997. As will be discussed later in this paper, this should be achievable, primarily through miniaturization and reduced power consumption in support electronics. Magnetic storage with HDD's will remain a mainstay for notebooks, but other read/write and read-only storage technologies such as Flash EEPROM and CD-ROM will find increasing application as well [9]; Flash EEPROM for smaller capacity files, and CD-ROM for large-capacity read-only storage.

5) Input Technologies: Today, input technologies (e.g., keyboards) are generally a negligible consumer of power

in notebook computers. In the future, as new modes of input become popular and are in some cases connected by wireless to the portable computer, their energy consumption will be a key design consideration. Input technologies will be discussed briefly later in this paper.

6) System Power Management: Power management strategies are used to conserve power by monitoring system activity and reducing power levels in parts of the system when they are inactive, either by slowing or stopping the clock to the block in question or by powering it down. The effectiveness of such strategies will be greatest for systems in which the major power-consuming blocks can be operated with low duty factor. Note that for the 1997 notebook of Table 1, 50% of the power is consumed in the color display. Probably little can be done to reduce the display duty factor while the system is in use, since it is difficult to determine when an operator is looking at the display.

7) dc Power Supplies: In order to achieve the battery-life objectives shown in Table 1, there is a strong need for batteries with increased specific energy [1]. Li-polymer technology has the potential to provide about 3–4× improvement in specific energy when compared to today's NiCd or NiMH batteries, and we have assumed 3.5× for our year–1997 portable in Table 1. Design of dc-dc converters is complicated by the large number of different voltages present in a system which combines power-optimal logic, memory, storage and display technologies. Progress in efficient power conversion for low-voltage CMOS applications has been reported [10]. High-efficiencies in switched capacitor systems also requires matching frequencies to current load supplied. Claimed efficiencies (delivering 3.3 V) of greater than 90% across nearly two orders of magnitude of output current are now seen in commercially available standard power converters [11]. Converter efficiencies of >90% over three orders of magnitude of output current are a good target for future designs.

8) Dense Packaging: Packaging is an important system integration issue for portables that has power implications as well [12]. The very dense packaging that will be required for future portable computers implies not just small-form-factor package parts, but dense functional packaging of the system as a whole. Also, there are thermal constraints in portables not found in other systems. In portables, all of the power-dissipating components are enclosed in a small chassis in which little or no airflow can be provided. A third issue is the reduction of weight, which requires effort on case design and materials.

E. Summary

The portable computer market has grown rapidly and this growth is expected to continue for the foreseeable future. The trend is toward smaller, lighter-weight, more power-efficient products with enhanced performance and function. Significant reductions have been made in the power consumption of portable computers over the last few years, with today's color notebook computers typically dissipating about 10 W, roughly the same amount of power as the monochrome notebooks of just two or three years ago. To obtain longer battery life in future notebooks with greater performance and function (e.g., added communication capability), there will be a need for further power reduction in most technology areas. Significant improvements will also be needed in battery energy storage capability, and may be possible with Li-polymer battery technology. In the following sections, we turn to a more detailed look at some of the key technologies required for future notebook computers: display technology, low-power CMOS technology, wireless communications, hard-disk-drive technology, input technologies, and system power management.

III. Display Technology

For many notebook computers, the largest single power drain is the flat panel display. All emissive displays (i.e., electroluminescent or plasma) or backlit liquid crystal (LC) are rather inefficient at converting electrical power (watts) into visible radiation (lumens). The conversion factors are typically in the range of one lm/W, while a fluorescent light bulb produces about 50 lm/W. A perfectly efficient conversion would produce 680 lm/W if the light were peaked narrowly in the green (550 nm).

Today's most popular displays are backlit LC displays, in particular the passive matrix super-twisted nematic (STN) and the active matrix thin film transistor (TFT) LC display [13]. The latter display uses a transistor to control each LC picture element (pixel) in order to obtain higher contrast and better interactivity at high resolution. Both displays can produce color images when they are fabricated with an integrated array of color filters, one filter aligned to each pixel. Because of the continuing demand for better image quality, the focus of this section will be on the increasingly important TFT/LC color display. Coincidentally, this display also poses some of the greatest challenges in managing its present power demand and reducing its future power consumption. For a recent overview of displays more generally, the reader is referred to the lecture notes by Pleshko [14].

A. Power Consumption in Color TFT/LC Displays

Power is consumed in color TFT/LC displays by both the peripheral driver circuitry and the backlight. The drivers are of two types, gate drivers which select individual rows to be addressed and the column or data drivers which provide the analog voltages to store at each picture element along a selected row. Together, these circuits consume 1–2 W for a typical 10-in diagonal 640 × 480-pixel video graphics adapter (VGA) display refreshed at 72 Hz. The largest power drain, however, is actually the backlight. To understand the power demand of a backlight in a typical color panel, it is helpful to consider the path of light from source to viewer using the cross section in Fig. 4. A fluorescent tube (shown schematically) is the light source and its luminous flux is captured and redirected by a light guide diffuser system (not shown) to distribute

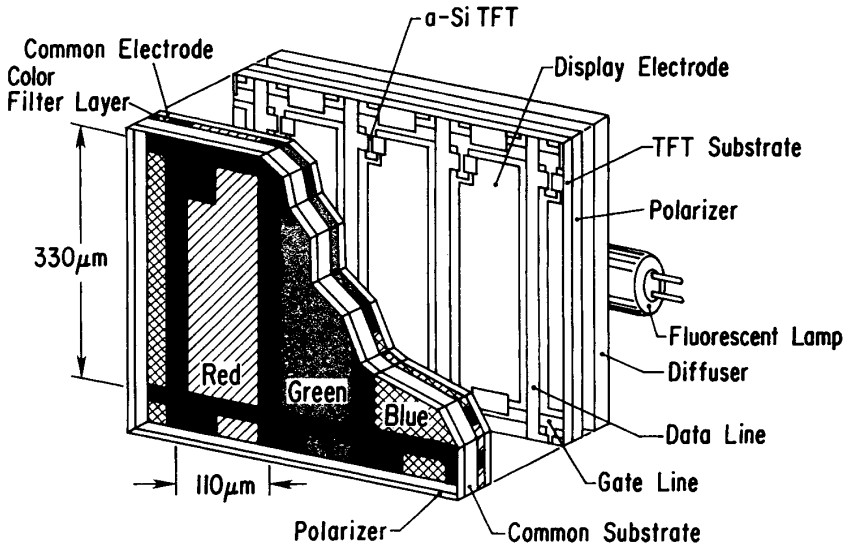

Fig. 4. Schematic diagram showing a portion of the structure of a backlit color TFT/LC display panel. Each 330-μm square pixel contains 110-μm-wide red, green, and blue subpixels as indicated. The pixel size is that required for a 10.4-in 640 × 480-pixel display.

light uniformly over the viewable area of the display. The efficiencies (E_D) of current diffuser designs are in the range of 30–50%, a significant achievement considering that the designs require the fluorescent bulb to be placed at the edge of the diffuser to reduce overall display thickness. With the drive to produce even thinner displays, it will be very challenging to substantially increase diffuser efficiency in the near future.

After leaving the diffuser, the light passes through the display panel with its layers and structures that modulate the intensity and color at every point; unfortunately, much of the light is lost along the way. The first layer encountered is a polarizer and would theoretically pass 50% of the unpolarized light from the diffuser, but actually achieves about 43–45% owing to absorptive losses. As the light passes through the individual light valves formed by each TFT and its associated transparent conductors, the opaque gate lines, data lines and transistor shadow a considerable fraction of the available area. Together with the black matrix, this leaves an "aperture ratio" of only 40–50% for most current VGA designs. Further along are the color filters which together could only pass 33% of the light in a perfect system. However, there are mismatches between the color filters and the spectral bands emitted by the fluorescent tube phosphors as well as absorption in the pass bands of the filters. The result is an actual efficiency in the range of 25%. This number, however, applies to color filters with a large color gamut like that of a cathode-ray tube (66% of the National Television System Committee (NTSC) standard). For many applications, it is possible to reduce the gamut somewhat (e.g., 40% of the NTSC standard), sacrificing color saturation for increased light throughput. In this case, the efficiency can reach 32% or even higher since each filter passes some of the other colors. Further small losses occur owing to reflections at each interface and absorption in the various materials, e.g., the transparent conductors. The cumulative effect is not insignificant, however, and multiplies the overall efficiency by a further factor of a little over 90%. The exit polarizer, again due to absorptive losses, passes only 87–89% of the correctly polarized light. As a result of all these factors, the transmission (T) through the display, excluding the diffuser, is only about 5%.

We can now calculate the power required by the backlight of a TFT/LC display given its size and its specification for luminance or brightness (B). For typical environments, a luminance of 40-100 nits (candelas per m^2) at normal viewing is necessary. The relation is

$$P = \frac{BA\pi}{E_L E_D T G}$$

where P is the power in watts, B is in nits, A is the display area in m^2, E_L is the luminous efficiency of the light source in lm/W, E_D, and T are unitless and G is the on-axis gain above a Lambertian distribution. Lambertian implies that light emitted from an area of the surface falls off as cos θ where θ is the viewing angle measured from the perpendicular. Gain comes from deliberately designing the backlight or adding elements to it which peak the light in the forward direction where the performance of the LC light valve is best. Typical on-axis gains are in the range of 1.5.

Putting all the typical values together for a 10.4-in display (0.034 m^2) yields

$$P = \frac{(70\,\text{nits})(0.034\,\text{m}^2)(\pi)}{(50\,\text{lm/W})(0.4)(0.05)(1.5)} = 5.0\,\text{W}$$

which is an average power requirement with today's technology. Actually the power requirement may be slightly greater to account for the 80–90% efficiency of the inverter which provides ac power to the fluorescent bulb(s).

A good measure of backlight progress is on-axis nits/W, which in the example above is 14.1. Two years ago typical values were 6–8 and future products are being designed to exceed 20. The keys to future reduction in backlight power are larger aperture ratios through higher conductivity metallurgy and new designs for the TFT array, as well as more efficient light source/diffuser combinations. There is little room to decrease the color gamut further as good multimedia performance requires saturated colors. Likewise, further peaking of the light distribution in the forward direction runs counter to the trend for wider viewing angle.

The next opportunity for significant progress in display power reduction will likely be in the driver circuits, which are consuming a larger percentage of the power as the backlight improves. New low power designs with lower supply voltages should appear soon, but progress will be paced by the simultaneous need for a greater number of gray levels and for higher speed as resolution increases. In general the reduction of display power will become progressively more difficult since the future trends in displays involve increased resolution, size, thinness, viewing angle, grayscale, and color gamut; all of which require more power. This suggests that there will be opportunities for breakthroughs which dramatically change the present boundaries, like the need for polarizers or color filters.

B. Reflective LC Color Displays

Although the most popular displays today are backlit LC displays, an interesting alternative under active investigation is the reflective LC display. Rather than selectively modulating transmitted light as discussed above for backlit displays, these displays selectively reflect ambient light, eliminating the need for the backlight power. To date these displays do not have the screen quality of backlit LC displays, but there are possible applications where one or more of the qualities discussed above, such as contrast, color gamut, screen luminance, number of grayscales, etc., can be sacrificed to reduce power. Small monochrome (or limited color) prototypes have been exhibited that can display alphanumeric and limited graphical information. For example, Sharp has exhibited a 5-in diagonal four-color (black, white, red, cyan) 320 × 240-pixel display using a guest/host LC that dissipates only 50 mW [15]. It is possible that one or more types of reflective LC displays will be sufficiently developed in the foreseeable future to find a market, at least for niche applications.

C. Summary

The growing popularity of backlit TFT/LC color displays for portable computers has led to major efforts to reduce their power consumption. Most of the power consumption in today's versions is in the backlight. A good measure of backlight progress is on-axis nits/W. Today's backlit 10.4-in TFT/LC color displays provide about 14 nits/W (which corresponds to 5 W backlight power dissipation at 70 nits luminance) compared to about 6–8 nits/W two years ago and about 20 nits/W for future products being designed now. Most of the known opportunities for significant backlight power reduction are being exploited, and the rate of future progress is likely to be less than in the past few years. The next area for significant progress in reducing total display power will probably be the driver circuitry, but there are limits here also, given the trends toward increased resolution and grayscale. Overall, getting much below about 3 W of dissipation for 10.4-in backlit TFT/LC color displays (including driver power) will be a difficult challenge. One direction of future effort will be in the development of reflective LC color displays, thus eliminating the need for backlighting. Versions which have been demonstrated up to now cannot match backlit versions in terms of front-of-screen quality, although they may be suitable for niche applications.

IV. Low-Power CMOS Technology

Another major area for power reduction in portable computers is in CMOS logic and memory circuitry. Battery-powered portable systems lend themselves to the optimization of CMOS circuitry for low-voltage, low-power operation [3]. Because of the self-contained nature of portable systems, their design can be decoupled to a large degree from the voltage levels and loading conditions of other equipment. The exploitation of this property of portable systems is in its infancy. The treatment of low-power CMOS technology in this paper will be brief. Many opportunities for power reduction in both logic and memory are treated in detail elsewhere in this issue and in the literature (see, for example, the papers by Davari *et al.* [16], [17] and by Itoh *et al.* [18]). Here, we will concentrate on illustrating the leverage available for CMOS power reduction through scaling to lower supply voltages and reduced ground rules by considering the particular case of CMOS microprocessors. We will very briefly review the principles of scaling and then examine microprocessor power dissipation in the light of those principles.

A. CMOS Logic Scaling

There are two aspects to power dissipation in CMOS logic circuitry; active power and inactive or standby power. We will be concerned here with active power dissipation. The active power dissipation (neglecting contributions from dc and noncapacitative currents, such as short-circuit currents [19]) is given by the well known equation

$$P = CV_{DD}^2 f$$

where P is power, C is load capacitance, V_{DD} is the power-supply voltage, and f is the switching frequency at which the circuit is operated. Clearly, reducing V_{DD} has the biggest effect on reducing active power. Simple dc CMOS circuits will operate over a wide range of power supply voltage–in theory, right down to the threshold voltage of

the devices in the circuit. Thus, a major direction for power reduction is to reduce the supply voltage. However, for a given circuit, reducing power supply voltage results in lower performance. To first order, for a given circuit and technology, delay goes as [20]

$$\tau \propto \frac{1}{V_{DD}}$$

and since τ is a measure of the maximum switching speed of a circuit, there is accordingly a tradeoff between power and performance: reducing V_{DD} decreases power but slows down the circuit.

In contrast, devices and circuits scaled according to constant-field scaling theory [21] improve in both power and delay. With constant-field scaling, horizontal dimensions, vertical dimensions, and all voltages are scaled down by the same reduction factor $k(k > 1)$, in order to maintain constant electric fields. An additional benefit of constant-field scaling is the avoidance of reliability problems associated with high electric fields as technology moves to finer dimensions. Constant-field scaling yields

$$P \propto \frac{1}{k^2}$$
$$\tau \propto \frac{1}{k}$$
$$P\tau \propto \frac{1}{k^3}.$$

Another quantity of interest with regard to microprocessors is the ratio of computing capability (or performance) to power. This ratio, for a given microprocessor design, should scale inversely as CV_{DD}^2 as ground rules and supply voltage are reduced [22]. For constant-field scaling, this ratio will scale as k^3.

While constant-field scaling results in more favorable power/performance than simply lowering the supply voltage in a given technology, it does require a more advanced device structure and technology at every scaled design point. Improvements in interconnection technology will also be necessary for long cross chip wires on large, very high performance chips to keep up with advances in device speed, due to the nonscalability of RC delays for lossy interconnections. In addition, scaling to supply voltages under about 1 V will be difficult, due to the inability to scale threshold voltages to arbitrarily low values: thresholds must be maintained high enough to prevent excessively high subthreshold leakage and standby power. Nevertheless, projections indicate that a scaling scenario very close to constant-field scaling should be appropriate for low-power applications down to supply voltages of about 1.5 V [16]. We turn next to an examination of CMOS microprocessor power dissipation, within the framework of constant-field scaling as outlined above.

B. Power Dissipation in CMOS Microprocessors

The power dissipation of performance-optimized microprocessors has increased greatly as clock rates have increased, and has reached 15–30 W at 100–200 MHz clock rates [23]. Recently, however, there has been considerable progress in the design of microprocessors with workstation-level performance but with power consumption of a few watts or less; low enough to be considered for portable applications [7], [8]. Reduced power consumption in these microprocessors has been achieved in part by taking advantage of reduced ground rules and lower supply voltage, but also by logic design, circuit design, and power management techniques aimed at reducing chip power consumption. Today's most power-efficient microprocessors provide about 30 SPECint92/W. Yeung *et al.* [7] have reported the achievement of 55 SPECint92 performance at 1.8 W power consumption for a RISC microprocessor built in 0.64-μm technology and operating with 80-MHz internal clock rate at 3.3 V.

Scaling of such a microprocessor design to still lower voltages and smaller ground rules offers the prospect of much lower power consumption and higher ratio of microprocessor performance to power. To estimate the improvement possible, assume that voltage is scaled from 3.3 V to 1.8 V, and that all dimensions are scaled by the same factor, so that minimum lateral dimensions are scaled from 0.64 μm to 0.35 μm. The performance-to-power ratio should increase by a factor of $(3.3/1.8)^3$, or 6.2 times. The maximum operating frequency will increase by the factor of 3.3/1.8 also; some of the additional speed can be traded back for a further power reduction. Reducing the supply voltage to 1.5 V would increase the performance-to-power ratio by an additional factor of $(1.8/1.5)^2$. Recall that the starting assumption was a microprocessor at 0.64 μm and 3.3 V with 55 SPECint92 performance at 1.8 W and 80 MHz. The result of reducing the supply voltage and ground rules as described would be a microprocessor at 0.35 μm and 1.5 V which should be capable of about 80 SPECint92 performance at 0.3 W and 120 MHz. Similar design points with increased computing capability (at higher power) should be possible, given the increased circuit density available in 0.35-μm technology. We have used 100 SPECint92 at 0.4 W for our 1997 microprocessor estimate in Table 1.

There are other aspects of advanced technology and circuit design which may be exploited to reduce power further; for example, by trading back some of the additional speed capability realized through scaling for lower power, by reducing device widths [19], [24]. This will make possible a lower capacitance design with reduced power consumption, a potentially fruitful design direction for PDA's and other handheld systems in which extremely low power dissipation is desirable. These and other potential tradeoffs require examination from the points of view of margins, technology complexity, standby power, cost, etc., to determine the best approach for a particular application.

C. Summary

In summary, we have shown the leverage which is available in advanced CMOS technology for reducing the power dissipation in the processors used in portable computers. Power is reduced because of the reduced capacitance

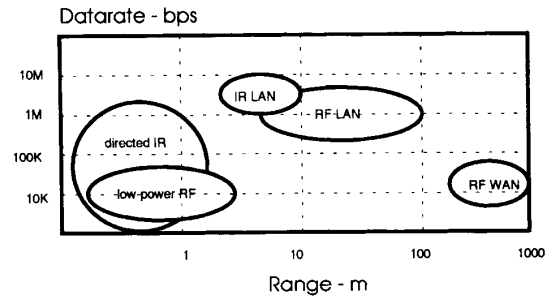

Fig. 5. Data rates and ranges of various wireless communications technologies.

of smaller devices and wires, because lower operating voltages can be used without greatly affecting performance, and because the intrinsically higher performance available in advanced technologies can be traded off for further reductions of power. Thus the leverage in low-voltage CMOS technology is real and large, and lends itself to exploitation in portable systems.

V. Wireless Communications for Mobile Computing

In addition to reducing the power consumption in the display and system electronics, low-power technology will be needed for the growing communications content in future portable computers. While wireless communication technology is still evolving, broadly speaking, it can be categorized according to usage as *cable replacement*, *local area*, or *wide area* communication technology. This view is somewhat biased toward the computer industry perspective, and the perceived needs of mobile computer users. *Cable replacement* technologies are used for short-range, point-to-point links. *Local area* technologies refer to office or campus environments, and *wide area* technologies generally include metropolitan and nationwide services such as cellular telephone and paging networks. Fig. 5 illustrates typical data rates and cell sizes for IR and RF technology. Cable replacement covers ranges up to a few meters. LAN technologies have cell sizes from a few meters to 100 m, and WAN technologies extend upward from 100 m.

A. Cable-Replacement Technologies

Cable replacement technologies are driven by the end-user requirement for simplicity and ease of use. Today, cabling between computers (mobile and fixed), and to their peripherals, is one of the aspects of computing most disliked by users. The obvious applications for wireless technology in this category are all-wireless peripherals, including the mouse, printer, keyboard, and perhaps audio-input devices (see also the discussion in the section on input technologies). This application has very strict power-dissipation requirements since users will demand long battery life. Battery life requirements for this class of application are from weeks to months, depending on usage.

An example of this kind of product is the wireless docking station, used to synchronize files between a mobile computer and a desktop. The communication device uses RF to transmit digital file information. The portable radio unit typically operates from a single battery pack and provides connectivity at a range up to 30 ft. This class of device is typically a narrowband, analog transmission system designed for minimal cost.

The Infrared Data Association (IRDA) has been created to promote the use of IR communications within the computer industry for connectivity applications, and to promote standards allowing for interoperability between a wide variety of devices, including PC's, electronic organizers, and PDA's. The basic data rate has been established at 115 kbps with a range of 1 m (point and shoot). Component cost and power dissipation are very low.

B. Local Area Wireless Communications

Networks are distinguished from the point-to-point applications just described by their use of multiple-access protocols. Future wireless communication networks, both local and wide area, are expected to carry data as well as voice and multimedia (image and video) information to the end user. For example, it is expected that as asynchronous transfer mode (ATM) networks are deployed, wireless access points into the ATM backbone will be needed. This expectation provides one driving force behind the need for increasing bandwidth over time. Current-generation RF wireless networks support data rates in the range of 200 kbps to a few Mbps. As ethernet and token-ring local-area environments move higher in speed (for example to 100 Mbps ethernet), wireless network access point bandwidth will need to increase. Since each access point defines one cell, and cell bandwidth is shared amongst all users of the cell, high-performance wireless LAN's will depend on high-data-rate radio technology. Unfortunately, achieving higher data rates is constrained not only by technology, but by available spectrum.

The most commonly employed wireless communication techniques for LAN systems are pulse-modulated infrared (IR) and spread-spectrum radio. IR, unlike RF, can be used in either a diffuse mode or in a directed (point-and-shoot) mode. The choice has an impact on power dissipation, as well as on the protocols employed. Diffuse systems have the advantage of not requiring the user to establish a line of sight to communicate, but they carry a generally higher power requirement since more optical power is needed to fully "illuminate" a typical office. Diffuse products are already on the market, and other companies are expected to field diffuse IR communication devices over the next few years. The so-called directed class of IR links will also emerge both for LAN access as well as simple printer connections and portable/desktop synchronization.

Wireless LAN adapters using RF technology are now available, with a flood of new products expected. Typical products using spread-spectrum technology include the RangeLAN/PCMCIA from Proxim and the WaveLAN/PCMCIA from NCR [25]. Frequency-hopping and direct-sequence are two forms of spread-spectrum radio in use. The advantage of spread-spectrum technology

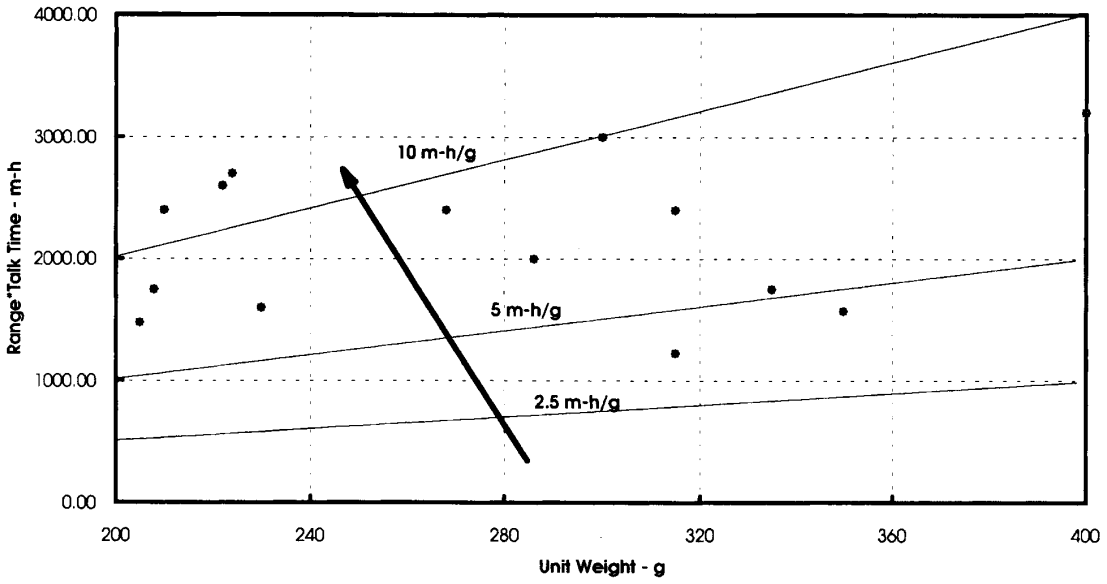

Fig. 6. A figure of merit—talktime multiplied by range—is plotted versus handset weight for currently available mobile telephone handsets.

is its inherent tolerance to noise sources and its built-in security (transmissions are very difficult to intercept). Federal Communications Commission (FCC) rules permit spread-spectrum operation with output power up to 1W in the industrial, scientific, and medical (ISM) bands: 902–928 MHz, 2.4–2.4835 GHz, and 5.725–5.850 GHz. Of course, one of the major challenges associated with RF technology for mobile computing is in making adapters in the PCMCIA form-factor, especially the preferred Type II cards, which are only 5 mm thick. Major difficulties in designing and debugging a radio which operates reliably in a credit-card-sized package, with large sources of emissions in the form of CPU, video, and other clocks in the mobile computer, are commonly encountered by developers. Radio integration and low-power operation are key to making these devices commonplace. Integration of the antenna can also become a major design challenge.

C. Wide Area Wireless Communications

The dominant form of wide-area wireless communications in North America is analog cellular telephony. The Advanced Mobile Phone System (AMPS) is a first-generation analog system with a current user base of about 16 million subscribers. The Electronic Industry Association (EIA) has announced Interim Standard 54 (IS-54), which adds digital voice transmission capability to the network [26]. In Europe, the Groupe Special Mobile (GSM) standard is emerging as the preferred digital cellular system, and is significantly further along in terms of deployment than the US market. There are also networks dedicated to data communications, including the Advanced Radio Data Information Service (ARDIS) network, and the many paging and messaging networks. The specialized mobile radio (SMR)

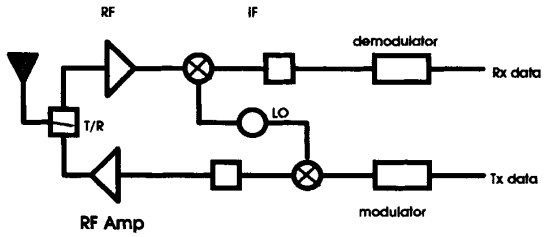

Fig. 7. Block diagram of a typical RF communication adapter.

and enhanced specialized mobile radio (ESMR) service providers are also well positioned to provide low-data-rate wireless services. Perhaps the most intriguing development in the US market is the FCC action to make spectrum available for so-called Personal Communication Services (PCS). PCS is a micro-cellular, digital network which has the potential to make wireless personal communications truly ubiquitous.

With all of the competing wide-area systems offering or preparing to offer data services today, it is not surprising that the wide-area wireless data market is characterized by fragmented offerings and interim solutions. Data-over-cellular is the most straightforward wide-area solution, but is hampered by high connection fees. Cellular Digital Packet Data (CDPD), a packet data service to be offered by the cellular providers, is designed to lower data transmission costs while utilizing the existing cellular infrastructure. Early CDPD adapters typically have radiated powers of about 600 mW, with a total adapter power of about 4–5 W peak during transmission. Receive-mode power levels are typically about 2 W. Reasonable expectations for CDPD adapters in the 1997 timeframe would be about 2 W in

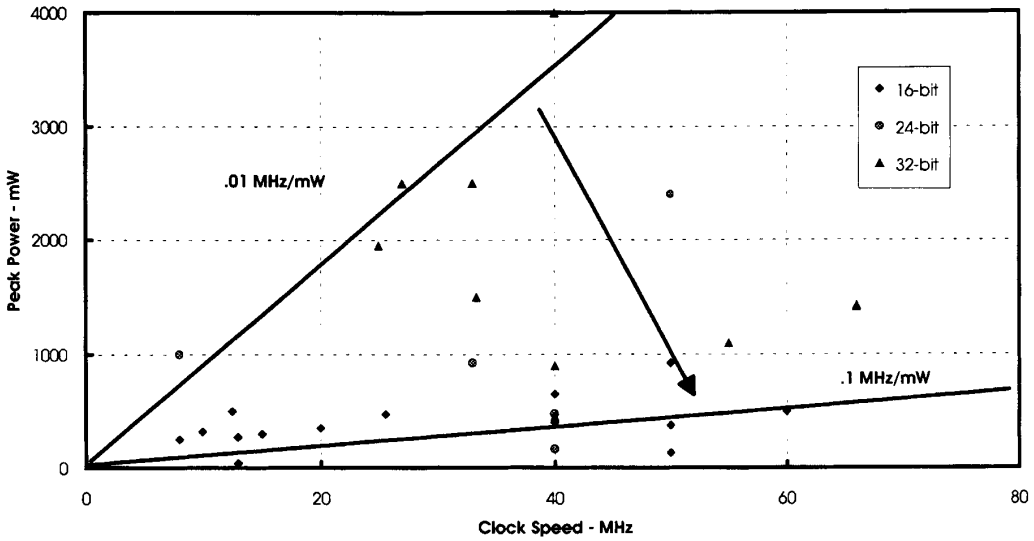

Fig. 8. Peak DSP power versus clock speed for currently-available 16-, 24-, and 32-b DSP's.

transmit mode (600 mW radiated), and perhaps 1 W in receive mode. Since the transmit duty cycle is low, most of the power is consumed while in receive mode. We have used this 1 W receive-mode estimate in the 1997 projected notebook computer power budget in Table 1. One possible option for further reducing receive-mode power is to provide an energy detection or "listen" capability without powering up the entire receiver. This may make it possible to reduce the 1 W estimate still further.

Low-power design for wide-area communications is already an important issue for the mobile telephone industry, and continued improvements in power efficiency can be observed. For example, Fig. 6 is a plot for selected cordless handsets showing a figure of merit, the talktime in hours multiplied by the range in meters, as a function of handset weight in grams [27]. Range and talktime can be increased by adding weight to the handset in the form of additional batteries; trend lines are shown for the range multiplied by talktime (m-h), divided by the handset weight (m-h/g). This metric is an indicator of overall power efficiency of the handset design, and can be seen for this set of devices to be exceeding 10 m-h/g.

D. Key Technology Elements for Mobile Communications

Fig. 7 is a highly simplified block diagram of a typical RF communication subsystem. It shows the major functional blocks from the digital interface out to the antenna. This section discusses trends in power dissipation for the key building blocks of an RF transceiver.

The key technology elements important for future wireless communications systems are the transceiver front-end, the modulator/demodulator stage, which may be implemented using a digital signal processor (DSP), and the microcontroller. The radio architecture may be analog or digital. A digital approach is more difficult to implement, as it requires oversampling the data at high-speed. Therefore, for most transceivers the front-end stage is analog, while the IF and demodulator blocks may be a mix of analog and digital design. Proponents of all-digital designs point out that digital processing is more flexible than standard analog techniques, but it is unclear that an all-digital design would result in the lowest power dissipation. Typical DSP's in use today can have relatively large power dissipation and large current requirements, even when they have power management modes. The high power and current demands can be troublesome for PCMCIA implementations. Fig. 8 shows a plot of peak DSP power for 16-, 24-, and 32-b DSP's available as of this writing [28]. Power dissipation must be driven down in these devices for portable applications, since, as more signal processing and control functions are loaded onto the DSP, more processing capability is required.

The transceiver front-end, on the other hand, is typically difficult to power manage, other than through careful analog design and proper selection of biasing conditions. Powering-down sections of the transceiver, such as buffer and amplifier stages, can result in system stability problems, as time is required for the circuits to settle once power is restored. Transceiver turnaround time may be impacted in this case.

The average transceiver power dissipation is a function of how often the transmitter is operating, and of the transmit and receive mode power. Since a transceiver is required to have the receiver constantly operational and "listening" for transmissions intended for it, the receive mode power is the most important factor in determining overall transceiver power dissipation. Duty cycles are typically low, as Fig. 9 illustrates.

The total power dissipation in transmit mode is dominated by the power radiated, commonly referred to as the

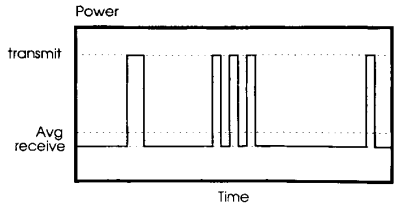

Fig. 9. Schematic diagram of the time dependence of transceiver power, indicating that receive-mode power is the most important factor in determing average transceiver power dissipation.

effective isotropic radiated power (EIRP). The RF amplifier in the transmit chain supplies power to the antenna, where typical maximum allowed EIRP values range from 10mW for CT-2 (cordless telecommunications, second generation [26]) to 1 W for ISM band, spread spectrum, to as high as several watts for cellular systems. Generating these power levels often creates high-current problems in a PCMCIA environment. Excessive power dissipation in the rest of the transceiver can create thermal problems as well.

System issues, such as the choice of modulation, also can have a major impact on transceiver power dissipation. For example, for a pulse-modulated IR transceiver, power efficiency can vary depending on which type of pulse modulation is used. In Fig. 10, typical pulse trains are depicted for pulse-width modulation (PWM) and pulse-position modulation (PPM). In the latter case (PPM), pulses are of constant width but vary in position, whereas PWM uses pulses of varying widths. The duty cycle for PPM will be lower, resulting in lower overall power dissipation, but with a slightly more complicated demodulator.

E. PCMCIA and Small Form-Factor Constraints

For the class of portable systems which provide for "upgradable" communications, PCMCIA is emerging as the adapter standard of choice. Technology suppliers and adapter manufacturers all are driven by the need to package the adapter hardware into the common Type II version of the PCMCIA specification.

The electrical interface for PCMCIA today is a 68-pin connection, with symmetrically-loaded power and ground pins provided in case of inverted card operation. The standard calls out four ground pins and two V_{CC} pins, with each pin limited to less than 0.5A of current. More importantly, however, the packaging form-factor itself typically limits the power dissipation supportable in a PCMCIA implementation to a few watts. In this environment, low-power design is essential for meeting thermal specifications. The PCMCIA standard is likely to move to higher performance in the future, supporting a higher speed, 32-b interface to the same physical package over the same connector. This means that the adapter interface will not be a bottleneck for communications. However, the faster interface electronics will consume more of the limited power budget in the adapter, leaving even less for the communication function itself.

F. Summary

Most future portable computers in all but the smallest sizes will require some form of wireless connectivity provided by optional communications adapters or built into the system, both to replace wired connections to devices such as an optional keyboard and to combine existing office-computer functions with new applications to be used away from an office by mobile users. Accordingly, as Fig. 3 clearly shows, the power required for on-board wireless communications is expected to be an increasing fraction of the total power budget for future portable systems. There are different wireless I/O technologies which are likely to find significant application: cable-replacement technologies (both RF and IR), RF and IR local-area network (LAN) technologies, and RF wide-area network (WAN) technologies. When transmitting, RF and IR devices for LAN's and WAN's will consume significant power in order to cover the desired ranges. Receive-mode power needs to be minimized for long battery life in portable devices, with the ideal low-power transceiver dissipating power only while actually receiving or transmitting data. With reasonable assumptions about transmit and receive duty cycles, it should be possible to reduce average power to levels that will be tolerable in future notebook computers.

VI. Hard Disk Drive Technology

Magnetic storage technology provides the major memory function in portable computer systems. In spite of the emergence of nonvolatile semiconductor memory chips, hard disk drives (HDD's) have persisted largely because of cost and capacity advantages. Historically, HDD's evolved through aerial density improvements to meet needs for greater capacity [29]. In the 1980's, with the advent of personal computers, volumetric efficiency became an important design feature. The growth in the market for portable computer systems has now fueled more innovations in HDD design. The extraordinary physical (shock) and environmental (temperature and humidity) exposure of portable systems has challenged HDD reliability, while limited battery capacity has demanded efficient power usage. The new generation of HDD's has been adapted to these challenges much more rapidly than other subsystems of a portable computer. In this section the design methodologies and tradeoffs to achieve optimum power-performance characteristics in HDD's are addressed.

A. Key Parameters of Hard Disk Drives

Table 2 shows a representative parameter set and power requirements of various form factor drives that came to market in 1993. The diameter of the disk-platter is a measure of an HDD's physical dimension, and is generally referred to as its form factor.

One may observe that disk diameter decreases by a factor of two every two generations of form factor. The 3.5-in and 1.3-in products are not as important as the 2.5-in and 1.8-in HDD's in the portable computer arena. The 3.5-in HDD's, however, dominate the storage market because of

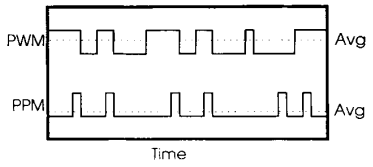

Fig. 10. Typical pulse-modulated IR transceiver pulse trains for pulse width modulation (PWM) and pulse position modulation (PPM).

Table 2 Key Parameters and Power Requirements of Hard Disk Drives (HDD's)

Form-Factor	3.5-in	2.5-in	1.8-in	1.3-in
Key Parameters				
Capacity (MB)	720	340	85	40
Weight (g)	500	180	75	28
Height (mm)	20	12.7	10.5	10.5
Disk Diameter (mm)	95	65	48	34
Number of Disks	2	2	2	2
RPM	4500	3800	4500	5400
Power Requirements				
Start-Up (W)	8.0	4.7	3.0	2.2
Seek (W)	4.5	2.0	1.5	1.7
Read/Write (W)	4.2	2.8	1.725	1.6
Idle (W)	2.5	1.35	0.75	1.0
Standby (W)	0.8	0.4	0.035	0.5
Sleep (W)	0.1	0.2	0.02	0.015

their importance for desktop and high performance workstations. Within the 2.5-in and 1.8-in market, obviously, the 2.5-in plays a critical role in portable systems with embedded (fixed inside) storage. Capacity growth, driven by versatile disk operating systems and image intensive application offerings, has necessitated the use of a 2.5-in HDD. By 1995 2.5-in products are expected to be 20% of the total HDD units shipped worldwide [30]. The detachable storage system has gained increasing user acceptance. This has implications for both 2.5-in and 1.8-in HDD's. Notably, the 1.8-in HDD, with the establishment of the PCMCIA standards, is anticipated to be above 5% of total units shipped worldwide by 1995. An increase in the subnotebook computer market can be expected to spur the growth of 1.8-in HDD's both in embedded and detachable applications. The 1.3-in HDD, because of its limited capacity at present, is targeted for special applications including palm-size computers and personal digital assistants. Assuming a 60% annual growth in areal density, a 1.8-in PCMCIA HDD, as shown in Table 1, is expected to have about 700 MB capacity by year 1997.

B. Hard Disk Drive Subsystems

To help discuss the power saving trends in portable HDD's, a brief discussion of HDD subsystems is in order. Fig. 11 shows a schematic representation of an HDD with its major electronic and electromechanical subsystems. The electromechanical subsystem is composed of an actuator, a spindle motor, and a latch. The read/write head transport mechanism is actuated by a voice coil motor (VCM), and is typically of a rotary type. Rotary actuator systems

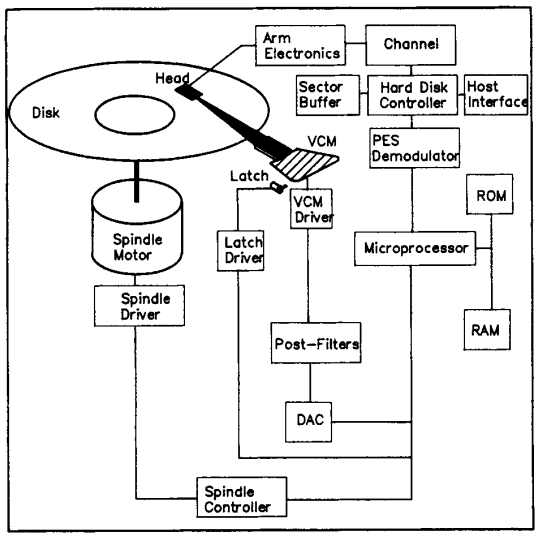

Fig. 11. Schematic diagram of the electronic and electromechanical subsystems of a hard disk drive (HDD).

are common in HDD's designed for portable applications, largely because of their compactness and low cost.

The spindle motor is a brushless type dc motor that supports and drives the disk-platters directly at constant speed. The power dissipation in an HDD is a strong function of spindle motor system performance as discussed later in this section. The design of a compact spindle is a difficult compromise between torque generation efficiency and ruggedness. Larger bearings provide shock resistance, but lesser magnetic volume limits torque generation capability. For spindle speed control, the rotor position measurement is achieved indirectly using back-EMF instead of traditional Hall sensors. As a result, part count and cost are reduced, but the methods to initialize rotation under uncertain stiction torque have become increasingly complex.

The latch that keeps the actuator at a secure position under power-off conditions has become an important component in providing ruggedness. Latch designs have varied from product to product. 3.5-in products traditionally have had solenoid type devices, but attempts have been made to incorporate aerodynamically activated latch mechanisms to save cost. They are, however, not very reliable because of the marginal activating force presented by internal airflow. In 2.5-in products, either passive or active magnetism-based latches can be found. In a recent 1.8-in product, an inertially activated latch is employed in which the cost and power is expected to be the lowest achieved so far.

The electronic subsystem provides servo control of the electromechanical system and facilitates the magnetic recording/readback process. A microprocessor unit (MPU) is at the core of all servo-control and data manipulation functions. The digital servo for VCM is driven by the MPU at a sampling rate of about 3 ksamples/s. The position error signal (PES) is read by the MPU and digital servo

control signals are generated and applied through a digital-to-analog converter (DAC). In certain HDD's the DAC output is further conditioned by post-filters such as a notch to suppress system resonance typically located around 2 kHz. Due to the relatively low sampling rates used in portable HDD's, post-filtering is provided in the analog domain as shown in Fig. 11. Some HDD designs have avoided altogether the use of post-filtering to minimize cost and power use. The spindle-servo, which is of a relatively low-bandwidth compared to the VCM servo, is either commanded directly by the MPU or by a specialized spindle controller supervised by the MPU.

The magnetic recording/readback process is also supervised by the MPU. The MPU commands the hard disk controller to manipulate data through the host interface. A sector buffer provides volatile temporary storage for the hard disk controller to facilitate data movement. The arm-electronics module (thus called because it used to be mounted on the actuator arms) provides read signal amplification or write current to the heads. It is multiplexed between heads when several heads are used. The channel optimally decodes/encodes and drives the digital/analog signals for read/write functions. As the read/write heads move from the outer-diameter (OD) to the inner-diameter (ID) of the disks, the read-back signal characteristics change and channel electronics, for example, equalizes such variations. In the case of multizone recording, where a fixed linear bit density pattern is used, the channel electronics are designed to drive the heads at different frequencies.

The seemingly unimportant issue of power-up and power-down operations creates unique compromises in the integration of the subsystems. Of particular importance is the mandatory requirement that the read/write heads are never allowed to land on the data-zone as the HDD is powered down. In conventional HDD's, the design approach known as contact start stop (CSS) is employed and the heads are rapidly moved to the ID zone of the disk at power down. As an alternative to CSS, the heads can also be unloaded from the disk surface and held without contacting the disk [31]. This method requires the loading and unloading of heads (L/UL) during the power-up and power-down operations of an HDD. A ramp-like design is used to assist the head L/UL operation. The power-up sequence is effectively achieved using the intelligence of the microprocessor, whereas the power down sequence is delicately achieved using the kinetic energy stored in the rotating disk in conjunction with simple transistor logic. Use of a L/UL mechanism to achieve ruggedness and unlimited start-stop cycles has added further importance to the power-down sequence because of the increased energy levels required to move the heads onto the L/UL ramp.

C. Power Savings in a Portable Hard Disk Drive

Power savings in a portable HDD are achieved by introducing operational modes as defined in Table 2 (under Power Requirements). Each mode is a trade-off between recovery time (to read/write) and energy saved from the battery. Conventional operational modes are:

- Startup: Spindle accelerated from rest to rated speed.
- Seek: VCM is actively moved to a new position.
- Read/Write: Data is transferred between media and head while on a track.
- Idle: Spindle and actuator are under normal control. No read/write of data. Only PES information is read by the demodulation circuit.
- Standby: Spindle is at rest, heads are parked, and interface can receive commands from host.
- Sleep: In addition to "Standby" condition the interface is powered off and a single logic line is active to sense a reset signal from host.

From Table 2 it can be seen that the peak power is required at the starting phase of an HDD. Generally the peak power drain is limited to less than 5 s for all form factor designs. Start-up power drops with the decrease in form factor. The power and energy (time-integral of power) reduction with the form factor is largely due to a decrease in the rotational inertia of the two disks, which is a function of (r^4t) where r is the radius and t is the thickness of the disk. In a later section, we will discuss the impact of reduced start-up (or spin-up) energy on power management strategies for small disk files. Use of a glass disk, which is stiffer than aluminum, is becoming common. This trend not only improves the ruggedness of a drive at the head-disk interface, but potentially further decreases start-up energy through the reduction of disk mass. Seek power is related to target access time and actuator inertia. As the form factor is reduced, the seek power for a given access time is expected to decrease, but around 1.3-in the advantage of the form factor on seek power is not so visible. Use of light weight actuator material has been an avenue for seek power reduction. For portables, the seek and read/write duty cycles are relatively low compared to the idle mode. Thus the idle mode power dissipation is generally the most important component of the average power dissipation of an HDD, and its minimization is critical for achieving significant power savings.

The distribution of the idle mode power dissipation is shown in Fig. 12. The power values correspond to a 2.5-in HDD with 2 disks spinning at 3600 RPM. It can be observed that the spindle system consumes about 50% of the idle power. Even though the start-up power decreases with the form factor the idle power levels off around 1 W in Table 2. This important trend deserves attention and interpretation. The use of thin-film heads (where the read signal is velocity dependent) or the need to provide acceptable data-rates requires a minimum linear velocity for the media. In order to maintain a minimum linear velocity, the spindle speeds have to grow at least inversely proportional to the form factor. It can be observed from Table 2 that the spindle speeds indeed show an increasing trend with smaller form factor HDD's.

Traditionally, power use for the spindle had a strong higher order (nonlinear) dependence on disk-diameter and spindle speed. However, for a given small form factor HDD, measurements show a linear relationship of power with spindle rotation speed. Fig. 13 shows the effect of speed on total spindle power for a 2.5-in HDD, where

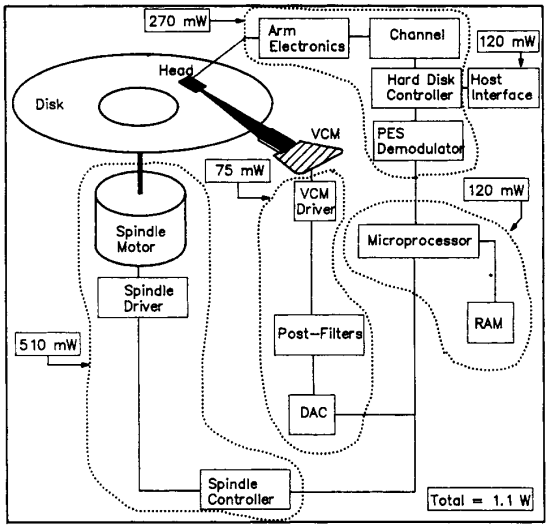

Fig. 12. Schematic diagram showing the contributions to idle-mode power dissipation.

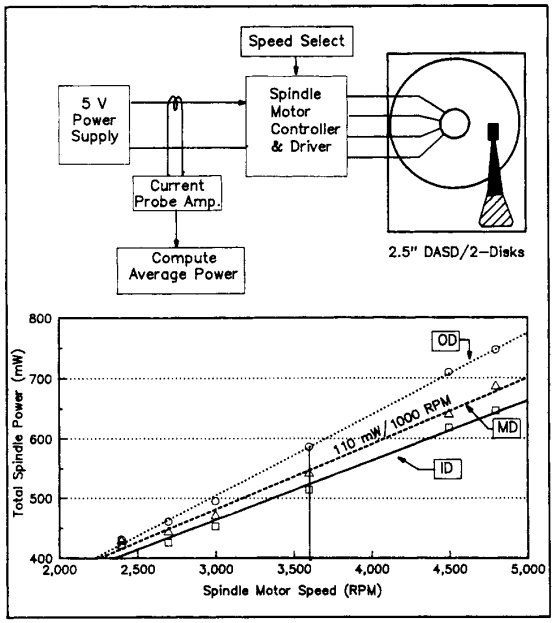

Fig. 13. Spindle motor power.

the linear dependence is observable [32]. Controlled experiments are required to study the form factor dependence in which the power dissipation mechanism as well as the spindle motor efficiency are subject to change. Regardless, it can be argued that an increase in speed and a decrease in disk diameter will have a counter-balancing effect. Therefore the idle power, which is dominated by spindle power does not show a strong correlation to the form factor. Further measurement and analysis of the power dissipation mechanism within a 2.5-in/2-disk spindle motor subsystem show the following components (and estimated losses): Driver (30%), sliders (at OD) + disks (30%), Ferro-fluid seals (15%), eddy current loss (15%), and bearings (10%) [32]. As pointed out later, the sliders contribute to aerodynamic drag, which is as high as the losses due to disks spinning freely in a volume filled with air. A trend towards smaller low-preload sliders and less elaborate ferro-fluid sealing protection may allow about 25% of spindle system power reduction. Use of efficient (low on-resistance) FET drivers and high energy magnets will reduce both copper loss as well as eddy current loss. Optimistically, for a given rotation speed, a total of 50% reduction in idle mode power is possible. Frictional and viscous components are the fundamental limiters to power savings in the spindle motor system.

Use of magnetoresistive (MR) heads, which are insensitive to speed, may produce a different power trend in the future. Because of MR head insensitivity to media speed, idle power optimization can be flexibly traded off with other performance parameters. Improvements in high energy magnetic materials, spindle motor switching algorithms and the introduction of efficient drive circuits still hold promise for reducing the idle power usage of the electromechanical subsystem.

Fig. 13 also shows the effect of slider aerodynamic drag on average spindle power. The OD head position consumes about 75 mW more than ID in a 4-head/2-disk configuration. Nevertheless, this feature is not exploited in the idle mode due to head-disk interface concerns at ID. Further issues for design optimization are considered in [33]. The VCM system alone takes less than 10% of total idle power, where the VCM current typically has both AC and dc components. Slider aerodynamic drag, VCM pivot friction, and electronic cable bias forces contribute to the net dc current. Since track-follow accuracy is not critical during the idle mode, the servo system can be configured to operate at a low sampling rate. By skipping PES sectors the electronic circuits that provide read, as well as PES demodulation functions, can be deactivated periodically. The power dissipated in channel/PES electronics is therefore limited to the PES generation phase only. In fact, recent 2.5-in drives and follow-on products achieve minimum idle mode power by "skip-sector" servo. In the example of Fig. 12 a passive latch is used and hence no power dissipation occurs. A solenoid-based latch will, however, consume less than 5% of the total idle power.

In idle mode, as defined before, the HDD must be ready to receive host commands. The interface circuit therefore remains active and dissipates about 10% of the idle power. The MPU and its memory consume another 10%. The electronic circuits that include hard disk controller, channel, arm-electronics and PES demodulator consumes less than 30% of the power. In the electronics the frequency dependent switching losses play a dominant part in the dissipation mechanism. Demand for high data rate or improved command response time in general will force higher power dissipation. Migration to a CMOS technology with a reduced voltage (3.3 V) chip set is already in the horizon [34], and a major reduction in electronic

power (as much as 2x) is derived. Eventually, innovative system design, technology improvement and intelligently partitioned circuits will contribute to idle mode power reduction. Standby and sleep mode powers have already reached fractional watt levels, and any variation seen in Table 2 from one form factor to another is mostly due to the level of electronic integration used. Newly released products in 1.8-in and 1.3-in form factor tend to have better standby and sleep performance because of the use of a new generation of chips.

D. Summary

In today's hard disk drives, the electromechanical system (including drivers) and electronics dissipate about equal amounts of power during idle mode operation. Spindle motor power reduction, however, has a fundamental limit. Electronic integration and partitioning of electronic functions, and an introduction of low voltage electronics, coupled with power efficient semiconductor technology is expected to contribute the most to power savings in the future. Taken together, form factor reduction, novel power management techniques, lower voltage chips, and silicon technology improvements can be expected to contribute to a factor-of-two reduction in idle power requirements, with a goal of 0.5 W (total idle power) likely being attainable by year 1997.

VII. Input Technologies

Today, input devices take up a very small portion of the power budgets of notebook computers (they do not even show up in the power budgets displayed in Fig. 2 and Fig. 3). They are of greater relative importance in personal communicators, but are still overshadowed by the transceiver. Here we will briefly review input technologies as they are likely to evolve through this decade.

A. Keyboards

Today, keyboards are typically an attached part of portable computers, and the main issue is reducing their thickness and making them fit ever smaller formats without sacrificing usability. Similarly, today many notebooks incorporate pointing devices—mice, trackballs, and isometric joysticks [35]. In the future, these technologies are likely to diverge. Pointing devices will likely be increasingly integrated, so that they can be incorporated even in communicators (to permit selection by the hand holding the communicator, leaving the other hand free). But the keyboard will become increasingly an option, needing some form of wireless for connection, rather than cables. The primary power challenge will be the dissipation associated with gracefully connecting the computer via wireless to office-grade keyboard input devices.

B. Pen-Input and Touch Technologies

Three technologies are competing in this arena; it is not clear at this point which will dominate. They are electromagnetic, electrostatic, and elastomeric technology. Electromagnetic technology gives the greatest precision, requires a moderately expensive pen and does not support touch. Power reductions are achieved by the use of an active pen, rather than systems in which the digitizer interrogates the pen. Electrostatic systems sense touch, pen, and presence of other objects, and so require more complex logic to manage. Elastomeric technology (involving use of a transparent layer with resistance affected by stylus or finger pressure) permits a very inexpensive pen, but has not appeared to offer as good resolution or stability against manufacturing variations. Nevertheless, it is being used successfully in current products [4]. In the future, the use of sensor input is likely to increase–detecting a user's presence, attention, etc. There are opportunities for putting more function in the pen, such as user identification and personal data, even all the strokes that the pen has made while writing, waiting for the opportunity for wireless download, or letting the pen also become the holder of the microphone one uses for dictation. In this case one will need a good encoding to make wireless transmission of the information to the system affordable from a power point of view.

C. Eyetracking Technology

Today, eyetracking is used in some expensive cameras to detect the user's gaze direction in a viewfinder. There are some exciting uses to which this could be put in a heads-up display or virtual reality viewer to make selection a hands-free operation, in the future. Success in this area could make feasible systems with much smaller displays and thus less display power.

D. Summary

Clearly, the minimum energy-consumption design points for each of these input devices will need to be understood. In some cases they can possibly even be self-powered. A keyboard can generate a tiny current from the force with which a finger presses down; a microphone can do the same from the acoustic energy input. As we move toward input devices as separate components which identify themselves and their availability to the systems through wireless, each such device needs to be designed for something like a year's battery life, such as the electronic key/alarm transmitter for a luxury car offers today.

VIII. Power Management in Portable Systems

Power management techniques are used to conserve power by monitoring system activity and reducing power levels in parts of the system when they are inactive, either by slowing or stopping the clock to the block in question or by powering it down. We begin this section by describing power management techniques which can be implemented at the chip, subsystem, or system level, and then discuss briefly what can be done to further improve their effective-

ness beyond what is achieved today. We discuss subsystems separately, since at present software does not exist to coordinate the different subsystems, or to use information from one subsystem to force activity or inactivity in another. For example, we cannot yet manage a total budget of estimated power consumption in software. Nonetheless, significant savings can be achieved by reducing the duty cycle of each subsystem separately. Following this discussion of power management techniques, we turn next to an examination of the problem of power management in small hard files in order to give a concrete example of the power management tradeoffs possible in a system or subsystem. Finally, we give a brief summary of the key points discussed.

A. Power-Management Techniques

Power management is a design problem that couples hardware and several layers of software. Today, power management architectures have been established [36]–[37], but in general they have not yet been fully exploited on any of the layers. Levels of activity are defined for blocks of the hardware: the CPU-memory complex (or parts thereof), the hard disk, the display backlight, etc. which translate into a set of states with different power dissipation which can be implemented at the chip, subsystem, and/or system level. The following is a listing of some of the possible power states of a system:

- **Full-On:** The system is operating at its highest clock rate with all devices powered. The "full-on" power consumption shown in Fig. 2 corresponds to this mode of operation for a notebook computer.
- **Power-Managed:** This is the normal operating mode of a notebook computer with power-management features enabled. Unused devices are powered down and powered devices are operated at a reduced clock rate except when higher-frequency operation is needed; e.g., when input requiring execution is being received from the keyboard or peripherals, or when the LC display screen is being redrawn. The "power-managed" case of Fig. 2 corresponds to this mode of operation; for the example shown there, about 25% power savings is achieved compared to full-on operation.
- **Standby:** Clocks to the system are stopped, so the system is inactive and dissipates no ac power. The system can respond to clock startup instantaneously. For a notebook computer in standby mode with LC display backlight off and hard disk drive spun down, the power savings in this case is typically 75% [36], similar to the case shown in Fig. 2.
- **Suspend:** The system is inactive and takes some number of cycles to become active, but when resumed the activity commences where it left off when it was suspended. An example would be a system which is depowered except for the DRAM which is maintained in retention mode, with the state of the system stored in it. In many notebook computers, suspend mode is entered when the LC display panel lid is closed. The system turns off the processor, disk drives, and standard I/O, greatly reducing power consumption. When the lid is opened, full operation is resumed where it was left off. Power savings in this case can be 99% or more; enough to allow many days or weeks of battery life.
- **Hibernation:** Many modern systems have what is referred to as a hibernation mode. If the system is left on but inactive for a user-determined length of time, it will automatically save all memory contents to a file on the hard disk and turn itself off, including the DRAM. When it is restarted (e.g., a year later) it will return exactly where it left off. The restart time for hibernation mode is longer than for suspend/resume, but shorter than a full reboot.
- **Off:** The system is shut down, no power is dissipated, and its active programs and data are not retained. To resume from an off state one performs a reboot, usually by rebuilding state and data structures with comprehensive checking.

Each level of activity is characterized by a power consumption, a service delay, and an energy cost for transitions to the level above or below it. Switching between the various levels of activity can be done under hardware control (e.g., the hard disk may have a user-settable timeout which depowers the disk drive when it is not used for a predefined period of time), or it can be done under software or firmware control (the instruction stream is monitored and analyzed and portions of a system are depowered when they are not needed). At the lower levels of firmware, timers detect that a subsystem has been idle and initiate a transition to a lower level of activity; at higher levels in software, policies are set, of which idle time thresholds are the simplest examples. A successful policy achieves the response time of the higher level of activity at an average power dissipation characteristic of the lower level of activity. This is only possible if the activity is characterized by long periods of idleness separating periods in which the rapid response is needed.

This ideal picture exaggerates the degree of "management" presently exercised. In principle, a power-optimized application could request a service in time for it to awaken from a deep sleep and be ready when needed. In practice, applications today are completely unaware of the needs of power-managed services, and policies such as idle thresholds are extracted from the user when the system is first set up, then left mostly unchanged during its use.

Two trends are improving the effectiveness of power management. One is a trend to fine-grained objects that can be separately managed, as a consequence of distributing power regulation across different chips, across different voltage levels present, and perhaps ultimately down to the individual logic blocks. The second is the use of circuit design approaches in which the blocks only consume power when they have useful work to do. Once these steps are taken, software design at higher levels can focus on avoiding calling for costly services that may not be needed.

It is also worth noting that the effectiveness of power management will be greatest in systems in which the

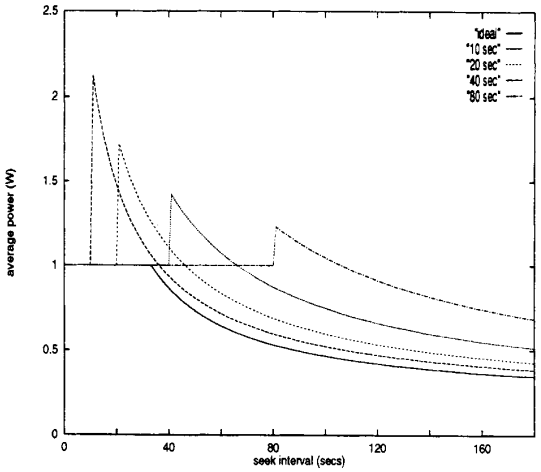

Fig. 14. Model calculations of average power dissipation for a 2.5-in hard file as a function of interval before seek. Results assuming a timer-based power-management strategy with timeout settings ranging from 10 to 80 s (dashed curves) are compared with an optimal strategy which gives a lower bound for average power dissipation (solid curve).

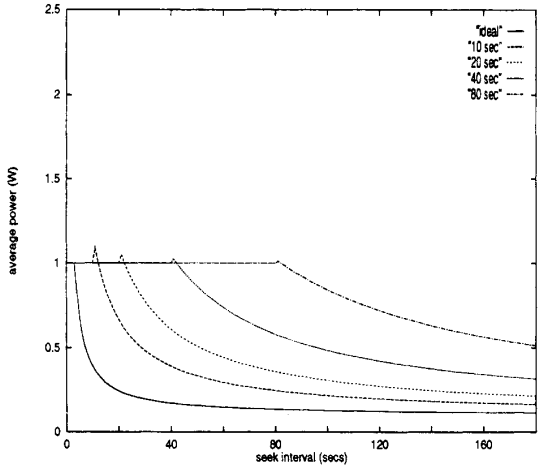

Fig. 15. Model calculations of average power dissipation for a 1.8-in hard file as a function of interval before seek. Results assuming a timer-based power-management strategy with timeout settings ranging from 10 to 80 s (dashed curves) are compared with an optimal strategy which gives a lower bound for average power dissipation (solid curve).

major power-consuming elements can be operated with small "duty factor" in most applications. In systems with monochrome or reflective (nonbacklit) displays, much of the power consumption is in electronics and peripherals which can likely be operated with low duty cycle in most applications. Thus there is room for significant enhancement of battery life through the use of power-saving strategies of the types described earlier in this section. On the other hand, in color notebooks most of the power consumption is in the color LC display backlighting. Today, the duty factor of the display will probably be near unity in most applications. Thus without innovation in power-saving display modes, the fractional improvement in battery life achievable through power management in backlit color notebooks will likely be much less than in notebooks with monochrome or nonbacklit displays.

B. Simple Model of Power Management for Hard Disk Files

In order to demonstrate the application of power management principles, we now turn to a generic situation in device power management–controlling the power dissipated by a small hard file by selectively placing it into a sleep state when there is no work to be done. To simplify the analysis we consider a two-state model in which the file is either active with the disk idling (idle state), or inactive with the disk no longer spinning (sleep state). The parameters characterizing this model are P_A, the idle state power consumption; P_B, the sleep state power consumption; τ_B, the time required to awake from sleep to idle state; and P_{BA}, the power required to spin up the disk and return to idle state from sleep state.

Using advertised parameters of small files on the market today (similar to the parameter values given in Table 2), we see that the industry is making steady progress in reducing sleep power, and achieving dramatic reductions in the spin-up time and power requirements, because of the smaller, lighter disks coming into use. For a 2.5-in hard file, such as is used in today's notebooks, $P_A = 1$ W, $P_B = 0.2$ W, $\tau_B = 7$ s, and $P_{BA} = 4$ W. A 1.8-in file from the same vendor shows no change in the idle-state power, P_A, but offers $P_B = 0.1$ W, $\tau_B = 1$ s, and $P_{BA} = 3$ W. Thus there has been a two-fold reduction in the sleep power of the smaller disk, but about a 10-fold reduction in the energy required to spin up. The objective of power management is to offer effectively the power dissipation of the file at sleep, P_B, while hiding the delay τ_B by incurring it as infrequently as possible or by clever scheduling. But we will see that even optimal strategies cannot fully accomplish this, while ineffective or ill-tuned power management strategies could have the effect of increasing both power dissipation and seek delay.

Figs. 14 and 15 show the average power dissipation during seek intervals ranging from a few seconds to two minutes for these two simplified models of hard files. We analyze two management strategies. The optimal strategy is to put the file in its sleep state at the conclusion of each seek, then turn it on just in time to finish spinning up when the next seek is required. If the amount of energy saved by spinning down the file is less than the spin-up energy cost ($P_{BA}\tau_B$), the file remains spinning but idle until the next seek. Obviously, this is an unrealistic strategy, but it provides a useful lower bound, shown by the solid lines in Figs. 14 and 15. Note that there is no power savings for spinning down the 2.5-in file (Fig. 14) until the interval between seeks exceeds about 35 s, but the lighter 1.8-in file (Fig. 15) can usefully be spun down during seek intervals as short as 4 s.

The power management strategy that has been used in practice is timer-based. If the file has no activity for some period of time, t_0, it is placed in the sleep state and must spin up when the next seek occurs. In this strategy, both the spin-up energy and the cost of remaining in idle state while deciding that there is no work to be done are costs offsetting the savings of placing the file in sleep state. In Fig. 14, the spin-up energy causes a big increase in the power dissipated when a seek request occurs just after the timeout interval. This spike is greatly reduced in the 1.8-in file (Fig. 15). But in both files the power dissipation averaged over a seek interval of up to several minutes is significantly greater than the lower bound because of the cost of detecting inactivity. In Figs. 14 and 15 the dashed curves show average power dissipation for timeout intervals of 10, 20, 40, and 80 s. Setting the timeout to a long interval (as owners' manuals for many notebook computers advise) will waste the advantages of the newer 1.8-in files, and would only be advisable if fast seek response is essential to an application. Software such as file caches may tend to produce a workload with seeks tightly grouped as an application accesses a new piece of data, separated by longer pauses in disk activity. We see from Figs. 14 and 15 that this will have favorable consequences for power consumption.

To go beyond Figs. 14 and 15 to estimate the power consumption during realistic use requires activity measurements that have not been done extensively. We have some data from measurements on office PC's that suggests that in a typical hour, half an hour will go by with no more than one seek per 10 min, about 15 min will consist of rapid seeks, less than 10 s apart, and that in the remainder of the time there are seek intervals of 10 s to a few minutes. Making up a plausible data set for an hour's activity with these characteristics, we obtain the estimates of energy consumption for the two files, shown in Fig. 16. (Note that the extra power required to read and write is not included in this analysis, since it is independent of power management strategy.) To appreciate the energy units, it helps to recall that one AA-size alkaline battery contains about 1 Wh, while a rechargeable 0.6-kg NiCd or NiMH battery holds about 30 Wh. The lower bound, following the ideal optimal strategy, is an average energy consumption of 0.52 Wh for the 2.5-in file, while timer-managed consumption varies from 0.6 to 0.75 Wh, and would be higher if timer settings greater than 5 min were used. The relative penalty of the timer strategy is roughly 25%. For the 1.8-in file, the timer-managed energy consumption ranges from 0.4 to 0.7 Wh for timer settings from 10 to 300 s, while the lower bound is 0.33 Wh. The relative penalty ranges from 25% to almost 100%!

For the newer, lighter files, the best naive power management strategy seems to be to set a very short timeout interval, or even to turn off the file after every seek, and accept occasional slow seek response times.

This analysis can be improved in some obvious ways. Files under development have additional states, such as a standby mode with a fast, low-energy recovery. Also, the energy cost and delay to bring a disk back to full speed can potentially be significantly reduced when the sleep time is brief enough that the disk is only partially spun down during this time period. Finally, more sophisticated management strategies, perhaps based on spotting patterns of disk activity, should be explored. But all these extensions first require systematic gathering of better workload data.

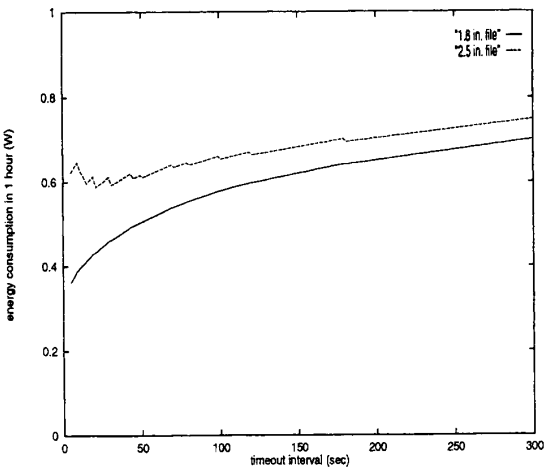

Fig. 16. Estimated energy consumed in 1 h, as a function of timeout setting (0–5 min). Upper curve is the model for the 2.5-in file, lower curve describes the 1.8-in file.

C. Summary

Power management techniques are used to conserve power by reducing power levels in parts of the system when they are not needed. Power management architectures have been defined, but in general are not yet being fully exploited. In principle, applications could request services in time for them to awaken and be ready when needed. The trends toward finer-grained power-managed blocks and new circuit design approaches which reduce unneeded switching activity should enhance the effectiveness of power management. The effects of power management on battery life will be greatest in those systems in which the major power-consuming elements can be operated with reduced "duty factor" in most applications. A simple model of power management for small disk files demonstrates how power management strategies can be defined, but the development of optimal strategies for particular cases will require measurements which give a better understanding of activity patterns.

IX. Conclusions

The advent of and market demand for battery-powered portable systems is driving the development of many low-power technologies. Technology enhancements are making it possible to capture increasing function and performance in small, highly portable systems with long battery life. Added capabilities such as speech and handwriting recog-

nition and wireless communications capability will draw new classes of users to the portable marketplace.

The term "portable computer" covers a broad range of system types, including palmtop and notepad computers, sub-notebook and notebook computers, and larger laptop and ac-powered units. Of these, the greatest market growth so far has been for the A4-format battery-powered notebook computer. Today's 486-generation notebook computers with a monochrome display and a 30-Wh battery typically weigh about 2.5–3 kg, are about 4 cm thick, and dissipate about 8 W of power when running at full speed with HDD and LC display on, thereby having about 3–4 h of battery life in continuous operation. There has been significant improvement in battery life of notebook computers in the last few years. This is attributable in part to the introduction of 3.3-V CMOS components, and, in the case of color notebooks, to very significant reductions in the amount of power required for backlighting. Almost all of the increased battery life is due to reduced power dissipation, as battery energy capacities have not changed appreciably.

In the future, portable systems will continue to evolve in the direction of increased capabilities and battery life with decreased size, weight, and cost. Two key directions of evolution for the A4 notebook computer are envisioned. One direction is the evolution of the notebook toward reduced size and weight at "constant" function and performance, leading toward notepad and palmtop computers with significant computing capability. The second direction, which we have chosen to emphasize in this paper, is the incorporation over time of greater performance and function in A4-format notebooks while continuing to improve portability (decreased weight and thickness) and battery life. By about 1997, it should be possible to produce A4 active-matrix color notebooks with 100-SPECint92 performance. If Li-polymer batteries become available in that timeframe, then greatly enhanced battery life for such a system should be possible, even with batteries which weigh much less than today's batteries.

The power budget of future color notebooks will differ in important respects from today's power budgets. The power dissipation associated with CMOS logic and memory will likely take up a smaller fraction of the total system power than in today's system, even with the greatly enhanced computing capability projected for future systems. This is because CMOS electronics offers tremendous potential for power reduction, primarily through voltage reduction and device scaling to smaller ground rules, and also through design optimization for low power. The effective utilization of the added computing capability will entail the incorporation of greater communications capability. It is therefore expected that communications capability will become a relatively larger power consumer in these future systems. LC display backlight power dissipation will likely take up a similar or greater fraction of the total power when compared to today's systems, primarily because there are few unexploited opportunities for large power reductions in this area. LC display backlight power has been reduced significantly for color displays over the last few years, but further large reductions will require breakthroughs in light management. It is anticipated that HDD power reductions of about 50% should be possible in this timeframe, primarily through miniaturization and electronics power reduction. Overall, even with high-energy-capacity Li-polymer batteries, the total power dissipation of future color notebook computers will have to be reduced from today's values if significantly enhanced portability and battery life are to be simultaneously achieved.

REFERENCES

[1] J. F. Freiman, "Portable computer power sources," in *Proc. 9th Annu. Battery Conf. on Appl. and Advances, IEEE/Cal. State*, sponsored by IEEE and Dept. of Elec. Eng., Cal. State Univ.–Long Beach, Jan. 1994, Long Beach, CA, pp. 152–158.

[2] E. Vittoz, "Micropower techniques," in *Design of MOS VLSI Circuits for Telecommunications*, Y. Tsividis and P. Antognetti, Eds. Englewood Cliffs, NJ: Prentice-Hall, 1985.

[3] S. Mahli and P. Chatterjee, "1-V microsystems–scaling on schedule for personal communications," *IEEE Circ. and Syst.*, pp. 13–17, Mar. 1994.

[4] M. Culbert, "Low power hardware for a high performance PDA" in *IEEE Spring COMPCON'94 Dig. of Papers*. Los Alamitos, CA: IEEE Comp. Soc. Press, pp. 144–147, Mar. 1994.

[5] R. Nagaraj, "Designing a 3 V notebook with the Intel 386™ SL architecture," *Proc. 1992 Silicon Valley Personal Compu. Design Conf.*, vol. I, pp. 307–315, Santa Clara, CA, Aug. 1992.

[6] L. R. Rabiner, "Applications of voice processing to telecommunications," *Proc. IEEE*, vol. 82, pp. 199–228, 1994.

[7] N. K. Yeung, Y.-H. Sutu, T. Y.-F. Su, E. T. Pak, C.-C. Chou, S. Akki, D. D. Yau, and R. Lodenquai, "The design of a 55SPECint92 RISC processor under 2W," *ISSCC Dig. Tech. Papers*, pp. 206–207, Feb. 1994.

[8] D. Pham, *et al.*, "A 3.0 W 75SPECint92 85SPECfp92 superscalar RISC microprocessor," *ISSCC Dig. Tech. Papers*, pp. 212–213, Feb. 1994.

[9] J. Stockton, "Portable electronic storage systems," *IEEE Micro*, pp. 69–76, Feb. 1994.

[10] A. Stratakos, R. Broderson, and S. Sanders, "High-efficiency low-voltage dc-dc conversion for portable applications," *Proc. 1994 Int. Workshop on Low Power Design*, sponsored by ACM-SIGDA and IEEE-CAS, Napa, CA, Apr. 1994, pp. 105–110.

[11] R. Dobkin and S. Pietkiewicz, "Power supply considerations for PCMCIA compliant systems," *Proc. 1993 Silicon Valley Personal Compu. Design Conf.*, Santa Clara, CA, July 1993, pp. 351–364.

[12] "Device cools Pentiums in notebooks," *PC Week*, vol. 10, no. 4, p. 25, Dec. 13, 1993.

[13] W. E. Howard, "Thin-film transistor/liquid crystal display technology–an introduction," *IBM J. Res. Dev.*, vol. 36, pp. 3–10, 1992.

[14] P. Pleshko, "Overview and status of information displays," *Society for Information Display (SID) Seminar Lecture Notes*, vol. 1, Seminar M-0, pp. 3–76, May 18, 1992.

[15] "Sharp color TFT-LCD panel operates on 50 mW," *Electron.*, vol. 67, no. 8, p. 8, Apr. 1994.

[16] B. Davari, and R. H. Dennard, "CMOS scaling, the next ten years," *Nikkei Microdevices*, pp. 142–154, Sept. 1994, and *IBM Research Report RC-19682*.

[17] B. Davari, R. H. Dennard, and G. Shahidi, "CMOS scaling for high performance and low power–the next ten years" *Proc. IEEE*, this issue.

[18] K. Itoh, K. Sasaki, and Y. Nakagome, "Trends in low-power RAM circuit technologies," *Proc. 1994 IEEE Symp. on Low Power Electron.*, Oct. 10–12, 1994, San Diego, CA, pp. 84–87.

[19] J. M. Rabaey, "Low power digital design," *1994 IEEE Int. Symp. on Circ. and Syst.*, London, UK, May/June 1994, Circuits and Systems Tutorial Vol., Tutorial 8, pp. 373–386.

[20] N. Weste and K. Eshragian, *Principles of CMOS VLSI Design—A Systems Perspective*. Reading, MA: Addison-Wesley, 1985, ch. 4.

[21] R. H. Dennard, F. H. Gaensslen, H. N. Yu, V. L. Rideout, E. Bassous, and A. R. LeBlanc, "Design of ion-implanted MOSFETs with very small physical dimensions," *IEEE J. Solid State Circ.*, vol. SC-9, pp. 256–268, 1974.
[22] D. Liu and C. Svensson, "Trading speed for low power by choice of supply and threshold voltages," *IEEE J. Solid State Circ.*, vol. SC-28, pp. 10–17, 1993.
[23] L. Geppert, "Not your father's CPU," *IEEE Spectrum*, pp. 20–25, Dec. 1993.
[24] A. Chandrakasan, S. Sheng, and R. W. Broderson, "Low-power CMOS digital design," *IEEE J. Solid State Circ.*, vol. SC-27, pp. 473–484, 1992.
[25] "New LAN gear snaps unseen desktop chains," *Data Commun.*, pp. 75–80, Mar. 21, 1994.
[26] D. J. Goodman, "Trends in cellular and cordless communications," *IEEE Commun. Magazine*, vol. 29, pp. 31–40, June 1991.
[27] "Test 21: Schnurlos telefone. meter freiheit," *Funkschau*, pp. 10–17, Feb. 4, 1994.
[28] "EDN's DSP chip directory," *EDN*, Sept. 1993.
[29] J. M. Harker, D. W. Brede, R. E. Pattison, G. R. Santana, and L. G. Taft, "A quarter century of disk file innovation," *IBM J. Res. Dev.*, vol. 25, pp. 677–689, 1981.
[30] *1993 Disk/Trend Rep.-Rigid Disk Drives*, Oct. 1993.
[31] T. Parrish and G. Kelsic, "Dynamic head loading technology Increases drive ruggedness," *Computer Tech. Rev.*, pp. 51–55, Winter 1992.
[32] M. Sri-Jayantha, H. Tokumitsu, and A. Aoyagi, "HDA configuration for minimum idle mode power by optimum read/write head positioning," *IBM Tech. Disclosure Bull.*, vol. 36, no. 06A, pp. 181–184, June 1993.
[33] K. Sun, "Laptop hard disk drive specs require design optimization today," *Computer Tech. Rev.*, pp. 49–54, Spring 1990.
[34] F. Goodenough, "Lone 3-V rail powers IC disk-drive read channel," *Electron. Design*, pp. 43–48, June 25, 1992.
[35] D. H. Straayer, "Pointing solutions for notebook computers: The isometric alternatives," *Proc. 1993 Silicon Valley Pers. Compu. Design Conf.*, Santa Clara, CA, July 1993, pp. 553–558.
[36] S. C. Ellis, "Power management in notebook computers," *Proc. 1991 Silicon Valley Pers. Compu. Design Conf.*, Santa Clara, CA, July 1991, pp. 749–754.
[37] S. Datta, "Advanced power management (APM) BIOS implementation on the Intel 386™ SL microprocessors," *Proc. 1992 Silicon Valley Pers. Compu. Design Conf.*, vol. 2, Santa Clara, CA, Aug. 1992, pp. 319–324.

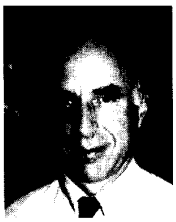

Erik P. Harris (Member, IEEE) received the Bachelor of Engineering Physics degree from Cornell University, Ithaca, NY, in 1961 and the M.S. and Ph.D. degrees in physics from the University of Illinois in Urbana-Champaign in 1963 and 1966, respectively.

In 1967 he joined IBM at the Thomas J. Watson Research Center in Yorktown Heights, NY, where between 1967 and 1988 he held various technical and management positions in microelectronics, primarily in Josephson technology and in GaAs MESFET technology. In 1988 he joined the IBM Advanced Semiconductor Technology Center in East Fishkill, NY, where he was manager of lithography. From 1990 to 1992, he was a member of the Corporate Technical Strategy Development staff at IBM Headquarters in Armonk, NY. He returned to the T. J. Watson Research Center in 1992, where he is currently manager of exploratory low-power systems.

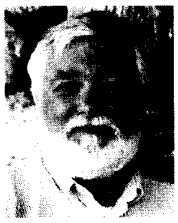

Steven W. Depp (Member, IEEE) received the Ph.D. in physics at the University of Illinois in 1972.

His first job was with the Los Alamos Scientific Laboratory where he worked on high temperature electronics and advanced telemetry systems. In 1977 he joined IBM as a Research Staff Member at the San Jose Research Lab. His work there centered on advanced display technologies (such as thin film transistor/liquid crystal displays) with particular emphasis on the materials and device physics. In 1982, he became a staff member of the Director of Research at the T. J. Watson Research Center, Yorktown, NY, where he monitored the areas of Input/Output and Storage. Following that assignment, he remained in Yorktown and became leader of the Exploratory Devices and Processes group in the Gallium Arsenide program. In 1984, he transferred to the Input/Output Technology Department and initiated the Input/Output Systems Department. The following year he was promoted to Manager of the Interactive I/O Technologies Department, which was responsible for display technologies as well as I/O systems. He helped initiate the current program in active matrix liquid crystal displays and in 1990 was promoted to Department Group Manager of I/O Technologies. His current position is Director of the Subsystem Technologies and Applications Laboratory.

William E. Pence (Member, IEEE) received the B.S. degree in physics from the University of Virginia in 1984, and the Ph.D. degree in electrical engineering from Cornell University in 1989. His dissertation was on electrical performance of multi-chip modules and modeling tools for system optimization.

Since 1989, he has been a Research Staff Member at the IBM T. J. Watson Research Center in Yorktown Heights, NY. He has worked in applications of electronic packaging to low-end systems, design tools for EMI suppression, and prototyping of advanced personal systems. In 1993, he became senior manager for Communication Technology, where he directs programs in advanced circuit and subsystem designs for fiberoptic, wireless (IR & RF), and broadband (ATM) communications.

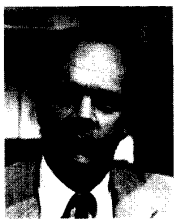

Scott Kirkpatrick (Senior Member, IEEE) received the A.B. degree from Princeton University and the Ph.D. degree from Harvard, both in physics.

He has been Research Staff Member at the IBM Research Division's T. J. Watson Research Center, Yorktown Heights, NY, since 1971. His initial research was in the physics of amorphous or disordered materials, such as spin glasses. This led to work on design automation for VLSI and optimization by simulated annealing. He has managed activities in the design, measurement, and analysis of hardware for advanced personal computers in the Computer Sciences Department. He has also been a Visiting Professor at the State University of New York at Stony Brook, the University of Paris, and the Hebrew University in Jerusalem.

Dr. Kirkpatrick is a fellow of the American Physical Society and the American Association for the Advancement of Science.

M. Sri-Jayantha (Member, IEEE) was born in Jaffna (Tamil Eelam), Sri Lanka, on January 24, 1954. He received the B.Sc. degree (first class) in 1976 from the University of Sri Lanka and the M.S. degree in 1979 from The Pennsylvania State University, both in mechanical engineering. He later received a Guggenheim fellowship to Princeton University where he completed the M.A. and Ph.D. in mechanical and aerospace engineering.

During 1977–1979, he did research on electric vehicle propulsion system modeling and simulation at the Pennsylvania Transportation Institute. At Princeton University he conducted research in the areas of aircraft instrumentation, digital data acquisition and aerodynamic parameter identification using a Schweizer 2-32 sailplane equipped with Z80 based microcomputer system. He has been with IBM for about 10 years and is a Manager for the Disk System Integration group. His current research is in the area of ruggedization and servo control of electromechanical systems with a particular focus on disk mechanisms for portable computer applications. He holds several patents issued or pending.

Dr. Sri-Jayantha has received Outstanding Innovation and Outstanding Technical Achievement Awards from IBM. He was a recipient of a UNESCO team gold medai from University of Sri Lanka and an IBM fellowship for the year 1982–1983.

Ronald R. Troutman (Fellow, IEEE) received the B.S. degree in electrical engineering from the Massachusetts Institute of Technology in 1962 and the M.S. and Ph.D. degrees in electrical engineering from New York University in 1963 and 1966.

He served as Project Officer in the Electronics Branch of the Office of Naval Research from 1966 to 1968. In 1969 he joined the IBM Corporation at Essex Junction, VT, where he participated in the development of many semiconductor technologies for memory and logic applications and was instrumental in initiating IBM's current CMOS technology. He has made many contributions to random access memory design, to the understanding and modeling of short channel effects, of subthreshold behavior, and of hot electron phenomena in IGFET's, and to the modeling and control of latchup in CMOS. In 1987 he became Manager of the Thin Film Transistor Materials and Devices group at IBM's T. J. Watson Research Center, where he has managed a variety of research activities on TFT liquid crystal displays, including the development and commercialization of the TFT/LCD array tester. He is the author or co-author of over 60 publications, including a book on CMOS latchup, and is the holder or co-holder of 14 patents.

Dr. Troutman is a member of Sigma Xi and the New York Academy of Sciences. During 1983–1984 he was a Fellow at the Center of Advanced Engineering Studies at MIT.