Adaptive Power Management for Mobile Hard Drives

Adaptive Battery Life Extender™

—

A Self-Managed Approach to Saving Energy

April, 1999

IBM Corporation
Storage Systems Division
5600 Cottle Rd
San Jose CA 95193
(408) 256-8000

ADAPTIVE POWER MANAGEMENT FOR MOBILE HARD DRIVES

1. INTRODUCTION

Notebook computers have shown tremendous improvements within the last years. Today there exist an amazing variety of models, from light weight ultra-portables to desktop replacements. Notebooks have increased performance, larger and higher resolution screens, more memory and greater storage capacity (up to 14 GB HDD are available today) compared with a few years ago. They are also sporting a larger number of peripherals, including high-speed CD-ROM drives, DVD drives, integrated connectivity, fax/modem functions and a plethora of plug-in PC cards. Once esoteric functions such as wireless communications and GPS functions are becoming more widespread.

All of these wonderful new capabilities have placed increasing pressure on system battery life. Advances in battery design and the use of more advanced power saving technology have allowed battery lifetime to remain in the 1 – 4 hour range under normal usage conditions. However, every system component needs to improve power consumption just to stay even.

Increasing system battery life therefore requires an unconventional approach to power management. IBM, the leader in mobile hard drive technology, was the first company to introduce adaptive power management for hard drives in 1995. IBM continues to advance the state of the art in low power technology, through cooperation between the Research and Storage Systems Divisions.

What's New?

Adaptive Battery Life Extender™ version 2 introduced enhanced intelligence and pattern recognition, and self-management of Standby mode. This intelligence allows each mode to examine only the appropriate accesses. This allows for the use of each mode to be independently optimized for the current activity. The self-managed Standby feature can be enabled and the behavior adjusted via ATA/ATAPI-4 commands. The activity monitoring has been further improved in version 3 to provide better characterization of the access patterns. The pattern recognition has been enhanced to identify more common types of activity, such as word processing. The drive can detect certain "signatures" which are characteristic of such applications, and can use this information to adjust its behavior.

The result is that IBM offers the highest capacity mobile hard drives in the industry with the lowest energy consumption, providing the longest system battery lifetimes.

2. CONVENTIONAL POWER MANAGEMENT

All mobile hard drives today support some form of power management. Most drives follow the conventional approach of defining a number of specific power modes. The power modes save energy by reducing performance. A typical mode configuration is shown in Figure 1.

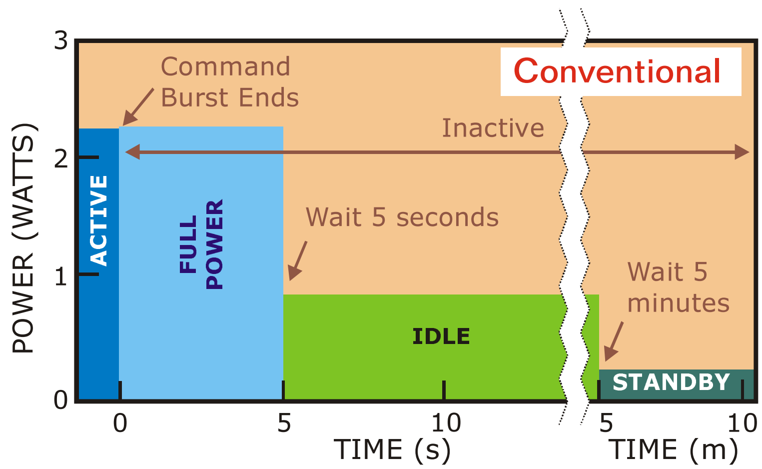

Figure 1. Conventional hard drive power management. The vertical axis shows the typical power consumption in watts. The horizontal axis shows the time elapsed since a command burst ends, and the time scale is compressed on the right. The inactive interval indicates that the drive is not receiving any commands from the host system. The areas of the rectangles represent the energy consumed in each mode.

Power modes

The typical hard drive operates in one of four modes:

Active mode: In Active mode, the hard drive reads, writes, seeks and processes host commands. Typical hard drives consume 2.0-2.5 watts in Active mode.

Idle mode: Idle mode reduces power consumption by turning off some of the drive electronics at the expense of a short recovery time. In one common implementation, the head is moved to a parking position, and the servo tracking function either is turned off, or operated at a reduced level of control. The disk remains spinning, the interface electronics remain ready to accept commands, and the drive will return to Active mode when a new command is received. The drive power is reduce to slightly less than 1 watt. The drive will require about 40 msec to return to Active mode.

Standby mode: In Standby mode, the spindle motor is stopped, and most of the electronics are powered off. Power consumption is lowered to the range of 0.3 watts, but recovery time from Standby mode to Active mode is increased to a few seconds. The interface electronics remain ready to accept commands, and the drive will return to Active mode when a new command is received. The drive will require between 1.5 and 5 seconds returning to Active mode. (Desktop and server drives may require up to 30 seconds recovering from Standby mode.)

Sleep mode: Sleep mode is entered by a specific command, and is used for long periods (hours) of system inactivity. All electronics are powered off except those needed to respond to a wake-up command, typically requiring about 0.1 watt. Recovery time is several seconds.

Mode control

Conventional hard drives don't enter power save modes immediately following the completion of a command. Rather, they wait a while to see if further commands are forthcoming. The typical approach is to wait for a fixed period. Most drives will remain in Active mode for 5 seconds after a command is received before entering Idle mode. They approach for entering Standby mode is similar, but the waiting time is much longer, typically 5 minutes. This long time is due to a number of factors, including the user-perceptible startup delay, and to limit the number of times the heads land on the disk. The host computer or the end user may adjust the Standby entry time, while the idle entry time tends to be fixed. This provision allows for some tailoring of the drive behavior.

Weaknesses

The conventional approach clearly leaves a lot to be desired. Even a cursory observation of hard drive activity will show that the fixed entry time approach is not optimal. The conventional approach works on the assumption that if there is a period of inactivity of at least the fixed time setting, then entering Idle mode or Standby mode will save energy. However, this assumption is very weak. While it can detect that the user has left to get a cup of coffee, it can't really determine anything about the access patterns. In fact, it can make some pretty egregious errors: if the Standby time is set slightly shorter than an application autosave time, then the drive will waste a lot of energy jumping in and out of Standby mode. (The fixed time method doesn't account for energy expended in switching between modes.) It is very difficult to save energy without knowing this information. Clearly, there must be a better way.

3. ADAPTIVE BATTERY LIFE EXTENDER™

The goal of Adaptive Battery Life Extender is to reduce the hard drive power consumption *during actual user workloads*. It is possible to dramatically reduce power consumption with an adaptive control design, while maintaining HDD performance. IBM's Adaptive Battery Life Extender technology improves the conventional approach in two important ways. First, Adaptive Battery Life Extender employs a more sophisticated set of power modes; second, it intelligently manages the transition between these modes. This technology also takes advantage of IBM's load/unload design, which prevents the heads from landing on the disk surfaces. This represents a dramatic departure from conventional power management, which is unable to effectively use the power save modes during user activity.

More power levels

Adaptive Battery Life Extender drives operate using the following modes:

Active mode: Same as conventional.

Performance Idle mode: Performance idle mode is entered immediately following the completion of command processing in Active mode. Unlike conventional Idle mode, there is no entry delay. In Performance Idle, full servo performance is maintained, but some of the electronics are powered down. Subsequent commands are processed with no delay. Performance Idle power consumption is about 1.5 – 2.0 watts. The duration of this mode is self-managed, as described below.

Fast Idle mode: In Fast idle mode, power consumption is similar to conventional Idle mode. The head is moved to parking location parked and the servo control turned off.

Fast idle power is in the 0.8 watt range. Recovery time to Active mode is about 40 ms. The entry and duration of this mode is self-managed.

Low Power (LP) Idle mode: In LP idle mode, the power consumption is reduced by 25% compared with Fast idle. The heads are unloaded from the disk, reducing power consumption to the 0.6 watt range. In LP idle mode, the drive has improved shock tolerance since the heads are not flying over the disk surfaces. Recovery time to Active mode is about 400 ms. The entry and duration of this mode is self-managed.

Standby mode: Similar to conventional, except entry may be self-managed adaptively by the drive. The spindle motor is stopped, and most of the electronics are powered off. Recovery time is less than 2 seconds.

Sleep mode: Same as above.

Adaptive Self-management

Adaptive Battery Life Extender determines the appropriate mode based on the actual access pattern, and the internal level of drive activity.

Figure 2 illustrates a typical power sequence using Adaptive Battery Life Extender technology. Once the command burst ends, the drive immediately enters Performance Idle mode. At this point, the drive performs calculations to determine the most appropriate power save mode, based on the characteristics of the command burst, the command history, the desired performance level, and the energy costs associated with each mode.

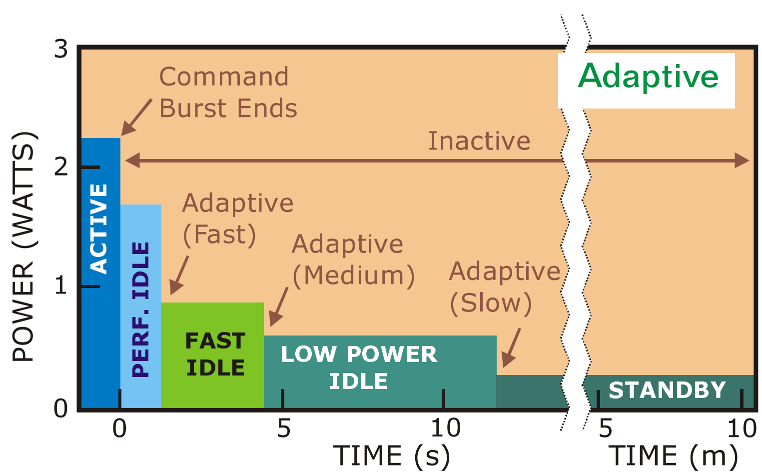

Figure 2. The power consumption is shown as a function of time, for a typical access pattern. The active rectangle shows the power consumption while actively processing commands. Time 0 indicates the completion of the last command in a sequence. The next rectangles illustrate a progression through various power levels. The horizontal time scale is shown compressed on the right. The mode durations shown are examples. Adaptive Battery Life Extender determines the appropriate amount of time to spend in each mode based on the access activity.

The drive enters low power modes rapidly when it is confident that the command burst has ended, and more slowly (or not at all), when it determines that the command activity has not ceased. As shown in Figure 2, on average the drive will enter Fast Idle in less than 1 second. Low Power Idle is typically entered in a few seconds. If enabled, Standby will be entered in less than 20 seconds. The specific sequence and times to enter the modes will vary according to the drive activity. The net result is that disk drives employing Adaptive Battery Life Extender will generally enter power saving states much more quickly than conventional drives, reducing power consumption and increasing system battery life.

The key to the power efficiency of Adaptive Battery Life Extender is the intelligent management of transitions between the various power modes. Adaptive Battery Life Extender dynamically manages the transitions based on the user's recent access pattern, instead of fixed times.

Software applications generally generate commands in bursts rather than randomly. This effect tends to be enhanced by system cache activity. Adaptive Battery Life Extender achieves power savings by exploiting the command burst patterns. The command frequency distribution is measured to characterize the activity. This information is maintained in a history buffer, and is used to judge how long the current burst of commands will continue. The algorithm calculates the probability that the current burst of commands is complete based on the statistics of the distribution. This information is coupled with the energy characteristics of the modes to determine which mode, if any, is best suited to the current conditions.

Each mode is treated independently. Therefore, the intermediate modes may be skipped at times.

Power / Performance Control

Reducing the power consumption of the hard drive (or any other system component) is useful only when it increases the amount of work that can be performed. For example, if the power consumption is reduced by 20% while the performance is also reduced by 20%, then nothing has really been gained. In such a case, it just takes longer to get the work done, but the amount of work completed remains unchanged. Therefore, a well-designed system should reduce the power consumption faster than it reduces the performance. A useful figure of merit is the amount of energy required to complete a task. The Joules used per Job, or JpJ, is one way to measure it. Therefore, in a well-designed system, the JpJ decreases when the power consumption is reduced. It may take longer to perform a task, but more tasks will be completed on the same battery.

Unfortunately, operating at the minimum JpJ isn't always suitable, due to the performance impact. Adaptive Battery Life Extender addresses this problem by using a parameter to adjust behavior between maximizing performance, and minimizing power consumption. This is called the power/performance control.

The Advantages of Self-Management

Self-management of power modes provides better power savings than can be achieved by external control. This arises from the detailed knowledge of the hard drive required to make a mode entry decision. Such a decision can't be based simply on the power consumption of a mode. It also requires knowledge of the energy required to enter and exit a mode, the current power level and the current drive activity. The drive activity can include drive-initiated operations as well, such as read ahead and write behind functions.

This information is only available to the hard drive itself. Therefore, the drive needs to manage its own power state transitions.

A self-managed hard drive can trade power consumption for performance, since it can measure both effects. This is not possible using the conventional approach, as neither the performance nor power impacts of a fixed time can be known in advance. A shorter time will actually increase the power consumption at some point, when the command intervals are close to the fixed time value. Here, the drive will use more energy returning to active than it would have saved during the low power mode. With self-management, the host system or user can still control the hard drive behavior, but is freed from the burden of trying to predict the impact of a fixed time setting.

Self-Managed Standby

IBM hard drives using Adaptive Battery Life Extender version 2 and higher supports self-managed Standby mode. The use of Standby mode is very advantageous, since it has such lowest power consumption (about half that of Low Power Idle). However, aggressive use of Standby is likely to have some impact on user performance, as recoveries from standby become more common. IBM's load/unload technology reduces wear on the disks from excessive takeoffs and landings.

4. RESULTS

Power Consumption

Benchmark studies consistently show a reduction in hard drive power consumption, and a resulting increase in system battery life with Adaptive Battery Life Extender. Such studies have shown that total drive power consumption can be reduced 21% to 40% compared to the same hard drive without Adaptive Battery Life Extender.

The power consumed by the drive will vary with user, workload and the power/performance setting. In this section, we consider 3 power levels. Level 0 is the maximum performance setting. Level 1 is an intermediate setting, with minimum power consumption without allowing Standby mode. Level 2 is minimum power consumption allowing Standby mode.

Figure 3 shows the results of a benchmark study of the drive power consumption using these settings. The difference between level 1 and level 0 is not pronounced here. This is primarily due to the power savings level 0 achieves without impacting performance. The power consumption is reduced by over 20% when Standby mode is used.

Benchmark testing shows that system battery life is increased from 2% to more than 8%. Figure 4 shows the impact of Adaptive Battery Life Extender settings on system battery life.

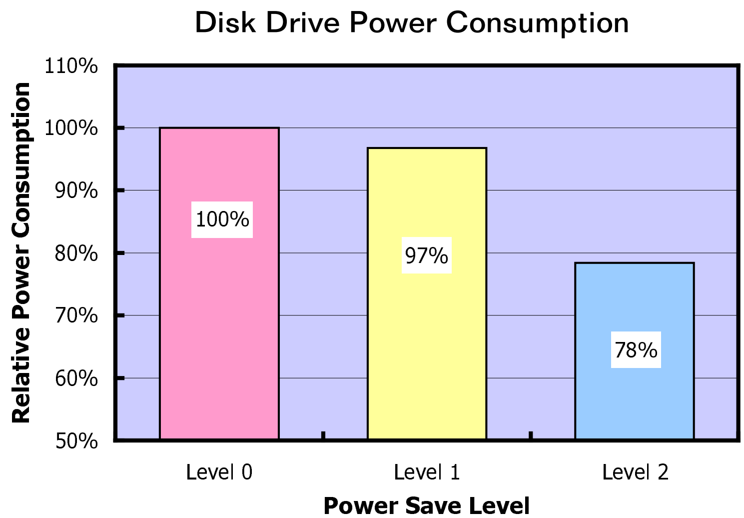

Figure 3. Hard drive power consumption vs. power/performance level. The vertical axis measures the power consumption relative to level 0, which is maximum performance. The power consumption decreases as the level increases. The power consumption was measured using Ziff-Davis Battery Mark 2. The power consumption decreases dramatically at level 2.

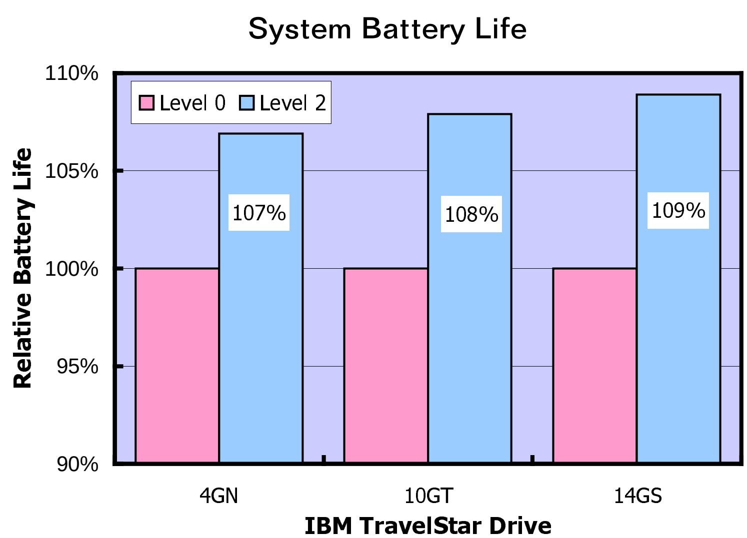

Figure 4. Effect of power/performance level on system battery life. The vertical axis measures the battery life for each drive at level 2 relative to level 0. The results are shown for 3 different drive models. These tests were made using Ziff-Davis Battery Mark 2. The measured lifetime increase is significant in each case.

Performance

Benchmark studies have shown that the performance of Adaptive Battery Life Extender in most cases is equal to or better than the same drive using conventional power management. The studies have shown that at the high performance settings, performance is almost always better than conventional power management.

The real test is how performance is impacted at the minimum power settings. The results of such a benchmark study are shown in Figure 5. The measured performance impact is less than 2% in all cases. This implies that at level 2, the Joules per Job ratio has decreased, and more computing can be accomplished on a fixed battery.

User Experience

The benchmark data may not reflect actual user experience, since they use the hard drive more intensively than a typical user workload. It is anticipated that the power savings will increase under real-world conditions, but that the performance impact will increase as well. This behavior occurs since power save modes are used more frequently in real world applications than in benchmarks. Empirical studies have shown that at level 2, the drive enters Standby mode an average of once per minute. From a worst-case perspective, the drive has lost 2/60, or 3.3% of throughput. However, the energy savings more than make up for this.

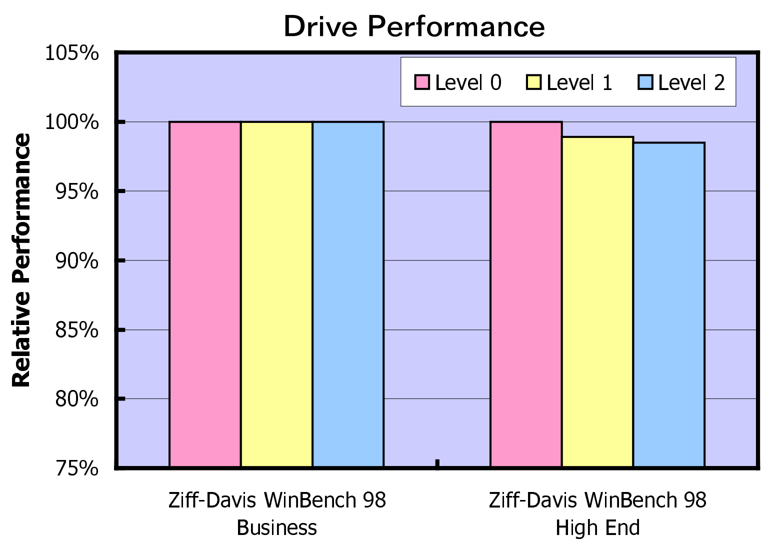

Figure 5. Hard drive performance impact of power/performance level. The vertical axis measures the drive performance at level 1 and 2 relative to level 0. Two different benchmarks were used, as indicated. The measured performance impact is minimal in all cases.

5. ACTIVATION OF POWER MANAGEMENT

IBM's TravelStar products with Adaptive Battery Life Extender version 2 or greater ship with a power/performance setting of level 1. This setting may be adjusted through the

ATA/ATAPI-4 advanced power management interface (8.37.11). This ATA interface provides a sliding scale between maximum performance and minimum power.

Control program

IBM has developed a control program for Windows 95 and 98 that allows the user to adjust the power/performance setting. This program is available free on the IBM alphaWorks web site at:

http://www.alphaworks.ibm.com

With this program, users can adjust the Adaptive Battery Life Extender power/performance behavior. The program can be set to automatically use the maximum performance setting whenever the computer is plugged in. Note: this program only functions with drives supporting Adaptive Battery Life Extender version 2 and higher. The user interface is shown in Figure 6.

Figure 6. Power management control program user interface.

ACPI and Device States

IBM mobile hard drives are shipped with the self-managed Standby feature turned off, due to a conflict between the Advanced Configuration and Power Interface (ACPI) specification and the ATA/ATAPI-4 specification. ACPI was developed with conventional power management in mind, and assumes that the operating system will control the power state of the disk drive (state-based control). However, as we have shown, there isn't enough information at the OS level to provide significant energy savings. We feel that ACPI should be modified to support self-managed disk drives, and to control them through the ATA/ATAPI-4 interface. IBM is working to change the ACPI specification so that the end user is provided with longest battery life possible.

Realization

The IBM control program is one example of how an OEM might integrate self-managed devices into a system. It may be preferred to integrate this function into the existing control structure, such as the storage device driver.

6. FUTURE DIRECTIONS

The HDD portion of overall system power consumption has been significantly reduced in the past few years, from 25% to less than 10% today. At the same time, drive capacity and performance have improved dramatically.

IBM is working continuously to further improve the power consumption of hard drives. For example IBM's new microdrive uses 3.3V technology.

More Device States

IBM is researching further enhancements to Adaptive Battery Life Extender. This is likely to result in the addition of new low power modes, more advanced algorithms and improved self-tuning. One result of this research is that the order of power mode entry can change under certain conditions. This situation can occur when one power mode becomes more energetically favorable than another mode. Situations such as this are easily handled in a self-managed system, but present serious difficulties for state-based control systems.