



SATA Power Management: It's Good to Be Green
April 8, 2009



Introduction

Due to the rising costs of powering and cooling computer systems and storage, power consumption has become an increasingly important concern in all computing markets. It is now a de facto requirement that system components, such as peripherals and host controllers, support some form of power management.

Power management typically includes removing or reducing power to circuitry or electromechanical elements during times when they are not being used, thereby reducing the overall power consumption of the system. Power management protocols have been included in many interface standards, including Serial Attached SCSI and Serial ATA.

The Serial ATA (SATA) interface is pervasive in the PC market and has a significant presence in server and enterprise computing environments. SATA is also ubiquitous in the battery-operated notebook computer market, in part due to its aggressive approach to power minimization. This article describes the industry-leading power management features of SATA products.

SATA Standards

Before delving into the details of power management, it is necessary to provide some background about SATA standards. The Serial ATA specification is developed by SATA-IO, the Serial ATA International Organization (<http://www.sata-io.org>). This document defines the physical interface, control/status protocol, and data/command transmission.

While SATA-IO specifies power management for the SATA interface itself, T13 specifies power management within the device as a whole. A second document is generated by the INCITS T13 (<http://www.t13.org>) standards organization. The AT Attachment - ATA/ATAPI Command Set (ATA8-ACS) standard defines the command set.

These documents address two different areas of SATA power management:

1. On the SATA interface, as defined by the SATA specification
2. Within a SATA product, as defined by the ATA8-ACS standard



Power Management of the SATA Interface

The SATA physical interface (PHY) consumes a significant portion of the total power budget of the host or peripheral electronics. SATA provides the means to place the PHY into reduced power modes. As portions of the PHY are shut down to conserve power, the ability of the SATA device to respond to commands is impacted. The SATA protocol therefore enables the tracking of the power modes of devices and makes allowances for the added latencies required to wake up from reduced power modes.

The SATA specification defines three SATA interface power modes:

1. PHY Ready (PHYRDY) - the SATA PHY is ready to send/receive data
2. Partial - the PHY is in a reduced power mode; exit time can be up to 10 microseconds
3. Slumber - the PHY is in a reduced power mode (lower power than Partial mode); exit time can be as much as 10 milliseconds

Entry into the Partial or Slumber power modes can be initiated by the host or by the SATA peripheral itself, typically after some period of inactivity.



The SATA PHY is awakened, or placed back into PHYRDY mode, upon receipt of a wake-up signaling sequence sent by the host or device. The exit time is the maximum time allowed by the SATA specification for the SATA product to enter PHYRDY mode after receiving a wake-up signal.

Power Management Within a SATA Product

Even greater power savings can be achieved by shutting down portions of a SATA device other than the PHY, such as spinning down the disks in a hard drive.

The ATA8-ACS standard describes four modes of power consumption for SATA products

1. Active – The device is fully powered up and ready to send/receive data.
2. Idle – The device is capable of responding to commands but the device may take longer to complete commands than when in the Active mode. Power consumption of the device in this state is lower than that of Active mode. If a hard drive is present, it is spun up.
3. Standby – The device is capable of responding to commands but the device may take longer (up to 30 seconds) to complete commands than in the Idle mode. Power consumption is reduced from that of Idle mode. If a hard drive is present, it is spun down.
4. Sleep – This is the lowest power mode. The device interface is typically inactive and, if a hard drive is present, the drive is spun down. The device will exit the Sleep mode only after receiving a reset. Wake up time can be as long as 30 seconds.

Here again, the tradeoff of increased power savings is a longer initial response time from the device. So spinning down a hard drive should only be done when it is known that the hard drive will not be accessed for a significant period of time.

SATA Power Management Controls

SATA allows PHY Power Management to be Host Initiated (HIPM) or Device Initiated (DIPM), thus providing the flexibility to optimize the SATA components for a wide range of usages and applications.

A SATA peripheral may contain a Standby Timer, which is programmed by the host, and will put the peripheral into Standby mode after a predetermined period of inactivity.

The host may also have the ability to put SATA peripherals directly into Idle, Standby, and Sleep modes, and may be able to confirm the current power management mode of SATA peripherals.

Power Management Applications

The simplest application of SATA power management is a hard drive in a notebook computer. The notebook is usually running power management software which allows the user to optimize the hard drive behavior for their application while operating on battery power. The software allows the user to select from multiple modes, to pick the best compromise between performance and battery life.

Newer desktop PCs also support SATA power management, typically spinning down the hard drive and putting the system into sleep mode after some period of inactivity.

Storage servers use multiple levels of power management, but the focus is typically on performance, so Sleep mode may not often be used, due to the extended wake up time.





At the high end of the power management spectrum are MAID systems (Massive Array of Idle Disks) which may consist of hundreds to thousands of hard drives and are used for storage of reference data that is only occasionally read and very seldom changed. The hard drives in a MAID system spend most of their time in Sleep mode, waking up only when the data is accessed.

SATA – The Low Power Solution

When selecting a data storage component such as a hard drive or host controller, the focus has traditionally been on specifications such as capacity and performance. More recently, however, power consumption has become just as important. To ensure that the system power consumption is minimized, take advantage of the power management capabilities provided by SATA technology.

