



US 20050223255A1

(19) **United States**

(12) **Patent Application Publication**
Watts, JR.

(10) **Pub. No.: US 2005/0223255 A1**

(43) **Pub. Date: Oct. 6, 2005**

(54) **METHOD FOR IMPLEMENTING THERMAL AND POWER MANAGEMENT IN A PROCESSOR AND/OR APPARATUS AND/OR SYSTEM EMPLOYING THE SAME**

division of application No. 09/727,597, filed on Dec. 1, 2000, now Pat. No. 6,427,211, which is a division of application No. 08/395,335, filed on Feb. 28, 1995, now Pat. No. 6,158,012, which is a continuation-in-part of application No. 08/023,831, filed on Apr. 12, 1993, now Pat. No. 6,006,336, which is a continuation of application No. 07/429,270, filed on Oct. 30, 1989, now Pat. No. 5,218,704.

(76) Inventor: **La Vaughn F. Watts JR.**, Austin, TX (US)

Correspondence Address:
TEXAS INSTRUMENTS INCORPORATED
P O BOX 655474, M/S 3999
DALLAS, TX 75265

Publication Classification

(51) **Int. Cl.7** **G06F 1/28**
(52) **U.S. Cl.** **713/322**

(21) Appl. No.: **11/137,032**


(57) **ABSTRACT**

(22) Filed: **May 25, 2005**

A method for detecting temperature and a work load level associated with a processor, results of the detecting being used for controlling power dissipation associated with the processor and/or apparatus and/or system employing the same.

Related U.S. Application Data

(60) Continuation of application No. 10/106,261, filed on Mar. 26, 2002, now Pat. No. 6,901,524, which is a






Fig. 1

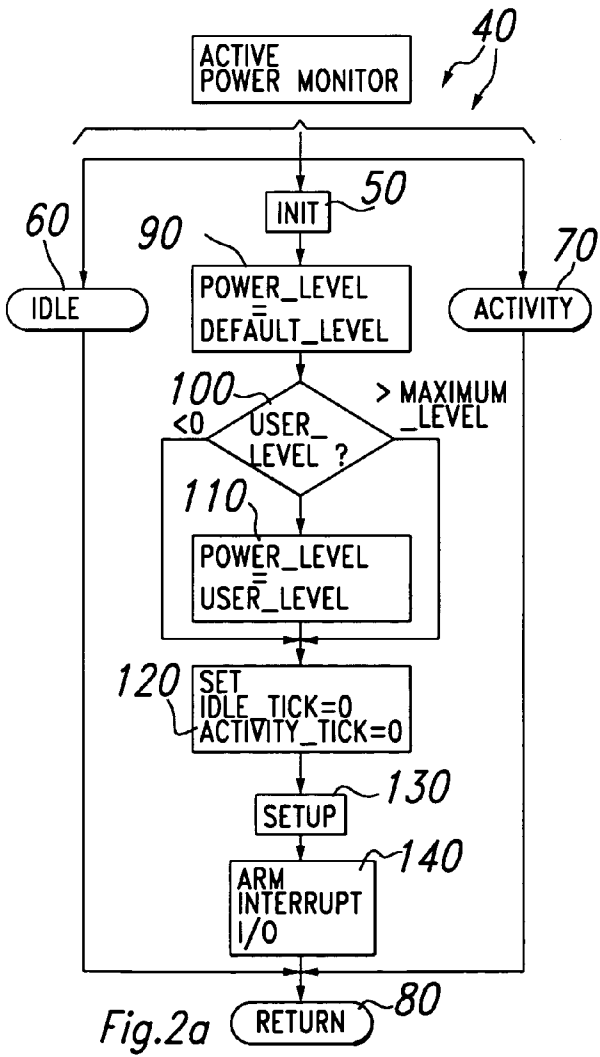


Fig. 2a

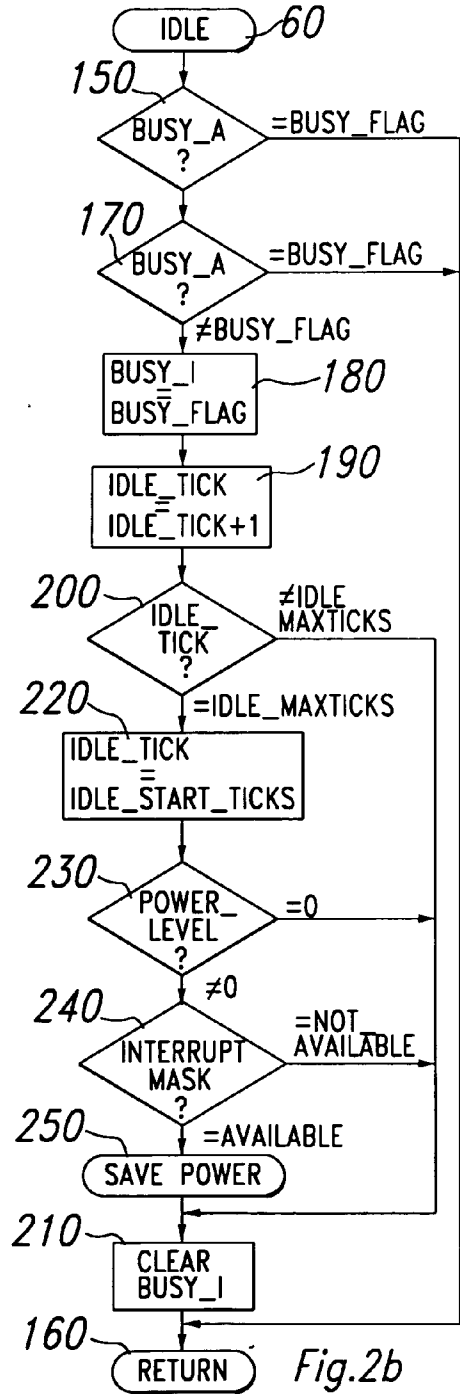

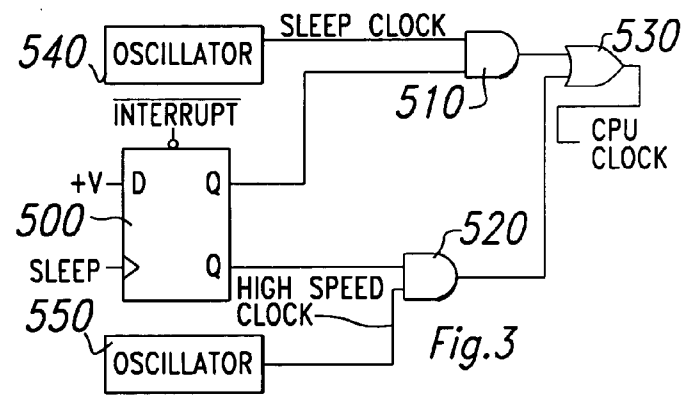

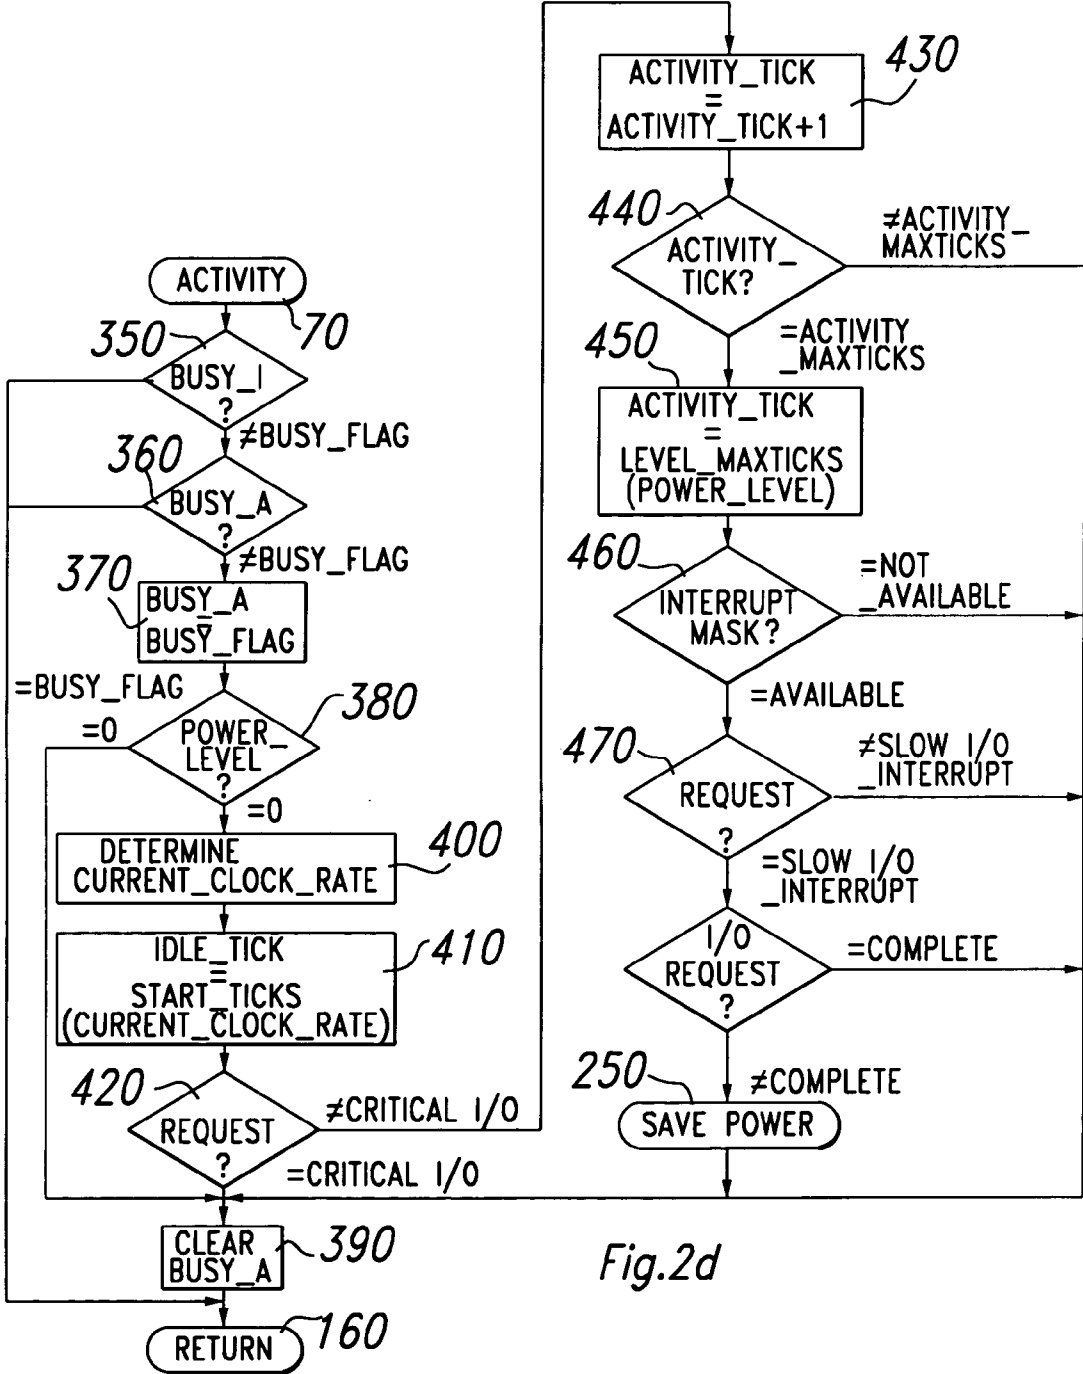


Fig. 2b





**METHOD FOR IMPLEMENTING THERMAL AND
POWER MANAGEMENT IN A PROCESSOR
AND/OR APPARATUS AND/OR SYSTEM
EMPLOYING THE SAME**

BACKGROUND OF THE INVENTION

[0001] 1. Field of the Invention

[0002] This invention relates to real-time computer thermal management and power conservation, and more particularly to an apparatus and method for decreasing and increasing central processing unit (CPU) clock time based on temperature and real-time activity levels within the CPU of a portable computer.

[0003] 2. Description of the Related Art

[0004] During the development stages of personal computers, the transportable or portable computer has become very popular. Such portable computer uses a large power supply and really represents a small desktop personal computer. Portable computers are smaller and lighter than a desktop personal computer and allow a user to employ the same software that can be used on a desktop computer.

[0005] The first generation "portable" computers only operated from an A/C wall power. As personal computer development continued, battery-powered computers were designed. Furthermore, real portability became possible with the development of new display technology, better disk storage, and lighter components. Unfortunately, the software developed was designed to run on desk top computers without regard to battery-powered portable computers that only had limited amounts of power available for short periods of time. No special considerations were made by the software, operating system (MS-DOS), Basic Input/Output System (BIOS), or the third party application software to conserve power usage for these portable computers.

[0006] As more and more highly functional software packages were developed, desktop computer users experienced increased performance from the introductions of higher computational CPUs, increased memory, and faster high performance disk drives. Unfortunately, portable computers continued to run only on A/C power or with large and heavy batteries. In trying to keep up with the performance requirements of the desk top computers, and the new software, expensive components were used to cut the power requirements. Even so, the heavy batteries still did not run very long. This meant users of portable computers has to settle for A/C operation or very short battery operation to have the performance that was expected from the third party software.

[0007] Portable computer designers stepped the performance down to 8083- and 8086-type processors to reduce the power consumption. The supporting circuits and CPU took less power to run and therefore, lighter batteries could be used. Unfortunately, the new software requiring 80286-type instructions, that did not exist in the older slower 8088/8086 CPUs, did not run. In an attempt to design a portable computer that could conserve power, thereby yielding longer battery operation, smaller units, and less weight, some portable computer designers proceeded to reduce power consumption of a portable computer while a user is not using the computer. For example, designers obtain a reduction in power usage by slowing or stopping the disk

drive after some predetermined period of inactivity; if the disk drive is not being used, the disk drive is turned off, or simply placed into a standby mode. When the user is ready to use the disk, the operator must wait until the disk drive spins up and the computer system is ready again for full performance before the operator may proceed with the operation.

[0008] Other portable computer designers conserve power by turning the computer display off when the keyboard is not being used. However, in normal operation the computer is using full power. In other words, power conservation by this method is practical only when the user is not using the components of the system. It is very likely, however, that the user will turn the computer off when not in use. Nevertheless, substantial power conservation while the operator is using the computer for meaningful work is needed. When the operator uses the computer, full operation of all components is required. During the intervals while the operator is not using the computer, however, the computer could be turned off or slowed down to conserve power consumption. It is critical to maintaining performance to determine when to slow the computer down or turn it off without disrupting the user's work, upsetting the third party software, or confusing the operating system, until operation is needed.

[0009] Furthermore, although a user can wait for the disk to spin up as described above, application software packages cannot wait for the CPU to "spin up" and get ready. The CPU must be ready when the application program needs to compute. Switching to full operation must be completed quickly and without the application program being affected. This immediate transition must be transparent to the user as well as to the application currently active. Delays cause user operational problems in response time and software compatibility, as well as general failure by the computer to accurately execute a required program.

[0010] Other attempts at power conservation for portable computers include providing a "Shut Down" or "Standby Mode" of operation. The problem, again, is that the computer is not usable by the operator during this period. The operator could just as well turned off the power switch of the unit to save power. This type of power conservation only allows the portable computer to "shut down" and thereby save power if the operator forgets to turn off the power switch, or walks away from the computer for the programmed length of time. The advantage of this type of power conservation over just turning the power switch off/on is a much quicker return to full operation. However, this method of power conservation is still not real-time, intelligent power conservation while the computer is on and processing data which does not disturb the operating system, BIOS, and any third party application programs currently running on the computer.

[0011] Some attempt to meet this need was made by VLSI vendors in providing circuits that either turned off the clocks to the CPU when the user was not typing on the keyboard or woke up the computer on demand when a keystroke occurred. Either of these approaches reduce power but the computer is dead (unusable) during this period. Background operations such as updating the system clock, communications, print spooling, and other like operations cannot be performed. Some existing portable computers employ these circuits. After a programmed period of no activity, the

computer turns itself off. The operator must turn the machine on again but does not have to reboot the operating system and application program. The advantage of this circuitry is like the existing "shut down" operations, a quick return to full operation without restarting the computer. Nevertheless, this method only reduces power consumption when the user walks away from the machine and does not actually extend the operational life of the battery charge.

[0012] Thermal over-heating, of CPUs and other related devices is another problem yet to be addressed by portable computer manufacturers. CPUs are designed to operate within specific temperature ranges (varies depending on CPU type, manufacturer, quality, etc). CPU performance and speed degenerates when the limits of the operation temperature ranges are exceeded, especially the upper temperature range. This problem is particularly acute with CPUs manufactured using CMOS technology where temperatures above the upper temperature range result in reduced CPU performance and speed. Existing power saving techniques save power but do not measure and intelligently control CPU and/or related device temperature.

SUMMARY OF THE INVENTION

[0013] In view of the above problems associated with the related art, it is an object of the present invention to provide an apparatus and method for real-time conservation of power and thermal management for computer systems without any real-time performance degradation, such conservation of power and thermal management remaining transparent to the user.

[0014] Another object of the present invention is to provide an apparatus and method for predicting CPU activity and temperature levels and using the predictions for automatic power conservation and temperature control.

[0015] Yet another object of the present invention is to provide an apparatus and method which allows user modification of automatic activity and temperature level predictions and using the modified predictions for automatic power conservation and temperature control.

[0016] A further object of the present invention is to provide an apparatus and method for real-time reduction and restoration of clock speeds thereby returning the CPU to full processing rate from a period of inactivity, which is transparent to software programs.

[0017] These objects are accomplished in a preferred embodiment of the present invention by an apparatus and method which determine whether a CPU may rest (including any PCI bus coupled to the CPU) based upon CPU activity and temperature levels and activates a hardware selector based upon that determination. If the CPU may rest, or sleep, the hardware selector applies oscillations at a sleep clock level; if the CPU is to be active, the hardware selector applies oscillations at a high speed clock level.

[0018] The present invention examines the state of CPU activity and temperature, as well as the activity of both the operator and any application software currently active. This sampling of activity and temperature is performed real-time, adjusting the performance level of the computer to manage power conservation, CPU temperature and computer power. These adjustments are accomplished within the CPU cycles and do not affect the user's perception of performance.

[0019] Thus, when the operator for the third party software of the operating system/BIOS is not using the computer, the present invention will effect a quick turn off or slow down of the CPU until needed, thereby reducing the power consumption and CPU temperature, and will promptly restore full CPU operation when needed without affecting perceived performance. This switching back into full operation from the "slow down" mode occurs without the user having to request it and without any delay in the operation of the computer while waiting for the computer to return to a "ready" state.

BRIEF DESCRIPTION OF THE DRAWINGS

[0020] The novel features believed characteristic of the invention are set forth in the appended claims. The invention itself, however, as well as other features and advantages thereof, will be best understood by reference to the detailed description with follows, read in conjunction with the accompanying drawings, wherein:

[0021] FIG. 1 is a flowchart depicting the self-tuning aspect of a preferred embodiment of the present invention.

[0022] FIGS. 2a-2d are flowcharts depicting the active power conservation monitor employed by the present invention.

[0023] FIG. 3 is a simplified schematic diagram representing the active power conservation associated hardware employed by the present invention.

[0024] FIG. 4 is a schematic of the sleep hardware for one embodiment of the present invention.

[0025] FIG. 5 is a schematic of the sleep hardware for another embodiment of the invention.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

[0026] If the period of computer activity in any given system is examined, the CPU and associated components have a utilization percentage. If the user is inputting data from the keyboard, the time between keystrokes is very long in terms of CPU cycles. Many things can be accomplished by the computer during this time, such as printing a report. Even during the printing of a report, time is still available for additional operations such as background updating of a clock/calendar display. Even so, there is almost always spare time when the CPU is not being used. If the computer is turned off or slowed down during this spare time, then power consumption is obtained real-time. Such real-time power conservation extends battery operation life and lowers CPU temperature.

[0027] According to one embodiment of the present invention, to conserve power and lower CPU temperature under MS-DOS, as well as other operating systems such as OS/2, XENIX, and those for Apple computers, requires a combination of hardware and software. It should be noted that because the present invention will work in any system, while the implementation may vary slightly on a system-by-system basis, the scope of the present invention should therefore not be limited to computer systems operating under MS-DOS.

[0028] Slowing down or stopping computer system components reduces power consumption and lowers CPU tem-

perature, although the amount of power saved and CPU temperature reduction may vary. Therefore, according to the present invention, stopping the clock (where possible as some CPUs cannot have their clocks stopped) reduces power consumption and CPU temperature more than just slowing the clock.

[0029] In general, the number of operations (or instructions) per second may be considered to be roughly proportional to the processor clock:

$$\text{instructions/second} = \text{instructions/cycle} * \text{cycles/second}$$

[0030] Assuming for simplicity that the same instruction is repeatedly executed so that instructions/second is constant, the relationship can be expressed as follows:

$$Fq = K_1 * Clk$$

[0031] where Fq is instructions/second, K_1 is constant equal to the instructions/cycle, and Clk equals cycles/second. Thus, roughly speaking, the rate of execution increases with the frequency of the CPU clock.

[0032] The amount of power being used at any given moment is also related to the frequency of the CPU clock and therefore to the rate of execution. In general this relationship can be expressed as follows:

$$P = K_2 + (K_3 * Clk)$$

[0033] where P is power in watts, K_2 is a constant in watts, K_3 is a constant and expresses the number of watt-second/cycle, and Clk equals the cycles/second of the CPU clock. Thus it can also be said that the amount of power being consumed at any given time increases as the CPU clock frequency increases.

[0034] Assume that a given time period T is divided into N intervals such that the power P is constant during each interval. Then the amount of energy E expended during T is given by:

$$E = P(1) \text{ delta } T_1 + P(2) \text{ delta } T_2 \dots P(N) \text{ delta } T_N$$

[0035] Further assume that the CPU clock "CLK" has only two states, either "ON" or "OFF". For the purposes of this discussion, the "ON" state represents the CPU clock at its maximum frequency, while the "OFF" state represents the minimum clock rate at which the CPU can operate (this may be zero for CPUs that can have their clocks stopped). For the condition in which the CPU clock is always "ON", each P(i) in the previous equation is equal and the total energy is:

$$\frac{E(\text{max})}{P(\text{on}) * T} = (\text{delta } T_1 + \text{delta } T_2 \dots \text{delta } T_N)$$

[0036] This represents the maximum power consumption of the computer in which no power conservation measures are being used. If the CPU clock is "off" during a portion of the intervals, then there are two power levels possible for each interval. The P(on) represents the power being consumed when the clock is in its "ON" state, while P(off) represents the power being used when the clock is "OFF". If all of the time intervals in which the clock is "ON" are [is] summed into the quantity "T(on)" and the "OFF" intervals are summed into "T(off)", then it follows:

$$T = T(\text{on}) + T(\text{off})$$

[0037] Now the energy being used during period T can be written:

$$E = [P(\text{on}) * T(\text{on})] + [P(\text{off}) * T(\text{off})]$$

[0038] Under these conditions, the total energy consumed may be reduced by increasing the time intervals T(off). Thus, by controlling the periods of time the clock is in its "OFF" state, the amount of energy being used may be reduced. If the T(off) period is divided into a large number of intervals during the period T, then as the width of each interval goes to zero, energy consumption is at a maximum. Conversely, as the width of the T(off) intervals increase, the energy consumed decreases.

[0039] If the "OFF" intervals are arranged to coincide with periods during which the CPU is normally inactive, then the user cannot perceive any reduction in performance and overall energy consumption is reduced from the E(max) state. In order to align the T(off) intervals with periods of CPU inactivity, the CPU activity and temperature levels are used to determine the width of the T(off) intervals in a closed loop. FIG. 1 depicts such a closed loop. The activity level of the CPU is determined at Step 10. If this level is a decrease over an immediately previous determination (Step 22), the present invention increases the T(off) interval (Step 20) and returns to determine the activity level of the CPU again. If, on the other hand, this activity level is an increase over an immediately previous determination (Step 22), a determination is made as to whether or not the temperature of the CPU is a concern (Step 24). If CPU temperature is not a concern, the present invention decreases the T(off) interval (Step 30) and proceeds to again determine the activity level of the CPU. If, on the other hand, CPU temperature is a concern, a determination is made as to whether or not the CPU is processing critical I/O, a critical function or a critical real-time event (Step 26). If critical I/O or critical function or a critical real-time event are being processed, the present invention decreases the T(off) interval (Step 30) and proceeds to again determine the activity level of the CPU. If no critical I/O is being processed, the present invention increases the T(off) interval (Step 20) and proceeds again to determine the activity level of the CPU. Thus the T(off) intervals are constantly being adjusted to match the system activity level and control the temperature level of the CPU.

[0040] Management of CPU temperature (thermal management) is necessary because CPUs are designed to operate within a specific temperature range. CPU performance and speed deteriorates when the specified high operating temperature of a CPU is exceeded (especially in CMOS process CPUs where temperatures above the high operating temperature translate into slower CPU speed). The heat output of a CPU is directly related to the power consumed by the CPU and heat it absorbs from devices and circuitry that immediately surround it. CPU power consumption increases with CPU clock speed and the number of instructions per second to be performed by the CPU. As a result, heat related problems are becoming more common as faster and increasingly complex CPUs are introduced and incorporated into electronic devices.

[0041] In any operating system, two key logic points exist: an IDLE, or "do nothing", loop within the operating system and an operating system request channel, usually available for services needed by the application software. By placing logic inline with these logic points, the type of activity request made by an application software can be evaluated, power conservation and thermal management can be activated and slice periods determined. A slice period is the number of T(on) vs. T(off) intervals over time, computed by

the CPU activity and thermal levels. An assumption may be made to determine CPU activity level: Software programs that need service usually need additional services and the period of time between service requests can be used to determine the activity level of any application software running on the computer and to provide slice counts for power conservation according to the present invention. Another assumption that may be made is that each CPU has a temperature coefficient unique to that CPU-CPU temperature rise time, CPU maximum operating temperature, CPU temperature fall time and intervention time required for thermal control. If this information is not provided by the CPU manufacturer, testing of the CPU being used (or another of the same make and type tested under similar conditions) is required to obtain accurate information.

[0042] Once the CPU is interrupted during a power conservation and thermal management slice (T(off)), the CPU will save the interrupted routine's state prior to vectoring to the interrupt software. Of course, since the power conservation and thermal management software was operating during this slice, control will be returned to the active power conservation and thermal management loop (monitor 40) which simply monitors the CPU's clock to determine an exit condition for the power conservation and thermal management mode thereby exiting from T(off) to T(on) state. The interval of the next power conservation and thermal management state is adjusted by the activity level monitor, as discussed above in connection with FIG. 1. Some implementations can create an automatic exit from T(off) by the hardware logic, thereby forcing the power conservation and thermal management loop to be exited automatically and executing an interval T(on).

[0043] More specifically, looking now at FIGS. 2a-2d, which depict the active power conservation and thermal management monitor 40 of the present invention. The CPU installs monitor 40 either via a program stored in the CPU ROM or loads it from an external device storing the program in RAM. Once the CPU has loaded monitor 40, it continues to INIT 50 for system interrupt initialization, user configurational setup, and system/application specific initialization. IDLE branch 60 (more specifically set out in FIG. 2b) is executed by a hardware or software interrupt for an IDLE or "do nothing" function. This type of interrupt is caused by the CPU entering either an IDLE or a "do nothing" function. This type of interrupt is caused by the CPU entering either an IDLE or a "do nothing" loop (i.e., planned inactivity). The ACTIVITY branch 70 of the flow chart, more fully described below in relation to FIG. 2d, is executed by a software or hardware interrupt due to an operating system or I/O service request, by an application program or internal operating system function. An I/O service request made by a program may, for example, be a disk I/O, read, print, load, etc. Regardless of the branch selected, control is eventually returned to the CPU operating system at RETURN 80. The INIT branch 50 of this flowchart, shown in FIG. 2a, is executed only once if it is loaded via program into ROM or is executed every time during power up if it is loaded from an external device and stored in the RAM. Once this branch of active power and thermal management monitor 40 has been fully executed, whenever control is yielded from the operating system to the power conservation and thermal management mode, either IDLE 60 or ACTIVITY 70 branches are selected depending on the type of CPU activity: IDLE branch 60 for power conservation and thermal man-

agement during planned inactivity and ACTIVITY branch 70 for power conservation and thermal management during CPU activity.

[0044] Looking more closely at INIT branch 50, after all system interrupt and variables are initialized, the routine continues at Step 90 to set the Power_level equal to DEFAULT_LEVEL. In operating systems where the user has input control for the Power_level, the program at Step 100 checks to see if a User_level has been selected. If the User_level is less than zero or greater than the MAXIMUM_LEVEL, the system used the DEFAULT_LEVEL. Otherwise, it continues onto Step 110 where it modifies the Power_level to equal the User_level.

[0045] According to the preferred embodiment of the present invention, the system at Step 120 sets the variable Idle_tick to zero and the variable Activity_tick to zero. Under an MS/DOS implementation, Idle_tick refers to the number of interrupts found in a "do nothing" loop. Activity_tick refers to the number of interrupts caused by an activity interrupt which in turn determines the CPU activity level. Tick count represents a delta time for the next interrupt Idle_tick is a constant delta time from one tick to another (interrupt) unless overwritten by a software interrupt. A software interrupt may reprogram delta time between interrupts.

[0046] After setting the variables to zero, the routine continues on to Setup 130 at which time any application specific configuration fine-tuning is handled in terms of system-specific details and the system is initialized. Next the routine arms the interrupt I/O (Step 140) with instructions to the hardware indicating the hardware can take control at the next interrupt. INIT branch 50 then exits to the operating system, or whatever called the active power and thermal management monitor originally, at RETURN 80.

[0047] Consider now IDLE branch 60 of active power and thermal management monitor 40, more fully described at FIG. 2b. In response to a planned inactivity of the CPU, monitor 40 (not specifically shown in this Figure) checks to see if entry into IDLE branch 60 is permitted by first determining whether the activity interrupt is currently busy. If Busy_A equals BUSY_FLAG (Step 150), which is a reentry flag, the CPU is busy and cannot now be put to sleep. Therefore, monitor 40 immediately proceeds to RETURN I 160 and exits the routine RETURN I 160 is an indirect vector to the previous operating system IDLE vector interrupt for normal processing stored before entering monitor 40. (I.e., this causes an interrupt return to the last chained vector.)

[0048] If the Busy_A interrupt flag is not busy, then monitor 40 checks to see if the Busy_Idle interrupt flag, Busy_I, equals BUSY_FLAG (Step 170). If so, this indicates the system is already in IDLE branch 60 of monitor 40 and therefore the system should not interrupt itself. If Busy_I=BUSY_FLAG, the system exits the routine at RETURN_I indirect vector 160.

[0049] If, however, neither the Busy_A reentry flag or the Busy_I reentry flag have been set, the routine sets the Busy_I flag at Step 180 for reentry protection (Busy_I=BUSY_FLAG). At Step 190 Idle_tick is incremented by one. Idle_tick is the number of T(on) before a T(off) interval and is determined from IDLE interrupts, setup interrupts and

from CPU activity and temperature levels. Idle_tick increments by one to allow for smoothing of events, thereby letting a critical I/O activity control smoothing.

[0050] At Step 200 monitor 40 checks to see if Idle_tick equals IDLE_MAXTICKS. IDLE_MAXTICKS is one of the constants initialized in Setup 130 of INIT branch 50, remains constant for a system, and is responsible for self-tuning of the activity and thermal levels. If Idle_tick does not equal IDLE_MAXTICKS, the Busy_I flag is cleared at Step 210 and exits the loop proceeding to the RETURN I indirect vector 160. If, however, Idle_tick equals IDLE_MAXTICKS, Idle_tick is set equal to IDLE_START_TICKS (Step 220). IDLE_START_TICKS is a constant which may or may not be zero (depending on whether the particular CPU can have its clock stopped). This step determines the self-tuning of how often the rest of the sleep functions may be performed. By setting IDLE_START_TICKS equal to IDLE_MAXTICKS minus one, a continuous T(off) interval is achieved. At Step 230, the Power_level is checked. If it is equal to zero, the monitor clears the Busy_I flag (Step 210), exits the routine at RETURN I 160, and returns control to the operating system so it may continue what it was originally doing before it entered active power monitor 40.

[0051] If, however, the Power_level does not equal zero at Step 240, the routine determines whether an interrupt mask is in place. An interrupt mask is set by the system/application software, and determines whether interrupts are available to monitor 40. If interrupts are NOT_AVAILABLE, the Busy_I reentry flag is cleared and control is returned to the operating system to continue what it was doing before it entered monitor 40. Operating systems, as well as application software, can set T(on) interval to yield a continuous T(on) state by setting the interrupt mask equal to NOT_AVAILABLE.

[0052] Assuming an interrupt is AVAILABLE, monitor 40 proceeds to the SAVE POWER subroutine 250 which is fully executed during one T(off) period established by the hardware state. (For example, in the preferred embodiment of the present invention, the longest possible interval could be 18 ms, which is the longest time between two ticks or interrupts from the real-time clock.) During the SAVE POWER subroutine 250, the CPU clock is stepped down to a sleep clock level.

[0053] Once a critical I/O operation forces the T(on) intervals, the IDLE branch 60 interrupt tends to remain ready for additional critical I/O requests. As the CPU becomes busy with critical I/O, less T(off) intervals are available. Conversely, as critical I/O requests decrease, and the time intervals between them increase, more T(off) intervals are available. IDLE branch 60 is a self-tuning system based on feedback from CPU activity and temperature interrupts and tends to provide more T(off) intervals as the activity level slows and/or the CPU temperature becomes a concern. As soon as monitor 40 has completed SAVE POWER subroutine 250, shown in FIG. 2c and more fully described below, the Busy_I reentry flag is cleared (Step 210) and control is returned at RETURN I 160 to whatever operating system originally requested monitor 40.

[0054] Consider now FIG. 2c, which is a flowchart depicting the SAVE POWER subroutine 250. Monitor 40 determines what the I/O hardware high speed clock is at Step 260.

It sets the CURRENT_CLOCK_RATE equal to the relevant high speed clock and saves this value to be used for CPUs with multiple level high speed clocks. Thus, if a particular CPU has 12 MHz and 6 MHz high speed clocks, monitor 40 must determine which high speed clock the CPU is at before monitor 40 reduces power so it may reestablish the CPU at the proper high speed clock when the CPU awakens. At Step 270, the Save_clock_rate is set equal to the CURRENT_CLOCK_RATE determined. Save_clock_rate 270 is not used when there is only one high speed clock for the CPU. Monitor 40 now continues to SLEEPLOCK 280, where a pulse is sent to the hardware selector (shown in FIG. 3) to put the CPU clock to sleep (i.e., lower or stop its clock frequency). The I/O port hardware sleep clock is at much lower oscillations than the CPU clock normally employed.

[0055] At this point either of two events can happen. A system/application interrupt may occur or a real-time clock interrupt may occur. If a system/application interrupt 290 occurs, monitor 40 proceeds to interrupt routine 300, processing the interrupt as soon as possible, arming interrupt I/O at Step 310, and returning to determine whether there has been an interrupt (Step 320). Since in this case there has been an interrupt, the Save_clock_rate is used (Step 330) to determine which high speed clock to return the CPU to and SAVE POWER subroutine 250 is exited at RETURN 340. If, however, a system/application interrupt is not received, the SAVE POWER subroutine 250 will continue to wait until a real-time clock interrupt has occurred (Step 320). Once such an interrupt has occurred, SAVE POWER subroutine 250 will continue to wait until a real-time clock interrupt has occurred (Step 320). Once such an interrupt has occurred, SAVEPOWER subroutine 250 will execute interrupt loop 320 several times. If however, control is passed when the sleep clock rate was zero, in other words, there was no clock, the SAVE POWER subroutine 250 will execute interrupt loop 320 once before returning the CPU clock to the Save_clock_rate 330 and exiting (Step(340)).

[0056] Consider now FIG. 2d which is a flowchart showing ACTIVITY branch 70 triggered by an application/system activity request via an operating system service request interrupt. ACTIVITY branch 70 begins with reentry protection. Monitor 40 determines at Step 350 whether Busy_I has been set to BUSY_FLAG. If it has, this means the system is already in ACTIVITY branch 70 and cannot be interrupted. If Busy_I=BUSY_FLAG, monitor 40 exits to RETURN I 160, which is an indirect vector to an old activity vector interrupt for normal processing, via an interrupt vector after the operating system performs the requested service.

[0057] If however, the Busy_I flag does not equal BUSY_FLAG, which means ACTIVITY branch 70 is not being accessed, monitor 40 determines at Step 360 if the BUSY_A flag has been set equal to BUSY_FLAG. If so, control will be returned to the system at this point because ACTIVITY branch 70 is already being used and cannot be interrupted. If the Busy_A flag has not been set, in other words, Busy_A does not equal BUSY_FLAG, monitor 40 sets Busy_A equal to BUSY_FLAG at Step 370 so as not to be interrupted during execution of ACTIVITY branch 70. At Step 380 the Power_level is determined. If Power_level equals zero, monitor 40 exits ACTIVITY branch 70 after clearing the Busy_A reentry flag (Step 390). If however, the Power_level does not equal zero, the CURRENT_CLOCK_RATE of the

I/O hardware is next determined. As was true with Step 270 of FIG. 2C, Step 400 of FIG. 2d uses the CURRENT_CLOCK_RATE if there are multiple level high speed clocks for a given CPU. Otherwise, CURRENT_CLOCK_RATE always equals the CPU high speed clock. After the CURRENT_CLOCK_RATE is determined (step 400), at Step 410 Idle_tick is set equal to the constant START_TICKS established for the previously determined CURRENT_CLOCK_RATE. T(off) intervals are established based on the current high speed clock that is active.

[0058] Monitor 40 next determines that a request has been made. A request is an input by the application software running on the computer, for a particular type of service needed. At Step 420, monitor 40 determines whether the request is a CRITICAL I/O. If the request is a CRITICAL I/O, it will continuously force T(on) to lengthen until the T(on) is greater than the T(off), and monitor 40 will exit ACTIVITY branch 70 after clearing the Busy_A reentry flag (Step 390). If, on the other hand, the request is not a CRITICAL I/O, then the Activity_tick is incremented by one at Step 430. It is then determined at Step 440 whether the Activity_tick now equals ACTIVITY_MAXTICKS. Step 440 allows a smoothing from a CRITICAL I/O, and makes the system ready from another CRITICAL I/O during Activity_tick T(on) intervals. Assuming Activity_tick does not equal ACTIVITY_MAXTICKS, ACTIVITY branch 70 is exited after clearing the Busy_A reentry flag (Step 390). If, on the other hand, the Activity_tick equals constant ACTIVITY_MAXTICKS, at Step 450 Activity_tick is set to the constant LEVEL_MAXTICKS established for the particular Power_level determined at Step 330.

[0059] Now monitor 40 determines whether an interrupt mask exists (Step 460). An interrupt mask is set by system/application software. Setting it to NOT_AVAILABLE creates a continuous T(on) state. If the interrupt mask equals NOT_AVAILABLE, there are no interrupts available at this time and monitor 40 exits ACTIVITY branch 70 after clearing the Busy_A reentry flag (Step 390). If, however, an interrupt is AVAILABLE, monitor 40 determines at Step 470 whether the request identified at Step 420 was for a SLOW I/O_INTERRUPT. Slow I/O requests may have a delay until the I/O device becomes "ready". During the "make ready" operation, a continuous T(off) interval may be set up and executed to conserve power. Thus, if the request is not a SLOW I/O_INTERRUPT, ACTIVITY branch 70 is exited after clearing the Busy_A reentry flag (Step 390). If, however, the request is a SLOW I/O_INTERRUPT, and time yet exists before the I/O device becomes "ready", monitor 40 then determines at Step 480 whether the I/O request is COMPLETE (i.e., is I/O device ready?). If the I/O device is not ready, monitor 40 forces T(off) to lengthen, thereby forcing the CPU to wait, or sleep, until the SLOW I/O device is ready. At this point it has time to save power and ACTIVITY branch 70 enters SAVE POWER subroutine 250 previously described in connection with FIG. 2C. If, however, the I/O request is COMPLETE, control is returned to the operating system subsequently to monitor 40 exiting ACTIVITY branch 70 after clearing Busy_A reentry flag (Step 390).

[0060] Self-tuning is inherent within the control system of continuous feedback loops. The software of the present invention can detect when CPU activity is low and/or CPU temperature is high enough to be of concern and therefore

when the power conservation and thermal management aspect of the present invention may be activated. To detect when CPU temperature is high enough to be of concern, the power and thermal management software monitors a thermistor on the PWB board adjacent the CPU (or mounted directly on or in the CPU if the CPU includes a thermistor). In one embodiment of the present invention, the software monitors the thermistor 18 times/sec through an A/D converter. If no power is being conserved and the temperature of the thermistor is within acceptable parameters, then monitoring continues at the same rate. If, however, the temperature of the thermistor is rising, a semaphore is set to tell the system to start watching CPU temperature for possible thermal management action. Each CPU has a temperature coefficient unique to that specific CPU. Information on how long it takes to raise the temperature and at what point intervention must occur to prevent performance degradation must be derived from information supplied with the CPU or through testing.

[0061] According to one embodiment of the invention, a counter is set in hardware to give an ad hoc interrupt (counter is based on coefficient of temperature rise). The thermal management system must know how long it takes CPU temperature to go down to minimize temperature effect. If the counter is counting down and receives an active power interrupt, the ad hoc interrupt is turned off because control has been regained through the active power and thermal management. The result is unperceived operational power savings. The ad hoc interrupt can be overridden or modified by the active power interrupt which checks the type gradient i.e., up or down, checks the count and can adjust the up count and down count ad hoc operation based on what the CPU is doing real time. If there are no real time interrupts, then the timer interval continually comes in and monitors the gradual rise in temperature and it will adjust the ad hoc counter as it needs it up or down. The result is dynamic feedback from the active power and thermal management into the ad hoc timer, adjusting it to the dynamic adjustment based on what the temperature rise or fall is at any given time and how long it takes for that temperature to fall off or rise through the danger point. This is a different concept that just throwing a timer out ad hoc and letting it run.

[0062] For example, assume that the CPU being used has a maximum safe operating temperature of 95 degrees C. (obtained from the CPU spec sheet or from actual testing). Assume also that a thermistor is located adjacent the CPU and that when the CPU case is at 95 degrees C., the temperature of the thermistor may be lower since it is spaced a distance from the CPU (such as 57 degrees C.). A determination should be made as to how long it took the CPU to reach 95 degrees. If it took an hour, the system may decide to sample the thermistor every 45 minutes. Once the CPU is at 95 degrees, CPU temperature may need to be sampled every minute to make sure the temperature is going down, otherwise, the temperature might go up, i.e., to 96 degrees. If 5 minutes are required to raise CPU temperature from 95 to 96 degrees, CPU temperature sampling must be at a period less than 5 minutes—i.e., every 3 or 1 minutes. If the temperature is not going down, then the length of the rest cycles should be increased. Continual evaluation of the thermal read constant is key to knowing when CPU temperature is becoming a problem, when thermal management intervention is appropriate and how much time can be

allowed for other things in the system. This decision must be made before the target temperature is reached. Once CPU temperature starts to lower, it is O.K. to go back to the regular thermal constant number because 1) you have selected the right slice period, or 2) the active power portion of the active power and thermal management has taken over, so the sampling rate can be reduced.

[0063] Examples of source code that can be stored in the CPU ROM or in an external RAM device, according to one embodiment of the invention, are listed in the COMPUTER PROGRAMS LISTING section under: 1) Interrupt 8 Timer interrupt service—listed on pages _____ to _____; 2) CPU Sleep Routine—listed on pages _____ to _____; 3) FILE=FORCE5.ASM—listed on pages _____ to _____; and 4) FILE=Thermal.EQU—listed on pages _____ to _____.

[0064] Utilizing the above listed source code, and assuming that Interrupt 8 Timer interrupt service is the interrupt mask called at Step 240 of IDLE loop 60 or at Step 460 of ACTIVITY loop 70, the procedure for thermal management is set up “Do Thermal Management if needed” after which the system must decide if there is time for thermal management “Time for Thermal Management?”. If there is time for thermal management, the system calls the file “force_sleep” if there is time to sleep (which also sleeps any PCI bus coupled to the CPU), or alternatively, could do a STI nop and a halt—which is an alternate way and does not get PCI devices and does not have a feedback loop from the power and temperature management systems. The “force_sleep” File gets feedback from other power systems. Force_sleep does a jump to force5.asm, which is the PCI multiple sleep program. Are there speakers busy in the system? Is there something else in the system going on from a power management point of view? Are DMAs running in the system? Sleeping may not be desirable during a sound cycle. It needs to know what is going on in the system to do an intelligent sleep. The thermal management cares about the CPU and cares about all the other devices out there because collectively they all generate heat.

[0065] There are some equations in the program that are running—others that may or may not be running. “tk” is the number of interrupts per second that are sampled times the interval that is sampled over. “it” represents a thermal read constant and the thermal read constant in the present embodiment is 5. In the code, the thermal read constant is dynamically adjusted later depending on what the temperature is. Thus, this is the starting thermal read interval, but as the temperature rises, reading should be more often and the cooler it is, reading should be less often than 5 minutes—e.g., 10 minutes. The thermal read constant will adjust. TP1 or TP2 represents what percentage of the CPU cycles do we want to sample at—for example, TP7 set at 50=the number of interrupts that have to occur over some period of time such that if we take that number that going to represent every so many clock cycles that go by before we sample and sleep the CPU. These equations are variable. Other equations can also be used

[0066] Thus, one concept of the present invention is that there are various levels of temperature that require testing in relationship to the hottest point to be managed. The sample period will change based on temperature and active feedback. Active feedback may be required even though thermal

management has determined that the CPU temperature is too high and should be reduced (by slowing or stopping the CPU clock). CPU clock speed may not be reduced because other system things are happening—the result is intelligent feedback. The power conservation and thermal management systems asks the CPU questions such as are you doing something now that I cannot go do? If not, please sleep. If yes, don’t sleep and come back to me so that I can reset my count. The result is a graduated effect up and graduated effect down and the thermal read constant time period adjusts itself in response to CPU temperature. Performance taken away from the user during power conservation and thermal management control is balanced against critical I/O going on in the system.

[0067] Active power and thermal management cooperates with standard CPU power management so that when standard power management gets a chance to take over the active feedback can start degrading even though the temperature has not. Existing power/thermal management systems turn on and stay on until the temperature goes down. Unfortunately, this preempts things in the system. Such is not the case in the environment of the present invention. The same sleep manager works in conjunction with power conservation and thermal management—the sleep manager has global control. As an example, while CPU temperature may be rising or have risen to a level of concern, the system may be processing critical I/O, such as a wave file being played. With critical I/O, the system of the present invention will play the wave file without interruption even though the result may be a higher CPU temperature. CPUs do not typically overheat all at once. There is a temperature rise gradient. The system of the present invention takes advantage of the temperature rise gradient to give a user things that affect the user time slices and take it away from him when it is not affected.

[0068] Thermal management can be also be achieved using a prediction mode. Prediction mode utilizes no sensors or thermistors or even knowledge as to actual CPU temperature. Prediction mode uses a guess—i.e. that the system will need the ad hoc interrupt once every 5 seconds or 50 times/second (=constant) and then can take it up or down based on what the system is doing with the active power and thermal management. The prediction theory can also be combined with actual CPU temperature monitoring.

[0069] Once the power conservation and thermal management monitor is activated, a prompt return to full speed CPU clock operation within the interval is achieved so as to not degrade the performance of the computer. To achieve this prompt return o full speed CPU clock operation, the preferred embodiment of the present invention employs some associated hardware.

[0070] Looking now at FIG. 3 which shows a simplified schematic diagram representing the associated hardware employed by the present invention for active power conservation and thermal management. When monitor 40 (not shown) determines the CPU is ready to sleep, it writes to an I/O port (not shown) which causes a pulse on the SLEEP line. The rising edge of this pulse on the SLEEP line causes flip flop 500 to clock a high to Q and a low to Q_̄. This causes the AND/OR logic (AND gates 510, 520, OR gate 530) to select the pulses travelling the SLEEP CLOCK line from SLEEP CLOCK oscillator 540 to be sent to and used

by the CPU CLOCK. SLEEP CLOCK oscillator **540** is a slower clock than the CPU clock used during normal CPU activity. The high coming from the Q of flip flop **500** ANDed (**510**) with the pulses coming from SLEEP CLOCK oscillator **540** is ORed (**530**) with the result of the low on the Q_̄ of flip flop **500** ANDed (**520**) with the pulse generated along the HIGH SPEED CLOCK line by the HIGH SPEED CLOCK oscillator **550** to yield the CPU CLOCK. When the I/O port designates SLEEP CLOCK, the CPU CLOCK is then equal to the SLEEP CLOCK oscillator **540** value. If, on the other hand, an interrupt occurs, an interrupt—value clears flip flop **500**, thereby forcing the AND/OR selector (comprising **510**, **520** and **530**) to choose the HIGH SPEED CLOCK value, and returns the CPU CLOCK value to the value coming from HIGH SPEED CLOCK oscillator **550**. Therefore, during any power conservation and/or thermal management operation on the CPU, the detection of any interrupt within the system will restore the CPU operation at full clock rate prior to vectoring and processing the interrupt.

[0071] It should be noted that the associated hardware needed, external to each of the CPUs or any given system, may be different based on the operating system used, whether the CPU can be stopped, etc. Nevertheless, the scope of the present invention should not be limited by possible system specific modifications needed to permit the present invention to actively conserve power and manage CPU temperature in the numerous available portable computer systems. For example two actual implementations are shown in **FIGS. 4 and 5**, discussed below.

[0072] Many VLSI designs today allow for clock switching of the CPU speed. The logic to switch from a null clock or slow clock to a fast clock logic is the same as that which allows the user to change speeds by a keyboard command. The added logic of monitor **40** working with such switching logic, causes an immediate return to a fast clock upon detection of any interrupt. This simple logic is the key to the necessary hardware support to interrupt the CPU and thereby allow the processing of the interrupt at full speed.

[0073] The method to reduce power consumption under MS-DOS employs the MS-DOS IDLE loop trap to gain access to the “do nothing” loop. The IDLE loop provides special access to application software and operating system operations that are in a state of IDLE of low activity. Careful examination is required to determine the activity level at any given point within the system. Feedback loops are used from the interrupt **21H** service request to determine the activity level. The prediction of activity level is determined by interrupt **21H** requests, from which the present invention thereby sets the slice periods for “sleeping” (slowing down or stopping) the CPU. An additional feature allows the user to modify the slice depending on the activity level of interrupt **21H**. The method to produce power conservation under WINDOWS employs real and protect modes to save the power interrupt which is called by the operating system each time WINDOWS has nothing to do.

[0074] Looking now at **FIG. 4**, which depicts a schematic of an actual sleep hardware implementation for a system such as the Intel 50386 (CPU cannot have its clock stopped). Address enable bus **600** and address bus **610** provide CPU input to demultiplexer **620**. The output of demultiplexer **620** is sent along SLEEP_{CS}- and provided as input to OR gates **630**, **640**. The other inputs to OR gates **630**, **640** are the I/O

write control line and the I/O read control line, respectively. The outputs of these gates, in addition to NOR gate **650**, are applied to D flip flop **660** to decode the port. “INTR” is the interrupt input from the I/O port (peripherals) into NOR gate **650**, which causes the logic hardware to switch back to the high speed clock. The output of flip flop **660** is then fed, along with the output from OR gate **630**, to tristate buffer **670** to enable it to read back what is on the port. All of the above-identified hardware is used by the read/write I/O port (peripherals) to select the power saving “Sleep” operation. The output “SLOW_̄” is equivalent to “SLEEP” in **FIG. 2**, and is inputted to flip flop **680**, discussed later.

[0075] The output of SLEEP CLOCK oscillator **690** is divided into two slower clocks by D flip flops **700**, **710**. In the particular implementation shown in **FIG. 4**, 16 MHz sleep clock oscillator **690** is divided into 4 MHz and 8 MHz clocks. Jumper **J1** selects which clock is to be the “SLEEP CLOCK”.

[0076] In this particular implementation, high speed clock oscillator **720** is a 32 MHz oscillator, although this particular speed is not a requirement of the present invention. The 32 MHz oscillator is put in series with a resistor (for the implementation shown, 33 ohms), which is in series with two parallel capacitors (10 pF). The result of such oscillations is tied to the clocks of D flip flops **730**, **740**.

[0077] D flip flops **680**, **730**, **740** are synchronizing flip flops; **680**, **730** were not shown in the simplified sleep hardware of **FIG. 2**. These flip flops are used to ensure the clock switch occurs only on clock edge. As can be seen in **FIG. 4**, as with flip flop **500** of **FIG. 2**, the output of flip flop **740** either activates OR gate **750** or OR gate **760**, depending upon whether the CPU is to sleep (“FASTEN_̄”) or awaken (“SLOWEN_̄”).

[0078] OR gates **750**, **760** and AND gate **770** are the functional equivalents to the AND/OR selector of **FIG. 2**. They are responsible for selecting either the “slowclk” (slow clock, also known as SLEEP CLOCK) or high speed clock (designated as 32 MHz on the incoming line). In this implementation, the Slow clock is either 4 MHz or 5 MHz, depending upon jumper **J1**, and the high speed clock is 32 MHz. The output of AND gate **770** (ATUCLK) establishes the rate of the CPU clock, and is the equivalent of CPU CLOCK of **FIG. 2**. (If the device includes a PCI bus, the output of AND gate **770** may also be coupled to the PCI bus if it is to utilize the clock signal.)

[0079] Consider now **FIG. 5**, which depicts a schematic of another actual sleep hardware implementation for a system such as the Intel 50286 (CPU can have its clock stopped). The Western Digital FE3600 VLSI is used for the speed switching with a special external PAL **780** to control the interrupt gating which wakes up the CPU on any interrupt. The software power conservation according to the present invention monitors the interrupt acceptance, activating the next P(i)deltaTi interval after the interrupt.

[0080] Any interrupt request to the CPU will return the system to normal operation. An interrupt request (“INTRQ”) to the CPU will cause the PAL to issue a Wake Up signal on the RESCPU line to the FE3001 (not shown) which in turn enables the CPU and the DMA clocks to bring the system back to its normal state. This is the equivalent of the “Interrupt_̄” of **FIG. 2**. Interrupt Request is synchronized to

avoid confusing the state machine so that Interrupt (INT-DET) will only be detected while the cycle is active. The rising edge of RESCPU will wake up the FE 3001 which in turn releases the whole system from the Sleep Mode.

[0081] Implementation for the 386SX is different only in the external hardware and software power conservation loop. The software loop will set external hardware to switch to the high speed clock on interrupt prior to vectoring the interrupt. Once return is made to the power conservation software, the high speed clock cycle will be detected and the hardware will be reset for full clock operation.

[0082] Implementation for OS/2 uses the “do nothing” loop programmed as a THREAD running in background operation with low priority. Once the THREAD is activated, the CPU sleep, or low speed clock, operation will be activated until an interrupt occurs thereby placing the CPU back to the original clock rate.

[0083] Although interrupts have been employed to wake up the CPU in the preferred embodiment of the present invention, it should be realized that any periodic activity within the system, or applied to the system, could also be used for the same function.

[0084] While several implementations of the preferred embodiment of the invention has been shown and described, various modifications and alternate embodiments will occur to those skilled in the art. Accordingly, it is intended that the invention be limited only in terms of the appended claims.

Computer Programs Listing

[0085] 1) Interrupt 8 Timer interrupt service—pages 27 to 32. Interrupt 8 Timer interrupt service is loaded onto the CPU ROM or an external RAM and is an interrupt mask that may be called at Step 240 of IDLE loop 60 or at Step 460 of ACTIVITY loop 70.

[0086] 2) CPU Sleep Routine—page 33. CPU Sleep Routine is loaded onto the CPU ROM or an external RAM and is a file that may be called at Step 250 of IDLE loop 60 or ACTIVITY loop 70.

[0087] 3) FILE=FORCE5.ASM—pages 34 to 38. FILE=FORCES.ASM is a PCI multiple sleep program that is loaded onto the CPU ROM or an external RAM and is a file that may be called at Step 250 of IDLE loop 60 or ACTIVITY loop 70.

[0088] 4) FILE=Thermal.EQU—listed on page 39. FILE=Thermal.EQU is loaded onto the CPU ROM or an external RAM and is a file that may be called at STEP 240 of IDLE loop 60 or at Step 460 of ACTIVITY loop 70.

1-23. (canceled)

24. A method, comprising the steps of:

determining temperature and a work load level associated with a processor; and

using results of said determining for reducing power consumption associated with said processor as said work load level decreases.

25. The method of claim 24, wherein an amount of said reducing power consumption is proportional to the decrease of said work load level.

26. The method of claim 24, wherein said reduction in power consumption is accomplished in incremental steps.

27. The method of claim 24, wherein said reduction in power consumption continues until no decrease in work load level is detected over a previous determination of work load level.

28. The method of claim 24, wherein said reducing power consumption continues until one of: a) no decrease in work load level is detected over a previous determination of work load level; b) said processor has reached its minimum power consumption level; and c) said temperature is at and/or above a reference temperature.

29. The method of claim 24, wherein said power consumption is reduced by lowering a clock frequency.

30. The method of claim 29, wherein said clock frequency is lowered in incremental steps.

31. The method of claim 24, wherein said power consumption is accomplished by lowering a clock speed.

32. The method of claim 31, wherein said clock speed is lowered in incremental steps.

33. The method of claim 24, wherein said reduction in power consumption is accomplished while said processor is processing data.

34. The method of claim 33, wherein said data is part of a program being run on said processor.

35. A method, comprising the steps of:

determining temperature and a work load level associated with a processor; and

using results of said determining for increasing power consumption associated with said processor as said work load level increases.

36. The method of claim 35, wherein an amount of said increasing power consumption is proportional to the increase of said work load level.

37. The method of claim 36, wherein said power consumption is accomplished in incremental steps.

38. The method of claim 35, wherein said increasing power consumption continues until no increase in work load level is detected over a previous determination of work load level.

39. The method of claim 35, wherein said increasing power consumption continues until one of: a) no increase in work load level is detected over a previous determination of work load level; b) said processor has reached its maximum power consumption level; and c) said temperature is at and/or above a reference temperature.

40. The method of claim 35, wherein said power consumption is accomplished by raising a clock frequency.

41. The method of claim 40, wherein said clock frequency is raised in incremental steps.

42. The method of claim 35, wherein said power consumption is accomplished while said processor is processing data.

43. The method of claim 42, wherein said data is part of a program being run on said processor.

44. A method, comprising the steps of:

determining temperature and a work load level associated with a processor; and

using results of said determining for reducing power consumption associated with said processor as said work load level decreases and/or said temperature is at and or above a reference temperature and increasing power consumption associated with said processor as said work load level increases.

45. A method, comprising the steps of:
determining temperature associated with a processor;
determining processing demand on said processor; and
using results of said determining temperature and determining processing demand for reducing power consumption associated with said processor.

46. A method, comprising the steps of:
determining temperature associated with a processor;
determining processing demand on said processor; and
using results of said determining temperature and determining processing demand for increasing power consumption associated with said processor as said processing demand increases.

47. A method, comprising the steps of:
determining temperature associated with a processor;
determining processing demand on said processor; and
using results of said determining temperature and determining processing demand for reducing power consumption associated with said processor as said processing demand decreases and/or said temperature is at and or above a reference temperature and increasing power consumption associated with said processor as said work load level increases.

* * * * *