



US 20050204179A1

(19) **United States**(12) **Patent Application Publication****Watts, JR. et al.**(10) **Pub. No.: US 2005/0204179 A1**(43) **Pub. Date: Sep. 15, 2005**(54) **METHOD FOR CONTROLLING POWER CONSUMPTION ASSOCIATED WITH A PROCESSOR**(76) Inventors: **LaVauchn F. Watts JR.**, Temple, TX (US); **Steven J. Wallace**, Temple, TX (US)

Correspondence Address:

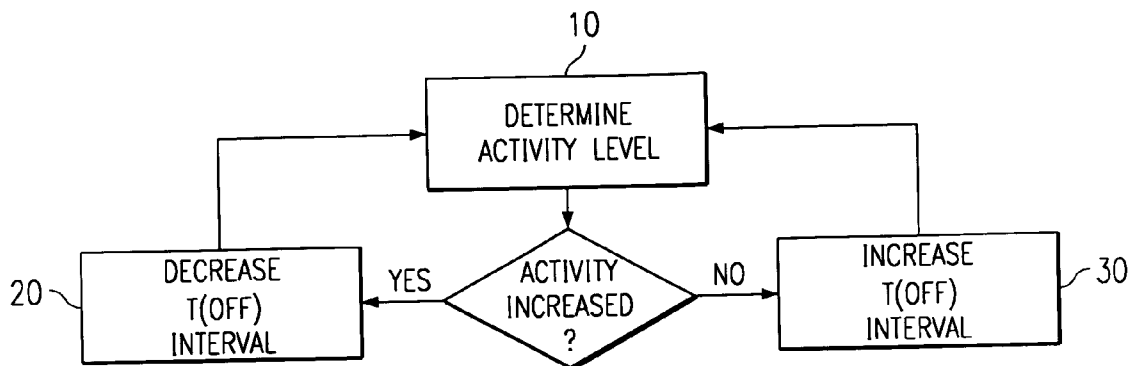
**TEXAS INSTRUMENTS INCORPORATED  
P O BOX 655474, M/S 3999  
DALLAS, TX 75265**

6,633,988, which is a continuation of application No. 09/756,838, filed on Jan. 9, 2001, now Pat. No. 6,397,340, which is a continuation of application No. 09/392,205, filed on Sep. 8, 1999, now Pat. No. 6,173,409, which is a continuation 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.

**Publication Classification**(21) Appl. No.: **11/123,464**(51) **Int. Cl.<sup>7</sup>** ..... **G06F 1/26**(52) **U.S. Cl.** ..... **713/322**(22) Filed: **May 6, 2005**(57) **ABSTRACT****Related U.S. Application Data**

(63) Continuation of application No. 10/837,172, filed on Apr. 30, 2004, which is a continuation of application No. 10/375,982, filed on Feb. 28, 2003, now Pat. No. 6,732,283, which is a continuation of application No. 10/074,739, filed on Feb. 11, 2002, now Pat. No.

A method for controlling power consumption associated with a processor in which, depending on the respective embodiment, a relative amount of idle time, activity time, or idle time and activity time associated with the processor are measured or detected, results of the measuring being used by the processor for controlling a clock speed.



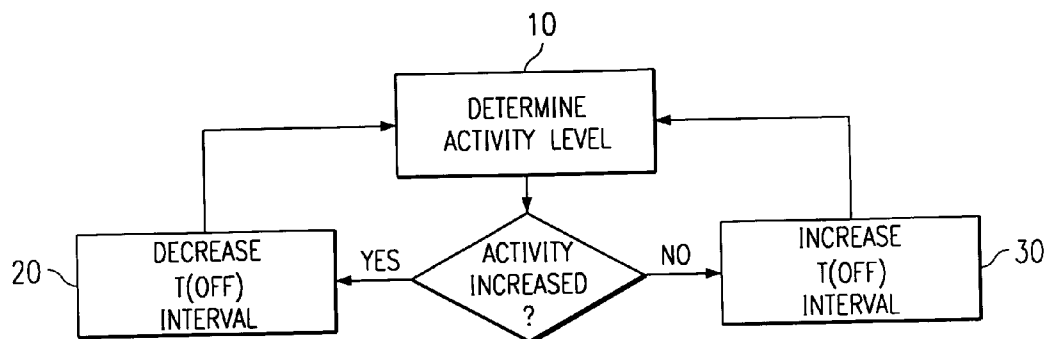


FIG. 1

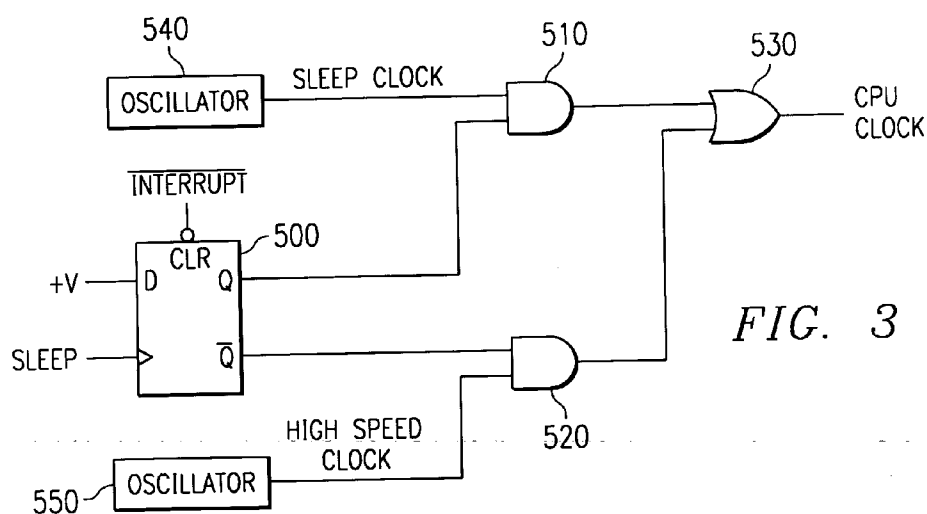


FIG. 3




FIG. 5


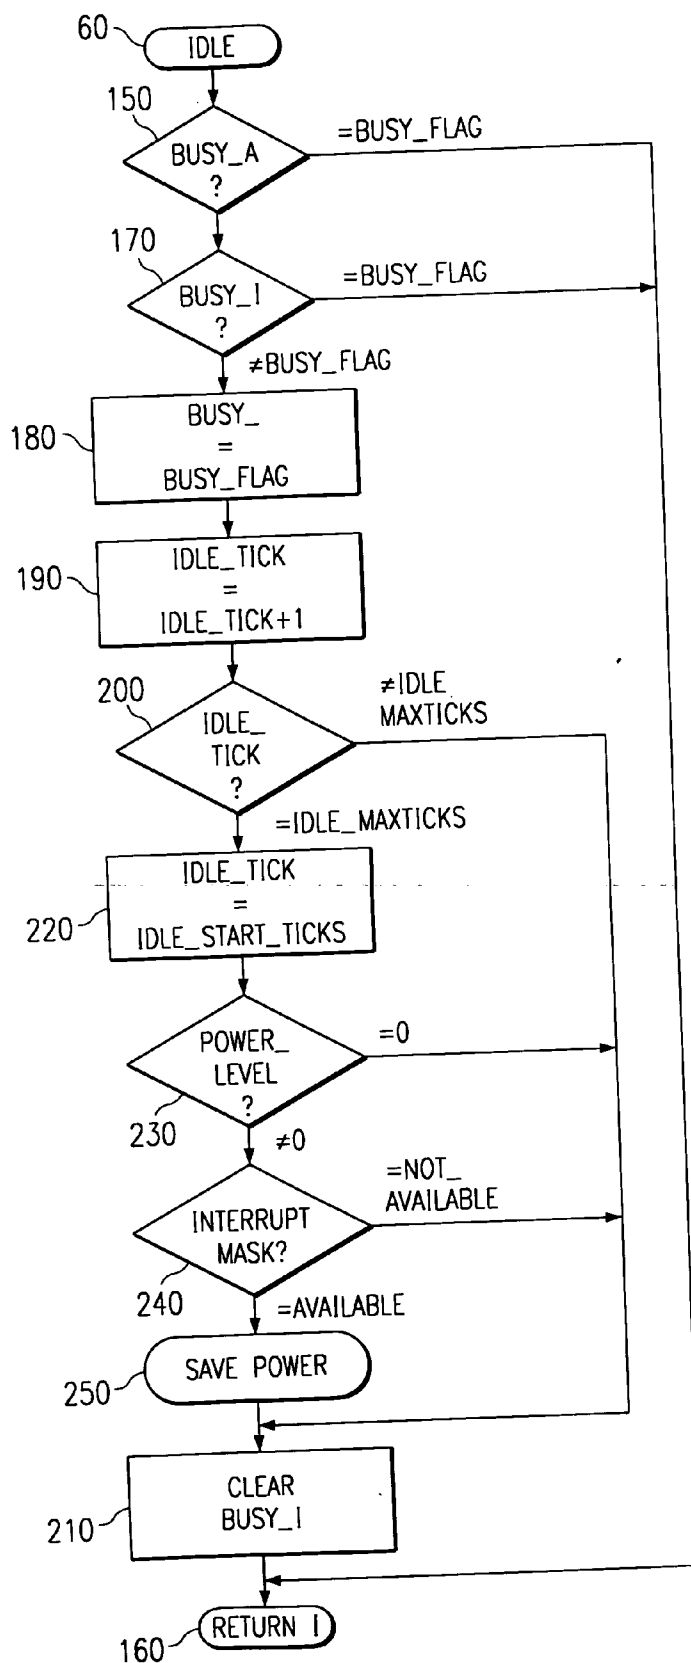


FIG. 2a

FIG. 2b






FIG. 2c

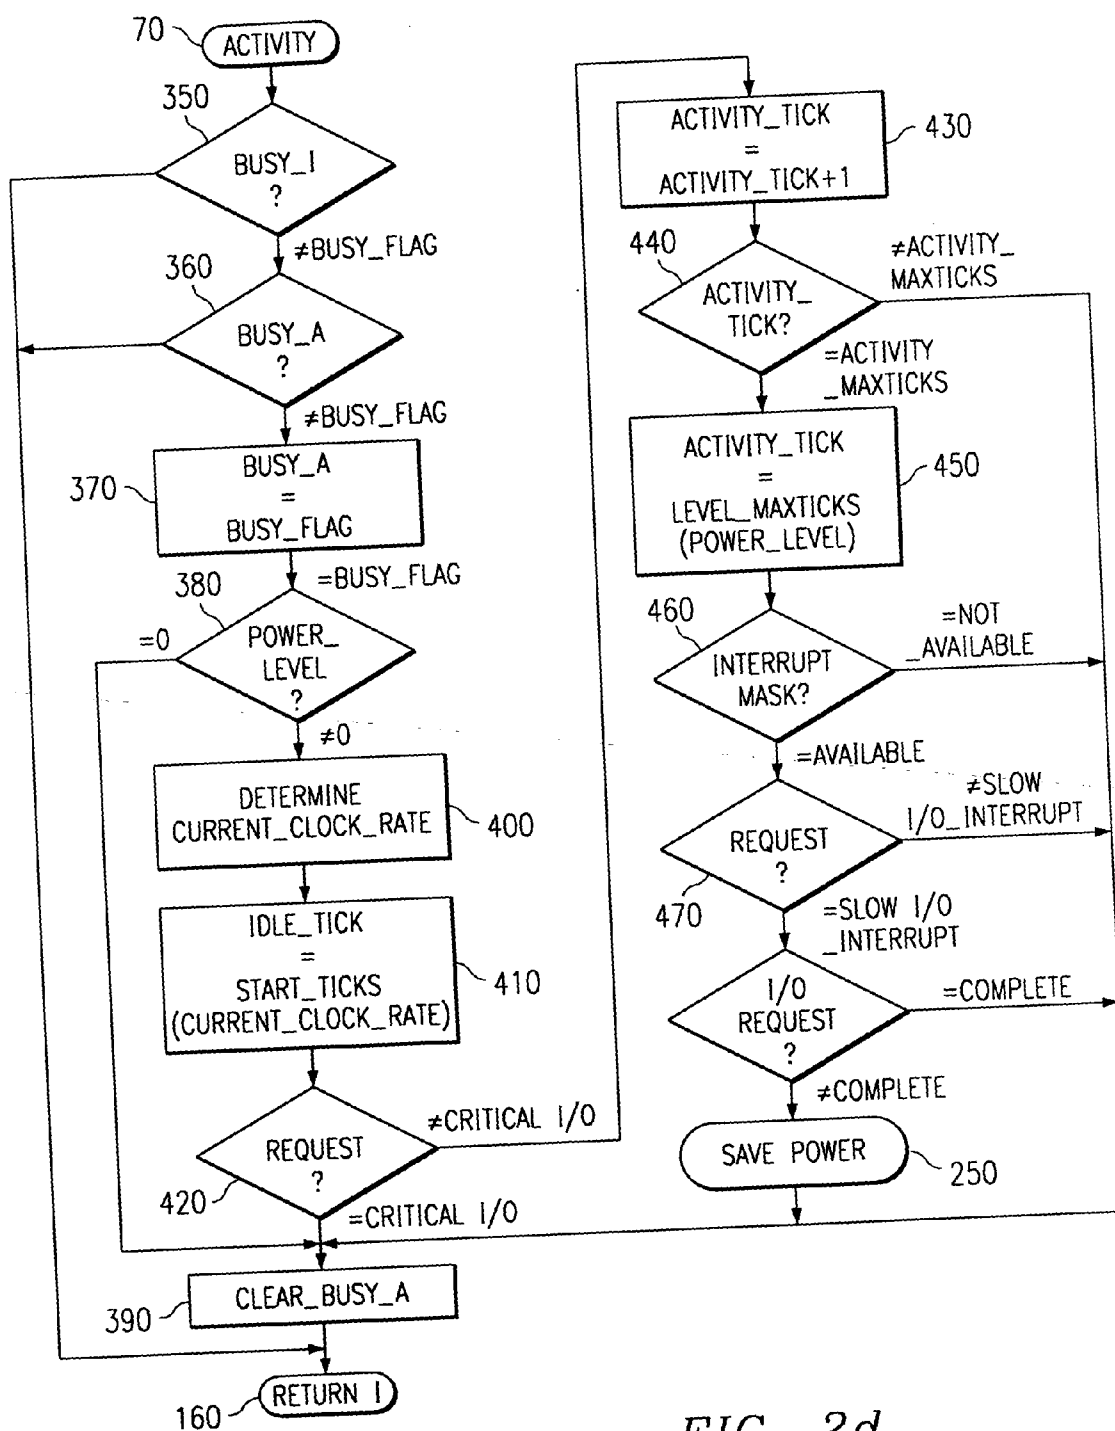
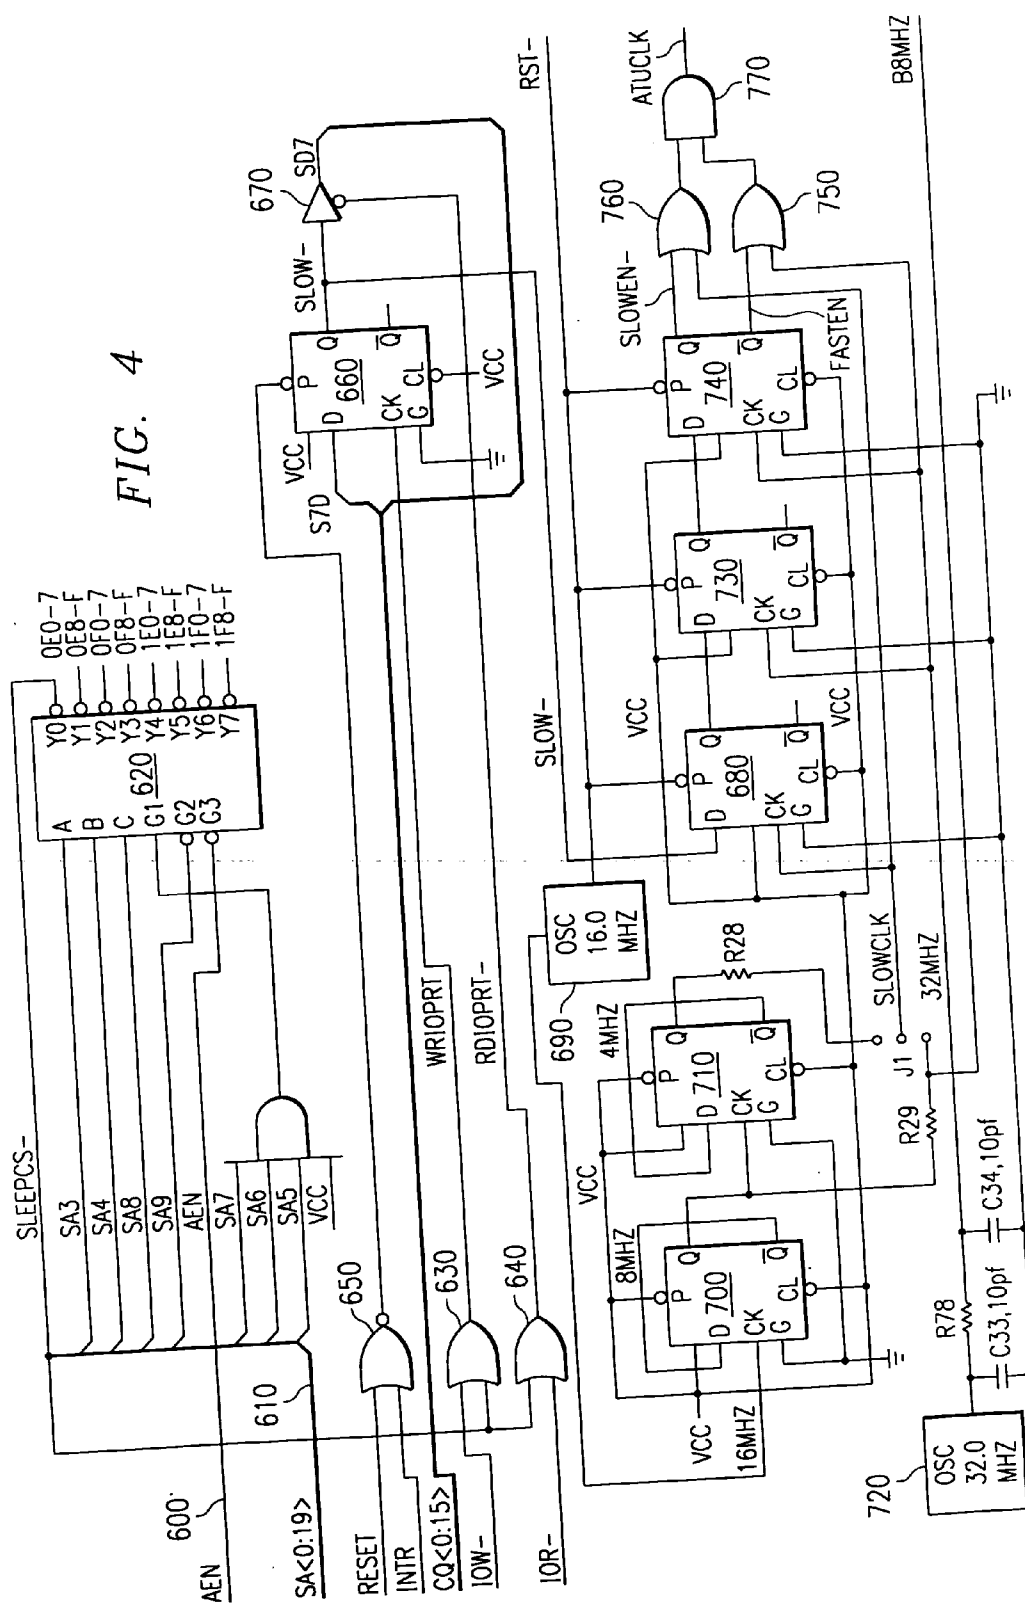


FIG. 2d

FIG. 4



## METHOD FOR CONTROLLING POWER CONSUMPTION ASSOCIATED WITH A PROCESSOR

### BACKGROUND OF THE INVENTION

#### [0001] 1. Field of the Invention

[0002] This invention relates to real-time computer power conservation, and more particularly to an apparatus and method for reduction of central processing unit (CPU) clock time based on the real-time activity level 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 an 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.

[0006] However, the software developed was designed to run on a desk top personal computers, with all the features of a computer, 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.

[0007] As more and more highly functional software packages were developed, desk top computer users experienced increased performance from the introductions of higher computational CPUs, increased memory, and faster high performance disk drives.

[0008] 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 had to settle for A/C operation or very short battery operation to have the performance that was expected from the third party software.

[0009] Portable computer designers stepped the performance down to 8088- 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.

[0010] In an attempt to design a portable computer that could conserve power, thereby yielding longer battery operation, smaller units, and less weight some possible computer designers proceeded to reduce power consumption

of a portable computer while an 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 is spinned up and the computer system is ready again for full performance before the operator may proceed with the operation.

[0011] 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.

[0012] 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.

[0013] Furthermore, although an 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.

[0014] 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 or the programmed length of time. The advantage of this type of power conservation over 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.

[0015] 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, communica-



tions, print spooling, and other like operations cannot be performed. Some existing portable computers employ these circuits. After 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.

#### SUMMARY OF THE INVENTION

[0016] 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 for computer systems without any real-time performance degradation, such conservation of power remaining transparent to the user.

[0017] Another object of the present invention is to provide an apparatus and method for predicting the activity level within a computer system and using the prediction for automatic power conservation.

[0018] Yet another object of the present invention is to provide an apparatus and method which allows user modification of automatic activity level predictions and using the modified predictions for automatic power conservation.

[0019] 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.

[0020] These objects are accomplished in a preferred embodiment of the present invention by an apparatus and method which determine whether a CPU may rest based upon the CPU activity level 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.

[0021] The present invention examines the state of CPU activity, as well as the activity of both the operator and any application software program currently active. This sampling of activity is performed real-time, adjusting the performance level of the computer to manage power conservation and computer power. These adjustments are accomplished within the CPU cycles and do not affect the user's perception of performance.

[0022] 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 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.

[0023] These and other features and advantages of the invention will be apparent to those skilled in the art from the

following detailed description of a preferred embodiment, taken together with the accompanying drawings, in which:

#### DESCRIPTION OF THE DRAWINGS

[0024] FIG. 1 is a flowchart depicting the self-tuning aspect of a preferred embodiment of the present invention;

[0025] FIGS. 2a-2d are flowcharts depicting the active power conservation monitor employed by the present invention;

[0026] FIG. 3 is a simplified schematic diagram representing the active power conservation associated hardware employed by the present invention;

[0027] FIG. 4 is a schematic of the sleep hardware for one embodiment of the present invention; and

[0028] FIG. 5 is a schematic of the sleep hardware for another embodiment of the present invention.

#### DETAILED DESCRIPTION OF A PREFERRED EMBODIMENT

[0029] 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.

[0030] According to the preferred embodiment of the present invention, to conserve power 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.

[0031] Slowing down or stopping the computer system components according to the preferred embodiment of the present invention, reduces power consumption, although the amount of power saved may vary. Therefore, according to the present invention, stopping the clock (where possible as some CPUs cannot have their clocks stopped) reduces the power consumption more than just slowing the clock.

[0032] In general, the number of operations (or instructions) per second may be considered to be roughly proportional to the processor clock:

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

[0033] Assuming for simplicity that the same instruction is repeatedly executed so that instructions/second is constant, the relationship can be expressed as follows:

$$Fq = K_1 * \text{Clk}$$

[0034] where Fq is instructions/second, K<sub>1</sub> is constant equal to the instructions/cycle, and Clk equals cycles/sec-

ond. Thus, roughly speaking, the rate of execution increases with the frequency of the CPU clock.

[0035] 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:

$$P=K_2+(K_3*\text{Clk})$$

[0036] where P is power in watts,  $K_2$  is a constant in watts,  $K_3$  is a constant and expresses the number of watt-seconds/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.

[0037] Assume that a given time period T is divided into N intervals such that the power P was constant during each interval. Then the amount of energy extended during T would be given by:

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

[0038] 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:

$$\begin{aligned} E(\text{max}) &= P(\text{on}) * (\text{delta} T_1 + \text{delta} T_2 \ldots + \text{delta} T_N) \\ &= P(\text{on}) * T \end{aligned}$$

[0039] 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" is summed into the quantity "T(on)" and the "OFF" intervals are summed into "T(off)", then it follows:

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

[0040] Now the energy being used during period T can be written:

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

[0041] 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.

[0042] 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 level is 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 an increase over an immediately previous determination, the present invention decreases 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 a decrease over an immediately previous determination, the present invention increases the T(off) interval (Step 30) and proceeds to again determine the activity level of the CPU. Thus the T(off) intervals are constantly being adjusted to match the system activity level.

[0043] 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 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 activity level. 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.

[0044] Once the CPU is interrupted during a power conservation 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 software was operating during this slice, control will be returned to the active power conservation loop (monitor 40) which simply monitors the CPU's clock to determine an exit condition for the power conservation mode, thereby exiting from T(off) to T(on) state. The interval of the next power conservation state is adjusted by the activity level monitored, 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 loop to be exited automatically and executing an interval T(on).

[0045] More specifically, looking now at **FIGS. 2a-2d**, which depict the active power conservation 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" loop (i.e., planned inactivity). The ACTIVITY branch 70 of the flowchart, 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 monitor **40** has been fully executed, whenever control is yielded from the operating system to the power conservation mode, either IDLE **60** or ACTIVITY **70** branches are selected depending on the type of CPU activity: IDLE branch **60** for power conservation during planned inactivity and ACTIVITY branch **70** for power conservation during CPU activity.

[0046] 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 uses the DEFAULT\_LEVEL. Otherwise, it continues onto Step **110** where it modifies the Power\_level to equal the User\_level.

[0047] 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.

[0048] 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 monitor originally, at RETURN **80**.

[0049] Consider now IDLE branch **60** of active power 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.)

[0050] 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**.

[0051] 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 level. Idle\_tick increments by one to allow for smoothing of events, hereby letting a critical I/O activity control smoothing.

[0052] 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 level. 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**.

[0053] 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.

[0054] 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.

[0055] 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 activity interrupts and tends to provide more T(off) intervals as the activity level slows. 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**.

[0056] 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 CURRNT\_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 SLEEPCLOCK 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.

[0057] 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 establishes the CPU at the stored Save-clock-rate. If the sleep clock rate was not stopped, in other words, the sleep clock rate was not zero, control is passed at a slow clock and SAVE POWER 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).

[0058] 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 IDLE branch 60 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.

[0059] If however, the Busy\_I flag does not equal BUSY\_FLAG, which means IDLE branch 60 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.

[0060] 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 380.

[0061] 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).

[0062] 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 therefore when the power conservation aspect of the present invention may be activated. Once the power conservation 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 to full speed CPU clock operation, the preferred embodiment of the present invention employs some associated hardware.

[0063] Looking now at FIG. 3 which shows a simplified schematic diagram representing the associated hardware employed by the present invention or active power conservation. When monitor 40 (not shown) determines the CPU is ready to sleep, it writes 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 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.

[0064] It should be noted that the associated hardware needed, external to each of the CPUs for any given system, may be different based on the operating system used, whether the CPU can be stomped, 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 in the numerous available portable computer systems. For example two actual implementations are shown in FIGS. 4 and 5, discussed below.

[0065] 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.

[0066] 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 or 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.

[0067] Looking now at FIG. 4, which depicts a schematic of an actual sleep hardware implementation for a system such as the Intel 80386 (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.

[0068] 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".

[0069] 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.

[0070] 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-").

[0071] 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 8 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.

[0072] Consider now FIG. 5, which depicts a schematic of another actual sleep hardware implementation for a system

such as the Intel 80286 (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  $P(i)\Delta T_i$  interval after the interrupt.

[0073] 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 (INTDET) 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.

[0074] 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.

[0075] 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.

[0076] 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.

[0077] 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.

1-21. (canceled)

22. A method, comprising the steps of:

determining 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.

23. The method of claim 22, wherein an amount of said reducing power consumption is proportional to the decrease of said work load level.

24. The method of claim 23, wherein said reduction in power consumption is accomplished in incremental steps.

25. The method of claim 22, wherein said reduction in power consumption continues until no decrease in work load level is detected over a previous determination of work load level.

26. The method of claim 25, 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; and b) said processor has reached its minimum power consumption level.

27. The method of claim 22, wherein said power consumption is accomplished by lowering a clock frequency.

28. The method of claim 27, wherein said clock frequency is lowered in incremental steps.

29. The method of claim 22, wherein said power consumption is accomplished by lowering a clock speed.

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

31. The method of claim 22, wherein said reduction in power consumption is accomplished while said processor is processing data.

32. The method of claim 31, wherein said data is part of a program being run on said processor.

33. A method, comprising the steps of:

determining 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.

34. The method of claim 33, wherein an amount of said increasing power consumption is proportional to the increase of said work load level.

35. The method of claim 34, wherein said power consumption is accomplished in incremental steps.

36. The method of claim 33, wherein said increasing power consumption continues until no increase in work load level is detected over a previous determination of work load level.

37. The method of claim 33, 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; and b) said processor has reached its maximum power consumption level.

38. The method of claim 33, wherein said power consumption is accomplished by raising a clock frequency.

39. The method of claim 38, wherein said clock frequency is raised in incremental steps.

40. The method of claim 33, wherein said power consumption is accomplished by raising a clock speed.

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

42. The method of claim 33, 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 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 increasing power consumption associated with said processor as said work load level increases.

\* \* \* \* \*