

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
22 April 2010 (22.04.2010)

PCT

(10) International Publication Number
WO 2010/045124 A2

(51) International Patent Classification:
G06F 1/20 (2006.01)

(21) International Application Number:
PCT/US2009/060253

(22) International Filing Date:
9 October 2009 (09.10.2009)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
12/250,498 13 October 2008 (13.10.2008) US

(71) Applicant (for all designated States except US): **APPLE INC.** [US/US]; 1 Infinite Loop, Cupertino, CA 95014 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **COX, Keith** [US/US]; 1234 Colleen Way, Campbell, CA 95008-1712 (US). **KAPOOR, Gaurav** [IN/US]; 1901 Halford Avenue, Apt49, Santa Clara, CA 95051-7404 (US). **ARNOLD, Vaughn** [US/US]; 621 Canepa Drive, Scotts Valley, CA 95066-4702 (US).

(74) Agents: **VINCENT, Lester, J.** et al.; **BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP**, 1279 Oakmead Parkway, Sunnyvale, CA 94085-4040 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— without international search report and to be republished upon receipt of that report (Rule 48.2(g))



WO 2010/045124 A2

(54) Title: A METHOD FOR ESTIMATING TEMPERATURE AT A CRITICAL POINT

(57) Abstract: Methods and apparatuses are disclosed to estimate temperature at one or more critical points in a data processing system comprising modeling a steady state temperature portion of a thermal model at the one or more critical points using regression analysis; modeling the transient temperature portion of the thermal model at the one or more critical points using a filtering algorithm; and generating a thermal model at the one or more critical points by combining the steady state temperature portion of the thermal model with the transient temperature portion of the thermal model. The thermal model may then be used to estimate an instantaneous temperature at the one or more critical points or to predict a future temperature at the one or more critical points.

A METHOD FOR ESTIMATING TEMPERATURE AT A CRITICAL POINT

FIELD OF THE INVENTION

[0001] At least certain embodiments of the invention relate generally to data processing systems, and more particularly but not exclusively to the estimation of temperatures in data processing systems.

BACKGROUND OF THE INVENTION

[0002] Computer systems are designed to be able to continuously operate under varying degrees of thermal loading. As the operating power consumption of computer systems creeps upwards, the thermal budgets of these systems have become tighter. It is now a challenge to design a computer system to operate continuously under worst-case thermal loading for extended periods of time, while at the same time both maintaining a desired range of operating performance and maintaining the temperature of the computer housing within safety limits.

[0003] Further, computer systems are being packed into smaller profile housing to improve portability. As a result, many electronic devices have “hotspots” where there is a potential to overheat and cause injury to the person using the electronic device. These hotspots are the result of computer systems such as laptop computers having both compact housing (also called skin) and limited cooling capacity. The UL and other government bodies establish safety standards for products. As a result, manufacturers of computer systems and other electronic products are required to limit the skin temperature to levels within government approved safety limits so that they are free from reasonably foreseeable risks of fire, burn and other related hazards.

[0004] To resolve thermal issues, several proactive measures have been developed including thermal-throttling, engaging cooling fans, providing heat exhaust pipes, managing and redistributing power, switching from external graphics to internal graphics, and etcetera. Many systems designed today require these proactive measures to provide adequate performance under various thermal

conditions. Temperature sensors may be used to determine the ambient temperature of the working environment of computer systems. Thermal models may then be developed based on data collected from these temperature sensors. [0005] However, temperature sensors are often implemented as integrated circuits which can only be placed at a limited number of locations around the computer system and its skin. Also, thermal sensors can be integrated circuits, so using a large number of thermal sensors may in some cases be cost-prohibitive. **FIG. 1** illustrates an example of a location to measure ambient temperature for thermal control. In the illustrated embodiment, temperature sensor 101 is positioned in the middle portion of the display panel of a portable computer 100. Temperature sensors may be placed in various places within a computer system such as, for example, on the microprocessor (CPU) die, in the proximity of the CPU die (CPU Prox), on known hotspots such as a motherboard, on a hard drive device, on a heat pipe, on a battery, on a track pad, and etcetera.

[0006] Often, a hotspot may be located on a particular point where it is difficult to put a sensor such as the bottom case of a laptop computer system. These points may be generally referred to as critical points. Thermal models may be generated to accurately estimate the temperature at these critical points based on data received from temperature sensors to estimate the temperature at the critical point. These thermal models may be generated so that proactive measures may be taken to mitigate the thermal behavior at hotspots such as turning on the system fans and so on.

SUMMARY OF THE DESCRIPTION

[0007] At least certain embodiments of the invention disclose methods and apparatuses to estimate temperature at one or more critical points or locations in a data processing system including generating a representation of a thermal model at the one or more critical points, the thermal model using a filtering algorithm to filter at least one temperature sensor at a plurality of frequencies to compensate for time-based under estimation and over estimation in the thermal model. The thermal model can then be stored as one or more equations which

can be used at runtime to predict a current and/or future temperature at a critical point on a data processing system. The filtering algorithm (or alternative algorithm) is used to create a model of a transient temperature curve (or other representation) at the one or more critical points. The thermal model can be generated by merging a model of a steady state temperature curve (or other representation) with the thermal model of the transient temperature curve (or other representation). The merged model may be used for a variety of different data processing systems in a class of data processing systems (e.g., a set of laptop computers having a certain housing and display, but different microprocessors or other internal differences); the merged model may be stored as one or more equations in a memory of such a system.

[0008] Additionally, some embodiments include methods and apparatuses to generate a future thermal model based on the thermal model to predict future temperatures at the critical point based on the future thermal model. These future predictions may be used to perform proactive measures to reduce temperature at the critical point based on the future temperature model. These proactive measures may include any one or more of thermal-throttling the data processing system, engaging cooling fans, managing and redistributing power, and switching from external graphics to internal graphics.

BRIEF DESCRIPTION OF THE DRAWINGS

[0009] A better understanding of at least certain embodiments of the invention can be obtained from the following detailed description in conjunction with the following drawings, in which:

[0010] **FIG. 1** illustrates an example of a location to measure ambient temperature for thermal control.

[0011] **FIG. 2** illustrates an autocorrelation of a CPU die sensor and a skin temperature of a computer or other data processing system.

[0012] **FIG. 3A** illustrates a graph of a curve generated by actual data from a thermocouple device connected to the critical point.

- [0013] **FIG. 3B** illustrates a Biplot of the principal components of a set of temperature sensors and the critical point.
- [0014] **FIG. 4** illustrates the problem of underestimation and overestimation using simple exponential smoothing.
- [0015] **FIG. 5A** illustrates the three frequencies algorithm according to an exemplary embodiment of the invention.
- [0016] **FIG. 5B** illustrates the respective temperature curves at each of the three frequencies with respect to the actual temperature curve of the critical point according to an exemplary embodiment of the invention.
- [0017] **FIG. 5C** illustrates an estimation of the actual critical point temperature curve using the three frequencies algorithm according to an exemplary embodiment of the invention.
- [0018] **FIG. 6A** illustrates predicting future values of temperature at a critical point using the three frequencies algorithm according to an exemplary embodiment of the invention.
- [0019] **FIG. 6B** illustrates a method of estimating temperature at a critical point using the three frequencies algorithm according to an exemplary embodiment of the invention.
- [0020] **FIG. 7** illustrates an exemplary data processing system upon which the methods and apparatuses of the invention may be implemented.

DETAILED DESCRIPTION

[0021] Throughout the description, for the purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of the present invention. It will be apparent, however, to one skilled in the art that the present invention may be practiced without some of these specific details. In other instances, well-known structures and devices are shown in block diagram form to avoid obscuring the underlying principles of embodiments of the present invention.

[0022] Methods and apparatuses for estimating temperature at a critical point (e.g., the bottom outside housing of a laptop computer) are described. At least

certain embodiments of the invention disclose a method and apparatus to estimate temperature at a critical point in a data processing system including receiving temperature data from temperature sensors, separating the estimation of the steady state temperature curve from the estimation of the transient temperature curve at the critical point, and merging the model of the steady state temperature curve with the model of the transient temperature curve to generate a thermal model of the temperature at the critical point. These embodiments may then estimate the instantaneous temperature at the critical point based on the thermal model generated from the merging.

[0023] In addition, some embodiments generate a future thermal model based on the thermal model to predict future temperatures at the critical point. These future temperature predictions may then be used to initiate proactive measures to reduce temperature at the critical point or to run the system at higher performance levels while still keeping the temperature at the critical point below an acceptable level. These proactive measures may include any one or more of thermal-throttling the data processing system, engaging cooling fans, providing heat exhaust pipes, managing and redistributing power, and switching from external graphics to internal graphics.

[0024] It is often desired to determine the temperature of critical points like the bottom case of a laptop computer or portable Personal Digital Assistant (PDA) device or other device such as a handheld computer or wireless telephone where it is difficult to put a temperature sensor. Thus, the temperature at the critical point (where there is no temperature sensor) may be estimated based on the temperature at a known point such as a temperature sensor. Direct and indirect thermal models may be developed to model temperature in a data processing system and to determine the desired temperature at the critical point based on the thermal behavior at a plurality of temperature sensors that may be located in various locations within the data processing system. Direct modeling uses the instantaneous power, ambient temperature, and air flow, whereas indirect modeling requires estimating the critical temperature with the help of the temperature sensors placed in feasible places. The presence of compounding

factors like the distribution of the power in the system, inability to measure room temperature accurately, etcetera, result in a far too complex model resulting in inaccuracy in the direct modeling case. Additionally, the models developed using direct modeling fail to generalize well from one data processing system to other data processing systems.

[0025] One method of performing indirect thermal modeling includes performing regression analysis. Regression analysis may be used to estimate the thermal behavior at the particular critical point based on the known thermal behavior at the thermal sensors. Data regression analysis is a technique used for modeling and analysis of numerical data consisting of values of a dependent variable (response variable) and of one or more independent variables (explanatory variables). The dependent variable in the regression equation is modeled as a function of the independent variables, corresponding constants plus an error term. The error term is treated as a random variable representing unexplained variation in the dependent variable. The parameters are estimated so as to give a "best fit" of the data. Most commonly the best fit is evaluated by using the least squares method, but other criteria have also been used. Regression can be used for prediction (including forecasting of time-series data), inference, hypothesis testing, and modeling of causal relationships.

[0026] These uses of regression rely heavily on the underlying assumptions being satisfied. However, the errors between the temperature at a particular sensor and the temperature at the critical point violate the underlying assumptions of regression analysis because the errors are correlated in time.

FIG. 2 which illustrates the autocorrelation of a CPU die sensor and the skin temperature of a computer or other data processing system. As can be seen, the errors between the two autocorrelation functions vary with time and a fundamental assumption of regression analysis is that errors do not correlate in time. Thus, linear regression analysis yields unsatisfactory results since the fundamental assumption of linear regression is violated.

[0027] Another option is to perform techniques such as Multivariate Time Series Analysis like vector autoregression functions, and etcetera. However,

these generally result in models that are computationally prohibitive to use in real time.

[0028] At least certain embodiments of the invention include separating the steady state temperature estimation from the transient temperature estimation so that regression analysis techniques may be performed to model the steady state temperature because the steady state temperature does not vary in time; and as a result, standard regression analysis techniques may be used. These embodiments may then use a method referred to as the “Three Frequencies Analysis” or a filtering algorithm or other frequency or time based approaches to approximate the results of the transient response to power fluctuations within the data processing system. Finally, the two approximations are combined to generate a complete thermal model of the temperature at the critical point. This model may then be used to measure the temperature of the critical point at any time instant. Additionally, this thermal model may be modified to generate a fairly accurate prediction of temperature at the critical point. This information may then be leveraged to take proactive measures to mitigate the temperature at a critical point. This is especially important when the critical point is a hotspot on the data processing device such as the bottom of a laptop or PDA.

[0029] FIG. 3A illustrates a graph of a curve generated by actual data from a thermocouple device connected to the critical point. In this figure, power is applied at various levels to a hypothetical data processing system and then subsequently removed. This graph shows the temperature in response to the applied power over time. At time t_0 , power is removed from the system and the temperature begins decreasing toward the steady state value when no power is applied. This steady state is reached at time t_1 where all power has been removed from the system for a long period of time. This is the first steady state portion of the graph. Power is subsequently applied to the system and the temperature rises accordingly until it reaches a different steady state temperature at time t_2 . Transition portion 321 illustrates this increase in temperature over time. The steady state temperature at time t_2 may be, for example, the steady state at full operation of the data processing system including CPU power, hard

drives, monitor, and etcetera. After this time, power is once again removed from the system and the temperature decreases to a third steady state value at time t3. This steady state value may be different than the steady state value at time t1. For example, the steady state value may correspond to power applying power to only the hard drives. After time t3, power is reapplied to the system until the temperature reaches another steady state value at time t4. This steady state value may correspond, for example, to powering up the Graphics Processing Unit (GPU) only. After time t4, power is partly removed from the system until the temperature reaches the steady state value at time t5. This steady state value at time t5 may correspond, for example, to powering only one-half of the CPU. Alternative power levels are considered to be within the teachings of this description.

[0030] Thus, **FIG. 3A** demonstrates the actual temperature variation with time of the critical point. However, as discussed above, it may not be feasible to place a sensor at the critical point. In this case, the temperature of the critical point may be approximated using the modeling techniques described herein.

I. Steady State Temperature Estimation

[0031] At least certain embodiments describe performing steady state temperature estimation separately from the transient temperature estimation because the errors (ϵ) are not correlated in time at the steady state. The first step may be to pick the minimum number of temperature sensors with which to perform the steady state analysis. This minimum number of temperature sensors may be determined by performing a principal component analysis to select the best set of temperature sensors from which to perform the regression analysis at the critical point. That is, principal component analysis may be performed to determine which temperature sensors are most likely to produce accurate results using regression analysis at the critical point.

[0032] **FIG. 3B** illustrates a Biplot of the principal components of a set of temperature sensors and the critical point. The temperature sensors in the illustrated embodiment include CPU die 302, CPU Prox 303, Heat Pipe 304, Battery 305, and Trackpad 306. However, this is given by way of illustration

and not of limitation as any temperature sensors within the data processing systems are contemplated. The direction of each of the vectors in the Biplot corresponds to how the particular temperature sensor or critical point behaves under the application of power to the system. The Biplot indicates that there are two groups of temperature sensors that exhibit similar response to the application of power to the system. The first group consists of the CPU die 302, CPU Prox 303, Heat Pipe 304. The heat sensors in the first group are relatively correlated based on the Biplot because their thermal behavior is similar manner when power is applied to the system. The second group consists of the Battery 305 and Trackpad 306. The heat sensors in the second group are relatively correlated based on the Biplot because their thermal behavior is similar manner when power is applied to the system.

[0033] Based on the Biplot, it is clear that the thermal behavior of any of the sensors in the first group is relatively uncorrelated with the thermal behavior of any of the sensors in the second group. In embodiments, the minimum number of sensors may be a pair of sensors such that the pair explains the maximum variance with respect to the critical point. This may yield the best results using regression analysis. In the illustrated embodiment, the pair sensors may include the CPU die 302 and the Track Pad 306 since they are the most divergent sensors in the Biplot. If required, other sensors may be added such that the addition maximizes the adjusted R-squared values when the critical point is regressed on the set of sensors. This may be an iterative process where additional sensors continue to be added until the gain in the adjusted R-squared is negligible.

[0034] Once the temperature sensors are chosen, the regression analysis is performed to generate a model of the steady state temperature curve in time. On the steady state response, the generated multivariate linear regression equation may be given by the following equation:

$$\text{CriticalTemp} = \beta_0 + \beta_1 * \text{sensor1} + \beta_2 * \text{sensor2} + \beta_3 * \text{sensor 3} + \dots \beta_N * \text{sensorN}$$

where β_0 is an initial value and $\beta_1 - \beta_N$ represent data point weight factors for each of the temperature sensors.

II. Transient Temperature Estimation

[0035] Whenever there is a change in applied power, ambient temperature, or airflow to a computer or other data processing system, temperature at each point will require a different time to reach its steady state value. For example, CPU Die 302 sensor's temperature reaches its steady state value far earlier than the laptop bottom case. Hence, if the CPU Die 302 sensor is used for estimating the laptop bottom case (e.g., critical point), then simple exponential smoothing should be performed on its values so that both the CPU die sensor 302 and the critical point follow the same temperature curve characteristics. Unfortunately, simple exponential smoothing does not perform well. Simple exponential smoothing tends to underestimate the temperature during the start of the temperature transition and overestimate near the middle of the temperature curve. This is demonstrated in **FIG. 4**. **FIG. 4** illustrates the problem of underestimation and overestimation using simple exponential smoothing. In the figure, the results of the estimation using exponential smoothing are compared to the temperature curve of the actual critical point such as laptop case bottom for example. At time t1, the divergence is readily apparent since the estimated curve lags the actual curve of the critical point temperature. This results in an underestimation of critical point temperature and a time lag which may be as much as thirty (30) minutes. At time t2, the two graphs cross and the estimated curve begins to overestimate the temperature at the actual critical point. At time t3, the temperature transition begins to reach the steady state value (where errors do not correlate in time), so the overestimation begins to be reduced, and at time t4, the steady state is almost reached and the exponential smoothing begins to accurately predict the actual temperature of the critical point.

[0036] At least certain embodiments describe a new curve fitting algorithm referred to as the "Three Frequencies Algorithm" to estimate the transient temperature curve at a critical point in a data processing system. The new algorithm is based on analyzing the temperature curve using, in one embodiment, the Three Frequencies algorithm. **FIG. 5A** illustrates the three frequencies algorithm according to an exemplary embodiment of the invention. In the figure,

each of the temperature sensors picked for the steady state analysis are used to perform the transient analysis. The sensors 501 are input into three filters at three different frequencies including a maximum correlation frequency filter 502, a slow-moving frequency filter 503, and a fast-moving frequency filter 504. In the preferred embodiment, these filters may be infinite impulse response filters (IIR). However, other frequency filters are contemplated within the teachings of this description.

[0037] The maximum correlation frequency provides the necessary shape of the overall temperature curve, whereas the fast-moving frequency removes the underestimation from the start of the temperature curve and the slow-moving frequency removes the overestimation at the middle of the temperature curve. The respective temperature curves at each of the three frequencies with respect to the actual temperature curve of the critical point according to an exemplary embodiment of the invention are illustrated in **FIG. 5B**. In the figure, the maximum correlation frequency temperature curve 522 provides the necessary shape of the overall temperature curve, the fast-moving frequency temperature curve 524 removes the underestimation from the start of the temperature curve, and the slow-moving frequency temperature curve 523 removes the overestimation at the middle of the temperature curve. The maximum correlation frequency may correspond to the skin of the data processing system that has some thermal mass. The slow-moving frequency may correspond to the table or other external item upon which the data processing system such as a laptop computer is located. Finally, the fast-moving frequency may correspond to localized heating and is direct and immediate.

[0038] The three filtered values in **FIG. 5A** include Sensor_Max_Correlation 505, Sensor_SlowSpeed 506, and Sensor_FastSpeed 507. These three filtered values are then combined in a weighted manner by coefficients α_1 , α_2 and α_3 to obtain weighted_filtered sensor 515 as shown in **FIG. 5A** resulting in a temperature curve that accurately models to the actual critical point temperature curve with at least some of the underestimation and overestimation in the

resulting temperature curve removed. The three frequencies and their corresponding weights may be determined using an algorithm as follows:

1. Find $\alpha_{\text{max_correlation}}$ such that it maximizes the correlation between the filtered sensor value and the actual thermocouple at the critical point.

The filtered sensor value is calculated as:

$$\text{Sensor_Max_Correlation}_t = \alpha_{\text{max_correl}} * \text{sensor}_t + (1 - \alpha_{\text{max_correl}}) * \text{Sensor_Max_Correlation}_{t-1};$$

2. Find $\alpha_{\text{slow_freq}}$ ($>\alpha_{\text{max_correlation}}$) such that it maximizes the adjusted R-squared when the actual thermocouple at the critical point is explained using **Sensor_Max_Correlation** and **Sensor_SlowSpeed**.

The filtered value may be calculated as:

$$\text{Sensor_SlowSpeed}_t = \alpha_{\text{slow_freq}} * \text{sensor}_t + (1 - \alpha_{\text{slow_freq}}) * \text{Sensor_SlowSpeed}_{t-1};$$

3. Find $\alpha_{\text{fast_freq}}$ ($>\alpha_{\text{max_correlation}}$) such that it maximizes the adjusted R-Squared when the actual thermocouple at the critical point is explained using **Sensor_MaxCorrelation**, **Sensor_SlowSpeed** and **Sensor_FastSpeed**.

The filtered value may be calculated as:

$$\text{Sensor_FastSpeed}_t = \alpha_{\text{fast_freq}} * \text{sensor}_t + (1 - \alpha_{\text{fast_freq}}) * \text{Sensor_FastSpeed}_{t-1}.$$

However, any error gradient descent algorithm may be used and is contemplated within the teachings of this description. The weights $\alpha_{\text{max_correl}}$, $\alpha_{\text{slow_freq}}$, and $\alpha_{\text{fast_freq}}$ of the three frequencies respectively are, in at least certain embodiments, in proportion to the coefficients obtained after fitting a regression equation on the actual thermocouple data and the three frequencies values (Sensor_Max_Correlation, Sensor_SlowSpeed, Sensor_FastSpeed). This resulting curve is shown in **FIG. 5C**, which illustrates an estimation of the actual critical point temperature curve using the three frequencies algorithm according to an exemplary embodiment of the invention. In the figure, it is noted that the estimation using the three frequencies algorithm 532 provides a much more

accurate model of the actual critical point (e.g., laptop case bottom) temperature curve 531 than the temperature estimation curve 533 using simple exponential smoothing (e.g., linear regression analysis).

[0039] The resulting equation using the three frequencies algorithm may include the equation:

$$\text{Sensor1_Filter} = \alpha_1 * \text{sensor1 @ Freq1} + \alpha_2 * \text{sensor1 @ Freq2} + \alpha_3 * \text{sensor1 @ Freq3}.$$

This equation may then be merged with the equation for the temperature curve obtained using the steady state analysis to yield the overall thermal model equation as follows:

$$\text{CriticalTemp} = \beta_0 + \beta_1 * \text{sensor1_Filter} + \beta_2 * \text{sensor2_Filter} + \beta_3 * \text{sensor3_Filter} + \dots \beta_N * \text{sensorN_Filter}.$$

[0040] The teachings of this description may also use the Three Frequencies Algorithm to predict the “future” temperature at a critical point. This can be achieved either by simply changing the weight of the three frequencies generated by the algorithm or by refitting the whole time-shifted temperature curve. **FIG. 6A** illustrates predicting future values of temperature at a critical point using the three frequencies algorithm according to an exemplary embodiment of the invention. In the figure, the future predicted value of temperature at the critical point is obtained after changing the weight of the three frequencies. The actual critical point temperature curve 631, the temperature curve estimated using the Three Frequencies Algorithm 632, and the future predicted temperature curve 633 are each marked in **FIG. 6A** respectively. In the illustrated embodiment, the future predicted value curve 633 reaches the stable point far earlier than the actual temperature curve 631. This fact can be exploited to take proactive measures like bringing on the cooling fans a bit earlier, redistributing the power, or switching from external graphics to internal graphics, and so on.

[0041] **FIG. 6B** illustrates a method of estimating temperature at a critical point using the three frequencies algorithm according to an exemplary embodiment of the invention. Process 600 begins at operation 601 where the estimation analysis

of the steady state and transient temperature curves at the critical point are separated. Process 600 continues with modeling the transient temperature curve using the three frequencies algorithm described above (or using other algorithms to create a model of the transient temperature curve or change) at operation 603 and modeling the steady state temperature curve using regression analysis at operation 605. Process 600 continues at operation 607 where the results of the separately estimated steady state analysis and the transient analysis are combined to generate the complete thermal model at the critical point. Optionally, process 600 concludes with using the thermal model of operation 607 to predict the temperature of the critical point at a future time. Operations 601, 603, 605 and 607 can be performed in a research and development phase of a product's development to produce a complete or merged thermal model for one or more data processing systems; this phase can include building prototypes of such systems and taking temperature measurements of the prototypes at points other than the critical points (and optionally also at the critical point using a mechanical thermocouple connected to the critical point since the critical point may not have an attached temperature sensor) and then generating both the transient temperature portion of the thermal model and the steady state portion of the thermal model from the measured temperatures. Once the merged thermal model is generated, it can be stored in memory in the one or more data processing systems which are manufactured for distribution to customers so that the stored representation of the merged model can be used at runtime of these manufactured systems.

[0042] The complete thermal model may be created for a class of data processing systems (e.g., a set of laptop computers having a similar housing and display but with a different microprocessor or other internal differences) and represented in one or more equations stored in a memory of such systems; a microprocessor coupled to the memory in such systems can use current temperature measurements, at points other than the critical point, to predict the current and/or future temperatures at the critical point.

[0043] **Figure 7** shows one example of a typical data processing system, such as data processing system 700, which may be used with the present invention. Note that while **Figure 7** illustrates various components of a data processing system, it is not intended to represent any particular architecture or manner of interconnecting the components as such details are not germane to the present invention. It will also be appreciated that network computers and other data processing systems which have fewer components or perhaps more components may also be used. The data processing system of **Figure 7** may, for example, be a workstation, or a personal computer (PC) running a Windows operating system, or an Apple Macintosh computer.

[0044] As shown in **Figure 7**, the data processing system 701 includes a system bus 702 which is coupled to a microprocessor 703, a ROM 707, a volatile RAM 705, and a non-volatile memory 706. The microprocessor 703, which may be a processor designed to execute any instruction set, is coupled to cache memory 704 as shown in the example of **Figure 7**. The system bus 702 interconnects these various components together and also interconnects components 703, 707, 705, and 706 to a display controller and display device 708, and to peripheral devices such as input/output (I/O) devices 710, such as keyboards, modems, network interfaces, printers, scanners, video cameras and other devices which are well known in the art. Typically, the I/O devices 710 are coupled to the system bus 702 through input/output controllers 709. The volatile RAM 705 is typically implemented as dynamic RAM (DRAM) which requires power continually in order to refresh or maintain the data in the memory. The non-volatile memory 706 is typically a magnetic hard drive or a magnetic optical drive or an optical drive or a DVD RAM or other type of memory systems which maintain data even after power is removed from the system. Typically, the non-volatile memory 706 will also be a random access memory although this is not required. While **Figure 7** shows that the non-volatile memory 706 is a local device coupled directly to the rest of the components in the data processing system, it will be appreciated that the present invention may utilize a non-volatile memory which is remote from the system, such as a network storage device which is

coupled to the data processing system through a network interface such as a modem or Ethernet interface (not shown). The system bus 702 may include one or more buses connected to each other through various bridges, controllers and/or adapters (not shown) as is well known in the art. In one embodiment the I/O controller 709 includes a USB (Universal Serial Bus) adapter for controlling USB peripherals, and/or an IEEE-1394 bus adapter for controlling IEEE-1394 peripherals.

[0045] It will be apparent from this description that aspects of the present invention may be embodied, at least in part, in software, hardware, firmware, or in combination thereof. That is, the techniques may be carried out in a computer system or other data processing system in response to its processor, such as a microprocessor, executing sequences of instructions contained in a memory, such as ROM 707, volatile RAM 705, non-volatile memory 706, cache 704, or a remote storage device (not shown). In various embodiments, hardwired circuitry may be used in combination with software instructions to implement the present invention. Thus, the techniques are not limited to any specific combination of hardware circuitry and software or to any particular source for the instructions executed by the data processing system 700. In addition, throughout this description, various functions and operations are described as being performed by or caused by software code to simplify description. However, those skilled in the art will recognize that what is meant by such expressions is that the functions result from execution of code by a processor, such as the microprocessor 703.

[0046] A machine readable medium can be used to store software and data which when executed by the data processing system 700 causes the system to perform various methods of the present invention. This executable software and data may be stored in various places including for example ROM 707, volatile RAM 705, non-volatile memory 706, and/or cache 704 as shown in **Figure 7**. Portions of this software and/or data may be stored in any one of these storage devices.

[0047] The invention also relates to apparatus for performing the operations herein. This apparatus may be specially constructed for the required purposes, or

it may comprise a general purpose computer selectively activated or reconfigured by a computer program stored in the computer. Such a computer program may be stored or transmitted in a machine-readable medium. A machine readable medium includes any mechanism that provides (i.e., stores and/or transmits) information in a form accessible by a machine (e.g., a computer, network device, personal digital assistant, manufacturing tool, any device with a set of one or more processors, etc.). For example, a machine readable medium includes recordable/non-recordable media such as, but not limited to, a machine-readable storage medium (e.g., any type of disk including floppy disks, optical disks, CD-ROMs, and magnetic-optical disks, read-only memories (ROMs), random access memories (RAMs), EPROMs, EEPROMs, flash memory, magnetic or optical cards, or any type of media suitable for storing electronic instructions), or a machine-readable transmission medium such as, but not limited to, any type of electrical, optical, acoustical or other form of propagated signals (e.g., carrier waves, infrared signals, digital signals, etc.).

[0048] Additionally, it will be understood that the various embodiments described herein may be implemented with data processing systems which have more or fewer components than system 700; for example, such data processing systems may be a cellular telephone or a personal digital assistant (PDA) or an entertainment system or a media player (e.g., an iPod) or a consumer electronic device, etc., each of which can be used to implement one or more of the embodiments of the invention.

[0049] Throughout the foregoing specification, references to “one embodiment,” “an embodiment,” “an example embodiment,” etc., indicate that the embodiment described may include a particular feature, structure, or characteristic, but every embodiment may not necessarily include the particular feature, structure, or characteristic. Moreover, such phrases are not necessarily referring to the same embodiment. When a particular feature, structure, or characteristic is described in connection with an embodiment, it is submitted that it is within the knowledge of one skilled in the art to bring about such a feature, structure, or characteristic in connection with other embodiments whether or not

explicitly described. Various changes may be made in the structure and embodiments shown herein without departing from the principles of the invention. Further, features of the embodiments shown in various figures may be employed in combination with embodiments shown in other figures.

[0050] In the description as set forth above and claims, the terms “coupled” and “connected,” along with their derivatives, may be used. It should be understood that these terms are not intended to be synonymous with each other. Rather, in particular embodiments, “connected” is used to indicate that two or more elements are in direct physical or electrical contact with each other. “Coupled” may mean that two or more elements are in direct physical or electrical contact. However, “coupled” may also mean that two or more elements are not in direct contact with each other, but yet still co-operate or interact with each other.

[0051] Some portions of the detailed description as set forth above are presented in terms of algorithms and symbolic representations of operations on data bits within a computer memory. These algorithmic descriptions and representations are the means used by those skilled in the data processing arts to most effectively convey the substance of their work to others skilled in the art. An algorithm is here, and generally, conceived to be a self-consistent sequence of operations leading to a desired result. The operations are those requiring physical manipulations of physical quantities. Usually, though not necessarily, these quantities take the form of electrical or magnetic signals capable of being stored, transferred, combined, compared, and otherwise manipulated. It has proven convenient at times, principally for reasons of common usage, to refer to these signals as bits, values, elements, symbols, characters, terms, numbers, or the like.

[0052] It should be borne in mind, however, that all of these and similar terms are to be associated with the appropriate physical quantities and are merely convenient labels applied to these quantities. Unless specifically stated otherwise as apparent from the discussion as set forth above, it is appreciated that throughout the description, discussions utilizing terms such as “processing” or “computing” or “calculating” or “determining” or “displaying” or the like, refer to the action and processes of a computer system, or similar electronic computing

device, that manipulates and transforms data represented as physical (electronic) quantities within the computer system's registers and memories into other data similarly represented as physical quantities within the computer system memories or registers or other such information storage, transmission or display devices.

[0053] The invention also relates to apparatus for performing the operations herein. This apparatus may be specially constructed for the required purposes, or it may comprise a general purpose computer selectively activated or reconfigured by a computer program stored in the computer. Such a computer program may be stored or transmitted in a machine-readable medium, such as, but is not limited to, a machine-readable storage medium (e.g., any type of disk including floppy disks, optical disks, CD-ROMs, and magnetic-optical disks, read-only memories (ROMs), random access memories (RAMs), EPROMs, EEPROMs, flash memory, magnetic or optical cards, or any type of media suitable for storing electronic instructions), or a machine-readable transmission medium such as, but not limited to, any type of electrical, optical, acoustical or other form of propagated signals (e.g., carrier waves, infrared signals, digital signals, etc.).

[0054] The algorithms and displays presented herein are not inherently related to any particular computer system or other apparatus. Various general purpose systems may be used with programs in accordance with the teachings herein, or it may prove convenient to construct more specialized apparatuses to perform the method operations. The structure for a variety of these systems appears from the description above. In addition, the invention is not described with reference to any particular programming language. It will be appreciated that a variety of programming languages may be used to implement the teachings of the invention as described herein.

[0055] Embodiments of the invention may include various operations as set forth above or fewer operations or more operations or operations in an order which is different from the order described herein. The operations may be embodied in machine-executable instructions which cause a general-purpose or special-purpose processor to perform certain operations. Alternatively, these operations may be performed by specific hardware components that contain

hardwired logic for performing the operations, or by any combination of programmed computer components and custom hardware components.

[0056] Throughout the foregoing description, for the purposes of explanation, numerous specific details were set forth in order to provide a thorough understanding of the invention. It will be apparent, however, to one skilled in the art that the invention may be practiced without some of these specific details. Accordingly, the scope and spirit of the invention should be judged in terms of the claims which follow as well as the legal equivalents thereof.

CLAIMS

What is claimed is:

1. A machine-implemented method comprising:

retrieving a representation of a thermal model for one or more points of a data processing system from a memory in the data processing system, the thermal model having a first component which estimates a steady state temperature behavior of the data processing system at the one or more points and having a second component which estimates a transient temperature behavior of the data processing system at the one or more points;

receiving temperature data from at least one temperature sensor not located at the one or more points; and

estimating a temperature at the one or more points using the thermal model.

2. A machine-readable medium storing executable instructions, which when executed, causes a system to perform a method comprising:

retrieving a representation of a thermal model for one or more points of a data processing system from a memory in the data processing system, the thermal model having a first component which estimates a steady state temperature behavior of the data processing system at the one or more points and having a second component which estimates a transient temperature behavior of the data processing system at the one or more points;

receiving temperature data from at least one temperature sensor not located at the one or more points; and

estimating a temperature at the one or more points using the thermal model.

3. A data processing apparatus comprising:

at least one temperature sensor located on the data processing apparatus at a first position;

a memory configured to store a representation of a thermal model for a location on the data processing apparatus at a second position which is different from the first position, the thermal model having a first component which estimates a steady state temperature behavior at the second position and having a second component which estimates a transient temperature behavior at the second position;

a processor coupled to the memory and to the at least one temperature sensor, the processor configured to estimate a temperature at the second position using the representation of the thermal model.

4. A method comprising:

storing a representation of a thermal model at one or more points of a data processing system in a memory, the thermal model using a filtering algorithm to filter at least one temperature sensor at a plurality of frequencies to compensate for time-based under estimation and over estimation in the thermal model;

applying power to the data processing system;

receiving temperature data from the at least one temperature sensor; and

estimating an instantaneous temperature at the one or more points using the thermal model.

5. The method of claim 4, wherein the plurality of frequencies includes a set of three frequencies including a maximum correlation frequency, a slow moving frequency, and a fast moving frequency.

6. The method of claim 4, wherein the thermal model of the temperature at the one or more points is generated by merging a transient temperature portion of the thermal model with a steady state portion of the thermal model, wherein there are no temperature sensors located at the one or more points, and wherein the filtering algorithm filters temperatures received

from the at least one temperature sensor which is not located at the one or more points.

7. The method of claim 4, further comprising:
storing a representation of a future thermal model at the one or more points of the data processing system in the memory; and
predicting a future temperature at the one or more points using the future thermal model.

8. The method of claim 7, further comprising initiating proactive measures to reduce temperature at the one or more points of the data processing system using the future thermal model, the proactive measures including one or more of the following:

thermal-throttling the data processing system,
engaging cooling fans,
managing and redistributing power, and
switching from external graphics to internal graphics.

9. A computer-readable storage medium that provides instructions, which when executed by a computer system, cause the computer system to perform a method to estimate temperature at one or more points of the computer system, the method comprising:

storing a representation of a thermal model at the one or more points of the computer system in a memory, the thermal model using a filtering algorithm to filter at least one temperature sensor at a plurality of frequencies to compensate for time-based under estimation and over estimation in the thermal model;

applying power to the computer system;
receiving temperature data from the at least one temperature sensor; and
estimating an instantaneous temperature at the one or more points using the thermal model.

10. The computer-readable storage medium of claim 9, wherein the plurality of frequencies includes a set of three frequencies including a maximum correlation frequency, a slow moving frequency, and a fast moving frequency.

11. The computer-readable storage medium of claim 9, wherein the thermal model of the temperature at the one or more points is generated by merging a transient temperature portion of the thermal model with a steady state portion of the thermal model, wherein there are no temperature sensors located at the one or more points, and wherein the filtering algorithm filters temperatures received from the at least one temperature sensor which is not located at the one or more points.

12. The computer-readable storage medium of claim 9, further comprising:

storing a representation of a future thermal model at the one or more points of the computer system in the memory; and

predicting a future temperature at the one or more points using the future thermal model.

13. The computer-readable storage medium of claim 12, further comprising initiating proactive measures to reduce temperature at the one or more points of the computer system based on the future thermal model, the proactive measures including one or more of the following:

thermal-throttling the computer system,

engaging cooling fans,

managing and redistributing power, and

switching from external graphics to internal graphics.

14. A data processing apparatus comprising:

at least one temperature sensor;

a memory to store a representation of a thermal model at one or more critical points of the data processing apparatus, the thermal model using a filtering algorithm to filter at least one temperature sensor at a plurality of frequencies to compensate for time-based under estimation and over estimation in the thermal model; and

a controller coupled with the at least one temperature sensor and the memory, the controller to perform the following:

apply power to the data processing apparatus;

receive temperature data from the at least one temperature sensor;

and

estimate an instantaneous temperature at the one or more critical points using the thermal model.

15. The data processing apparatus of claim 14, wherein the plurality of frequencies includes a set of three frequencies including a maximum correlation frequency, a slow moving frequency, and a fast moving frequency.

16. The data processing apparatus of claim 15, wherein the thermal model of the temperature at the one or more critical points is generated by merging a transient temperature portion of the thermal model with a steady state portion of the thermal model, wherein there are no temperature sensors located at the one or more critical points, and wherein the filtering algorithm filters temperatures received from the at least one temperature sensor which is not located at the one or more critical points.

17. The data processing apparatus of claim 14, wherein the memory is further to store a representation of a future thermal model at the one or more critical points of the data processing apparatus in the memory.

18. The data processing apparatus of claim 17, wherein the controller is further to predict a future temperature at the one or more critical points using the future thermal model stored in the memory.

19. The data processing apparatus of claim 18, wherein the controller is further to initiate proactive measures to reduce temperature at the one or more critical points of the data processing apparatus based on the future temperature model, the proactive measures including one or more of the following:

thermal-throttling the data processing apparatus,
engaging cooling fans,
managing and redistributing power, and
switching from external graphics to internal graphics.

20. A method to estimate temperature at one or more critical points in a data processing system comprising:

generating a transient temperature portion of a thermal model at the one or more critical points including:

filtering at least one temperature sensor at a plurality of frequencies to obtain a corresponding plurality of filtered values; and
generating a set of weighted values by applying weights to each of the plurality of filtered values; and

combining the set of weighted values to generate the transient temperature portion of the thermal model at the one or more critical points to remove at least some time-based under estimation and over estimation in the transient temperature portion of the thermal model.

21. The method of claim 20, wherein the plurality of frequencies includes a set of three frequencies including a maximum correlation frequency, a slow moving frequency, and a fast moving frequency.

22. The method of claim 20, further comprising generating a steady state temperature portion of the thermal model at the one or more critical points by performing a regression analysis of the one or more critical points at steady state temperatures.

23. The method of claim 22, further comprising choosing additional temperature sensors to perform the regression analysis including:

applying power to the data processing system;

monitoring thermal behavior of a plurality of temperature sensors over time; and

choosing a minimum number of the plurality of temperature sensors to perform the regression analysis.

24. The method of claim 23, wherein the minimum number of temperature sensors includes a pair of the temperature sensors having a maximum variance with respect to the one or more critical points.

25. The method of claim 22, further comprising generating the thermal model at the one or more critical points by merging the transient temperature portion of the thermal model with the steady state temperature portion of the thermal model.

26. The method of claim 25, further comprising estimating an instantaneous temperature at the one or more critical points based on the thermal model.

27. The method of claim 26, further comprising:

generating a future thermal model at the one or more critical points based on the thermal model;

predicting a future temperature at the one or more critical points based on the future thermal model; and

initiating proactive measures to reduce temperature at the one or more critical points.

28. The method of claim 27, wherein generating the future temperature model at the one or more critical points includes modifying the weights of the set of filtered values used to generate the transient temperature portion of the thermal model.

29. A computer-readable storage medium that provides instructions, which when executed by a computer system, cause the computer system to perform a method to estimate temperature at one or more critical points, the method comprising:

generating a transient temperature model at the one or more critical points including:

filtering at least one temperature sensor at a plurality of frequencies to obtain a corresponding plurality of filtered values; and
generating a set of weighted values by applying weights to each of the plurality of filtered values; and

combining the set of weighted values to generate the transient temperature portion of the thermal model at the one or more critical points to remove at least some time-based under estimation and over estimation in the transient temperature portion of the thermal model.

30. The computer-readable storage medium of claim 29, wherein the plurality of frequencies is a set of three frequencies including a maximum correlation frequency, a slow moving frequency, and a fast moving frequency.

31. The computer-readable storage medium of claim 29, further comprising generating a steady state temperature portion of the thermal model at the one or more critical points by performing a regression analysis of the one or more critical points at steady state temperatures.

32. The computer-readable storage medium of claim 31, further comprising choosing additional temperature sensors to perform the regression analysis including:

applying power to the computer system;

monitoring thermal behavior of a plurality of temperature sensors over time; and

choosing a minimum number of the plurality of temperature sensors to perform the regression analysis.

33. The computer-readable storage medium of claim 31, further comprising generating the thermal model at the one or more critical points by merging the transient temperature portion of the thermal model with the steady state temperature portion of the thermal model.

34. The computer-readable storage medium of claim 33, further comprising estimating an instantaneous temperature at the one or more critical points based on the thermal model.

35. The computer-readable storage medium of claim 33, further comprising:

generating a future thermal model at the one or more critical points based on the thermal model; and

predicting a future temperature at the one or more critical points based on the future thermal model.

36. The computer-readable storage medium of claim 35, wherein generating the future temperature model at the one or more critical points includes modifying the weights of the plurality of filtered values used to generate the transient temperature portion of the thermal model.

37. A method to estimate temperature at one or more critical points in a data processing system comprising:

generating a steady state temperature portion of a thermal model at one or more critical points of the data processing system using regression analysis;

generating a transient temperature portion of the thermal model at the one or more critical points of the data processing system using a filtering algorithm;

generating the thermal model at the one or more critical points of the data processing system by combining the steady state temperature portion of the thermal model with the transient temperature portion of the thermal model; and

estimating an instantaneous temperature at the one or more critical points based on the thermal model.

38. The method of claim 37, wherein the filtering algorithm comprises:

filtering at least one temperature sensor at a plurality of frequencies to obtain a corresponding plurality of filtered values;

generating a set of weighted values based on applying weights to each of the plurality of filtered values; and

combining the set of weighted values to generate the transient temperature portion of the thermal model at the one or more critical points to remove at least some time-based under estimation and over estimation in the transient temperature portion of the thermal model.

39. A computer-readable storage medium that provides instructions, which when executed by a computer system, cause the computer system to perform a method to estimate temperature at one or more critical points, the method comprising:

generating a steady state temperature portion of a thermal model at one or more critical points of the data processing system using regression analysis;

generating a transient temperature portion of the thermal model at the one or more critical points of the data processing system using a filtering algorithm;

generating the thermal model at the one or more critical points of the data processing system by combining the steady state temperature portion of the thermal model with the transient temperature portion of the thermal model; and
estimating an instantaneous temperature at the one or more critical points based on the thermal model.

40. The method of claim 39, wherein the filtering algorithm comprises:

filtering at least one temperature sensor at a plurality of frequencies to obtain a corresponding plurality of filtered values;

generating a set of weighted values based on applying weights to each of the plurality of filtered values; and

combining the set of weighted values to generate the transient temperature portion of the thermal model at the one or more critical points to remove at least some time-based under estimation and over estimation in the transient temperature portion of the thermal model.




Fig. 1
Prior Art

Autocorrelation for CPU Die Temp

Autocorrelation for Skin Temp





Fig. 2




Fig. 3A




Fig. 3B




Fig. 4




Fig. 5A




Fig. 5B




Fig. 5C


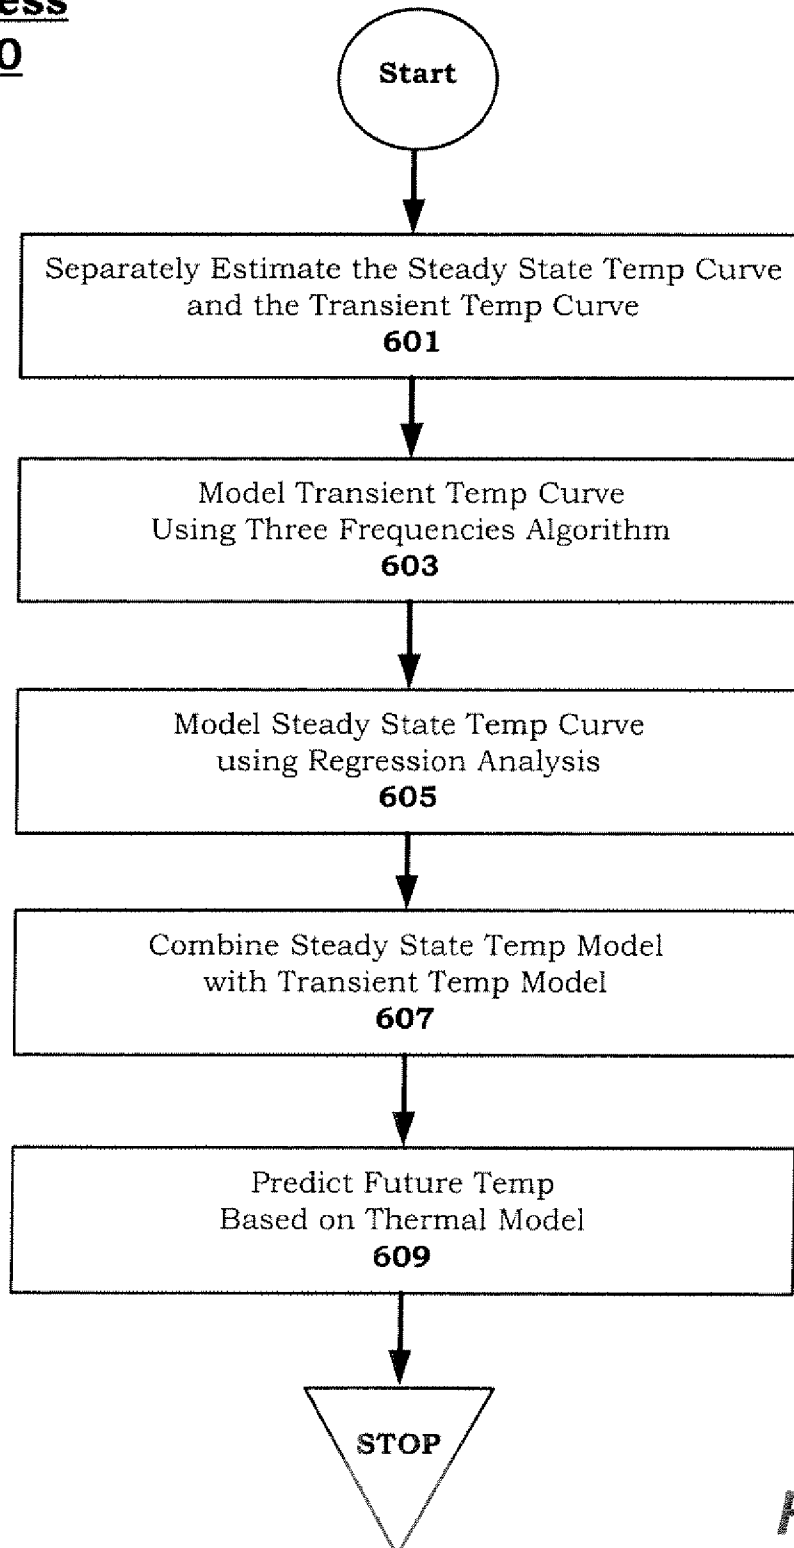


Fig. 6A

10/11

Process
600*Fig. 6B*

Data
Processing
System
700

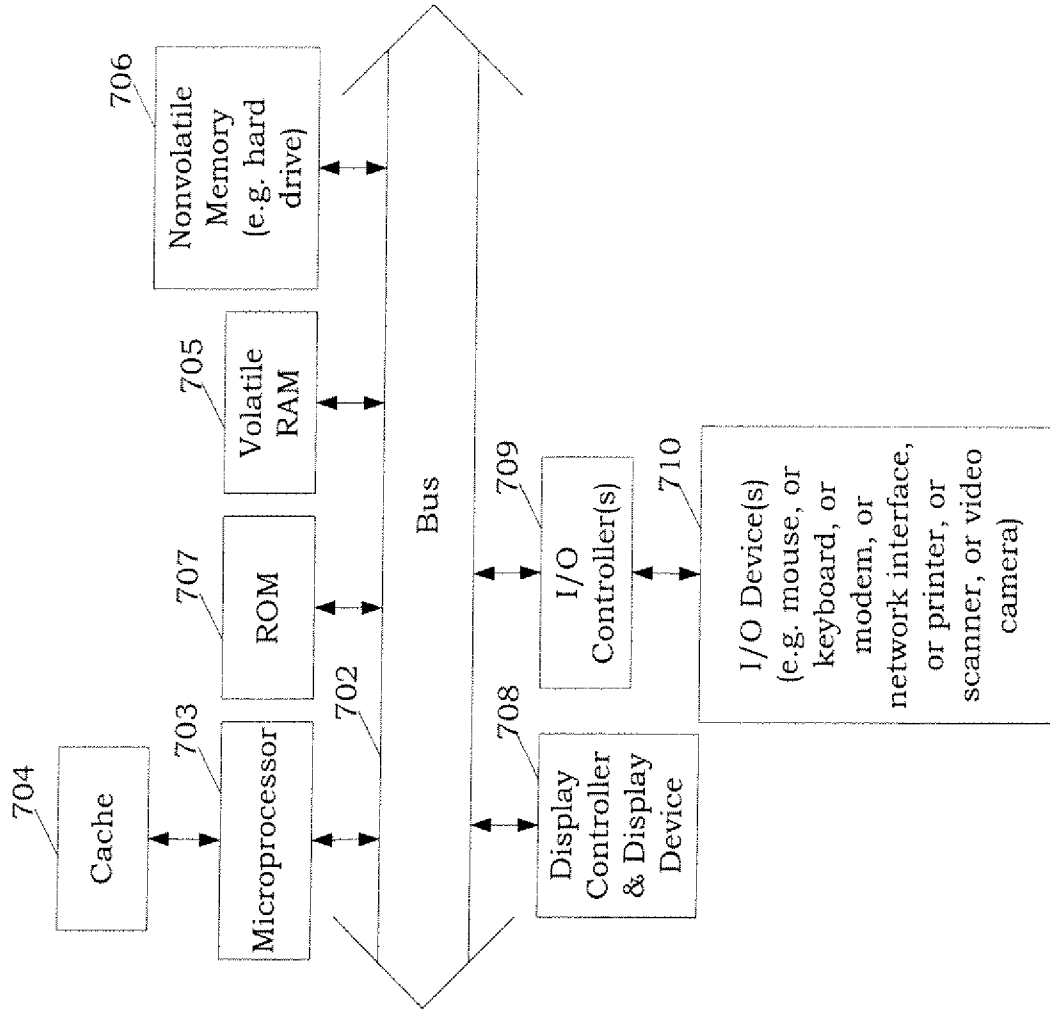


Fig. 7