

(19) World Intellectual Property Organization  
International Bureau



(43) International Publication Date  
6 July 2006 (06.07.2006)

PCT

(10) International Publication Number  
**WO 2006/072106 A2**

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

(21) International Application Number:  
PCT/US2005/047688

(22) International Filing Date:  
29 December 2005 (29.12.2005)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:  
11/025,140 30 December 2004 (30.12.2004) US

(71) Applicant (for all designated States except US): **INTEL CORPORATION** [US/US]; 2200 Mission College Boulevard, Santa Clara, California 95052 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **ROTEM, Efraim** [IL/IL]; 8 Vizo Street, 34400 Haifa (IL). **HERMERDING, Jim** [US/US]; 3385 Cavendish Drive, San Jose, California 95132 (US). **DISTEFANO, Eric** [US/US]; 1535 Frankfurt Way, Livermore, California 94550 (US). **COOPER, Barnes** [US/US]; 13425 SW Nahcotta Drive, Tigard, Oregon 97223 (US).

(74) Agents: **O'DOWD, Shawn W.** et al.; KENYON & KENYON, 1500 K Street, Suite 700, Washington, D.C. 20005 (US).

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

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: DEVICE AND METHOD FOR ON-DIE TEMPERATURE MEASUREMENT

(57) Abstract: A system for measuring and managing thermal operations of a processor core on a semiconductor die using a sensor positioned in a hotspot of the processor core. A measured temperature reading is determined based upon a temperature sensed by the sensor. Interrupt signals and a software readable register indicating temperature information provide feedback about the thermal environment to the processor. Based upon the measured temperature reading, the interrupt signals direct the processor to modify operation.

WO 2006/072106 A2

## DEVICE AND METHOD FOR ON-DIE TEMPERATURE MEASUREMENT

### Technical Field

[0001] Embodiments of the present invention relate to devices and methods for  
5 temperature measurement of processor cores, and more particularly to on-die circuits and  
methods for temperature measurement and reporting.

### Background

[0002] Modern central processing units (“CPUs”), with increasing processor core  
frequencies and power densities, are rapidly reaching a point in which the CPU  
10 performance is limited by the amount of heat that can be extracted from the CPU by  
cooling technology.

[0003] High performance CPUs sometimes integrate an on-die thermal diode that  
converts junction temperatures to some electric value. External analog-to-digital (“A/D”)  
devices can be found on several computer platforms, converting the electric value into  
15 usable information for temperature control.

[0004] Due to routing and accuracy limitations, the thermal diode is placed in locations  
which are not at a hotspot of the CPU. A hotspot is a localized region having a higher  
temperature than regions of the processor adjoining the hotspot. Temperature differences,  
as high as 15°C have been observed between the diode location and the hotspot. The  
20 result is inaccurate temperature readings and discrepancy between the reported value and  
the actual temperature used for device specification and control. Inaccurate temperature  
reading leads to inclusion of higher design margins, increased cooling costs and limited  
functionality.

[0005] Another thermal protection mechanism uses a separate sensor to detect a  
25 maximum threshold temperature at the hottest hotspot on the die. Once the threshold is  
reached, a single trigger occurs and thermal protection is initiated. Other than this single  
trigger, the sensor provides no other feedback for use in thermal or acoustic management.

[0006] If this second sensor is external to the device, the external sensor suffers from  
time delay and accuracy problems resulting from the physical separation of the sensor  
30 from the hotspot.

[0007] There is a need for improved on-die temperature monitoring to improve  
performance in a thermally constrained environment. There is also a need for faster and  
more accurate reporting, eliminating gaps between thermal measurement and control.

**Brief Description of the Drawings**

[0008] FIG. 1 is a schematic diagram of a first embodiment of the invention.

[0009] FIG. 2 is a schematic diagram of a second embodiment of the invention.

[0010] FIG. 3 is a schematic diagram of a third embodiment of the invention.

5 [0011] FIG. 4 is a schematic diagram of a hybrid of the first and third embodiments.

[0012] FIG. 5 is a schematic diagram of a hybrid of the first and second embodiments.

[0013] FIG. 6 illustrates the operation of an embodiment of a thermal event generator.

[0014] FIG. 7 illustrates an alternative measurement circuit which can be interchanged with the measurement circuit of the second and third embodiments

10 [0015] FIG. 8 illustrates a supplemental circuit for the first embodiment of the invention when plural hotspots are monitored, to detect when the hot spots exceed the threshold temperature.

[0016] FIGS. 9 and 10 illustrate alternative measurement circuits which can be interchanged with the measurement circuits of the second and third embodiments to

15 measure temperature for multiple hotspots.

[0017] FIG. 10 illustrates an alternative measurement circuit which can be interchanged with the measurement circuit of the second and third embodiments to measure temperature for multiple hotspots.

20 [0018] FIG. 11 illustrates a fourth embodiment of the invention, providing centralized thermal management for plural processor cores.

[0019] FIG. 12 illustrates an example of the Central Management System that appears in FIG. 11.

[0020] FIG. 13 illustrates a thermal management system including both on-die and off-die thermal management.

25

**Detailed Description**

[0021] Embodiments of the present invention integrate a sensor at a hotspot of the die for the measurement and reporting of temperature values, and for the generation of notifications. Using a single sensor to provide both temperature-responsive triggers and temperature measurement improves the speed, consistency, and accuracy of temperature reporting. Processor operation can be optimized, maximizing performance in a thermally constrained environment.

30

[0022] The temperature responsive triggers may be utilized to direct both on-die and off-die temperature control, including directing the processor to modify operation, activating external fans or cooling systems, modifying the output of an external power supply, and/or putting a portion or an entirety of the processor into a standby state. Examples of  
5 modifying processor operation include modifying clock speeds of the processor and modifying operating voltages of the processor. Reporting of a measured temperature may be utilized, for among other things, to provide feedback to the processor.

[0023] For processors utilizing software control mechanisms triggered by management systems such as Advanced Configuration and Power Interface (“ACPI”; Revision 3.0,  
10 September 2004; Intel Corp., et al.) or System Management Mode (“SMM”), direct on-die temperature reporting is particularly advantageous, since the ability to report temperatures through the processor to native software improves response time, and provides system designers and manufacturer increased flexibility for adapting processor operation. Flexibility can be further increased by using software-programmable temperature-  
15 responsive triggers, which may, for example, be programmed by software and/or firmware instructions executed on the processor (e.g., by third-party software, the operating system, and/or the Basic Input/Output System (“BIOS”)) or be programmed by fuse or other firmware on the die (e.g., programmed by a manufacturer or distributor).

[0024] Prior to fabrication, at least one hotspot of the processor is identified by either  
20 modeling, empirical data, analysis of the package, and/or by other conventional techniques. A temperature sensor is fabricated in the hotspot, modifying the layout of the circuit to provide for the sensor, if necessary. The sensor may be any temperature sensitive device, such as a p-n junction device (e.g., a diode or transistor) or a temperature sensitive resistor. Advantages of P-N junction devices, besides being easy to integrate  
25 onto the processor, are a very quick response to temperature changes and that P-N junction devices do not need a reference temperature (i.e., in comparison to thermocouples). In the embodiments herein, the sensors shown are diodes. However, it should be understood that diodes are only used as an example and that other sensor-types may be also be used.

[0025] A first embodiment of the invention is illustrated in FIG. 1. A processor 110 and  
30 temperature control circuitry are integrated on a semiconductor die 100. A diode 130 is located within hotspot 120. The voltage drop across the diode 130 changes with the temperature of the hotspot 120.

[0026] The voltage across diode 130 is provided to a thermal trigger circuit 140 and a measurement circuit 160. In the thermal trigger circuit 140, the diode voltage is compared with a critical point voltage 142 by analog comparator 144. The critical point voltage 142 corresponds to a threshold temperature, such as a maximum thermal operating limit for the processor. By utilizing a simple circuit such as an analog comparator to determine temperature, threshold reporting is near-instantaneous.

[0027] Critical point 142 can be fixed or adjustable. A way to provide critical point 142 is by using a digital-to-analog converter to read a register holding a temperature value (not illustrated). The output of the digital-to-analog converter is then provided as input to comparator 144. By programming the register, the threshold can be adjusted.

[0028] The thermal trigger signal, output by comparator 144, is provided to thermal control circuits 112 and/or 113. The thermal trigger signal may be provided as an output 101 from the die 100. When the temperature signal (i.e., voltage) from the diode 130 equals or exceeds the critical point 142, the thermal control circuits 112/113 initiate on-die thermal protection, such as a reduction of operating voltage or a reduction of operating frequency. Additionally, the thermal trigger signal may be output from the die to control external thermal control apparatus.

[0029] Measurement circuit 150 also receives the temperature signal from the diode 130. An analog-to-digital converter 160 of measurement circuit 150 converts the diode voltage into an M-bit digital value. This M-bit digital value serves as a raw measured temperature reading. The processor 110 may include a software readable register 111, to receive the measured temperature reading. An example of the software readable register 111 is model-specific register (MSR) address space. Alternatively, a code converter 195 can be included to include the raw measured temperature reading into a formatted temperature, which is then provided to the software readable register 111. Use of the code converter 195 allows for compensation of any non-linearity in the raw value, and allows the measured temperature value to be scaled (e.g., Celsius, Fahrenheit). A simple implementation for code converter 195 is a non-volatile memory look-up table.

[0030] The measured temperature reading from measurement circuit 150 is also provided to a thermal event detection circuit 180, which directs the processor to modify processor-operation based upon the measured temperature reading. As shown in FIG. 6, the thermal event detection circuit 180 comprises a plurality of programmable trip points (683, 685, 687, and 689), each of which are compared (682, 684, 686, and 688) to

measured temperature reading. The results of the comparisons are provided to the processor 110 by means of interrupts, and can also be used for other on-die and off-die thermal protection control. One way to generate interrupts is through a programmable interrupt controller (“PIC”) such as an advanced programmable interrupt controller (“APIC”). The PIC can be included in the thermal control circuits 112/113, or can be a separate circuit. One or more of the results of the comparisons by events circuit 180 may be provided as an output 102 from the die 100.

[0031] As can be understood from the time delays inherent in the circuitry involved, even if the critical point 142 and a trip point (683/685/687/689) are set to a same threshold value, the thermal trigger 140 will respond more quickly than event circuit 180.

Accordingly, the critical point 140 may be set for a maximum temperature threshold value the processor can tolerate, since the speed of response can impact processor survivability.

[0032] Thermal event detection circuit 180 also may include a circuit 690 directing the processor according to trends of the measured temperature reading over time, and a circuit 692 directing the processor based on an average temperature over time. An example of circuit 690 is at least one memory element to register temperature at a previous time or times and a comparison circuit to compare this previous temperature or temperatures to the current temperature. An example of circuit 392 is at least one memory element to accumulate an average result and at least one computation element to modify the average result in accordance to current temperature, such as a differentiator sampling the value every period N and calculating  $T_{rand}=T(n)-T(n-1)$ .

[0033] FIG. 2 illustrates a second embodiment of the invention. Semiconductor die 200 again includes processor 110, code converter 195, and events circuit 180, as in the first embodiment; and one or more of the comparison results by events circuit 180 may be provided as an output 202 from the die 200. However, the thermal trigger circuit 240 and measurement circuit 250 are different.

[0034] Measurement circuit 250 utilizes digital feedback to determine the temperature of the hotspot. Number generator 256 generates a number corresponding to an interim temperature value within a range of temperatures. The interim temperature value is converted into an analog voltage by digital-to-analog converter 252, and comparator 254 compares the analog value with the diode voltage. If the hotspot temperature is approximately equal to the interim temperature value, the interim temperature value is stored in latching register 258, and the number generator 256 can either be reset, or

number generation can continue, depending upon the particular technique used. Otherwise, another interim temperature is generated.

[0035] This can be better understood by considering particular implementations for number generation, such as linear searches and successive approximation (e.g., a binary search). The simplest example is a linear search using an up-down counter.

[0036] For example, consider a counter sweeping the temperature range, with the hotspot temperature being fixed at a temperature within the range. If the hotspot temperature lies between two interim temperatures produced by the number generator 256, the results are not exact. On an up-sweep of the temperature range, the comparator output 254 will undergo a transition when the hotspot temperature equals or exceeds the interim temperature values. Thus, the register 258 will latch a temperature value that is slightly higher than the actual temperature. Additionally, if a down-sweep is performed for the same diode temperature, the result is different than for the up-sweep (off by one). The accuracy is therefore directly dependent upon the interim temperature-granularity provided by number generator, such that the measured temperature is approximate.

[0037] Using the output of comparator 254 to latch register 258 is appropriate with search techniques which will result in comparator transitions at approximately the hotspot temperature. However, if a non-linear search is used, such as a binary search, latching the register in this manner will result in errors, since a binary search can under-shoot and over-shoot the actual temperature during the search. In this case, relying on the output of comparator 254 as the trigger of latching register 258 results in inaccurate readings. Instead, as shown in FIG. 7, measurement circuit 350 includes a number generator 356 that itself latches the latching register 258.

[0038] Any number of circuits can be used to perform a binary search. For example, a binary search can be performed utilizing two registers, a subtractor, and some basic logic. One register stores a high temperature value, whereas the other register stores a low temperature value. An intermediate value between high and low is determined by subtracting the low temperature value from high temperature value in the subtractor and dividing by 2 (e.g., dropping the lowest order bit subtraction result). This intermediate value is output as the interim temperature value. If the comparator 254 indicates that the hotspot is hotter than the interim temperature value, then low temperature value is set to the intermediate temperature and the process is repeated. Likewise, if the comparator 254 indicates that the hotspot is cooler than the interim temperature value, then high

temperature value is set to the intermediate temperature and the process is repeated. The process ends with a triggering of the latching register when the low register and high register lock (i.e., are equal or approximately equal), such that the search is over. Other types of searches can be performed in a similar manner, and lookup table or an arithmetic logic unit (“ALU”) can also be used to determine the intermediate value, instead of the binary subtractor.

[0039] Measurement circuits 250 and 350 may also signal the processor and/or thermal control circuitry if the temperature sensed by diode 130 is outside the range of temperatures searched. Examples of how this can be accomplished include: independently comparing the temperature to the high and low values of the range; for a linear search, providing logic utilizing an overflow flag signal of the number generator (e.g., counter) in combination with the comparator 254 output if the full range is swept but no result found; and for a non-linear search, providing logic to signal if an intermediate search temperature value reaches the high or low value, in combination with the comparator 254 output. In such a case, the measured temperature reading is set to be a limit of the range of temperatures closest to the temperature sensed by the diode.

[0040] Returning to FIG. 2, another feature of the second embodiment is thermal trigger circuit 240. Digital comparator 244 compares a critical point 242 with the interim temperature value output by the number generator 256/356. Latch 248 latches the output of comparator 244 as the thermal trigger signal. Latch 248 is controlled by the same latching signal used for register 258. The critical point 242 may be programmable. The thermal trigger signal may also be provided as an output 201 from the die 200.

[0041] FIG. 3 illustrates a third embodiment of the invention. Semiconductor die 300 includes processor 110, measurement circuit 250 (or 350), and code converter 195, as described above. In this embodiment, the thermal trigger is incorporated into thermal events circuit 380. This is accomplished by utilizing one of the N-trip points (683, 685, 687, 689) as the critical point. The thermal trigger and/or one or more of the results of the comparisons by events circuit 380 may be provided as output 303 from the die 300.

[0042] FIG. 4 illustrates a hybrid of the first and third embodiments. Semiconductor die 400 includes processor 110, measurement circuit 150, code converter 195, and events circuit 380, as described above. The thermal trigger and/or one or more of the results of the comparisons by events circuit 380 may be provided as output 403 from the die 400.



[0043] A disadvantage of generating the thermal trigger using events circuit 380 or thermal trigger 240, instead of thermal trigger circuit 140, is a slight delay in response time. An advantage, however, is a reduction analog circuits, which are more sensitive to temperature and transients than digital circuits, usually operate at higher voltages, and generally require more die space. Other advantages are the relative ease with which digital logic can be programmed by software (e.g., without requiring a digital-to-analog converter (“DAC”) to program critical point 142 in thermal trigger circuit 140), and that software-programmable logic can be used instead of hard-wired circuits (thereby increasing die design flexibility).

10 [0044] FIG. 5 illustrates a hybrid of the first and second embodiments. Semiconductor die 500 includes processor 110, measurement circuit 250 (or 350), code converter 195, and events circuit 380, as described above. The thermal trigger may be provided as an output 501 from the die 500. One or more of the comparison results by events circuit 180 may be provided as an output 502 from the die 500. Here, the advantage of thermal trigger 140 is the critical point response time.

[0045] FIG. 8 illustrates a thermal trigger 440 as a variation on the first embodiment of the invention when plural hotspots are monitored. Thermal trigger 440 provides a unified thermal event signal, which can simplify thermal management. If the comparators 144/144' transition from high-to-low when the hotspot temperature exceeds critical point 142, AND-gate 359 transitions to low, thereby signaling that one of the hotspots has exceeded the critical point. If the comparators 144/144' transition from low-to-high when the hotspot exceeds critical point 142, an OR-gate can be used instead of AND-gate 359 to accomplish the same function (not shown). Likewise, the same circuits can be used to determine when all of the hotspots have exceeded the critical point by interchanging AND-for-OR and OR-for-AND in the above description.

25 [0046] FIGS. 9 and 10 illustrate supplemental measurement circuits which can be interchanged with the measurement circuit from the second and third embodiments to measure a temperature of multiple hotspots. Depending upon whether comparators 254 and 254' are configured to switch from low-to-high or high-to-low when the interim temperature value exceeds the respective hotspot temperature, measurement circuits 450 and 550 either determine the highest or lowest temperature among all of the diodes. As with the embodiment in FIG. 8, an OR-gate can be used instead of an AND-gate, depending upon the result desired. Referring to FIGS. 9 and 10, the number of hotspots

monitored can be increased with minimal increase in the size and complexity of the measurement circuit. In addition, central reporting of temperature lends itself to centralized thermal management, which can be performed, for example, by consolidating on-die off-processor thermal control in a single circuit (e.g., one thermal control circuit  
5 113 for plural processors). Further, if multiple hotspots and diodes are included on one processor, this configuration simplifies reporting to the one processor.

[0047] FIG. 11 illustrates a semiconductor die 1110 comprising a plurality of processors, each processor being monitored by a thermal trigger circuit, a measurement circuit, and an events circuit, as described in the other embodiments. Additionally, central management  
10 system 710 comprises a routing matrix that accepts temperature indications and thermal event indications and routes them to the individual processor cores based on a decision table stored in routing matrix. The decision table can be held, for example, in a volatile memory or non-volatile memory, and either be fixed or loaded (e.g., by the BIOS). Central management system 710 can also generate interrupts to either one or both cores,  
15 and/or can generate interrupts through a programmable interrupt controller ("PIC") or controllers. A PIC can be provided for each processor, or a centralized PIC can be provided to handle interrupts for multiple processors. One type of PIC is an advanced programmable interrupt controller ("APIC"; e.g., Intel® 82093AA I/O Advanced Programmable Interrupt Controller). A PIC can be provided in the processor-internal or  
20 external thermal control circuits, can be included as a separate circuit, or can be integrated into the central management system 710.

[0048] As an example of operation, a thermal event can initiate a change of the processor frequency if the temperature is high. A policy of the central management system 710 could be to reduce both core's frequency if either core is high. Another, more complex,  
25 policy would be to first reduce one core frequency only. If after some time temperature does not go down, the other core frequency and operating voltage would also be reduced. Implementation of this second policy can be accomplished by having a timer feedback to the routing matrix. When the decision table indicates to reduce operations in one processor but not the other, the timer begins counting. When the timer expires, a signal is  
30 fed back to the input of the routing matrix, resulting in a reduced operation for the second processor. The delay time can either be fixed, pre-stored in a register, or output by the routing matrix.

[0049] Logic circuit 715 compares the measured temperature readings from the plurality of measurement circuits, and provides central management system 710 with a highest of the measured temperature readings as a maximum measured temperature value of the die.

[0050] In addition, the central management system 710 can provide central reporting capability. Each core normally can access only its own register set. Central management system 710 can report to each core the package state such as the maximum measured temperature, status bits, etc. Additionally, lookup table entries within central management system 710 can provide some or all of the values used as programmed thresholds for both critical point and thermal events, adapting system thresholds during operation. Central management system 710 can also provide control signals at output 1104 for off-die thermal management.

[0051] While the above embodiments utilize a diode for sensing a temperature within a hotspot, any p-n junction device that is sensitive to temperature can be used.

[0052] In the embodiments, interrupts used to direct the processor or processors include hardware interrupts, software interrupts such as operating system interrupts, and combinations thereof. Among other things, processor interrupts direct to the operating system may be used to generate native instructions.

[0053] FIG. 12 illustrates an example of Central Management System 710. One-to-N timers 802, 812 provide feedback to routing matrix 800. The routing matrix 800 provides each timer a reset signal, a delay time, and a start signal. By including plural timers, plural possibilities can be simultaneously pending.

[0054] FIG. 13 illustrates a thermal management system including both on-die and off-die thermal management. Outputs 101, 102, 201, 202, 303, 403, 501, 502, and/or 1104 from the dies of the various embodiments, are used to control off-die thermal management 1310 elements, including fans, cooling systems, and power supplies. Control of such elements may include activating external fans or cooling systems, or modifying the output of the external power supply.

[0055] The p-n junctions are biased using a band-gap based device to maintain consistent operation as package temperatures change. An example of a band-gap current source for biasing the diodes is a band-gap comparator.

[0056] As can be understood from the above description, components of the various embodiments can be interchanged as a matter of design. It is contemplated that numerous

modifications may be made to the present invention without departing from the spirit and scope of the invention as defined in the following claims.

**What is claimed is:**

1. A system comprising:  
a processor having a hotspot during operation, the hotspot being a localized region  
5 having a higher temperature than regions of the processor adjoining the hotspot;  
means for sensing a temperature within the hotspot;  
means for generating a measured temperature reading based upon the temperature  
sensed by said means for sensing;  
means for directing the processor to modify processor-operation based upon the  
10 measured temperature reading;  
a semiconductor die comprising said processor, said means for sensing, said means  
for generating, and said means for directing; and  
means for thermal management, external to the semiconductor die and responsive  
to output from the means for directing the processor.  
15
2. The system of claim 1, the semiconductor die further comprising means for  
determining whether the temperature sensed by said means for sensing is equal to or above  
a threshold value, said means for determining comprising:  
means for comparing the threshold value with the temperature sensed by said  
20 means for sensing.
3. The system of claim 1, said means for generating a measured temperature  
reading comprising:  
means for comparing an interim temperature value with the temperature sensed by  
25 said means for sensing; and  
means for performing a search of a range of temperatures to provide the interim  
temperature value to said means for comparing.
4. The system of claim 3, said semiconductor die further comprising means  
30 for indicating that the temperature sensed by said means for sensing is equal to or above a  
threshold value, when the comparison by said means for comparing during the search of  
the range of temperatures indicates that the temperature sensed by said means for sensing  
is equal to or above the threshold value.

5. The system of claim 1, wherein the means for directing the processor to modify processor-operation directs the processor according to a magnitude of the measured temperature reading relative to a plurality of threshold values.

5

6. The system of claim 5, wherein the plurality of threshold values are accessible to, and modifiable by, instructions executed on the processor.

7. The system of claim 1, wherein the means for directing the processor to modify processor-operation directs the processor according to trends of the measured temperature reading over time.

10

8. The system of claim 1, wherein the means for directing the processor to modify processor-operation directs the processor based on an average temperature over time.

15

9. The system according to claim 1, wherein the die further comprises a register readable to instructions executed on said processor to store the measured temperature reading, or store a temperature value based on the measured temperature reading.

20

10. A device comprising:

a processor having a hotspot during operation, the hotspot being a localized region having a higher temperature than regions of the processor adjoining the hotspot;

25

a p-n junction within the hotspot, a junction voltage of the p-n junction to change as a function of the temperature of the hotspot of the processor;

a comparator to compare a voltage from of the p-n junction with a voltage representing a first threshold temperature value;

30

an analog-to-digital converter to convert the voltage from the p-n junction into a measured temperature reading;

a logic circuit to monitor the measured temperature reading to determine occurrence of particular thermal events to direct the processor to modify processor-operation when a thermal event occurs; and

a semiconductor die comprising the processor, the p-n junction, the comparator, the analog-to-digital converter, and the logic circuit.

11. The device according to claim 10, wherein the logic circuit comprises at least one digital comparator to compare the measured temperature reading with at least a second threshold temperature value.

12. The device according to claim 11, wherein said at least a second threshold temperature value is accessible to, and modifiable by, instructions executed the processor.

13. The device according to claim 10, wherein the logic circuit comprises a trend circuit to calculates trends of the measured temperature reading over a period of time, the trend circuit comprising:

at least one memory element to register temperature at a previous time and a comparison circuit to compare this previous temperature to the current temperature.

14. The device according to claim 10, wherein the logic circuit comprises an averaging circuit to calculate an average of the measured reading over a period of time, the averaging circuit comprising:

at least one memory element to accumulate an average result and at least one computation element to modify the average result in accordance to current temperature.

15. The device according to claim 10, wherein an output from the logic circuit indicating occurrence of a particular thermal event is provided as an output from the semiconductor die.

16. The device according to claim 10, the semiconductor die further comprising a register readable to instructions executed on said processor to store the measured temperature reading, or store a temperature value based on the measured temperature reading.

17. The device according to claim 16, wherein the register readable to instructions is model-specific register (MSR) address space of the processor.

18. The device according to claim 16, the semiconductor die further comprising a memory containing a look-up table, the measured temperature reading to be input into the memory and a scaled temperature reading to be output by the memory, the look-up  
5 table programmed to compensate for non-linearity in the measured temperature reading, wherein the scaled temperature reading to be stored in the register readable to instructions.

19. The device according to claim 10, the die further comprising:  
10 another processor having another hotspot, with another p-n junction, comparator, analog-to-digital converter, and logic circuit; and  
a routing matrix comprising memory storing a decision table, the routing matrix to receive the thermal event indications from the logic circuits, and receive outputs from the comparators, to route thermal event directives to the processors in accordance with the  
15 decision table.

20. The device according to claim 10, the die further comprising:  
another hotspot, with another p-n junction and another comparator, the another comparator to compare a voltage from the another p-n junction with the voltage  
20 representing the first threshold value; and  
a logic gate to receive outputs from the comparators and to signal if at least one of the hotspots exceeds the first threshold value.

21. A device comprising:  
25 a processor having a hotspot during operation, the hotspot being a localized region having a higher temperature than regions of the processor adjoining the hotspot;  
a p-n junction within the hotspot, a junction voltage of the p-n junction to change as a function of the temperature of the hotspot of the processor;  
a control circuit comprising:  
30 a number generator to generate a binary number from a plurality of binary numbers corresponding to temperature values within a temperature range;  
a digital-to-analog converter to convert the binary number of the plurality of binary numbers into a voltage;



a first comparator to compare a voltage from the p-n junction with the voltage from the digital-to-analog converter; and

a register for holding a measured temperature reading, the register to be set to the binary number as the measured temperature reading if the voltage from the p-n junction approximately equals the voltage from the digital-to-analog converter,

a logic circuit to monitor the register to determine occurrence of particular of thermal events to direct the processor to modify processor-operation when a thermal event occurs; and

a semiconductor die comprising the processor, the p-n junction, the control circuit, and the logic circuit.

22. The device of claim 21, the semiconductor die further comprising:

a digital comparator to compare the binary number generated by the number generator with a threshold temperature value; and

a latch for holding the output of the digital comparator, set when the register of the control circuit is set.

23. The device according to claim 21, the semiconductor die further comprising

a second comparator to compare the junction voltage of the p-n junction with a voltage representing a threshold temperature value.

24. The device according to claim 23, the die further comprising:

another hotspot, with another p-n junction and a third comparator, the third comparator to compare a junction voltage of the another p-n junction with the voltage representing the threshold temperature value; and

a logic gate to receive outputs from the second and third comparators and to signal if at least one of the hotspots exceeds the threshold temperature value.

25. The device according to claim 21, wherein the logic circuit comprises at

least one digital comparator to compare the measured temperature reading with a threshold temperature value.

26. The device according to claim 25, wherein the threshold temperature value is accessible to, and modifiable by, the processor.

27. The device according to claim 21, wherein the logic circuit comprises a trend circuit to calculates trends of the measured temperature reading over a period of time, the trend circuit comprising:

at least one memory element to register temperature at a previous time and a comparison circuit to compare this previous temperature to the current temperature.

28. The device according to claim 21, wherein the logic circuit comprises an averaging circuit to calculate an average of the measured reading over a period of time, the averaging circuit comprising:

at least one memory element to accumulate an average result and at least one computation element to modify the average result in accordance to current temperature.

15

29. The device according to claim 21, wherein an output from the logic circuit indicating occurrence of a particular thermal event is provided as an output from the semiconductor die.

30. The device according to claim 21, the semiconductor die further comprising a register readable to instructions executed on said processor to store the measured temperature reading, or store a temperature value based on the measured temperature reading.

31. The device according to claim 30, wherein the register readable to instructions is model-specific register (MSR) address space of the processor.

32. The device according to claim 30, the semiconductor die further comprising a memory containing a look-up table, the measured temperature reading to be input into the memory and a scaled temperature reading to be output by the memory, the look-up table programmed to compensate for non-linearity in the measured temperature reading,

wherein the scaled temperature reading to be stored in the register readable to instructions.

33. The device according to claim 21, the die further comprising:  
another processor having another hotspot, with another p-n junction, control  
circuit, and logic circuit; and  
5 a routing matrix comprising memory storing a decision table, the routing matrix to  
receive the thermal event indications from the logic circuits, to route thermal event  
directives to the processors in accordance with the decision table.

34. A method comprising:  
10 sensing a temperature of a hotspot of a processor on a semiconductor die, the  
hotspot being a localized region having a higher temperature than regions of the processor  
adjoining the hotspot;  
determining whether the sensed temperature equals or exceeds a threshold  
temperature;  
15 generating a temperature reading in accordance with the sensed temperature; and  
controlling internal operations of the processor to regulate temperature based upon  
the temperature reading,  
each of said sensing, said determining, said generating, and said controlling being  
performed by structures internal to the semiconductor die.

20 35. The method of claim 34, wherein controlling the operation of the processor  
comprises generating interrupts according to a magnitude of the temperature reading  
relative to a plurality of threshold values.

25 36. The method of claim 35, wherein the interrupts are generated through a  
programmable interrupt controller included on the semiconductor die.

37. The method of claim 35, wherein said interrupts are selected from the  
group consisting of hardware interrupts, software interrupts, operating system interrupts  
30 and a combination thereof.

38. The method of claim 35, further comprising the processor modifying at  
least one of the plurality of threshold values.

39. The method of claim 34, wherein controlling the operation of the processor comprises generating an interrupt according to trends of the temperature reading over time.

5

40. The method of claim 34, wherein controlling the operation of the processor comprises:

calculating an average temperature over time; and  
generating an interrupt based upon the average temperature.

10

41. The method of claim 34, wherein controlling the operation of the processor comprises varying operating voltage in accordance with the temperature reading.

42. The method of claim 34, wherein controlling the operation of the processor comprises varying operating frequency in accordance with the temperature reading.

15

43. The method of claim 34, further comprising instructions executed on the processor reading the temperature reading, or reading a temperature value based on the temperature reading.

20

44. The method of claim 43, further comprising compensating for non-linearity in the temperature reading to generate a scaled temperature reading, the instructions executed on the processor reading a scaled temperature reading.

25

45. A device comprising:

a plurality of processors, each processor having at least one hotspot during operation, each hotspot being a localized region having a higher temperature than regions of the processor adjoining the hotspot; and

a routing matrix comprising memory storing a decision table, the routing matrix to receive thermal event indications, each thermal event indication to indicate that one of the hotspots has reached a threshold temperature, to route thermal event directives to the processors in accordance with the decision table.

30

46. The device of claim 45, wherein if a first thermal event indication is for a first processor of the plurality of processors, at least one of the thermal event directives corresponds to an instruction the first processor to initiate a change of operating frequency.

5

47. The device of claim 46, wherein at least one of the thermal event directives corresponds to an instruction to a second processor of the plurality of processors to initiate a change of operating frequency responsive to the first thermal event indication.

10

48. The device of claim 45, further comprising a timer having an input connected to an output of the routing matrix, and having an output connected to an input of the routing matrix.

15

49. The device of claim 48, wherein if a first thermal event indication is for a first processor of the plurality of processors, the timer is started and at least one of the thermal event directives corresponds to an instruction the first processor to initiate a change of operating frequency, and

20

wherein if the timer reaches a set period of time and the first thermal event indication has not reset, at least one of the thermal event directives corresponds to an instruction to a second processor of the plurality of processors to initiate a change of operating frequency, a change of operating voltage, or a change of operating frequency and voltage.

25

50. The device of claim 49, wherein a value of the set period of time is fixed, stored in a programmable timer register, or output by the routing matrix.

30

51. The device of claim 45, further comprising a logic circuit to compare respective measured temperature readings for the hotspots and provide the routing matrix with a highest of the measured temperature readings.

52. The device of claim 51, wherein the routing matrix reports the highest of the measured temperature readings to the plurality of processors, each processor including a register to receive the highest reading.




FIG. 1

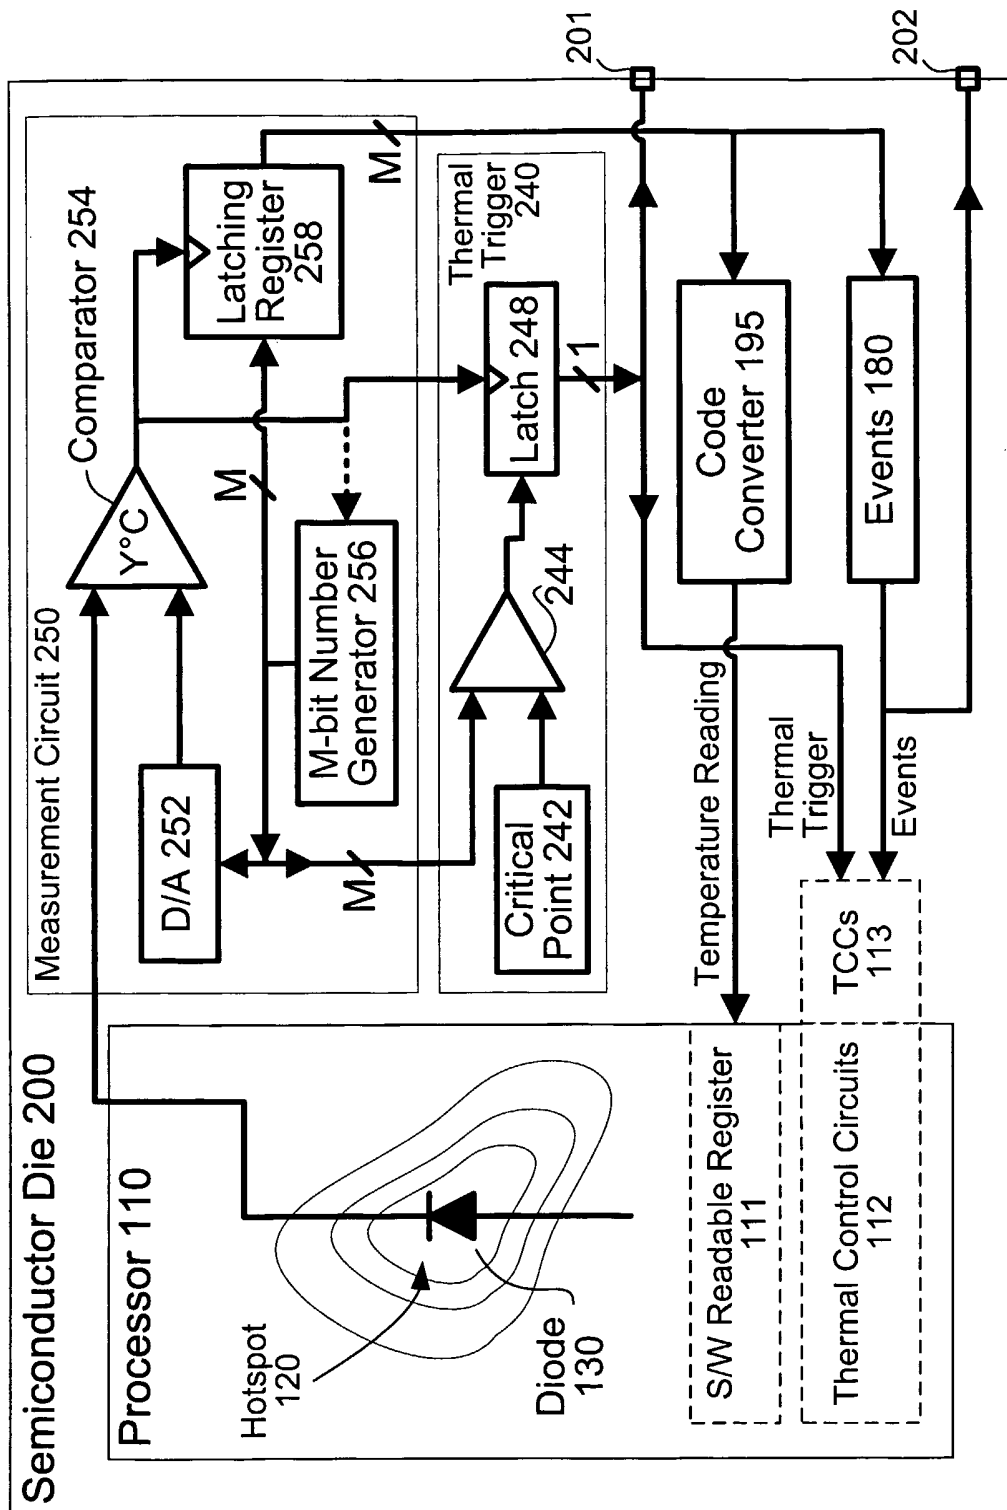


FIG. 2

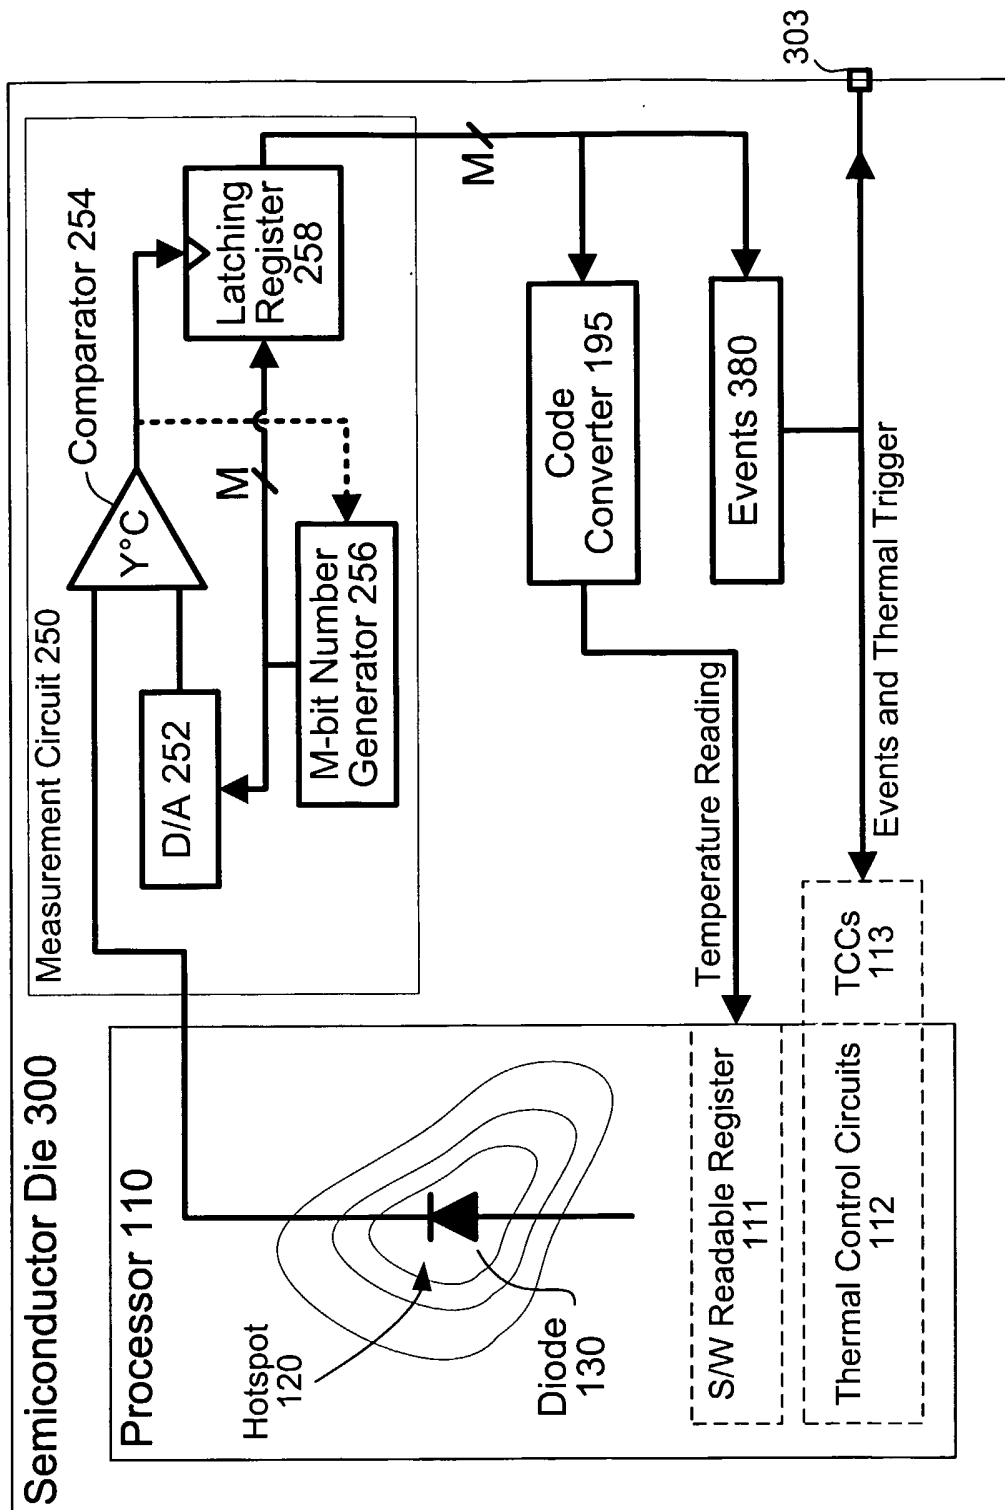


FIG. 3






FIG. 4

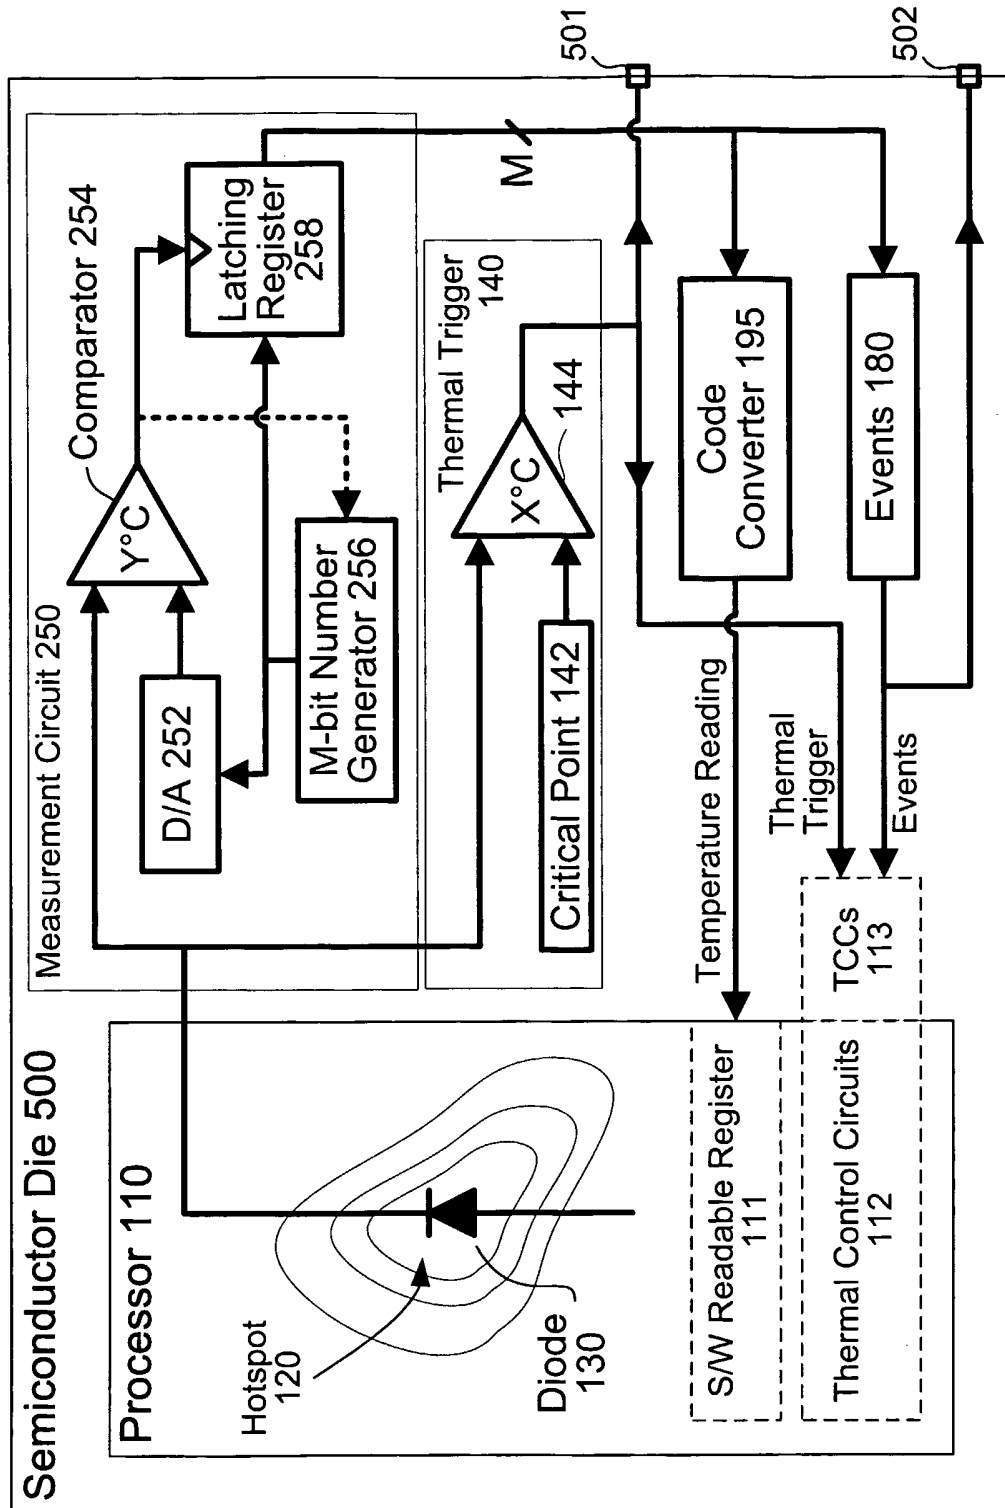


FIG. 5




FIG. 6

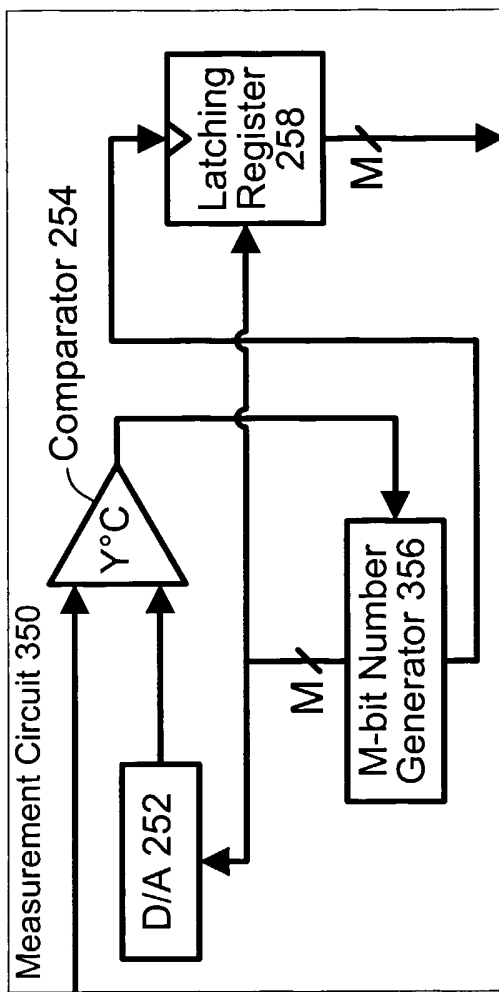


FIG. 7

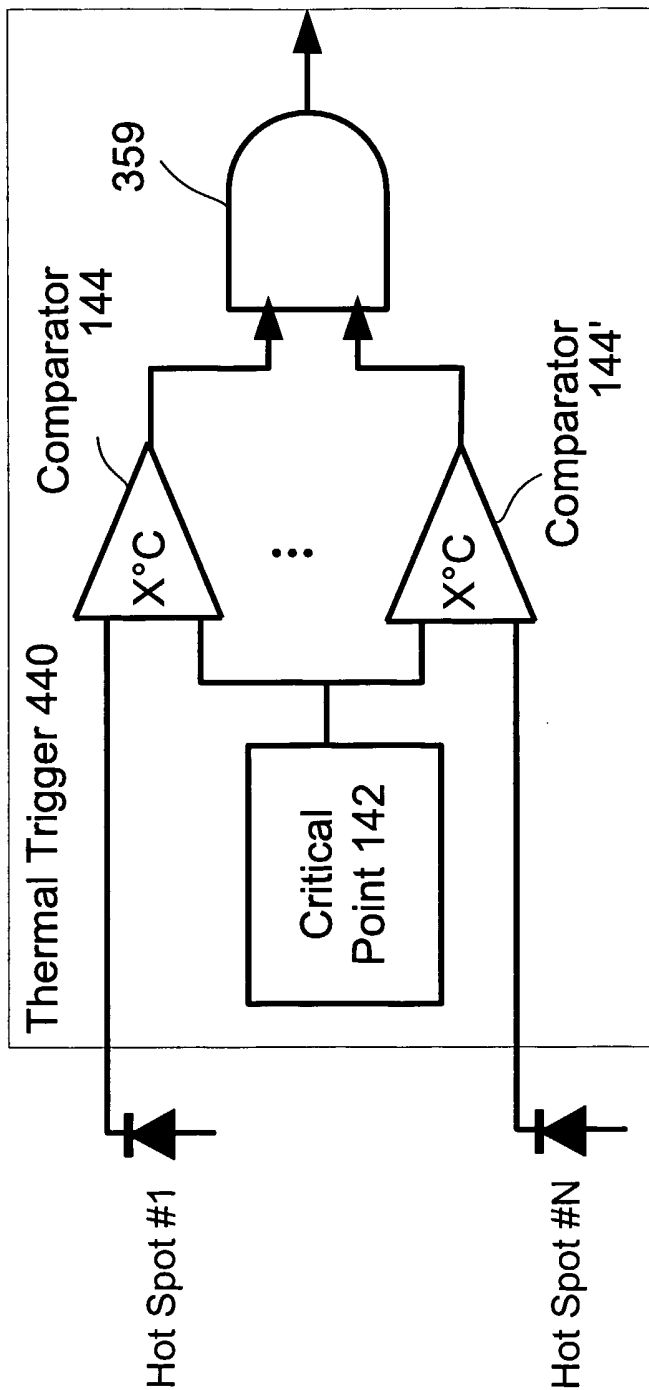


FIG. 8




FIG. 9




FIG. 10




FIG. 11






FIG. 12




FIG. 13