



US 20100077236A1

(19) **United States**

(12) **Patent Application Publication**
Trautman et al.

(10) **Pub. No.: US 2010/0077236 A1**

(43) **Pub. Date: Mar. 25, 2010**

(54) **METHOD, SYSTEM, AND APPARATUS FOR
DYNAMIC THERMAL MANAGEMENT**

(76) Inventors: **Mark A. Trautman**, Aloha, OR
(US); **Muralidhar Tirumala**,
Beaverton, OR (US)

Correspondence Address:
INTEL CORPORATION
c/o CPA Global
P.O. BOX 52050
MINNEAPOLIS, MN 55402 (US)

(21) Appl. No.: **12/592,302**

(22) Filed: **Nov. 23, 2009**

Related U.S. Application Data

(63) Continuation of application No. 11/476,955, filed on
Jun. 28, 2006.

Publication Classification


(51) **Int. Cl.**
G06F 1/00 (2006.01)

(52) **U.S. Cl.** **713/300**

(57) **ABSTRACT**

A method, apparatus, article of manufacture, and system, the method including, in some embodiments, processing a computational load by a first core of a multi-core processor, and dynamically distributing at least a portion of the computational load to a second core of the multi-core processor to reduce a power density of the multi-core processor for the processing of the computational load.

200







FIG. 1

**FIG. 2**




FIG. 3





FIG. 4

505

**FIG. 5**




FIG. 6







FIG. 7

**FIG. 8**

**FIG. 9**

**FIG. 10**

METHOD, SYSTEM, AND APPARATUS FOR DYNAMIC THERMAL MANAGEMENT

[0001] This application is a continuation of U.S. patent application Ser. No. 11/476,955 filed on Jun. 28, 2006.

BACKGROUND

[0002] A device, system, platform, or operating environment may include more than one processor or a processor having more than one core (i.e., a multi-core processor). The security, reliability, and efficient operation of such a device, system, platform, or operating environment may be enhanced by the inclusion and use of the multi-core processor. For example, a multi-core processor may provide the processing performance of multiple processors by executing multiple threads of instruction in parallel while consuming less power, costing less, and using less space than multiple single-core processors.

[0003] Operationally, the die of a single core processor may have a power density that is higher in some regions of the die (i.e., hot spots) as compared to other regions of the die. Hot spots may present challenges to efficiently managing thermal and power dissipation aspects of the processor. In some instances, a multi-core processor may have a tendency to have a greater number or intensity of hot spots as compared to a single core processor.

BRIEF DESCRIPTION OF THE DRAWINGS

[0004] FIG. 1 is an exemplary depiction of operational aspects of an apparatus, in accordance with some embodiments herein;

[0005] FIG. 2 is a flow diagram of a process, in accordance with some embodiments herein;

[0006] FIG. 3 is a flow diagram of a process, in accordance with some embodiments herein;

[0007] FIG. 4 is an exemplary depiction of characteristics relating to some embodiments herein;

[0008] FIG. 5 is an exemplary depiction of characteristics relating to some embodiments herein;

[0009] FIG. 6 is an exemplary depiction of multi-core processors, in accordance with some embodiments herein;

[0010] FIG. 7 is an exemplary depiction of multi-core processors, in accordance with some embodiments herein;

[0011] FIG. 8 is an exemplary depiction of characteristics relating to some embodiments herein;

[0012] FIG. 9 is an exemplary depiction of characteristics relating to some embodiments herein;

[0013] FIG. 10 is an exemplary depiction of a system, according to some embodiments herein.

DETAILED DESCRIPTION

[0014] The several embodiments described herein are solely for the purpose of illustration. Embodiments may include any currently or hereafter-known versions of the elements described herein. Therefore, persons skilled in the art will recognize from this description that other embodiments may be practiced with various modifications and alterations.

[0015] An apparatus may include a multi-core processor having more than one processor core (also referred to herein as a “core”) on a die. The multiple cores may provide a power efficient device, particularly with regard to processing parallel or multithreaded tasks. In some embodiments, the die of a

multi-core processor may have one or more regions of increased power density as compared to other regions of the die. The regions of increased power density may be referred to herein as ‘hot spots’ since the thermal temperature of the die at the regions of increased power density is greater than other regions of the die.

[0016] FIG. 1 is an exemplary depiction of certain aspects of an apparatus having a multi-core processor including two cores, core 1 and core 2. Graph 105 is a depiction of the multi-core processor with core 1 operating at a first power 110 and core 2 operating at a different, lower power 115. A higher operating power may be an indication of a greater processing function being performed by the core. Graph 120 is a depiction of exemplary temperatures for cores 1 and 2 shown in graph 105. Based on the power of cores 1 and 2, core 1 exhibits a higher operating temperature 125 relative to an operating temperature 130 of core 2. Accordingly, the multi-core processor including core 1 and core 2 may have an increased power density (i.e., hot spot) in a region corresponding to core 2 due to the increased temperature of core 2, as compared to core 1.

[0017] In some embodiments, the die of the multi-core processor may have a non-uniform power density due to the variance in temperatures of the cores therein. A non-uniform power density may tend to limit an overall power dissipation from the processor. Further, the occurrence of a non-uniform power density may increase as the number of cores increase for a given die size.

[0018] In some embodiments herein, a method, apparatus, system, and article of manufacture may provide mechanisms to distribute the processing power of a multi-core processor across the cores of the multi-core processor die. In general, the mechanisms to distribute the processing power of a multi-core processor across the cores of the multi-core processor may be referred to herein as dynamic thermal management (DTM). By distributing the processing of the multi-core processor between the cores at a sufficiently fast rate, the effective power density on the die may be reduced through a thermal capacitance effect of the cores. The reduced power density may result in a lower die temperature. In some embodiments, the power density may be reduced by a factor that is proportional to the number of cores included in the distribution process.

[0019] Referring again to FIG. 1, graph 135 is an exemplary depiction of core 1 and core 2 operating in accordance with some DTM embodiments herein. In graph 135, core 1 and core 2 operate at power 140 and power 145, in a time varying manner. Processing is dynamically distributed between core 1 and core 2. During certain periods of time, core 1 operates at the higher power 140 while core 2 operates at the lower power 145 and during other periods of time core 2 operates at the higher power 140 while core 1 operates at the lower power 145.

[0020] In some embodiments, there may be a disparity in a computing time constant of a core (e.g., 1e-9 seconds to 1e-6 seconds) and a thermal time constraint of the core (e.g., 1e-04 seconds to 1e-1 seconds). Accordingly, a core may be switched on, perform processing operations, and then turned off in a time less than it takes for the core to thermally heat as a result of the processing.

[0021] It is noted that a finite period of time is needed for a mass to heat due to a thermal capacitance of the mass. Accordingly, a large power spike for a core over a relatively short

period or interval of time will not typically translate to a corresponding large increase in temperature of the core due to the thermal capacitance of the core. In some embodiments, DTM mechanisms herein dynamically distribute processor power of a multi-core processor across multiple cores of the multi-core processor to effectively reduce a heat flux density of the die of the multi-core processor.

[0022] The frequency at which core 1 and core 2 alternate or swap between operating at powers 140 and 145 is faster than a time needed for either of cores 1 and 2 to heat to the temperature 125, as shown in graph 150. As shown, the maxi-

exist or if the core count included in the processing distribution process is increased (i.e., the denominator in equation 3 is proportional to the number of cores involved in the DTM sequence).

[0030] Applicant(s) have realized a DTM control mechanism using, for example, a simulation using a time-varying finite element analysis of a dual core processor package. Certain aspects of such a DTM control mechanism (e.g., a control algorithm) in accordance with embodiments herein may be expressed by the following exemplary programming code.

```

While (DTM = ON)
    HC = index to hottest core;
    CC = index to coolest core;
    If (Power(HC) > Power(CC)) // this condition could be removed and the
        Swap core loads; //cores swapped regardless of power state and still be
    End //an effective controller
    Else
        While (counter < Migration_Period) // built in timer that defines
            //migration frequency
            Check for overtemperature, break if occurs;
            Check DTM Status, break if changes to off
        End
    End
    Check DTM status, break if changes to off
End // while DTM

```

imum temperature of core 1 (line 155) and core 2 (line 160) is lower than temperature achieved in graph 1.

[0023] The distribution of processing between core 1 and core 2 may result in a reduction of the maximum die temperature in a region of core 1 and core 2 due to the power of the multi-core processor being more evenly distributed across the cores of the die. The reduction of the maximum die temperature is depicted in graph 150 as a temperature refund 165 (e.g., for graph 150, temperature refund=temperature 125-temperature 127).

[0024] Temperature refund 165 may be a result of a lower power density of the multi-core processor die. On a time averaged basis, a maximum power density may be expressed by the following equations:

[0025] No DTM:

$$\text{Maximum die power density} = Pd_1 = (P1)/(A) \quad (1)$$

[0026] With DTM:

$$\text{Maximum die power density} = Pd_2 = (P1+P2)/(2A) \quad (2)$$

[0027] where Pd_1, P1, and A are the power density, power, and core area for the NO DTM case, respectively. Also, Pd_2, P1, P2, and A are the power density, core 1 power, core 2 power, and core area for the DTM case, respectively.

[0028] In some embodiments, a potential relative reduction ratio of the power density between the instances of NO DTM (e.g., graphs 105, 120 and eq. 1) and active DTM (e.g., graphs 135, 150 and eq. 2) may be expressed as a Power Density Reduction Potential (PRPD), as follows:

$$\text{PRDP} = (Pd_1 - Pd_2)/Pd_1 = [1 - P2/P1]/2 \quad (3)$$

[0029] Equation 3 suggests that smaller ratios of P2/P1 lead to greater reductions of the power density with a limit of one-half (1/2) reduction in the power density for a dual core processor. This indicates that a greater benefit may be obtained when a larger disparity between core power states

[0031] In some embodiments, FIG. 3 is an exemplary flow diagram corresponding to the coded instructions above. It is noted that FIG. 3 may be extended to cover other control schemes in addition to and different than those relating to the above code listing.

[0032] At operation 305, a DTM controller or other mechanism is invoked. At operation 310 at least a portion of a computational load being processed by a multi-core processor is routed from a core having a temperature higher than other cores of the multi-core processor to a core having a lower temperature than the other cores. In some embodiments, operation 310 routes processing of the highest temperature core to the lowest temperature core.

[0033] Operation 310 includes a basis for the routing of the computational load 305 between the cores of the multi-core processor. In some embodiments, the basis for the routing of the computational load (e.g., an algorithm, basis, relationship, etc.) may include more, fewer, and different factors than the temperature illustrated at operation 310.

[0034] For example, the migration period provided in the corresponding code above may be based on a fixed time period (i.e., a fixed rate), may be based on a variable time period that is a function of a core temperature (i.e., a migration frequency that is temperature adaptive), and combinations thereof. In some embodiments, multiple migration frequencies may be used. The migration frequencies may vary in a linear or a non-linear manner from a possible low rate(s) to a high rate(s).

[0035] At least a portion of the computational load may be distributed to at least one of the cores 315, 320, 325. In some embodiments, at least a portion of the computational load may be distributed from the core having the highest temperature to the one core having the lowest temperature. The basis for the routing may vary in accordance with the various embodiments herein.

[0036] At operation 330, a determination is made whether there is an over-temperature condition for the core(s) processing at least a portion of the computational load. Also, a determination may be made at operation 330 to determine whether the DTM process is still active. In an instance there is an over-temperature condition or the DTM process is not active, process 300 proceeds to exit 335.

[0037] In an instance there is not an over-temperature condition and the DTM process is still active, process 300 proceeds to operation 310. At operation 310, at least a portion of the computational load is again routed from a core having a temperature higher than other cores of the multi-core processor to a core having a lower temperature than the other cores. In some embodiments, the rate at which operation 310 is performed is the same as the migration frequency of the DTM process.

[0038] FIG. 4 is an exemplary depiction of two cores, core 1 and core 2, operating under various conditions per a DTM mechanism, in accordance herewith. In the examples shown in graphs 405, 410, 415, and 420 a dual core processor is depicted. Core 1 is operating at 65 watts (W) and core 2 is in an idle state of 35 W. The difference in core power is referred to herein as the migration amplitude (MA) of the cores and represents an opportunity to exchange power between the two cores. In this instance, the MA is 30 W. For each of graphs 405, 410, 415, and 420, a temperature response for the cores under an active DTM condition is shown. For graph 405 the MF (migration frequency)=1 Hz, for graph 410 the MF=10 Hz, for graph 415 the MF=100, and for graph 420 the MF=1000 Hz.

[0039] A temperature refund is equal to the overall reduction in the instantaneous peak temperature that may be achieved by initiating the DTM mechanism at time=0 seconds. As illustrated, a higher load MF is more effective at distributing the heat over the two cores. For example, at the MF=100 Hz there is about a 4.5° C. temperature reduction in the maximum die temperature and at MF=1000 Hz the temperature reduction increases to nearly 6° C. It is noted that the thermal budget for a processor package may be, for example, about 25° C. to 30° C. (at 130 W). Thus, a temperature refund provided in accordance herewith by DTM mechanisms may represent an effective 20%-25% improvement in the thermal performance of an processor package.

[0040] FIG. 5 is an exemplary summary representation 505 illustrative of how a temperature refund scales with a load migration amplitude (MA) and a load migration frequency (MF). In particular, the temperature refund increases with both migration amplitude and migration frequency. In instances where large migration amplitudes are available (e.g., 50 W-120 W), a reduction of about over 20° C. may be achieved using migration frequencies between 100-1000 Hz. Accordingly, DTM mechanisms in accordance with some embodiments herein may be used to enhance scalar computing performance of multi-core processors. Furthermore, DTM mechanisms in accordance with some embodiments herein may be used without impacting parallel applications performance that may use both (multiple) cores simultaneously to obtain the highest throughput.

[0041] In some embodiments, a DTM mechanism in accordance herewith may provide improved scalar computing. In some embodiments, an individual core frequency may be increased, thereby providing improved performance on scalar tasks. Also, the multi-core architecture of the processor may still be utilized for high throughput in applications with, for

example, high levels of parallelism. That is, an adaptive nature of the DTM mechanisms herein may enable high scalar performance without impacting high throughput during parallel applications.

[0042] In some embodiments, the DTM mechanisms in accordance herewith may be adaptive in the sense that such features may be selectively activated. For example, a hardware implemented DTM control may be selectively turned on and turned off by an operating system (O/S) of a device or system.

[0043] FIG. 6 is an exemplary depiction illustrating adaptive aspects of a DTM mechanism, in accordance with embodiments herein. FIG. 6 includes three configurations of a multi-core processor or array of cores. In particular, there is shown a multi-core processor 605 having uniform array of cores 610, all cores therein operating at frequency, f.

[0044] Also shown is a multi-core processor 615 including an array of cores 630 wherein a number of cores 630 are grouped into two clusters 620 and 625. Clusters 620 and 625 may, under the control of a DTM mechanism in accordance herewith, operate as superscalar clusters that are selected from adjacent cores 630. Clusters 620 and 625 may operate at a higher frequency than the remaining cores of multi-core processor 615 not included in clusters 620 and 625. For example, clusters 620 and 625 may operate at a frequency= $f+\Delta f$, while the cores not included in clusters 620 and 625 operate at a frequency= f . Clusters 1 and 2 may operate at the higher frequency ($f+\Delta f$) without increasing the power dissipation of multi-core processor 615 in accordance with the DTM mechanism disclosed herein.

[0045] Multi-core processor 635 may operate under control of a DTM mechanism in accordance herewith to form a cluster 640. Cluster 640 is formed by a grouping of non-adjointing cores 645. Cores of cluster 640 may be operated at a higher frequency than the cores not included in the cluster since the cores of cluster 640 have the computational load being processed by the cluster dynamically distributed amongst the cores of the cluster, in accordance with embodiments herein.

[0046] In some embodiments, clusters 620, 625, and 640 may operate as a superscalar core. When, for example, the need for the superscalar cores 620, 625, and 640 are no longer needed (i.e., no longer processing scalar tasks), the DTM functionality associated with multi-core processors 515 and 520 may be turned off and the clustered cores returned to the collective array of core.

[0047] In some embodiments, the number of cores included in a cluster may vary. For example, a cluster may include at least two cores. The clustered cores may be adjoining, non-adjointing, and combinations thereof. In some embodiments, the configuration or groupings of cores may be predetermined or vary in accordance with operational contexts. For example, the number of cores included in a cluster(s) may depend on the number of cores available for clustering, the power to be dissipated, the computational tasks and/or computational load being processed, and other factors.

[0048] FIG. 7 is an exemplary depiction illustrating adaptive aspects of DTM mechanisms, in accordance with embodiments herein. For example, multi-core processor 705 having a plurality of cores 720 includes two clusters of cores grouped as cluster 710 and cluster 715. Regarding clusters 710 and 715, a DTM mechanism in accordance herewith may be invoked to dynamically distribute the processing of a computational load among cores within a cluster (indicated by the

arrows within the cluster). In some embodiments, DTM mechanisms may be used to distribute processing amongst the cores within cluster 610 and DTM mechanisms may also be used to control to distribute processing amongst the cores within cluster 715. The cores in clusters operate 710 and 715 may operate at a substantially higher frequency than the cores not in the cluster, in accordance with the aspects of DTM disclosed herein.

[0049] Multi-core processor 725 having a plurality of cores 740 includes two clusters of cores, cluster 730 and cluster 735. For clusters 730 and 735, a DTM mechanism in accordance herewith may be invoked to dynamically distribute the processing of a computational load across the clusters (indicated by the arrows between the clusters). In some embodiments, DIM mechanisms may be used to distribute processing between clusters 730 and 735. Here, the cores in clusters 730 and 735 may operate at a substantially higher frequency than the cores not in the cluster.

[0050] Thus, a DTM mechanism may be applied to clusters cores in a variety of manners, including amongst cores within clusters (multi-core 605), between clusters (multi-core processor 625), and a combination thereof (not shown).

[0051] In some embodiments, aspects of the DTM mechanisms herein may be used as a throttle mechanism to correct for over temperature event. Over temperature events may occur in a device or system due to, for example, poor platform thermal management solutions otherwise employed in the device or system. In some embodiments, the migration frequency associated with a DTM mechanism may act as a throttle control. The temperature of the multi-core processor may be reduced as migration frequency increases, even though there may typically be more computation overhead at a higher migration frequency. FIG. 8 illustrates an exemplary comparison of the relative performance of a DTM mechanism and voltage/frequency throttling, over a range of available migration ratios, to reduce a temperature of a multi-core processor.

[0052] Note that the migration ratio (MR) represents the amount of heat that is available to be migrated between cores, and $MR=1-(\text{Low power core Watts})/(\text{High power core Watts})$. Also, at $MR=0$, there is no opportunity to migrate heat between cores and at $MR=1$ there is full (100%) opportunity to migrate heat between cores. For some multi-core processors, MR =about 0.6 to about 1.0 may be typical.

[0053] Referring to FIG. 8, data in graph 700 are for equivalent levels of temperature reduction. The data shows that a DTM mechanism operating at 100 Hz has a lower performance penalty than voltage/frequency (v/f) scaling for a wide range of MR values. The graphed data suggests that DTM mechanisms in accordance with some embodiments herein may be a more effective throttle mechanisms than a v/f scaling approach.

[0054] In some embodiments, a temperature refund obtained through the use of a DTM mechanism may be used to lower an acoustic emission of a device or system having a multi-core processor and a cooling device that produces acoustic emissions (e.g., a fan). An example of such a system may include a personal computer having a multi-core processor and at least one cooling fan. The lowering of the acoustic emissions may be the sole purpose for invoking the DTM mechanism and, in some embodiments, invoking the DTM mechanism may also at least contribute to increasing the power and performance of the multi-core processor.

[0055] In some embodiments, a temperature refund is used to lower the revolutions per minute (RPM) of a cooling processor fan until a desired temperature is reached. The desired temperature may be equivalent to the temperature the processor would achieve in the absence of activating or including the DTM mechanism. In this manner, the processor is not allowed to operate at a temperature any worse than it would normally operate (i.e., within design specifications). The lower fan RPMs may significantly lower a noise signature of the device or system including the multi-core processor and cooling fan.

[0056] FIG. 9 is an exemplary graph 900 illustrating that there is a small degradation in performance due to penalties accrued by switching the core context (ST penalty) and warming up the core (WUT penalty). However, a significant temperature reduction may be achieved. It is expected that a 5° C. temperature translates to about a 0.5 BA of acoustic noise. As illustrated, the temperature reduction is reached with the DTM mechanism operating at about 100 HZ, suggesting for example that acoustic emission reduction using DIM mechanism may be used to implement a “whisper” mode of operation for processor based devices and systems.

[0057] In some embodiments, the DTM mechanism control may be applied by a user (e.g., end-user, technician, etc.). In some embodiments, the DTM mechanism for acoustic reduction could be invoked for a “whisper” mode of operation, turned off for typical processing applications, and invoked to increase power or performance of the multi-core processor in a “turbo” mode of operation.

[0058] In some embodiments, a DTM mechanism in accordance herewith may be used to reduce leakage power of a processor. This aspect of some DTM mechanism herein may be particularly suited, though not limited to, mobile applications where battery life is highly valued.

[0059] In an instance a multi-core processor is operating without DTM mechanisms in accordance with embodiments herein, an active core may produce (severe) hot spots in the region of the active core. Accordingly, the leakage power of the active core is reflected in a higher temperature field.

[0060] It is noted that leakage power may be a highly non-linear function of temperature. Thus, a hot spot caused by an active core may result in a large or significant leakage power.

[0061] In an instance a multi-core processor is operating with DTM mechanisms activated and processing of a computational load is dynamically distributed among multiple cores in accordance with embodiments herein, active cores may avoid producing hot spots. The resultant heat spreading may produce a lower temperature field. Accordingly, the leakage power for the multi-core processor may correlate to a lower temperature environment. Also, due to the temperature dependence of the leakage power the overall leakage power may be lowered, thereby extending, for example, battery life of a mobile device. In some embodiments, a leakage power savings on the order of about 5 to about 10 watts may be expected.

[0062] In some embodiments, a dynamic distribution of processor power of a multi-core process across multiple cores of the multi-core processor is accomplished at a frequency (e.g., a migration frequency) sufficiently fast to distribute the power over the cores and reduce the power density, and yet only increases a computational overhead a relatively small amount.

[0063] FIG. 10 is an exemplary depiction of a system 1000 including an apparatus, for example a multi-core processor 1005 in communication with a controller 1010. A memory

1015 is attached to controller **1010** by a conductor and other electrical connections. Cooling device **1020** may be provided to at least cool multi-core processor **1005**.

[0064] Controller **1010** may include a hardware implemented DTM mechanism, in accordance herewith. In some embodiments, code or program instructions may be stored in controller **1010** and further executed by the controller to effectuate the DTM mechanisms herein. In some embodiments, at least a portion of memory **1015** may be used to store code or program instructions used by controller **1010**, an operating system, and other information.

[0065] Those in the art should appreciate that system **1000** may include additional, fewer, or alternative components to multi-core processor **1005**, controller **1010**, memory **1015**, and cooling device **1020**.

[0066] In some embodiments, cooling device **1020** may include a fan. Memory **1015** may comprise any type of memory for storing data, including but not limited to a Single Data Rate Random Access Memory, a Double Data Rate Random Access Memory, or a Programmable Read Only Memory.

[0067] It should be appreciated that the drawings herein are illustrative of various aspects of the embodiments herein, not exhaustive of the present disclosure.

1. A method comprising:
processing a computational load by a first cluster of at least two associated cores of a multi-core processor;
dynamically distributing at least a portion of the computational load to a second cluster of at least two associated cores of the multi-core processor to reduce a power density of the multi-core processor for the processing of the computational load; and
processing the dynamically distributed portion of the computational load by the second cluster, wherein the dynamically distributing and processing occurs at a frequency greater than a thermal time constant associated with a temperature response of the first cluster and the second cluster.
2. The method of claim 1, wherein the dynamic distributing occurs at a predetermined frequency.
3. The method of claim 1, wherein the dynamic distributing occurs at an adaptive frequency.
4. (canceled)
5. The method of claim 1, wherein the temperature response includes at least one of a predetermined operating temperature of the multi-core processor, an adaptive operating temperature of the multi-core processor, and combinations thereof.
6. (canceled)
7. The method of claim 1, wherein a temperature of the second cluster of at least two associated cores is lower than a temperature of the first cluster of at least two associated cores.
8. The method of claim 1, further comprising dynamically distributing at least a portion of the computational load from the second cluster of at least two associated cores back to the first cluster of at least two associated cores.
9. The method of claim 1, wherein the dynamic distributing is further based on relationship between a migration frequency and the temperature response of the first and second clusters of the multi-core processor.
10. An apparatus comprising:
a multi-core processor including a plurality of cores, the plurality of cores including a first cluster of at least two

associated cores to process a computational load and a second cluster of at least two associated cores; and

a controller to dynamically distribute at least a portion of a the computational load to the second cluster of at least two associated cores of the plurality of cores to reduce a power density of the multi-core processor for the processing of the computational load by processing the dynamically distributed portion of the computational load by the second cluster of at least two associated cores, wherein the dynamically distributing and processing occurs at a frequency greater than a thermal time constant associated with a temperature response of the first cluster and the second clusters.

11. The apparatus of claim 10, wherein the dynamic distributing occurs at a predetermined frequency, an adaptive frequency, and combinations thereof.

12. The apparatus of claim 10, wherein the dynamic distributing is based on a temperature response of the multi-core processor.

13. The apparatus of claim 12, wherein the temperature response includes at least one of a predetermined operating temperature of the multicore processor, an adaptive operating temperature of the multi-core processor, and combinations thereof.

14. (canceled)

15. The apparatus of claim 10, wherein the controller recursively and dynamically distributes at least a portion of the computational load based on a power density response of the multi-core processor.

16. The apparatus of claim 10, wherein the controller dynamically distributes at least a portion of the computational load to more than one of the plurality of the cores.

17. An article of manufacture having program instructions embodied thereon, the article of manufacture comprising:

program instructions for processing a computational load by a first cluster of at least two associated cores of a multi-core processor having a plurality of cores;

program instructions for dynamically distributing at least a portion of the computational load to at least one a second cluster of at least two associated cores of the plurality of cores to reduce a power density of the multi-core processor for the processing of the computational load; and

program instructions for processing the dynamically distributed portion of the computational load by the second cluster, wherein the dynamically distributing and processing occurs at a frequency greater than a thermal time constant associated with a temperature response of the first cluster and the second cluster.

18. The article of manufacture of claim 17, wherein the dynamic distributing occurs at a predetermined frequency, an adaptive frequency, and combinations thereof.

19. The article of manufacture of claim 17, wherein the dynamic distributing is based on at least one of a predetermined operating temperature of the multi-core processor, an adaptive operating temperature of the multi-core processor, and combinations thereof.

20. The article of manufacture of claim 17, wherein the dynamic distributing occurs at a frequency greater than a thermal time constant associated with a temperature response of the multi-core processor.

21. A system comprising:

a multi-core processor including a plurality of cores, the plurality of cores including a first cluster of at least two

associated cores processing a computational load and a second cluster of at least two associated cores;

a controller to dynamically distribute at least a portion of the computational load to the second cluster of at least two cores of the plurality of cores to reduce a power density of the multi-core processor for the processing of the computational load by processing the dynamically distributed portion of the computational load by the second cluster, wherein the dynamically distributing and processing occurs at a frequency greater than a thermal time constant associated with a temperature response of the first cluster and the second clusters; and

a double data rate memory.

22. The system of claim **21**, wherein the dynamic distributing occurs at a frequency greater than a thermal time constant associated with a temperature response of the multi-core processor.

23. The system of claim **21**, wherein the dynamic distributing occurs at a predetermined frequency, an adaptive frequency, and combinations thereof.

24. The system of claim **21**, wherein the dynamic distributing is based on a temperature response of the multi-core processor.

25. A method comprising:

processing a computational load by a first core of a multi-core processor; and

throttling a processor pipeline of the multi-core processor to support dynamic power and performance optimizations by dynamically distributing at least a portion of the computational load to a second core of the multi-core processor to reduce a power density of the multi-core processor for the processing of the computational load in order to trade levels of activity and power of the first core of the multi-core processor for increased activity in the second core of the multi-core processor.

26. The method of claim **25**, wherein the dynamic distributing occurs at a predetermined frequency.

27. The method of claim **25**, wherein the dynamic distributing occurs at an adaptive frequency.

28. The method of claim **25**, wherein the dynamic distributing is based on a temperature response of the multi-core processor.

29. The method of claim **28**, wherein the temperature response includes at least one of a predetermined operating temperature of the multi-core processor, an adaptive operating temperature of the multi-core processor, and combinations thereof.

30. The method of claim **25**, wherein the dynamic distributing occurs at a frequency greater than a thermal time constant associated with a temperature response of the multi-core processor.

31. The method of claim **25**, wherein a temperature of the second processor core is lower than a temperature of the first processor core.

32. The method of claim **25**, further comprising dynamically distributing at least a portion of the computational load from the second core back to the first core.

33. The method of claim **25**, wherein the dynamic distributing is based on a relationship between a migration frequency and a temperature response of the multi-core processor.

34. A method comprising:

measuring a temperature of at least one portion of a multi-core processor; and

dynamically reducing an activity of a first core of the multi-core processor and increasing an activity of a second core of the multi-core processor in response to the measuring.

35. The method of claim **34**, wherein the dynamic reducing and increasing occurs at a predetermined frequency.

36. The method of claim **34**, wherein the dynamic reducing and increasing occurs at an adaptive frequency.

37. The method of claim **34**, wherein the dynamic reducing and increasing is based on a temperature response of the multi-core processor.

38. The method of claim **37**, wherein the temperature response includes at least one of a predetermined operating temperature of the multi-core processor, an adaptive operating temperature of the multi-core processor, and combinations thereof.

39. The method of claim **34**, wherein the dynamic reducing and increasing occurs at a frequency greater than a thermal time constant associated with a temperature response of the multi-core processor.

40. The method of claim **34**, wherein a temperature of the second processor core is lower than a temperature of the first processor core.

41. The method of claim **34**, further comprising dynamically distributing at least a portion of the computational load from the second core back to the first core.

42. The method of claim **34**, wherein the dynamic reducing and increasing is based on a relationship between a migration frequency and a temperature response of the multi-core processor.

* * * * *