Heat Transfer Principles

Chapter Eight

Multi-Core CPU Air Cooling

Unit Introduction

Electronic gadgets that operate at high speeds produce greater heat as compared to other equipment. The issue of air cooling for portable electronic devices is covered in this chapter. Air cooling restrictions impact portable electrical gadgets. In the upcoming years, multi-core CPUs will predominate in mobile handset platforms. Solutions for the DTM (dynamic thermal management) of the central processor unit (CPU) are provided by advanced control techniques. This chapter aims to employ fuzzy logic control to reduce the effect of air cooling constraints and ensure sustainable CPU usage. The construction of a DTM regulator depending upon fuzzy logic control is the main component of the suggested remedy for the air cooling constraint. Because this is independent of the Central processing unit and its air conditioning system transmit functions, this technique lowers issue design time. Thermal gradients or differences in operational temperature along a design are calculated and reported by on-chip thermal analysis. For modern digital integrated circuits, this analysis is becoming increasingly crucial (ICs). Putting cores into a CPU processor with today's 65 nanometers and 45 nanometers technology enhances power density and causes thermal throttling. The Central processing unit thermal throttling issue is resolved using modern control methods. To that end, a thermal model resembling an actual IBM CPU chip with eight cores is constructed.

Such a thermal model has been incorporated into a thermal simulator for semiconductors. The CPU chip's open-loop output is collected. That thermal profile of a CPU chip depicts CPU thermal throttling. The DTM regulator design described depends upon 3-dimensional fuzzy logic. The CPU chip contains several cores, each of which is a heat source. The link between such core temperatures and their operating frequencies enhances DTM response and diminishes the influence of air cooling constraints. The three-dimensional fuzzy controller takes such associations into account. This chapter introduces the "Thermal Spare Core (TSC)" algorithm, a novel DTM approach. The TSC is a brand-new DTM algorithm. TSC function by reserving CPU cores during periods of low CPU load and triggering them during thermal emergencies. The designation of specific CPU cores as TSC has no impact on the total CPU use. Because of air cooling limits, these cores aren't engaged concurrently.

The CPU chip can accommodate additional cores thanks to semiconductor progress. Therefore, in the case of TSC, there isn't any chip area wastage. The TSC provides a workaround for the problems with multi-core CPU air cooling.

Learning Objectives

By the end of this chapter, students will be able to learn the following:

1. Functions of multi-core CPU air cooling
2. Thermal spare core
3. Analysis based on simulation

Key Terms

1. Multi-Core CPU
2. Thermal spare core
3. Thermal throttling problem
4. Simulation Analysis

8.1. The CPU air cooling limitations

Our time is computer-controlled, and our lives are so frequently dependent on machines and their programming that we are unaware of it. For instance, in our daily lives, we rely on digital processors for mobile phones, portable electronics, computers, medical equipment, and several other gadgets. There is no question that the portability of this portable technology is being hampered by its size and weight. Unfortunately, a lot of things can affect how portable electrical systems are. The battery is being impacted by power usage. Because of air cooling restrictions, effective cooling of portable electronic gadgets is becoming an issue.

Did you know?

Air cooling is a method of dissipating heat. It works by expanding the surface area, increasing the flow of air over the object to be cooled, or both.

Designing for the on-chip temperature difference is difficult. Numerous technological aspects influence the gradients in chip temperature. With each successive technology node, power density (power per unit area) has risen from a technical perspective (ITRS, 2006). After all, smaller geometries allow for the integration of more functionality into a chip's same space, which can lead to large temperature gradients (Huangy et al., 2006). The overall power usage of the CPU chip increases as more cores are added, as illustrated in Figure.1A, and Figure.1B depicts the maximum amount of cores per chip and its optimum operating frequencies (D.D.Kim et al., 2008).

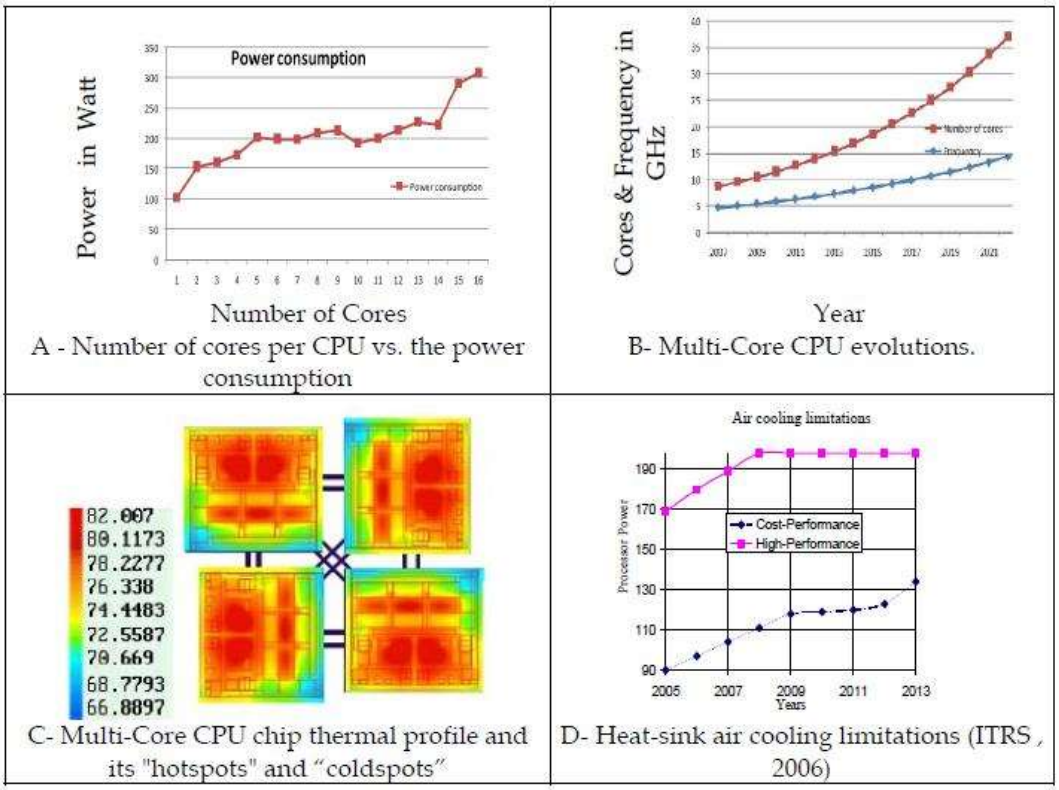

Figure 8.1. Evolutions of Multi-Core CPU (Source: Pranav Tendulkar, Creative Commons License).

On-chip memory often runs cold, whereas the CPU cores are typically exceptionally hot. Consequently, depending upon the method of operation, there are constantly changing combinations of "cold" and "hot" regions. A good illustration of such a design seems to be a cell phone. Text message creation would activate particular functionality that generates a particular thermal profile. However, sending this message would activate diverse functions, which creates a new profile. The same applies to utilizing a cell phone to capture a picture, play an audio file, or make a voice call. A typical range for the temperature fluctuation across a chip is ten degrees Celcius to fifteen degrees Celcius. If such temperature distribution is not controlled, temperature differences could reach thirty to forty degrees Celsius (Mccrorie, 2008).

Leakage and Dynamic power work together to dissipate the CPU's power (Kim et al., 2007). Buffer strengths, logic toggle rates, and parasitic loading affect dynamic power. The device and technology characteristics affect the leakage power. Both sources of power should be taken into consideration in thermal analysis solutions. Figure 1C's thermal profile illustrates the temperature change across a CPU chip's surface. This phenomenon results from the power density changing depending upon every function block design. On the surface of the CPU chip, such power density division creates "coldspots" and "hotspots" (Huangy et al., 2006). The higher operating temperature of

the CPU lowers transistor efficiency by increasing leakage current, lowering electromigration limits, and increasing interconnection resistance (Mccrorie, 2008). Leakage current also raises power usage.

8.2. The CPU thermal throttling problem

The CPU chip's ability to accommodate more cores while maintaining a smaller size and faster speed is made possible by the fabrication technique. However, adding multiple cores to a CPU chip improves power density and presents new problems for dynamic power management. The amount of transistors that may be installed on an integrated circuit (IC) has expanded dramatically since its conception, doubling roughly every two years (Moore, 1965). Gordon E. Moore, a co-founder of Intel, noted the tendency for the first time in a 1965 article. Moore's law has persisted for nearly fifty years! Moore's 1965 discussion of the heat problem was no accident: "Is it conceivable to remove the heat created by a large number of elements on a silicon chip?" (Moore, 1965). Relative to the dynamic power for CMOS technology, the static power usage of the IC had been disregarded. Static power is currently a design issue. The CPU chip's millions of transistors emit more heat than previously. The capability of the CPU unit air conditioning system restricts the number of cores on the Central processing unit chip (ITRS, 2008).

A group of professionals from the semiconductor industry created a set of documents known as the International Technology Roadmap for Semiconductors (ITRS). The upper limit for high-efficiency heat-sink air cooling, as specified by ITRS, is 198 Watts (ITRS, 2006). The level capacity of the air conditioning system sets a restriction on the chip power usage layout, according to Figure. 1D, we already passed the air cooling limit in 2008.

The CPU reaches its maximum operating temperature after a specific time because of the highest CPU usage, as depicted in Figure 8.2A. As a result, the CPU usage is decreased to a safe level so as not to surpass. CPU thermal throttling is the term for this occurrence. The comparison of the best scenario, "no thermal limitations," the best scenario having "low energy usage with thermal limitations," and the best case with "high power usage having thermal limitations" is shown in Figure 8.2B. CPU usage doesn't increase with the number of cores added to the CPU chip. The curve drifts toward lower CPU usage due to the CPU's thermal limitation in cases of low power usage. By attaching extra cores to the CPU chip, the CPU utilization falls in cases of high power usage. As a result, the increase in CPU usage isn't related to the core count.

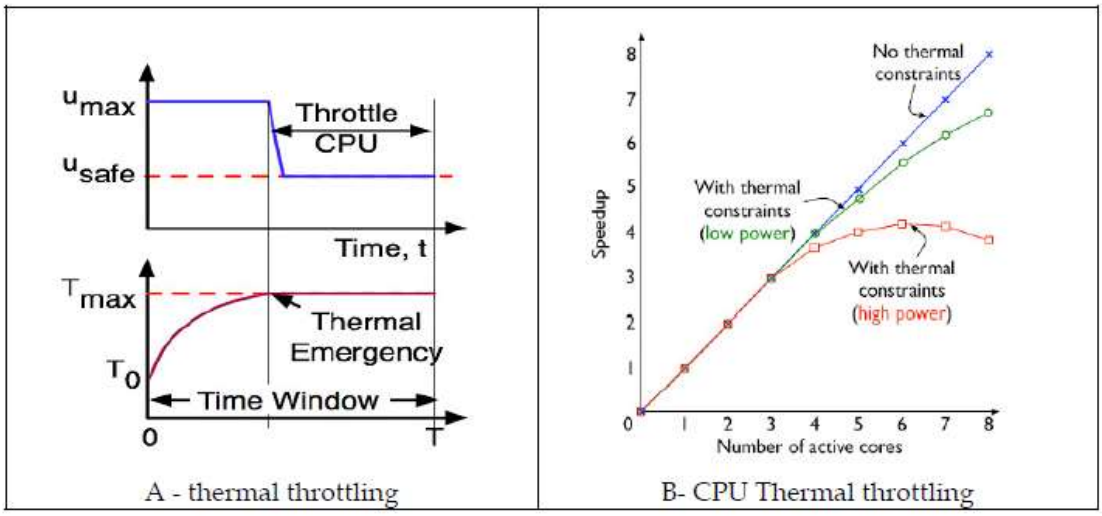

Figure 8.2. CPU thermal throttling (Source: novatech.co.uk, Creative Commons License).

8.3. The advanced DTM controller design

To stop CPU thermal throttling, modern dynamic thermal management solutions are essential. The fuzzy control offers a practical way to build nonlinear controllers using heuristic information. This heuristic data could originate from a process operator who has functioned as a "human-in-the-loop" regulator. A set of guidelines for controlling the process are written down using the fuzzy control design methodology. Once a fuzzy controller that simulates decision-making has been created, put these rules into it.

The consumer-friendly and high-efficiency control offered by fuzzy control is independent of the source of the control knowledge (Patyra et al., 1996). The DTM approach must utilize all of the CPU's resources. The DTM prevents thermal throttling for portable devices and maintains battery life. The speed of each core's "working frequency" is determined by the DTM controller's measurement of the temperatures of the CPU cores. Temperature and operation frequency both affect how much power is used. The relationship between temperature and power loss is that of the temperature change.

A DTM approach called DVFS (dynamic voltage and frequency scaling) modifies a core's operating frequency while running (Wu et al., 2004). The stop-go strategy, often known as clock gating (CG), involves stopping all vigorous activities (Donald & Martonosi, 2006). To halt development till the thermal emergency has been ended, CG disables the clock signals. The state of the processor, comprising branch predictor tables, registers, and local caches, is retained while dynamic operations have been frozen (Chaparro et al., 2007). Therefore, less dynamic power was used while waiting. GC is more akin to a suspend or sleep switch than an off switch. The real-time Operating system-based DTM approach is referred to as TM (thread migration), also referred to as core hopping. By moving core task "threads" from an extra heated core to a core with a low temperature, TM lowers the CPU

temperature. To execute DVFS, the current standard DTM regulator employs P (proportional) or PID (proportional-integral-derivative), or PI (proportional-integral) controllers (Donald & Martonosi, 2006; Ogras et al., 2008).

Lotfi A. Zadeh first proposed fuzzy logic in 1965 (Trabelsi et al., 2004). The classic fuzzy set has two dimensions, one for the variable's domain of discussion and the second for the membership level. A non-linear system can be handled by this two-dimensional fuzzy logic controller (FC) without the need to identify the system transfer function. On the other hand, this two-dimensional fuzzy set cannot handle a system with a spatially dispersed variable. A traditional fuzzy set with an additional dimension for spatial information makes up a 3D (three-dimensional) fuzzy set. Unlike the conventional two-dimensional FC, the three-dimensional FC combines numerous sensors to give three-dimensional fuzzy inputs. The three-dimensional FC creates the "spatial membership function, "which also has three-dimensional information. The two-dimensional Fuzzy rules and the three-dimensional rules are equivalent. The quantity of spatial sensors has no bearing on the number of rules. This three-dimensional FC's computation is appropriate for use in practical situations.

8.4. Thermal spare core

Since a CPU isn't always fully employed, a few of the cores may be set aside for heat emergencies. Take a look at Figure. 8.3A. the cooling system may dissipate the expelled heat outside the chip once a core reaches steady state temperature T_1. The cooling mechanism cannot expel the heat from this core if it becomes overheated. As a result, the core temperature rises until it hits T_3, which is the thermally throttling temperature (Rao & Vrudhula, 2007).

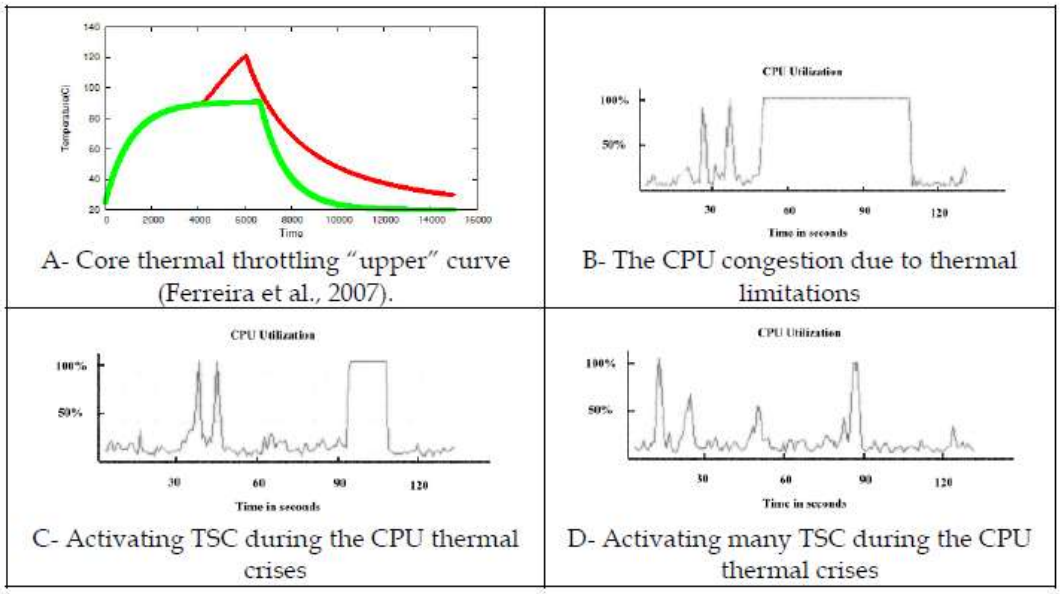

Figure 8.3. TSC Diagram (Source: Raghad Hameed, Creative Commons License).

The identical thermal events depicted in Fig. occur due to cooling system flaws (Ferreira et al., 2007). The CPU chip can accommodate additional cores thanks to semiconductor progress. ITRS states that the overall chip area overhead is equal to 27.9 percent (ITRS, 2009). Therefore, in the situation of TSC, there isn't any chip area wastage. Thus, retaining cores as TSC has no effect on the usage of the CPU as a whole. Because of temperature restrictions, such cores are not triggered simultaneously. "Serial sections restrict parallel speedups," states Amdahl's law (Gustafson, 1988). As a result of the serial part limitations, putting extra cores in the CPU chip doesn't increase speed. As a result, if parallelism is absent, not all cores are effectively utilized or even used at all. The TSC concept takes advantage of the chip area that is already available thanks to semiconductor technology. Approximately thirty percent of the heat transmission from the Central processing unit chip occurs along the horizontal heat transfer channel (Stan et al., 2006). The CPU area's TSC, a sizable cold spot, manages the horizontal heat transmission channel. As the TSC core is turned off, the cool TSC lowers the static power. The TSC is also used in conjunction with other DTM methods. The quantity of TSC cores is determined by equation.

TSC illustration

The CPU is fully used for approximately fifty seconds in this segment to demonstrate the TSC strategies indicated in Figure.4B. The OS is aware of the CPU overload. The CPU works slowly to complete tasks. The processor experiences thermal throttling. From an operating system perspective, this CPU usage curve demonstrates CPU congestion brought on by thermal constraints.

As depicted in Figure.4C, the excessive CPU temperature was discovered by the DTM controller. Therefore, the TSC algorithm is carried out by the DTM controller. A TSC core substitutes a hot core at forty seconds into the timeframe. A CPU peak results from the handover between the hot core and the TSC core. However, the CPU speeds up after that peak because the TSC is still relatively cold and runs at a faster frequency. The CPU achieves thermal throttling once more at eighty-six seconds. As a result, CPU congestion occurs once more. As a result, the time it takes for the CPU to degrade from fifty to fifteen seconds when a TSC core is activated during a CPU thermal crisis. According to Figure. 4D, the CPU usage rises when three TSC cores are activated during thermal crises at timelines of twenty-five seconds, forty-five seconds, and eighty-five seconds, correspondingly. Instead of seeing specific CPU peaks, the CPU runs its tasks regularly without congestion. Due to the large number of extra cores on this CPU chip, the DTM controller triggers the necessary TSC during CPU heat emergencies. Therefore, the Central Processing Unit potentially prevents thermal throttling.

 **Tip:** Air-cooled engines rely on the circulation of air directly over heat dissipation fins or hot areas of the engine to cool them in order to keep the engine within operating temperatures.

DTM controller in three-dimensional Fuzzy

The interaction between the various variable of a distributed variable system may be handled by the three-dimensional fuzzy control (Li & Li, 2007). As a result, three-dimensional fuzzy logic may process the correlation data from the Multi-Core Central Processing Unit. The three-dimensional fuzzy control shows how it can be used in various technical applications. The real-world applications are achievable for the three-dimensional fuzzy control (Li & Li, 2007). A distributed variable system underlies the heat management procedure. Nonlinear partial difference equations depict the thermal management procedure (Doumanidis & Fourligkas, 2001).

Simulation results

Simulation has been utilized to verify the designed three-dimensional fuzzy DTM controller. The quantity of published has been used to determine which CPU chip to use. Depending on publicly available data, comprising floor plan, technology, TDP (thermal design power), chip area, and operating frequencies, the IBM POWER CPU family has been chosen. The chosen chip is the IBM POWER4 MCM chip. The POWER4 processors and the MCM's floor plans are available as images. The CPU floor plan and associated power density map are considered to be secret information by all processor makers. As a result, creating a thermal model depending on actual CPU chip data is quite challenging. Only outdated thermal data for CPU chips are published. The power density map and floor plan for the MCM POWER4 are available. The reverse engineering of IBM MCM POWER4 is the only method for creating a CPU thermal model. Reverse engineering was a time- and labor-intensive technique. As the POWER4 chip has been constructed using the outdated ninety nanometers technology, the extracted MCM POWER4 chip has been scaled into forty-five nanometers technology (Sinharoy et al., 2005).

Depending upon the simulator's capabilities and the online assistance offered by the Virginia University Hotspot team, the Virginia Hotspot simulator was chosen. The Hotspot 5 simulator models silicon heat transfer by combining RC circuits with thermal systems. The differential equations governing the operation of the thermal RC circuit are resolved by the Hotspot 5 simulator using Runge-Kutta (fourth-order) numerical approximations (LAVA, 2009).

8.5. Simulation analysis

All simulations begin at 814 seconds since the CPU thermal model needed 814 seconds to achieve $T_{Control}$ seventy degrees Celsius. We will assume that the output response of the CPU will follow the

open-loop curve till it hits seventy degrees Celsius. The DTM controller output decides the fundamental operating frequency when you get to $T_{Control}$. The temperature of each core adjusts itself accordingly based on the frequency at which it is functioning. Without making any alterations to the fuzzy rules, every DTM fuzzy design's tuning is determined solely by tuning its output membership functions (MSF). The DTM assessment index considers simulation times ranging from 814 seconds to 1014 seconds. Such simulation experiments 3D-FC1, FC1, FC2, FC2, 3D-FC3, and FC3 execute DVFS and TSC simultaneously. However, only DVFS is used in these experiments FC4, 3D-FC4, 3D FC5, and 3D-FC6. The temperature and frequency are the only two inputs l= 2 to the DTM regulator assessment index (4). Its intended value is ζ_t = 2 or near 2.

This subsection will look at two different DTM implementations of evaluation indexes. The initial DTM implementation used the assumption that the CPU must operate at its highest frequency for twenty percent of the time, at its specific rate for fifty percent of the time, at its medium frequency, and at its lower frequencies for ten percent of the time. Additionally, the CPU must spend thirty percent of its time operating at high temperatures, forty percent at medium temperatures, and thirty percent at lower temperatures. The 1st DTM demand assessment results are as follows when compared to the DTM controller designs: The amount of time the CPU spends operating at every frequency range is displayed in Table 8.1.

Table 8.1 the comparisons of the frequencies during the initial implementation

Controller Name	Frequency Ranges % σ_{1j}^{Actual}				Frequency Ranges Values σ_{1j}				λ_1
	(M) j=1	(H) j=2	(m) j=3	(L) j=4	(M) j=1	(H) j=2	(m) j=3	(L) j=4	
$\sigma_{1j}^{Desired}$	20%	50%	20%	10%	1.0	1.0	1.0	1.0	1.00
Switch	0%	100%	0%	0%	0	2	0%	0	0.500
P	10%	0%	22%	22%	2.7	0.0	1	2	1.528
FC1	12%	22%	44%	22%	0.5	0.4	2	2.2	1.315
3D-FC1	0%	10%	33%	11%	0.0	1.1	1.7	1.1	0.972
FC2	0%	100%	0%	0%	0.0	2.0	0.0	0.0	0.500
3D-FC2	0%	89%	11%	0%	0.0	1.8	0.6	0.0	0.123
FC3	22%	22%	10%	0%	1.1	0.4	2.8	0.0	1.083
3D-FC3	0%	78%	22%	0%	0.0	1.6	1.1	0.0	0.667
FC4	0%	66%	33%	0%	0.0	1.3	1.7	0.0	0.750
3D-FC4	22%	10%	22%	0%	1.1	1.1	1.1	0.0	0.833
3D-FC5	0%	10%	33%	11%	0.0	1.1	1.7	1.1	0.972
3D-FC6	0%	78%	0%	22%	0.0	1.6	0.0	2.2	0.944

Table 8.2 demonstrates the proportion of time the CPU performs at every temperature range.

Controller Name	Temperature Ranges % σ_{2j}^{Actual}			Temperature Ranges Values σ_{2j}			λ_2
	(H) j=1	(m) j=2	(L) j=3	(H) j=1	(m) j=2	(L) j=3	
$\sigma_{2j}^{Desired}$	30%	40%	30%	1.0	1.0	1.0	1.00
Switch	0.0%	100%	0.0%	0.0	2.5	0.0	0.83
P	78%	0%	22%	2.6	0.0	0.7	1.11
FC1	11%	89%	0%	0.4	2.2	0.0	0.86
3D-FC1	22%	78%	0%	0.7	1.9	0.0	0.90
FC2	67%	33%	0%	2.2	0.8	0.0	1.02
3D-FC2	10%	44%	0%	1.8	1.1	0.0	0.99
FC3	67%	33%	0%	2.2	0.8	0.0	1.02
3D-FC3	33%	67%	0%	1.1	1.7	0.0	0.93
FC4	44%	10%	0%	1.5	1.4	0.0	0.96
3D-FC4	33%	67%	0%	1.1	1.7	0.0	0.93
3D-FC5	0%	100%	0%	0.0	2.5	0.0	0.83
3D-FC6	33%	10%	11%	1.1	1.4	0.4	0.96

Bold is used to highlight the most significant outcomes. The optimum DTM controller designs according to the DTM assessment index were FC3 and 3D-FC6, as depicted in Table 4. Bold is used to highlight the most significant outcomes. The only controllers that perform well in both temperature and frequency assessment criteria are FC3 and 3D-FC6. Both DTM controllers' frequency change responses constantly fluctuate, as seen in Fig. 10A. The frequency and amplitude oscillations of the 3D-FC6 controller are smaller. The FC3 controller runs at its highest frequency for 1014 and 1100 seconds before turning off. The 3D-FC6 controller never shuts off and runs at a higher frequency, though not at the highest possible frequency. Both controllers' temps fluctuate from a temperature perspective. According to Fig. 10B, the 3D-FC6 controller exhibits minimum temperature amplitudes between 970 and 1070 seconds. When compared to the FC3 controller, the 3D-FC6 always runs cooler. A better controller than the FC3 controller is the 3D-FC6 controller.

Summary

The air conditioning module is used in consumer electronics such as autos and communication devices. Consumer electronics are growing more intricate and intelligent, and the transition is happening faster than before. To recall the author's early experience with numerous consumer-electronic goods,

heat/thermal issues have played a vital role during the last two decades. This chapter explores all approaches in the consumer electronics market for Personal Computer (PC), Note Book (NB), Server including a central processing unit (CPU) and, graphic processing unit (GPU), and LED lighting bulbs of smaller areas and higher power. This technique is designed to assist them in making correct, rational, and systematic decisions about the lifetime and reliability of their products.

The authors hope to contribute to the LED industry, government, and academics in the field of green energy-saving lamps. The author's future efforts could be devoted to constructing a system of LED green energy-saving lamps. It is also hoped that the thermal module evaluation approach will be expanded to other categories of consumer LED products, such as home appliances, workplace automation, personal communication devices, automotive interior design, and so on. Finally, the authors highlight a few points as the study's contribution. This can be used as a resource for future studies.

Review Questions

1. What are the functions of a multi-core CPU?
2. What are the limitations of CPU air cooling?
3. Explain the thermal spare core process.