COOLING SYSTEM DESIGNING FOR HPC

RAKOTOARISOA Hasina Patrick[1], RATSIMBA Mamy Nirina[2], RASTEFANO Elisée[3]

[1] *PhD Student, Embedded System – Instrumentation – Electronic System Modeling Laboratory,*
[2] *PhD, Embedded System – Instrumentation – Electronic System Modeling Laboratory,*
[3] *Professor, Embedded System – Instrumentation – Electronic System Modeling Laboratory,*

Doctoral School of Science and Technics of Engineering and Innovation
University of Antananarivo, Madagascar

ABSTRACT

Energy consumption is a big problem in HPC (High Performance Computing). The cooling system plays dual role in HPC. Not only does it consume energy, but also it keeps electronic components in the appropriate functioning conditions. So the design of a cooling system is very important in HPC. The size of the cooling system to be used in an HPC is function of the clock frequency of the CPU (Central Point Unit). From this expression, we can have the TDP (Thermal Design Power) of one processor and also that of a HPC. The correct temperature of HPC cooling can be obtained from TDP.

Keyword: *TDP (Thermal Design Power), HPC (High Performance Computing), design of cooling system, clock frequency, CPU, processor*

1. INTRODUCTION

Currently, many researches are done towards building exaflopic machine, a machine that can execute 10^{18} Flops (Floating-point Operations per second). One major obstacles in achieving this goal is energy consumption because according to DARPA (Defense Advanced Research Projects Agency), the exaflopic machine can't exceed energy consumption of 20 MW [1]. Much of the energy consumption is due to the heating of the components and its cooling systems. This paper investigates the design of an adequate cooling system for HPC. It is organized as follows. Section 2 presents related work on the cooling system design. Then, the methodology used to obtain results is exposed. The tests and results are then shown. The last section is devoted to conclusion.

2. RELATED WORK

Essentially, microprocessors (CPU), memories, network and storage devices are the main part of an HPC. Microprocessors are the basis of an HPC because they do the calculation and they are the element that consume considerable energy as dissipated heat. Yewan has studied the impact of temperature on the power of a machine [2]. In this work, he demonstrated that the temperature of the microprocessor can induce a significant variation in the power of the servers. Thus by increasing the temperature of the microprocessor, the power of the server gradually decreases. During the experiments, he used high-power processors.

Several works were also done, and several tools were used to measure the temperature of a microprocessor for different usages. The goal of his studies is to calculate the proper cooling system for the machine. But until now, there is no precise formula for doing this calculation. This article gives then an example of a cooling system calculation for HPC. Moreover, one of the solutions to optimize the use of a HPC is the control of the consumed energy.

3. METHODOLOGY

For our investigation, an HPC dedicated for making intensive computing is shown. So high-power processors have been used. An HPC is essentially composed of microprocessors. And in Shuaiwen's investigation of the power dissipated by HPC components, power dissipated by other components has been neglected [3]. In his paper, microprocessors are responsible for heating HPC. Thus, the studies focus on the heat dissipation of microprocessors. Then, we multiply the result obtained by the number of microprocessors in the HPC.

In order to do study a cooling system of a microprocessor, it is necessary to know its characteristics. The manufacturers' data were therefore taken and adopted. In the specifications of each microprocessor, there are several values that determine its characteristics. Among them, there is the TDP (Thermal Design Power) value which is the dissipated power of the integrated circuits. Above this value, the operation of the integrated circuits is no longer guaranteed.

To calculate the TDP, we have the Matrix Formula which is written:

$$TDP = f * V_{Core}^2 * k \quad (1.1)$$

where f is the clock frequency, V_{core} *the voltage of the core microprocessor and k a specific constant for each microprocessor model*

This TDP value can also be calculated from the temperature of the work environment. To do this, the technology used in the microprocessor must be shown. Most of the microprocessors used in HPC are manufactured using CMOS technology. So that formula from the study of the CMOS transistor is to be used:

$$TDP = \frac{T_{case} - T_{ambient}}{R_{th}} \quad (1.2)$$

where T_{case} *is the temperature of the microprocessor body,* $T_{ambient}$ *the ambient temperature of the machine and* R_{th} *the thermal resistance of the microprocessor.*

According to equation 1.1, the TDP is function of the clock frequency of the microprocessors. Thus, the more the frequency of the microprocessor is increased, the more it heats up and the higher the value of this TDP can be reached.

To reduce the heating of the microprocessor, it is necessary to put a cooling system. The sizing of this system is very important and essential because in HPC, there are thousands of microprocessors that work at high frequencies and need to be cooled off.

The idea in this paper is to calculate the size of the cooling system adequate for a microprocessor and for HPC.

4. TEST AND RESULT

For the test, two HPC were selected. The first HPC is named "Curie" and the second is called MiHawk.

4.1 Curie Supercomputer in France

The Curie Supercomputer, owned by GENCI and operated in the TGCC (Très Grand Centre de Calcul) by CEA (Commission de l'Energie Atomique), is the first French Tier0 system open to scientists through the French participation into the PRACE (Partnership for Advanced Computing in Europe) research infrastructure. We took the "Curie Thin Nodes" with 10080 microprocessors Intel Xeon E5-2680 -8 cores - 2,700 GHz [4] [5].

During our work, the following parameters (data coming from the manufacturer) are taken:

$$T_{case} = 85\ °C\ ,\ T_{ambient} = 40\ °C$$

To calculate *k* the specific constant of a processor, the expression in the following equation is used:

$$k = \frac{Typical\ power\ consumption}{(f_{base} * V_{core}^2)}$$

According to the processor characteristics given in reference [5], the typical power consumption for Intel Xeon E5-2680 is 105.63 W. The clock frequency f_{base} and V_{core} values are respectively 2.7 GHz and 1.35 V. So, the specific constant for the Intel E5-2680 8-core 2,700 GHz processor is **0.0214**. From this value, we have the equation of the TDP with respect to the frequency of the microprocessor:

$$TDP = 0{,}03912 * f$$

Figure 1 shows the relationship between the TDP and the clock frequency of the Intel Xeon E5-2680 microprocessor.

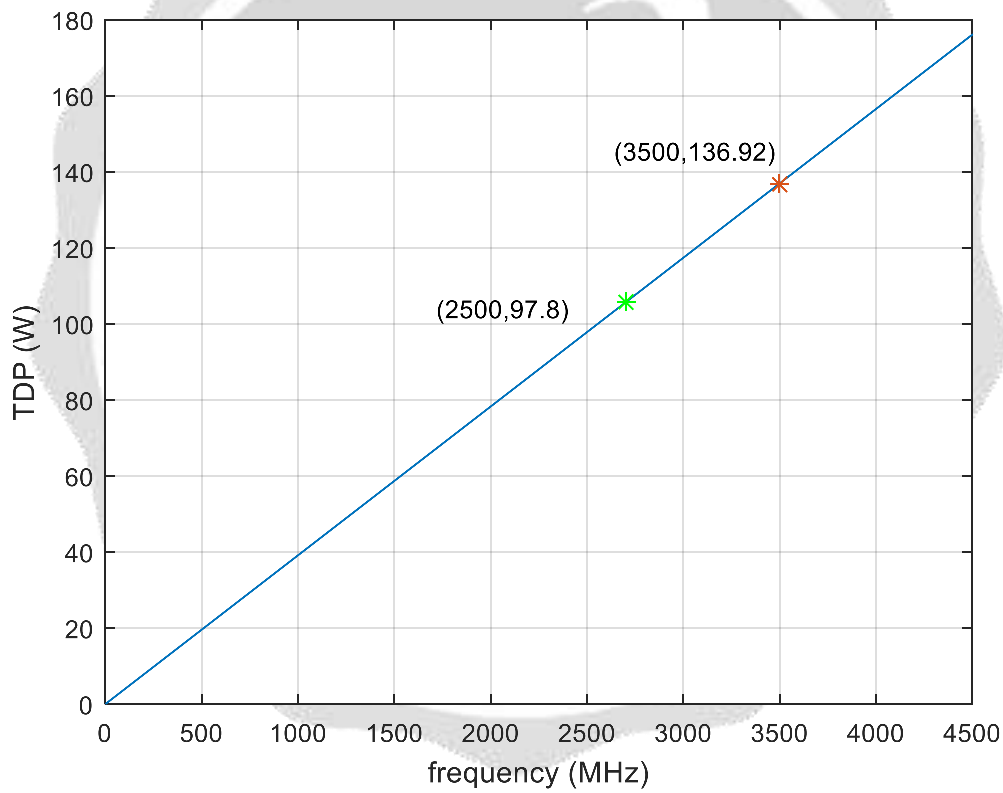

Fig -1: Relationship between TDP and clock frequency for Intel Xeon E5-2680 microprocessor.

The curve makes it possible to calculate the dissipated power of the microprocessor with respect to the frequency. The value indicated in green is the basic clock frequency and the one in red is the maximum frequency that the microprocessors can support. Thus, the thermal resistance of the Intel Xeon E5-2680 8c 2.7 GHz is:

$$R_{th} = 0{,}328\ °C/W$$

Then, for Curie Thin Nodes Supercomputer, we are the global TDP:

$$TDP_{Curie\ Thin\ Nodes} = 10080 * 136{,}92 = 1.380.153{,}6\ W\ or\ 1{,}38\ MW$$

A cooling system capable of evacuating 1.38 MW is therefore necessary for the proper functioning of the Curie Supercomputer. Here are some cooling systems that can be used: air cooling, watercooling, adiabatic cooling ... Figure 2 represents an adiabatic cooling system for Curie.

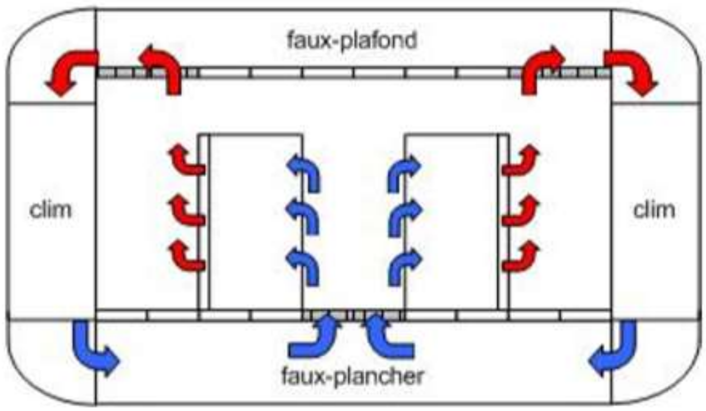

Fig -2: Adiabatic cooling system for Curie.

4.2 MiHawk server in USA

P93D2-2P (MiHawk) is a 2U2S Power9 server made by Wistron Corporation. MiHawk have two Power9 LaGrange processors in it. Because of the high-speed Input/Output (I/O) and flexibility features, MiHawk has the capability of fulfilling any kind of intensive work such as Artificial Intelligence, Big Data, and HPC [6].

The manufacturer's data on the processor Power9 LaGrange are:

$$T_{case} = 85\ °C\ ,\ T_{ambient} = 40\ °C$$

To calculate the *k* which is specific constant of the IBM Power9 Lagrange, we use also the expression:

$$k = \frac{Typical\ power\ consumption}{(f_{base} * V_{core}^2)}$$

From Bo’s work [7], the typical power consumption of Power9 Lagrange is 117,7 W. The others values are: f_{base} = 2,75 GHz et $V_{core} =$ 1,2 V. The specific constant *k* for IBM Power9 LGrange is **0,0297**. Then, we are the TDP formula:

$$TDP = 0{,}0427 * f$$

The curve in Figure 3 then gives the variation of TDP as a function of frequency for the IBM Power9 LaGrange processor.

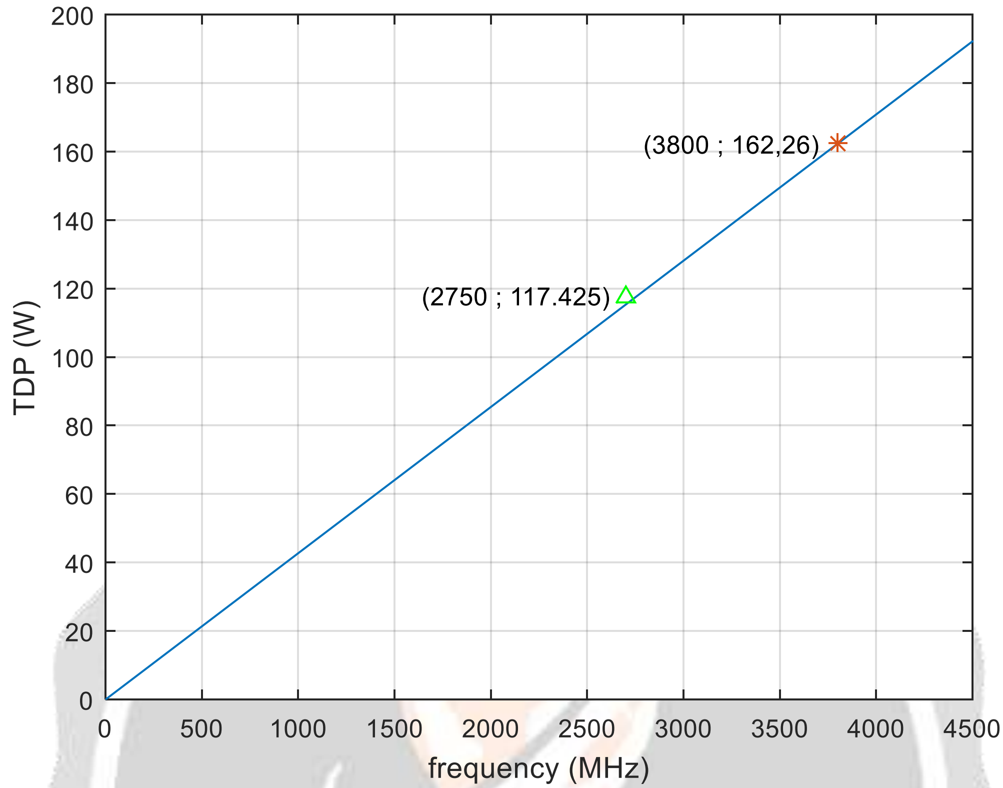

Fig -3: Relationship between TDP and the clock frequency of IBM Power9 LaGrange processor.

The thermal resistance for the Power9 LaGrange is:

$$R_{th} = 0{,}237\ °C/W$$

For MiHawk server, the dissipated power are:

$$TDP_{MiHawk} = 2 * 162{,}26 = 324{,}52\ W$$

So MiHawk server needs air cooling for example to evacuate 324,52 W of dissipated power.

5. CONCLUSION

According to the curves of Fig-1 and Fig-3, the TDP of a microprocessor is therefore a function of the clock frequency. While playing with the value of the frequency of the microprocessors, one can calculate the dimension of the adequate cooling system configuration for the machine. One perspective is to study the arrangement of microprocessors on the HPC. Proper management of the cooling system can optimize the use of an HPC and its energy consumption. Because uncontrolled heating can lead to a rapid shortening of microprocessors lifetime.

6. REFERENCES

[1]. Bergman K., Borkar S., Campbell D., and others (2008). ExaScale Computing Study : Technology Challenges in Achieving Exascale Systems. In DARPA Information Processing Techniques Office, page pp. 278, Washington, DC.

[2]. Yewan W., David N., Stéphane L. M., Jean-Marc M. « Etude de l'influence de la température du processeur sur la consommation des serveurs » Compas'12018 ENSEEIHT – Toulouse, France

[3]. Shuaiwen S., Rong G., Xizhou F., Kirk W. C. “Energy Profiling and Analysis of the HPC Challenge Benchmarks”
[4]. http://www-hpc.cea.fr/en/complexe/tgcc-curie.htm, 15/01/2019
[5]. http://cpuboss.com/cpu/Intel-Xeon-E5-2680, 15/01/2019
[6]. https://openpowerfoundation.org/?resource_lib=wistron-corp-p93d2-2p-mihawk, 15/01/2019
[7]. Bo L., Edgar A. L., Kirk W. C. “Understanding Power Measurement Capabilities on Zaius Power9”, INFOCOMP 2018 : The Eighth International Conference on Advanced Communications and Computation