# Analyzing Temperature Data in a High-Performance Computing Cluster

**Jarilyn M. Hernández Jiménez**
**Polytechnic University of Puerto Rico**
**Research Alliance in Math and Science**
**Computer Science and Mathematics Division**
**Oak Ridge National Laboratory**
**Mentors: Dr. Line Pouchard and Dr. Houssain Kettani**
**http://sites.google.com/a/g.ornl.gov/jarilyn-hernandez/**

## Abstract

Clusters are being used by scientists and engineers to process data-intensive applications due to their multiprocessing capability. Processors in a high-performance environment can be in the order of hundreds of thousands and, due to their complexity, the entire system is prone to failure. The aim of this project is to analyze temperature data in order to establish a corresponding signature of the failure of a processor and to validate that the CPU temperatures behave as they are supposed to when they are strained. This project focused on the temperature of processors and CPU usage. Accordingly, as the processor gets hotter, it may be necessary to avoid running further jobs on it so that it does not fail. This is intended to establish a criteria for detecting a failing processor before migrating the job to a healthier processor.

## Clusters

A cluster is defined as a group of computers working all together as one. Clusters use MPI to communicate. MPI is defined as a specification of an API that allows many computers to communicate with one another. The goal of the MPI is to provide high performance, scalability and portability. The cluster used to gather the data uses MPI to interact with the nodes. The word node refers to the computer located in each blade of the cluster. Each computer is a HP Proliant DL385 G5. The cluster used to gather the temperature data is called Trident.

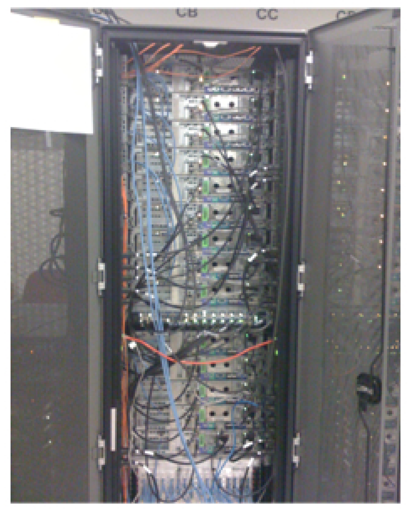

**Figure 1. Trident cluster cabinet with door open to show processors.**

## Trident

Trident is a 32-node cluster with HP motherboards. Each node has two quad-cores AMD Opteron 2356 processors and 16GB of RAM. Every node contains sensors to measure the temperature. Trident 29 was the node selected to be analyzed, because this node is the one used by the scheduler.

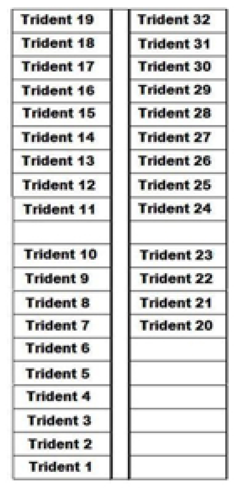


**Figure 2. Distribution of the racks in Trident.**

## Data Collection

The variables taken in consideration for this research are: CPU usage time, temp2 and temp3. Only CPU temperature is being analyzed, and that is why temp2 and temp3 were selected. The temperature unit is in Celsius and the unit for the CPU usage time is in *jiffys*. A jiffy is the duration of one tick of the system timer interrupt. The tool used to gather the data is called DCAT. DCAT is an open source tool developed at ORNL that gathers sensor data in a high-performance computing environment. It was also necessary to use a subroutine in the basic linear algebra subprogram to heat up the system and to validate that the CPU temperatures behave as they are supposed to do. This subroutine is called DGEMM.

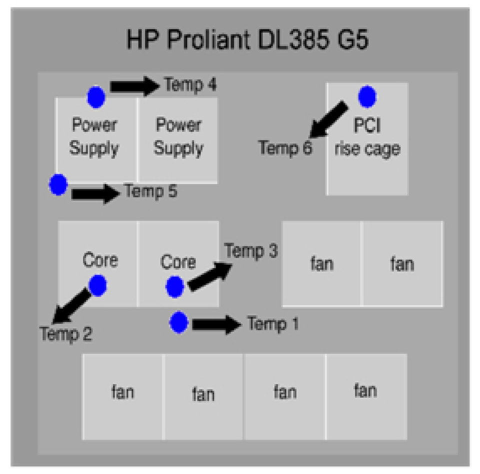


**Figure 3. Physical diagram of the HP Proliant computer.**

## Results

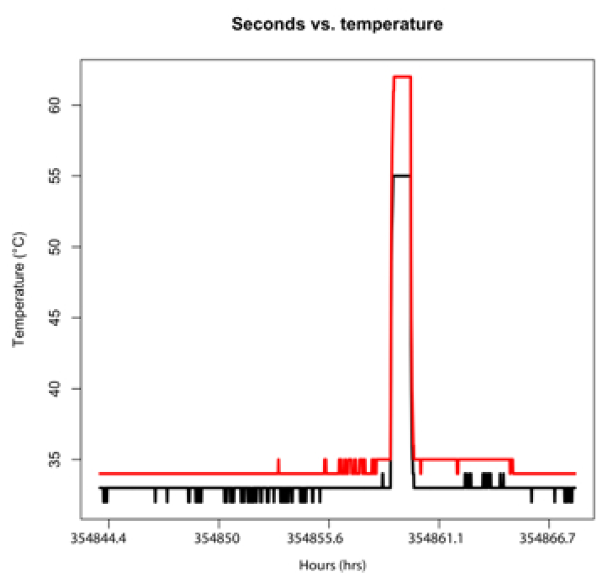


The tool used to display and analyze the data is called R. R is a free software environment for statistical computing and graphics. The graph to the left illustrates a time-series for Temp2 and Temp3 at the time of the experiment.

**Figure 4. Difference between temp2 and temp3.**

**Figure 5. Jiffys vs. temperature for sensor temp2.** **Figure 6. Jiffys vs. temperature for sensor temp3.**

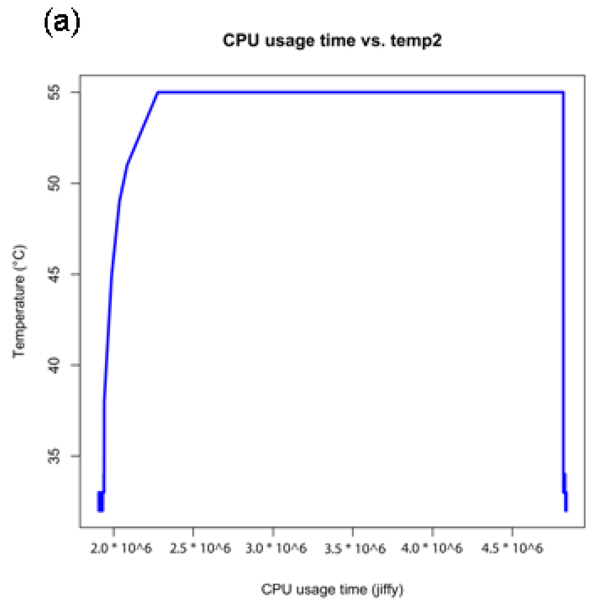


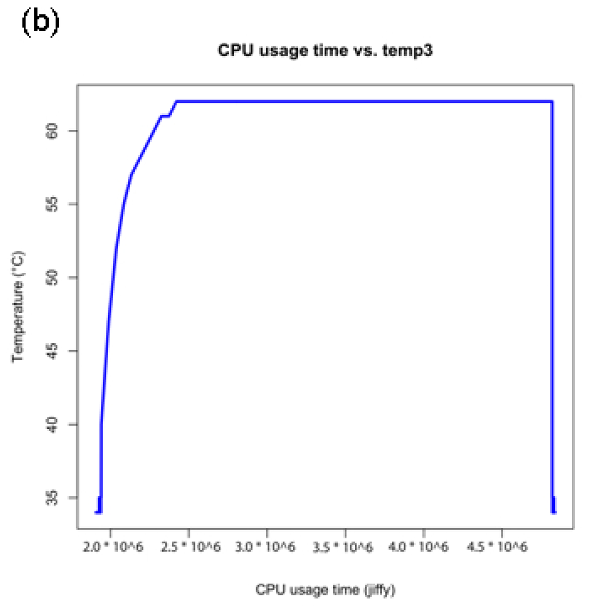


## Conclusions

After analyzing the data, it was concluded that once a CPU intensive code starts running in both CPU cores, the temperature will start to rise. At idle, baseline temperatures are as follows:

- ☑ 32C given by sensor Temp2
- ☑ 34C given by sensor Temp3

As the process continues to run, the temperature in both cores starts to rise until it reaches:

- ☑ 55C given by sensor Temp2
- ☑ 62C given by sensor Temp3

These temperatures remained constant for about 2 hours until the process was terminated. At this time, the temperature returned to its initial state. From this data, it can be concluded that if a similar process is running for long enough, the temperature of the cores may rise to a point where it may cause a failure.

## Future Research

Detecting anomalies aids to failure identification and eventually to the development of a failure prediction capability. Analyzing temperature data can help us understand one of many factors that affect the performance of clusters and supercomputers. The data of this project can be used to implement mechanisms that allows predicting when a failure will occur. This mechanism will help in the prevention of loss of data since before the failure occurs, the data can be migrated to a healthier processor or the job can be shared with other cores to minimize the load on the failing core.

## References

L. Pouchard, J. Dobson, S. Poole, "Collecting Sensor Data for High Performance Computing: A Case Study," 2009.
L. Pouchard, S. Poole, J. Lothian, "Open Standards for Sensor Information Processing, ORNL/TM-2009/145," July 2009.
F. Capello, "Fault Tolerance in Petascale/Exascale Systems: current knowledge, challenges and research opportunities, The International Journal of High Performance Computing Applications," Vol. 23, No. 3, Fall 2009, pp. 212-226.

U.S. DEPARTMENT OF ENERGY
Office of Science

The Research Alliance in Math and Science program is sponsored by the Office of Advanced Scientific Computing Research, U.S. Department of Energy. The work was performed at the Oak Ridge National Laboratory, which is managed by UT-Battelle, LLC under Contract No. De-AC05-00OR22725.
This work has been authored by a contractor of the U.S. Government, accordingly, the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or allow others to do so, for U.S. Government purposes.
I would like to thank Dr. Line Pouchard for the opportunity to work on this project. Also, I would like to thank Dr. Houssain Kettani for his continued support. Finally, special thanks go to Mrs. Debbie McCoy, who made provisions for this research experience.

POLYTECHNIC UNIVERSITY PUERTO RICO
UT-BATTELLE
Oak Ridge National Laboratory