

LA-UR-15-24365

Approved for public release; distribution is unlimited.

Title: Experimental Design and Comparative Testing of a Hybrid-Cooled
Computer Cluster

Author(s): Bonnie, Amanda Marie

Intended for: Masters Thesis

Issued: 2015-07-06 (rev.2)

Disclaimer:

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the Los Alamos National Security, LLC for the National Nuclear Security Administration of the U.S. Department of Energy under contract DE-AC52-06NA25396. By approving this article, the publisher recognizes that the U.S. Government retains nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

Amanda Bonnie

Candidate

Electrical and Computer Engineering

Department

This thesis is approved, and it is acceptable in quality and form for publication:


Approved by the Thesis Committee:

PAYMAN ZARKESH—HA, chairperson

JAMES PLUSQUELLIC

TARIEF ELSHAFIEY

DEPARTMENT OF ELECTRICAL AND COMPUTER ENGINEERING



SCHOOL OF ENGINEERING UNIVERSITY OF NEW MEXICO

Experimental Design and Comparative Testing of a Hybrid-Cooled Computer Cluster

A Thesis Presented to
the Graduate School of the University of New Mexico
in Partial Fulfillment of the Requirements
for the Degree Master of Science in Computer Engineering

by

Amanda M. Bonnie ¹
Department of Electrical and Computer Engineering
The University of New Mexico
Albuquerque, NM 87131
e-mail: abonnie@unm.edu

~

Bachelors of Science in Mechanical Engineering
Clemson University
Clemson, SC 29631

Graduation: July 15th 2015
LANL LA-UR-15-24365

Report Date: July 5, 2015

¹Supported by Los Alamos National Laboratory.

Dedication

To my loving husband, Dave, who introduced me to these evil computers...

errr... I mean, for his incredible patience and unconditional love and support over the course of my research, and all of our days together. I love you.

Acknowledgments

This work was performed using facilities and resources at Los Alamos National Laboratory and was funded by the United States Department of Defense.

The author acknowledges her direct management, Jeffery Kuehn, Daryl Grunau, Parks Fields, and Craig Idler, for allowing time to complete this work, as well as all the support and efforts provided to make it all happen.

The author would also like to acknowledge Gary Grider for all his guided support and random discussions in bettering this study; as well as the motivation to complete a degree in a year and a half.

The author thanks Peter Deangelis, Steve Harrington, Chris Nykamp, and others at Chilldyne Inc, for their support and efforts that allowed for this cluster to be retrofit with water.

The author thanks Ron Velarde, Cindy Martin, and the LANL HPC Facilities team for their collaboration in allowing this cooling system to be installed in the 341 Data Center.

The author thanks Bob Villa for his efforts in helping with the preliminary install with the Chilldyne vendors. His effort allowed us to complete the install in two days, despite the hiccups.

The author thanks Al Guerra for taking time out of his day to provide the fantastic photographic evidence of this project.

The author would like to thank all the various persons at LANL from HPC who were available at strange hours, and willing to give up their personal time to assist this research. Far too many of you to be named, but you know who you are!

The author thanks her committee Dr. Payman Zarkesh-Ha, Dr. James Plusquellic, and Professor Tarief Elshafiey for their assistance.

The author again thanks her husband, David Bonnie, for his continued support of this crazy adventure of study.

Thank You All!

Abstract

With water cooling becoming an affordable option both at home and at scale, it is important to consider the possible benefits over air cooling. There are several methods of liquid cooling, notables include: immersion, cold water cooling, and warm water cooling. The total cost of ownership is difficult to determine with these options as each has a different impact on the data center. Considering retrofit, over a new data center, introduces unforeseen variables that make cost analysis a challenge. Besides the added costs of additional infrastructure, and the cost to remove old, the upfront costs could be daunting. Therefore a cost analysis would be a study of its won. This study however hopes to reveal the resulting tradeoffs in temperature, performance, and power usage presented in the case between classical airflow based heat sink mechanisms to water provided directly at the heatsink. Having control over a discrete chiller will provide answers to the CPU temperatures, power usage, and performance at various inlet water temperatures. To water or to air?

Keywords

water cooling, power savings, varied temperatures, performance changes, water vs. air cooling, data center, compute cluster, CPU throttling

Contents

List of Figures	viii
List of Tables	viii
1 Introduction	1
1.1 Contributions	2
2 Background	4
2.1 Water Cooling Is Not A Novel Idea: Immersion & Other Liquids	4
2.1.1 CRAY-2	5
2.1.2 Green Revolution Cooling	5
2.2 Current Examples of Water Cooling	6
2.2.1 Gamers to Clusters	6
2.2.2 NREL: Warm Water Cooling	6
2.2.3 Top500 with Warm Water Cooling	7
2.2.4 Performance Gains with Water Cooling	7
3 Theoretical Backbone	8
3.1 Costs: Running the Water Cooling System	8
3.2 Reducing Power Use: Fans Small but Notable	10
3.3 Electronic Reliability and Longevity	11
3.4 Jitter and Noise	13
4 Testbed Overview	14
4.1 Full System Specs	15
4.1.1 Software	15
4.1.2 Hardware	16

4.2	Chilldyne Setup	18
4.2.1	Standalone Chiller	19
4.2.2	Cooler Distribution Units & Vacuum Pumps	20
4.2.3	Rack Mainfolds	21
4.2.4	Waterblocks	22
4.2.5	Chilldyne Monitoring & Failsafe Features	25
4.2.6	Waterblock Installation	26
5	Testing Setup	28
5.1	Testing Configurations	28
5.2	Stress Testing Nodes	28
5.2.1	Pavillon	29
5.2.2	HPL	29
5.2.3	Systemburn	30
5.2.4	Running the Tests	31
5.3	Monitoring	32
5.3.1	LDMS	33
5.3.2	Temperature: Im_sensors & LDMS	33
5.3.3	Power: Power Distribution Units	33
5.3.4	Power: Running Average Power Limit	34
5.3.5	CPU Clock Speed	35
6	Results	36
6.1	Performance	36
6.1.1	HPL	36
6.1.2	DGEMM	37
6.1.3	DSTREAM	37
6.1.4	PV2	37
6.2	Temperature	37
6.2.1	HPL	38
6.2.2	DGEMM	39
6.2.3	DSTREAM	39
6.2.4	PV2	39

6.3	RAPL	40
6.3.1	HPL	40
6.3.2	DGEMM	41
6.3.3	DSTREAM	41
6.3.4	PV2	41
6.4	Power from PDU	42
6.4.1	HPL	42
6.4.2	DGEMM	42
6.4.3	DSTREAM	43
6.4.4	PV2	43
6.4.5	PDU Data Summary	45
7	Discussion	46
7.1	Performance	46
7.2	Temperature	47
7.3	RAPL	48
7.4	PDU	48
7.5	BONUS: Additional Discovery on Air Cooled Nodes	49
8	Conclusions	52
9	Future Work	55
9.1	Warmer Water Cooling	55
9.2	Tightly Coupled Applications	56
9.3	Looking at Scale	56
A	Chiller Unit: ThermoFlex 7500 Spec Sheet	57
B	Pavillon Test Configuration File Example	58
C	HPL.dat	59
D	Systemburn Load Files	60

List of Figures

1	TAMIRS rack layout diagram.	15
2	Picture Showing TAMIRS Socket Population [31]	17
3	Image of the CDUs with vacuum pumps (left), and standalone chiller (right).	19
4	Image of the manifolds mounted in rack. Orange and black rubber caps block off the stub outs not in use.	21
5	Chris from Chilldyne measuring the tubing for rack to rack connections.	22
6	Chilldyne, Inc. designed water blocks (left) to replace the R920 heatsink (right)	23
7	Final water block design for the Dell R920.	24
8	Left: Internal portion of the hot-swap connection. Right: External portion of the hot-swap connection to allow a node to be serviced while the water system is live [31].	25
9	Screen captures showing available data from the CDU.	26
10	Thermal paste application on the CPU for the water block	27
11	Temperature plots during HPL; air cooled (left), 65°F water (center), 75°F (right).	38
12	Temperature plots during DGEMM; air cooled (left), 65°F water (center), 75°F (right).	39
13	Temperature plots during DSTREAM; air cooled (left), 65°F water (center), 75°F (right).	39
14	Temperature plots during PV2; air cooled (left), 65°F water (center), 75°F (right).	39
15	RAPL plots during HPL; air cooled (left), 65°F water (center), 75°F (right).	40
16	RAPL plots during DGEMM; air cooled (left), 65°F water (center), 75°F (right).	41
17	RAPL plots during DSTREAM; air cooled (left), 65°F water (center), 75°F (right).	41
18	RAPL plots during PV2; air cooled (left), 65°F water (center), 75°F (right).	41

19 PDU plots during HPL; air cooled (left), 65°F water (center), 75°F (right). . 42

20 PDU plots during DGEMM; air cooled (left), 65°F water (center), 75°F (right). 43

21 PDU plots during DSTREAM; air cooled (left), 65°F water (center), 75°F (right). 43

22 PDU plots during PV2; air cooled (left), 65°F water (center), 75°F (right). . 44

23 HPL Test run showing per CPU temperatures. 51

List of Tables

I	Intel Xeon Processor E7-4870v2 Specifications [29]	16
II	80 PLUS Certification 115V Internal Non-Redundant [32]	18
III	Systemburn module summary, note the “size” parameter described will be used in the LOAD configuration file [36].	31
IV	Run times for each test suite, recall 24 hours = 1440 minutes	32
V	HPL test results, given as average values for, minimum, mean, and maximum.	36
VI	Systemburn:DGEMM test results, given as average values for, minimum, mean, and maximum.	37
VII	Systemburn:DSTREAM test results, given as average values for, minimum, mean, and maximum.	37
VIII	Systemburn:PV2 test results, given as average values for, minimum, mean, and maximum.	37
IX	PDU Mean Power Data Summary.	45

Glossary

CDU	Cooler Distribution Unit; A closed loop system that dissipates the return heat load from the system side.
DOE	Department of Energy.
drain	State of a node provided by SLURM indicating that the node reached critical temperature during a scheduled job. The node was throttled and able to complete the job, but was put in a “drain” state so it could not be allocated by the scheduler until reviewed by an admin.
down	State of a node provided by SLURM indicating that the node is not on. It is either powered off, still at POST, or possibly in an unknown state to the system. This notifies the admin that there may be a problem if this status was not expected.
HPC	High Performance Computing —
HPL	High Performance Linpack; A benchmark implementation popular in HPC which solves a uniformly random system of linear equations and reports time and a floating-point execution rate generally in GFLOP.
LDMS	Lightweight Distributed Metric Service.
node	A single compute component of a cluster. A cluster is often comprised of multiple nodes, which may house several CPUs, PCI cards, and RAM.
PDU	Power Distribution Unit; used for providing power throughout a data center and or cluster.
PUE	Power Usage Effectiveness; a measurement of how efficiently a data center uses energy. It is specifically the ratio of Total Facility Energy to IT Equipment Energy.
systemburn	A software package developed for methodically created load testing.
SLURM	Simple Linux Utility for Resource Management; An open source, fault-tolerant, highly scalable cluster management and job scheduling system for linux clusters.
TAMIRS	Tiered Active Multi-dimensional Indexed Record Store; A compute cluster at LANL.
TDP	Thermal Design Power; represents the average power the processor dissipates when operating at be frequency with all cores active under an intel-defined, high-complexity workload.
thermal paste	A paste like substance that improves the contact between the CPU and water block filling any voids in the surfaces, thus improving thermal conductivity and promoting heat removal.
water block	A heatsink designed to have water channels run through for additional cooling beyond the standard copper heatsink. Often are still made from copper, with the addition of water passages in and out of the block.

Chapter 1

Introduction

The driving factors for this research are: 1. HPC is growing, and the race to exascale is here; the machine room and/or data center size is expanding; 2. Cluster density is increasing producing more to cool; 3. The drive to maintain a low PUE, mandated by the DOE; 4. Total cost of ownership concerns, as cooling data centers is expensive; nearly 30% of a data center electricity bill is spent on cooling [1].

There are many facets of water cooling, with at least two at odds with one another. One is driven by resiliency. Cooler electronic components last longer, so keeping them colder will increase their lifespan [2]. A thought is that cooler components, such as the CPU, can also produce better results through maintaining higher clock speeds longer [3].

Another means of water cooling is warm water cooling. Warm water cooling allows for warm (usually around 45°C) to be delivered via direct to chip methods. This means less work needs to be done by chillers and cooling towers to cool the water down. In some cases it can provide what is known as “free cooling” in which no work needs to be done to cool the water going into a system [1]. “Hot-water” cooling also falls in this “free cooling” regime at much higher inlet temperatures [4]. Some means of warm water cooling also use the waste heat from the computer side water loop, also known as the process loop, to effectively return a heat

source for heating buildings and or water.

Saving energy is a desirable outcome from water cooling, but running more, and or longer, cycles is desirable as well. One must rationally consider this tradeoff based on individual driving factors of a clusters purchase. These factors can include, the workload of the cluster, the concerns for down time, the expected lifetime the machine must exist, and the environment of the current data center it will live in.

Power Usage Effectiveness (PUE) is a term coined by The Green Grid [5] used to measure how efficiently a data center uses energy. The value is given by a ratio of the Total Facility Energy to the IT Equipment Energy [5]. PUE is often a value for argument as it is often difficult to accurately determine. Human error such as not remembering to include the computer room lighting, as well as the difficulty to obtain specific power measurements for IT or facility components, are two common issues associated with determining PUE. Nonetheless, the Department of Energy still has a mandate for all federal data centers to be at or below 1.4 [6].

1.1 Contributions

This work is comprised of the development of user-space LDMS daemons; the deployment, configuration, and support of the TAMIRS cluster at LANL; contract management and installation of the Chilldyne hybrid water cooling system; and the integration of a test suite for monitoring and benchmarking the system for comparative analysis. This Thesis is the culmination of over 1.5 years of effort on the part of the author in preparation for this analysis.

The novelty of this research was perhaps not realized until measured results were achieved. One of the important concepts of applying High Performance Computing (HPC) to physical simulations is the requirement for tightly coupled applications in which the slowest core brings down the performance and increases time to completion. This reduction in variation, known as “jitter”, could significantly increase the performance of typical HPC jobs. The larger the

scale, the larger the jitter; therefore wasting more machine time, and more power.

Chapter 2

Background

Liquid cooling is not a novel concept even in the world of HPC. Early systems such as the CRAY-2 even used an immersion technique, utilizing liquid to cool components to enable faster computations in a densely packed system. Even today liquid cooling is being used to remove heat from systems as well as provide unique heat sources for building heating. Liquid cooling is by no means a novel concept, and this research does not try to make it be.

2.1 Water Cooling Is Not A Novel Idea: Immersion & Other Liquids

Immersion of clusters is neither a new idea or one that has been completely abolished either. The CRAY-2 was immersion cooled along with other machines since the 1980's [7]. Several factors besides the concern of safety and costs kept the immersion idea from flourishing. The development of personal computing delayed the growth of highly dense computing and the design of microprocessors that could be air cooled took things into the massively parallel and beyond era. The demise of Moore's Law as well as the breakdown of MOSFET scaling drove the concepts of scaling frequency and scale to new architectures [8], leaving immersion as a less needed solution. However, as the increase in density comes again, Green Revolution

cooling has a white paper out on its new “ElectroSafe” non-proprietary mineral oil blend [7] which may provide resolution to these old concerns if it becomes a solution once again.

2.1.1 CRAY-2

One of the first computers to be cooled with fluid was the CRAY-2 in 1985 [9]. The CRAY-2 used liquid immersion cooling. The densely packed components resulting in shorter signal paths and higher speeds resulted in increased operating temperatures. This limitation was alleviated through the use of liquid immersion cooling. It placed the cooling medium in direct contact with the components to be cooled [10]. The mainframe operated in a cabinet filled with a colorless, odorless, inert fluorocarbon fluid. It’s high thermal stability, good heat transfer properties along with it being non-toxic, and nonflammable made the liquid safe to work with while providing a cooling benefit [10]. Coolant would flow through the module circuit boards at one inch per second and had direct contact with the integrated circuit packages and power supplies [10]. The valveless system of 200 gallons operated at room temperature cooling ranges, providing enough heat removal to remove the temperature limitations for faster speeds [10].

2.1.2 Green Revolution Cooling

As density increases and a desire for better cooling continues to be demanded, past methods for cooling seem to be resurfacing. The idea of immersion, like the CRAY-2, is being redeveloped with “ElectroSafe” a non-toxic, clear, odorless, dielectric mineral oil blend [7]. Being a readily available, low cost, liquid allows for it to show up in today’s data centers. Besides Green Revolution Cooling in-house testing, Green Revolution Cooling has immersion installations at Intel, the Department of Defense, the Texas Advanced Computing Center (TACC), and even at the National Security Agency [11].

2.2 Current Examples of Water Cooling

2.2.1 Gamers to Clusters

Asetek is an Original Equipment Manufacturer (OEM) for many consumer electronics companies. It provides “all-in-one” water cooling systems for desktop computers and workstations [12]. These solutions allow for increased thermal density but do nothing to reduce the thermal load of the system. Similarly, they have begun to offer data-center/cluster-centric solutions. These solutions again ease the problems with cooling systems with high thermal density and their most recent solution follows the “free cooling” concept [13]. Unfortunately, these “drop in” replacements still require specific space tolerances for deployment as they are often used as a retrofit option to OEM, which generally has no need for leaving enough space for such components.

2.2.2 NREL: Warm Water Cooling

The U.S. Department of Energy’s National Renewable Energy Laboratory (NREL) has been awarded 52 R&D 100 Awards for its groundbreaking work in HPC [14]. When NREL was designing its new data center, the Energy System Integration Facility (ESIF), ideas for maintaining NREL’s mission of being a living laboratory for energy efficiency and sustainability were of high interest [15]. The concept of Peregrine was thus born, a means of using waste-heat from an HPC cluster to provide office space heating.

Peregrine is NREL’s flagship warm water cooled cluster. Built from 6,912 Intel Xeon E5-2670 Sandy Bridge processors and 24,192 Xeon E5-2695v2 Ivy Bridge processors, Peregrine is capable of 1.19 PetaFLOP performance [16]. The key element of Peregrine is its warm water cooled design. Hewlett Packard (HP) aggressively bid on the project as the HP Apollo system was already under development in house [15]. The targeted availability for the HP Apollo 8000 system was originally a year out from NRELs proposed deadline, HP accelerated

the program and was awarded the contract for Peregrine [15].

The final delivery of Peregrine in August of 2013 [17] was a full 11 IT racks and six CDUs. Warm water (75°F) is used to operate the servers while the waste-heat comprised of the 95°F water exiting the system, which would normally need to be cooled back down for use again, is recovered as a primary heat source for the ESIF's office space and laboratory areas [14]. It provides a sufficient heat source to meet the needs of the ESIF at 182,500 square feet [15]. During the hotter summer months, the warm water is fed to the loop for the cooling system that lowers the buildings temperature (via evaporative cooling towers) [14]. Because of this, NREL combined with the energy efficient data center, saves about \$1 million a year in energy cost, and consume about 74% less energy than the national average for office buildings [15]. With the help of Peregrine, the ESIF is able to achieve an annualized average Power Usage Effectiveness (PUE) rating of 1.06 or better [17].

2.2.3 Top500 with Warm Water Cooling

Lawrence Livermore National Laboratory (LLNL) reached the top of the TOP500 in June of 2012 with Sequoia [18], a warm water cooled cluster running at more than 16 petaflops, and was 2.49 times more energy efficient than the next fastest super computer in the world. It used only 7,890kW to reach an Rpeak value of 20,132 TFlop/s nearly doubling the performance of second in line at 12,000kW [18]. So it goes to show that big, fast, computers can be cooled with warm water cooling as well as be more efficient than its competitors.

2.2.4 Performance Gains with Water Cooling

On IBM 2U chassis it has been shown that a 34% increase in processor frequency resulted in roughly 33% increase in performance over an air cooled node. This same processor frequency increase could not be reached in an air-cooled chassis due to temperature limits. [3]

Chapter 3

Theoretical Backbone

3.1 Costs: Running the Water Cooling System

There is still an associated cost with running the water cooling system. Even if “free cooling” was obtained through not chilling the water on inlet, there would still be a cost associated with running the pumps to move water through the system. This section is specific to the TAMIRS cluster and data center explained in Section 4.

The Cooler Distribution Units (CDU) and vacuum pumps (further discussed in Section 4.2.2) use power to move water through the TAMIRS water system. Considering a total cost evaluation, the power usage of these devices becomes a necessity. Each vacuum pump uses 400W and each CDU uses 75W when in primary use. Only one CDU and one vacuum will ever be in primary use; however, the standby power for the others is still measurable. The vacuum pump shuts off when not in use, but it uses .5 watt for the pilot light. The secondary CDU runs reserve as backup to a failover on the first CDU and consumes 40W. Therefore running properly, it can be said that the cooling system consumes 515.5 W. While the CDU and vacuum pumps may contribute a small amount of heat load to the room, it is not substantial compared to the heat produced by the CRAC unit motors.

The conventional way to cool a data center is through Computer Room Air Condition (CRAC) units [3]. These units often are most efficient with warmer return air temperatures [19]. The electricity costs of cooling systems can account for 30% of the total electricity bill for operating a data center [20]. The CRAC units provide a simple example of cost savings to be had in both power and monetary value if less air cooling was needed.

Through another study by the author, performed in the same data center as TAMIRS, it was determined that each of the 18 CRAC units consumes 19.46kW of power, for a total of just over 350kW. A CRAC unit also contributes an additional heat load to the data center through its motors. This amount is about 300kW of heat load to the room for all 18 units [19]. Using the ratings from the CRAC units, Equations 1, and 2 shows the calculations for the CRAC units power consumption and heat production respectively.

$$kVA = (.746kW/hp)(10hp)(0.87^2) = 5.646 [kVA] \quad (1a)$$

$$kW = PF(kVA) = 0.87(5.646)(3motors) = 14.23 [kW] \quad (1b)$$

$$(0.746kW/hp)(7.5hp)(3motors)(18units) = 302.13[kW] \quad (2)$$

Each unit has the ability to cool between 114-120kW of load. With 18 units in the room, the total cooling capacity is 2069kW. If the room is running at full capacity the cost to run the CRAC units, not factoring water costs, the motors in the CRAC units pull nearly 260kW. In a day, at \$0.1256/kW/h (the national residential average) [21], approximately \$780 is spent just running the CRAC units per day. In a year that is nearly \$286,100 spent on electricity.

Assuming the environment could be reduced to a 70:30 environment, 70% water cooled and 30% air cooled, a significant savings could be had. Approximately 620kW of the 2096kW load would need to be cooled by air and the CRAC units. This would allow for 12 CRAC

units to be turned off, both reducing the cost of power use and reducing the heat load on the room. The cost of running only 6 of the 18 units at \$0.12.56/kW/h reduces to a cost of \$256 per day and only \$93,520 spent in a year. That is a savings of approximately 67%. However, not knowing the cost associated for the water system to remove that heat, this savings does not represent a net savings for the data center, rather just the savings associated with turning off the CRAC units which represents the upper bound in terms of savings.

3.2 Reducing Power Use: Fans Small but Notable

Not only are fans a noise hindrance in the data center, they are also an additional source of power use in a chassis. Server nodes often have multiple fans for cooling both the CPU and other onboard components. These fans often ramp up and down with temperature, which allows them to save a little power when not running full speed. However, because some jobs tend to spike in temperature faster than some controllers can adjust, some clusters are set to run full speed all the time.

The Dell R920 nodes in the TAMIRS cluster (further detailed in Section 4) have six relatively large, five inch square, fans. The particular node settings do not require the fans to run at full speed all of the time; however, when an intensive CPU job is run, the fans do spin up to full. The Dell R920 has six Nidec UltraFlow 12VDC, 2.31A, fans capable of moving 158 CFM each. At full speed, these fans pull nearly 166 W per node. That totals to 3 kW for just 20 nodes. A typical HPC cluster can have thousands of nodes; Trinity, LANL's next super computer will have over 19,000 nodes [22]. Even with only one fan in each, that could still easily be 526kW of power just for fans. Though this amount seems small in the megawatts the compute system will pull, every little bit counts.

One hope from this study is to show that with direct water cooling power can be saved by not needing to spin up these fans, and can help reduce the number of fans needed in the system.

Essentially if one could reduce the number of fans needed (or the amount a fan needs to run at full speed), a savings for both power and ones ears could be attainable.

3.3 Electronic Reliability and Longevity

Temperature affects both the reliability and longevity of electronic devices. As the density of compute increases the chip temperature is also increasing. An increased temperature leads to higher power consumption via leakage current and results in higher power consumption over all [2].

The reliability of a chip is said to reduce exponentially as temperature increases as a function of $e^{(-E_a/kT)}$, where E_a is the activation energy for the failure mechanism being accelerated by increased temperature, k is Boltzmann's constant ($8.617\text{e-}5 \text{ eV/K}$), and T is temperature in degrees Kelvin. At elevated temperatures silicon devices can fail catastrophically [2].

The life of an electronic device is directly related to its operating temperature. It is said that every 10°C temperature increase, from the operating conditions, reduces a component's life by 50% [23]. The Arrhenius Equation is typically used to model the acceleration of temperature dependent physical processes that lead to a function of wearout [24]. The Arrhenius Equation given in Equation 3 shows the process rate coefficient is a result of primary material properties, and temperatures; where M is an experimentally determined constant specific to the materials used.

$$C_R = M * e^{(-E_a/kT)} \quad (3)$$

The general case is used to determine a components operating life time at room temperature. Therefore the comparative form of the Arrhenius Equation is used to determine an Acceleration Factor for a change in temperature [24]. Equation 4 gives the Acceleration Factor, where T_1 is

the reference temperature (often room temperature 25°C) and T_2 is the actual use temperature.

$$\text{Acceleration Factor} = e^{[(\frac{E_a}{k})(\frac{1}{T_1} - \frac{1}{T_2})]} \quad (4)$$

Even if silicon devices do not fail catastrophically at higher temperatures the electrical characteristics are still experiencing frequent intermittent and even permanent changes [2]. Manufactures of computer hardware often specify a maximum operating temperatures; most devices are rated to function properly up to that specified temperature [25]. Exceeding these temperatures in a chassis can easily occur on a loaded down compute node. Although the CPU itself might not be reaching its thermal limit to throttle, the other components in the system could be reaching their manufacturing limits. This can result in memory errors (though easily found with ECC memory), disk read-write errors, and other problems [25].

It seems logical to conclude cooler components will result in higher reliability, greater longevity, and lower power use over the entire node. It can also be thought that removing heat through water rather than air also means the node temperature itself could be reduced greatly by not pushing hot air removed from the CPUs through the node. This leaves the fans capable of both moving cooler air, as well as more of it, across the entire chassis. This could increase the lifetime of all of the components.

Directly associated with temperature is leakage current. Leakage current is seen as an increase power use and is often caused by hotter component temperatures [2]. Therefore reducing the temperature of a chip will also result in less leakage power and hence better power performance [3].

3.4 Jitter and Noise

Variations in performance in a cluster can be a major contributor to performance issues for HPC applications. Jitter, also known as noise, can be introduced through various means, including the operating system, CPU, and any component that can introduce a variance [26]. HPC applications in general are more susceptible to noise because they tend to be tightly coupled across the entire job. This means that a single slow core out of an entire cluster can cause the entire cluster to run at the speed of that slow core, thus slowing down the job time and performance of the overall job running across the cluster [27]. Therefore reduction of this variance can result in an improvement in overall performance and power efficiency.

Chapter 4

Testbed Overview

Tiered Active Multi-dimensional Indexed Record Store (TAMIRS) is a Los Alamos National Laboratory (LANL) testbed cluster for use in the exploration of next generation tiered storage technologies. This testbed is composed of 22 Dell PowerEdge R920 servers, 2 custom SuperMicro servers, 4 Dell PowerEdge R720 management nodes and an assorted networking/-management/storage equipment. This modest sized cluster presents an interesting opportunity to explore the use of direct-to-chip cooling technology while allowing for experimentation without much interruption of other testing.

The rack configuration of the TAMIRS cluster is shown in Figure 1. There are ten R920 nodes in Rack 2 and ten R920 nodes in Rack 4. These two racks contain only the R920 nodes and thus are desirable for direct comparison of air cooled versus water cooled. Rack 2 will remain “stock” and will run with the original heatsinks shipped from Dell. Because this cluster is in production and will need to provide cycles to other users, only four nodes will be retrofit to run the water blocks in the lower section of Rack 4. Nodes 11, 12, 13, and 14 will be water cooled and directly compared to its equal air cooled nodes (same height and location in Rack 1), nodes 1, 2, 3, and 4.




Figure 1: TAMIRS rack layout diagram.

The R920 provides a unique heatsink design from factory (shown later in Section 4.2.4). Tight clearances and tolerance requirements for design of the water block design lead to the need for a custom solution. LANL collaborated with Childdyne, Inc. [28] to achieve a workable water block design. Section 4.2 covers this design in detail.

4.1 Full System Specs

4.1.1 Software

The operating system (OS) running on TAMIRS is Tri-Lab Operating System Stack (TOSS) 2.2-3, a common grown OS between the tri-labs (LANL, Lawrence Livermore National Lab (LLNL), and Sandia National Lab (SNL) as a derivative of Red Hat Enterprise Linux (RHEL)).

It has similarities between LANL production cluster in the sense that users obtain jobs via the Moab scheduler, and have access to module files and a shared NFS file system for access to users home and project directories. Additional software utilized for monitoring or testing will be discussed in detail in Section 5.

4.1.2 Hardware

The R920 is unique in that it is a high-density node that is capable of supporting 4 CPU sockets and up to 6 TB of RAM (in the form of 96 - 64GB DIMMS). This makes the chassis for the R920 extremely dense, in spite of its 4U chassis size and a good facsimile to the expected density for an exascale class compute node. The R920 nodes in TAMIRS are half populated with only two of the four sockets populated. The two processors are Intel Xeon E7-4870 v2 at 2.30 GHz with a Turbo clock of 2.90 GHz. Each has 15 cores with 30 total threads, for a total of 60 threads per node. The Thermal Design Power (TDP) is 130W [29]. Table I has more information on the CPU, including thermal threshold values. The Dell labeled sockets “Processor 1” and “Processor 2” [30] are populated as shown in Figure 2. These will later be referred to as CPU0 and CPU1 respectively.

Table I: Intel Xeon Processor E7-4870v2 Specifications [29]

Processor Number	E7-4870v2	Processor Base Frequency	2.3 GHz
Cache	30MB	Max Turbo Frequency	2.9 GHz
Intel QPI Speed	8 GT/s	Maximum Case Temperature	73°C
# QPI Links	3	High Temperature Threshold	73°C
Lithography	22nm	Critical Temperature Threshold	83°C
# of Cores	15	TDP	130W
# of Threads	30		




Figure 2: Picture Showing TAMIRS Socket Population [31]

Each node is comprised of eight 16GB (1600MHz RDIMM dell#331-4428) sticks of RAM for a total of 128GB per node. Divided out among four channels, each CPU receiving two, there are two sticks per RAM riser.

Each node also has three 1100W (dell#450-AAVK) power supplies in a 2+1 redundancy configuration. This means that two power supplies provide power to the node normally, while an extra is ready to go online if a power supply failure occurs. These power supply units (PSU) are rated 80 PLUS Titanium on the 80 PLUS Certification. The 80 PLUS Certification requires that PSUs be 80% or greater energy efficient at 10, 20, 50, and 100% of rated load with a true power factor of 0.9 or greater [32]. Table II shows how the Titanium rating compares in efficiency to other 80 PLUS Certifications.

There are additional hardware components in each node, but are not relevant to this document as they will not be utilized during testing.

Table II: 80 PLUS Certification 115V Internal Non-Redundant [32]

% of Rated Load	10%	20%	50%	100%
80 PLUS	—	80%	80%	80% / PFC .90%
80 PLUS Bronze	—	82%	85% / PFC .90	82%
80 PLUS Silver	—	85%	88% PFC .90	85%
80 PLUS Gold	—	87%	90% / PFC .90	87%
80 PLUS Platinum	—	90%	92% PFC .95	89%
80 PLUS Titanium	90%	92% / PFC .95	94%	90%

4.2 Chilldyne Setup

The Chilldyne system is unique in the field of water-cooling due to the fact that it works under vacuum (negative pressure) rather than positive pumping pressure. This means that the water is pulled through the system, through the manifolds, and water blocks, rather than pushed. Because of this design the entire system is less prone to leaks or spills on equipment. If a line is cut or removed from a manifold or node on purpose or by accident, the vacuum pulls the water back into the system loop and at most a droplet of water results. Traditional systems under pressure will spray out water from a cut or unplugged line and not stop until the pump is stopped or out of water. This makes the Chilldyne a preferred choice in a data center where leaking water could have an effect on several system besides the cluster itself. It also provides assurance that the hardware is safe even when the system is not under visual inspection 24x7.

The following configuration is Chilldyne’s smaller configuration deemed the “demo unit”. The full scaled system is capable of cooling up to 200 kW while consuming only 3 kW of power. This would effectively be a 1.5% power consumption for the amount of cooling power produced. The “demo unit” on the other hand uses, as explained prior, 515.5 W of power. Though the exact cooling capability of the “demo” has not been released, it is more than enough to easily cool the entire Tamirs cluster at over 6 kW full load.

4.2.1 Standalone Chiller

The standalone chiller unit is a ThermoFlex 7500. Its standard operating temperature range is between 5°C and 40°C. Due to the dew point (55°F or ~12°C) and the desire to avoid condensation within the nodes, a safe minimum setting was determined to be 65°C. At 20°C the standard cooling capacity of the ThermoFlex 7500 is 7500W, more than the expected load from the water cooled rack. The full specification sheet for the chiller can be found in Appendix A. This chiller unit could (and ideally would) be replaced with a connection to facility water; a standalone chiller was used for this study to enable the control of water temperature set points outside the available facility water lines. Regardless of where the external water comes from, the standalone chiller or facility water, the system effectively has two loops. One external that provides the chilling (which can follow less strict water requirements) and a second internal loop that includes the CDUs and the cluster itself. Figure 3 shows both the chiller and CDU configuration.



Figure 3: Image of the CDUs with vacuum pumps (left), and standalone chiller (right).

4.2.2 Cooler Distribution Units & Vacuum Pumps

Childdyne has provided a full rack with two Cooler Distribution Units (CDU), and two vacuum pumps. This is run in a 1+1 redundancy configuration, with one vacuum pump per CDU, and currently being run with manual fail-over valve to switch to the secondary CDU and pump if a failure were to occur. An automatic failover valve is still under development at Childdyne.

Because the failover is currently manual, several additional failsafe features were put in place to protect the cluster hardware. These failsafe features include forcing the node to stay shutdown if a thermal shutdown is reached (rather than reboot); and emails are being sent from the master node if nodes are in either the “down” or “drain” state.

The pump in the CDU was developed by Childdyne’s sister company, Flometrics. It was originally designed as a rocket fuel pump. It uses two pumping chambers which are alternately filled with fluid and pressurized in sequence to maintain a steady flow of pressured fluid [33]. This is what allows the system as a whole to run at negative pressure.

Because this is considered a closed loop system (by which water is only added/replaced for maintenance) there are specific water requirements to prevent growth and or build up from occurring. On the inner loop and additive from Childdyne was used in combination with distilled water. Because Los Alamos is known for high silica in the water, distilled water was also used on the chiller side to prevent buildup of scale which could cause blockages and or reduced flow.

As previously stated this deployment is a scaled down version of a larger system where the CDU and vacuum pumps are capable of handling a 200 kW thermal load while dissipating 3 kW of power themselves. The small units in use for this testing have not been evaluated for an absolute thermal load.

4.2.3 Rack Manifolds

There are water lines from the chiller to the CDU, but the CDU still needs a route out to the racks. In each of the three racks (Rack 2, Rack 3, and Rack 4) there are two aluminum $\sim 1/2$ " square manifolds. Each has a total of 12 stub outs to connect tubing to nodes. One manifold is considered the "cold" side while the other is called the "hot" side. Since the entire system works under vacuum, the water is being pulled from the cold side to the hot side by the vacuum in the Chillydyne rack; as mentioned before this reduces the chance of leaks. Because the system is run under low pressure, the manifold stub outs not in use can be simply capped off with rubber caps to maintain the system vacuum. These can be seen on the manifolds in Figure 4.



Figure 4: Image of the manifolds mounted in rack. Orange and black rubber caps block off the stub outs not in use.

These manifolds are connected to each other between racks via waterlines beneath the floor. Figure 5 shows these lines being measured for install between the racks.



Figure 5: Chris from Chillydyne measuring the tubing for rack to rack connections.

4.2.4 Waterblocks

The Dell R920 posed a challenge for a water block design. The original heatsink for the nodes was a copper block with heat pipes extruding out away from the CPU into a set of copper fins. Although it appears as a traditional heatsink, the space to access the CPU was small and had tight tolerances to allow for additional hardware to fit in the node. Because of this Chillydyne had to custom make the water block from an iterative stage. Figure 6 shows the original heatsink on the right, and plastic prototypes for the new water block design to take its place. There was also experimentation with hybrid blocks, in which the fins from the original block were implemented into the water block as well, but this is beyond the scope of this paper.




Figure 6: Chillydyne, Inc. designed water blocks (left) to replace the R920 heatsink (right)

After fit was verified the design was sent out to be fabricated from copper blocks. The resulting blocks that are actually installed in the TAMIRS cluster are shown in Figure 7. Because each node has two CPUs populated, the water blocks were chained in series to allow for one inlet and one outlet to reduce necessary routing space for tubing.




Figure 7: Final water block design for the Dell R920.

Because in a cluster environment nodes may need to be serviced from time to time, the ability to remove all connections from the back of the node is desired. This means that the connection to the water manifolds needs to be removable as well. Again, because the Chillydyne system works under negative pressure, this task can be completed by removing all the water from the node, and providing a disconnect. The receptacles designed by Chillydyne allow for a “hot-swap” connection, allowing for a disconnection at the back of the node without turning of the cooling system. This allows for maintenance that would normally occur on a live system to still occur. The connector when removed pulls the water from the node and allows the node to be removed from the front of the rack with no required connection to the water loop. This connection can be seen in Figure 8.




Figure 8: Left: Internal portion of the hot-swap connection. Right: External portion of the hot-swap connection to allow a node to be serviced while the water system is live [31].

4.2.5 Chillydyne Monitoring & Failsafe Features

The standalone chiller unit can be controlled at the physical unit itself. The temperature set point is set on the front of the chiller and the unit works to maintain the set point based on the water in its closed loop and the dissipation of heat from the CDU side.

The CDU has a web interface accessible through IP access. The GUI interface provides information on the state of the CDU and provides control for turning the unit on or off as well as filling procedures, drawing, vacuum testing, and evacuation for both install and maintenance settings. The text version of the page has updating values for vacuum pressures, inlet and outlet temperatures, and current flow rate of the system. This data is being logged by a curl cron job every minute. Figure 9 shows this available data, which refreshes on a frequency higher than once per second. This data provides insight to the functionality of the CDU.

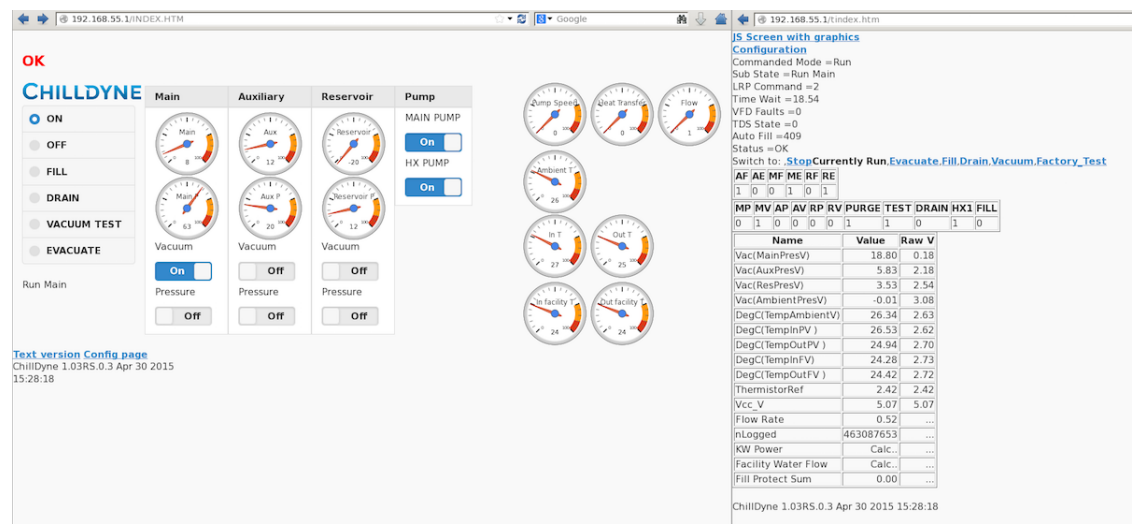


Figure 9: Screen captures showing available data from the CDU.

Because this is all that the Chillydyne system can provide several failsafe features are in place on the cluster side to handle the failure of a CDU or the entire water side system. The main failsafe feature implemented is through the Dell DRAC system in which the option for reaching critical shutdown temperature is to shutdown and stay down. This prevents the cycling of a node on and off if a problem has occurred. It protects the cluster from damages to hardware caused by heat without cooling. When one of the compute nodes is in the “down” state, email messages are sent to system admins who can look into the issue(s) further.

4.2.6 Waterblock Installation

In order to maintain warranty with Dell a specific procedure needed to be followed upon installing the water blocks to the TAMIRS nodes. As per some of these requirements thermal paste from ShinEtsu MicroSi, Inc. (part no. X23-7853W1A-S) was used in a specified application. This application of the paste was as shown in Figure 10. Full documentation of the install process can be obtained upon request.




Figure 10: Thermal paste application on the CPU for the water block

Chapter 5

Testing Setup

5.1 Testing Configurations

The testing programs and configurations are detailed in the following sections. One integral part of the experimentation was testing the variance in the water supply temperature from the chiller unit. Because the chiller was able to provide chilled water, cooler temperatures were tested with 65°F and a warmer temperature of 75°F.

5.2 Stress Testing Nodes

It was important during testing execution to provide a stress level on the nodes comparable to a representative computational task to provide usable results. It was also important to design controlled and repeatable test(s) to be performed across testing runs as water-cooling parameters are varied. Pavillion is a test harness under development at LANL which allows the same configurations of a test suite to be run again and again with minimal efforts.

5.2.1 Pavillon

Pavillon is a test harness under development at LANL. It allows for a test suite to be configured to help insure the same performance metrics are measured between different testing configurations as well as different testing runs. Through this testing harness two tests have been built and configured, HPL and Systemburn. The appropriate configuration files for both Pavillon and the individual tests are included in the subsequent sections. Once Pavillon is configured, the different test configuration can be run using the following command:

```
pav run_test_suite configurationFile.yaml -m
```

This command invokes Pavilion (pav) to run the test suite defined by a given configuration file (which is written in yaml). The -m option toggles the metrics option which allows for collection of data via the Lightweight Distributed Metric Service (LDMS, see Section 5.3.1). This starts the job when the requested nodes are available through the Moab scheduler. STDOUT and STDERR from the jobs run is located in a subdirectory specified by date, test name, and the specific Moab job id. The LDMS data is sent to another directory which is also specified by the Moab job id. This provides a logical placement of output data that can later be correlated. An example Yaml configuration file is provided in Appendix B.

5.2.2 HPL

The High Performance LINPACK (HPL) benchmark, also referred to as LINPACK, is common for testing CPU stability as well as cluster performance in the HPC world. The Top 500 uses the LINPACK benchmark as a yardstick of performance to compare some of the top rated Supercomputers [34]. The benchmark reflects a clusters “performance of a dedicated system for solving a dense system of linear equations [34].”

Because HPL is a CPU intensive test, it tends to use more power and get the CPU hotter

than many other tests. Therefore it was chosen as a focused test that would show the direct effects of the single change from air cooling to water cooling.

The parallel implementation of HPL can be configured to run on multiple nodes over a variety of cluster structures. The HPL.dat file is configuration file in which parameters can be tweaked for best GFLOP performance values. Tuning this file can take several runs to perfect. The main factors that effect how the file is configured are based on the number of nodes desired to run on, the number of cores per node, and the amount of memory per node. Using a simple HPL.dat file tuner available online from Advanced Clustering Technologies Inc. [35] provided an adequate configuration for testing. The HPL.dat used can be found in Appendix C.

5.2.3 Systemburn

Because HPC clusters tend to use more than their CPUs for computation, relying on performance of other hardware components as well, it was important to test “real world” examples to show what effects water cooling may have under typical use. Therefore configurations were built using Systemburn to provide representation of utilization of more than the CPUs on the node(s).

Sytemburn is a software package developed at Oak Ridge National Laboratory (ORNL) designed to allow methodical creation of system loads. Testing “modules” include: CBA, DGEMM, DSTREAM, DSTRIDE, FFT1D/2D, GUPS, LSTREAM, LSTRIDE, PV4, RDGEMM, SLEEP, TILT, and WRITE [36]. These different tests can be configured with different options and combined together to form “LOAD” files for Systemburn to run. Detailed descriptions are available in the documentation for these modules. Only the four used for testing will be explained for sake of brevity.

DGEMM, DSTREAM, PV2, and SLEEP were the four modules used with Systemburn. Table IX provides descriptions for these tests. The “size” parameter will refer to a given parameter in the Systemburn LOAD file.

Table III: Systemburn module summary, note the “size” parameter described will be used in the LOAD configuration file [36].

DGEMM	A double precision matrix multiplication benchmark which will run to consume “size” bytes of memory
DSTREAM	Streaming double precision floating point vector operations turn to consume “size” bytes of memory
PV2	A power hungry streaming computational algorithm on one array of 64bit values, which will operate with a memory footprint of “size” bytes. This load was tuned to a quadcore Intel “Nehalem” processor, but may be suitable for loading multiple x86-64 cores until the memory system is saturated. It is intended to be run with a footprint large enough to require main memory access.
SLEEP	Puts a thread to sleep for N seconds at a time

These load files are dependent on the amount of cores and memory available per node. It was important to think about the configurations for these different tests in detail. For example, two test configurations were created for DGEMM; DGEMM_LARGE* and DGEMM_SMALL for runs with larger amounts of ram and smaller runs respectively. If done with a small enough amount of memory, such to remain in cache (30MB in the case of these nodes) would ultimately result in a warmer running job as more time is spent at the CPU rather than fetching from RAM. The PV2 test takes advantage of this scheme as well. Likewise the larger RAM size assignment meant the job had to hit RAM, which while isn’t getting the CPU as hot, is providing example of more real application use. *Note: the larger RAM amounts for DGEMM_LARGE required longer periods of time to complete due to the memory access.

Configuration for each of these .LD files is included in Appendix D

5.2.4 Running the Tests

Due to the test run lengths getting in multiple iterations of each test was difficult. To maintain statistical relevance, each of the tests was run a total of 5 times for each of the valued temperatures. Each full test set was comprised of: 5 iterations of HPL, 5 iterations of DGEMM, 5 iterations of DSTREAM, and 5 iterations of PV2. The test set was repeated on each of the configurations (air and both water configurations).

Since the time length of the tests were long, the node was able to reach an equilibrium

temperature and run there for most of the duration of the test to provide best representation of temperature, power use, and performance. Each of the tests were also followed by a sleep period in which the node was allowed to reach idle temperature before having to run the next test. The runtime values are shown in Table IV.

Table IV: Run times for each test suite, recall 24 hours = 1440 minutes

Test Name	What is Run	Time to Run [MIN]	x5
HPL	HPL.dat	~84	420
DGEMM	DGEMM_LARGE & DGEMM_SMALL & SLEEP	105	525
DSTREAM	DSTREAM & SLEEP	45	225
PV2	PV2 & SLEEP	45	225
	TOTAL TIME	279	1395

5.3 Monitoring

Measuring temperature would have required an additional script to launch with the jobs collecting data to later be deciphered. Fortunately prior work by the author in development of LDMS@LANL [37] provided an adequate tool for collecting temperature data during a job run in an easily parseable format with date and time. This tool is described in a following subsection.

Measuring power was one of the more difficult metrics to measure. In a data center many clusters may be run from one Power Distribution Unit (PDU). Trying to measure the power of a particular cluster from that PDU can be difficult to factor out. However, rack mounted PDUs are available in metered options. TAMIRS was purchased with per outlet PDUs that allow for individual plug power collection. Running Average Power Limit (RAPL) is an Intel tool used in power capping; it also has the ability to calculate an estimated power use value. These two tools will be used and compared jointly. Details on these monitoring tools is described in the following subsections.

5.3.1 LDMS

The Lightweight Distributed Metric Service (LDMS) tool was developed at Sandia National Lab (SNL) as a data collection and transport system. It provides capacities for lightweight runtime collection of high-fidelity data. Data can be accessed or transported off node via CSV files. This tool is now publicly available on Github under the ovis-hpc project [38]. The LANL version of LDMS (developed to run as a user interface [37]) was used to collect temperature data on the interval of every second during the duration of the tests run by Pavillion.

5.3.2 Temperature: lm_sensors & LDMS

The package “lm_sensors” is installed on the TAMIRS cluster. This allows for the command line program “sensors” to be called on any node. This outputs the currently reported core temperatures from all 30 physical cores on the node separated by CPU0 and CPU1. The Physical id outputs the maximum of these temperatures at the top of each of the cores. This provides a quick reference for how temperature looks on a node.

LDMS uses plugins, often written in C, to determine where to get data for collection. The plugin for “procsensors”, which is used for temperature collection, also uses the same system locations as lm_sensors for collecting the temperature data. This also means that this maximum value per node is also collected. Because of this it was important to be sure to remove the duplicate data when plotting the resulting data.

5.3.3 Power: Power Distribution Units

The cluster configuration allows for power monitoring at the individual plug level with the American Power Conversion Corporation (APC) per outlet metered Power Distribution Units (PDU); therefore the power usage of each node is available. There are a total of six PDUs for the cluster, two per rack. The specific model: AP8641 has a web interface in which all the

command line parameters can be set from a web page accessible to the configured IP address. After setting up the interface through serial port connection access to the PDUs was available from the cluster via IP. The PDUs were set to collect power data on a minute interval. This data was sent to the master node of the cluster via an FTP server. This data was collected for each of the PDUs and stored for access by date and time. Parsing of this data for the specific nodes was required.

5.3.4 Power: Running Average Power Limit

Intel's Running Average Power Limit (RAPL) driver was designed for power capping [39], however, it also provided a means for power metering. RAPL provides a set of counters for energy and power consumption information. It is a software power model that estimates energy usage by using hardware counters and I/O models [39].

Because power usage is a desirable metric, several scripts have been written for collection of data from the RAPL driver. A program from the University of Maine was written to read the RAPL msr registers containing the power usage and other relevant counter data for Sandy Bridge machines [40]. Once compiled, one could run and obtain the following output with root or sudo privileges:

```
Found Ivybridge-EP CPU
Checking core #0
Power units = 0.125W
Energy units = 0.00001526J
Time units = 0.00097656s

Package thermal spec: 130.000W
Package minimum power: 73.000W
Package maximum power: 230.000W
Package maximum time window: 0.045898s
Package power limits are unlocked
Package power limit #1: 130.000W for 0.043945s (enabled, not_clamped)
Package power limit #2: 156.000W for 0.002930s (enabled, not_clamped)
```

```

Package energy before: 2374.695236J
Accumulated Package Throttled Time : 0.603516s
PowerPlane0 (core) for core 0 energy before: 64566.172470J
PowerPlane0 (core) for core 0 policy: 0
PowerPlane0 (core) Accumulated Throttled Time : 0.603516s
DRAM energy before: 12250.594055J

Sleeping 1 second

Package energy after: 2412.593170 (37.897934J consumed)
PowerPlane0 (core) for core 0 energy after: 64591.276184 (25.103714J consumed)
DRAM energy after: 12257.118683 (6.524628J consumed)

Note: the energy measurements can overflow in 60s or so
      so try to sample the counters more often than that.

```

A script was written to call this program every 30 seconds in an infinite loop. This was launched with the HPL and Systemburn tests to collect power usage data during the lifetime of the job. Any overhead inserted by this collection of data would be invoked equally upon each of the jobs to be compared and thus is said to be negligible. This data was later parsed for plotting.

5.3.5 CPU Clock Speed

Throttling by temperature is a reason a job may take longer to complete than expected. Thermal throttling reduces the clock speed of the CPU when a thermal limit is reached. It is a protective method to protect the integrity of the hardware. Using counters from `/proc/cpuinfo` was not adequate in measuring the changes in CPU speed as the frequency for collection did not capture changes.

Chapter 6

Results

6.1 Performance

Performance data between air-cooled and water cooled results are shown in the following subsections. The raw result values were presented for HPL while the Systemburn modules presented results as minimum, mean, and maximum for the run. These resulting values presented are averages over the five runs each.

6.1.1 HPL

Table V: HPL test results, given as average values for, minimum, mean, and maximum.

Cooling Method	RESULT [GFLOPS]	STDEV	% Improved
air	1247	9.34	—
water (65°F)	1257	11.04	0.80%
water (75°F)	1260	14.13	1.04%

6.1.2 DGEMM

Table VI: Systemburn:DGEMM test results, given as average values for, minimum, mean, and maximum.

Cooling Method	MIN. [MFLOPS]	MEAN [MFLOPS]	MAX. [MFLOPS]	% Improved (Mean)
air	356.19	367.44	381.63	---
water (65°F)	371.11	378.93	393.16	3.12%
water (75°F)	366.11	375.54	387.00	2.20%

6.1.3 DSTREAM

Table VII: Systemburn:DSTREAM test results, given as average values for, minimum, mean, and maximum.

Cooling Method	MIN. [MFLOPS]	MEAN [MFLOPS]	MAX. [MFLOPS]	% Improved (Mean)
air	354.19	360.32	366.58	---
water (65°F)	348.81	360.19	367.922	-0.04%
water (75°F)	353.94	361.81	367.47	0.45%

6.1.4 PV2

Table VIII: Systemburn:PV2 test results, given as average values for, minimum, mean, and maximum.

Cooling Method	MIN. [MTRIPS/s]	MEAN [MTRIPS/s]	MAX. [MTRIPS/s]	% Improved (Mean)
air	22.49	23.47	24.13	---
water (65°F)	23.92	24.09	24.186	2.66%
water (75°F)	23.90	24.08	24.192	2.60%

6.2 Temperature

The average data center room temperature was about 66°F (~19°C) during testing. Variance in this temperature was not controllable due to being a production environment, thus comparative tests between air and water were run during the same time frames.

A key thought to remember while working in both °F and °C, is that the changes in temperature do not relate one to one. For example a 10° change in °F is not the same as a 10° change in °C. In fact the expected difference from a delta in fahrenheit can be expressed as about 1.8x

the celsius value. Therefore a difference of 10°F correlates to ~5.55°C difference. Ideally one would maintain the same units throughout a study, unfortunately the differences in compute hardware and facilities components also disagree in units (°C and °F respectively) and thus leads to the discussion of both units.

Running with a 65°F (~18°C) water inlet temperature the temperatures during jobs were cooler on average of ~20°C. At idle the same difference between an air-cooled node and water cooled node was about ~20°C. In both cases about a 7° variance between cores was observed between the coldest and hottest core.

Running with a 75°F (~24°C) water inlet temperature the expected increase about the expected 5-6°C increase was seen over the 65°F water.

Selected plots, representing the overall data, are provided for each of the testing sections representing the general temperature summary above. Because the data was collected at a high interval (once per second) a rolling average was used over a 30 second interval to reduce noise in the plots.

6.2.1 HPL




Figure 11: Temperature plots during HPL; air cooled (left), 65°F water (center), 75°F (right).

6.2.2 DGEMM




Figure 12: Temperature plots during DGEMM; air cooled (left), 65°F water (center), 75°F (right).

6.2.3 DSTREAM




Figure 13: Temperature plots during DSTREAM; air cooled (left), 65°F water (center), 75°F (right).

6.2.4 PV2




Figure 14: Temperature plots during PV2; air cooled (left), 65°F water (center), 75°F (right).

6.3 RAPL

The RAPL data showed little difference between air cooling and water cooling during the different job runs. This difference could be due to the fact that RAPL uses a calculation to determine expected power usage based on clock speed and the processor details. It is unclear whether temperature is taken into consideration.

RAPL data was collected on 30 second intervals. The data needed to be collected more frequently than 60 second intervals due to the rollover of the counters RAPL uses. Unfortunately, the rollover of counters still incurred during some test runs. Because of this these large negative values were rejected as outliers in the data and removed for plotting and statistical purposes.

6.3.1 HPL




Figure 15: RAPL plots during HPL; air cooled (left), 65°F water (center), 75°F (right).

6.3.2 DGEMM




Figure 16: RAPL plots during DGEMM; air cooled (left), 65°F water (center), 75°F (right).

6.3.3 DSTREAM




Figure 17: RAPL plots during DSTREAM; air cooled (left), 65°F water (center), 75°F (right).

6.3.4 PV2




Figure 18: RAPL plots during PV2; air cooled (left), 65°F water (center), 75°F (right).

6.4 Power from PDU

The PDU data was collected every minute, highest frequency for the PDU collection. The data was available as per plug watt values. Each node had three connections to the PDUs and thus had to be summed over three plugs to obtain total power consumption for the node itself. The resulting per node watt value is plotted over time for each of the jobs run. Representative plots for each of tests are shown the the following figures. A blanket statement can be said that a small difference in power was seen between the 65°F inlet water and 75°F inlet water temperatures.

6.4.1 HPL

Measurements during the HPL testing runs showed that the power consumption under water cooling was reduced on average by about 25W. It was also noted that the maximum power reached during the run was reduced by nearly 30W.




Figure 19: PDU plots during HPL; air cooled (left), 65°F water (center), 75°F (right).

6.4.2 DGEMM

Power usage during the DGEMM testing showed the mean power consumption reduced by nearly 30W by water cooling. The maximum power reached during the run was also reduced by approximately 40W.




Figure 20: PDU plots during DGEMM; air cooled (left), 65°F water (center), 75°F (right).

6.4.3 DSTREAM

During the DSTREAM runs the average power consumption was reduced by almost 20W from air cooling to water cooling. The maximum power reached was reduced by almost 30W.




Figure 21: PDU plots during DSTREAM; air cooled (left), 65°F water (center), 75°F (right).

6.4.4 PV2

The mean power use while running PV2 under water was reduced by just over 20W. The maximum power reached was reduced by just over 30W.




Figure 22: PDU plots during PV2; air cooled (left), 65°F water (center), 75°F (right).

6.4.5 PDU Data Summary

Table IX: PDU Mean Power Data Summary.

Test	Mean Power Savings Per Node	Total Mean Power Savings
HPL	~25 [W]	~100 [W]
DGEMM	~30 [W]	~120 [W]
DSTREAM	~20 [W]	~80 [W]
PV2	~20 [W]	~80 [W]

Chapter 7

Discussion

The data may not have matched up exactly with expectations presented prior to the results of this study. The following sections detail the observations.

7.1 Performance

Performance data between air cooled nodes and water cooled nodes was exactly what was desired. A performance increase would have shown that cooler transistors are able to perform better. However, the particular architecture the study was performed on did not have the ability to turn off power limits to enable these longer bursts of turbo. In this situation the hardware was power limited not thermal limited.

However, in some BIOS on some motherboards, it is possible to not be power limited, and rather be thermal limited for sustaining higher clock speeds. In these situations, water cooling could (and would be expected) to have an effect on performance data.

For the case of this study, having the same exact performance showed that water cooling, even at the warmer temperatures did not negatively effect performance; which was ultimately the desired result.

The surprising result shown in the performance data is actually observed in the minimum performance points shown in the Systemburn data. The minimum performance in the calculation intensive tests, DGEMM and PV2 the minimum performance was increased over air by 4.19% and 6.27% respectively. This, while clearly contributing to the increased mean, provides encouraging results on overall cluster performance. Considering that a job is always waiting on the slowest core to finish, and cluster wide, the slowest node to finish, this improvement on the slowest one shows better performance overall. On a much larger job, over multiple nodes, this could prove to be a substantial increase in speed, and could result in more than minutes or hours in difference of job time to completion.

7.2 Temperature

The difference in temperature between the air cooled nodes and the water cooled nodes was better than expected. Having to have custom blocks designed to fit in such a tight tolerance region, it was unknown exactly what performance would be expected compared to other water block designs from Chillydyne.

Peak temperatures were reduced by about 20°C at a 65°F inlet temperature and about 15°C at at 75°F inlet temperature. In this particularly cool data center, the 75°F water temperature was actually warmer than room temperature (approx. 66°F); however, even the idle dissipation of the nodes was enough to maintain this water temperature.

Relating this back to the prior section, this further shows that on the right setup, thermal limitations would likely not be the cause for CPU throttling, and rather could remain in turbo mode longer.

Because water is removing most of the thermal load from the node, warmer room temperatures could be obtained without causing temperature issues that may be observed in the air cooled nodes. This warming of the room would reduce facilities costs and increase CRAC

efficiency.

7.3 RAPL

Not knowing exactly the method used beneath the covers of RAPL, it is thought the RAPL data provided somewhat inconclusive results. The estimated power use between air cooled and the water cooled nodes varied little, though the estimated power usage did seem smoother for the water cooled nodes. Based upon these results it is thought that RAPL may factor in the temperature of the CPU for its calculation. These estimations may or may not be accurate as described in the next section.

7.4 PDU

The actual power data collected by the in rack PDUs provided proof of reduced power use between air cooled and water cooled nodes. It is hard to argue whether the power reduction was caused by a reduction in fan speed or perhaps from the CPUs running cooler themselves. If fan voltages could be measured, this could probably be reduced to a root cause for reduction. Nonetheless a power reduction was observed.

Although the 166W power reduction from maximum fan speed was not observed a substantial 30W per node was observed. It was expected that the fans would still be running, so a 166W decrease was an unreasonable to be expected. 30W does not seem like a lot of savings, but taken to scale does add up. Over the four nodes in this testing scenario it was a total of about 120W saved. Over a 24 hour period this is a savings of 2.88kWh/day. Though this only adds up to a savings of \$131.61 in a year on power alone, this does not factor in the effects taken at scale.

Scaled up to the full cluster of a small 20 nodes saves about 600W, the equivalent of ten

60W lightbulbs left on year round. This savings results in 14.4kWh/day adding up to \$658.05 in the course of a year. It is easy to see how this can continue to increase in cost savings with scale.

Though the cost of the Chilldyne system can not be released, the cost for it is not recouped by these power savings alone. Additional savings however could be acquired if this was further taken to the scale of the data center. Returning to the initial thoughts for a 70:30 divided data center significant cost savings could be obtained in reduced facilities costs.

This PDU data, combined with the RAPL results, suggests that the power savings incurred was not due to the CPUs reduction in energy consumption but rather other components of the node. It is thought that perhaps the fans were not spinning up as much during the job. This could be validated through measurement of fan voltages during testing scenarios.

7.5 BONUS: Additional Discovery on Air Cooled Nodes

During early testing each of the sockets was observed for temperature values before averaging the data to produce readable plots. What was observed was interesting and perhaps revealing of the node layout. On the air cooled nodes there was always a split between CPU0 and CPU1 on the order of about 10°C. This was observed on every air node within the cluster. Once water cooled, as shown in Figure 23, the temperatures between sockets became much closer in temperature.

After having removed the air cooled heatsinks to install the water blocks, two explanations are possible for the observed behavior. The first explanation could be due to the physical location of CPU0 and CPU1. CPU0 is the furthest left on the board. This sits behind a populated disk enclosure, thus likely receiving less airflow despite having the same amount of cooling fans pushing through the sink. CPU1 on the other hand sits a socket to the right of CPU0 and does not have a similar obstruction in front of it, but is rather more open. Unfortunately the

other two sockets, CPU2 and CPU3 are not populated. This would have provided more conclusive data as to whether this layout was part of the cause for difference as CPU2 and CPU3 have the same obstructions mirrored on the other side.

The second explanation for the temperature difference, though thought to be less likely, was the variance in production quality of the heatsink from factory. While removing the air cooled heat sinks to install the water blocks it was observed that several of the mating surfaces with the cpu had variance in flat-ness that could hinder thermal conductivity between it and the CPU even with the addition of thermal paste. Out of the eight observed heat sinks only one had what was thought to be a significant difference in surfaces. Thus, it is thought that the first explanation is more likely responsible for the large variations seen during testing.

Considered for an explanation was also the application of the thermal paste. However, the thermal paste application from factory was uniform and likely applied in a mechanical fashion. The variation of thermal paste applied by a human for the water cooled nodes did not show this significant difference observed.

What might be most notable by looking at the specific CPU0 and CPU1 break down between a air-cooled and water cooled node is the significance in temperature spread. What is seen as nearly a 20°C between the coldest and hottest CPU on the air cooled node is reduced to about 7°C on the water cooled side. Even at warmer water temperatures this temperature variance across the same set of nodes is reduced significantly over the air cooled nodes. This shows that the cooling is pretty evenly distributed to the nodes, whereas in the air-cooled case the warmest CPU is going to be the one to slow the entire job down. Once it has to throttle the entire job is again waiting on that core to finish. The even distribution shown in the water cooled nodes means that the performance should likely be even among the cores as well, resulting in better performance over all.




Figure 23: HPL Test run showing per CPU temperatures.

Chapter 8

Conclusions

Through the culmination of over a year and a half of work; LDMS tool development TAMIRS deployment; configuration, and on going support; Chillydyne contract management and installation; integration of a test suite and monitoring; the results of this study were accomplished.

It was determined, at the temperatures tested, that no performance gain was obtained from water cooling alone. However, water cooling did provide cooler CPU temperatures which could allow for higher clock speeds to be obtained. The idle and load temperatures were significantly lower than the air cooled nodes. This provides increased resiliency and lifetime of the components as even running 10°C cooler increases the life time of electronic components by double. With HPC clusters often having a home in a data center for at least 5 years, it is important to have resilient components to keep them running while minimizing down time. The easiest way to reduce down time, is to not have the hardware fail in the first place; thus having longer lasting components by providing them with a cooler working environment seems like a start. However, the testing performed did not confirm nor deny extra longevity.

Considering the inlet water temperatures, the warmer water temperature of 75°F, the CPUs were still able to run 10°C cooler than when cooled with 55°F air temperatures. The real savings here, which could not be observed at this small scale, would have been the cost savings

of the running chiller units, the CRAC units, as well as the cooling towers. It is important to consider the scalable impacts that this smaller study has shown plausible. Reducing the need for cool air cooling in a data center significantly reduces the costs associated with creating the cool air. Large chiller units, and cooling towers are used for producing the 55°F air to the data centers. It is possible that chiller units would not even have a part in the process if a data center were to be comprised of water cooled clusters. There may be a need for some comfort cooling; but this would be a minimal cost compared to the costs of maintaining the cold room.

The RAPL data suggested that little to no power efficiency was gained at the chip level. The concepts of leakage current at the temperatures tested did not show in the RAPL results. It is unclear whether or not the RAPL calculations factor in the chip temperature or not; however while under water, the RAPL data did show smoother curves. Assuming the RAPL data was accurate, the power savings at the PDU can then likely be explained by reduced fan speeds in the water cooled nodes. At nearly 30W per node in power reduction the reduction of power use over a scaled cluster would be significant. Even at the full capacity of the TAMIRS cluster, this would be over a 600W savings. Scaled to a cluster with thousands of nodes, the savings would be even more significant.

Through this testing it was also discovered that the air cooled nodes had a particular notable difference in temperature per socket based on socket layout. Due to the hardware configuration CPU0 had less air available to it through blockages in the chassis. CPU1 had more air flow from the front of the chassis. This was discovered through the close temperature monitoring of the nodes. It was determined that CPU0 ran 10°C hotter than CPU1. Because many HPC jobs are tightly coupled this difference in temperature could create a significantly slower CPU that the entire job would have to wait for. This combined with the variance in heat sink surface quality could cause a real slowdown for a typical HPC job. The water cooled nodes show similar temperatures between CPU0 and CPU1 providing a better consistence of speed/performance between chips.

The warmer water temperatures can be seen as a gateway to “free cooling” with a great savings on the facilities side. In efforts to maintain a specific PUE value while driving towards exascale and its demanding power requirements on the compute and facilities side, the reduction in facility cost via water cooling could provide more power for where it counts; the compute side.

In a data center limited by CRAC unit capacity, systems such as this can be used to expand data center compute capacity far beyond the limitations of the available air cooling. In a full scale system each CDU can handle approximately 200 kW of thermal load while only using 3 kW of power. With an external water to air heat exchanger the entire compute thermal load could be removed from the CRAC units and other resources (additional storage, networking, etc) could be deployed to utilize the excess thermal capacity. Alternatively, the thermal load savings could be utilized in a way that fewer CRAC units are required which directly lowers the cost of maintaining the data center and its systems.

Chapter 9

Future Work

Due to time constraints and the desire to graduate, this experimentation and study was limited to the above documentation. However, there are several tests and areas to still be explored. This work will continue despite the completion of this document. A few of the planned testing areas the author wishes to explore are described in the following subsections.

9.1 Warmer Water Cooling

Warm water cooling is essentially free cooling. If the water does not need to be chilled before entering the system a lot of money can be saved. Being able to warm up the water and maintain a warmer than room temperature water temperature requires a little more work.

When more water blocks are installed on the system, a “dummy” load can be run on the non-testing nodes. This will provide additional heat in the water loop to help maintain the water temperature between tests as well as during tests that may not produce enough heat to maintain a water temperature.

The plan is to measure warmer water temperatures to see just how warm the water can be before causing issues with throttling at the CPU level. Based on the data presented, it is

thought that there is a 20°C temperature between the throttling point and where the nodes are running at a 65°F water temperature. This relates to a 36°F delta in water temperature resulting in 101°F water temperature. This is on the upper end of what is supported by the Chiller unit (limited to 104°F) but is still an interesting goal.

9.2 Tightly Coupled Applications

Many of the synthetic benchmarks presented here are designed to maximize power usage as well as maximize throughput. Because of this, much of the workload is dominated by compute power and not by synchronized computation. Applications that tend to run on clusters of this design are inherently more coupled and require various synchronization points during the run. These synchronization points depend on the performance of the slowest node, and as such, if the environment for these systems was not as ideal as it was, thermal throttling of an air cooled node could present itself in the form of much reduced system performance for these tightly coupled workloads.

9.3 Looking at Scale

Unfortunately due to time constraints and other considerations the entire cluster was not able to be completely water cooled for this particular research. However, there are plans to install water blocks on the other 20 compute nodes. Additional testing will be performed to see how well this four node test scales out to the entire cluster.

A Chiller Unit: ThermoFlex 7500 Spec Sheet

3500 to 10000 series

3500 to 10000
watts of cooling



Specifications	ThermoFlex 3500	ThermoFlex 5000	ThermoFlex 7500	ThermoFlex 10000
Standard Temperature Range	5 °C to 40 °C (41 °F to 104 °F)			
Optional Temperature Range	5 °C to 90 °C (41 °F to 194 °F)			
Ambient Temperature Range	10 °C to 40 °C (50 °F to 104 °F)			
Temperature Stability	±0.1 °C			
Standard Cooling Capacity				
60 Hz at 20°C	3500 W / 11953 BTU	5000 W / 17076 BTU	7500 W / 25575 BTU	10000 W / 34100 BTU
50 Hz at 20°C	3050 W / 10416 BTU	4400 W / 15027 BTU	6425 W / 21910 BTU	8500 W / 28985 BTU
Reservoir Volume	7.2 Liters (1.9 gallons)		17.9 Liters (4.8 gallons)	
Refrigerant	R407C			
Physical Dimensions (HxWxD)				
Air-Cooled	38.9 x 19.3 x 30.9 in (98.7 x 48.8 x 78.4 cm)		52.3 x 25.2 x 33.8 in (132.7 x 63.9 x 85.6 cm)	
Water-Cooled	38.9 x 19.3 x 30.9 in (98.7 x 48.8 x 78.4 cm)		45.9 x 25.2 x 33.8 in (116.6 x 63.9 x 85.6 cm)	
P1 — Positive Displacement Pump				
60 Hz	2.1 gpm @ 60 psig (7.9 lpm @ 4.1 bar)	—		
50 Hz	1.7 gpm @ 60 psig (6.4 lpm @ 4.1 bar)	—		
P2 — Positive Displacement Pump				
60 Hz	4.0 gpm @ 60 psig (15.1 lpm @ 4.1 bar)			
50 Hz	3.3 gpm @ 60 psig (12.5 lpm @ 4.1 bar)			
T1 — Turbine Pump*				
60 Hz	3.5 gpm @ 60 psid (13.2 lpm @ 4.1 bar)	—		
50 Hz	2.5 gpm @ 60 psid (9.5 lpm @ 4.1 bar)	—		
T5 — Turbine Pump*				
60 Hz	—	7.3 gpm @ 60 psid (27.6 lpm @ 4.1 bar)		
50 Hz	—	6.2 gpm @ 60 psid (23.5 lpm @ 4.1 bar)		
P3 — Centrifugal Pump*				
60 Hz	10 gpm @ 32 psid (37.9 lpm @ 2.2 bar)			
50 Hz	10 gpm @ 20 psid (37.9 lpm @ 1.4 bar)			
P4 — Centrifugal Pump*				
60 Hz	15 gpm @ 57 psid (56.8 lpm @ 3.9 bar)	—		
50 Hz	15 gpm @ 34 psid (56.8 lpm @ 2.3 bar)	—		
P5 — Centrifugal Pump*				
60 Hz	—	20 gpm @ 60 psid (75.7 lpm @ 4.1 bar)		
50 Hz	—	20 gpm @ 35 psid (75.7 lpm @ 2.4 bar)		
Chiller Weight (with P2 pump)	264 lb (120 kg)		356 lb (161.5 kg)	
Voltage Options				
208-230 V/60 Hz & 200 V/50 Hz ^{1,2}	Available	—		
230 V/50 Hz ¹	Available	—		
200-230 V/50-60 Hz Global Voltage ^{1,2}	Available	—		
208-230 V/60 Hz/3 phase & 200V/50 Hz ^{1,2}	—	Available		
400 V/50 Hz/3 phase	—	Available		
460 V/60 Hz/3 phase & 400 V/50 Hz/3 phase ^{1,2}	—	Available		
Compliance				
T5 Pump CSA compliance pending	CE ¹ CE compliant ² CSA compliant			

Specifications obtained at sea level using water as the recirculating fluid, at a +20 °C process setpoint, +25 °C ambient condition, at nominal operating voltage. Cooling capacity based on chillers with P2 pumps with no backpressure. Other pumps will affect cooling capacity performance. Specifications subject to change.
¹Pressure values for centrifugal and turbine pumps are differential pressures between the inlet and the outlet of the chiller.

B Pavillon Test Configuration File Example

```
hpl_single :
  name: hplAir
  source_location: '/nfz/home/noranyk/pvTamirs/test_exec/hplFour/'
  run :
    cmd: runHPL
    scheduler: moab
    test_args: '240' #nodes * 60
    count: 1
  moab:
    num_nodes: 4
    procs_per_node: 60
    time_limit: 00:02:30:00
    node_list: "ta01+ta02+ta03+ta04"
```

C HPL.dat

```

HPLinpack benchmark input file
Innovative Computing Laboratory , University of Tennessee
HPL.out      output file name (if any)
6            device out (6=stdout,7=stderr,file)
1            # of problems sizes (N)
204800      Ns
1            # of NBs
128         NBs
0           PMAP process mapping (0=Row-,1=Column-major)
1            # of process grids (P x Q)
15          Ps
16          Qs
16.0        threshold
1            # of panel fact
2           PFACTs (0=left, 1=Crout, 2=Right)
1            # of recursive stopping criterium
4           NBMINs (>= 1)
1            # of panels in recursion
2           NDIVs
1            # of recursive panel fact.
1           RFACTs (0=left, 1=Crout, 2=Right)
1            # of broadcast
1           BCASTs (0=1rg,1=1rM,2=2rg,3=2rM,4=Lng,5=LnM)
1            # of lookahead depth
1           DEPTHs (>=0)
2           SWAP (0=bin-exch,1=long,2=mix)
64          swapping threshold
0           L1 in (0=transposed,1=no-transposed) form
0           U  in (0=transposed,1=no-transposed) form
1           Equilibration (0=no,1=yes)
8           memory alignment in double (> 0)

```

D Systemburn Load Files

```
#DGEMM_LARGE
```

```
LOAD.START
```

```
    RUNTIME 3600
```

```
    SCHEDULE BLOCK
```

```
    SUBLOAD 2
```

```
    [
```

```
        PLAN 30 DGEMM 512MB
```

```
    ]
```

```
LOAD.END
```

```
#DGEMM_SMALL
```

```
LOAD.START
```

```
    RUNTIME 1800
```

```
    SCHEDULE BLOCK
```

```
    SUBLOAD 2
```

```
    [
```

```
        PLAN 30 DGEMM 0.5MB
```

```
    ]
```

```
LOAD.END
```

```
#DSTREAM
```

```
LOAD.START
```

```
    RUNTIME 1800
```

```
    SCHEDULE BLOCK
```

```
    SUBLOAD 2
```

```
    [
```

```
        PLAN 30 DSTREAM 1GB
```

```
    ]
```

```
LOAD.END
```

```
#PV2
LOAD.START
  RUNTIME 1800
  SCHEDULE BLOCK
  SUBLOAD 2
  [
    PLAN 30 PV2 512MB
  ]
LOAD.END
```

```
#SLEEP
LOAD.START
  RUNTIME 900
  SCHEDULE BLOCK
  SUBLOAD 2
  [
    PLAN 30 SLEEP 1
  ]
LOAD.END
```

References

- [1] Henry Coles, Michael Ellsworth, and David Martinez. “hot” for warm water cooling. *ACM*, 978-1-4503-0771-0, November 2011.
- [2] Yongkui Han, Israel Koren, and Csaba Andras Moitz. Temperature aware floorplanning.
- [3] M.J. Ellsworth, L.A. Campbell, R.E. Simons, M.K. Iyengar, R.R. Schmidt, and R.C. Chu. The evolution of water cooling for ibm large server systems: Back to the future. *IEEE*, 978-1-4244-1701-8:266–274, August 2008.
- [4] Christian Conficoni, Andrea Bartolini, Andrea Tilli, Giampietro Tecchiolli, and Luca Benini. Energy-aware cooling for hot-water cooled supercomputers. *EDAA*, 978-3-9815370-4-8:1353–1358, 2015.
- [5] The green grid. Available: <http://www.thegreengrid.org>, Accessed: June 2015.
- [6] U.S. Department of Energy. 1.4 pue – the new mandate for federal data center energy efficiency. Available: http://www.efcog.org/wg/ei/events/Aug_22-23_2012_meeting/2012%20ESSG%20Training%20-%20Data%20Center%20Profiler.pdf, Accessed: June 2015.
- [7] Dhruv Varma. Oil submersion cooling for today’s data centers - an analysis of the technology and its business implications. white paper, Green Revolution Cooling, 2014.
- [8] Mark Horowitz, Elad Alon, Dinesh Patil, Samuel Naffziger, Rajesh Kumar, and Kerry Bernstein. Scaling, power, and the future of cmos. *IEEE*, 0-7803-9269-8, May 2005.
- [9] Cray history. Available: <http://www.cray.com/company/history>, Accessed: June 2015.
- [10] Inc. Cray Research. The cray-2 computer system. Available: <http://archive.computerhistory.org/resources/text/Cray/Cray.Cray2.1985.102646185.pdf>, MP-0201, 1985.

- [11] Featured customers. Available: www.grcooling.com/customers, Accessed: June 2015.
- [12] Benefits of asetek liquid cooling for desktops. Available:
asetek.com/desktop/technology/desktop-benefits.aspx, Accessed: June 2015.
- [13] Faster, denser, greener, quieter servers and data centers. Available:
asetek.com/data-center/data-center-coolers.aspx, Accessed: June 2015.
- [14] New ultra-efficient hpc data center debuts. Available:
http://www.nrel.gov/news/features/feature_detail.cfm/feature_id=2133, March 2011.
- [15] Hp supercomputer at nrel garners top honor. Available:
http://nrel.gov/news/features/feature_detail.cfm/feature_id=15403, October 2014.
- [16] Peregrine: Information for users of peregrine. Available:
<http://hpc.nrel.gov/users/systems/peregrine>, February 2015.
- [17] Rich Brueckner. Nrel's peregrine: Hp's warm water cooled hpc platform. Available:
<http://www.insidehpc.com/2014/02/nrels-peregrine-hps-warm-water-cooled-hpc-platform>, February 2014.
- [18] Top500. Available: <https://www.top500.org/lists/2012/06>, June 2012.
- [19] Amanda Bonnie and Doug Hefe. Ta03-1498-341:optimizing cooling. LANL LA-UR-13-26011, July 2013.
- [20] Yanwei Zhang, Yefu Wang, and Wang Xiaorui. Testore: Exploiting thermal and energy storage to cut the electricity bill for datacenter cooling. *IFIP*, 978-901882-48-7:19–27, 2012.
- [21] Pacific Power. Available: <https://www.pacificpower.net/about/rr/rpc.html>, August 2014.

- [22] LANL Presentation. Trinity advanced technology system overview. Available: http://www.lanl.gov/projects/trinity/_assets/docs/trinity-overview-for-web.pdf, Accessed: June 2015.
- [23] Shaoming Chen, Yue Hu, and Lu Peng. Optimization of electricity and server maintenance costs in hybrid cooling data centers. *IEEE*, 978-0-7695-5028-2:526–533, 2013.
- [24] Steve Wetterling and Pat Barrett. Modeling temperature driven wearout rates for electronic components. Available: servernrger.com/Resources/Modeling_Temperature_Driven_Wearout_Rates_for_Electronic_Components_b.pdf, Accessed: June 2015.
- [25] Operating temperature vs. system reliability. Available: www.firepower-technology.com/support-resources/system-integrator-resources, Accessed: June 2015.
- [26] Pete Beckman, Kamil Iskra, Kazutomo oshii, and Susan Coghlan. The influence of operating systems on the performace of collective operations at extreme scale. *Argonne National Laborator*, W-31-109-ENG-38.
- [27] Alessandro Morari, Roberto Gioiosa, Robert Wisniewski, Bryan Rosenburg, Todd Inglett, and Mateo Valero. Evaluating the impact of tlb misses on future hpc systems. *IEEE*, 1530-2075(12):1010–1021, 2012.
- [28] Chillydyne. Liquid cooling technology [online]. Available: <http://www.chillydyne.com/>, 2015.
- [29] Intel xeon processor e7-4870v2 specifications. Available: http://ark.intel.com/products/75250/Intel-Xeon-Processor-E7-4870-v2-30M-Cache-2_30-Ghz, Accessed: June 2015.

- [30] Dell poweredge r920 systems owner's manual, rev. a00, 2014-03.
- [31] Alfred Guerra. Photographic images, April 2015.
- [32] PLUG LOAD Solutions. 80 plus certified power supplies and manufactures. Available: <http://www.plugloadsolutions.com/80PlusPowerSupplies.aspx>, May 2015.
- [33] Pistonless pump technology. Available: <http://www.rocketfuelpump.com>.
- [34] The linpack benchmark. Available: <http://www.top500.org/project/linpack/>, June 2015.
- [35] How do i tune my hpl.dat file? Available: <http://www.advancedclustering.com/act-kb/tune-hple-dat-file/>, September 2014.
- [36] Jeffery Kuehn, Stephen Poole, Stephen Hodson, Josh Lothian, Matthew Baker, Johnathan Dobson, David Reister, Nicholas Lewkow, Ross Glandon, and Jacob Peek. Systemburn. Available: <https://github.com/jlothian/systemburn>, April 2015.
- [37] Amanda Bonnie. Ldms@lanl. Unpublished. Internal source at LANL, May 2015.
- [38] James Primary: Brandt and Ann Gentile. ldms. Available: <https://github.com/ovis-hpc/ovis/tree/master/ldms>.
- [39] Srinivas Pandruvada. Running average power limit – rapl. Available: <https://01.org/blogs/tlcounts/2014/running-average-power-limit---rapl>, June 2015.
- [40] Vince Weaver and Romain Dolbeau. rapl-read.c source code. Available: <http://web.eece.maine.edu/~vweaver/projects/rapl/rapl-read.c>, November 2013.