

FAILURE ANALYSIS OF DIRECT LIQUID COOLING SYSTEM IN DATA CENTERS

Sami Alkharabsheh, Bharath Ramakrishnan, Bahgat Sammakia

Department of Mechanical Engineering
Binghamton University
Binghamton, NY, USA

ABSTRACT

In this paper, the impact of direct liquid cooling (DLC) system failure on the IT equipment is studied experimentally. The main factors that are anticipated to affect the IT equipment response during failure are the CPU utilization, coolant set point temperature (SPT) and the server type. These factors are varied experimentally and the IT equipment response is studied in terms of chip temperature and power, CPU utilization and total server power. It was found that failure of the cooling system is hazardous and can lead to data center shutdown in less than a minute. Additionally, the CPU frequency throttling mechanism was found to be vital to understand the change in chip temperature, power, and utilization. Other mechanisms associated with high temperatures were also observed such as the leakage power and the fans speed change. Finally, possible remedies are proposed to reduce the probability and the consequences of the cooling system failure.

NOMENCLATURE

CDM	coolant distribution module
CRAH	computer room air handler
CW	chilled water
DAQ	data acquisition
DLC	direct liquid cooling
DRAM	dynamic random access memory
FO	fail open
HDD	hard disk drive
IPMI	intelligent platform management interface
NO	normally open
NC	normally closed
PCV	proportional control valve
PDU	power distribution unit
PGW	propylene glycol water
PUE	power usage effectiveness
RTD	resistance temperature detectors

RU	rack unit
SNMP	simple network management protocol
SPT	set point temperature
T	temperature (°C)
U	CPU utilization

INTRODUCTION

Cooling was reported to consume 30%-40% of the total energy used in legacy data centers [1]. Despite the optimistic data center energy consumption trends in a recent survey [2], the average power usage effectiveness (PUE) was reported to be 1.8 mostly due to inefficiencies in the cooling system [3]. Several thermal management technologies are used in data centers cooling to address the inefficiency challenges [4]. The notion of cooling electronic systems using liquids is not novel; however, potential leaks and the capital cost have greatly restricted its application in real data centers [5-7].

In direct liquid cooling (DLC) technology a liquid cooled cold plate is situated on top of a chip, which reduces the thermal resistance between the chip junction and the cooling source. This provides an opportunity to enhance the thermal efficiency of the cooling system. A study showed that for computationally equivalent clusters, DLC can save 45% energy compared with standard hot/cold aisle air cooling [8]. The low thermal resistance promoted designs to operate at higher coolant temperatures (known as warm-water cooling) in which a water side economizer utilizing ambient temperature replaces a chilled water system. This leads to savings of more than 90% compared to conventional air cooling systems [9,10]. The big potential energy savings using DLC encouraged the industry to develop commercial products which may be used in data centers [11,12].

Few studies in literature have focused on system level analysis of DLC. These studies mainly investigated the thermal performance such as the effect of coolant inlet temperature, flow rate, ambient temperature, and chip power on the cooling of the

DLC system [13-15]. The thermal performance research is mainly focused on increasing the electronics power absorbed by the liquid cooling to increase the cooling system efficiency. The pressure drop of a DLC system is also investigated in literature to reduce the pumping power and the cooling system energy consumption [16]. The characteristics of DLC system under exceptional failure situations are crucial to estimate the downtime and provide intelligent remedies to diminish it. The topic of DLC system failure and its impact on the IT equipment have not been investigated in literature.

This research addresses the research gap in characterizing the DLC system experiencing failure for the first time. Collaborative thermal engineering and computer science experimental measurements are used to study the impact of the cooling failure on the IT equipment. The major outcome from this work is to determine the available time in case of failure to prevent losing the computing capability of the IT equipment. The factors which are anticipated to affect the failure behavior are: the server CPU utilization, server type and cooling set point temperature (chilled water cooling versus warm water cooling). These factors are addressed in this study.

DLC SYSTEM DESCRIPTION

The tested system is a rack level direct liquid cooling (DLC) solution for data centers [17]. This system consists of two main loops: the primary loop and the secondary loop. The purpose of the primary loop is to carry the heat from the rack and to dispose of it in the environment outside the data center. It consists mainly of a flow regulating valve, a facility side piping and a dry cooler (for chillerless systems) or a chiller (for chilled water systems). The primary loop in this study is connected to an existing building chilled water system. The supply/return lines from the main building chilled water system to the rack under investigation were modified by placing a novel control system. The control system provides the flexibility to experimentally simulate the DLC behavior under different set point temperatures (SPT).

The purpose of the secondary loop is to carry the heat from the chips inside the servers (via direct contact cold plates) and to dispose of it in the primary loop. The coolant in the secondary loop is propylene glycol (15%) water (85%) mixture. The secondary loop consists of three modules: (a) coolant distribution module (CDM) (b) supply and return manifolds module (c) server module. The CDM in Figure 1a contains a liquid to liquid heat exchanger, two recirculation pumps in series (for redundancy), coolant reservoir, flow meter, and temperature and pressure sensors at the supply and return lines.

The manifold module in Figure 1b is six feet long with square cross section and is made of a stainless material. The manifold can accommodate 42 server modules connected in parallel using dry-break quick disconnect sockets. The server module in Figure 3c consists of corrugated hoses, fittings, dry-break quick disconnect plugs and the microchannel cold plate component. The cold plate component consists of a plastic cover and a copper microchannel heat sink, as shown in Figure 1d. The copper heat sink contains a V-groove to split the impinging jet between two

sink contains a V-groove to split the impinging jet between two sets of parallel channels.

EXPERIMENTAL FACILITY

The experimental facility is located inside the Binghamton University data center laboratory. The data center contains 41 racks divided between three cold aisles with a total area of 215 m². A rack level DLC solution is installed in one of the racks to perform this experiment, as shown in Figure 2.

The DLC rack is equipped with 14 liquid cooled servers of three different types, as indicated in Table 1. Type A-C is used as a notation for the servers in this paper for simplicity. From a thermal perspective, the major difference between these servers is the thermal design power (TDP) for each server. The TDP of Type A, B and C are 160W, 95W and 80W, respectively.

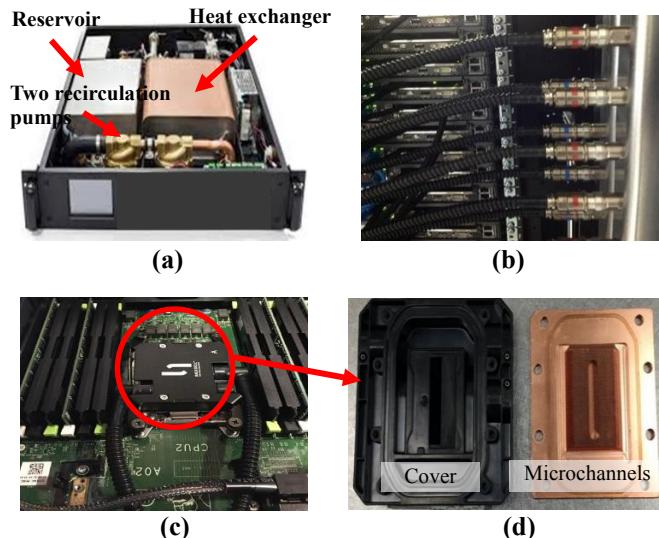


Figure 1: The three modules of the secondary loop: (a) coolant distribution module (b) manifold module (c) server module (d) microchannel cold plate component

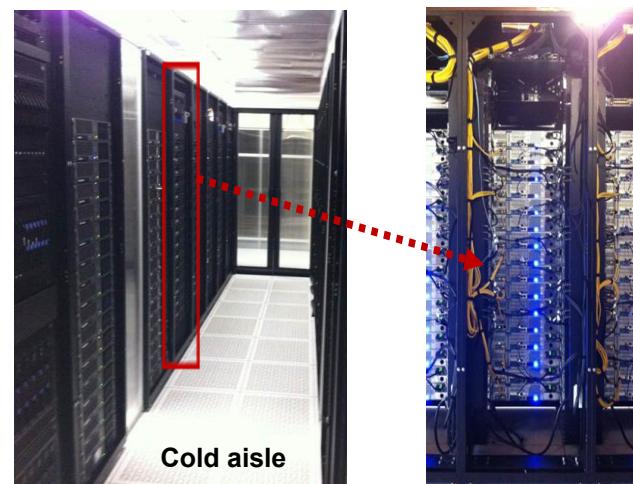


Figure 2: Liquid cooled rack inside the Binghamton University data center laboratory

Therefore, for a certain CPU utilization the dissipated power from each server model will be different. The servers are connected with the cooling system in parallel, which indicates that each server is connected independently with the cooling system distributor (manifold). The liquid system is a retrofit cooling solution that was installed on air cooled servers. The cold plates were used to cool the CPUs only while the remaining components of the IT equipment (DRAM, PSU, HDD, etc.) were maintained air cooled.

Table 1: Specs of the servers in the DLC rack

Server	Processor	Quantity	Type
Dell PowerEdge R730	Intel Xeon (E5-2687WV3)	2	A
Dell PowerEdge R520	Intel Xeon (E5-2440)	9	B
Dell PowerEdge R520	Intel Xeon (E5-2430V2)	3	C

A novel design for the primary side of the DLC system is used to conduct this research. A proportional control valve (PCV) is used to control the primary side inlet temperature, as shown in Figure 3. For instance, if the DLC rack primary loop inlet temperature is sought to be higher than the building chilled water (CW) supply temperature, a signal is initiated to the actuator of the PCV. The PCV motor gradually closes the valve, forcing the return flow from the DLC rack to merge with the supply flow. This leads to an increase in the supply temperature of the rack based upon the flow rate of the forced return flow to the supply line.

The fluid temperature in the primary loop is measured using immersion liquid temperature sensors (i.e. in direct contact with the fluid) with a Platinum RTD element. The temperature range for the temperature sensors is -1°C to 121°C (30°F to 250°F) and the accuracy can be calculated using $[\pm (0.3^{\circ}\text{C} + 0.005 \times |T^{\circ}\text{C}|)]$ as indicated by the manufacturer. This results in accuracy of $\pm 0.4^{\circ}\text{C}$ and $\pm 0.53^{\circ}\text{C}$ at 20°C and 45°C , respectively. The flow rate is measured using an inline direct beam path wetted ultrasonic sensors utilizing differential transit time velocity measurement. The flow rate range is 0.6-15 GPM with $\pm 1\%$ accuracy. The proportional control valve (PCV) assembly consists of a two-way valve of Powermite 599 Series type plus an MT Series SSC electronic valve actuator. The PCV is chosen such that the valve is NC (normally closed) and FO (fail open) to ensure that the PCV will not halt the flow of the primary side to the DLC racks if the valve fails. The output signals of the temperature and flow sensors, and the PCV are connected to a central building management system for data logging and control.

The fluid temperature in the secondary loop is measured using Measurement Specialties 10K3D682 temperature sensors. The accuracy of these sensors is 2.45% according to the manufacturer specs. The flow rate is measured using an inline Adafruit flow meter. The flow meter has a magnet attached to a pinwheel and a magnetic sensor on the flow meter tube. The flow rate is measured by the number of spins the pinwheel makes. The

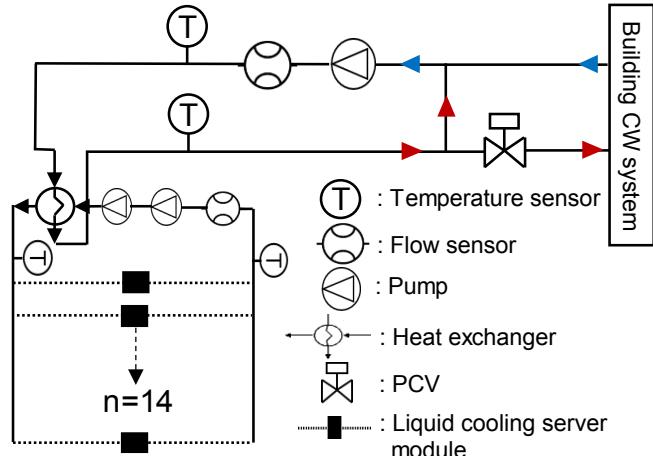


Figure 3: Schematic of the cooling system and the installed sensors

flow meter is calibrated for this application with 3% accuracy. The secondary loop sensors are connected to an integrated control and monitoring system inside the CDM.

Internal sensors of the servers and performance counters are used to measure the fans speed, chips power and temperature, and CPU utilization. The readings of the internal sensors and performance counters are retrieved at the server level using the intelligent platform management (IPMI) interface and transmitted to a network management system via TCP/IP. Power distribution units (PDU) are used to measure the total power of servers. The PDUs are connected to the data center network, therefore the readings can be retrieved using a network management system.

A Linux based tool is used to retrieve the data from the building management system, the intelligent platform management (IPMI) interface and the PDU network interface [18]. The tool utilizes a simple network management protocol (SNMP) and runs on an administrative machine in the data center. The data is collected in less than a second-time step and eventually becomes available to the user in an Excel sheet format.

FAILURE SCENARIOS

Failure can occur in the primary or secondary loops of the DLC system, which leads to a loss of cooling, loss of flow, or loss of flow and cooling combined. Power outage, mechanical failure and human error are the primary reasons responsible for failures.

This study focuses on investigating a complete and partial loss of flow in the secondary loop. The pumps failure due to mechanical failure or power outage is the main cause that is responsible for loss of flow. The secondary loop consists of two recirculation pumps in series, as shown in Figure 1a. The complete loss of flow is simulated experimentally by shutting down the two recirculating pumps, which declines the flow to zero, as shown in Figure 4a. The partial loss of flow can occur due to failure in one of the two recirculating pumps. It is simulated experimentally by shutting down one of the pumps and

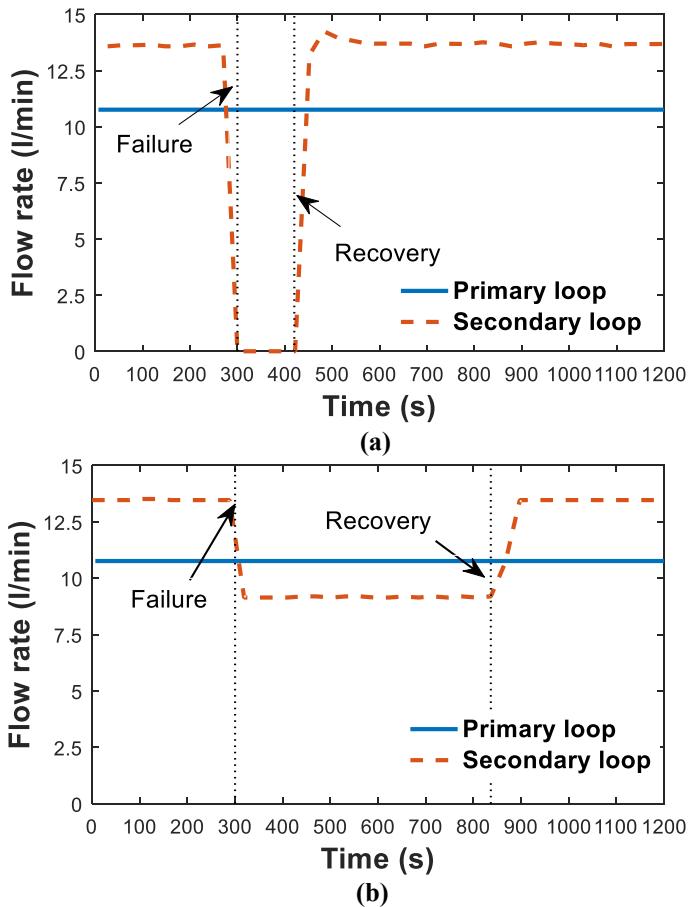


Figure 4: The flow rate scheme simulating failure in the secondary loop (a) complete loss of flow scenario (all pumps failure) (b) partial loss of flow scenario (single pump failure) in the secondary loop

retaining the other pump operative declining the flow to approximately 68% of the original flow rate, as shown in Figure 4b. The recirculating pumps were turned back on after the chip temperatures reach a certain limit to avoid causing damages to the IT equipment due to overheating. The recirculating pump in the primary loop was maintained operative at all times during the experiments.

The factors that are anticipated to influence the system response due to the loss of flow failure are the CPU utilization (affects the chip power), cooling set point temperature (SPT) and the server model. These factors are varied in this study to understand their influence on the IT equipment response. The IT equipment response is determined by measuring the rate of change of the chip temperature, change in the server computing capability (utilization), change in chip power, and change in the server total power.

RESULTS AND ANALYSIS

This section presents the experimental results for the complete failure and partial failure modes. The presented results show the effect of each failure mode on the chip temperature and power, CPU utilization, fans speed, and the total server power.

Complete failure

The complete failure of the pumps leads to a complete loss of flow. The presented experimental results are conducted at different CPU utilization, SPT, and different type of servers. At each coolant SPT of 20°C and 45°C, the CPU utilization was changed to 100%, 25% and 0% (idling state). The chosen coolant SPTs are intended to compare the response of a chilled water DLC system (20°C) with a warm water DLC system (45°C) experiencing a failure incident. Additionally, the DLC rack is equipped with three different models of servers, which allows for studying their behavior under the failure mode.

The simulated failure mode starts at time of 300 seconds for all cases by shutting down the recirculation pumps. This leads to completely halting the flow in the cold plates that are cooling the CPUs inside the servers. During failure, the forced convection cooling through the cold plate microchannels turns into natural convection in the liquid side and conduction through the PCB board to air. This change in the cooling mechanism increases the thermal resistance and thus the chip junction temperature.

The chip temperature of Type A server increases momentarily after failure then stabilizes at a certain temperature, as shown in Figure 5a. The chip state at which the temperature stabilizes is called the throttling state. The chip throttles its frequency when it undergoes an excessive temperature increase to reduce the generated power, which lowers the chip temperature. This can be noticed in Figure 7b by observing the reduction in the chip power (marked as 1) during failure. The throttling state on the other hand reduces the CPU utilization which makes the CPU computing capability inefficient. This is demonstrated experimentally in Figure 7c (marked as 1). It should be noted that the CPU utilization exceeds the 100% in the presented data due to the Intel turbo boost technology. Chip power and CPU utilization drops are barely noticeable for the idling state since the utilization at idling is almost 0%.

Figure 5a shows that the chip temperature rate of change depends on the CPU utilization. The time that the chip takes to start throttling after failure is 23s, 56s and 305s for 100%, 25% and idling utilization, respectively. This indicates that the available time before the CPU frequency throttles at 25% utilization is 1.4 times that at 100% utilization and at idling state is 4.4 times that at 25%. The coolant SPT is found to have a small impact on the server response experiencing failure, as shown for SPT of 20°C in Figure 6a. The time that the chip takes to start throttling after failure in this case is 34s, 82.1 and 353s for 100%, 25% and idling utilization, respectively. This indicates 48%, 47% and 16% increase in the available time before throttling starts for 100%, 25% and idling utilization, respectively.

It can be concluded from the aforementioned results that the CPU utilization has a more significant impact on the server response in case of failure than the cooling SPT. However, ultimately if the server has a computing load more than 25% utilization, it takes the CPU less than a minute to throttle. At throttling state, the server practically is not effective as its processing capability deteriorates. Also, the high chip temperature affects the reliability of the chip eventually if it is maintained for a long time. Therefore, the solution for this mode

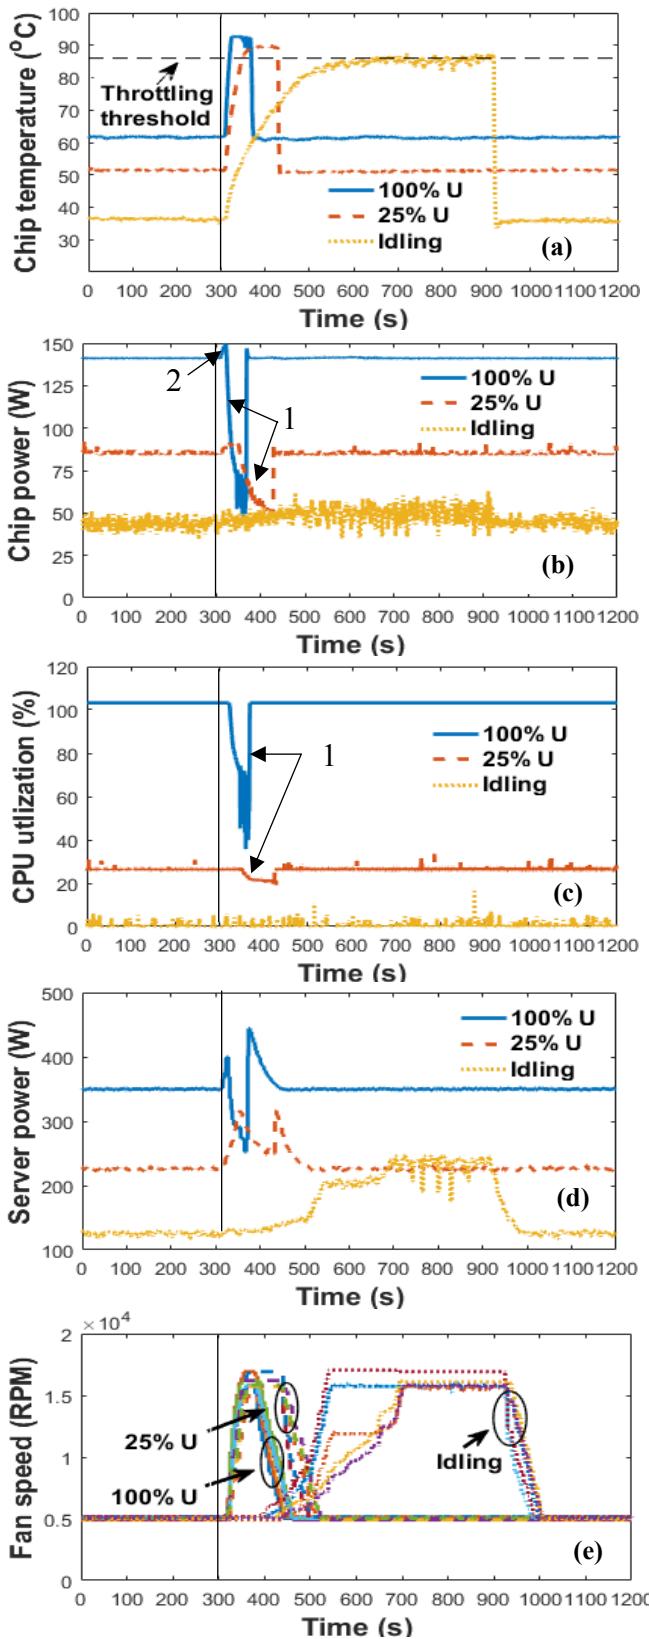


Figure 5: Effect of failure at 45°C SPT on (a) chip temperature (b) chip power (c) CPU utilization (d) fans speed (e) server power (line indicates when failure was initiated)

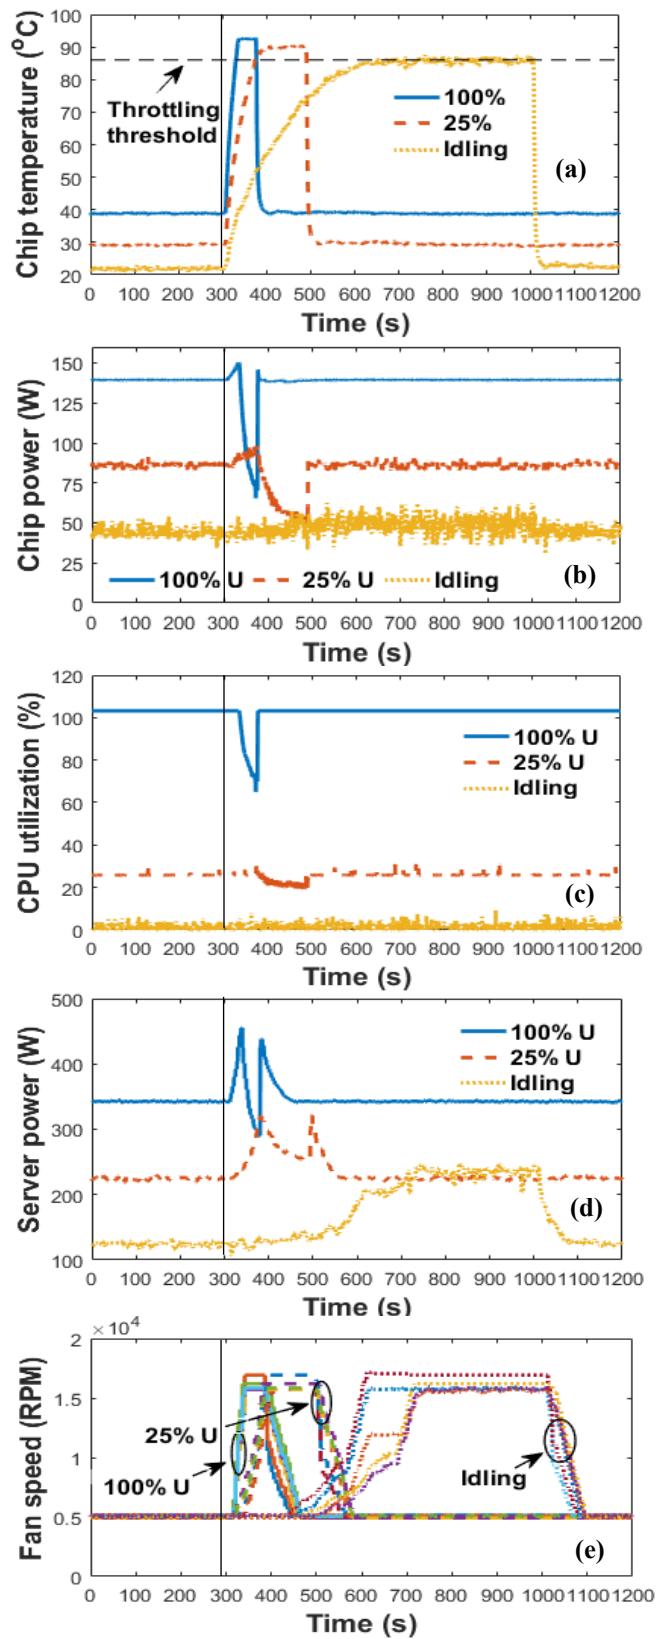


Figure 6: Effect of failure at 20°C SPT on (a) chip temperature (b) chip power (c) CPU utilization (d) fans speed (e) server power (line indicates when failure was initiated)

of failure should be proactive, meaning that momentarily sensing the change in temperature and acting in less than 23s for the worst-case scenario (100% utilization, 45°C SPT). Remedies for this mode of failure are proposed later in this paper.

In addition to the chip temperature and the associate throttling behavior, leakage power is observed in the experimental data. The leakage power occurs when the chip exhibits high temperatures, which further increases the chip power and temperature. Figure 5b shows the increase in the power while the server is experiencing failure (marked as 2). The chip power keeps rising until the CPU throttling mechanism initiates. Since the power drops when the CPU throttles, the leakage current decreases as well. Whether the CPU utilization is 100% or idling, leakage current is observed after failure. The leakage power increases the chip power by 6%, 10% and 15% for 100%, 25% and idling state utilization, respectively. The leakage power is also observed at 20°C as well, as shown in Figure 6b. The leakage current in this case increases the chip power by 7.5%, 12.8% and 37% for 100%, 25% and idling state utilization, respectively.

The server total power is measured in this analysis as well. The PDU readings are assumed to represent the total power consumed by the various server's components, however, the server's power supply unit (PSU) has an efficiency associated with it. Figure 5d shows the total server power for different CPU utilizations. There are three distinct changes in the server power at 100% and 25% CPU utilizations. The server power increases after failure because of the leakage power and the increase in the fans speed. The increase in the fans speed during failure is demonstrated in Figure 5e. When the CPU throttles, the chip power drops as shown in Figure 5b, which reduces the total server power. The server power spikes when cooling is recovered since the fans are still ramped up and the chip power returns to the original level. When the chip cools down to the original level, the total server power stabilizes at the original steady state power. Since the servers does not experience throttling at idling state, the server power never drops during failure.

In the previous analysis, Type A server is used to understand the effect of failure on the IT equipment behavior. We used a Type A server since it has the highest TDP thus presents the worst cooling failure scenario. The behavior of Type B and Type C servers along with Type A is shown in Figure 7.

During a 60 second failure incident, it is noticed that servers of Type A and Type B only reach a temperature at which the CPU frequency throttles. These types are characterized with a higher TDP than Type C. This can be observed by the chip temperature and utilization response in Figure 7a and Figure 7c. The chip power increases then decreases after failure for Type A and Type B servers. This behavior occurs due to the leakage power and then the CPU throttling mechanism. Since Type C server does not reach the throttling state, the chip power only increases due to the leakage power. The server power for Type C server is also distinct compared with Type A and Type B servers. The server power for Type C does not drop during failure since the chip power does not decrease during failure. The increase in

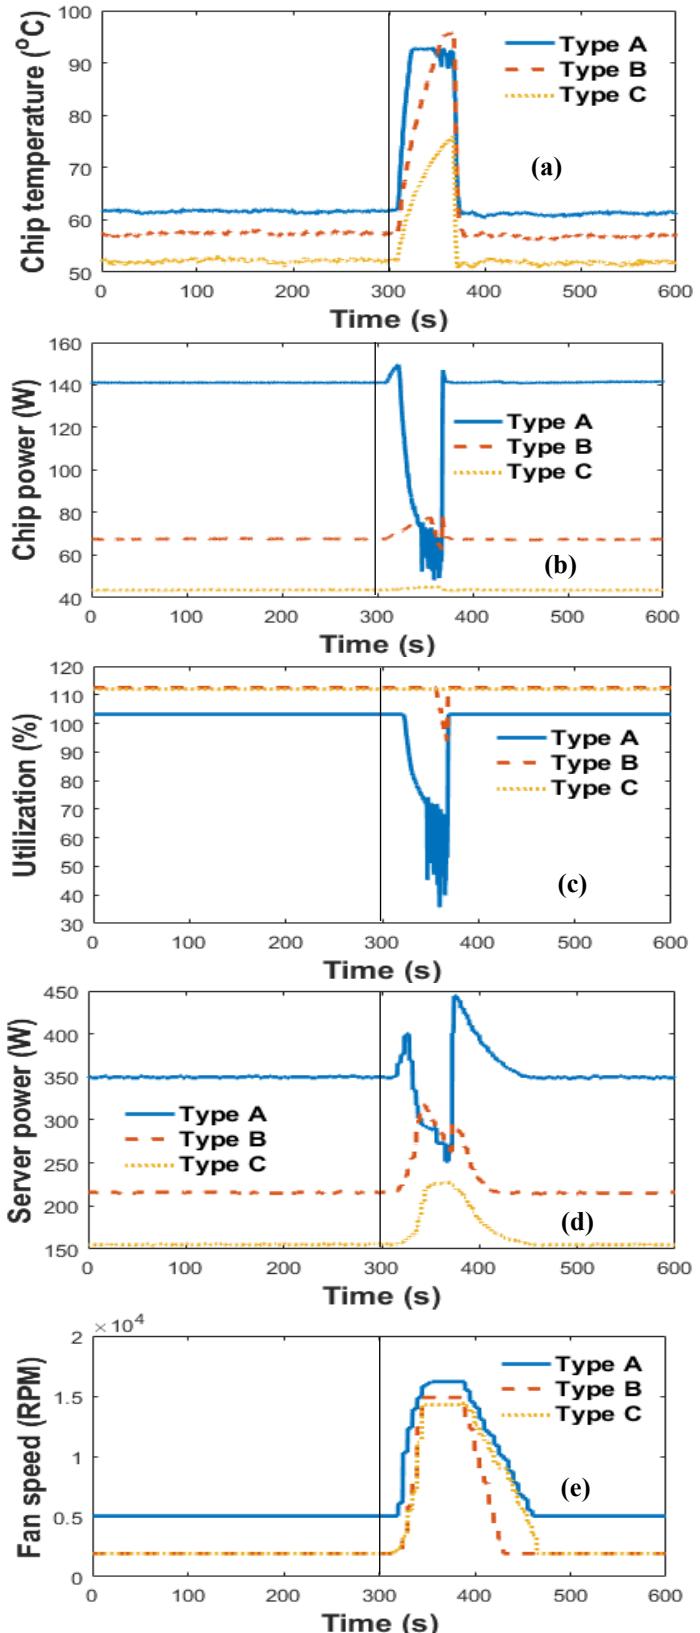


Figure 7: Effect of failure on different types of servers (SPT=45°C, 100% utilization)

the server power is due to the increase in the fans speed, as shown in Figure 7e.

The previous results showed the IT equipment response when the secondary loop of the cooling system fails. The main factors that are anticipated to affect the IT equipment response are the CPU utilization, coolant SPT and the server type. These factors are varied experimentally and the IT equipment response is studied in terms of chip temperature and power, CPU utilization, and total server power. The CPU utilization and the type of servers were found to have a more noticeable effect on the IT equipment during failure than the coolant SPT. Additionally, the CPU frequency throttling mechanism was found to be crucial in understanding the change in chip temperature, power, and utilization. Other mechanisms associated with high temperatures were also observed such as the leakage power and increasing the fans speed.

It is evident that this failure mode is hazardous. There is a high potential that a data center will go offline almost instantaneously if this mode of failure occurs. Data center designers should be aware of the sequences of this failure mode and attempt to avoid it from happening. In a later section, we propose possible remedies that can reduce the possibility of this mode of failure and a solution to cope with it if it occurs.

Partial failure

Two recirculating pumps are used in the secondary loop to provide cooling for the servers. The two pumps are connected in series and are permanently operative. In the partial failure mode, one of the pumps only failed to simulate a partial loss in flow. The partial failure leads to a 32% reduction in the flow rate of the secondary loop, as opposed to a complete loss of flow in the complete failure mode. The failure is initiated at time 300 seconds. The entire servers in the rack are at 100% CPU utilization and the coolant SPT is 45°C, representing the worst-case scenario.

Interestingly, it is found that the partial failure does not affect the IT equipment, as shown in Figure 8. The chip temperature, CPU utilization, chip and server power, and fans' speed are stable before and after failure. The chip temperature is maintained well below the throttling temperature of 87°C, that is also proven by the absence of the CPU utilization drop. The unchanged chip temperature also leads to a constant fans' speed and chip power.

An experiment is conducted to understand the effect of the coolant flow rate in the secondary loop on the cooling of a server, as shown in Figure 9a. In this experiment, the coolant inlet and outlet temperatures are measured for a Type A server using J-Type thermocouple probes. The coolant flow rate in the server module is measured using OMEGA FTB-314D micro-flow meter with accuracy of ± 0.18 l/min. The junction temperature and chip power are measured using the server registry data via IPMI. The experiment is conducted at 45°C coolant SPT and 100% CPU utilization for a Type A server, while the coolant flow rate is varied.

Figure 9b presents the calculated sensible power using the measured coolant inlet and outlet temperatures, and flow rate.

The sensible power (power removed by the liquid coolant) is calculated using the energy balance equation on the coolant side. During normal operation, the sensible power is 105.1W indicating that the liquid coolant removes 75.1% of the total chip power. It is anticipated that the air side removes the remaining chip power, as the server fans were not removed when the liquid cooling module was installed in the server. The chip power can be transferred to the air side via conduction from chip to board then convection from board to air or through the exposed area of the cold plate to the air. Since 75.1% of the chip power is removed by the liquid coolant, the thermal resistance from chip junction to liquid coolant inlet is small compared with chip junction to the air side. The chip junction temperature at normal operation is 61.1°C, as shown in Figure 9b.

After failure of a single pump, the flow rate for the tested server coolant loop drops from 1.1 l/min to 0.75 l/min. By varying the liquid coolant flow rate from 0.3 l/min to 1.15 l/min, it can be noticed in Figure 9b that the sensible power varies approximately from 67% to 75% of the total chip power. The reduction in the flow rate to 0.75 l/min due to partial failure will have a small impact on the cooling of the chip, which explains why the partial failure has a negligible impact on the IT equipment in Figure 8. Furthermore, the negligible impact of partial failure on the IT equipment can be clearly observed by the

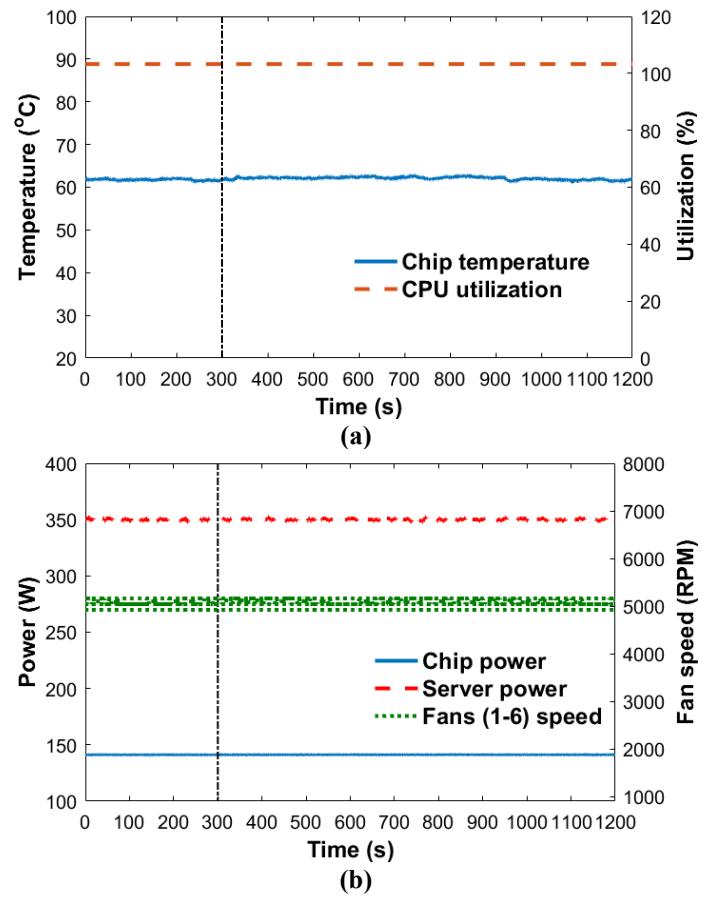
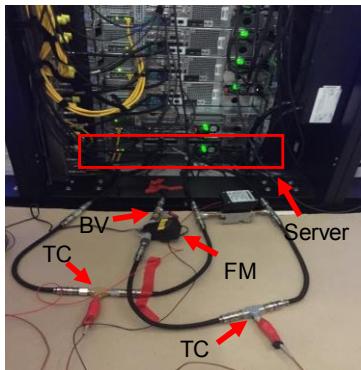
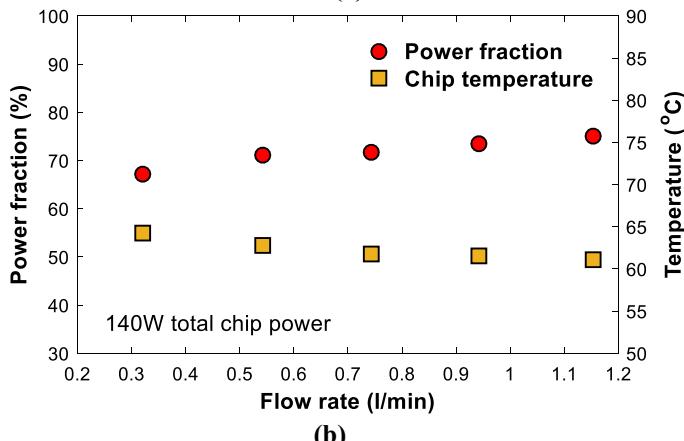


Figure 8: Effect of partial failure on (a) chip temperature and utilization (b) chip power, server power and fans speed



(a)



(b)

Figure 9: (a) Sensible power measurement (TC: thermocouple, FM: micro flow meter, BV: ball valve) (b) The fraction of chip power removed by the liquid coolant and the chip junction temperature at different liquid coolant flow rates

chip temperature values of 61.1°C and 61.8°C at 1.1 l/min and 0.75 l/min, respectively, as shown in Figure 9b.

The observation that the single pump failure does not have an impact on the cooling of the servers is not generic. In this study, 15 server modules are used to cool 14 2RU-servers in the experimental setup, however, the DLC system is equipped with outlets that are sufficient to accommodate 42 server modules. In a hypothetical case of using 42 server modules, each server module would receive 0.32 l/m during normal operation and 0.22 l/min during partial failure. The flow rate values are calculated by assuming that the pumps maintain the same operational point when 42 server modules are used instead of 15 and the flow is evenly distributed between the 42 server modules. The flow rate drop from 0.32 l/m to 0.22 l/m is expected to increase the chip temperature in a more noticeable manner than the flow rate drop from 1.1 l/min to 0.75 l/min based on the data shown in Figure 9b. Nevertheless, the chip temperature will still be below the throttling temperature of 87°C and the partial failure will not affect the IT equipment computing capability.

In summary, the partial failure does not affect the IT equipment computing performance even when the cooling system is at maximum capacity of server modules.

PROPOSED REMEDIES

In this section, possible remedies are proposed to reduce the possibility of the cooling system failure and the sequences in case failure occurs.

The proposed remedies are stemmed from the root cause of failure. The loss of flow due to pump failure can be due to power outage. Data center IT equipment are connected to an uninterrupted power supply system (UPS) to ensure continuous operation in case of a power outage. On the other hand, the cooling system is connected to a backup generator, which takes a longer time than a UPS system to kick in. A big interest to data center designers is to estimate the time that IT equipment can be maintained operational on a UPS system without overheating until the backup generator reactivates the cooling system. From the results of this research, if a DLC system is used for cooling, such a power system scheme (IT equipment on a UPS system and cooling system on backup generators) will lead to data center shut down. Therefore, it is recommended to connect the pumps of the secondary loop of a DLC system to a UPS system to avoid this scenario from happening.

Mechanical failure of the pumps is another root cause that could lead to loss of flow in the secondary loop of a DLC system. The recommended action to diminish this issue is pump redundancy. If a pump fail, another pump becomes available to provide the required flow. Also, continuous maintenance for the pumps would reduce the possibility of mechanical failure.

Load migration is a mechanism used by computer scientists to move the computing load between different physical machines for different purposes [19]. This idea can be used to cope with cooling failure if it happens. The presented results are encouraging to use this technique, as the results showed that despite the complete loss of liquid cooling at idling state, the chip does not go through excessive temperature. For instance, if a server is at a 100% utilization and the cooling systems fails, the computing load can be moved (migrated) to a server that is under normal operational conditions. This scenario is simulated experimentally in Figure 10. The chip temperature increases drastically after failure. If the load is removed, meaning that the CPU utilization changed from 100% to almost 0% (idling), the chip temperature operates within the safe limits. The challenge however is in migrating the load in a short time. The other advantage of load migration is that the load can be moved from server to server, rack to rack and from a data center to a data center.

SUMMARY AND CONCLUSIONS

This study presented an experimental investigation of the impact of failure in a DLC cooling system on the IT equipment. The CPU utilization, coolant SPT, and the server type are varied and the IT equipment response is studied in terms of chip temperature and power, CPU utilization and total server power.

The CPU utilization and the type of servers were found to have a more significant impact on the IT equipment during failure than the coolant SPT. Additionally, the CPU frequency throttling mechanism was found to be crucial in understanding the change in chip temperature, power, and utilization. Other

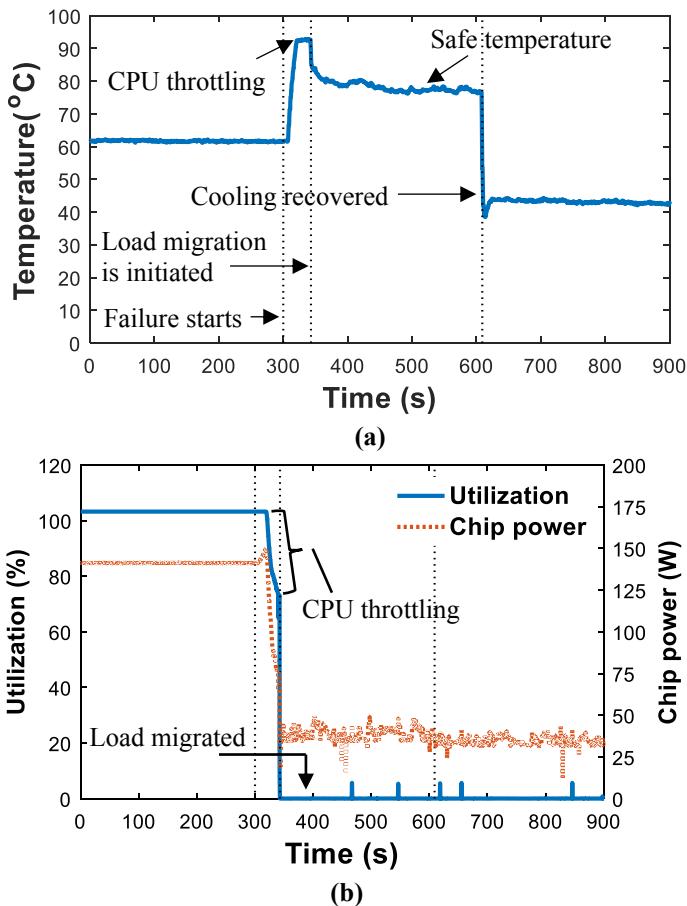


Figure 10: Proposed remedy for failure using load migration
(a) chip temperature (b) CPU utilization and chip power

mechanisms associated with high temperatures were also observed such as the leakage power, and increasing the fans speed.

Ultimately, it is evident that the DLC system failure is hazardous. There is a high potential that a data center will go offline almost instantaneously if this failure occurs. Data center designers should be aware of this consequence and take the proper precautions. We proposed possible remedies that can reduce the possibility of this mode of failure and a solution to cope with it if it occurs. It is recommended to connect the pumps of the secondary loop of a DLC system to a UPS system and use pump redundancy to avoid failure due to power outage and mechanical failure. Additionally, proactive implementation of load migration technique can protect the computing data if this failure actually happens.

ACKNOWLEDGMENTS

We would like to thank Dr. Kanad Ghose and Udaya L N Puvvadi for their valuable assistance. This work is supported by NSF IUCRC Award No. IIP-1134867.

REFERENCES

- [1] M. Iyengar, "Energy Consumption of Information Technology Data Centers," *Journal of Electronics Cooling*, vol. 16, no. 4, 2010.
- [2] A. Shehabi, S. Smith, D. Sartor, R. Brown, M. Herrlin, J. Koomey, E. Masanet, N. Horner, I. Azevedo and W. Lintner, "United States Data Center Energy Usage Report," Ernest Orlando Lawrence Berkeley National Laboratory (LBNL), 2016.
- [3] M. Stansberry and J. Kudritzki, "Uptime Institute 2012 Data Center Industry Survey," white paper, Uptime Institute, 2013.
- [4] S. Alkharabsheh, J. Fernandes, B. Gebrehiwot, D. Agonafer, K. Ghose, A. Ortega, Y. Joshi and B. Sammakia, "A Brief Overview of Recent Developments in Thermal Management in Data Centers," *Journal of Electronic Packaging*, vol. 137, no. 4, p. 040801 (19 pages), 2015.
- [5] D. L. Beaty, "Liquid Cooling--Friend or Foe," *ASHRAE Transactions*, vol. 110, no. 2, 2004.
- [6] M. J. Ellsworth, L. A. Campbell, R. E. Simons, M. Iyengar and R. R. Schmidt, "The Evolution of Water Cooling for IBM Large Server Systems: Back to the Future.," in *Thermal and Thermomechanical Phenomena in Electronic Systems (ITHERM)*, 2008 IEEE Intersociety Conference on, Lake Buena Vista, Florida, 2008.
- [7] "2014 Data Center Industry Survey," UptimeInstitute, 2014.
- [8] M. Ellsworth and M. Iyengar, "Energy efficiency analyses and comparison of air and water cooled high performance servers," in *ASME 2009 InterPACK Conference* collocated with the *ASME 2009 Summer Heat Transfer Conference* and the *ASME 2009 3rd International Conference on Energy Sustainability*, San Francisco, CA, USA, 2009.
- [9] M. Iyengar, M. David, P. Parida, V. Kamath, B. Kochuparambil, D. Graybill, M. Schultz, M. Gaynes, R. Simons, R. Schmidt and T. Chainer, "Extreme energy efficiency using water cooled server inside a chiller-less data center," in *IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems*, San Diego, CA, USA, 2012.
- [10] M. David, M. Iyengar, P. Parida, R. Simons, M. Schultz, M. Gaynes, R. Schmidt and T. Chainer, "Impact of operating conditions on a chiller-less data center test facility with liquid cooled servers," in *IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems*, San Diego, CA, USA, 2012.
- [11] R. Miller, "Rise of Direct Liquid Cooling in Data Centers Likely Inevitable," *Data Center Knowledge*, 2014.
- [12] "New Liquid Cooled HPC Cluster Launched," *Electronics Cooling*, 2016.

- [13] L. Parnell, D. Demetriou and E. Zhang, "Combining Cooling Technology and Facility Design to Improve HPC Data Center Energy Efficiency," in Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), 2016 15th IEEE Intersociety Conference on, Las Vegas, NV, USA, 2016.
- [14] H. Coles and S. Greenberg, "Direct Liquid Cooling for Electronic Equipment," Ernest Orlando Lawrence Berkeley National Laboratory, 2014.
- [15] T. Cader, L. Westra, A. Marquez, H. McAllister and K. Regimbal, "Performance of a Rack of Liquid-Cooled Servers," ASHRAE Transactions, vol. 13, no. 1, pp. 101-114, 2007.
- [16] S. Alkharabsheh, B. Ramakrishnan and B. Sammakia "Pressure Drop Analysis in Direct Liquid Cooled (DLC) Rack," in Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), 2017 16th IEEE Intersociety Conference on, Orlando, FL, USA, 2017.
- [17] "Direct Contact Liquid Cooling for the Datacenter- Can it be Simple, Low Cost, High Performance and Efficient?," Electronics Cooling Magazine, 2012.
- [18] T. J. Stachecki and K. Ghose, Short-Term Load Prediction and EnergyAware Load Balancing for Data Centers Serving Online Requests.
- [19] V. Shrivastava, P. Zerfos, K. W. Lee, H. Jamjoom, Y. H. Liu & S. Banerjee, "Application-aware virtual machine migration in data centers," in IEEE INFOCOM, Shanghai, China, 2011.