Flow Disruptions and Mitigation in Virtualized Water-Cooled Data Centers

Udaya Puvvadi, Anuroop Desu, Tyler Stachecki, Sami Alkharabsheh, Kanad Ghose and Bahgat Sammakia
State University of New York, Binghamton, NY 13902
{upuvvad1, adesu1, tstache1, salkhar1, ghose, bahgat}@binghamton.edu

Abstract—**Recent availability of warm water cooling systems that can be easily retrofitted to stock server by replacing the heatsinks with coldplates have made it possible to use such cooling for non-HPC cloud/data center servers. These cooling systems use internal pumps in rack-level heat exchangers as well as external pumps that can fail. We present a systematic study of the pump failures that disrupt flow in the cooling system, propose and experimentally evaluate techniques for reducing service disruptions during failures while avoiding damage to the servers where water cooling has failed.**

Keywords—water-cooled servers; virtual machines; dependable systems; virtual machime migration; gang migration

I. INTRODUCTION

A significant part of the operational energy consumption by a data center is in operating the cooling system. In some data centers that use chilled air cooling, this can be about half of the total energy expenditures [17]. Liquid-based cooling for data centers have emerged to provide an alternative cooling mechanism.

Current liquid cooling solutions include racks with rear door heat exchangers that circulate chilled water [5], in-line coolers that use independent air to water heat exchanger units that stand alone adjacent to high heat flux racks [25] and immersion cooling systems [2, 12] that immerse IT equipment in electrically inert fluids. In some high-performance computing platforms *custom designed* cold plates with chilled water circulation are used to cool IT equipment that are thermally connected to the coldplates [10, 11].

Direct liquid cooling systems (DLCS) have appeared recently in the market at competitive pricing levels. *Historically, the emphasis of DLCS has been to cool HPC systems, but recent product offerings [6] aim to bring this cooling solution the non-HPC cloud/data center market.* DLCS bring liquid directly to the hottest components inside a server such as the CPU chips, GPUs and DRAM DIMMs. These DLCS, including solutions from Asetek [3] and CoolIT [6], use coldplates attached to the top of the CPUs, GPUs and the DRAM and circulate water at normal environmental temperatures ("warm water") through these cold plates to take out the heat. Water with its high thermal conductivity and well-managed flow can be 2000 to 3000 times more effective at removing heat compared to air-cooling solutions.

In a typical retrofittable DLCS, each server has two connections for the circulating water: one to bring in the warm water and another to take out the heated water after it goes through the coldplates, each coldplate having its own built-in pump. A pair of manifolds at the back of the rack provides short connections to the intake and outtake plumbing lines at each server and to/from a rack level heat exchanging unit (**HE**). Cooled water supplied from the rack-level HE to the coldplates via a supply manifold, circulates through the coldplates and is returned to the HE for cooling via a return manifold.

The DLCS system from CoolIT is designed to be retrofitted into conventional servers. Heat sinks used for air cooling on the CPU chips are replaced with coldplates that circulate water. These coldplates are capable of removing well over half of the heat generated by the CPU, with the rest of the CPU heat dissipated into the motherboard through heat conduction via the CPU connections even at water supply temperatures as high as 45 degrees Celsius. Consequently, significant savings can be realized in the cooling cost through a reduction in need for chilled air cooling. The hot water carried off from the CPU chips reject heat via a rack level heat exchanger (Section 2) that houses three pumps. The heat exchanger, in turn, uses water supplied by the facility to transfer the heat to an evaporative cooling tower and thus to the environment. Pump failures in the heat exchanger can lead to a disruption in flows in the coldplates or the facility side and can cause the CPU core temperatures to go up, causing the CPUs to throttle down and ultimately shut down.

The goal of this study is to experimentally evaluate the implications of various pump failure scenarios, systematically classify pump failures into different classes and propose software solutions for dealing with these failures in a virtualized environment to minimize service disruptions through the use of controlled virtual machine migration.

II. BACKGROUND

Fig. 1(a) and 1(b) depicts a DCLS system available currently from CoolIT systems that can be easily retrofitted to a stock 1U or larger form factor server. Existing heatsink assemblies for air cooling are replaced with coldplates that circulate water, permitting cooling to be efficiently directed to the hot CPU chips.

Figure 1(a) shows the two coldplates connected in series for two CPUs and the associated plumbing as used in a 2-socket Dell server. One of these plumbing connections come from an intake manifold at the back of the rack that bring in warm water from a rack-mounted heat exchanger. The other plumbing connection from the series connected server delivers the water, after circulating through the coldplates, to a second manifold - the outtake manifold - that takes the heated water to the rack-level heat exchanger (Figure 1 (b)). The connections through the coldplates, the intake and outtake manifolds and their flow path in the heat exchanger form a loop called **coldpate loop, C-loop**.

978-1-7281-2927-3/19/$31.00 ©2019 IEEE

Fig. 1. Components of the warm water cooling system

Within the rack level heat exchanger, water comes in from the facility side via an external pump, goes through the heat exchanger, taking off heat from the heated water flowing out of the coldplates via the outtake manifold. The heated water from the heat exchanger flows back to the facility side where another facility level heat exchanger such as an evaporative cooler, rejects the heat to the environment. This facility side flow path is called the **facility-side loop, F-loop**.

In the coldpate system used, the flow in the F-loop is maintained by an external pump (Fig. 2). Inside the rack-level heat exchanger, two pumps are used in series for fault-tolerance in the C-loop. Fig. 2, shows a configuration similar to the one used in our assessments. Here, rack-scale heat exchangers, located in a rack can cool multiple racks (two racks in this example). Across an entire aisle (or row, depending on the configuration), a single pump maintains the F-loop flow for all heat exchangers in the aisle (or row).

Warm water DLCS can provide all or the bulk of the necessary cooling for the servers and constitutes the basis of a "compressor-less" cooling system. In general, it has the potential of decreasing the cooling expenditures in existing installations that use chilled air cooling, through a retrofit of the servers, reducing demands on the chilled air cooling system dramatically, which, in turn reduces installation and operational energy costs.

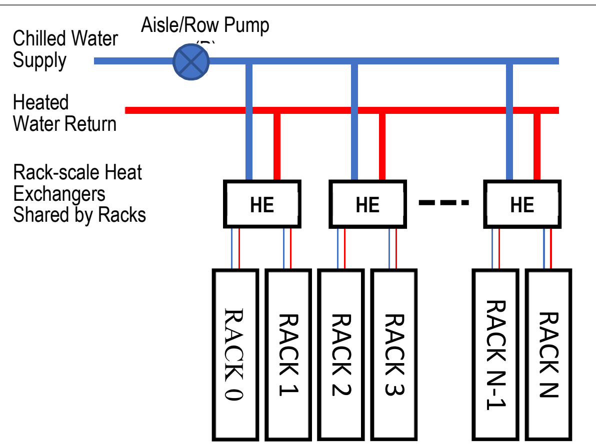

Fig. 2. Aisle pump and rack-scale shared heat exchangers (HEs) – two racks share a HE in this example

III. PUMPS AND FAILURE MODES

A. *Pumps in the Cooling Loops*

Fig. 3, shows a schematic for the flow paths for the F-loop and C-loop and the pump locations. There is a single pump circulating the water from the facility side in the F-loop. On the C-loop, there are two pumps in series to mitigate the impact of the failure of a single pump – if one of these pumps fail, the other pump can maintain sufficient circulation for cooling. The rack-level heat exchanger used supports two racks using two C-loops in parallel, one C-loop for each rack. The studies presented here use a single rack of 16, 2U, 2-socket servers and are extrapolated to two racks as described below.

For the studies reported here, we use a single water-cooled rack of Dell 520 and 730 servers, each with two Intel Xeon CPU chips. Of these, the 730s represent the hottest servers, dissipating about 490 Watts at full utilization of the CPUs, while the peak dissipation of the 520s are about 224 Watts. The rack mounted heat exchanger has a Simple Network Management Protocol (SNMP) interface that is used to read out flow rates, inlet and outlet temperatures for the loops and other statistics. The CPU temperatures, fan speeds, clock rate and other server statistics are read out using a daemon inside each server, sampling the data at intervals of 500ms.

B. *Pump Failure Scenarios*

We now examine the scenarios that can cause water flow failures in the DCLS.

Complete flow disruption in the C-loop: There are two pump failure scenarios that affect flow in the C-loop – complete C-loop disruption happens when both pumps (which are in series) fail, completely interrupting water flow through the coldplates. We simulated this scenario by shutting off the C-loop via signals sent on the SNMP interface to the rack-level heat exchanger and the results are shown in Fig. 4 (a). In this case, CPU throttling is initiated within 35 seconds at 100% CPU utilization and 100% clock rate. (At lower clock frequencies, this time is stretched out gradually as the CPU clock rate is decreased but throttling occurs in less than one minute at the lowest clock frequency – this is not shown in Fig. 4 (a)). Fig. 4 (b) shows how the CPU clock frequency is throttled by the

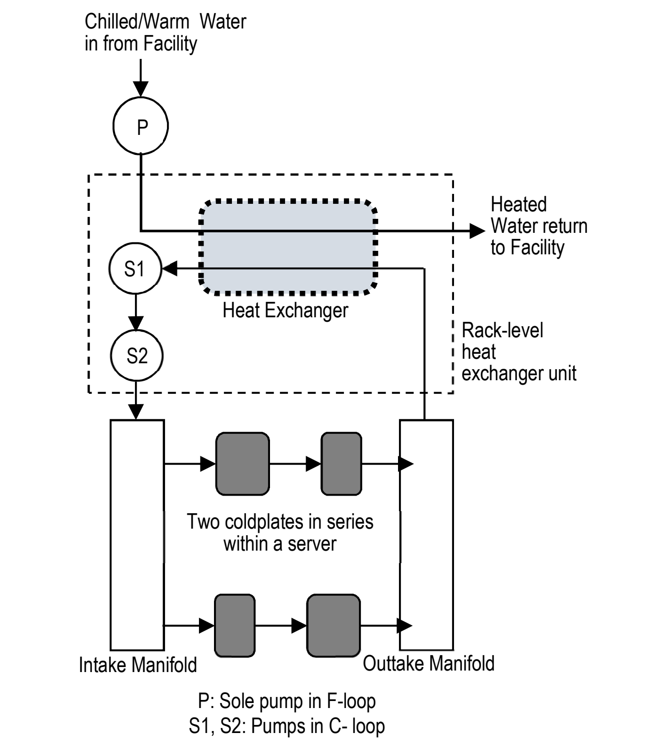

Fig. 3. Water flow paths and pump locations

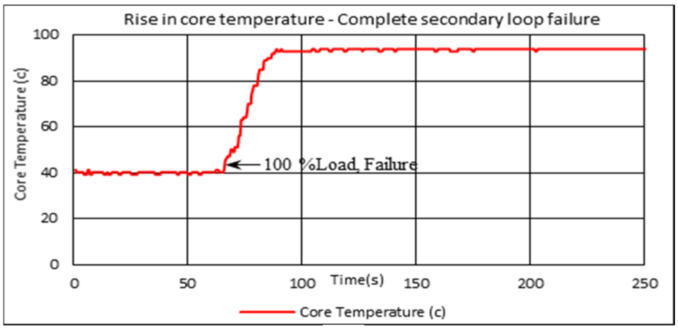

(a)

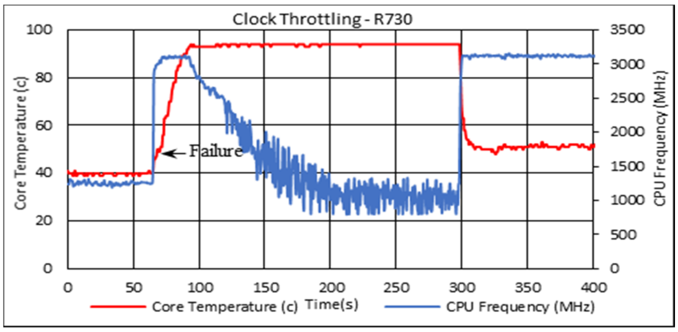

Fig. 4. CPU throttling on complete C-loop failure

(b)

temperature sensors from 3.1 GHz to 700 MHz., when both pumps in the C-loop fail. This occurs when the core temperature hits 93 degrees Celsius in about 35 seconds after the point of failure. At 300 seconds, both pumps in C-loop are turned on and CPU frequency is back to 3100MHz levels. This also corroborates why CPU clock throttling was not observed on pump failure in the F-loop and on the failure of a single pump in the C-loop. Note that the probability of both pumps in the C-loop is extremely unlikely, so this is not a realistic flow disruption scenario and is thus not addressed further in the paper. These results also corroborate the need to have redundant pumps in the C-loop, as heat needs to be taken off the CPU chips quickly through the maintenance of water flow in the C-loop.

Partial flow disruption in the C-loop: When one of the two pumps fail in the C-loop, a similar experimentation as above shows that sufficient flow rates are maintained by the remaining pump that is operational and sufficient C-loop flow and cooling is always provided. The impact of the failure of a single pump on the C-loop side is thus not catastrophic.

Complete flow disruption in the F-loop: We consider a scenario where the pump in the F-loop fails, stopping circulation in the F-loop. The C-loop still continues to run and circulate water through the coldplates and the heat exchanger. As a consequence of the F-loop failure, progressively lower amounts of heat will be taken off the by the rack-mounted heat exchanger and if the server activities continue, the core temperatures will rise and cause the CPU clock frequency to be throttled down by the Dynamic Voltage and Frequency Scaling mechanism (DVFS).

Fig. 5 (a) depicts what happens when synthetic application that exercise the cores at 100% utilization level at the full 100% clock rate is run on *all of the 16 servers in the water cooled rack* and the F-loop circulation failure is introduced by shutting off the F-loops inlet with an external solenoid valve. As seen from this figure, the CPU core temperatures go up steadily but slowly. However, they still do not go up enough to hit the temperature threshold at which clock throttling starts (around 90 degrees Celsius). Instead, the core temperature goes up immediately following the F-loop pump failure to a level where the server fans run at full speed (about 15,000 rpm) and stays there for a long time (Fig. 5 (b)). Starting from the time of induced failure (about 1000 Secs.) till this happens (at 2872 Secs. where *all fans* are ramped up) on the highest-powered server, we have 1872 Secs. (that is, ~30 minutes) of cooling left in the DCLS to provide cooling to the CPU. This primarily comes from the high heat capacity of water and the cold water left in F-loop of the heat exchanger and its internal tank and the two manifolds and all plumbing lines on both loops. It is interesting to note that the server fans ramp up to cool the CPU on sensing an increasing core temperature, even though the fans are not the primary means of cooling the CPUs. This is clearly seen in Fig. 5 (a) and Fig. 5 (b), where the core temperature goes up steadily even when the fans are running at full speed. The fans only provide indirect cooling to the CPU by taking away the heat that is conducted to the motherboard by the CPU pins and, in turn, cooling the CPU a bit but not enough to fully counter the core temperature increase. (Air blowing over the coldplates provide little cooling as the coldplate is thermally insulated with a hard plastic jacket.) As of now, fan speed control and DCLS are not unified, so the motherboard-based server control systems use their existing algorithms based on sensed CPU core

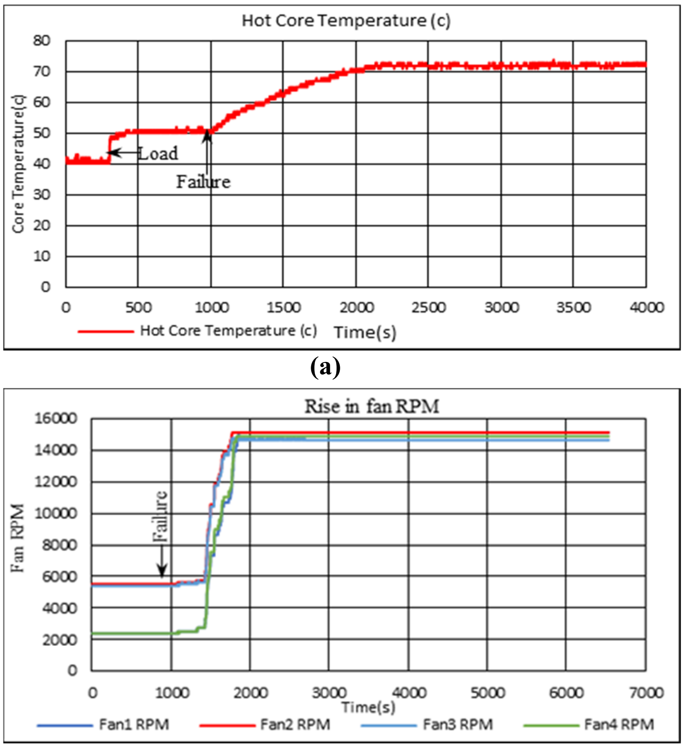

(a)

(b)

Fig. 5. Rise in core temperature in different failure modes

temperatures. In fact, server vendors do not expose the fan control APIs to end-users for a variety of reasons, so unifying fan speed control to take advantage of DCLS is not a commonly available option. Running fans at higher RPM for prolonged duration has several disadvantages:

1. The typical power consumption by fans at lower RPM is 6 Watts but at higher RPMs it increases to 14 Watts.
2. As a consequence of higher RPM, the fan coils and motor heat up which leads to reduced fan lifetime. [28, 29]
3. In an enclosed cabinet, because of the higher fan RPMs, it can lead to back pressure inside the server cabinet and recirculate the hot air within the server. [15]

It is imperative to shift the load away from the server where the F-loop has failed – certainly before the fans ramp up to full speed to avoid damage to the fan bearings from prolonged running at full speed.

When two racks of servers are connected to the rack-level heat exchanger, the time from F-loop pump failure till the fans ramp up to full speed will be, ideally, half of the corresponding time we see on a single rack in the ideal situation, that is, 15 minutes. The actual time seen may be slightly lower due to some small heat radiation from the cooling system to the air, even though the manifolds are well insulated. Note that the 15-minute limit applies to scenarios where the CPUs see 100% sustained utilization. For most applications, which do not see a sustained CPU utilization of 100% throughout the run, the time from F-loop pump failure to the time that the fans ramp up to full speed is likely to be higher than the 30 and 15 minute limits observed for a single and double rack, respectively. However, from a worst-case standpoint, whatever corrective actions are taken to deal with this failure must complete within the 30 and 15 minute limits with one and two racks, respectively. With a single heat exchanger, if each of these two racks have enough high- powered servers to double the CPU heat load that has to be handled by the DCLS, the time to fan ramp-up to full speed comes down to 7.5 minutes under ideal situations without any heat loss in the DCLS.

From the Fig. 5 (a) and Fig. 5 (b), the **residual cooling energy** left in the system for running and migrating Virtual Machines (VMs) before fans ramp up to their full speed is ~2541KJ.

Redundant pumps for the F-loop can avoid flow disruptions, but providing such redundancy is expensive in terms of plumbing costs, pump energy costs and/or pump control system costs. Specifically, where the control system uses two pumps in parallel for redundancy, the redundant pump needs to be activated when a single pump fails and diverter valves have to be activated to bypass the failed pump, all adding to the installation and maintenance costs. When redundancy is provided with two pumps in series, additional energy is expended in running both pumps at all times. As we will see later, in Section IV, the need for redundant pumps in the F-loop can be avoided with a software solution that capitalizes on the thermal capacity of the F-loop segment from the pump to the heat exchangers. This solution avoids damage to the CPUs affected by the pump failure and simultaneously permits graceful degradation of service.

IV. Dealing with Pump Failures

Of the 3 pump failure scenarios discussed in Section III, the two that need to be addressed in practice are the failure of the pump in the F-loop and the failure of one of the pumps in the C-loop. The third scenario, where both pumps in the C-loop fail, is very unlikely.

A. Dealing with Single Pump Failure in the C-Loop

For partial disruptions in the C-loop, no corrective actions are needed as the enough flow is maintained in the C-loop to provide the necessary cooling. However, to avoid fan speed ramp-ups to sustained full speed and avoiding fan damage that can result from sustained use at full speeds, VM migration as in the case of F-loop flow failure can be triggered on sensing fan speed increase beyond a pre-specified threshold *after* detecting the failure of a single pump (which is sensed via the remote access interface of the rack-level heat exchanger).

B. Dealing with F-Loop Pump Failure

The obvious way of dealing with this failure in the virtualized environment is to migrate the VMs on the affected racks to other racks where cooling is available before the fans ramp up to full speed. If the migration cannot be completed, the CPU clock can be reduced to lower the power dissipation and increase the time to ramp-up to full fan speed from the point of

TABLE I. Grouped Migration vs Parallel Migration of VMs

Benchmark	Average latency (s)		Top 95% latency (s)		Top 99% latency (s)		Avg. Migration time per VM(s)	
	Staged	Parallel	Staged	Parallel	Staged	Parallel	Staged	Parallel
Crypto.rsa	82.5	94.3	110	155	120	158.8	10.6	426.3
Mpegaudio	210.2	241.9	273	376.2	281.2	381.2	11.1	628.7
Scimark.monte_carlo.large	128	133.7	172.3	224.2	191.1	240.7	10.7	206.4
Serial	221.6	310	363.7	445	499.3	452.2	37	1169.3
XML	6.3	7.5	10.5	14.3	13.5	18.3	12.6	598.5

Legend: Staged: Grouped migration of VM's in batches of 4; **Parallel:** Simultaneous migration of VM's which causes resource contention;

failure. This is a viable solution as the migration activities are mostly IO-bound and migrations times are not substantially affected by operating the CPU at a lower (but not too low) clock rate.

A second consideration is to speed up the migration to reduce contention over a wide-variety of resources during migration of the VMs [9]. Such contention increases the overall migration time and increase the average request service latencies, particularly the latencies in the upper 95th to 99th-percentiles of request served sorted by service time. As cloud and data center operators strive hard to limit the 95-th percentile latency compared to the average latency, it is imperative to reduce contentions during VM migration.

Examples of resources over which contention can occur during VM migration include access to the local disk for dirty pages, access to remotely mounted file system, contention at the network interfaces used for migration and other local resources such as the servers IO bus. A solution to reduce but not altogether avoid the contention is gang migration that systematically divides up the VM into groups and migrate the groups serially. Contention only occurs among the VMs within a group and by keeping the group size small, we can lower the overall migration time.

The solution proposed and evaluated in this paper does the following:

- Sense F-loop pump failure by reading the F-loop flow rate from the rack-level heat exchanger.
- On detecting a sudden reduction in the flow rate or a complete absence of flow in the F-loop, reduce the CPU clock frequency to an acceptable lower frequency that does not affect migration time appreciably. In our experimental studies, we used a clock rate of 100% as the maximum clock rate possible (without turbo boosting). In general, the clock rate to use can be determined empirically for the class of services run on the VM.
- Initiate VM migration simultaneously within a group of servers and sequentially to complete VM migration within the time it takes for the server fans to ramp up to full speed. If this is not possible, migrate as many VM groups as possible till the fan speed ramps up and reduce the CPU clock further to ensure that all VMs can be migrated within the time to ramp up to full speed. Again, a priori calibration runs can be done to determine the right clock rate to use.
- Shut down the servers in the rack once all VMs have been migrated successfully.

When VM migration is necessary from an affected server on sensing the F-loop failure, a back-end daemon running on the affected server is responsible for triggering the VM migration. BED uses libvirt [18] VM migration APIs to migrate the VM from an affected server to a non-affected server.

V. EXPERIMENTAL RESULTS

We evaluate a prototype implementation of our technique on a set of 16 heterogeneous servers on a water-cooled rack, consisting of 13 Dell R520 servers (dual Xeon E5 2440 @ 2.4GHz), 2 Dell R730 servers (dual Xeon E5 2687W v3 @3.1GhHz) and one Dell R730 server (dual Xeon E5 2687W v4 @3.00 GHz). A medium sized VM configuration was chosen with 4 virtual CPUs, and 4GB of memory for the experiments. A total of 94 VMs are hosted in the 16-server rack with water cooling. The servers and VMs are booted with Ubuntu 16.04 disk image, running Linux kernel 4.4.0-112., and qemu-kvm 1:2.5 version. 10 Gbps network links are used for network connectivity across the servers. F5 Networks BIG-IP 4000s LTM load-balancer is used to load balance the requests.

In the solution studied experimentally in this paper, we used SPECjvm2008 benchmark suite. SPECjvm2008 benchmark suite was released by Standard Performance Evaluation Corporation (SPEC) [28] in 2008 and it replaces SPECjvm98 which has been used for 10 years.

SPECjvm2008 contains 38 benchmark programs in total that

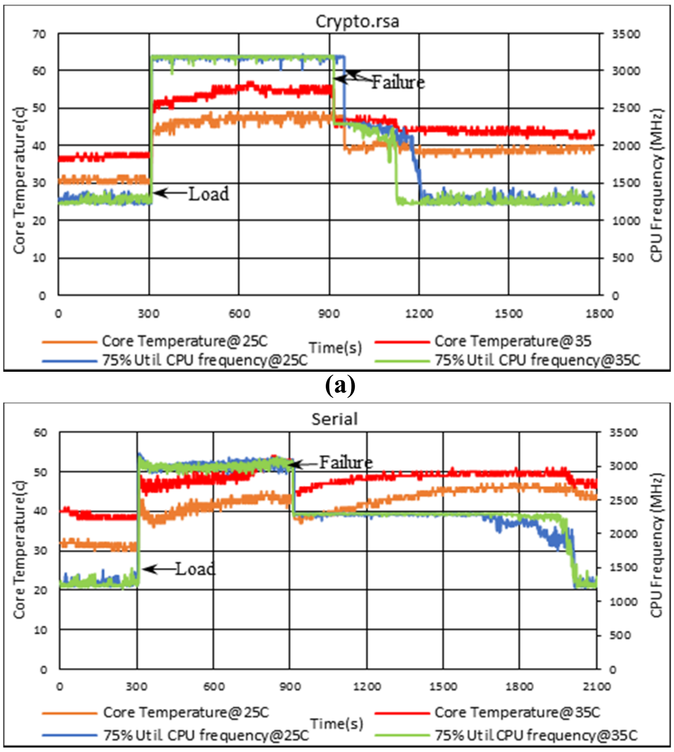

(b)

Fig. 6. Impact of inlet coolant temperature on core temperature

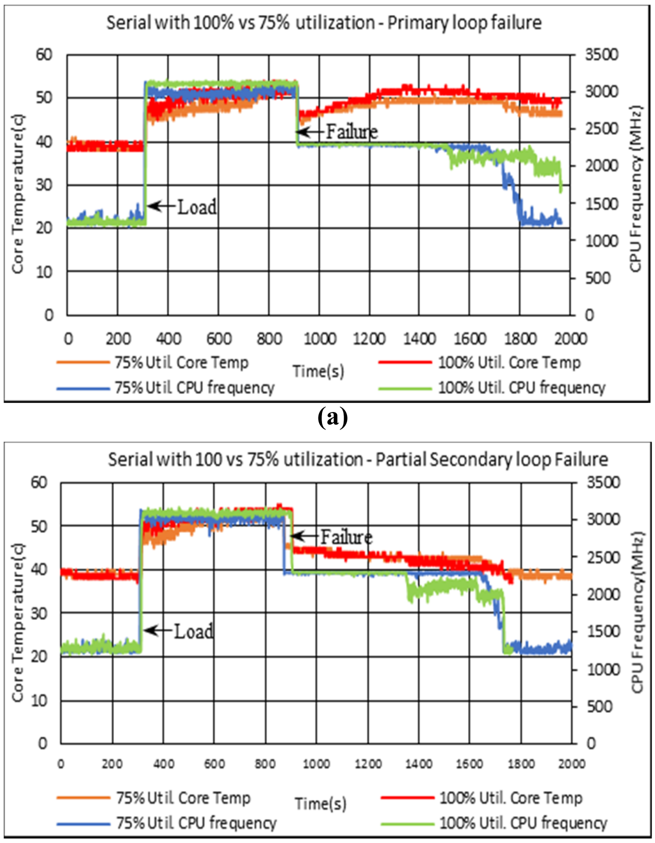

(b)

Fig. 7. Impact of CPU utilization on core temperature

evaluates client-side Java Runtime Environment. Among those 38 benchmarks, we used five representative benchmarks from SPECjvm2008 benchmark suite that had substantial per VM migration time. These benchmarks included the following: *crypto.rsa, mpegaudio, scimark.monte_carlo.large, serial, xml.transform.*

Crypto.rsa does encryption and decryption of input data using RSA protocol. The input data is 100Bytes and 16KB. *mpegaudio* is floating point heavy benchmark, which evaluates mp3 decoding. The mp3 library has been replaced with JLayer, an LGPL mp3 library [21]. *SciMark* is a computational workload which contains different benchmarks and are divided into two categories [14]. *SciMark.X.large* works with input data size of 32MB and *scimark.X.small* work with input data size of 512KB. *SciMark* contains 5 workloads fft, lu, monte_carlo, sor and sparse, each with large and small input data size. *Serial* works in a producer-consumer scenario, where producer sends serialized objects via socket to consumer and consumer deserialize them on the same system. Xml benchmark workload contains two sub-benchmark programs, *xml.transform* and *xml.validate.* xml.transform evaluates java.xml.transform of the JRE under test and xml.validate compares the XML files and XML specification in .xsd files.

Table I shows the advantages of doing grouped migration ("staged") instead of concurrent ("parallel") VM migration. As stated above, concurrent migration causes resource contentions and will lead to an increase in the per-VM migration time. In some cases, these increases are substantial, as seen in Table 1. The increase in the average latency with concurrent migrations against the average latency with grouped migration ranges from 4.3 % in the case of scimark.monte_carlo.large to 39.8% in case of serial benchmark. Most datacenter operators guarantee the upper 95th percentile tail latency, and as seen in Table 1 grouped migration does well in limiting the increase in this percentile compared to the average latency. Against average latency, the latency for the 95th percentile is between 22.3% to 40.9% over the average. The 95th-percentile increase in latencies compared to average latency are substantial when concurrent migration is used.

The inlet coolant temperature plays a critical role in the increase of core temperature during the failure scenario. Fig. 6, shows the effect of inlet coolant temperature on core during F-loop failure with inlet 25C, 35C and 75% utilization for *crypto.rsa* and *serial*. With lower inlet temperature, the increase in core temperature is slower and provide us with ample amount of time to migrate VM's. Its effect on tail latency can be contained. During the start of an experiment, core temperatures are at 31c and 36c with respective to inlet coolant temperature of 25c and 35c. All VMs are loaded to 75% after 300 seconds from the start of the experiment. F-loop failure is simulated around 900 seconds into the experiment which leads to an increase in core temperature. When cooling failure is detected, all hosts are explicitly lowered to 75% of CPU clock frequency in-order to have sufficient amount of time to migrate all the VMs to another rack with ample amount of

TABLE II. TOTAL MIGRATION TIME AT DIFFERENT INLET TEMPERATURES AND UTILIZATION (SECONDS)

Benchmark	25C – 50 % Utilization	25C – 75% Utilization	35C – 50% Utilization	35C – 75% Utilization
Crypto.rsa	534	552	426	452
Mpegaudio	472	532	654	488
Scimark.monte_carlo	510	456	526	528
Serial	934	1828	1102	1758
XML	654	646	590	560

TABLE III. TOTAL MIGRATION TIME AT DIFFRERENT CPU FREQUENCY (SECONDS)

Benchmark	100 % frequency	75 % frequency
Crypto.rsa	514	452
mpegaudio	500	488
Scimark.monte_carlo.large	550	528
Serial	1668	1758
XML	600	560

cooling and of different cooling type. This ensures us that, we're not migrating VMs back onto affected hosts.

Crypto.rsa is CPU intensive benchmark and VM migration time is lower compared to *serial* and its effect can be seen on the core temperature. When failure is detected, CPU frequency is throttled down to 75%, which lowers CPU temperatures by 7C and allows us to migrate VMs. When all hosted VMs are migrated, CPU frequency, core temperature goes down and host can be powered off for maintenance. *Serial* which works by serializing and deserializing objects sent over socket, requires more amount of time to complete VM migration. Hence the increase in the core temperature of host even after throttling of CPU frequency when failure is detected. Around 2000th seconds, VM migration is completed and core temperature drops. As stated earlier, prior calibration is required to decide on clock rate to use for different workloads.

Another factor that contributes to higher core temperatures during normal environment or failure scenario is CPU utilization. It is optimal and ideal to run servers at highest ratio of CPU utilization to power drawn. But during a failure scenario this should be averted and try to limit damage to hardware and impact on request latencies. Fig. 7, compares running VMs at 75% utilization and 100% utilization. Fig. 7 (a) shows the effect on host during F-loop failure and Fig. 7 (b) during partial C-loop failure.

In the F-loop failure scenario, the average latency at 100% utilization increased by 24.4% and its upper 95-th percentile got affected by 10.4% when compared to 75% utilization. In the partial C-loop failure scenario, the average latency at 100% utilization is greater than the latency at 75% utilization by 19.7% and its upper 95-th percentile by 6.9%. This implies that a complete F-loop failure can cause more damage due to increase in the CPU core temperature and because its impact on tail latency is more than what is encountered on partial C-loop failure.

Table II presents the total amount of time required to migrate all VMs at different inlet coolant temperature and CPU utilization. All VM migrations are completed well below the time required for fans to ram up by throttling CPU frequency to 75%. VM migration time while running Serial is longest because of its larger memory foot print.

Table III shows total elapsed time for migrating all affected VMs when hosts are running at different CPU frequencies. Since we are explicitly reducing the CPU frequency to 75% when failure is detected, we're able to migrate VMs without any fan ramp up and maintaining lower core temperatures. When CPU frequency is lowered to 75% after the detection of F-loop failure, fans will stay in the range of 6500-8000 RPM, since core temperature stays below 67 degrees Celsius.

VI. RELATED WORK

An exhaustive experimental study related to inefficiencies in a warm-water DLCS from Asetek is presented in [6]. The benefits of warm-water DLCS are explored using mode ling in [4, 8, 13, 26] but most of these do not systematically identify the inefficiencies or address reliability issues. To address the inefficiency of datacenter cooling, several thermal management technologies are used [1]. Potential energy savings from DLCS has paved the path for commercial product development which can be used for datacenter deployment beyond the HPC segment [20, 22].

Different VM migration techniques are surveyed in [17] and categorized based on energy efficiency, load balancing capability and fault tolerance. The Collective project has explored the option of VM migration, which stops OS execution for transfer [27], optimized for slow links and longer time spans. In a mobile environment, the work of [16] has explored the ideas of migration by stopping and transferring which leads to long migration time. To reduce the amount of data to be transferred during migration, the idea of partial OS virtualization was proposed in [23] based on the use of a customized Linux Kernel to isolate all file handles, sockets into a name space and migrate the name space. Alternative VM migration techniques are proposed in [27]. The migration technique presented here is thermally-aware with aim of to provide dependable service in the case of cooling system failure.

VII. CONCLUSIONS

Warm-water cooling systems, previously used in cooling HPC servers, are now available from at least one vendor for retrofitting into off-the-shelf non-HPC servers for cloud data center applications. The particular system studied in this paper replaces air cooled heatsinks on servers with coldplates that circulate water to take the heat off the CPU chips. We studied two realistic failure scenarios using a rack of 16 servers equipped with the warm water cooling system. On the failure of the pump that supplies cold water to the rack-level heat exchanger using the F-loop, CPU temperatures go up slowly but fan speeds ramp-up to the maximum value well before the CPU core temperature exceed safe limits. This is due to the thermal capacity of water in the loops and the heat exchanger, which continue to provide adequate cooling for a while after the F-loop pump failure. Since sustained running of the server fans at the maximum speed can reduce the fan lifetime (while continue to consume higher power than at lower speeds), the server workload must be moved off the server when the F-loop's pump failure is detected and before the fans ramp up to full speed. When one of the two pumps on the C-loop fails, sufficient circulation is still available to provide cooling to the CPUs to permit them to continue running. When both pumps in the C-loop fail, an unlikely scenario, CPU core temperatures and fan speeds ramp up.

We proposed and demonstrated techniques for dealing with the F-loop pump failure in a virtualized environment that migrate tasks away to other servers on failure before the fan speed ramps up. Our technique relies on migrating tasks in groups to mitigate the effect of resource contentions during VM migration and, when necessary, to use a reduced clock frequency during the migration (which are IO bound in general). Reducing the CPU clock frequency makes the available residual cooling to last for a longer time, permitting the VM migrations to complete. The same techniques can be used to deal with the failure of a pump in the C-loop. Techniques for dynamically estimating the residual cooling energy and throttling the CPU clock only when necessary, to minimize average and tail latency increases for the served requests, are presented in a forthcoming paper [24].

ACKNOWLEDGEMENT

This work is supported in part through NSF awards 1040666, 1134867, 1738793 and a NYSERDA award under the PON 2846 competition.

REFERENCES

[1] Alkharabshe S *et al.*,"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.

[2] Allied Control Immersion Cooling System Product Pages at: http://www.allied-control.com/immersion-cooling.

[3] Asetek Corpn., Data Center Cooling Solutions Product Pages for Rack DCU and Internal Loop available at: http://asetek.com/data-center/data-center-coolers.aspx.

[4] Chu, R. C., Simons, R. E., Ellsworth, M. J., Schmidt, R. R., & Cozzolino, V. "Review of cooling technologies for computer products". IEEE Transactions on Device and Materials Reliability,4(4), 2004, pp. 568-585.

[5] Coolcentric Inc., RDHx Rear-door Heat Exchanger Product web pages at: http://www.coolcentric.com/products/rear-door-heat-exchangers.php

[6] CoolIT Systems, Rack DCLC Product Guide, Product Document No. 761-00012, Rev. 04, December 2013.

[7] Coles, Henry (Lawrence Berkeley National Laboratory) and Steve Greenberg (Lawrence Berkeley National Laboratory), "Demonstration of Direct Liquid Cooling for Electronic Equipment.* California Energy Commission". Publication number: CEC-(report number pending), 2004. Also available at: http://datacenters.lbl.gov/resources.

[8] David, M. P., Parida, P. R., Schultz, M. D., Simons, R. E., Gaynes, M., Schmidt, R., & Chainer, T. J.,"Dynamically Controlled Lon Duration Operation of a Highly Energy Efficient Chiller-Less Data Center Test Facility". In ASME 2013 International Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems, July 2013 (pp. V002T09A014-V002T09A014). American Society of Mechanical Engineers.

[9] Deshpande. U. et al., "Live Gang Migration of Virtual Machines," High-Performance Parallel and Distributed Computing, June 2011.

[10] Ellsworth Jr, Michael J., Roger R. Schmidt, and Dereje Agonafer. "Design and analysis of a scheme to mitigate condensation on an assembly used to cool a processor module." IBM Journal of Research and Development 46.6, (2002), pp. 753-761.9.Ellsworth, M. J., & Simons, R. E. (2005). "High powered chip cooling-air and beyond". Electronics Cooling, 11(3), 2005, pp. 14-22.

[11] Ellsworth, M. J. Simons, R. E. (2005). "High powered chip cooling-air and beyond"; Electronics Cooling, 11(3), 2005, pp. 14-22.

[12] Green Revolution Cooling Submersion Cooling product web pages at: http://www.grcooling.com

[13] Iyengar, M., David, M., Parida, P., Kamath, V., Kochuparambil, B., Graybill, D., & Chainer, T., "Server liquid cooling with chiller-less data center design to enable significant energy savings". In Semiconductor Thermal Measurement and Management Symposium (SEMI-THERM), 2012 28th Annual IEEE, March 2012 (pp. 212-223).

[14] Java SciMark 2.0 at: http://math.nist.gov/scimark2/

[15] Khalili S, Alissa H, Nemati K, et al. Impact of Internal Design on the Efficiency of IT Equipment in a Hot Aisle Containment System: An Experimental Study. In ASME 2018 InterPACK Conf., August 2018.

[16] Kozuch M and Satyanarayanan M. Internet suspend/resume. In Proceedings of the IEEE Work-shop on Mobile Computing Systems and Applications,2002.

[17] Leelipushpam and Sharmila J., "Live vm migration techniques in cloud environment: A survey," in ICT, 2013.

[18] Libvirt web pages at: www.libvirt.org.

[19] Masanet E.R., Brow R.E., Shehabi A., Koomey J.G., and Nordman B., "Estimating the Energy Use and Efficiency Potential of U.S. Data Centers". Proceedings of the IEEE 99(8), 2011, pp. 1440-1453.

[20] Miller. R, "Rise of Direct Liquid Cooling in Data Centers Likely Inevitable," Data Center Knowledge, 2014.

[21] MP3 library for the Java Platform at: http://www.javazoom.net/javalayer/javalayer.html

[22] "New Liquid Cooled HPC Cluster Launched," Electronics Cooling, 2016.

[23] Osman S, Subhraveti D, Su G, and Nieh J. "The design and implementation of zap: A system for migrating computing environments". In Proc. 5th USENIX Symposium on Operating Systems Design and Implementation (OSDI-02), pages 361–376, December 2002.

[24] Puvvadi., U., Desu., A., Stachecki, T., Ghose., K. and Sammakia, B., " An Adaptive Approach for Dealing with Flow Disruption in Virtualized Water-Cooled Data Centers", to appear in Proc. IEEE Cloud 2019.

[25] Rittal In-line cooler product overview at: http://www.42u.com/rittal-lcp.htm.

[26] Sharma C. S., Tiwari M. K., Michel B., & Poulikakos D.., "Thermofluidics and energetics of a manifold microchannel heat sink for electronics with recovered hot *computing* (HPDC '11). ACM, New York, NY, USA, 135-146.

[27] Sapuntzakis C. P., Chandra R, Pfaff B, Chow J, Lam M.S., and Rosenblum M. "Optimizing the migration of virtual computers". In Proc. of the 5th Sym- posium on Operating Systems Design and Implementation (OSDI-02), December 2002.

[28] SPECjvm2008, http://www.spec.org/jvm2008/ .

[29] Tian, X., "Cooling fan reliability: failure criteria, accelerated life testing, modeling and qualification", In Proc.. Annual Reliability and Maintainability Symp. 2006, pp. 380-384.