ENERGY-BASED MODELLING AND SIMULATION OF LIQUID IMMERSION COOLING SYSTEMS

Michele Lionello[1,*], Mirco Rampazzo[1], Alessandro Beghi[1], Damiano Varagnolo[2], and Mattias Vesterlund[3]

[1] Department of Information Engineering, University of Padova, via Gradenigo 6/B, I-35131 Padova, Italy
[2] Department of Engineering Cybernetics, Norwegian University of Science and Technology, Norway
[3] Swedish Institute of Computer Science North, Research Institutes of Sweden AB, Sweden
*Correspondence author. Email: lionello@dei.unipd.it

Keywords: Liquid Immersion Cooling, Graph-based modelling, Simulation

ABSTRACT
Currently, most of the existing data centers use chilled air to remove the heat produced by the servers. However, liquids have generally better heat dissipation capabilities than air, thus liquid cooling systems are expected to become a standard choice in future data centers. Designing and managing these cooling units benefit from having control-oriented models that can accurately describe the thermal status of both the coolant and the heat sources.

This preliminary research aims at deriving a control-oriented model of liquid immersion cooling systems, i.e., systems where servers are immersed in a dielectric fluid having good heat transfer properties. More specifically, we derive a general lumped-parameters gray box dynamical model that mimics energy and mass transfer phenomena that occur between the main components of the system. The proposed model has been validated against experimental data gathered during the operation of a proof-of-concept immersion cooling unit, showing good approximation capabilities.

INTRODUCTION
Nowadays, the demand for data center computing is continuously rising, increasing the total energy consumption of data centers worldwide. In particular, a large portion of the electrical power provided to a data center is drawn by cooling systems, which have to guarantee safety thermal requirements for the correct operation of the Information Technology (IT) equipment [1]. Thus, improving the energy efficiency of the data center cooling infrastructure while guaranteeing the thermal constraints is imperative, [2], and it can be obtained by combining the use of new cooling technologies and the chances offered by the design of advanced control systems.

Currently, most of existing data centers use chilled air to remove the thermal energy produced by the IT equipment. However, air-based cooling suffers from many inefficiencies, like hot air recirculation and cold air bypass [3]. Also, effective cooling in a conventional air-cooled data center requires a lot of space to locate air conditioners and server racks. Furthermore, the air is basically a low efficient cooling medium due to its low density and heat removal capacity [4]. A promising alternative is given by liquid cooling solutions, since liquids used in cooling have generally higher heat transport capacity than air. In particular, Liquid Immersion Cooling (LIC) is arousing interest as a possible method to cool high heat flux electrical components in data centers. Unlike the water-cooled cold plate solutions which utilize physical walls to separate the coolant from the chips, immersion cooling brings the coolant in direct physical contact with the chips resulting in a lower internal thermal resistance, [5]. More precisely, in LIC systems servers are directly immersed within a dielectric liquid bath which ensures electrical insulation. During servers operation, thermal energy is transferred from the chips to the dielectric liquid and, as a result, servers electrical components are cooled down and the thermal energy is stored in the liquid. A water-based cooling circuit which comprises one or more immersed cooling plates can be used to extract the stored thermal energy by exploiting natural convection: the coolant nearby the heat sources heats up and by thermal expansion rises towards the top of the bath, while the liquid in the proximity of the immersed plates falls from the top to the bottom of the bath. Note that the thermal energy extracted by the cooling circuit can then be recovered, thus a data center which uses LIC technology can be used, for example, for pre-heating water in a district heating scenario, [6].

In the context of improving cooling and heat recovery capabilities of LIC systems, the availability of mathematical models can be of great support during both hardware and advanced control systems design. For example, they are essential in deriving Computer Aided Control Systems Design (CACSD) software tools, which are cheaper and safer than conducting experiments [7]. Also, a model of the system can facilitate the design and the assessment of different control strategies, and it can be included in model-based control schemes. It is of paramount importance to obtain a model that is accurate enough in reproducing the behaviour of the real system and, as far as possible, easily achievable. For this purpose, in this work we derive the structure of the model by using the First-Principle laws of the physics to guarantee the physical interpretability of the model. Then, we calibrate some of their parameters by exploiting experimental data that have been gathered from an experimental LIC testbed. More specifically, the model structure has been derived by using a graph-based modelling framework, which relies on the conservation of mass and energy and which captures the structure and interconnections in the system, [8]. The effectiveness of graph-based approaches to modelling power flow systems has been shown, for example, in [9] for building thermal dynamics, in [10] for process systems, and in [11] and [12] for vehicle energy management systems.

This manuscript is organized as follows. First, we describe the experimental setup for LIC that we considered in this work. Then, we provide a brief overview of the graph-based modelling technique and we derive the proposed graph-based model for LIC systems. After that, we describe the model calibration phase and we report the results of the model validation. Finally, some conclusions and some future research directions are presented.

THE EXPERIMENTAL SETUP

In this manuscript, we consider an experimental testbed for LIC that is physically located in the Infrastructure and Cloud research & test Environment (ICE) facility at Research Institutes of Sweden (RISE) Swedish Institute of Computer Science (SICS) in Luleå, Sweden, a facility dedicated to testing innovative technologies for data centers. The experimental setup consists of a 0.84m side cubic vessel, built with an aluminium frame and sealed acrylic glass walls and filled with a mineral oil having high dielectric strength, low viscosity, and good heat transfer properties. The vessel includes also a 2mm thick metal fully welded sheet, to avoid oil leakages. Pictures of the vessel before and after sealing it are reported in Fig. 1.

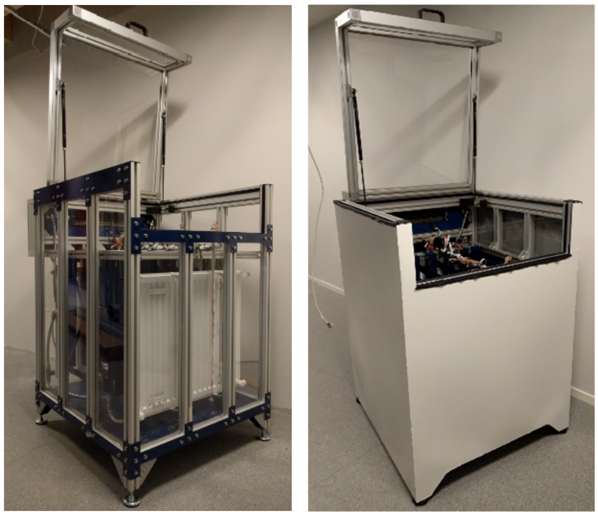

Figure 1
Pictures of the vessel before and after sealing it.

Four Open Compute Windmill V2 servers donated from Facebook are directly immersed in the mineral oil. More precisely, the servers are mounted in two chassis, each of which is fixed to the aluminium frame and can contain two servers. Before immersing them into the oil, their local cooling fans have been removed and the thermal paste commonly used in air-based cooling has been replaced by thin aluminium foils. While the servers are operating, part of the thermal energy produced by the chips is transferred to the mineral oil. In order to extract the thermal energy stored in the oil, the testbed comprises a water-based closed cooling system which consists in: two plate coils that are immersed in the oil; a pump used to circulate the cooling water and driven by an inverter; an external water-to-air heat exchanger which disperses the extracted thermal energy into the ambient. It is worth highlighting that, in real applications, the external heat exchanger could be substituted by a possible heat recovery unit, in order to reuse the waste heat produced by the servers. A schematic representation of the experimental testbed is depicted in Fig. 2.

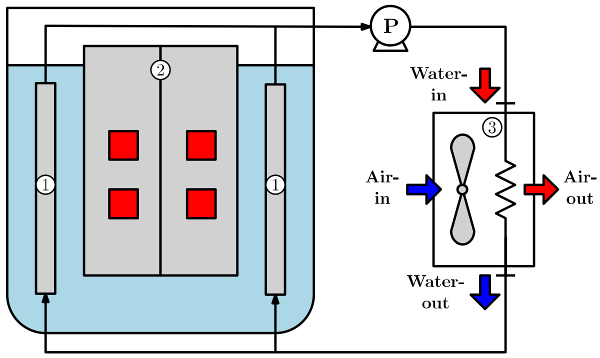

Figure 2
Schematic representation of the experimental testbed, which comprises: the two immersed heat exchangers (1), the immersed servers (2) each endowed with two CPUs (depicted by the red squares), the water pump (P), and the external water-to-air heat exchanger (3).

The testbed comprises also a monitoring system, which consists of a set of sensors and a dedicated Supervisory Control And Data Acquisition (SCADA) system. In particular, sensors are installed in order to measure the flow rate and the temperature of the cooling water and of the ambient air in both inlet and outlet sections of the external heat exchanger. Moreover, the testbed comprises a set of thermocouples that can be immersed in the mineral oil, in order to measure its temperature in different zones of the vessel. By using the Intelligent Platform Management Interface (IPMI), we also monitor the status of each server, i.e. the temperatures of its Central Processing Units (CPUs) and its electrical power consumption. Finally, a set of Python scripts, which are based on the Modbus communication protocol and the stress-ng tool, allows to change remotely the water pump speed and the computational load of each server.

SYSTEM ANALYSIS AND MODELLING

Graph-based modelling: The thermal model proposed in this work relies on the graph-based modelling technique proposed in [4]. The main idea behind this approach is to model the thermal system as a directed graph $G = (V, E)$ that captures the energy storage and power flow throughout a system S. The graph consists in a set of vertices $V = \{v_i : i \in \{1, 2, \dots, N_v\}\}$ and a set of edges $E = \{e_i : i \in \{1, 2, \dots, N_e\}\}$. The vertices represent the capacitive elements of the system where energy can be stored, while the edges represent the transport of energy between two adjacent vertices. The orientation of each edge e_j indicates the positive direction of associated power P_j from the tail vertex $v_{j,tail}$ to the head vertex $v_{j,head}$. The set of edges directed into the i-th vertex is denoted by $E_{i,in} = \{e_j : v_{j,head} = v_i\}$, while the set of edges directed out of the i-th vertex is denoted by $E_{i,out} = \{e_j : v_{j,tail} = v_i\}$. The graph includes also sources that can add power to the system and sinks where power can be rejected from the system. Both sources and sinks are represented by dashed vertices and include, for example, thermal loads and atmospheric conditions. Let $V_s \in \mathbb{R}^{N_s}$ and $V_t \in \mathbb{R}^{N_t}$ denote the source and sink vertices, respectively, such that $V_s \subset V$ and $V_t \subset V$. Finally, we denote the N_d dynamic vertices such that $V_d \subset V \backslash (V_s \cup V_t)$.

For each dynamic vertex $v_i \in V_d$, the dynamic state x_i represents its temperature and C_i represents its thermal capacity. For each vertex, the energy conservation law must hold, i.e.:

$$C_i \dot{x}_i = \sum_{e_j \in E_i^{in}} P_j - \sum_{e_j \in E_i^{out}} P_j, \qquad (1)$$

where the right-hand side represents the summation of all the power flows entering the i-th vertex minus the summation of all the power flows leaving that vertex. By collecting the states of all the dynamic vertices in the vector $\boldsymbol{x} = [x_i] \in \mathbb{R}^{\mathrm{N_d}}$, the dynamical graph system S can be written as:

$$C\dot{\boldsymbol{x}} = -M\boldsymbol{P} + D\boldsymbol{P^{in}}, \qquad (2)$$

where C is the diagonal matrix containing all the vertices capacities, i.e. $C = diag([C_i]) \in \mathbb{R}^{\mathrm{N_d} \times N_d}$, matrix $M = [m_{ij}] \in \mathbb{R}^{(\mathrm{N_d}+N_t)\times(N_e - N_s)}$ is the incidence matrix which captures the structure of the graph and which is given by:

$$m_{ij} = \begin{cases} +1, & \text{if } v_i \text{ is the tail of } e_j, \\ -1, & \text{if } v_i \text{ is the head of } e_j, \\ 0, & \text{otherwise}, \end{cases} \qquad (3)$$

vector $\boldsymbol{P} = [P_i] \in \mathbb{R}^{(N_e - N_s)}$ contains all the edge power in G that is not from a source vertex, and $\boldsymbol{P^{in}} = [P_i^{in}] \in \mathbb{R}^{N_s}$ is a vector of power flows from source vertices. Finally, matrix $D = [d_{ij}] \in \mathbb{R}^{\mathrm{N_d} \times N_s}$ captures the input of power from the source vertices:

$$d_{ij} = \begin{cases} +1, & \text{if } v_i \text{ is the head of } P_j^{in}, \\ 0, & \text{otherwise}. \end{cases} \qquad (4)$$

Each power flow P_i contained in $\boldsymbol{P}$ can be written as a function f of the vertex states and an input u_i:

$$P_i = f_i(x_i^{tail}, x_i^{head}, u_i). \qquad (5)$$

In particular, we consider both power flows due to mass transfer from one element to another:

$$P = \dot{m}c_pT, \tag{6}$$

and power flows via heat transfer:

$$P = hA\Delta T, \tag{7}$$

where h is the heat transfer coefficient and A is the heat transfer area. Thus, power flows in (5) can be modelled with the following equation:

$$P = (a + \dot{m})(bT_1 + cT_2), \tag{8}$$

that is a generalization of Eqs. (6) and (7).

Graph-based model for LIC systems: The graph model for the LIC system considered in this work is represented in Fig. 3. In the following, we describe the vertices of the graph model and the power flows represented by the edges, which are obtained as particular cases of Eq. (8).

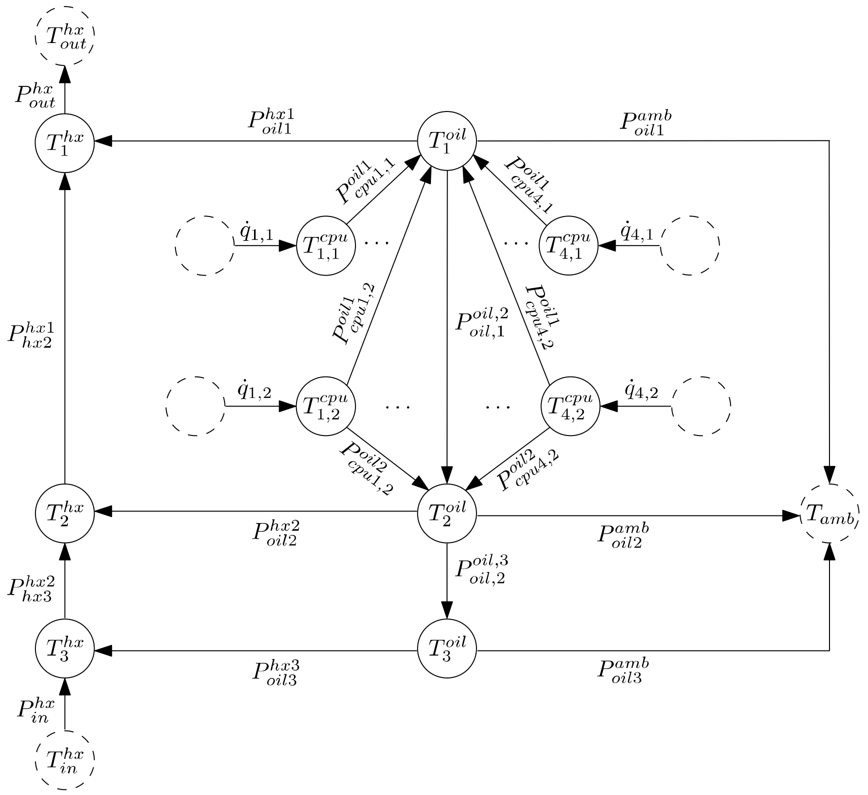

Figure 3
Graph representing the LIC model.

Vertices: In order to take into account the possible thermal stratification, we consider three vertices to represent the state of the dielectric oil in the vessel at three different heights: the first vertex represents the portion of oil at the top of the vessel, where the first level of processing units are immersed; the second vertex represents the portion of oil which contains the second level of CPUs; the third vertex represents the remaining oil below the processing units. The temperature and the thermal capacity of each oil layer are denoted by T_i^{oil} and C_i^{oil}, respectively, with $i \in \{1, 2, 3\}$. Moreover, we model as vertices of the graph the state of the processing units of the four immersed servers. As it is schematized in Fig. 2, each of the four servers contains two CPUs, which are placed at different height. With $T_{i,j}^{cpu}$ and $C_{i,j}^{cpu}$ we denote the temperature and the thermal capacity of the j-th CPU in the i-th server, with $i \in \{1, \dots, 4\}$ and $j \in \{1, 2\}$. As regards the immersed heat exchangers, we assume for simplicity to have only one fictitious immersed heat exchanger having exchange area and volume equal to the sum of those of the two real ones. Then, we represent as vertices of the graph the state of the water flowing in the portions of the fictitious immersed heat exchanger in contact with each of the three oil zones. We denote with T_i^{hx} and C_i^{hx} the temperature and the thermal capacity of each water portion, respectively, with $i \in \{1,2,3\}$. The inlet conditions of the water flow are represented by the dashed source vertex which is denoted by the inlet water temperature T_{in}^{hx}. The outlet conditions of the water flow are represented by the dashed sink vertex denoted by the outlet water temperature T_{out}^{hx}, instead. The graph model considers also the thermal state of the ambient air, which is represented by the sink vertex denoted by the ambient air temperature T^{amb}. Finally, the blank dashed source vertices represent the sources that provide the electrical power to the processing units during their operation.

Edges: The heat exchanged between adjacent oil vertices is assumed to be mainly due to conduction, thus we express it as:

$$P_{oil,i}^{oil,i+1} = g^{oil}\frac{A_i^{oil}}{d_i^{oil}}\left(T_i^{oil} - T_{i+1}^{oil}\right), \tag{9}$$

where g^{oil} is the heat conductivity of the mineral oil, A_i^{oil} is the contact area between the two adjacent oil layers, and d_i^{oil} is the distance between the centers of the two layers. The heat flowing from the i-th oil layer and the ambient air through the vessel walls is given by:

$$P_{oil,i}^{amb} = h^{amb}A_i^{amb}\left(T_i^{oil} - T^{amb}\right), \tag{10}$$

where h^{amb} is the heat transfer coefficient and A_i^{amb} is the exchange area. The power flows $\dot{q}_{i,j}$ are the electrical power provided to the CPUs. The heat exchanged between the higher oil layer and the cpus immersed in it is given by:

$$P_{cpu,i,1}^{oil,1} = h^{cpu}A_{i,1}^{cpu}\left(T_{i,1}^{cpu} - T_1^{oil}\right), \tag{11}$$

where h^{cpu} is the heat transfer coefficient and $A_{i,1}^{cpu}$ is the contact area. Similarly, the heat exchanged between the middle oil layer and the processing units immersed in it is:

$$P_{cpu,i,2}^{oil,2} = h^{cpu}A_{i,2}^{cpu}\left(T_{i,2}^{cpu} - T_2^{oil}\right). \tag{12}$$

In order to take into account also the effect of the oil flowing from the middle layer to the top one due to the natural convection, we consider also a fictitious heat flow which goes from the lower CPUs and the top oil layer:

$$P_{cpu,i,2}^{oil,1} = h^{conv}A_{i,2}^{cpu}\left(T_{i,2}^{cpu} - T_1^{oil}\right), \tag{13}$$

where h^{conv} is the fictitious heat transfer coefficient and $A_{i,2}^{cpu}$ is the fictitious contact area. The heat flowing from the i-th oil layer and the portion of the fictitious heat exchanger immersed in it is expressed by:

$$P_{oil,i}^{hx,i} = h^{hx}A_i^{hx}\left(T_i^{oil} - T_i^{hx}\right), \tag{14}$$

where h^{hx} is the heat transfer coefficient and A_i^{hx} is the contact area. The thermal power flowing between two consecutive portions of the immersed heat exchanger is due to mass transfer:

$$P_{hx,i+1}^{hx,i} = \dot{m}c_p\left(T_i^{hx} - T_{i+1}^{hx}\right), \tag{15}$$

where $\dot{m}$ is the mass flow rate of the cooling water and c_p is its specific heat at constant pressure. Similarly, the inlet and the outlet power flows are given by:

$$P_{in}^{hx} = \dot{m}c_p\left(T_3^{hx} - T_{in}^{hx}\right), \tag{16}$$

$$P_{out}^{hx} = \dot{m}c_p\left(T_{out}^{hx} - T_1^{hx}\right), \tag{17}$$

respectively.

MODEL CALIBRATION AND VALIDATION

The equations (9)-(17) that represent the power flows that occur in the LIC system depend on both geometrical parameters and heat transfer coefficients. While all the geometrical quantities have been obtained by direct inspection of the experimental setup, the heat transfer coefficients have been estimated from experimental data acquired by conducting an ad-hoc experimental campaign. More specifically, we have designed two different experimental tests where we have varied the computational load of the servers and the flow rate of the cooling water in order to excite the system in different operating conditions, as shown in Fig. 4. By exploiting the monitoring system included in the testbed, the data gathered during the two experiments have been collected with sampling time $\tau_s = 15s$ in two different datasets, which contain the following measured quantities:

$$\boldsymbol{m} = \left[T_1^{oil}\;\; T_2^{oil}\;\; T_3^{oil} \middle| T_{1,1}^{cpu}\;\ldots\; T_{4,2}^{cpu} \middle| T_{out}^{hx}\right]^T. \tag{18}$$

By following the standard cross-validation procedure to calibrate the model and assess its ability to predict new data, we have split the data acquired during the two experiments in training and validation sets, as highlighted in Fig. 4. Then, we have found the optimal value of the heat transfer parameters vector:

$$\boldsymbol{\vartheta} = [g^{oil}\quad h^{amb}\quad h^{cpu}\quad h^{conv}\quad h^{hx}]^T \tag{19}$$

by solving the following minimization problem:

$$\boldsymbol{\vartheta}^{opt} = arg \min_{\boldsymbol{\theta} \in \boldsymbol{\Theta}} \sum_{k=1}^{K} \|\widehat{\boldsymbol{m}}^{tr}(k; \boldsymbol{\theta}) - \boldsymbol{m}^{tr}(k)\|^2 \qquad (20)$$

where $\boldsymbol{m}^{tr}$ is the vector of measured data contained in the two training sets, K is the number of training samples, and $\widehat{\boldsymbol{m}}^{tr}$ is the vector of simulated data obtained by propagating the model equations initialing the states at $\widehat{\boldsymbol{m}}^{tr}(1; \boldsymbol{\theta}) = \boldsymbol{m}^{tr}(1)$. Moreover, the domain $\boldsymbol{\Theta}$ is chosen in order to preserve the parameters physical meaningfulness.

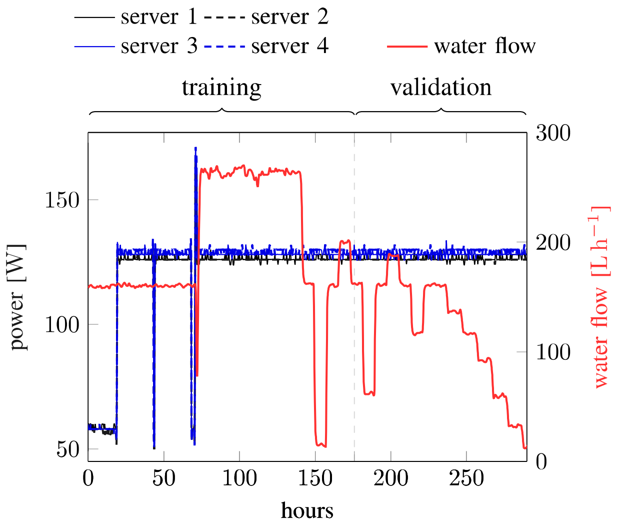

(a) Experimental test 1

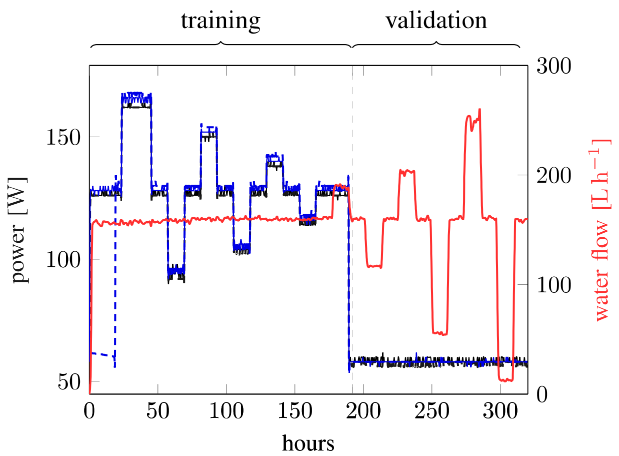

(b) Experimental test 2

Figure 4
Servers electrical power and water flow rate designed for the first (a) and the second (b) tests.

The solution of problem (20) leads to parameters values reported in Tab. 1. Once estimated the values of the model parameters, we have validated the capabilities of the model in predicting the outflowing water temperature data contained in the two validation sets.

Table 1
Values of the model parameters

Parameter	Estimated value	Unit
g^{oil}	1.12	$Wm^{-1}K^{-1}$
h^{amb}	4.04	$Wm^{-2}K^{-1}$
h^{cpu}	3.28	$Wm^{-2}K^{-1}$
h^{conv}	0.40	$Wm^{-2}K^{-1}$
h^{hx}	8.59	$Wm^{-2}K^{-1}$

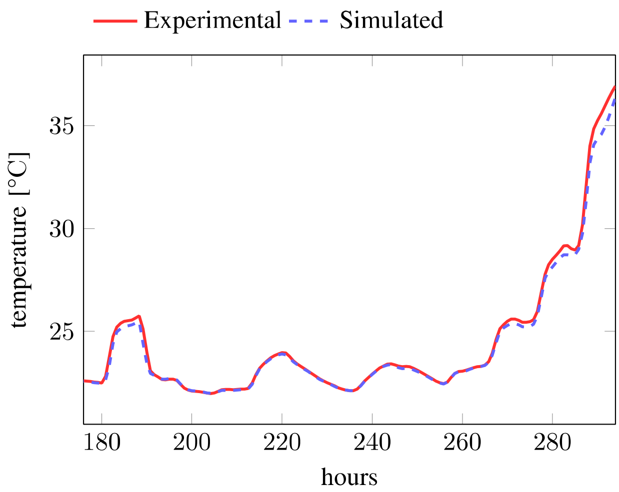

Figure 5
Comparison between measured and simulated outlet water temperature in the first validation dataset.

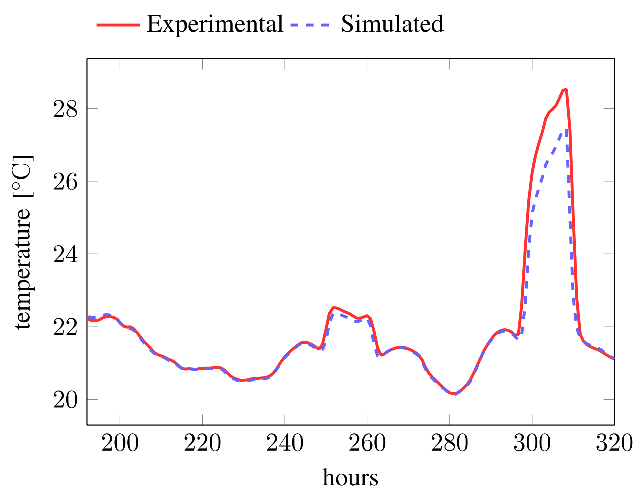

Figure 6
Comparison between measured and simulated outlet water temperature in the second validation dataset.

Fig. 5 and Fig. 6 report the comparison between experimental and simulated data contained in the first and second test datasets, respectively. As it can be observed, the model approximates satisfactorily the measurements of the outflowing water temperature, especially those contained in the first validation set. A slight underestimation of the water temperature can be observed just in the second validation set at around 300 hours. To improve even more the accuracy of the proposed model, more training data may be collected by conducting more meaningful experimental tests, where the system is excited in other different operating conditions. These considerations are confirmed by computing the Root Mean Squared Error (RMSE) between mesured and simulated outlet water temperatures: 0.26°C for the first validation dataset and 0.46°C for the second one.

CONCLUSIONS

In this work, we propose a control-oriented model for Liquid Immersion Cooling systems for data center applications. In particular, we have considered a graph-based modelling approach, that allows representing, through a directed graph, the energy storage and the power flows that occur within a physical system. The graph obtained by using this approach is easy-to-read and allows to derive straightforwardly the energy balance equation of each of the most relevant components of the physical system. By exploiting the availability of the unusual experimental testbed for Liquid Immersion Cooling, we have designed two different experiments to calibrate and validate the proposed model, and the results have shown the effectiveness of the considered modelling approach. Future directions will include the design and test of advanced control strategies for the considered LIC system and other more complex configurations including heat recovery units.

ACKNOWLEDGMENTS

We would like to thank all the staff at RISE SICS North for supporting this work with their time and technical expertise.

REFERENCES

1. Parolini, L., Sinopoli, B., Krogh, B. H., and Wang, Z., 2011, A cyber–physical systems approach to data center modeling and control for energy efficiency, Proceedings of the IEEE, 100, pp. 254-268.
2. Dai, J., Ohadi, M. M., Das, D., and Pecht, M. G., 2016, Optimum Cooling Of Data Centers, Springer.
3. Alkharabsheh, S., Fernandes, J., Gebrehiwot, B., Agonafer, D., Ghose, K., Ortega, A., and Sammakia, B., 2015, A brief overview of recent developments in thermal management in data centers, Journal of Electronic Packaging, 137.
4. Khalaj, A. H., and Halgamuge, S. K., 2017, A review on efficient thermal management of air-and liquid-cooled data centers: from chip to the cooling system, Applied energy, 205, pp. 1165-1188.
5. Chu, R. C., Simons, R. E., Ellsworth, M. J., Schmidt, R. R., and Cozzolino, V., 2004, Review of cooling technologies for computer products, IEEE Transactions on Device and Materials Reliability, 4, pp. 568-585.
6. Ebrahimi, K., Jones, G. F., and Fleischer, A. S., 2014, A review of data center cooling technology, operating conditions and the corresponding low-grade waste heat recovery opportunities, Renewable and Sustainable Energy Reviews, 31, pp. 622-638.
7. Chin, C. S., 2012, Computer-Aided Control Systems Design: Practical Applications Using MATLAB® and Simulink®. CRC Press.
8. Koeln, J. P., Williams, M. A., Pangborn, H. C., and Alleyne, A. G., 2016, Experimental validation of graph-based modeling for thermal fluid power flow systems, In ASME Dynamic Systems and Control Conference.
9. Moore, K. L., Vincent, T. L., Lashhab, F., and Liu, C., 2011, Dynamic consensus networks with application to the analysis of building thermal processes, IFAC Proceedings Volumes, 44, pp. 3078-3083.
10. Preisig, H. A., 2009, A graph-theory-based approach to the analysis of large-scale plants, Computers & Chemical Engineering, 33, pp. 598-604.
11. Koeln, J. P., Williams, M. A., and Alleyne, A. G., 2015, Hierarchical Control of Multi-Domain Power Flow in Mobile Systems-Part I: Framework Development and Demonstration, ASME Dynamic Systems and Control Conference.
12. Williams, M. A., Koeln, J. P., and Alleyne, A. G., 2015, Hierarchical control of multi-domain power flow in mobile systems-Part II: Aircraft application, ASME Dynamic Systems and Control Conference.