Thermo-Fluid Modeling Framework for Supercomputer Digital Twins: Part 1, Demonstration at Exascale

Vineet Kumar[1] Scott Greenwood[1] Wesley Brewer[2]
David Grant[3] Nathan Parkison[3] Wesley Williams[1]

[1]Fusion and Fission Energy and Science Directorate, Oak Ridge National Laboratory, USA,
{kumarv, greenwoodms, williamswc}@ornl.gov
[2]Computing & Computational Sciences Directorate, Oak Ridge National Laboratory, USA, brewerwh@ornl.gov
[3]Facilities & Operations Directorate, Oak Ridge National Laboratory, USA, {grantdr, parkisonjn}@ornl.gov

Abstract

A thermo-fluid modeling framework is being developed for ExaDigiT—an open-source framework for developing comprehensive digital twins of liquid-cooled supercomputers. The work is being conducted in two parts, and discussion is divided into two companion papers. The work documented in this paper focuses on the development of a cooling system library in Dymola for the Frontier supercomputer at Oak Ridge National Laboratory. The second part, outlined in a companion paper, focuses on a templating structure called *Auto-CSM* for easily creating model-agnostic, physics-based thermo-fluid cooling system models for liquid-cooled supercomputers using a text-based schema. The cooling model is being developed using primarily the open-source Transient Simulation Framework of Reconfigurable Models (TRANSFORM) library. The library follows the templating architecture developed within the TRANSFORM library for modeling subsystems. A full-system validation was performed to validate a very simple model that is integrated with the system controls, and the results are presented herein.

Keywords: cooling system, liquid-cooled supercomputers, Frontier, Modelica

1 Introduction

The electricity consumption of data centers is projected to increase in the United States from around 200 TWh in 2022—which represents about 4% of the country's total electricity demand—to 260 TWh, which is expected to be around 6% of the total electricity demand (IEA 2024). Additionally, water consumption is expected to be significant for both direct and indirect liquid-cooled supercomputing clusters. It is projected that approximately 15–27% of the energy consumed by data centers can be reduced via advanced cooling technologies, such as natural and liquid cooling techniques (Zhu et al. 2023). Therefore, there is an acute need for dynamic modeling of liquid-cooled supercomputing clusters using open-source tools like Modelica. The potential use cases for such a model could include the design and commissioning phase of new liquid-cooled data centers, as well as for operational optimization of existing facilities (Todd et al. 2021). ExaDigiT is a comprehensive open-source framework under development at Oak Ridge National Laboratory (ORNL) that focuses primarily on liquid-cooled supercomputers. Figure 1 shows the high-level architecture of the ExaDigiT framework, which consists of three main modules: (1) the cooling model discussed here, (2) a resource allocator and power simulator (RAPS), (3) a visual analytics module consisting of both an augmented reality component for 3D interactive visualization and a web-based dashboard for launching experiments and creating 2D plots of power and cooling behavior. The development of ExaDigiT is currently centered around the 2 exaflop Frontier supercomputer, which was deployed in 2022 at ORNL's Oak Ridge Leadership Computing Facility (OLCF) (Atchley et al. 2023).

Most of the work on data center cooling has been focused on air-cooled systems—see, for example, (Lee and Chen 2013; Ham and Jeong 2016; Fu, Wetter, and Zuo 2018)—especially on cooling efficiency. Zhang et al. (Zhang et al. 2022) built a digital twin for air-cooled data centers using a combination of computational fluid dynamics (CFD) using 6SigmaDC (*DataCenter Design Software* n.d.) and an AI-based XGBoost model. They used AI to optimize the control parameters of the air conditioning system so as to optimize the power usage effectiveness (PUE) of a data center. Heydari et al. (Heydari et al. 2022) performed extensive analysis of secondary flow loops to deploy liquid-cooled systems in air-cooled data centers by a combination of numerical modeling and experimental testing of four different cooling loops. The numerical modeling for different cold plate designs was performed using a commercial CFD solver 6SigmaET (*Data-*

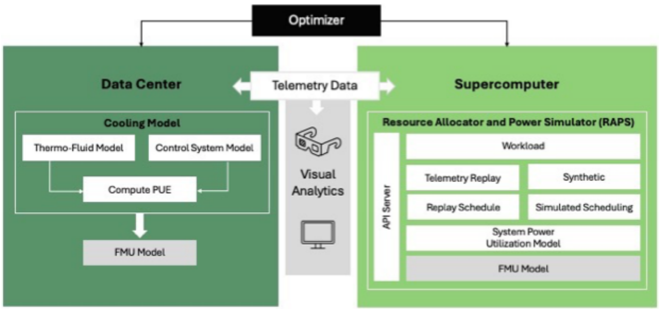

Figure 1. ExaDigiT architecture overview (Brewer et al. 2024)

Center Design Software n.d.), and flow network modeling of a liquid-cooled rack was performed with a custom system model and the commercial CFD solver 6SigmaRoom (*DataCenter Design Software* n.d.). Modi et al. (Modi et al. 2023) performed transient CFD simulations to optimize different flow configurations for rack-level models using the commercial CFD solver 6SigmaET. Modelica has been previously used by Fu et al. (Fu, Wetter, and Zuo 2018) to model air-cooled data center systems, which also used the Modelica buildings library (MBL) (Wetter et al. 2014). Leva et al. (Lee and Chen 2013) developed an open-source Modelica library using an object-oriented modeling (OOM) framework to model both air-cooled and liquid-cooled supercomputing clusters, and it can couple with 3D-ICE for chip simulations.

The current study primarily uses the Transient Simulation Framework of Reconfigurable Models (TRANSFORM) library, which is a Modelica-based open-source library developed at ORNL to enable rapid development of dynamic, advanced energy systems with an extensible system modeling tool (M. S. Greenwood 2017). TRANSFORM is organized as a series of packages, each of which has a general application; the library sub-categorizes models within each package, which helps users easily locate a component (M Scott Greenwood et al. 2020). Additional details on the TRANSFORM library, including model templates and the supervisory control system, can be found in previous work by the authors (M. S. Greenwood 2017; Michael Scott Greenwood et al. 2017). TRANSFORM was developed using the commercial integrated development environment (IDE) Dymola by Dassault Systèmes (Systèmes 2022) but should be compatible with other IDEs that are compatible with the Modelica specification 3.4+ (M. S. Greenwood 2017). In the current study, Dymola was used as the IDE. The current library has an additional dependency on the Buildings library (Wetter et al. 2014) for the variable speed cooling tower model.

This study is divided into two parts. The objectives of this work that form part one of this study are (1) to demonstrate a use case of the templating structure that is being laid out for modeling liquid-cooled supercomputing clusters in part 2—documented in a companion paper—and (2) to perform a validation exercise of the overall model using telemetry data. The supercomputing cluster chosen for validation is that of Frontier at ORNL. The validation exercise is divided into two parts: a component validation effort and the overall validation effort. The validation for the overall model serves to demonstrate how this library can be used to create very simple system models upon which additional complexity can be layered.

2 Frontier Model Description

2.1 Physical Facility Description

Frontier consists of 74 liquid-cooled HPE Cray EX supercomputing cabinets, which hold a total of 9472 compute nodes (Atchley et al. 2023; Choi 2022). Each Frontier node, designated as Cray EX 235a, contains one AMD 7A53 EPYC^{TM} 64-core "Trento" processor and four AMD Instinct MI250x GPUs. Each cabinet of Frontier consists of four shelves, each shelf has two chassis, and each chassis contains four active rectifiers and eight compute blades—a total of 64 blades and 32 rectifiers per cabinet (Brewer et al. 2024). Each cabinet is directly supplied with three-phase power from the distribution transformer switchboard, which is converted from 480 V AC to 380 DC voltage using AC-DC rectifiers and is subsequently stepped down to 48 V DC using super intermediate voltage converters (SIVOCs). The 48V DC is what supplies power to the node. Since each blade contains two nodes, there are two SIVOC converters per blade. Each HPE Cray EX cabinet can support up to 400 kW of power. Further details regarding the Frontier compute architecture can be found in (Atchley et al. 2023), whereas details regarding the Frontier system power conversion, including modeling conversion losses, are covered in forthcoming papers (Brewer et al. 2024; Wojda et al. 2024).

A simplified schematic of the overall cooling system layout for Frontier is shown in Figure 2; locations are marked in the figure to indicate the current cooling model outputs. The cooling system can be divided into three cooling loops, referred to in this paper as the cooling distribution unit (CDU) loop, the intermediate or high-temperature water (HTW) loop, and the cooling tower water (CTW) loop. The CDUs are used to remove heat from the compute nodes via forced convective liquid cooling. The CDUs deployed for the Frontier system can remove approximately 1.6 MW of heat, which translates to four cabinets. In practice, each CDU serves three cabinets at most, and some CDUs serve two cabinets. In total, there are 27 CDUs installed in the data center to serve the compute cabinets; of these, 25 are in operation, serving the 74 cabinets. The data center room contains piping underneath the raised floor which distribute primary flow to the CDUs. Additionally, the room also contains air handling systems and other auxiliary systems; these systems are not covered here, as they were not modeled. These systems serve a number of functions, including maintaining the dew point temperature in the data center room within an adjustable setpoint. Therefore, they will be considered in a more detailed modeling effort in the future.

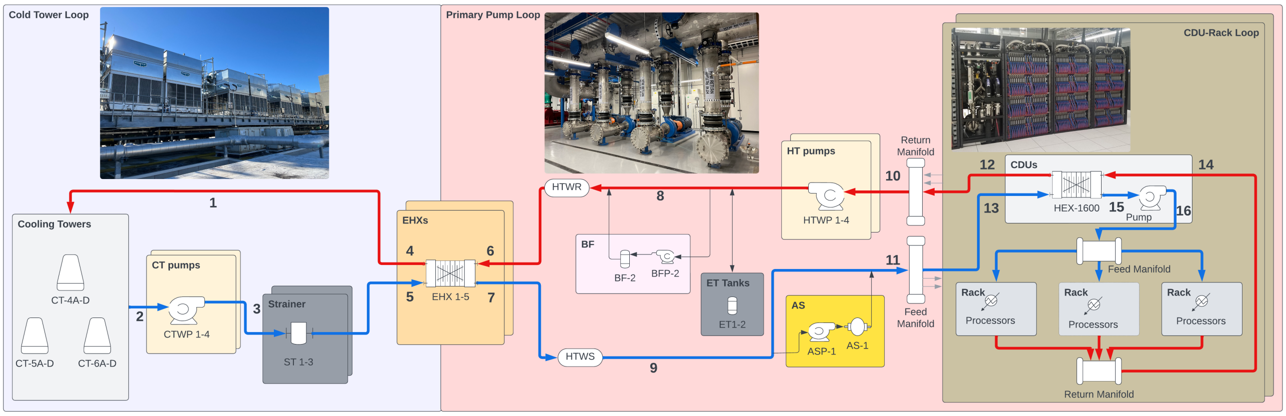

Figure 2. Simplified schematic of Frontier cooling system. Locations at which the cooling model predicts pressures, temperatures, and flow rates are numbered

From the current modeling perspective, the primary flow pumped from the HTW loop is used to remove the CDUs heat load generated primarily by the compute nodes in the CDU heat exchanger (shown as HEX-1600 in Figure 2). The secondary side of the CDU supplies pressurized liquid coolant via two pumps in parallel with a flow rate of $\sim$ 250–300 gpm at a system (gauge) pressure of $\sim$ 3.5 bar. This supply coolant is distributed in parallel to the cabinets/racks. In each rack, the flow passes through the 64 compute blades and the 32 rectifiers, with both actively cooled. Each compute blade consists of CPU and GPU cold plates, as well as cooling channels for peripheral components from a heat load perspective, such as memory, network interface cards (NICs), and SIVOCs. The hot water from the return side of the secondary CDU loop is cooled in the heat exchanger, which completes the loop. There is a tank on the secondary side to maintain the system pressure, as well as instrumentation that consists principally of temperature sensors, pressure gauges, and flowmeters. The primary side has a control valve that regulates the flow into the CDU based on the supply-side temperature, which is discussed further in Section 2.3.

The intermediate loop (or the HTW loop) provides up to 40 MW of process cooling at around 12–32 °C. The intermediate loop mainly consists of four variable-speed HTW pumps (HTWPs), expandable to 8 pumps; five economizer heat exchangers (EHXs) (four are considered in the model), expandable to 8 heat exchangers; and associated piping, which connects to the data center room for supply and return. The intermediate loop also consists of air separator pumps (denoted as ASP in Figure 2), HTW water bag filter (denoted as BF in Figure 2), and other systems that are peripheral to the current modeling perspective. The HTWPs supply water to the data center room at approximately 5000–6000 gpm at a gauge (system) pressure of $\sim$ 6.2 bar. The majority of this flow is directed toward the CDUs. About 10–15% of the flow is diverted through the bypass flow valve when system flows are low, and is mixed with the hot return, which results in a slightly decreased bulk return temperature. Some of the HTW flow also provides cooling for the Orion file system ($\sim$ 10%) when operating with supply temperatures less than 18 °C. This aspect was not modeled but would be considered in the future. The piping network in the central energy plant and the data center room are vast and complex. Here, considerable simplifications were made to capture the essence of the piping network; these simplifications are discussed in Section 2.3.

The CTW loop principally consists of five counter flow cooling towers, expandable to ten cooling towers, and four variable-speed CTW pumps (CTWPs), expandable to ten pumps. Each cooling tower (CT) has four independent cells with individual control valves and four corresponding variable-speed CTs fans. The CTWPs supply water at approximately 6000–10000 gpm at a gauge (system) pressure of approximately 1.5 bar. The CTW loop also consists of strainers for cooling water blowdown as well as other systems associated with chemical treatment of the cooling tower water, which are not shown in Figure 2. These systems are currently ignored in the model. The main flow path in the CTW loop (see Figure 2) starts from the hot water from the primary return side of the EHX, which flows to the cooling towers located on the roof of the central energy plant; the cold return from the cooling towers is pumped to the primary supply side of the EHXs. In the current model, only four cooling towers were modeled (i.e., 16 independent cells). This is a reasonable approximation, as generally 9–15 cells were in operation for the range of data that was analyzed. However, the current model could be easily expanded to include the additional CT or any of the other components with the templating system that is being put into place to easily generate Modelica models for liquid-cooled supercomputing clusters (S. e. a. Greenwood 2024).

2.2 Library Structure for the Frontier Model

The Modelica model library for the Frontier system is currently hosted in an internal ORNL Git repo: `https:`

//code.ornl.gov/exadigit/coolingModel. It is expected to be open-sourced within the next few months. The library relies primarily on components from the TRANSFORM model, as previously discussed. An additional dependency is the Buildings library from Lawrence Berkeley National Laboratory (Fu, Zuo, et al. 2019) for the variable-speed CT model. The library is being developed in Dymola (Systèmes 2022) as the TRANSFORM library was developed in Dymola. In the future, the library will be extended to work with Open Modelica. The library structure for the Frontier cooling model follows from the subsystem templating approach used in TRANSFORM library. This structure allows for ease of modeling integrated systems, and it is extensively discussed in previous work by Greenwood (M. S. Greenwood 2017). A sample of the subsystem templating is shown in Figure 3 for the CDU subsystem.

In the figure, the left-hand side shows the structure of the library for the CDU package, which opens to a directory structure with the following packages: 'Examples', 'Components', 'Data', 'ControlSystems', and 'BaseClasses'. The physical models that are located in the top-level directory extend from the 'Partial_SubSystem model' within the 'BaseClasses' package. The 'BaseClasses' directory also defines signal buses and actuator buses for the input and output signals, respectively, from the 'Partial_SubSystem model' to the 'Partial_ControlSystem model'. The models within the ControlSystems package are inherited from the 'Partial_ControlSystem' model. Similarly, the data records within the 'Data' package extend from the 'Record_Data' record within the 'BaseClasses' package. The 'Examples' directory contains example tests for the physical models. This type of layout, which exploits the inheritance and replaceability features of Modelica, allows for easily layering complexity. As an example, the compute cabinet model shown in Figure 3 is inherited from 'PartialCabinetModel' and is a simple model. This simple model can easily be replaced with a better-resolved model which also inherits from 'PartialCabinetModel', without requiring the modification of any higher-level models. More details on the layout can be found in (M. S. Greenwood 2017). A user of this library could easily copy an existing model, the CDU model and adapt it to their supercomputing facility as long as the input structure to their model remain unchanged. As noted earlier, the AutoCSM model is specifically being developed for this purpose.

The three main loops—that is, the CDU loop, the HTW loop, and the CTW loop—follow a similar structure to those of the CDU model and are part of the systems package. The reason why the systems package does not contain subsystem-specific models for blades and cabinets is that the control systems operate only at the level of the CDU. Beyond the 'System' package, there are four other top-level packages in the library: 'SubComponents', 'Examples', 'Icons', and 'FMUs'. The 'Examples' package consists of integrated subsystem model tests. The 'Icons' package defines the icons used for the various subsystems, as well as some of the components. The 'Controls' package hold the sub-models used in the control system models in the systems package, as well as unit tests. The 'SubComponents' package mainly contains packages for the 'Media' used throughout the library and a 'Fluid' package that holds the major components used in the different loops—CTs, heat exchangers, and pumps—along with unit tests for each. A cold plate model is currently under development within the 'Fluid' package.

Two fluid media have been used with the models in the library, specified in the 'Media' package within the 'SubComponents' package. The first, taken from the Modelica Standard Library, is water with constant properties, and the second is water with linear properties, which was taken from the TRANSFORM library. The principal disadvantage of the constant properties medium is that a tank must be modeled that fixes the pressure at the location of the tank. Therefore, the linear properties model was used in the current study as a balance between the constant properties model and the more comprehensive water model from the Modelica Standard library, which often presents convergence issues. In the actual system, the coolant has additives that inhibit bacterial growth. Here, it is assumed that the thermal properties of the medium are largely unchanged with the addition of a small concentration of additives.

The system model that was built for the current study was simplified by replacing pipes with a combination of fluid volumes and hydraulic resistances to model fluid mass and pressure drop, respectively. The hydraulic resistances were tuned with telemetry data. The heat exchangers were also approximated as a combination of fluid volumes and hydraulic resistances in series for each stream, with thermal data obtained from the manufacturer either as performance curves of overall heat transfer coefficient as a function of flow rates or constant values of overall heat transfer coefficient. For the pump models, pump curves were obtained from publicly available manufacturer data for all the pumps at their corresponding nominal pump speeds. The system controls were incorporated into the simple model, as the focus was to accurately capture the system dynamics. It must be noted that in the absence of pipes, the model lacks accuracy in capturing fluid flow dynamics, especially during sharp transitions. However, the compromise was a model with a shorter run time that can reasonably capture system dynamics for the range of data tested. Future extensions to this model could easily be made to incorporate pipe flow models with the templating structure in place.

The present model only simulates to the level of the CDU, and the cabinets are represented as a combination of simple volume components and hydraulic resistances from the TRANSFORM library (shown as 'Compute Cabinets' in Figure 3). Similar approximations have been used to model the heat exchangers, whereby the fluid flow in the primary and secondary streams is replaced by volumes

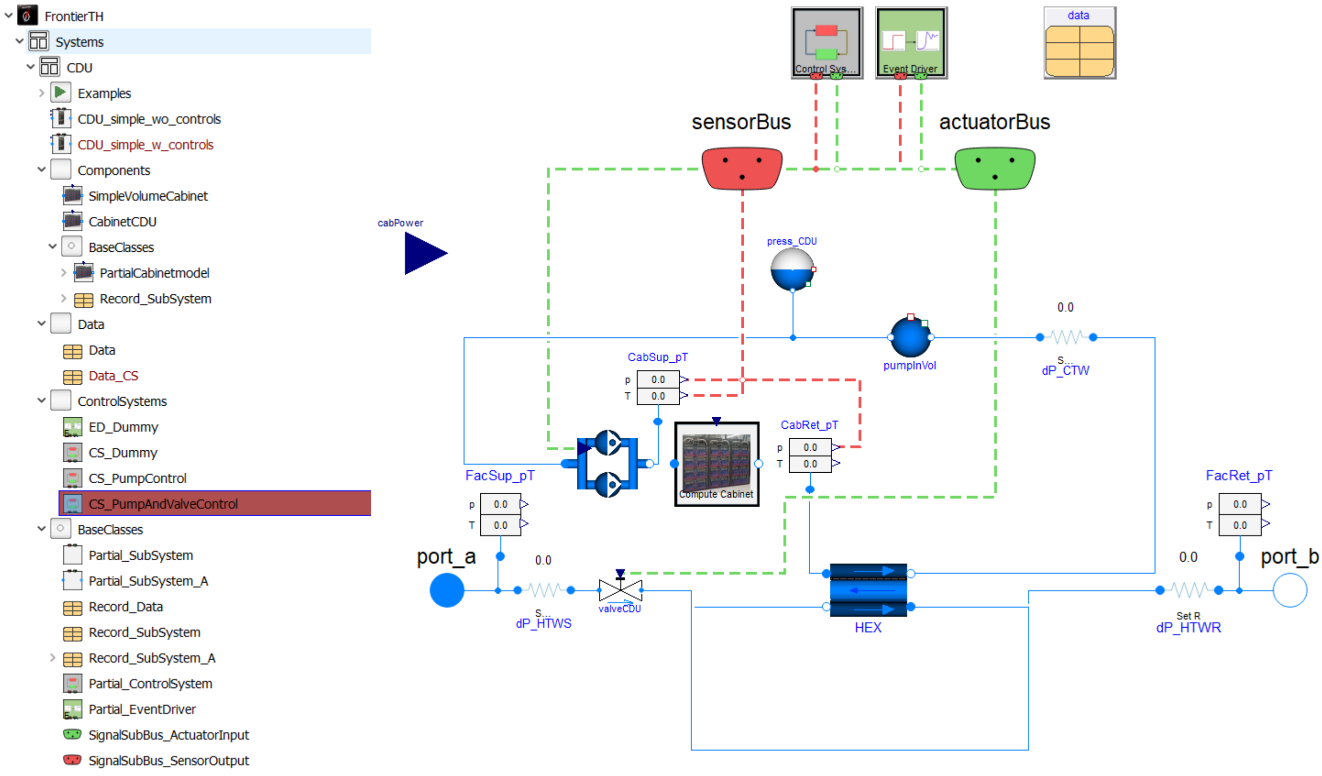

Figure 3. CDU model in Dymola based on the subsystem specific TRANSFORM library template (M. S. Greenwood 2017)

and resistances, and the overall heat transfer coefficients are taken from data. Lastly, publicly available data were used for all the pump curves.

2.3 Controls System Modeling

Figure 4 depicts the simplified control system logic currently implemented for the cooling model. The control system logic is divided into the central energy plant and the data center. A detailed overview of the control system logic is beyond the scope of this paper. The Modelica model captures the essentials of the control logic, which activates once the auto-operation of the physical cooling system has commenced after the start-up sequence has been completed. At this stage, the fail-safe configuration has not been implemented into the control logic, but the existing model could easily be extended to include it. The control logic for the system can be described as follows from Figure 4. Any disturbance in the CDU loop in terms of changes in the load, the HTW inlet pressure, or the HTW supply temperature would trigger the control system to bring the system back to an operational set point. Any given CDU can regulate its primary valve as its compute demand changes, and, consequently, the demand for more or less coolant flow in the primary side is regulated by the speed of the HTWPs via the differential pressure setpoint in the HTW loop. HTWPs stage up or down at a given moment depending on the % relative speed of the pumps currently in operation. A change in the primary supply temperature, on the other hand, is regulated by the CT loop by staging the number of CTs up or down. CTs stage up when the CTW return (CTWR) header pressure is at its maximum boundary and the HTW supply (HTWS) temperature is increasing, and, consequently, they stage down when the CTWR header pressure is at its minimum set point and the HTWS temperature is decreasing. Additionally, EHXs are staged up or down depending on the number of CTs in operation. Therefore, the criteria to achieve HTWS temperature stability inform both the staging of the CTs directly and the EHXs indirectly. Finally, the CTWPs maintain the header pressure in the CT loop within certain bounds by regulating its speed and staging the number of pumps in operation. Once the HTWS temperature has stabilized within certain bounds and the differential pressure setpoint in the HTW loop is met, the compute CDU is satisfied, as shown in Figure 4.

A proportional-integral-derivative (PID) controller is used to regulate the CDU relative % pump speeds based on the CDU loop differential pressure in the current model. Both the CDU pumps are assumed to be in operation at all times with the same speeds. This is a reasonable approximation based on telemetry data. A control value on the primary side, as just discussed, is used to regulate the primary coolant flow based on the secondary supply temperature setpoint. A snapshot of the controls used for the CDU loop is shown in Figure 5. At the start of the simulation, the CDU pumps and the control valve are fixed for numerical considerations and a 1.0 second delay clock is used to switch to the PID mode for both the pumps and the valve. Additional low-pass filters are employed to filter the output from the PIDs to prevent very high oscillatory

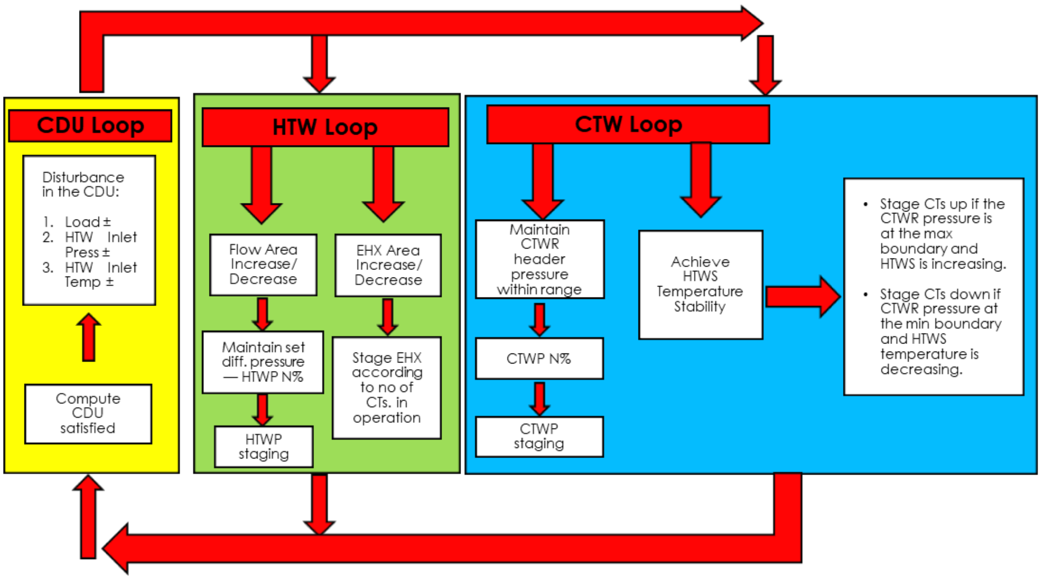

Figure 4. Frontier cooling model controls summary

behavior and to ensure numerical convergence.

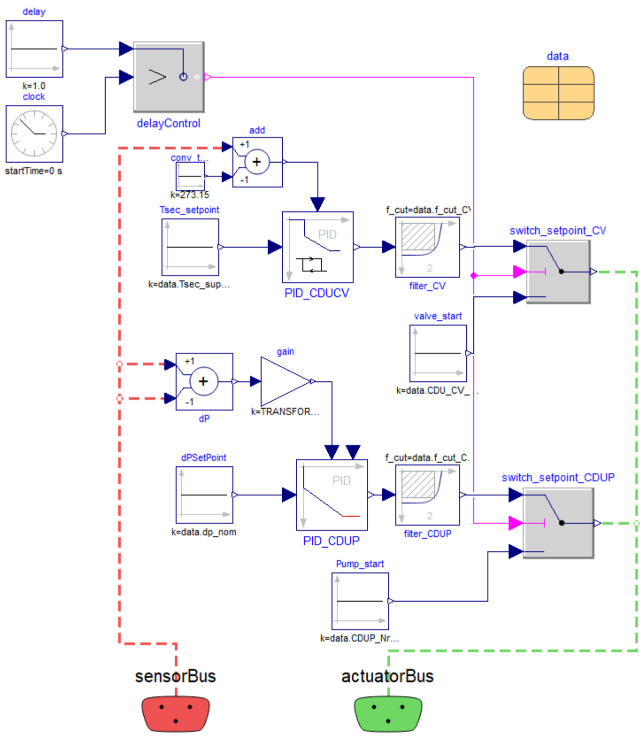

Figure 5. CDU controls model snapshot in Dymola

The intermediate loop similarly has a PID controller for the four HTWPs. The HTWPs are staged up or down depending on the relative % pump speeds of the running pumps, as previously discussed. The controls logic has been templated where possible, such as for the staging of the pumps. This should allow for easier extensions to other systems by users of the library. It must be noted that the HTW loop relies on the number of CTs in operation for staging the EHXs, and the staging of the CTs, in turn, relies on the HTWS temperature. This cross-flow of information exchange among subsystems is handled in the model by using a delay transfer function between the intermediate loop and the CT loop. A future scope would be to optimize the control parameters to achieve better system stability as well as respond quickly to a surge in the load.

3 Validation of the cooling model for Frontier

3.1 Component validation

Component validation was conducted as a first step before undertaking the full system validation. Component Verification & Validation (V&V) tests are important to ensure that model performs as expected—firstly, to validate the range of the expected telemetry data, and secondly to validate the components individually using telemetry data as boundary conditions before integration into the larger system. These tests can help to identify input errors in component parameters such as pump curves and heat exchanger data. Verification tests were not explicitly performed because the majority of the components used in the library were taken primarily from the TRANSFORM library and one component from the Buildings library, with modifications made in the margins which do not warrant a thorough verification process.

An example of component validation for the counterflow CDU heat exchanger is shown in Figure 6. This validation was performed using 22000 seconds of telemetry data for the primary flow rate, primary return pressure, primary supply temperature, secondary flow rate, secondary supply temperature, and secondary return pressure. As shown in the figure, the model can predict the primary (Facility) and the secondary (Cabinet) return temperatures with reasonable accuracy. The over-prediction in the facility return temperature is most likely because of the secondary flowmeter instrument uncertainty. The lower end of the prediction uncertainty shown in Figure 6, matches the prediction data well for the majority of the ∼ 6 hour snapshot.

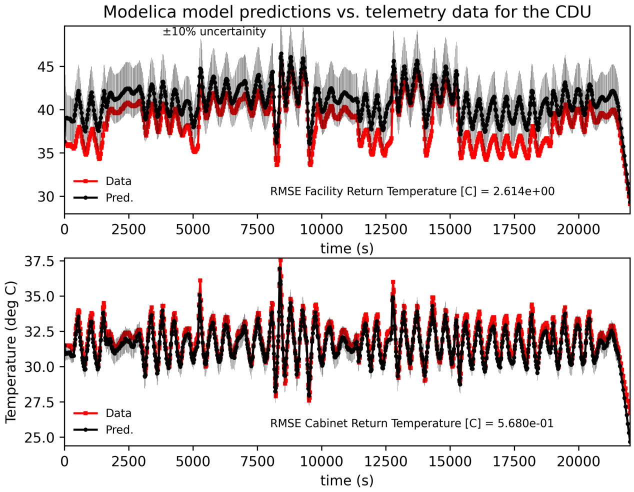

Figure 6. Modelica model predictions vs. telemetry data for the CDU return temperature for a ∼ 6 hour snapshot

Another validation test, shown in Figure 7, was used to test the control logic for the CTWP speed control, which is regulated by the CTWR header pressure setpoint. Similar to the CDU heat exchanger validation, telemetry data of the CTWR header pressure was used to test the model. The model can respond to the sharp changes in the pump speed after approximately 17 hours for this particular dataset but essentially filters the header pressure for smaller variation. Further improvements are required in the model to match the CTWP speed variation for changes of smaller magnitude the CTWR header pressure.

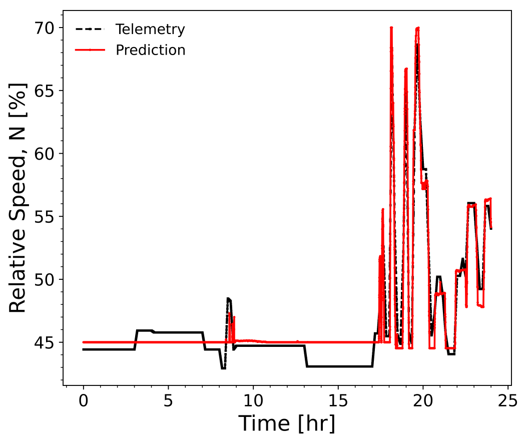

Figure 7. Modelica model predictions vs. telemetry data for the CTWP relative % pump speed for a 1 day snapshot

3.2 Complete system validation

A validation study of the entire *cooling model* (shown in Fig. 8) was conducted using ∼15 hours of telemetry data for a given day from the central energy plant and the data center down to the level of the CDU. The only inputs to the model were the power to the 25 CDUs and the wet bulb temperature as a function of time. The run time of the model exported as an an Functional Mock-up Unit (FMU) to simulate ∼15 hours of telemetry data on a standard windows workstation was about ∼20 minutes. Significant improvements have been made to the model to reduce the number of nonlinearities, and the current model runs 2-3x faster than the model discussed here. The translated model discussed here has 912 states with ∼14000 time varying variables and ∼13000 alias variables. The model outputs for the model exported as an FMU, have been listed in Table 1. The power to the CDUs for the validation exercise was calculated from the heat removed by the cooling water using telemetry data. This was found to closely match the power sensor data, which were obtained after the validation study was performed. After the Resource Allocator and Power Simulator (RAPS) module being developed as part of the ExaDigiT digital twin framework has been thoroughly validated with power sensor data, it will be utilized as an input to the cooling model.

Parameter	Description
Pressure	at locations shown in Figure 2
Temperature	at locations shown in Figure 2
Flow rate	at locations shown in Figure 2
Pump speed	% speed of HTWPs & CTWPs
Pump staging	operational HTWPs & CTWPs
EHX staging	operational EHXs
CT staging	operational CTs
Power consumption	HTWPs, CTWPs & CT fans
PUE	for the entire facility

Table 1. Outputs for the cooling model exported as an FMU.

The first second of predictions was removed from the analysis of ∼15 hours because the model predictions are biased by initial conditions at the very start of the transient. The resolution of the telemetry data varied from 15 seconds at the level of the CDU to 10 minutes for some of the facility telemetry data. For consistency, all measurements and predictions were interpolated to 15 seconds intervals for comparison. The annotations on the plots in Fig. 8 (a–e) correspond to stations in Fig. 2. A few observations can be made when comparing model predictions with telemetry data. The Frontier system was idle for about half a day because of system upgrades, which is why the cooling system load is at a minimum beyond ∼ 10 hours. This coincidentally proved to be a good transient test to see how the model performs in transition from a loaded system state to an idle state. For most of the predicted parameters, some of which are shown in Fig.

8, the trend-wise predictions are good up to about ∼ 10 hours, after which some deviation occurred in the predicted facility parameters when the power to the CDUs is at a minimum. This can be confirmed by the good agreement in the primary flow rates (shown in Fig. 8(a)) up to ∼ 10 hours. However, in the physical system, the primary flow in the intermediate loop is maintained at a minimum of ∼3000 gpm, whereas the model predicted a minimum flow of ∼2300 gpm (not shown). This flow difference is attributed to an additional bypass flow of ∼700 gpm in the actual system. Improving this aspect of the model should improve the prediction when the system transitions to an idle state. The model predictions for the CDUs secondary supply temperatures (not shown) show greater fluctuation than the physical system, which does a good job maintaining the temperature at the setpoint. This result suggests that the controls for the primary valve in the CDU must be further investigated. The staging of the HTWPs is predicted to occur earlier than it does in the actual system. The staging of the CTWPs and the CTs (not shown) must also be improved. It must be noted that there are manual overrides within the system, such as those for manually staging the CTs, which were deployed during this particular transient. This is a feature that could be introduced into the model in the future. Overall, the root mean square error (RMSE) of the parameters shown in Fig. 8 are within reasonable bounds, and a future study would focus on model uncertainty.

Finally, the comparison between the PUE predicted by the model and that calculated from telemetry data is shown in Fig. 8(f). The predicted PUE is within four percent of the calculated PUE, within the range of data tested (∼8.3 hours). It must be noted that in both the calculations, the auxiliary systems considered for power consumption are the following: CDU Pumps (CDUPs), HTWPs, CTWPs, and CT fans. Other auxiliary systems such as the air-handling system are not considered in the calculation, as they were not modeled. Therefore, it is expected that the actual PUE would be higher.

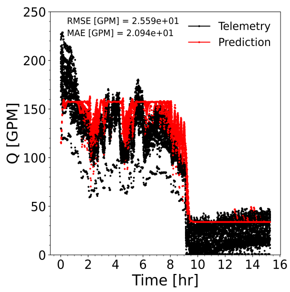

(a) Primary CDUs flow rate predictions (Station 12 in Fig. 2)

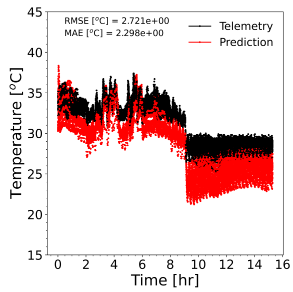

(b) Primary CDUs return temp. predictions (Station 12 in Fig. 2)

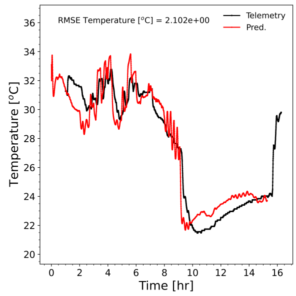

(c) HTW return temperature prediction (Station 10 in Fig. 2)

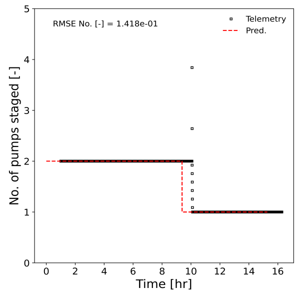

(d) HTWP staging (Station 10 in Fig. 2)

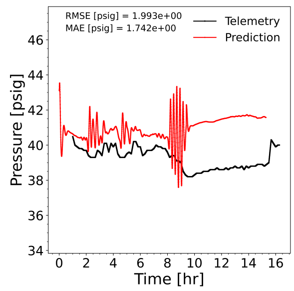

(e) HTW return pressure predictions (Station 10 in Fig. 2)

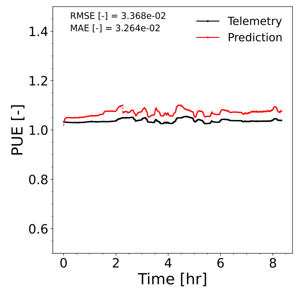

(f) PUE Modelica model predictions

Figure 8. Modelica model predictions (exported as an FMU) vs. telemetry data for the CDU and the central energy plant.

4 Conclusions

This paper presents the cooling model that is being developed in Modelica using Dymola as part of the ExaDigiT project to develop digital twins for liquid-cooled exascale supercomputers. The cooling model is being developed using primarily the open-source TRANSFORM library developed at ORNL, with the cooling tower model from the Buildings library. The overall goal is to develop a templating structure, Auto-CSM, for creating physics-based thermo-fluid cooling system models. Auto-CSM seeks to streamline the creation of cooling system model (CSM) for integration into the ExaDigiT framework, and that work is covered in Part 2 of the study and is documented in a companion paper. While Auto-CSM is being developed, the current study (Part 1) focuses on the cooling system library that is being developed and demonstrated on the cooling system of the 2 exaflop Frontier supercomputer at ORNL. The library follows the templating architecture developed within the TRANSFORM library for modeling subsystems and integrating them to quickly model complex systems. Although the library is currently hosted in an internal Git repo, it is expected to be open-sourced within the next few months.

The subsytems used for the cooling system library to model Frontier are the following: the CDU loop, the HTW loop, and the CT loop. It remains to be seen how generalizable these subsystems are to other supercomputing clusters. The simplified model, which makes use of fluid volumes and hydraulic resistances in place of pipes, extends only to the level of the CDU and is integrated with system controls. Extending the model to the level of the compute blade would result in a more accurate thermal response prediction with the downside of an increased run time. A component validation was conducted before performing the full model validation with telemetry data for an approximately 15 hour snapshot that was provided for a given day. The model performed reasonably well, especially when the system was loaded, and significant improvements have been made to improve model performance in terms of runtime and robustness. Future use cases for such a model could be both in the design phase when designing new systems or optimizing the operation of existing systems.

Acknowledgments

This research was sponsored by and used resources of the Oak Ridge Leadership Computing Facility (OLCF), which is a DOE Office of Science User Facility at the Oak Ridge National Laboratory (ORNL) supported by the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

References

Atchley, Scott et al. (2023). “Frontier: Exploring Exascale”. In: *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, pp. 1–16.

Brewer, Wesley et al. (2024-11). “A Digital Twin Framework for Liquid-cooled Supercomputers as Demonstrated at Exascale”. In: *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*.

Choi, Charles Q. (2022). *Beneath the Hood of the First Exascale Computer*. https://spectrum.ieee.org/frontier-exascale-supercomputer. [Online; accessed 4-April-2024].

DataCenter Design Software (n.d.). Available at https://www.cadence.com/en_US/home/resources/datasheets/datacenter-design-software-ds.html. Accessed: 2024-03-12.

Fu, Yangyang, Michael Wetter, and Wangda Zuo (2018). *Modelica models for data center cooling systems*. Tech. rep. Univ. of Colorado, Boulder, CO (United States).

Fu, Yangyang, Wangda Zuo, et al. (2019). “Equation-based object-oriented modeling and simulation of data center cooling systems”. In: *Energy and Buildings* 198, pp. 503–519.

Greenwood, M Scott et al. (2020). “Demonstration of the advanced dynamic system modeling tool TRANSFORM in a molten salt reactor application via a model of the molten salt demonstration reactor”. In: *Nuclear Technology* 206.3, pp. 478–504.

Greenwood, M. S. (2017-09). *TRANSFORM - TRANsient Simulation Framework of Reconfigurable Models*. Computer Software. Accessed on August 30, 2023. DOI: 10.11578/dc.20171025.2022. URL: https://github.com/ORNL-Modelica/TRANSFORM-Library.

Greenwood, Michael Scott et al. (2017-08). “A Templated Approach for Multi-Physics Modeling of Hybrid Energy Systems in Modelica”. In: DOI: 10.2172/1427611. URL: https://www.osti.gov/biblio/1427611.

Greenwood, Scott et al. (2024-10). “Thermo-fluid Modeling Framework for Exascale Supercomputing Digital Twins: Part 2, Automated Cooling Models”. In: *Proceedings of the American Modelica Conference 2024*. Under review.

Ham, Sang-Woo and Jae-Weon Jeong (2016). “Impact of aisle containment on energy performance of a data center when using an integrated water-side economizer”. In: *Applied Thermal Engineering* 105, pp. 372–384.

Heydari, Ali et al. (2022). “Liquid to Liquid Cooling for High Heat Density Liquid Cooled Data Centers”. In: *International Electronic Packaging Technical Conference and Exhibition*. Vol. 86557. American Society of Mechanical Engineers, V001T01A007.

IEA (2024). *Electricity 2024*. Tech. rep. INTERNATIONAL ENERGY AGENCY.

Lee, Kuei-Peng and Hsiang-Lun Chen (2013). “Analysis of energy saving potential of air-side free cooling for data centers in worldwide climate zones”. In: *Energy and Buildings* 64, pp. 103–112.

Modi, Himanshu et al. (2023). “A Transient CFD Study on Implementation of Dynamic Liquid Cooling for Series and Parallel Arrangement of Components in a Server at Rack Level”. In: *2023 39th Semiconductor Thermal Measurement, Modeling & Management Symposium (SEMI-THERM)*. IEEE, pp. 1–6.

Systèmes, Dassault (2022). *Dymola*. URL: https://www.3ds.com/products-services/catia/products/dymola/ (visited on 2023-02-10).

Todd, Austin et al. (2021). *Artificial Intelligence for Data Center Operations (AIOps)*. Tech. rep. National Renewable Energy Lab.(NREL), Golden, CO (United States).

Wetter, Michael et al. (2014). “Modelica buildings library”. In: *Journal of Building Performance Simulation* 7.4, pp. 253–270.

Wojda, R et al. (2024-10). “Dynamic modeling of power conversion stages for an exascale supercomputer”. In: *Proceedings of the IEEE Energy Conversion Congress & Exposition (ECCE)*. Under review.

Zhang, Ziting et al. (2022). “Smart DC: an AI and digital twin-based energy-saving solution for data centers”. In: *NOMS 2022-2022 IEEE/IFIP Network Operations and Management Symposium*. IEEE, pp. 1–6.

Zhu, Hongyu et al. (2023). “Future data center energy-conservation and emission-reduction technologies in the context of smart and low-carbon city construction”. In: *Sustainable Cities and Society* 89, p. 104322.