A Supercomputing Center Experience With Cooling Control Design

Michael Kercher
Computational Information Systems Lab (CISL)
National Center for Atmospheric Research (NCAR)
Cheyenne, WY U.S.A.
michaelk@ucar.edu

Gary New
Computational Information Systems Lab (CISL)
National Center for Atmospheric Research (NCAR)
Cheyenne WY U.S.A.
garyn@ucar.edu

***Abstract*—After designing and implementing an automated control system for a new HPC center, the National Center for Atmospheric Research (NCAR) elected to use a simpler operator-based solution. The solution has proven successful, and this case study documents the reasons for both the decision and the process used to choose it. Additional refinements to the cooling system controls are also documented and their adoption explained.**

Keywords— data center, controls, cooling

I. Introduction

The National Center for Atmospheric Research (NCAR) initially planned for an automated sequence to transition from free-cooling to chiller cooling, but ended up implementing required operator intervention for making this transition. This paper will describe how and why NCAR made this major change to the design of their mechanical cooling control system. The audience is primarily those involved in the initial design or upgrade of high performance center (HPC) data centers, which includes site operations and facility managers as well as engineering firms. This case provides some key considerations that all sites should consider for cooling controls design.

The NWSC (NCAR Wyoming Supercomputing Center) was constructed in 2011 after the National Science Foundation (NSF) and the Computational and Information Systems Laboratory (CISL), operated through NCAR, decided that the current HPC compute facility located in Boulder CO could no longer sufficiently support the HPC mission. Capacity Management is an essential part of the efficient planning and operation of data centers. The need for capacity management grows with the density, size, and complexity of the data center [1]. These requirements were at the forefront in the minds of the NCAR team. Planning began in 2003, groundbreaking in Cheyenne took place in June 2010,the facility was completed in December of 2011, and computing operations began in October 2012.

The Yellowstone supercomputer was installed in the NWSC as its inaugural HPC resource. Yellowstone has a peak performance of 1.504 petaflops. It debuted as the world's 13th fastest computer in the November 2012 ranking by the TOP500 organization.

II. Site Background and Design Requirements

The NWSC building is located in a remote site near Cheyenne WY at roughly 6142 feet above sea level. Its climate can be described as cold, semi-arid which makes it ideal for free cooling. The temperature ranges from -7 °C to 28 °C with humidity averaging 42.4% [2]. These favorable conditions are also seen on the Green Grid Free Cooling Map (Figure 1) [3]. These key weather attributes were some of the main reasons that Cheyenne WY was chosen as the site for the NWSC.

FIGURE 1:

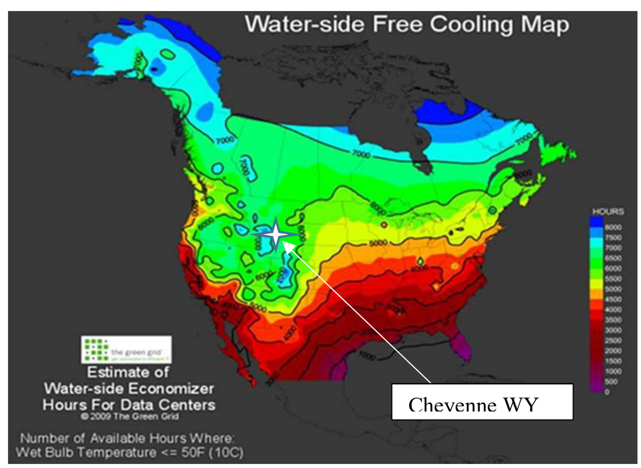

The design goal was to exceed the Green Grid projections and target less than 80 hours per year of compressor based cooling. In addition the NCAR site had to be capable of cooling both air-cooled and water-cooled computer equipment inside of the data center. The use of a water-side economizer and back-up compressor cooling is the base of the cooling plant used to supply chilled water to the NWSC Information Technology equipment.

Through the evaluation process of site selection, and as the data center characteristics were developed it was also

978-1-7281-6677-3/20/$31.00 ©2020 IEEE
DOI 10.1109/CLUSTER49012.2020.00073

determined that the NWSC facility staff would operate in a 24x7 rotating shift structure to guarantee that an operator was always available in case of a major electrical or mechanical event. These two systems are significant enough that an outage or upset could result in serious damage or downtime. The electrical system is comprised of 480v, 5,000 amp main breaker substations, along with a dual feed from the utility provider rated at 24.9kV (phase to phase) at roughly 1200 amps. The IT mechanical system is comprised of a piping system and chilled water plant that contains roughly 250,000 gallons of treated water with water flows that could reach 3600 GPM, allowing the system to handle an initial design load of 13.5 million BTUs. The impact of unmanned events in either of these systems was deemed substantial enough to warrant 24x7 coverage. Additionally, unlike many HPC centers, the mission demanded high availability with required weather analysis for the Antarctic Scientific Mission, as well as other critical regional customers. Check-point / restart was not an option.

III. About the Case

A. System Design and Development

Considering the climate mission of NCAR it was an obvious goal to maximize the efficiency of the physical plant supporting the HPC operations. As a leader in atmospheric sciences, NCAR prioritized building a data center with a low target PUE. This process was extremely collaborative where multiple HPC data centers around the world were visited and input was gathered from multiple operation teams on lessons learned and what worked correctly in their center, in their specific climate. In addition NCAR enjoyed collaborative review and input from colleagues through the Energy Efficient High Performance Computing Working Group (EE HPC WG) [4]. NCAR used a novel engineering approach as well with multiple data center specific design teams performed some of the higher level conceptual designs and then more local / regional design firms adapted these concepts to the specific climate of Cheyenne and the specific mission of NCAR. Through this collaborative process the plant operational metrics were created for the NWSC (see Figure 2)

FIGURE 2:

NWSC Plant Operating Metrics

Description	Set point
HPC Equipment / Computer Room Chilled Water Supply Temp	18 DegC (65 DegF)
Condensing Water Supply Temp	Floating based on 18 DegC (65 DegF) Chilled Water Temp due to Water side Economizer, Average of 17 DegC (63 DegF)
HPC Equipment / Computer Room Chilled Water Flow	1400 – 3000 GPM
Condensing Water Pump Flow	1500 – 3600 GPM
De-Humidification Chilled Water Loop	7 Deg C (45 Deg F)
Computer Room Humidity	40%
Computer Room Air Temperature	22 DegC (72 DegF)

Ultimately the team at NCAR tried to keep the energy component of the total cost of ownership (TCO) in mind. It was an objective to encourage innovation with vendors where they are incentivized to reduce the total costs for energy and/or power-related expenditures while balancing the operational costs for energy.

In order to ensure the highest level of efficiency, building automation was known to be critical and a controls engineering specialist was added to the design team to help in this area. The amount of automation is always an interesting design discussion. There are still data centers in operation with the simplest control of a return air thermostat modulating a combined fan and chilled water valve in a CRAH unit. The other end of the range would be a full Artificial Intelligence (AI) based control system being explored by Google for their advanced data centers. [5]. While these may be the direction of the future, during NCARs design, AI was not a viable consideration

There are a number of trade-offs in this area of controls strategy and energy efficiency is only one of them.

First cost – Obviously the instrumentation for extensive monitoring and controls represent a significant capital cost. Will the energy efficiency gained be able pay for this hardware through a TCO process?

Operational Cost – Energy efficiency is a key part of this, but headcount needs to be considered as well. Does more manual control add the need for head count and additional operating staff? Do all aspects of a controls system need to be automatic to get the best energy efficiency?

Operational Risk – Generally, we expect an automated control system to be better at minimizing risk. Certainly a control loop isn't going to accidentally turn a valve the wrong way, but can an automated control loop be aware of enough of the operational parameters and specific circumstance to properly change a significant portion of a system to an alternate operational mode?

Reconfiguration – Control loops that are regularly used will be retuned as systems are changed or modified largely thru the efforts of a competent hands-on operational staff. But what happens when the control sequence is infrequently used. How much of the loop dynamics may have changed and will it still be able to make the required changes properly? What about system expansions and modifications on the IT side; what must be accomplished to ensure the control loop will be able to function as needed based on the new loads?

Our goal was a control sequence where the most efficient equipment available at that moment could function together in the most efficient way possible through control design methodologies. However, this goal was formulated during the conceptual design process and final equipment selections were yet to be determined so some parts of the sequence were more generalized to encompass the use of many different types of equipment that could possibly be used.

Through the design process it was determined that a back-up chiller was needed to be installed to maintain the required 18 °C Chilled Water Supply Temperature (CHWS-T) set point. The design team determined that during the summer months there are enough afternoon thunderstorms in the Cheyenne area that raise the outdoor air enthalpy to a point that the cooling towers would not be able to remove enough of the heat and the 18 °C chilled water could be unattainable through the use of water-side economizers alone [7]. The transition between the chiller and free-cooling modes needed to be seamless, especially due to the fact that water cooled HPC systems can overheat quickly if the mechanical systems have issues. Initially it was anticipated that this would be an automated function.

As the design and equipment list were being developed, a parallel piped condensing water system was designed to feed the two 750 Ton 18 °C CHWS-T free cooling HXs and the single 1000 Ton 18 °C CHWS-T centrifugal chiller that are responsible for cooling the NWSC computer equipment (see Figure 3). It is interesting to note that the centrifugal chiller was not redundant due to the low demand (80 hours out of 8760 annual hours).

There are also two 90T 7 °C CHWS-T scroll chillers used for comfort cooling for the facility, as well as de-humidification when necessary. These are not associated with cooling the NWSC IT equipment

The large 1000T °C centrifugal chiller was designed to utilize the large 135,000 gallon 18 °C chilled water tank during the transition from HX to Chiller.

FIGURE 3:

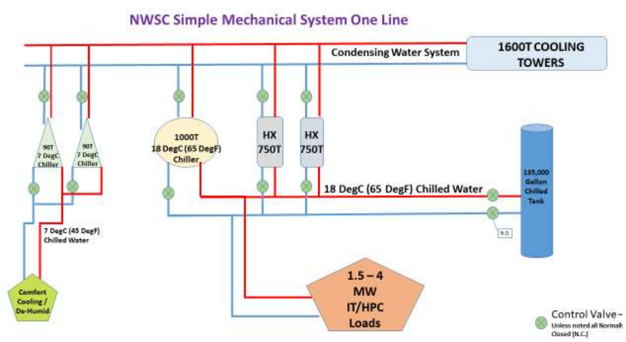

As the NWSC facility was being constructed the design/construction team along with the NWSC operation staff were trying to find the least complicated and lowest risk way to switch our cooling plant from flat plate HX (free cooling) to compressor (Chiller) cooling.

The original design and control system called for an automatic switch between these two sources of cooling that involved a complex sequence of operations. The original sequence looked at data such as; current load, outdoor wet bulb temperature, and the temperatures from our Chilled Water, Condensing Water, and the temperature of our large Chilled Water Storage Tank.

Without fully knowing what kind of chiller the NWSC would have installed through original construction, the design aspects needed to cover the worst-case scenario. This worst-case scenario involved the need for raising the condensing water temperature from 17 °C to 21+ °C in order to allow the chiller to properly start-up. This worse case meant that we would have to decouple our cooling from the condensing water system and use our reserve Chilled Water Tank to supply the 18 °C water for 10-20 minutes while the condensing water temp was raised. Once it was raised to the determined set point, the chilled water system would transfer to the chiller and that would continue cooling our load. In order to accomplish this task on a chilled water system that moves 1500 – 3600 GPM through 24” piping, a large three way valve was installed to be able to shift the flow of water through the Chilled Water Tank. The above sequence was a concern for the design team being able to properly establish the control loops to complete such a transition.

B. System Commissioning

During commissioning of the system, a major incident occurred within the chilled water piping system. A control loop driven change in the three-way valve position that allowed flow through the reserve chilled water tank caused a significant event that could have been either a cavitation or water hammer. This event was felt throughout the entire facility and staff determined that some parts of the piping system physically displaced as much as six inches. The root cause of the problem was not determined specifically. Hindsight tells us we should have, but the demands of the schedule caused it to become a second tier problem. In any case, the team was now aware and concerned that there could be issues with the transition between the two cooling modes that could possibly cause damage. Manual switchovers between the modes had been accomplished with no issues resulting.

With the previous decision to operate in a 24x7 fashion, it gave the NWSC operations staff the ability to make this switchover a manual process vs. an automatic one. We have three different steps of status notifications and alarms; 1) switch to chiller possible 2) switch to chiller likely and then 3) switch to chiller required (see Figure 4). So throughout the different states of outdoor temperature and humidity and indoor temps and conditions we will get those three levels of state warnings / alarms. That helps cue the operator to pay attention to what is going on and to check on conditions for making a decision as to whether or not to switch to compressor assisted cooling. The control programmers had done the required work figuring out all the psychometrics that this state variable (possible, likely, required) became the primary signal to the operators to ensure proper transitioning of the system between free and compressor cooling.

FIGURE 4:

Table for Switch to Chiller

State Warning / Alarm Name	Trigger	Decision	Control Point Operation	Percentage of Year System Runs in this Mode
Switch to Chiller Warning Possible	All four cooling tower cell fans reach 100%, both HX valves are open, Supply / Return Delta is at least 5 degF	Operator verifies system running properly, and checks weather forecast	None unless system issues exist	3%-5% (262 – 438 hours per year)
Switch to Chiller Warning Imminent	All states satisfied above and the 18 DegC (65 DegF) Chilled Water Temp Cannot reach set point (within 1.3 DegC (2.5 DegF) for 10 minutes	Operator checks system operation for anomalies, if no anomalies are present switch to chiller engaged	Control point “mechanical / free” cooling selector is toggled to mechanical	1% - 2% (87 – 175 hours per year)
Switch to Chiller Required	All states satisfied above and the 18 DegC (65 DegF) Chilled water system reaches 20.3 DegC (68.5 DegF) for 5 minutes	Operator immediately switches to Chiller Operation	Control point “mechanical / free” cooling selector is toggled to mechanical	1% (87 hours per year)

C. Operational Learnings

The NWSC was designed and built to be one of the most efficient data centers in the world. Creating this kind of efficiency can also create complexity, but it also creates some opportunities. This is especially true in an environment/industry like HPC that has a higher priority on data collection and reporting compared to other more traditional industries. These data requirements are due to the cutting edge nature of the technology deployed and the high density racks where new cooling technologies are being used every few years. The NWSC was designed to be an extremely flexible and adaptable facility to this ever-changing environment. NCAR decided that in order to implement the adaptions needed over time data needed to be collected and evaluated.

The complexity of the control sequence was drastically reduced when it was decided to make the switch from free cooling to chiller a manual process done by the operator. This manual process only involves one control point (free or mechanical cooling). The operator can simply choose this point between the two states. Once this manual selection is completed, the control system takes over, performing the needed valve operations and altered pump control from GPM in free cooling mode to DP control in Chiller mode for the condensing water system. The chilled water pump sequence did not need to be altered between the two modes.

To date this manual process has worked well for the NWSC. The system currently only needs to operate the chiller about 80 hours per year, and we are able to stay in free-cooling mode longer because it is a manual process and not automated. This approach is very similar to recommended practices that relate to job scheduling on an HPC system where an operator will manually queue specific jobs based on electrical grid capabilities, demands, and charge rates. [6]

The energy efficiency of NCAR has not suffered from making this transition a manual operation. During normal operation with free cooling the PUE average is 1.2. Obviously the 80 hours of compressor based cooling is NWSCs lowest efficiency but even in this time the PUE is 1.3. NCAR does not believe that automating this switchover would measurably change the annual PUE in either direction, particularly since this is generally 80 out of 8760 hours of the year.

D. HPC system expansion

After approximately six years of production with a HPC system that utilized rear door heat exchangers the NWSC plant was operating very well. In 2017 the second HPC system was deployed that utilized a CDU cooling sub loop and cooled the HPC equipment directly at the chip level via liquid. This computer is named Cheyenne and provides a significant computational boost past Yellowstone to 5.34 petaflops. Cheyenne is roughly the same power draw in one third the floor space. With an average load of 1.3 MW over 26 racks vs. the Yellowstone average load of 1.2 MW over 79 racks. This new system and its inherent change in HPC cooling type caused a much more varied load in chilled water cooling demand over a shorter amount of time when flow rates were lowered, or if the 18 °C chilled water supply temperature varied, or when the HPC machine workload changed. This change in the system exposed some internal hardware issues relating to the HPC cooling system and the HPC equipment ran warmer than design and some temperatures were close to upper thresholds.

The new system installation and operations, when combined with Yellowstone's loads were less than ideal. Through the use of constant collaboration and a view that all systems, especially mechanical control systems can be improved and simplified over time the NWSC formed strong relationships in support of the needed improvements. The first example of this type of collaboration was joining various working groups - specifically the EE HPC WG. The EE HPC WG has membership relating to the running of complex HPC facilities which included facility staff, IT staff, vendors, and engineers. The second collaborative approach was with control vendors, including the much more frequent onsite training and troubleshooting with the onsite staff in the use and programming of the NWSC's complex Building Automation System.

Working with the EE HPC WG allowed the NWSC to be able to share experiences that were gained through the operation of the facility. One of the discussions circled around the issues the NWSC had regarding the first half of this case study and how we were able to simplify the design and operation of the system. Through these discussions it was identified that there are other ways to transition a chilled water plant between free and mechanical cooling operations. Some of these ideas included running all of the systems in parallel with each other and then slowly closing the HX valves when transitioning to Chiller, or slowly shutting the Chiller valve when transitioning to HX.

This idea of running things in parallel was very intriguing to the NWSC team, and it was brought forth as a solution and was discussed with the control contractor used onsite. After some simulation testing and sequence identification it was decided that the Condensing Water Pumps were sized large enough to provide the requisite total flow. There was also some sequence changing associated with the valves and the flow control to the computer was also investigated. The concept of flexibility and the ability to reconfigure the control system was paramount to our success. The automated control loops (previously abandoned for the manual change over) originally in place, might have been salvageable had the site not expanded, but with the new parallel system operations during transition, that design work was no longer of any use.

After all components were vetted and a new sequence was developed to run all systems in parallel during the transition a shutdown of the NWSC chilled water plant was scheduled to implement the changes. During the scheduled outage the system's programming was updated and the HPC administration group ran the Linpack code on the system to give it a large false load for testing. Through the original sequence of transferring from HX to Chiller the system temperature would vary from the 18 °C set point to roughly 21 °C during a 600 second transition period. After the sequence changes the transition between HX and Chiller now only takes roughly 90 seconds, and the Chilled Water temperature spikes to 19.1 °C. These changes have hardened the NWSC mechanical system even more, and it has made things simpler for the operators. These changes have also made the NWSC HPC systems run more efficiently as there are no more major temperature swings caused by the transition and the nodes, switches, and other liquid cooled equipment within the system stay cooler. The equipment staying cooler causes the Cooling Distribution Units (CDUs) associated with the HPC system not have to work as hard and less energy is needed by these devices to ensure the liquid cooled equipment within the HPC system stays within operating parameters.

IV. CONCLUSIONS, LIMITATIONS AND FUTURE INVESTIGATION

This problem could relate to other sites as supply side chilled water temperature set points rise and free cooling becomes an option for areas that have higher humidity than what is normally thought of as ideal for free-cooling practices.

When looking at designing the control system for a complex HPC facility many different factors play into deciding what level of automation should be implemented into the system. Some of these factors include mission, budget, available staffing, location, and available utility supply. If there are approaches to the datacenter design where multiple cooling resources are deployed and used understanding how often these modes will be needed is critical in understanding and deciding how to transition.

As discussed in the introduction these many aspects of the decision process for the controls loop all came into play. To summarize:

NWSC originally planned on a fully automated system but found a primarily automated with infrequent manual mode changes to be the best set up for the site.

First cost – NWSC may have been able to save some design and programming costs had they had a crystal ball but there would not have been significant other capital savings as the control elements and the instrumentation were still required and being used to track NWSCs efficiency and low PUEs

Operational Cost –With NCAR's commitment to staff the facility 24x7 from the beginning options were available to make things more manual when switching between free and mechanical cooling based on the limited number of times this process would be needed.

Operational Risk – From the initial physical system bump during commissioning the risk to a system upset was real and considered. NWSC's change to manual transitions is felt to have provided a more stable system. In hindsight we also feel that system growth and IT load changes have made the original automatic transition obsolete. The amount of tuning and testing to ensure proper sequencing would have taken more time and presented more risk than the annual summer transitions for compressor based cooling.

Reconfiguration – NCARs IT loads have doubled since the initial start-up and the site continues to go through changes and upgrades across the full IT hardware kit. Simplifying operations and including a few manual transitions has supported these reconfigurations well.

Also understanding the best time to bring in experts and what kind of solutions are working in the industry is critical when making these types of decisions. Having the assistance of competent engineering and design teams, as well as belonging to long term working groups can help make these types of complex systems more adaptable and robust.

A good understanding of the overall procurement process associated with building a large scale facility and anticipating equipment that may not even be designed yet can help expose some of the limitations to this process. Facility construction takes a long time to get approvals, create designs, and then implement the plans designed. This process can sometimes force teams to make decisions out of the preferred order and limitations become exposed. Good risk management and collaboration can help mitigate these limits.

The NWSC plans to continue this research and to continue to make the systems better over time. Continuing strong partnerships with vendors and making sure that staff are well trained and feel ownership in the processes help make sure change and adaption is embraced. Another large part of making sure the NWSC continues to be successful is to continue to be involved in working groups like the EE HPC. This collaboration has been a huge help in understanding the industries best practices as it relates to control systems and HPC cooling. The extensive cross industry collaboration means that if a member of the team experience new problems first they can share with the others so mistakes are not made again and again and the HPC world as a whole is made better.

REFERENCES:

[1] Rasmussen, N. "Power and Cooling Capacity Management for Data Centers". Schneider Electric – Data Center Science Center White Paper #150 Revision 3, 2012

[2] .A., "Cheyenne WY Weather Data". Collected from website: https://weatherspark.com/y/3765/Average-Weather-in-Cheyenne-Wyoming-United-States-Year-Round, retrieved on May 30th 2020.

[3] S. Strutt, C. Kelley, H. Singh, V. Smith,. "Data Center Efficiency and IT Equipment Reliability at Wider Operating Temperature and Humidity Ranges". The Grjk een Grid White Paper #50, N.D. (no date Found)

[4] Energy Efficient High Performance Computing Workiing Group (EE HPC WG) Infrastructure Sub-Group Controls Teams. https://eehpcwg.llnl.gov/infra_ctrls.html

[5] J. Gao "Machine Learning Applications for Data Center Optimization". Collected from website: https://static.googleusercontent.com/media/research.google.com/en//pubs/archive/42542.pdf

[6] A. Clausen, G. Koenig, S. Klingert, G. Ghatikar, P. Schwartz, N. Bates. An Analysis of Contracts and Relationships between Supercomputing Centers and Electricity Service Providers. Proceedings of the Workshop on Energy Efficient State of the Practice Workshop as part of the 48th Annual International Conference on Parallel Processing. Kyoto, Japan. 2019

[7] A. Anderson. NCAR-Wyoming Supercomputing Center (NWSC) – New Industrial Facility. ASHRAE Technology Awards. August 29th 2014