David Sellers

Control System Design

It Really Is About Lags

BY DAVID SELLERS, P.E., MEMBER ASHRAE

David St. Clair closes Chapter 5 of *Controller Tuning and Control Loop Performance, A Primer* with a huge text box that states, "it all depends on the lags."[1] I did not fully appreciate what it meant until I blew up a duct. Then, it was obvious.* Lags are a major part of the dynamics that fascinate and challenge us in our building system endeavors. They continuously revise load profiles and perform real-time testing of control processes. A theoretically perfect control system using the best hardware available can run amuck if you don't address the lags, a concept we will explore via this case study.

Overview

Our focus is a large condenser water system serving the central plant for a research cleanroom and office tower at a San Francisco Bay area university (*Figure 1*). Note the rooftop cooling tower location relative to the chillers and the length of large pipe in-between. The dynamics discussed here are intimately related to the physical and schematic system arrangement.

Loads

Heat rejection is provided for:

- A 60 ton (211 kW) modular chiller that serves a mission-critical load.
- A 600 ton (2110 kW) absorption chiller, operated when the campus cogeneration system benefits from having a heat sink for waste heat.
- A 600 ton (2110 kW) centrifugal chiller, operated when the cogeneration system does not require a heat sink.

Given a peak facility load of 600 tons (2110 kW), the two large machines back each other up. Thus, the concurrent operation of all chillers represents a transient condition, occurring only when transitioning to or from the absorption chiller. While the system can see a peak flow rate of 3,640 gpm (230 L/s), most of the time the targeted flow rate will be 2,520 gpm (159 L/s) (absorption plus modular) or 1,320 gpm (83 L/s) (centrifugal plus modular). Occasionally, it can be as low as 200 gpm (13 L/s).†

Heat Rejection Equipment

Identical induced draft, cross flow cooling tower cells are provided. The cells are piped in parallel and operated together under all conditions (no isolation valves).‡

Temperature Control Mechanisms

Temperature control mechanisms include:

- A bypass valve allowing water discharged from the loads to be recirculated to maintain a minimum entering water temperature;
- Partial to full system flow directed over both tower cells, providing modest heat rejection without fan operation via natural draft; and
- Variable speed fans (one per cell), which are staged on and ramped up to achieve design heat rejection.

The bypass valve and variable speed fans represent the potential for simultaneous heating and cooling if the

---

*I talk more about this here: https://tinyurl.com/LagExample.

David Sellers, P.E., is senior engineer at Facility Dynamics Engineering's office in Portland, Ore.

control process managing them is not properly structured. The goal of the bypass valve is to warm up the supply water temperature to the loads, while the goal of the fan operation is to cool the water down. The fans should never be active if the bypass valve is open.

Ultimately, the goal is to deliver condenser water at a temperature appropriate for the online machinery (Point 4, *Figure 1*). That temperature is regulated by sequencing the fans and bypass valve. The setpoint is a function of the requirements of the online chiller(s).

Temperature Dynamics

Measuring Temperature

We encounter our first lag in the mass of the temperature sensor and well used as an input to the process. The mass creates a lag between the time the temperature of the water surrounding the well changes and the time the control process "sees" it.[§] The magnitude can be 10 to 20 seconds or more, depending on the temperature difference driving the energy transfer and the mass involved.

FIGURE 1 System diagram.

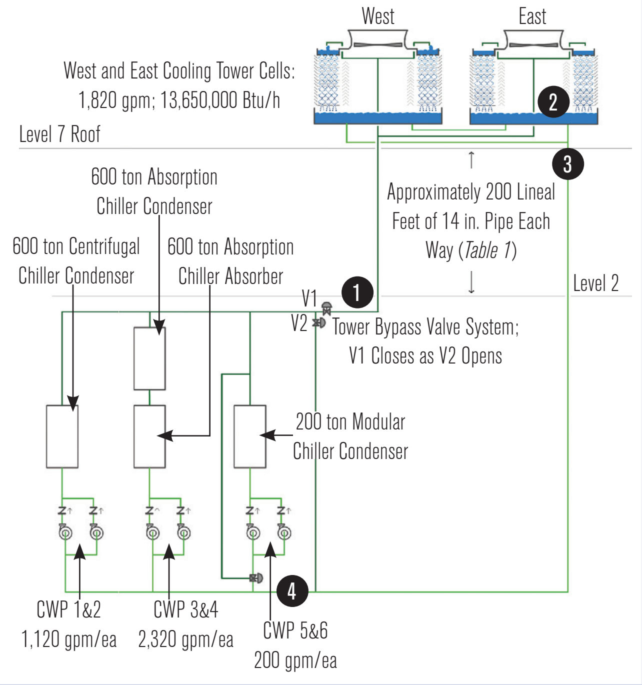

Absorption Chiller Requirements

The efficiency of the absorption cycle (a chemical reaction) is not meaningfully impacted by the condensing temperature. What is impacted is the machine's ability to operate. Absorption chillers typically require condenser water no cooler than 80°F to 85°F (27°C to 29°C); otherwise, the lithium bromide solution driving the cycle solidifies, causing the machine to cease operation.

Centrifugal Chiller Requirements

In contrast, vapor compression chiller efficiency will improve at lower condensing temperatures. While the exact relationships are machine specific, condenser water reset is a common efficiency strategy. There are multiple ways to do this.[2]

Centrifugal Chiller Limitations

Practical limits exist for entering condenser water temperature. For most machines, the lubrication system is inside the refrigeration system, and refrigeration pressures impact lubrication pressures. If the condenser water is too cold, nuisance oil pressure safety trips can occur. Condenser water temperature also impacts the "lift" (pressure difference) across the compressor. For centrifugal machines, a combination of refrigerant flow rate, inlet guide vane position, impeller speed and lift can allow flow to reverse through the impeller, which is called surge. Surge typically manifests as vibration when the flow through the compressor stalls and becomes unstable. This imposes significant stress on impeller bearings, blades and seals, as well as the motor, and should be avoided.[3]

Bypass Valve to the Rescue

Properly implemented, the bypass valve provides a remedy for these challenges.[#] For the absorber, it can warm up condenser water when transitioning from the centrifugal machine. For the centrifugal chiller, it provides

---

[†] The variation in targeted flow rate sets up an interesting pump interaction when the system transitions between modes. With no variable frequency drives (VFDs) on the pumps, the head variation in the common piping associated with the flow changes causes the targeted flows to not be achieved. This impacts the lags we are discussing, but they would exist even without a flow management issue.

[‡] The flow variation described previously,* combined with the lack of tower isolation valves, leads to issues with adequately wetting the tower fill under some conditions. You will find a "sneak peak" of this future column subject at https://tinyurl.com/CTFlowVariation.

[§] The videos here illustrate this phenomenon: https://tinyurl.com/LagExampe.

[#] "Properly implemented" is crucial. Valve must be properly sized under all flow conditions.

a way to bring the chiller online if it is starting with tower basins full of cold water due to their tendency to approach the ambient wet-bulb temperature when idle. For both machines, it allows supply temperature to be tempered under low ambient wet-bulb conditions when the natural draft of the tower provides too much cooling.

Setpoint Dynamics

The setpoint itself is frequently in a state of flux to optimize the plant. This moving target becomes another variable impacting control process stability, introducing lags via the resetting sensing mechanism and the network dynamics transmitting the resetting information to the primary control process. The absorption chiller adds a "wrinkle" to the plant in our example. With its mix of technologies, this plant requires a different setpoint strategy that varies with the mix of chillers online.

Efficiency Dynamics

Condenser water reset strategies are a trade-off between improving compressor kW per ton at the cost of increased tower fan energy. On a day with design heat rejection and wet-bulb temperature, the cooling tower fans will operate at full speed to deliver the design water supply temperature. If the wet-bulb temperature drops, the towers could deliver colder water if the fan speed remained constant (*Figure 2*).

Reset strategies leverage this characteristic, spending fan energy to save compressor energy on off-design days, which is not a "one-size-fits all" affair. For example, in a recent life-cycle cost analysis I worked through, a tower selected to reject 9,365,000 Btu/h (2745 kW) and deliver 85°F (29°C) water with a 74.2°F (23.4°C) wet-bulb temperature, targeting lowest first cost, had a 40 hp (30 kW) fan motor. A tower selected for the same conditions but targeting best life-cycle cost had a 15 hp (11 kW) fan motor.

The cost premium had a 6.9 year simple payback, which was acceptable to the owner given the 15-plus year anticipated remaining plant life.[II] But because of the difference in fan energy spent to save compressor energy, the reset schedule for the plant with the lowest first-cost tower will be different from the one used for the best life-cycle cost tower.

FIGURE 2 Typical cooling tower performance (blue dot = design).

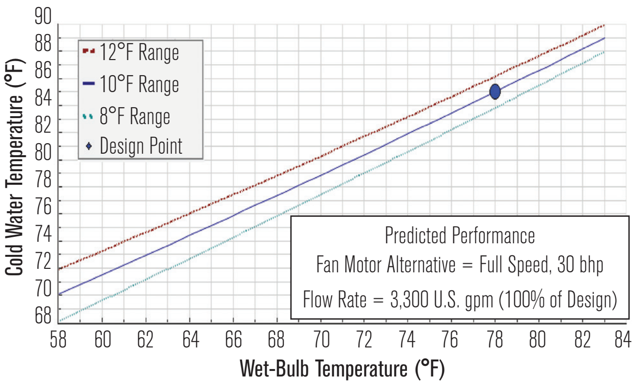

These Performance Curves are Based on Constant Fan Power

Wait, There's More!

The chiller kW/ton profile for the life-cycle cost analysis varied with condenser water temperature, was nonlinear and was machine-dependent (similar, but not identical chillers [*Figure 3*]).

The amount of fan energy you could afford to spend to save chiller energy varied with the tower selection, the chiller load,the combination of machines online and the ambient wet-bulb temperature. A reset schedule that was crafted without taking all the variables into consideration might not deliver as intended.

For math-phobics like me, that's really intimidating. Fortunately, the building systems know the answer; you just need to ask, which was the plan for this project. We provided the structure for a reset schedule in the control logic and then used functional testing and trend data to establish the optimal setpoints after verifying the basic functionality of the new tower cells.

Control Strategies–Option 1

Returning to the case study, a very common approach to controlling a tower bypass valve is to sequence it with the fans based on the temperature supplied to the plant. This approach was specified for the plant discussed here. This seems extremely logical: you need to maintain the desired temperature at Point 4 (*Figure 1*), and you have three mechanisms to accomplish it. Sequencing the bypass valves and fans to maintain setpoint seems to make sense.

[II] In this case, the space and structure required for the physically larger tower associated with best life-cycle cost were not a concern. That is not always the case.

FIGURE 3 Chiller kW per ton profiles for the two similar but not identical chillers served by the tower in the life-cycle cost analysis example.

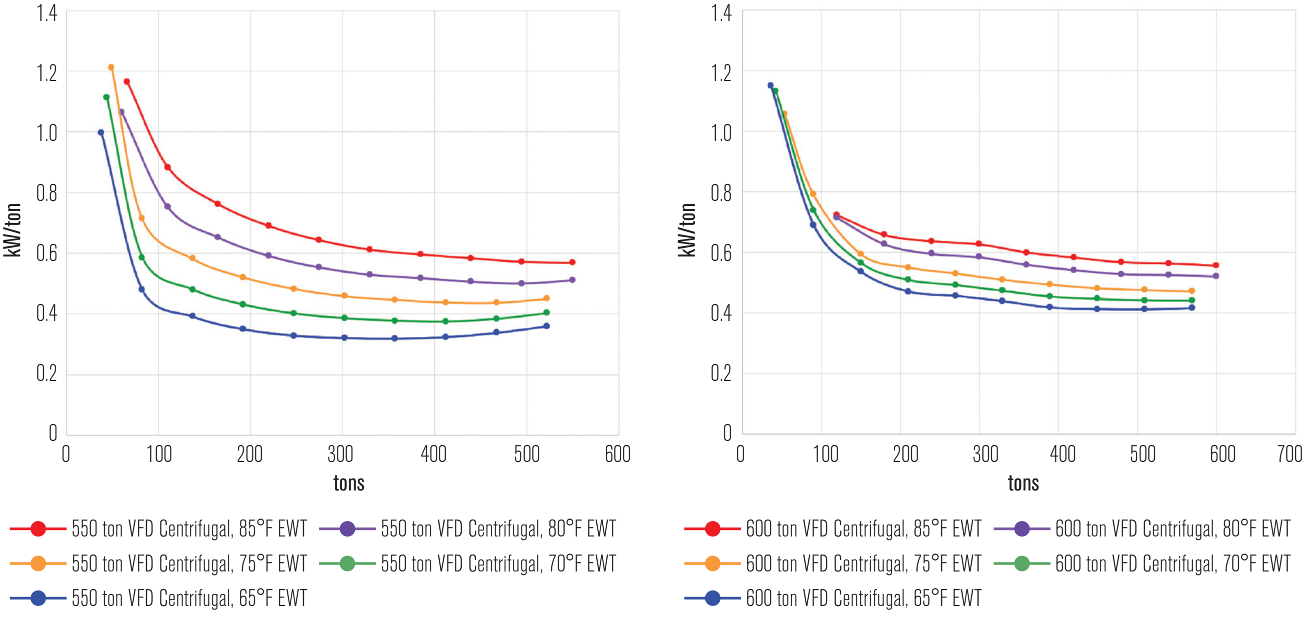

Enter Transportation Lags

The variables discussed to this point may seem obvious. But the transportation lag may not "jump out" until you consider the large volume of water in the piping and tower basins. That is how it was for me the first time I recognized this phenomenon while commissioning one of my first central plant control system designs.

TABLE 1 Transportation times for the different operating conditions associated with the system illustrated in *Figure 1.*

CONDITION	FLOW RATE (GPM)	TRANSPORTATION TIME, MINUTES			
		BYPASS VALVE DISCHARGE TO TOWER TO REMOTE SENSOR (FIGURE 1, POINTS 1 TO 4, 5,282 GALLONS)	BYPASS VALVE DISCHARGE TO TOWER TO COLD BASIN (FIGURE 1, POINTS 1 TO 2, 2,109 GALLONS	TOWER BASIN TO REMOTE SENSOR (FIGURE 1, POINTS 2 TO 4, 3,173 GALLONS)	TOWER BASIN TO LOCAL SENSOR (FIGURE 1, POINTS 2 TO 3, 1,797 GALLONS)
Absorption Chiller Only	2,320	2.28	0.91	1.37	0.77
Centrifugal Chiller Only	1,120	4.72	1.88	2.83	1.60
Modular Chiller Only	200	26.41	10.55	15.86	8.98
Absoprtion Chiller + Modular Chiller	2,520	2.10	0.84	1.26	0.71
Centrifugal Chiller + Modular Chiller	1,320	4.00	1.60	2.40	1.36
All Chillers	3,640	1.45	0.58	0.87	0.49

Table 1 illustrates the approximate transportation times for a parcel of water moving between the indicated points in our system, which varies with operating mode. As a result, a tuning solution that works well for one mode might not work for others.

Focusing our discussion on the most common operating mode (centrifugal chiller plus modular chiller), we discover that if the tower fan were to cycle off, it might take 2.4 minutes for a parcel of water that had just fallen off the fill into the cold basin (Point 2 in *Figure 1*) to reach the location of the temperature sensor controlling the process using the Option 1 strategy (Point 4 in *Figure 1*).

That means the cooling tower fan cycled off based on a condition that existed more than 2.4 minutes ago—more because those numbers do not take into account the lag introduced by the mass of the sensor and well.**

Let's assume the current setpoint is steady at 65°F (18°C) and that the system is steady state at a load condition, generating a 4°F (2.2°C) rise across the condenser (69°F [21°C] return temperature). Because the system is steady state, the fan speed will be in balance at a point that causes the 69°F (21°C) water flowing into the hot basins to drop to 65°F (18°C) by the time it falls off the fill into the cold basin. Thus, the cold basin, and the piping

---

**This discussion is also conservative because *Table 1* focuses on the common piping and does not include the volume contained in the connections between the common pipe and the chiller and the condenser barrel.

to the chiller, are full of 65°F (18°C) water.

Now, imagine the load drops. Since the entering temperature to the condenser is a steady 65°F (18°C), the drop in load will show up as a reduction in leaving water temperature. We'll assume the leaving water temperature drops and stabilizes at 68°F (20°C). This would happen gradually in the real world. To make things easier to visualize, we'll assume it shows up as a step change, i.e., if you were watching the temperature with a magical, massless thermometer, you would see the water leaving the chiller drop 1°F (0.6°C) in a fraction of a second.

When the load reduction occurs, the pipe between the chiller and the tower is full of 69°F (21°C) water, and it will take about 1.6 minutes for the first parcel of cooler, 68°F (20°C) water to reach the tower and fall through the fill to the cold basin. Since the sensor controlling the process does not know the load changed (because at the instant it happened, the cold basin and piping to the sensor location were full of 65°F [18°C] water), it has continued to operate the fans at a speed that generates a 4°F (2°C) temperature drop. As a result, the parcel of water that entered the hot basins at 68°F (20°C) is cooled to 64°F (18°C).

It will be another 2.4 minutes before the parcel of cooled water reaches the temperature sensor controlling the process. System reaction will be further delayed by the thermal mass impact discussed previously.

By the time the control system realizes the load has dropped, several minutes will have passed, and the process will have over cooled and undershot the targeted setpoint. Recovery will take a while because all that time the cooling tower was filling the piping leaving the cold basin with colder than desired water.

The transportation lag will make the departure below setpoint appear to persist, causing the proportional plus integral (PI) control process to compensate by shutting down the fans and opening the bypass valve. For the system under discussion, that happened frequently.

Because the bypass valve was oversized, the reaction produced was out of proportion to the need, triggering further instability, which rippled out to the system's chilled water side, where load changes produced would feed back into the condenser water system (*Figure 4*).

FIGURE 4 Cascading instability.

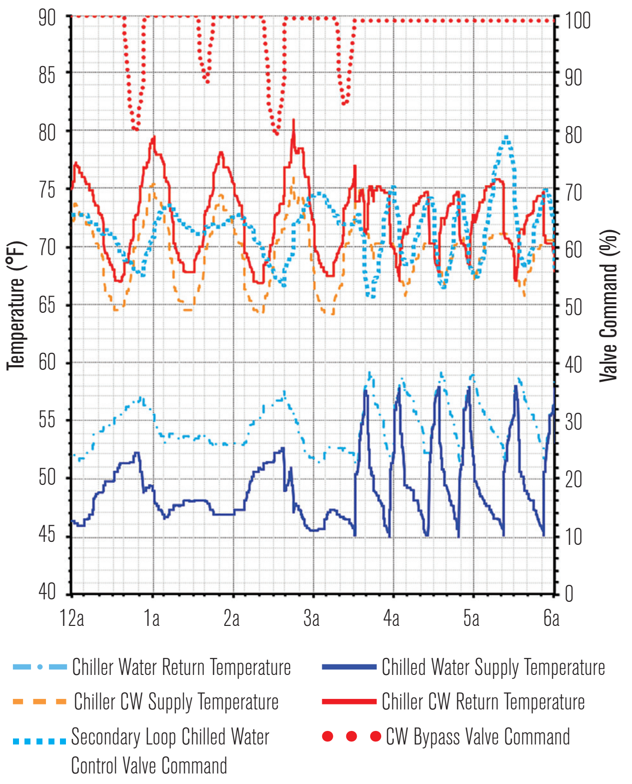

Control Strategy–Option 2

Having recognized transportation lag, a solution becomes obvious: control the fans based on their cold basin leaving water temperature. This cuts the transportation lag in half for a signification portion of the system (*Table 1*).

But, like most engineering solutions, solving one problem leads to others.

- Two control loops = two setpoints. The setpoints need to be coordinated so if one is changed, so is the other.
- With two temperature sensors involved, the "out of the box" accuracy comes into play.

Even the best sensors have a manufacturing tolerance; two sensors with identical accuracy specifications may not read the same value when subjected to the same condition. For applications like our case study, the relative accuracy of one sensor compared to the other becomes critical. Ensuring relative calibration is performed along with training the operating team in that regard will be an important part of the commissioning process for this option.

Making it Happen

Conceptually, the logic required by the two options is illustrated in *Figure 5*. For those interested in the details, pneumatic and DDC logic diagrams for both options are at https://tinyurl.com/ASHRAECWDetails, along with a more detailed version of the system diagram.

FIGURE 5 One vs. two loop functional block diagrams. Note how Option 1 has one loop, setpoint and input, while Option 2 has two loops with different inputs but a common setpoint.

A More Informed Solution

If you study *Table 1* in the context of Option 2, you will recognize a significant lag in the process still exists due to the volume of water in the tower cold basins. Using a strategy based on the basin temperature seems desirable.

Experience indicates this is more complicated than it seems. One lesson learned in my experience is that water falling through a cooling tower is like a gentle rain; it stratifies. More specifically, if a tower fan is off and then starts, there will be an immediate impact on the temperature of the water falling off the fill relative to the temperature of the water entering the fill. But, because this water is falling gently into the cold basin, initially the temperature of the water in the top "layer" will be different from that in the bottom "layer."

On that project, tower cells installed over a 4 ft (1.2 m) deep basin taught me this. When I probed its depths by attaching a weight to my temperature sensor, I discovered stratification.

As a result, I recommend that

- Each cell be controlled based on its cold basin temperature;
- Measuring the cold basin temperature using a rigid averaging probe with length matched to the basin's minimum operating level; and
- If possible, finalize sensor location based on field testing.

The logic required is like the Option 2 logic, but with a loop for each cell. All contingencies mentioned for Option 2 apply.

Conclusion

There are a lot of dynamics to consider when developing logic to control a large condenser water system. That is part of the fun and fascination associated with control system design and facility =operations. But as David St. Clair and Mother Nature point out, success ultimately involves recognizing and addressing the lags.

References

1. St. Clair, D. 1993. *Controller Tuning and Control Loop Performance, a Primer, 2nd Edition*, P. 56. Self-Published. http://www.straightlinecontrol.com.

2. Taylor, S. 2012. "Optimizing design & control of chilled water plants, part 5." *ASHRAE Journal* (6):56.

3. Peterson, K. 2018. "Avoiding centrifugal chiller surge." *ASHRAE Journal* (11). ■