providing insights for today's hvac system designer

Engineers Newsletter

volume 47–3

impacts of

Chilled-Water System Design Decisions

This *Engineers Newsletter* walks through a number of design decisions, with discussion and examples to explain how and why those decisions are made.

While designing a chilled-water system, a myriad of decisions must be made. Experienced engineers often make these decisions "automatically" as they have in the past, based on what they have learned from experience. Today, all engineers likely use an internet search engine to get direction, but when there are differing—or even conflicting—recommendations a decision must be made.

Common decisions regarding chilled-water system designs include:

- bypass line sizing in variable flow systems
- dynamically varying condenser water flow
- number of chilled-water pumps to operate
- series chillers and power consumption
- whether to use pressure-independent control valves

Bypass line sizing

This seemingly simple decision can have significant consequences if not done correctly.

- *In a **primary-secondary system**, the bypass pipe should be the same diameter as the pipe going into the largest chiller. Its length should be about 8-10 pipe diameters long or have an equivalent pressure drop.*
- *In a **variable primary flow (VPF)** system, the bypass line should be sized for the largest **minimum** flow rate and it will have a control valve.*

The reasons behind this guidance follow.

Primary-secondary system bypass sizing. The premise of a primary-secondary system is to hydraulically separate (decouple) the primary (chiller) flow from the secondary (system) flow. This decoupling prevents the operating pressures of chiller pumps from impacting the operating pressures of the system pumps. This is accomplished by installing a bypass line with a small pressure drop. The bypass allows water to flow in either direction and is used to indicate chiller and primary pump sequencing in this system arrangement (Figure 1).

Chillers and constant flow primary pumps are enabled in pairs, making the primary flow rate a step function. As system load and flow increase, the excess flow rate (from supply to return) in the bypass line decreases. At the point that system flow rate exceeds chiller flow rate, deficit flow

Figure 1. Primary-secondary system

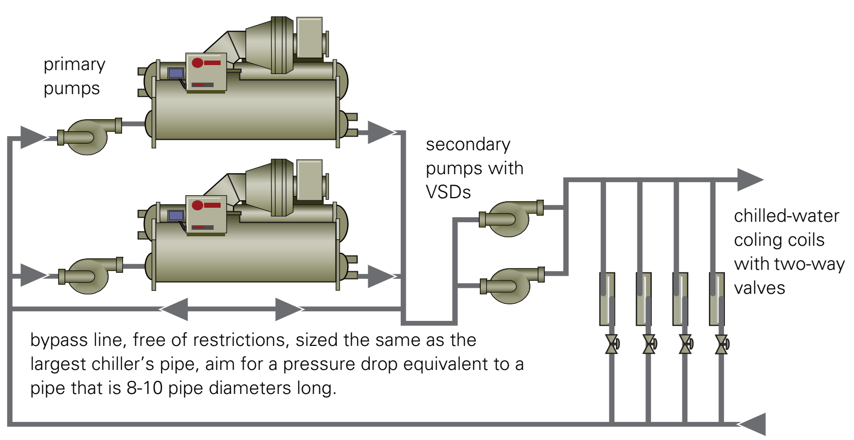

(from return to supply) in the bypass line occurs. This warm, deficit water flow increases the system supply-water temperature and at some point an additional chiller and primary pump are enabled.

A chiller and pump are disabled when excess flow rate in the bypass line is high enough to still have excess flow after a primary pump is turned off. Therefore the maximum flow rate the bypass line ever experiences is a little higher than the design flow rate of the largest chiller. Often designers wait for 10-15 percent excess flow, to ensure that chillers are not cycled on and off rapidly. Therefore the bypass pipe should be sized for 110 to115 percent of the largest chiller's flow rate—which most designers simplify to the same as the pipe size going into the largest chiller.

What happens if bypass sizing is not right? Too little or too much pressure can cause issues...

Too small, pressure too high? An undersized bypass line can result in high fluid velocity, which in extreme conditions may cause pipe erosion, vibration and acoustic issues. The high fluid velocity can also result in high enough pressure drop so that flow is restricted through the bypass, and may actually cause the primary and secondary pump pressures and flow rates to affect each other.

Too large, pressure too low? A benefit of an oversized bypass pipe is a very low pressure drop, however if it's too low, the water may not flow as intended. Figure 2 illustrates the issue:

- Cold water leaves the chillers and flows into the chiller supply manifold.
- A portion of that cold water flows toward the secondary pump. However due to the very low pressure drop in the bypass, some of the cold water seeks the path of least resistance and flows through the bypass line in the excess direction.
- Similarily in the return-water side, warm return water flows in the system return manifold.

Figure 2. Oversized bypass can allow simultaneous flow in both directions

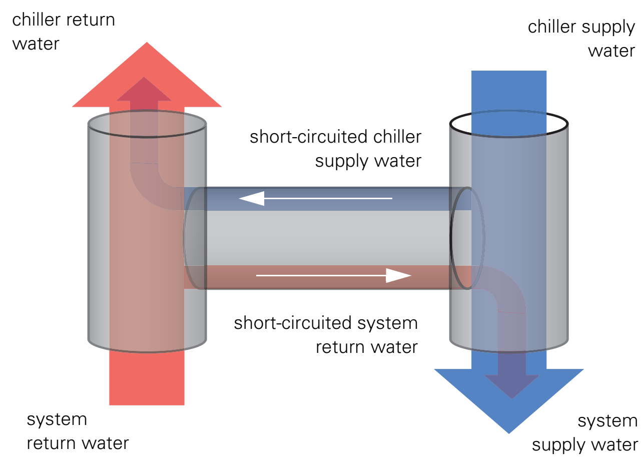

- Much of that water continues to the chillers but some of the warm water (due to low pressure drop in the bypass) flows in the deficit direction in the bypass line (from return to the supply).

As a result there is simultaneous flow in opposite directions because the bypass line has become so large it functions like a tank!

The diluted (higher) system supply-water temperature results in more pumping energy. The short-circuited supply water mixing with return water results in reduced chiller return-water temperature. This can restrict the chillers' ability to fully load.

If the pipes are already installed how can the situation be improved? A resolution for too low of a pressure drop is to impose a modest restriction to keep water from short-circuiting. Aim for the equivalent pressure drop of a pipe that is 8 to 10 pipe diameters long.

Often the first solution considered is to add a valve. Given that the bypass line is oversized, it's likely the additional valve will be big and expensive. It's also likely that at some point a well-meaning but uninformed operator will close the valve too much. Or, they might open it all the way and defeat this solution. Recall that the bypass line in a primary-secondary system should allow water to flow freely, in either direction, as needed.

Add pressure drop. A better option may be to place an orifice in the line. This can simply be a plate with a hole in it. This imposes a pressure drop, but allows water to flow freely in either the surplus or deficit direction. In a few extreme cases it has been neccessary to completely block off the oversized/short bypass pipe and install a properly sized pipe of sufficient length to eliminate the mixing condition.

The best way to avoid issues in a primary-secondary system is to size the bypass line properly during the design process. Simply check the drawings and ensure that the bypass line is smaller than the manifold, and the same size as the pipe going into the largest chiller and 8 to 10 equivalent pipe diameters long. If the pipe is less than 8 to10 pipe diameters long, using elbows to form a "U" adds an appropriate pressure drop.

Variable-primary-flow system bypass line sizing. Conceptually a variable-primary-flow system is simpler to get right. The valve in the bypass line only opens when the system flow rate approaches the minimum flow rate of the operating chiller(s). So the bypass pipe and valve only need to be sized for the largest minimum flow rate. Usually that's the largest chiller's minimum flow rate. However, depending on chiller selections, the largest minimum flow rate might not be for the largest chiller in the plant. Also consider the combined minimum flow required when two chillers are operating at part load, just before sequencing off one chiller.

Should we dynamically vary the condenser water flow?

Yes, savings are available in existing systems designed between 2.5 to 3.2 gpm/ton (12 to 9.4°F ΔT) condenser water flow rate. Controls complexity should be accounted for and the system needs to be properly commissioned.

No, in new systems designed for 1.8 to 2.2 gpm/ton (16.6 to 13.6°F ΔT) condenser water flow rate since they already achieve almost all the savings and reduce system complexity—a lot. Systems designed at these flow rates do not require varying the condenser water flow rate.

Why? Varying condenser water flow rate can be complicated. There are several limit conditions and setpoints to manage:

- The flow rate has to stay above the minimum flow rate as defined by the highest of minimum tower flow, minimum chiller flow, or to produce the static lift in the open portion of the condenser water system. When any of these are close to the system design flow, variable flow should not be attempted.
- At each operating point during the year, determine the optimal condenser water pump and cooling tower speed.
- Make sure not to reduce the condenser flow at conditions and operating points that cause the chiller to surge.
- Ensure the sequence is documented and properly commissioned.

This complexity requires that the system be commissioned, the controls remain operational and future changes are accommodated. Careful consideration should be made for system changes such as chiller, pump, and/or tower replacements.

Table 2. Industry recommendations for condenser water design flow rates

Source	ΔT (°F)	Flow rate(s) (gpm/ton)
Historical practice	9.4	3.0
Today's Industry recommendations		
ASHRAE GreenGuide[1]	12-18	2.3 - 1.7
Kelly and Chan[2]	15	2.0
Taylor[3]	15	1.9

Table 1. Plant annualized kW/ton and percent savings compared to two-chiller base for three alternatives

Alternative	Chiller type	Cooling tower fan	Condenser water flow rate (gpm/ton)	Condenser water flow type	Tower control method*	Plant annualized kW/ton	Savings
Base	VS	VS	3	CF	Opt	0.5462	NA
Variable CW flow	VS	VS	3	VF	Opt	0.5260	3.7%
Reduced design CW flow	VS	VS	2	CF	Opt	0.5255	3.8%
Reduced design CW flow and variable CW flow	VS	VS	2	VF	Opt	0.5252	3.8%

*Near optimal control (Opt) is minimum the sum of chiller + cooling tower fan kW at each operating point during the year

Example. Let's compare the condenser water flow options.

It's imperative that the sum of chiller, condenser water pump and cooling tower fan power is considered. Less-than-optimal results can happen if decisions are based on only one component.

The measure of efficiency is the annual kWh of chiller, cooling tower fans and condenser water pumps divided by the annual ton-hours of cooling. This results in ***annualized performance*** in terms of kW/ton for those components.

Table 1 illustrates a comparison of four alternative two-chiller systems:

- Base design of 3 gpm/ton and constant condenser water flow rate (CF)
- Design of 3 gpm/ton and varying condenser water flow rate (VF)
- Design of 2 gpm/ton with constant condenser water flow rate
- Design of 2 gpm/ton and varying condenser water flow rate

Assumptions. These alternatives assume the pipe size remains unchanged. And that near optimal tower control resets the tower water setpoint to achieve the minimum sum of the chiller plus cooling tower fan kW at each operating point during the year. Two gpm/ton is shown in our example but projects have varying optimal design flow rates.

Effect of flow rate on equipment. With reduced design condenser water flow rate based on present industry recommendations (Table 2):

- Pumps are smaller with significantly lower power
- Chiller power rises marginally
- Cooling towers become more effective as heat exchangers, since warmer water is sent to the cooling tower, resulting in
  - Reduced tower fan power
  - Reduced tower cost
  - Possible reduction of tower size (which further reduces tower cost).

Observations. In all cases the operating savings (Table 1) are very similar, so what guidance can be provided?

- Many existing systems were designed using historical practices and a condenser flow rate of 3 gpm/ton. In existing systems, annualized plant energy can be reduced by varying the condenser flow rate. Be sure to properly commission the system and ensure control is performed correctly long-term.
- For new systems or when all existing chillers are being replaced, using low design condenser flow rates and constant flow reduces installed costs, saves the same amount of energy (or more) as if condenser flow rate is varied and and keeps system control both simple and understandable.

Does operating two pumps at lower speed save more energy than operating one pump?

Perhaps a little, but not nearly as much as many people think.

Why? This came up during an ASHRAE conference conversation with respect to the affinity laws, and a wise consulting engineer raised his eyebrows and said one word: "Think."

As it turns out, the reason the question is asked is because of a misunderstanding, or perhaps lack of *thinking*—about the pump, or affinity, laws. The pump laws describe the relationship between flow rate and power and, in a straight pipe that relationship is cubic. So the thought process is, if an additional pump is enabled, the flow rate per pump will go down and the kW per pump will go down with the cube of that flow rate.

This just isn't true - and is an improper understanding of the pump laws for this situation. What did we miss?

As the pump power equation below shows, pump power is dependent on the flow rate and system pressure drop irrespective of the number of pumps that are operating. The 0.746 and 3960 terms are conversion factors. In the denominator are the pump, motor and drive efficiencies. They may improve or worsen at different pump operating conditions.

$$\text{Pump kW} = \frac{\text{Flow (gpm)} \times \Delta P(\text{ft. } H_2O) \times 0.746}{3960 \times \text{Pump eff.} \times \text{Motor eff.} \times \text{Drive eff.}}$$

Figure 3. Pressure drops in a VPF system

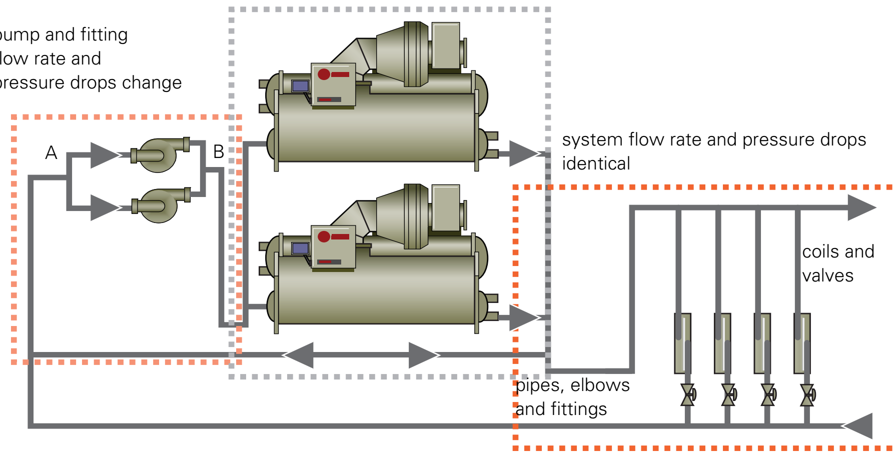

Example. Consider the system in Figure 3. At a given point in time, the system flow rate is identical—no matter how many pumps operate—because it's dependent on the required coil flow rates. The flow rate is the same through the coils and valves, as well as the pipes, elbows and fittings. Meaning their pressure drop is the same. The flow rate through the chillers also remains identical, as do their pressure drops. The only change in the system is that there are now two paths for water flow across the pump manifold.

When the flow rate through the operating pump is reduced, so is the pressure drop through its fittings - and the pump itself. So there is a pressure drop reduction - but it is only the pressure drop from the return manifold (A) to the supply manifold (B), and through the pump.

How about pump, motor and drive efficiencies with two pumps operating? Burt Rishel's ASHRAE Journal article[4] titled, "Wire-to-Water Efficiency of Pumping Systems," explains that there are some operating conditions that allow the combined pump, motor and drive efficiency to rise when more pumps operate. While it takes some calculation time, if the combined efficiency is higher, pump power is lower.

So, is operating an additional pump beneficial?

- The system flow rate does not change,
- Almost the entire system pressure drop is identical.
- There is only a small change in system pressure drop - between the pump return and supply manifold.
- And without study, we don't know if the pump, motor and drive efficiencies are better or worse at reduced pump speed. Some pump manufacturer's selection programs can provide manifolded pump/drive efficiency ratings with various numbers of pumps in operation.

As long as combined efficiency is the same or better, the result is pump power that is a little lower - due to the system pressure drop being a little lower. But it is nowhere near the "cubic" pump savings that some assume.

How much power can be saved by piping chillers in a "series counterflow" arrangement compared to piping chillers in parallel?

Almost 13 percent chiller power reduction is available by piping in a series-counterflow configuration.

- *Using Trane Duplex chillers, more that 19 percent can be saved. Increased pumping power reduces these savings a little, but can be mitigated.*
- *Series evaporators and condensers should be considered when system temperature differences are 14°F or larger.*

Why? Before we examine the system configuration, recall that the two major impacts on chiller performance are:

- The cooling load the chiller must satisfy.
- The refrigerant "lift" the compressor must develop.

Chiller power is proportional to load multiplied by lift. So when either load or lift are reduced, so is chiller power.

Chiller power (kW) $\propto$ *Load* x *Lift*

So, what is "lift"? Refrigerant is compressed from its evaporator pressure to its condenser pressure. This difference is referred to as the *lift*. So measuring these pressures allows lift to be determined. But measuring temperatures is simpler.

Since refrigerant is saturated in both vessels, at a specific refrigerant pressure the refrigerant has a specific temperature. The colder the evaporator refrigerant and the warmer the condenser refrigerant, the higher the lift. To simplify, lift is often approximated to the difference between leaving condenser water temperature and leaving evaporator water temperature.

With this background, Figure 4 shows the lift calculation for a chiller used in a plant where both the evaporators and condensers are piped in parallel. These temperatures are based on a system installed in a convention center and detailed in an ASHRAE Journal[5] article. Each chiller is designed to produce 37°F chilled water and return water back to the cooling tower at 98.9°F. So the lift of each chiller is 61.9°F.

Figure 4. Chiller lift

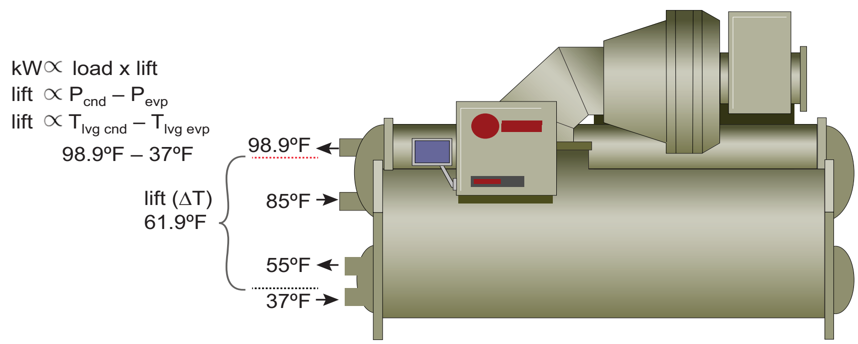

The chillers could also be piped with both evaporators and condensers in *series* (Figure 5). To simplify chiller selection, the chiller making the coldest chilled water receives the coolest tower water. This is referred to as "counterflow" since the condenser water flows counter to the chilled water. By installing the chillers in series, two levels of thermal staging are created—reducing each chiller's lift.

The upstream chiller receives 55°F return water and satisfies about half of the load, cooling the water to 45.1°F. The downstream chiller then cools the water to the desired 37°F.

On the condenser side, the flow goes in the opposite (counter) direction. The downstream chiller receives 85°F water and raises it to 91.3°F. The upstream further increases the condenser water temperature to 98.9°F.

So the respective lifts are:

- Downstream chiller: 54.3°F
- Upstream chiller: 53.8°F
- Average: 54.05°F

The average lift reduction compared to chillers piped in parallel is 7.85°F. Since power is proportional to lift, the power production is reduced by 12.7 percent.

When designing a chilled-water system, optimized *system* performance is the goal.

While the chillers become more efficient, there is additional pump power due to the higher pressure drop of pumping all the water through both chillers. In order to minimize pumping energy, it's common to design series systems with ΔTs of 14°F or larger. ANSI/ASHRAE/IESNA 90.1-2016 requires a 15°F chilled-water ΔT, so series chillers are expected to become more common. One may also consider utilizing single-pass evaporators and condensers to reduce pressure drop and pump power.

Figure 5. Chiller lift reduction savings series-series (or series counterflow)

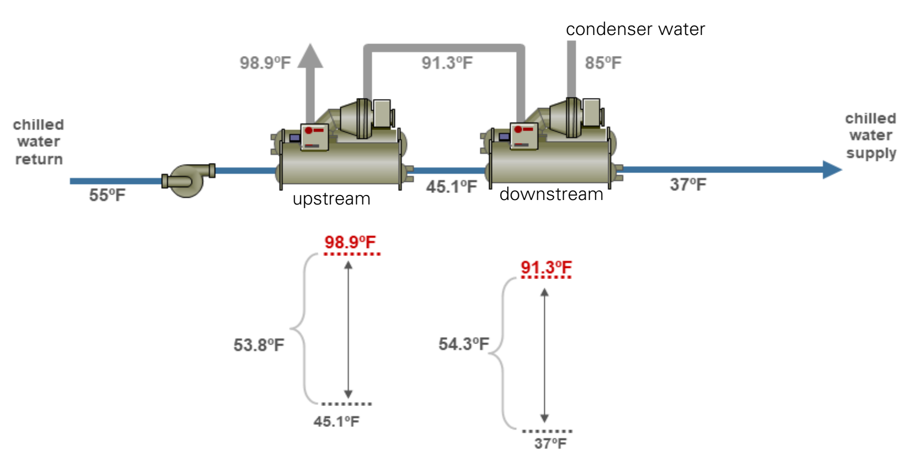

Now let's take the same design using two Trane Duplex® chillers (Figure 6). Each Duplex chiller is essentially a packaged series-counterflow chiller. Installing two of these chillers in series results in four levels of thermal staging. The lift is reduced to an average of 50.1°F. This staging results in over 19 percent reduction in lift compared to the base parallel configuration.

To summarize, chiller power is directly proportional to compressor lift. A simple way to reduce chiller power by almost 13 percent is to pipe the chillers in a series-counterflow arrangement. Two Duplex chillers in a series-counterflow arrangement save over 19 percent chiller power (Table 3).

Table 3. Lift versus chiller power

Configuration	Lift	Reduction (%)
Parallel	61.9ºF	baseline
Series counterflow	54.05ºF	12.7
Series counterflow duplex	50.1ºF	19.1

Figure 6. Chiller lift reduction savings series-series-series (or series-series counterflow)

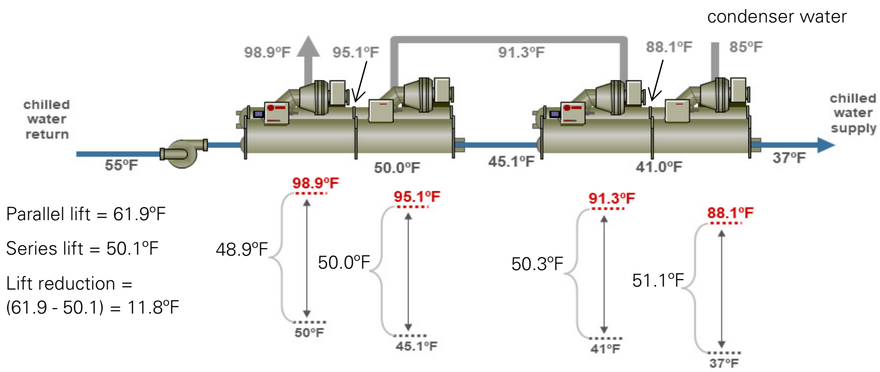

Figure 7. Series arrangement of evaporators and condensers

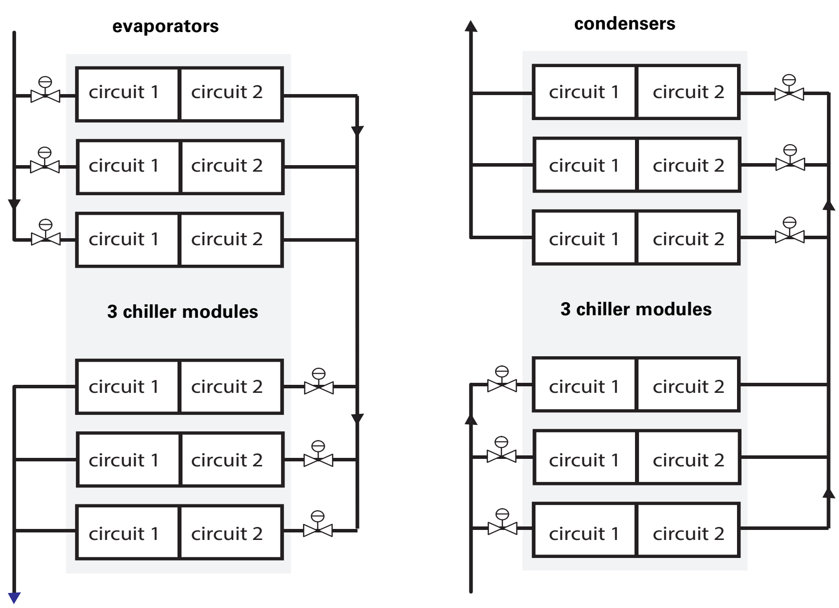

Additional questions about series chillers.

What about redundancy? To allow either chiller to operate alone if the other chiller is being serviced, manual bypass lines and valves are encouraged. When more than two chillers are installed, chillers can be piped in upstream and downstream "pods" with a manifold between them (Figure 7). This allows any upstream chiller to operate with any downstream chiller and provides the same redundancy as if all chillers were piped in parallel. Trane Applications Engineering can provide options for series-counter flow system configurations and their operating characteristics.

Chiller selection capacity. A system can be designed with each chiller handling 50 percent of the load, but it can be more efficient if the capacity split is optimized. A typical optimized (efficiency and cost) split has the upstream chiller designed to meet 53 percent of the load and the downstream the remaining 47 percent. Another benefit is that both chillers can make the system supply-water temperature in the event the downstream chiller is down or being serviced.

Should pressure-independent control valves be used?

Pressure-independent control valves (PICVs) have come down in price, are readily available and help maintain system ΔT, but are higher priced than pressure-dependent valves. Is the additional cost worth it?

Definition. Let's begin by explaining what a pressure-independent control valve is. Those who have sized a conventional control valve are probably familiar with the flow coefficient calculation shown below.

To select the appropriate valve, one must first determine the required flow coefficient; given a flow rate through the valve and a desired pressure drop across the valve. The valve is selected to have an "authority" that will provide accurate and stable control.

$$C_V = Q \times \sqrt{\frac{SG}{\Delta P}}$$

Where,

Cv = valve flow coefficient
Q = flow rate (gpm)
SG = Fluid specific gravity (water = 1.0)
ΔP = Valve pressure drop (psi)

Rearranging this equation, we see that flow is dependent on the pressure differential for a given valve.

$$Q = C_V \times \sqrt{\frac{\Delta P}{SG}}$$

Valve performance would be fine if the pressure differential across the valve was always at the selection condition. But during operation, other parts of the system are constantly changing. This causes the pressure at the valve to change also. With a fixed coefficient, the flow therefore has no option but to also vary.

Even if nothing has changed in a particular space, the control valve must now modulate to adjust flow in response to these system pressure changes. Thus, the "controllability" of a conventional control valve is significantly affected by variations in pressure.

What if there were a separate device that absorbed the system pressure variations so that the control valve always experienced its selected pressure differential? Or what if the valve control gains could be dynamically recalibrated to compensate for the changes in system pressure? In other words, what if we could make the control valve pressure independent? The valve would always have its selected control authority, it would always pass only its design flow when wide open, and it would offer control stability.

A pressure-independent control valve yields a number of advantages:

- Valve control becomes much more stable. To maintain space temperature, the actuator no longer has to adjust for varying system pressures. This will extend valve and actuator life. Also, stability positively impacts system efficiency and coil performance.
- Accuracy improves because the valve is now controlling flow directly.
- Selection is easier. Simply choose a valve that provides the needed flow rate (gpm)—a flow coefficient calculation is no longer needed upfront.
- Installation is easier because pressure-independent valves automatically balance the system. No separate balancing valves are required.

Stability and accuracy are particularly important. Anything that causes a valve to lose accurate and stable flow control, under any operating condition, can result in lower than desired average waterside ΔT. Due to a subsequent increase in required flow rate, lower ΔT results in an increase in pump power and a decrease in chiller plant efficiency. Sometimes the impacts are very significant.

There are two different technologies used to implement pressure-independent control; *mechanical* pressure regulation across the control valve and *electronic* gain modulation for the control valve using flow measurement and valve characteristics.

The mechanical variety basically combines two valve types into one:

- First, a pressure regulating section. As system pressures change, the pressure regulating section automatically adjusts to keep a constant pressure across the control section of the valve.
- Next, the control section of the valve is modulated by the control system to adjust the flow through the valve as space conditions change.

Mechanical pressure-independent valve advantages include:

- Compact size when compared to a conventional control valve plus flow limiting valve package, or to their electronic counterparts,
- Depending on the manufacturer, capability to be paired with any rotary actuator,
- Easy and straightforward to select,
- No additional controls or programming required,
- Near instantaneous response to changes in system pressure which provides optimal control stability.

On the other hand, an electronic pressure-independent valve doesn't actually maintain a constant differential pressure across the control valve surface. Instead, it achieves independent control similar to a pressure-independent VAV air valve: it includes a flow meter in series with a standard control valve. The electronics calculate the instantaneous flow and pressure across the valve and continuously recalibrates the control coefficients to provide stable and accurate control.

Likewise, the electronic variety has a number of advantages:

- Potential for lower hardware costs,
- Provision for actual load measurement,
- Programmable for alternative operation methods. For example, the valve could be set up to limit delta T, not flow.

Note: The 2016 ASHRAE Handbook - Systems and Equipment p.47 states that, "...an authority between 0.25 and 0.5 usually provides the right balance between controllability and energy performance."

$$\text{authority} = \frac{\text{differential pressure of valve}}{\text{differential pressure of valve + differential pessure of branch}}$$

Both mechanical and electrionic pressure-independent valves have communication capability to enable data sharing and trending.

Advantages should be weighed against the additional complexity. For example, special software is needed to setup and maintain the valves for the life of the product. Also, operators need to be trained in this software and software must be kept current.

Example office building. So now that we understand what a pressure-independent control valve can do for a building system, let's look at case study.

A high rise building in Atlanta was identified as a good candidate because it had existing control problems. Once those problems were identified and resolved, one floor was retrofitted with a pressure-independent control valve and was compared to a floor that kept the existing conventional control valve (Figure 8).

First notice how both valves are able to control to a high chilled-water ΔT. In other words: a pressure independent valve isn't *required* to achieve higher ΔTs. However, the conventional control valve is less stable than the pressure independent valve. In particular, notice the considerable variation in valve position for the conventional valve.

As previously discussed, stability improves ΔT. Also, less action on the actuator should improve actuator reliability.

Figure 8. Pressure-independent versus conventional valve operation in Atlanta (3rd and 4th floor performance)

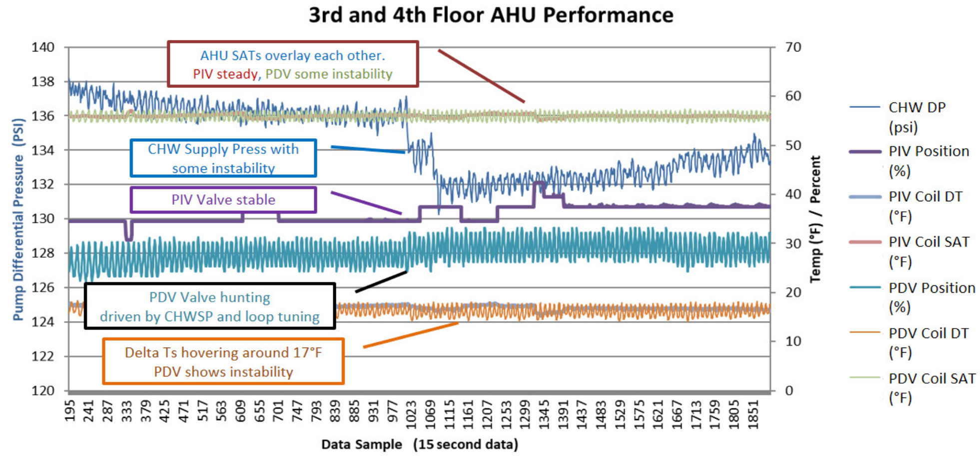

Summary.

- Pressure-independent valves are more stable and more accurate, which improves system ΔT.
- Pressure-independent valves make valve selection much easier. A lot of conventional valves are poorly selected, or one valve C_V is selected for the whole building, even though the pressure across valves varies significantly by distance from the system pumps. It's highly unlikely a poorly selected valve will provide good control.
- Pressure-independent valves are easier to install since they eliminate the need for balancing valves.
- Considering the price premium has been steadily dropping, pressure-independent valves may be cost neutral if all costs, including balancing, are considered.

So are pressure-independent valves worth it? If you're not sure you can get high quality, properly selected, location-specific valves installed on a job; then by all means, specify quality pressure independent valves. It should be well worth the customer's investment.

By Trane Applications Engineering. To subscribe or view previous issues of the *Engineers Newsletter* visit trane.com/EN. Send comments to ENL@trane.com.

References

[1] ASHRAE GreenGuide (third edition), Atlanta, GA: ASHRAE, 2010.

[2] Kelly, D. W. and Tumin Chan, "Optimizing Chilled-water Plants," *Heating/Piping/Air Conditioning*, January, 1999.

[3] Taylor, S. "Optimizing Design & Control Of Chilled Water Plants." *ASHRAE Journal*, (December 2011): 22-34. Available at www.ashrae.org.

[4] Rishel, Burt, "Wire-to-Water Efficiency of Pumping Systems," *ASHRAE Journal*, April, 2001.

[5] Schwedler, M. "Series-Series Counterflow for Central Chilled-Water Plants," *ASHRAE Journal*. (June 2002):pp. 23-29

Trane,
A business of Ingersoll Rand

For more information, contact your local Trane office or e-mail us at comfort@trane.com

ADM-APN067-EN (October 2018)