The Demise of the Primary-Secondary Pumping Paradigm for Chilled Water Plant Design

Accepting that low ΔT chilled water plant syndrome exists in almost all big distributed chilled water systems and recognizing the need to seek design solutions that can cope with or prevent it

By WAYNE KIRSNER, PE,
Kirsner, Pullin & Associates, Atlanta, Ga.

I've got some bad news—the 20-odd-year experiment with primary-secondary design of chilled water plants hasn't panned out. If you've designed a large distributed chilled water system and monitored the operation of the central plant, you already know about the problems: the Δ*T* of the chilled water (CHW) returning to the campus plant is below the design value for which the chillers and pumps were selected—in fact, it's way below; the secondary CHW flow doesn't vary a hoot; and the expensive variable-speed drive (VSD) purchased to vary the flow of the secondary pumps was great for test and balance but hasn't done much since (besides heating up the central plant building). The low Δ*T* at the plant causes the operators to run extra pumps and chillers to meet the load, which, in addition to reducing the plant's cooling output capacity, wastes energy. The system may be keeping the campus cool, but you know it's inefficient and idling a lot of chiller capacity.

The problem described above has come to be known as "low Δ*T* central plant syndrome." To my knowledge, every large chilled water plant serving distributed loads is afflicted with it to some degree. The article "Troubleshooting Chilled Water Problems at the NASA Johnson Space Center" (*HPAC*, February 1995)[1] describes a typical situation. A central plant originally designed for a 16 F Δ*T* between the chilled water return (CHR) and chilled water supply (CHS) could only develop an 8 F Δ*T* because of low CHR temperature from the campus. This meant not only that twice as much CHW as originally intended had to be pumped around the 5-mile campus piping loop but also that the seven 2000-ton chillers in the central plant couldn't be loaded much beyond half their capacity. Thus, operators were usually forced to run twice as many chillers to meet the campus load, and the frictional loss in the mains due to the excessive CHW flow made it

[1] *Wayne Kirsner authored the February 1995 article cited above as well as the article "What Caused the Steam System Accident that Killed Jack Smith?,"* HPAC, *July 1995.*

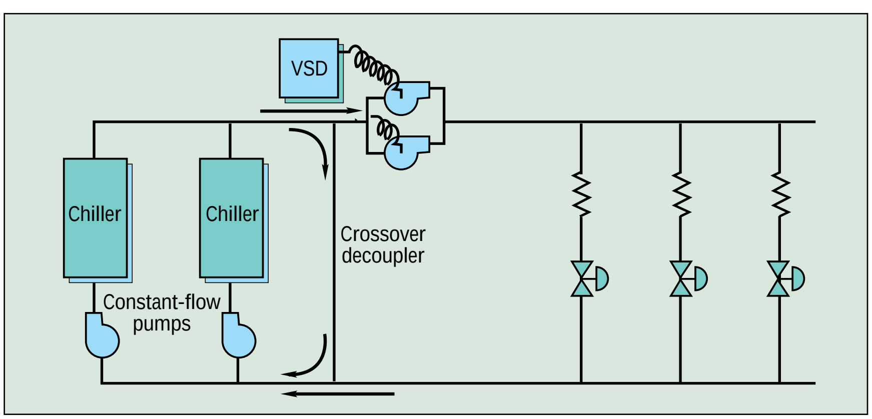

1 Archetypal primary-secondary CHW plant design.

tough to deliver sufficient CHW to hydraulically distant buildings.

The causes of low ΔT syndrome are not mysterious, but they are often pervasive and thus can be hard to remedy. Low ΔT can be caused by dirty cooling coils, throttling valves with insufficient shutoff capability, reset CHS temperature, poorly controlled blending stations, and of course, CHW bypassing out in the system. But most often, low system ΔT is the result of faulty controls and improperly adjusted set points. This article, however, is not about the causes of central plant syndrome. It's about *accepting* that the problem exists in virtually every big distributed chilled water system and then *recognizing* the need to seek design solutions that can cope with it, if not prevent it.

So why can't a standard primary-secondary chilled water design cope with low CHW ΔT?

Problem #1

The primary-secondary control scheme is "blinded" by low ΔT central plant syndrome. Fig. 1 depicts what I would describe as the archetypal primary-secondary chilled water schematic configuration. The primary feature of the configuration is the decoupled primary and secondary loops, which allow constant flow through the chillers while permitting varying flow in the system to save pumping energy. Chillers are staged on and off based on CHW flow through the crossover bridge (although the sensor may be elsewhere). The *sole indicator* of system load, upon which control of the chillers and pumps depends, is *chilled water flow*.

In a plant with low ΔT syndrome, CHW flow is no longer much of an indicator of load. The amplitude of flow variation is just a fraction of the amplitude of load variation. Fundamentally then, a primary-secondary control scheme that depends on system flow to gauge system load is virtually blind to load variation.

Problem #2

The primary loop is constant flow. Constant flow through chillers is a highly desirable feature of primary-secondary chilled water plant design, and most chiller manufacturers still prefer and recommend it. I've been convinced, however, that most modern chiller controls no longer require constant flow to keep the chillers out of trouble. Let me explain.

When chiller vanes were controlled by conventional pneumatic proportional controls, response time to changes in load was necessarily slow and gradual to prevent overshoot and hunting as the chiller controls tried to achieve leaving CHS set point. Hence, chiller capacity controls would lag behind a sudden load change. If the change was a drop in load, the chiller would overcool the leaving CHW, dropping it below set point until capacity control vanes could react to reduce chiller refrigerating capacity. If the drop in load was sharp enough, the chiller's low evaporator temperature safety would knock the chiller off line, requiring a manual reset to restart the chiller. This is a situation to be avoided.

Now consider the response of a chilled water plant designed for constant flow versus one designed for variable flow in the event that load across a fully loaded chiller suddenly dropped in half. (This is a severe upset, but it's not far-fetched at all. Starting a second chiller in a two-chiller plant, where identical chillers operate in parallel, typically results in the load to the active chiller being halved.[2]) In a constant-flow primary loop designed to chill, say, 55 F CHR to 45 F CHS, a 50 percent drop in load would manifest itself in CHR temperature rising to 50 F. (This might occur because approximately half the primary flow of 45 F CHS is recirculating through the crossover bridge to mix with the 55 F CHR from the system.) The 50 F CHR entering the formerly fully loaded active chiller would initially be subjected to the full cooling capacity of the chiller until its controls could respond to decrease capacity. The chiller would thus tend to drive the entering 50 F CHW down toward 40 F.

Compare this upset condition to a variable-flow configuration. Starting a second equal CHW pump could cut CHW flow through the active chiller roughly in half.[3] The active chiller would initially continue to try to apply its full output capacity to half the mass flow, thereby doubling the ΔT of CHW passing through it—*i.e.*, it would tend to drive 55 F CHR down to 35 F. This is pretty close to freezing. If the design ΔT was larger, the CHW would be driven down below freezing. In either case, a simple low evaporator temperature sensor would likely cause the chiller to trip off line to protect it from freezing. The constant-flow chiller, in comparison, whose leaving CHS temperature dips only half as far, would probably remain on line. For this reason alone, one can easily understand why chiller manufac-

[2] *For example, in a variable-flow plant, flow through the active chiller will be cut roughly in two as the second chiller's pump instantly usurps half the flow. In a primary-secondary plant, approximately half the total primary CHW flow recirculates through the crossover bridge and, once the second chiller's compressor has begun outputting CHW, mixes with system secondary CHR to halve its ΔT.*

[3] *Assuming immediate system control valve response and the absence of central plant syndrome. If control valve response were slow or the system were afflicted with central plant syndrome, flow would not suddenly fall to half, and thus the upset condition would be far less traumatic.*

turers would prefer constant flow through chillers.

So what's changed to invalidate this argument? The low evaporator temperature control is more sophisticated, for one. It's no longer simply a low-temperature safety cutout. The Trane Company's microprocessor-based control, for example, integrates (*i.e.*, sums) the number of degree-seconds (deg-sec) below the low evaporator temperature set point. Don Eppleheimer of Trane tells me that if this sum remains below 50 deg-sec, the control logic will not initiate a safety shutdown. This means:

◆ Evaporator temperature may drop *below* freezing momentarily.

◆ The chiller's capacity controls are allowed time to catch up with the load change.

In fact, the sophistication of the control logic is such that Trane feels confident in setting its low evaporator temperature set point as low as 30 F.

The upshot of this improvement, and to a lesser extent the capacity control improvements, is that chillers can survive a severe upset condition in CHW flow without tripping off line. In fact, and this is really the proof of the pudding, Mr. Eppleheimer says that Trane routinely tests its chillers to insure they can withstand a 50 percent drop in CHW flow without tripping the low evaporator temperature safety cutout. (Other manufacturers have different control strategies for handling upsets in applied load. York's chillers, for example, can accept an input that delays powering down of the chiller compressor upon a large drop in leaving CHS temperature as long as the temperature does not fall below 36 F. Carrier's low evaporator temperature control overrides the chiller capacity controller to close compressor vanes should evaporator temperature approach 33 F. At 33 F, the safety shuts down the machine.)

There is another benefit to constant flow through chillers. The possibility of laminar flow through the evaporator due to low CHW flow is eliminated. This condition can easily be avoided, however, in a variable primary flow system. Burt Rishel of Systecon suggests the best way to do this is the old-fashioned way—with a bypass from CHS to CHR opened via a signal from a flow meter or differential pressure controller across the chiller. This might seem to replicate the expense of a primary-secondary crossover bridge, but the difference is that the recirculation pipe is sized to handle no more than about half a chiller's design flow, and its function is less likely to confuse the operator.

So the rationale for avoiding variable flow through chillers is, in my opinion, no longer compelling. But even if constant flow through chiller evaporators is no longer essential for stable chiller operation, *what's wrong with it?*

A constant-flow primary CHW system with one nonvarying pump per chiller cannot respond effectively to low ΔT syndrome. If CHR temperature returning from the system is below design and cannot be raised, a central plant operator's only option in responding to a call for more cooling capacity is to energize more pumps and more chillers.[4] What would be preferable, of course, would be to increase the CHW flow through individual chillers to load them more fully. But with the primary-secondary configuration shown in Fig. 1, this is not possible. Primary-secondary systems can be retrofitted, of course, in response to low ΔT. New pumping capacity can be added and flow through chillers can be increased up to the manufacturer's recommended maximum rate. Perhaps evaporators can even be converted from three-pass to two-pass, reducing pressure drop through the chillers. But obviously, the constraints of the existing equipment limit the flexibility to cope with the need to force more water through the chillers. Furthermore, retrofitting and adding equipment to accommodate a lower ΔT within the context of the existing constant-flow design results in permanently locking in the higher CHW flow rate. A retrofit of this kind effectively throws in the towel on the quest to find and fix the root causes of low ΔT out in the system.

So if not constant flow through the chillers, then what? A variable-flow CHW pumping scheme can respond to low system ΔT if the pumps are selected for excess capacity. In fact, in my view, designers need to consider the design ΔT for which they select the

continued on page 77

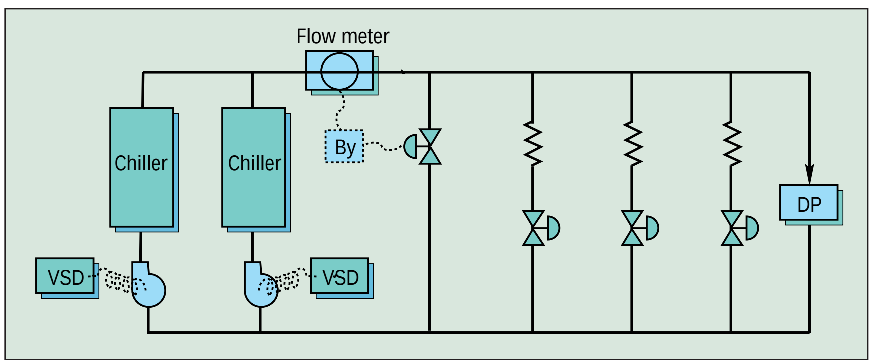

2 Typical variable-flow CHW plant design.

[4]*Assuming pumps are not ganged in a common header, of course.*

[5]*The maximum tube velocity recommended in the* ASHRAE Equipment Handbook *is 7 fps; most manufacturers recommend 11 fps.*

continued from page 75

chillers as a *target value* and then provide for the *eventuality* that extra flow beyond the design target may be required. In practice, this means:

◆ Selecting chiller evaporator tubes for tube velocities not more than about 5.5 fps at design so that flow can be increased up to twofold if necessary.[5]

◆ Selecting pumps to overpump the chillers. The best scheme is to bank the pumps and provide them with VSDs.[6]

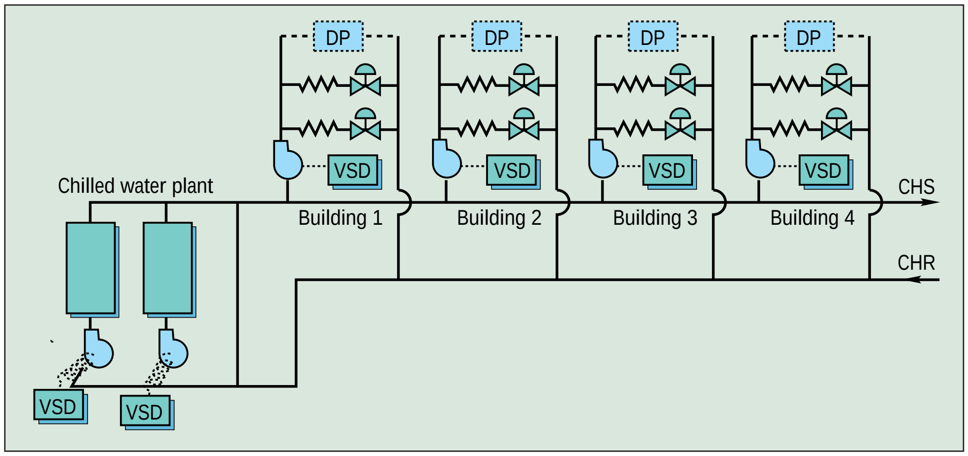

3 Distributed campus CHW pumping.

With variable flow pumping through the chillers, the crossover bridge and the secondary pumps can be dispensed with, so a typical schematic layout for a simple building can look like Fig. 2. Chillers are staged based on leaving CHS temperature. When a chiller can't hold leaving CHS temperature set point, a second chiller is energized. Pump speed is controlled by a differential pressure sensor situated across the hydronically farthest coil. A flow meter and smart controller open a bypass valve should flow through the chillers fall below the manufacturer's recommended minimum.

The advantages of this system are:

◆ It automatically responds to low ΔT by increasing flow through chillers.

◆ There's only one set of pumps.

◆ Minimum chilled water flow is pumped.

◆ The system is simpler; there's no decoupling bridge.

Problem #3

Secondary pumping is not the most efficient pumping distribution scheme. In Fig. 2, a single set of pumps handles the job of both the primary and secondary pumps. But in a big system, a single set of pumps is not always desirable. If the pressure needed to pump an entire campus is large, it's advantageous to place a second set of pumps, and perhaps even a third set, out in the system to avoid imposing high pressure on the equipment close to the pumps' discharge. Secondary pumping, as shown in Fig. 1, can achieve this objective, but *it's not the most efficient* pumping scheme. That's because the same head is imparted to all CHW passing through the secondary pumps, whether it's making the short trip through the closest building or the longest trip through the hydraulically most distant building. The extra head imparted to CHW passing through the closer buildings must be wasted across balancing valves and/or throttling valves at those buildings. Only the small fraction of the total CHW flow going to the most distant building is produced without wasted energy.

A better way to pump distant loads is via distributed pumping, as illustrated in Fig. 3. Burt Rishel of Systecon gives credit to Wilber Shuster of Cincinnati for first proposing this pumping scheme. Distributed building pumps assume the function of the secondary pumps. Each pump is sized to deliver its building's flow at just the head needed to pump the building hydronic loads *and* draw the CHW through the mains from the central plant. There are no decoupling loops at the buildings, so no CHW is bypassed. No balancing valves are needed to eat up excess head since there is none. Pump speeds are controlled by VSDs receiving signals from differential pressure switches at the end of the loop in each building. Pumping horsepower saving equals the sum:[7]

$$\sum_{i=1,\,N-1} \text{CHW flow}_i \times (H_{\text{building }N} - H_{\text{building }i}) / (3960 \times \eta_p)$$

The primary pumps are VSD-controlled, as before, and can operate in series with the distributed pumps or be decoupled as shown in Fig. 3. If decoupled, the VSDs would be controlled to maintain slightly positive flow from CHS to CHR in the crossover bridge and not let flow through any chiller go below its minimum recommended value. Chillers are

[6]*Of course, oversizing the pumps and balancing them down with a throttling valve is not an option unless you routinely wear your shoes on the wrong feet and when you tighten your belt, cut off your windpipe.*

[7]*Assuming equal pump efficiencies for building pumps and hypothetical secondary pumps.*

staged based on their ability to maintain leaving CHS temperature.

The advantages of this system, besides minimizing pumping power, are:

◆ It minimizes the potential for low system ΔT by eliminating crossover bridges at the buildings, but if it does become a problem, the pumps and chillers can effectively deal with it.

◆ It reduces head pressure imposed on equipment.

◆ It's simple and, more importantly, looks simple to the operators who run it.

The only unusual aspect of distributed pumping is that it reverses the typical pressure gradient in the system. The CHS main is negative with respect to the pressure in the CHR main. Thus, every load must be pumped.

In conclusion . . .

The traditional arguments for desiring constant flow through chiller evaporators no longer carry much weight; most modern microprocessor-based chiller controls can effectively deal with upsets due to variable flow. Moreover, constant-flow primary designs cannot respond to the need to put more CHW through chillers in the event that the distribution system returns low CHW ΔT to the central plant.

A variable-flow design with pumps either oversized and controlled by VSDs or banked can respond to low ΔT central plant syndrome. Thus, for the same reason that we as HVAC designers provide freezestats upstream of cooling coils, nonoverloading motors to drive pumps and fans, and tube pull space at chillers, boilers, and air-handling units, we need to design chilled water plants that can anticipate the possibility of low CHW ΔT and respond to it. Therefore, I believe it's time to put primary-secondary pumping back into our tool bag of applications to address specific design situations and adopt a new paradigm for chilled water system design. HPAC

Bibliography

University of Wisconsin's seminar on Chilled Water Plant Design, organized by Harold Olsen.

Burt Rishel's seminar speech, "Current Trends in HVAC Water System Design."

Gil Avery's article "Designing and Commissioning of Variable Flow Hydronic Systems," *ASHRAE Journal*, July 1993.

Al Utesch's seminar speech at the University of Wisconsin.