

Primary/Secondary Type Variable Flow Chilled Water System Operation

The two sketches included later in this attachment are intended to illustrate the difference in plant operating characteristics for a primary/secondary variable flow chilled water plant piped with the bypass in the “inside the plant” location vs. the “outside the plant” location. The following concepts apply to both illustrations and are illustrated in Figure 1 below. (Note that color is used in the figures to illustrate some of the concepts so viewing this document in color will be helpful.) A related PowerPoint® demonstrates these concepts graphically and may be a useful reference.




Figure 1 - Typical primary secondary type variable flow chilled water system components

Operating Principles

- 1 Variable flow systems are implemented based on the concept that chiller plant efficiency can be optimized if the plant flow rate can be made to follow the load profile; i.e. at low load conditions, the plant flow rate is low and at high load conditions the plant flow rate is high. Generally this is accomplished by installing two-way valves at the loads so that the flow through the loads varies with load condition. The loads are then served by a variable flow distribution system. Pumping energy is saved by designing the distribution pumping system so that the pump flow rate is modulated with load in response to the changing two way valve positions. Several approaches are possible ranging from simply allowing the valve modulation to push the pump(s) up its curve to measuring system differential pressure and varying pump speed as required to maintain a constant differential regardless of valve position. When properly designed, the temperature difference between the supply and return mains on this type of system will tend to be constant, regardless of the loads while the flow varies with the loads. This is in contrast to a constant flow system where the flow remains constant regardless of the loads and the temperature difference between the mains varies with load.
- 2 Many chillers can not tolerate major variations in evaporator flow rate without operating difficulties ranging from control problems and nuisance safety trips to equipment damage or failure due to tube frosting or freeze-up. These problems are primarily related to the low evaporator pressures and temperatures that occur when the machines are operated with very low evaporator flow rates and loads relative to their design capacity. As a result, most variable flow chilled water system must be designed to maintain relatively constant flow rate through the chiller evaporators.¹ This is accomplished by providing independent an evaporator pump(s) that moves water out of the variable flow loop, through the evaporator(s) and then returns it to the variable flow loop.
- 3 The section of pipe between where the evaporator pump(s) removes water from the variable flow low and where it returns it to the variable flow loop is commonly called the plant bypass line. The direction of flow in this line and the temperature in it will change as a function of load. This will be discussed in subsequent bullets. Ideally, the bypass piping will be configured in a manner that results in very little pressure drop between the points where the chiller evaporator loop(s) connects to the system. As a result, the head that the evaporator pump(s) will see will remain constant, regardless of the flow condition in the variable flow distribution loop. Due to the nature of centrifugal pumps, if the head on the evaporator pump(s) remains constant, then the flow produced by the pump will remain constant. The bottom line is that a properly sized bypass line hydraulically decouples the pumped evaporator loop(s) from the distribution loop,

¹ Fairly recent technology advances are allowing manufactures to provide machines that can actually tolerate wide variations in evaporator flow. As a result, system configurations are beginning to evolve to eliminate the independent pumping circuit through the evaporator. Instead, the chiller evaporators are being piped directly into the variable flow circuit, thereby saving first cost and operating cost. First costs are reduced due to less piping complexity and fewer pumps. Operating costs tend to be lower because even though the head of the chiller is added to the variable flow loop, the energy burden associated with moving water through the chiller varies with load rather than being constant regardless of load.

thereby allowing constant flow to be maintained through the evaporator(s) while the distribution flow varies with load.

- 4 Due to the physical principle of conservation of mass, the flow in the bypass line will vary and change direction as a function of the relationship between the flow in the distribution loop and the flow in the evaporator loop(s).

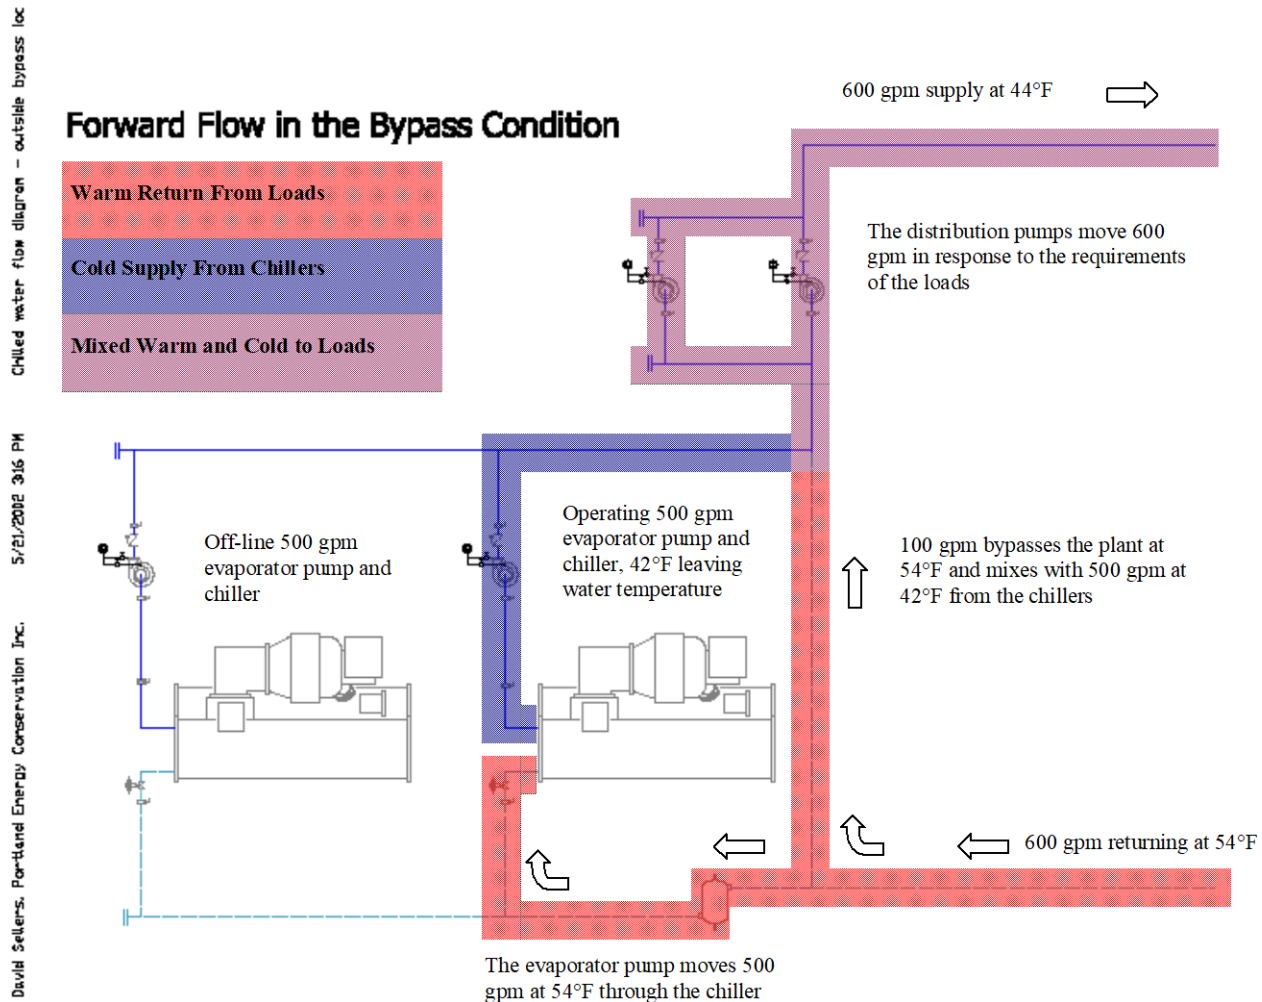


Figure 2 - Forward bypass flow condition. Note how the supply temperature is higher than the chiller leaving water temperature due to the mixing of the return water with the chiller leaving water. Operationally, it is probably about time to start a second chiller unless the excess load condition will be of short duration.

- 5 When the flow in the distribution loop exactly matches the flow in the evaporator loop(s), there will be no flow in the bypass.² This is a theoretically interesting condition that seldom occurs for any length of time in an operating plant.

² This can be a difficult concept to accept if you don't fully understand the principle of conservation of mass. Basically, this principle states that in a Newtonian model of the universe, we can't really create or destroy matter. We can alter its state (solid, liquid, or gas), break it down into more fundamental components (atoms

- 6 If the variable flow loop is moving more water than the constant flow loop(s) through the chiller evaporators, then the flow through the bypass line will be from the return end towards the supply end and the water in the variable flow loop will bypass the chillers (hence the name “bypass pipe”). This condition is often called “forward bypass flow” and is illustrated in Figure 2.
- 7 On the other hand, if the variable flow loop is pumping less water than the constant flow loop(s) through the chiller evaporator(s), then the flow in the bypass will be from the supply end towards the return end. This is commonly called “reverse bypass flow” and is illustrated in Figure 3.
- 8 Under conditions of forward bypass flow, the supply water temperature to the loads will be higher than the water temperature leaving the chillers because warm return water bypasses the chillers and blends with their leaving water as the flows combine for distribution to the loads. The amount that the supply temperature to the loads increases above the leaving water temperature from the chillers is a function of the amount of water bypassing the plant relative to the total flow. If the bypass percentage is small relative to the total flow, the effect will be minimal. If the bypass percentage is high relative to the total flow, then the effect can be significant.
- 9 Under conditions of reverse bypass flow, the supply temperature to the loads will be identical to the chiller leaving water temperature and the chiller entering water temperature will be depressed below the return temperature due to the mixing of the reverse flow in the bypass with the return flow at the tee where the bypass pipe connects to the return system and to the entering side of the chiller evaporator(s). This causes the chiller(s) to unload. The magnitude of the depression in temperature is again a function of the ratio of reverse bypass flow to total system flow.
- 10 In general one of the goals of operating a variable flow plant of this type is to minimize the forward bypass flow so that the plant leaving water temperature is nearly identical to the chiller leaving water temperature, thus ensuring that the system is supplied with water at or near the chiller leaving water temperature set point. In fact, most plants are controlled to always maintain some reverse flow in the bypass line. The minor, short term excursions that often occur when chillers are cycled can usually be tolerated by most systems since the thermal fly wheel of the piping circuit and loads tends to cancel them out. For example, on one site

vs. compounds), put it together into more complex molecules, etc. but when all is said and done, the mass of the elements going into the process will be the same as the mass of the elements coming out of the process. In the case of our variable flow system, this means that the volumes of water flowing into and out of the tees that connect the evaporator pumps to the distribution system and the bypass line must all add up. Consider the tee in the piping circuit where the evaporator pump(s) connect back into the distribution loop after directing the water through the chillers. If the flow in the distribution system is greater than the flow through the evaporator, then the distribution system will recirculate some water from the return system through the bypass line and mix it with the water from the evaporator. If distribution flow exactly matches the evaporator flow, all of the water coming back from the distribution system will be picked up by the evaporator pump(s) at the tee where the return line and bypass line connect. It will then be moved through the chillers and directed to the tee where the bypass line and the supply line connect. At that point, the distribution pumps will pick up all of the water from the evaporator pumps and move it out to the system. The water in the bypass itself will be stagnant/un-moving under this very specific condition.

where a 4,000 ton system served a semiconductor clean room, if all of the chillers in the plant shut down but the distribution pumps remained in operation, there was a time delay of 10-20 minutes between the time the chillers failed and the time the first space temperature warning alarm from the most critical, temperature sensitive section of the clean room showed up.

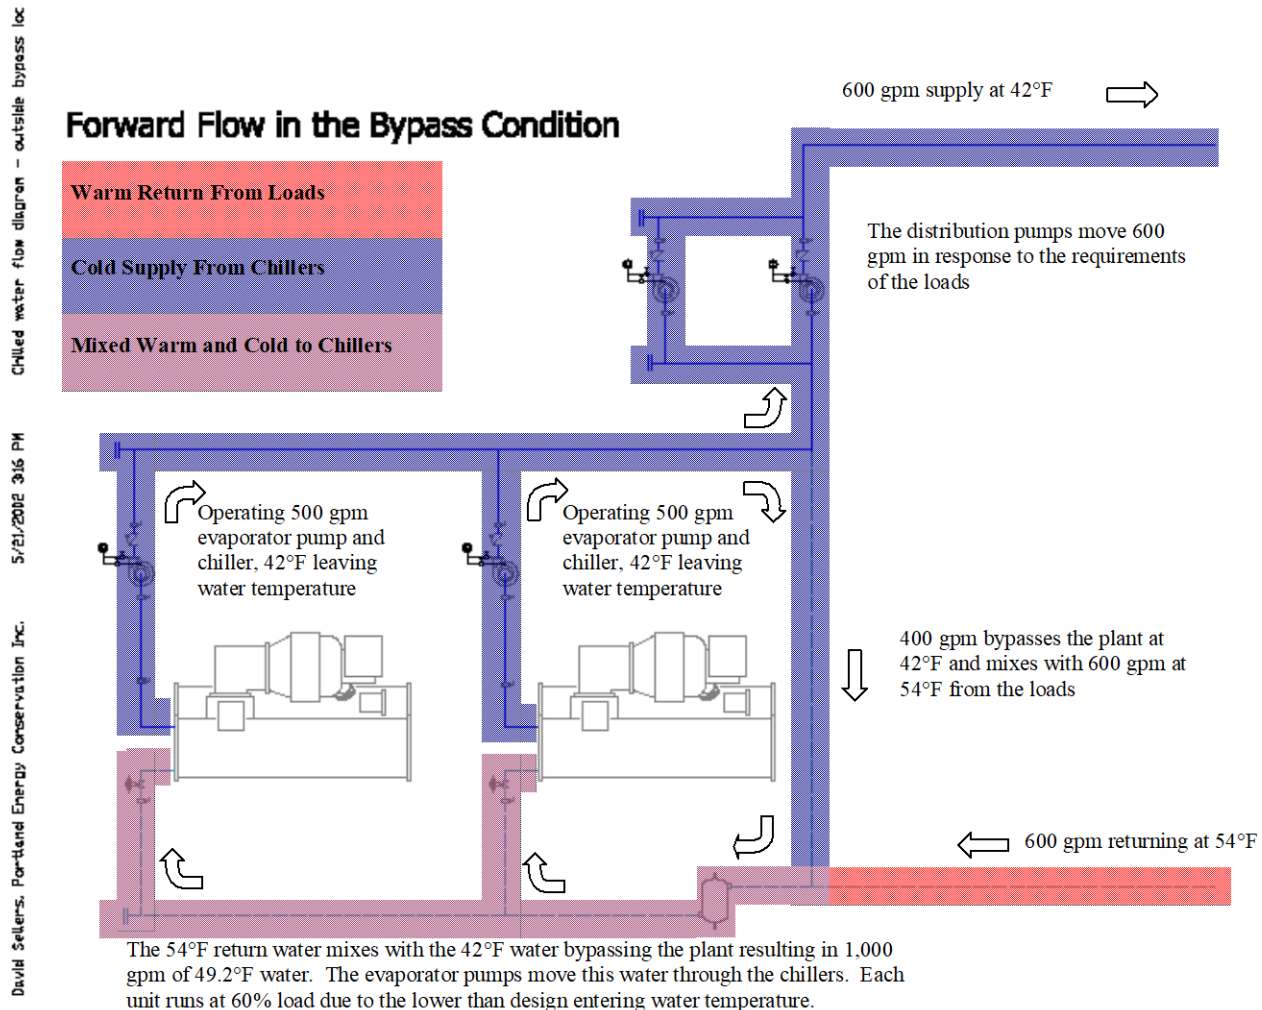


Figure 3 - Reverse bypass flow condition. This is what the plant in Figure 2 would look like if a second chiller was started. Note how the supply temperature is the same as the chiller leaving water temperature and the chiller entering water temperature is depressed below the return temperature due to the mixing of the return water with the chiller leaving water.

- 11 If forward flow in any significant quantity is allowed to persist for an extended period of time in this type of plant, significant operating problems can result. The elevated supply temperature to the loads causes them to open their control valves further in order to satisfy the requirements of their control systems. This further aggravates the forward flow problem and a viscous cycle results. To solve the problem the operating staff must either shed load or place additional chillers on line. Adding chillers often has an electrical demand penalty associated with it that can be very significant.

12 This type of variable flow system depends on the relationship between load and flow mirroring each other. For systems of this type to function correctly, the loads need to be selected so that the temperature difference across them matches the design temperature difference across the chillers. In addition, the heat transfer elements serving the loads need to be selected and circuited so that the temperature difference across them tends to remain constant as the flow is varied in response to load changes. System designs that do not address these important issues can be very difficult to operate because the flow conditions may not reflect the actual load conditions. Consider a system that was designed to handle loads with a 12°F temperature rise where the chillers were selected to produce a 12°F temperature drop. If a major circuit is added to this system that serves loads with a low temperature rise capability then there could be problems.³ On a design day, the 12°F loads will be at their design flow rate producing a 12°F temperature rise. The low temperature rise loads would also be at their design flow rate, but would not be producing a 12°F temperature rise. The return water from the low temperature rise loads will mix with the return water from the 12°F loads in the return system and the result will be that the chiller plant will see a design flow rate, but the temperature difference between the mains will be less than the design difference of 12°F. How much different will be a function of the ratio of the flow from the low temperature rise loads to the total system flow, but if the low temperature loads are a significant portion of the total flow, the temperature difference on a design day could be several degrees less than what the chillers were selected for. As a result the chillers cannot fully load because their entering water temperature is below design. If the pumping rates between the distribution system and chiller evaporators were not matched, there will also be forward flow in the bypass and the supply temperature will spiral up, making a bad situation worse. If extra capacity is available, placing it on line can solve the problem, but plant efficiency will suffer because the operators are running three chillers (and their auxiliary equipment) to serve a two chiller load. If excess capacity doesn't exist, then the chiller LWT set point can be lowered which will lower the system supply temperature and may change the performance of the loads enough to cause them to throttle and reduce the over-all flow rate. But again, this is at the expense of plant efficiency since lowering the chiller LWT generally will reduce the efficiency of the machine for a variety of reasons.

13 The temperature of the water in the bypass line is a very good indicator of the plant operating mode. Unless there is no flow in the bypass (a rare occurrence in a real time operating environment) it will either be at the system return temperature or the system supply temperature. If it is at the return temperature, then there is forward flow in the bypass and additional capacity may be desirable. If it is at the supply temperature, then the flow in the bypass is reversed and the current on-line capacity exceeds the requirements of the loads.

³ One example of a low temperature rise loads that are quite common is a fan coil unit circuit serving a hospital patient care wing where the fan coil units have 3 or 4 row coils because they are primarily intended to offset space sensible load; ventilation loads are handled by an independent system. Fan coil units serving a load of this type often will have only a 6 or 8°F temperature rise under design conditions.

Seattle Court Building Chiller Plant at Part Load with an "Inside the Plant" Bypass (As Drawn Configuration)

5/21/2002 3:16 PM

David Sellers, Portland Energy Conservation Inc.

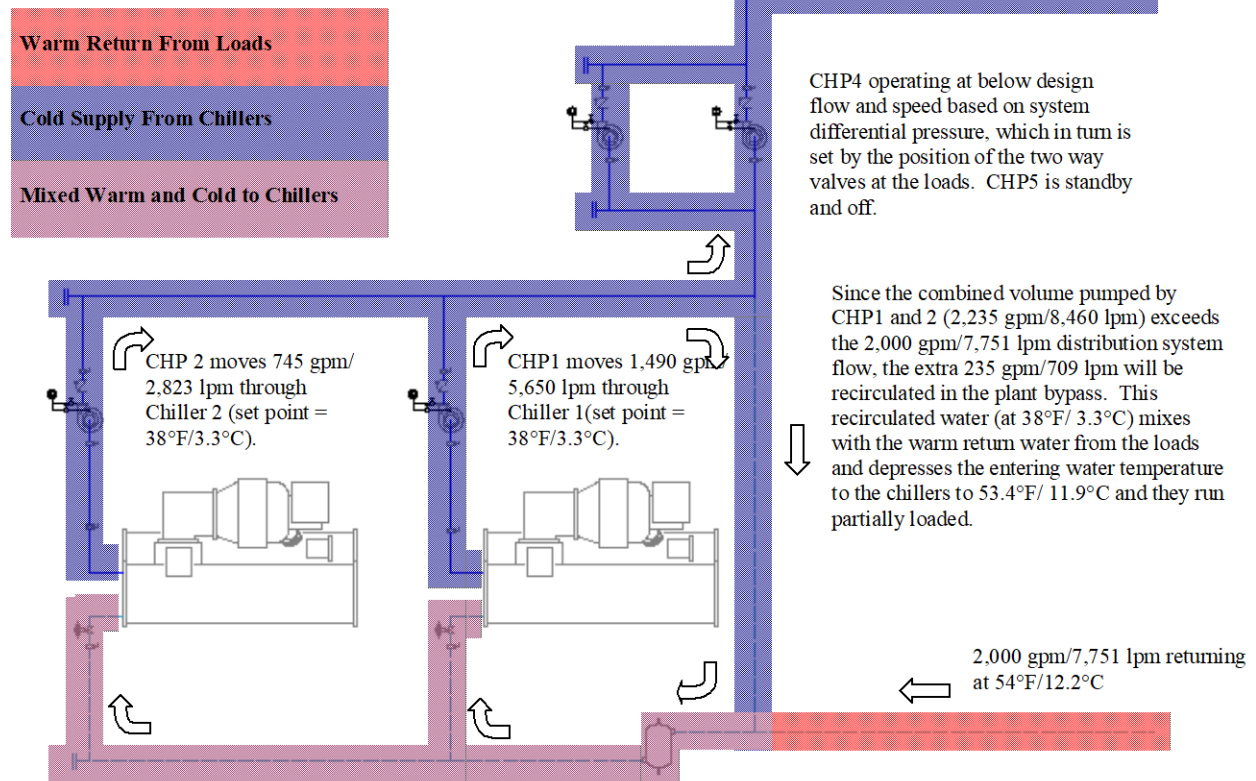


Figure 4 - Chiller plant as currently piped based on the contract documents. Notice how the inside the plant bypass location causes both chillers to see the same entering water temperature any time both of them are operating. As a result they will both load to about the same percentage of their total capacity.

Bypass Location

The exact location of the bypass line has a significant impact on how this type of plant will operate. The Figures 4 and 5 illustrate this in the context of a typical variable flow chiller plant. Figure 4 illustrates how the plant would function at a part load condition with a load in excess of 1,000 tons and the bypass piping arrangement shown on the construction documents. In this configuration, both chillers will see the same entering water temperature and load to the same percentage of capacity.

Seattle Court Building Chiller Plant at Part Load with an "Outside the Plant" Bypass (As Drawn Configuration)

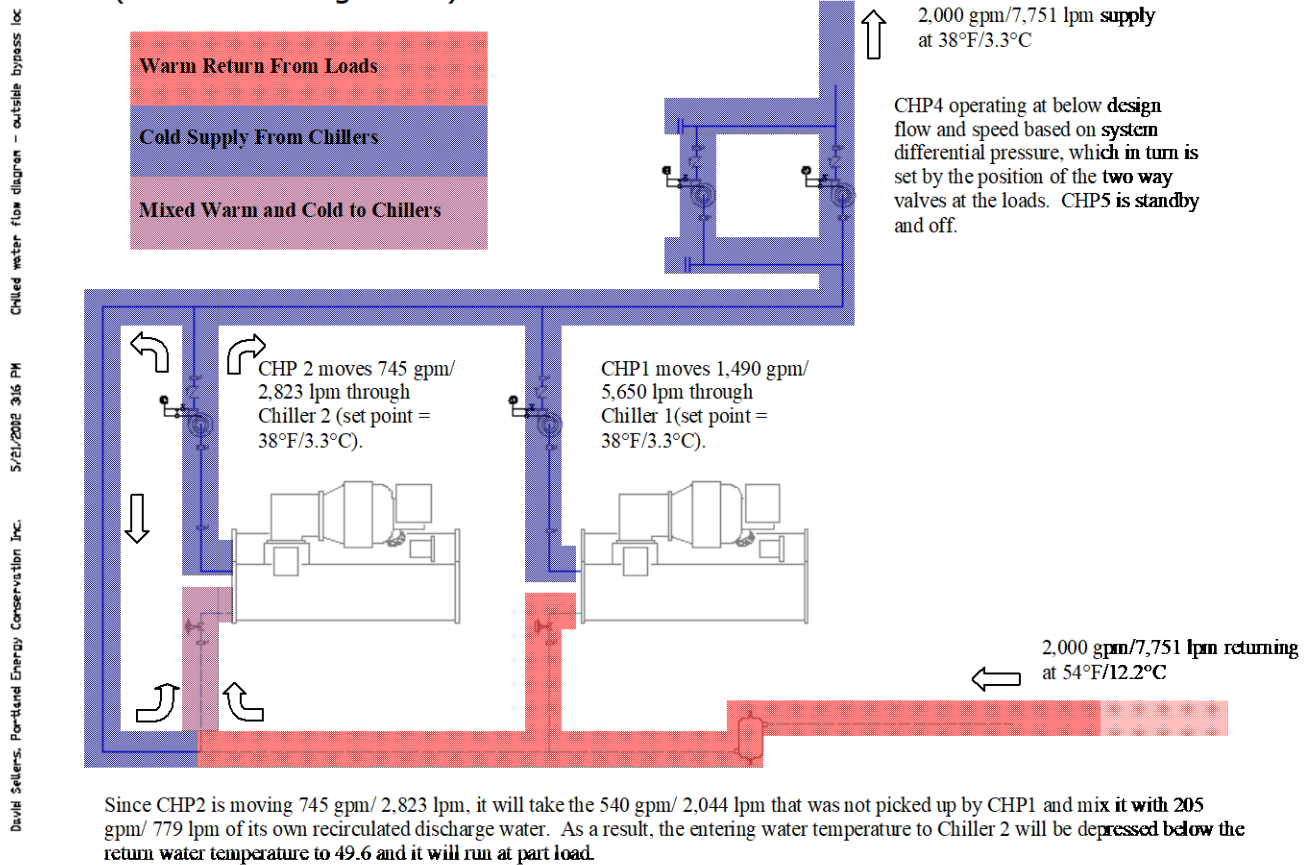


Figure 5 - Chiller plant piping required to meet the operating sequence that may be implied by the JCI energy model report. Notice how the outside the plant bypass location causes the 1,000 ton chiller to see the system return water temperature and thus, run fully loaded. At the same time, the 500 ton chiller sees a depressed entering water temperature because it recirculates some of its own discharge water back to its inlet and runs partially loaded.

Figure 5 illustrates how the plant would operate with an outside the plant bypass. This arrangement would allow the larger chiller to run fully loaded while the smaller chiller absorbed the load swings. This arrangement has the advantage of allowing the larger machine to be optimized for full load operation and the smaller machine to be optimized for part load operation. The wording of the JCI energy model report may imply that this was a premise upon which the model was based and thus could be a significant factor in achieving the proposed energy consumption baseline.

Systems configured with an outside the plant bypass can be a little trickier to operate as compared to systems that are configured with an inside the plant bypass. The biggest problem occurs when the chiller closer to the bypass is operating and a chiller closer to the return connection is started. Starting the second chiller creates a nearly instantaneous load change on the operating chiller because as soon as its pump starts, it diverts a significant portion of the return flow to its evaporator. As a result a significant portion of the supply flow from the chiller that was already operating and

pipled closest to the outside bypass begins to recirculate in the bypass line.⁴ Since the bypass line is usually very short (often, it consists of two tees welded together), the operating chiller sees a sudden and dramatic drop in its entering water temperature. The chiller's control system begins to unload the machine, but typically can't unload the machine as quickly as the actual load on the machine changes.⁵ As a result, evaporator temperatures and pressures plummet. In the best of cases, the machine will trip itself off on an operating limit control because the leaving water temperature has suddenly met the set point. It will then restart after the flows and temperatures stabilize. In and of itself, this is not the end of the world, but it is an additional start operation and start sequence that may not have been necessary. In some cases, the load change happens so quickly that a nuisance low evaporator temperature or pressure safety trip occurs even though the operating limit control has shut down the machine. This is because the evaporator pressure and temperature continues to fall as the impeller spins down with the vanes partially open and driving closed. There are several ways to avoid this problem.

- 1 The most obvious solution is to stop the on line machine prior to starting the second machine. This approach is not an option in a plant where all of the machines are the same size but is viable if the machine that is running is a smaller, low load machine and the machine that is being started exceeds its capacity by a reasonable margin. In the latter case the second machine can be started and will pick up the current load plus some additional load without requiring the operation of the machine that was originally on line. When the load increases to the point where the larger machine can no longer handle it and the increment of excess load is large enough that the smaller machine can turn down to match it and remain in operation, then the second machine can be restarted.
- 2 There is a second option, which minimizes the disruption of plant supply temperature and can be used for plants in which both/all of the machines are the same sizes. The first step in implementing this approach is to unload the machine that is operating.⁶ This can be accomplished manually or automatically through a variety of techniques including:
 - Lowering the set point of the demand limiter.
 - Resetting the set point of the vane controller up to within a couple of degrees of the current return water temperature to the machine.

⁴ If the load exceeds the on line chiller capacity by a large margin, there will generally not be a problem starting the second machine.

⁵ The load change associated with the flow change can happen in a matter of 10-15 seconds. Since the chiller was probably operating at full load or near full load when the other machine was started (that's probably why the other machine started), to track the load change, the vanes would have to be able to move from wide open to nearly closed in the 10-15 seconds it takes for the flow pattern in the chilled water system to change. Most inlet vane actuators take a minute or more to drive from the full load position to the unloaded position.

⁶ For plants with more than two machines and where several machines are already operating, the operating machine that is connected closest to the outside the plant bypass is the machine that needs to be stopped or unloaded because it is the machine that will experience the sudden load change when the next machine is started. Ideally, the set point change should be targeted at adjusting the machine's capacity to match the incremental load difference it will have to handle once the second machine is started.

This will cause the machines capacity control system to unload the machine. After waiting for the machine to unload and perhaps verifying the unloaded condition via observation of the vane position, temperature differential, amperage or kW, the second machine can be started and placed on line. Once the second machine is on line and stable, the set points of the first machine (the one that was running and connected closest to the bypass) can be released/returned to normal. It may take some experimentation to discover the right combination of set points, time delays, and incremental load increases that will allow the second machine to be reliably started and placed on line with out knocking the first machine off line.

Both of these approaches will create a small, short term increase in the supply water temperature leaving the plant. However, as noted previously, the combined thermal inertia of the chilled water system and the loads it serves generally masks the effect and prevents it from producing a significant or detectable impact on the areas served, even in sensitive applications like clean rooms or surgeries. If the system absolutely can not tolerate a supply temperature increase above the chiller leaving water temperature set point, then an inside the plant bypass arrangement should be considered. Or, if an outside the plant bypass arrangement is desirable from a chiller optimization standpoint, then the plant needs to be configured with different sized chillers. The capacity increment between the largest chiller and the other machines should be at least as large as the lowest capacity that the large chiller can be expected to turn down to and operate reliably at. The largest chiller should be piped closest to the bypass and optimized for part load operation. This will allow a smaller machine to be started when the connected load approaches the larger chiller's rated capacity and load up fully while the larger machine throttles back to its minimum operating capacity. Some temporary set point manipulation similar to what was describe in the preceding paragraph may be required for the larger machine when the smaller machine is started to allow it to accommodate the load change that will be introduced when the small machine is placed on line. Regardless of the exact approach and plant configuration selected, diligent operating practice and well written and tested sequencing software will be required to ensuring that the plant never runs with any significant forward flow in the bypass.

Regardless of the exact arrangement of the plant and bypass line, having at least one small chiller coupled with several larger machines (as compared to having all of the machines equally sized) provides better turn down and part load capability. Most centrifugal chillers can be expected to reliably unload and run at 20-25% of their rated capacity.⁷ This means that a 500 ton chiller can probably run against a 100 to 125 ton load and a 1,000 ton chiller can probably run against a 200 to 250 ton load. For a 1,500 ton plant, if two 750 ton chillers are installed, then the plant can probably run against a 150 to 190 ton load with out undue difficulty. If on the other hand, the plant has a 1,000 ton chiller and a 500 ton chiller instead of two 750 ton machines, it will be able to run reliably at a lower load condition. This could be an important factor for many projects, especially those in relatively mild climates with low temperature air systems that result in a need to run the chilled water system until it is

⁷ Some manufacturers indicate that they can operate at lower percentages of full load than this, but as a practical matter, our experience has been that this is difficult to achieve reliably and repeatably in the field with out a lot of operator attention and occasional nuisance shut downs.

in the low to mid 40°F range outdoors due to the low temperature air system requirements.

Check Valves in the Bypass

The intent of the bypass line, from a design stand point, is to decouple the constant volume evaporator pumping circuits from the variable volume distribution system pumping circuits. Properly designed, the bypass will have very little pressure drop, thus the pressure difference between the supply and return mains will be insignificant and, what ever pressure difference does exist is easily overcome by the springs in a typical check valve. In a typical variable flow plant, flow in the bypass line can be bi-directional. When the load on the system is less than the operating capacity in the plant, the flow in the bypass will be from the supply mains to the return mains. Generally, this will be the condition that exists more frequently and will unload the chillers by mixing cold discharge water with warm return water from the system, depressing the entering water temperature to the chiller. As the system load matches then exceeds the on-line capacity, flow through the bypass drops through 0 gpm, then increases in the opposite direction (from the return to the supply). The flow from the returns to the supply tends to elevate the supply temperature to the loads, with the magnitude of the elevation being a direct function of the ratio of the bypass flow to the total supply flow. Some facilities see this condition as totally intolerable because, if left unchecked, the plant will experience a condition called over-flow that can spiral out of control and upset the system thermal balance. Other facilities simply see it as a way to stage chillers and bring another chiller on line when the total system flow will support it and before the elevation in plant leaving water temperature becomes critical. The bottom line is that the bypass allows the variable flow distribution system and constant flow evaporator loops to operate in an independent manner from each other from a flow standpoint while allowing the chiller capacity to vary as required by the loads.

As long as the flow in the distribution system is a function of flow, then this arrangement will work well assuming attention has been paid to matching the ΔT requirements of the loads and the chillers. However, things fall apart when the load/flow relationship is upset. Such upsets can be quite common in a real world operating environment and include things like:

- *Significant loads with three way valves:* Variable flow systems require that most of the loads be equipped with two way valves to create the flow as a function of load relationship. Three-way valves are often installed on small loads at the end of the mains in extensive distribution systems to guarantee some flow for water treatment and response purposes, which is a valid consideration. But for the system to be successful, the flow through these valves must represent a very minor part of the total system flow.
- *Hydraulic short circuits:* Frequently, cross connections are made between the supply and return mains on a system to facilitate flushing during start-up. Some designers install these connections with the intent of having them balanced for a modest flow as an alternative to the technique of using three way valves on small loads at the end of mains to ensure circulation and responsiveness. In either case

these connections are not removed, valved out, or balanced properly when the variable flow system is placed in operation, they will represent a hydraulic short circuit breaking the relationship between flow and load on the system and creating an operational problem at the chiller plant.

- *Loads with discharge temperature set points that are unachievable:* If, in an effort to quickly handle a complaint, the operating team sets the discharge temperature for a system to a value that simply can not be achieved, then the valve serving the load will be driven fully open, regardless of the actual load condition. An example of this is setting an AHUs leaving air temperature controller set point to 42°F on a system that is served with 45°F chilled water. With the control valve driven fully open regardless of the load, the relationship between flow and load is broken and thus an operational problem for the chiller plant is created.

Compromises like those listed in the preceding bullets can drive a variable flow system in to an overflow condition, making it difficult to operate and load chillers and difficult to satisfy the loads. Some designers install a check valve in the bypass with the intent of preventing forward bypass flow and the potential for an overflow condition. Unfortunately, while generally effective at preventing over-flow, the check valve defeats the intended hydraulic decoupling of the loops under some load conditions. If the distribution pumps serving the loads are trying to move more water than the evaporator pumps (for what ever reason) then check valve in the bypass closes and places the distribution pumps (which are usually high head pumps) in series with the evaporator pumps (which are usually low head pumps). Generally, this reduces the predictability of the systems operation and can create some unusual operating conditions.

- *Water flows from the return header to the supply header through inactive evaporator pumps:* Because the distribution pumps are in series with the evaporator pumps, they will force water through inactive pumps unless the pumps are automatically or manually valved out of the system when they are on line. The check valves on the inactive pumps will not prevent this since they are oriented to prevent reverse flow that might be induced by a pump in parallel with them but allow forward flow.
- *The distribution pumps over-power the evaporator pumps:* In some situations, its possible for the distribution pumps to overpower the evaporator pumps due to the mismatch in their operating curves. This is In essence, the high head, high flow (relative to the evaporator pumps) distribution pumps force water through the evaporator pumps. This reduces the slip between the rotors in the evaporator pump motors and the rotating field created by the three phase utility power,⁸

⁸ Bear in mind that the torque generated by the squirrel cage induction motors typically applied on HVAC equipment is created by the difference between the frequency of the rotating field created by the three phase utility power and the speed of the armature in the motor as it is “dragged” by the rotating field. The frequency of the frequency of the rotating field will always be a multiple of the utility distribution frequency, which is 60 cycles per second in the United States. Thus, a motor that is rated 1750 rpm actually has a rotating field of 1800 rpm and generates its shaft power via the torque created by the difference in its speed relative to the rotating field. This difference in speed is termed “slip” and varies with load. A motor that is rated 1750 rpm will operate at that speed at full load. As the load is reduced, the speed of the motor actually will increase and approach (but never reach) the speed of the rotating field since less energy is being taken out of its shaft. This

causing the evaporator pumps to become generators, albeit inefficient ones. As a result, they become restrictions to flow and increase the power required at the motors on the distribution pumps

can be demonstrated in the field by obtaining a strobe-tach, measuring the speed of a pump operating under load, and then observing what happens as the discharge valve on the pump is throttled.

Document3