David Sellers

From: David Sellers
Sent: Sunday, July 15, 2018 6:36 PM
To: 'Lane Burt'; Stroupe, Ryan
Cc: 'abram.mission@flysfo.com'; 'andrew.gustafson@lmco.com'; 'arturo.piceno@ableserve.com'; 'dale.tutaj@dnvgl.com'; 'darin.bernstein@presidio.edu'; 'david.chiu@p2scx.com'; 'Hambalek, Steve'; 'john.rosendo@ucsf.edu'; 'karkis@gene.com'; 'laura.bello@ucsf.edu'; 'mendonsad@saccounty.net'; 'mingc1728@gmail.com'; 'raniel.camacho@flysfo.com'; 'raymond@verdafero.com'; 'rhollingsworth@harvestproperties.net'; 'steve.eng@ableserve.com'; 'theo.pujianto@p2scx.com'; 'toddzilla40@gmail.com'; 'wenhan@stok.com'
Subject: RE: EBCx-Yr13 reminder
Attachments: Right sizing pumps.pdf

Hi Lane,

First off, I noticed the attachment was fairly large and in my experience, something that large may have bounced from Ryan's e-mail account. So a "head's up on that. You may want to resend the e-mail to Ryan with out the attachment if you have not heard from him.

Ryan if you didn't get the e-mail with attachment, let me know and I will upload it via Hightail.com.

Having said that, I have been on a vacation trip in PA this past week, but was working between things to keep the balls up in the air. So, I was trying to get this done before now by using spots of time that showed up here and there. But, doing what I thought I needed to do to really answer the question about the bypass ended up taking a while, thus the delayed response. In any case, what follows are my thoughts on what you sent, mostly focused on explaining the bypass arrangement.

Also, before I forget, I have copied your classmates because there are things in here that may be useful to them also.

And, after finishing this, I realize it really would make a good blog post since it explains a bunch of hydraulic concepts and then shows how, by understanding the system, you have identified potentially significant energy savings opportunities. So, I am going to do some tweaks and put it up on the blog. When I do that, I would like to use the picture you sent and one of your system diagrams as part of the illustration I would acknowledge you as the source of the photo of course, but just acknowledge that the system diagram was from a student and I wanted to use it to illustrate some potential improvements.

So my question is would you be O.K. with that? If not, I can do it other ways, but in particular, I would like to be able to use the photo, so let me know what you think.

Also, I know this building has been really challenging for you as a project because of its age and the complexity of the systems. But in the end, that is a gift I think because if you get to where you understand and can assess something like this, then you will find yourself in a good place to assess other systems that you run into, especially when you consider how many different system configurations there might be out there in the world.

$$N_{Config_{Sys}} = \left[\left(\sum_{All} HVAC\ Engineers \right)^2 \times K_{Climate} \times K_{BuildingType} \right] + \frac{\partial Y_{Earth-Moon}}{\partial Z_{Sun-Saturn}}$$

Where:

$N_{Config_{Sys}}$ = The number of potential HVAC system configurations

$\sum_{All} HVAC\ Engineers$ = The number of HVAC engineers

$K_{Climate}$ = Climate coefficient; adjusts for the climate type at the system location

$K_{BuildingType}$ = Building type coefficient; adjusts for the building type that the system serves

$\frac{\partial Y_{Earth-Moon}}{\partial Z_{Sun-Saturn}}$ = Planetary alignment compensation factor

One of the first buildings I was exposed to as an HVAC technician was of a similar vintage and complexity and the lessons I learned there really provided a very solid foundation for a lot of the things that came up later in my career, especially because they were supplemented by mentoring from some really experienced people, which in your case, is the role Ryan and I are endeavoring to fulfill.

But bottom line, there are some potentially significant opportunities to save pump energy and improve performance associated with what I discuss in the following ramble. But to understand that, I think that the information in the ramble is required. Don't let that overwhelm you. If you submitted/presented your project just as it is in the attachment you sent, I think you would "check all of the boxes" in terms of graduating from the EBCx program; Ryan has final say on that of course, but I think it is generally true.

At a minimum all you would need to do was add the targets that come up out of what follows to your findings list so they are captured (assuming you agree with them). You could then follow up on them via the EBCx Project Review workshops if you wanted to. But at a minimum they would be potentially "on the radar" for your client as long as you got them on the findings list.

So, don't let this overwhelm you or divert you from your final presentation goal. I think you can fine tune what you have and be just fine. What follows is mostly an answer to your question about the bypass configuration, which is actually a fairly complex topic, along with some added opportunities that emerge once you understand that.

<u>System Diagram Improvements – General Ideas</u>

Anyway, I think the reason you find that the bypass is confusing is that it is actually a pretty complex thing when you start to analyze it. As a result, this gets a bit long. But there are a few things that I need to clarify to make sure you fully understand the arrangement and what you are seeing. Some of what I will explain is related to improvements you could make to your system diagram to clarify and add a few things, so I touch on those topics also.

Plus, once you understand what is going on, you will realize there is a significant, relatively low cost opportunity to save energy and improve performance. So, it is important to understand the details so you can make good assessments of the savings potential and costs.

One thing that is not clear from your diagram is which pipe connects to the pump suction and which pipe connects to the pump discharge, which basically tells you the direction of flow in a given portion of the piping. Another thing that might help would be to include any valves associated with the pump, especially a check valve if there is one. Typically, there just about always is a check valve and frequently, if there really is not one, then you found a potential problem.

Other details like the pump service valves and the service valves for the AHU coil that you see in your picture are also important, but second steps when you are developing your drawing. But at a minimum as part of the first step, you should include anything that can cause the flow to change automatically or quasi-automatically, which would include:

- Pumps
- Control valves
- Flow regulation valves
- Check valves
- Valves that need to be specifically positioned manually in certain operating modes.

An example of the latter is that sometimes one pump is put in as a back-up pump for two different systems and there are manual isolation valves on the suction and discharge of the pump that allow the operators to use the pump with either system by manually opening or closing one set or the other. I recall someone in the class showing a system where this was the case, but can't remember who.

Finally, adding performance metrics like pump flows, heads, bhps, efficiencies, motor hps, etc., and similar information for the prime movers and loads would be desirable.

And for the chilled water diagram, adding the central plant configuration just like you did for the HHW diagram would be a really good plan.

System Diagram Improvements - Mixing vs. Diverting Valves

Another detail that stuck me as I was developing the diagram below is that, if you do not know the direction of flow produced by the pump (and the system could kind of work if the pump was pumping either way but one way is better than the other, which I will explain in a minute) then the three way valve would be seen by many as a diverting valve, meaning all of the water comes into one port and then is diverted to one or the other or both ports.

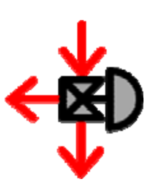

Note that as the valve positions to allow more flow to one of the output ports, it simultaneously reduces flow to the other output port. So, if the valve is properly selected and sized, as it modulates, the total flow will tend to remain constant and the flow out of the two ports will vary proportionally with one port going from 0-100% flow as the other port goes from 100% to 0% flow assuming the differential pressure across it remains constant.

But if the valve is not properly sized (which happens more frequently than you would like to believe), the total flow will vary and the proportionality will not be linear. For instance you might go from 0% flow to 20% flow on one port and the flow at the other pot might only drop from 100% to 95% of what it was, meaning the net flow increased.

In any case, a diverting configuration is definitely a valid three-way valve configuration, but not very common, especially in our industry. The more common arrangement is a mixing valve, where water leaves by one port and comes in by one or the other or both of the other two ports.

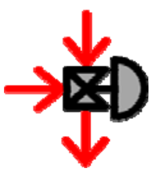

The discussion above about how the valve functions and proper sizing applies to this configuration as well.

For a typical coil, without a pump, if you put the valve ahead of the coil, you would need to use a diverting valve; the supply comes in and you either send it to the coil or divert it around the coil.

If you put the valve after the coil, you would need to use a mixing valve; i.e. the water that bypasses the coil and the water that went through the coil come in and combine into one stream that leaves the valve as the return.

The reason this matters is that since both valves have three ports, you could physically pipe them at either location. But in terms of getting the valve plug to move, especially the plug in a pneumatically actuated valve, the relationship between the balance of forces moving the plug needs to be considered. Specifically, if you did what is called a free body diagram (you may remember that from a mechanics/kinematics class) of the valve plug, you would find three forces come into play.

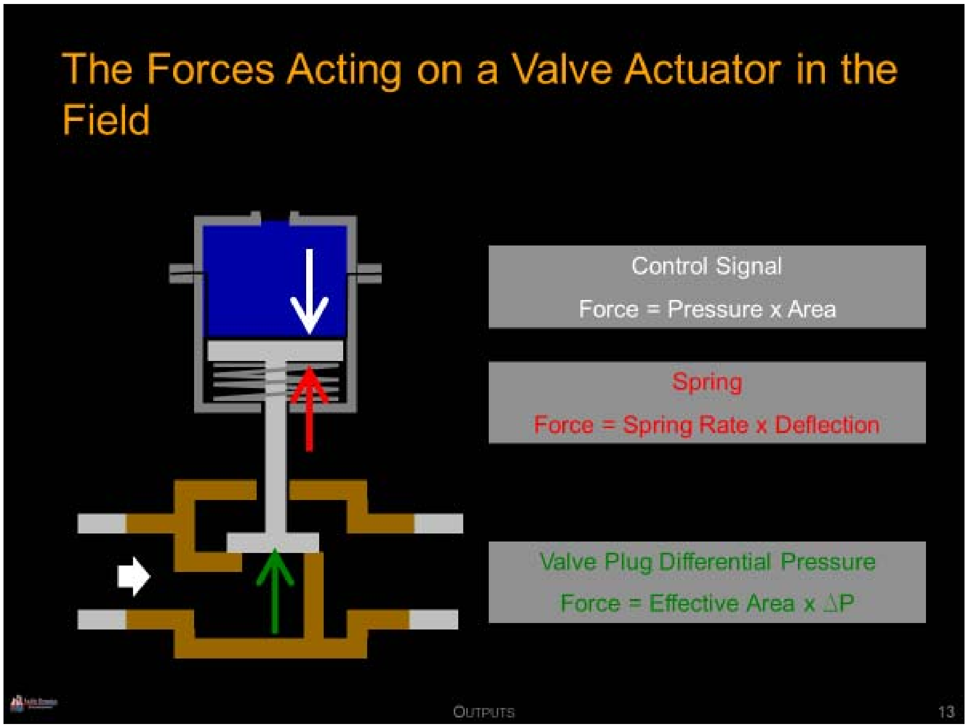

In the example to the left, which is a normally open valve (i.e with no air pressure available, the spring would cause the valve to open),

- The spring acts to open the valve,
- The air pressure on the actuator piston acts to close the valve, and
- The pressure created by water pushing on the valve plug also acts to open the valve.

The picture to the left is a two way valve, but similar concepts apply to a three way valve. And what that means is that if you install a mixing valve so that water only comes in one of the input ports and actually flows out of the other input ports, then pressure drops associated with the flow paths through the valve plug to the input port with the backwards flow that is being used as an output port are reversed from what the designer anticipated.

In turn, that shifts the balance of forces in a way that can mean that the spring is not strong enough to do its job, which can cause the valve to "chatter". In other words, the spring is trying to close the valve, but at some point, the combination of pressure drops through the valve plug result in a force that opposes the spring, vs. being in the same direction as the spring, and the valve plug starts to move the other way. But when it moves the other way, the combination of forces favors the spring again and it starts to close the valve again, which sets up the reverse situation.

If this happens near the position that would have closed the plug, then the valve starts opening and closing rapidly, which sets up rapid water hammer cycles, thus the term "chatter". I have seen this happen with larger valves and the result is more like "valve heated exchange"; very loud and a lot of force involved; in one case, enough force to make a 4 inch pipe move a couple of inches at about 20-30 cycles per minute. Very "eye catching".

So bottom line, it would be good to verify, if you can, if the valve is a mixing or a diverting valve (you would need the submittals or the valve model number, which often can be found on a metal band wrapped around the neck of the valve) and also to verify it is installed in the system as shown in terms of the order of connection principle. I suspect it is, and you may have already verified this. But if not, it would be good to check.

System Diagram Improvements – Balance Valves

Three other things to verify are:

1. Is there a balance valve in the bypass C-D-E? That is important because if there is not, then there is an energy conservation opportunity as well as a performance improvement opportunity. Note that the valves at B or F could serve this function if they were the right type of valve (a valve that could throttle, like a butterfly valve or globe valve or plug valve or ball valve, but not a gate valve).

   If the balance valve is not there, the circuit from B-C-D-E-F is a hydraulic short circuit across the pump at A. What that means is that the loads closest to the pump will "steal" all of the flow. But since the pumps at the loads (the pump at G in the diagram) is only sized for it's load, a lot of the flow and the heating or cooling potential it represents will be "wasted".

   For instance, if the system was a hot water system, then you would bypass a lot of hot water around the loads, warm up the return and minimize the temperature difference across the heat exchanger. So the loads closest to the heat exchanger would likely be just fine but the remote loads would be cold, all while the heat exchanger exhibited the characteristics associated with a lightly loaded piece of heat transfer equipment (a low temperature change relative to the design metric).

   On a variable flow primary/secondary system, a situation like that has even broader, negative implications because it is one of the things that will drive the system into low delta t syndrome.

2. Is there a balance valve in the bypass H-I-J? If there is, that is also a similar opportunity to the one I allude to above, but not quite as big of a deal.

3. Your diagrams show what appears to be a butterfly valve at the end of the system. It is important to verify that, including the size and how far open it is. Given what the analysis below reveals, it is likely that you could completely close that and not loose any functionality. Doing that would also save energy and improve performance since the only function is likely to keep some flow up to improve response times and ensure uniform water quality and the piping configurations at the loads will do that by their nature.

For the drawing I am going to use to explain how the bypasses work, I started with your diagram and I assumed the pump suction was connected to the coil, which would mean the valve is a mixing valve, even though it is on the entering side of the coil. That will become clear when I explain how this all works. I also assumed there would be a check valve and used my triple duty valve symbol because it was faster for me and includes a check valve. But if the check valve is not there, then this is one of those rare instances when the system would probably work just fine without it. I also used the HW system to illustrate the concept because it allowed me to show the heat exchanger and pump and I was not sure of how your chilled water system was piped.

In any case, with those assumptions/modifications, I think your diagram would look like this.

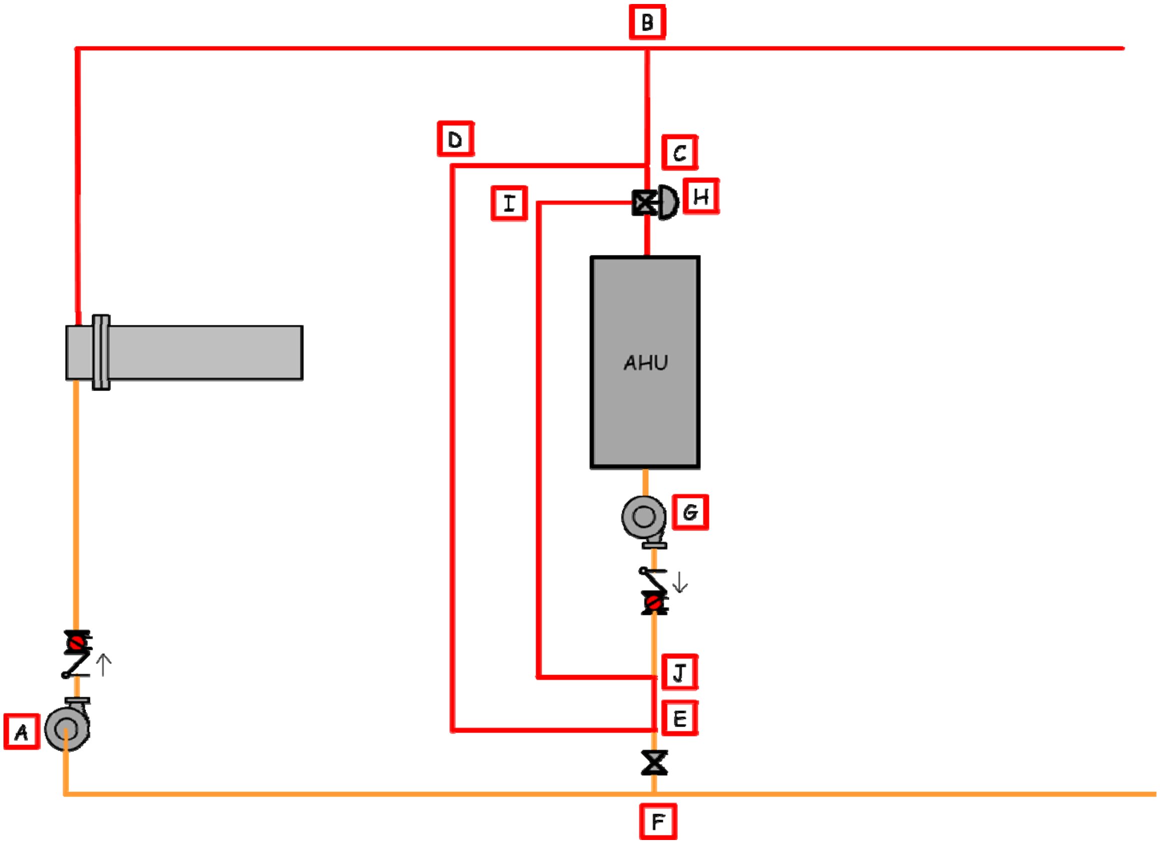

System Diagram Improvements – A Bit More Untangling

But, having looked at it a bit as I developed this, I think there is one more "untangling" step that you could take, which I think will make it a bit easier to understand the connection. Basically, what I am suggesting is as if you dragged/stretched the diagram until the line B-C-D-E-F became straight and rotated the three-way valve and pump, you would get something like this.

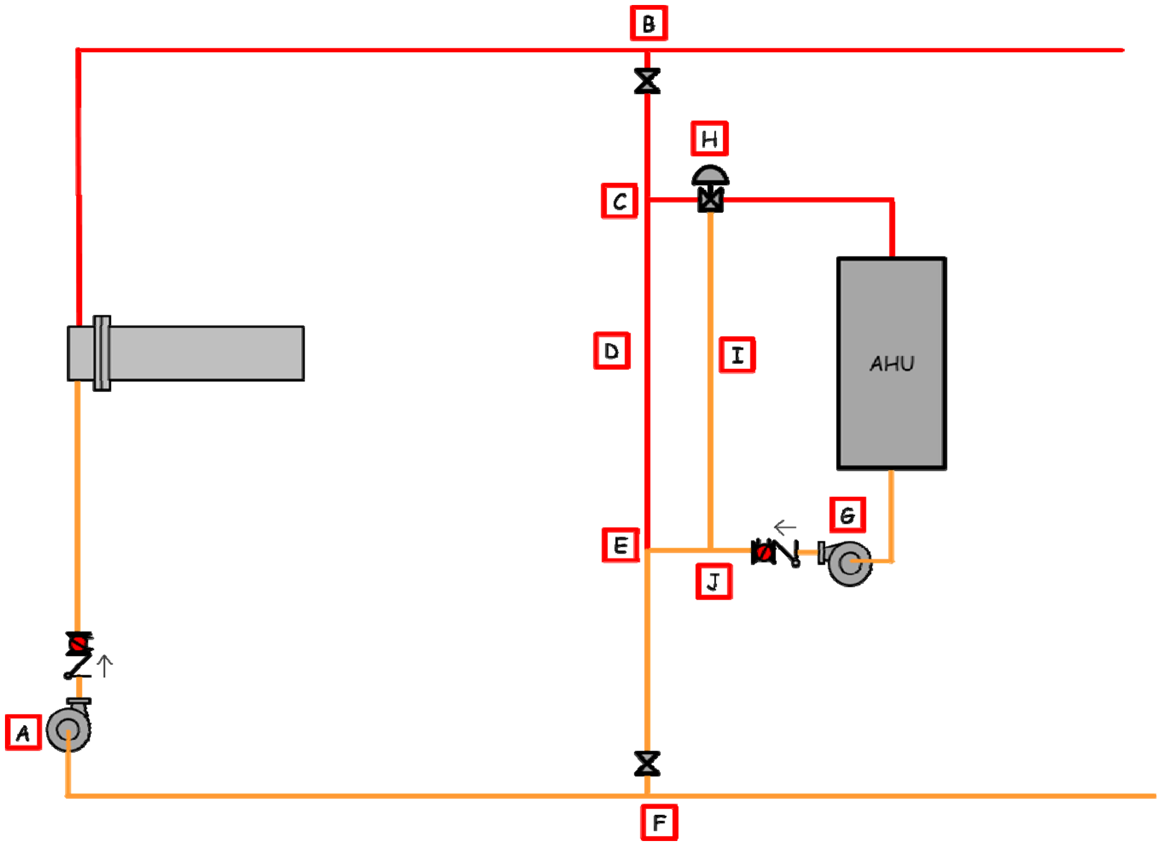

Primary-Secondary Piping Arrangement - Overview

If study this for a bit, you will perhaps discover that it is just a cleverly disguised version of the primary/secondary system in the variable flow chilled water plant we work with in the SketchUp model.

- The pump at A and its heat exchanger would be analogous to the evaporator loop. Specifically, the flow loop would be from A to B to C to D to E to F and back to A. Note that in your system, there would be multiple parallel B to C to D to E to F loops, all in parallel, one per load.

- The pump at G and its three-way valve and coil would be analogous to the distribution loop. Specifically, the flow would be from G to J to I to H and back to G. In your system, each load would have a loop like this.

- The B-C-D-E-F pipe is the decoupling bypass and in your system, you would have multiple decoupling bypasses, one for each loop.

Primary-Secondary Piping Arrangement – Why Use It?

There are a number of reasons a designer would use this configuration.

1. For a heating water application or a chilled water application, the bypass B-C-D-E-F would keep a constant source of water at the plant leaving water temperature available to the load; thus the response to demand from the load would be immediate and predictable, no matter what the operating condition was (low load, high load, start-up, etc.).

2. For a heating water application or a chilled water application, pumping the coil keeps the coil tube velocity at the value associated with the design flow rate, which was likely fully developed turbulent flow.

   a) Fully developed turbulent flow provides better heat transfer than what would happen if the flow dropped to the point where there was a transition to laminar flow in the tubes, which can happen with a coil served by a two-way valve where the flow varies in the tubes.

   b) A consistent flow rate keeps the heat transfer characteristic more consistent and predictable, which can make the control loop easier to tune. This is where the balance valve (or lack of one) in the coil bypass (J to I to H) becomes important. With out the balance valve, the head the pump will see pumping through the coil will likely not be the same as the head it would see pumping only through the bypass. That would make the flow somewhat variable and could defeat one of the intended functions of this piping arrangement. It also makes the interaction between the pump at A and the pump at G more complex.

3. For just about any application, providing a pump for the load makes the pump head for the distribution pump independent of the pump head required to move water for the load. In other words, the pump at A only has to provide the energy to get the water to B, through C, D, E, and F, and back again. The pump serving the load (the pump at G), provides the energy to move the water to and from the load and through the control valve.

   In this particular application, where the load, as I understand it, is only a single coil in an AHU, that probably does not matter that much unless the different coils in the system have significantly different head requirements (for instance one or two required 20 feet of head and the rest only required 5 feet of head).

   If that was the case, then if the system was configured so that the pump at A also provided the energy to move water through the loads, then it would provide at least 15 more feet of head than was required to most of the loads (20 minus 5), probably more because the loads with higher pressure drops probably will need control valves sized for higher pressure drops for the valves to have a reasonable control characteristic.

   A more common application for independent pump at the various loads is a large central plant serving multiple buildings where, due to the nature of the building configuration and size, the piping circuits in a given building will have significantly different pumping head requirements. So, pumps at the central plant, which are analogous to the pump at A, are selected to move water to the hydraulically most remote load and through the interconnecting piping (analogous to the pipe B-C-D-E-F) and back to the central plant. The pumps serving the load, which is likely a number of coils for a large building (analogous to the pump at G), are optimized for the characteristics of the building's coils and piping circuit.

4. For heating water applications, where the coil is a preheat coil (first in the system and intended to be able to safely handle sub-freezing air), this arrangement is one of the approaches use to make the coil "freeze-proof". To some extent, this is leveraging the old adage that "moving water won't freeze"

(but FYI, it will if it is not moving fast enough or it gets cold enough). But the physics behind it are that if you supply a really high flow rate with water that is fairly hot, then for a given load, you will have a very low temperature drop, and thus are less likely to have the water flowing through the coil get cold enough to freeze. Of course, if the pump stops, all bets are off and if you don't shut down the AHU when that happens, you will in fact, freeze the coil.

Primary-Secondary Piping Arrangement – How It Works – Step 1 – The Pump A Circuit

As a next step in understanding this, consider what would happen if you only ran the pump at A. Bear in mind that I have assumed the three way valve is a mixing valve, which means it is designed for water to come in on one port and go out on the other two ports.

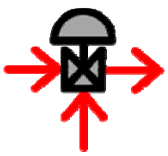

So, depending on the valve position, water could:

- Come in the left port and go out the right port with no flow in from the bottom port,
- Come in the bottom port and go out the right port with no flow in from the left port, or
- Come in both the left port and bottom port and go out the right port if the valve was in a modulated position.

I should also point out that in this discussion, when I say left port, right port, bottom port, it is in the context of the symbol for the valve on the system diagram. The physical ports on the valve may or may not match the symbol, depending on the type of valve and its design. For globe type mixing valves …

… usually the output port is on either the left or right side and the bottom port will just about always be an input port. But if the valve is a three way butterfly valve that is actually a pipe tee with a butter fly valves with interconnecting linkages on two of the three connections (called a butterfly tee) …

… then the output could be the branch or either one of the run connections.

Given all of that, with only the pump at A running, water would flow from F to A to B and then split to the various loads. How much went to each load would depend on the position of the valve at the loads, the position of any balance valves in the bypass and the pressure drop due to flow that existed in the network for each segment.

But, upon reaching a given load, the flow has a number of options.

1. If the three way valve at the load had the bottom port fully open and the right port fully closed, the water would flow from B to C and then would have to options:

   a) Flow through the decoupling bypass; C to D to E, or

   b) C to H to I to J to E.

   How much went either way would simply depend on the pressure drops of the two parallel paths relative to each other. If there were balance valves in either of them, then those would come into play.

   But no matter which path was taken, when the flow streams reached E, they would combine, flow to F, combine with the flow from the other loads, and return to the pump at A.

   Note also that if this condition were to exist, there would be flow in the wrong direction with regard to the mixing valve; i.e. the bottom port would have flow leaving it when it was designed to have flow entering it. Thus, it could chatter, depending on the valve design and the amount of flow and rated pressure drop.

2. If the three way valve at the load had the bottom port closed and the right port fully open, then the water would also have two options once it got to C:

   a) Flow through the decoupling bypass; C to D to E, or

   b) Flow from C to H to G to J to E.

   Note that in this configuration, the pump at A will open the check valve on the pump at G because the pressure ahead of the check valve is higher than the pressure leaving the check valve. As a result, the pump at A would be pushing water through the pump at G if the pump was not running. As a result, the pump might actually be spinning, perhaps, even backwards due to the dynamics of flow through an impeller that is not being rotated in the intended direction by the motor. It probably would not be spinning fast though. And the motor and air moved through it by the fan and coming out the air vents would be cool.

   The fact that the pump at A can spin the pump at G would be something that would need to be taken into account in terms of how the control system started the pump, especially if it had a VFD on it since starting against a spinning load, especially a reverse spinning load can set up a number of problems, especially with a VFD. I will not get into that right now because this is already long and our focus is on explaining the bypass. But I am happy to explain it if you want to know more; just let me know. This is something we get into in the VSD class, which is currently scheduled for December 12th.

   As was the case before, how much water went either way would simply depend on the pressure drops of the two parallel paths relative to each other. If there were balance valves in either of them, then those would come into play.

   But no matter which path was taken, when the flow streams reached E, they would combine, flow to F, combine with the flow from the other loads, and return to the pump at A, just as was the case before.

3. If the three way valve was modulated with both the bottom and right port open to the left port, then there would actually be three parallel paths. I suspect that you can figure that out with out me going into the details and am going to assume that, but if you have a question or are not sure, please ask.

<u>Primary-Secondary Piping Arrangement – How It Works – Step 1 – The Pump G Circuit</u>

Now consider what would happen if you only ran the pump at G with out running the pump at A. Again, there would be a number of optional flow paths that depend on the position of the three-way valve.

1. If the valve had the left port closed and the bottom port open, the flow would be from G to J to I to H and back to G. Note that now, the flow is in the correct direction through the bottom port, assuming the valve is a mixing valve.

2. If the valve had the left port open and the bottom port closed, then flow would be from G to J to E to D to C to H and back to G.

3. If the valve was modulated so that both the left and bottom ports were open, there would be flow through both of the parallel paths J-I-H and J-E-D-C-H, with the amount of flow in each path a function of the flow vs. pressure drop characteristics of the path and the valve position and characteristics.

In case 2 and 3, there in theory could be some flow through the pump at A. But for that to happen the pressure drop due to flow from E to D to C would be what drove it and it would need to be enough pressure drop to open the check valve on pump A in addition to overcoming any resistance due to flow created by the flow in that path. Unless the bypass was significantly undersized or restricted by a valve between E and C, it would likely be inconsequential.

Running At Full Load – How it Works

Now let's look at what would happen if both pumps were running with the load at full, design capacity. We will assume that each load has been balanced for its design flow using the valve where the load connects to the return main and that the pump at A is a constant volume pump delivering the total design flow rate to all of the balanced loads. Note that conservation of mass says that if 100 gpm leaves the central system at point B and flows into the load, then 100 gpm must also leave the load at point F and return to the central system.

And we will assume that the load in the illustration is a 100 gpm load with a properly sized control valve (near linear characteristic) and a pump that has been balanced to the design flow rate.

I have also attempted to use color as an indication of temperature with red being hot and orange being cold; the lighter the shade of orange, the cooler the water is. A large change in temperature would be associated with a large load and a small change in temperature would be associated with a small load.

If the load in the illustration was at full capacity, things would look like this.

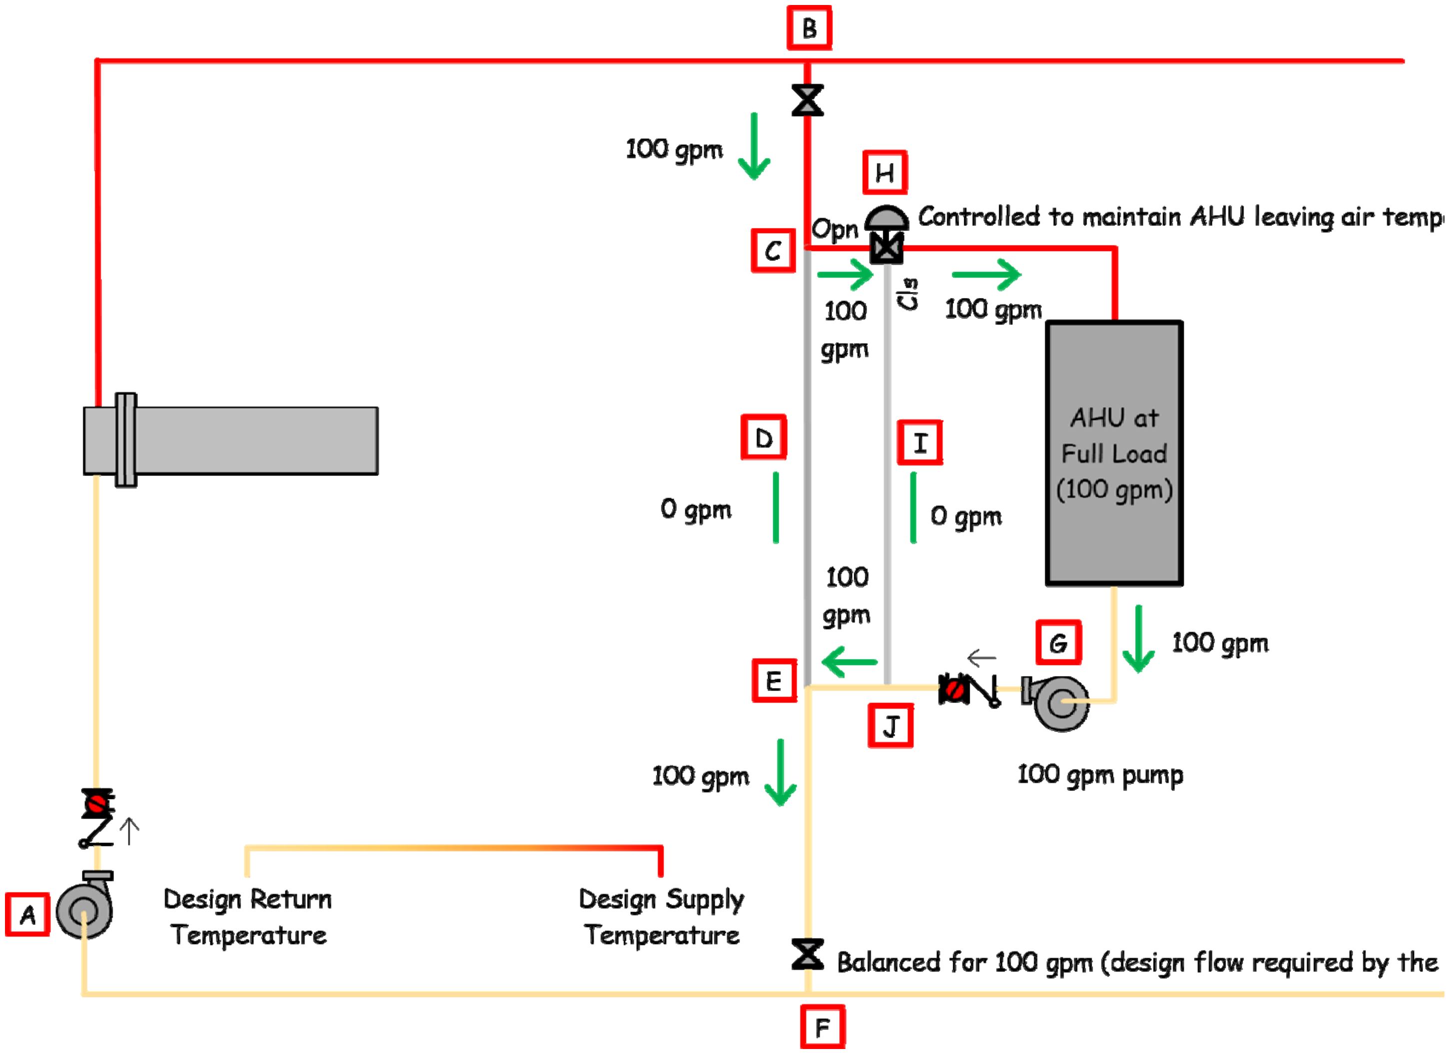

Note the following:

1. Since the load is at the design condition, and because we have assumed everything is properly sized and functioning properly, the control process associated with the AHU has positioned the control valve so that no water is being recirculated to the load. Because the system is at full load, to achieve the design leaving air temperature, the control process needs to position the valve so that it only allows flow from the central system into the loop associated with the load.

   In other words, we are at the design condition and the coil has been selected to deliver the design leaving air temperature when supplied with the design supply water temperature. If the load was less than full load, the design flow rate (which in this piping arrangement, is always circulated by the pump at G) would have done more heating than necessary if served with the design supply water temperature. As a result, the temperature leaving the coil would start to rise above set point and the control process, upon seeing this, if properly designed and tuned, would modulate the control valve to partially close the left port and partially open the bottom port.

   That would result in some of the cooler water leaving the load being mixed with the supply water, which would lower the supply temperature to the coil relative to the supply temperature from the central hot water system. A lower supply temperature to the coil means there is a smaller temperature difference between the water inside the tubes and the air outside the tubes. Since it is the temperature difference that that drives the heat transfer, the capacity of the coil would be reduced and when things are properly sized, adjusted, and tuned, the result is for the control process to modulate the valve until everything comes into balance; i.e. the set point for the control process is equal to the supply air temperature generated by the coil.

We will further explore the part load condition next. For now, I want to make a few more points about what is going on at full load.

2. All of the water that is delivered to the load from the central loop at point B (the loop circulated by the pump at A) is returned from the load to the central loop at point F after passing through the load. From a pumping energy standpoint, the pump at A provided the energy to get the water from E to F to A to B to C. The pump at G provided the energy to move the water from C to H to G to J to E.

3. As a result of item 2, there are two pieces of pipe that have no flow in them (the gray pipes in the diagram). This is simply a result of conservation of mass. Points H and J are "nodes" in the system; i.e. points where a number of streams converge or diverge and the mass and energy in must equal the mass and energy out. So, for example, if 100 gpm is delivered by pump A to the left port of the control valve and pump G is removing 100 gpm from the right port and circulating it through the coil, then there is no mass taken from the bottom port of the valve. A similar condition prevails at point J.

4. The temperature of the water delivered to the load is the same as the temperature of the water leaving the heat exchanger (other than for losses through the pipe insulation, which would be relatively minor).

5. The temperature drop through the load is the design temperature drop. If this was the only load on the system, or if all of the loads were identical, that temperature drop at design would match the temperature rise at design conditions provided by the heat exchanger.

<u>System vs. Load Return Water Temperatures</u>

The last observation merits some additional discussion, especially if you contemplate making changes to the system and/or are working with an existing system and trying to understand if it is working per the design intent or not.

One of the things that needs to be considered when you are designing a with multiple loads, each of which may have a different temperature rise characteristic, is that the design temperature rise that needs to be provided by the "prime mover" in the system (in our case, the heat exchanger, but for a chilled water system, a chiller or for a heating system served by boilers, the boilers) would need to match the design *<u>system</u>* temperature drop associated with the loads. In other words, the heat exchanger on a design day, would need to heat the water from the temperature in the pipe just after F to the design supply temperature, i.e. the temperature in the pipe just ahead of B (where F and B are the connection to the first load; with insulated piping, those temperatures are virtually identical to the temperatures at the inlet and outlet connections of the heat exchanger).

The *system* flow and temperature difference is just an amalgamation of the flows and temperature differences associated with the various loads, and if you go through the math, which is based on conservation of mass and energy and the steady flow energy equation, you would discover that you could calculate it using a relationship that looks like this, which assumes the supply temperature to all of the loads is the same.

$$t_{Return_{System}} = \left[\frac{\left(t_{Return_{Load1}} \times Q_{Load1}\right) + \left(t_{Return_{Load2}} \times Q_{Load2}\right) + \left(t_{Return_{Load3}} \times Q_{Load3}\right) + \ldots + \left(t_{Return_{Loadn}} \times Q_{Loadn}\right)}{Q_{System}}\right]$$

Where:

$t_{Return_{System}}$	=	System return temperature in consistent units
Q_{System}	=	System flow in consistent units
$t_{Return_{Load1}}$	=	Return temperature from the first load in the system in consistent units
Q_{Load1}	=	Flow rate associated with the first load in the system in consistent units
$t_{Return_{Load2}}$	=	Return temperature from the second load in the system in consistent units
Q_{Load2}	=	Flow rate associated with the second load in the system in consistent units
$t_{Return_{Load3}}$	=	Return temperature from the third load in the system in consistent units
Q_{Load3}	=	Flow rate associated with the third load in the system in consistent units
$t_{Return_{Loadn}}$	=	Return temperature from the n^{th} load in the system in consistent units
Q_{Loadn}	=	Flow rate associated with the n^{th} load in the system in consistent units

The term "n^{th}" is a way to mathematically refer to an un-specific/general member of a long series of of items, usually the last member in the series. For instance, if this equation was being applied to a system with 8 loads in it, then n would equal 8 and the numerator of the equation would have 8 of the $\left(t_{Return_{Load}} \times Q_{Load}\right)$ expressions in it.

(Incidentally, the derivation of that is essentially the same as the derivation of the equations we used to figure out mixed air conditions and if you want to see the details, I derive those relationships in a blog post titled Economizers–The Physics of a Mixed Air Plenum).

The temperature change that you get through a given heat transfer element is very much a function of the physical characteristics of the element (for coils, the number and size of the tubes, the length of the tubes, how many passes the water makes through the coil, the number of circuits, the fin spacing, the number of rows, the materials of construction, etc.) and the physical characteristics of the media on either side of the heat transfer surface (for a coil the entering air and water temperature, the flow rates, etc.). And while you can target a certain temperature rise for all of the loads, it may not be possible to achieve it.

For instance, it is much more challenging to get a large waterside temperature change from a coil in a typical fan coil unit (like the unit you would find in your hotel room) relative to the coil in the make up air system serving your hotel room simply because the coil in the make up air system likely has longer tubes because the AHU it is in is much wider than the fan coil unit in the hotel room, has more rows because it is handling 100% outdoor air vs. the fan coil unit in your hotel room, which is handling air that has already been preheated/precooled/dehumidified, etc..

A very common problem that we run into in existing buildings all of the time, especially buildings with a mix of fan coil loads and large air handling unit loads like a hotel or hospital or even an office building that uses induction units or fan coil units to handle the perimeter load, is that the designer assumed all of the loads would generate the same temperature drop (or rise for a cooling application) and selected the central plant equipment (in our case, the heat exchanger, but in the case of a chilled water plant, the chillers) to produce a commensurate temperature drop (or rise for a cooling application). But since the real loads may not create that condition in terms of a system temperature rise (or drop), there is a miss-match which for a variable flow primary secondary chilled water system, leads to low delta t syndrome, which is the Achilles heel of that particular design concept and is a kind of "death spiral" in terms of making the plant work efficiently.

For a heating system served by a variable flow primary secondary plant, a similar, less discussed condition can happen. But, since the temperature differences (both water to air and across the loads) are usually bigger for heating systems, and because heating systems only are doing sensible energy transfer, they are generally more forgiving. So, when the condition occurs, it is usually not a "death spiral" in terms of meeting the loads. But if you recognize it and can correct it, you will probably improve the over-all system efficiency by optimizing the pumping energy.

For the system in this example, since it is constant volume and a heating system, if the designer did not take the potentially different load temperature drop characteristics into account, it is probably not a huge concern simply because of the constant volume, steam heat exchanger based system design.

Running At 50% Load – How it Works

Now let's look at what would happen if both pumps were running with the load at 50% design capacity. And, the same assumptions I made for the full load condition also will apply; i.e. each load has been balanced for its design flow, the design load in the illustration is a 100 gpm load with a properly sized control valve, etc.

Given those assumptions, if the load in the illustration was at 50% capacity, things would look like this.

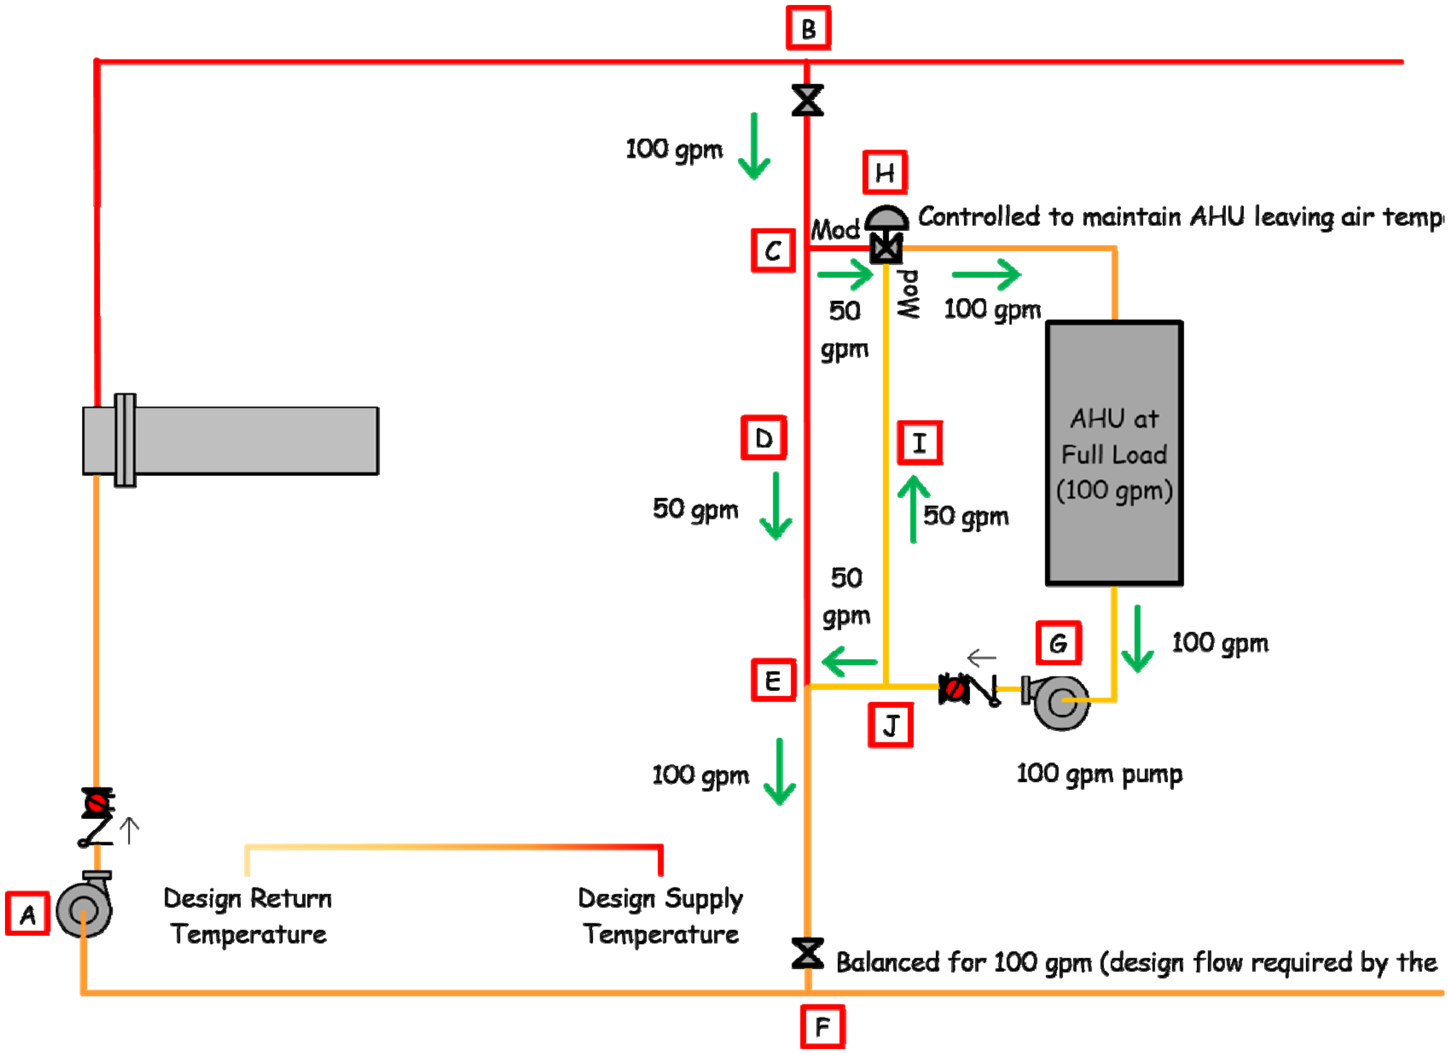

Note the following:

1. When the load started to drop, the design flow rate at the design temperature provided more capacity than was necessary, as discussed previously. As a result, are well designed and tuned control process modulated the control valve, reducing incoming flow from the central system on the left port, and allowing some return flow from the loop served by the pump at G to enter on the bottom port until the system became steady state again at the reduced load condition (supply temperature = set point with a stable leaving condition).

2. Focusing on the pump G loop;

   a. Since the load condition is 50% of design, only 50% of the design flow at the design supply water temperature is required, as can be seen from the amount of flow entering the pump G loop through the left port of the mixing valve.

   b. Since the pump G loop is a constant volume loop by design, with only 50% of the design flow entering from the left port of the mixing valve, conservation of mass requires that the other 50% of the design flow be recirculated via the bottom port of the mixing valve, as can be seen in the illustration.

c. Conservation of mass and energy also requires that at node J, since only 50% of pump G's flow is leaving the node to go to the mixing valve, the other 50% of the flow must leave the loop and return to the loop served by pump A.

d. Since the pump G loop is a constant volume loop by design, conservation of mass and energy require that if the flow rate is held constant, then the temperature change across a heat transfer element will vary directly with the load. The mathematical statement of this is as follows, which comes from the steady flow energy equation.

$$Q_{Btu/Hr} = 500 \times Flow_{gpm} \times \left(t_{Entering,°F} - t_{Leaving,°F}\right)$$

Where:

$Q_{Btu/Hr}$ = Load in Btu/hr

500 = Units conversion constant, good for water between 30 and 200°F

$Flow_{gpm}$ = Flow through the heat exchanger in gallons per minute

$t_{Entering,°F}$ = Temperature entering the heat exchanger in °F

$t_{Leaving,°F}$ = Temperature leaaving the heat exchanger in °F

As you can see, if the flow is held constant and the load reduced by 50%, then the temperature change across the heat transfer element will be 50% of what it was at full load.

e. Since the water stream leaving the mixing valve represents the mixture of water at the design supply temperature with water coming back from the load at 50% capacity, the supply temperature to the load will be reduced. In fact, numerically, it will be 50% of the way between the temperatures of the two streams that are being mixed, in this case, the supply temperature from the system and the return temperature from the load.

So, for instance, if the supply temperature from the heat exchanger was 180F and the return temperature from the load 150F, the blending of 50% water at 180F with 50% water at 150F would result in 100% water at 165F. This can be calculated by the following relationship, which is very similar to the relationship for figuring out the blended return water temperature from multiple loads and comes from the same assumptions and derivations.

$$\left(Q_{1\%} \times t_1\right) + \left(Q_{2\%} \times t_2\right) = \left(Q_{Total\%} \times t_{Mix}\right)$$

Where:

$Q_{1\%}$ = Entering fluid stream 1 flow as a percentage of the total leaving the node

t_1 = Entering fluid stream 1 temperature in consistent units

$Q_{2\%}$ = Entering fluid stream 2 flow as a percentage of the total leaving the node

t_2 = Entering fluid stream 2 temperature in consistent units

$Q_{Total\%}$ = Leaving fluid stream flow as a percentage; i.e. 100% for one leaving fluid stream

t_{Mix} = Leaving fluid stream temperature in consistent units

The blog post I mentioned above shows the derivation if you are interested and there is a spreadsheet tool available on our Cx resources web site that has several of the useful mixing relationships built into it that you can download if you want to. It is set up for mixing air, but it applies to any place where there is a node that mixes to fluid streams.

f. The lower supply temperature to the coil created by the blended water streams at the mixing valves lowers the temperature difference between it and the air stream. Since the temperature difference is what drives the heat transfer, the capacity of the coil is reduced, even though the flow rate did not change.

g. The exact temperatures associated with steady state operation at 50% capacity will ultimately be set by the physical characteristics of the coil (size, fluid velocities, depth, circuiting, etc.). As a side note, if you needed to figure that out, you could make a model of the coil at design conditions in a software tool like Greenheck's free coil selection program, and then lock down the physical characteristics (rows, fins per inch, etc.) and then change the other variables until you achieved 50% of the design capacity at the design water flow rate. For me, when I was first learning about this, doing that exercise was worth the effort simply because it reinforced my understanding and got me more comfortable with things, so in terms of self-study, you may want to consider giving it a shot.

3. Focusing on the pump A loop:

    a. Since only 50 gpm of the flow delivered to node C was taken by the pump G loop, and the flow through pipe B-C-D-E-F has been balanced for the design flow rate of 100 gpm, conservation of mass and energy requires that the remaining 50 gpm move through the pipe from C to D to E. This water will be at the full supply water temperature.

    b. At node E, the 50 gpm of water that did not flow into the pump G loop at node C mixes with the 50 gpm of water from the pump G loop that did not recirculate to the mixing valve. The water coming from the pump G loop will be at the return water temperature leaving the coil.

    c. Conservation of mass and energy requires that the total flow leaving node E be equal to the sum of the flows entering it; thus 50 gpm plus 50 gpm = 100 gpm. And it also requires that the leaving water temperature from node E is between the supply water temperature entering it from the supply side of the pump A loop and the return water temperature from the coil, entering it from the pump G loop.

    d. Since the pump A loop is a constant volume loop, as it unloads, conservation of mass and energy require that the temperature difference across the heat exchanger varies in direct proportion to the load, as shown by the water slide load equation above. Thus the return water temperature to the heat exchanger will be higher than it was under full load operation. In fact, if the system truly is constant volume and the load truly is 50% of the design, then if the design supply temperature was 180F with a 40F temperature drop across the loads, then at 50% load, with a 180F supply water temperature, the return water temperature will be 160F (180F – (50%*40))

Other Part Load Conditions

The concepts I discussed in the preceding section would also apply at any other part load condition. So hopefully, at this point, you have the information you need to understand how the load connection works. That all, of course, assumes the assumptions I made regarding how the pump is piped, etc. are correct. But even if you discover something different when you verify the connections, the same analysis technique and principles are the approach you would use to figure out how it works. So my hope is that by going through the example in detail, I have illustrated the approach you would use and shown you how to think about this sort of thing. That is the real trick to all of this; getting familiar with fundamental physical principles and then learning how to use them to think about the different system configurations we run into.

Things to Explore – Overview

So, having gotten you this far (I hope), I wanted to point out a few opportunities that you might want to explore.

You may or may not have time to get into them for your final presentation; I certainly am not suggesting that you need to do that.

But I am suggesting that you capture them on your findings list, assuming that you agree with them and they make sense to you. That way, ideally, someone can follow up on them, maybe even you if you continue to work with this facility, either due to your relationship with it as a client or as a way to keep on learning via the ongoing EBCx project review workshops that you will be invited to as a graduate of the EBCx class.

Bear in mind that this is kind of first pass/off the cuff, so you should make sure you agree and are comfortable with executing anything I suggest. And, I am happy to discuss any of this with you, but I suspect a GoToMeeting would be the way to go for that.

Things to Explore – System Diagram

Just about all of this hinges on the system diagram being accurate. So, if you are uncertain about any of the assumptions I made or any of the questions I asked, verifying those things would be a good first step.

Once you know that is right, you can more on to the other stuff. Initially, you could use a sampling approach to this; i.e. pick one load to start out with and check it out pretty thoroughly. Then, based on what you find out, pick another at random and see if it matches the pattern, then maybe another. For the HW system, checking one unit represents a 25% sample. If you check a second unit and find similar results, then it is not unreasonable to assume they apply to the other units you have not looked at in detail. Obviously, the more units you check out, the more certain you are. But you have to balance time with certainty in some instances, so in Cx we sometimes use a sampling approach and engineering judgement to apply those results to other systems we have not looked at as closely.

I have tried to list these ideas in the order that they should be explored; i.e. they build off of each other, so bear that in mind. Also, remember that all I really know is what you have shared in class tempered by past experience and having seen similar things before. That is why the system diagram verification is a really important first step.

Things to Explore – Pump Performance and Coil Bypass Balancing

As you know, pumps can be significant energy users and are often oversized. The most obvious indicator of that is a throttled valve. I am not sure if you have seen those, but if the valves are not throttled, the second obvious indicator is an estimated pump head requirement that is much lower than the rated head. We did that exercise in one of the labs at the UCSF campus. If you need a memory jogger, then there is an article I wrote that I have attached that goes into how you do it.

If there is not a balance valve in the line that bypasses the coil (the pipe from J to I to the bottom port of the mixing valve) then the pressure drop at design flow through the bypass could be significantly less than through the coil. That means that even if the pump is about right pumping through the coil, it will over-pump anytime the bypass valve is modulated. So, verifying pump performance in full flow to the coil position as well as full flow to the bypass position is important.

If the pump flow in bypass in excess of design, you should throttle the bypass to get to the design value. If there is no balance valve there, then the data you gather will help you understand the implications and may justify the addition of the valve if the flow is significantly high and/or there are a lot of part load hours (hours in full or partial bypass).

Also, simply shutting the pumps down when they coil is not doing anything would save a lot of energy I suspect if that is not already being done.

All of these things might help justify that new energy management system that the operator is hoping for.

Things to Explore – Load Decoupling Bypass Balancing

The load decoupling bypass is basically a hydraulic short circuit, as you can see from the untangled system diagram. Thus in terms of performance and energy, it is important to make sure the line is balanced so that only the design flow required by the load it is associated with can pass through it. I can't tell from the picture if there is a balancing valve in the line or if the valves that are there could be used for balancing. Like I said, if they are the right type of valve, they could be. The advantage of a balancing valve is that there frequently are pressure taps on it that can be used in conjunction with a chart for the valve to determine the flow for any given valve position.

If there is no balancing valve and if the other valves are not suitable for balancing (i.e. gate valves), then the first problem/challenge would be getting a balancing valve installed. If you can come up with the flow rates, there may be enough pump energy being wasted to actually justify the cost, especially if you combine it with the potential for poor flow distribution to the loads downstream and related comfort and other performance issues.

Lacking a balance valve, there are number of options for measuring and setting up the flow.

1. Borrow an ultrasonic flow meter from Ryan and use it to measure the flow while you set it up with one of the valves. This could be the most labor intensive and invasive approach because you would have to cut the insulation and then patch it back. But it may be the most accurate and direct approach depending on the information available about the pumps and the trim on the pumps.

2. Use a pump test and temperatures to set up the flow. This is how we would have done it in the olden days lacking a balance valve and the ultrasonic flow meter technology. The idea is to place the three-way valve in the position that closes off the coil bypass so that you know all the water going to the coil is coming from the central system. If this is in fact true, the temperature leaving the valve/entering the coil should be pretty close to the central system supply temperature. If it isn't then the three-way valve bypass port could be leaking by or the flow in the decoupling bypass could be lower than the design flow.

    a. The temperature of the two bypass lines is a really useful indication of what is going on there, so even if there are not thermometers in the lines, cutting a little hole in the insulation so you can take a surface temperature can be really useful.

        - For the decoupling bypass (C-D-E) the temperature just about has to be either the supply temperature from the central system or the return temperature from the coil. If there really was no flow, then it would kind of float someplace in between but would also be influenced by conduction a bit. But literally having no flow in the line in a field situation would probably be a transient, short duration event. So bottom line the temperature is a pretty good indicator of there being flow or not.
        - That means, that if you are seeing a supply temperature out of the mixing valve that is lower than the central system supply temperature and the surface temperature in the decoupling bypass is at or near the central system supply temperature, and you have the three way valve commanded to fully close its bypass port (the bottom port in the system diagram), then that port just about has to be leaking water.

        - In contrast, if the bypass temperature is closer to the leaving water temperature from the coil than not, that says you are reversing flow through the bypass, which either means the pump is over-pumping or that the flow across the decoupling bypass is less than the design flow.

    b. As a next step, use a pump test to figure out what the current flow rate is. If it is high, then for the time being, throttle the pump to get the design flow rate. And, if it is high, you have found an opportunity.

    c. Once you think you have the flow established at design, check the decoupling bypass temperature. If the pump was originally flowing more water than design and the decoupling bypass was at the coil return water temperature but now is at the supply water temperature, it is pretty likely that the recirculation was via the decoupling bypass.

If the decoupling bypass temperature is at or near the central system supply temperature and you still have a lower temperature supplied to the coil, then it is likely that the three-way valve is leaking by, so make note of that. N

d. Now, repeat the pump test with the three-way valve in full bypass. When the valve is in full bypass, the decoupling bypass temperature should get up to the central system supply temperature for sure since, in theory, you are not taking any of that flow into the pump G loop. And, the coil supply temperature should drift down and equalize with the coil return water temperature over time since you are not adding any heat from the supply loop. If that does not eventually happen, then the valve is probably leaking by on the port that connects to the supply loop.

   In any case, document the flow the pump now produces and if it is in excess of design and there is a balancing valve in the bypass, then throttle it to get the design flow. If there is not a balancing valve in the bypass, note the flow since you can use that to assess the energy savings associated with adding one.

e. Go back to full flow through the coil. If the decoupling bypass temperature is close to the central system supply temperature, then there is more than the design flow through the decoupling bypass. So, you can gradually throttle the flow with the valve at the return or supply connection assuming they can be use for throttling) until you see the bypass temperature drop and start to approach the leaving water temperature from the coil.

f. When that happens, you know you have started to recirculate water in the decoupling bypass. So, open the balancing valve back up slightly until you see the temperature start to come back up. At that point, you know that the decoupling bypass is set up for a flow rate that is at or near the design flow rate, assuming you balanced the pump for that flow rate

Once you have done those things, you probably have the system as well set up as you can get it given the existing equipment. But there are two additional measures you could implement that have some cost associated with them but which would also save additional energy and improve performance.

Convert the Central Loop to Variable Flow

Frequently, it is quite viable to take systems of this vintage that are configured the way they are and convert the central plant loop to variable flow (the loop served by the pump at A). For your building, this is probably more viable for the HHW loop than the chiller loop since steam heat exchangers can generally tolerate a wide flow variation where-as that is less the case for chillers.

For the chilled water system, how big a constraint that is relates to the configuration of the central plant and the nature of the chillers.

In any case, the "trick" for making the conversion is to install a two way control valve in the decoupling bypass that modulates to control the leaving water temperature from the bypass at the design leaving water temperature for the load. Here is what that would look like.

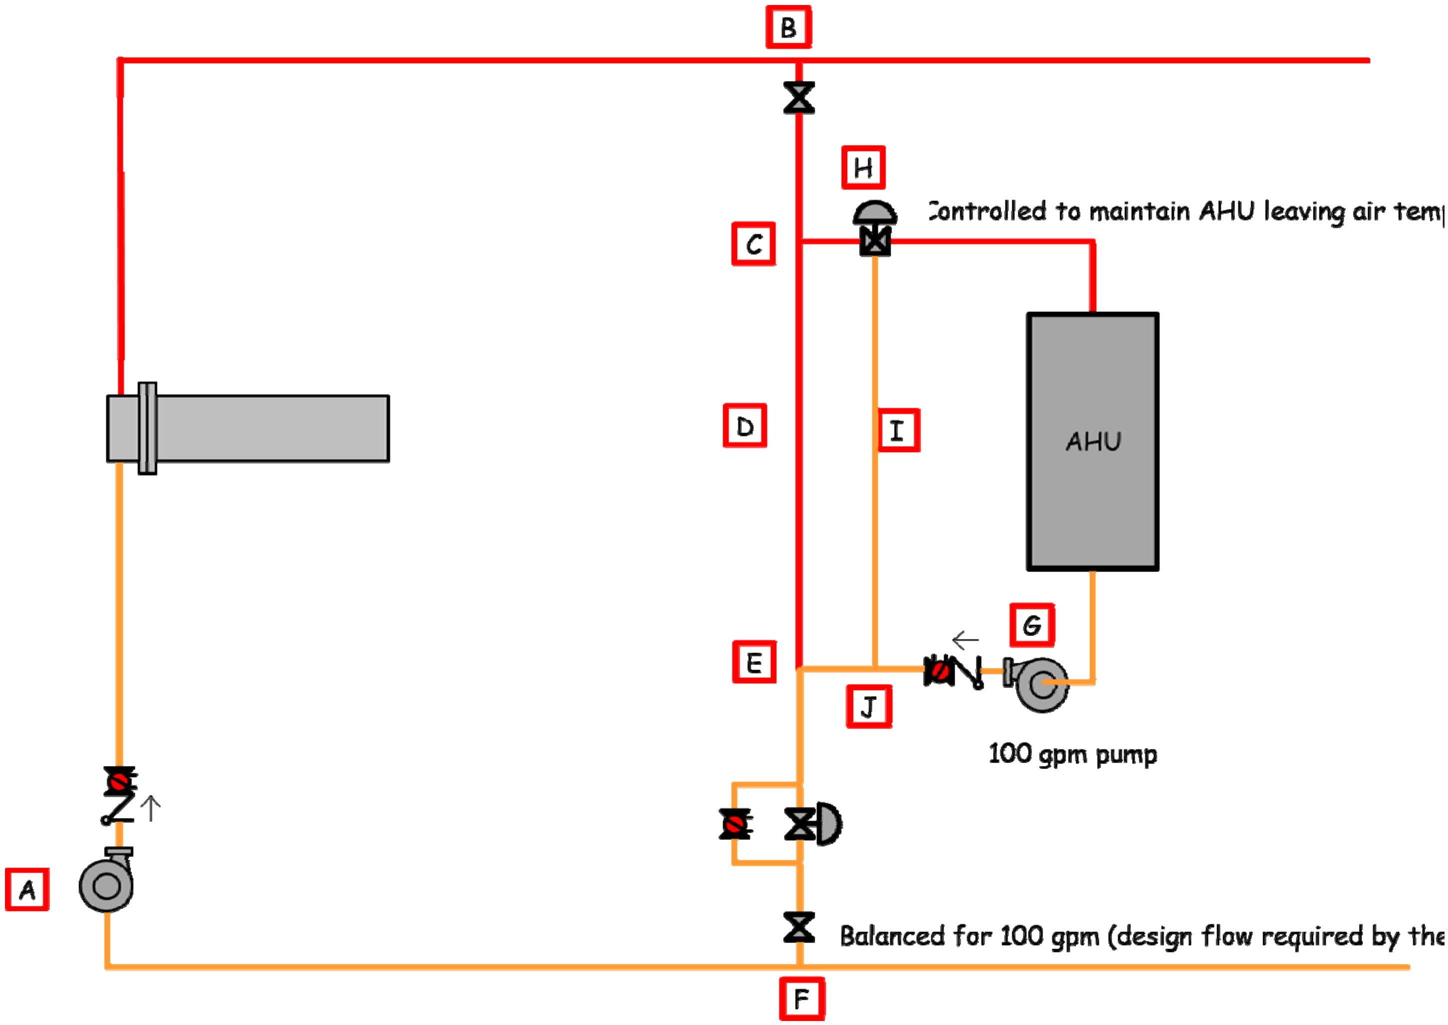

You probably noticed the bypass line with a balance valve around the two way valve that I added. The reason for that is that if you are controlling the two way valve for its own leaving water temperature, then if it closes, you don't have valid data to work with; the pipe would have to warm up via losses to the ambient environment to get the valve back open. So, if you add a small bypass and balance it for a couple of gpm, then you keep circulation up, which is good anyway, and the valve will open back up if it drives itself fully closed.

Just adding the valve will save energy by pushing the pump at A up its curve and that is how the system worked in the olden days before we could afford VFDs as easily as we can now. The trick was to make sure the actuator on the two way valve was sized to close the valve against the peak pressure the pump could produce.

But adding a VFD will provide additional savings by moving the operating point down the system curve instead of up the pump curve.

Convert the Load Loop to Variable Flow

It would also be possible to convert the load loop (the loop served by the pump at G) to variable flow. For that to happen, the control valve needs to be changed from a three-way valve to a two way valve at a minimum. Adding a VFD to the pump would add to the savings you can achive.

If there is a balance valve in the coil bypass line, you can approach the performance of a two way valve by almost but not fully closing the balance valve in the bypass. If the balance valve is not there, then you probably need to do the piping work to replace the three-way valve with a two-way valve.

Developing Cost-Benefit Numbers for the Variable Flow Conversions

To establish the cost/benefit for either of the variable flow conversions, you would use a procedure similar to what we discussed in our last lab session. The first pass approach would be to use the SCE coefficients, which would be fairly quick as long as you had a load profile.

Of course, as we discussed in the lab, getting the load profile is the trick. If you have enough time and data, the building can tell you that. But you may not have enough data from the system to establish a load profile that way.

Having said that, I notice you have interval data, at least I think you do based on your slides, at least for the electrical side of things. And if you do have that data (I think Ryan could help you get the thermal data if you don't already have it), then as a starting point, if you were to look at it as a time series plot for week days and week ends, then the shape of the curve is probably a reasonable way to project the load profile for the chiller plant if you look at the kW vs time curve for a hot day.

Similarly, the shape of the thermal curve for a cold day is a reasonable way to project the load profile for the HHW plant.

In either case, would assume the valley of the curve was at the base load condition and thus represented near 0 tons or therms in terms of load on the heat exchanger or chillers. And I would assume that the peak on the curve represented the full load condition. So, if you then correlated that with the outdoor temperature, I think you could generate a reasonable load profile curve based on OAT and use that for your first pass assessment.

If you want to go there, I would be happy to help you with it.

But for now, I will close, since I am out of time and also since I have given you a lot to think about I suspect.

David

Senior Engineer
Facility Dynamics Engineering
Northwest Satellite Office
8560 North Buchanan Avenue
Portland, Oregon 97203
Office - 503-286-1494
Cell - 503-320-2630
DSellers@FacilityDynamics.com
http://www.facilitydynamics.com/

Visit my Commissioning Resources website at www.Av8rDAS.com
View *A Field Perspective On Engineering* *and past posts from*
A Field Guide for Engineers at http://av8rdas.wordpress.com/
View The Other Side of Life at https://av8rdaslife.wordpress.com/

From: Lane Burt <lburt@emberstrategies.com>
Sent: Wednesday, July 11, 2018 11:28 AM
To: Stroupe, Ryan <R2S2@pge.com>; David Sellers <dsellers@facilitydynamics.com>
Subject: Fwd: EBCx-Yr13 reminder

Hi Ryan and David,

See attached updated slides on my building. I've added a savings calculation based on the 4 weeks of limited BMS data I was able to get (slides 24 to 27). I also added a picture of one of the confusing bypasses around the three-way valves on the air handlers (slide 9) but I still don;t understand it.

Can you glance at these and let me know if I need to do more on the system diagrams, savings calcs or otherwise? I suspect it would be good to have clarity on the bypass issue before the presentation, but with travel I only have one more possible day where I could visit the building before Aug 2nd.

Cheers,

Lane

Lane Burt, PE | Managing Principal, Ember Strategies
(646) 801-5263 | lburt@emberstrategies.com

---------- Forwarded message ----------
From: **Stroupe, Ryan** <R2S2@pge.com>
Date: Tue, Jul 10, 2018 at 9:41 AM
Subject: EBCx-Yr13 reminder
To: "toddzilla40@gmail.com" <toddzilla40@gmail.com>, "laura.bello@ucsf.edu" <laura.bello@ucsf.edu>, "lburt@emberstrategies.com" <lburt@emberstrategies.com>, "raniel.camacho@flysfo.com" <raniel.camacho@flysfo.com>, "mingc1728@gmail.com" <mingc1728@gmail.com>, "david.chiu@p2scx.com" <david.chiu@p2scx.com>, "steve.eng@ableserve.com" <steve.eng@ableserve.com>, "andrew.gustafson@lmco.com" <andrew.gustafson@lmco.com>, "Hambalek, Steve" <s3h6@pge.com>, "rhollingsworth@harvestproperties.net" <rhollingsworth@harvestproperties.net>, "karkis@gene.com" <karkis@gene.com>, "mendonsad@saccounty.net" <mendonsad@saccounty.net>, "abram.mission@flysfo.com" <abram.mission@flysfo.com>, "arturo.piceno@ableserve.com" <arturo.piceno@ableserve.com>, "wenhan@stok.com" <wenhan@stok.com>, "john.rosendo@ucsf.edu" <john.rosendo@ucsf.edu>, "raymond@verdafero.com" <raymond@verdafero.com>
Cc: "David Sellers (dsellers@facilitydynamics.com)" <dsellers@facilitydynamics.com>, Tony Pierce <tonyp@facilitydynamics.com>

Hi folks,

This is a gentle reminder to wrap up you EBCx project work and integrate all your findings into a single PowerPoint presentation for our final meeting on August 2. Here is the intended outline for this final presentation:

1. Building description (abbreviated) [1]
2. Benchmark data [1-2]
3. Interval data analysis (cloud graphs) [1-4]
4. Issues log/list of measures found [1]
5. System diagram(s) [1-4]
6. Monitoring points list [1]
7. Annotated trend/data-logger data graphs [1-4]
8. Functional test results [1-2]
9. Energy savings calculations [1-4]
10. List of recommendations/Next steps [1-2]

The number in [] is the desired number of slides for each section. Try to rehearse you presentation so you can cover all your slides in 20 minutes. Each of you will be allocated an additional 10 minutes for Q&A.

If you want some guidance as you wrap up your project work, please reach out to me. I am happy to set up a go-to-meeting with any of you.

If you are in a panic about finishing by August 2, you are welcome to present your final project work at one of the classes intended for EBCx graduates. Here is the schedule for these with registration links through the end of the year:

EBCx Workshop and Project Reviews

September 21, 2018, 8:30 am to 4:30 pm at PEC, 851 Howard Street in SF.

http://66.198.243.12/event-details?EventID=19395

December 14, 2018, 8:30 am to 4:30 pm at PEC, 851 Howard Street in SF.

http://66.198.243.12/event-details?EventID=19396

And please plan to join me, David and Tony for a drink and conversation at the Chieftain after class on August 2.

Ryan Stroupe

Building Performance Program Coordinator

PG&E Pacific Energy Center

851 Howard Street

San Francisco, CA 94103

P: 415.973.7257

E: r2s2@pge.com

www.pge.com/pec