From: David Sellers
To: Ryan Stroupe (R2S2@pge.com); Patrick Testoni (testoni@ucsc.edu); Mike Dolder - UC Santa Cruz Physical Plant; Sarah Gilchrist (sgilchri@ucsc.edu)
Subject: Follow up on UC Santa Cruz Site Visit
Date: Sunday, January 26, 2014 7:49:00 PM
Attachments: 2013-01-16 CW System Diagram Overview (1).pdf
CW Flow Diagram v1 Layout1 (1).pdf
2013-01-16 CW System Diagram.dwg
image001.png
image002.png
Thimann AHUs.dwg
Typical AHU v2.pdf
Music Building AHU v1.dwg
Music Building AHU v1-Overview.pdf
2014-01-20 Piano Tuner"s Association Humidity Control.pdf
2014-01-20 Piano Technician"s Guild Humidity Control.pdf
image024.png
Humidifiers.pdf
image025.png
image026.png
image027.png
image028.png
image031.png
image032.png
image033.png
image034.png
image029.png
image030.png
image005.png
image008.png
image016.png
image018.png
image022.png
image023.png

Hi all,

I thought I would take a minute (actually longer than that, but you get the idea) to send some notes that follow up on some of the observations I had from the week we spent on site in case they might be useful for moving forward with projects. I had meant to get this out sooner, but it took a while to put it all together; I was so sick that week that I didn't get much done in the evenings. Since then, I was off a couple of days up in Canada with my grandson and his dad for their birthday. And, I have needed to juggle it in with other stuff. So it's just taken a while.

Your site is very "rich" (to borrow Ryan's term) in the context of diverse systems with interesting operating challenges. I am revising my upcoming economizer class to include two of the things I observed there (your wet economizer cycle - if you could send a screen shot of the Sin Lab piping, that would be really helpful since I never got that graphic to work - and the Music Building AHUs).

Ryan, I think there could be a mutually beneficial sort of thing in here; specifically, maybe some sort of applied EBCx class called *Saving Energy and Improving Performance at UC Santa Cruz; Case Studies in Applied Existing Building Techniques and Implementations*. Or maybe it's a string of classes that starts out with how the targets were scoped and quantified and then follows up with a second class about the implementations and verifications. Point being that if some sort of arrangement were developed were I worked with the PEC and UCSC to move these things forward but at the same time, captured the content we would need for a pretty nitty-gritty class, everyone could win.

But even without that, I think there is a lot of potential for "wins" for the UCSC team, both in their current projects and for future projects, which I am hoping there might be seeds for in what follows.

In any case, I think we touched on all of this stuff during the course of the week, but it's a lot. So I thought it might be helpful if I wrote down the key points. These links will take you to particular topics so you don't have to read the whole thing if you don't want to or don't have the time. The "Back to Contents" link at the end of each sections will bring you back here.

I've also attached some of the stuff I developed as we worked through things over the course of the week; mostly first draft system diagrams.

Condensing Boilers

This is something that I think might apply to any building on your site that is not served from the cogen plant; i.e. any building with boilers. Patrick mentioned his insight regarding condensing boilers; specifically, that for them to deliver the extra efficiency they are capable of, they need to be operating in condensing mode, which means that the return water temperatures need to be lower than 130-140°F or lower for a gas fired unit. That is an important insight to have and since the campus HHW standard is 160°F return and 180°F supply, it may not seem like condensing boilers are an option.

But I think they may actually be something to consider, at least for the buildings that are not on the cogeneration system. I say that because for one thing, what you need on the design day (less than 0.4% of the hours in a year) is a lot different from what you need the rest of the time. Plus, I am not sure what drove that design standard in the first place, especially for systems that are not on the cogen plant HHW system other than it is a common rating point for systems vs. the numbers you might need to deliver heat, reheat, and preheat in the Santa Cruz environment.

Radiant Heat and Hot Water Supply Temperature

Finned tube radiation usually needs 140°F or so supply water to start delivering meaningful amounts of heat. But, in Santa Cruz, you may not need much more than the FTR would deliver with water at that temperature, at least most of the time. And, if you have (or targeted) using radiant slabs, you actually don't want the temperature of the supply water to be much more than 85°F for a bunch of reasons. So, moving towards designs that did space heating with radiant slabs could be desirable. LEED really likes this approach too, so if you did it, you would be "all the rage" with the LEED folks I think.

Preheat and Hot Water Supply Temperature

A significant number of your buildings 100% outdoor air, meaning those buildings probably have a significant preheat load. If you look at preheat coil selections, you will discover that even a 2 row coil can do a lot of preheat with relatively cool water, mostly because you are not looking for a leaving air temperature that is that high.

For example, if you look at what a 2 row coil could do on the extreme day for the area (33°F outside), you can get 60°F leaving air with water supplied a t 110°F, which will come back at 86.4°F, both of which put you at a pretty good efficiency in a condensing boiler.

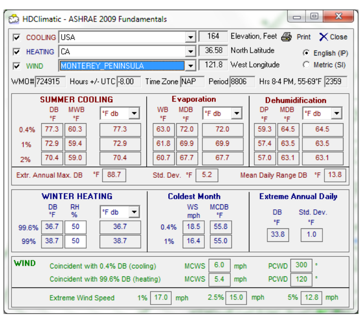

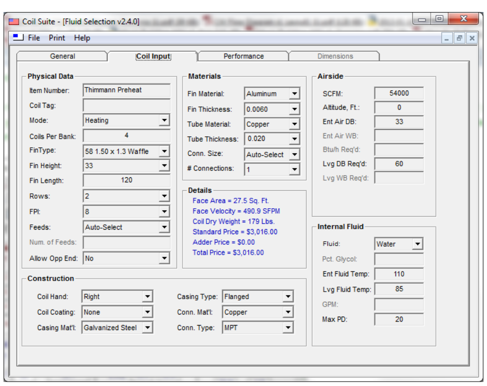

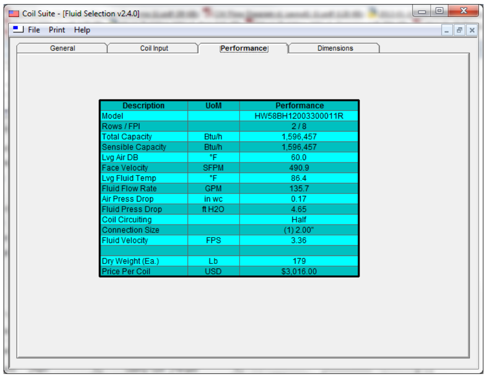

Description	UoM	Performance
Model		HW58BH12003300011R
Rows / FPI		2 / 8
Total Capacity	Btu/h	1,596,457
Sensible Capacity	Btu/h	1,596,457
Lvg Air DB	°F	60.0
Face Velocity	SFPM	490.9
Lvg Fluid Temp	°F	86.4
Fluid Flow Rate	GPM	135.7
Air Press Drop	in wc	0.17
Fluid Press Drop	ft H2O	4.65
Coil Circuiting		Half
Connection Size		(1) 2.00"
Fluid Velocity	FPS	3.36
Dry Weight (Ea.)	Lb	179
Price Per Coil	USD	$3,016.00

If you look at that same coil at a more average condition (say 45°F outside), you can get 60°F leaving air with water supplied at 88°F, which will come back at 74.6°F, even better if you are a condensing boiler in terms of being able to operate at a very high efficiency point.

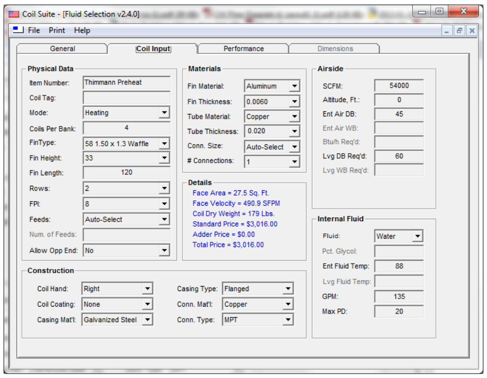

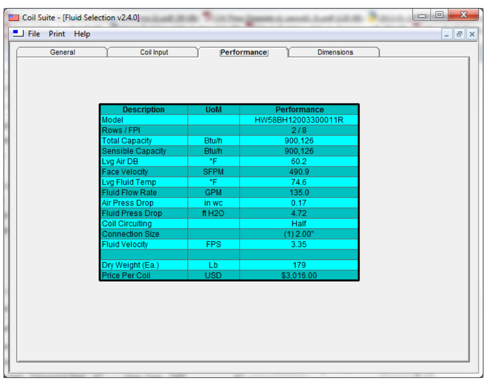

Description	UoM	Performance
Model		HW58BH12003300011R
Rows / FPI		2 / 8
Total Capacity	Btu/h	900,126
Sensible Capacity	Btu/h	900,126
Lvg Air DB	°F	60.2
Face Velocity	SFPM	490.9
Lvg Fluid Temp	°F	74.6
Fluid Flow Rate	GPM	135.0
Air Press Drop	in wc	0.17
Fluid Press Drop	ft H2O	4.72
Coil Circuiting		Half
Connection Size		(1) 2.00"
Fluid Velocity	FPS	3.35
Dry Weight (Ea.)	Lb	179
Price Per Coil	USD	$3,016.00

In fact, if you model what is there in Thimmann (can't quite model the spiral fins, but this should be close), and limit the gpm to what the pipe that is there could reasonably deliver (about 130 gpm) here is what you get on the nominal day.

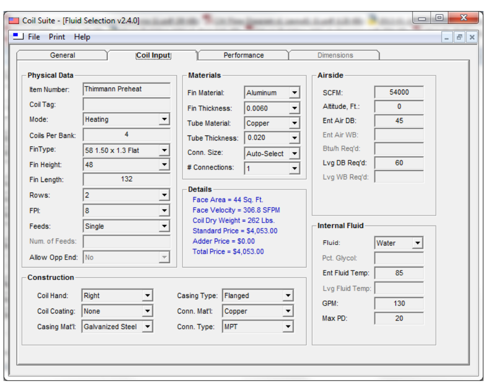

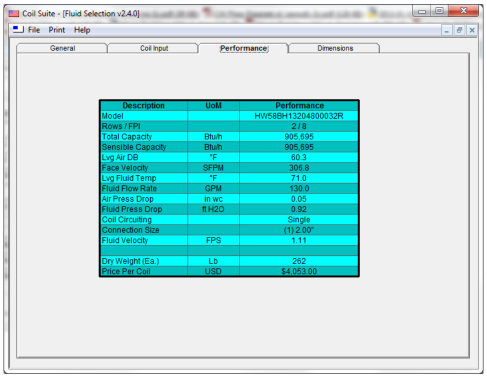

Description	UoM	Performance
Model		HW58BH13204800032R
Rows / FPI		2 / 8
Total Capacity	Btu/h	905,695
Sensible Capacity	Btu/h	905,695
Lvg Air DB	°F	60.3
Face Velocity	SFPM	306.8
Lvg Fluid Temp	°F	71.0
Fluid Flow Rate	GPM	130.0
Air Press Drop	in wc	0.05
Fluid Press Drop	ft H2O	0.92
Coil Circuiting		Single
Connection Size		(1) 2.00"
Fluid Velocity	FPS	1.11
Dry Weight (Ea.)	Lb	262
Price Per Coil	USD	$4,053.00

Here is what you get on the design day.

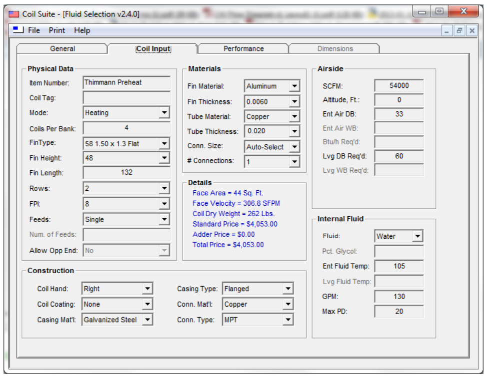

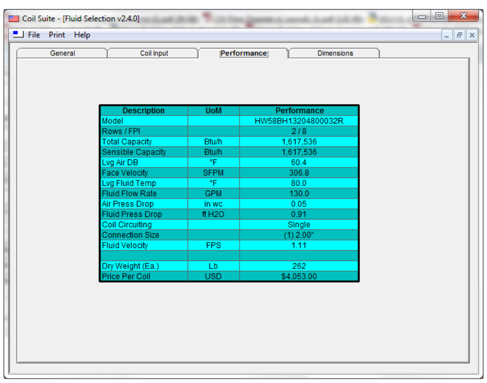

Description	UoM	Performance
Model		HW58BH13204800032R
Rows / FPI		2 / 8
Total Capacity	Btu/h	1,617,536
Sensible Capacity	Btu/h	1,617,536
Lvg Air DB	°F	60.4
Face Velocity	SFPM	306.8
Lvg Fluid Temp	°F	80.0
Fluid Flow Rate	GPM	130.0
Air Press Drop	in wc	0.05
Fluid Press Drop	ft H2O	0.91
Coil Circuiting		Single
Connection Size		(1) 2.00"
Fluid Velocity	FPS	1.11
Dry Weight (Ea.)	Lb	262
Price Per Coil	USD	$4,053.00

Meaning if the existing coil performed anywhere in the ballpark of this, it would be pretty good conditions for condensing boilers if you look at the return water temperatures.

Reheat and Hot Water Supply Temperature

I think I have mentioned this in class, but wanted to re-iterate it here. It is important to realize that a reheat process is just using up unnecessary cooling that is being supplied to the zone in question in the form of colder air than necessary as a result of the cooling or dehumidification requirements of other zones served by the air handling system. Meaning, that the air is still coming out of the diffuser for the zone in question at a temperature below the zone set point (providing cooling capacity to the zone) but it is not as cold as the supply air temperature.

For a building that heats (supplies air above the zone temperature to offset envelope losses) using an air system with reheat coils, you will discover that frequently, once the ambient outdoor temperatures get into the range of the balance point for the building, the air temperature supplied to the zone will need to be at some temperature below the required set point (providing cooling) vs above the zone set point (providing heating). And, you will discover that a reheat coil that has been selected to deliver air above the zone temperature on a design heating day (when the zone has a net loss of energy due to the envelope losses) can deliver air that is below the zone temperature but above the supply air temperature if it is supplied with water in the 85-95°F range (pretty near perfect condensing boiler temperatures).

So, while the system may need to run in the 140-160°F range (or higher) on a design heating day (less than 0.4% of the hours in a year), it can often run with supply temperatures that are near perfect supply and return temperatures for high efficiency condensing boiler operation when the ambient temperatures are near the building balance point (a lot of the time in Santa Cruz I supect). Tom Stewart and I wrote a paper about this for ACEEE, which is available via a link from my blog.

My point in bringing it up here is that you have a lot of constant volume reheat zones in a climate that is fairly mild most of the time. As a result, it would not be a huge surprise to discover that you could run your HHW systems at supply temperatures in the 90-100°F range (or lower) for a significant number of hours. Meaning that if the systems were served by condensing boiler, they would be running at pretty high efficiency levels and that the parasitic losses would be minimized. So condensing boilers really may be a viable option for the campus if you allow yourselves to deviate from the HHW supply temperature design standards on non-design days.

Having said that, it is important to realize there are a couple of constraints.

1. The HHW piping needs to be able to tolerate the lower temperatures with out leaking. This can be an issue for system with mechanical couplings, as we have discussed in class.
2. The existing, non-condensing boilers will be ruined if you operate them with EWTs below about 140°F. That means that if you go this route, you need to totally replace them with condensing boilers or you need to make a decoupled system that lets you run the boiler loop in a safe range for the non-condensing boilers while the distribution loop (and any added condensing boilers) are operated in the condensing range.

Back to Contents

Sinsheimer Labs Chiller Plant

- I really think your "wet economizer" cycle is clever; I e-mailed Liz Fischer (executive director for BCA and the person who basically runs the National Conference on Building Commissioning) to see if she was looking for a case study to fill a spot in the agenda at NCBC this year and suggested it as a candidate if she is. Even though the call for papers is closed, a lot of times, they end up needing something like that as they put the agenda together, so you never know.

- I never was able to get the graphics for the Sin Lab piping to work. Could one of you just send me a screen shot? I would like to mention what you have done (with acknowledgement of course) as another option for doing a wet economizer and I would like to have a system diagram to contrast with the other options I talk about and I think I could come up with a generic one if I had the piping graphic. The class is coming up in a little over a week so if you could send it sooner rather than later, that would be great.

- There were a number of throttled valves on evaporator and condenser pumps for the chillers. Based on a rough estimate of the actual head required vs. the nameplate data for one of the pumps, I suspect the energy savings potential is in the 3-4 kW range for the larger pumps; less for the smaller ones. So how much you can afford to do may be limited by the fact that you probably don't have to run the pumps that much. Meaning that throttling could be the best option if the pumps remain in place.

- With regard to the closing line in the preceding bullet, I think it may be possible to eliminate the condenser water pumps entirely. I did a quick sketch of a system diagram for the connection and added it to the central condenser water plant system diagram that I made the last time I was on site (copy attached and screen shot below).

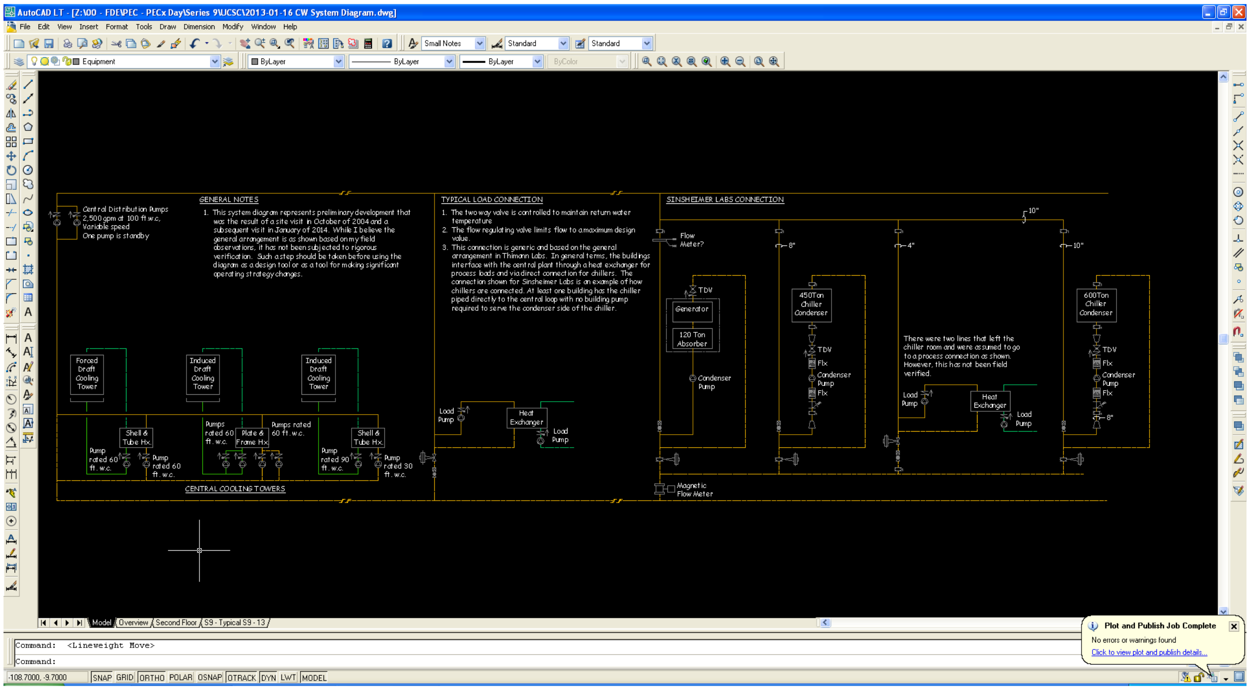

It seems to me that if you know you can get design flow to the chillers, then on a design day, that flow is being delivered by the central plant pumps. If they have enough capacity to pick up the head through the condenser barrels, then I think it might be possible to eliminate the condenser pumps and modulate the control valve based on head pressure or a target leaving water temperature. This would total eliminate the condenser pump energy in the building. If the central plant CW pumps were delivering this flow anyway and delivering it at or above the head required (meaning that on a design day, the control valves were not wide open), then there would be no additional pump energy required and you could further reduce the central plant pump energy by optimizing the pump speed(s) so that the delivered pressure caused the critical valve in the system to be almost fully open.

Back to Contents

Thimmann Labs

- This is a screen shot of a very elementary system diagram I made of the air side system in the building. It needs a lot of work to get it fully developed, but it helped me think things through while I was there.

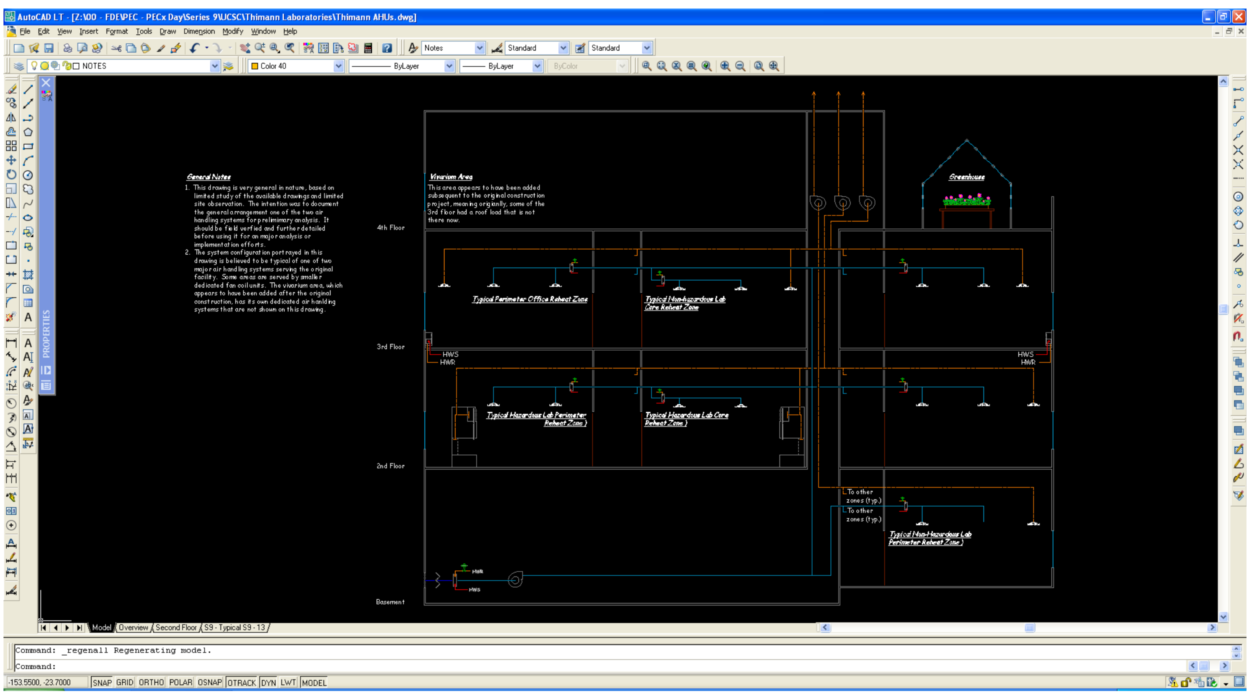

I am pretty sure that the diagram is typical for two systems, each of which serve about half of the building.

I think the primary opportunity that I saw was the zone level scheduling idea that we discussed during the tour. But you need to be sure of a couple of things if you do it. One is that at the zone level, you don't violate and Environmental Health and Safety rules, which typically go back to OSHA and code requirements, but also can be set by policies in the organization or by policies in associations that support research at the site. My experience is that you can do it for labs like the physics lab we were in.

But if you do it at labs like that, it can still create an issue at the other end of the system, which we alluded to but did not discuss much due to time constraints. Specifically, from what I can tell from fairly limited exposure to the project drawings you provided and from your own knowledge of the building, the exhaust fans are likely dedicated to floors (one fan per floor per make up AHU or two exhaust fans per floor total) and that they serve all exhaust loads (hoods, general exhaust, toilet exhaust, etc.).

That means that the effluent from the fans could contain materials that are considered hazardous (in the state of CA and otherwise). Frequently, when that is the case, it means that the discharge velocity from the fan must be maintained above a minimum velocity in order to ensure that the plume is ejected far enough into the atmosphere that any hazardous fumes are dissipated.

In my experience, most of the time, the velocity is either dictated prescriptively by some governing standard, or it is derived from modeling. So, the point is that you would need to determine what, if any standard applies in this regard and then make sure that when you drop the flow, you don't violate it. There are a number of ways to engineer your way out of that problem, but the first step is to understand what the requirements are. Typically, the fan energy, reheat, and preheat savings associated with a change like this are significant enough that you can justify some pretty significant modifications to allow them to happen. Of course, the more complex your strategy, the harder it may be to ensure persistence. But you guys seem to be pretty on top of stuff so that is probably not as much of an issue on your campus as it might be other places.

- It occurred to me after I was on site that on the 3rd floor, there could be a couple of other opportunities. From our discussions and based on comparing drawings that you provided, I think the vivarium area was built after the original construction project. That means that for much of the floor, there originally was a roof load that is no longer there. From a cooling standpoint, the zone flows would have been set up for sufficient cooling to offset that load in addition to the other loads. So, if that load went away, then, in a constant volume reheat system, the unneeded cooling would be absorbed as a reheat load unless the system was rebalanced at the time. This is a phenomenon similar to the issue that comes up when you reduce the lighting load in a constant volume reheat system without rebalancing; the electrical consumption goes down as a function of the reduced electrical energy used by the lights, but the thermal load goes up due to the added reheat burden in the cooling season and the need to offset losses through the envelope during the heating season that were originally offset by the internal gains from the lighting.

- I alluded to another item but didn't really discuss it in the context of your building. Specifically, on the third floor, the convectors/Finned Tube Radiation (FTR) elements around the perimeter are there to offset the envelope loses (which have been reduced by the elimination of the roofing load incidentally). That means at some point, that means that the cool (relative to the space) air being provided by the reheat air handling system is working against you since it is trying to cool and, if it is not needed, the system reheats the air to be neutral, and, if the FTR is not on or has its valve closed (meaning the FTR is not picking up the perimeter heating load), then it is heated above the space temperature to offset the envelope losses. A more efficient approach at that point would be to reduce the follow from the reheat system to what is required for ventilation and simply allow the FTR to offset the perimeter loads.

  So, I think the opportunity is to understand where that balance point is and then install some two position dampers for the office spaces that reduce the air flow to the ventilation requirement when you are below that point and operate the FTR system to offset perimeter loads. The two position dampers may provide for another savings opportunity in that you could use them to do zone level scheduling for the office areas by closing them total when the offices were empty.

- Mike, one thing that you might want to follow up on more is the chilled water connection we were looking at on the roof. Something about it just didn't seem right but to really understand that, you would need to take the connection back to the central plant and look at it in the context of the rest of the system (i.e. make a system diagram). But it seemed like the connection, as it was configured and controlled, could be a contributing factor to the shortage of chilled water that you mentioned. Meaning you might not be getting the capacity you need because of issues with the configuration and control of the secondary connection, not parasitic losses (temperature rise) in the piping to the point of connection. In my experience, a degree or two of temperature rise in the distribution system is not unusual. But 4, 5, or 6 degrees – which I think is the number you mentioned in terms of temperature rise between the plant and the building - is a lot and could mean you are measuring a temperature that is a blend of supply water and return water, which is something that could happen with the type of connection we observed.

Back to Contents

Music Building

- One of the striking things about the music building is the air handling system design. It's an interesting system and I have not seen this arrangement before.

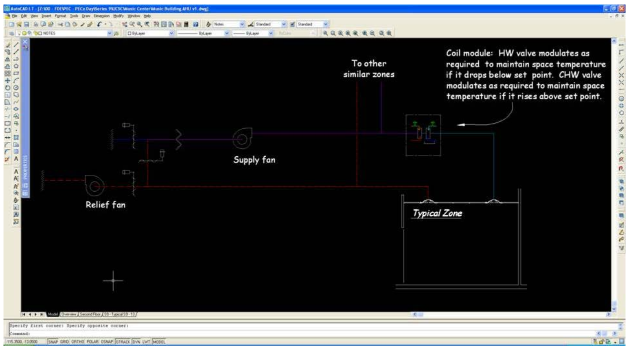

  Systems that use a central 100% outdoor air MAU to provide the ventilation air and smaller AHUs to pick up the actual loads using a wet economizer cycle look similar to this, but each set of coils would have a fan that recirculated the total required air flow while the main fan only provided the ventilation air.

  The Music Building system configuration seems simple, but when I start thinking about how it would work and how to optimize it, it's actually a lot more complicated than you would think, mostly because the AHU is delivering all of the air required, not just the ventilation air and it has an economizer.

  From an energy perspective, I suspect it has the potential to be very efficient since it provides economizer cooling for all zones and will only use heat or cooling if a zone needs it and then only use it on a zone by zone basis. But if the zones are diverse in terms of their load profile, I think that efficiency potential might fall away or become difficult to optimize.

  For instance, if all of the zones needed about the same amount of cooling, then they would all need about the same supply temperature and you could pretty easily optimize the economizer to deliver it and, since it would be the temperature needed by all of the zones, none of them would be using their heating coil, nor would they be using their cooling coil, assuming you could provide the desired temperature with outdoor air. And, if you could not, you could still optimize the change-over set point for the integrated economizer pretty easily given the similarity between the requirements in the zones.

  But, if you consider an extreme condition I think that things become more challenging. For instance, for the sake of discussion, if you had a core zone with a lot of heat generating equipment that made heat all of the time combined with some perimeter zones that had significant envelope losses during cold weather due to a lot of glass, and then a number of zones that were lightly loaded but did not have much of an envelope loss, then in cold weather, you would have to decide:

  - Should you use the economizer to make the air cold enough to satisfy the heavily loaded core zone, which would mean all of the other zones would be using their heating coil. The perimeter zones would not need any cooling, meaning that they would need air coming out of the diffusers at a condition that was above the space temperature to add energy to the space and make up for the losses through the envelope. So not only would they need to heat the air to offset the envelope losses, they would have to heat the air from the supply temperature to the space temperature too (a reheat process).

    Meanwhile the lightly loaded zones might need cool air (air below their space temperature) but not as cold as the heavily loaded zone,, so they would be using their heating coil to offset the unneeded cooling provided by the economizer, which was set based on the requirements of the heavily loaded zone; i.e. they would be doing reheat.

    Or

  - Should you satisfy the lightly loaded zones with the economizer, which would minimize the reheat they would use, and would also minimize, but not eliminate the reheat at the perimeter zones, but would also require that you run a chiller or increase the air flow to the heavily loaded zone to satisfy it.

    Or

  - Should you go to minimize outdoor air to minimize the amount of reheat the perimeter zones need, which would mean you needed to run the chiller to satisfy all of the other zones, but would not be doing any reheat (all though if it was cold enough, you would be doing preheat at the zone level if the mixed air temperature on MOA was less than the space temperature at the perimeter or less than the supply temperature you needed for the lightly loaded zones). Incidentally, this particular option would likely minimize your humidification problem since you would be keeping the moisture in the building by recirculating most of the air. (More on this in a minute.)

  The correct decision is very much a number of the ratio of the different zone types and the economics of the cost to provide heating and cooling. For instance, if the system was served by electricity and heat from the cogen plant (I know it isn't, at least on the hot water side, but I guess it could be part of the electrical load – not sure about that, but am saying this just for the sake of discussion) then:

  - If the plant was being operated to meet an electrical load and, as a result, was making more heat than could be used by the connected hot water loads, then using the economizer to satisfy the heavily loaded core zone and the lightly loaded zones (meaning you would generate no additional electrical demand by running a chiller) might be a good plan since the heat you needed would essentially be free since, if you didn't use it in the building, it would have to be rejected at the plant's heat dump cooling tower, which would actually cost energy.

    But

  - If the plant was meeting the electrical demand and you had to fire boilers to supplement the waste heat it was generating to satisfy the heating load, then minimizing the need for heat in the music building by going into a recirculation mode and running a chiller, which would increase electrical demand on the cogen plant and thus, increase the waste heat that was created. In a perfect world, you might bring the plant back into balance, meaning the waste heat generated by meeting the electrical requirements would exactly meet the heating load. While perfection is unlikely, operating the added electrical load would minimize the need for boilers.

    But

  - If the building is served totally from the public utility, the decision has more to do with the cost of gas, the cost of electricity, and the conversion efficiencies of the chillers and boilers. Meaning the answer would be different for a building that made its heat with condensing boilers and used hot gas bypass when it ran a chiller at low load vs. a building that had an old atmospheric burner type boiler but a modular chiller that could turn down to a very low load condition very efficiently.

  So, while optimizing the supply temperature set point from the AHUs is a good target, figuring out exactly what to do is complicated. But the good news is that the building knows the answer and you can probably use data logging and trend analysis techniques to figure it out.

- Another striking thing to me about the Music Building was that, in terms of the whole purpose of HVAC being to provide a safe, clean, productive environment, (where in this case, a prime consideration is the "health" of the instruments), we are failing. I did some research to understand the temperature associated with the RH levels that were discussed during our meeting with the staff and the psych chart below has an ideal envelope and state point on it, based on information from what I believe would be reputable organizations with regard to the topic.

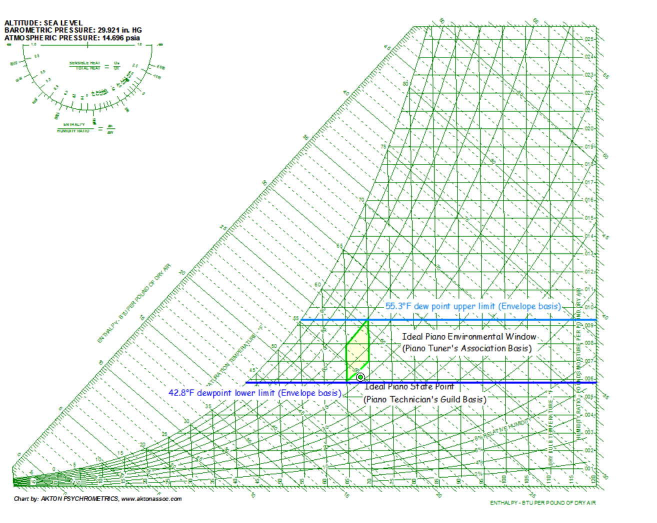

This is that same chart along with the Santa Cruz environment as a bin plot.

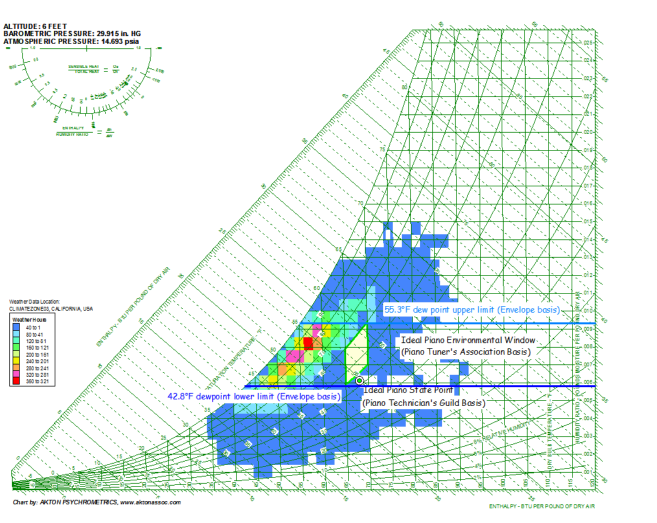

If one were to target the ideal envelope vs. the ideal state point, and assume that:

- o The air (and moisture) in the building comes from the outdoor air and thus, the outdoor air dew point sets the minimum dew point you will see in the facility, and
- o The loads in the facility are not significant contributors or removers of moisture on the average (the exception is a full auditorium for a performance),

then much of the time, the air from the outside would provide the required indoor environment if you use the area inside the ideal envelope as a target.

The problem then becomes what to do about the exceptions to those conditions, and there are two cases to consider.

- o Conditions when the outdoor air would drive the humidity above the envelope
- o Conditions where the outdoor air would drive the humidity below the envelope

Given the underlying goals behind HVAC systems, specifically, in this case, to protect and preserve valuable instruments, it would seem like saying "just let things drift" is not the right answer, despite the fact that it is probably the least energy intensive alternative. So, I think our goals should be to understand how we can protect the instruments and preserve their tuning as efficiently as possible.

I pulled back a couple of years of hourly weather data from Watsonville (it was the closest station I could get hourly data from’ there is a Citizen's Weather Observer site in Santa Cruz, but no long term records available from it that I could find) and generated the hourly dew point pattern, which is illustrated below along with the target window illustrated in the psych chart. First, the pattern for the entire two years that I pulled back.

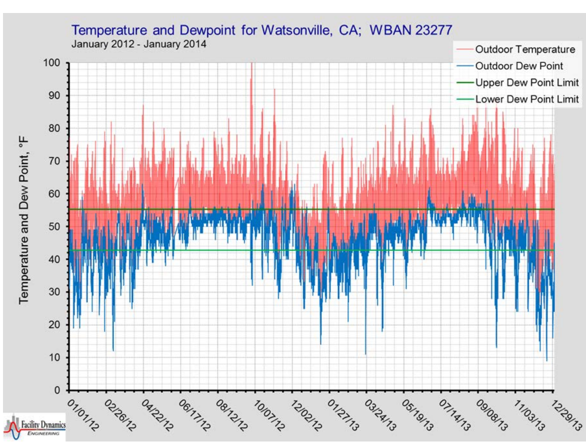

What struck me about this is that it says the real issue is more a lack of humidity than excessive humidity. That sort of matches the conclusion you would reach from the bin plot on the psych chart, which is nominal/average weather data vs. specific weather data. Meaning that on the average, there are more blue areas (hours) when the outdoor moisture conditions are below the lower limit on the ideal envelope than there are blue areas (hours) when the outdoor moisture content is above the upper limit.

If you look at one of the low dew point weeks, you see something like this.

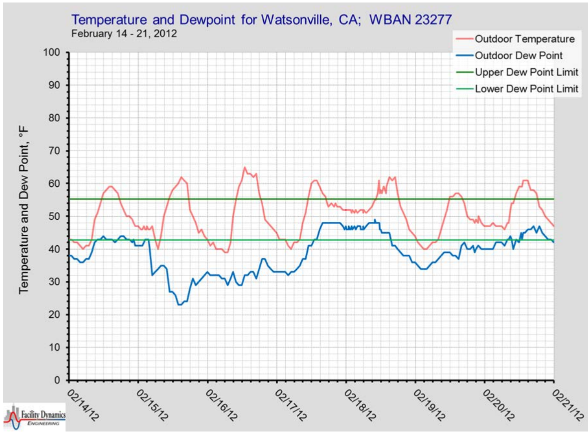

A week with high dew point events looks something like this.

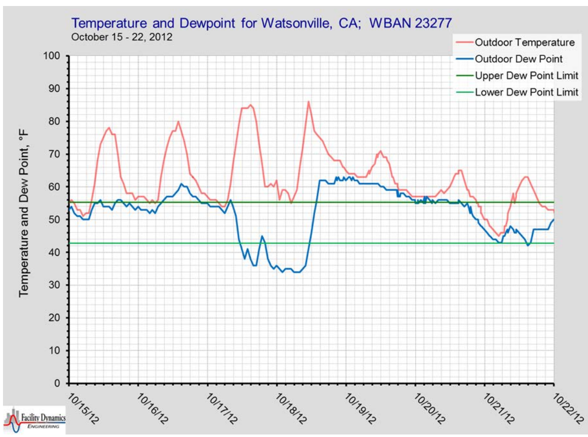

… which is interesting because the same week has a low dew point event. There are other instances that look more like this.

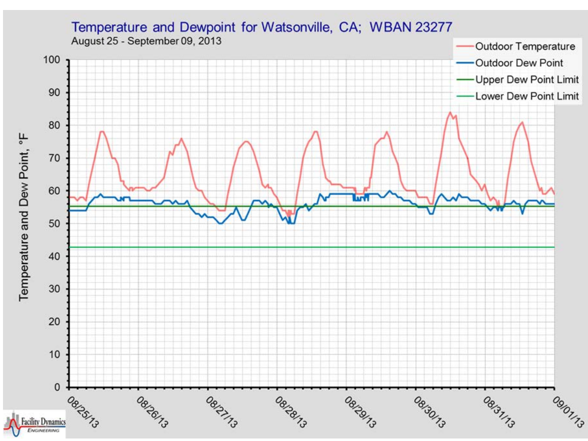

In all of the instances, if there is a deviation, it tends to last a while vs. being a one or two hour event. That makes sense because from what I have observed watching weather data, dew point is something that is associated with an air mass; meaning that it will tend to persist at a fairly steady level until a new air mass (weather system/front) comes through.

The reason that mattered to me was that initially, when I looked at the pattern, I thought a mitigation strategy might be to go to minimum outdoor air for the duration of the event. If you watch building moisture levels, you come to realize that there is actually moisture inertia, just like thermal inertia. (I did that some on some archival storage projects where we were trying to dry out a space to a fairly low specific humidity level. I've also seen it as buildings come on line and there is a lot of drywall compound drying out and stuff like that.) Basically, even though you might fill a space with moist or dry air (relative to what had been there) it takes some time for the moisture levels in the air to equalize with the moisture levels in the materials that make up the space.

So my thought was that maybe going to a recirculation mode bringing in only the amount of OA needed for ventilation would ride you through a couple of hour event, and I think it might. But the events we would need to deal with are days long, not hours long, so a control strategy alone will probably would not solve the problem.

That gets us back to needing to humidify a lot and dehumidify occasionally. Sadly, the system you have, at least as I understand it currently, is not set up to do either.

One way to look at this is to say "you know what, there are a lot of piano's sitting around the Bay Area (and the world) that see environmental swings that are as bad or worse than what we are talking about here and they are doing just fine, thank you very much" and not worry about it. But, given a number of things:

- The goal of HVAC systems alluded to previously;
- The University setting;
- The quality of the instruments that are likely in the building;
- The value of the instruments that are likely in the building and the cost to repair or replace them when soundboards crack, glue joints fail, felt dries out, pins and strings rust; etc.;

I think we need to at least consider some options.

If the pianos are really the primary focus of concern, then I suspect the least expensive alternative in terms of first cost and energy might be too install one of the piano environmental control systems that are on the market.

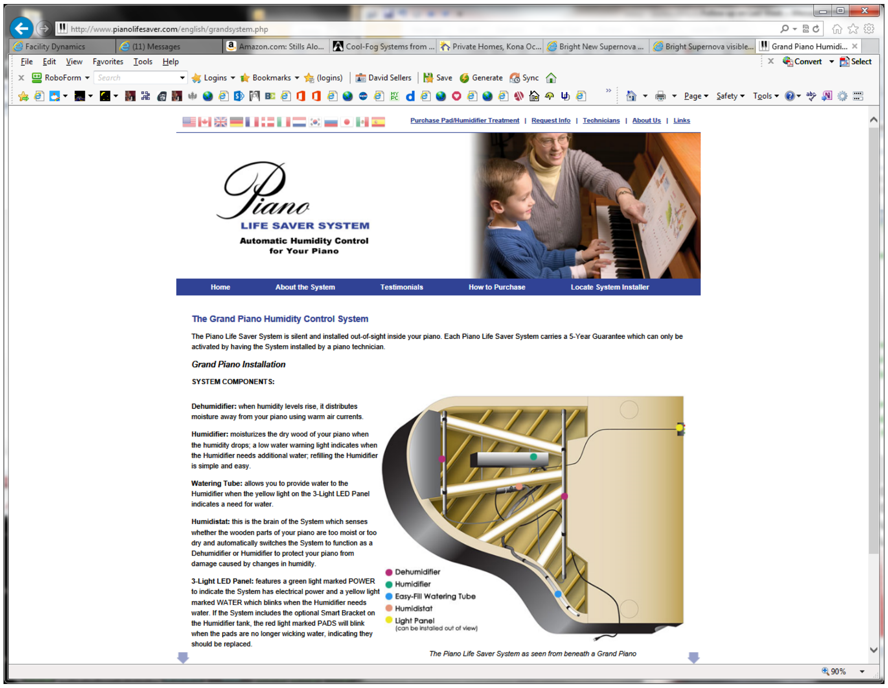

These systems install totally inside the piano and target maintaining that local environment rather than maintaining the environment of the room that contains the piano. So energy and water consumption would be minimized relative to the room because the air in the room is always being changed by the action of the AHU whereas the air inside the piano is somewhat isolated, especially if the lid is kept closed when the piano is not in use.

The next level of effort might involve trying to control the local environment for the rooms with piano's and other instruments in them. I think this is what you would have to do if the concern was broader than just the health of the piano, which it might be, especially if there is an instrument storage room. Probably the lowest first cost alternative would be to flip the coil modules to put the heating coil in the reheat vs. preheat condition for the zones where this is a concern. This step, along with a control sequence change, would allow you to dehumidify. Supplementing this with a zone level humidifier would allow you to add moisture when needed.

One of the down sides to this approach would be that the humidity you added would not be retained in the zone due to the air changes provided by the air handling system. In other words, probably 4-6 times an hour, the air in the room is replaced with air that comes from the central system. Since the central system will likely be running an economizer process that is driven by the need to cool other zones, not the need to maintain humidity, then it will be dryer than you need and you will constantly have to add moisture relative to what would happen if the zone with the piano was self-contained and the only air you added or removed was what was needed for ventilation purposes. So bottom line is the constant air changes would tend to drive the resource costs up. Exactly what they would be would be a function of the humidification technology you used. (I have attached an article I did that looked at the common technologies out there).

You could mitigate the resource consumption a bit by taking some added first cost steps. Specifically, if you converted the piano zones to little self-contained AHUs, either by replacing the coil module with a fan powered terminal unit with heating and cooling coils and a humidifier, or by adding a fan, return connection, filter, humidifier, and volume damper on the incoming air to the existing zone coil modules along with flipping the module (probably more cost effective).

This would let you retain a lot of the moisture you added, especially if you used demand controlled ventilation to manage the primary air from the main system. It would add a bit of complexity and potentially some fan energy, but the latter may not be that much of an addition since you would reduce the fan energy requirement from the central system since the flow it would be required to deliver to the zone would simply be the ventilation air vs. the total flow. That would be a bigger deal in an environment where the economizer process was on MOA a lot of the time (meaning you would have to up the flow to the zone as a function of the OA % when you were not on 100% OA). But in your environment, you are likely at or near 100% OA (i.e. looking like a make-up AHU) a lot of the time.

This is also the approach you would use for the auditorium, where the only zone is the area with the piano; i.e. you would modify this system to include the ability to humidify and dehumidify. An interesting option in this regard would be to look and see if this system has the capacity and physical location that would let it, once modified, also serve the other areas with piano's and/or instrument storage. It just may be that it does, especially if you made it variable volume given the fact that the auditorium is mostly unoccupied. Meaning, most of the time, its capacity could be diverted to the other zones while also holding the auditorium where it needed to be. When the auditorium was in use, it would become the favored zone. But since you would likely only be talking about under-serving the other zones for a couple of hours at most, then the thermal/moisture inertia would probably carry you through, especially since the air that was circulated would be air from the humidified/dehumidified areas only, thus tending to preserve the targeted dew points. As I type this, it seems to me that this could be the most viable option as long as the other zones are not totally spread out from the auditorium area which they might be). Of course, you maybe could address that by shifting things around so the areas with piano's and instrument storage were in the same vicinity as the auditorium.

If you were starting from scratch, you probably would set up a separate AHU to handle the zones where you were storing instruments and had pianos, which is what that last option is basically looking at doing in a retrofit manner.

No matter what, if you choose to manage the environment for the critical instruments more closely, then there will be some added maintenance cost, mostly associated with the humidifiers. Of course, there would technically be less time required to maintain the instruments, so it could be a net win, especially if the piano technician were willing to undertake the maintenance of the humidifiers. I think that would be a good plan since he would have a vested interest in their operation and thus, be likely to make it a priority.

Back to Contents

Social Sciences 2

- Clearly, there were some O&M opportunities in the Social Sciences building in terms of the condition of the relief dampers in the one AHU. I'm not sure how significant the actual energy savings associated with fixing it might be since the system is seldom on MOA I suspect and if the relief damper had been fully closed and the system needed the OA, it would have just pulled it through the OA intake. That said at some point, if you are going to have a building, you probably need to do the basic maintenance required to provide the intended level of performance, which is why you elected to put in an HVAC system in the first place. If the over-arching goal of the HVAC system was to save energy, then you would not have put one in because having one will use more energy than not having one.
- The negative relief plenum may or may not have been a clue regarding the performance of the fan tracking strategy. In other words, depending on a number of factors, the discharge of the return fan could actually be negative relative to the outside under certain operating conditions depending on the pressure drop through the return dampers, the OA requirement and the pressure drop at those dampers, etc. So the EBCx target may be to functionally test the fan tracking strategy to see if it is working as intended. If you want to understand the nuances of the various strategies, there is a link to a guideline on the topic that were are involved with developing.

  https://docs.google.com/folder/d/0B_dloliwnT1ySmhHcmNTUzd1b0k/edit?usp=sharing

- The chiller short cycling issue was an interesting intellectual exercise for the class and exposed them to hot gas bypass and unloaders. But it sounds like you have a "plan B" in the works that will deal with it. If you didn't then I think testing to see how much of a flywheel existed and maybe doing a flywheel cycle like the one I have discussed for Le Conte Hall at UCB would be an option. And it may be an option in other buildings on the campus if you have that issue. I discuss the test strategy and an overview of the results in my NCBC paper from last year, which you can link to from the blog.
- We may have talked about it and I just don't remember, or you may already be doing it, but if not, zone level scheduling may be a viable opportunity in this building. Meaning that you would totally shut down the flow with the VAV terminal damper when a zone was not in use vs. letting it drive itself to minimum flow. If you did it, you would need a set-back/set-up strategy to pop the damper back open if the zone got out of specs. And, if a lot of zones were unoccupied, you would be pushing the central system and the utility systems serving it to a place they may not have been before, which could uncover some other things that need attention.

Back to Contents

Central Plant

- I was surprised when Mark's lab group only had 4 gpm of flow on the central CW system side of the heat exchanger with the control valve wide open. That could actually be a good thing, meaning that the control system is managing the main pump(s) at the central plant so the head generated is just what is needed to satisfy a load with its valve wide open. But I think someone said that the valve was manually locked wide open, so that could be an indication of an issue with the main plant pumping capacity or control algorithm.

  My notes on the original system diagram I did (nearly 10 years ago now; I'm feeling a bit old) say that the design intent was to deliver 2,500 gpm at 100 feet of head with one pump. The second pump was 100% standby. But, if my notes are correct, the mains are 14" diameter, which means they could support about 6,000 gpm of flow as illustrated by this screen shot from my iPhone pipe sizing application.

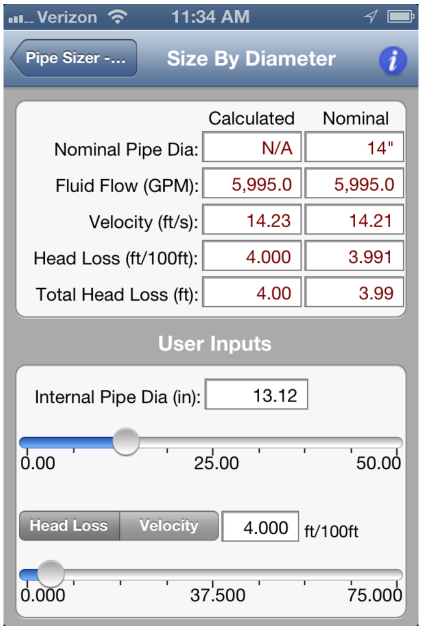

So, I think that means you have significantly more capacity available than 2,500 gpm. You may already have realized that and be running both pumps some of the time; I forgot to look while I was there.

Granted, that is not redundant capacity, meaning if you had 6,000 gpm of critical loads and lost a pump, you would be in trouble. But I suspect that since the head required by the system varies with the flow, you have significantly more than 2,500 gpm of redundant capacity available if one pump were to fail since the other pump would run out its curve and deliver more than the design flow, just at less head. But then, at less than design flow, you need less than design head.

All of that could be established by pump test. The bottom line in the context of my original point is that there is probably an opportunity of some sort that is being revealed by the lack of flow with a wide open valve that Mark saw in his lab.

- Patrick showed me the humidity sensor that you are using as the input to your optimization algorithm. I believe it was one of the nicer Vasalia devices. But I just wanted to be sure you realized that you still need to do pretty regular maintenance on that device to ensure its accuracy. And, if it is of the level of quality I think it might be doing a field calibration may be difficult. So, it may make sense to actually have two devices that you swap in and out so you can send one back for factory calibration. You may already be doing this or something like that, but I just thought I should mention it since if the sensor is not accurate, you could end up "shooting yourself in the foot" in terms of energy savings.

Back to Contents

Hope this is useful information vs. too much information.

Thanks again for sharing your site with the class,

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/
View *A Field Perspective On Engineering* and past posts from
A Field Guide for Engineers at http://av8rdas.wordpress.com/