An IT Approach to Optimization and Diagnosing Complex System Interactions

By Interval Data Systems, Inc.

Bill Gnerre, CEO
Greg Cmar, CTO
Kevin Fuller, EVP

Abstract

Most facilities personnel do not know what is really happening with their building operations. They can't because they see only a tiny fraction of the operational data that would give them real insights. Today, optimization happens at a component or sub-system level with little or no information on how the parts work together. Diagnosing complex problems is often impossible because you cannot see how all the affected systems interact. Actual hourly operating cost data, and any ability to see how control changes, weather, occupancy, etc. influence costs is simply beyond most organizations.

This case study of a hospital in the southeastern U.S. illustrates how an IT approach, which makes all operational information available, quickly uncovers the interdependencies of the chilled water plant and building systems. The diagnosis explores VAV box, AHU, secondary pump, and chiller operations. Interval data provides an in-depth look at the unintended side effects of control engineering, operator decisions, and resultant energy waste caused by addressing the wrong issue—and the verified savings from addressing the correct issues.

Critical to identifying and solving challenges such as these is the availability of continuous, historical operational data from both the plant and buildings. This paper describes how data is used to show cause-and-effect relationships between (and within) building and plant operations, and unnoticed system instability. It also demonstrates how continuously available data reduces the time demands for operations and maintenance, making such analysis both feasible and cost-effective.

Introduction

A recent news story told how UPS is saving millions of dollars by reducing gas consumption of its delivery vehicles. They use data to analyze dispatch plans and optimize delivery routes and times, saving UPS almost 14 million gallons of fuel annually. It's the data that makes the improvements possible, and enables them to measure and verify results (such as reducing mileage by 100 million miles) as well. Mike Eskew, UPS's chairman and CEO commented, "At UPS, we're never satisfied with the status quo."

What does UPS's gas consumption have to do with your HVAC system? A lot. It's all about using the data to improve operations.

1. Both UPS and your facilities operations group are service businesses. Whether formally structured that way or not, facilities operations is a services business, and the more it is run like a business, the better it will perform.
2. Both strive to provide a high level of service (package delivery or building comfort) that consumes a substantial amount of energy to deliver.

3. Both have customers that take the service somewhat for granted, expect nearly perfect reliability, and give little thought to the costs incurred to make it happen.
4. Both employ technicians/mechanics to keep their equipment (delivery trucks or HVAC systems) running, and benefit if the personnel can work more efficiently and proactively leading to labor savings and extended equipment life.
5. Both can optimize performance, using data and information technology (IT) to calculate detailed operational costs. But, in all probability, UPS does a much better job of this than your facility.

Yes, there are some differences. UPS uses IT to leverage their operational data to great benefit. Most facilities operations groups don't. Without the data, you're guessing, and optimization is not possible on any meaningful scale. Without the data, you're accepting the status quo.

Background

Before we get started, here's some background information to help give context to the diagnostic information presented later in the paper.

The Facility

The hospital in this case study is a 2,000,000 square foot facility. The hospital campus includes a half dozen buildings with a total of 115 air handlers served by the physical plant. The plant, partially shown in Figure A, has a total of 5,700 tons of chilling capacity, a bi-directional bypass, and a secondary loop with four 75hp pumps that service seven zones. The chiller plant also includes three cooling towers and a 400 ton absorber with its own cooling tower.

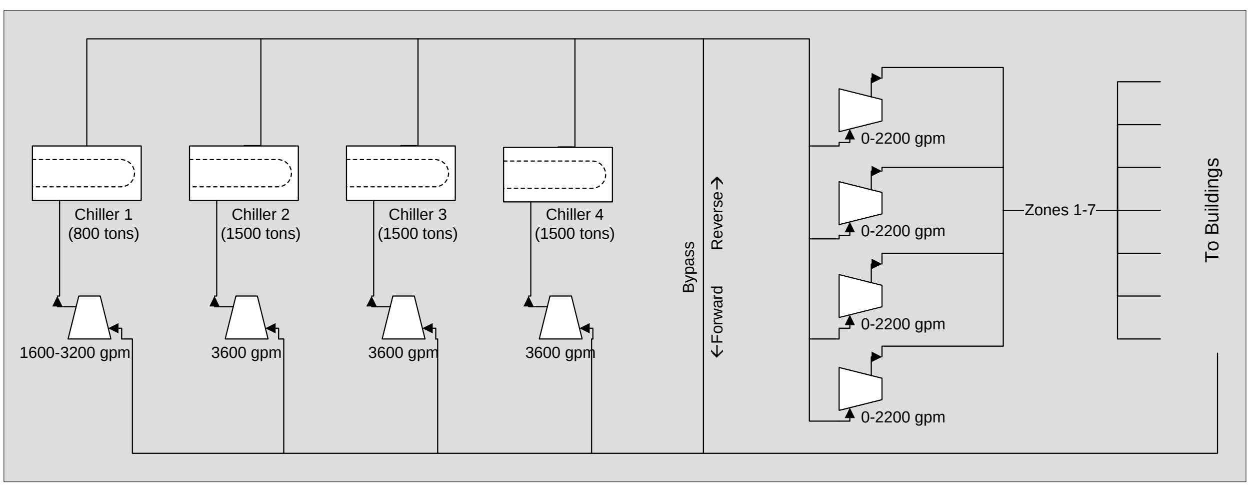

Figure A: Chilled Water Plant—Chillers with Primary and Secondary loop.

In Figure B you can see the standard configuration for an air handling unit. A VFD fan provides the air flow and two cooling coils chill the air. Dampers regulate return air, exhaust, and outside air. This AHU has 24 VAV boxes, each with a reheat coil.

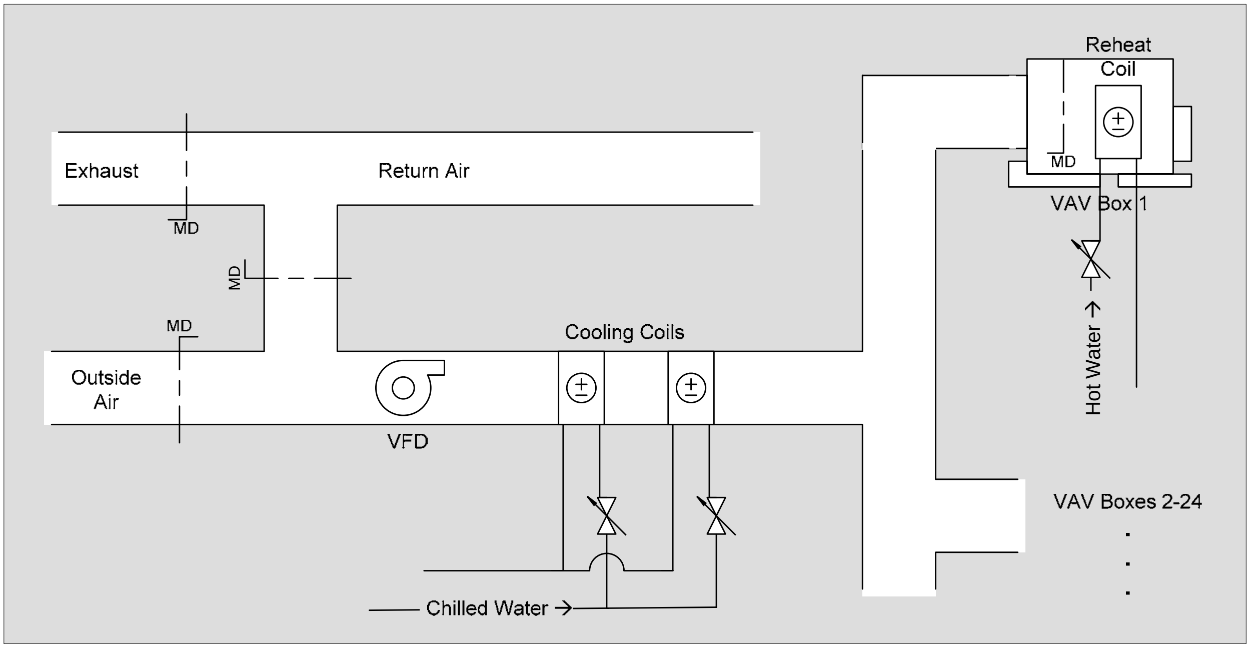

Figure B: Air Handling Unit and VAV boxes.

Data Collection

A building automation system (BAS) will contain thousands of monitoring and control points (over 100,000 for larger campuses). You need them all. The hospital is collecting data from about 18,000 points (see Table 1 for a partial listing). Ideally, you want data not just from the BAS (although it's a better place to start than meters), but from every data source:

- Building automation systems (all of them if you have more than one)
- Advanced meter systems
- Manually read meters
- Utility interval data
- Utility billing data (including all rate information)
- Weather (from local airport as BAS weather sensors are often faulty)

The hospital collects each point at standard 15-minute time intervals to enable synchronized views, for example, showing chilled water supply temperature, AHU discharge air temperature, and outside wetbulb at the same time. Saving the data indefinitely (for the life of the equipment, 20 years or so) allows the review of past operations and provides a historical record of all the equipment in the facility.

Table 1: Sample of Interval Data Points from Hospital Systems

System	Partial Points List
Chillers & Absorber	tons output; CHW flow, setpoint, supply temp, return temp, ΔT*; CW flow, supply temp, return temp, ΔT*; % full load amps; efficiency*; kW/ton*; production balance*; pump brake horsepower*
Cooling Towers	CW supply temp, setpoint, low limit, high limit; CW return temp; CW flow, % full load amps; fan speed, load, kW*, kW-hour; fan motor speed, volts, amps, estimated make-up water*; estimated make-up water BTU*

System	Partial Points List
Secondary Pumps	CHW supply, return, & differential pressure; zone differential pressures & setpoint; average valve % open; secondary CHW supply temp, return temp, flow; bypass flow & direction; kW*;
AHUs	return air humidity, humidity setpoint, temp; discharge air temp, setpoint; cooling valve positions; outside air CFM, CFM setpoint, damper position; supply fan static pressure, setpoint; average VAV heating valve position
VAVs	zone temp, heating limit, cooling limit; discharge air temp; heating valve position; CFM min, max, heating, calc & actual; damper position
Weather	outside air temp, humidity, wetbulb

* – calculated point

The Impact of a Single Plant Control Decision

There's a saying in this business that chiller plants are run by legend. Unfortunately there is a lot of truth to it. Instead of a disciplined engineering process, based on complete operational data, facilities rely on estimates based on a tiny data sample and educated guesswork.

Nowhere else are decisions with such a large financial impact based on so little information. Even small decisions get more thorough analysis. Imagine buying a computer only knowing the processor speed and hard drive size—it would be a crap shoot. You wouldn't make a $2,000 PC decision not knowing if the screen was large enough to be readable, if there was enough memory to handle your workload or just enough to boot the operating system, if it came with an operating system at all, or even if it were a laptop or desktop. But facilities groups make decisions with 100 times the financial consequences of buying the wrong PC every day as part of standard operating procedure.

Facilities business decisions don't have to be left to random chance. You can put the engineering back into your operations and make informed decisions based on a complete set of facts. You do it with data and an IT approach that turns the raw system data into actionable information.

Inside the Chilled Water Distribution Loop

During an investigation of the chiller plant, it was found that the secondary chilled water flow exceeds the primary flow quite frequently. When this occurs, the bypass flow reverses and the secondary pumps, of course, work harder creating the additional flow. This one situation has an impact across a wide assortment of air handling units—this paper focuses on just one.

Experience has shown that *every* time you collect data and thoroughly examine it, you will find scores of problems. It has also shown that building HVAC systems are complex, interactive, and self-compensating, requiring a systemic view to really identify root causes of problems. Once you have a starting point, diagnostics is a stream-of-consciousness process that requires immediate interactive access to data or else the effort gets derailed. You'll see this in practice in the sections that follow.

Note: the starting point, in this case chilled water flows in the distribution loops, is not particularly important. It's just where this diagnosis started based on what was noticed first. The data exploration through a business-oriented IT system is the important part, and would have led to the same conclusions if the process had begun inside the air handler or with the VAV boxes.

Figure C illustrates the flow problems that are commonplace. You can see the primary chilled water flow drop suddenly at the same time the secondary flow jumps. At the same time, the bypass flow increases and reverses direction.

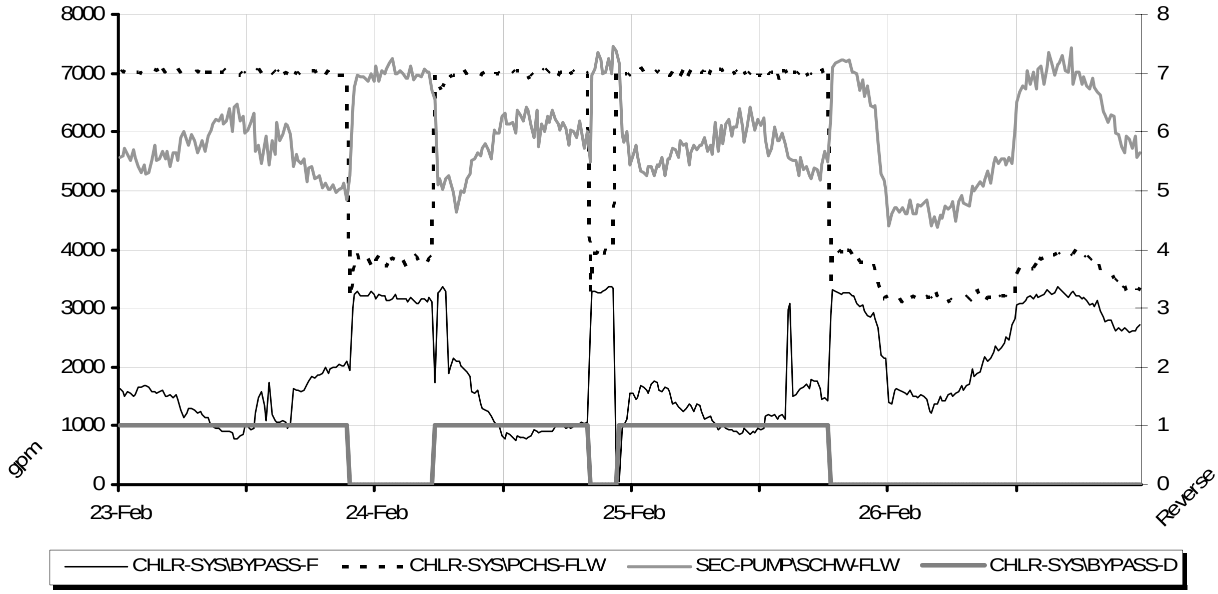

Figure C: Coincident Flows in System—the Primary CHW Flow (dotted line) drops as a chiller is shut down, and at the same time, both the Secondary CHW Flow (gray line) and Bypass Flow (thin black line) jump. You can also see the Bypass direction reverse (1 is forward, 0 is reverse).

There is a cost associated with this behavior—that of the secondary pumps that are now working overtime. The pumps are maxing out as can be seen by the total pump kW. This isn't as costly as it might be, as this is all happening during off-peak hours when electric rates are lower, but the consumption is notably up.

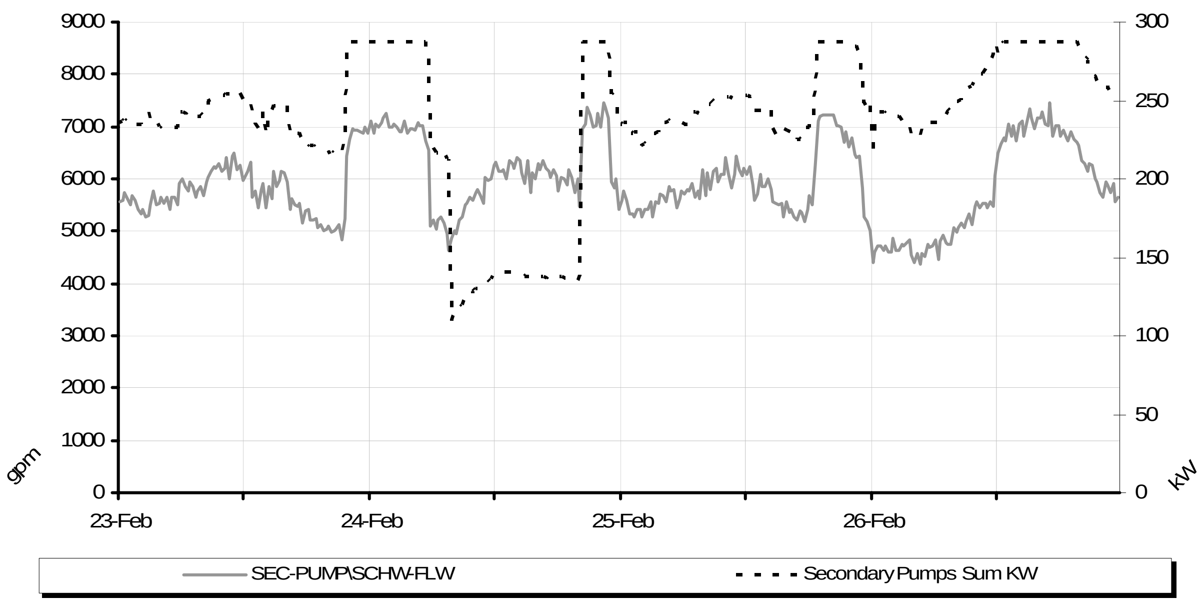

Figure D: Pump Electrical Consumption Maxes Out—each time the Secondary CHW Flow (gray line) jumps to 7000 gpm, the total kW (dotted line) of the pumps reaches their maximum.

What is causing the demand for greater flow? The air handlers start to tell the story.

Cooling Valves and Discharge Air

Looking at data for one of the hospital's rooftop units during the same time period shows more details. The cooling valve opens to 100% instead of the 45 – 60% range where one sees it at other times (Figure E). This demands more chilled water, making the secondary pumps work harder to deliver the increased flow. The correlation between this single valve and total secondary system water flow indicates that the issue is systemic.

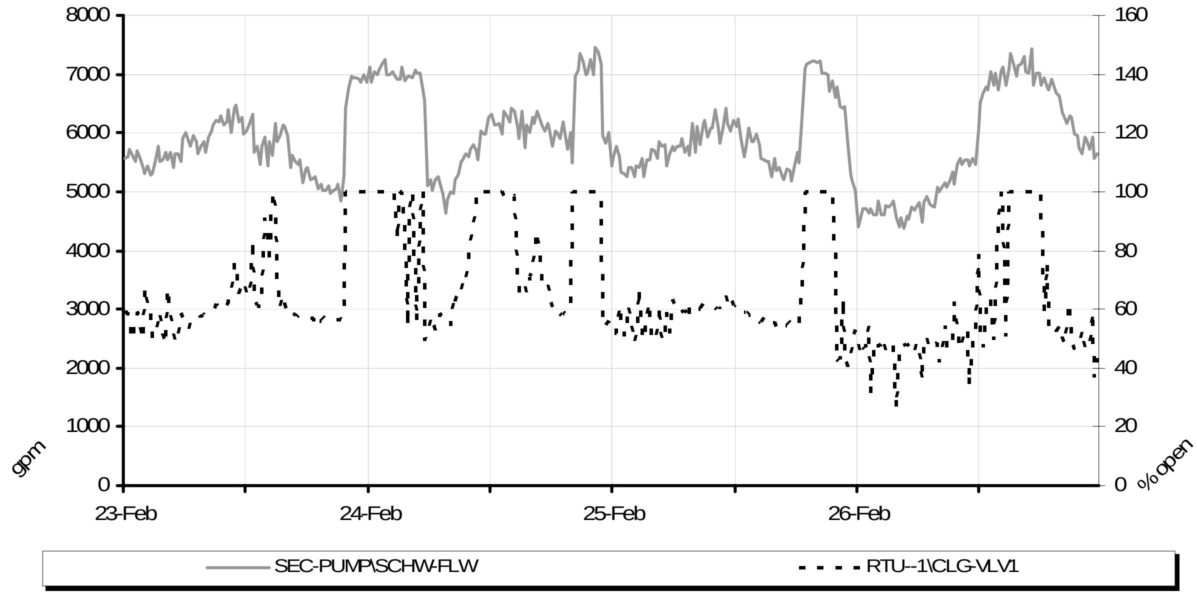

Figure E: Cooling Valve Opening 100%--the Secondary CHW Flow (gray line) is peaking because the Cooling Valve (dotted line) is 100% open.

The cooling valves need a reason to open up, and that reason is the discharge air temperature (Figure F). It fails to meet its 55°F setpoint at the times in question, and calls for more cooling, causing the cooling valves to open to meet that request.

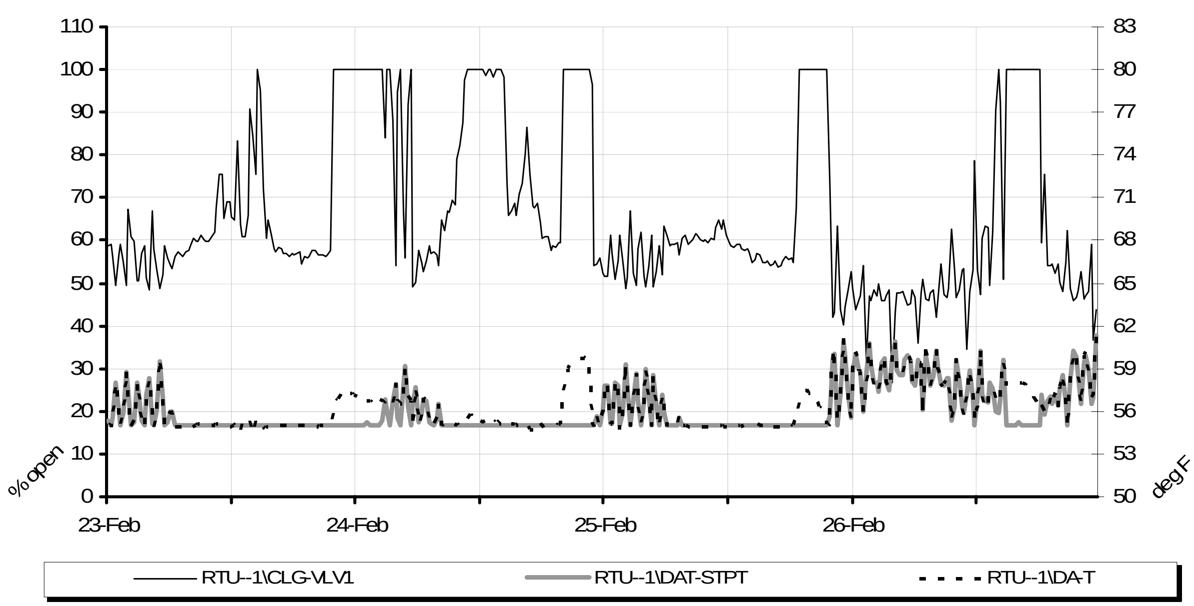

Figure F: AHU Discharge Air Can't Make Setpoint—the Cooling Valve (thin black line) is at 100% open because the Discharge Air Temperature (dotted line) does not maintain Setpoint (gray line).

All of this activity is happening at night or on the weekend. The outside temperature and the building load are low most of these times. So why can't discharge air meet setpoint?

Chilled Water Temperature

In order to determine what was causing the discharge air to exceed setpoint the chilled water temperature was examined. Figure G shows the discharge air temperature and setpoint and the secondary chilled water supply temperature. At the times in question, the secondary system is delivering 47 – 50°F water instead of around 42.5°F, where it stays most of the weekday hours.

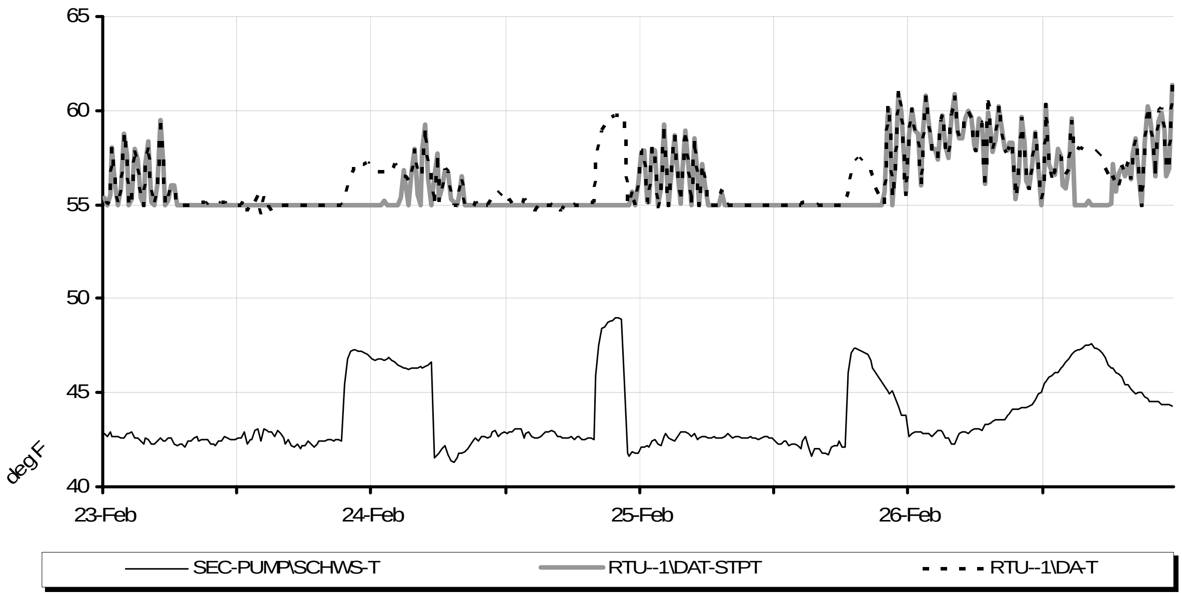

Figure G: Secondary CHW Supply Temp Rises—the Discharge Air Temperature (dotted line) does not maintain Setpoint (gray line) when the Secondary CHW Supply Temperature jumps.

Chiller Operations

February in Florida does require cooling, but the load is not high, especially nights and weekends. During those off hours, the hospital ran only one chiller, as it was assumed it would save money. Figure H shows that every time the system cuts back to one chiller, all the other issues start to arise.

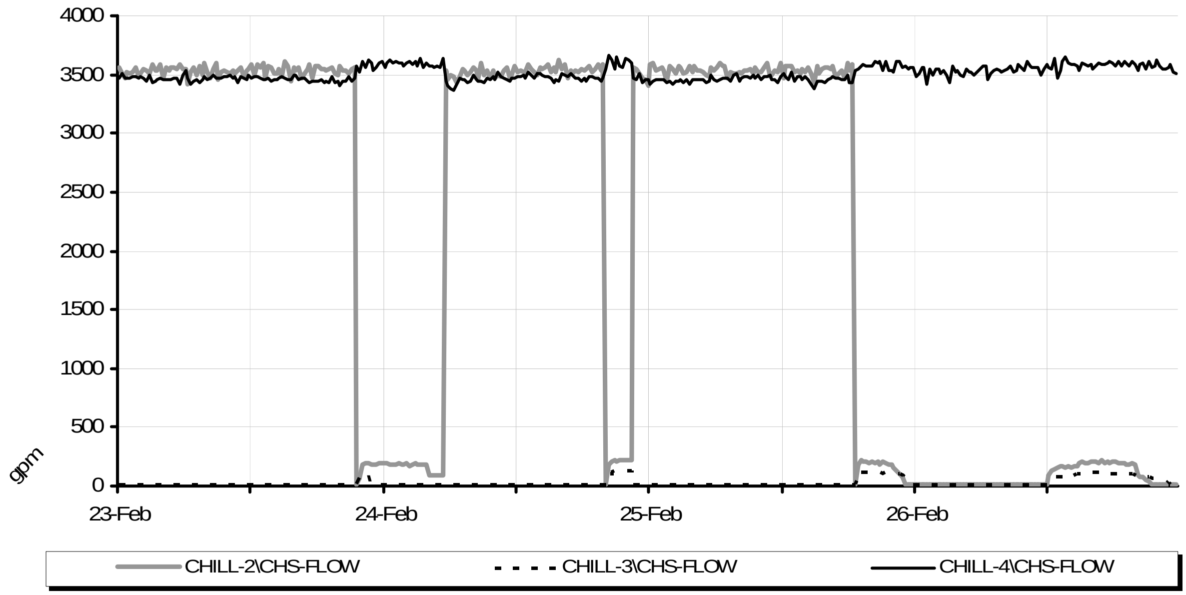

Figure H: Chiller Flows—Chiller 4 (black line) stays on at a fairly constant 3,500 gpm flow, while Chiller 2 (gray line) is turned off during the times in question. Note the small flow amounts for Chiller 2 and Chiller 3 (dotted line) as there is some drag flow through those chillers.

Now look at the supply temperatures in Figure I. The chiller 4 and primary loop supply temperatures hold tight, as they should, except in our problem areas. At those times the primary loop loses about 1.5°F. This is due to the drag flow through chillers 2 and 3 (Figure H), where 200gpm of warm return water is mixing with the properly chilled output from chiller 4.

The additional 2°F jump between primary and secondary supply temperatures is the result of the reverse bypass flow. When the secondary loop flow is greater than the primary loop, the bypass reverses direction (Figure B) mixing return water with supply water, increasing the temperature.

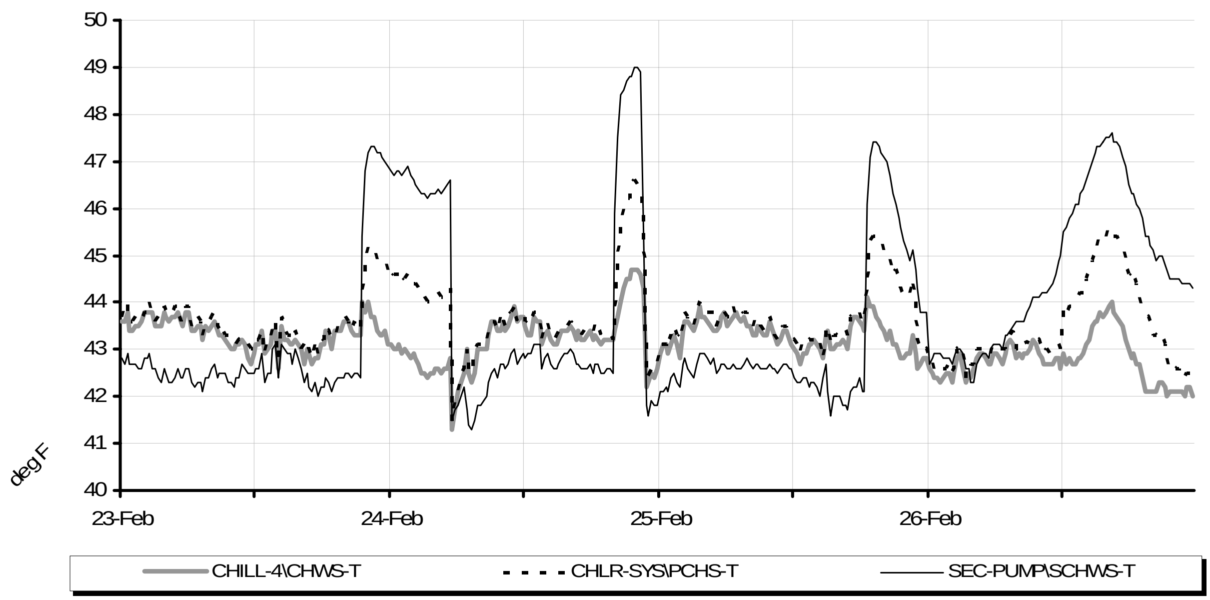

Figure I: Chilled Water Supply Temps—there is an increase in chilled water temperature as it circulates through the system. The Chiller 4 Supply Temperature (gray line) is the lowest, with the

Primary CHW Supply Temperature (dotted line) above it, and finally the Secondary CHW Supply Temperature has the highest temperature.

That circles us all the way back to Figure C where you can see, of course, that going from running two chillers to one cuts the primary flow in half since chillers 2, 3, and 4 are all fixed-capacity 3,600gpm units (Figure A). When the chiller shutdown occurs, even though the space load is low, there is a chain reaction due to the interdependencies of the chiller plant, chilled water loop, and air handlers.

Did the Strategy Really Save Money?

Throughout this little ride, the VAV boxes never noticed anything going on, and operate steadily (Figure L). One could argue that saving $34 – $37/hour (the cost of running one of these chillers during off-peak hours) is worth a little instability as long as space comfort isn't affected. But the truth is that the savings are only a fraction of that.

Figure J shows the hourly cost of operations for various plant components and the total plant cost. While you do cut the cost of one chiller, and you get an extra $2/hour in reduced cooling tower operating cost, the remaining chiller is working harder, costing $20/hour more to run than when running as one of two (A). Then when the secondary pumps start working harder to compensate, they increase another $3/hour (a 25% jump). In total, there is a small operational savings of about $10/hour (B), but on a ton-hour basis, costs actually increased 7% (C). Factor in the instability introduced and the added wear on equipment—is it worth it?

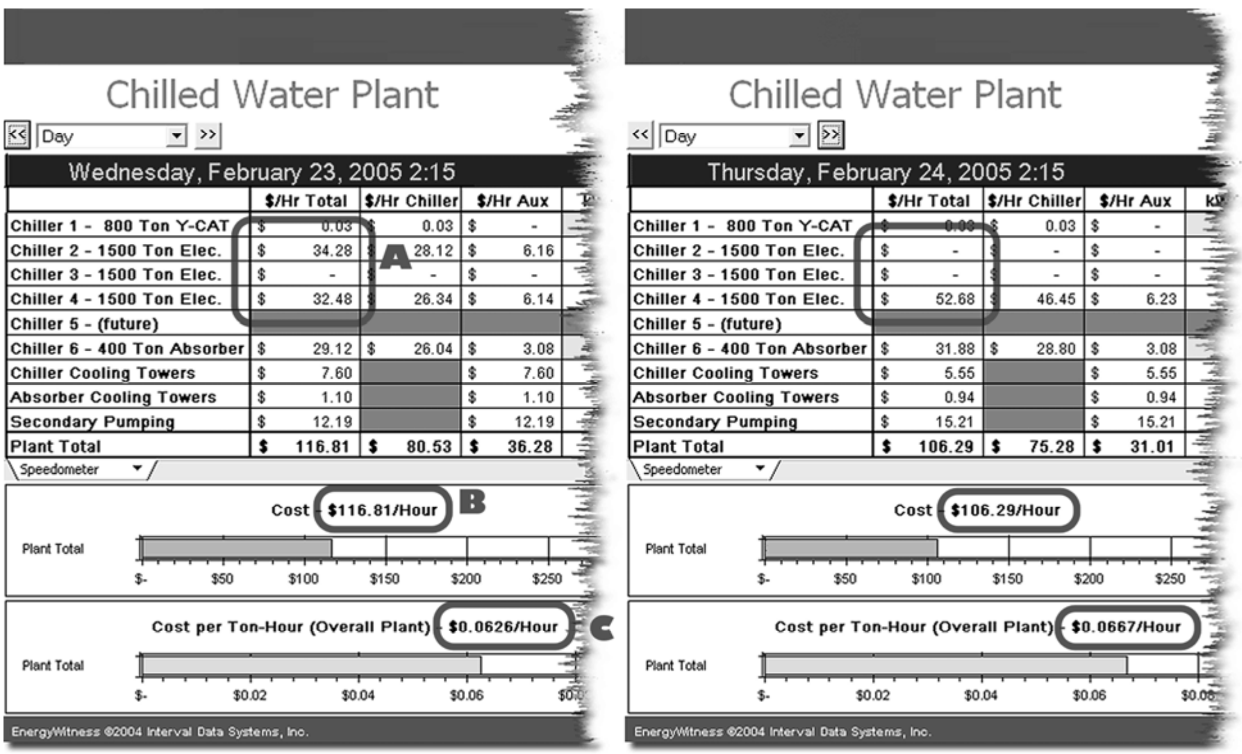

Wednesday, February 23, 2005 2:15

	$/Hr Total	$/Hr Chiller	$/Hr Aux
Chiller 1 - 800 Ton Y-CAT	$ 0.03	$ 0.03	$ -
Chiller 2 - 1500 Ton Elec.	$ 34.28	$ 28.12	$ 6.16
Chiller 3 - 1500 Ton Elec.	$ -	$ -	$ -
Chiller 4 - 1500 Ton Elec.	$ 32.48	$ 26.34	$ 6.14
Chiller 5 - (future)			
Chiller 6 - 400 Ton Absorber	$ 29.12	$ 26.04	$ 3.08
Chiller Cooling Towers	$ 7.60		$ 7.60
Absorber Cooling Towers	$ 1.10		$ 1.10
Secondary Pumping	$ 12.19		$ 12.19
Plant Total	$ 116.81	$ 80.53	$ 36.28

Thursday, February 24, 2005 2:15

	$/Hr Total	$/Hr Chiller	$/Hr Aux
Chiller 1 - 800 Ton Y-CAT	$ 0.03	$ 0.03	$ -
Chiller 2 - 1500 Ton Elec.	$ -	$ -	$ -
Chiller 3 - 1500 Ton Elec.	$ -	$ -	$ -
Chiller 4 - 1500 Ton Elec.	$ 52.68	$ 46.45	$ 6.23
Chiller 5 - (future)			
Chiller 6 - 400 Ton Absorber	$ 31.88	$ 28.80	$ 3.08
Chiller Cooling Towers	$ 5.55		$ 5.55
Absorber Cooling Towers	$ 0.94		$ 0.94
Secondary Pumping	$ 15.21		$ 15.21
Plant Total	$ 106.29	$ 75.28	$ 31.01

Figure J: Hourly Cost of Chiller Plant Operations

This is great. The analysis has shown what occurred, how the plant and building systems interact, how to take the instability out of the operations, and what the cost implications are of these choices. There's only one problem…

…none of this actually fixes anything.

“As Designed” versus “As Needed”

What do we mean that it didn’t fix anything? This is a classic case of treating the symptoms instead of the real root cause. The one-chiller-or-two decision affects behavior as far out as the air handler. But the real problem with the discharge air not maintaining setpoint isn’t too little chilled water; it’s that the AHU’s discharge air setpoint is too low in those circumstances. The issue isn’t originating at the plant, but at the other end of the system… at the VAV boxes.

Space Comfort and Air Flow

Looking at this system from the space end, at first glance everything seems fine—temperatures hold within defined comfort ranges (Figure K) and air flows have steady CFM readings (Figure L). Most engineers/technicians (including some at the hospital) would report that the VAV boxes are operating right on design spec—no problems—since the VAV is supplying the space with proper temperature and air flow. Gathering more data (there are 288 points associated with the VAV boxes for this one air handler) without probable cause is usually too labor intensive to look further.

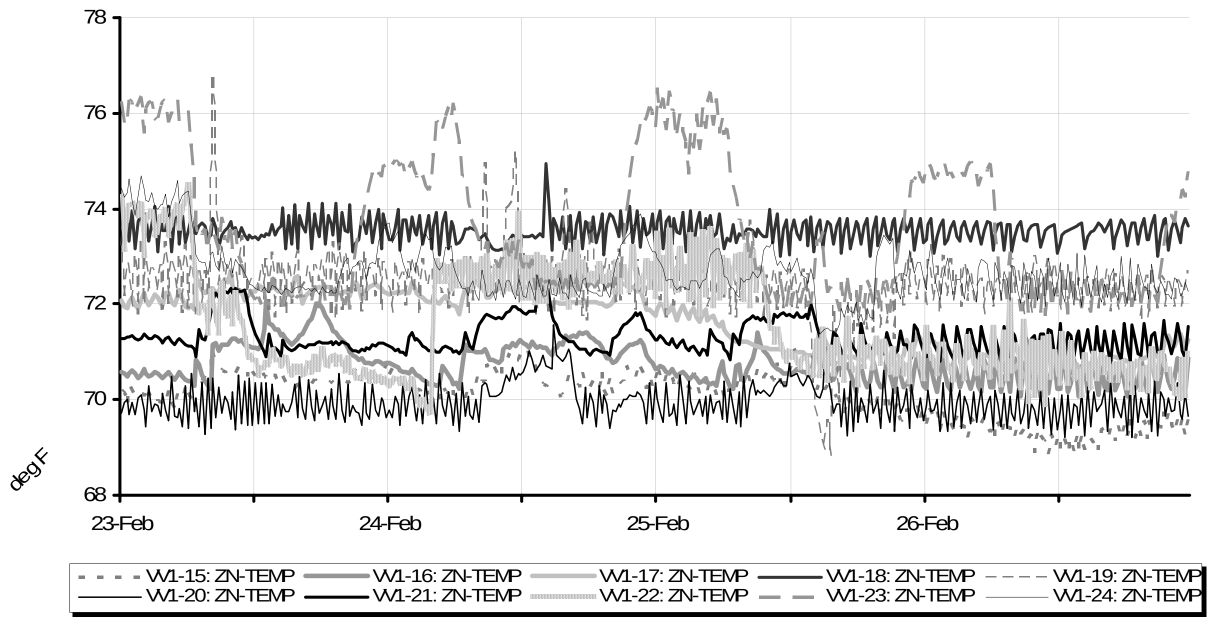

Figure K: Space Temperatures at 10 VAV Boxes—the Zone Temperatures are all fairly constant and within a comfortable range of 69 – 74°F. The one temperature that jumps up to 76°F does so in response to occupants adjusting the thermostat.

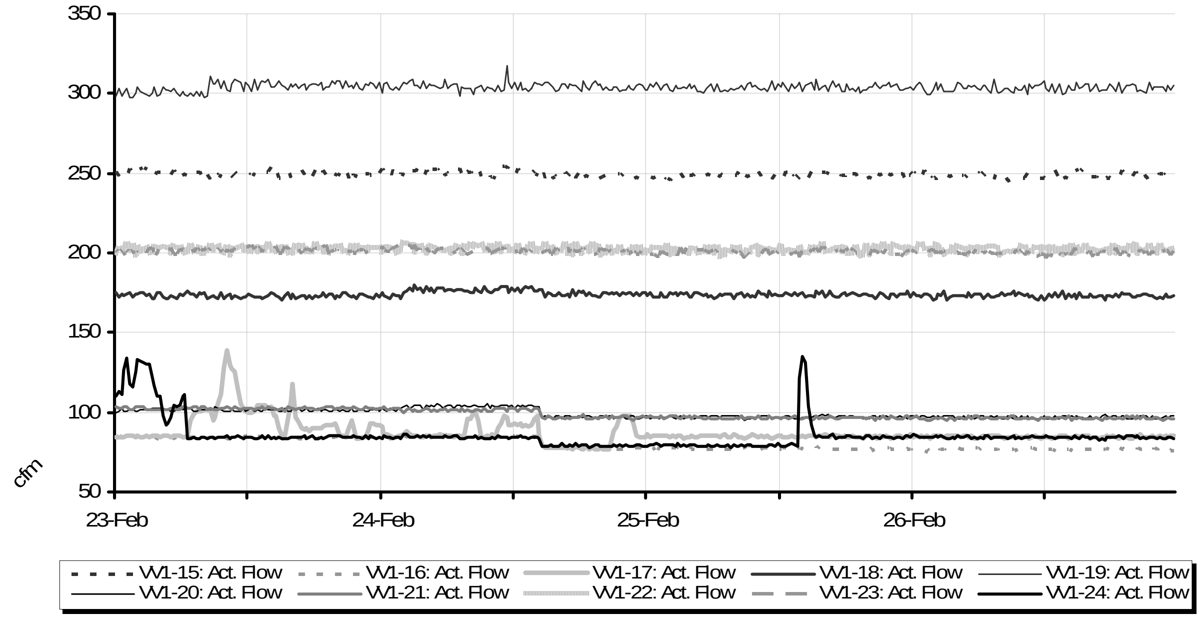

Figure L: Air Flows (CFM) for 10 VAV Boxes—only two of the VAV boxes ever operated above their minimum CFM levels.

"No problems" is an understandable conclusion without looking at the data in detail. Unfortunately, it's also wrong. A properly built IT system has the information to allow a closer analysis of the VAV information, which leads to a better picture of how things should operate.

Too Much Reheat

As can be seen by looking at the discharge air temperatures for each VAV box (Figure M) and the VAV box heating valve (Figure N), there is a substantial amount of reheat happening throughout the building. Of the ten VAV boxes shown in Figure N, only two are not supplying a significant amount of reheat.

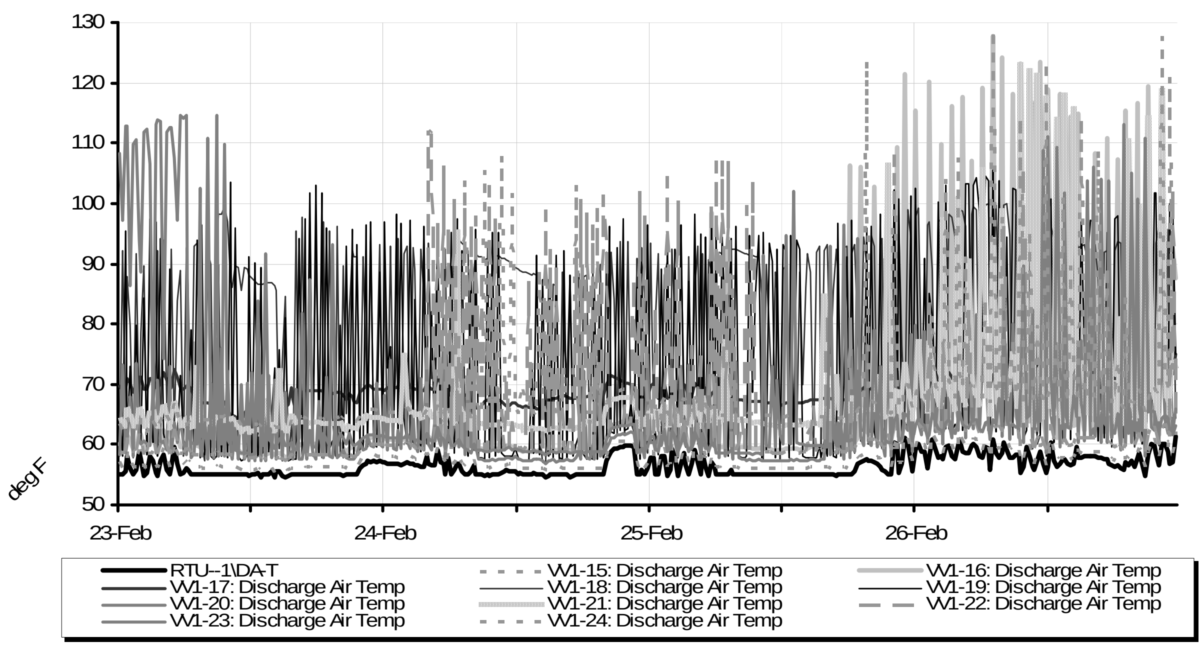

Figure M: Discharge Air Temperatures for 10 VAV Boxes—only a couple of the VAV boxes Discharge Air Temperatures track near the air handler Discharge Air Temperature (black line across bottom), the others are performing significant reheat, but tend to spike rather than provide steady temperatures.

The chiller plant is operating under the false load created by reheating the air that was just chilled. The hospital is paying to cool the air to 55°F in the AHU, then paying again to warm that air to 70 – 100+°F in the VAV boxes.

Heating valves are frequently opening to 100%, even during daytime hours for some VAV boxes.

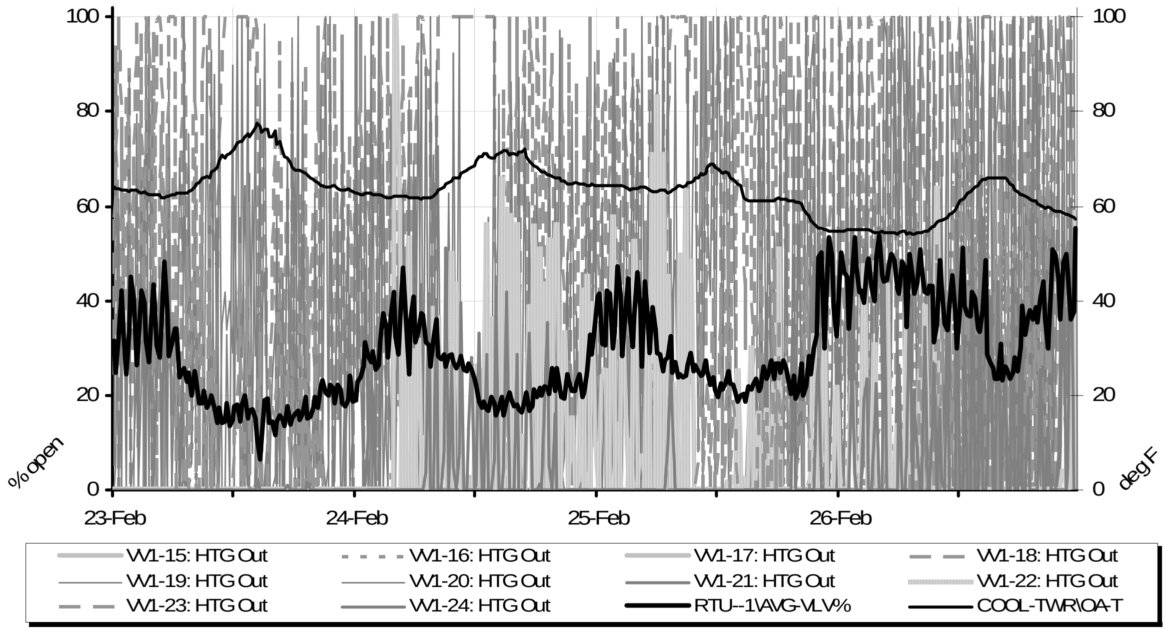

Figure N: Reheat Valve Position on 10 VAV Boxes—many VAV box Reheat Valves are opening to 100% and cycle widely between 0 – 100%. The chart also shows Outside Air Temperature (upper black line) and the Average Heating Valve Position (lower thick black line).

VAV Influence on the Whole System

The VAV behavior, specifically the frequent reheat, is the real culprit in the system instability shown earlier.

There is a lot of information in Figure O that serves to pull the story together. From the AHU there is the discharge air temperature and setpoint, and the cooling valve setting (there are two, operating in unison). In addition you can see the average reheat valve setting from the VAV boxes and the outside temperature.

Even though there is reheat occurring during the day, the outdoor temperature is low enough that the AHU discharge air setpoint stays at its 55°F minimum. The need for 55°F air is based on meeting summer daytime cooling demands, but this data is for February. Even with air flow at minimum CFM, discharge air at 55°F causes reheating to occur, creating a feedback loop demanding more cooling.

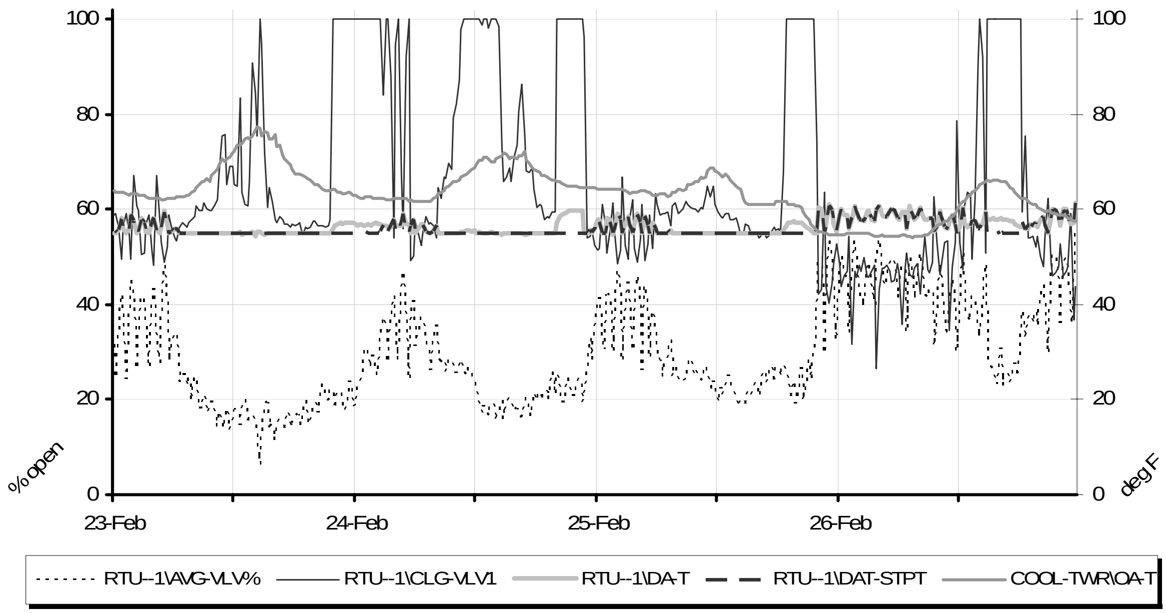

Figure O: Reheat Causes Systemic Instability from VAV to AHU to Plant—Average Reheat Valve Position (bottom, dotted line); Discharge Air Setpoint (dashed line); Discharge Air Temperature (light gray line); Outside Air Temperature (medium gray line); and Cooling Valve Position (top thin black line).

One chiller could not deliver enough cooling to keep the AHU discharge air at setpoint, causing the increased secondary loop flow and bypass reversal shown earlier. Eventually, the average reheat valve positions impact the discharge air setpoint, moving it up 3 – 5°F, although oscillating by about 3°F due to the cycling reheat valve positions. The discharge air temperature reaches setpoint again either when a second chiller comes online or when reheat pushes the setpoint high enough so that the cooling valves don't need to be 100% open and driving the demand on the chilled water loop.

Implementing the Solution and Measuring the Savings

The AHU and VAV boxes needed rebalancing based on actual needs for operating the space properly (as opposed to original design specs). Balancing is an iterative process—make changes, re-examine the data, make more changes as necessary, etc.

The hospital changed the control program on the AHU's discharge air setpoint to stop using the average reheat setting. They employed a simpler approach based on outside air temperature and a higher baseline setpoint. We monitored the affected spaces and also adjusted air flow settings for several of the VAV boxes.

It's worth noting that this whole scenario and strategy is based on winter conditions. The summer will bring a new set of circumstances and needs, which the data will make clear when it happens.

The results of the optimization didn't change space comfort (which was fine and never an issue). It did, however, have a significant impact on the operating costs.

Table 1: Cost Savings Calculation

Temp Range	Differential DA Temp	Hours at Temp	Delta Enthalpy	Avg Tons Saved	Cooling Savings	Heating Savings	Fan Savings	Total Savings
45 - 50°F	4.64°F	314	1.13	967	$70	$80	$17	$172
50 - 55°F	4.74°F	371	1.15	1,170	$90	$97	$20	$208
55 - 60°F	3.55°F	617	0.86	1,457	$113	$121	$33	$267
60 - 65°F	2.54°F	854	0.62	1,442	$112	$119	$46	$277
65 - 70°F	2.29°F	1,005	0.56	1,531	$118	$127	$54	$300
Totals		**3,160**		**6,567**	**$ 508**	**$ 544**	**$171**	**$1,223**

Total system operating cost: $6612 **Percent reduction: 18.5%**

All of the adjustments made were within the buildings, to the air handler unit and the VAV boxes. Yet as you can see from the chart above, the majority of the savings are realized in the central plant, with approximately equal portions coming from cooling and heating savings. While much of the optimization work happens in the buildings, most of the dollar savings occurs in the plant.

Repeat the Process 115 More Times

That optimization did a lot to improve operations and lower costs for the air handler. In the big scheme of things, it is only a start at reducing the hospital load enough so that next winter they could indeed run with just one chiller. Another 115 air handlers need optimization as well.

The implementation of the changes to the rest of the AHUs is still underway as this paper is being written. However, the hospital’s facilities director and staff have a plan in place to prioritize and systematically address each one. The director suspected that issues existed for some time, but prior to this analysis didn’t have the evidence and facts to put an action plan into place.

Using the data we were able to evaluate every AHU, determine if space comfort requirements were being met, analyze the operations (whether meeting comfort levels or not), and make recommendations where needed. The plan exists because the data exists.

Over one third of the 115 air handlers had issues. A partial list includes:

- AHU discharge air temperature is 50°F but cooling valve is closed.
- AHU discharge air temperature setpoint is 50°F. Cooling valve at 100% all the time
- Discharge air temperature resets to 49°F based on humidity, but humidity rises because space is over-cooled.
- Unit provides 51°F discharge air temperature, but not making return air temp setpoint.
- Unit is not making discharge air temperature. Outside air pre-cooling has capacity.
- Unit is not making pre-cooling temperature, but is making discharge air temp. Pre-cooling valve is 100% open.
- Unit is not making return air temperature, set at 72°F, despite cooling valve @100%.
- Cooling valve is locked at 20% open and space is below setpoint.
- Zone temperature is 68°F but cooling valve is closed.

Oh, did I forget to mention, one diagnostician did the entire analysis and report in less than three (count ’em, 3) days.

The Answers Are In the Data

The exploration of the hospital's operations shows two main points:

- First, it showed in detail how the system operates; how changes have a systemic impact since the overall system is self-compensating; and how "as designed" is not "as needed" regarding how the buildings should operate.
- Second, the exercise of looking at interval data, in detail, was the key to unlocking real insights about this facility in an extremely efficient manner.

It doesn't really matter what the questions are, the answers are somewhere in the data. This case—identifying where problems existed, tracking them throughout the facility, measuring the cost impact, finding the true root cause, determining a plan to fix the problem, and measuring the results—took only *three hours* of engineering/diagnostician time. Another four to five hours went into communicating with the hospital over the course of a week. And, it was done without ever stepping foot in the hospital—completely diagnosed from 1,000 miles away using the data.

The optimization discussion in this paper is only the tip of the iceberg in terms of what you can accomplish with an IT approach to optimization and diagnostics. Effective IT systems are driven by data. Data collection is not a waste of time, but does represent a lot of time wasted. Let me explain. Collecting data and having it at your fingertips is invaluable. However, in most facilities today, data collection consumes a huge amount of time by senior staff, a situation that is either not realized or accepted as "part of the job." Conservatively speaking, an engineer is likely to spend 5 – 8 times as long collecting data as doing real engineering analysis. Think of the value your organization would realize if the data collection to analysis ratio were reversed.

Today's HVAC systems are complex and the components are highly interdependent. The self-compensating control logic often masks problems so they go unnoticed or show up as symptoms in an entirely different part of the facility. A complete view of the data from the plant all the way to the terminal boxes, which can track and show the interdependencies, is the only way to effectively manage and diagnose these systems. The artificial walls that organizations create separating the plant from the buildings, or operations from utilities, are worse than counterproductive, they impede the facilities from operating at peak efficiency. Business-oriented IT systems that sit above the engineering/control systems break down those walls of inefficiency.

Operational and cost data are the cornerstones to getting a complete understanding of how your facilities operate. It provides you with the engineering facts to establish a systematic plan to make real progress. And once you look at the data, you will want to establish a new plan—any that you had before will be tossed once you actually see what is going on. Data creates a different world for operations that completely changes what you can do and how long it takes.

Or, as Yoda might say, "The answers you seek, in the data they are."

Citations

Cmar, G.; Gnerre, W. “Defining the Next Generation Enterprise Energy Management System” In *Web-Based Energy Information and Control Systems: Case Studies and Applications*; Capehart, B., Capehart, L., Eds.; Fairmont Press: Atlanta, GA, 2005; Chapter 32.

Cmar, G.; Gnerre, W.; Rubin, L. “Diagnosing Complex System Interactions (from 1000 miles away),” Proceedings from the 13th National Conference on Building Commissioning. New York, 2005.

Cmar, G.; Gnerre, W. “Optimizing Building HVAC through Data and Costs,” Proceedings from the 2005 World Energy Engineering Congress. Austin, TX, 2005.

Mills, E., Friedman, H., Powell, T., Bourassa, N., Claridge, D., Haasl, T., Piette, M., 2004, “The Cost-Effectiveness of Commercial-Buildings Commissioning: A Meta-Analysis of Energy and Non-Energy Impacts in Existing Buildings and New Construction in the United States”

Interval Data Systems, Inc.
104 Summer St. | Watertown, MA 02472 | www.intdatsys.com
tel: 617-744-1091 | fax: 617-744-1092 | e-mail: info@intdatsys.com