



## **Case Study: Diagnosing Complex System Interactions (from 1,000 miles away)**

**Gregory Cmar and Bill Gnerre**  
*Interval Data Systems, Inc.*

**Larry Rubin, CEM, CPE, CHFM, CBCP**  
*St. Joseph's Hospital*

### **Synopsis**

When commissioning existing buildings it is important to extend the scope of ongoing analysis to understand the interrelationship among central plant, secondary system, and terminal system operations. Reacting only to isolated components of the overall system can lead to compensating for poor operations elsewhere within the facility, masking the real root of the problem. It is very easy for plant operations to impact AHU performance, or VAV controls to ultimately affect the central plant. In the worst cases, efforts to reduce energy can lead to a net increase in consumption and cost due to the impact on related systems.

This case study of St. Joseph's Hospital in Tampa, Florida illustrates the interdependencies of their chilled water plant and building systems. The presentation 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.

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, the energy cost impact of control changes, and unnoticed system instability. This case study also demonstrates how continuously available data reduces the time demands for commissioning, and later operations and maintenance, making such analysis both feasible and cost-effective.

### **About the Authors**

Greg Cmar, cofounder and CTO of IDS, is an expert on energy systems and interval data. He has 30 years of experience in facility operations, energy conservation, energy analytics, monitoring and control systems, and utility billing, as well as database, application and automation technologies. Prior to IDS, Greg was a cofounder and director of engineering at ForPower, an energy conservation consulting firm; and held senior roles at Coneco, Enertech Systems, Johnson Controls, the Massachusetts Energy Office, and Honeywell. Greg holds patent #5,566,084 for the process for identifying patterns of electric energy, effects of proposed changes, and implementing such changes in the facility to conserve energy.

Bill Gnerre brings twenty-plus years of technology entrepreneurial experience to his role as cofounder and CEO of IDS. His previous positions include being a vice president at Circadian Software, cofounder of ChannelWave Software, and various management responsibilities at Formtech and Computervision, both CAD technology vendors. Earlier in his career Bill worked in a range of engineering positions and holds a degree in mechanical engineering from Northeastern University.

Larry Rubin has been the Director of Facilities Management at St. Joseph's-Baptist Health Care, which is part of the BayCare Health System, for the past eight years. As the chairman of the eight-member hospital Value Analysis Team for Facilities, Energy and Construction, he has led efforts that have saved the network well over \$1,000,000. He is a member of AFE, AEE, ASME, ASHRAE, and ASHE, and was named *Engineer of the Year* for AFE's Tampa Bay chapter. He holds CEM, CPE, CHFM, and CBCP designations. Larry has a strong power plant/central plant background, which he is utilizing to optimize the new \$10 million central energy plant at St. Joseph's Hospital. Larry also holds a Mass. First Class Engineer's license and has a Masters degree in education.

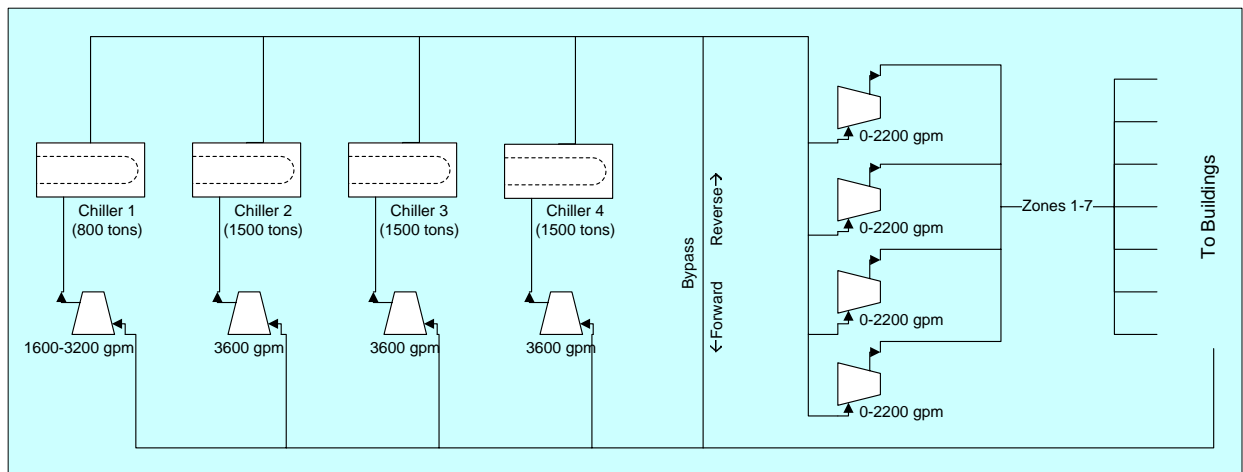
## Background and Introduction

This case study represents chapter two in an ongoing project to redefine how the operations staff at St. Joseph's Hospital in Tampa, Florida manages the facility. The research, analysis, and findings of this paper occurred over a one-week period in March 2005. The implementation and verification of prescribed changes are ongoing at the time this paper is being written (preliminary results will be shown during the conference session).

First, there is a short overview of the facility, and then a brief recap of the previous work and findings. (For a complete case study of "chapter one," see the *Hospital Case Study* citation at the end of this paper.)

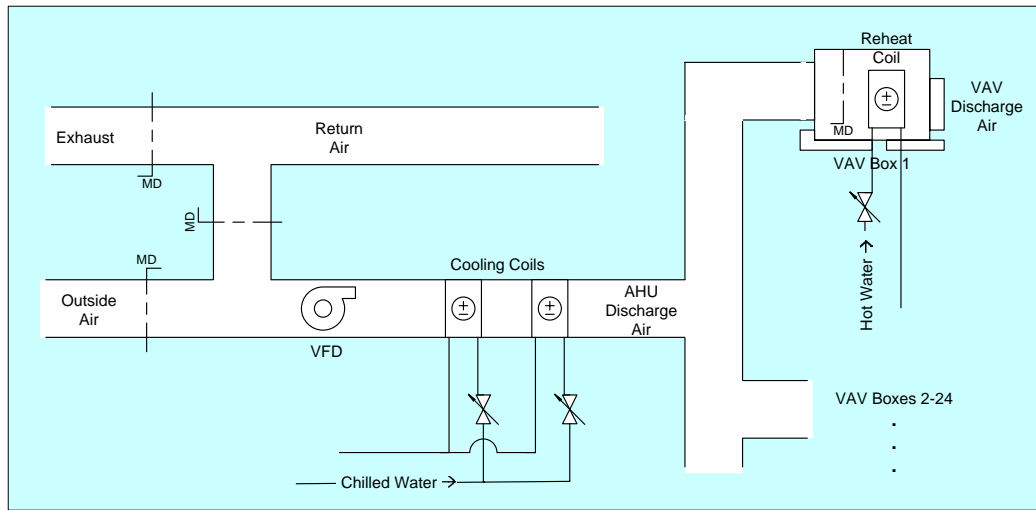
### The Facility

St. Joseph's Hospital (SJH) is a 2,000,000 square foot facility. In addition to the main building there are separate cancer center, children's hospital, and medical arts buildings with a total of 165 air handlers served by the physical plant. The plant, partially shown in Figure A, has a total of 5,700 tons of total chilling capacity, a bi-directional bypass, and a secondary loop with four 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**

## Previous Analysis and Resolutions

The initial project focused on a couple areas: providing diagnostics and data to support or refute a redesign of the chilled water system, and provide a mechanism to show the hourly cost of operations for the plant (which now exists). It took place in two phases. The first, described in the *Operational Instability and Over-Control* section below, was spread out over six weeks during April and May of 2004. Phase two occurred over a four week period in August/September 2004, and is described in the *Engineering Design and Space Comfort Issues* section below.

## Data Collection

The first step was to get the data—15 minute interval data collected 24x7. Initially data was collected from all 800 building automation system (BAS) points within the chiller plant. That number grew to 5,000 points (see Table 1 for a sample listing) as additional data collection began for air handlers, and should exceed 20,000 points once data is collected for all 165 AHUs.

**Table 1: Sample of Interval Data Points from SJH Systems**

System	Partial Points List
Chillers & Absorber	tons output; CHW flow, setpoint, supply temp, return temp, $\Delta T^*$ ; CW flow, supply temp, return temp, $\Delta T^*$ ; % 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*
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*;

System	Partial Points List
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

## Operational Instability and Over-Control

As expected, the first thing viewing the data shows is that some sensors are broken or out of calibration. These were corrected. With properly calibrated sensors reporting, the data showed significant system instability. Output from the 1,500 ton chillers were often cycling by 500 tons every 15 minutes. There was similar instability in the chilled and condenser water temperatures. It was determined that the chilled water reset control was causing the setpoint to oscillate between 42°F and 45°F. Changing the setpoint to a constant 42°F stabilized the system immediately, as was shown by the data.

Another area addressed was the cooling towers. Fans were running 100% all the time. Reviewing the data showed that the condenser water setpoint was set to equal the outside air wetbulb. Since there was no approach, the tower could never actually reach wetbulb, forcing the fans to overwork. This was resolved by ultimately setting a 4°F approach to wetbulb. The hospital also repaired a broken VFD drive in one of the tower cells, spreading the load over six fans instead of only five. These changes resulted in \$40K/year as the fans were able to run at lower speeds and back off during the cooler nights.

## Engineering Design and Space Comfort Issues

This earlier project went on to identify various issues with primary/secondary return temperatures, underutilized gas chiller capacity, primary pump efficiency, primary/secondary flows, and differential pressure. In all, it quantified another \$226K/year in lost savings due to apparent design flaws.

The issues related to the primary/secondary flows, temperatures, and zone differential pressures led to finding space comfort problems. Although alerted to it, tracking them down and resolving this new set of concerns was beyond the scope of this part of the project.

## The Impact of a Single Plant Control Decision

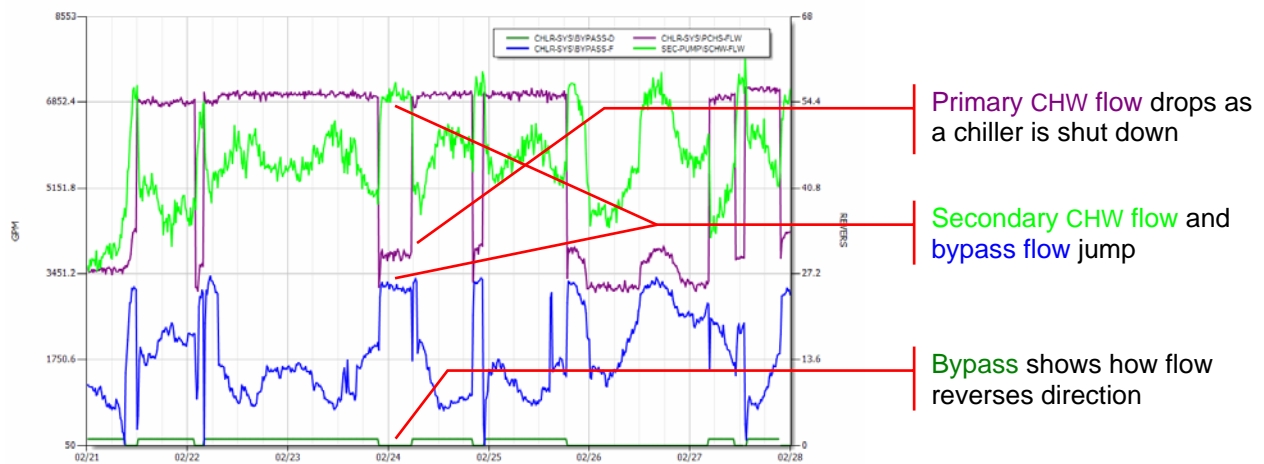
During the investigation of the chiller plant, the investigation found that the secondary chilled water flow exceeds the primary flow quite frequently. When this occurs, the bypass flow reverses and the secondary pumps, of course, work harder creating the additional flow. This one situation has an impact 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 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.

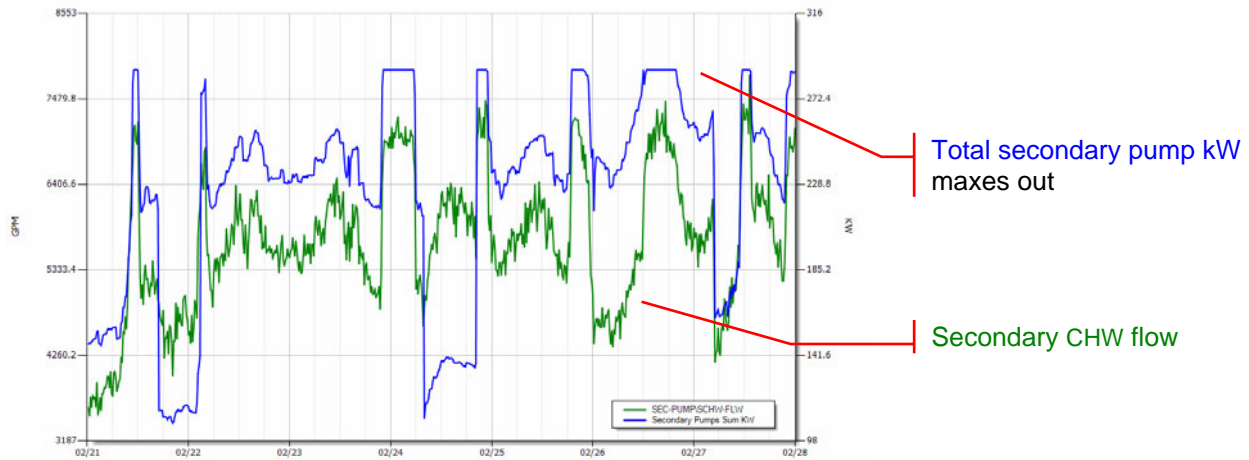
### Inside the Chilled Water Distribution Loop

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**

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, but the consumption is notably up.

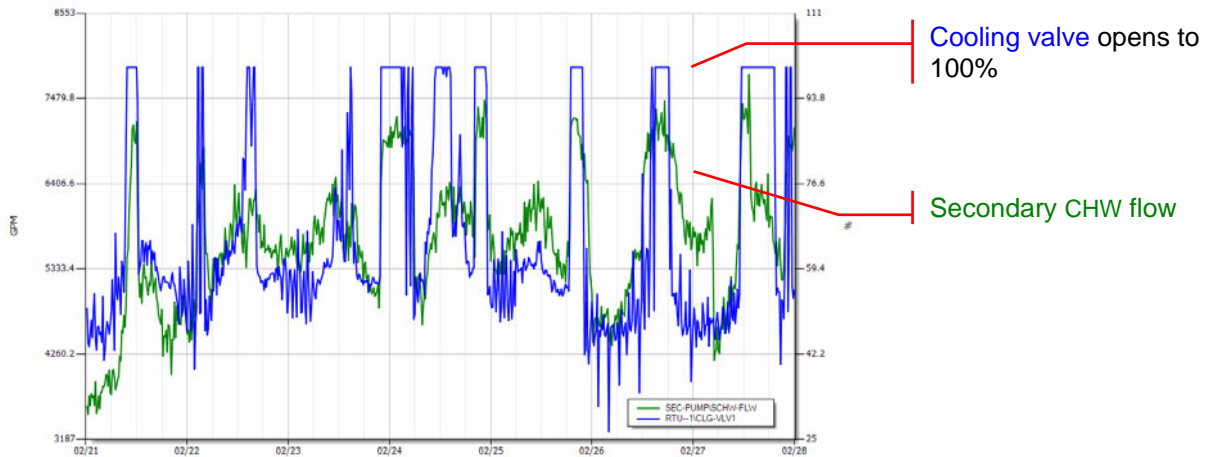


**Figure D: Pump Electrical Consumption Maxes Out**

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

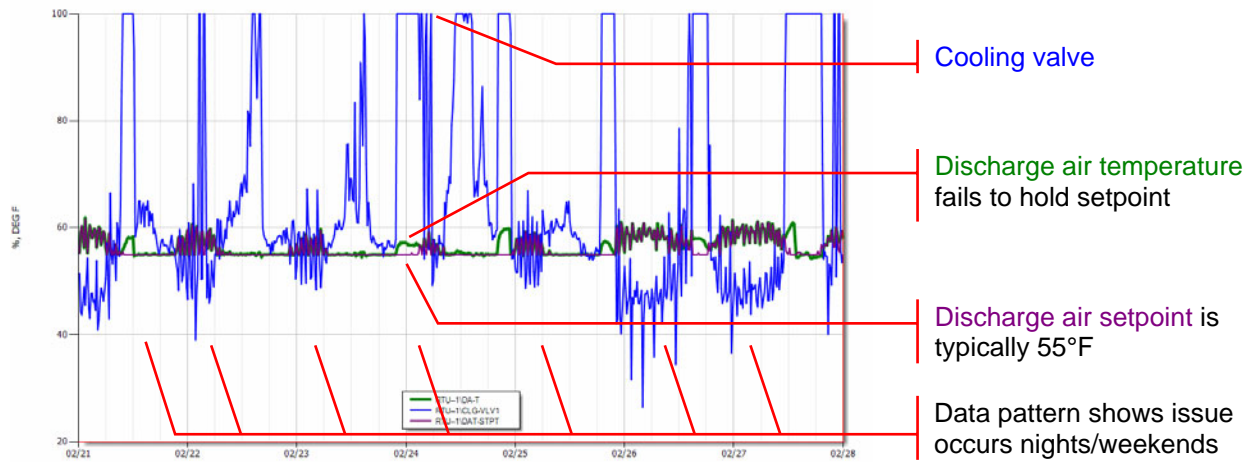
### Cooling Valves and Discharge Air

Looking at data for one of SJH’s roof-top 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 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**

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.

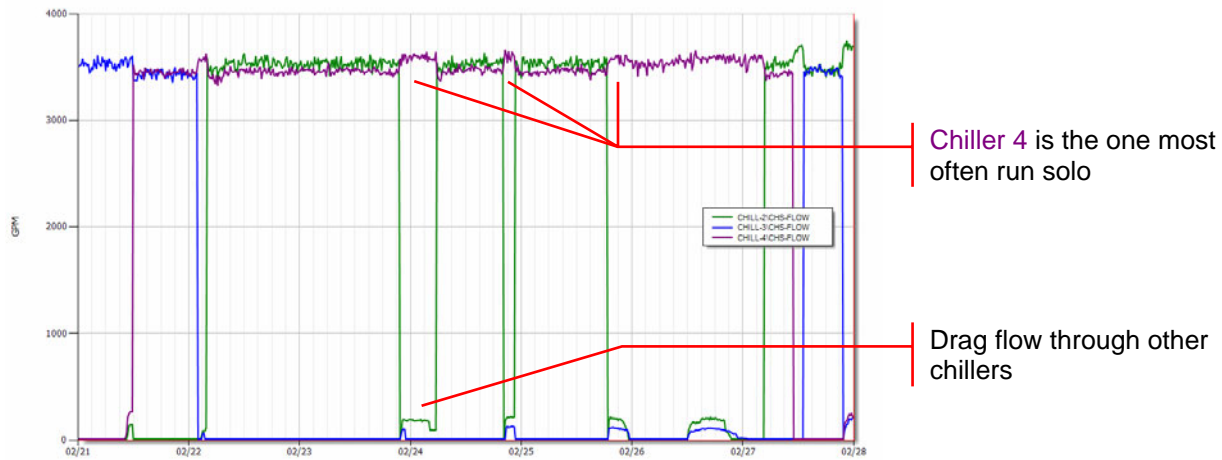


**Figure G: Secondary CHW Supply Temp Rises**

### Chiller Operations

February in Florida does require cooling, but the load is not high, especially nights and weekends. During those off hours, SJH ran only one chiller, as it was assumed it would save

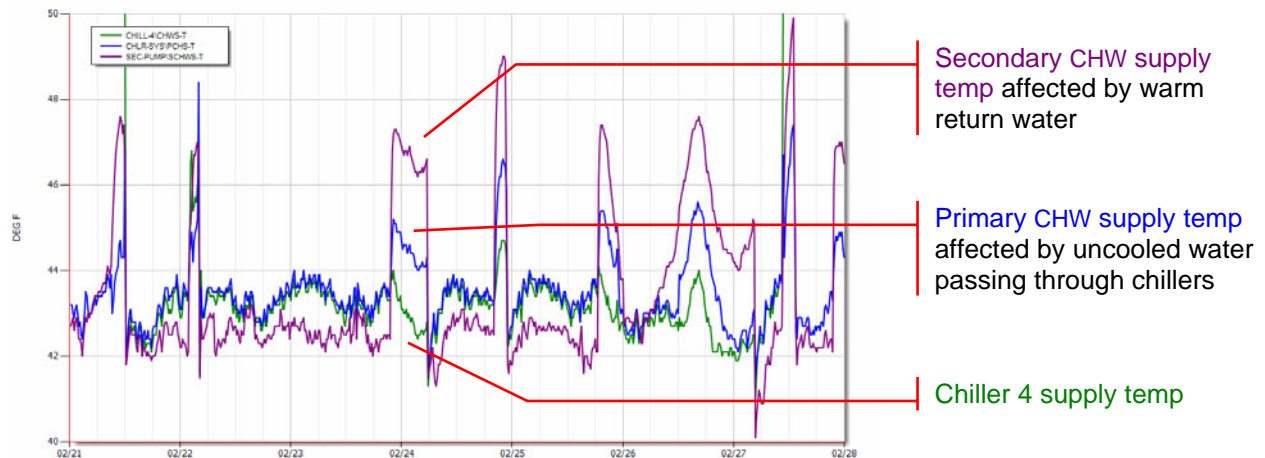
money. Figure H shows that every time the system cuts back to one chiller, all the other issues start to arise.



**Figure H: Chiller Flows**

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**

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 also the total plant cost. (Outside air temperature and bypass direction are also included for reference.) 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 \$18/hour more to run than when running as one of two. This wipes out 53% of the anticipated savings of turning a chiller off. Then when the secondary pumps start working harder to compensate, you lose another \$3/hour. While there is a small operational savings (about \$10/hour), factor in the instability introduced, the added wear on equipment—is it worth it?



**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...

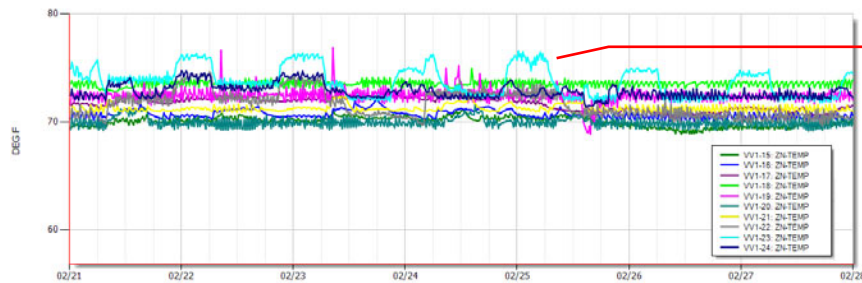
...none of this actually fixes anything.

## “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.

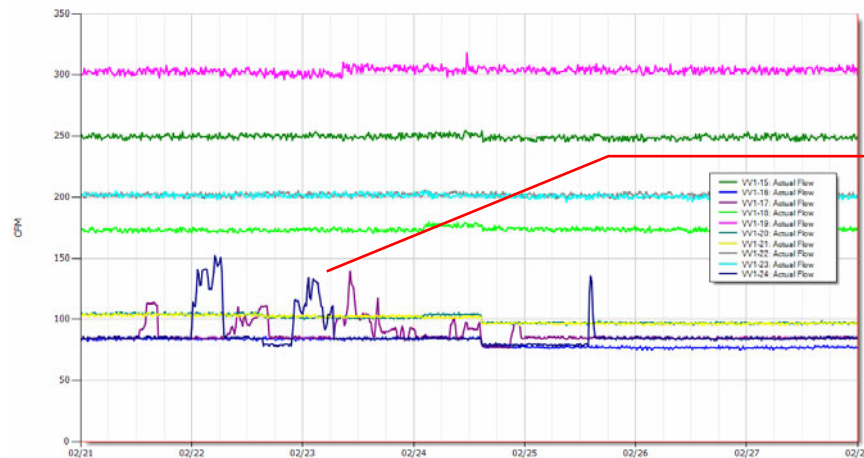
## Space Comfort and Air Flow

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 commissioning agents/engineers/technicians (including some at SJH) 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, and gathering more data (there are 288 points associated with the VAV boxes for this one air handler) without probable cause is usually labor intensive.



The inconsistency of **VAV 1-23 zone temp** is mostly due to changing setpoints as occupants adjust their thermostats up and down

**Figure K: Space Temperatures at 10 VAV Boxes**



Only two VAVs ever operated above their minimum CFM levels

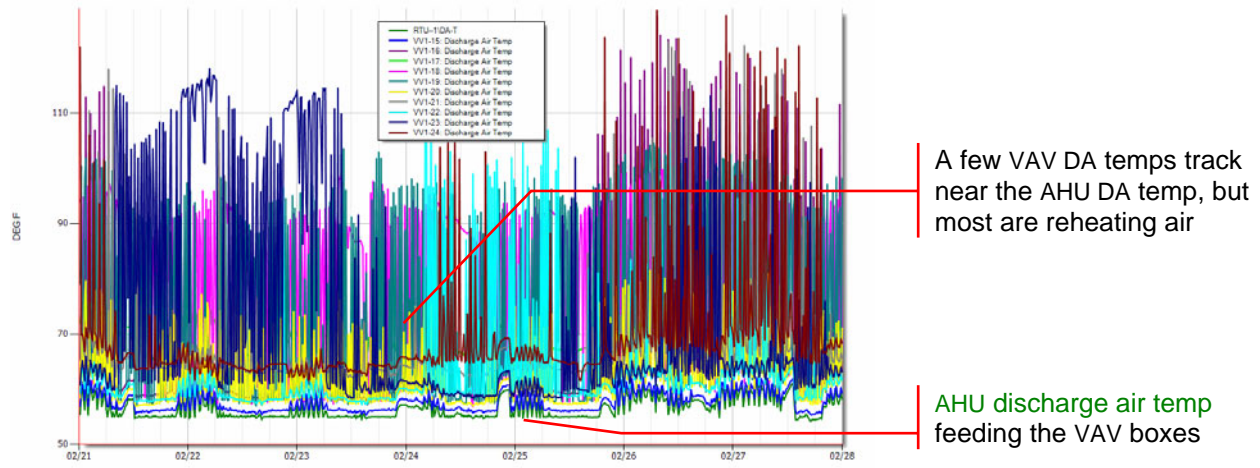
**Figure L: Air Flows (CFM) for 10 VAV Boxes**

“No problems” is an understandable conclusion without looking at the data in detail. Unfortunately, it’s also wrong. A closer analysis of the VAV information 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

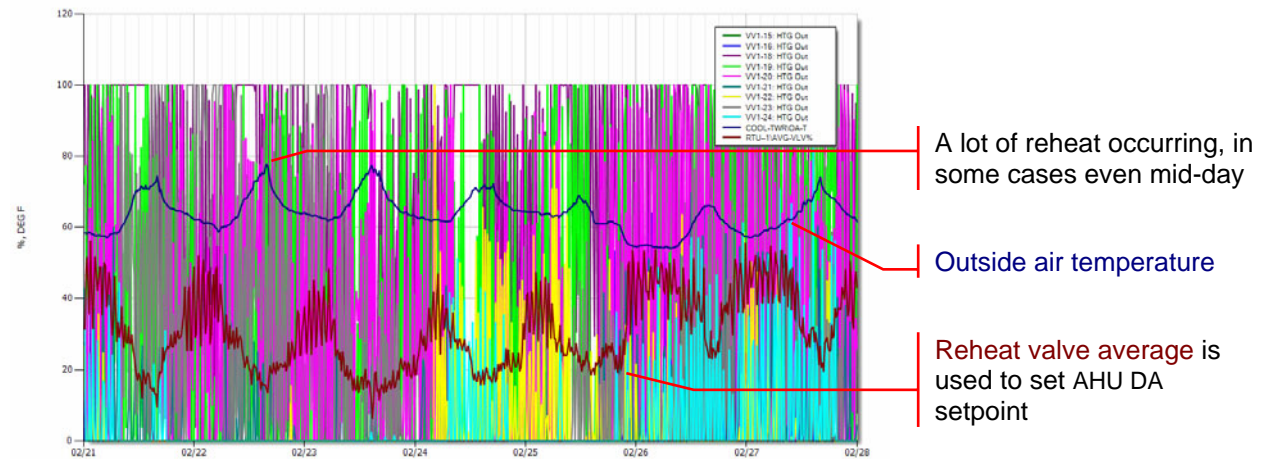
throughout the building. Of the ten VAV boxes shown in Figure N, only two are not supplying a significant amount of reheat.



**Figure M: Discharge Air Temperatures for 10 VAV Boxes**

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 air 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**

### 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**

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 (as can be seen on the weekend), so that the cooling valves don't need to be 100% open and driving the demand on the chilled water loop.

## Recommendations

The whole air handling system and VAV boxes need to be rebalanced based on what is actually needed to operate the space properly instead of what the design specifies. Balancing is an iterative process—make changes, re-examine the data, make more changes as necessary, etc.

The system should not calculate the discharge air setpoint based on the average reheat setting. While considered to be a best practice in the days of pneumatic controls, digital controls often respond too dramatically to fluctuations, as seen in the cycling behavior of the reheat valves and VAV discharge air temperatures rather than smoother operations.

The AHU discharge air setpoint should be raised until no reheating occurs. Given the time of year, this may mean a slightly higher baseline with a 3°F jump at night and on weekends (tied to outside air temperature rather than average reheat position). The operations must balance the elimination of reheat with the increased air flow demands if some spaces are getting too warm in order to minimize the fan horsepower used.

Once balanced, system-wide stability will result, and during these winter nights the facility can operate with just one chiller, saving on chiller operations (and boiler operations given the dramatic reduction of reheating required), without the side effects of secondary pumping and other system instabilities.

## The Answers Are In the Data

The exploration of SJH'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 whole case—identifying where problems existed, tracking them throughout the facility, measuring the cost impact, finding the true root cause, and coming up with recommendations—took only *three hours* of engineering/diagnostician time. Another four to five hours went into communicating with SJH over the course of a week. And, it was done without ever stepping foot in the state of Florida—completely diagnosed from 1,000 miles away using the data.

Commissioning projects have proven many times in the past that it is possible to lower costs while maintaining or improving comfort. The Continuous Commissioning® projects done by the Energy Systems Laboratory found a typical savings of 20%, with examples reaching 45%. However, labor costs to implement this level of gain are significant—enough so that continual retro-commissioning projects are far from commonplace.

As illustrated in this case study, 15 minute interval data provides all the information needed to resolve the most complex operational issues. There is still a place for a physical inspection as part of a commissioning project (the data won't tell you if air ducts are starting to corrode, for example, although it could tell you if they are actually leaking), but a data-oriented approach improves the accuracy and thoroughness, and is also self-documenting and self-verifying. Effective use of data is the key to maximizing commissioning results while minimizing the labor hours/cost. When SJH saw hard engineering facts, they were sometimes surprised, but always confident in the information and resulting decisions.

One last comment on interval data collection: it needs to be complete and constant. Collect data from every point in the BAS (meters and utility data too, if possible). There is no way to know ahead of time which points you'll need, so have them all (this could exceed 100,000 points for a large campus). 15 minute intervals provide a good balance between providing enough detail and not overburdening the system with finer granularity that adds little value. Keep the data indefinitely—it's common to want to look back in time. Just having data isn't sufficient—to reap

the benefits of this approach, the data must be as fast an easy to access as your e-mail. This is a very different approach to data and trend information than BAS trend logs.

### Maximizing the O&M / Commissioning ROI

There is a relationship that exists between the operating costs of the facility (energy and operations & maintenance) and the interactive use of operational data (for performing analysis, diagnostics, monitoring, verification, and other tasks related to commissioning). Simply put, the more labor hours you put into O&M and commissioning, the more you can improve system efficiency and drive down energy costs. However, it does reach a point of diminishing returns where energy savings do not compensate for the effort expended.

Figure P shows this inflection point as it exists for standard industry O&M practices (in blue). Energy costs drop as systems are tuned to peak operating efficiency. At the same time costs for ongoing O&M labor rise to make these improvements, with costs dramatically increasing the nearer the approach to ideal system efficiency. The combined energy and O&M costs create a curve showing the optimal point to balance the expense necessary to achieve the maximum budget reduction. The only way to significantly change where this balance point occurs is to employ a system that increases productivity by directing O&M efforts exactly where they are needed. To accomplish this goal, operational data must be continually collected and available for interactive use.

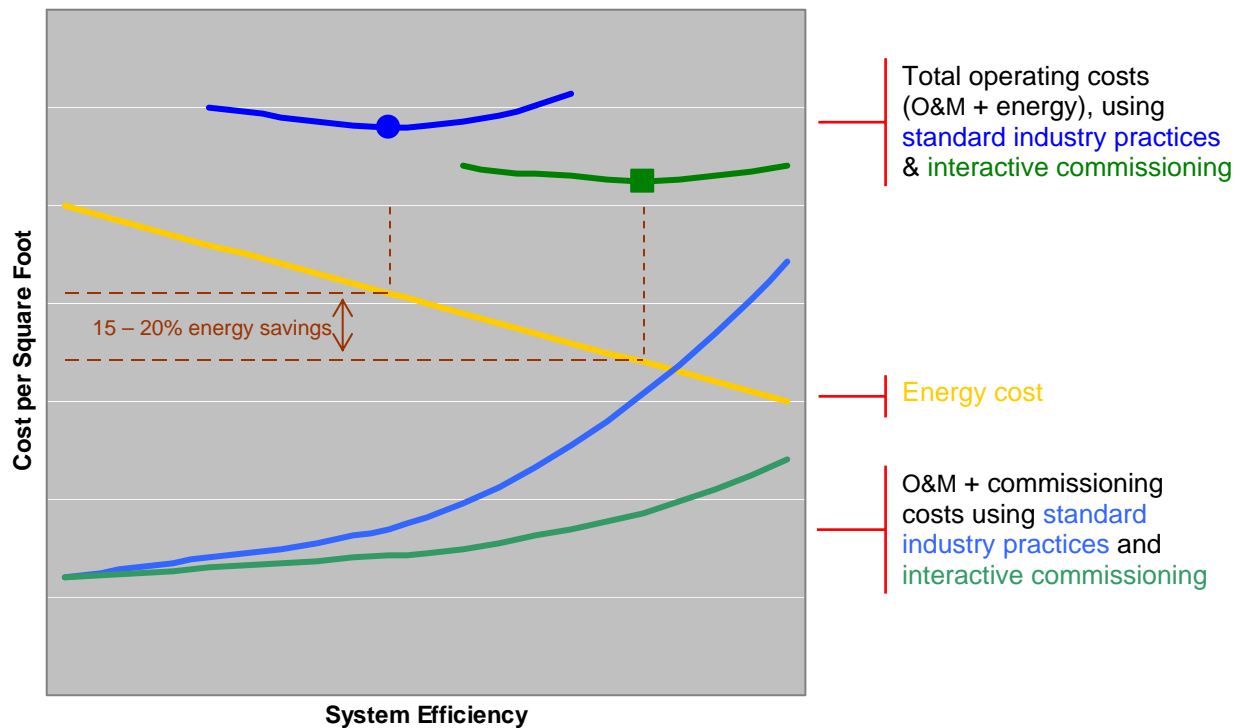


Figure P: Correlation between Operating Cost and Data Usage

Automated data collection eliminates the costly process of manually gathering data. Using that data allows staff to focus on actual issues—working smarter and using less labor. Interactive data access makes the already well-defined procedures for commissioning cost-effective, and enables the industry to take the next big step towards operational efficiency and energy savings.

Future work will provide further data to better define and quantify the impact and labor savings of the interactive commissioning approach. Additional results will be published as they become available.

## Citations

Culp, C.H., Turner, W.D., Claridge, D.E., Haberl, J.S., “Continuous Commissioning<sup>SM</sup> in Energy Conservation Programs”  
<http://tmpwebesl.tamu.edu/programs/morecc.htm>

Interval Data Systems, Inc., 2004, “Hospital Case Study”  
<http://www.intdatsys.com/publications.htm>

Interval Data Systems, Inc., 2004, “Defining the Next Generation Enterprise Energy Management System,” 2<sup>nd</sup> Edition  
<http://www.intdatsys.com/publications.htm>

Liu, M., Claridge, D.E., Turner, W.D., 2002, “Continuous Commissioning<sup>SM</sup> Guidebook: Maximizing Building Energy Efficiency and Comfort,” Energy Systems Laboratory, prepared for FEMP, U.S. DOE  
[http://www.eere.energy.gov/femp/operations\\_maintenance/commissioning\\_guidebook.cfm](http://www.eere.energy.gov/femp/operations_maintenance/commissioning_guidebook.cfm)

Matrikon Inc., 2005, “Matrikon Solutions for the Power Industry”  
<http://www.matrikon.com/download/presentations/thePowerBlueprint.pdf>

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”  
<http://eetd.lbl.gov/emills/PUBS/Cx-Costs-Benefits.html>

**Interval Data Systems, Inc.**  
104 Summer St. | Watertown, MA 02472 | [www.intdatsys.com](http://www.intdatsys.com)  
tel: 617-744-1091 | fax: 617-744-1092 | e-mail: [info@intdatsys.com](mailto:info@intdatsys.com)  
Copyright © 2005 by Interval Data Systems, Inc.

