

CHILLED WATER SYSTEM OPTIMIZER

**A White Paper by
Steve Tom, P.E., Phd.**



**Carrier Corporation
Farmington, Connecticut**

July, 2017

INTRODUCTION

When it comes to measuring HVAC energy use in buildings, it's no surprise to learn there is a wide diversity in measurements. After all, building design, climate, building usage and many other factors vary widely from one site to another. The Center for Climate and Energy Solutions estimates that 32% of the energy used by a commercial building is from providing heating, cooling, and ventilation. By comparison, an Australian government study estimated that 70% of a building's energy use was consumed by the HVAC system, 25 – 35% of which was attributed to the chillers. The reason for the disparity in numbers is not relevant to this paper. What matters for our purposes here is that widely differing studies agree, worldwide, commercial HVAC consumes a tremendous amount of energy. If cooling is provided by a chilled water system, the chiller consumes a significant portion of the HVAC energy, and quite possibly consumes more energy than any other single component in the system. When building owners try to reduce their energy consumption, they often focus on optimizing their chillers. This can be a difficult task and there is also the danger that optimizing the chillers alone might not provide the best overall energy reduction. This paper will describe a new algorithm designed to optimize the entire chilled water system.

The Australian study provided an important warning: "When engineers and controls specialists focus on improving chiller efficiency, it is often at the detriment of the energy consumption of associated equipment such as cooling towers, air distribution fans, chilled water, and condenser water pumps. Sometimes the net result is an increase in total energy consumption. It is important that a more holistic systems type approach is used when looking to improve chiller efficiency." To understand why this is true, we need to take a very basic look at the dynamics of a simple chiller system. Consider how the energy used by a chiller varies with the chilled water supply temperature. (Fig. 1)

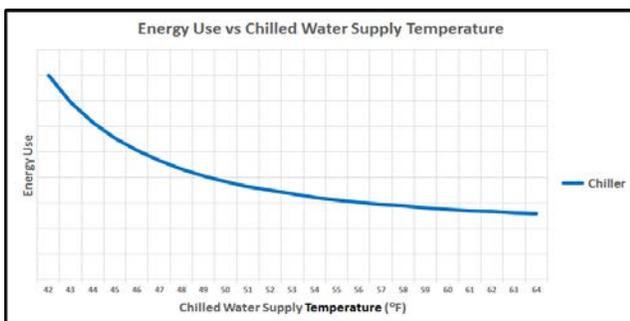


Figure 2: Simple Chiller Performance

The shape of a chiller curve largely depends on the type of chiller, operating speed, and many other parameters. In general, the energy used by the chiller decreases as the chilled water supply temperature increases. This result leads to the conclusion that to optimize chiller performance, you should raise the chilled water supply temperature as high as possible. Ignoring for the moment the fact that this warmer chilled water might not provide sufficient cooling and could cause problems in the chiller; the warmer supply water temperature could actually increase the total energy consumed by the chilled water system. This occurs because, as the chilled water temperature rises, variable flow chilled water supply pumps have to pump a higher volume of chilled water through the system to satisfy the same cooling load. As a general rule, the volume of water moved by a pump varies directly with the pump speed, while the energy used by the pump varies with the cube of the pump speed. This leads to a pump curve as shown in Figure 2.

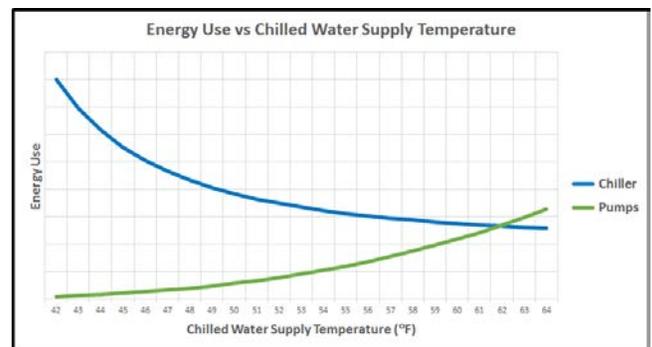


Figure 3: Simple Chiller and Pump Performance

Figure 2 shows that as the chilled water supply temperature increases, the chiller uses less energy, but the chilled water supply pumps use more energy. The total energy used by the chilled water supply system can be found by adding the energy used by the chiller and the pumps. This is shown in Figure 3:

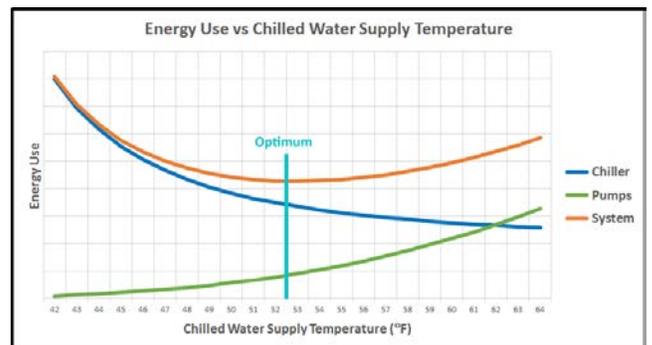


Figure 1: Total Chilled Water Supply System Energy

Figure 3 shows that combining the two curves produces a “bowl” shaped system curve. On the left side of the curve, an increase in chilled water supply temperature reduces the energy used by the chiller by a greater amount than the increase in the pump energy, so the system energy use decreases. On the right side of the curve, the increase in pumping energy outweighs the savings in chiller energy, so further increasing the chilled water supply temperature increases the overall energy use. The optimum operating point is at the “bottom of the bowl,” which is 52.5 °F in this hypothetical example.

So far, we’ve only looked at the energy used by the chiller and the primary chilled water supply pumps, sometimes called the “plant energy” because the primary pumps are often located in the chiller plant. Chiller plant optimization schemes typically focus on minimizing this energy. In a small chilled water system, with no secondary pumps and no variable speed fans, this may in fact represent the total chilled water system energy use. Larger chilled water systems are more complex. If there are multiple buildings in the system, there may be secondary loops and secondary loop pumps in the individual buildings. Sometimes there are tertiary loops and pumps as well. If the chilled water supply temperature is raised, cooling coil valves will open because it takes a greater flow of the warmer chilled water through the coils to meet the cooling load. This means the secondary and tertiary pumps will work harder to provide the increased flow rate. Fan coil units may need to cycle the fans on more often or run them longer to meet the cooling load and Variable Air Volume systems may similarly need to increase the fan speed and airflow. This adds a “building energy component” to the chilled water system energy analysis. The chiller and primary pumps no longer represent all of the energy being consumed by the chilled water cooling system. This additional load is shown in Figure 4:

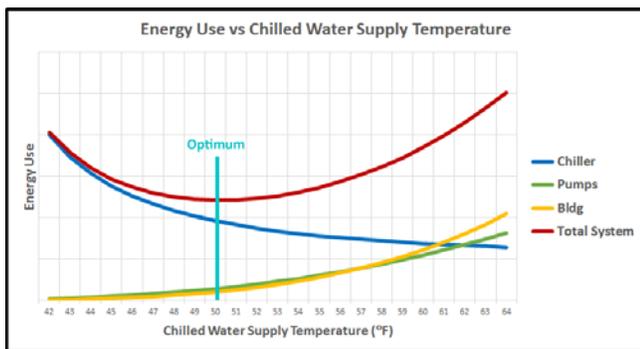


Figure 4: Chilled Water Energy Use including building load

In Figure 4, the energy used by the chilled water system components in the building has been added to the energy used by the chiller and primary pumps to create the Total System energy curve. This is also a bowl-shaped curve, but since both the building energy and pump energy increase as the chilled water temperature increases, the right hand side of this curve rises faster than the curve in Figure 3, which does not include the building energy. The net effect is to shift the optimum chilled water supply temperature to the left, which is 50° F in this hypothetical example.

This is a very simplistic description of chilled water system optimization. The efficiency of the chillers, use of variable speed chillers and pumps, staging of chillers and pumps in larger plants, line losses through the piping system, and many other factors will influence the shape of these curves and have a significant effect on the overall energy use. Similarly, there are many constraints that limit the range of chilled water supply temperatures that may be used. The optimum temperature from an energy standpoint may not be cool enough to satisfy the cooling requirements of the buildings. Dehumidification requirements may dictate a cooler than optimum supply water temperature, and the need to prevent operational problems in the chiller will definitely limit the range of acceptable supply water temperatures. When everything is considered, the net result will be a bowl-shaped curve like that shown in Figure 4. The bottom of this bowl represents the supply water temperature that will result in the minimum energy usage under the current operating conditions. As long as that temperature does not cause operational problems in the chiller, and is cool enough to not cause temperature or humidity problems in the controlled spaces, you want to control the chiller plant to supply water at that temperature. The goal is to find that optimum temperature.

One way to determine the optimum supply water temperature is to model the chillers, pumps, and building systems. These models can then be used to calculate the curves shown in Figure 4. Once you have the equations for these curves, you can calculate the optimum temperature. While simple in theory, developing accurate models can be extremely difficult. Modeling pumps is fairly straightforward. Manufacturers typically provide operating curves for new pumps so the measurements needed to calculate the system curve and to determine the actual pump operating point under different conditions are not overly difficult. Modeling the chiller, on the other hand, can be much more difficult. Factory curves may

be available for some new chillers, but these curves don't account for aging, fouling, and other changes that occur over time in real-world chillers. It's possible to adjust for these factors, but that requires costly field measurements and expert analysis. Multiple chiller plants are even more complicated to model. These measurements and analysis need to be repeated as the system ages or as changes are made to the system.

Modeling the building systems is even more difficult since there are no "factory" baseline models to use as a starting point. The possible combinations are endless and in large campuses, the buildings are almost continuously being modified. Most commercial chiller optimization programs do not attempt to model the buildings. Instead, they simply model the chiller plant, and their models may only be applicable to specific plant configurations. In essence, they are modeling the curves shown in Figure 3.

An alternative way to determine operating temperature, without modeling the system, is to apply an "intelligent reset" algorithm. Since no modeling is involved, the algorithm does not know the equation for the "Total System" curve shown in Figure 4, but it knows it is bowl-shaped. The algorithm will make a small adjustment to the supply water temperature, say, by raising the temperature slightly, and measure the resulting changes in energy consumed by the chiller, pumps, and building systems. If the previous supply water temperature was to the left of the optimum point, the total energy used by these systems will decrease. This tells the algorithm that raising the supply water temperature made an improvement, so it will raise the temperature a little bit more. If, on the other hand, the previous supply water temperature was to the right of the optimum point, the total energy used by the system will increase when the supply water temperature is increased. This tells the algorithm that raising the supply water temperature was the wrong thing to do, so its next adjustment will be to decrease the temperature. Eventually, it will find the optimum temperature, where either increasing or decreasing the supply water temperature increases the total energy used. In essence, it's like dropping a marble into a bowl. The marble may roll back and forth a bit, but eventually, it will settle to the bottom of the bowl.

The optimum chilled water supply temperature is not a fixed point. It will vary throughout the day, from day to day, and from month to month, as the load on the building changes, as the weather changes, and as other factors affecting the performance of the chiller, pumps,

and building change. These changes will alter the position and curvature of the Total System curve shown in Figure 4, but they do not change the fact that it is essentially bowl-shaped. The intelligent reset algorithm will, therefore, react to these changes and find the new optimum operating temperature. Carrier's new chilled water system optimizer uses such an intelligent reset algorithm as stated above.

Of course, controlling a real-world chiller plant is much more complicated than the simplified explanation provided in this paper. Remember the analogy of a marble in a bowl? Imagine that bowl is on the deck of a small ship, which is being tossed about by a violent storm. The marble will no longer sit calmly at the bottom of the bowl. The performance of a real-world chilled water system is not quite as extreme as a ship in a storm, but there are many factors that can affect the energy used by the chiller, pumps, and buildings. Building loads may change abruptly as scheduled occupancy and operations vary throughout the day. Single speed pumps and fans cycle on and off. Solar loads may change as clouds drift in front of the sun. The intelligent reset algorithm needs to respond to the effect these changes have on the sustained energy use of the system, but not to the abrupt discontinuities. Thus the patent-pending algorithm used by Carrier includes time delays, filters, and "predict and verify" adjustments to minimize or eliminate the effects of these short-term discontinuities, allowing the algorithm to focus on the long-term changes that affect the optimum operating temperature.

The discussion so far has centered on the chilled water supply side, but if water cooled chillers are used, there is a similar trade-off taking place on the condenser water side. In general, if you lower the temperature of the condenser water as it's supplied to the chiller, the amount of energy used by the chiller will be reduced. In essence, the cooling tower is doing some of the work for the chiller. As you might expect, this means the cooling tower has to work harder. The cooling tower can't provide water that's any cooler than the outside air wet bulb temperature, and the closer the condenser water supply temperature approaches to this wet bulb temperature, the harder the fans and circulating pumps in the cooling tower need to work. This may be achieved by speeding up the fans and pumps, bringing on more tower stages, or both. The net result is a bowl-shaped curve for the cooling tower and chiller system as shown in Figure 5:

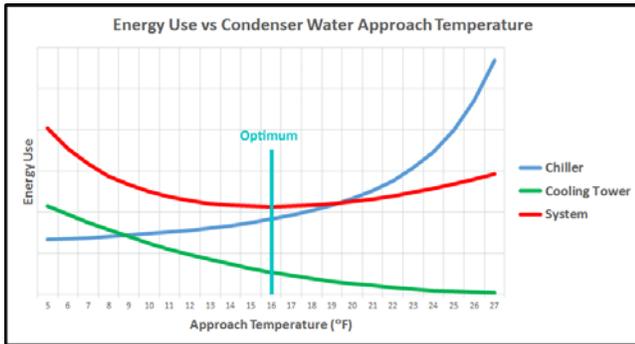


Figure 5: Cooling Tower and Chiller Energy Use

Since this is another bowl-shaped curve, an intelligent reset algorithm comparable to the one used to find the optimum chilled water supply temperature can be used to find the optimum cooling tower approach temperature. The approach temperature has no effect upon the chilled water pumps and the building system, as long as the chiller is able to supply the desired chilled water supply temperature. The optimization loop for the cooling tower only has to consider the power consumed by the chiller and the cooling tower fans and pumps. This loop reacts more quickly than the chilled water supply loop, as chilled water needs to circulate throughout the entire building (or campus) water system before you can accurately determine the effects of changes to the chilled water supply temperature.

To optimize both the chilled water system and the condenser water system, we run the two optimization programs sequentially. First, we make a change to the chilled water temperature, wait for that change to affect the energy used by the chilled water loop, and then see whether this reduced or increased the energy used by the chilled water loop. If it reduced the energy use, we continue to adjust in the direction of the new chilled water temperature. If it increased the energy use, we adjust back toward the previous temperature. Then we make a change to the condenser water supply temperature, wait for that change to affect the energy used by the chiller, the cooling tower fans, and pumps, and then see whether this reduced or increased the energy used by the condenser water loop. If it reduced the energy use, we continue to adjust in the direction of the new condenser water temperature. We then turn our attention back to the chilled water loop, make a change to the chilled water temperature, and repeat that loop. After a few cycles of adjusting the chilled water temperature, condenser water supply temperature, chilled water temperature, etc., we will be zeroing in on the optimum temperature for both loops. The Carrier chilled water system optimizer includes

intelligent algorithms for the cooling towers as well as for the chilled water supply system, and finds the optimum temperatures for both loops.

The overall optimization system includes some key safeties. First, the chilled water supply temperature and condenser water supply temperature are not allowed to vary into ranges that could cause surging in the chiller. Also, provisions are made to allow buildings to request cooler water from the chiller in the event that the building HVAC systems are not able to maintain acceptable temperature or humidity conditions. (A discussion of the “cost” of comfort is beyond the scope of this paper, but if the people in the building aren’t comfortable, the cost of lost productivity will greatly outweigh any energy savings that may have resulted from raising the supply water temperature.)

It should also be noted that a complete implementation of this chilled water system optimization requires measuring the energy used by every component in the chilled water system – chillers, primary pumps, secondary pumps, HVAC system fans, cooling tower fans, and cooling tower pumps. Very few systems have individual meters on all those components. Fortunately, energy metering can be done through individual meters, variable frequency drives, or even chiller panels, if each measures and provides its kW power value. The good news is that the algorithm will optimize whatever is metered. Thus, if only the chiller and the primary system pumps are metered, the algorithm will find the operating point that minimizes the energy used by these two components. Essentially, it will optimize the curves shown in Figure 3. If the cooling tower is also metered, it will find the optimum cooling tower approach temperature as well, as shown in Figure 5. If a few of the larger HVAC systems are metering electrical energy used by the fans and/or the secondary loop pumps, the optimum chilled water temperature determination will be refined to take the energy used by these components into effect and approach the performance shown in Figure 4. Not everything needs to be metered. More meters provide a more refined optimization, but the system will optimize based upon what is metered.

In developing the intelligent reset algorithm, we tested the algorithm against some simple modeled systems. The reason for this is to calculate the optimum chilled water temperature for that modeled system and see whether or not our algorithm found this optimum temperature. We first tested it against a fixed optimum temperature or “sweet spot.”



Figure 6: Optimizing to fixed sweet spot

While Figure 6 is not particularly exciting, it does show that the intelligent reset algorithm did correctly reset the chilled water supply temperature until it found the optimum temperature. During the same test, we calculated the energy used by the modeled components and compared that use to the energy that would have been used with a fixed chilled water supply temperature, i.e. with no reset. Not surprisingly, it showed the intelligent reset algorithm used less energy than a system with a fixed setpoint:



Figure 7: Energy savings compared to a fixed sweet spot

In a real system this optimum temperature would not remain fixed, but would instead vary throughout the day as the outdoor conditions changed and as the building use varied. The next test of the reset algorithm required it to control a simulated system where the optimum temperature varied sinusoidally over time. We also introduced some random “noise” into the energy feedback signal from the simulated building system to simulate the short term effects of pumps switching on and off, as if the sun were going behind a cloud, or as if the building-load changes, as discussed previously:

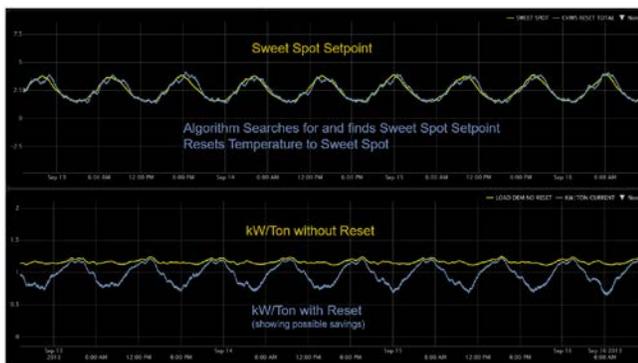


Figure 8: Optimizing a system where the sweet spot varies over time.

Figure 8 shows the intelligent reset algorithm did a good job of tracking a moving “sweet spot,” even when that optimum temperature cycled through two highs and two lows in a 24 hour period and even when there was “noise” in the energy feedback. Realistically, most chilled water systems only experience one high and one low per day, so this was a fairly aggressive test. Figure 8 also shows that the energy saved by an intelligent reset algorithm compared to the energy used by a system without reset varied sinusoidally with the optimum temperature. It should be noted that when comparing an optimized system to a system with a fixed chilled water temperature, the magnitude of the savings depends upon the value chosen for the fixed chilled water temperature. In the test shown in Figure 8, the fixed temperature chosen for comparison was within the range of the optimized temperature. Just as a stopped clock is right twice a day, the fixed chilled water temperature periodically was equal to the optimum chilled water temperature and the optimized system briefly showed no energy savings compared to the fixed system. In a real-world chiller system, the operators typically do not know the range of optimum chilled water temperature. If they choose a chilled water supply temperature that is not within the range of optimum temperatures, the energy savings will look like Figure 9:

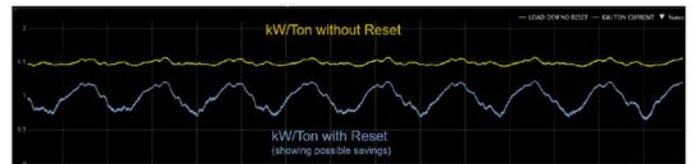


Figure 9: Energy savings compared to a non-optimum fixed chilled water temperature

Figure 9 shows what the results would have been if the fixed chilled water temperature had not been within the range of optimized chilled water temperatures. In this case the fixed system never would have operated as efficiently as the optimized system. The optimized system would always show a savings compared to the fixed system, and the amount of savings would have varied as the optimized chilled water temperature varied closer to or farther away from the fixed temperature. When comparing the performance of an optimized system to a non-optimized baseline system, it’s important to realize that the savings achieved depends upon how well the baseline system was configured. A poorly configured baseline system will make the optimized system look very good by comparison. In all our testing, including our field

testing, we compared the optimized system to a baseline system which had fixed setpoints within the range of the optimized setpoints. This was a conservative approach that minimized the energy savings reported, but it was the only way to avoid the bias that would have resulted if we had used a poorly configured system as a baseline. It's also worth pointing out that operators typically do not know what the "range" of optimum chilled water temperatures is unless they have some type of chiller optimizer algorithm. Thus, Figure 9 could be more indicative of the savings to be gained in a real-world chilled water system than Figure 8.

Field Testing and Verification

Field testing was performed at two academic campuses in the northeast: Site A, a four-year preparatory school (grades 9 – 12), and Williams College, a liberal arts college in Williamstown, Massachusetts. Site A typically enrolls 650 students, 90% of whom reside on campus full time. Williams College has approximately 2,250 students, almost all of whom reside on campus. (Having students reside on campus means there may be a cooling load 24 hours a day, on weekends and on weekdays.) Installation and configuration of the optimization program were both performed by the local controls office using instructions and online documentation provided by the factory. Installation was performed over the summer and did not require shutdown or disruption of the existing chiller plant controls. The optimization package acted as a "supervisory control system" to the existing controls. It adjusted the chilled water supply temperature and condenser water supply temperature setpoints, but did not interfere with the existing on/off, staging, or safety controls.

The chilled water system that was optimized at Site A included two parallel 225-ton variable flow chillers, one pair of variable speed primary chilled water pumps, three pairs of variable speed secondary chilled water pumps (six total), and six variable speed Air Handling Units in three buildings. Due to weather conditions during the test, only one chiller was running at any one time. The condenser water system at Site A consisted of two open cooling towers with variable speed fans and two variable flow condenser water pumps.

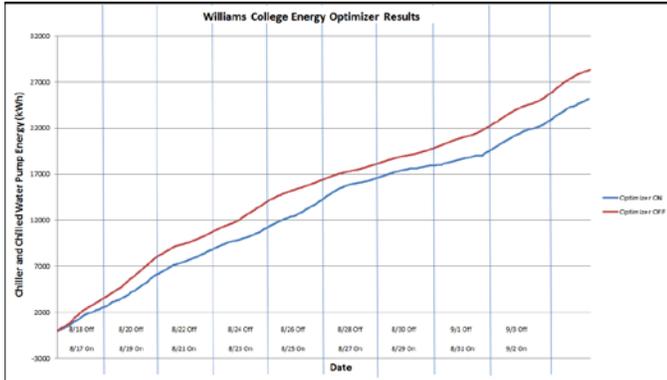
At Williams College, only the chillers and primary chilled water pumps were metered, so the supervisory controls essentially optimized the chiller plant

operation without regard for the operation of the building HVAC systems. The portion of the Williams chilled water system that was optimized included three 500 ton variable flow parallel chillers and three parallel variable speed primary loop pumps. The condenser water controls were not designed to allow an external system to change the condenser water setpoint, so only the chilled water supply temperature was optimized. Due to the weather conditions, the Williams system ran on one chiller most of the time and occasionally brought on a second chiller as well. The plant did not run all three chillers during this test. The Williams system also included a local routine to raise the chilled water supply temperature during certain load conditions. This local routine did not activate very often, and it did not reset the chilled water temperature as much as the optimization algorithm did. It undoubtedly did save some energy, however, so the savings achieved by the optimizer routine at Williams College were probably less than what they would have been had the optimizer been compared to a fixed setpoint.

At both campuses, the summer installation left only a limited number of days of "normal" operations (classes in session) before the weather cooled off to the point that regular chiller operation was not required. To provide a meaningful comparison during this test period, the chiller optimization was enabled one day and disabled the next. When the optimization was disabled, the chillers and cooling towers operated with fixed setpoints. (In the case of Williams College, the chilled water supply temperature setpoint was occasionally reset as described in the previous paragraph.) These fixed setpoints (or the setpoint range at Williams) were within the range of optimized setpoints as discussed previously. While the weather and building operating schedules were not identical from day to day, they were similar enough that over the test period minor daily variances were assumed to average out so a meaningful comparison could be made.

Figure 10 shows the cumulative energy consumption at Williams College with the optimizing algorithms controlling the chillers every other day, from August 17th through September 2nd. The performance curve for this test is superimposed over the performance curve for operations with a fixed setpoint, which occurred every other day from August 18th through September 3rd. The net result of this schedule was that the optimizer was on for 9 days and off for 9 days. Two of the days when the optimizer was off were weekends, while only one of the days when the

optimizer was on was a weekend, which might have given a slight advantage to the “optimizer off” energy use. It’s clear that overall the chiller and pumps consumed less energy when the optimizer was active than when the system was operating with fixed setpoints.



The total energy consumed during nine days of optimized operation was 25,148 kWh while the energy consumed during nine days of fixed setpoint operation was 28,329 kWh. The optimizer reduced energy consumption by 11.23%.

A similar graph for the cumulative energy consumption at Site A is given in Figure 11:

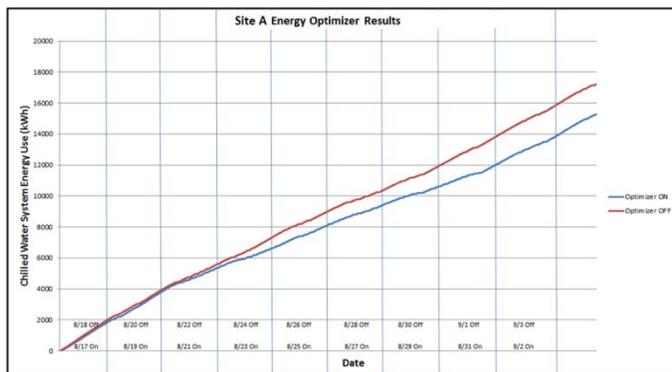


Figure 11: Cumulative energy consumption at Site A

Figure 11 shows the performance at Site A over the same time period. Again, there are 9 days when the optimizer was off and 9 days when the optimizer was on, with the “off” period having the advantage of two weekend days versus one weekend day for the “on” period. The primary difference between the two systems, besides the fact that Williams College is a larger system overall, is that the optimization program at Site A included energy feedback from some of the

larger building systems. The optimizer did not appear to have provided much advantage during the first two days of operation, but after that, it consistently outperformed the non-optimized system. Total energy use with the optimizer on was 15,272 kWh while the total use with fixed setpoints was 17,215 kWh. Thus, the optimizer reduced energy consumption by 11.28%. At first glance, it might appear that including the energy used by building systems in the optimizing algorithm made very little difference, but the percentages are being calculated against two different baselines. The total energy used by the chiller plant alone will always be less than the energy used by the total chilled water system, since the total system includes the chiller plant as well as the building chilled water HVAC systems. Saving 11% of the total chilled water system therefore saves more energy (and more energy dollars) than saving 11% of the chiller plant energy alone. One cannot always expect the same results each time as systems, loads, and equipment do vary from site to site, and from season to season. However, these tests do show a possible range of savings.

In Summary:

- Carrier has developed a chilled water optimization program that resets the chilled water supply temperature and the condenser water supply temperature to provide optimum performance.
- This is an adaptive algorithm that uses actual system measurements.
 - It does not require detailed modeling and set up.
 - It works with a wide variety of chillers, pumps, and configurations.
 - It automatically adapts as conditions change in the system.
- This algorithm has the ability to minimize energy used by the entire chilled water system, not just the chiller plant.
 - It will optimize what is measured. If only the chiller plant is metered, it will optimize the chiller plant. If major chilled water components in the buildings are metered (secondary loop pumps, HVAC fans, etc.), it will provide chilled water at a temperature which minimizes energy use by the entire system.
- This algorithm is essentially an overlay to the existing chiller plant controls. It adjusts the chilled water and condenser water setpoints, but it does not interfere with the existing stop/start, staging, and safety controls.
- The algorithm includes provisions to override the setpoints to ensure minimum temperature and humidity conditions are met and to prevent surging.
- Initial field tests of this algorithm showed that it reduced chiller plant energy consumption by 11.23% at one installation and total chilled water system energy consumption by 11.28% at another installation.
- As seen from the test data, one cannot always expect the same results each time as systems, loads, and equipment do vary from site to site, and from season to season. However, the tests do indicate a possible range of savings.