How High Can You Go?

New approaches for cooling system efficiency

Ben Erpelding, PE, Director of Engineering, Optimum Energy LLC

The volatility of energy costs and emphasis on reducing carbon footprints have increased the urgency of improving district cooling plant energy performance. At the same time, facility managers and owners face increased pressure to reduce operating costs, take on new loads and continue to provide 100 percent reliability and zero downtime for their chilled-water services. This poses a dilemma for plant operators, as energy efficiency usually ranks lower as a priority than reliability, availability and even construction time schedules.

It is possible, however, for chilled-water plants to achieve new levels of efficiency – what Optimum Energy calls "ultra-high performance" – without compromising reliability. It's already being done in plants from California to the United Arab Emirates using approaches that represent a new paradigm in system optimization.

What Is Ultra-High Performance?

In district cooling plants, ultra-high performance (a term first introduced in a 2001 ASHRAE article by Tom Hartman) is defined as the delivery of chilled water at an average annual efficiency of 0.45 to 0.70 kW/ton (annual kilowatt-hours divided by annual ton-hours) [7.8 to 5.0 coefficient of performance (COP)], including all chiller, chilled-water distribution pumping, condenser water pumping and cooling tower fan energy. [This annual average efficiency is based on an all-electric plant and excludes hybrid (gas- or steam-driven) plants and deep lake/ocean water cooling plants; these can achieve electrical efficiencies of less than 0.30 kW/ton.]

While 0.45 kW/ton is achievable, plant design, equipment, load profile and climate (wet-bulb) push these levels higher. The transitional climates, such as New York, allow free cooling (plate-and-frame heat exchanger) strategies throughout the winter; temperate climates like Los Angeles have year-round cooling loads at relatively low wet bulbs; the arid subtropical climate of the Middle East (e.g., Dubai) is known for extreme wet-bulb designs (89 degrees F); and tropical climates, such as Singapore, have a nearly constant outdoor wet-bulb temperature year-round. Though the efficiency variance in ultra-high performance plants is minimal (0.45 to 0.7 kW/ton), engineers need to understand how to mitigate each climate's effect on potential cooling plant efficiency.

A Winning Combination

Achieving ultra-high performance in district cooling plants requires a combination of best-of-class plant design and sophisticated optimization algorithms, as well as a fully automated plant and extensive control system. Only this combination makes it possible to sustain consistent operation of a chiller plant at levels less than 0.6 kW/ton for extended periods across different facility types and global locations.

In the past, district cooling plants held strong to the control theory that because their plants have multiple chillers, equipment can be sequenced accordingly to keep them as fully loaded as possible. Control strategies were written into the plant design and specification by the engineer of record, and custom programming was required by the control contractor. Plant efficiency, therefore, was limited by the equipment full-load efficiency and the technical expertise of the on-site programmer. This old paradigm has resulted in average annual plant operating efficiencies of 0.8 to 1.1 kW/ton.

Today higher levels of efficiency are attainable by employing (1) variable-frequency drives throughout the plant – all chillers, chilled-water and condenser-

© 2009 International District Energy Association. ALL RIGHTS RESERVED.

water pumps and cooling tower fans; (2) fully automated network controls; (3) new relational control algorithms to optimize all system components for energy efficiency; and (4) optimization software. Striving for ultra-high performance using this combination of technology represents a new paradigm in district cooling plant optimization.

Variable-Frequency Drives

Variable-frequency drive retrofits for large-tonnage chillers are becoming more and more commonplace. Retrofits have been done on both 4,160-V and 12-kV chillers by major chiller manufacturers and local contractors in almost all major world markets. Variable-speed medium-voltage new chillers have been available for a few years now and are offered by most manufacturers. Variable-speed dual-compressor chillers are also available from at least two major manufacturers.

There is nothing new in primary-only variable-speed chilled-water pumping. Variable primary-only plants are being implemented around the world with plant tonnages well exceeding 50,000 tons. The key to successful variable-speed chilled-water pumping is getting away from solely relying on the 'blind' differential-pressure control and blaming the problem on the 'building guys' and their low delta T. New relational control methodologies with extensive networked controls need to employed. An energy management and control system network down to the building connections – bridges, booster pumps, plate-and-frame heat exchangers and even air-handler chilled-water valve positions – is mandatory. Open conversations with building owners and managers about improving delta T are also strongly encouraged (or put into the chilled-water contract language).

Variable-frequency drives on condenser water pumps significantly contribute to lowering a plant's annual kilowatts per ton. Variable-frequency drives mitigate condenser pump energy at low chiller capacities and also allow new control strategies for optimal chiller, tower and condenser pump sequencing (add/shed). The key to success in employing variable-frequency drives on condenser water pumps is not robbing Peter (the chiller energy) to pay Paul (condenser pump energy).

Although there is nothing new about the concept of employing variable-frequency drives on tower fans, unfortunately there is a new design and control trend not to include them, especially in tropical climates. This design method assumes the objective is to 'make the cooling tower water as cold as possible,' which would require the tower fans to run at full speeds anyway. Variable-frequency drives, again, mitigate tower energy at low chiller capacities – without compromising the amount of tower cells kept on line – and allow new control strategies for optimal chiller, pump and tower cell sequencing (add/shed).

Fully Automated Networked Controls

Many district chilled-water plants depend on operators to make efficiency control decisions. Optimization is deployed manually based on operator experience, or trial and error, and can be complaint-driven (i.e., 'tweak something until a customer complains'). This type of optimization is unsustainable and can never achieve the full potential for energy savings.

For truly effective optimization, the district plant should have as much cooling demand visibility as possible. Cooling demand visibility is defined as energy management system points that extend to individual building booster chilled-water pump variable-frequency drive speeds (0 percent to 100 percent), building chilled-water flow (gpm), building chilled-water supply and return temperatures, commanded heat exchanger valve positions (0 percent to 100 percent) and commanded building air-handler chilled-water valve positions (0 percent to 100 percent). Networking all this information back to the plant control system can hugely affect the algorithms' ability to automatically control and significantly improve system performance and delta T.

Figure 1 shows delta T results before and after improved cooling demand visibility at the University of North Carolina (UNC) at Chapel Hill. Jim McAdam, PE, energy engineer with UNC Energy Management, further explains, "We have begun operating the bridge pumps in series with the mains pretty much everywhere now and are very pleased with the energy savings, reliability improvements and delta T improvements we are seeing. We have also piloted resetting each building's end-of-line differential-pressure setpoint using building chilled-water valve positions. This has significantly lowered building pumping heads and allowed us to run the buildings off the campus differential pressure much more than we thought possible. The available pressure increases nicely as building load increases."

Relational Control Algorithms

Relational control algorithms are defined as control sequences that use mathematical relationships between subsystems [as opposed to proportional-integral-derivative (PID) independent temperature or pressure setpoint control]. Examples of relational control would be using orifice position from air-handling unit valves to control chilled-water distribution pump speed or using chiller load to control cooling tower fan speed.

Figure 1. Chilled-Water Delta T Improvement at the University of North Carolina at Chapel Hill, 2004 vs. 2007.

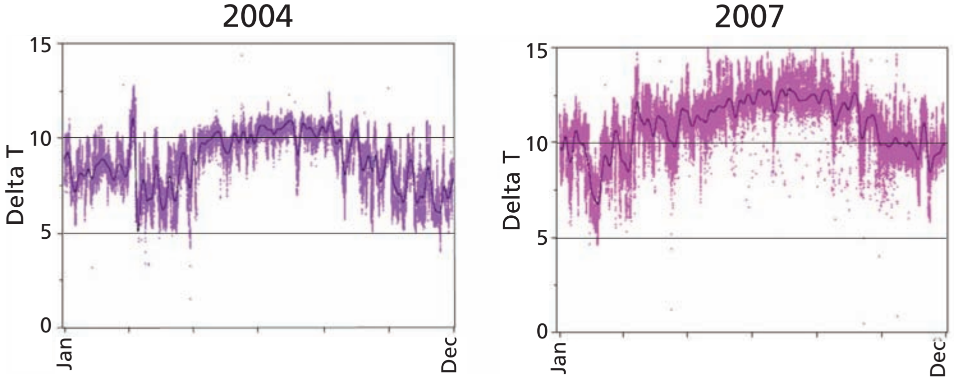

Source: Jim McAdam, University of North Carolina at Chapel Hill.

© 2009 International District Energy Association. ALL RIGHTS RESERVED.

The most important difference between relational controls and PID control is energy optimization. Energy optimization with PID control is accomplished by continually adjusting the various controlled setpoints. But continually changing setpoints leads to additional stability issues and is a reason that operators often discontinue setpoint reset functions. Relational controls are far more effective, providing simpler and more direct energy optimization relationships than what are required when attempting to optimize for energy efficiency by adjusting temperature and pressure setpoints.

Optimization Software

In district energy plants, the priority has always been 100 percent reliability/availability of services. In the past, complex energy efficiency improvement projects were difficult to implement because of the months and months they required of custom programming, commissioning and continuous tweaking of the control algorithms. This process could sometimes take more than one year to complete – with finger pointing, manual overrides and customer complaints weakening the importance and intent of the original goal. Unfortunately, the project still is often at the mercy of the expertise of the local control contractor. While there are many skilled control programmers, many of us never seem to get these guys on our jobs!

Tradeoffs among reliability, complexity and efficiency no longer have to be made with this new optimization approach. The fourth important piece of the paradigm shift enabling ultra-high performance is the introduction of chiller plant optimization software to mitigate the challenges above. Optimum Energy has now packaged complex optimization logic using relational controls in software that is integrated into the existing plant control system (programmable logic control or direct digital control).

The benefits of this optimization software include the following:

- **Reliability** – Software algorithms are prefabricated and can be pretested prior to implementation. The software model allows the complex algorithms to be turned on or off manually (using the existing control system front end) or automatically if there is a communication failure. The ability to 'turn optimization off' offers huge flexibility in troubleshooting and fundamental control sequence redundancy.
- **Scalability** – In the old paradigm, highly custom optimization programming is implemented on nearly every central plant job. Though the relational control algorithms are customized for the specific plant, the patented demand-based software modules allow operators to eliminate PID control and associated loop tuning and not to reinvent the high-performance modules each and every time.
- **Decreased implementation time schedule** – Because of software scalability and the ability to operate the plant without the optimization software, troubleshooting, commissioning and cut-over are much quicker.
- **Performance verification** – Measurement and verification (M&V) of overall plant performance is crucial to ensuring desired efficiency levels have been achieved, and just as importantly, are maintained. The new optimization software offers M&V 'dashboards' for monitoring, trending and documenting savings over time.

Ultra-High Performance Examples

Implementing high-performance solutions in large district and campus cooling plants can save money and help reduce carbon footprints. Figure 2 provides examples of two ultra-high performance district cooling plants – one plant in the United Arab Emirates and the other on a college campus in the desert of Southern California. These facilities present two distinctly different examples of all-electric central plant profiles, average wet-bulb temperatures and total plant kilowatt-per-ton performance, including all chillers, chilled-water pumps, condenser pumps and cooling tower fans. The annual average ultra-high efficiency goal in the United Arab Emirates would be less than 0.7 kW/ton and in California less than 0.52 kW/ton.

When originally designed and constructed, both the UAE and California plants used conventional designs and control methodologies. Based on the one year of trend data, the UAE plant

Figure 2. Ultra-High Performance Profile for District Cooling Plants in the United Arab Emirates and on a Southern California College Campus (Central Plant Load Profile, Kilowatts Per Ton and Average Outdoor Wet-Bulb Temperature).

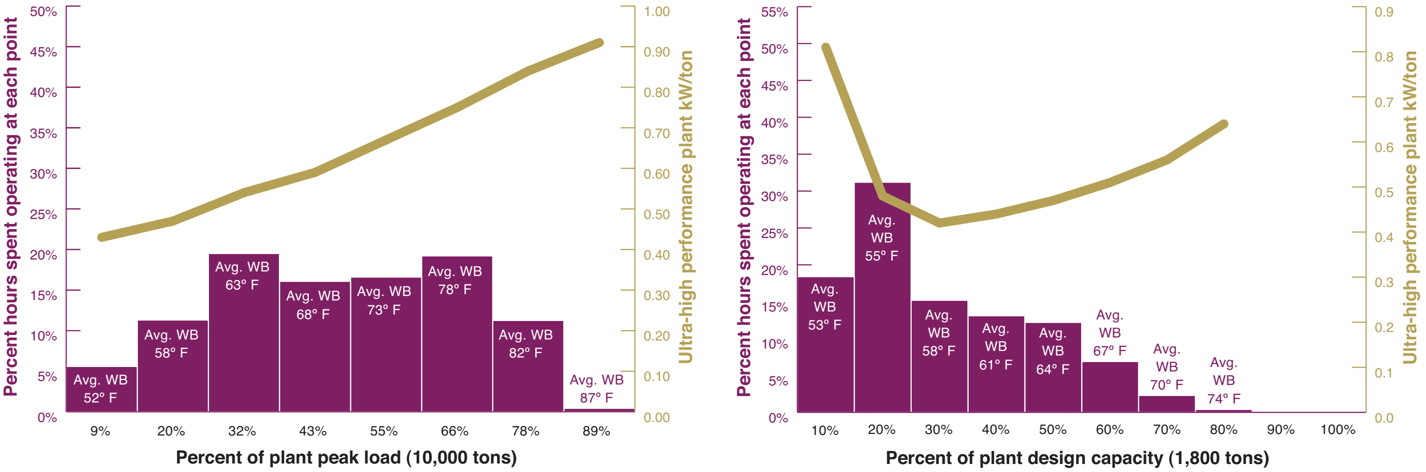

Source: Optimum Energy LLC.

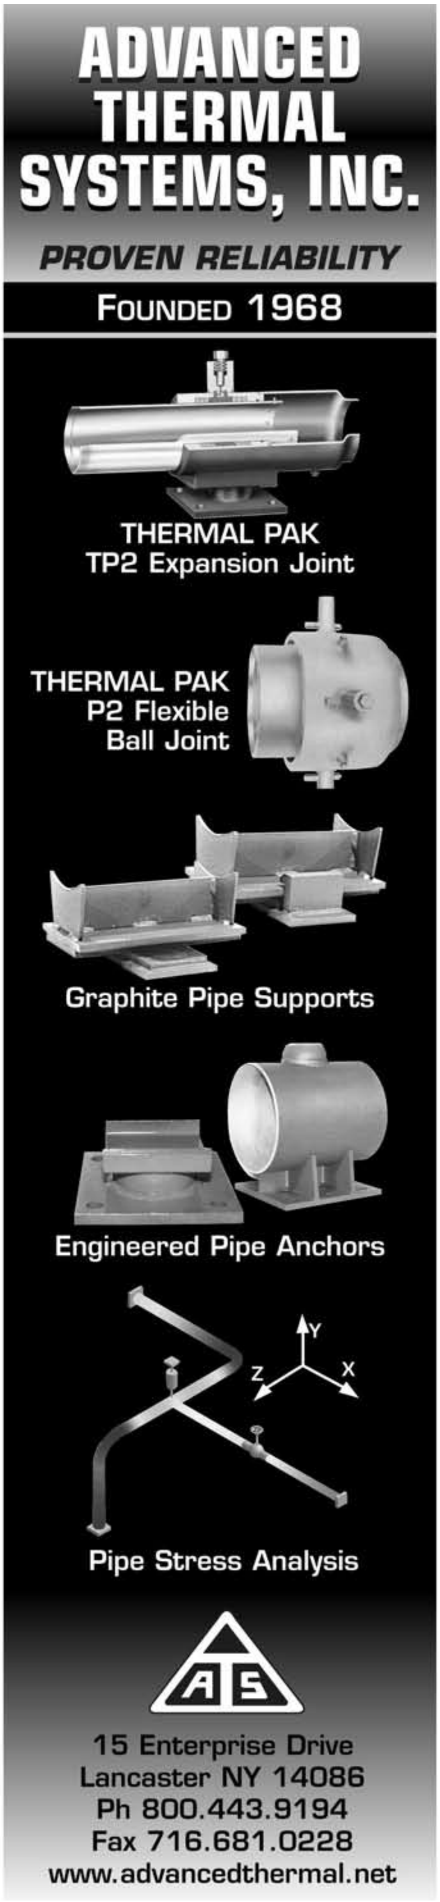

15 Enterprise Drive
Lancaster NY 14086
Ph 800.443.9194
Fax 716.681.0228
www.advancedthermal.net

had an average annual total plant efficiency between 0.90 kW/ton and 1.2 kW/ton; the California plant had an average annual total plant efficiency between 0.80 kW/ton and 1.0 kW/ton.

Even in very humid climates such as the UAE (with design wet bulbs greater than 89 F), district cooling is still a part-load and part-lift industry with a large portion of the hours occurring at less than 74 F wet bulb and 60 percent capacity. One mistake many engineers make is confusing "load" and "lift" in centrifugal chiller plants. "Lift" can be simply defined as the difference between the exiting condenser water temperature (out of the chiller) minus the chilled-water supply temperature. When talking about large district cooling plants, the use of variable-speed and high-performance control algorithms, one must talk first in terms of lift optimization and, secondly, load optimization. Load optimization directs the plant control sequences to run everything as close to full load as possible, while lift optimization lends itself to the use of variable-speed chillers and components.

> One mistake many engineers make is confusing "load" and "lift."

As can be seen in figure 2, the sharp drop in plant efficiency occurs at very low loads (10 percent and below) in the California example vs. the UAE example. This is a result of a two-chiller plant serving the California campus vs. eight chillers in the UAE plant. When one chiller is operating at capacities below 20 percent load, chiller performance falls, and auxiliary equipment must maintain the minimum flow requirements of the chiller condenser and evaporator bundles. Low load is often a problem in newer systems because the plants are built for future growth as the campus builds out. But, without variable speed mitigating these low-load conditions, performance with a standard constant-speed plant at capacities below 20 percent can be anywhere between 2 and 5 kW/ton.

In contrast to conventional control methodologies, relational controls take advantage of the ability to optimize the operation of all variable-speed HVAC-system components in relation to one another and in response to real-time building loads. Additionally, software solutions available today are delivering ultra-high performance across a wide range of plants without compromising reliability.

> A paradigm shift is leading to significantly improved kilowatts per ton.

These solutions can be quickly installed, and they provide persistent, reliable optimization and real-time measurement and verification capabilities that are unprecedented in the industry. Along with M&V reporting, new Web-based software dashboards provide real-time monitoring of plant performance metrics and operator alerts if the system is not performing as expected, thus helping to maintain savings over time. This represents a paradigm shift that is leading to significantly improved kilowatts per ton (COP), providing enhanced automatic plant control and meeting the goals of long-term reduced energy consumption and operating costs.

Ben Erpelding, PE, CEM, is director of engineering for Optimum Energy LLC. He has more than 11 years' experience in energy efficiency, distributed generation, renewable energy and demand response. Prior to joining the firm, he measured and verified actual performance and cost savings for energy efficiency retrofits, photovoltaic installations, demand response audits, and combined heat and power projects while at the nonprofit San Diego Regional Energy Office. Erpelding received a master of science degree, with an emphasis in combined heat and power, and a bachelor of science degree in mechanical engineering from San Diego State University. His email address is ben.erpelding@optimumenergyco.com.

© 2009 International District Energy Association. ALL RIGHTS RESERVED.