LO-09-013

High Performance Computing with High Efficiency

Steve Greenberg, PE | **Amit Khanna** | **William Tschudi, PE**
Member ASHRAE

ABSTRACT

High Performance Scientific Computing typically involves many "clusters" of processors that are close connected. This results in high energy dissipation in tightly compacted areas creating high heat intensity. As these "machines" continue to evolve, cooling requirements become more challenging and the total electrical power requirements more resemble large industrial facilities than typical buildings. Compounding the complexity of the HVAC design is the fact that these computers may be designed for air or liquid cooling. A new computational facility under design for the University of California in Berkeley, CA is such a center that is being designed to accommodate either air or liquid cooling.

This paper describes the unique design features of this center whose goals included both being a model of high performance computing and a showcase for energy efficiency. The mild climate in Berkeley provides an ideal opportunity to minimize energy use through the use of free cooling but traditional data center approaches could not fully take advantage of the mild climate to save energy. A design that utilizes outside air for cooling for all but a few hundred hours per year is described. But in addition, there was a desire to provide for the eventual transition to liquid cooling–in various possible configurations. This capability is also described.

INTRODUCTION

A new supercomputer facility at the University of California – the Computational Research and Theory Facility (CRTF) was designed to incorporate energy efficiency strategies while providing flexibility for a wide range of supercomputer cooling strategies. One of the primary goals of this facility was to provide a design that not only demonstrated world-class computational ability but also demonstrated best practices and novel solutions to the energy use of high performance computers. This one building, with office space, computer room, and infrastructure is expected to more than double the energy use for the campus that it is associated with. An arbitrary power budget of 7.5 MW for the initial buildout and 17 MW for the ultimate buildout was established by management decision. Since the computing sciences group that will occupy the building is judged on computational output, there was strong incentive to maximize the amount of energy available for computational work and to minimize the infrastructure loading.

As a result, a design target for data center infrastructure efficiency (DCiE) was established. This established the ratio of IT energy to energy use of the total facility at 0.83. (Power Usage Effectiveness (PUE) – the inverse of DCiE – target = 1.2). With this goal, the design team was challenged to seek out and implement best practices and push the limits of existing technologies. As can be seen in Figure 1, this would clearly place this center above the centers previously benchmarked by LBNL. While the design is not finalized, DCiE for peak power is predicted to be in the range of 0.83–0.90, and the DCiE for energy is predicted to be in the range of 0.90–0.95.

FLEXIBILITY FOR AIR OR LIQUID COOLING

A key design concept for the CRTF is to accommodate many generations of supercomputer over the course of several decades. While air is the dominant cooling scheme for such machines at present (and thus the first iteration of computers is most likely to be air-cooled), there is a general trend in the industry toward liquid cooling as power densities increase (ASHRAE 2008), and it is anticipated that future equipment

Steve Greenberg is an energy management engineer and **William Tschudi** is a program manager in the Building Technologies program at the Lawrence Berkeley National Laboratory, Berkeley CA. **Amit Khanna** is a senior consultant at Arup North America Ltd., San Francisco, CA.

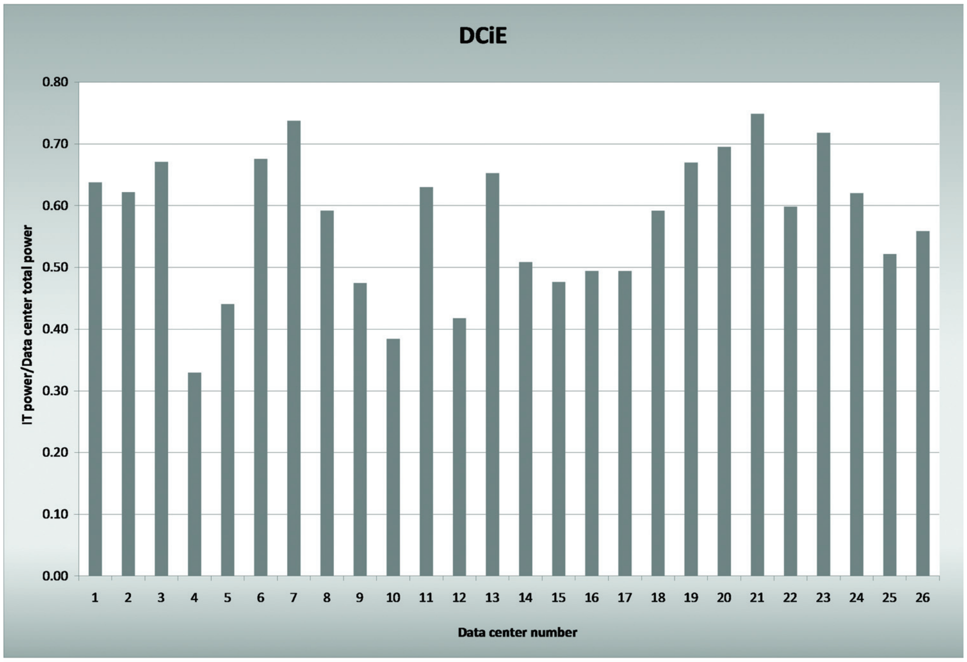

Figure 1 *LBNL Benchmark Results.*

will be largely or completely liquid-cooled. The lower density areas housing memory and network equipment will likely be air-cooled for a longer period than the scientific computing machines. Thus it is imperative to maintain flexibility throughout the facility to cool the IT equipment with air, liquid, or a combination. Adding to the challenge is to maintain this flexibility with maximum energy efficiency and minimum first cost.

USE OF LARGE EFFICIENT AIR HANDLERS

For air cooling, the required airflow is determined primarily by the IT equipment, as well as by how effectively the flow is managed (see "Air Management" section below). With a given flow, the power and energy requirements of the air-handling equipment are determined by the fan and motor efficiencies and by the total pressure drop in the system. All of these strategies are facilitated in the CRTF design by using central AHUs.

The AHUs are located in a level below the computer floor, which frees up expensive raised-floor space and allows for maximum IT placement flexibility in the high-performance computing (HPC) space. See Figure 2.

The AHU configuration is modular, with unit sizes of 100,000 cfm (2800 m^3/min) each, and each bay (20′ (6.1 m) wide) can be equipped with one or two AHUs. (This is a maximum flow, adjusted to meet load using variable-speed fans). If a bay requires two AHUs, they will be vertically stacked in the basement area. The ductwork from the AHU(s) in each bay feeds supply air into the 4′ (1.2 m)-high raised floor plenum in multiple locations (Figure 3). The air is then delivered to the IT equipment either through a typical cold aisle arrangement or directly into the bottoms of the racks, depending on the computer design. The hot air discharged from the equipment is either exhausted via exhaust fans located high on the east wall of the HPC, returned to the AHUs through ductwork down the west wall of the HPC, or most commonly, a combination of the two (see Air-Side Economizer and Evaporative Cooling section below). The modular AHU scheme allows maximum flexibility in intial and staged construction in terms of supplying the right amount of air to the right place with a minimum of excess capacity.

The large cross sectional area of the AHUs results in low face velocities of the filters, cooling coils, and evaporative cooling media of approximately 500 fpm (150 m/min). Supply air from the air handlers is supplied to the plenum via a short ductwork which is designed at 1500 fpm (450 m/min) (max.) at full design flows. These velocities, careful attention to other internal AHU pressure drops, and low pressure drops in the ductwork and air distribution result in total initial static pressure of 1.5″ (380 Pa) at design flow (Table 1). The user group understands the value of timely replacement of filters and direct media pads and is committed to follow an appropriate maintenance schedule.

MODULAR DESIGN

To help achieve an energy-efficient design and control capital cost, a modular approach to the design was incorporated. This approach provides the desired flexibility for future uncertainty while allowing systems to operate more efficiently. Space and other design considerations were provided so that as the facility load is increased additional capacity can easily be added. This approach reduced first cost of the facility but it also allowed components to be sized to better match the load requirements.

Table 1. Initial Pressure Drops Across Major HVAC Components

Components	Initial Pressure Drop in w.g (Pa)
OA louvers	0.15 (38)
OA dampers	0.20 (50)
Filters	0.35 (90)
Direct Evaporative Media Pad	0.25 (62)
CHW Coil	0.25 (62)
Ductwork + Plenum + Outlets	0.3 (75)
Total	1.5 (380)

LIQUID COOLING CAPABILITY

Because the industry is moving toward liquid cooling for IT equipment, the CRTF is designed to accommodate the distribution of cooling water for direct or indirect use at or in the computer racks. A four-pipe distribution scheme is planned, including chilled water from the chiller plant (using water-cooled, electrically driven centrifugal chillers) and closed-loop cooling water ("treated water") from the cooling towers (via plate and frame heat exchangers). Mixing valves will allow individual computing systems to use 100% chilled water, 100% treated water, or anything in between as needed to satisfy the entering water temperature requirement (see Figure 4). Chilled water and treated water temperature setpoints and reset schedules will be established to meet requirements in the most energy-efficient manner.

Since no water-cooled IT equipment is anticipated in the initial configuration of the CRTF, the treated water system will be accommodated by appropriate headers, valves, blank-off plates, and space for pipe runs. The chilled water system will initially run only to the AHUs, but taps with valves and blank-off plates will be installed also for future water cooling requirements.

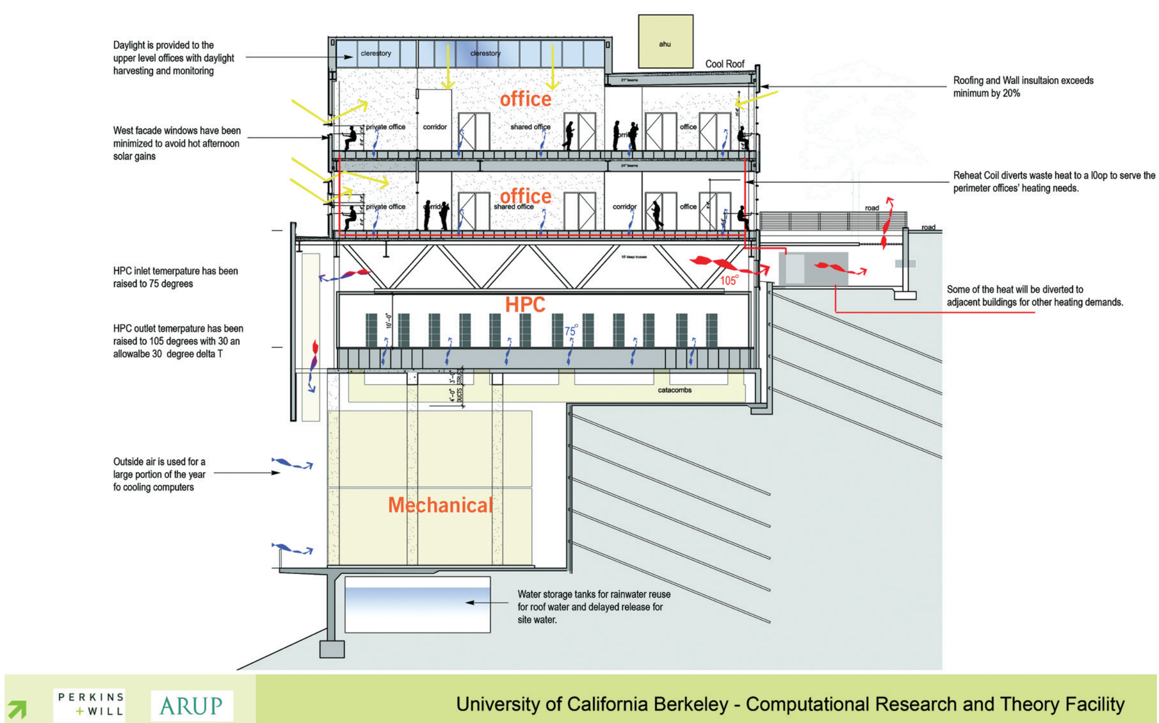

Figure 2 Building section showing AHU located below computer floor, exhaust air path (return is to left and down to AHUs), and office floors above.

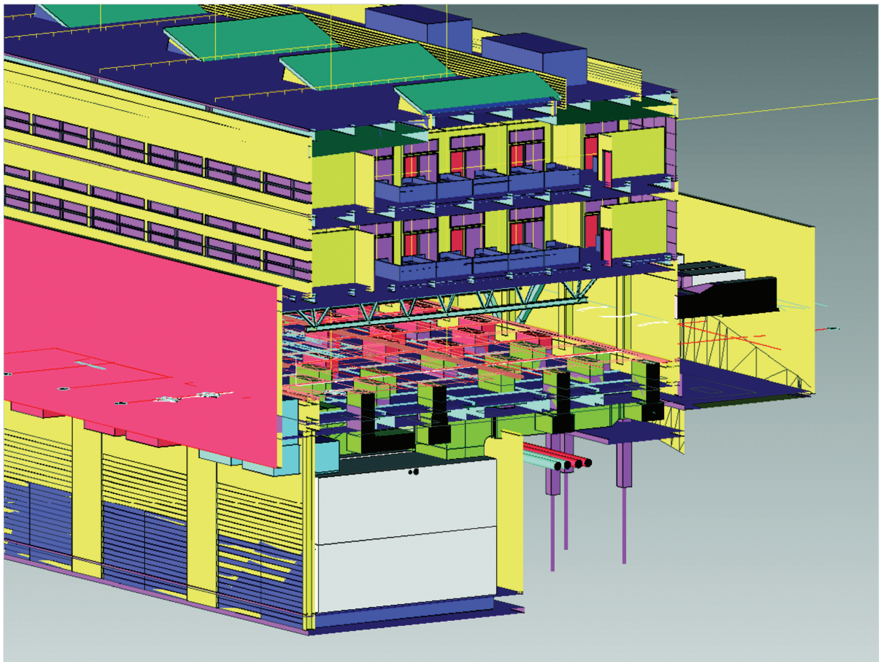

Figure 3 *3-D image illustrating air movement from outside air louvers through to underfloor plenum via multiple 3′× 8′ (0.9 m × 2.4 m) punctures at the structural slab. Design documents were produced using 3-D Revit models for better coordination between disciplines. [Source: Arup]*

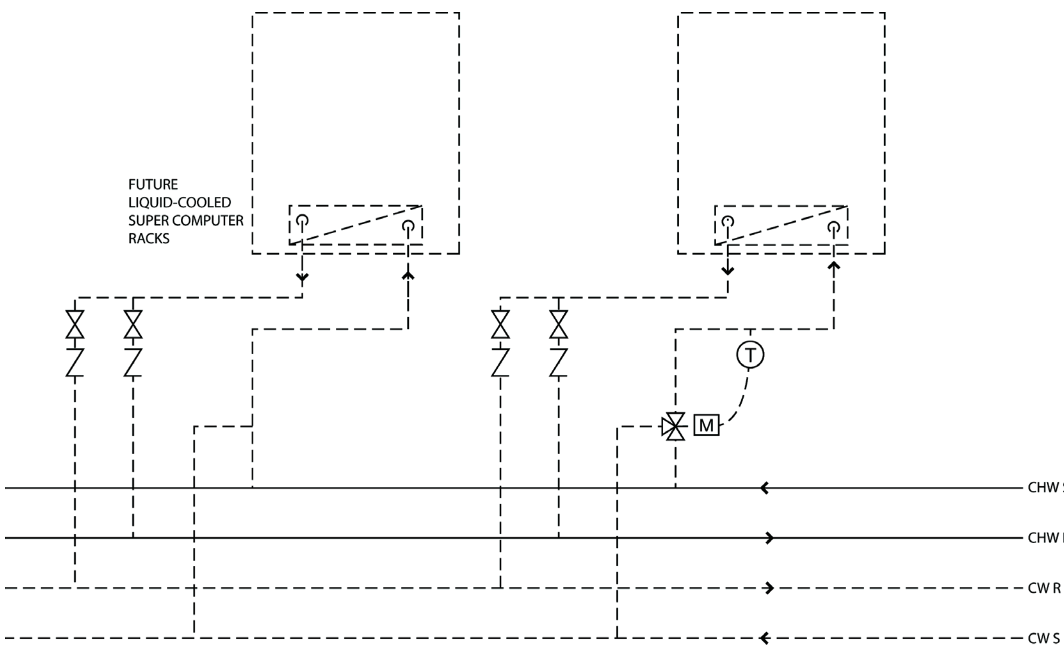

Figure 4 *Four-pipe cooling water system. The chilled water supply and return are the solid lines; the dashed lines are future closed-loop cooling water supply and return (via cooling towers and heat exchangers; see Figure 7).*

PART-LOAD MODULATION

The CRTF load is expected to grow from an initial load of 7.5 MW to at least 17 MW over the course of several years. The load will vary over shorter time periods as computing systems are added, changed, and turned on and off for maintenance. In addition, the weather variation will result in diurnal and seasonal load changes. It is key to the operation of the facility that all of these load variations be met in a way that provides uninterrupted service, but that modulates in an efficient manner. To the latter end, the cooling plant will be modular, and all of the significant loads in the plant and system (tower fans; chiller compressors; chilled, tower, and treated water pumps; and AHU and exhaust fans are all designed with variable-frequency drives. Part-load curves will be integrated into the building automation system so that overall energy and power use are minimized at any combination of cooling load and outdoor conditions.

ENVIRONMENTAL CONDITIONS

The project team debated whether ASHRAE TC9.9 recommended environmental conditions (ASHRAE 2004) could be used as a design basis for the facility since some of the supercomputers on the market required more stringent conditions. To resolve whether ASHRAE recommended ranges could be specified, a workshop was held with all of the major supercomputer vendors where it was agreed that all of the vendors would agree to using the recommended ranges. Subsequent to this meeting the TC9.9 committee voted to broaden the recommended ranges even further. With these assurances, the team agreed to use a maximum of 77°F (25°C) as the design temperature for the inlet of the IT equipment. Of course for much of the year in Berkeley by using outside air for cooling the temperatures could be lower than 77°F (25°C). A broad design humidity range was also established at 30–60% RH at the inlet to the IT equipment.

AIR-SIDE ECONOMIZER AND EVAPORATIVE COOLING

The location of the CRTF in Berkeley, California (across the bay from San Francisco), and the design indoor conditions, allows nearly all of the air cooling to be provided by outside air. The CRTF design indoor conditions are 60 to 77°F (16 to 25°C) dry-bulb and 30–60% RH, as noted above. Because the facility needs to be able to meet the indoor conditions at all times, outdoor temperature extremes (beyond normal summer design temperature) were assumed with a 100°F (38°C) dry-bulb and 65°F (18°C) coincident wet bulb chosen as the design condition.

Given the above design conditions, analysis of the psychrometric data (see Figure 5) shows that the system can meet the requirements by operating in one of four modes, as noted in Table 2.

For over 90% of the hours in a year, the indoor conditions can be met by mixing outside and return air (the psychrometric process of this mode is shown by the arrows in Figure 5, which is to first order along lines of constant absolute humidity, since there is negligible latent load in the HPC). Direct evaporative cooling (with a mix of return air as needed) brings the humidity into the proper range when outdoor conditions are too dry, which occurs less than 1% of the year, as does the condition where a combined use of direct evaporative cooling and the chilled water cooling coil is indicated. Approximately 500 hours per year require the chilled water coil alone.

Table 2. HVAC Operating Modes

Operating Mode	Number of Hours Per Year	Percent of Year
Mix of outside and return air	8200 hours	93%
Direct evaporative cooling	45	0.5%
Direct evaporative cooling and chilled water coil	38	0.4%
Chilled water coil	510	6%

By using a wetted media-type (using the sensible heat in either the outside air or the return air to evaporate the water), the CRTF avoids the energy use of steam or infra-red humidifiers. A direct-spray system was considered, but the extra pressure drop caused by the wetted media didn't justify the first and operating cost of the reverse-osmosis or deionization system required for the make-up water needed for the direct-spray system. Figure 6 illustrates the life cycle cost performance of spray nozzle type strategy compared to pad-media for CRTF.

Other alternate strategies were explored such as floor mounted and plenum located humidifiers but were discarded due to concerns regarding non-uniform distribution, and owners preference to keep the plenum and floor clear for maintenance, accessibility, and flexibility. It must be noted that the facility will have installed multiple super computers (with a variety of rack configurations) at a given time which will be replaced by new-generation super computers every 5 years, so flexible use of the space is very important.

WATER-SIDE ECONOMIZER

When water-based IT cooling is implemented at the CRTF, close-approach cooling towers and plate-and-frame heat exchangers will be used to supply as much of the cooling as possible without operating the chillers. We anticipate that most of the cooling will be provided without the chillers, though until the IT equipment cooling requirements are known, no prediction can be made (Figure 7).

AIR MANAGEMENT

In order to reduce fan energy and ensure adequate cooling for the high intensity computing equipment, it is necessary to separate hot and cold air streams. To accomplish this, the design provides for physical separation of hot and cold sides of the IT equipment. The detailed design will be finalized once

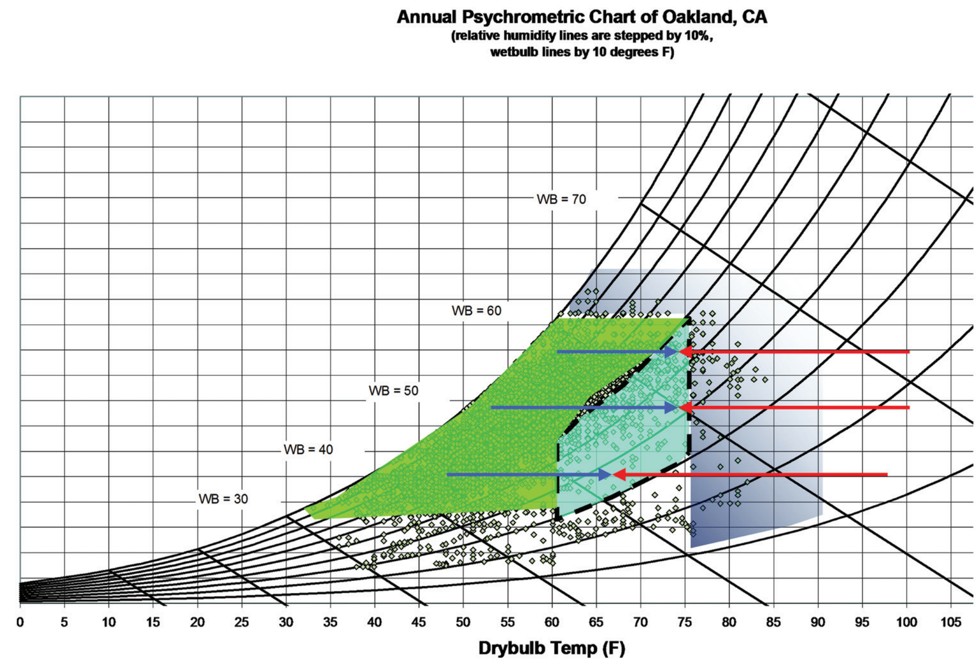

Figure 5 *Psychrometric data for Oakland (adjacent to Berkeley). The dashed line of 60–77°F (16–25°C) dry-bulb and 30-60% RH is the designed supply air condition. The arrows represent mixing of outside air and return air.*

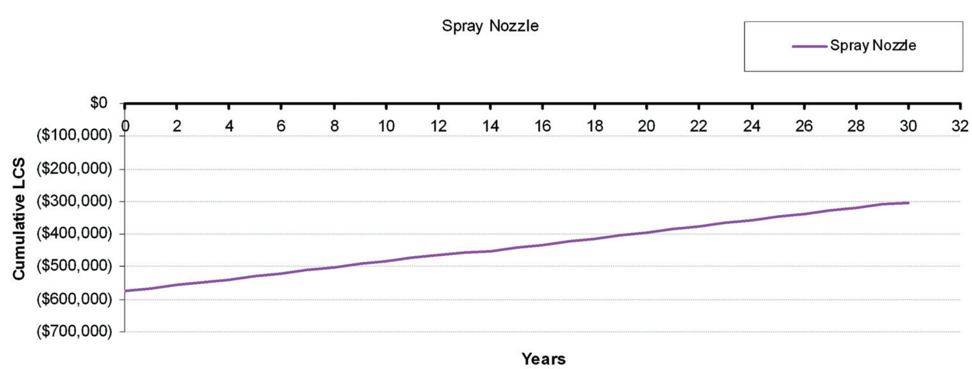

Figure 6 *Life Cycle Cost performance of 'spray nozzle-type' system compared to base-case pad-media type direct evaporative cooler.*

the IT equipment set is defined. By providing a high ceiling area, hot exhaust from the IT equipment is collected and flows to the building exhaust or is recirculated for cooling or mixing with outside air. To confirm that the design would perform as intended, Computational Fluid Dynamics (CFD) models were used. CFD modeling confirmed that the air system would function as designed (Figure 8).

Since the IT equipment environmental condition requirements and heat intensity may vary widely, zoning of the HPC area is required, both above and below the raised floor. Zoning will allow matching of airflow and environmental conditions to the particular needs of the equipment in that zone. For example, a central data storage area is planned that will operate at significantly lower load density compared to the supercomputers. Zoning will allow this area to be operated at much different conditions than the main computing areas. Also mixing liquid and air cooling solutions could require varying amounts of air, so providing zoning capability will allow optimized cooling for the various conditions.

HEAT RECOVERY

Given that over 80% of the over 7 MW entering the CRTF will be converted to heat by the IT equipment, and with

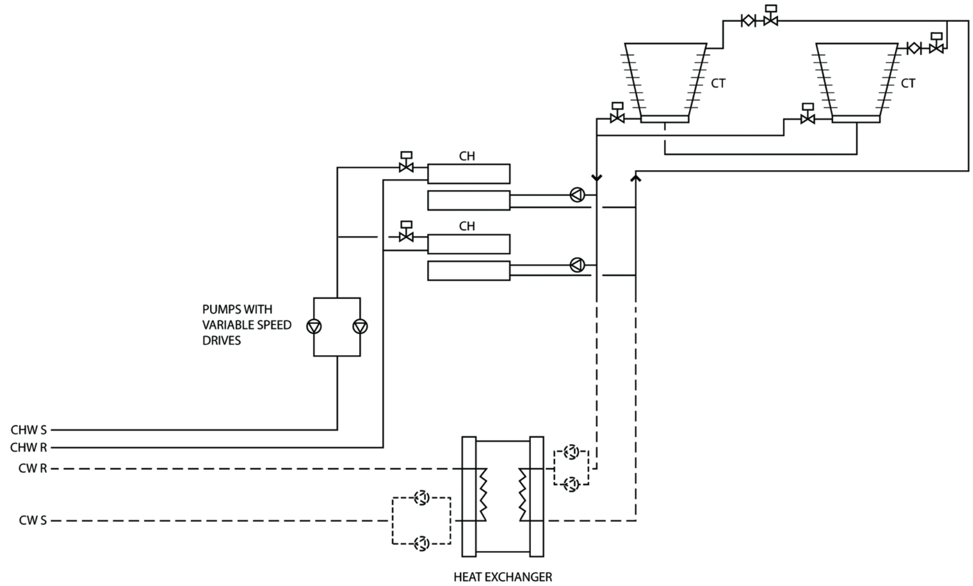

Figure 7 *Water-Side Economizer.*

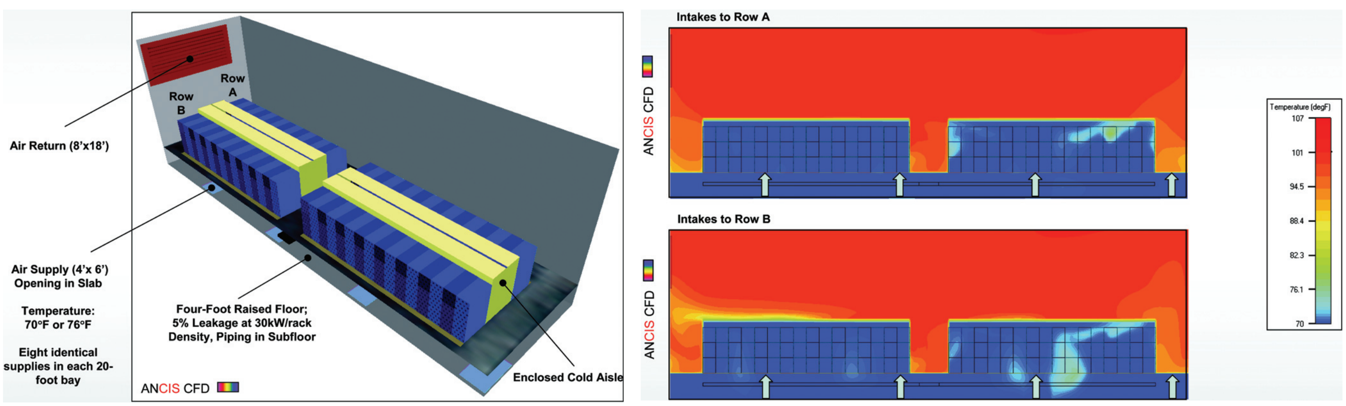

Figure 8 *Computational Fluid Dynamics study was performed to optimize under floor plenum design. [Source: ANCIS Incorporated]*

discharge air temperatures of about 100°F (38°C), there is significant opportunity to recover this heat. The office floors of the building are being designed to high standards of energy efficiency, but will require some heating, making heat recovery within the building an obvious opportunity. There is so much waste heat available from the CRTF that a whole cluster of nearby buildings, including laboratories that need high volumes of outside air on a continuous basis and thus their thermal loads are dominated by heating demand in the Berkeley climate, could be completely heated in lieu of using their boilers fired by natural gas. The feasibility of these opportunities (using such technologies as run-around air-to-water coils and a local district heating system) is being investigated.

POWER DISTRIBUTION

Another opportunity for data center energy and power savings is in the electrical distribution system. The CRTF is fortunate in that most of the IT equipment, since it is used for scientific computing and is thus not as critical as many computing services, does not need to be on uninterruptible power supplies (UPSs). Of the initial computer load of over 6 MW, only 500 kW will be on UPS power. The topology of the

UPS is still being decided, but high efficiency across a wide range of loading is a prime selection criterion.

Another area in which the CRTF minimizes distribution losses is the use of IT equipment that takes 480-volt power directly at the racks. Thus, for the main computing load, there is only one voltage transformation (from the 12-kV site distribution system to 480 volts). Using 480 rather than the normal 208-volt distribution means further savings due to lower currents in the building wiring.

CONCLUSION

The University of California's CRT facility is designed to be a model for energy efficiency as well as a leading scientific computing facility. Best practices identified through previous investigations and benchmarking are incorporated into the design resulting in a very energy-efficient design with a DCiE above the design goal of 0.83. Through aggressive use of the upper end of the ASHRAE recommended temperature range coupled with direct use of outside air in the favorable Berkeley climate, most of the cooling for the center does not require use of chillers for air cooling. If liquid cooling emerges as the preferred vendor cooling solution, then additional energy and power efficiency gains are possible, though with added water use.

ACKNOWLEDGEMENTS

Special thanks are extended to the California Institute for Energy Efficiency (CIEE), the California Energy Commission Public Interest Energy Research Industrial program, the US Department of Energy, and Pacific Gas and Electric Company for sponsoring work performed by Lawrence Berkeley National Laboratory. Thanks to Arup and Perkins+Will for providing design data and details. ANCIS Incorporated performed the Computational Fluid Dynamics (CFD) modeling analysis for the project.

REFERENCES

ARUP 2008. “LCCA Systems Discussion: Progress Update, CRT LBNL”. ARUP and Partners California, Ltd. San Francisco, June 25, 2008.

ASHRAE 2004. *Thermal Guidelines for Data Processing Environments.* American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc. Atlanta, 2004.

ASHRAE 2008. *Best Practices for Datacom Facility Energy Efficiency,* p. 131. American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc. Atlanta, 2008.