DOI:10.1145/1721654.1721671

Article development led by **acmqueue**
queue.acm.org

What can be done to make cooling systems in data centers more energy efficient?

BY ANDY WOODS

Cooling the Data Center

POWER GENERATION ACCOUNTS for about 40% to 45% of the primary energy supply in the U.S. and the U.K., and a good fraction is used to heat, cool, and ventilate buildings. A new and growing challenge in this sector concerns computer data centers and other equipment used to cool computer data systems. On the order of six billion kilowatt hours of power was used in data centers in 2006 in the U.S., representing about 1.5% of the country's electricity consumption. Of this power demand, much more than 20% is typically used for cooling the computer equipment, but some newer installations have managed to reduce consumption through a series of innovations in the

design of data-center cooling systems, as well as improvements in the software and hardware.

The need to control power consumption is of increasing importance as computing power grows and large data centers house increasingly dense arrays of servers. A number of different systems can be adopted to provide cooling for these large, energy-intense buildings, and opportunities continue to emerge for improving the energy efficiency of cooling schemes.

This article reviews some of the generic approaches to cooling and identifies opportunities for further innovation. In assessing the energy demand for cooling, there are both external considerations, relating to the climate and the associated heat losses between the building and the exterior, and internal considerations, relating to the method of cooling to be adopted, and the associated challenges of achieving efficient heat exchange.

Reference Model

In the models and simple calculations presented here, I have adopted some simplified assumptions about power generation in data centers, although the numbers can be scaled to different system sizes. The engineering design of servers includes both the horizontal rack server and the vertically stacked rack servers, which can have energy densities of 10 kW per square meter of floor area. Hence, for a 4,000-square-meter data center, this represents a tremendous heat load up to 350GW hours per year.

Two approaches can be taken to cool the servers: water or air. Air has the advantage that it can be passed directly through the space containing the equipment, and hence may be easy to circulate. Unless the airflow path is targeted to the regions of high heat production, however, air-cooled systems will likely involve spatial variations in temperature. Therefore, to maintain the equipment at the correct operating temperature, the main volume of air will need to be somewhat colder. Of-

PHOTOGRAPH BY SIMON NORFOLK/INSTITUTE

CHILLED WATER SUPPLY

CHILLED WATER RETURN

SUPPLY

RETURN

ten there is not only a main supply of air to the space, or racks of servers, but also small fans adjacent to equipment with high heat-generation rates that drive large volumes of the main-space air over these specific pieces of equipment. The generation of temperature gradients in the space will still tend to happen if the servers are arranged in clusters. This can lead to a less efficient cooling system if the exterior temperature is higher than the temperature required in the main space; the inefficiency arises from the increase in heat transfer through the fabric of the building as the interior temperature is reduced, which in turn requires additional cooling.

There have been numerous attempts to provide cooling that is targeted more directly to the hot equipment, resulting in smaller temperature gradients. One approach to help achieve such targeted cooling is through the use of water as the heat-exchange fluid since it has a large specific heat and very high density relative to air, thus requiring much less volume flux than air per unit of heat flux being transferred. By locating a cooling coil adjacent to hot equipment, the water can directly convect away the heat load. If the coil has some cooling fins, air passed through the fins will also be cooled by the water circulating through the coil. This provides the opportunity for hybrid cooling, whereby small fans may blow chilled air from the chilling coil to the remainder of the equipment, while the water coil itself is responsible for the heat exchange with the high-heat load equipment. One caveat about water heat exchange, however, is that it requires careful construction so as not to damage the IT systems in case of failure or leakage.

Simple considerations illustrate the typical volume flow rates of air through the system. For example, the American Society of Heating, Refrigerating, and Air Conditioning Engineers suggests that servers be supplied with cooling air in the temperature range of 20°C–25°C (68°F–77°F), while the heated air vented from the systems may be as warm as 35°C (95°F). A heat load of 10kW per square meter requires an airflow rate of 1–2 cubic meters per second per square meter of the floor area. Scaling this up to a 4,000-square-meter data center would require an air-circulation rate

There are a number of energy-efficient ways of rejecting the heat from the building. One critical issue to consider is exchanging the interior air directly with the exterior air.

of 4,000–8,000 cubic meters per second. This air needs to be directed carefully through ducts or plena to ensure there are no recirculation or stagnant regions that will then overheat. With a water-cooling system, the specific heat is much higher, so for the same temperature change we would need only 1–2 cubic meters per second of water distributed through the 4,000-square-meter data center. Here the challenge is in distributing this water effectively throughout the system to limit the possibility of overheating; this might involve a novel rack design with built-in cooled water panels, for example.

Once the air or water has passed through a part of the server system and heated up, a heat-exchange or heat-rejection system needs to transfer this heat flux out of the data center and provide a new source of cooled air or water, which can then circulate through the server system again. Let's now explore some possible approaches for rejection of this heat flux from the data center. First I assess designs for air-to-air cooling, assuming that either the system is air-cooled or the water is cooled by a water-to-air heat-exchange system within the building, with this heat then rejected with an air-to-air device.

Air Cooling: Climatic Challenges

There are a number of energy-efficient ways of rejecting the heat from the building. One critical issue to consider is exchanging the interior air directly with the exterior air; owing to the challenges of dust removal and humidity conditioning of the air, it may be that such direct exchange of air is restricted; air-to-air heat transfer between the interior and exterior may then be achieved using heat exchangers. If external air can be brought directly into the data center, then the heat exchanger can be simplified to, or combined with, the direct exchange of air; the merits of heat exchange or air exchange compared with refrigeration/heating lies in the different energy costs for pumping air through a space or heat exchanger compared with the energy required to cool the interior air using a refrigeration cycle.

Noting the complexity of temperature gradients within the data center, we assess the ideal situation in which the cooling air is able to access directly

the regions of high heat load, thereby minimizing the large temperature gradients that can develop. One key issue relates to the geographical location of the data center and the local climate. A building located in a region with a temperate climate may have considerable energy savings available compared with a building in a hot climate in which the temperature exceeds the desired temperature of the inflow air.

In essence, if the external temperature is lower than $20^{\circ}\text{--}25^{\circ}\text{C}$ ($68^{\circ}\text{--}77^{\circ}\text{F}$), it may be possible to use a direct heat-exchange system to reduce the temperature of the interior air, which is recirculated through the space by passing through a heat exchanger also connected to the external air (see Figure 1). Such direct heat exchange can provide an effective means of cooling the air, since the heat exchanger requires only mechanical power to drive the air through the system. Furthermore, if the external air can be brought directly into the interior, some of the mechanical work associated with driving the air through a heat exchanger (for example, a thermal wheel) can be reduced, although in cold external conditions, it may be necessary to preheat the incoming air to a comfortable or acceptable temperature. If this is not achieved through a heat exchanger with the heated interior air, it can be achieved by direct mixing with the interior air.^a

If the exterior air temperature rises above the temperature range, $20^{\circ}\text{--}25^{\circ}\text{C}$ ($68^{\circ}\text{--}77^{\circ}\text{F}$), then the recirculated air may need some direct cooling through a heat pump or other refrigeration system. As long as the exterior temperature is not in excess of the temperature of the hot outflow air, however, it may be possible to use direct heat exchange to reduce the temperature of the air in the space and thereby complement the use of the heat pump, which would then provide the last phases of cooling (Figure 2). The seasonal fluctuation in the temperature of the exterior can therefore be enormously significant, since it limits the range of conditions for which direct heat exchange rather than refrigeration is required. In a predominantly winter climate, the direct heat-exchange approach may be very effective and can lead to an efficient sys-

Figure 1. Schematic of a direct heat-exchange system for cooling a data center in winter.

Figure 2. Schematic of a hybrid heat-exchange system, using some refrigeration in mild summer conditions to complement the direct heat rejection in the air-to-air heat exchanger.

tem, which may need a backup cooling system only for very hot days. One of the challenges of such a system is the capital cost of including both a heat pump and a heat-exchange system, although the operating costs can be substantially reduced if the heat pump is used for only a small fraction of the year.

For example, the climate in Berkeley, CA, typically has only 330 hours when the temperature exceeds 77°F (25°C), as seen in Figure 3; and 730 hours with a temperature in excess of 72°F (22°C), for which a simple heat exchanger would need supplementary cooling capacity. Since data centers typically run continuously, this represents less than 8% to 9% of the year, and so a direct heat-exchange system (Figure 1) may have considerable merit.

With a direct heat-exchange system, the interior air transfers its heat to the exterior, for example, by passing through a series of parallel plates for air-to-air heat exchange or by using a rotating heat-exchange wheel. There is then no thermodynamic work required for cooling, although there is mechani-

cal work driving the air through the heat-exchange system. In contrast, in hotter conditions this approach would make it possible to some exchange heat with the exterior, but then additional cooling would be required to bring the temperature of the inflowing air to $20^{\circ}\text{--}25^{\circ}\text{C}$ ($68^{\circ}\text{--}77^{\circ}\text{F}$).

In an air-to-air heat exchanger the main energy consumed is in the work done to drive the air through the heat exchanger. Such heat exchangers can be relatively efficient. For example, a rotary-wheel heat exchanger may involve a pressure loss of about 200–300 Pa (pascal), and with a flow rate of 4,000 cubic meters per second and hence many such heat exchangers, this would involve a work load of about 0.8MW–1.2MW (or perhaps more), depending on the efficiency and number of fans (see Figure 4).

This would then allow for cooling loads of 10MW to be achieved using about 10% of this cooling load in mechanical work for the heat exchange—although it is important to recognize there would be additional losses in

^a A. Woods et al. *Energy and Buildings*, Elsevier, 2009.

the overall efficiency of the system, perhaps associated with losses in ducts and with the efficiency of the air-pumping system. It is worth noting that we would expect this headline energy-consumption figure to be greater than but of comparable magnitude to the energy required to supply and mix air directly from the outside in the case that direct air-exchange ventilation is possible without a heat exchanger.

Given the inefficiencies in practical implementation of this approach, we may expect the energy cost of direct heat-exchange cooling to be in excess of 10%–20% of the cooling load; this can be compared with the energy loss in a cooling device—for example, a heat

pump, which might operate with a COP factor of 1:4 or 1:5, representing about 20%–25% of the cooling load (Figure 5). This system of direct mechanical air recirculation, with heat rejection to the exterior through a heat-pump or chiller system, will need to be used when the exterior is warmer than the interior (that is, $> 25^{\circ}\text{C}$, 77°F), so that direct exchange is no longer viable. In running a refrigeration system, there is a balance between the amount to which the air temperature is reduced and the flow rate of the air, so as to achieve the same overall cooling flux. The optimal balance between the reduction in temperature and the flow rate depends on the energy required by the fans, which

increases as the flow rate increases, and the efficiency of the cooling heat pump, which typically decreases as the temperature difference across the heat pump increases. By optimizing the couple system, the most energy-efficient mode of operation of the mechanical cooling may be found.

These simple observations, based on the external climatic conditions and the internal heat loads, identify the main challenge in data centers to be the circulation of large volumes of air (or cooling fluid) to remove the heat; this uses a considerable amount of energy, and in hot conditions, the need to chill the recirculated air in addition to the pumping work increases this energy load considerably.

The typical savings that may be achieved by adopting the above ideas, relative to use of a refrigeration cycle throughout the year, are possibly very substantial; using direct heat exchange when external conditions are colder could represent a large percentage of savings in the energy for cooling. The ideas presented here may be able to reduce that energy consumption substantially, but this depends on the ambient temperature; some schemes that adopt this general approach have been installed, including those by Intel and the KyotoCooling system, but there seems to be tremendous potential for much more widespread adoption of the design principle.

This analysis, however, does point to the need to locate data centers in cooler climatic zones, where the exterior temperature may allow for direct heat exchange for a greater part of the year. Indeed, as the number of hours for which cooling is required increases, the energy consumed in the data center increases. If the number of hours increases by a fraction x of the total year, then the energy consumption may increase by an amount on the order of $0.2x$ of the total energy consumed in powering the IT equipment, assuming the use of direct cooling accounts for about one-quarter of the energy consumption. This points to the benefits of locating data centers in more northerly or colder climatic zones.

Interior Design Considerations for Air Cooling

Let's now turn to the design of the in-

Figure 3. Example data on the number of hours the temperature exceeds different values per year in Berkeley, California.

Figure 4. Illustration of the fan power required for a given flow rate of air through a building. This is relevant for the power consumption in a direct heat-exchange system, where the power to drive air through the heat exchanger is the main load on the system.

terior to help achieve more energy savings. One of the main challenges within the data center is the distribution of cooled air to achieve the cooling of each server. As noted earlier, if not all of the chilled air reaches the equipment with high heat load, then temperature gradients become established as heat is transferred from this equipment to the main airstream; although in equilibrium the outflowing mixed air stream will be the same, the equipment may be hotter than in a “well-mixed” model for a given temperature. The main challenge with such stratification is that in order to keep the equipment within the desired temperature, the surrounding air needs to be cooler than the desired equipment temperature. In hot external conditions, maintaining the main space at lower temperatures, in order to keep the equipment at the required temperature, may lead to overcooling; this is because the heat gains to the space through the insulated envelope of the building will increase as the interior temperature falls.

Depending on the heat transfer across the walls and ceiling, with 10°C (50°F) of additional cooling, this may lead to an additional heat gain of up to 20–30 watts per square meter of the data center. In a 4,000-square-meter data center, this represents nearly 0.1MW of power for cooling. Although this is only a small percentage of the total load, it is nonetheless significant. In addition, any air that is exchanged with the exterior for ventilation purposes will need to be cooled a further 10°C (50°F), again increasing the cooling load.

Often underfloor air-distribution systems are used to provide a uniform supply of air through the equipment being cooled. One approach is to supply the air through a “cold” corridor and extract the hot air from the adjacent “hot” corridor, between successive rows of servers. This creates, therefore, a cross-flow across the rows of cooling equipment, transferring heat to the circulating air; the idea is to recycle air through the cooling plant below the floor. This configuration, however, sets up temperature gradients.

Figure 6 shows a data center with rows of servers, with cold air entering through the floor and passing through the racks of servers and then exiting through the next hot corridor back to

Figure 5. A schematic approach to the cooling of a data center using an air-to-air refrigeration cycle.

Figure 6. Schematic of an underfloor cooling system.

Figure 7. Schematic of a pure upflow system with a cooling unit in the floor of the data center.

the cooling system below the floor. The hot return air may be cooled by a refrigeration plant located below the floor, or by using a direct heat exchanger system with the exterior air, depending on the external temperature, as discussed earlier. The yellow arrow in Figure 6 shows how the heat-exchange system vents the heat to the exterior, perhaps through a wet-coil system.

This design, which involves a series of inflow and outflow corridors, is already in use in some data centers, but

there may be opportunities to enhance the energy efficiency of the system.

One challenge is that the flows are running against gravity in the hot corridors if the air is extracted through the floor, and any mixing of the hot and cold air zones then requires an even greater flow rate past the servers to avoid overheating. One resolution of this problem may be to supply air at a low level to the server rack and extract directly above the top of the server rack. The air can then be ducted into

the refrigeration unit and resupplied to the space (Figure 7). This can lead to substantial stratification in that the air entering the equipment racks is cold, while that leaving at a high level has been heated by heat exchange; such a configuration is beneficial from an energy perspective, since the cold air passes over the hot equipment and warms up, venting from the space at a much higher temperature. The vertical extent of the heat load requires careful limits, however, so that the equipment at the high level does not become unsustainably hot owing to the ascent of heat from the lower level.

This design reduces the chance of heating the cold-supply air, except by heat exchange from the server, and hence can lead to reduced airflows. Ideally, this system would be combined with the direct heat exchange to the exterior (figures 1 and 2) so that the refrigeration unit is used as little as possible for cooling (such as, only in very hot external temperature conditions).

Another area for potential enhancement of the energy efficiency is the geometrical arrangement of the equipment itself, and the associated control of the use of the machines depending on the computing load. One of the key messages stated earlier related to the development of temperature gradients between the heat-generating equipment and the air in the main space; this situation may require lowering the main-space temperature to maintain the heat transfer from the equipment so it can operate within the correct range of operating temperatures. Stacking and localizing equipment may enhance the development of such temperature gradients or may require larger local air circulation rates. Using an upflow displacement scheme and distributing the equipment that is generating heat across the floor space will minimize the buildup of such temperature gradients, and, hence, either the degree of cooling required at any time or the flow rate of the cooling air through specific supply pathways. The integration of the server switching and virtualization software with the energy efficiency of cooling the system offers the potential to develop strategies of hardware use that can lead to substantial savings in the cooling loads and the airflow-rate requirements for a given

In an energy-efficient world, the primary use of the servers would be minimized so as to reduce the primary heat generation.

activity of the hardware, especially in conditions where some latency exists in the data-center hardware capacity.

Summary

This simplified picture of computer data centers has illustrated the considerable variation in energy performance of the center based on the design of the cooling system for the servers. In an energy-efficient world, the primary use of the servers would be minimized so as to reduce the primary heat generation. That is a software challenge that can be addressed using virtual server technology and other approaches to minimize the use of computing resources without compromising on the speed of the system. Given a heat load, however, it is clear that with careful design the convective flow patterns of the air in the data server centers, coupled with the use of a hybrid heat-exchange/refrigeration system can substantially reduce the additional energy required to cool the servers in operation.

In addition to the design of the direct heat-exchange system and the fans, there are many options for optimizing the energy efficiency of the refrigeration system, including the use of ground-source heat pumps and air-to-air heat pumps. Given the large expansion in data centers for Internet and bank service providers, optimizing the design of their cooling systems has the potential for major energy savings, as well as significantly reducing the operating costs of the centers.

Related articles on queue.acm.org

A Conversation with Steve Furber

<http://queue.acm.org/detail.cfm?id=1716385>

Power-Efficient Software

Eric Saxe

<http://queue.acm.org/detail.cfm?id=1698225>

Powering Down

Matthew Garrett

<http://queue.acm.org/detail.cfm?id=1331293>

Andy Woods is head of the BP Institute at the University of Cambridge, England. His research interests lie in developing mathematical and experimental models of fluid flow and heat transfer in natural systems, including energy efficiency and natural ventilation in buildings and the thermodynamics of heat pumps, as well as the dynamics of flows in porous rocks for geothermal power generation, oil and gas production, and carbon dioxide sequestration.

© 2010 ACM 0001-0782/10/0400 \$10.00