Energy Conversion and Management 58 (2012) 26–34

Contents lists available at SciVerse ScienceDirect

Energy Conversion and Management

journal homepage: www.elsevier.com/locate/enconman

Thermodynamic feasibility of harvesting data center waste heat to drive an absorption chiller

Anna Haywood [b,*], Jon Sherbeck [b], Patrick Phelan [b,c], Georgios Varsamopoulos [a], Sandeep K.S. Gupta [a]

[a] School of Computing, Informatics, and Decision Systems Engineering, Arizona State University, P.O. Box 9309, Tempe, AZ 85287-9309, USA
[b] School for Engineering of Matter, Transport & Energy, Arizona State University, P.O. Box 9309, Tempe, AZ 85287-9309, USA
[c] National Center of Excellence on SMART Innovations, Arizona State University, P.O. Box 9309, Tempe, AZ 85287-9309, USA

ARTICLE INFO

Article history:
Received 12 January 2011
Received in revised form 28 December 2011
Accepted 28 December 2011
Available online 2 February 2012

Keywords:
PUE
Data center
CRAC
Absorption chiller
Blade server
Green computing

ABSTRACT

More than half the energy to run a data center can be consumed by vapor-compression equipment that cools the center. To reduce consumption and recycle otherwise wasted thermal energy, this paper proposes an alternative cooling architecture that is heat driven and leads to a more efficient data center in terms of power usage effectiveness (PUE). The primary thermal source is waste heat produced by CPUs on each server blade. The main challenge is capturing enough of this high-temperature heat to energize an absorption unit. The goal is to capture a high fraction of dissipated thermal power by using a heat capture scheme with water as the heat transfer fluid. To determine if the CPU temperature range and amount of heat are sufficient for chiller operation, we use server software, validation thermocouples, and chip specifications. We compare these results to required values from a simulator tool specific to our chiller model. One challenge is to simultaneously cool the data center and generate enough exergy to drive the cooling process, regardless of the thermal output of the data center equipment. We can address this by adding phase change latent heat storage to consistently deliver the required heat flow and, if necessary, a solar heat source. Even with zero solar contribution, the results show that the number of CPUs we have is sufficient and our PUE indicates a very efficient data center. Adding solar contribution, the steady-state model proposed leads to a potentially realizable PUE value of less than one.

© 2012 Elsevier Ltd. All rights reserved.

1. Introduction

The rise in demand for the important work of data centers has created a noticeable impact on the power grid. The efficiency of data centers has become a topic of concern as the densely packed, energy intensive computer equipment inside a data center are creating power demands that are much higher than those of a standard residence or commercial office building [1]. In fact, data centers can be 40 times more energy intensive than a standard office building and require higher levels of power and cooling [2]. Furthermore, direct power consumption for these datacom facilities is increasing due to growing demand for the services they provide, particularly internet and intranet services [3]. The rise in demand for the important work of data centers has created a noticeable impact on the power grid.

US data centers consume a rising portion of the US electricity supply. According to the August 2, 2007 US Environmental Protection Agency (EPA) report on server and data center energy efficiency to Congress, the US server and data center sector used 61 terawatt hours (TW h) of electricity in 2006 (double the amount consumed in 2000). In 2006, this 61 TW h of electricity represented 1.5% of total US electricity consumption and cost $4.5 billion [3]. Under current efficiency trends, US energy consumption by servers and data centers could nearly double again by 2011 to 107.4 TW h or $7.4 billion annually [4].

Data centers are growing larger in size, consuming more electricity and dissipating more heat. According to a survey administered by the Association for Computer Operation Managers (AFCOM) and InterUnity Group, data center power requirements are increasing by 8% per year on average, and 20% per year in the largest centers [5]. This trend is leading to more problems with temperature control. As processor manufacturers such as Intel, AMD, IBM, and others continue to deliver on Moore's Law – doubling the number of transistors on a piece of silicon every 18 months – the resulting power density increase within the chips leads to dramatically rising temperatures inside and around those chips [6]. As servers become more power dense, more kilowatts are required to run and to cool them [7].

In 2007, The Uptime Institute's whitepaper entitled *Data Center Energy Efficiency and Productivity* stated that the 3-year operational and capital expenditures of powering and cooling servers can be 1.5 times the cost of purchasing server hardware [8]. Three years

* Corresponding author. Tel.: +1 480 965 3199; fax: +1 480 965 2751.
E-mail addresses: anna.haywood@asu.edu, bluetoolteam@impact.asu.edu (A. Haywood).

0196-8904/$ - see front matter © 2012 Elsevier Ltd. All rights reserved.
doi:10.1016/j.enconman.2011.12.017

Nomenclature

COP	coefficient of performance
CRAC	computer room air conditioner
CF	captured fraction of board heat
DCiE	data center infrastructure efficiency
HHF	high heat fraction
PUE	power usage effectiveness
$\dot{Q}_{CRAC}$	heat load on the CRAC, W
$\dot{Q}_{C}$	heat removed by chiller, W
$\dot{Q}_{IT}$	computer equipment heat, W
$\dot{Q}_{HiT,S}$	high-temperature heat from computer racks in data center to thermal storage, W
$\dot{Q}_{Air,C}$	low-temperature heat flow of hot air in room removed by heat-driven chiller, W
$\dot{Q}_{EXT,S}$	heat flow from external source of thermal energy to thermal storage, W
$\dot{Q}_{ext}$	heat from an external supplemental source of thermal energy, W
$\dot{Q}_{DC,CRAC}$	heat flow from the data center removed by the CRAC, W
$\dot{Q}_{S,C}$	heat flow from thermal storage to heat-driven chiller, W
$\dot{Q}_{DC,CT}$	heat flow from data center to cooling tower when utilizing ambient air, W
$\dot{Q}_{C,CT}$	Heat flow from heat-driven chiller to cooling tower, W
$\dot{Q}_{CT,A}$	heat flow from cooling tower to ambient, W
$\dot{Q}_{TOT,DC}$	total heat flowing from data center, W
$\dot{W}_{TOT,DC}$	total electric power into data center, W
$\dot{W}_{IT}$	electric power into racks holding computer equipment, W
$\dot{W}_{Loss}$	electric power loss from PDU to racks, W
$\dot{W}_{Lights}$	electric power into lights, W
$\dot{W}_{Cool}$	electric power to operate the data center cooling system, W
$\dot{W}_{CRAC}$	electric power into CRAC unit, W
$\dot{W}_{AC_FAN}$	electric power into condenser fans for CRAC unit, W
$\dot{W}_{CP}$	electric power into condenser fan for chiller unit, W

was the chosen period as this is the functional life of most servers [5]. Projections out to 2012 show powering and cooling at three times the cost of the hardware under even the best conditions and up to 22 under worst-case assumptions [8].

Of the total power consumed by a typical data center, about half is attributed to conventional cooling [3]. Therefore, an effective way to implement energy efficiency in data centers is to address the cooling power necessary for operations. Since approximately 100% of electrical power into a computer device ($\dot{W}_{IT}$ is dissipated as heat ($\dot{Q}_{IT}$), the central idea of this project is to transfer directly that available thermal energy from the highest power components on a server blade (i.e., the CPUs) to drive a heat-activated cooling process, thereby lessening the heat load on and subsequent electrical grid power consumed by the standard vapor-compression refrigeration typical in data centers [9].

A similar effort in CPU heat reuse is the "zero-emission" data center by IBM [10]. Although in both cases, the type of heat we are targeting specifically is the concentrated heat directly from each CPU, the difference in our approach is how we plan to reuse that heat to drive a cooling process [11].

2. Primary objective and approach

Electrically-driven computer room air conditioners (CRACs) are the standard for cooling data centers, but these vapor-compression units are responsible for consuming up to 50% of power that a typical data center uses [10]. Since cooling is such a significant portion of data center grid power consumption, the primary objective is to reduce the grid power consumption of the cooling system for our data center (DC).

Our primary approach is to employ an absorption refrigeration unit that is energized by thermal energy to provide cooling [11]. Such an approach satisfies the primary objective by reducing DC power consumption in several ways. First, since DC waste heat will be used to drive the chiller, this reduces the amount of cooling power needed to cool that otherwise dissipated waste heat. At the same time, the absorption system provides additional cooling which lessens the load on the CRAC unit.

3. Absorption chiller and performance models

As shown in Fig. 1, the unit proposed is a 10-ton single-effect lithium bromide–water (Li–Br) absorption refrigeration system for our application to utilize data center waste heat and to reduce the cooling load on the CRAC. The rationale behind choosing this specific model is its availability as a donation from the local utility company. This heat-driven system uses lithium bromide salt as the absorbent and water as the refrigerant. According to standard specifications (Table 1), at its design point of 88 °C and 50.2 kW_{th} of heat input and a pump power input of 0.21 kW_e the cooling capacity is 35.2 kW_{th} and its coefficient of performance (COP_C) is 0.7. Here, kW_e is taken to mean kilowatt electrical or electric power whereas kW_{th} is kilowatt thermal or heat power. Below the standard specifications, the unit can be energized at lower values of temperature and heat input, favorable to fluctuating output values of CPUs.

To explore the range of possible efficiencies and cooling capacities within the rated conditions, Fig. 2 is devised from Table 1 and a simulator tool provided by Yazaki Energy Systems, Incorporated. Also, the graphs of Figs. 2 and 3 these graphs can enable us to determine if the CPUs can meet each set of temperatures and heat input requirements to successfully operate the chiller.

The absorption chiller performance depends on many factors including the cooling water inlet temperature, the heat medium flow rate, and the temperature of the input heat, which determine the resulting cooling capacity. Fig. 2 shows the Yazaki WFC-SC10 absorption chiller's performance under standard conditions defined as a heat medium flow of 100% or 2.4 l/s and cooling water temperature of 31 °C. To energize the chiller at this level requires a minimum heat input of 14.1 kW_{th} at 70 °C for a cooling capacity of 8.8 kW_{th} ($COP_C = 0.62$). Although 8.8 kW_{th} provides low realizable cooling under these conditions, it does demonstrate that an absorption machine can be energized using a minimal amount of waste heat and it will be shown that the CPUs can easily reach these possible requirements.

Notice that cooling capacity significantly increases to 35.2 kW_{th} at its design point of $COP_C = 0.7$ and 88 °C. Heat medium flow remains at 100%. Also notice in Fig. 2 that the COP_C reaches a maximum value of 0.86 at lower temperatures. This value corresponds to the least amount of heat input necessary to achieve the highest efficiency. The tradeoff is a lowered cooling capacity of 25.0 kW_{th} at 80 °C. However, this point may be the best compromise to achieving a reasonable cooling capacity for a temperature at which the CPUs can safely operate.

Under maximum conditions, we can achieve a maximum cooling capacity above the WFC-SC10 chiller's rated value. However, as

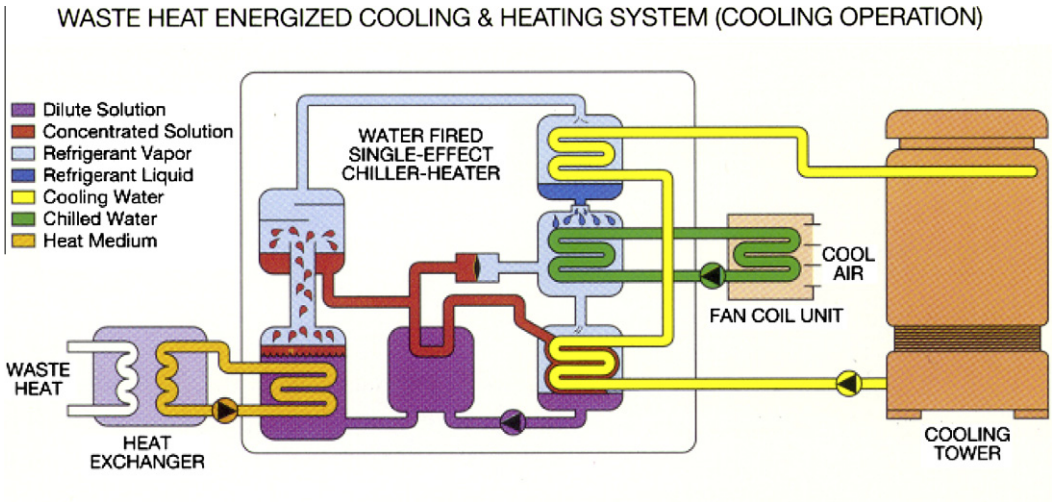

Fig. 1. Model WFC-SC10 cooling cycle with permission from Yazaki Energy Systems, Inc. [12].

Table 1
Specifications in SI units for the Yazaki Water Fired SC10 Chiller under standard conditions [12].

Model	WFC-SC10	Standard
Cooling	Capacity (kW_{th})	35.2
	Chilled water temp. (°C)	7 Outlet, 12.5 inlet
Chilled water	Rated water flow (l/s)	1.5
	Evap. press drop (kPa)	55.8
	Water retention volume (l)	17.0
Cooling water	Heat rejection (kW_{th})	123.9
	Inlet temperature (°C)	31 (standard)
	Rated water flow (l/s)[a]	5.1
	Cond./abs. press. drop (kPa)	84.8
	Water retention volume (l)	65.9
Heat medium (HM)		
	Input (kW_{th})	50.2
	Inlet temperature (°C)	Temperature range
	88 (Standard)	70 (min)–95 (max)
	Rated water flow (l/s)	2.4
	Generator press. drop (kPa)	90.3
	Water retention volume (l)	20.8
Electrical	Power supply	208 V, 60 Hz, 3 ph
	Consumption (W_e)	210

[a] Minimum cooling water flow.

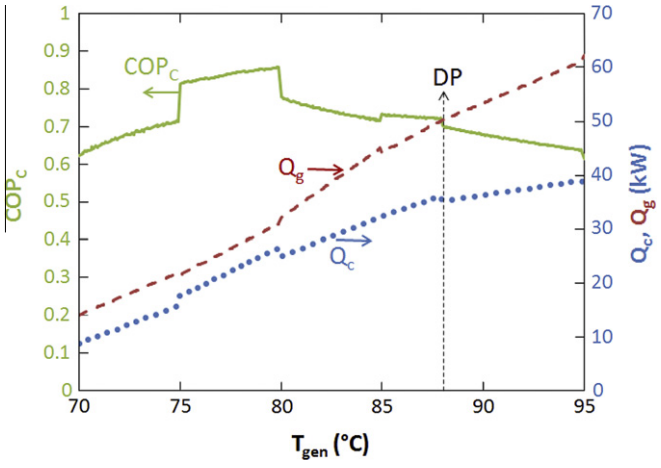

Fig. 2. WFC-SC10 performance under standard conditions. DP denotes its design point of $COP_C = 0.7$ and cooling capacity 35.2 kW_{th} at 88 °C, requiring heat input of 50.2 kW_{th}.

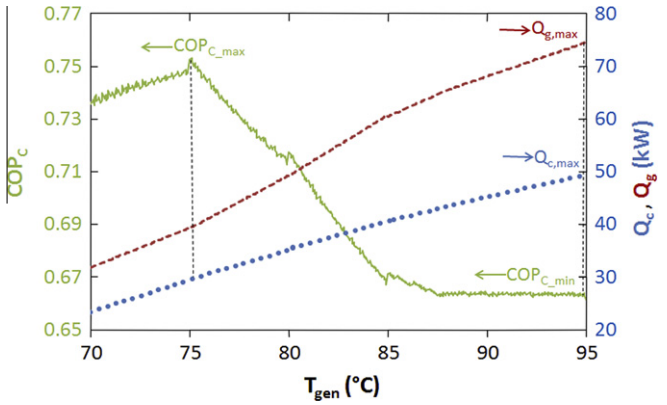

Fig. 3. WFC-SC10 performance at maximum conditions, i.e., cooling water reduced to 26.7 °C.

shown in Fig. 3, there is a distinct tradeoff between maximum efficiency and maximum cooling capacity. For a cooling water temperature lowered to 26.7 °C, maximum COP_c of 0.753 occurs at 75.1 °C but only delivers a cooling capacity of 29.6 kW_{th}. Compare this cooling capacity with the maximum possible of 49.4 kW_{th} at a minimum COP_c of 0.66. In general then, the hotter the heat medium and the lower the cooling water temperatures, the more cooling capacity will be generated. Also, the greater the heat medium flow rate, the more cooling capacity will be generated.

As expected, the COP_C is low at the lowest inlet temperature, but it then peaks in the middle and considerably falls to a minimum value as more input temperature is necessary to generate higher cooling capacities. The ratios of the surface areas of the various tube bundles is such that it is more efficient at releasing refrigerant with the lower temperature up to its design point limit. The COP_C climbs in the middle higher than the design point due to the surface area of the tube bundles. The ratios of surface area of the various tube bundles is such that it is much more efficient at releasing refrigerant with the lower temperature, but the lower temperature itself cannot release as much total refrigerant so the capacity falls below design capacity. In essence, the system is better at releasing refrigerant with the amount of heat it delivers, but it delivers significantly less total heat. The heat medium generator is made of a particular grade

of stainless steel tubing wound into three parallel-flow rows, meaning an inside, middle, and outside tubing coil. At the bottom of the tube bundle, all three of the coils tie into a header. Each coil is wound in a circular pattern and all three are rejoined together at the top of the tube bundle, rejoining in the same configuration. The tube on the outside at the bottom header winds around to be the outside coil in the generator and rejoins the top header at the outside position. The different diameters of each of these bundles means there will be slight differences in how each of these individual coils transfer heat.

At higher temperatures, there is so much input heat that transfer is again limited by the surface area and heat transfer properties of the heat exchange materials and COP_C falls. Then, at the cost of lowered efficiency, the cooling capacity rises to a maximum under each set of environmental conditions. This is advantageous when waste heat is recycled thermal energy; when the heat source is "free", capacity is the main consideration.

4. CPU waste heat potential

The energy sources to run the absorption chiller will be the waste heat dissipated from the IT equipment inside the data center as well as any external thermal supplementation as necessary [13]. Therefore, the benefits realized will be twofold and interconnected: removing a main source of data center waste heat and, consequently, lessening the load on the CRAC. The data center waste heat to drive the chiller will be originated from the main heat producing components (processors) on each of the computer server blades. The server blades of choice are the Dell PowerEdge 1855 blades and the processors are two Intel Xeon Nocona CPUs on each blade.

The main challenge is capturing enough high-temperature heat from the processors on each server blade, and then transporting that heat effectively and efficiently to power a Li–Br absorption chiller [14]. Since the energy sources to run the chiller will be the CPUs, the performance graphs in Fig. 2 and Fig. 3 set the temperature and heat requirements that our CPUs should deliver. From benchtop tests, we discovered that our CPU can withstand more than the required temperature range and multiple processors can deliver the necessary heat to energize the absorption machine.

4.1. CPU temperature data

After installing Dell's Open Manage System Administrator (OMSA) software, we discovered that each processor can operate at temperatures ranging from 10.0 °C to 120.0 °C. The maximum possible temperature is 125.0 °C as shown in Fig. 4. According to OMSA, failure would not occur until 125.0 °C. Therefore, the required heat medium temperature range of 70–95 °C for chiller operation is well within the temperature range of the processors' capabilities.

We tested the CPU temperatures by tasking the CPUs at different levels and modulating the fans. Fig. 4 shows the reported CPU temperature at a high tasking level and significantly reduced fan air flow (all but one of the chassis fans removed). Shutting off the fans yielded temperatures above 100 °C and allowed for conservation of power and reduced noise levels. Also, with liquid cooling proposed, fans would not be used and adjusting coolant flow would be the primary modulation method.

To make certain the OMSA readings were reporting accurately, we further tested the CPU by employing a National Instruments FieldPoint data acquisition device and Labview to measure the response from a validation thermocouple affixed directly onto the CPU case. The results over a 5-min test run are shown in Fig. 5.

Fig. 5 is a validation test of the OMSA software using a K-type thermocouple (TC) until we reached a target temperature of 110 °C. The Open Manage System Administrator (OMSA) software response lags behind the thermocouple (TC) but shows consistency.

From Table 1, the useable range for chiller operation is between 70 °C and 95 °C. Figs. 4 and 5 agree that the CPUs are able to operate within the required chiller temperature range. Furthermore, the CPU temperatures can be reasonably stabilized by modulating the fans as shown in Fig. 6. To observe the smaller temperature resolution in Fig. 6, we used the TC and Labview setup.

Taking into consideration a ΔT_{CH} loss between the CPU case and a cold plate interface, we use the general principle that the temperature drop ΔT across a given absolute thermal resistance R_θ is the product of said resistance and a given heat flow Q through it [15]. Applying this product to our application, $\Delta T_{CH} = Q \cdot R_{\theta CH}$, where Q is the average maximum heat dissipation of the CPU (103 W) and $R_{\theta CH} = 0.1$ K/W (a typical value for a heat-transfer pad) the expected ΔT_{CH} loss is then 10.3 °C [15]. Doubling this value, we arrive at an overall estimated system ΔT loss of about 20 °C.

Assuming a maximum system ΔT loss of 20 °C between the CPU and the heat medium inlet to the chiller, Fig. 6 suggests this level of CPU temperature would deliver the minimum T_{gen} of 70 °C to energize the Yazaki WFC-SC10 absorption chiller. This ΔT loss of 20 °C is a worst case scenario and we anticipate our actual temperature loss to be much lower, around half of that value. Although, higher CPU temperatures can be achieved by reducing or eliminating the fans and by tasking the CPU, we plan to run the CPUs within their standard operating range of 70–90 °C in order to ensure that no reliability issues should occur. If necessary, additional external heat supplementation from a sustainable energy source such as solar thermal can contribute to increasing the temperature of the required input heat [16].

4.2. CPU power

Since nearly 100% of CPU electrical power is dissipated as heat and to determine if we could meet the chiller's heat input requirements, we investigated the thermal design point (TDP) and maximum theoretical power possible. According to Intel specifications, each Intel Xeon Nocona CPU has a TDP of 103 W [17].

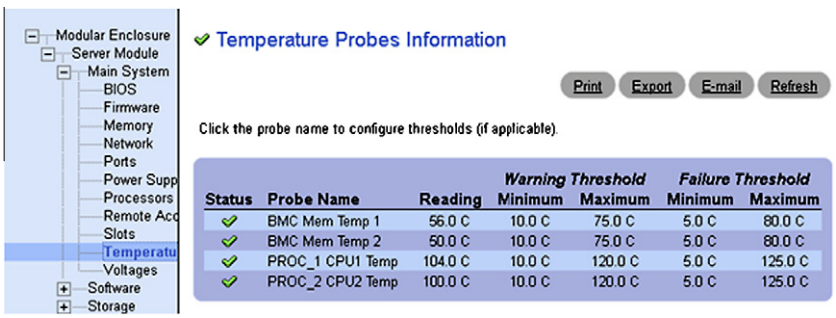

Fig. 4. OMSA indicates CPUs capable of up to 125.0 °C.

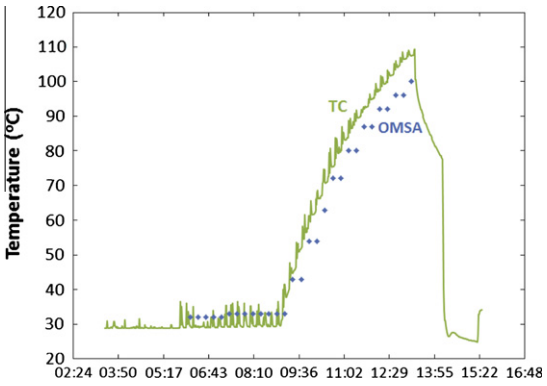

Fig. 5. CPUs can operate within required chiller temperature range and beyond.

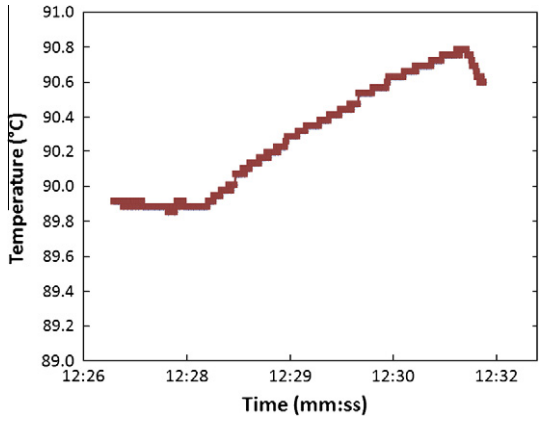

Fig. 6. CPU stabilized to an average temperature of 90.2 °C over a 5-min test run.

While TDP is used as a target point for design solutions, both Intel and AMD agree that TDP is not the maximum power the CPU may draw or generate as heat but rather the *average* maximum power that it draws when running typical user applications [18]. There may be periods of time when the CPU dissipates more power than designed, such as during strenuous demands by engineering or scientific applications, or more commonly, video games. Under such loads, the CPU temperature will rise closer to the maximum and the CPU power will surpass that of the TDP. TDP is primarily used as a guideline for manufacturers of thermal solutions (heatsinks/fans, etc.) [19]. This guideline ensures the computer will be able to handle essentially all applications without exceeding its thermal envelope, or requiring a cooling system for the maximum theoretical power (which would cost more but in favor of extra headroom for processing power) [18].

Maximum power dissipation is always higher than TDP [20]. Maximum power dissipation by the CPU occurs at the maximum core voltage, maximum temperature and maximum signal loading conditions that the system can handle. CPU maximum power dissipation is usually 20–30% higher than the rated TDP value [19]. With the TDP rating of each CPU specified by Intel to be 103 W, this would equate to a maximum possible power dissipation of between 123 W and 144 W per CPU. Referring to Fig. 3 and assuming no temperature difference or heat loss, the maximum chiller capacity "ideally" could be reached by using 74.8 kW/0.144 kW = 520 CPUs (rounding up). However, our approach is more conservative as it takes into account an estimated temperature fall of 20 °C and a system heat loss of 15% [10].

According to the Dell Datacenter Capacity Planner v3.04 configuration tool, the system heat is 294.0 W for each PowerEdge 1855 Blade [21]. With two CPUs per blade, we can assume an average maximum power of 206 W. This equates to exactly 70% of the system board being dominated by the heat coming from the CPUs. This 0.7 value is hence denoted the highest heat fraction or HHF of the system board.

In contrast, the capture fraction (CF) refers to the maximum *heat* that we estimate can be recovered. With a liquid cooling system in place, we estimate that we can recover 85% of the HHF from each CPUs or 59% of each system board. This estimate is based on a similar cold plate system currently in use by IBM [10]. Similar to IBM, our heat extraction system will house a heat transfer fluid of water but the extracted heat will be used for cooling rather than for heating. Since the system heat is 294.0 W for each PowerEdge 1855 blade, the bottom line is that we anticipate being able to capture $0.59 * 294\ \text{W} = 174\text{W}_{\text{th}}$ per blade.

Taking the most conservative value in the chiller standard performance graph of Fig. 2, the chiller must be energized with a minimum of 14.1 kW_{th} at 70 °C. This conservative stance would call for 14.1/0.174 = 81 server blades and we have more than twice this amount in our data center. For a higher efficiency scenario of $\text{COP}_{\text{C}} = 0.86$ at $T_{\text{gen}} = 80$ °C in Fig. 2, we would need 144 server blades to deliver 25.0 kW_{th} worth of cooling. Finally, at the design point of 88 °C, the heat medium requirement is 50.2 kW_{th} or 289 server blades. Since we have access to more than 300 server blades for our data center and since the CPUs can be run near 120 °C (Fig. 4), we can theoretically accommodate all three scenarios.

5. Overall system design for system analysis

5.1. System level design and description

Fig. 7 illustrates the heat flow from the data center to various cooling components.

High-temperature, high-quality heat, denoted as $\dot{Q}_{\text{HiT,S}}$ in Fig. 7, is dissipated by approximately 70% of the components on a server blade inside a standard 42U (200 cm tall) rack. To clarify, each cabinet-style rack can house up to six chassis and each chassis box can hold up to ten server blades [21]. Each rack has a footprint of 0.6 m^2 [21]. To leave plenty of room for our heat extraction equipment as well as switchgear and any PDUs, there can be four chassis per rack equivalent to 80 CPUs per rack. We plan to have two rows of two racks or 640 CPUs.

As illustrated by lines 1 and 11 in Fig. 7, the captured high quality heat $\dot{Q}_{\text{HiT,S}}$ from each server blade can be "piped" directly to the Li–Br absorption chiller or, along with the auxiliary heat source, can be sent to a phase change thermal storage unit for later use [22]. Due to its ability to provide a high energy storage density at a nearly constant temperature, this latent heat storage technique holds the key to the provision of energy when the intermittent sources of energy are not available [23].

A large number of organic and inorganic substances are potential PCMs, based on their melting points and heats of transition. Paraffins, some plastics and acids such as stearic acid are the most practical compounds for thermal storage systems because they are relatively inexpensive, tend to have broad melting ranges, and their properties do not change over time [24]. The main disadvantages of organic PCMs include low thermal conductivity, low durability and high volume change during melting [24]. An ideal PCM may be a combination of organic and inorganic substances [23]. Another idea is to use high thermal conductivity materials within the PCM to increase its apparent thermal conductivity. For example, using finned tubes in which the PCM is placed between the fins can significantly improve heat transfer rates for liquid-based systems [24].

Once enough heat has been supplied to the chiller, it can function to reduce the heat load on the CRAC. As shown by line 2 in

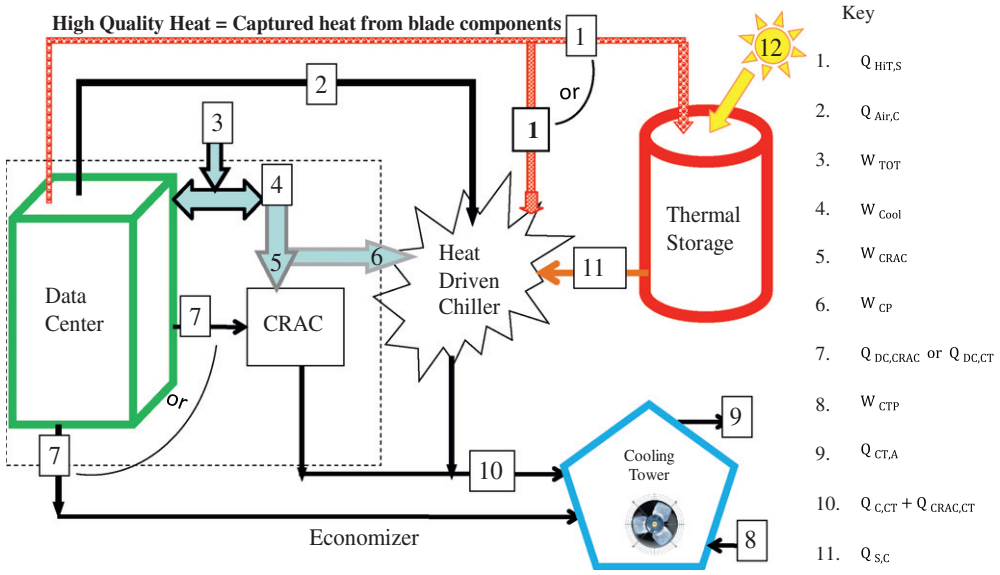

Fig. 7. Overall system level diagram showing the work and heat flow paths.

Fig. 7, a portion of the heated air in the data center, denoted as $\dot{Q}_{Air,C}$, can be removed by the absorption chiller, which lessens the workload on the CRAC. This ambient heat inside the data center is lower temperature, lower quality heat relative to the higher temperature, higher quality heat of the computing components inside each server blade. This lower temperature air can be cooled by the absorption chiller. Any remaining heat, as illustrated by line 7, is removed by the CRAC or by the Economizer and rejected to the cooling tower.

One of the most challenging aspects of the proposed system is transferring the heat from the blade components to the generator of the Li–Br chiller. Our approach is for the CPUs to be cold-plated, i.e., in thermal contact with water-cooled heat sinks. The cold plates contain water as the heat-exchange fluid which can then transfer a fraction of high-temperature blade heat to thermal storage to run the chiller. A capture fraction (CF) of 0.85 of this HHF (0.70) is expected to be recovered or about 174 W_{th} per server blade.

The advantages of using water cooling over air cooling include water's higher specific heat capacity, density, and thermal conductivity [25]. For example, the thermal conductivity of liquid water is about 25 times that of air. Even more significant, water is approximately 800 times denser than air, and this greater density allows water to transmit heat over greater distances with much less volumetric flow and reduced temperature difference than that of air. For cooling CPU cores, water's primary advantage is its increased ability to transport heat. Furthermore, since the Yazaki absorption system also runs on water, more efficient water-to-water heat exchangers are used. Also, a smaller flowpath means less area to insulate and less pump power work as compared to using air transport and fans.

5.2. PUE: power usage effectiveness metric

Besides creatively re-using data center waste heat and reducing the power grid consumption of the CRAC, an additional benefit of applying the system level design of Fig. 7 is an improvement in data center efficiency. The most widely accepted benchmarking standard to gauge data center efficiency is power usage effectiveness (PUE). Introduced by The Green Grid in 2006, PUE is useful for understanding the total amount of power consumed by IT equipment relative to total facility power [26].

As shown in Eq. (1), PUE is the ratio of power into the data center, measured at the utility meter, to the power required to run the IT equipment for computing [26]:

$$\text{PUE} = \frac{\text{Electrical utility power}}{\text{IT equipment electric power}} = \frac{\dot{W}_{TOT}}{\dot{W}_{IT}} \quad (1)$$

For example, a PUE of 2.0 means that for every 1 W of electric input to the IT equipment, an additional 1 W is consumed to cool and distribute power to the IT equipment. Thus, the intent of the PUE metric is to characterize the fraction of the total data center power consumption devoted to IT equipment. It compares how much power is devoted to driving the actual computing IT components versus the ancillary support elements such as cooling and PDUs (power distribution units). With PUE, the focus is on maximizing the power devoted to the equipment running applications and minimizing the power consumed by support functions like cooling and power distribution [26]. Thus, the lower the PUE, the more efficient is the data center.

As PUE is defined with respect to the "power at the utility meter," it leaves room to re-use on-site energy generation from waste heat and to use external alternative energy sources such as solar thermal energy. These previously unaccounted-for energy sources make it possible to lower the PUE to a number below one. That is, the facility power use can actually be a negative number. This is an intriguing prospect as the data center would be generating energy, or put another way, producing exportable power. Therefore, as is suggested by the following steady-state analysis, a very efficient PUE is possible by re-using data center waste heat and a PUE < 1 is possible with an external supplemental source of alternative energy.

5.3. Steady-state analysis

The total facility power or total data center electric power consumption ($\dot{W}_{TOT}$) is comprised of three major categories: IT computer rack power ($\dot{W}_{IT}$), power delivery loss from the power distribution units (PDUs) to the racks ($\dot{W}_{Loss}$), and the electric power consumed to operate the data center cooling system ($\dot{W}_{Cool}$). Since we are choosing a "lights out data center," electrical power to the lights ($\dot{W}_{Lights}$) will be off during normal operation and can be neglected compared to the other three terms [27]. Thus, the numerator in Eq. (1) can be expanded as

$$\dot{W}_{TOT} = \dot{W}_{IT} + \dot{W}_{Loss} + \dot{W}_{Cool} \quad (2)$$

where

$$\dot{W}_{Cool} = \dot{W}_{CRAC} + \dot{W}_{AC_FANS} + \dot{W}_{CP} \quad (3)$$

Also, the electric power driving the cooling system ($\dot{W}_{Cool}$) is itself composed of three main terms: the power to run the vapor compression computer room air conditioner ($\dot{W}_{CRAC}$), whose electrical power requirements are determined by the demands of the data center at any given moment; the power to run the CRAC supply and condenser fans ($\dot{W}_{AC,FANS}$) and any electrical power needed to run the heat-driven chiller ($\dot{W}_{CP}$). The CRAC of choice is an 88 kW_{th} (25-ton) unit which utilizes a supply fan at 5 hp and two condenser fans at 1 hp each for a total $\dot{W}_{AC_FANS}$ of 7 hp or 5.22 kW_e [28]. The chiller power $\dot{W}_{CP}$ is broken down into the solution pump and condenser fans. Although the electrical power consumption by the Li–Br pump is negligible compared to the power consumed by the vapor compressor of the CRAC, the condenser fans cannot be neglected [29]. Compared to a compression system, the cooling tower will be between 1.5 and 2.5 times larger on an absorption system, which translates into more condenser fan energy at an increase of about 1.25 hp for the fan motor. For example, a 10-ton cooling tower for a vapor compression system may use a 3/4 hp fan motor, but a 10-ton WFC-S10 Yazaki LiBr absorber requires a 25-ton cooling tower. As it turns out, a 25-ton cooling tower uses 1-hp fan motor or 0.7456 kW_e. For the absorption system, this increase of electrical consumption equates to about 25% above what the compression system uses.

From Eqs. (1)–(3), the PUE can now be expressed as

$$PUE = \frac{\dot{W}_{TOT}}{\dot{W}_{IT}} = \frac{\dot{W}_{IT} + \dot{W}_{CRAC} + \dot{W}_{AC_FANS} + \dot{W}_{CP} + \dot{W}_{Loss}}{\dot{W}_{IT}} \quad (4)$$

Then, to relate the electric power supplying the CRAC compressor, $\dot{W}_{CRAC}$, to the heat load on the CRAC, $\dot{Q}_{CRAC}$, we can make use of the standard vapor compression COP (coefficient of performance) equation:

$$\dot{W}_{CRAC} = \frac{\dot{Q}_{CRAC}}{COP_{CRAC}} \quad (5)$$

The $\dot{Q}_{TOT}$ total heat load on the CRAC can be reduced by any supplemental chiller cooling capacity, $\dot{Q}_C$,

$$\dot{Q}_{CRAC} = (\dot{Q}_{TOT} - \dot{Q}_C) \quad (6)$$

Substituting Eq. (6) into Eq. (5) and then into Eq. (4), Eq. (4) now becomes

$$PUE = \frac{\dot{W}_{IT} + \frac{(\dot{Q}_{TOT} - \dot{Q}_C)}{COP_{CRAC}} + \dot{W}_{AC_FANS} + \dot{W}_{CP} + \dot{W}_{Loss}}{\dot{W}_{IT}} \quad (7)$$

where $\dot{Q}_{TOT}$ is the total heat flow from the data center to be removed by the cooling equipment. A portion of this total heat flow is removed right away by the heat extraction equipment on each server blade, and is denoted as $\dot{Q}_{HiT,S}$:

$$\dot{Q}_{TOT} = \dot{Q}_{IT} - \dot{Q}_{HiT,S} + \dot{Q}_{Loss} \quad (8)$$

However, since the heat output of an electronic device in steady operation is approximately equal to the power input to the device, we can express Eq. (8) in terms of power as

$$\dot{Q}_{TOT} = \dot{W}_{IT} - [\dot{W}_{IT} \cdot HHF \cdot CF] + \dot{W}_{Loss} \quad (9)$$

in which $\dot{Q}_{HiT,S} = [\dot{W}_{IT} \cdot HHF \cdot CF]$ and is the portion of rack heat removed from the racks by the heat recovery scheme and sent to the chiller.

As previously stated, HHF = 70% of each server blade provides the highest quality heat. An estimated CF = 85% of HHF or 59% can be captured and transferred from the blades to thermal storage to drive the LiBr absorption chiller. With a high enough CF, a significant portion of this high-temperature heat can be removed from the original heat load on the CRAC and utilized to drive the chiller.

To continue with the rest of Eq. (6), the cooling provided by the absorption chiller, $\dot{Q}_C$, is also driven by the heat recovered from the server blades, $\dot{Q}_{HiT,S}$, along with any supplemental external heating, $\dot{Q}_{EXT,S}$:

$$\dot{Q}_C = COP_C(\dot{Q}_{HiT,S} + \dot{Q}_{EXT,S}), \quad (10)$$

And, since $\dot{Q}_{HiT,S} = [\dot{W}_{IT} \cdot HHF \cdot CF]$,

$$\dot{Q}_C = COP_C\left(\dot{W}_{IT} \cdot HHF \cdot CF + \dot{Q}_{EXT,S}\right) \quad (11)$$

Substituting Eqs. (9) and (11) into Eq. (7) and simplifying, yields the final general expression for a PUE taking into account parasitic power, recycled thermal energy and alternative energy:

$$PUE = 1 + \frac{\dot{W}_{AC_FANS} + \dot{W}_{CP} + \dot{W}_{Loss}}{\dot{W}_{IT}} + \frac{1}{COP_{CRAC}}\left[1 + \frac{\dot{W}_{Loss}}{\dot{W}_{IT}} - HHF \cdot CF - COP_C\left(HHF \cdot CF + \frac{\dot{Q}_{EXT,S}}{\dot{W}_{IT}}\right)\right] \quad (12)$$

Note that any useful heating power extracted from the data center and supplemented by an external source (e.g., solar) is subtracted from the total electric utility power. Also, note that, depending on the magnitude of the term

$$\frac{COP_C}{COP_{CRAC}}\left(HHF \cdot CF + \frac{\dot{Q}_{ext}}{\dot{W}_{IT}}\right)$$

PUE can become less than one, and the data center can actually export cooling. Here, $\dot{Q}_{ext}$ is essentially the same as $\dot{Q}_{EXT,S}$ but before it reaches any thermal storage. This cooling (removal of heat) is divided by COP_{CRAC} so that it represents an equivalent exported electric power. In other words, external heating can generate excess cooling that can be "exported," i.e., used to cool adjacent rooms or facilities. The equivalent exported power can be expressed generally as

$$W_{Exported} = \frac{Q_{Excess}}{COP_{CRAC}}, \quad (13)$$

where

$$Q_{Excess} = COP_C\left(HHF \cdot CF + \frac{\dot{Q}_{ext}}{\dot{W}_{IT}}\right). \quad (14)$$

6. Results and discussion

To specify working values for each term in the final Eq. (12), we can apply the values in Table 2 from our 15 m^2 experimental data center running at full capacity. At full load, we have 320 blades that consume 294 W_e per blade and deliver a HHF of 70% of that value or 205 W_{th}, and we expect to capture 85% of that dissipated heat or 59% from each blade (174 W_{th}). This captured heat is denoted in Table 2 as $\dot{Q}_{HiT,S}$ and its value for our case of 320 blades is 52.2 kW_{th}. We assume a 10% power delivery loss of the total power delivered to the data center and its components of $\dot{W}_{IT} + \dot{W}_{CRAC} + \dot{W}_{AC_FANS} + \dot{W}_{CP}$. The parasitics of $\dot{W}_{AC_FANS} + \dot{W}_{CP}$ are calculated assuming a maximum running load of 5.22 kW_e.

For the COP_{CRAC}, we selected a standard value from a thermodynamics text for a typical vapor compression unit [30]. Then, for COP_C, we chose three scenarios from Fig. 2: at design point, at minimum chiller capacity, and at maximum chiller COP. To isolate how

Table 2
Values and calculations for our 15 m^2 data center model running at full load.

Power/heat	Data center total	Units
$\dot{W}_{IT}$(294 W * 320 blades)	94.08	kW_e
$\dot{Q}_{HiT,S}$ (capture of 59%)	55.51	kW_{th}
$\dot{W}_{Loss}$ (power loss of 10%) [25]	9.66	kW_e
$\dot{W}_{AC_FANS} + \dot{W}_{CP}$	5.22	kW_e
Coefficients of performance		
COP_{CRAC}	4	[30]
COP_C (T_{gen} = 88 °C)	0.7	Fig. 2
COP_C (T_{gen} = 80 °C)	0.86	Fig. 2
COP_C (T_{gen} = 70 °C)	0.63	Fig. 2
PUE results for CF of 85%		
PUE (COP_C = 0.86)	1.16	Fig. 8
PUE (COP_C = 0.7)	1.18	Fig. 8
PUE (COP_C = 0.63)	1.19	Fig. 8
PUE ($\dot{Q}_{ext} = 1.05\dot{Q}_{IT}$)	0.997	Fig. 8
PUE ($\dot{Q}_{ext} = 0.75\dot{Q}_{IT}$)	0.996	Fig. 8

Table 3
Industry benchmarked data center efficiencies Error! Reference source not found.

PUE	Level of efficiency
3.0	Very inefficient
2.5	Inefficient
2.0	Average
1.5	Efficient
1.2	Very efficient

PUE responds to chiller performance under standard conditions, we eliminated any external supplemental heat until the last two rows of Table 2 ($\dot{Q}_{ext}$) Note that even with no external supplementation, the PUE is near the "very efficient" value in Table 3.

As shown in the last two rows of Table 2, to realize a PUE value of less than one at a CF of 85%, an amount of supplemental external heat equating to 1.05 times that of the IT heat is necessary at the chiller design point. If the maximum COP_C of 0.86 is selected instead, this 1.05 value can be lowered to 0.75. However, again the tradeoff would be chiller capacity.

With only rack heat driving the chiller, there is a reduction of heat load on the CRAC. Even with no external thermal contribution and driven solely on a portion of data center IT equipment heat, there is a significant reduction of PUE. Fig. 8 represents the range of possible PUE values for a high-temperature heat capture fraction, CF, ranging from 0% to 100%. CF = 0% corresponds to the "business as usual" approach and yields PUE ~ 1.43. As expected, PUE decreases with increasing CF, and with increasing COP_C, leading to a PUE as low as 1.16 at our maximum expected capture fraction of 85%. According to the benchmarks shown in Table 3, these PUE values suggest a "very efficient" data center is possible even with zero solar contribution. By employing an external heat source such as solar thermal, we can reach beyond the limitations of Table 3. With an amount of external heat equating to 1.05 times that of the IT heat, the PUE ranges from 1.25 at zero CF down to 0.95 at 100% CF.

Of course, the effectiveness of this approach depends on the CF and the factors that affect it. One possible way to achieve an increase in CF is to insulate the hotter components from the data center (e.g., using an insulating gel such as Aerogel®). Another possible way to increase the CF would be to cold plate more blade components. Also, the size and surface of the cold plate areas may impact the CF percentage, as could the thermal conductivity of the materials of construction of each cold plate.

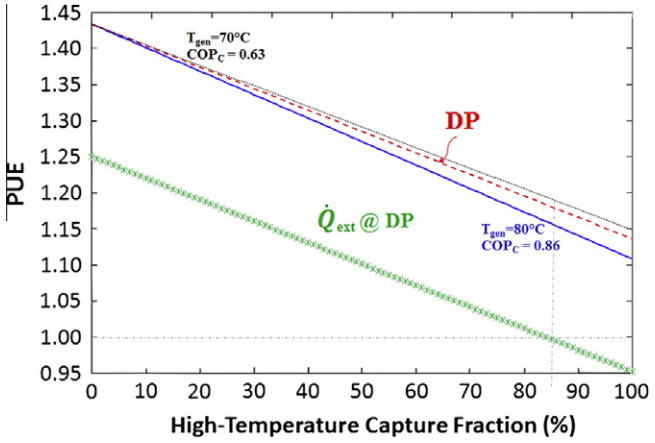

Fig. 8. Power usage effectiveness (PUE) versus high-temperature capture fraction (CF). $\dot{Q}_{ext}$ is the external heat supplied to the generator at design point conditions.

Finally, to capture as much of the dissipated heat as possible, it is important to lower as much as possible the ΔT of the heat capturing scheme. The ΔT_{CH} between the metal case surrounding the CPU die, T_C, and the cold plate, T_H, is especially variable and the only interface under our control. A lower ΔT_{CH} can be achieved by reducing the thermal resistance between the device metal case and cold plate (heat sink) interface [15]. One way to reduce these values is to increase thermal contact (or decrease thermal contact resistance) between the two interfaces. Although the top surface of the device case and the bottom surface of the heat sink appear flat, there are many surface irregularities with air gaps that can prevent intimate thermal contact and inhibit heat transport. For a near perfect thermal contact, the two surfaces would need to be match ground and polished to better than mirror quality [31]. Although ideal, this 100% thermal contact would be very difficult to achieve. Instead, a thermal interface material (TIM) can be applied to help reduce both the contact and thermal resistances between the device case and the heat sink.

The heat removal capacity of a TIM is also determined by the TIM's ability to create an intimate contact with the relevant surface. A highly electrically insulating but thermal conductive TIM is essential to good heat capture performance. Fortunately, there are a number of TIM alternatives for electrically insulating the transistor from the heat sink while still allowing heat transfer. Wet-dispensed TIMs include adhesives, encapsulants, and non-curing thermal compounds. Adhesives can create intimate surface contact, resulting in low contact resistance [32]. Encapsulants can adopt any thickness and offer higher mechanical strength than other wet dispensed TIMs [32]. Thermal compounds such as Arctic-Silver 5 provide relatively high thermal conductivity and low thermal resistance. Other TIM categories include pads, gap fillers, and phase change materials. An advantage of pads is that they can be applied without dispensing and therefore can be easily reworked [32]. For example, Bergquist gap filler pads readily accommodate irregular surfaces with minimal pressure and offer a thermal conductivity as high as 5.0 $W\,m^{-1}\,K^{-1}$, which is comparable to Arctic Silver 5 thermal paste of 8.88 $W\,m^{-1}\,K^{-1}$. Finally, for a permanent solution, phase change materials have low thermal resistance, require low mounting force, and can accommodate irregular surfaces as well [32].

7. Conclusions

Data center electric power consumption is an acknowledged problem that is likely to get worse in the future. The potential ex-

ists to utilize the waste heat generated by data centers to drive absorption chillers, which would relieve the cooling loads on the conventional computer room air conditioner (CRAC) and, ultimately, reduce the grid power consumption of the data center. By effectively capturing at least 85% of the heat dissipated from the CPUs on each blade server and efficiently transporting that thermal energy to drive a heat-activated lithium bromide absorption chiller, it is possible for our design to achieve a very efficient power usage effectiveness (PUE) ratio. With external heat supplementation such as that offered by solar thermal, the value of PUE can fall to that of less than one.

Although conventional thinking defies a PUE < 1, the Green Grid has already begun receiving reports of some data centers actually achieving this possibility. For instance, the Finnish public energy company, Helsingin Energia, with their 2 MW eco-efficient datacenter underneath the Uspenski Cathedral in Helsinki, uses geothermal cooling with the heat being re-used by local businesses [33]. According to The Times of London, "This centre's power usage effectiveness – the central measurement of data centre efficiency – will be an unprecedented figure of less than one." [33].

All this points, ultimately, for the need for more metrics to take into account the new eco-efficient data center designs and reward waste heat re-use and the use of renewable energy [33]. Fortunately, the Green Grid, the National Renewable Energy Laboratories and Lawrence Berkeley National Laboratory are responding by developing a new metric, tentatively termed ERE or Energy Reuse Effectiveness [34]. The aim of this new metric is to allow for recognizing and quantifying energy re-use efforts whose benefits are felt outside of the data center's boundaries.

Acknowledgment

The authors gratefully acknowledge the support of the National Science Foundation through award 0855277.

References

[1] Patterson MK. The effect of data center temperature on energy efficiency. In: Thermal and thermomechanical phenomena in electronic systems, 2008 ITHERM 2008 11th intersociety conference on Orlando, FL; 2008. p. 1167–74.
[2] Mitchell-Jackson JD. Energy needs in an internet economy: a closer look at data centers. Energy and resources group. Berkeley: University of California; 2001.
[3] Koomey JG. Worldwide electricity used in data centers. Environ Res Lett. 2008;3:1–8.
[4] Report to congress on server and data center energy efficiency public law 109-431. US Environmental Protection Agency; 2007. pp. 1–133.
[5] Brill KG. Moore's law economic meltdown. Forbes.com; 2008.
[6] Brill KG. The invisible crisis in the data center: the economic meltdown of Moore's Law. The Uptime Institute White Paper. The Uptime Institute; 2007.
[7] Vanderbilt T. Data center overload. The New York Times Magazine; 2009.
[8] Brill KG. Data center energy efficiency and productivity. The uptime institute symposium 2007: the invisible crisis in the data center: how IT performance is driving the economic meltdown of Moore's Law, Orlando, FL; 2007.
[9] Rabah G. Investigation of the potential of application of single effect and multiple effect absorption cooling systems. Energy Convers Manage. 2010;51:1629–36.
[10] Brunschwiler T, Smith B, Ruetsche E, Michel B. Toward zero-emission data centers through direct reuse of thermal energy. IBM J Res Dev 2009;53:11:1–3.
[11] Gupta Y, Metchop L, Frantzis A, Phelan PE. Comparative analysis of thermally activated, environmentally friendly cooling systems. Energy Convers Manage 2008;49:1091–7.
[12] Water Fired Chiller/Chiller-Heater. In: SB-WFCS-1009.pdf, Acrobat 913, 2nd ed. Yazaki Energy Systems, Incorporated; 2009. <www.yazakienergy.com/docs/SB-WFCS-1009.pdf>.
[13] Joudi KA, Abdul-Ghafour QJ. Development of design charts for solar cooling systems. Part I: computer simulation for a solar cooling system and development of solar cooling design charts. Energy Convers Manage. 2003;44:313–39.
[14] Ali AHH, Noeres P, Pollerberg C. Performance assessment of an integrated free cooling and solar powered single-effect lithium bromide–water absorption chiller. Sol Energy 2008;82:1021–30.
[15] Thermal resistance. Wikipedia; 2011.
[16] Mazloumi M, Naghashzadegan M, Javaherdeh K. Simulation of solar lithium bromide–water absorption cooling system with parabolic trough collector. Energy Convers Manage. 2008;49:2820–32.
[17] 64-bit Intel® Xeon® Processor 3.60 GHz, 1M Cache, 800 MHz FSB. Intel Corporation. <http://ark.intel.com/products/27088/64-bit-Intel-Xeon-Processor-3_60-GHz-1M-Cache-800-MHz-FSB>.
[18] CPU power dissipation. Wikipedia; 2011.
[19] Thermal Design Power (TDP). CPU World; 2011.
[20] Minimum/Maximum power dissipation. CPU World; 2011.
[21] Dell PowerEdge Rack Systems. Dell, Inc.; 2003. <http://www.verinet.dk/dell_rack_system.pdf>.
[22] Aghareed M T. A simulation model for a phase-change energy storage system: experimental and verification. Energy Convers Manage 1993;34:243–50.
[23] Aghareed M T. Organic–inorganic mixtures for solar energy storage systems. Energy Convers Manage 1995;36:969–74.
[24] Farid MM, Khudhair AM, Razack SAK, Al-Hallaj S. A review on phase change energy storage: materials and applications. Energy Convers Manage 2004;45:1597–615.
[25] Water cooling. Wikipedia; 2011.
[26] Best practices for datacom facility energy efficiency, 2nd ed. ASHRAE, Atlanta, GA; 2009.
[27] High density data centers: case studies and best practices. ASHRAE, Atlanta, GA; 2008.
[28] Lu L, Cai W, Soh YC, Xie L, Li S. HVAC system optimization––condenser water loop. Energy Convers Manage 2004;45:613–30.
[29] Yu FW, Chan KT. 06/02239 Modelling of the coefficient of performance of an air-cooled screw chiller with variable speed condenser fans. Build Environ 2006;41:407–17.
[30] Cengel YA, Boles M. Thermodynamics: an engineering approach. 6th ed. Boston, MA: McGraw Hill; 2006.
[31] Elliot R. The design of heatsinks. Elliot sound products; 2007.
[32] Harris J. Empfasis – tech tips: thermal interfaces & materials. In: Frederickson MD, editor. National Electronics Manufacturing Center of Excellence, EMPF; 2006.
[33] Lawrence A. PUEs of less than 1. N plus one: all about available, scalable and sustainable IT. Uptime Institute, WorldPress.com; 2010.
[34] Tschudi B, Vangeet O, Cooley J, Azevedo D. In: Patterson M, editor. ERE: a metric for measuring the benefit of reuse energy from a data center; 2010.