A SUSTAINABLE DATA CENTER WITH HEAT-ACTIVATED COOLING

Anna Haywood*, Jon Sherbeck*, Patrick Phelan*†, Georgios Varsamopoulos§, Sandeep K. S. Gupta§
§School of Computing, Informatics, and Decision Systems Engineering Computer Science and Engineering
*School of Mechanical, Aerospace, Chemical & Materials Engineering
†National Center of Excellence on SMART Innovations
Arizona State University
P.O. Box 9309
Tempe, Arizona, U.S.A., 85287-9309
Phone: (480) 965-1730
Fax: (123) 987-6543
Email: bluetoolteam@impact.asu.edu

ABSTRACT

Technological and economic trends in data centers push toward facilities operated at higher ambient temperatures and at higher power densities to meet ever-increasing computational demands. Conventionally, data centers are cooled with vapor-compressor equipment which requires extra power to be driven. This paper proposes an alternative and sustainable data center cooling architecture that is heat driven. The thermal source is the heat produced by the data center room's equipment. A major challenge is providing both enough cooling to the data center and enough exergy to drive the cooling process, regardless of the thermal output of the data center equipment. This challenge is addressed by the use of organic heat storage and a sustainably powered (i.e. solar-powered) heat source, leading potentially to a PUE (Power Usage Effectiveness) value of less than one.

KEY WORDS: PUE, data center, CRAC, absorption chiller, blade server, chassis, rack, green computing, sustainability

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
$\dot{Q}_{DC,CRAC}$	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
$\dot{Q}_{\mathbf{TOT,DC}}$	Total heat flowing from data center, W
$\dot{W}_{\mathbf{TOT,DC}}$	Total electric power into data center, W
$\dot{W}_{\mathbf{IT}}$	Electric power into racks holding computer equipment, W
$\dot{W}_{\mathbf{Loss}}$	Electric power loss from PDU to racks, W
$\dot{W}_{\mathbf{Lights}}$	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
$\dot{W}_{CP}$	Electric power into absorption pump, W
$\dot{W}_{CTP}$	Electric power into cooling tower pump, W

INTRODUCTION

Data centers are running into limits related to power, cooling, and space. Today's average data center is 12 to 15 years old, reaching end of life, and unable to meet today's computing power demands [1]. 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 an important topic of global discussion among end users, policy makers, technology providers, facility architects, and utility companies. 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 [1]. By 2010, data center consumption from the U.S. power grid is expected to jump to 4% [2]. Extrapolating forward suggests U.S. grid power consumption

978-1-4244-5343-6/10/$26.00 ©2010 IEEE

to be near 10% by 2011. According to *The Economic Meltdown of Moore's Law*, the three-year operational and capital expenditures of powering and cooling high-end servers can be 1.5 times the cost of purchasing server hardware. Projections out to 2012 show this multiplier at three times the cost of the hardware under even the best conditions and up to 22 under worst case assumptions [3].

Demands for computing power will keep rising as computing technology continues to progress and evolve. Creating an energy-efficient data center is paramount to curbing runaway power consumption and accommodating greater data center capacity. Since conventional cooling power alone can be as much as four times [2] that of computing power, 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 is dissipated as heat ($\dot{Q}_{IT}$) [4] [5], the central idea is to transfer directly that available thermal energy from the highest power components on a blade to drive a thermally activated cooling process.

A typical data center is based on air-cooled technologies. It is laid out in equipment rows resting on a *raised floor*. The equipment in the rows is placed so that air inlets of the equipment in one row face the air inlets in another row (and likewise for the outlets). The raised floor in the aisles of air inlets features perforations which allow cool air to enter the room; perforations or other contraptions above the air outlet aisles gather the hot air, which is then passed to the computer room air conditioner (CRAC) (also known as HVAC). The above physical configuration is referred to as "***hot aisle / cold aisle***" layout. The CRAC is conventionally of vapor-compression technology, based either on the outdoor air temperatures or on water supply, many times chilled, to perform the cooling cycle.

There are three major categories of research on saving energy in data centers:

- ***Software-based approaches***. This category involves saving energy by means of a software-controlled operation. Research has shown that by changing the runtime behavior of software, one can induce a change in the performance and the energy consumption. A viable option for this thermal, power and energy control is via dynamic voltage and frequency scaling (*DVFS*), [6] [7] [8] and energy-saving scheduling policies [9]. Research has shown that a combination of separate energy-saving techniques can yield synergistic energy-saving benefits [10].
- ***Hardware-based approaches***. This category involves improving the energy efficiency of any energy-consuming hardware used in the data center, that being mostly IT equipment.
- ***Data center design approaches***. [10] [12]. This category includes efforts such as proper physical configuration, location, and construction.

The proposed data center cooling configuration belongs in the latter category. Most notable and innovative efforts include: the "zero-emission" data center by IBM [13], chillerless data centers (Google's in Belgium and Microsoft's in Ireland), and the warehouse-sized data centers in Oregon [14].

The proposed configuration is more closely related to IBM's zero-emission data center, as both approaches try to reuse the data center's produced heat for some useful purpose instead of discarding it into the environment. The difference of our approach is in using the heat to drive the cooling of the data center itself, an effort into creating a sustainable data center.

Electrically-driven CRACs (computer room air conditioners) are the standard for cooling data centers, and they are responsible for much of the grid power consumption. Therefore, we propose a thermally-driven cooling system to lessen the load on the CRAC, thereby reducing workload on the grid. An additional benefit would be an improvement in power usage effectiveness (PUE), the widely accepted benchmarking standard which gauges energy effectiveness within a data center by characterizing the fraction of the total data center power used for IT work [15], [16].

As shown below, PUE is the ratio of power in to the data center measured at the utility meter to the conditioned power out to run the IT equipment for computing [17]. For example, a PUE of 1.6 means that for every 1.6W of electric input at the utility meter, 1W of electric output is delivered to the connected IT load:

$$PUE = \frac{Electric\ Utility\ Power}{IT\ Equipment\ Electric\ Power} = \frac{\dot{W}_{TOT}}{\dot{W}_{IT}} \qquad [17]\ (1)$$

The intent of the PUE metric is to highlight 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 lighting. 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 use alternative energy sources or even exporting power[1]. Therefore, a PUE <1 is difficult but possible with on-site generation from waste heat [18], or with external heating provided, for example, via solar energy.

PRIMARY OBJECTIVE & APPROACH

The overall objective of this project is to reduce the grid power consumption of the cooling system for data centers. A heat-driven, lithium-bromide (Li-Br) absorption chiller is proposed to reduce the cooling load on a typical vapor compression CRAC. The heat to drive the chiller will be originated from the data center itself, more specifically, from the blade components inside each chassis within each rack.

[1] http://perspectives.mvdirona.com/2009/06/15/PUEAndTotalPowerUsageEfficiencyTPUE.aspx

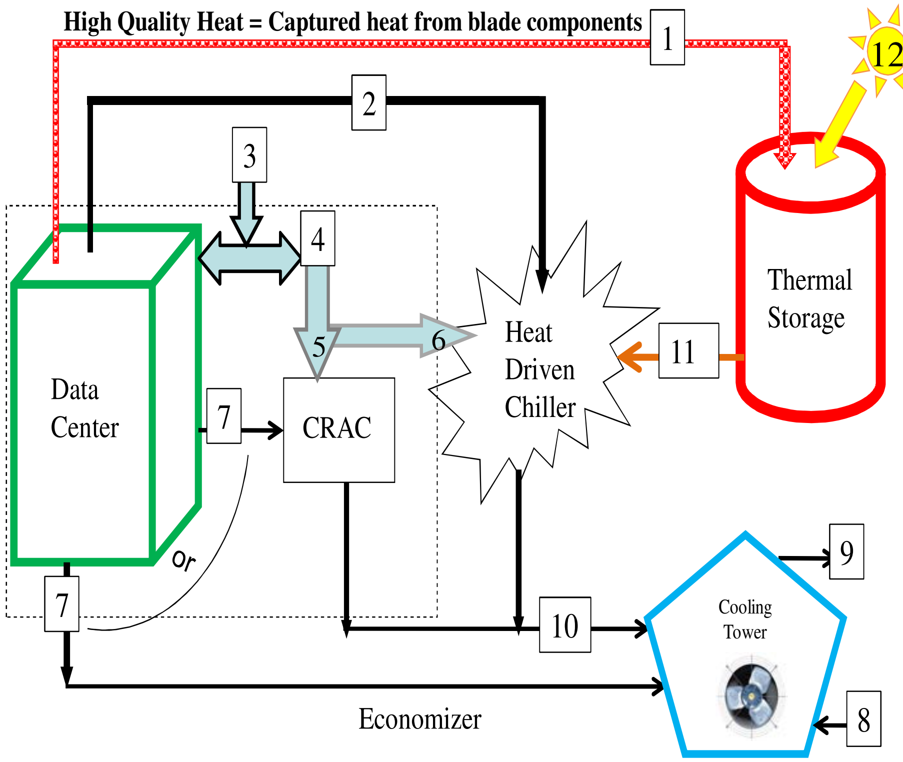

Key

1. $\dot{Q}_{HiT,S}$ (High quality heat flow from server blades to thermal storage)
2. $\dot{Q}_{Air,C}$ (Heated ambient air in data center removed by chiller)
3. $\dot{W}_{TOT} = \dot{W}_{IT} + \dot{W}_{Cool} + \dot{W}_{Loss} + \dot{W}_{Lights}$ (Total electric power into data center)
4. $\dot{W}_{Cool} = \dot{W}_{CRAC} + \dot{W}_{C}$ (Power to run data center cooling system)
5. $\dot{W}_{CRAC}$ (Electric power to run CRAC)
6. $\dot{W}_{CP}$ (Electric power to run chiller pump)
7. $\dot{Q}_{DC,CRAC}$ or $\dot{Q}_{DC,CT}$ (Data center heat removed either by CRAC or economizer)
8. $\dot{W}_{CTP}$ (Electric power to run cooling tower pump)
9. $\dot{Q}_{CT,A}$ (Heat flow rejected from cooling tower)
10. $\dot{Q}_{C,CT} + \dot{Q}_{CRAC,CT}$ (Heat flow rejection from cooling system to cooling tower)
11. $\dot{Q}_{S,C} = \dot{Q}_{HiT,S} + \dot{Q}_{EXT,S}$ (Heat flow from thermal storage to heat-driven chiller)
12. $\dot{Q}_{EXT,S}$ (Heat flow from an external source of heat energy to thermal storage)

Figure 1. Overall system level diagram showing the work and heat flow paths

The main challenge addressed is generating enough high-temperature heat from the blade components inside the data center, ($\dot{Q}_{HiT,S}$), and then capturing and transporting that high-quality heat effectively and efficiently to the Li-Br absorption chiller.

System Level Description

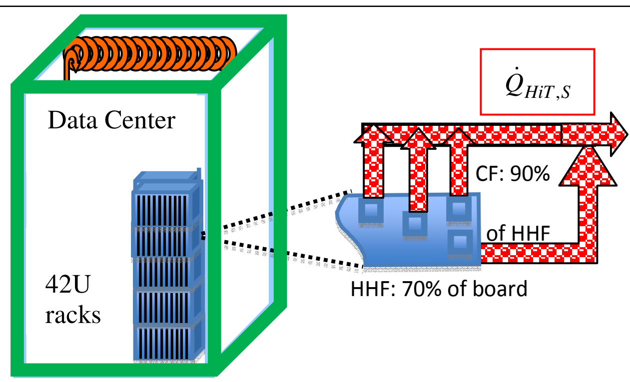

Figure 2. Detail view of heat capturing process from the server blades. A fraction of the heat from the hottest components on each blade is removed by liquid cooling to drive the LiBr absorption chiller.

Figure 1 illustrates the heat flow from the data center to various cooling components. High-temperature, high-quality heat of about 100°C and denoted as $\dot{Q}_{HiT,S}$ is given off by approximately 70% of the components of a server blade inside a typical 42U rack. The server blades of choice are the Dell PowerEdge 1855 blades. To clarify, there are 10 server blades in each Dell PowerEdge 1855 chassis [19] [20], and there can be up to 6 chassis per 42U rack [20]. Each rack has a footprint of 0.6 m^2 [21].

The captured high quality heat from each rack drives a lithium bromide absorption chiller. Additionally, a solar array can be a supplemental external heat source for the thermally driven chiller.

The lower temperature, lower quality heat is denoted as $\dot{Q}_{Air,C}$, and is the heated ambient air in the data center removed by the absorption chiller. The remaining heat is removed by the CRAC or by the Economizer and rejected by the cooling tower.

As shown in Figure 2, an assumed 70% of the blade components are dissipating the most heat. This 0.70 estimation is the fraction of blade components providing the highest quality heat and is referred to as the high heat fraction (HHF). These highest power components are to be in thermal contact with the boiler-end of a thermosyphon, a heat pipe with a boiler and condenser. An expected capture fraction (CF) of 0.90 of this HHF (0.70) is to be captured by each thermosyphon. Each thermosyphon contains heat-exchange fluid that transfers this captured fraction of heat of the HHF to thermal storage to consistently run the absorption chiller.

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 PDUs to the racks ($\dot{W}_{Loss}$), electrical power to the lights ($\dot{W}_{Lights}$), and the electric power consumed to operate the data center cooling system ($\dot{W}_{Cool}$):

$$\dot{W}_{TOT} = \dot{W}_{IT} + \dot{W}_{Loss} + \dot{W}_{Lights} + \dot{W}_{CRAC} + \dot{W}_{CP} \tag{2}$$

and

$$\dot{W}_{Cool} = \dot{W}_{CRAC} + \dot{W}_{CP}\ . \tag{3}$$

The electric power driving the cooling system is, in turn, comprised of the power to run the computer room air conditioner ($\dot{W}_{CRAC}$), assumed to be a vapor compression refrigeration cycle, as well as the power to run the pump of the heat-driven chiller ($\dot{W}_{CP}$). The electrical power consumption by the pump is negligible as it is several orders of magnitude smaller compared to the rest of the power consumptions listed in Table 1.

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

$$PUE = \frac{\dot{W}_{TOT}}{\dot{W}_{IT}} = \frac{\dot{W}_{IT} + \dot{W}_{Cool} + \dot{W}_{Loss} + \dot{W}_{Lights}}{\dot{W}_{IT}} \tag{4}$$

or, since $\dot{W}_{CP}$ is negligible,

$$PUE = \frac{\dot{W}_{TOT}}{\dot{W}_{IT}} = \frac{\dot{W}_{IT} + \dot{W}_{CRAC} + \dot{W}_{Loss} + \dot{W}_{Lights}}{\dot{W}_{IT}} \tag{5}$$

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 equation:

$$\dot{W}_{CRAC} = \frac{\dot{Q}_{CRAC}}{COP_{CRAC}}\ . \tag{6}$$

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

$$\dot{Q}_{CRAC} = (\dot{Q}_{TOT} - \dot{Q}_C). \tag{7}$$

Now, equation (5) becomes

$$PUE = \frac{\dot{W}_{IT} + \frac{\left(\dot{Q}_{TOT} - \dot{Q}_C\right)}{COP_{CRAC}} + \dot{W}_{Loss} + \dot{W}_{Lights}}{\dot{W}_{IT}}\ . \tag{8}$$

The total heat flow from data center, $\dot{Q}_{TOT}$, to be removed by the cooling equipment, is

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

in which $[\dot{W}_{IT} \cdot HHF \cdot CF]$ is the portion of rack heat driving the absorption chiller. Recall that approximately 100% of the electricity into the IT equipment, wires and lights is dissipated as heat.

Furthermore, the heat removed by the absorption chiller, $\dot{Q}_C$, is enhanced by heat in from the server blades, $\dot{Q}_{HiT,S}$, and any supplemental external heating, $\dot{Q}_{EXT,S}$:

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

where

$$\dot{Q}_{HiT,S} = \dot{W}_{IT} \cdot HHF \cdot CF \tag{11}$$

As previously stated, 70% of each server blade provides the highest quality heat. This 0.70 fraction is denoted HHF. Another fraction of this 0.70 will be captured (*CF*) and transferred from the blades to thermal storage to drive the lithium bromide 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.

Substituting equations (9), (10), and (11) into equation (8),

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

where the electric power consumed to drive the absorption chiller pump, $\dot{W}_{CP}$, has been neglected. 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

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

PUE can become less than one, and actually export cooling. This cooling 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.

In order to stay consistent with the definition of PUE as being composed of entirely electric power terms, we can convert any excess cooling to the equivalent power a vapor-compression unit would have to generate, i.e., the exported power would be

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

RESULTS AND DISCUSSION

With rack heat driving the chiller, there is a reduction of heat load on the CRAC. Note that even with no external thermal supplementation and driven solely on a portion of data center IT equipment heat, there is significant cooling contribution by the chiller as evidenced by the reduction of PUE in Table 1.

Of course, the effectiveness of this approach depends on the CF and the factors that affect it. Even without an external heat source, the CF may still be increased. One possible way to achieve an increase is to insulate the hotter components from the data center (e.g., using an insulating gel). Another possible way to increase the CF would be to link more blade components to a thermosyphon (heat pipe) by conduction or convection. This could be accomplished by lowering as much as possible the delta T of the heat capturing scheme. Finally, the size and surface of the condenser and boiler areas may impact the CF percentage, as could the thermal conductivity of the materials of construction of each thermosyphon.

Table 1. Values and calculations for a 100 m^2 data center with 50 racks at an expected capture fraction, *CF*, of 0.90. Here, kWe is taken to mean kilowatt electrical or electric power whereas kWth is kilowatt thermal or heat power.

$\dot{W}_{TOT}$	1,306.7	kWe
$\dot{W}_{IT}$	1,176.0	kWe
$\dot{W}_{Loss}$	130.7	kWe
$\dot{W}_{Lights}$	1.4	kWe
$\dot{Q}_{HiT,S}$ (70% of $\dot{Q}_{IT}$)	740.9	kWth
$\dot{Q}_C$ (COP=0.7)	518.6	kWth
PUE (CF=0.90)	1.12	
Coefficients of performance		
COP_{CRAC}	3.9	*typical COP_R
COP_C	0.7 [22]	LiBr chiller

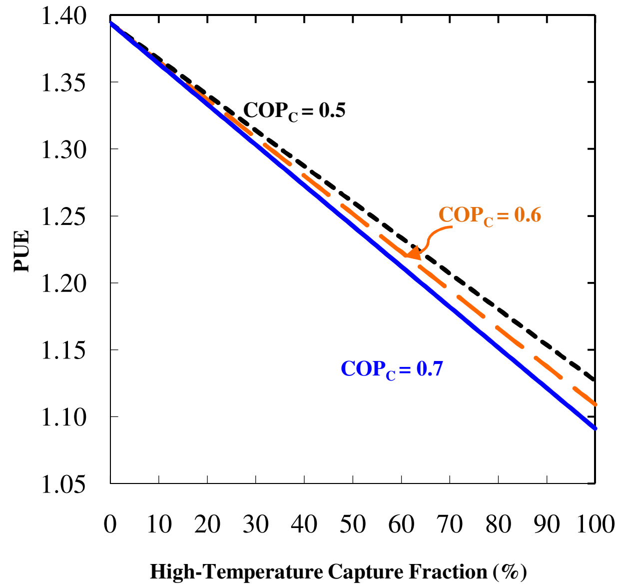

Figure 3. Power usage effectiveness (PUE) versus high-temperature capture fraction (*CF*) for three different values of heat-driven chiller coefficient of performance (COP_C). These results assume that 70% of $\dot{Q}_{IT}$ is high-temperature heat ($\dot{Q}_{HiT,S}$). A potentially realizable (CF=0.90) PUE is 1.12 from varying only the capture fraction percentage and excluding external heat source (zero solar) contribution.

Figure 3 presents 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 ~ 1.4 As expected, PUE decreases with increasing CF, and with increasing COP_C, leading to a PUE as low as PUE ~ 1.1. These PUE values can be compared to the benchmarks shown in Table 3 [23], suggesting a "very efficient" data center for PUE ≤ 1.2.

Table 2. Illustrating the effect on PUE by fixing the capture fraction at 90% and varying the external heat source.

$\dot{Q}_{EXT,S}$ (kW)	$\dot{Q}_{EXT,S}/\dot{W}_{IT}$	PUE vs $\dot{Q}_{EXT,S}$ (COP_C=0.5)	PUE vs $\dot{Q}_{EXT,S}$ (COP_C =0.6)	PUE vs $\dot{Q}_{EXT,S}$ (COP_C =0.7)
0	0.000	1.154	1.138	1.121
5	0.007	1.153	1.137	1.121
10	0.013	1.152	1.136	1.119
. .	. .	. .	. .	. .
1170	0.995	1.027	0.986	0.944
1175	0.999	1.026	0.985	0.944
1180	1.003	1.026	0.984	0.943

The situation, not surprisingly, is improved if we can supplement the high-temperature heat removed from the data center with an external source of heat ($\dot{Q}_{EXT,S}$), such as from solar energy. The values in Table 2 and Figure 4 show that PUE can be reduced to less than one, meaning that cooling can potentially be exported from the data center.

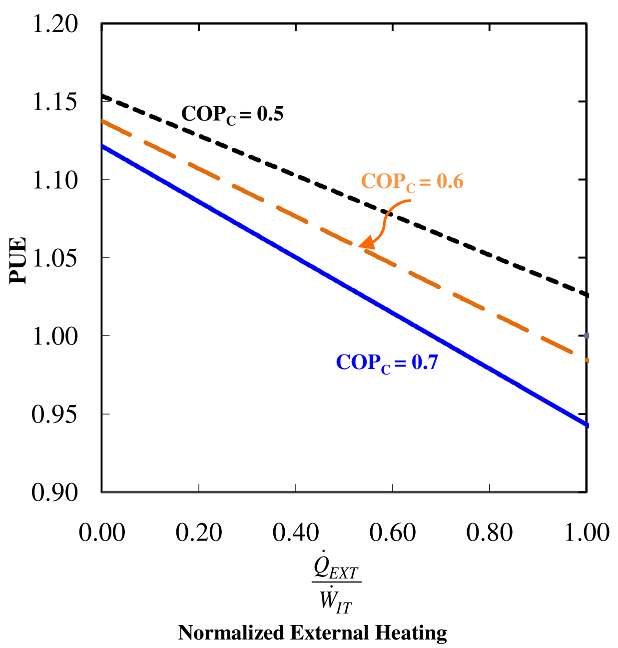

Figure 4. Power usage effectiveness (PUE) versus normalized external heat source for three different values of heat-driven chiller coefficient of performance (COP_C). With both a CF percentage of 90% and 100% solar contribution, the most efficient realizable PUE is 0.94.

CONCLUSIONS

Data center electric power consumption is an acknowledged problem that it likely to get worse in the future. The potential exists to utilize some of the waste heat generated by data centers to drive absorption chillers, which would then relieve some of the cooling load on the conventional computer room air conditioner (CRAC). By supplementing the high-temperature heat captured from the data center racks with an external source of heating, such as from solar energy, it is theoretically possible to generate a PUE (Power Usage Effectiveness) ratio of less than one.

ACKNOWLEDGMENTS

The authors gratefully acknowledge the support of the National Science Foundation through award 0855277. Special and continual thanks to the Arizona State University Thermal Engineering and Bluetool Teams for their participation and support. Also, much appreciation sent to SRP and Yazaki for their generous sponsorship and technical expertise. Finally, thanks to Pamela Dunn and all the supportive staff of the School of Computing, Informatics, and Decision Systems Engineering Computer Science and Engineering at Arizona State University.

Table 3. Industry benchmarked data center efficiencies [23].

PUE	Level of Efficiency
3.0	Very Inefficient
2.5	Inefficient
2.0	Average
1.5	Efficient
1.2	Very Efficient

REFERENCES

[1] American Power Conversion. (2009). “Energy Efficient IT: The 21st Century Brings A Shift In Data Center Efficiency” EBook Retrieved from http://manageddatacenter.searchdatacenter.com/document;5141550/research-center.htm.

[2] BlueTool: A research infrastructure project toward cyber-physical data center management. PowerPoint presentation and CRI-II-EN-BlueTool--Description.pdf, 2009. Arizona State University.

[3] Kenneth G. Brill, “Data Center Energy Efficiency and Productivity,” *The Uptime Institute White Paper presented at The Uptime Institute Symposium 2007*: *The Invisible Crisis in The Data Center: How IT Performance is Driving the Economic Meltdown of Moore’s Law.*

[4] Qinghui Tang; Sandeep. K. S. Gupta; Daniel Stanzione; Phil Cayton, “Thermal-Aware Task Scheduling to Minimize Energy Usage of Blade Server Based Data centers,” *Proceedings of the 2nd IEEE International Symposium on Dependable, Autonomic and Secure Computing (DASC'06),* March 2006.

[5] American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc. Thermal guidelines for data processing environments., Atlanta, GA, 2004.

[6] Hanson, Heather, Stephen W. Keckler, Soraya Ghiasi, Karthick Rajamani, Freeman Rawson and Juan Rubio. “Thermal response to DVFS: analysis with an Intel Pentium M,” *International Symposium on Low Power Electronics and Design. Proceedings of the 2007 international symposium on Low power electronics and design.* New York, NY, USA, 2007, pp. 219-224.

[7] Y. Chen, A. Das, W. Qin, A. Sivasubramaniam, Q. Wang, and N. Gautam, “Managing server energy and operational costs in hosting centers,” in *SIGMETRICS ’05: Proceedings of the 2005 ACM SIGMETRICS International conference on measurement and modeling of computer systems.* New York, NY, USA: ACM Press, 2005, pp. 303–314.

[8] B. Khargharia, S. Hariri, and M. S. Yousif, "*Autonomic power and performance management for computing systems*," Cluster Computing, vol. 11, no. 2, pp. 167–181, Jun. 2008.

[9] Q. Tang, S. K. S. Gupta and Georgios Varsamopoulos. "Energy-Efficient, Thermal-Aware Task Scheduling for Homogeneous, High Performance Computing Data Centers: A Cyber-Physical Approach." Transactions on Parallel and Distributed Systems, Special Issue on Power-Aware Parallel and Distributed Systems (TPDS PAPADS), 19(11):1458–1472, November 2008.

[10] T. Mukherjee, A. Banerjee, G. Varsamopoulos, S. K. S. Gupta, and S. Rungta, "*Spatio-temporal thermal-aware job scheduling to minimize energy consumption in virtualized heterogeneous data centers*," Computer Networks, vol. 53, no. 17, pp. 2888–2904, 2009.

[11] ASHRAE Technical Committee 9.9, Mission Critical Facilities, Technology Spaces and Electronic Equipment. "Best Practices for Datacom Facility Energy Efficiency," 2nd Edition. ASHRAE, 2009.

[12] R. R. Schmidt, E. E. Cruz, and M. K. Iyengar, "Challenges of data center thermal management," IBM Journal of Research & Development, vol. 49, no. 4/5, pp. 709–723, 2005.

[13] T. Brunschwiler, B. Smith, E. Ruetsche, B. Michel. "Toward zero-emission data centers through direct reuse of thermal energy." IBM Journal of Research and Development. 53(3), pp. 11:1-11:13, 2009.

[14] R. H. Katz, "*Tech titans building boom*," Spectrum, IEEE, vol. 46, no. 2, pp. 40–54, Feb. 2009

[15] Belady, C.L. and Malone, C.G., "Metrics and an infrastructure model to evaluate data center efficiency," Proceedings of the ASME InterPACK Conference 2007. IPACK2007, 751-5, 2007

[16] Neil Rasmussen, "Implementing Energy Efficient Data Centers," *American Power Conversion White Paper #114*, 2006.

[17] Ashrae. 2009 *Best practices for datacom facility energy efficiency*. Atlanta, GA: American Society of Heating Refrigerating and Air-Conditioning Engineers, Inc.

[18] Google. (2009). "Going green at google | clean energy initiatives,"http://www.google.com/corporate/green/datacenters/measuring.html.

[19] Dell PowerEdge 1855 Blade Server. Specifications, http://www.newservers.com/1855-specs.pdf.

[20] Dell™, Dell Datacenter Capacity Planner (DCCP). http://www.dell.com/html/us/products/rack_advisor_new/index.html?~lt=popup.

[21] Dell™, Dell™PowerEdge™Rack Systems. POWEREDGE RACK 4210 42U CABINET Specifications, p.4. December 2003. http://www.verinet.dk/dell_rack_system.pdf.

[22] Yazaki Energy Systems, I. (2002-2009). *Yazaki Energy - Gas Fired Double Effect Chiller-Heaters*. Retrieved December 2009, from Yazaki Energy Systems, Inc.: http://www.yazakienergy.com/waterfiredspecifications.htm.

[23] Green, Grid. (2009). "PUE and DCiE calculator: how to calculate your PUE/DCiE." Retrieved from http://www.42u.com/measurement/pue-dcie.htm