

Cold Air Containment

Mikko Pervilä
Department of Computer Science
University of Helsinki
pervila@cs.helsinki.fi

Jussi Kangasharju
Helsinki Institute for Information Technology
University of Helsinki
jakangas@cs.helsinki.fi

ABSTRACT

This article describes two benchmark studies involving the cooling technique known as cold aisle containment (CAC). One test case studies a 26U server rack operating on unconditioned outside air only in a carefully controlled setup. The other examines a server room with a power draw of over 80 kW during normal operation. In both cases we measure how incorporating CAC changes the air flow, electricity consumption, operating temperatures, and cooling requirements. Our results show how the air flow separation affects the temperatures in the server room and verify that using CAC can reduce CRAC power by roughly a fifth.

Categories and Subject Descriptors

B.8 [Hardware]: Performance and Reliability; B.8.1 [Performance and Reliability]: Reliability, Testing, and Fault-Tolerance

General Terms

Experimentation, Reliability

Keywords

Sustainable computing, cooling, empirical system reliability

1. INTRODUCTION

For the past 20 years, the standard method of building data centers (DC:s) is the so called hot aisle / cold aisle layout. Server racks are placed in rows so that their front panels face the same direction, and a gap, called an *aisle*, is left between the rows. If the DC does not use in-row air conditioning units, the floor is raised so that an air pathway forms underneath it. The aisle between the front panels uses perforated floor tiles. Computer-room air conditioning (CRAC) units pump cold air into the raised floor space, so that the air exits upwards through the perforated tiles and into the front panels of the server racks. The aisles with perforated tiles are called cold aisles.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GreenNet'11, August 19, 2011, Toronto, Ontario, Canada.
Copyright 2011 ACM 978-1-4503-0799-4/11/08 ...\$10.00.

Exhaust air is vented through the rear panels of the racks. The rear panels oppose the next row's rear panels, again with an aisle between them. The tiles between the rear panel rows are not perforated. The hot air rises naturally upwards towards a clear pathway between the tops of the racks and the ceiling. Thus, the exhaust aisle is called the hot aisle. CRACs use powerful fans to draw in the exhaust air and cool it into new supply air, which is pumped back under the floor.

Figure 1(a) shows the motion of air under these ideal conditions. In real data centers, multiple factors render the air circulation cycle sub-optimal. The major issues are obstructions in air flows, unmatched supply and intake air flows, cold air leakages, localized hotspots, and hot air recirculation.

Much of the work concerning the effects of air flows on data center cooling needs and electricity consumption has been done through emulation [5, 7], not real equipment [1]. In this paper, we push the envelope further by experimenting with cold air containment using real servers in both a controlled setup as well as an operating data center.

1.1 Related work

Data centers are under constant scrutiny to improve cooling efficiency as the amount of power used by the ICT industry increases. The big players in the field are hardware and software manufacturers that employ the largest data centers. Not all of their solutions are known [4], since cooling improvements yield direct reductions of the power requirements of a data center. These reductions translate directly as economic gains and thus become competitive advantages.

As early as 1991, Nakao et al. [7] considered the effects of hot air circulation caused by exhaust air short circuiting back as supply air. Recently, white papers describing new ideas surface at an increasing pace. Space constraints force us to we prune this section to those articles that are most relevant to air stream containment, with a strong emphasis on measured experiments.

Intel's white paper by Atwood and Miner [2] is attributed as the pioneer work in using natural air for data center cooling. Their technique was later verified by Microsoft's [3] tent experiment, while we extended the technique's feasibility into a much harsher climate [9].

Natural air cooling can be impossible to retrofit to an existing DC, either due to the ambient climate or building limitations. A complementary solution is aisle containment. Hilss and Iyer [5] describe HP's results with CAC. They performed an emulation-based study using ten racks equipped with 2.4 kW load banks. Using CAC yielded a maximum




Figure 1: Separation of cold and hot air streams

CRAC efficiency improvement of 41% when the racks were loaded with an average of 9 kW. Adams [1] measured the effects of the complementary technique, Hot Aisle Containment (HAC). The study omits a comparison to the situation before HAC was installed.

Patterson [8] analyzed the tradeoff between centralized cooling through CRAC:s vs distributed cooling with server fans. Both are modelled diligently, with the conclusion that turning up the supply temperature can result in negative savings as server fans will need to ramp up their rotations per minute. Thus, cooler supply air could reduce the total power draw of the servers.

In this paper, we verify the results of Hills and Iyer [5] in a real data center environment. We extend the work of Adams [1] by measuring the opposite CAC technique and also include data before and after our changes. Section 2 describes our initial prototype with some preliminary results. In Section 3 we study the development of a real-world, small scale DC with a power draw of 80 kW during baseline operation and over 110 kW during high-performance cluster computations. Section 4 concludes with lessons learned while building a home-brew CAC setup, and what changes can be expected by those willing to duplicate our methods.

2. COLD AISLE CONTAINMENT

Both aisle containment techniques are simple in their key idea: either the hot or the cold aisle is covered at the top and edges of the aisles. This forces the hot and cold air streams to separate. The shaded areas in Figures 1(a) and 1(b) show how CAC limits the flow of the cold air stream so that it must pass through the equipment racks.

In both cases the aisles must be refurbished so that leakages are minimized. Reasonably airtight doors are required at the edges to allow for operator access, and cable ducts must be isolated to prevent leakages. Different vendors' solutions range from purchasing entirely new racks to installing plastic curtains constraining the air flows. Obviously, replacing the racks is a very time-consuming and delicate operation, which makes retrofit-capable solutions more desirable.

Our solution consists of see-through plastic blankets, a handful of PVC tubes, and plenty of duct tape. We decided to verify whether a solution constructed as cheaply as possible would yield benefits for energy consumption, cooling requirements, or operating temperatures of an existing set of servers. The major difference to [5] is that we use real servers: the combined effect of hundreds of temperature-driven fans might yield additional benefits when supplied with colder air.

Our key research questions were as follows. If one or more of the answers proved positive, other DC:s in the small to medium range could improve their existing installations' energy efficiency with a shoe-string budget.

1. Can CAC reduce the cooling load of the CRAC:s?
2. Will CAC yield a more uniform supply air temperature?
3. Will the servers use less energy in total when given more uniform supply air?
4. Does CAC reduce the electricity consumption of our DC?

The experiments take a black-box approach at the rack level. We measure the supply and exhaust air temperatures as well as the aggregate energy consumption of the data center, including CRAC:s but excluding the chiller plants and cooling towers. These exclusions are caused by practical limitations at our installation. Finally, we do not monitor internal server temperatures, but assume that they are linear to inlet temperatures. Despite these limitations, we show that CAC is an efficient technique and can be retrofitted with very low additional costs to existing DC installations.

2.1 Prototype Helsinki Chamber

As a first step, we constructed a small, controllable prototype setup to verify how CAC would help in stabilizing inlet air temperatures. This prototype was designed to employ both CAC and unconditioned natural air cooling.

Continuing on our previous, experimentation-based studies on air-based free cooling [9], we now operate 14 rack servers in a modified enclosure on the roof terrace of the Department of Computer Science. Two of our servers were used in our previous experiment as well; they have been cooled with unconditioned outside air for 14 months at the time of writing.

Our custom-built enclosure contains a 26U 19" Rittal rack. We have named this type of enclosure a *Helsinki Chamber* (HC), following the style of previously existing Kyoto Wheels [6]. Figure 2 shows the schematics of our own construction. A full evaluation of the merits and flaws will follow in a later work, when we have gathered data for the summer months as well.¹

¹A view of the setup is available from <http://www.cs.helsinki.fi/Exactum-kamera/>.

We designed the HC not only to employ CAC during the warmer months of the year, but to intentionally recirculate exhaust heat during winter. However, due to an unexpected very sharp temperature drop in February 2011, we have now employed CAC with supply air as low as -23°C . Our prototype uses duct tape to isolate all leakages, but a temperature-driven vent would be straight-forward to add. The ideal solution would be that the servers themselves would recirculate the warm exhaust air when the inlet temperature drops too low.

2.1.1 Construction and instrumentation

The 26U of servers running in the HC have a combined power draw of over 3.4 kW. A full 42U rack would imply a power draw of 5.5 kW. The servers draw in supply air from the lower front of the HC, through two standard, Euro-type pallets elevating the HC about 30 cm from the roof of the building. After periods of heavy snowfall, some manual labour is required to ensure that the intake vents do not entirely block.²

The HC is instrumented with four Lascar EL-2-USB-LCD data loggers³ that record temperatures, relative humidities, and dew points. The sensors are marked in Figure 2 as $A - D$. Sensor A is positioned at the highest rack server and B is positioned at the intake of the HC. The difference $A - B$ shows how much supply air is warmed by exhaust air recirculation combined with heat radiated by the servers. Sensor D is positioned at the bottom of the exhaust section, while C is positioned at the top of the exhaust. The difference $D - C$ shows the buildup of exhaust heat. Finally, the pairwise differences $A - C$ and $B - D$ show how much exhaust heat the servers are venting.

2.1.2 Restricted air flow

We isolated the perforated holes of the 19" rack using duct tape on February 10th, 2011. Before this, we had allowed hot air to recirculate and mix with the supply air. Because of the sharp temperature drop just after Feb. 10th, intake air temperatures dropped below -20°C . In order to verify how much the isolation affected inlet and exhaust temperatures, we did a pairwise search for matching ambient temperatures from Feb. 1st to March 1st, 2011. As the HC air intake is indirect, we considered the wind factor effect to be negligible. A visual inspection of the graphs showed no evidence of skew due to sunlight.

Due to space concerns, we have omitted the graphs from this article. They all follow the same pattern: there is a sharp increase in sensors A , C , and D before CAC installation caused by the increased computational load on the servers. After the change, readings from sensor A drop by a regular -10°C , whereas readings from sensor D increase by 10°C . In other words, exhaust recirculation is now blocked and the temperature difference $A - B$ lowers from around 20°C to 8.5°C on average. The remainder was caused by exhaust recirculation through the servers, and eliminated later.

The experiment on our prototype HC gave confidence that cold air containment was a feasible technique for stabilizing the inlet temperatures of a 26U rack.

²This turned out to be much less work than we expected. The first author spent a total of $3 * 30$ minutes shoveling snow during last winter.

³See <http://www.lascarelectronics.com/> for full details.




Figure 2: Helsinki Chamber schematics, side view, and the locations of the data loggers $A - D$

3. EXACTUM DATA CENTER

As the follow-up, we moved from cold air to cold aisle containment in the server room of our CS department building (named “Exactum”). The server room is a small-scale data center housing the combined IT equipment of the Department of Computer Science and the Helsinki Institute for Information Technology (HIIT).

The DC has undergone a very natural evolution from using mid-tower PC:s to extensive virtualization and state-of-the-art blade enclosures. Due to its history, our set of hardware is extremely heterogeneous. The latest addition to the palette is the Ukko High-Performance Computing Cluster, which consists of 240 blade computers and a total of 1920 physical cores. Ukko by itself is specified for a maximum active power draw of 74 kW.

For the CAC experiment, we chose a before and after type measurement. Both are further subdivided into an idle phase, where the Ukko cluster idles with almost no load, and an active phase, where Ukko is given a computationally intensive task that raises its active power draw from under 34 kW to a steady 68 kW.

3.1 Different power lines

For energy measurements we used Hager EC-370 power consumption meters⁴ read through their LED light outputs. We soldered two BPW85A phototransistors to the DTR-DTS and RTS-CTS signal pairs of a RS-232 cable, and then used the RXTX Java library to read the RS-232 signals. Each LED signal was translated into fractional kWh:s and then recorded into a RRDtool database.

The power consumption meters were connected before UPS devices and power distribution units (PDU:s) in order to capture the actual power usage of the DC. The measure-

⁴See <http://www.hager.com/> for full details.

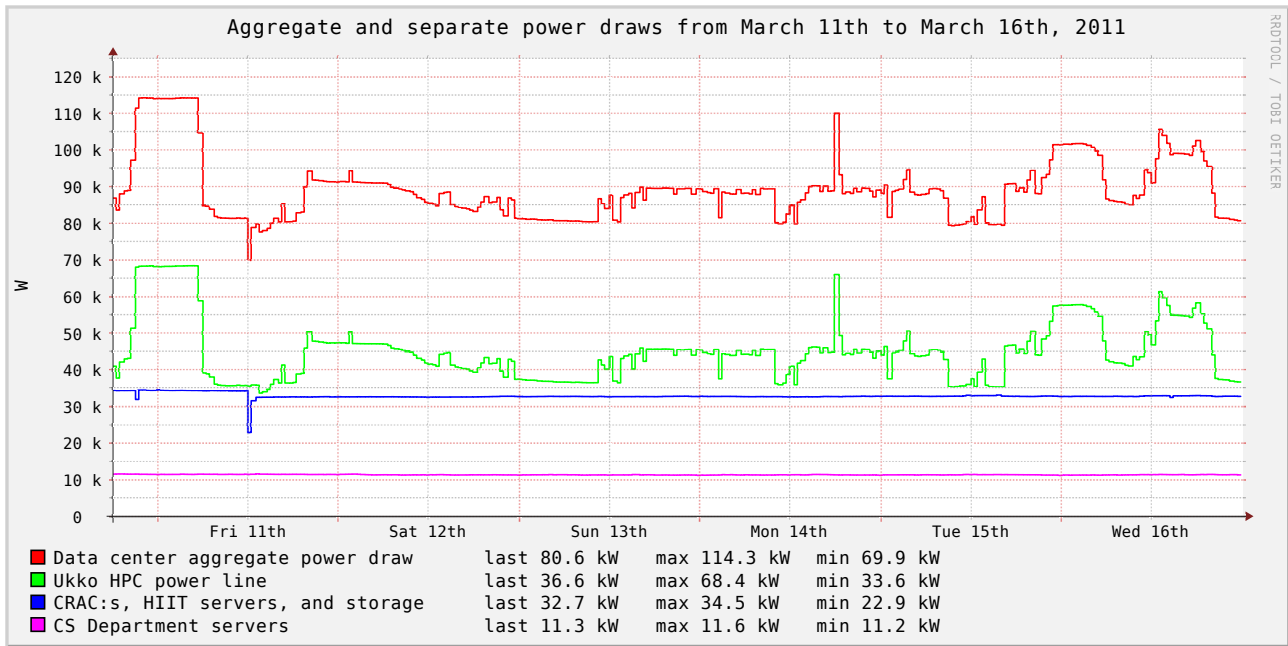


Figure 3: Aggregate and separate power draws of the Exactum data center

ments contain the CRAC units but exclude chiller plants and cooling towers.

Due to historical reasons, Exactum’s DC electricity is supplied by three different power supply lines. This was very beneficial for our experiment, as it allowed us to separate the loads of the infrastructure services (two line) from the power-intensive loads of the Ukko cluster (one line). Most importantly, we discovered that the infrastructure services had very stable power draws during our experiments, whereas the total power consumption of the DC was dominated by Ukko’s computational load. This allowed us to measure both a baseline and a pathological power draw of the DC by controlling only the CPU loads of Ukko.

Figure 3 depicts the power draw of the three power supply lines and their sum as the aggregate power draw of the entire DC. Our infrastructure services are divided into two categories: CS Department servers and the rest, i.e., CRAC:s, HIIT servers, and storage units. As can be seen from the figure, the infrastructure power draws vary within a few hundred watts normally. The sharp drop visible during Friday, March 11th was caused by human error, as the power meter’s phototransistor got accidentally dislodged from the LED.

The EC-370 meters were able to measure both active and reactive power readings. We monitored the power factors before and after CAC installation, both during peak and normal computational loads. Probably due to the UPS devices, the power factors of the power lines connected to the infrastructure servers remained stable at 0.99 and 0.89. For Ukko, there was a small change from 0.96 to 0.97 during periods of heavy computation. During these periods, the DC’s combined power draw remained steady at 114 kW.

3.2 Temperature changes

Even without the peak 114 kW loads, our CRAC:s had started to struggle. As the DC’s physical size is only just over 70 m², we were operating at idle and peak power den-

sities of over 1.1 and almost 1.6 kW/m². An earlier inlet monitoring system reported a steady increase of about 4°C within the last year. According to our vendor, the cooling system had reached its maximum capacity and could not be further optimized.

We instrumented the data center using five Lascar EL-2-USB-LCD data loggers. As we were mainly interested in the mixing of cold and hot air streams, not all of the inlet temperatures of the servers, the low number of measurement points was deemed sufficient.

Figure 1(b) shows our rack configuration and data logger placement. The first aisle consists of four racks for the Ukko cluster and three racks for storage devices. The second aisle consists of four racks for the CS Dept.’s systems, and four racks for HIIT’s systems.

Data loggers were placed as follows. Sensor *B* was placed under the raised floor in front of the middle CRAC. Sensors *A* and *C* were placed near the tops of HIIT’s and Ukko’s racks. The differences $C - B$ and $A - B$ indicate supply air temperature elevation from the CRAC:s to the server inlets nearest to the rack tops. Sensor *D* was positioned near the middle of the Ukko’s exhaust panels. The difference $A - D$ shows the maximum temperature elevation for inlet to exhaust air, as Ukko was known to be the most power-intensive of the servers. Finally, sensor *O* was positioned at the return intake of the middle CRAC, and $O - B$ shows the temperature delta at the CRAC.

3.3 Cooling fluid loops

Exactum’s DC was retrofitted with additional cooling in early 2010 when Ukko was installed. Previously, the DC ran on two CRAC:s supplied by the neighboring building’s chiller plant. With the increase in heat load, a new, dedicated chiller plant and four new CRAC:s were installed to Exactum. The previous two CRAC:s were left running, bringing the total to two old and three new units in the data

center, plus one new unit in the adjacent room housing the UPS devices.

Both chiller plants are connected to their own cooling towers via secondary, closed loops. Both towers use free cooling whenever outside temperatures are below 3°C. Using free cooling simplifies the analysis, for we can study the baseline effects of CAC without considering the average times and amount of power direct expansion systems spend running.

Sadly, the chiller plants proved impossible to instrument conclusively. We were forced to compromise due to the lack of line valves in the pathways of the cooling fluids. The only possibility was to measure the intake of primary cooling fluid (water) for the older CRAC:s at two different valves, and the return flow of cooling liquid (35% ethylene alcohol) from the tower to the newer chiller plant at Exactum.

Exactum’s secondary cooling liquid loop did not reveal much data. The temperatures followed ambient temperatures closely, which is not surprising considering the effectiveness of free cooling during the winter in Finland. The older chiller plant’s primary cooling loop proved more interesting, since it shows the chiller plant reacting to Ukko’s computational load. Before CAC, the chiller plant would periodically drop down the cooling fluid temperature by about 2°C. After CAC, these drops are no longer required, meaning that less work is required of the chiller plant, and thus, less energy is used in general.

3.4 Structural changes

We used duct tape and plastic cut-outs to isolate openings between our racks and the servers. Gaps between the side walls and openings underneath the racks were covered in a similar manner. All cracks in the raised floor were carefully covered. The front panels of the Ukko cluster and storage units were sufficiently isolated, requiring no further work.

On top of the racks we installed PVC tubes to support the roof of the CAC aisle. We covered the PVC tubes and the edges of the cold aisles using two long sheets of see-through plastic. Our original plan was to build a continuous cover for both aisles, but in practice it was easier to build two aisles with their own doorways. In hindsight, this was a mistake, for the CAC aisles pressurized unequally. We have later reconstructed the two into a single, continuous aisle.

After the covers had been secured, both aisles were pressurized by the supply air supplied from the CRAC:s. An unforeseen benefit of using the plastic covers is that they flex along with the air pressure. The DC operators can easily verify that each aisle is provided with enough supply air: if the demand exceeds the supply, the covers are sucked inwards by the negative pressure. In the opposite case, the covers blow outwards. We try to avoid this for maximum efficiency and also due to structural reasons, since there are limits to how much tension the duct tape can withhold.

The CAC installation took less than three hours for three persons.⁵ The material costs amount to 29.50 € for the plastic covers, 37.14 € for the PVC tubes, and 54.01 € for the duct tape, for a total of 120.65 €.

3.5 Results

In order to compare temperature changes we selected two time frames consisting of 16 hours each. The time frame before CAC runs from March 4th, 12:00 to March 5th, 03:56.

⁵A video of the finished installation is available from <http://www.youtube.com/watch?v=jZIUJIVYsDs>

The first 7.5 hours consist of normal operating load, whereas for the following 8.5 hours the Ukko cluster calculates near full capacity. The time frame after CAC runs from March 10th, 20:00 to March 11th, 12:00. Ukko runs near full capacity for the first 9 hours, after which the DC resumes normal operation.

Figure 4(a) shows the supply temperature measured by sensor *B*. All CRAC:s had already been set to operate at full cooling capacity, and their fans set to maximum power. The supply temperature was just above 15°C. Figure 4(c) shows that when this supply air reached the highest server inlets, the temperature had already risen to above 19°C. This elevation was caused by exhaust recirculation. When Ukko operated near full capacity, the elevation kept on rising above 21°C.

Figures 4(b) and 4(d) show the changes after CAC installation. We were able to return the three newer CRAC:s back to temperature-driven fan speeds and set new supply temperatures to 25°C. The older CRAC models were left to operate at the same setting as before the change. Mixing the different CRAC supply air streams caused both Ukko’s and HIIT’s (not shown) inlet temperatures to stabilize near 18°C, about 1°C less than previously. Further, running Ukko near full capacity no longer had visible effects on its inlet temperatures as exhaust recirculation was eliminated. This was also verified by plotting the inlet temperature differences $A - C$ (not shown).

Return temperatures rose only about 2°C. The elevation was reduced by the fact that we removed all perforated tiles from the hot aisles. They had been installed in order to mitigate exhaust hotspots near heavily loaded servers.

After the changes, the CRAC:s kept operating at very minimal capacities. We continued to experiment and were later able to **turn one newer CRAC off completely** while still maintaining the same inlet temperatures as before CAC installation. The other CRAC:s did not have to compensate: the end result was a 2 kW reduction in total DC energy consumption. The drop is just barely visible in Fig. 3, above the label “Fri 11th”. As can be expected, the drop is minimal when compared to the aggregate power draw of the DC. The real benefits can be reaped by avoiding the purchase of an additional CRAC unit.


Finally, both our IT staffs have reported that the DC working conditions *feel* better, perhaps due to the exhaust heat being removed more effectively. As our operators are by necessity quite conservative, this reaction to the change seemed surprisingly positive.

4. CONCLUSION AND FUTURE WORK


Depending on how the individual CRAC:s consume power, our low-cost CAC has reduced our cooling requirements by roughly one fifth. This directly verifies Hilss and Iyer’s [5] findings and answers our first research question positively. It also means that we can now install more servers and thus increase our thermal load further.

Inlet air temperatures for our computing cluster are no longer affected by exhaust recirculation which means a more uniform operating temperature. This answers our second research question positively. Since our cooling equipment no longer operates over its capacity, Ukko can be fully loaded for longer periods of time.


We have noticed no changes in the server fans as their power draw remained the same. It is not surprising, as




(a) Supply temperatures before CAC



(b) Supply temperatures after CAC



(c) Ukko's inlet temperatures before CAC



(d) Ukko's inlet temperatures after CAC

Figure 4: Temperature changes before and after CAC installation

the inlet temperatures did not reach above 22°C even before CAC. Thus, our third research answer is a tentative no. Additional experimentation is necessary to measure how steadily raising the supply temperature will affect the aggregate power draw of the DC.

As for the fourth research question, CAC has been able to reduce our DC's total power consumption, although with only 2 kW as the computational loads remain unchanged. At the time of writing it seems entirely plausible that we might be able to reduce a second CRAC with further changes, e.g., better supply air pressure stabilization, CRAC management based on CPU loads, and better under-floor air flow management.

5. ACKNOWLEDGMENTS

The authors would like to specially thank Ville Hautakanen and Onni Koskinen from the CS department's IT staff for their assistance while planning and implementing the actual CAC construction. Also from the IT staff, Pekka Niklander was invaluable while planning the power measurements. Olli Moisio from the technical department helped us with the cooling fluid flow and power measurements. Last but not least, Mikko Rantanen designed and implemented the phototransistor connection to the power measurement units.

6. REFERENCES

- [1] B. J. Adams. Performance Evaluation for Modular, Scalable Cooling Systems with Hot Aisle Containment in Data Centers, 2009. <http://escholarship.org/uc/item/3536r5nv>.
- [2] D. Atwood and J. G. Miner. Reducing Data Center Cost with an Air Economizer, 2008. http://www.intel.com/it/pdf/Reducing_Data_Center_Cost_with_an_Air_Economizer.pdf.
- [3] C. Belady. Intense Computing or In Tents Computing?, 2008. http://blogs.msdn.com/b/the_power_of_software/archive/2008/09/19/intense-computing-or-intents-computing.aspx.
- [4] Google. Efficient computing, 2011. <http://www.google.com/corporate/datacenter/efficient-computing/>.
- [5] R. Hilss and S. Iyer. Improved data center efficiency—incorporating air stream containment, 2009. <http://h10134.www1.hp.com/insights/whitepapers/data-center-efficiency.aspx>.
- [6] R. Miller. Heat Wheel Could Cut Data Center Cooling Bills, 2008. <http://www.datacenterknowledge.com/archives/2008/11/14/heat-wheel-could-cut-data-center-cooling-bills/>.
- [7] M. Nakao, H. Hayama, and M. Nishioka. Which cooling air supply system is better for a high heat density room: underfloor or overhead? In *Telecommunications Energy Conference*, pages 393–400, Kyoto, 1991.
- [8] M. K. Patterson. The effect of data center temperature on energy efficiency. In *2008 11th Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems*, pages 1167–1174. Ieee, May 2008.
- [9] M. Pervilä and J. Kangasharju. Running Servers around Zero Degrees. *Computer Communication Review*, 41(1):96–101, 2011.