Two-phase Cooling of Datacenters: Reduction in Energy Costs and Improved Efficiencies

Jonathan Albert Olivier, jonathan.olivier@epfl.ch
Jackson Braz Marcinichen, jackson.marcinichen@epfl.ch
John Richard Thome, john.thome@epfl.ch
Laboratory of Heat and Mass Transfer (LTCM)
École Polytechnique Fédérale de Lausanne (EPFL)
Station 9, CH-1015 Lausanne, Switzerland
Phone: +41 21 693 5894, Fax: +41 21 693 5960

Abstract. *Cooling of datacenters is estimated to have an annual electricity cost of 1.4 billion dollars in the United States and 3.6 billion dollars worldwide. Currently, refrigerated air is the most widely used means of cooling datacenter's servers. According to recent articles published at the ASHRAE Winter Annual Meeting at Dallas, typically 40% or more of the refrigerated airflow bypasses the server racks in datacenters. The cost of energy to operate a server for 4 years is now on the same order as the initial cost to purchase the server itself, meaning that the choice of future servers should be evaluated on their total 4 year cost, not just their initial cost. Based on the above issues, thermal designers of data centers and server manufacturers now seem to agree that there is an immediate need to improve the server cooling process. Especially considering that modern datacenters require the dissipation of 5 to 15MW of heat, and the fact that 40-45% of the total energy consumed in a datacenter is for the cooling of servers. This makes the way in which servers are cooled and the recovery of the dissipated heat all the more attractive. This would also lead to a reduction of the overall* CO_2 *footprint of the datacenter. The present study develop a case study considering a vapour compression and a liquid pumping cooling systems applied on a datacenter and exploring the application of energy recovered in the condenser on a feedwater heater of a coal power plant. Aspects such as minimization of energy consumption and* CO_2 *footprint and maximization of energy recovery and power plant efficiency are investigated. The results showed for the added heat, when the pressure of the feedwater heater is optimized, an increase of the overall plant efficiency up to 2.2%. For a 500MW coal power plant, this translates to a savings of roughly 195 000 tons of* CO_2 *per year.*

Keywords*: datacenter, power plant, energy recovery, vapour compression and liquid pumping systems, micro-evaporator, two phase flow.*

1. INTRODUCTION

Reduction of primary energy consumption is strongly required to mitigate global warming. To achieve this objective, the use of waste heat thermal energy or renewable energy should be increased in energy conversion processes, reducing fossil fuel consumption. Under the current efficiency trends, the energy usage of datacenters in the US is estimated to be more than 100 billion kWh by 2011, which represents an annual cost of approximately $7.4 billion (EPA, 2007). With the introduction of a proposed carbon tax in the US (Larson, 2009), the annual costs could become as high as $1.4 trillion by 2012, increasing annually.

Cooling of datacenters can represent up to 45% (Koomey, 2007) of this total consumption using current cooling technologies (air cooling). In the US, this relates to an estimated 45 billion kWh usage by 2011 with an annual cost of $3.3 billion, or $648 billion with the inclusion of carbon tax... just for cooling. A problem with a datacenter is that all the energy consumed is converted into heat, which with current technologies is rejected as waste into the atmosphere. Therefore, reusing this waste heat can potentially not only reduce the overall operating costs, but also the carbon footprint of the datacenter.

Most datacenters make use of air-cooling technologies to ensure the correct running of the servers contained within. The limits of air-cooling, however, are being approached due to the performance increase in microprocessors, which will have heat fluxes in the order of 100 W/cm^2. It was shown that air has a maximum heat removal capacity of about 37 W/cm^2 (Saini and Webb, 2003). The problem is made worse with servers being more densely packed, such as blade centers of IBM (IBM, 2010), which could see racks generating in excess of 60 kW of heat. Today's datacenters are designed for cooling capacities in the order of 10-15 kW per rack (Samadiani *et al.*, 2008). Hence, if datacenters want to take advantage of increasing computer power having a smaller surface footprint, other solutions to air-cooling are required.

One solution is to make use of on-chip cooling. Recent publications show the development of primarily four competing technologies for the cooling of chips: microchannel single-phase flow, porous media flow, jet impingement cooling and microchannel two-phase flow (Agostini *et al.*, 2007). Leonard and Philips (2005) showed that the use of such new technology for cooling of chips could produce savings in energy consumption of over 60%. Agostini *et al.* (2007) highlighted that the most promising of the four technologies was microchannel two-phase cooling.

Single-phase microchannel chip cooling has its apparent advantage in that it is relatively easy to use. This technology is currently being used in the recently started Aquasar (Ganapati, 2009), which is an IBM blade center converted to make use of a water-cooled cycle. However, water has many disadvantages in that it requires a high pumping power to keep temperature gradients on the microprocessor to within acceptable limits. Furthermore, water presents a problem with its high freezing point, microbe formation, corrosive properties, electrical conductivity and erosive nature due to high fluid velocities.

Two-phase microchannel flow takes advantage of the latent heat of the fluid, which is much more effective in removing heat than when using the sensible heat of a single-phase fluid. The latent heat also implies that chip temperatures are much more uniform. It has been shown that heat fluxes as high as 300 W/cm^2 (Park, 2008) can be achieved using a refrigerant. It was further shown that heat fluxes of 180 W/cm^2 can be removed with a saturation temperature as high as 60˚C, while maintaining the chip temperature below 85˚C (Madhour *et al.* 2010). A refrigerant has the advantage of being a dielectric fluid with a long successful history in industrial applications, is inert to most engineering materials, readily available and relatively inexpensive. New refrigerants also have a negligible impact on the environment.

The main advantage of making use of on-chip cooling for the cooling of datacenters is that the heat gained from cooling the chips can be easily reused elsewhere. This is because the heat removal process is local to the server, thus minimizing any losses to the environment, which is prone to traditional air-cooling systems. The opportunity thus exists to reuse the datacenter's waste heat in secondary applications instead of being disposed of into the atmosphere. This not only has the potential to reduce datacenter energy costs, but also to reduce its carbon footprint and hence, its environmental impact.

A problem with recovering the heat of the datacenter is not in the quantity of heat available, but rather in the quality of heat. Currently, heat is being ejected into the atmosphere at temperatures of about 40˚C when using traditional air-cooling methods. This is because chip temperatures are being cooled at 15-20˚C. Due to the effective cooling of chips when using on-chip cooling, fluid approach temperatures of about 60˚C can be realized, while removing high heat fluxes and keeping chip temperatures below 85˚C (Madhour *et al.*, 2010). Although the temperature, and hence the quality of the heat is higher, the applications for using this heat are still limited.

Making use of a vapour compression cooling cycle, where temperatures in excess of 90˚C can be realized, can increase the quality of the heat even further. The potential for this kind of heat not only increases the number of applications that can use it, but also increases its value. Vapour compression cooling cycles for computer chips have been demonstrated successfully (Mongia *et al.*, 2006 and Trutassanawin *et al.*, 2006), although it was focused on a small system where the recovery of heat was not of concern.

The objective of this paper is to demonstrate the potential savings that can be made by implementing two-phase on-chip cooling within a datacenter. Two types of cooling cycle implementations will be investigated; one using a liquid pump as the main fluid driver (potential waste heat temperature of 60˚C) and the other using a vapour compressor as the fluid driver (potential waste heat temperature of up to 100˚C). The investigation will not be limited to the datacenter only, but also to a secondary application of the waste heat. For this paper a coal fired thermal power plant will be analysed since more than 80% of the world's energy is produced from these types of plants. This plant was also chosen as it produces most of the greenhouse gases, with any potential improvements in efficiency translating to savings for both the power plant and the datacenter.

2. CASE STUDY

Marcinichen *et al.* (2010) characterised a vapour compression and a liquid pumping cooling cycle using on-chip cooling with multi-microchannel evaporators. According to the authors the vapour compression cycle is characterized by a high condensing temperature (high heat recovery potential, i.e. high exergy) and a high range of controllability of the micro evaporator inlet subcooling (characteristic of systems with variable speed compressor and electric expansion valve), but it has a high energy consumption. The liquid pumping cycle is characterized as having a low condensing temperature (same as evaporating temperature), although its energy consumption is low.

The case study is performed where it is considered that the datacenter uses on-chip cooling to cool the servers, with the heat captured being redistributed to a power plant. The datacenter will be modeled as a cooling cycle consisting of an evaporator (on-chip cooling elements), a condenser, a fluid driver (compressor or liquid pump), and an expansion valve in the case of a vapour compression system. The power utility will be a thermal Rankine cycle consisting of a boiler, a high and low pressure turbine, a condenser, a low pressure and high pressure feedwater pump and a feedwater heater. The feedwater heater receives heat from steam tapped after the high pressure turbine. The optimal pressure for tapping the steam will be calculated to obtain maximum thermal efficiency. The datacenter waste heat will be injected into the Rankine cycle after the condenser and prior to the feedwater heater. This is shown schematically in Figure 1.

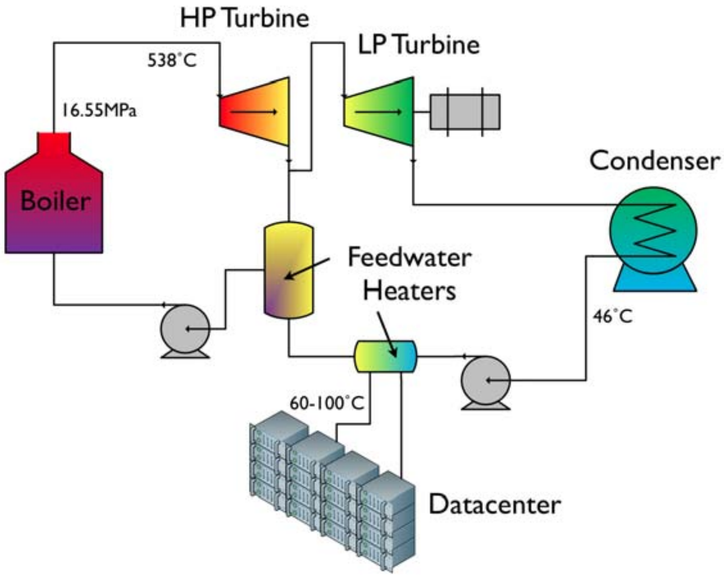

Figure 1. Datacenter integrated in a power utility.

The operating conditions for both cycles are listed in Table 1. A first-law analysis is performed on the two cycles, showing their overall performances. The following assumptions are made:

- no pressure drop in components;
- isentropic compression;
- isentropic pumping;
- isenthalpic expansion;
- 100% exchanger efficiency.

Table 1. Operating conditions for the power utility and datacenter cooling cycle

	Power Utility	***Datacenter Cooling***
Boiling/Evaporation Pressure	16.55 MPa	1.681 MPa
Turbine/Compressor		
Inlet Pressure	16.55 MPa	1.681 MPa
Inlet Temperature	538˚C	70˚C
Condensing Temperature	46˚C	60 – 100˚C

2.1. Power Utility

Heat captured in the datacenter can be reused by a power plant. Since the waste heat of the datacenter is of a low quality, it can only be inserted after the condenser of the power plant. This would then increase the temperature of the water leaving the condenser (46˚C typically) to a maximum temperature as defined by the condensing temperature of the datacenter cycle. Therefore, any additional heat added to the power plant's cycle will result in less fuel needing to be burnt, thus saving fuel and reducing the CO_2 footprint of the power plant.

Figure 2 shows the efficiency improvement of the power plant as a function of the datacenter cycle's condensing temperature. The figure shows that the higher the condensing temperature, the greater the efficiency improvements. The efficiency of the plant can be improved by up to 2.2% if the datacenter's waste heat is reused in the power plant. By using a liquid pumping cycle in the datacenter, condensing temperatures of 60˚C can be reached since this would imply that the evaporating temperature on the chip is also 60˚C. For higher condensing temperatures, a vapour compression cycle would be required.

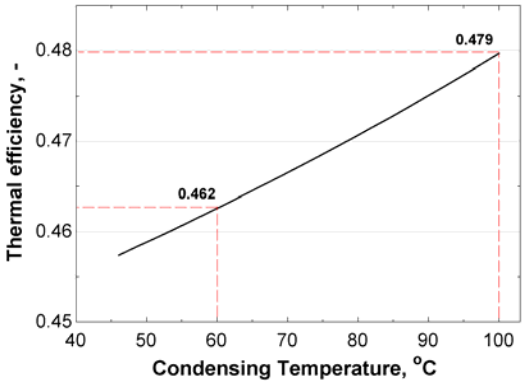

Figure 2. Thermal efficiency of power plant

In terms of CO_2 footprint, Figure 3 shows the reduction in the amount of CO_2 per kilowatt-hour output per year as a function of the datacenter condensing temperature. Also shown on the graph is the amount of CO_2 saved per kilowatt-hour output per year. These values assume that coal is used as the source of the power plant's energy and will be less for other types of fuel. Therefore, if a power plant with an output capacity of 500MW were considered, the savings in CO_2 would be in the order of 195 000 tons of CO_2 per year. This could potentially save about \$ 3 000 000 per year if a carbon tax of \$15 per ton of CO_2 (Larson, 2009) was considered.

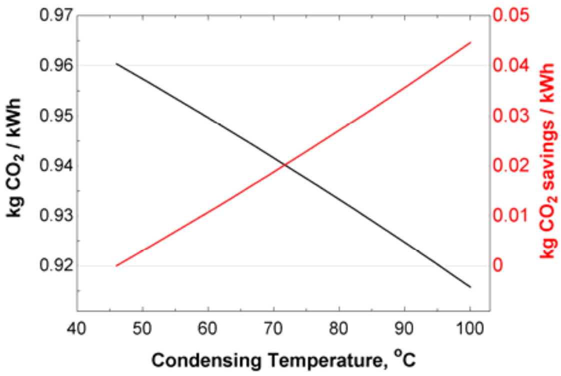

Figure 3. CO_2 footprint and savings per kilowatt-hour

2.2. Datacenter

Due to datacenter growth, it is more than likely to see datacenters in excess of 100 000 servers (*viz.* Figure 4). Therefore, the simulations following will be based on a datacenter containing 100 000 servers, with each server generating 325 W, which includes the main processor and auxiliary electronics (memories, DC-DC converters, hard drives, etc). Figure 5 shows a graph of the total power supply required by the datacenter to operate the IT equipment and the cooling system. Included in the graph is on-chip two-phase cooling methods using a vapour compression cycle and liquid pumping cycle. As a comparison, traditional air cooling units are also simulated, where it is assumed that their power consumption is the same as that required to operate the IT equipment (Koomey, 2007 and Ishimine *et al.*, 2009). The results are plotted as a function of the compressor or pump overall efficiencies (defined as the ratio between isentropic compression or pumping power and the electrical power for the drivers). All simulations were developed for an evaporating temperature of 60˚C (on the chip), with condensing temperatures being 60˚C and 90˚C for the liquid pumping and vapour compression cycles, respectively.

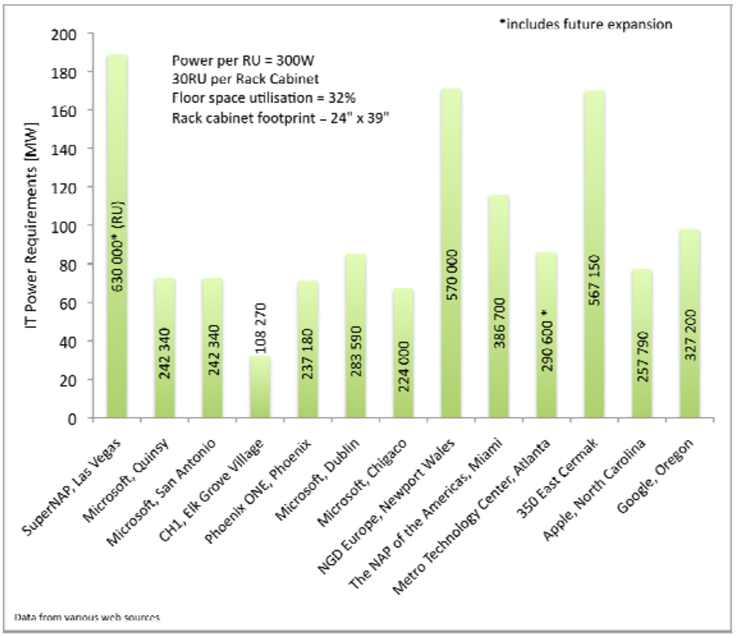

Figure 4. Datacenter size and information technology (IT) power requirements.

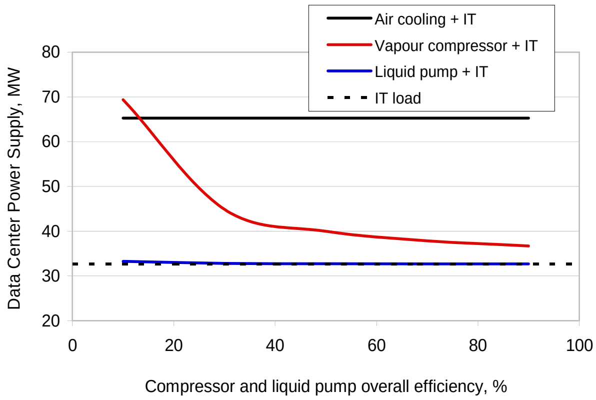

Figure 5. Datacenter power supply - IT and cooling system.

It is seen that the cycle using a vapour compressor has a strong function of the compressor's overall efficiency, up to a value of approximately 35% after which it becomes less dependent. Typically, compressors have an overall efficiency of about 60%. The liquid pumping cycle hardly shows any dependence on the pump efficiency. This is due to the power required to drive the pump being very low (Marcinichen and Thome, 2010). The higher power consumption of the compressor is due to the energy required to increase the pressure from a saturation temperature of 60˚C to 90˚C. What is noticed, though, is that the datacenter's power requirements are reduced considerably for overall efficiencies above 15% (vapour compression cycle) when compared to traditional air cooling. This reduction is in the order of 50% when using a liquid pumping cycle and 41% for a vapour compression cycle with a compressor having an overall efficiency of 60%.

Further savings can be made when energy recovery is considered. Since on-chip cooling is used, recovering this energy would be a simple process by just incorporating a condenser, where the energy absorbed by the fluid from the server is transferred to a secondary fluid, like water, in the condenser. Figure 6 shows a graph of the datacenter power supply for the three types of cooling technologies as a function of the efficiency with which energy is recovered. 100% efficiency implies that all the heat generated by the servers is recovered, while 0% means that none of the heat is recovered. Note that there is no change in power consumption for air cooling as it was assumed that the heat was not recovered, although this technology does exist, albeit not as effective as for liquid cooling. The plots for the liquid pumping and vapour compression cycles assume that the pump and compressor has an overall efficiency of 100% and 60%, respectively, although the choice of efficiency for the pump is negligible (*viz.* Figure 5). It should be noted that 0 MW datacenter supply does not mean the datacenter requires no power, but rather that all the power received as electricity is sold as heat. The financial implications would show this since the value of the heat sold would be different from the electricity purchased and will be a function on the application to which the heat is sold.

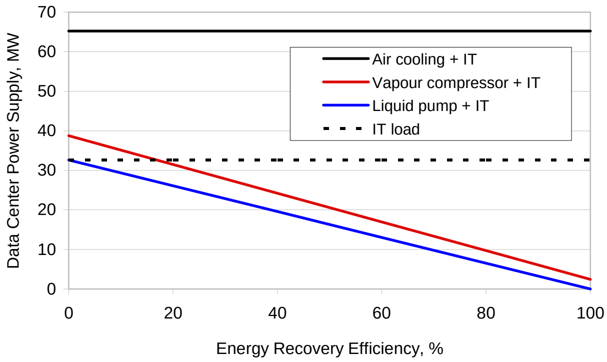

Figure 6. Datacenter power supply - IT and cooling system - Potential of energy recovery.

Figure 6 shows that practically all the power purchased in the form of electricity can potentially be sold as heat. This is especially the case for a liquid pumping cycle due to its energy consumption being very low. The vapour compression cycle always has some unrecoverable heat due to inefficiencies, with this heat being lost to the environment. At this point in stage it would appear as if the use of a liquid pumping cycle far outweighs the vapour compression cycle. However, what the graphs do not show is the quality of the heat being sold. The former can sell heat at a temperature of 60˚C, while the latter can sell it at 90˚C. The quality of heat is therefore important, not only due to the monetary value it adds, but also to the application to which it is sold. A limited number of applications can use 60˚C of heat, with the limits becoming less as the temperature is increased.

2.3. Carbon Footprint

For the calculation of the carbon footprint, only the contribution of the electricity used is considered. The effect of greenhouse gases (GHG) being formed by the manufacturing, transporting, storage and disposal of the components of the datacenter, as well as the datacenter building, fall under the Life Cycle Assessment, which falls outside the scope of the current paper. Further, of the greenhouse gases, only CO_2 will be considered as it contributes to more than 75% of all the greenhouse gases and contributes the most to the greenhouse effect. Figure 7 shows the reduction of the quantity of CO_2 for the three cooling technologies as a function of the efficiency with which the energy is recovered. The quantity of CO_2 is calculated with the assumption that the datacenter purchases its electricity from a power plant running on coal and that it is selling waste heat back to the power plant, as discussed earlier. This graph, therefore, takes into consideration the efficiency increase of the power plant, since the amount of CO_2 released is a function of the power plant's efficiency, which in itself is a function of the efficiency with which energy is recovered.

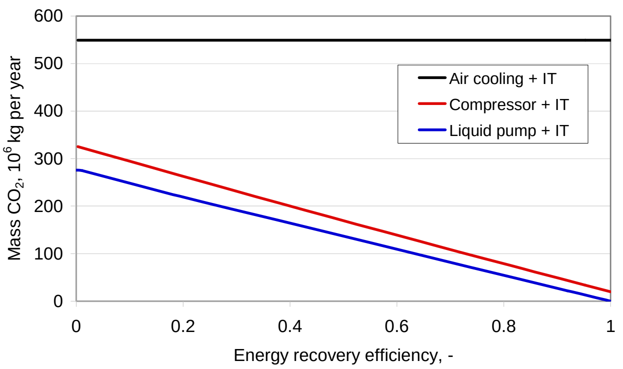

Figure 7. Reduction in CO_2 for three datacenter cooling technologies.

The figure shows that the datacenter could potentially have a zero carbon footprint regarding electricity usage when on-chip cooling using a liquid pumping or vapour compression cycle is used. The use of a liquid pumping or vapour compression cooling cycle without energy recovery (0% recovery efficiency), compared to traditional air-cooling, reduces the carbon footprint of the datacenter considerably, with a reduction of 50% for the former and 40% for the latter. This reduction is improved further with energy recovery, with a potential reduction of almost 100% and 96% being achievable.

As observed previously, Figure 7 does not show the quality of the heat being recovered, with the value and applicability of this heat not being clear. This can be seen more clearly when applying the waste heat to the thermal power plant, which is shown in Figure 8. This graph shows the CO_2 reduction of the datacenter due to energy recovery and the savings in CO_2 of the power plant due to efficiency improvements. Instead of plotting the carbon footprint as a function of the recovery efficiency it is plotted as a function of the condenser temperature, which is directly linked to the feedwater heater of the power plant. The effect is the same since a lower energy recovery efficiency would result in lower temperature increases of the power plant's feedwater and, hence, lower efficiency increases. The graph therefore shows the limit of each cooling cycle. Since the power plant's thermal efficiency improves with increase in condenser temperature, the amount of CO_2 saved by the power plant increases. However, when using a liquid pumping cycle for the datacenter, only 25% of the total potential savings in CO_2 can be achieved, amounting to approximately 17 000 tons per year for a 173 MW power plant. By making use of a vapour compression cycle, however, the potential savings in CO_2 can reach as high as 70 000 tons per year. Therefore, although the liquid pumping cycle was the better performing cooling cycle regarding energy usage and CO_2 reduction, due to the higher temperatures achievable the vapour compression cycle has a larger impact on the secondary application making use of the waste heat.

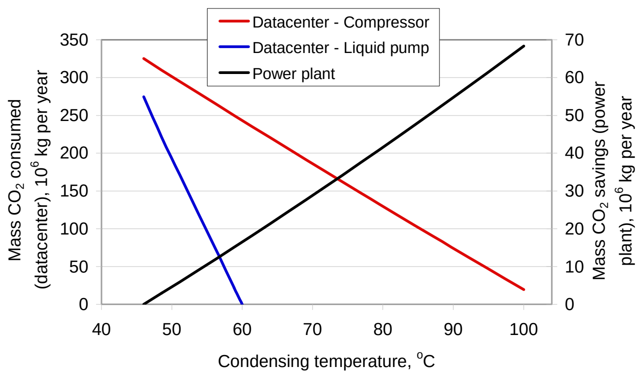

Figure 8. Carbon footprint of datacenter and CO_2 savings of power plant.

Of course, depending on the point of view taken, it may be argued that due to the savings the power plant is making in CO_2 due to the datacenter, the datacenter could claim those savings as part of its own reduction. This could potentially then lead to the datacenter having a negative carbon footprint regarding energy usage, which it can then use as a carbon offset, as shown in Figure 9. This offset could be used to compensate for other CO_2 emitting processes in the datacenter, or can even be sold to other organizations as is done in carbon-trading. This should, however, be viewed as a tentative idea as regulations would determine whether this is possible or not. However, if this were the case, by making use of a vapour compression cycle, the offset could be 167% more than when using a liquid pumping cycle.

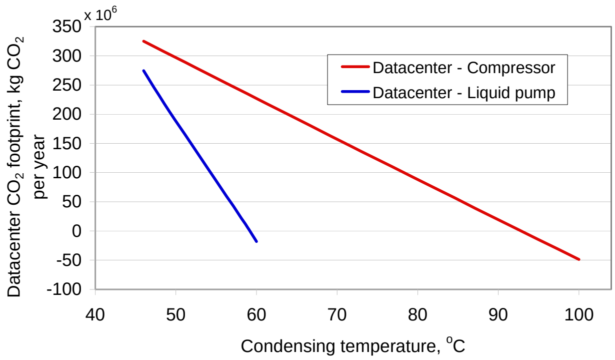

Figure 9. CO_2 footprint of a datacenter considering CO_2 reduction of power plant

2.4. Monetary Savings

Global warming is having a huge impact on the environment and on the livelihood of people regarding food production and natural resources. To counter this, a carbon tax is being introduced, which is aimed at helping the environment by not only reducing carbon emissions by forcing people and organizations to become more energy efficient, but also by raising funds to be used for clean energy research. The tax is levied on the carbon content of fuels, increasing the competitiveness of non-carbon technologies such as solar, wind or nuclear energy sources. Therefore, organizations using electricity produced from the burning of fossil fuels will pay a higher tax than those produced from non-carbon burning fuels. The probability also exists for taxing the utility generating the electricity. Carbon taxes have only been introduced in a few countries, with most European countries taking the lead, even though the way organizations are being taxed vary from country to country. In the United States the introduction of carbon tax has been made in California and the city of Boulder, Colorado, with taxes being in the order of 4 cents/ton of CO_2. European countries have been much more stringent with taxes in some countries, such as Sweden, being as high as $100 per ton of CO_2 (Carbon Tax, 2010). The Larson Bill (Larson, 2009) proposes to introduce a nationwide tax (US) of $15/ton CO_2 starting in 2012, increasing by $10/ton CO_2 every year. It also proposes to increase this increment to $15/ton CO_2 after 5 years if the US emissions stray from the Environmental Protection Agency's (EPA) glide-path prediction, which proposes to cut emission to 80% that of 2005 levels by 2050.

With a recommended price of $30/ton CO_2 (Nordhaus, 2008) could cost industries millions if efficiencies are not improved. Datacenters are also not exempt from these taxes, which will be introduced in the following years to come (Mitchel, 2010). Figure 10 shows the potential savings made by a datacenter with a size of 100 000 servers and a coal power plant with a size of 175MW if heat were captured from the datacenter and sold to the power plant. The savings not only includes that saved in energy costs by implementing a liquid pumping or vapour compression cycle instead of a traditional air cooling cycle, but also that saved in carbon tax. The savings of the power plant is in terms of fuel saved and the savings made in carbon tax. For fuel costs, a value of $90/ton of coal was used.

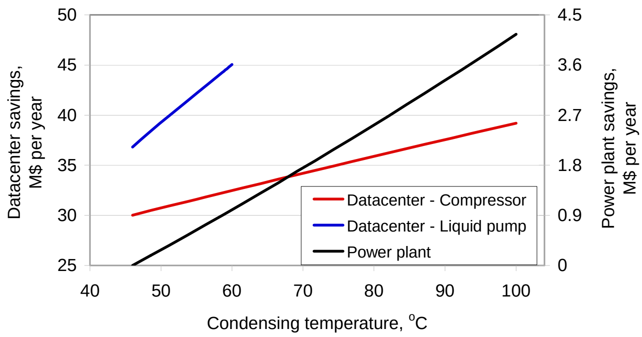

Figure 10. Datacenter and power plant savings

For the datacenter, the graph shows that most savings are made if a liquid pumping cycle was used, with the potential savings being in the order of $45 000 000 per year, while a vapour compression cycle would save in the order of $40 000 000 per year. The power plant, when recovering heat from a datacenter using a liquid pumping cycle, will only save about $1 000 000 per year, with savings reaching almost $4 500 000 a year if a vapour compression cycle were used.

The overall savings are given in Figure 11. This graph shows that the total savings, if a datacenter was to sell waste heat at 60˚C with a liquid pumping cycle, could be about $46 000 000 a year, while selling heat at 90˚C with a vapour compression cycle could save a total of $43 000 000 per year. These are savings that a customer would potentially not have to pay. Even though it seems as if the clear solution is to use a liquid pumping cycle within a datacenter, the incentive for a power plant to cooperate with a datacenter would be greater if a vapour compression cycle was used.

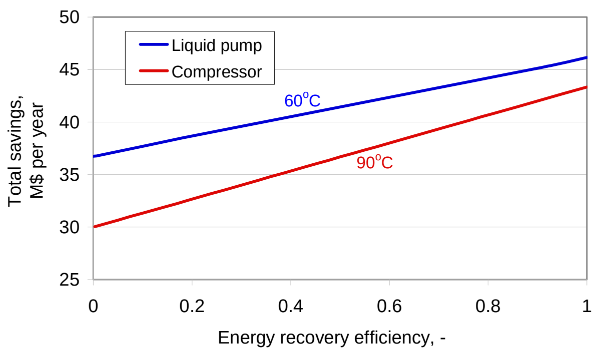

Figure 11. Total savings for liquid pumping and vapour compression cooling technologies.

3. CONCLUSION

This paper investigated the potential savings in energy a datacenter can make by implementing on-chip cooling with waste heat recovery as a replacement to traditional air-cooling. The investigation considered on-chip cooling cycles making use of a liquid pump and a vapour compressor as the main fluid driver. As an application for the waste heat, a coal-fired power plant was analyzed. Aspects such as energy consumption, energy recovery, carbon footprint and power plant efficiency were investigated.

The results showed that, when compared with traditional air-cooling systems, the energy consumption of the datacenter, without considering energy recovery, can be reduced by as much as 50% when using a liquid pumping cycle and 41% when using a vapour compression cycle. With the energy recovered, though, this value can be reduced even further. By applying this reused energy to a thermal power plant one can see thermal efficiency increases in the order of 2.2%.

By implementing on-chip cooling, the datacenter would see a 50% reduction of its carbon footprint. With the reuse of the waste heat, this footprint could tend to zero. With the efficiency improvements of the power plant, the carbon footprint of a 500MW plant can be reduced by as much as 195 000 tons of CO_2 per year. It could be argued that, since this reduction was due to the datacenter, the datacenter could claim the reduction for itself. This could imply that the datacenter will have a negative carbon footprint, which could potentially be used as a carbon offset.

The overall savings made, considering the datacenter and power plant as a whole, could be as much as $46 000 000 when using a liquid pumping cycle and $43 000 000 when using a vapour compression cycle. These savings include the reduction in energy usage, reduction in fuel usage and savings made due to reduced carbon emissions when considering carbon taxes. Though it appears as if most of the savings are made when using a liquid pumping cycle, its applicability is limited due to its low quality of energy. The incentive for a secondary application to use the waste heat will be greater when a vapour compression cycle is used due to the quality of its energy being higher.

4. REFERENCES

Agostini, B., Fabbri, M., Park, J.E., Wojtan, L., Thome, J.R. and Michel, B., 2007. "State-of-the-art of High Heat Flux Cooling Technologies", Heat Transfer Engineering, 28: p. 258-281.

Carbon Tax, 2010. In Wikipedia, The Free Encyclopedia. Retrieved 08:12, July 26, 2010, from http://en.wikipedia.org/w/index.php?title=Carbon_tax&oldid=374709823.

EPA, 2007. "Report to Congress on Server and Data Center Energy Efficiency Public Law", 109-431, U.S. Environmental Protection Agency.

Ganapati, P., 2009. "Water-Cooled Supercomputer Doubles as Dorm Space Heater", Available from: http://www.wired.com/gadgetlab/2009/06/ibm-supercomputer/.

IBM, 2010. Available from: ftp://public.dhe.ibm.com/common/ssi/ecm/en/blb03002usen/BLB03002USEN.PDF.

Ishimine, J., Ohba, Y., Ikeda, S., Suzuki, M., 2009. "Improving IDC Cooling and Air Conditioning Efficiency", FUJITSU Sci. Tech. J., Vol. 45, pp. 123-133.

Koomey, J.G., 2007. "Estimating Total Power Consumption by Servers in the U.S. and the World", Analytics Press, Oakland, CA, February 15 (http://enterprise.amd.com/us-en/AMD-Business/Technology-Home/Power-Management.aspx).

Larson J.B., 2009. "America's Energy Security Trust Fund Act of 2009", H.R. 1337.

Leonard, P.L. and Phillips, A.L., 2005. "The Thermal Bus Opportunity – A Quantum Leap in Data Center Cooling Potential", in ASHRAE Transactions, Denver, CO.

Madhour, Y., Olivier, J.A., Costa-Patry, E., Paredes, S., Michel, B. and Thome, J.R., 2010. "Flow Boiling of R134a in a Multi-Microchannel Heat Sink with Hotspot Heaters for Energy-Efficient Microelectronic CPU Cooling Applications", IEEE Transactions on Components and Packaging Technologies, Accepted for publication.

Marcinichen, J.B and Thome, J.R., 2010. "New Novel Green Computer Two-Phase Cooling Cycle: A Model for Its Steady-State Simulation", in Proceedings of the 23rd International Conference on Efficiency, Cost, Optimization, Simulation and Environmental Impact of Energy Systems - ECOS2010, Lausanne, Switzerland.

Marcinichen, J.B., Thome, J.R. and Michel, B., 2010. "Cooling of Microprocessors with Micro-Evaporation: A Novel Two-Phase Cooling Cycle", International Journal of Refrigeration.

Mitchel, R.L., 2010. "Carbon Tax could Whack Data Centers", Computerworld, http://blogs.computerworld.com/16128/carbon_tax_could_whack_data_centers\

Mongia, R., Masahiro, K., DiStefano, E., Barry, J., Chen, W., Izenson, M., Possamai, F., Zimmermann, A. and Mochizuki, M., 2006. "Small Scale Refrigeration System for Electronics Cooling within a Notebook Computer", in ITHERM 2006, San Diego, CA.

Nordhaus, W., 2008. "A Question of Balance – Weighing the Options of Global Warming Policies", Yale University Press, New Haven.

Park, J.E., 2008. "Critical Heat Flux in Multi-Microchannel Copper Elements with Low Pressure Refrigerants", PhD Thesis, École Polytechnique Fédérale de Lausanne, Switzerland.

Saini, M. and Webb, R.L., 2003. "Heat Rejection Limits of Air Cooled Plane Fin Heat Sinks for Computer Cooling", IEEE Transactions on Components and Packaging Technologies, 26(1): p. 71-79.

Samadiani, E., Joshi, S. and Mistree, F., 2008. "The Thermal Design of a Next Generation Data Center: A Conceptual Exposition", Journal of Electronic Packing, 130: p. 041104-1 - 041104-8.

Trutassanawin, S., Groll, E.A., Garimella, S.V. and Cremaschi, L., 2006. "Experimental Investigation of a Miniature-Scale Refrigeration System for Electronics Cooling", IEEE Transactions on Components and Packaging Technologies, 2006. 29(3): p. 678-687.