UNIVERSITY OF LEEDS

This is a repository copy of *An Experimental and Theoretical Investigation of the Effects of Supply Air Conditions on Computational Efficiency in Data Centers Employing Aisle Containment*.

White Rose Research Online URL for this paper:
http://eprints.whiterose.ac.uk/126520/

Version: Accepted Version

Proceedings Paper:

Tatchell-Evans, M, Burdett, D, Summers, J orcid.org/0000-0001-8266-5038 et al. (2 more authors) (2017) An Experimental and Theoretical Investigation of the Effects of Supply Air Conditions on Computational Efficiency in Data Centers Employing Aisle Containment. In: 33rd Thermal Measurement, Modeling & Management Symposium (SEMI-THERM 2017). SEMI-THERM 2017, 13-17 Mar 2017, San Jose, Calif., USA. IEEE , pp. 100-107. ISBN 978-1-5386-1531-7

https://doi.org/10.1109/SEMI-THERM.2017.7896915

(c) 2017, IEEE. This is an author produced version of a paper published in the proceedings of the 33rd Thermal Measurement, Modeling & Management Symposium. Uploaded in accordance with the publisher's self-archiving policy. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Reuse
Unless indicated otherwise, fulltext items are protected by copyright with all rights reserved. The copyright exception in section 29 of the Copyright, Designs and Patents Act 1988 allows the making of a single copy solely for the purpose of non-commercial research or private study within the limits of fair dealing. The publisher or other rights-holder may allow further reproduction and re-use of this version - refer to the White Rose Research Online record for this item. Where records identify the publisher as the copyright holder, users can verify any specific terms of use on the publisher's website.

Takedown
If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

White Rose
university consortium
Universities of Leeds, Sheffield & York

eprints@whiterose.ac.uk
https://eprints.whiterose.ac.uk/

An Experimental and Theoretical Investigation of the Effects of Supply Air Conditions on Computational Efficiency in Data Centers Employing Aisle Containment

Morgan Tatchell-Evans[1], Daniel Burdett[1], Jonathan Summers[1], Adam Beaumont[2], Geoff Fox[3]
[1]Institute of Thermofluids, School of Mechanical Engineering, University of Leeds, LEEDS LS2 9JT, United Kingdom, pm08mrte@leeds.ac.uk
[2]aql, 13-15 Hunslet Road, LEEDS LS10 1JQ, United Kingdom
[3]Digiplex Stockholm 1 AB, Smedbyvägen 6, 194 30 Upplands Väsby, Sweden

Abstract

Aisle containment is increasingly common in data centres, and is widely believed to improve efficiency and effectiveness of cooling. Investigations into the impacts of aisle containment on the behavior and power consumption of cooling infrastructure and servers have been limited. Nor has the impact of supply air conditions on these factors been extensively investigated. This work uses measurements of bypass in a test data centre and observations on server behavior in a wind tunnel, in conjunction with a system model, to investigate the efficiency with which computations can be undertaken in an aisle contained data centre, and how this is impacted by supply air conditions.

Keywords

Data centers, computational efficiency, aisle containment, bypass.

1. Introduction

Data centers (DCs) are energy intensive facilities which are estimated to account for 1.4% of global electricity consumption [1]. As demand for digital services expands, the sector's electricity consumption increases both in absolute terms and as a proportion of global consumption, trends which look set to continue [1], [2]. Accordingly, there is growing political and academic interest in improving energy efficiency in the sector [3]–[6].

Much of the research into DC energy efficiency focuses on cooling [6], which typically accounts for 20-50% of a DC's total electricity consumption, E_T [7]. Imperfect air management is a major barrier to efficient cooling. Specifically, poor air management leads to:

- **bypass** - the tendency for cold, supply air streams to bypass server inlets and return to the air conditioning units, and
- **recirculation** - the tendency for hot air from the server outlets to mix with cold, supply air streams and re-enter the server inlets.

Such behavior impairs the effectiveness of cooling by raising server inlet temperatures, and also leads to increased cooling energy consumption through the need to increase supply air flow rates and reduce supply air temperatures in order to achieve sufficiently low server inlet temperatures [8]–[12].

One increasingly common method for reducing bypass and recirculation is to install physical barriers separating the hot and cold air streams, known as aisle containment systems [9], [13], [14]. Experimental and numerical investigations have shown this approach to be effective in reducing electricity consumption and improving cooling effectiveness [8], [9], [15]–[18]. In DCs employing aisle containment, a positive pressure differential, Δp_{CH}, is usually maintained between the cold and hot aisles, in order to prevent recirculation [9] [17]. This tends to drive air through any gaps present in the containment system or server racks, hence bypass is not completely prevented by containment [9], [17], [18].

Academic investigations into the extent of bypass within DCs employing aisle containment and the potential to mitigate against it are limited. Numerical investigations have either neglected the potential for bypass within server racks (i.e. cold air passing through server racks whilst bypassing the server inlets) [8], [13], [15], [16], [19], or have incorporated this phenomenon without experimental calibration [17] [20]. One experimental study investigated the impact of aisle containment on cooling, but did not measure bypass within racks, effectively assuming that all air flow through racks passed through the servers [9]. A recent study by Tatchell-Evans et al. [18] showed that significant levels of bypass occur within the racks. The study also provided the first experimental investigation of the impact of Δp_{CH} on bypass [18]. However, this study only presented results of flow rates through racks containing no servers, with blanking panels being installed to prevent flow through the regions normally occupied by servers. Hence the impact of server fans on bypass flow rates was not investigated [18].

The impacts of bypass and recirculation on the operation of the DC depend on the responses of the various components of the DC to changing flow conditions. Various system models have been described in the research literature which seek to predict the impact of design and operation variables on DC power consumption and cooling effectiveness [12], [18], [21]–[26]. In addition, Zapater et al. [27] have investigated the impact of server fan control algorithms and supply temperature (T_{supply}) on computational efficiency of the server. Only Tatchell-Evans et al. have used a system model to investigate the impact of bypass on E_T [18]. This model used experimental data to determine the extent of bypass within the DC at different levels of Δp_{CH}. The impact of Δp_{CH} on server flow rates was based on experimental results undertaken with a server which was switched off, with a range of assumptions made about the response of server fans to changing Δp_{CH} and the resultant effect on power consumption. The temperature rise across the server was assumed to remain constant. The effect of supply flow conditions on computational efficiency was not considered in the model. It was estimated that measures undertaken to reduce bypass could reduce E_T by up to 8.8%, and that adjusting Δp_{CH} could reduce E_T by up to 16%. The impact of these factors on E_T was found to be affected strongly by the assumptions made about server fan behavior.

No system model has been described in the literature which can investigate the impact of supply air flow conditions on computing efficiency of the whole DC system. This is important since the true efficiency of the DC can only be measured if the DC's output, i.e. computation, is measured.

The bulk of research into DC efficiency focuses on the savings that can be made by looking at cooling at the room-scale, with relatively little research investigating the effect of temperature and pressure variations on computational efficiency in the servers. In this context, the ratio of E_T to the IT power consumption is commonly used as a metric for efficiency by DC managers, and is termed the Power Usage Effectiveness (PUE) [28]. However, the PUE does not consider the useful work done by the DC or the effect that variations in supply air conditions can have on the efficiency of the servers doing this work. Hence PUE does not measure the true efficiency of the DC, only the efficiency with which heat is rejected [29], [30].

Inside a server, electrical energy is consumed by a range of components, including processors, fans, DRAM and hard drives, but perhaps the most important, when considered from both a computational and thermodynamic perspective, is the CPU. This is where the processing of digital information by the DC is facilitated, and is also the greatest source of heat, due to electrical impedance in the increasingly densely populated microelectronics. These components are highly sensitive to the thermal environment and require constant cooling without which they would exceed their thermal envelopes, causing increased power consumption, failures and warranty breaches associated with exceeding temperature limits and loss of service [27].

DCs utilize a wide range of servers with competing architectures, and performance with server selection dependent on requirements, budget, and application. This means that effectively judging the degree to which performance varies with supply air conditions for any one DC is a complex problem, leaving most DC operators to simply adhere to generic guidelines provided by ASHRAE [30] for thermal envelope to avoid damaging components instead of considering computational efficiency.

2. Aims and objectives

This paper aims to further the existing knowledge on the variation of computational efficiency with Δp_{CH} and T_{supply} in DCs employing aisle containment.

This paper builds on work by Tatchell-Evans et al. [18] which used experimental results quantifying the extent of bypass within a test data center employing cold aisle containment. This data was used as an input to a system model which estimated the impact of Δp_{CH} on total facility power consumption (E_T).

This work has been expanded upon to utilize results from tests of a server in a wind tunnel, in which the power consumption of, flow rate through and computational performance of the server were investigated as functions of Δp_{CH} and T_{supply}. The model presented in [18] relied on making a range of assumptions about server flow rates and power consumptions, and did not consider the computational performance.

The updated system model predicts computational output and E_T as functions of T_{supply} and Δp_{CH}. This facilitates the quantification of the overall computational efficiency of the DC, i.e. the ratio between computational output and E_T. By contrast, the work presented in [18] effectively investigated only the efficiency of heat rejection.

3. Material and Methods

3.1. Bypass and recirculation in aisle contained data centers

The experimental work undertaken by Tatchell-Evans et al [18] to quantify the extent of bypass within aisle contained DCs has been extended, using similar methods. Specifically, tests were undertaken in the test data center described in [18], which houses a row of 4 racks, with the hot and cold aisles separated by a solid partition. The experiments used commercially available racks and aisle containment systems. Air was supplied to the cold aisle through a duct with the flow rate through the duct, V_{supply}, determined by taking a series of velocity measurements across a section of the duct using a hot wire anemometer. Velocity was assumed to vary linearly between measurement points, and was assumed to be zero at the walls of the duct. Velocity close to the duct walls was assumed to follow Karman’s conventional law for the variation of the fluid velocities in the peripheral zone. V_{server} was calculated by integrating between the measurement points, and between the measurement points and the wall. Δp_{CH} was measured using a manometer. A schematic of the test data center is shown in Figure 1, and the experimental methods are more fully described in [18].

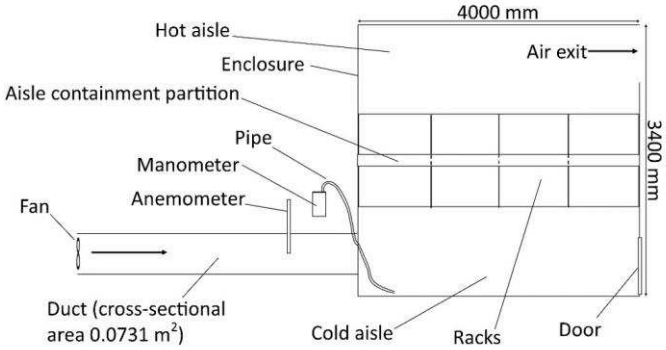

Figure 1: Plan view schematic of test data center, fully described in [18].

Hillstone HAC230-6R load banks (LBs) [31] were installed in the racks, and used to mimic servers in producing a heat load and driving flow from the cold to the hot aisle using their internal fans. Unoccupied slots within the racks were fitted with blanking panels, and efforts were made to seal potential paths for bypass and recirculation, to the extent that was deemed achievable in an operational DC setting. The intention of the tests undertaken with LBs installed was to determine the impact of fan-containing equipment within the racks on bypass and recirculation. The range of Δp_{CH} investigated was also expanded, to include the -5 to 2 Pa range, in addition to the range of 2 to 20 Pa which was studied in [18]. This allows the system model to make predictions of behavior where V_{supply} is less than the flow through the servers, V_{server}. Further tests were undertaken with a single 3U Sun Fire v40z server [32] [33] installed in the test data center, to investigate the extent to which bypass occurs at the interface between the server and the adjacent blanking panels and equipment rails.

In tests with LBs or servers installed, the flow rate through the server/LB ($V_{server/LB}$) was measured by attaching a rectangular duct to the rear of the server/LB, and taking 9 velocity measurements across the section of the duct using the anemometer. V_{server} was calculated from the velocity measurements in the same way as V_{supply}. The duct had a cross-sectional area of 0.016 m^2 at the point at which the measurements were taken. This measurement allowed bypass to be calculated according to $V_{bypass}=V_{supply}-V_{server/LB}$.

3.2. Generic Server Wind Tunnel tests

The effect of T_{supply} and Δp_{CH} on computational efficiency was studied utilizing the Generic Server Wind Tunnel (GSWT) situated in house at the University of Leeds (depicted in Figure 2). The GSWT enables servers to be stressed using a range of benchmarks, including SPECPower [34], Stresslinux [35], and a dynamic webserver simulation developed in house, while V_{server} and T_{supply} are controlled. Computational efficiency is determined by monitoring and logging server operations, CPU loading, and/or network requests compared with socket power consumption, with upstream and downstream temperatures, pressure drop across the server (Δp_{CH}), and flow rate in the wind tunnel also being recorded.

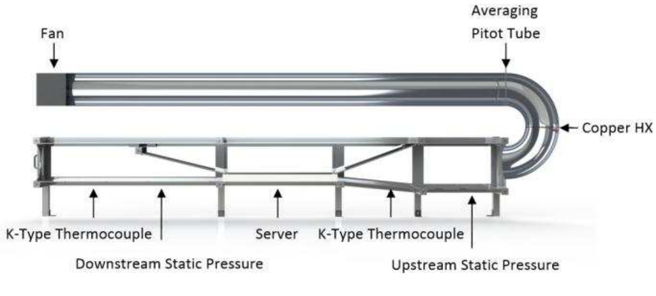

Figure 2: Schematic of GSWT tests.

The tests reported in this work were conducted on a 1U Sun Fire v20z server [36]. The server was stressed using SPECpower2008 [34] to simulate a CPU loading of 100%, whilst monitoring floating point operations per second (FLOPS), power consumption, and thermal environment for the duration of the test. The computational loading is achieved by an external control system sending a combination of different requests to the server, at varying frequencies and for varying durations. The tests were performed with T_{supply} ranging from 23°C to 28°C, and with Δp_{CH} ranging from 0 to 20 Pa. Power consumption was monitored by a script built into SPECpower which interfaced with a Voltech power analyzer. Temperatures were recorded using K-type thermocouples, and logged using a PICO TC-08 thermal datalogger. Δp_{CH} was determined as the difference between static pressure measurements taken upstream and downstream of the server, using a Digitron 2080p digital manometer [37].

The variation of flow rate through the server with Δp_{CH} was measured in the test data center described in section 3.1. The server was installed into one of the server racks, with no other equipment installed. The server inlets were initially sealed, and the relationship between V_{supply} and Δp_{CH} determined. The server inlets were then opened, and the CPU loaded as in the GSWT experiments. The new relationship between V_{supply} and Δp_{CH} was then determined. This allowed V_{server} to be deduced as a function of Δp_{CH}, by deducting V_{supply} as measured with the server inlet sealed from V_{supply} with the server inlet open. Note that measuring V_{server} by taking air velocity measurements in a duct attached to the rear of the server would not be appropriate, since such a duct would add resistance to the flow through the server.

3.3. The system model

The system model described in [18] has been modified to incorporate the findings of the experiments described in section 3.2. A schematic of the model is shown in Figure 3. The model assumes a DC employing aisle containment, with servers cooled using a closed loop of process air, which rejects heat to a chilled water loop in a Computer Room Air Handler (CRAH). The water is cooled in an economizer using ambient air, which is assumed to be at a temperature of 11°C, the average temperature for London [38]. An outline of the workings of the original model is given in the following, alongside a description of the changes made. A full description of the original model is given in [18].

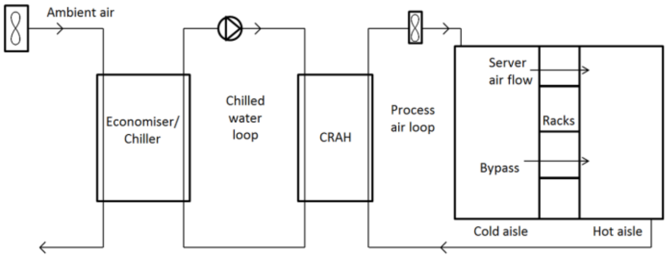

Figure 3: Schematic of the system model.

The model calculates the total power consumption of the DC (E_T) as the sum of the power consumptions of the servers ($E_{servers}$), the CRAH fans (E_{CRAH}), the chilled water pumps (E_{CW}) and the economizer fans (E_{eco}). E_{CRAH}, E_{eco} and E_{CW} are assumed to scale with the cube of the respective fan and pump flow rates, with reference flow rates and power consumptions based on manufacturer specifications.

The variation of heat transfer rates in the CRAH and economizer with air and water flow rates and temperatures are determined using empirically derived formulae from [39], [40] and [41], having first obtained reference heat transfers, fluid temperatures and flow rates from manufacturer data.

The relationship between V_{bypass} and Δp_{CH} was set to be consistent with the results of experiments described in section 3.1 and in [18]. Two scenarios were used: (i) a 'low bypass' scenario based on results from the test data center, in which bypass was minimized as far as reasonably practical, and (ii) a 'high bypass' scenario based on results of tests undertaken on a single rack in isolation, without paying special attention to minimizing bypass.

The behavior of the servers in the model has been modified so that V_{server} and $E_{servers}$ vary with T_{supply} and Δp_{CH} in a way which is consistent with the results of the experiments described in section 3.2. The DC being modelled was assumed to contain enough servers such that the IT power consumption was 12 kW per rack.

The model assumes that there is no recirculation, i.e. that the temperature of air supplied by the CRAH is equal to the temperature at the server inlets.

4. Results

4.1. Test data center experiments

The results of the original experiments undertaken in the test data center are available in [18]. The results of the new experiments in the test data center are presented in this section. V_{bypass} is presented as a percentage of the total flow required for cooling, with the flow rate through the servers calculated assuming a typical IT power load of 12 kW per rack [42] and temperature rise across the servers, ΔT, of 12.5 K [12] [43]. This allows the mass flow rate, $\dot{m}_{server}$, to be calculated by applying conservation of energy, i.e. $E_{server}=\dot{m}_{server}c_p\Delta T$. Here, c_p is the specific heat capacity of air, assumed to be 1.005 $kJkg^{-1}K^{-1}$. V_{server} can then be calculated using $V_{server}=\dot{m}_{server}\rho^{-1}$, assuming an air density of 1.2 kgm^{-3}.

Figure 4 shows the relationship between V_{bypass} and Δp_{CH} after the installation of the LBs. Note that the results represent the 'low bypass' scenario. The interfaces between the LBs and the rack were sealed with duct tape. The tests were undertaken under three conditions: (i) with the LB inlets sealed with duct tape, (ii) with the inlets open and the LB fans switched off (passive), and (iii) with the inlets open and the LB fans switched on. The results for all of these tests can be seen to be very similar at most pressures. With Δp_{CH} close to 0, tests with the LB fans switched on showed significant levels of bypass flow, which was not the case with the LB fans switched off. Observations made during these tests found considerable air velocities through the regions close to where the LBs were attached to the rack's equipment rails. Bypass through this region was not observed after removing the duct from the back of the LBs. It was concluded that the high air velocities prevalent at the outlet of the LB duct resulted in low pressures, causing air to be entrained from the cold to the hot aisle even in the absence of a prevailing pressure differential between the aisles. Since velocities at server outlets are likely to be significantly less than velocities at the outlet of the LB duct, it may be assumed that bypass is unlikely to be significantly elevated by the presence of servers in a real DC, although it highlights the potential impact of server exit velocity on thermal air management. Hence, the bypass measured with the LBs blocked provides a guide for an achievable level of bypass in a real DC, and has been used in the system model. The relationship between bypass and Δp_{CH} where $\Delta p_{CH}<0$ is similar to the relationship where $\Delta p_{CH}>0$.

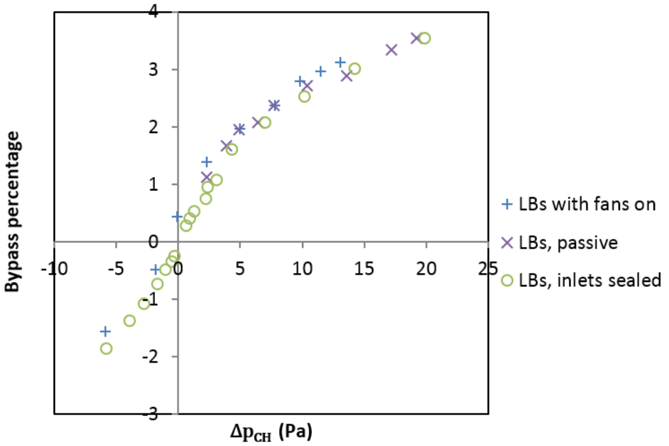

Figure 4: The effect on bypass of the introduction of LBs, with and without their fans in operation.

Figure 5 shows the results of the tests undertaken with the v40z server installed, compared with the results without the server installed. The tests with the server installed were undertaken firstly with the interface between the server and the rack left open, before sealing this interface with duct tape. The results show that allowing flow through this region significantly increases V_{bypass} at higher pressures, although this effect was less clear at lower pressures.

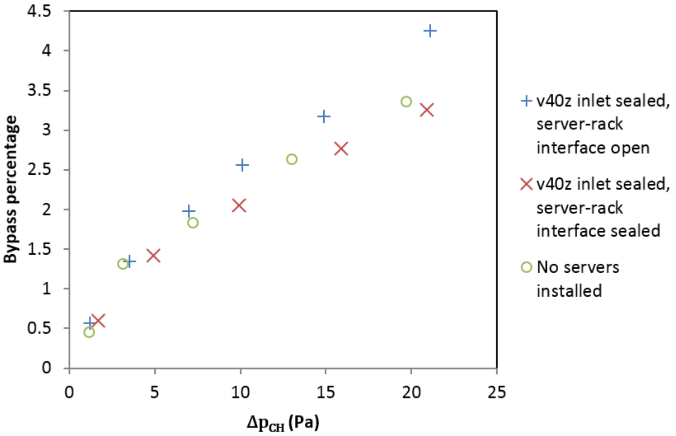

Figure 5: Relationship between V_{bypass} and Δp_{CH} with and without the v40z server in place.

4.2. GSWT experiments

Figure 6 shows the variation of the computational performance of the server (in GFLOPS/W) with T_{supply}. There is no clear trend between the two variables, showing that within this temperature range the computational performance of the server is consistent.

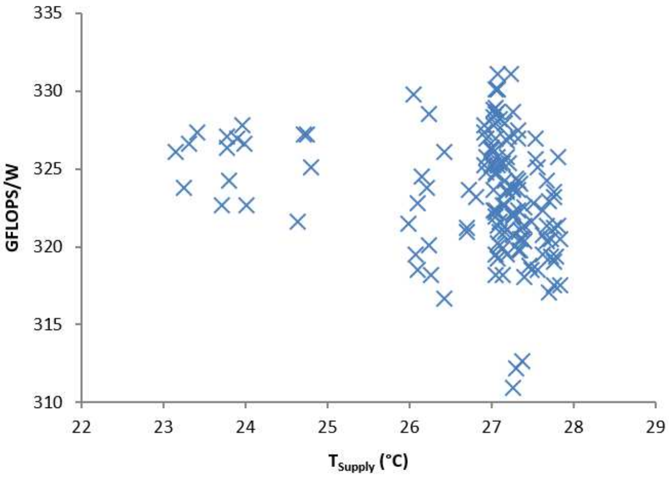

Figure 6: Variation of GFLOPS/W with T_{supply} for the v20z server.

Figure 7 shows the variation of the computational performance of the v20z server (in GFLOPS/W) with Δp_{CH}. There is no clear trend between the two variables, showing that, within this range of Δp_{CH}, the computational performance of the v20z server is consistent.

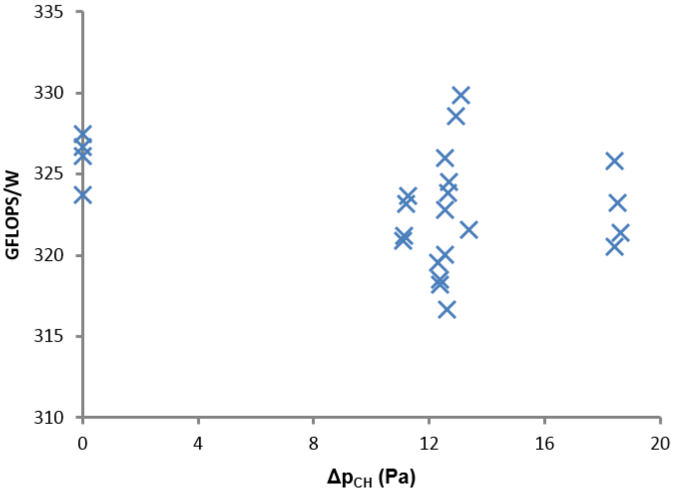

Figure 7: Variation of GFLOPS/W with Δp_{CH} for the v20z server.

Figure 8 shows the variation of V_{server} with Δp_{CH}. The flow rate can be seen to generally increase with pressure, although V_{server} falls as Δp_{CH} increases from 5 to 10 Pa. This could be due to V_{server} being determined as the difference between two measured flow rates, as described in section 3.2, leading to errors being compounded. However, the drop in V_{server} could also be in response to falling server fan speeds as the external static pressure makes an increasing contribution to flow through the server. Further investigation is required to establish the cause.

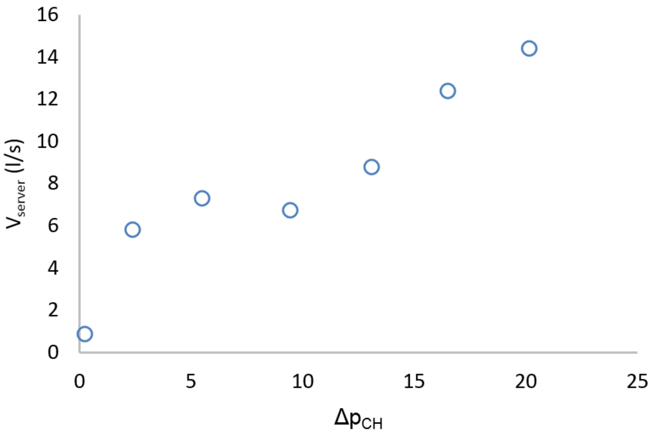

Figure 8: Variation of V_{server} with Δp_{CH} for the v20z server.

4.3. System model

The results shown in Figures 6-8 were used to determine flow rates through and computational efficiency of the servers in the system model, as functions of T_{supply} and Δp_{CH}. Since the computational efficiency was shown not to be affected by T_{supply} or Δp_{CH} over the ranges investigated, GFLOPS/W was maintained at the average level recorded.

Figure 9 and Figure 10 show the variation of partial PUE (pPUE) and GFLOPS/W_{DC}, respectively, with Δp_{CH}. Here, GFLOPS/W_{DC} is the number of GFLOPS per Watt of E_T, and is normalized such that the maximum value of GFLOPS/W_{DC} recorded is set equal to 1. pPUE is a term used by The Green Grid [28] to describe an imperfect PUE measurement, and here is given by $E_T/E_{servers}$. This could not be described as a true PUE since, for example, the electricity consumption of the power distribution system is not considered, and the results represent a snapshot of power consumption under certain ambient conditions. In each figure, the results are shown alternately with T_{supply}=23℃ and T_{supply}=28℃, and with the high and low bypass scenarios. The figures show that performance against both pPUE and GFLOPS/W_{DC} worsens as Δp_{CH} increases, and as T_{supply} falls (note that increasing efficiency is indicated by increasing GFLOPS/W_{DC} and decreasing pPUE). Since the performance of the v20z server was shown not to be affected strongly by supply air conditions, the increase in pPUE and decrease in GFLOPS/W_{DC} with increasing Δp_{CH} result from increasing power consumption in the cooling infrastructure as a result of increasing process air flow rate and decreasing return air temperature. The changes in pPUE and GFLOPS/W_{DC} are significantly greater in the high bypass scenario.

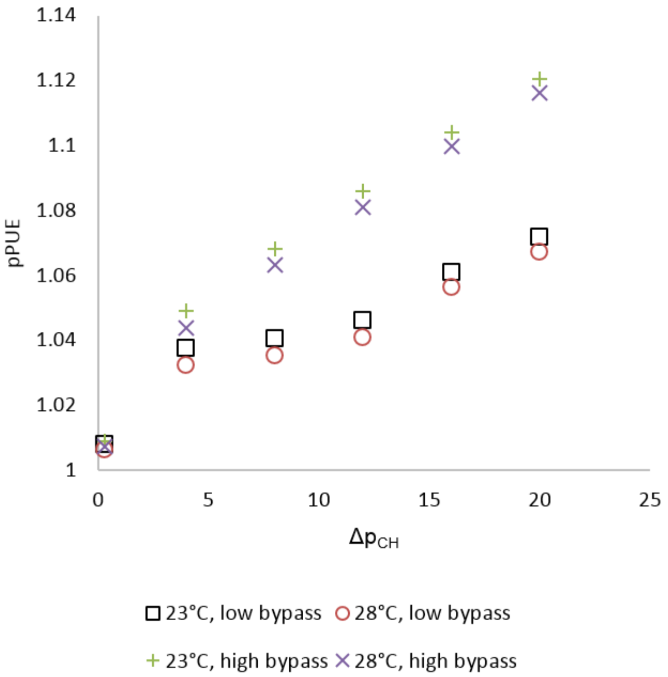

Figure 9: Variation of pPUE with Δp_{CH} for a DC populated with the v20z server.

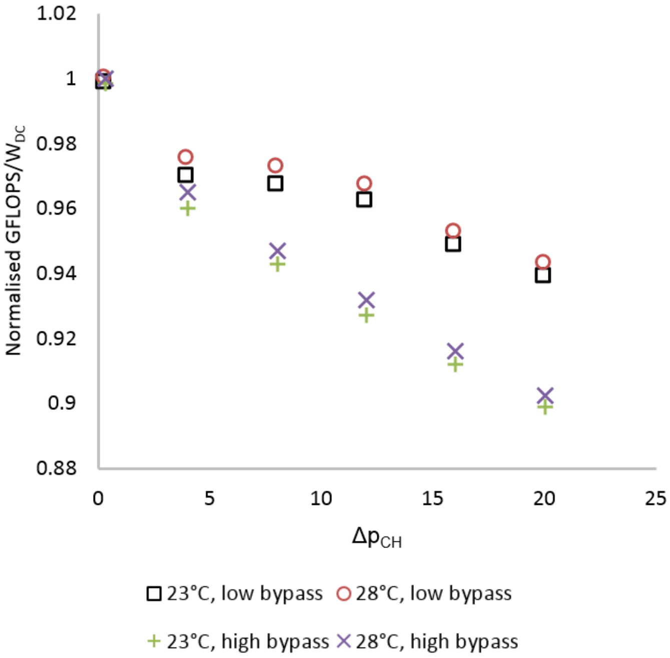

Figure 10: Variation of GFLOPS/W_{DC} with Δp_{CH} for a DC populated with the v20z server.

5. Discussion

The experimental results reported in this paper and in [18] have shown that V_{bypass} in aisle contained DCs is largely determined by Δp_{CH} and the extent to which leakage paths are sealed, and is not significantly affected by flow rates through IT equipment. Bypass at the interfaces between servers and racks may contribute significantly to V_{bypass}, particularly where $\Delta p_{CH} \geq 10$ Pa.

The results of the GSWT experiments show that, for the server investigated, computational efficiency is not significantly affected by T_{supply} or Δp_{CH} within the ranges investigated. This was unexpected, with results published by Zapater et al [27] demonstrating increasing server power consumption with increasing T_{supply} for an unspecified enterprise server. Similarly, the authors of the present paper note that preliminary tests using a more power dense server have shown computational efficiency of the server to be impacted significantly by Δp_{CH} and T_{supply}. Future work will explore the implications of these results.

The results of the system model show that pPUE increases by 6.3% and GFLOPS/W_{DC} falls by 5.9% as Δp_{CH} increases from 0.3 to 20 Pa, in the low bypass scenario, where T_{supply}=23°C. For the high bypass scenario, pPUE increases by 11.1% and GFLOPS/W_{DC} falls by 10.0% over the same pressure range and for the same T_{supply}, Note that increasing pPUE and decreasing GFLOPS/W_{DC} both represent worsening performance against the respective metrics. Increasing T_{supply} slightly increases the impact of Δp_{CH} on both metrics, and tends to slightly improve performance against both metrics regardless of Δp_{CH}.

The results build on those presented in [18], which showed that the relationship between E_T and Δp_{CH} depended on the assumptions made about the impact of Δp_{CH} on server flow rates. For example, allowing server fan speeds to reduce in response to increasing Δp_{CH} gave a minimum E_T with Δp_{CH} set at around 10-15 Pa, whereas if server fan speeds were held constant, lowering Δp_{CH} to 5 Pa reduced E_T further. Note that the results presented in the current paper are only representative of the behavior of the specific server upon which experiments were undertaken.

6. Conclusions

This work has furthered the existing knowledge relating to the efficient operation of DCs employing aisle containment. Laboratory experiments have been used to investigate the behavior of a server and aisle containment system under different supply air conditions. The experimental results have been used to develop a system model which quantifies the computational efficiency of the DC as a function of Δp_{CH} and T_{supply}. This is the first example of such a model presented in the literature, and the results presented constitute the first investigation of the impacts of Δp_{CH} and T_{supply} on computational efficiency in an aisle contained DC. Further experimental work is required to determine the impact of supply air conditions on computational efficiency with a wider range of server types installed.

Acknowledgments

The authors would like to thank DigiPlex Group of Companies, aql, and Hillstone Load Banks for funding the experiments detailed in this paper. Morgan Tatchell-Evans would also like to thank The Engineering and Physical Sciences Research Council for funding his research (Grant EP/G036608/1).

References

1. W. Van Heddeghem, S. Lambert, B. Lannoo, D. Colle, M. Pickavet, and P. Demeester, “Trends in worldwide {ICT} electricity consumption from 2007 to 2012,” Comput. Commun., vol. 50, pp. 64–76, 2014.
2. International Energy Agency, “World Energy Outlook 2014,” 2014.
3. L. Newcombe, M. Acton, J. Booth, S. Flucker, P. Latham, S. Strutt, and R. Tozer, “2012 Best Practices for the EU Code of Conduct on Data Centres,” 2012. [Online]. Available: http://re.jrc.ec.europa.eu/energyefficiency/pdf/CoC/Best Practices v3 0 8 _2_ final release Dec 2011.pdf. [Accessed: 12-Sep-2013].
4. U.S. Department of Energy and U.S. Environmental Protection Agency, “Energy Efficiency in Data Centers: Recommendations for Government-Industry Coordination,” 2008. [Online]. Available: http://www1.eere.energy.gov/manufacturing/tech_assistance/pdfs/data_center_wrkshp_report.pdf. [Accessed: 13-Jan-2014].
5. Environment Agency, “UMBRELLA CLIMATE CHANGE AGREEMENT FOR THE STANDALONE DATA CENTRES,” 2014. [Online]. Available: https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/336160/LIT_9990.pdf. [Accessed: 15-Sep-2014].
6. E. Oró, V. Depoorter, A. Garcia, and J. Salom, “Energy efficiency and renewable energy integration in data centres. Strategies and modelling review,” Renew. Sustain. Energy Rev., vol. 42, pp. 429–445, 2015.
7. M. Salim and R. Tozer, “Data Centers’ Energy Auditing and Benchmarking-Progress Update,” in ASHRAE Winter Conference, 2010, pp. 109–117.
8. J. Cho and B. S. Kim, “Evaluation of air management system’s thermal performance for superior cooling efficiency in high-density data centers,” Energy Build., vol. 43, no. 9, pp. 2145–2155, Sep. 2011.
9. V. K. Arghode, V. Sundaralingam, Y. Joshi, and W. Phelps, “Thermal Characteristics of Open and Contained Data Center Cold Aisle,” J. Heat Transf. - Trans. ASME, vol. 135, no. 6 Special Issue, 2013.
10. D. W. Demetriou and H. E. Khalifa, “Optimization of Enclosed Aisle Data Centers Using Bypass Recirculation,” J. Electron. Packag., vol. 134, no. 2, 2012.
11. Y. Joshi and P. Kumar, “Fundamentals of Data Center Airflow Management,” in Energy efficient thermal management of data centers, Y. Joshi and P. Kumar, Eds. New York: SPRINGER, 2012.
12. T. J. Breen, E. J. Walsh, J. Punch, A. J. Shah, and C. E. Bash, “From Chip to Cooling Tower Data Center Modeling: Influence of Server Inlet Temperature and Temperature Rise Across Cabinet,” J. Electron. Packag., vol. 133, no. 1, Mar. 2012.
13. K. Choo, R. M. Galante, and M. M. Ohadi, “Energy consumption analysis of a medium-size primary data center in an academic campus,” Energy Build., vol. 76, pp. 414–421, 2014.
14. Uptime Institute, “2013 Data Center Industry Survey,” 2013.
15. R. Schmidt, A. Vallury, and M. Iyengar, “Energy Savings Through Hot and Cold Aisle Containment Configurations for Air Cooled Servers in Data Centers,” in ASME 2011 Pacific Rim Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Systems, 2011, pp. 611–616.
16. S. K. Shrivastava, A. R. Calder, and M. Ibrahim, “Quantitative Comparison of Air Containment Systems,” in Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), 2012 13th IEEE Intersociety Conference on, 2012, pp. 68–77.
17. S. A. Alkharabsheh, B. Muralidharan, M. Ibrahim, S. K. Shrivastava, and B. G. Sammakia, “Open and Contained Cold Aisle Experimentally Validated CFD Model Implementing CRAC and Server Fan Curves for a Data Center Test Laboratory,” in ASME 2013 International Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems, 2013.
18. M. Tatchell-Evans, N. Kapur, J. Summers, H. Thompson, and D. Oldham, “An experimental and theoretical investigation of the extent of bypass air within data centres employing aisle containment, and its impact on power consumption,” Appl. Energy, p. -, 2016.
19. V. K. Arghode and Y. Joshi, “Room Level Modeling of Air Flow in a Contained Data Center Aisle,” J. Electron. Packag., vol. 136, no. 1, p. 11011, Feb. 2014.
20. S. A. Alkharabsheh, B. G. Sammakia, and S. K. Shrivastava, “Experimentally Validated Computational Fluid Dynamics Model for a Data Center With Cold Aisle Containment,” J. Electron. Packag., vol. 137, no. 2, Jun. 2015.
21. S. Pelley, D. Meisner, T. F. Wenisch, and J. W. VanGilder, “Understanding and abstracting total data center power,” in Workshop on Energy-Efficient Design, 2009.
22. T. J. Breen, E. J. Walsh, J. Punch, A. J. Shah, C. E. Bash, B. Rubenstein, S. Heath, and N. Kumari, “From Chip to Cooling Tower Data Center Modeling: Influence of Air-Stream Containment on Operating Efficiency,” J. Electron. Packag., vol. 134, no. 4, 2012.
23. A. P. Wemhoff, M. del Valle, K. Abbasi, and A. Ortega, “THERMODYNAMIC MODELING OF DATA CENTER COOLING SYSTEMS,” in Proceedings of the ASME 2013 International Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems (InterPACK2013), 2013.
24. C. D. Patel, R. K. Sharma, C. E. Bash, and M. Beitelmal, “Energy Flow in the Information Technology Stack: Coefficient of Performance of the Ensemble and its Impact on the Total Cost of Ownership,” 2006. [Online]. Available: http://www.hpl.hp.com/techreports/2006/HPL-2006-55.pdf. [Accessed: 11-Aug-2015].
25. A. Heydari, “Thermodynamics Energy Efficiency Analysis and Thermal Modeling of Data Center Cooling Using Open and Closed-Loop Cooling Systems,” 2007.

26. M. Iyengar and R. Schmidt, “Analytical Modeling for Thermodynamic Characterization of Data Center Cooling Systems,” J. Electron. Packag., vol. 131, no. 2, p. 21009, Apr. 2009.
27. M. Zapater, O. Tuncer, J. L. Ayala, J. M. Moya, K. Vaidyanathan, K. Gross, and A. K. Coskun, “Leakage-Aware Cooling Management for Improving Server Energy Efficiency,” IEEE Trans. Parallel Distrib. Syst., vol. 26, no. 10, pp. 2764–2777, 2015.
28. The Green Grid, “PUE: A Comprehensive Examination of the Metric,” 2012. [Online]. Available: http://www.thegreengrid.org/~/media/WhitePapers/WP49-PUE A Comprehensive Examination of the Metric_v6.pdf?lang=en. [Accessed: 01-Aug-2013].
29. J. Yuventi and R. Mehdizadeh, “A Critical Analysis of Power Usage Effectiveness and Its Use as Data Center Energy Sustainability Metrics,” Energy Build., vol. 64, pp. 90–94, 2013.
30. ASHRAE, Thermal Guidelines for Data Processing Environments. Atlanta: ASHRAE, 2012.
31. Hillstone Loadbanks, “Using the 6RM Server Simulator in IT Cabinets,” 2015. [Online]. Available: http://hillstone.co.uk/Download/HAC230-6RM.pdf. [Accessed: 06-Jul-2016].
32. Oracle, “Sun Fire V20z and Sun Fire V40z Servers--User Guide,” 2008. [Online]. Available: https://docs.oracle.com/cd/E19121-01/sf.v40z/817-5248-21/chapter1.html. [Accessed: 02-May-2016].
33. CNet, “Sun Fire V40z - Dual-Core Opteron 875 2.2 GHz.” [Online]. Available: http://www.cnet.com/products/sun-fire-v40z-dual-core-opteron-875-2-2-ghz-monitor-none-series/specs/. [Accessed: 02-May-2016].
34. Standard Performance Evaluation Corporation, “SPECpower_ssj2008.” [Online]. Available: https://www.spec.org/power_ssj2008/. [Accessed: 26-Sep-2016].
35. StressLinux, “StressLinux.” [Online]. Available: https://www.stresslinux.org/sl/. [Accessed: 26-Sep-2016].
36. Oracle, “Sun Fire V20z Server,” 2011. [Online]. Available: http://www.andovercg.com/datasheets/sun-sunfire-V20Z.pdf. [Accessed: 01-Apr-2016].
37. Digitron, “2000 Series Technical Data Sheet.” [Online]. Available: http://docs-europe.electrocomponents.com/webdocs/0785/0900766b8078507e.pdf. [Accessed: 23-Aug-2014].
38. The Met Office, “Mean temperature - Annual average: 1971-2000.” [Online]. Available: http://www.metoffice.gov.uk/public/weather/climate/gcpuckhb6#?region=southernengland. [Accessed: 24-Aug-2015].
39. C.-C. Wang and Y.-J. Chang, “Sensible heat and friction characteristics of plate fin-and-tube heat exchangers having plane fins,” Int. J. Refrig., vol. 19, no. 4, pp. 223–230, 1996.
40. T. RABAS, P. ECKELS, and R. SABATINO, “THE EFFECT OF FIN DENSITY ON THE HEAT-TRANSFER AND PRESSURE-DROP PERFORMANCE OF LOW-FINNED TUBE BANKS,” Chem. Eng. Commun., vol. 10, no. 1–3, pp. 127–147, 1981.
41. R. H. S. Winterton, “Where did the Dittus and Boelter equation come from?,” Int. J. Heat Mass Transf., vol. 41, no. 4, pp. 809–810, 1998.
42. R. Miller, “Emerson: 10 Common Data Center Surprises,” 2012. [Online]. Available: http://www.datacenterknowledge.com/archives/2012/03/21/emerson-10-common-data-center-surprises/. [Accessed: 13-Aug-2014].
43. D. W. Demetriou and H. E. Khalifa, “Energy Modeling of Air-Cooled Data Centers: Part II—The Effect of Recirculation on the Energy Optimization of Open-Aisle, Air-Cooled Data Centers,” in ASME 2011 Pacific Rim Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Systems, 2011.