University of Michigan – Ann Arbor

Electrical Engineering and Computer Science Department

DATA Center Modeling Analysis & Validation

Yifan YangGong

Thomas Wenisch

May 14th, 2009

I. Table of Contents

I. Introduction

A data center is a facility used to hold computer systems and associated components, such as data storage systems. Due to the need for fast internet connectivity and nonstop operation, data centers have grown rapidly in the past decade. With larger and denser data centers, energy consumption of data centers has increased sharply as well. It is projected to reach 100 billion kWh at an annual cost of $7.4 billion with two years [1]. Furthermore, the U.S. Environmental Protection Agency's (EPA) estimates that U.S. data center energy consumption will continue to grow by 12% per year [2].

As Figure 1 shows, much of the power within a data center is consumed by devices other than IT equipments which consist of computer systems and associated components [3]. Only 52% is actually consumed by IT equipment, whereas 38% of the energy is consumed by cooling equipment. Most existing efforts to improve data center energy efficiency focus on the energy-efficiency of IT equipment and cooling equipment, without considering global interaction across such subsystems. To minimize data center energy cost, we proposed the project Maelstrom.

Maelstrom consists of a real-time tracking system for data center thermal topology, which shows the temperature throughout the whole data center. During typical operation of a data center, servers produce heat that must be removed by the cooling equipment. Naturally, some servers can be more easily cooled than others (they may reside closer to the cooling equipment or contain newer equipment and thus be more efficient.) Thus, we may be able to increase energy efficiency by intelligently allocating computation tasks throughout the physical space of the data center. Maelstrom is a real-time dynamic computation allocation system based on thermal topology prediction, which will allocate computation evenly and eventually lessen/eliminate hot spots in data center. As Moore pointed out "For every 10^{o}C increase past 21^{o}C, equipment lifetime decreases by 50%" [4], with fewer hot spots, we can achieve both energy saving and longer equipment lifetime.

Since the airflow dynamics in a data center are complicated by many high-velocity air streams produced by IT equipment fans, Maelstrom uses computational fluid dynamics (CFD), a branch of fluid mechanics that uses numerical methods and algorithms to accurately solve for the pressure and heat topology of fluid flows. We chose Dolfyn, an open source CFD solver, to develop the real-time thermal topology tracking system.

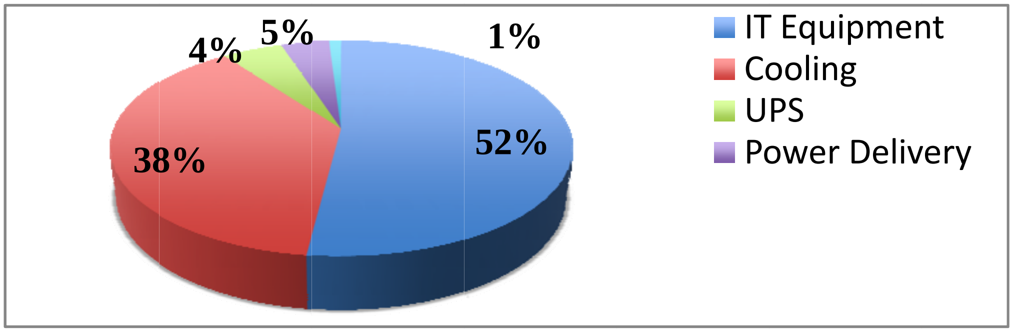

Figure 1 – Data Center Power Consumption Distribution

As a part of Maelstrom, my project focuses on analyzing the importance of CFD parameters using a current commercial data center modeling program, Flovent. Additionally, I was assigned to develop a tool to validate Dolfyn, by comparing its results with Flovent's.

II. MACC Data Center Modeling Analysis

Michigan Academic Computing Center (MACC) is a 8,500 square feet data center, and hosts a mix of systems across two rows with a total of 26 racks. Each rack can hold multiple servers stacked one above the other and typically houses a maximum of forty two servers. [5]

Flovent is a commercial CFD tool that predicts 3D airflow, heat transfer, and air contamination distribution in data centers. It was developed to help IT managers lay out the physical topology of data centers.

The MACC Data Center was previously modeled in Flovent with real-data settings on both IT equipments and cooling equipments by Alan Lee as shown in Figure 2. To simplify the model for analysis, computations allocation, fan speed and cooling temperature remains constant. By running the simulation with the MACC Data Center model in Flovent, we can get steady state pressure and thermal topologies (See Figure 3 for thermal topology). This result has been validated against experimental results.

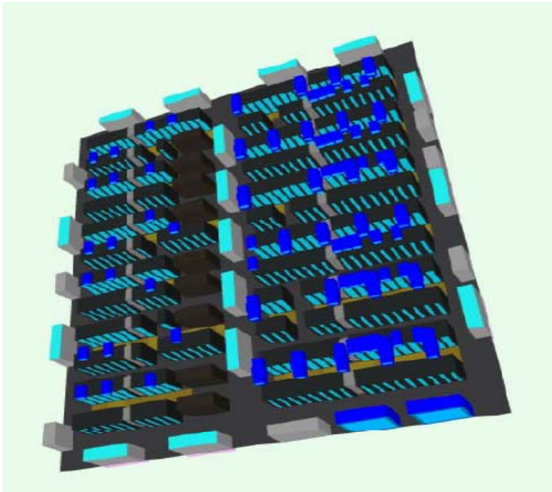

Figure 2 – MACC Data Center Model

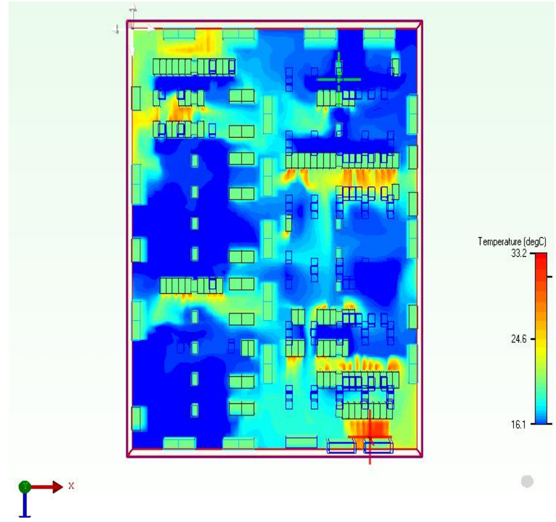

Figure 3 – MACC Data Center Thermal Topology

However, due to the complexity of the CFD solver, Flovent is only able to generate the steady state thermal topology with minimum 1 hour runtime for the MACC data center model. This is too slow to be useful for computation task scheduling, which requires the runtime to be no more than a few seconds. Thus, it is important to understand how various parameters of a CFD solver affect solution time and accuracy. My research investigates space gradient size, turbulence models, and locations of hot spots with respect to computing allocation (and thus heat generation).

A. Space Gradient Analysis

Modern CFD software uses the Finite Volume Method (FVM) to numerically calculate pressure and thermal topology. This method divides the solution domain into a number of small control volumes by a grid [6]. By solving the thermal equations independently for each grid, average temperature and pressure for each grid can be obtained. Since in reality there is some variation within each cell, smaller cells yield more accurate results. Hence, with more grids in a given region, more calculations are needed, but the pressure and thermal topologies can be represented more precisely.

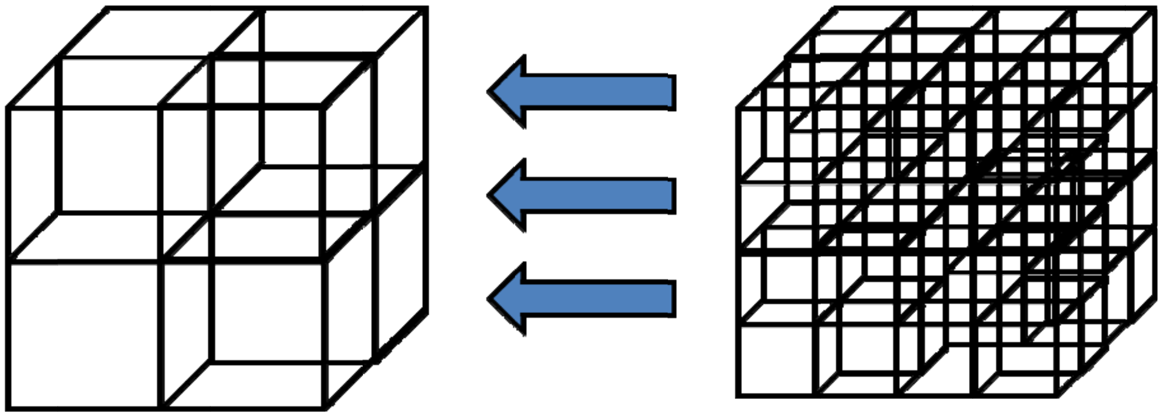

Figure 4 – Space Gradient Experiment Objective

In my project, I simulated the MACC with fewer numbers of grids except for the default setting (see Figure 4), and my objective was to find the coarsest grid that would still provide accurate results. Since some areas, where the temperature and pressure do not change dramatically from grid to grid, are less prone to error, Flovent allows different sizes of grids for different areas in a data center to speed the calculation (known as block-structured grids).

smallest grid size (m)	**Run time**	**Max (^oC)**	**Min (^oC)**	**Max diff (^oC)**	**Average diff (^oC)**
3×10^{-5}	1h14m48s	30.5638	16.1112	n/a	n/a
3×10^{-4}	1h14m54s	30.5638	16.1112	0	0
3×10^{-3}	1h14m16s	30.5638	16.1112	0	0
3×10^{-2}	59m42s	30.492	16.1113	4.9385	0.48153
3 X 10^{-1}	18m590s	30.492	16.1113	4.9385	0.48153

Table 1 – Space Gradient Simulation results

Flovent already has a function to optimize the size of the grid to get very accurate pressure and thermal topologies. Thus, I simply started with the default grid size settings, changed the smallest grid size to larger values and compared the results with the default result, 3×10^{-5} meter. As shown in Table 1, Flovent still gives the exact same result as the default grid size until the smallest grid size reaches 3×10^{-2} meter. Hence for a typical data center such as the MACC, 3×10^{-3} meter is a sufficiently small grid size to accurately solve for pressure and thermal topologies. Since the grid size is smaller than 3×10^{-3} meter in only a few areas in the MACC model, the run time doesn't decrease as much as expected in changing the grid size from 3 X 10^{-5} to 3 X 10^{-3}. However, its small decrease still shows that changing the grid size is a solution to speed up the pressure and thermal topologies calculation and we believe we can achieve more with other data center models.

B. Turbulence Analysis

Turbulence is the state of fluid motion characterized by apparent randomness and chaos [7]. The study of turbulence requires a firm grasp of applied mathematics and considerable physical insight into the dynamics of fluids [8]. This makes it difficult to decide whether modeling turbulence is necessary, even when turbulence only occurs near walls in data centers. Figure 5 shows the state of air motion in a simple data center.

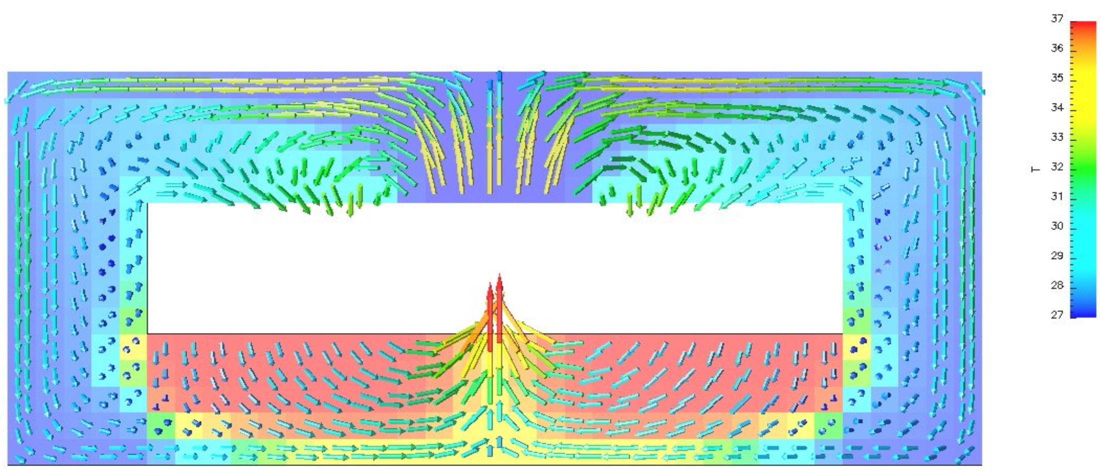

Figure 5—Turblence in Data Center

Since modeling turbulence takes much additional time, it is important to determine whether it is necessary. Thus, I carried out an experiment to numerically compare the thermal topologies generated with and without a turbulence model. This experiment is carried out in different space gradients for the MACC model in Flovent, as Table 2 shows.

Turbulence Mode	Min grid size (m)	Max (°C)	Min (°C)	Max diff(°C)	Average diff(°C)
On	3 X 10^{-5}	30.5638	16.1112	n/a	n/a
On	3 X 10^{-4}	30.5638	16.1112	n/a	n/a
On	3 X 10^{-3}	30.5638	16.1112	n/a	n/a
Off	3 X 10^{-5}	30.4927	16.1111	8.8862	0.5166
Off	3 X 10^{-4}	30.4927	16.1111	8.8862	0.5166
Off	3 X 10^{-3}	30.4927	16.1111	8.8862	0.5166

Table 2 – Turbulence Experiment Result

From the results, it is clear that modeling without turbulence gives drastically different results. In a typical data center model, only a 1-2 °C difference in temperature is acceptable and the maximum difference is 8.89 °C in this experiment. Therefore, we cannot simply say whether turbulence modeling is necessary at this point. If the difference only happens far away from IT equipments, we still can drop the turbulence modeling as it will have no effect on task scheduling. Thus, we need to keep the turbulence modeling, and further investigation is needed.

C. Computation Allocation Analysis

Hot spots not only decrease IT equipment lifetime, but also result in greater cooling costs. As current cooling equipment is designed to cool an area, not a single point, it is impossible to only cool hot spots without wasting energy cooling servers well within the desired temperature range. The whole data center needs to be cooler to keep hot spots cool enough so that it will not cause damage to the IT equipment. Since hot spots are where cold and hot air mix too early, not in the cooling system, they can frequently be eliminated by good design and maintenance. A good computation allocation design will ensure that hot air returns to the cooling system and will generate a more even thermal topology. Thus, eliminating hot spots by arranging the location of servers and computation load of servers is critical in data center design.

Figure 6 – Two Computation Allocation Experiment Result

In my project, I tested the relationship between workload location and the location of hot spots by placing various workloads throughout the MACC and observing the resulting temperature distributions. Unfortunately, as Figure 6 shows, even though the hot spots did move alongside the computation workload in the left graph, hot spots moved to seemingly random places in the right graph. Due to the flow behavior of air, it is frequently impossible to predict where hot spots will form giving a workload distribution using simple prediction schemes. Therefore, a CFD solver is necessary to accurately predict the locations of hot spots.

III. Dolfyn Validation

Dolfyn is an open source CFD solver. It was originally developed as an introductory tool to teach CFD. We chose Dolfyn to be the base for Maelstrom, as Dolfyn is widely used, such as the benchmark for FiPy developed by the American National Institute of Standards and Technology, and Dolfyn was proven to be bug free in 2008 [9]. Since Dolfyn is the base for Maelstrom, we continue to use Dolfyn as the name of the code of the real-time thermal tracking system.

As part of our modifications to Dolfyn, I am mainly responsible for proving its correctness, while Steven Pelley, a PhD student at the University of Michigan, is responsible for developing a real-time solver. To verify the thermal topology generated by Dolfyn, I constructed two simple data center models and designed an Auto-tester to compare the results from Flovent and Dolfyn.

A. Simple validation models

To simplify debugging, I constructed these following examples.

The first model is a $5m \times 5m \times 3m$ room with a single server placed in the middle. Since there is no cooling equipment in the room, the server is modeled as a fan, which only generates airflow without generating any heat. If heat were introduced, the temperature in the room will continuously increase, and never reach a steady state. By comparing the pressure topology against Flovent's results, it will show if Dolfyn works for this simple case. Figure 7 shows the pressure topology generated by Flovent.

The second model has the same size of the room as the first model, and has two servers symmetrically placed in the room. These two servers still only generate airflow, and do not generate heat. Figure 8 shows the pressure topology generated by Flovent.

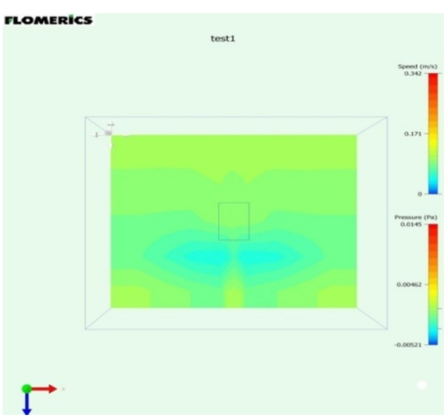

Figure 7 – One Server

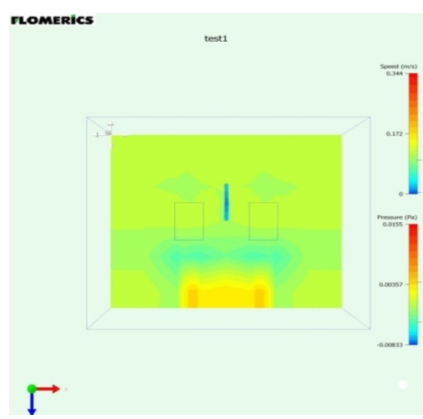

Figure 8 – Tow Servers

The third model has the same size of the room as previous two models, but has one server and one Computer Room Air Conditioner (CRAC), takes the place of the second server. In this model, the server will generate both airflow and heat, while the CRAC will remove the heat. To clearly show the difference in temperature throughout the room, tremendous power settings are applied to the server and CRAC. Figure 9 shows the thermal topology generated by Flovent.

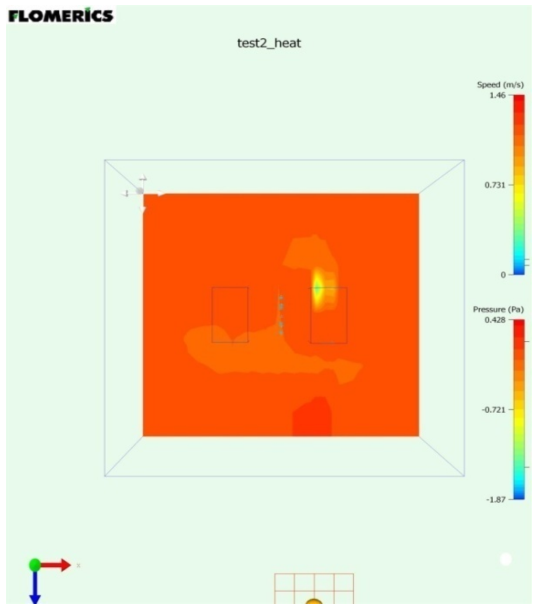

Figure 9 – One Server and One CRAC

B. Auto-tester

The results from Flovent and Dolfyn do not necessarily need to be exactly the same, but their difference needs to be within an acceptable range to show that they agree with each other. Therefore, I designed an auto-tester which uses statistical methods to compare the results, and display the output to show if they agree with each other.

Since CFD uses the finite volume method to calculate pressure and thermal topologies, the auto-tester's comparisons must be designed in the grid basis. Additionally, Flovent and Dolfyn may have different grid sizes. Therefore, the auto-tester's comparisons are designed to take all the grids from Flovent model within a cube, which is centered at the center of a grid from the model in Dolfyn. Then I applied statistic methods to compare them with the

Dolfyn grid at the center as shown in Figure 10. Finally, I applied this comparison to all grids from the model in Flovent.

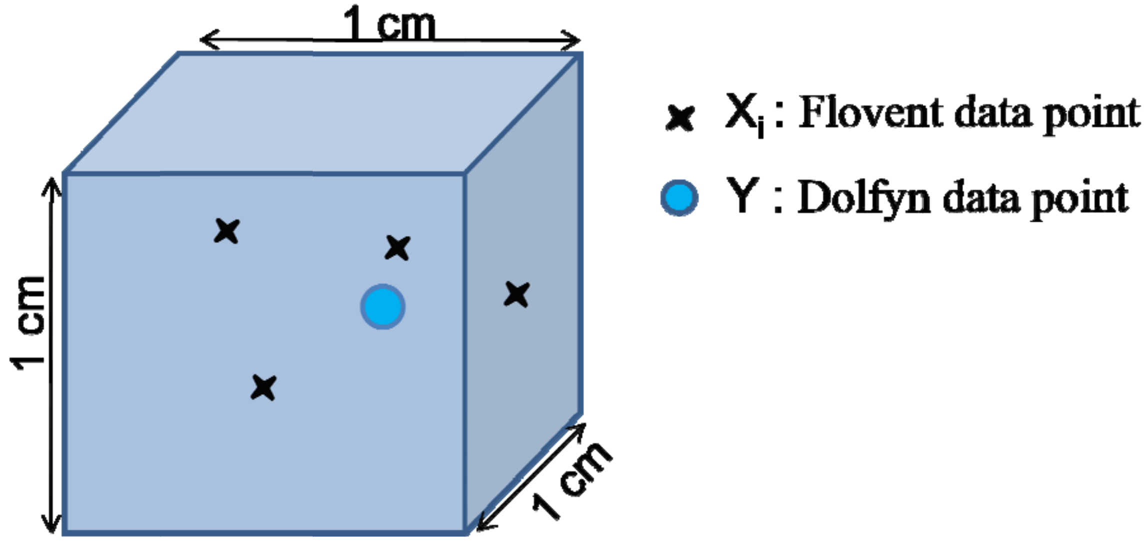

Figure 10 – Single Sample Comparison

The statistic methods used in Auto-tester are listed as follow:

- Arithmetic mean (default 2°C)

$$\Delta\beta = \frac{\sum X_i}{n} - Y$$

- Variance of the difference (default 0.25°C)\

$$\Delta\alpha = \frac{\sum (X_i - Y)^2}{n}$$

- Arithmetic mean of square difference (default 1°C)

$$\Delta\varepsilon = \frac{\sum \left(\frac{\sum (X_i - Y)^2}{n}\right)^2}{K}$$

K = number of Grids

Since these three statistic criteria have been set as parameters in the auto-tester, they can be changed based on different cases. If the parameters are not listed; when running the auto-tester, the default values will be used.

C. *Validation*

Validation is still in progress, as we continue implementing a real time solver.

IV. Conclusion

From the analysis of the MACC data center model in Flovent, it is clear that changing the size of the grid will speed up the thermal topology calculation, and the best solution for the MACC data center is 3 X 10^{-3}m. Additionally, since the results of the turbulence experiment show that turbulence significantly affects the thermal topology in some area, further research is needed to find out whether turbulence is critical for data center modeling. Finally, the computation allocation experiment demonstrates that the relationship between locations of hot spots and computation workload allocation is not straight forward, and the locations of hot spots are difficult to predict. This implies Maelstrom is beneficial and meaningful for energy saving in data center.

For Dolfyn validation, three simple models have been designed, and the pressure and thermal topologies have been solved for in Flovent. An auto-tester has been implemented and fully tested. Even though the auto-tester has not been used due to the incompletion of Dolfyn, I believe it will correctly catch differences in results from Flovent and Dolfyn, if there are any differences.

V. References

[1] U.S. EPA, “Report to congress on server and data center energy efficiency,” Tech. Rep., Aug. 2007.

[2] U.S. EPA, ENERGY STAR program, 2007, Report to Congress on Server and Data Center Energy Efficiency Public Law 109-431.

[3] S. Pelley, Y. YangGong, D. Meisner, T.F. Wenisch, and J.W. VanGilde, “Maelstrom: Real Time Data Center Temperature Modeling”, paper poster on ASPLOS, Mar. 2009.

[4] J. Moore, R. Sharma, R. Shih, J. Chase, C. Patel, and P. Ranganathan, “Going Beyond CPUs: The Potential of Temperature-Aware Data Center Architectures”, Workshop on Temperature-Aware Computer Systems, June 2004.

[5] MACC http://macc.umich.edu/

[6] J.H. Ferziger and M. Peric, *Computational Method for Fluid Dynamics*, 3rd ed. Springer, 2002.

[7] T. Bohr, M.H. Jensen, G. Paladin, and A. Vulpiani, *Dynamical Systems Approach to Turbulence*, 1st ed. Cambridge University Press, 1998

[8] P.A. Davison, *turbulence: AN INTRODUCTION FOR SCIENTISTS AND ENGINEERS*, 1st ed. Oxford University Press, 2004

[9] Dolfyn http://www.dolfyn.net/dolfyn/over_en.html