

**Cooling of Data Centers**

Akira Madono

## **Introduction to Problem**

The proliferation of mobile technology and internet traffic has increased needs for data transactions. This has resulted in a dramatic growth in data centers. Power consumption of these data centers is rapidly growing into an expensive issue. The main drivers of power consumption are high performance servers which have not decreased heat dissipation even when their processing speeds have increased. In fact, since high performance processors often involve more densely packing transistors into processor chips, their heat dissipation has increased. While the primary concern of a data center is ensuring continuous operation, cooling is certainly a concern given that cooling costs often constitute around 40% [8] of the power costs of a data center.

The most common configuration these days is one in which the server racks are arranged in rows such that the cooled inlet vents face one another and the hot exhaust sides face away from each other. The first data centers had this arrangement with cold air centrally vented into server rooms. While this arrangement allowed for the most flexibility in terms of server installation positions and adding additional servers, the servers at the bottom of racks suffered from insufficient flow and the top of the racks were vulnerable from hot air above. The raised floor configuration solves the floor server rack vulnerability while somewhat isolating incoming cold air supply from hot exhaust streams. In addition, the raised floor configuration maintained some degree of flexibility, allowing floor vents to be rearranged to accommodate new equipment or reconfigurations of rack layouts ad hoc. Nonetheless, the top server positions of racks remained vulnerable to exhaust air from the ceiling.

While the raised floor configuration remains the most common approach, further innovation has evolved to further consider isolating hot and cold air streams with partitions in the above floor space including dropping the ceiling and putting up vertical partitions blocking flow between aisles [1]. For some of the more extreme computing needs such as in supercomputing, individual air conditioning channels specific to each rack position are becoming more common.

Liquid cooling [7] is also being implemented with water pumped heat exchangers at the rear doors of the server racks. Nonetheless, installation of new racks or rearrangements are more cumbersome with these additions. Further, with liquid channels in close proximity to the server racks, leakage of the water is a potential hazard. The next generation of cooling solutions will better address efficiency and flexibility, while maintaining the reliability we have come to take for granted in our daily lives.

## Literature Review

*Source 1:* Pantakar's "*Airflow and Cooling in a Data Center*" [1] is a solid review of issues concerning cooling of data centers. His main analysis is of a conventional hot air aisle cold air aisle arrangement with an underfloor cooling chamber. The geometry of his arrangement is well described as 36'x14' with each rack occupying 6 ft<sup>2</sup>. A central cooling unit termed the CRAC provides cold air at 12.8°C at 4.72 m<sup>3</sup>/s to the underfloor plenum. The air then exits into two rows of square perforated tiles, each with 2'x2' area. Each rack is set to generate 2kW of continuous heat dissipation with 0.15 m<sup>3</sup>/s of forced convection air flow. The designed inlet temperature is 12.8°C and the designed outlet temperature is 24°C. The cooling is designated as exceeding tolerances when 35°C is exceeded.

His main argument throughout is that when insufficient airflow is found at the base of the rack, the inlet vents of the rack, particularly toward the top, become susceptible to hotter gases due to mixing of inlet air and exhaust air. He argues that if floor level airflow distributions are made as uniform as possible, cooling efficiency gains can be realized. Methods for achieving these uniform distributions are centered on plenum obstructions with the more optimal of these being inclined perforated planes. Above the floor, his simulations show that mixing above the racks and around the corners of the racks are obvious vulnerabilities. An often employed solution is to separately vent hot air and cold air channels, thus restricting inefficient mixing, further savings in cooling costs.

Due to the abundance of easily understood data given by the author, our simulations tried to match most of his parameters. These included the dimensions of the two-aisle data center as well as the flow temperatures and flow rates. We investigated the benefits of flow uniformity and venting. Because he did not explicitly specify the kinds of boundary conditions that would allow one to simultaneously model the heat generation and convection in the racks, we had to come up with our own strategy on this front.

*Source 2:* "*Effect of plenum chamber obstructions on data center performance*" by Fulpagare and others [2] focuses on the types of obstructions that might occupy underfloor plenum spaces and the effect they have on flow in the vicinity of the perforated tiles so often used in data centers. These obstructions, usually pipe obstructions are not usually engineered independently of cooling considerations. His work targets the rationale for regulations on piping building codes in data centers. His simulations show that when pipe obstructions are placed

under the floor tiles, temperatures rise and flow rates decrease with respect to those conditions at neighboring tiles. A further important result shows that the porous volume models assumed for perforated tiles often fail to accurately align with experimentally obtained velocity profiles.

Fulpagare's analysis provided us with some sense of caution regarding how to model the floor tiles, we found that his analysis mostly showed us the two main areas of contention among CFD modellers. Pantakar's paper argues that drop in pressure across the floor tiles is so small as to not impact heat and flow distributions. Fulpagare's paper, published in 2015 however, errs on the side of refuting that claim. Errors in estimating flow distribution's around any particular tile due to improper modeling could propagate to dramatic differences in velocity profiles. Another troubling finding of Fulpagare's paper was that plenum obstructions necessarily degrade flow quality. We decided to thus not explicitly model underfloor obstructions, observing that this would probably require a much more in-depth analysis of the variety suggested by Fulpagare's paper.

*Source 3: "Computational fluid dynamic investigation of liquid rack cooling in data centres"* by Ali Amoli and others [7] focuses on the calculations for efficiency factors involved in adding an active liquid cooling loop to the rear of server racks. The main appeal to considering this additional active element is the high specific heat capacity of the water. Given that the liquid cooling loop has its own pumping power requirements, the author's conclusion on efficiency gains is inconclusive. Nonetheless, the author does provide numerous tips for how to conduct a full-fledged sophisticated analysis of data centers. The author chooses grid size by iterating on the final aggregate temperature until variation is reduced. He finds that when the grid size is sufficiently large, the final aggregate temperature value becomes independent of grid size. The author further separates regions of the above floor simulation into numerous individual simulations with each simulations' outputs feeding the inputs of the neighboring simulation. This methodology of breaking the region was the principle motivation for why we chose to conduct separate simulations for below the floor and above the floor in our own simulations. For a more exhaustive procedure, we would further separately simulate inside the rack as the author has done. While we have attempted to model the porous perforated tile region in our model, we find that our approach pales in comparison to Almoli. He estimates the critical permeability term by estimating based on Hagen-Poiseuille flow in a pipe. Given more time, we could properly handle the source terms in the turbulence transport equations.

## Problem Setup

The start of our flow was given by the CRAC generating  $Q = 4.72 \text{ m}^3/\text{s}$  of flow through an area of  $A = 1.672 \text{ m}^2$  with hydraulic diameter of  $D_e = 1.219 \text{ m}$ . Thus assuming an initially uniform velocity profile, the average initial velocity was  $\frac{Q}{A} = u = 2.82 \text{ m}^2/\text{s}$ . The properties we assumed for air were the defaults given by FLUENT. Assuming a dynamic viscosity of air at room temperature of  $\mu = 1.7894e - 5 \frac{\text{kg}}{\text{ms}}$  and density  $\rho = 1.225 \text{ kg}/\text{m}^3$ , the Reynolds number  $Re = \frac{\rho u D_e}{\mu} = 2.36e5$ . The flow can thus be considered turbulent although not fully turbulent. For a turbulence model, we chose the reynolds averaged navier stokes (RANS) standard  $k - \epsilon$  model ( $k$  turbulent kinetic energy,  $\epsilon$  turbulent kinetic energy dissipation). This model expects fully conditions, which are not satisfied at the CRAC entrance. Nonetheless, given the popularity of the model, and because some areas of data centers may form fully turbulent flows we chose the  $k - \epsilon$  model. It is given by transport equations [5]:

$$\frac{\partial}{\partial t}(\rho k) + \frac{\partial}{\partial x_i}(\rho k u_i) = \frac{\partial}{\partial x_j} \left[ \left( \mu + \frac{\mu_t}{\sigma_k} \right) \frac{\partial k}{\partial x_j} \right] + G_k + G_b - \rho \epsilon - Y_M + S_k \quad (1)$$

$$\frac{\partial}{\partial t}(\rho \epsilon) + \frac{\partial}{\partial x_i}(\rho \epsilon u_i) = \frac{\partial}{\partial x_j} \left[ \left( \mu + \frac{\mu_t}{\sigma_\epsilon} \right) \frac{\partial \epsilon}{\partial x_j} \right] + C_{1\epsilon} \frac{\epsilon}{k} (G_k + C_{3\epsilon} G_b) - C_{2\epsilon} \rho \frac{\epsilon^2}{k} + S_\epsilon \quad (2)$$

Where the terms in the equation are defined as

|   |   |
|---|---|
| $G_k$   | Generation of kinetic energy due to mean velocity gradients                     |
| $G_b$   | Generation of turbulence kinetic energy due to buoyancy                         |
| $Y_M$   | Fluctuating dilation in compressible turbulence to the overall dissipation rate |
| $C_{1\epsilon}, C_{2\epsilon}, C_{3\epsilon}$ | Constants   |
| $\sigma_k, \sigma_\epsilon$                   | Prandtl numbers for $k$ and $\epsilon$  |
| $S_k, S_\epsilon$                             | User-defined source terms   |

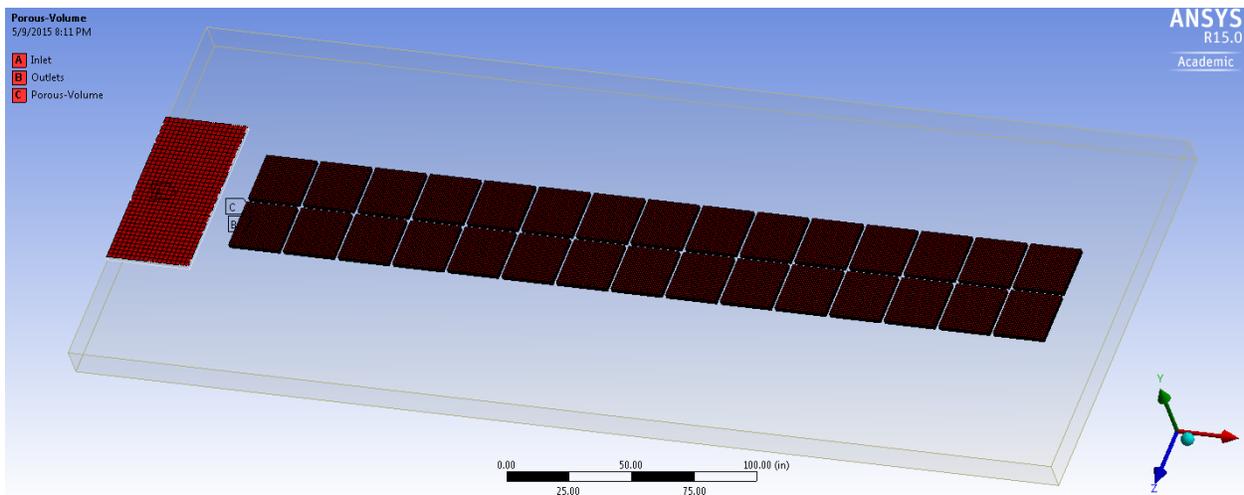
The solution method we chose was the SIMPLE scheme with first order upwind discretization for turbulent kinetic energy and dissipation rate. Momentum discretization was second order upwind.

Our boundary conditions were chosen with respect to those appropriate for compressible flow. Thus, for all inlets, mass-flow inlets were chosen. For all outlets, pressure-outlets were chosen. All other surfaces were modeled as no-slip zero flux walls. The only region of our flow

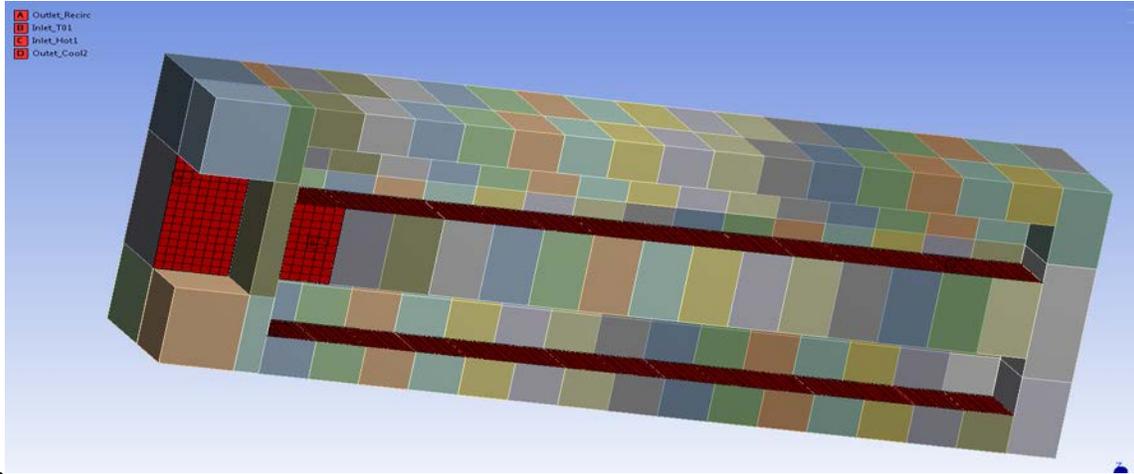
that we felt would need finer mesh sizes were the perforated tiles. We thus chose 0.75 inch element size meshes for the tiles. Otherwise, in order to accommodate a 510,000 element restriction and faster simulations, we chose an element size of 0.4 ft. In the context of a 36 ft x 14 ft x 15 ft total 3D simulation region, this element size resulted in a medium level mesh. Figure 1, Figure 2, and Figure 3 show schematics of our simulations. We determined incoming floor flow rates by converting Figure 13 of the Patakar data [1] to a tabulated numeric equivalent [6]. This was the basis for our uneven flow distributions and the data we strived to simulate in our underfloor, plenum simulations. The even flow distribution was based on the uniform  $0.15 \text{ m}^3/\text{s}$  volumetric flow rate deemed minimally sufficient by Pantakar [1]. In order to simulate temperatures, we invoked the energy conservation equation for energy  $E$ , velocity  $\vec{v}$ , enthalpy  $h_j$  and source term  $S_h$ .

$$\frac{\partial}{\partial t}(\rho E) + \nabla \cdot (\vec{v}(\rho E + p)) = -\nabla \cdot \left( \sum_j h_j J_j \right) + S_h \quad (3)$$

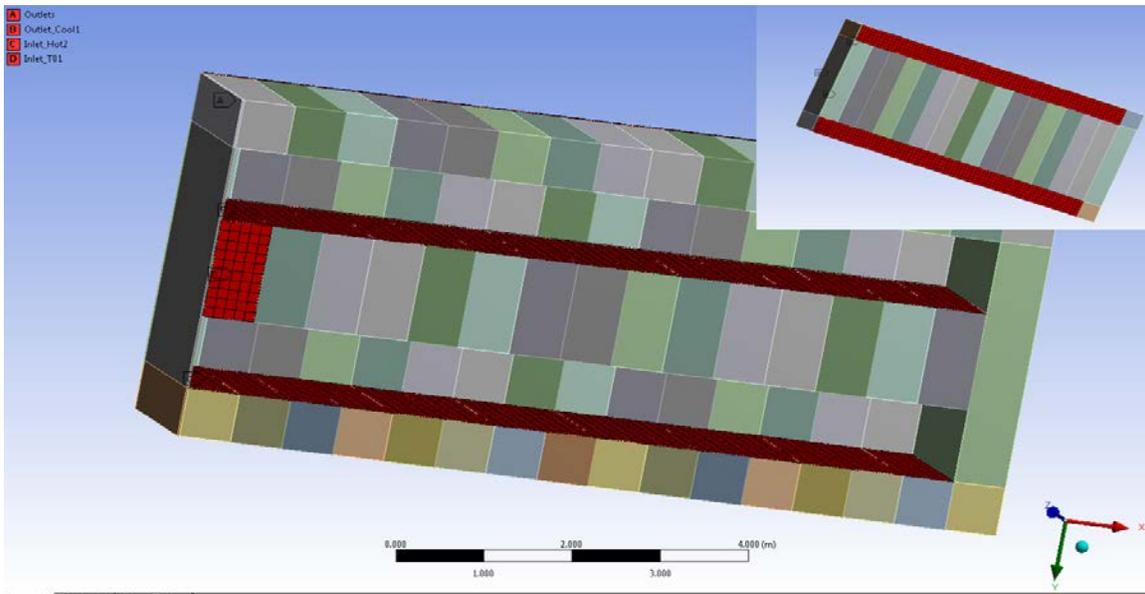
Nominal given values [1] of for incoming cool air of  $12.8^\circ\text{C}$  and exhausted hot air of  $24^\circ\text{C}$  were assumed.



**Figure 1.** Boundary Conditions: A (Inlet) Mass-flow inlet, B (Outlets) Pressure-outlet, C (Porous Volume) Mesh: 237,926 Nodes and 184,698 Elements.



**Figure 2.** *Boundary Conditions:* A (Outlet\_Recirc) Pressure-outlet, B (Inlet\_T1) Mass-flow inlet, C (Inlet\_Hot1) Mass-flow inlet, D (Outlet\_Cool2) Pressure-outlet. *Mesh:* 115,490 Nodes and 92,800 Elements.



**Figure 3.** *Boundary Conditions:* A (Outlets) Pressure-outlet, B (Outlet\_Cool1) Pressure-outlet, C (Inlet\_Hot1) Mass-flow inlet, D (Inlet\_T01) Mass-flow inlet. *Mesh:* 45,276 Nodes and 28,416 Elements.

## Discussion and Results

*Plenum:* When we simulated flow through the underfloor, plenum region, we found that the velocity profile down the length of the rack tiles was quite sensitive to tile porosity modeling. In order to facilitate discussion of our simulation results, we start with a description of the porous volume model we used. The mechanism of flow through the plenum is illustrated by the

pressure contour in Figure 4 which is a horizontal cross section just below the floor tiles. The sudden rush and expansion of air into the plenum creates a suction zone following the high pressure zone of the inlet. The suction region makes reverse flow possible and is also responsible for creating the uneven distribution of air. The extent of the suction zone is determined by the parameters of the porous, perforated floor tiles. The governing equation for pressure drop across the porous tile is given by

$$S_i = - \left( \frac{\mu}{\alpha} v_i + C_2 \frac{1}{2} \rho |v| v_i \right) \quad (4)$$

where  $S_i$  is the source term in the relevant momentum equation. We also knew that the pressure differences were modelled as

$$\Delta p = K \frac{1}{2} \rho V^2 \quad (5)$$

where  $K = 42.8$  for a 25% open tile [1]. Thus we set about to trying to determine parameters permeability  $\alpha$  and inertial resistance factor  $C_2$  that would provide us with the pressure differences sought after in Equation 5. The results of an attempt are given in Figure 5 and Figure 6. Difference from the literature velocities is evident in Figure 5. At a position of 10 m, the expected pressure difference based on a velocity of 0.9 m/s is 21.2 Pa based on Equation 5. Instead, the simulation gives a pressure difference of approximately 5 Pa as shown in Figure 6. The simulated velocity values given by Figure 7 shows velocity magnitudes closer to the literature values we could get by adjusting  $\alpha$ . At the same position of 10 m, the pressure difference was approximately 12 Pa as shown in Figure 8. The expected pressure difference was 7.93 Pa based on a velocity of 0.55 m/s, again by Equation 5. We thus proved that the velocity profile was sensitive to tile porous volume modelling as expected by [2].

At any given tile position, we also simulated non-uniform velocities as expected by literature [2]. Thus, the assumption of uniform velocity that we would carry over into boundary conditions for the above floor model would not be truly accurate. If we were to improve the above floor simulation, we would input the varied actual velocity values into the above floor simulation instead of assuming average constant values.

*Full ceiling above floor:* We did not need to model the perforated tiles as we had done in our first simulation because the above floor model begins with flow above the perforated tiles.

Furthermore, the above floor modelling could be accomplished independently of the below floor modelling because the flow values can be filled in manually at each tile outlet. Using the flow rates given by Pantakar and shown in Figure 5, we sought to determine sources of inefficient cooling in the above floor space.

The main difficulty to the above floor model was the need for setting appropriate thermal parameters with respect to the racks. This is because the racks provide both convection and thermal heat dissipation to the surrounding air flow. Unfortunately, we were unable to select boundary conditions in FLUENT that allowed temperatures to be unfixed, while fixing flow rates. We thus turned to design estimates from literature [1] of rack inlet temperature  $T_c = 12.8^\circ C$  and rack exit temperature  $T_h = 24^\circ C$ . From the first law of thermodynamics, we know that

$$\dot{Q} = \dot{m}c_p(T_h - T_c) \quad (6)$$

Each rack would receive half of the total flow such that  $\dot{m} = 2.86 \text{ kg/s}$ . The specific heat of air at room temperature,  $c_p = 1006 \frac{J}{kg K}$ . Thus one row of racks would dissipate 30kW of heat.

Given that our literature source stated a single rack heat dissipation of 2kW, and 15 racks per row, these initial design parameters were sensible to carry out throughout our simulations. If we were implementing our own solution scheme, we would certainly find a way to not fix the temperatures at the inlets and outlets of our racks.

Figure 9 shows the resulting overall pattern of temperature distribution in the data center for these parameters. The temperatures at the inlet of the rack are compromised while the temperatures on the exhaust side of the rack are left unperturbed. Because we did not want to bias flow between the rack and outlet, we also set the outlet to  $T_h$ ; because the highest temperature inputted into the simulation were  $T_h$ , no higher temperatures were possible. In reality, in the vicinity of the servers, higher temperatures could be produced if the servers are not provided sufficient cooling. Despite these setbacks, some clear trends are evident from the temperature contour in Figure 9. The inlet vent temperatures from Figure 9 are repeated in Figure 10(a) with respect to increasingly efficient configurations.

Three recirculation vulnerabilities are apparent in the full ceiling uneven flow simulation. Recirculation above the rack shown in Figure 11 (top) results in mixing of hot exhaust gases with cold air supply as shown in Figure 12 (top). This recirculation compromises cooling towards the top of the rack as seen in Figure 10(a) and Figure 10(b). Secondly, the lack of flow due to the

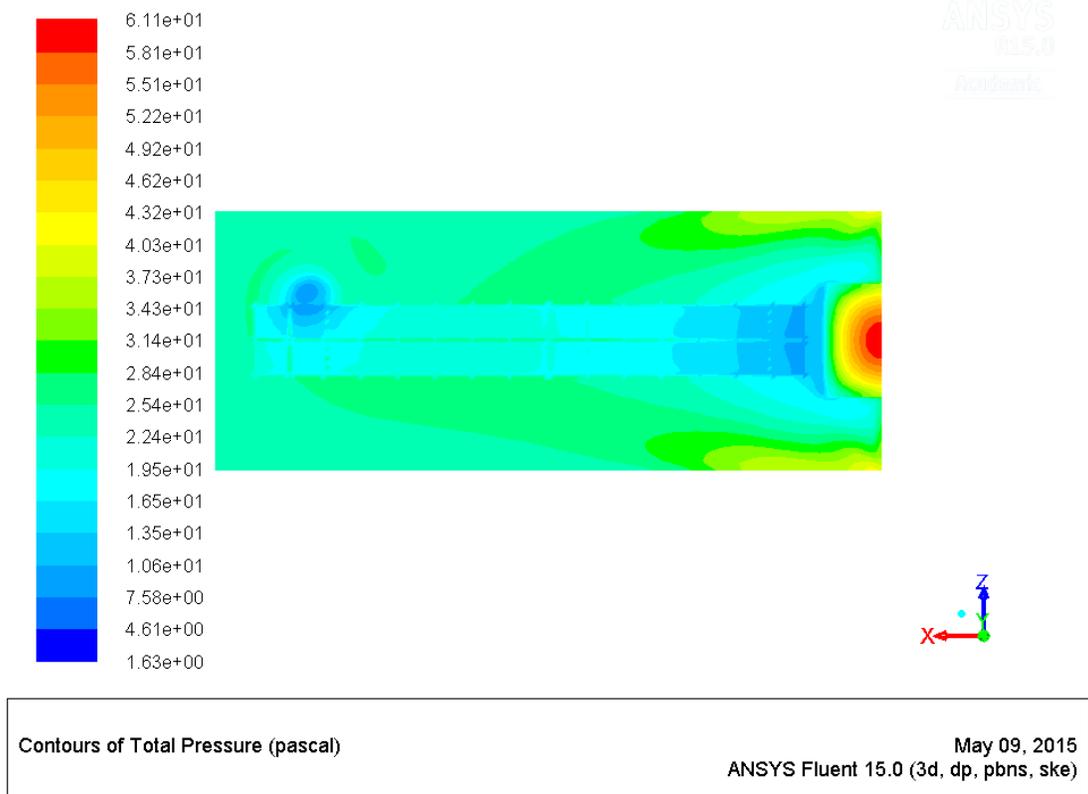
uneven flow makes cooling of the first rack positions vulnerable as shown in Figure 10(a) vs Figure 10(b). Third, recirculation at the corners of the rack row expose these regions to exhaust gas heating laterally. Our simulated designs did not directly address corner recirculation though restricting ceiling recirculation also seems to have decreased corner recirculation.

Figure 10 concisely shows the effects of each of these vulnerabilities. When the flow distribution is evenly set at  $0.15 \text{ m}^3/\text{s}$  the first few rack positions are better cooled. When the drop ceiling is vented, the eliminated recirculation volume largely eliminates higher temperatures. Further evening the flow subtly reduces the temperatures once again, though the effect is less pronounced. Of the two methods for reducing cooling inefficiencies, the drop ceiling is more effective. This counters the argument posed by Pantakar, that evening cooling flow distribution is the most effective improvement in cooling efficiency.

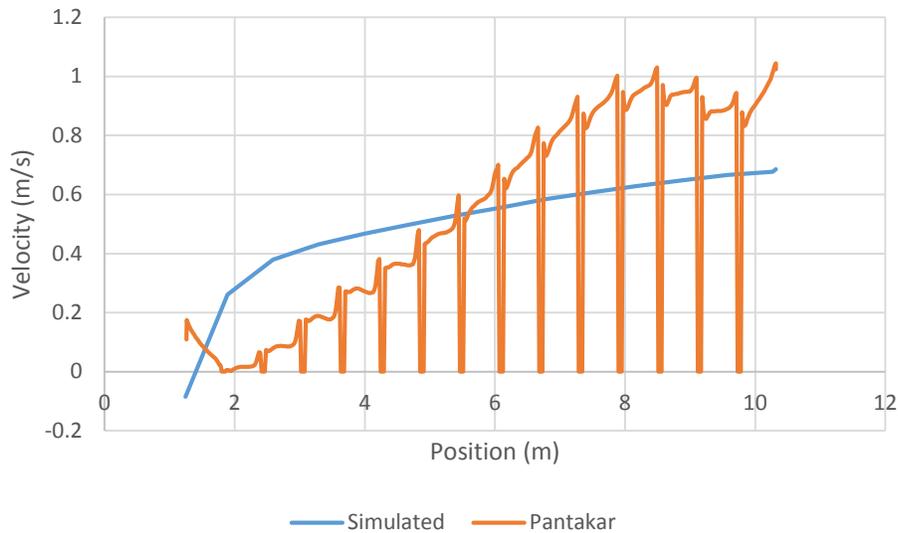
In order to simulate the above floor cooling at one go, we implemented the drop ceiling as one in which the ceiling was one inch above the top of the racks; otherwise a continuous flow cavity would not have been possible. Note that full implementation of a drop ceiling would completely eliminate flow between the cold air and hot air aisles; i.e. the ceiling would seal the top of the racks. Thus, we would follow an approach demonstrated by Amoli [7] in which the outputs of one model directly feed into the adjacent one. Nonetheless, Figure 11 (bottom) shows that the recirculation area, given by the corner eddy swirls in Figure 11(top), is effectively reduced by the drop ceiling configuration. The temperature contour in Figure 12 (bottom) shows that in fact minimal mixing occurs between the hot and cold air channels.

One additional caveat we should make about all our simulations is that convergence, especially with respect to continuity was not achieved as shown in all our convergence plots. This was a common difficulty we encountered when surveying fellow research groups. When we examined residual error in absolute terms however, we found that continuity was one of the lower residual errors. Nonetheless, residuals for other terms remained high. We believe that with further understanding of how to better specify turbulence models, we could nonetheless determine ways to better converge all relevant properties.

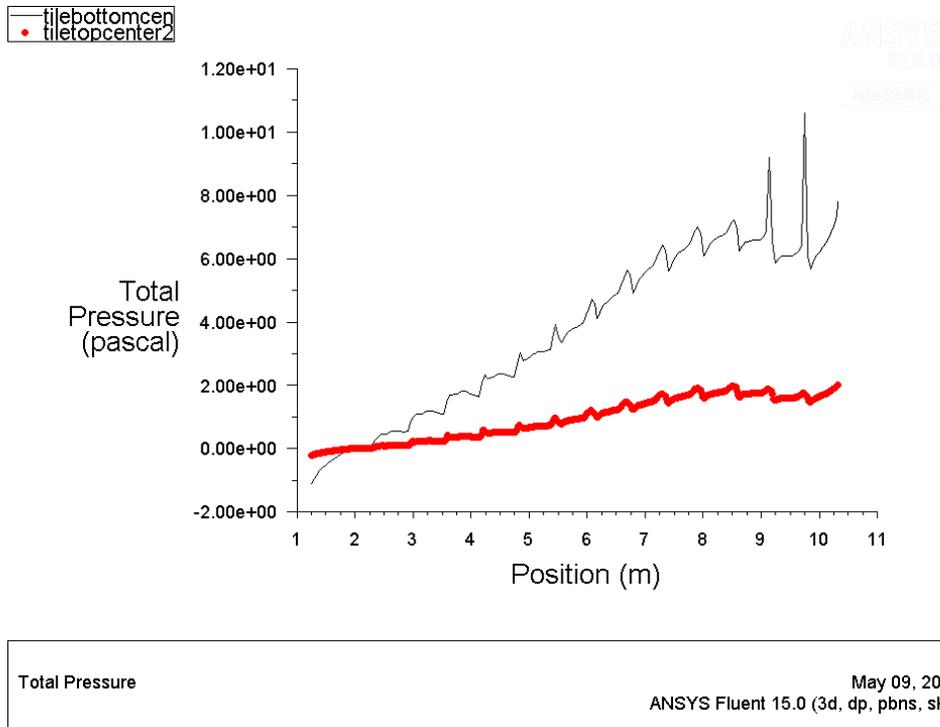
## Plots of Key Quantities



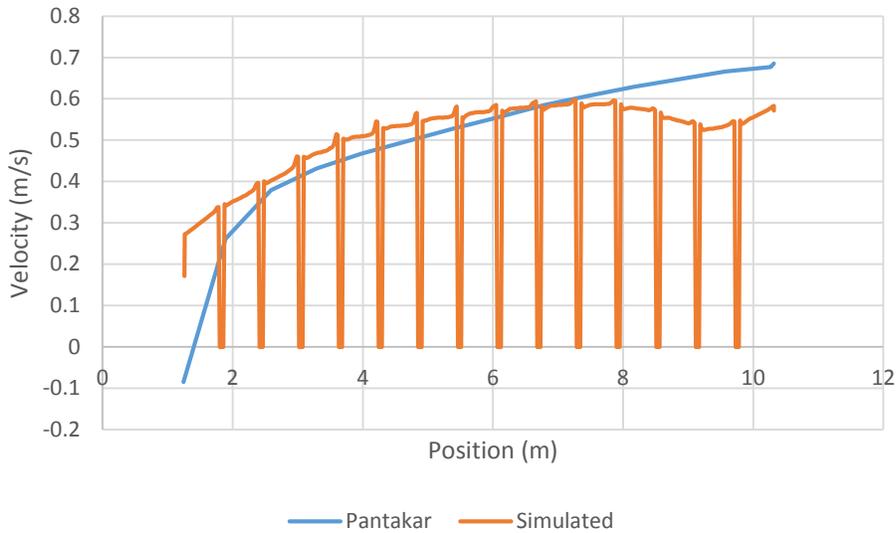
**Figure 4.** Pressure contour given by simulation.



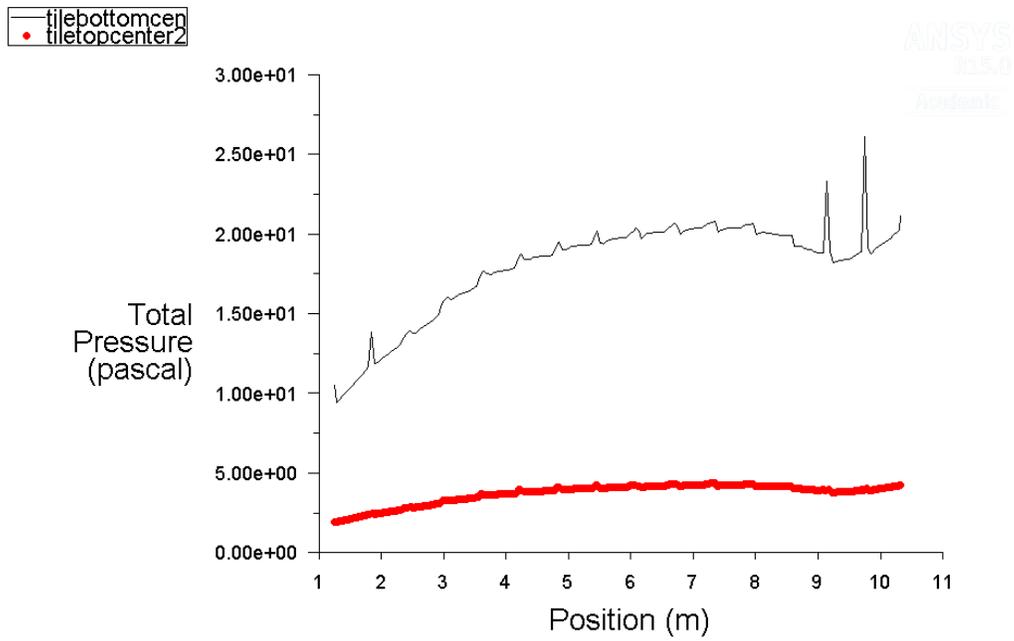
**Figure 5.** Plenum tile top center Simulated vs literature velocity values. [ $\alpha = 0.77e6 \text{ m}^{-2}$ ,  $C_2 = 20.414 \text{ m}^{-1}$ ].



**Figure 6.** Unadjusted Pressure difference top vs bottom of tiles for plenum. [ $\alpha = 0.77e6 \text{ m}^{-2}, C_2 = 20.414 \text{ m}^{-1}$ ].



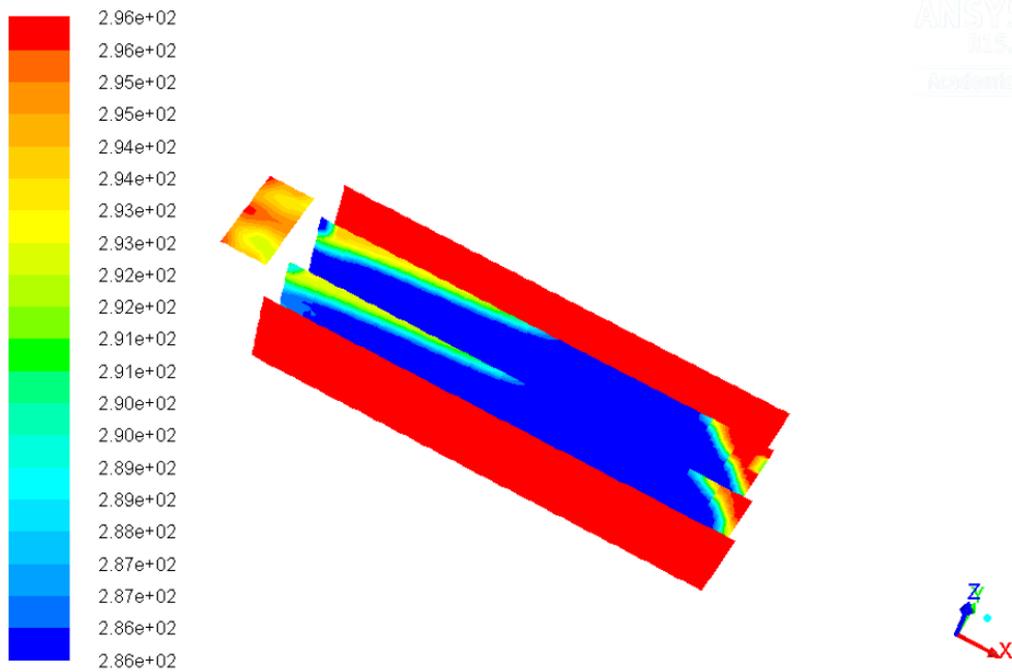
**Figure 7.** Plenum tile top center Simulated vs literature velocity values. [ $\alpha = 4.5e7 \text{ m}^{-2}, C_2 = 20.414 \text{ m}^{-1}$ ].



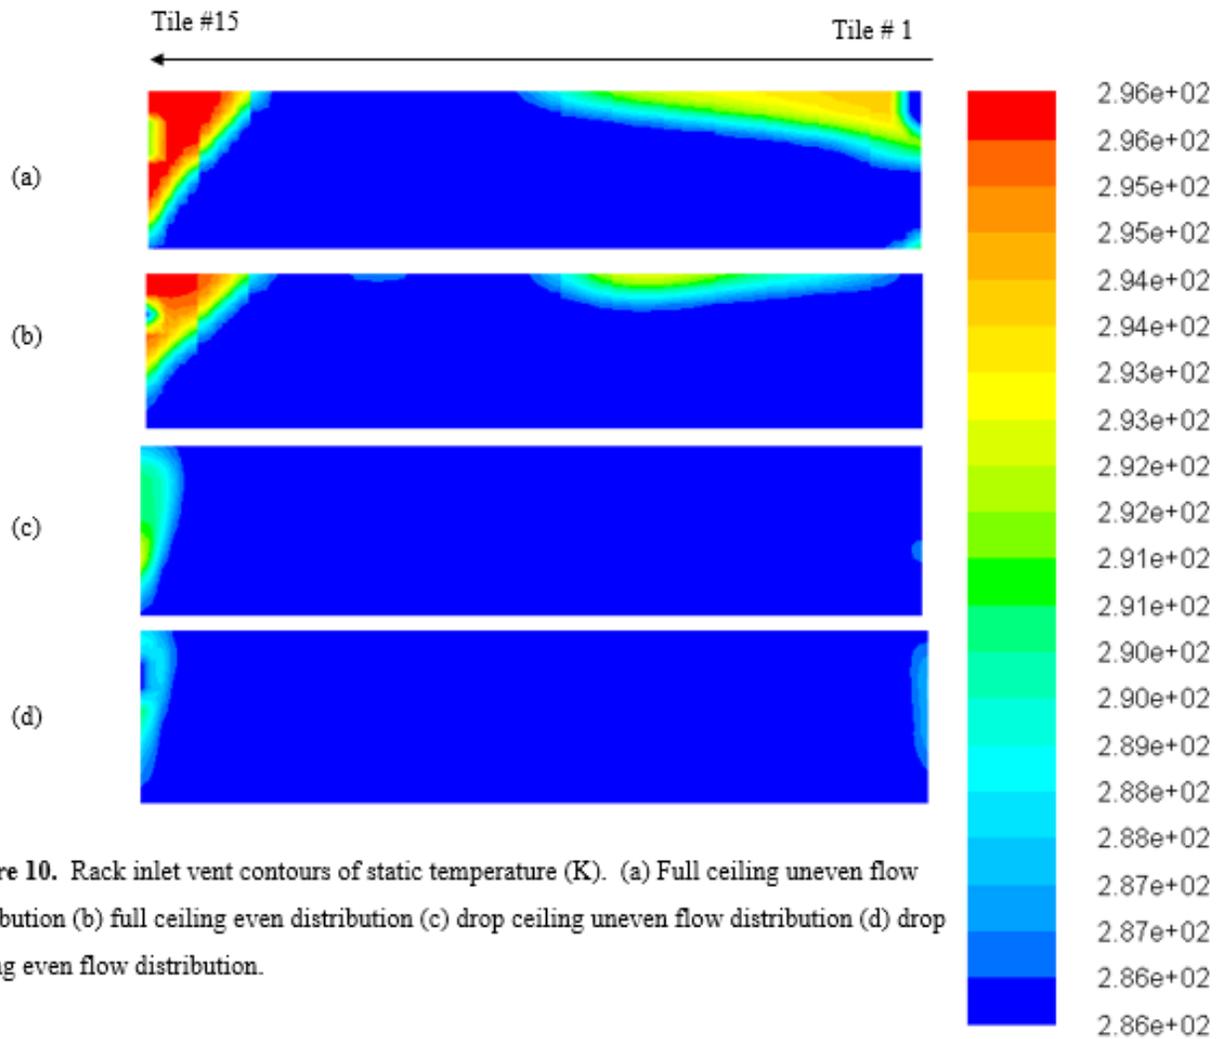
Total Pressure

May 09, 2015  
ANSYS Fluent 15.0 (3d, dp, pbns, ske)

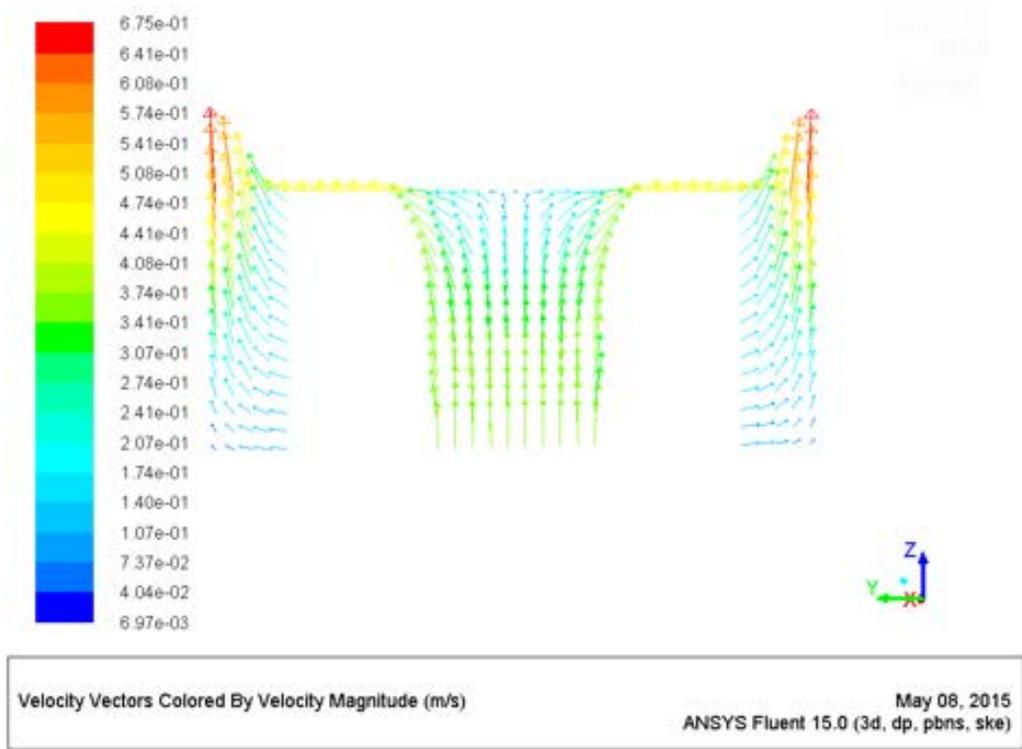
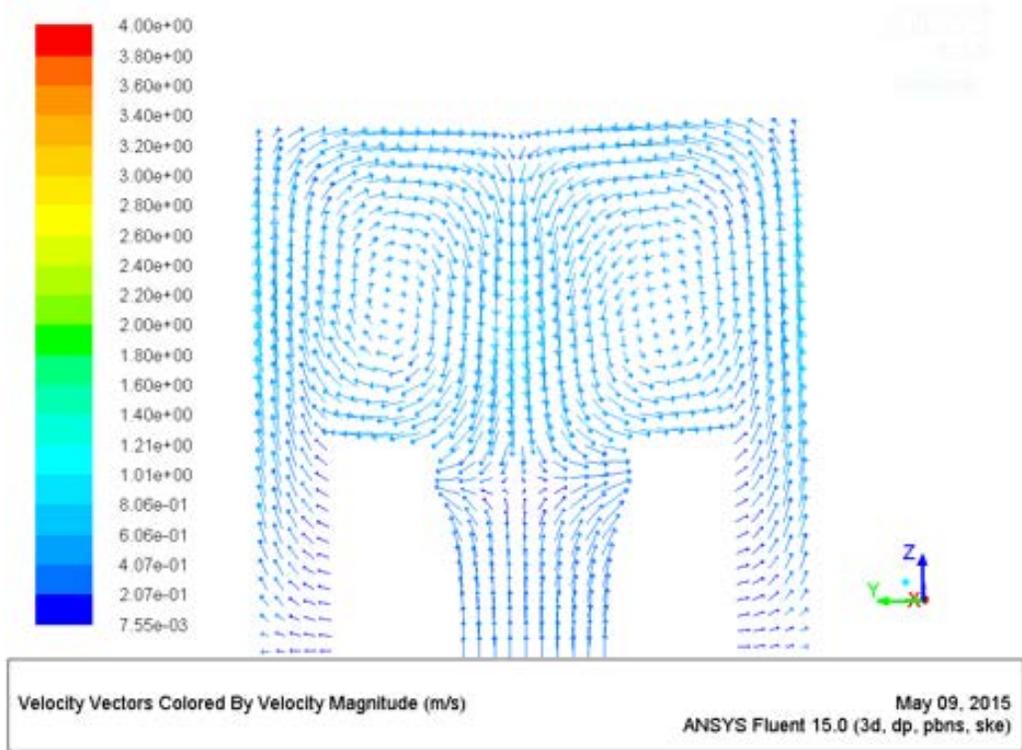
**Figure 8.** Adjusted Pressure difference top vs bottom of tiles for plenum. [ $\alpha = 4.5e7 \text{ m}^{-2}, C_2 = 20.414 \text{ m}^{-1}$ ].



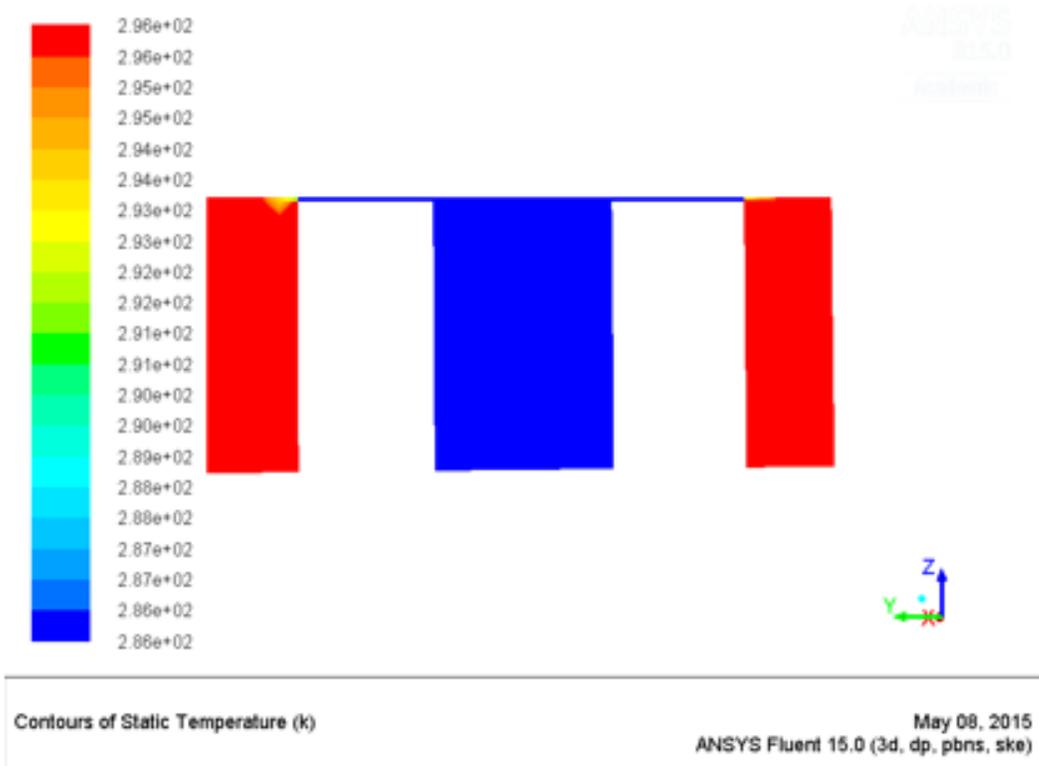
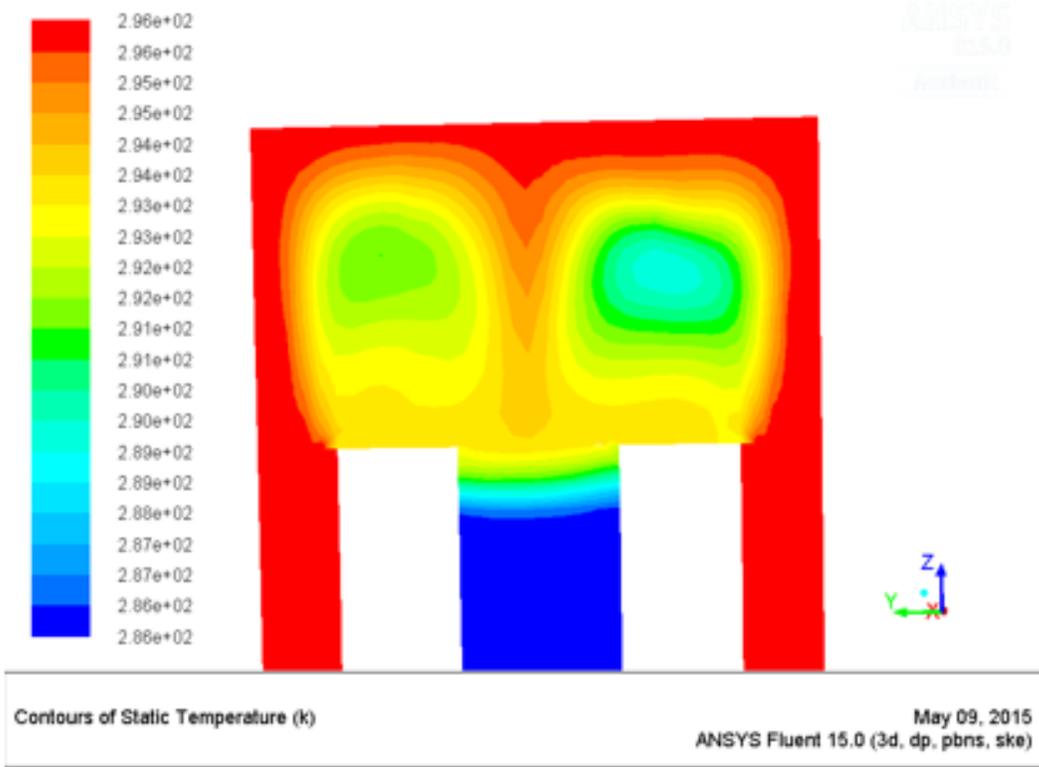
**Figure 9.** Contours of static temperature full ceiling uneven flow distribution.



**Figure 10.** Rack inlet vent contours of static temperature (K). (a) Full ceiling uneven flow distribution (b) full ceiling even distribution (c) drop ceiling uneven flow distribution (d) drop ceiling even flow distribution.

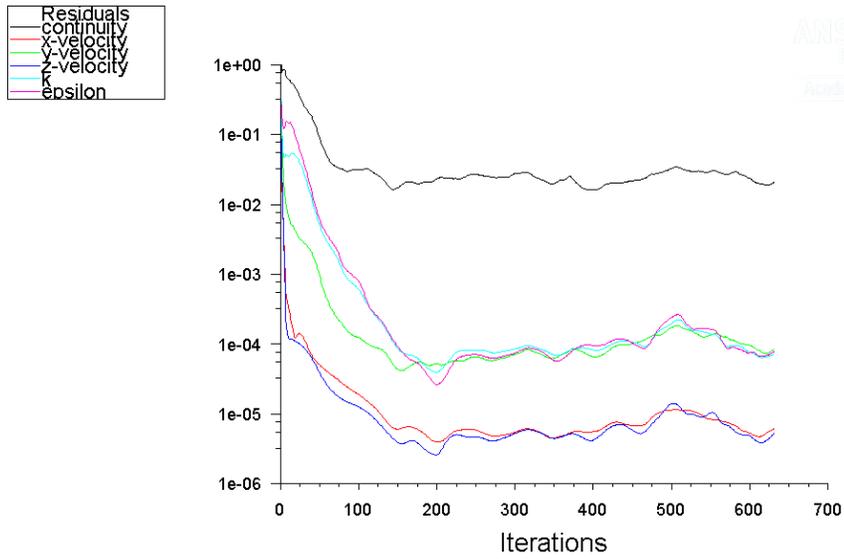


**Figure 11.** Ceiling flow for (top) full ceiling (bottom) drop ceiling simulations.



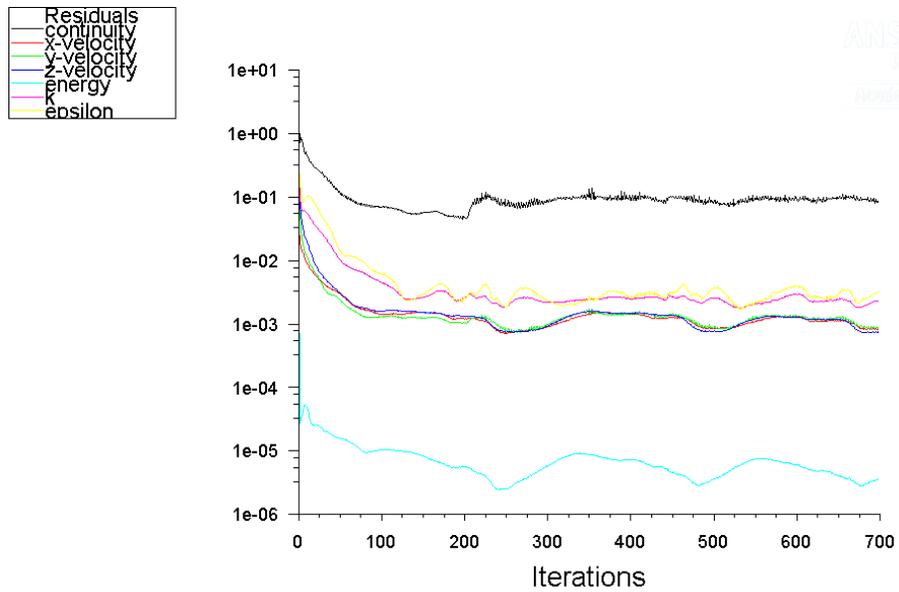
**Figure 12.** Temperature contours for (top) full ceiling (bottom) drop ceiling simulations.

## Convergence Plots



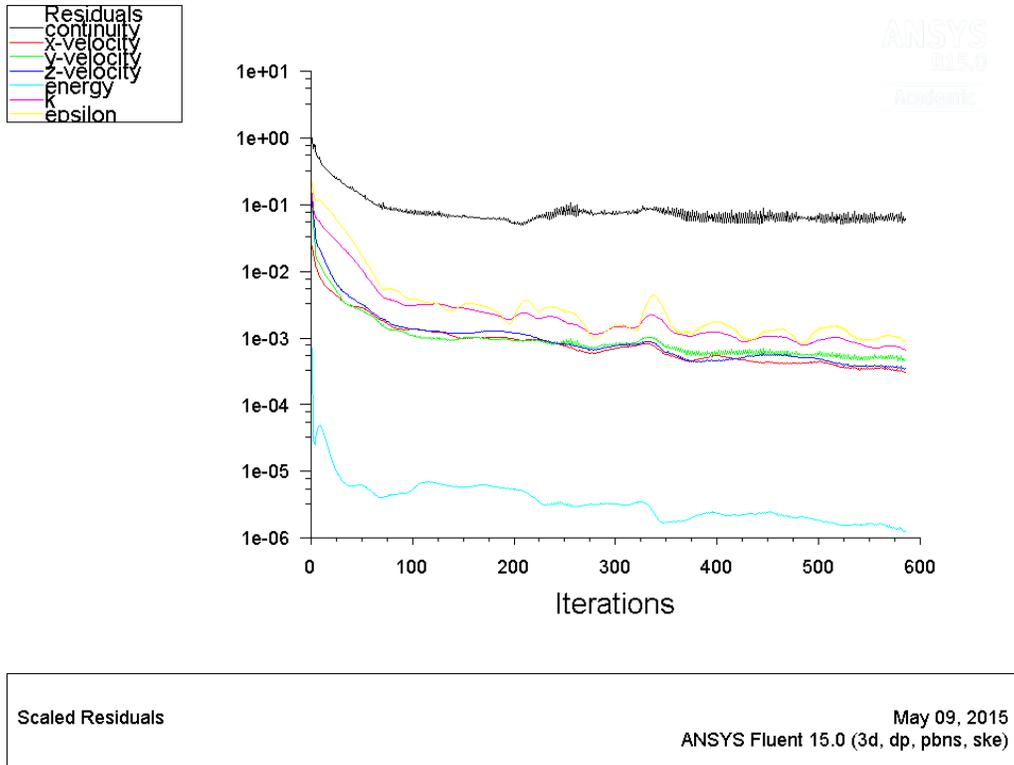
Scaled Residuals May 09, 2015  
ANSYS Fluent 15.0 (3d, dp, pbns, ske)

**Figure 13.** Convergence plot plenum simulation.

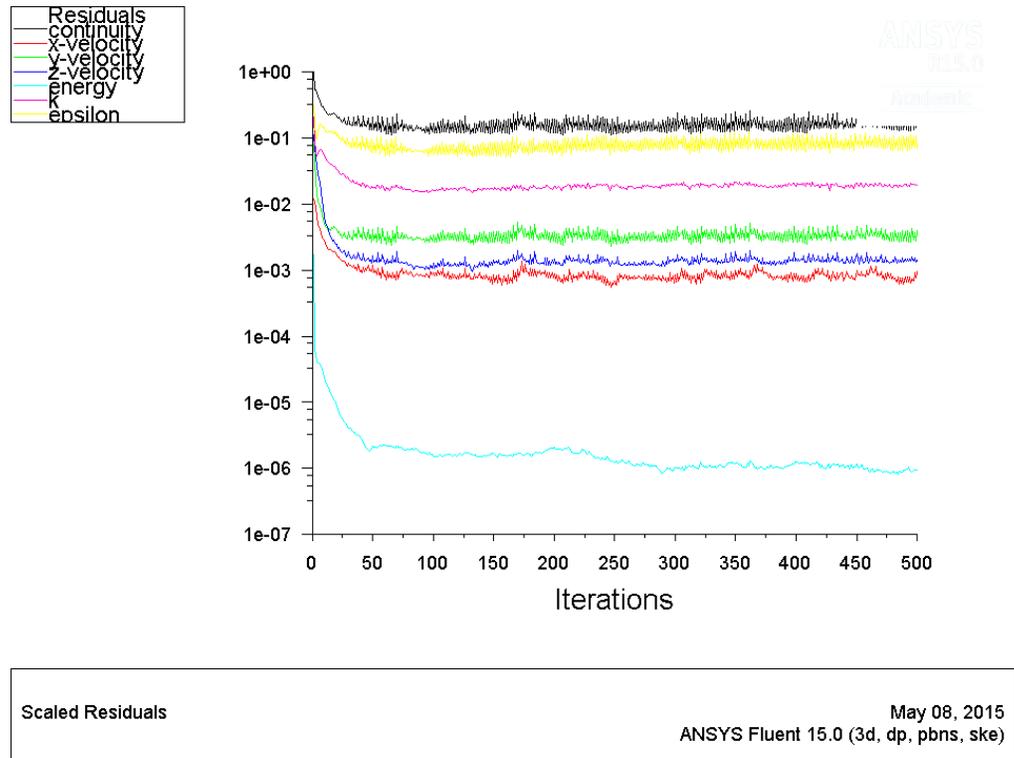


Scaled Residuals May 09, 2015  
ANSYS Fluent 15.0 (3d, dp, pbns, ske)

**Figure 14.** Convergence plot uneven flow distribution full ceiling.



**Figure 15.** Convergence plot even flow distribution full ceiling.



**Figure 16.** Convergence plot uneven flow distribution drop ceiling. Even flow convergence plot was very similar.

## Summary of Findings

The plenum and above floor region are the two regions we chose to simulate for the common raised floor configuration. The perforated tiles of the plenum presented difficulties in modeling pressure drops and velocity profiles across the floor vents. The above floor simulation also proved challenging with respect to thermal modelling of a heat generating source with internal convection. Nonetheless, we were able to show that by treating the perforated tile as a porous volume, we could tune the permeability factor  $\alpha$  to arrive at similar results as given by literature values. Further, the above floor simulations showed that a drop ceiling configuration could eliminate hot air and cold air mixing in above floor spaces, thereby significantly reducing server rack vulnerabilities to top rack vulnerabilities. Making flow uniform out of the perforated tiles, at a provided minimally sufficient value did not increase cooling benefits as much as the drop ceiling solution did. With respect to modeling, we validated the method of breaking the data center into smaller regions as suggested by [7], each with different modeling challenges.

## References

1. Patankar, Suhas V. "Airflow and cooling in a data center." *Journal of Heat Transfer* 132.7 (2010): 073001.
2. Fulpagare, Yogesh, Gaurav Mahamuni, and Atul Bhargav. "Effect of plenum chamber obstructions on data center performance." *Applied Thermal Engineering* 80 (2015): 187-195.
3. Radmehr, Amir, Kailash C. Karki, and Suhas V. Patankar. "Analysis of airflow distribution across a front-to-rear server rack." *ASME 2007 InterPACK Conference collocated with the ASME/JSME 2007 Thermal*
4. Zucker, Robert D., and Oscar Biblarz. *Fundamentals of gas dynamics*. John Wiley & Sons, 2002.
5. Fluent, A. N. S. Y. S. "12.0 Theory Guide." *Ansys Inc* 5 (2009).
6. Rohatgi, A. "Web Plot Digitizer. <http://arohatgi.info/WebPlotDigitizer/app/>" (accessed May 5, 2015).
7. Almoli, Ali, et al. "Computational fluid dynamic investigation of liquid rack cooling in data centres." *Applied energy* 89.1 (2012): 150-155.
8. Hemsoth, Nicole. "Mission Possible – Greening the HPC Data Center. [http://hpcwire.com/2009/08/31/mission\\_possible\\_-\\_greening\\_the\\_hpc\\_data\\_center/](http://hpcwire.com/2009/08/31/mission_possible_-_greening_the_hpc_data_center/)" (accessed May 5, 2015)