CWE-8
June 7-11, 2026
London, ON, Canada

Validation of chiller-to-chiller interactions in datacentre applications through CWE and Experiment

C.Rohr [a], A. Beyer-Lout [b], N. Coste [c], B. Cochran [d]

[a]CPP, Sydney Australia, crohr@cppwind.com
[b]CPP Inc, Windsor Colorado, USA abeyer-lout@cppwind.com
[c]Optiflow (CPP Inc), Marseille, Provence-Alpes-Côte d'Azur, France, ncoste@cppwind.com
[d]CPP Inc, Windsor Colorado, USA bcochran@cppwind.com

SUMMARY

Significant work has been produced over the last decades providing bounds on the accuracy of steady-state turbulence models as applied to external dispersion of emissions in wind flows both in the near and far-field.

However, the rapid expansion of large datacenter projects over the last five years in response to the needs of the AI sector has led to the widespread application of Reynolds Averaged Navier-Stokes (RANS) to the problem of external dispersion and convective heat management from thermal exhausts. This methodology can lead to significant disruption to datacenter operation and an underperforming asset when the design has not been adequately screened for its site and design specific thermal environment.

This paper evaluates the accuracy of such methods against wind tunnel measurements applying best practice methods from an adjacent but more established and validated field: laboratory and industrial exhaust design. The study includes numerical, RANS, wind tunnel, and best-practice LES modelling as applicable to datacenter exhaust dispersion with the aim of providing guidance to the industry on the appropriate methodology for this problem class.

Keywords: *CFD, CWE, Dispersion, RANS, LES, Chillers, Datacenter*

1. INTRODUCTION

Over several decades the reliability of steady state turbulence closures in Computational Fluid Dynamics has been rigorously tested and has set broad expectations for their accuracy when applied to external plumes in atmospheric boundary layer flows. In the present work buoyancy effects are not considered and temperature is treated strictly as a passive scalar. The intent is to assess the accuracy of modelling turbulent diffusion alone. Buoyancy can be introduced in later extensions of this work, but the dominant source of uncertainty at this stage remains the correct representation of scalar mixing. This study also focuses on thermal re-entrainment and does not consider pollutant ingestion, which is an equally important challenge for datacenter design. Many of the findings reported here apply equally to pollutant considerations.

The general accuracy of steady state turbulence closures is limited and can lead to over and under predictions of one or more orders of magnitude in concentration. Figure 1 shows modelled concentrations downstream of a wall mounted cube with a rooftop stack. Although simple, the case requires correct resolution of the rooftop separation bubble and the building wake. In several simulations the plume shape appears visually consistent with experiment, yet the concentrations can diverge substantially because of the exponential decay of concentration away from the source. This test highlights the strong sensitivity to both Schmidt number and turbulence closure and the problem specific nature of these parameters.

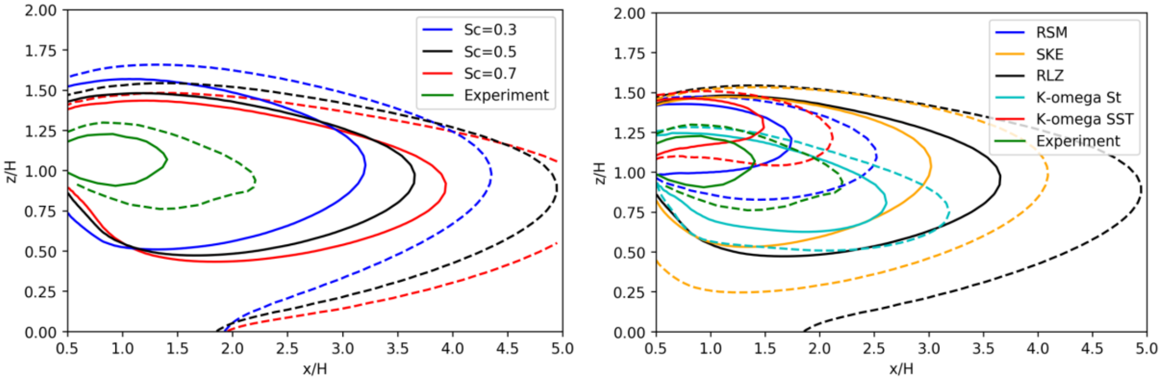

CFD results, Schmidt no var, k-epsilon realizable model.
Normalized gas concentration k.
Solid k=3.0, dashed k=2.0

CFD results, models comparison, Sc=0.5.
Normalized gas concentration k.
Solid k=3.0, dashed k=2.0

Figure 1: Comparison of various Schmidt numbers (L) and RANS turbulence closures (R) with experimental testing

The example in Figure 1 represents only one configuration. The broader literature presents contradictory findings across different geometries, closure models, and Schmidt numbers, which reinforces the strong dependence of this application on local flow features. Beyond predicting concentration at relevant receptors, steady state CFD remains limited by the number of ambient conditions that can be simulated within realistic project schedules. In standard practice hundreds of design cases are required. These include wind directions in increments of 5 degrees or finer and wind speed variations from near calm to strong winds to detect peak concentration at each receptor. Such scanning is necessary for integration with local wind climate statistics and for determining the probability of re entrainment events. Even with modern computation these requirements make comprehensive testing difficult for steady RANS and infeasible for LES.

All three approaches considered in this paper use a concentration based formulation where temperature is treated as a passive scalar. This choice aligns naturally with experimental practice in the wind tunnel where tracer concentration is used as the measurable quantity. The passive scalar method is valid for the present objectives which focus on turbulent diffusion rather than buoyancy driven behaviour. CFD can incorporate buoyancy directly if needed and the wind tunnel can approximate buoyancy through helium substitution or momentum ratio adjustment, but these effects lie outside the scope of this work.

Despite this consistent formulation, the ability to scan across the full range of wind directions and speeds remains the central difference between CFD and wind tunnel practice. The wind tunnel can test up to 72 wind directions and then vary wind speed upward and downward until concentrations are maximised at each receptor. This procedure quickly becomes impractical in CFD. Current datacenter projects therefore simulate only a small set of wind conditions which introduces a substantial risk of missing critical directions or speeds.

The rapid expansion of datacenter development has brought these issues to the forefront. RANS CFD has historically not been considered appropriate for plume dispersion and the established approach for this problem has been wind tunnel testing. The scale and speed of modern datacenter work has driven interest in computational screening despite its limitations. Continued growth in computational capability suggests that full envelope simulation may eventually become viable. This study examines the potential accuracy of these computational methods relative to controlled wind tunnel data ahead of their broad adoption.

2. COMPARISONS

Three approaches are evaluated: RANS, LES, and wind tunnel testing. All use the passive scalar formulation to represent thermal exhaust concentration, and all maintain consistent plume momentum and boundary layer conditions for comparison. The same geometry is tested for all cases and represents a generic datacenter with rooftop chillers. Concentrations are measured both at chiller intakes and elsewhere on the building to represent other types of intakes that can be affected by hot air re-entrainment. A single wind direction and speed is tested.

RANS reflects current industry pressure for rapid, cheap, early stage CFD screening. Two equation turbulence models such as k–ε and k–ω describe the momentum field while the scalar field is transported using an assumed turbulent Schmidt number. This parameter dictates the predicted entrainment and dilution of the plume. The passive scalar framework provides a consistent comparison with experiment, but the accuracy of scalar diffusion depends strongly on the turbulence model and the chosen Schmidt number. RANS remains steady and therefore cannot represent unsteady engulfment or intermittent recirculation.

While RANS setup is relatively straightforward for experienced computational wind engineers, simulating the turbulence produced motion in the context of ABL modelling requires specific care when LES is chosen, and these techniques are demonstrated. LES resolves the large turbulent structures that dominate scalar transport, giving access to both the peaks of concentration and the and its probability distribution. Temperature is also treated as a passive scalar and the subgrid model adjusts diffusivity based on local strain which provides a more realistic description of mixing than a fixed Schmidt number.

The wind tunnel provides the experimental baseline. Tracer concentration is used to represent temperature which aligns directly with the passive scalar formulation in CFD. The tunnel can examine the full directional and speed envelope that is required to identify peak concentrations and to integrate results with local wind climate. This sweep is conducted for the experimental testing to illustrate the wide variety of possible conditions which can be undesirable for datacenter operation, some of which would be missed by a limited sweep conducted computationally.

3. CONCLUSION

The primary goal of this work is to provide a stronger foundation to the work currently being conducted in response to the rapid increase in datacenter construction which is at risk of significant thermal mismanagement and subsequent under performance should the industry continue to rely on a limited set of simulations to validate its designs.

In this study a comparison between two numerical and one experimental technique has been conducted specifically as they apply to the field of modelling near-field dispersion around datacentres. The comparison is twofold: first the relative accuracy of computational and experimental methods is quantified. Secondly, the requirement for a full screening of all atmospheric conditions (wind directions, speeds, exhaust conditions) combined with the local wind climate is illustrated.

REFERENCES

Gousseau, P., Blocken, B., & Van Heijst, G. J. F. (2011). CFD simulation of pollutant dispersion around isolated buildings: On the role of convective and turbulent mass fluxes in the prediction accuracy. *Journal of Hazardous Materials*, *194*, 422-434.

Gousseau, P., Blocken, B., Stathopoulos, T., & Van Heijst, G. J. F. (2011). CFD simulation of near-field pollutant dispersion on a high-resolution grid: A case study by LES and RANS for a building group in downtown Montreal. *Atmospheric Environment*, *45*(2), 428-438.

Snyder, W. H. (1981). *Guideline for fluid modeling of atmospheric diffusion* (Vol. 81, No. 9). Environmental Sciences Research Laboratory, Office of Research and Development, US Environmental Protection Agency.

Stathopoulos, T., Hajra, B., Chavez, M., & Bahloul, A. (2014). A wind tunnel study of the effect of adjacent buildings on near-field pollutant dispersion from rooftop emissions.

Tominaga, Y., & Stathopoulos, T. (2007). Turbulent Schmidt numbers for CFD analysis with various types of flowfield. *Atmospheric environment*, *41*(37), 8091-8099.

Tominaga, Y., Iizuka, S., Imano, M., Kataoka, H., Mochida, A., Nozu, T., ... & Yoshie, R. (2013). Cross comparisons of CFD results of wind and dispersion fields for MUST experiment: evaluation exercises by AIJ. *Journal of Asian Architecture and Building Engineering*, *12*(1), 117-124.