CFDSC Bulletin

CFD Society of Canada | Société Canadienne de CFD

June. 2006 | Issue No. 17

In this Issue

PRESIDENT'S MESSAGE
by
Mahmood Khalid

The second Bulletin of the CFD Society of Canada is in your hands. We have received numerous congratulatory remarks from members for reviving the Bulletin. This year's annual conference of the CFD Society, as you know, will take place at Queen's University in Kingston. It is being organized by Prof. Andrew Pollard, who, as many of the earlier members might recall, was actually the founding president of this society. From all indications, it appears that this year's conference with its brand new CFD challenge and associated course will be a truly rewarding experience for the participants. So whether you have submitted a paper to the conference or not, make a point to attend the conference and have the opportunity network with your CFD colleagues from Canada and indeed from all over the world. Remember that we have had participants in the past from as far away as Korea, Iran, France, Germany, Malaysia, the United States, and so on. You will also have the opportunity to listen to invited talks and interact with leading experts in various CFD domains. From time to time, your CFD Society also recognizes the contributions of outstanding individuals in the form of awarding a Lifetime Achievement Award. Last year, as many of you will remember, this award was given to Dr. George Raithby. Although this is not meant to be a regular annual event, following the recommendation of the Awards Committee, we are making a rare exception and honoring another outstanding individual in our society with the same award at the CFD 2006 conference. You would also be pleased to learn that the CFD Society has evolved a very special relationship with the publishers of the International Journal of Computational Fluid Dynamics (IJCFD) edited by Dr. Fred Habashi. As for the last year, IJCFD has issued a special edition in February 2006 with papers selected from the CFD 2004 conference.

CFDSC Bulletin

CFD Society of Canada | Société Canadienne de CFD | June. 2006 | Issue No. 17

CFD Society of Canada

Objectives ● The Computational Fluid Dynamic Society of Canada (CFDSC) promotes the application of computational fluid dynamics (CFD) to provide a better understanding of fluid dynamic processes in all applicable areas and thereby supports the competitiveness of Canadian industry and advances knowledge in the field. These objectives are met by *(1)* establishing a network of CFD practitioners and developers from industry, government and universities, *(2)* identifying national needs and priorities for research in CFD, *(3)* promoting research in CFD and related areas, *(4)* enhancing industry capabilities in CFD by the education and training of professionals, *(5)* promoting the use and increasing the availability of computational tools for CFD, *(6)* promoting communication and exchange within CFD and related disciplines (e.g. environment, biology and chemistry), *(7)* organizing conferences and seminars, and *(8)* representing Canada at international forums.

Membership Annual membership fees of the CFDSC are $50 for regular members, $20 for student members, and $25 for members who are retired.

World Wide Web | www.cfdsc.ca

Address

CFD Society of Canada
P. O. Box 4871, Station E
Ottawa, Ontario, Canada KIS 5J1

Bulletin The Bulletin of the CFDSC is published about three times a year to provide members with information on current CFD issues. The published articles express the views and opinions of their respective authors and do not necessarily represent those of the CFDSC. The Managing Editors are:

Mahmood Khalid | NRC-IAR
Mobina Haider | Ottawa, Canada

mobina@inbox.com

Société Canadienne de CFD

Objectifs ● La Société canadienne de CFD (SCCFD) fait la promotion des applications de <<computational fluid dynamics>> afin de fournir une meilleure compréhension de la dynamique des fluides et contribue ainsi au développement de la compétitivité industrielle canadienne et à l'avancement des connaissances dans le domaine. Ces objectifs sont atteints en *1)* établissant un réseau d'utilisateurs et de développeurs de CFD dans les industries, les gouvernements et les universités, *2)* identifiant les besoins nationaux et les priorités de recherches en CFD, *3)* faisant la promotion de la recherche en CFD et les domaines reliés, *4)* rehaussant les capacités en CFD des industries par la formation et l'éducation des professionnels, *5)* faisant la promotion de l'utilisation et en augmentant la disponibilité d'outils informatiques en CFD, *6)* faisant la promotion de la communication et des échanges en CFD avec les disciplines reliées (par example, en environnement, en biologie et en chimie), *7)* organisant des conférences et des séminaires, *8)* représentant le Canada dans les forums internationaux.

Cotisation des membres Les frais annuels d'adhésion sont de 50$ pour les membres réguliers, 20$ pour les membres étudiants, et 25$ pour les membres à la retraite.

World Wide Web | www.cfdsc.ca

Adresse

Société canadienne de CFD
B. P. 4871, Station 'E'
Ottawa, Ontario,Canada K1S 5J1

Bulletin Le bulletin de la SCCFD est publié environ trois fois par année afin de fournir aux membres de l'information sur les grands courants reliés à la CFD. Les articles publiés représentent les opinions des auteurs respectifs et ne représentent pas nécessairement celles de la SCCFD. Les rédacteurs exécutifs sont:

Mahmood Khalid | CNRC-IAR
Mobina Haider | Ottawa, Canada

CFD 2006 | Current Status

CFD 2006 will be held on the campus of Queen's University in Kingston July 16-18, 2006. The conference is being supported by Sun Microsystems, IBM, MAYA Heat Transfer Technologies Ltd., Bombardier, Liquid Computing, Ensigh and Queen's University.

The conference organisation, which began in the Spring of 2005, got off to a sluggish start because the organisers decided to convert all operations to be web-based using a dedicated CFDSC server. For the most part, we have resolved most of the major logistic issues so that future organisers will have a set of rock-solid templates to handle the myriad interactions between the conference and the delegates, submission of paper and the allocation of reviewers to each paper and handling the receipt of the reviews to the authors. We have created a reviewer data base that we hope will be added to by other organisers, and a mailing list of possible delegates that spans many countries around the world. Thus, CFD 2006 will be the beginning, we hope, of a more international delegate base.

The invited speakers are superb: Fred Habashi from McGill and Ned Djilali from Victoria require no introduction to the CFDSC community. Gordon Mallinson from the University of Auckland and Roger Davis from the University of California at Davis are the other two invited speakers. Gordon is an expert in graphical visualisation methods and Roger is an expert in gas turbine simulations. Together, these four exceptionally qualified individuals will bring the latest advances in CFD technologies to the delegates.

The conference will feature a number of social events, including a Sunday afternoon/evening reception and a Monday evening banquet in the form of a 3-hour dinner cruise on a unique 'Bateau Mouche' style boat that will tour through the famous 1000 Islands followed by a late night "knee's up", with live music, at the Merchant Tap House pub. On Wednesday there will be a short course (additional cost) on Lattice Boltzmann methods.

The organisers look forward to welcome you to Kingston and to CFD 2006.

For more details, see **http://www.cfd2005.org**

CFDSC
Executive Committee & Directors

The Blue Gene/L Supercomputers
Top Performers in 2004 & 2005

Jose Castanos, George Chiu, Paul Coteus, Alan Gara, Manish Gupta
IBM T.J. Watson Research Center
Yorktown Heights, New York

Introduction

When we build a supercomputer with thousands to more than a hundred of thousands of chips, is it better to choose a few mighty high frequency processors, or is it better to start from many small low power processors to achieve the same computational capability? To answer this question, let us trace the evolution of computers.

The first general purpose computer, ENIAC (Electronic Numerical Integrator And Calculator), was publicly disclosed in 1946. It took 200 microseconds to perform an addition on ENIAC. ENIAC was built with 19,000 vacuum tubes. The machine was enormous, 30 m long, 2.4 m high and 0.9 m wide. Vacuum tubes had a limited lifetime and had to be replaced often. The system consumed 200 kW. ENIAC's cost the US Ordnance Department $486,804.22.

In December 1947, John Bardeen, Walter Brattain, and William Shockley at Bell Laboratories invented a new switching technology called the transistor. This device consumed less power, occupied less space, and was made more reliable than those of vacuum tubes. Impressed by these attributes, IBM built its first transistor based computer called Model 604 in 1953. By early 1960, transistor technology became ubiquitous. Further drive towards lower power, less space, higher reliability, and lower cost resulted in the invention of integrated circuits in 1959 by Jack Kilby of Texas Instruments. Jack Kilby made his first integrated circuit in germanium. Robert Noyce at Fairchild used a planar process to make connections of components within a silicon integrated circuit in early 1959; which became the foundation of all subsequent generations of computers. In 1966, IBM shipped the System/360 all-purpose mainframe computer, made of integrated circuits.

Within the transistor circuit families, the most powerful transistor technology was the bipolar junction transistor (BJT) technology rather than CMOS (Complementary Metal Oxide Semiconductor) transistor technology. However, compared to CMOS transistors, the bipolar ones, using the fastest ECL (emitter coupled logic) circuit, cost more to build, had a lower level of integration, and consumed more power. As a result, the semiconductor industry moved en masse to CMOS in early 1990s. From then on, the CMOS technology became the entrenched technology, and supercomputers were built with the fastest CMOS circuits. This picture lasted until about 2002 where CMOS power and power density rose dramatically to the point that they exceeded those of the corresponding bipolar numbers in the 1990's. Unfortunately, there was no lower power technology lying in wait to diffuse the crisis. Thus, we find ourselves again at a crossroad to build the next generation supercomputer. The "traditional" view to build the fastest and largest supercomputer is to use the fastest microprocessor chips as the building block. The fastest microprocessor is in turn built upon the fastest CMOS transistor switching technology that is available to the architect at the time of designing the chips. This line of thought is sound provided that there are no other constraints to build supercomputers. In the real world, there are many constraints (heat, component size…) that change this traditional straightforward reasoning.

In the mean time, portable devices such as PDAs, cellphones, and laptop computers, developed since the 1990's, all require low power CMOS technology to maximize the battery recharge interval. In 1999, IBM foresaw the looming power crisis, and asked the question whether we could architect

supercomputers using low power, low frequency, and inexpensive embedded processors to achieve a better collective performance than using high power, high frequency processors. While this approach has been successfully utilized for special purpose machines such as the QCDOC supercomputer, this counter-intuitive proposal was a significant departure from the traditional approach to supercomputer designs. However, the drive toward low power, and low cost remained a constant theme throughout.
We chose an embedded processor optimized for low power and low frequency design, rather than performance. Such a processor has a performance/power advantage compared to a high performance and high power processor. A simple relation is

performance/rack = performance/watt x watt/rack.

The last term in this expression, watt/rack, is determined by thermal cooling capabilities of a given rack volume. Therefore, it imposes the same limit (of the order of 25 kilowatts) for using either high- frequency, high-power chips or using low-frequency, low-power chips. To maximize the performance of rack, it is the performance/watt term that must be compared among different CMOS technologies. This clearly illustrates one of the areas in which electrical power is critical to achieving rack density. We have found that in terms of performance/watt, the low frequency, lower power embedded IBM PowerPC 440 core consistently outperforms high frequency, high power microprocessors by a factor of about 10 regardless of the manufacturers of the systems. This is one of the main reasons we chose the low power design point for our Blue Gene/L supercomputer. Figure 1 illustrates the power efficiency of some recent supercomputers. The data is based on total peak giga floating-point operations per second divided by total system power in watt, when that data is available. If the data is not available, we approximate it using Gflops/chip power (an overestimate of the true system Gflops/power number).

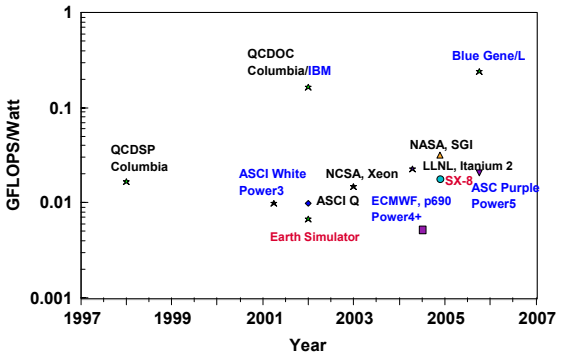

Figure 1 Power efficiencies of recent supercomputers. (Blue = IBM Machines, black = other U.S. machines, red = Japanese machine)

This chart presents empirical evidence of the fact that in the presence of a common power envelope, the collective peak performance per unit volume is superior with low- power CMOS technology. We now explain the theoretical basis of the superior collective performance of low power systems. Any performance metric such as flops (floating point operations per second), MIPS (millions instructions per sec), or SPEC benchmarks is linearly proportional to the chip clock frequency. On the other hand, the power consumption of the ith transistor is given by the expression:

P_i = switching power of transistor i + leakage power of transistor i
= ½ C_{Li} V^2 f_i + leakage power of transistor i,

where C_{Li} is the load capacitance of the ith transistor, V = VDD is the supply voltage, and f_i is the switching frequency of the ith transistor. Note that not every transistor participates in switching on every clock cycle f. Although the leakage power is increasingly important for 90nm, 65nm and 45nm technologies, we ignore the leakage power of the Blue Gene/L chips which, built in 130 nm technology, contributes less than 2% of the system power. The

switching power consumed in a chip is the sum of the power of all switching nodes. It can be expressed as:

P_{chip} = Σ switching power of transistor i = ½ C_{sw} V^2 f,

where the average switching chip capacitance is given by

C_{sw} = (Σ C_{Li} f_i)//f.

It is difficult to predict C_{sw} accurately because we seldom know the switching frequencies f_i of every transistor in every cycle, and furthermore f_i is different for each application. To simplify the discussion, we use an averaged value of C_{sw} obtained either from direct measurement or from power modeling tools. For high power, high frequency CMOS chips, the clock frequency f is roughly proportional to the supply voltage V, thus the power consumed per chip P_{chip} is proportional to V^2 f or f^3. Therefore, in the cubed-frequency regime, the power grows by a factor of 8, if the frequency is doubled. If we use 8 moderate frequency chips, each of them half the frequency of the original high frequency chip, we burn the same amount of power, yet we have a fourfold increase in flops/watt. This then is the basis of our Blue Gene/L design philosophy. One might ask if we can do this indefinitely. If 100,000 processors at some frequency is good, are not 800,000 processors at ½ the frequency even better? The answer is complex, because we must consider also the mechanical component sizes, power to communicate between processors, the failure rate of those processors, the cost of packaging those processors, etc. Blue Gene/L is a complex balance of these factors and many more. Moreover, as we lower the frequency, the power consumed per chip drops from cubic frequency dependence, quadratic dependence and finally to linear dependence. In the linear regime, both power and performance are proportional to frequency; there is no advantage in reducing frequency at that point.

Blue Gene/L Architecture

The Blue Gene/L supercomputer project is aimed to push the envelope of high performance computing (HPC) to unprecedented levels of scale and performance. The Blue Gene/L is the first supercomputer in the Blue Gene family. It consists of 65,536 high-performance compute nodes (131,072 processors), each node is an embedded 32-bit PowerPC dual processor. It has 33 Terabytes of main memory. Furthermore, it has 1024 I/O nodes, using the same chip that is used for compute nodes. A three-dimensional torus network and a sparse combining network are used to interconnect all nodes. The Blue Gene/L networks were designed with extreme scaling in mind. Therefore, we chose networks that scale efficiently in terms of both performance and packaging. The networks support very small messages (as small as 32 bytes) and include hardware support for collective operations (broadcast, reduction, scan, etc.), which will dominate some applications at the scaling limit. The compute nodes are designed to achieve a 183.5 Teraflops/s peak performance in the co-processor mode, and 367 Teraflops/s in the virtual node mode [1].

The system on chip approach used in the Blue Gene/L project integrates two processors, cache (Level 2 and Level 3), internode networks (torus, tree, and global barrier networks), JTAG and Gigabit Ethernet links on the same die. By using the embedded DRAM, we have enlarged the on-chip Level 3 cache to 4 MB, 4 to 8 times larger than competitive cache's made of SRAM and greatly enhancing the amount of realized performance of the processor. By integrating the inter-node networks, we can take advantage of the same generation technology, i.e., these networks scale with chip frequency. Furthermore, the off-chip drivers and receivers can be optimized to consume less power than those of industry standard networks. Figure 2 is a photograph of multi-rows of the Blue Gene/L system.

Figure 2. The Blue Gene/L system installed at the Lawrence Livermore National Laboratory.

One of the key objectives in the Blue Gene/L design is to achieve cost/performance on a par with the COTS (Commodity Off The Shelf) approach, while at the same time incorporate a processor and network design so powerful that it can revolutionize supercomputer systems.

Using many low power, power-efficient chips to replace fewer, more powerful ones succeeds only if the application users can realize more performance by scaling up to a higher number of processors. This indeed is one of the most challenging aspects of the Blue Gene/L system design and must be addressed through scalable networks along with software that will efficiently leverage these networks.

System Software

The system software for Blue Gene/L was designed with two key goals, familiarity and scalability. We wanted to make sure that high performance computing users could migrate their parallel application codes with relative ease to the Blue Gene/L platform. Secondly, we wanted the operating environment to allow parallel applications to scale to the unprecedented levels of 64K nodes (128K processors). It is important to note that this requires scaling not only in terms of performance but also in reliability. A simple mean-time-between-failure calculation shows that if the software on a compute node fails about once a month, under the assumption that failures over all nodes are independent, a node failure would be expected once every 40 seconds! Clearly, this shows the need for compute node software to be highly reliable.

We have developed a programming environment based on familiar programming languages (Fortran, C, and C++) and the single program multiple data (SPMD) programming model, with message passing supported via the message passing interface (MPI) library. This has allowed the porting of several large scientific applications to Blue Gene/L with a modest effort (often within a day).

We have relied on simplicity and a hierarchical organization to achieve scalability of software in terms of both performance and reliability. Two major design simplifications that we have imposed are:

- Strictly space sharing: only one parallel job can run at a time on a Blue Gene/L partition; we go one step further and support only one thread of execution per processor. This allows us to use efficient, user-space communication without protection problems (the Blue Gene/L partitions are electrically isolated). Furthermore, having a dedicated processor behind every application-level thread leads to more deterministic execution and higher scalability.

- No demand paging support: the virtual memory available on a node is limited to the physical memory size. This restriction, besides simplifying the compute node kernel, leads to a performance benefit that there are no page faults or translation lookaside buffer misses during program execution, leading to higher and more deterministic performance.

The software for Blue Gene/L is organized in the form of a three-tier hierarchy. A lightweight kernel, together with the runtime library for supporting user applications, constitutes the programming environment on the compute node. Each I/O node, which can be viewed as a

parent of a set of compute nodes (referred to as a processing set or pset), runs Linux, and supports a more complete range of operating system services, including file I/O and sockets, to the applications via offloading from the compute nodes. The Linux kernel on I/O nodes also provides support for job launch. Finally, the control system services run on a service node, which is connected to the Blue Gene/L computational core via a control network.

Results

In October 2004, an 8-rack Blue Gene/L system, which occupied less than 200 square feet of floorspace, and consumed about 200 KW in power, passed the Earth Simulator (which occupies an area of about 70,000 square feet, and consumes about 7 MW of power) in LINPACK performance. In the June 2005 TOP500 list [2], a 32 rack Blue Gene/L system, which was delivered to Lawrence Livermore National Laboratory, took the #1 spot with a LINPACK performance of 136.8 Teraflop/s. Blue Gene/L systems account for five of the top ten entries in that list. In the November 2005 TOP500 list [2], a 64 rack Blue Gene/L system, which has been accepted by Lawrence Livermore National Laboratory, again achieves the #1 spot with a LINPACK performance of 280.6 Teraflops/s, distancing the winning margin over all the competitors. In addition, the Blue Gene/L swept the competition of Class I High Performance Computing Challenge benchmarks.

More importantly, many scientific applications in the fields of classical molecular dynamics, quantum molecular dynamics, and computational fluid dynamics (CFD) have been successfully ported and scaled on the Blue Gene/L system. In particular, the 2005 Gordon Bell went to a classical molecular dynamics code called ddcMD (domain decomposition classical molecular dynamics code) [3], achieving 101.7 teraflops on 131,072 cpu's. This is the first time that >100TF was sustained on a general purpose platform. This is a 524 million atom phase change simulation, making it relevant for CFD problems as well. Another example is the porting of Miranda [4], a massively parallel spectral/compact solver for variable density incompressible flow, including viscosity and species diffusivity effects. Weak scaling to 65,536 processors is virtually perfect. Supernova explosion code FLASH with adaptive mesh refinement, weather code WRF, climate code HOMME, and CFD code Overflow were also scaled well on Blue Gene. Those results also represent the first proof point that MPI applications can effectively scale to over ten thousand processors.

Conclusions

In this paper, we described the main thrust of the Blue Gene/L supercomputer made of low power, low frequency processors. By exploiting the superior performance/watt metric, we can package 10 times more processors in a rack, thus it became the number 1 rated supercomputer since November 2004. In November 2005, three of the top ten supercomputers, including number 1 and number 2 spots, in the 26th TOP500 list were based on Blue Gene/L architecture. Blue Gene/L is currently producing unprecedented simulation in classical and quantum molecular dynamics, climate, weather, CFD, quantum chromodynamics, and the list is growing. The future is likely to be even more power constrained due to the slowing of the power-performance scaling of the underlying transistor technologies. This will likely drive systems to aggressively search for opportunities to build even more power efficient systems, likely driving to more Blue Gene/L-like parallelism. In the future, the low-power processors are likely to be active in nearly every area of computing.

Acknowledgement

The Blue Gene/L project has been supported and partially funded by the Lawrence Livermore National Laboratory on behalf of the United States Department of Energy under Lawrence

Livermore National Laboratory Subcontract No. B517522.

References

[1] IBM Journal of Research and Development, special double issue on Blue Gene, Vol.49, No.2/3, March/May, 2005.
[2] TOP500 Supercomputer Sites, http://www.top500.org.
[3] F. Streitz et al. 100+ Tflop Solidification Simulations on BlueGene/L, In Supercomputing 2005, Paper 307.
[4] Andrew W. Cook et al. Tera-Scalable Algorithms for Variable-Density Elliptic Hydrodynamics with Spectral Accuracy, In Supercomputing 2005, Paper 139.

Application of CFD to Security Science

Progress on the Development of a High-Fidelity Numerical Model for Hazard Prediction and Assessment in the Urban Environment.

fslien@uwaterloo.ca

F.S. Lien[1], E. Yee[2], H. Ji[3], A. Keats[1] and K.J. Hsieh[1]
[1] Department of Mechanical Engineering, University of Waterloo, Waterloo, Ontario, N2L 3G1, Canada
[2] Defence R&D Canada - Suffield, P.O. Box 4000, Medicine Hat, Alberta, T1A 8K6, Canada
[3] Waterloo CFD Engineering Consulting Inc., Waterloo, Ontario, N2T 2N7, Canada

1. Introduction

Atmospheric transport and diffusion models have played an important role in emergency response systems. These models have been developed to calculate the transport, diffusion, and deposition of contaminants released (either accidentally or deliberately) into the turbulent atmospheric boundary layer over relatively smooth and homogeneous surfaces. While this approach is reasonably well understood and useful for a landscape that is approximately flat and unobstructed, it is wholly inadequate for surface-atmosphere interactions over “complex” surfaces (i.e., most of the real world) such as cities and other built-up areas.

In the urban environment, dispersion is modified in a complex and non-trivial manner by interaction of the highly disturbed flow field and contaminant plume with the buildings. Consequently, a similar understanding of urban dispersion is currently lacking, in spite of the fact that it is the urban environment where human habitation is concentrated. The availability of atmospheric dispersion models for the urban environment is critically important to the security of Canada and its citizens owing to the fact that the release of a chemical, biological, radiological, or nuclear (CBRN) agent by terrorists or rogue states in a North American city (densely populated urban centre) and the subsequent exposure, deposition, and contamination is an emerging threat in an uncertain world. In this regard, the availability of a high-fidelity, time-dependent model for the prediction of a CBRN agent’s movement and fate in a complex urban environment can provide the strongest technical and scientific foundations for support of Canada’s more broadly based effort at advancing counter-terrorism planning and operational capabilities aimed at countering potential CBRN threats, and ensuring the security and prosperity of the nation.

To address the urgent problem of modeling of the dispersion of CBRN agents in the urban complex, characterized by extremely diverse length and time scales and complex geometries and interfaces, we require physically-based urban wind models that will be able to provide the needed spatial pattern of urban wind statistics, as well as dispersion models that can predict the transport and turbulent diffusion of contaminants released into this highly disturbed wind field. In this study, we report progress on the development of a high-fidelity multi-scale, multi-physics modeling system for the accurate

and efficient prediction of urban flow and dispersion of CBRN materials.

This article is organized as follows. Section 2 summarizes the dispersion models used to predict both continuous and instantaneous releases of contaminant material. Section 3 addresses work in progress which is summarized in two sub-sections which deal, respectively, with the inverse source determination problem and the coupling of urbanSTREAM with a mesoscale flow model. Finally, Section 4 contains conclusions.

2. Dispersion Modeling

The concentration, $\overline{c}$, and the concentration variance, $\overline{c'^2}$, are obtained from their transport equations (see, e.g., [1]). For example, the transport equation for $\overline{c'^2}$ can be written as

$$\frac{\partial \overline{c'^2}}{\partial t} + \frac{\partial}{\partial x_j}\overline{u}_j \overline{c'^2} = \frac{\partial}{\partial x_j}\left(D\frac{\partial \overline{c'^2}}{\partial x_j} - \overline{u_j' c'^2} \right) - 2\overline{u_j' c'} \frac{\partial \overline{c}}{\partial x_j} - \varepsilon_c, \quad (1)$$

where D is the (kinematic) molecular diffusivity of the scalar. The critical term in Eq. (1) is ε_c which represents the dissipation of $\overline{c'^2}$ by molecular diffusion in the fine-scale scalar structure. This is a term responsible for micromixing of the scalar and, as such, corresponds to a small-scale term controlled by the scalar gradient correlations. This term will be modeled algebraically using an expression with the general form $\varepsilon_c = \overline{c'^2}/t_d$, where t_d is an appropriately defined dissipation time scale that is characteristic of the decay time of the concentration fluctuations in the scalar field. In the present study, two models are used to describe the scalar dissipation time scale.

Model 1: In this model,

$$\varepsilon_c = C_\chi \frac{\varepsilon}{k}\overline{c'^2}, \quad (2)$$

where C_χ is a model constant that represents the ratio of the turbulence integral time scale to the scalar dissipation time scale. In the present study, we will use $C_\chi = 2$, as suggested by [2] in our simulations of the scalar variance field.

Model 2: In this model,

$$\varepsilon_c = C_\chi \frac{k^{1/2}}{\Lambda_d}\overline{c'^2}, \quad (3)$$

where Λ_d is the dissipation length scale in the dispersing plume or cloud (puff). Equivalently, in this model, t_d is proportional to $\Lambda_d/k^{1/2}$. In the present study, we propose to use a simple algebraic formulation for the dissipation length scale; namely,

$$\Lambda_d = [\sigma_y(x)\sigma_z(x)]^{1/2} \quad \text{(for a plume)}, \quad (4)$$

and

$$\Lambda_d = [\sigma_x(t)\sigma_y(t)\sigma_z(t)]^{1/3} \quad \text{(for a cloud or puff)}, \quad (5)$$

where σ_x, σ_y, and σ_z are the plume or puff spreads (standard deviations) of the mean concentration distribution in the streamwise x, vertical y, and crosswind z directions, respectively. The closure constant C_χ is equal to 1.25 and 2.1 for a plume and cloud (puff), respectively. These values for C_χ were chosen to give the best conformance with the available data for the MUST array, and in particular, were chosen to ensure the best fit to the decay of the plume or cloud centreline fluctuation intensity

(ratio of the concentration standard deviation to the mean concentration) with downstream distance from the source.

Fig. 1 shows the geometry of the Mock Urban Setting Test (MUST) array [3] and source location. This full-scale atmospheric experiment was replicated at a scale of 1:205 in a boundary-layer water channel by Coanda Research & Development Corporation (Burnaby, BC, Canada) [4], [5]. Sample results for the instantaneous release (puff) case only, obtained with a $203\times45\times43$ grid, are given in Fig. 2 in terms of the vertical profiles of streamwise velocity $\overline{u}$ at $z=0$ (vertical centre plane of the array) and the horizontal profiles of total dosage at $y/H=0.75$ (H is the height of the obstacles) at the 6.5 row location. Figure 2 also shows time histories of normalized mean concentration $(\overline{c}/c_s)$ and normalized concentration standard deviation (c'_{rms}/c_s) profiles at $z/H=0,\ y/H=0.75$ at the 6.5 row location. Here, c_s is the source concentration. The total dosage is defined as

$$\text{dose}=\int_0^T \overline{c}dt, \tag{6}$$

where T is a time interval that is chosen to be sufficiently long so as to include the duration (time interval) between the arrival time and departure time of the cloud (puff) at the given receptor location where the dosage is calculated (or, measured). As seen, excellent agreement between the predicted and measured $\overline{u}$ -profiles is achieved. Figure 2 also shows that the time history of $\overline{c}/c_s$ and the profile of horizontal dosage are reasonably well predicted by solving a simple advection-diffusion equation. In the case of the c'_{rms}/c_s prediction, Model 2 is clearly better than Model 1 in its prediction for the concentration standard deviation.

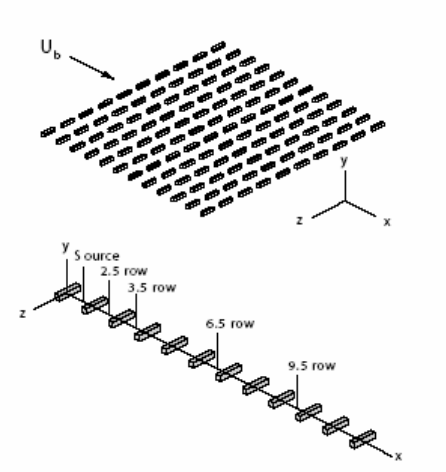

Figure 1: A 3-D perspective view of the MUST obstacle array.

3. Work In Progress

3.1 Inverse Source Determination

3.1.1 Mathematical Formulation

The inverse source determination problem is generally an ill-posed problem with no unique solution. To solve this ill-posed problem, it is necessary to formulate an approach that (1) incorporates *a priori* information on the possible solutions in a cogent manner and (2) incorporates quantitatively the lack of total precision in the problem due to measurement noise and model (and input) uncertainties. Hence, the inverse source problem is a problem of inductive logic (or inference) rather than deductive logic. Probability theory when interpreted as logic is a quantitative theory of inference, just as mathematics is a quantitative theory of deduction.

The rational framework for the formulation of the inverse source problem is to apply Bayesian probability theory. To this purpose, the application of Bayes' theorem to the inverse source determination problem takes the following form:

$$p(m \mid D, I) \propto \underbrace{p(m \mid I)}_{\text{prior}} \underbrace{p(D \mid m, I)}_{\text{likelihood}}, \quad (7)$$

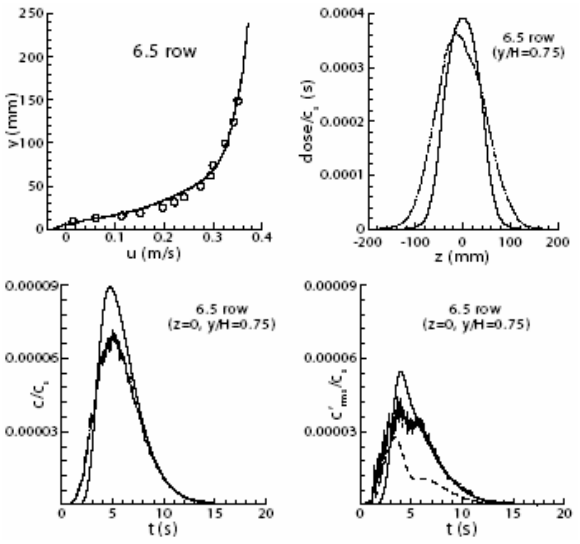

Figure 2: Vertical velocity profiles at *z/L*=0, horizontal dosage profiles at *y/H* = 0.75 and time histories of $\bar{c}/c_S$ and c'_{rms}/c_S, at *z/L* = 0, *y/H* = 0.75 at the 6.5 row location. For the top two panels and the bottom left panel:— · —, ○ Experimental data;—— Model prediction. For the bottom right panel: —·— Experimental data; – – – – Model prediction for concentration standard deviation using Model 1 for ε_c;—— Model prediction for concentration standard deviation using Model 2 for ε_c. Here, *L* and *H* are the streamwise length and height of the obstacles, respectively.

where $p(m \mid D, I)$ is the posterior probability density function (PDF) for m (source function), I denotes all the background (or, contextual) information available to the observer (e.g., meteorology, position of detectors, dispersion model), and D are the concentration data provided by the detectors. The posterior PDF for m is proportional to the product of the prior PDF for m, $p(m \mid I)$, and the likelihood function, $p(D \mid m, I)$. The prior PDF $p(m \mid I)$ encodes all the information available about the putative source before the receipt of the concentration data D. The likelihood function, $p(D \mid m, I)$, is the PDF for observing the data D (i.e., measured concentration data) under "hypothesis" m (given source distribution), accounting for (1) model errors and input errors in the theoretical source-receptor relationship between the observed concentration and source distribution, and (2) uncertainty arising from the measurement noise in the detector. Assume the i-th concentration datum is measured subject to an additive Gaussian noise with root-mean-square (RMS) experimental error $\sigma_{D,i}$ and that the modeling error for the concentration datum has a Gaussian RMS error $\sigma_{T,i}$ (expected error in the model prediction of the concentration datum given a source model m). Furthermore, assume that the measurement and modeling errors are statistically independent. Then the likelihood function reduces to

$$p(D \mid m, I) \propto \exp\left[-\frac{1}{2}(D - R(m))^t \Sigma^{-1} (D - R(m))\right], \quad (8)$$

where

$$\Sigma \equiv \text{diag}\left\{\sigma_{D,i}^2 + \sigma_{T,i}^2\right\}_{i=1}^{N}, \quad (9)$$

and the superscript t denotes matrix transposition. Here, R is the theoretical source-receptor relationship and N is the number of concentration data. In assigning the prior probability $p(m \mid I)$, exactly what is known about the source function m will have to be stated. Assuming that nothing is known about the source parameters (e.g., source location and strength) we have $p(m \mid I) = \text{constant}$, and the posterior probability density function for the source parameters in this case is given by

$$p(m \mid D, I) \propto \exp\left[-\frac{1}{2}\sum_{i=1}^{N}\frac{(C(\vec{x}_i)-R_i(m))^2}{\sigma_{D,i}^2+\sigma_{T,i}^2}\right], \quad (10)$$

where $C(\vec{x}_i)$ denotes the measured time-averaged concentration for the detector at location $\vec{x}_i$. To calculate the posterior probability, we must explicitly specify a model $R(m)$ (dispersion model) relating the source parameters to the concentration observed at the receptor locations; viz., the model concentration C_m at receptor location $\vec{x}_r$ is given by $C_m(\vec{x}_r) = R_r(m).$

In order to efficiently calculate the source-receptor relationship, the following duality relation is used:

$$(C, h) = (C^*, Q_s), \quad (11)$$

where the inner product of f and g is defined as $(f,g) \equiv \int_\Omega f\, g\, d\Omega$, h is the detector response function with units of m^{-3} which embodies the measurement characteristics of a detector (e.g., $h = \delta(\vec{x}-\vec{x}_r)$ for an ideal point detector at receptor location $\vec{x}_r$) and Q_s with units of kg m^{-3} s^{-1} is the source density function. Furthermore, C^*, the residence time density function with units of s m^{-3}, is the dual to the concentration function C and can be obtained by solving the adjoint of the steady-state advection-diffusion equation:

$$-\frac{\partial(\overline{u}_i C^*)}{\partial x_i} - \frac{\partial}{\partial x_i}\left(\left(D+\frac{\nu_t}{\sigma_c}\right)\frac{\partial C^*}{\partial x_i}\right) = h, \quad (12)$$

subject to the following boundary conditions:

$$\left(D+\frac{\nu_t}{\sigma_c}\right)\frac{\partial C^*}{\partial n} + \vec{u}\cdot\vec{n}C^* = 0 \quad (13)$$

at the inflow and outflow boundaries of the flow domain, and $\frac{\partial C^*}{\partial z} = 0$ at all solid surfaces (e.g., ground surface, walls and roofs of buildings, etc.) and in the far field. Here, ν_t is the turbulent (eddy) viscosity and $\sigma_c \approx 0.63$ is the turbulent Schmidt number for the scalar.

3.1.2. MUST Array Results

The source-receptor configurations outlined below in Table 1 is with reference to the MUST array. The posterior PDF was sampled using Metropolis-Hastings Markov chain Monte Carlo [6]. Figure 3 shows histograms of the source parameters constructed from the sample equivalents in the Monte Carlo output when the Markov chain has reached equilibrium after a sufficiently long "burn-in" of the chain. The statistical estimates of the source parameters, expressed in terms of the mean and standard deviation, and estimated from the samples drawn from the Markov chain are presented in Table 2. The source strength for the continuous release is 3.7799×10^{-3} kg m^{-3} s^{-1}. The combined RMS experimental and modeling errors, $(\sigma_{D,i}^2+\sigma_{T,i}^2)^{1/2}$, are assumed to be 10% of the time-averaged (mean) concentration observed at each detector.

	x_s	y_s	z_s
Source location	1.5 row	0.0H	0.0H
Detector positions	4.5 row	0.0H	0.0H
	6.5 row	0.0H	0.0H
	9.5 row	0.0H	0.0H
	12.5 row	0.0H	0.0H
	4.5 row	5.4H	-9.0H
	4.5 row	5.4H	-6.0H
	4.5 row	5.4H	-3.0H
	4.5 row	5.4H	3.0H
	4.5 row	5.4H	6.0H
	4.5 row	5.4H	9.0H

Table 1: Source and detector locations

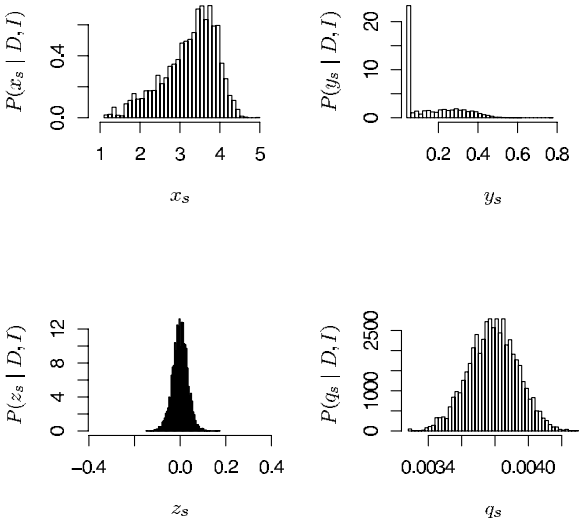

Figure 3: Histograms of source parameter samples

	actual	mean	std. dev.
x_s	3.665	3.245	6.684×10^{-1}
y_s	0.0	0.160	1.318×10^{-1}
z_s	0.0	0.000	3.335×10^{-2}
q_s	3.7799×10^{-3}	3.7880×10^{-3}	1.472×10^{-4}

Table 2: Source parameter estimates

It is clear from the results that the mean and standard deviations of the posterior distributions of the various source parameters are informative. Adaptive sampling strategies and extra prior information may improve the results, and remain to be investigated.

3.2 Coupling urbanSTREAM With a Mesoscale Flow Model

The interface between urbanSTREAM [7] and the "urbanized" Global Environmental Multiscale/Limited Area Model (GEM/LAM) [8], developed by Environment Canada, is demanding in that the information transfer between the two models must honor physical conservation laws, mutually satisfy mathematical boundary conditions, and preserve numerical accuracy, even though the corresponding meshes might differ in structure, resolution, and discretization methodology. Inter-grid communication allows the coarse mesh solution obtained by GEM/LAM to impose boundary conditions on the fine mesh of the urban microscale flow model (one-way interaction), and furthermore permits feedback from the fine mesh to the coarse mesh (two-way interaction).

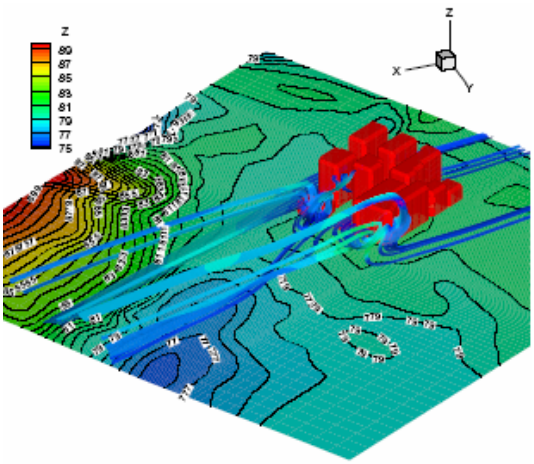

Figure 4: Flow over a block of buildings in the city of Ottawa.

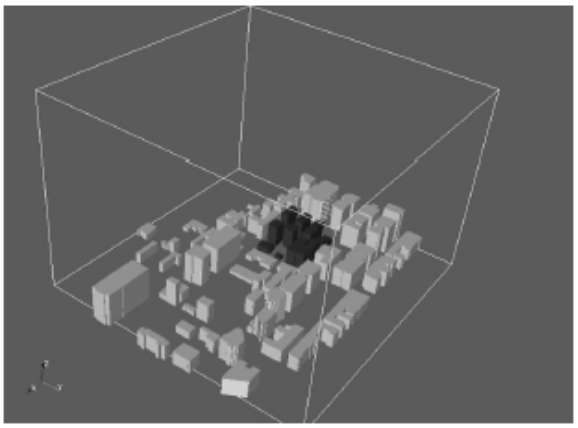

Figure 5: Virtual (gray) buildings surrounding a "high-resolution" sub-domain in the city of Ottawa.

Some very preliminary work has been undertaken to couple urbanSTREAM (urban microscale flow model) with urbanGEM/LAM (urban mesoscale flow model). In this initial effort, only a one-way coupling between the two models has been attempted; namely, the flow predicted by urbanGEM/LAM over the region of interest was used to provide the boundary conditions required to specify the flow over a small sub-domain of this region where a high-resolution urban flow simulation was undertaken by urbanSTREAM. This down-scaling of flow information from urbanGEM/LAM to urbanSTREAM was undertaken over a sub-domain of downtown Ottawa (see Fig. 4). In a future effort, the buildings surrounding this sub-domain (see gray buildings in Fig. 5) will be parameterized using the drag-force approach described in [9], [10] and [11].

4 Conclusions

Models for the prediction of the complex flow in the urban environment at the microscale have been developed, implemented, and validated against a number of comprehensive and detailed data sets obtained from wind tunnel and water channel simulations of flow over and through various building arrays (details can be found in [12]). The models are based on a Reynolds-averaged Navier-Stokes (RANS) approach with the hierarchy of turbulence correlation closure models based on a phenomenological two-equation model for the turbulence kinetic energy (k) and viscous dissipation rate (ε).

In the models described above, all the buildings in the cityscape were resolved explicitly in the sense that boundary conditions were imposed at all walls and roofs of every building. To reduce the computational cost of this approach, we investigated also the utility of representing groups of buildings in the cityscape in terms of a distributed drag force.

Finally, some effort was expended in interfacing the urban microscale flow model with an urban mesoscale model as well as with urban databases. A capability for ingesting ArcView ShapeFiles (which stores non-topological geometry and attribute information for building/obstacle features in a data set) has been developed. This capability permits building data in a ShapeFile to be ingested directly into the urban microscale flow model. An automatic scheme for generation of grids about the buildings imported within a specified flow domain has been designed and implemented. The grid information generated here can be imported directly into urbanSTREAM.

Acknowledgements

This work has been partially supported by Chemical Biological Radiological Nuclear Research and Technology Initiative (CRTI) program under project number CRTI-02-0093RD.

References

[1] K.J. Hsieh, F.-S. Lien and E. Yee, "Numerical Modeling of Scalar Dispersion in an Urban Canopy", submitted to *J. Wind Engineering and Industrial Aerodynamics* (2006) (in revision).

[2] C. Beguier, I. Dekeyser and B. Launder, *Phys. Fluids.* $21, 307-310(1978)$.

[3] E. Yee and C. Biltoft, *Boundary-Layer Meteorology* $111, 363-415(2004)$.

[4] E. Yee, R.M. Gailis, A. Hill, T. Hilderman and D. Kiel, "Intercomparison of Wind Tunnel and Water Channel Simulations Through a Large Array of Obstacles with a Scaled Field

Experiment", submitted to *Boundary-Layer Meteorology* (2006).
[5] R.M. Gailis, A. Hill, E. Yee and T. Hilderman, "Extension of a Fluctuating Plume Model to a Shear Boundary Layer and to a Large Array of Obstacles", submitted to *Boundary-Layer Meteorology* (2006).

[6] W. Hastings, *Biometrica* 57 , 97–109 (1970).

[7] F.-S. Lien and M.A. Leschziner, *Comput. Meth. Appl. Mech. Eng.* 114, $123-148(1994)$.

[8] J. Cote, S. Gravel, A. Methot, A. Patoine, M. Roch and A. Staniforth, *Monthly Weather Review* 126, $1373-1395(1998)$.

[9] F.-S. Lien and E. Yee, *Boundary-Layer Meteorology.* 112, $427-466(2004)$.

[10] F.-S. Lien, E. Yee and J.D. Wilson, *Boundary-Layer Meteorology.* 114, $245-285(2005)$.

[11] F.-S. Lien and E. Yee, *Boundary-Layer Meteorology.* 114, $287-313(2005)$.

[12] F.-S. Lien, E. Yee, H. Ji, A. Keats and K.J. Hsieh, "Development of a High-Fidelity Numerical Model for Hazard Prediction in the Urban Environment", *Proc. 13th Annual Conference of the CFD Society of Canada*, Jul. 31-Aug. 2, 2005, St. John's, NL, Canada.