FFI-RAPPORT

18/02185

Computational modeling of health effects from indoor dispersion of airborne biological and chemical agents

Hannibal Eie Fossum
Marius Dybwad
Anders Helgeland

Computational modeling of health effects from indoor dispersion of airborne biological and chemical agents

Hannibal Eie Fossum
Marius Dybwad
Anders Helgeland

Norwegian Defence Research Establishment (FFI) 20 March 2019

Keywords
CBRN-vern
Computational Fluid Dynamics (CFD)
Farlige stoffer
Matematisk mekanikk
Samfunnssikkerhet
Spredningsmodellering

FFI-rapport
18/02185

Project number
1443/1393

ISBN
P: 978-82-464-3136-9
E: 978-82-464-3137-6

Approvers
Janet M. Blatny, *Research Director*
Marius Dybwad, Anders Helgeland, *Research Manager*

The document is electronically approved and therefore has no handwritten signature.

(U) Summary

Numerous chemical and biological agents can be lethal or cause permanent human injury. Through evaporation or aerosolization, many such agents can be dispersed through air and subsequently be inhaled by people. The concentration of a toxic agent in air, required to estimate the inhaled dosage, is highly case-dependent and difficult to ascertain. For a given scenario, extensive experimental work or rigorous numerical simulations are needed to get a good estimate of how the concentration field varies in space and time, thereby enabling estimation of e.g. human casualty rates.

The present report aims to show how human casualty rates can be estimated by numerical simulations of dispersion of a released toxic agent in air. The goal is to increase the general understanding of realistic indoor dispersion processes, as well as to demonstrate the possibilites of advanced numerical simulation methods to assess consequences of specific threat scenarios. An indoor scenario at a conference center is used to exemplify how numerical simulations can quantify the potential consequences, such as casualty rates and time available for evacuation, of the release of a CB agent. Such results can help improve risk assessment and consequence evaluation within both civil emergency preparedness and military operations.

(U) Sammendrag

Biologisk eller kjemisk terror, der biologiske eller kjemiske stoffer benyttes for å skade eller ta livet av sivile, er en trussel med store potensielle konsekvenser. Flere kjemiske og biologiske trusselstoffer kan være dødelige eller forårsake varig helseskade. Ved fordamping eller aerosolisering kan slike stoffer spres gjennom lufta og derved pustes inn av mennesker. Konsenstrasjonen av et giftig stoff i luft er sterkt scenarioavhengig og vanskelig å fastslå, men kunnskap om dette er nødvendig for å beregne inhalert dose. Det kreves ofte omfattende eksperimentelt arbeid eller detaljerte numeriske simuleringer for å få et godt anslag over hvordan konsentrasjonsfeltet varierer i tid og rom, slik at for eksempel antall tapte menneskeliv kan estimeres.

Denne rapporten har som mål å synliggjøre hvordan populasjonsdødelighet kan beregnes ved hjelp av numeriske spredningssimuleringer for trusselstoffer. Målet er å øke den generelle forståelsen av realistiske spredningsprosesser innendørs. Samtidig ønsker vi å demonstrere mulighetene som ligger i avanserte numeriske beregningsmetoder for å vurdere konsekvenser av spesifikke trusselscenarioer. Vi har brukt et innendørs gassutslipp på et konferansesenter for å vise hvordan de mulige konsekvensene av utslippet kan kvantifiseres gjennom for eksempel populasjonsdødelighet og tilgjengelig evakueringstid. Slike resultater kan bidra til økt teknisk trusselforståelse og bedre konsekvensvurderinger både innenfor sivil beredskap og i militære operasjoner.

Contents

Appendix

Preface

This report is the result of work performed within two different FFI-projects, “Evaluation of biological threats IV” (Project 1443) and “Aerosols III” (Project 1393).

Abbreviations

The following abbreviations are introduced and used in the text:

AEGL	acute exposure guideline levels
CAD	computer-aided design
CB	chemical and/or biological
CFD	computational fluid dynamics
DNS	direct numerical simulation
FVM	finite-volume method
LD50	dosage corresponding to 50 % mortality rate
LES	large-eddy simulation
RANS	Reynolds-averaged Navier-Stokes

1 Background

1.1 Introduction

Numerous chemical and biological (CB) agents can be lethal or cause permanent human injury. Through evaporation or aerosolization, many such agents can be dispersed through air and subsequently be inhaled by people or animals. Such dispersion incidents can be caused by both accidental (e.g., Evensen and Olsen (2002) or Wenck et al. (2007)) or intentional (e.g., Tu (2007) or Pita and Domingo (2014)) releases of CB agents.

For most toxic agents, critical thresholds exist to estimate the health effect of a given dosage inhaled by humans; such measures include AEGL levels (Bruckner et al., 2004) or LD50 thresholds (Gill, 1982). However, the concentration of a CB agent in air, required to estimate the inhaled dosage, is highly case-dependent and difficult to ascertain *a priori*. For a given scenario, extensive experimental work or rigorous numerical simulations are needed to get a good estimate of how the concentration field varies in space and time, thereby enabling estimation of e.g. human casualty rates.

Moreover, by means of experimental or numerical work, it is also possible to investigate effects of agent source types, sizes, locations as well as other external conditons (such as changes in ventilation systems indoor or weather outdoor).

The present report aims to show how human casualty rates can be estimated by numerical simulations of dispersion of a released CB agent in air in a given scenario. Relevant data processing to assess the agent dispersion process is also demonstrated. The goal is to increase the general understanding of realistic indoor CB agent dispersion processes through an example, as well as demonstrate the possibilites of state-of-the-art numerical simulation methods to assess consequences of specific CB dispersion scenarios.

1.1.1 Objectives

The main objective of this study is two-fold, i.e.

- show how numerical simulations can be used to assess consequences of specific CB scenarios, and
- exemplify how indoor CB dispersion progresses in time and space.

1.2 Dispersion modeling

Several classes of numerical models exist to estimate dispersion of CB agents in air. For large-scale outdoor dispersion, operational models utilizing mean wind velocities and simple geometrical shapes like triangles or circles are the most simple. Gaussian models can include more complex phenomena, though at a higher computational cost. The most complex class of models comprise models based on computational fluid dynamics (CFD).

In a somewhat similar manner, indoor models range from simple assumptions of instantaneous uniform mixing to statistical methods and, furthermore, to the most complex high-fidelity models of CFD. In the following, indoor dispersion will be discussed, although many concepts apply equally well to outdoor dispersion modeling.

In CFD-based models, equations describing the physics of fluid flow are solved. Typically, CFD solvers aim to generate solutions which satisfy momentum and mass conservation at all points

within a specified geometry. By solving additional equations governing the dispersion of CB agents, reliable dispersion predictions can be obtained. Consult e.g. Versteeg and Malalasekera (1995) for an introduction to CFD. Due to increasing computational resources, CFD has become a popular tool for reasonably accurate dispersion modeling in recent years (Lien and Yee, 2004; Coirier et al., 2005; Lien et al., 2006; Santiago et al., 2007; Lateb et al., 2016).

Dispersion modeling within the CFD framework has been carried out at the Norwegian Defence Research Establishment (FFI) in many different contexts, both military and civilian, over the previous decade (Wingstedt et al., 2012a,b; Fossum and Petterson Reif, 2012; Fossum et al., 2012; Gjesdal et al., 2013; Vik et al., 2015; Aalbergsjø and Vik, 2016; Endregard et al., 2016; Wingstedt et al., 2017; Osnes et al., 2017). Applications range from simple Reynolds-averaged Navier-Stokes (RANS) models of passive, neutral gas dispersion in closed geometries to outdoor dispersion of dense gases in stratified environments and dispersion from improvised explosive devices (IEDs) utilizing so-called large-eddy simulations (LES).

The dispersion processes studied in the following are modeled by means of CFD.

1.3 Problem description

As an example case used throughout the remainder of this report, an indoor scenario at a conference center will be used to exemplify how numerical simulations can shed light on consequences of the release of a CB agent. Similar scenarios in the same conference center have been numerically investigated in previous studies (Endregard et al., 2010; Wingstedt et al., 2012b; Aalbergsjø and Vik, 2016), though with different parameters, models and data processing methods than will be used in the following.

Consider an indoor environment as shown in Figure 1.1a. The geometry consists of a large conference hall (colored blue in the figure), divided into eight subsections by temporary walls not reaching the ceiling, as well as a corridor (colored green in the figure) with two entrances into the conference hall. The ventilation system is also included in the figure; there are several vents where air enters or exits the indoor area.

Figure 1.1b shows the geometry from above, and the three horizontally different release locations are labeled A, B, and C. Simulated dispersion from different heights at each of these three locations will be discussed in the present report. From location A, release heights of 0, 1.5, and 8.8 m are used. From locations B and C, release heights of 0 and 1.5 m are used. Table 1.1 summarizes relevant aspects of the geometry as well as the simulated releases. Note that the release rate and type of CB agent are deliberately left out of the present report and are thus not included in Table 1.1, as this is classified information.

Given the geometry and releases described in the above, the remainder of this report will discuss how such dispersion processes can be simulated numerically (Part 2) and what may happen in the specific dispersion scenario that has been described (Part 3).

1.4 Document structure

This report is intended as a general documentation of the typical CFD dispersion methodology utilized by the group “Strømning og materialer” (“Flow and materials”) at FFI. As such, this document may serve as a useful citation in future FFI reports, thereby enabling future reports to treat the subject of CFD a little more superficially and instead focus more on problem-specific extensions and results.

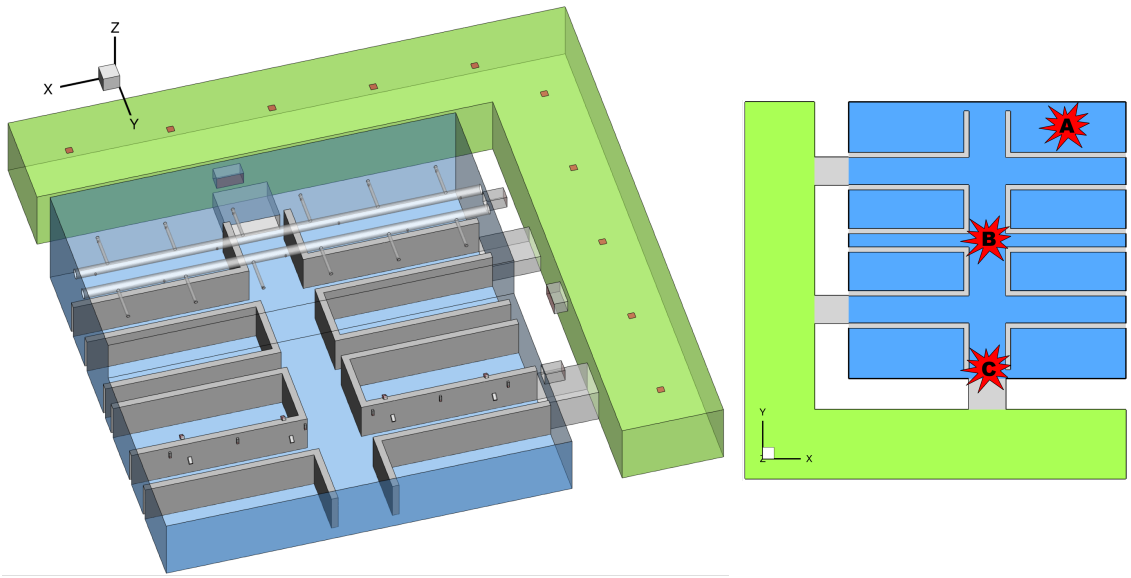

(a) Perspective view, including ventilation system.

(b) Top view (ventilation system not shown), including release locations A, B, and C.

Figure 1.1 The indoor geometry used in the dispersion simulations.

Table 1.1 Characteristics of the indoor geometry (above line) and releases (below line) used in the dispersion simulations.

Dimensions (l × w × h) [m]	82.4 × 82.4 × 16.6
Total volume [m^3]	$6.75 \cdot 10^4$
Air circulation volume [m^3/s]	34.7
Number of ventilation inlets	43
Number of ventilation outlets	4
Number of release positions	7
Release duration [min]	5

The second goal of this report is to present the results of the indoor dispersion scenario, as already discussed in Section 1.3. The results and conclusions from this particular case should be accessible even to readers without any prior knowledge of CFD, numerical methods or fluid mechanics.

In light of the above, it should be remarked that readers focused solely on the indoor dispersion case per se need only read Chapter 1, 3 (for details), and 4. The methodological details are generally contained within Section 2. This section treats the CFD framework in more detail, and are thus well-suited for readers with some knowledge of or interest in CFD.

2 Methodology

2.1 Introduction

This section treats the CFD framework in more detail, and are thus well-suited for readers with some knowledge of or interest in CFD. Readers purely interested in the indoor dispersion case may skip to Section 2.4 or, for the results of the simulations, to Chapter 3.

As discussed in Section 1.2, the dispersion process discussed in the present study is modeled by means of CFD. The general aim of CFD methods is to solve equations representing the physics of the fluid-dynamical phenomena that are relevant to the problem at hand.

In the following, the first chapter will describe the mathematical equations governing the fluid flow and dispersion, thus providing a mathematical framework for the CFD methodology. The second chapter will then focus on more practical aspects of the CFD methodology, i.e. the general workflow of CFD as well as how the specifics of the present dispersion problem have been accounted for.

2.2 Mathematical background

The governing equations of fluid motion state that mass and momentum are conserved, and they are thus referred to as conservation equations. The conservation equation for momentum follows from Newton's second law.

In the following, index notation is used in conjunction with Einstein's summation convention. Unless stated otherwise, free indices range from 1 to 3, and repeated indices imply summation.

A Cartesian coordinate system is assumed. The general spatial position vector is given by

$$\boldsymbol{x} = (x_1, x_2, x_3) = (x, y, z),$$

whereas time is denoted t. As seen in figure 1.1a, x and y are horizontal coordinates, whereas z is taken as the vertical direction. The corresponding three-dimensional instantaneous velocity vector field is denoted

$$\begin{aligned}\tilde{\boldsymbol{u}}(\boldsymbol{x}, t) &= (\tilde{u}_1(\boldsymbol{x}, t), \tilde{u}_2(\boldsymbol{x}, t), \tilde{u}_3(\boldsymbol{x}, t)) \\ &= (\tilde{u}(\boldsymbol{x}, t), \tilde{v}(\boldsymbol{x}, t), \tilde{w}(\boldsymbol{x}, t)),\end{aligned}$$

and the three-dimensional instantaneous pressure field is denoted $\tilde{p}(\boldsymbol{x}, t)$. Moreover, partial differentiation is abbreviated by

$$\begin{aligned}\partial_i &= \partial/\partial x_i = (\partial/\partial x, \partial/\partial y, \partial/\partial z), \\ \partial_t &= \partial/\partial t,\end{aligned}$$

for spatial and temporal gradients, respectively, when this is beneficial.

The conservation equations for momentum and mass for an incompressible, Newtonian fluid can be written

$$\partial_t \tilde{u}_i + \tilde{u}_k \partial_k \tilde{u}_i = -\frac{1}{\rho}\partial_i \tilde{p} + \nu \partial_k \partial_k \tilde{u}_i + F_{V,i}, \tag{2.1}$$

$$\partial_k \tilde{u}_k = 0, \tag{2.2}$$

respectively. Here, ρ is the (constant) fluid density and $\nu = \mu/\rho$ is the kinematic viscosity, with μ being the dynamic viscosity. The term $F_{V,i}$ represents any other volume forces affecting the fluid, such as buoyancy (gravitation) or rotation.

If scalars, such as gases, aerosols, or temperature, are transported passively in the fluid, each of them is also governed by its own transport equation, reading

$$\partial_t \tilde{c} + u_k \partial_k \tilde{c} = \gamma \partial_k \partial_k \tilde{c}, \tag{2.3}$$

in which $\tilde{c} = \tilde{c}(\boldsymbol{x}, t)$ is the transported scalar, and γ is its diffusivity.

Generally, for the system of partial differential equations to have a unique solution, initial and boundary conditions must be provided. In a few simplified cases, such as steady-state laminar pipe flow, only boundary conditions are required, and the equations can even be solved analytically. In more complex cases, numerical algorithms must be employed to obtain three-dimensional, time-dependent solutions.

The second term on the left-hand side of Eq. (2.1) is the advective term. Its nonlinearity makes the solution of this equation system highly susceptible to minimal perturbations in initial and boundary conditions, which in turn may lead to a turbulent solution. The nonlinearity generally makes the equations impossible to solve analytically. Even with the help of powerful computers, only a narrow range of real-life problems can be solved exactly by so-called direct numerical simulation (DNS). For most flow systems, approximate models related to this nonlinearity are utilized, so-called *turbulence models*.

2.2.1 Turbulence modeling: Large-eddy simulation (LES)

Presently, the flow turbulence arising from the nonlinear nature of Eq. (2.1) is modeled through so-called large-eddy simulation (LES).

The basis of LES is the decomposition of fluid flow fields into *resolved* and *subgrid* (unresolved) parts, e.g. $\tilde{u}_i^r(\boldsymbol{x}, t)$ and $\tilde{u}_i^s(\boldsymbol{x}, t)$, respectively, for the velocity field. Formally, the decomposition is obtained via a *filtering operation*, in which the field, say, $\tilde{u}_i(\boldsymbol{x})$, is filtered by a filter function $G(\boldsymbol{x})$ through the convolution

$$\tilde{u}_i^r(\boldsymbol{x}) = \int_{-\infty}^{\infty} G(\boldsymbol{x} - \boldsymbol{r}) \tilde{u}_i(\boldsymbol{x}) \, \mathrm{d}\boldsymbol{r},$$

to produce the resolved field, $\tilde{u}_i^r$, and, by extension, the subgrid field $\tilde{u}_i^s = \tilde{u}_i - \tilde{u}_i^r$.

The Fourier transform of a filter function is its associated transfer function, which represents the corresponding filtering in spectral space. The most common filter in LES algorithms based on the finite-volume method (FVM) is an implicit running-mean filter, where the spatial averaging size directly depends on the size of the local computational cell[1].

Regardless of the specifics of the filter function, the filtered – i.e. the resolved – conservation equations for momentum, mass, and scalars can be formulated

$$\partial_t \tilde{u}_i^r + \tilde{u}_k^r \partial_k \tilde{u}_i^r = -\frac{1}{\rho} \partial_i \tilde{p}^r + \nu \partial_k \partial_k \tilde{u}_i^r - \partial_k \tau_{ik} + F_{V,i}, \tag{2.4}$$

$$\partial_k \tilde{u}_k^r = 0, \tag{2.5}$$

$$\partial_t \tilde{c}^r + u_k^r \partial_k \tilde{c}^r = \gamma \partial_k \partial_k \tilde{c}^r - \partial_k f_k, \tag{2.6}$$

in which $\tau_{ik} = (\tilde{u}_i \tilde{u}_k)^r - \tilde{u}_i^r \tilde{u}_k^r$ and $f_k = (\tilde{c}\tilde{u}_k)^r - \tilde{c}^r \tilde{u}_k^r$ represent the unresolved and thus unknown *subgrid stresses* and *scalar subgrid fluxes*, respectively. These are the terms that require modeling.

[1] In spectral LES codes, a wave cutoff filter is commonly used instead.

In the indoor dispersion simulations, the dynamic Smagorinsky model is used to model the subgrid stresses and fluxes. A short description of the theoretical basis for this modeling approach can be found in Appendix A.

Note the resemblance of Eqs. (2.4)–(2.6) to Eqs. (2.1)–(2.3). In Eq. (2.4), it is assumed that the volume force $F_{V,i}$ passes through the filter unchanged. The solution of Eqs. (2.4)–(2.6) yields resolved velocity, pressure and scalar fields. These fields are generally three-dimensional time-dependent approximations to the full flow fields and can thus be subjected to Reynolds decomposition, e.g. $\tilde{u}_i^r = U_i^r + u_i^r$ for the velocity field, in order to compute flow field statistics (see Section 2.2.2).

Despite the advantages of LES in predicting details of turbulent flows, near-wall treatment at high Reynolds numbers remains a challenge. At lower Reynolds numbers, the viscous, high-shear region close to solid boundaries may be fully resolved, but for most real-life flows, this is unfeasible. If wall-generated shear is the main instigator of turbulence, near-wall modeling akin to Reynolds-averaged Navier-Stokes (RANS) wall models must be employed to obtain reasonable results. However, if other turbulence sources dominate, e.g. jets, vortex shedding, or shear from geometric roughness, LES may be well-suited. The latter is the case for the geometry considered in this study.

A final concern regarding LES, shared with DNS, is that of appropriate boundary conditions, particularly in the case of velocity inlets or free-stream boundaries. The prescription of physically realistic conditions consistent with the mathematical models and numerical implementation of a solver is still an area of active research, cf. e.g. Keating et al. (2004); Wu (2017). That said, exact boundary conditions are not essential to all studies of turbulent flow.

Note that in the following, the 'r' superscript notation is dropped for brevity. Hence, $\tilde{u}_i$, $\tilde{p}$, and $\tilde{c}$ always refer to instantanous LES-resolved fields.

2.2.2 Flow statistics

A turbulent flow field is random and stochastic in nature (Wyngaard, 2010, p. N). So-called Reynolds decomposition separates a turbulent field into two parts: By averaging the field, the *mean flow field* is obtained, and the difference between the full field and the mean field then comprises the *fluctuating flow field*. Mathematically, the Reynolds decompositions can be written

$$\begin{aligned}\tilde{u}_i &= U_i + u_i,\\ \tilde{p} &= P + p,\\ \tilde{c} &= C + c\end{aligned}$$

for the velocity field, pressure field and an arbitrary scalar field, respectively. Here, symbols marked by a tilde ($\tilde{\cdot}$) denote full four-dimensional fields – or, in the present case of LES fields, resolved (filtered) four-dimansional fields – whereas the uppercase and lowercase symbols denote mean and fluctuating fields, respectively. Note that the average of a mean field returns the same mean field, whereas the average of a fluctuating field is identically zero.

Formally, the averaging procedure is an ensemble average of infinitely many flow realizations. In practice, however, a finite spatial or temporal average (or a combination of both) is commonly used when possible. For example, for a temporal average of a statistically steady velocity field, $U_i(\boldsymbol{x}) = \langle \tilde{u}_i \rangle \approx \frac{1}{T}\int_0^T U_i(\boldsymbol{x}, t)\,\mathrm{d}t$, in which $\langle\cdot\rangle$ denotes the averaging process and T is the averaging period. Such an approximation to the true average is valid only if the flow is ergodic with respect to

the averaging dimension (time, t, in the previous example), i.e. homogeneous in one or more spatial directions (for spatial averaging) or statistically steady (for temporal averaging).

For finite agent releases in inhomogeneous geometries, the concentration field of the agent is not ergodic in time or space. Hence, an ensemble average of a finite number of realizations (i.e. "averaging several dispersion timelines") is the only feasible approximation to the true ensemble average. The velocity field, on the other hand, is statistically steady and thus lends itself to temporal averaging.

By utilizing Reynolds decompositions in Eqs. (2.1) and (2.2) (or, similarly, in Eqs (2.4) and (2.5)) and then applying an averaging procedure, the RANS equations are obtained. These equations, not shown here, describe the evolution of the mean flow fields and also form the basis for the RANS class of turbulence models. The transport equation for a mean scalar field, such as the concentration of a CB agent, can be derived in a similar manner.

The RANS equations are very similar to Eqs. (2.1) and (2.2), with the exception of an additional source term, $\partial_j r_{ij}$, in which r_{ij} is referred to as the Reynolds stresses. This term originates from the advection term in Eq. (2.1). The Reynolds-stress tensor can be written

$$r_{ij} = \langle u_i u_j \rangle,$$

i.e. each tensor component is a single-point correlation between fluctuating velocity vector components. The Reynolds stress tensor does not really represent physical stresses, but it has the same units. Physically, it is associated with momentum transport due to turbulent fluctuations, so-called turbulent momentum flux. Note also that the three normal components of the Reynolds stress tensor correspond to the statistical variances of the three flow velocity components, respectively, thus providing measures of the velocity fluctuations in the flow.

The *turbulence kinetic energy*, often used to estimate the flow turbulence level, is defined by an index contraction on the Reynolds stresses, i.e. $k = \frac{1}{2}\langle u_i u_i \rangle$.

2.3 Computational fluid dynamics

As discussed in Section 2.2, the governing equations of fluid motion cannot generally be solved analytically. The solution of Eqs. (2.1)-(2.2), as well as possible additional equations, by means of numerical algorithms on a computer is commonly called *computational fluid dynamics* (CFD). The methodology of CFD varies greatly from application to application and software to software, but the general workflow is often as outlined in Figure 2.1. The main steps in the workflow will be described briefly in the following.

2.3.1 Geometry and meshing

Firstly, the geometry of the problem needs to be supplied, whether this be a simple box or a complex urban area with lots of details. In the case of complex geometries, it is often necessary to clean up the geometry to reduce the demands for computational resources when solving the governing equations numerically. For example, details with relatively little impact on the flow field (such as window-frames or road signs in the case of an urban city center), can and should be removed.

In the present case, the geometry – already described in Chapter 1 – is well-suited for CFD simulations and can be meshed without further adjustments.

With rare exceptions, the geometry and the volume it encloses must be subdivided into computational *cells* for the CFD solver to give reasonable results. This is a process referred to as

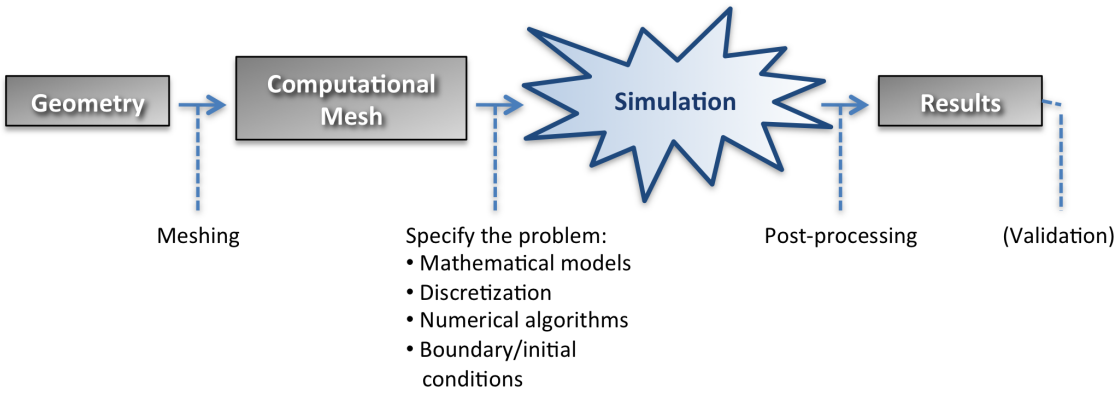

Figure 2.1 Typical CFD workflow (from left to right).

meshing or *gridding*. Depending on the application, everything from a few thousand to hundreds of millions of cells may constitute a computational mesh. The CFD code solves a linearized system of the governing equations (Eqs. (2.4) and (2.5)) simultaneously for each computational cell.

Usually, the mesh is not uniformly spaced; regions of high shear, for example, tend to require much smaller cells to capture the large velocity field gradients. Furthermore, the mesh must adhere to the shape of the geometry while retaining adequate quality; particularly, large aspect ratios, highly acute or obtuse angles, and rapidly varying cell sizes should be avoided.

In the context of implicit LES modeling, as used in this study, the mesh resolution, i.e. the range of cell sizes, also determines the size of the smallest turbulent scales that are directly resolved in the simulation, cf. Section 2.2.1.

2.3.2 Numerical discretization

The equations that are solved in a CFD program come in numerous variants (Gresho, 1991, cf. e.g.). In LES, Eqs. (2.4)-(2.5) or equivalent equations are implemented.

The equation set needs to be approximated and discretized so it can be solved numerically. For FVM solvers, this entails recasting the governing equations in conservative form so that each computational cell in the mesh can be treated as a discrete control volume. The solver thus calculates the solution to the transport equations by considering the fluxes through the surfaces of each cell in the computational mesh. As an example, consider the advection term of Eq. (2.1). Integrating over a volume and using Gauss' divergence theorem, the term can be rewritten

$$\iiint_V u_k \partial_k u_i \, \mathrm{d}V = \oiint_S u_i u_k n_k \, \mathrm{d}S,$$

where V and S are the volume and enclosing surface of a control volume – typically a computational cell – respectively, while n_k is the outward-pointing unit normal vector of the surface, S.

The resulting integral system can be discretized, i.e. converted into a system of algebraic equations, by employing a number of substitutions, which depend on the specific variety of FVM used, for the terms in the integrated equations. For incompressible flows, the discretized equations are then solved for each computational cell in the mesh simultaneously, time step by time step.

The solution for a given time step is computed by producing a linear system of equations of rank $N_{\text{tot}} \times N_{\text{tot}}$, where N_{tot} is the total number of computational cells, which is solved iteratively.

The quality of a numerical solver is often judged in terms of its convergence, consistence, and stability (Versteeg and Malalasekera, 1995, p. 6), but in practice versatility and ease-of-use might be equally important. In the context of LES, the numerical algorithm is particularly relevant since the mesh can be quite coarse compared to DNS meshes.

Conservation of kinetic energy is also of particular importance to LES solvers; dissipative numerical schemes tend to overwhelm the effect of the subgrid (and molecular) viscosity, whereas straight-forward, non-dissipative central-differencing schemes lead to numerical instability (Mahesh et al., 2004). Hence, the use of robust, non-dissipative numerical algorithms can be crucial to the success of an LES solver.

2.3.2.1 CDP – a multipurpose LES code

The bulk of numerical results discussed in the present report are based on numerical simulations performed with the FVM-based solver CDP v4.0.8 (Mahesh et al., 2002; Ham and Iaccarino, 2004; Mahesh et al., 2004; Ham et al., 2006), and the most important features of the solver will be summarized in the following. Some preliminary test simulations were also performed with the commercial FVM-based solver ANSYS Fluent 18.2 (ANSYS), which produced largely similar results.

The incompressible version of CDP, `Vida`, stores all flow fields in the mesh nodes. The code is second-order in space and up to second-order in time, using a Crank-Nicholson/Adams-Bashforth time-stepping algorithm. The pressure-coupling is based on a fractional-step approach.

The conceptual calculation of the subgrid viscosity in CDP's dynamic Smagorinsky model is described in Appendix A. CDP uses a test filter with twice the size of the ordinary filter width.

The advantages of CDP are similar to those of other FVM-based methods; in particular, the code uses unstructured meshes to be able to handle complex geometries, and it is relatively easy to extend the code by programming additional boundary conditions, source terms or data-processing algorithms. Additionally, the numerical algorithms in CDP are formulated so as to maximize conservation of kinetic energy without loss of numerical stability. Recent versions of the code is written in C++ and is highly parallelized.

The major drawback of FVM-based solvers is reduced accuracy compared to e.g. finite element methods, given similar mesh resolutions.

2.4 Particulars of the indoor dispersion scenario

2.4.1 Mesh

The geometry shown in Figure 1.1 was meshed using a tetrahedral mesh consisting of 987,428 computational nodes, illustrated in Figure 2.2.

To assess the mesh resolution at ground level, i.e. on the floor of the geometry, a common nondimensional quantity can be used, defined presently as

$$z_+ = \frac{z u_\tau}{\nu}, \quad \text{where}$$
$$u_\tau = \sqrt{\nu \left[\partial_3 U_t\right]_{z=0}}.$$

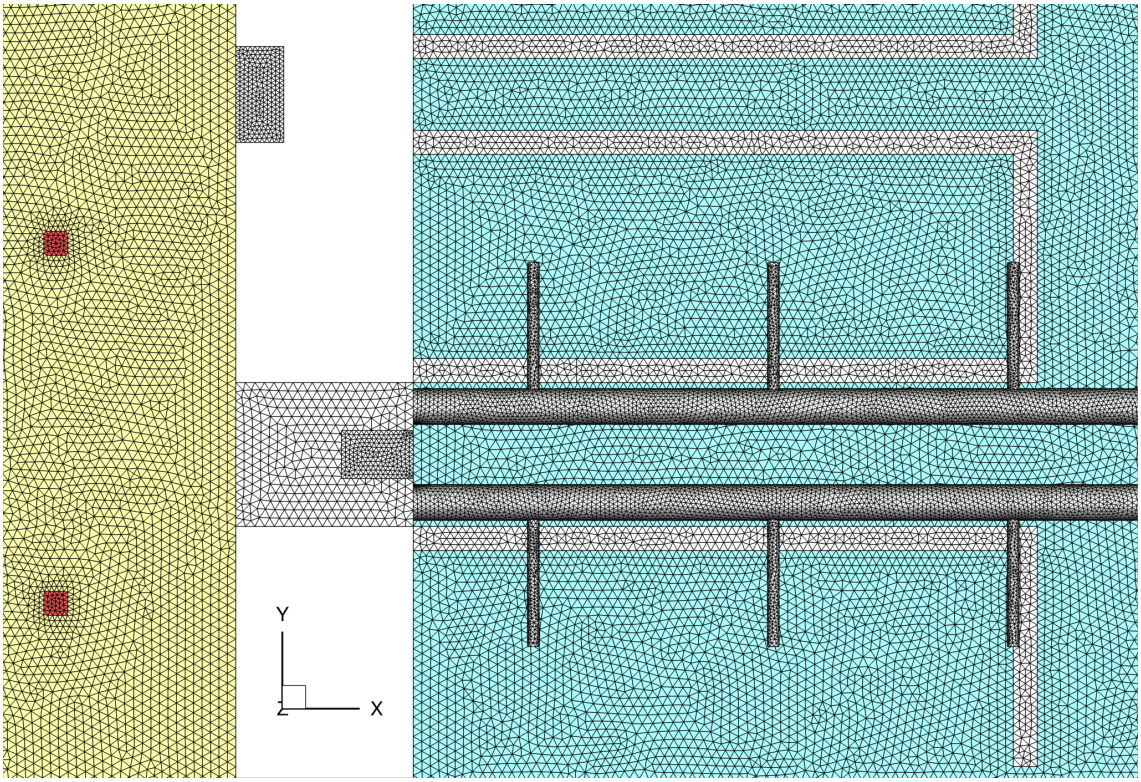

Figure 2.2 A portion of the numerical mesh, showing part of the corridor (left, yellow) and conference hall (right, blue), as well as parts of the ventilation system.

Here, u_τ is the friction velocity on the ground and U_t is the velocity tangential to the ground. z_+ values can be computed in all ground-adjacent points and thereby estimate how many viscous length scales are contained in (or, equivalently, unresolved by) the cell closest to the ground. A lower number implies higher mesh resolution, and for $z_+ \sim 1$, all turbulent scales (near the wall) are resolved, i.e. no turbulence model is generally required. Typical RANS wall-models require anything from $z_+ \sim 10$ to $z_+ \gtrsim 100$, depending on the model.

Low z_+ values are essential if realistic transition to turbulence or boundary-layer build-up will be simulated. In problems where wall-turbulence dynamics is less relevant, the requirement can be relaxed, often significantly. Higher Reynolds-number flows, i.e. flows with more turbulence, have smaller viscous scales. Hence, simulation of such flows require higher physical mesh resolutions than simulation of flows with lower Reynolds numbers to yield the same z_+ values.

The wall-normal z_+ metric is intended for estimates of near-wall resolution. A more general estimate of viscosity-based mesh resolution, which can be evaluated anywhere in the domain, is given by

$$l_+ = V_{\text{cell}}^{1/3} \sqrt{|S|/\nu},$$

in which $|S| = \sqrt{2S_{ij}S_{ij}}$ is the norm of the mean rate-of-strain tensor and V_{cell} is the computational cell volume. The strain-rate tensor is given by $S_{ij} = \frac{1}{2}(\partial_i U_j + \partial_j U_i)$. Similarly to z_+, l_+ indicates how many viscous lengths scales any given computational cell contains (resolves).

Table 2.1 summarizes the most important characteristics of the mesh. The relatively large z_+ values at the walls are acceptable here, since the turbulence is not instigated mainly by wall-shear. Rather, the turbulent flow structures are initiated from vortex formation at the edges of air jets from the ventilation fans, vortex formation due to jet impingement, and vortex shedding on the lee side of geometric structures.

Table 2.1 Characteristics of the computational mesh. z_+ and l_+ are nondimensional measures of mesh resolution, as described in the text.

Total number of computational nodes	987,428
Cell types	Tetrahedrons only
Cell volume; min, avg, max [m^3]	$5.1 \cdot 10^{-5}$, 0.11, 0.29
z_+ at ground; min, avg, max	$3.1 \cdot 10^{-2}$, 17, 40
l_+; min, avg, max	4.8, 51, 380

That said, the mesh resolution should ideally be higher, and a mesh sensitivity study should be performed to quantify the uncertainty related to resolution. However, the main objective of the present study is to prove a concept (fluid dynamical simulations as a means of assessing risk) and develop suitable data processing routines. Consequently, mesh sensitivity has not been studied systematically at present.

2.4.2 Ventilation

The indoor air circulation is instigated by the ventilation system. The volume flow of air of each individual inlet and outlet vent is taken from earlier simulations of the same building (Aalbergsjø and Vik, 2016). The total air circulation is 42.5 kg/s. Zero recycling of the extracted air is assumed. The air flow within the ventilation system itself has not been simulated.

Although the flows from the ventilation inlets are in fact turbulent, they have been modeled as laminar flows in the present study. This has negligible impact on the ambient velocity field in the room.

Prior to the release and dispersion of CB agents, the simulations reported presently were subjected to a run-up time of 27.5 minutes[2]. This ensured that the air flow field was statistically steady before any CB agent was released.

2.4.3 Release and dispersion

After a run-up period, CB agents were released from multiple locations. The release locations are listed in Table 2.2. General information about the releases can be found in Table 1.1. The release marked with an asterisk (*) in the table denotes the release for which an ensemble average was estimated (see Section 2.4.3.1).

Table 2.2 CB agent release locations.

Notation in Figure 1.1b	Coordinates
A	(48, 55, 0)
A	(48, 55, 1.5)
A	(48, 55, 8.8)
B	(30, 30, 0)
B*	(30, 30, 1.5)
C	(30, 2, 0)
C	(30, 2, 1.5)

The released gas is assumed to be neutrally buoyant, i.e. it is transported passively in the air without altering the air flow. This is often an acceptable assumption, provided that the concentration of gas is not too high.

2.4.3.1 Ensemble averaging

The turbulent nature of the flow field indicates that one release and dispersion event – one *realization* or *sample* – might not be representative of a typical such event. Since the dispersion process is transient and inhomogeneous, spatial or temporal averaging is not applicable. Thus, ensemble averaging must be applied to say something about the statistics of the dispersion process.

In practice, the ensemble average of muliple realizations is obtained by simulating all realizations after each other in time. In order to ensure that each realization is statistically uncorrelated with preceding and succeeding realizations, the time separation needs to be large enough for the flow field to be statistically uncorrelated.

Formally, the correlation time of the flow at any given location is found from the autocorrelation function of the velocity signal at that location. The correlation time is on the same order of magnitude as the integral turbulent time scale, the so-called eddy turnover time. In the present case, this can be estimated as $\tau \sim T_c = L_c/U_c \lesssim 6\ \mathrm{m}/0.5\ \mathrm{ms}^{-1} = 12$ s, in which τ is the integral time

[2]The time step was not constant during this process; the first 17.5 minutes of simulation had a time step of approx. 0.03 s, whereas subsequent simulation used a time step equal to 0.05 s.

scale, whereas T_c, L_c, and U_c are characteristic time, length and velocity scales of the problem. The characteristic scales are estimated from the height of the dividing walls and an assumed typical velocity based on the flow through the ventilation inlets. The estimate has been corroborated by inspection of simulation data, cf. Appendix B.

The estimated correlation time implies that the time separation between successive CB agent releases should exceed at least 12 s. A temporal separation of 60 s were used in the present study.

Assuming normal distribution, the number of samples required for the sample mean to be within a distance M of the true (population) mean must satisfy

$$N_{\text{realizations}} \geq \left(\frac{\sigma c_{\text{conf}}}{M}\right)^2. \tag{2.7}$$

Here, σ is the standard deviation and c_{conf} is a coefficient chosen based on the required confidence level. Typically, $c_{\text{conf}} = 2$, which implies that one can be 95 % confident that the sample mean differs no more than M from the population mean.

Note that σ in Eq. (2.7) is *a prioiri* unknown and needs to be estimated from simulation data or empirical data. However, by expressing M in terms of σ, an estimate of the required sample size can be found:

To be 95 % confident that the sample ensemble average is within half a standard deviation from the population ("true") ensemble average, $N_{\text{realizations}} = 16$ samples, i.e. 16 realizations, are required, seen by insertion of $c_{\text{conf}} = 2$ and $M = \sigma/2$ into Eq. (2.7). In fact, 40 realizations – i.e., 40 releases – were simulated from release location B at height 1.5 m and ensemble averaged in the present study, in order to assess the representativeness of one realization.

2.4.4 Health effects

One of the key aspects of a CB dispersion event is its potential consequences for public health. More specifically, given a concentration of a toxic agent, how does this concentration translate into potential health effects?

Different health effect models can be divided into two classes: The *concentration-dependent* models estimate health risks by assuming exposure to a given concentration of a CB agent for a specified period of time. In the context of numerical simulation, the concentration level is taken from the simulated concentration field at any given time. On the other hand, the *dosage-dependent* models estimate health risks by using the total dosage, i.e. concentration integrated over time, of CB agent for a specified period of time. Here, a time series of the simulated concentration field is used as input to the model. Both types of models rely heavily on empirical toxicity data.

The main advantage of the concentration-dependent models is their simplicity; no detailed time history of the concentration field is needed, no integral needs to be computed. Simply supply the concentration value and get the health effect estimate.

The major advantage of the dosage-dependent model is its ability to incorporate varying concentration levels in time, thus utilizing more information to get a potentially more accurate health effect estimate.

Regardless of the choice of model, assumptions about agent toxicity, person mass, respiration rate, etc., must also be incorporated into the model – explicitly or implicitly. Furthermore, the model output, i.e. the health effect estimate, can be given in a range of different ways, such as percentage mortality rates, classes of symptoms (e.g., ranging from annoying to lethal), or multiples of lethal dosages.

Two common health effect models are the Acute Exposure Guideline Levels (AEGL) model (Bruckner et al., 2004), which is a concentration-dependent model commonly used for chemical agents, and the 50% Lethal Dosage (LD50) model (Gill, 1982), which is a dosage-dependent model used both for chemical and biological agents. The latter has been used in the present indoor dispersion simulations and will be discussed in more detail in the following.

2.4.4.1 Dosage

In order to utilize the LD50 model, the dosage needs to be computed from the simulated time- and space-dependent concentration field for the CB agent. Given a time frame, T, and a path walked by a person in that time, the dosage to which the person is exposed can be calculated as

$$D_s = \int_S f\tilde{c}\ \mathrm{d}t,$$

where S is the path described by $\boldsymbol{r}(t)$, $0 \leq t \leq T$, and $\tilde{c} = \tilde{c}(\boldsymbol{x}, t) = \tilde{c}(\boldsymbol{r}(t), t)$ is the concentration field. The coefficient, f, is a product of the respiration rate and airway deposition efficiency. These are empirical parameters for which typical values of 10 liters/minute and 100 % efficiency are used in the present calculations. It is assumed that the integral's time period is short enough that no clearance of the CB agent from the body need to be considered.[3]

In the simpler case of a stationary person standing at location $\boldsymbol{r}_0$ for a time T, the integral becomes

$$D_s = \int_0^T f\tilde{c}(\boldsymbol{r}_0, t)\ \mathrm{d}t.$$

The dosage integral can be computed during simulation, but a more flexible solution is to save the entire concentration field at sufficient temporal resolution and then calculate the integral as a post-processing routine later. This enables the choice of position (or path) and time frame for exposure to be decoupled from the simulation per se. Since such post-processing calculations are rather quick, it is easy to "experiment" with how different human behavior models affect the outcome (e.g., number of casualties), using the same concentration field (i.e., the same dispersion event) as input every time. Note also that since the dosage integrals are linear operators on the concentration field, multiplying the amount of CB agent that is released, i.e. multiplying the concentration field, will multiply the dosages correspondingly.

2.4.4.2 The LD50 model

The LD50 model compares calculated dosage with a reference toxicity threshold in order to assess the health risk. Specifically, the LD50 threshold is the dosage for which lethality is expected in 50 % of the exposed subjects - hence the name 50 % *lethal* dosage. In the notation of the previous section, this means that if

$$D_s \geq D_{\mathrm{LD50}} * m_p,$$

where D_{LD50} is the LD50 threshold dosage (per body mass) and m_p is the body mass of the exposed person, then there is a statistical mortality probablity of 50 %. [4]

The LD50 threshold is based on empirical data and might be highly uncertain. In many instances, the thresholds are based on animal trials. Consult e.g. Gill (1982); Weil (1983); Randhawa (2009)

[3] The clearance rate could also be included in the integral, provided that the clearance rate of the CB agent is known.

[4] The criteria is obviously intended for use on a population scale and not on an individual level.

or references therein for more details relating to how LD50 thresholds are determined and examples of thresholds for different compounds.

In general, there is no linear relationship between dosage and mortality rates. In other words, multiplying the dosage to which a population is exposed will not multiply the mortality rate by a proportional factor. This means that dosages less or greater than D_{LD50} cannot be used to estimate mortality rates at those exact levels; rather, such data can only be used to say how far away the encountered dosage is from the LD50 threshold, e.g. by considering multiples of D_{LD50}.

When the LD50 threshold criterion is used in the present report, it has been assumed that the average weight of the exposed population is 70 kg, that their average breating rate is 10 liters/min and that most individuals are of sound health.

3 Results

In the following, the results of the numerical simulations will be discussed. Firstly, some information about the air flow in the building will be given, based on the simulated flow. Secondly, a collection of dispersion simulation results for the set of different release locations will be given. Finally, one of the release locations will be treated in more detail, ending in a practical example in which the potential for human casualties will be assessed.

3.1 Flow field

As previously discussed, the air flow within the simulated indoor geometry is driven by the ventilation system.

There are 43 relatively small ventilation inlets, most of which can be discerned from the velocity magnitude contours shown in Figure 3.1a. The ventilation outlets are fewer and larger than the inlets and, as opposed to the inlets, do not produce significant air jets.

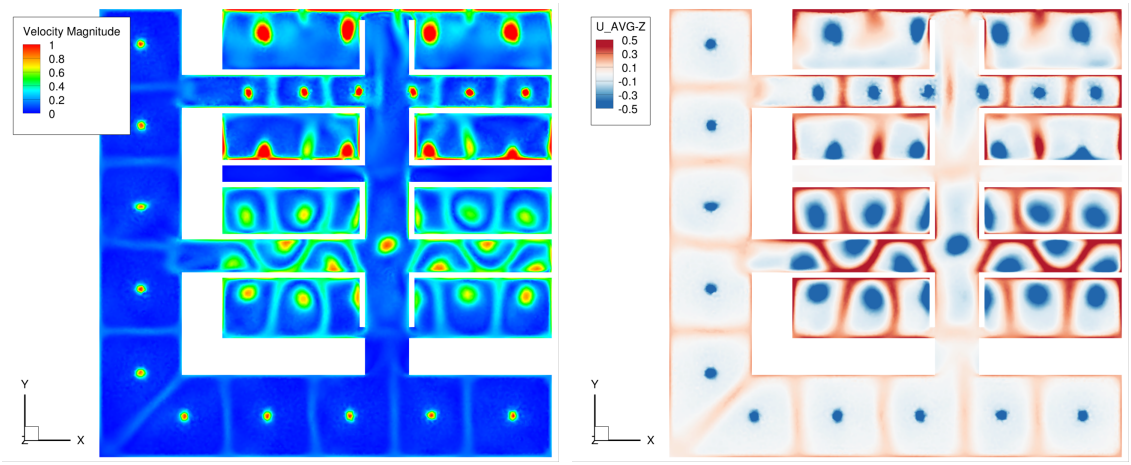

(a) Mean velocity magnitude (units are m/s).

(b) Mean vertical velocity (units are m/s).

Figure 3.1 Flow statistics. Contours shown in a plane at height $z = 1.5$ m.

Both the mean air flow and the turbulence in the indoor environment are caused only by the air circulation system. At around head-level height (1.5 m above floor), the main contributor to the peaks in velocity magnitude is the vertical motion which stems from inlets directed downward from the ceiling, as shown in Figure 3.1b.

The number of ventilation inlets, in combination with the geometry of the conference hall and corridor, implies that the mean flow will be rather complex, which is also evident from both subfigures of Figure 3.1. Moreover, the jets give rise to vortical motions in the flow, illustrated in Figure 3.2a.

The relatively high velocity magnitudes (as well as the combination with multiple inlets and complicated geometry) indicate that the flow will be highly turbulent, as is the case for almost all real-life flows. Indeed, Figure 3.2b shows significant levels of turbulence kinetic energy, i.e. energy contained in turbulent motion of air. Unsurprisingly, this is consistent with an estimated

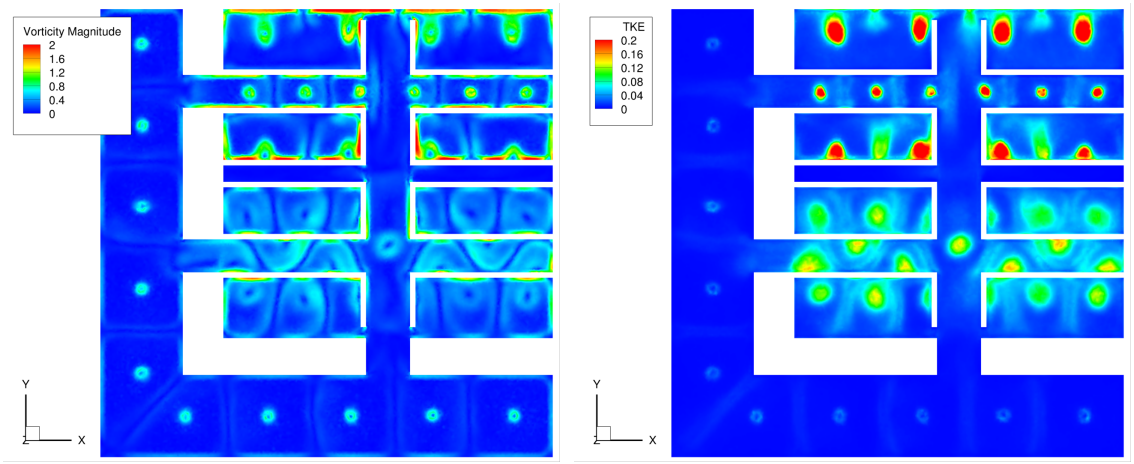

(a) Mean vorticity magnitude (units are s^{-1}).

(b) Turbulence kinetic energy (units are m^2/s^2).

Figure 3.2 Flow statistics. Contours shown in a plane at height $z = 1.5$ m.

bulk Reynolds number of

$$\mathrm{Re}_0 = \frac{U_c L_c}{\nu} \sim 10^5$$

in which U_c and L_c are characteristic velocity and length scales of the problem, and $\nu = 1.46 \cdot 10^{-5}$ is the kinematic viscosity of air. In the Reynolds number estimate above, typical fan diameters (approximately 0.5 m) and bulk exhaust velocities (typically around 5 m/s) are used.

The Reynolds number is the ratio between inertial and viscous forces in the flow; a higher Reynolds number implies higher levels of turbulence in the flow due to more inertial forcing. In circular pipe flow, a bulk Reynolds number higher than $\approx 2 \cdot 10^3$ generally implies a fully turbulent flow[5].

Figure 3.2b reveals that the turbulence kinetic energy is highest close to the ventilation inlets, which to be expected. Consequently, turbulent mixing (of a CB agent, for example) is expected to be most significant in these areas. That said, as the flow is fully turbulent everywhere, efficient mixing can be expected everywhere in the domain. This is an essential feature of the flow system in the context of dispersion of a CB agent.

3.2 Dispersion of the CB agent

Several numerical simulations were performed to investigate the release and dispersion of the CB agent, with different purposes.

A set of 8 releases, one from each source location (cf. Section 1.3), was studied to assess the effect of source location on the dispersion. The results of these simulations are discussed in Section 3.2.1.

One of the 8 cases – release from location B at 1.5 meters – was then selected for more extensive analysis, which is reported in Section 3.2.3. That section provides examples of more detailed data processing, as well as an example of how such simulation data can be used to assess the health risk for facility occupants.

[5]This is by no means an exact rule.

Finally, 39 additional dispersion events were simulated from location B (at different times) in order to compute an estimate of the ensemble average of a dispersion event. The result, discussed in Section 3.2.2, is used to assess to what degree one disperion simulation is representative of a "typical" (i.e. average) dispersion event. This gives an indication of the level of uncertainty inherent in the single-realization data discussed in the rest of the report.

3.2.1 Effect of source location

When designing dispersion scenarios for numerical simulations, only a finite number of source locations is typically used. Different criteria can be applied to select such locations depending on the objective of the study.

In any case, it is of interest to assess the importance of the particular choice of source location. In other words, how much does the dispersion outcome differ depending on source location, given the same flow field. To address this question, dispersion from sources at three different horizontal locations were simulated (A, B and C, as shown in Figure 1.1b). For these three locations, different source heights were also used, as described in Section 1.3. The effects of varying the source location will now be discussed.

Figure 3.3 shows two CB agent concentration fields 20 min after the releases have stopped. Figure 3.3a shows the concentration field from a source at location B and height 0 m (i.e. at the ground), whereas Figure 3.3b shows the concentration field from a source at the same location, but with height 1.5 m.

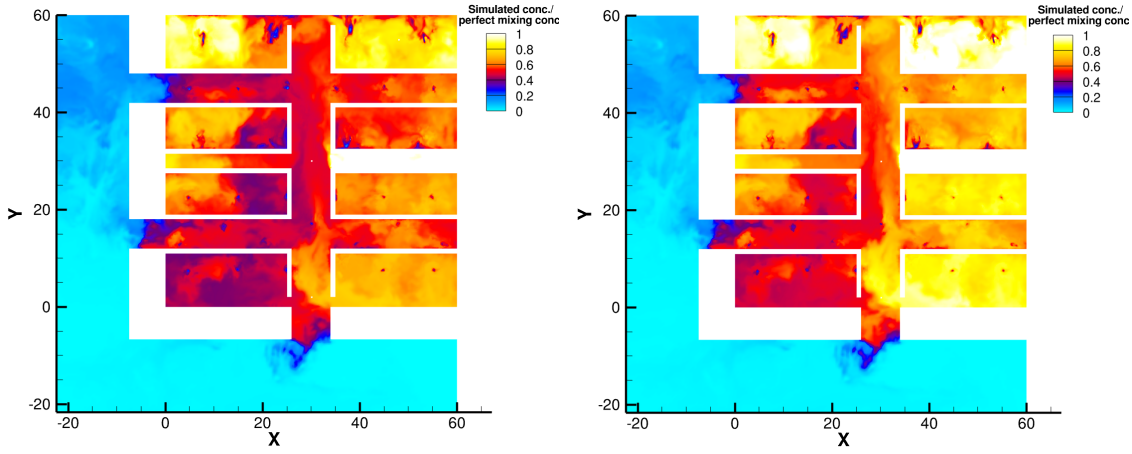

(a) Release from 0 m (at location B).

(b) Release from 1.5 m (at location B).

Figure 3.3 Countours of nondimensional concentration, $\tilde{c}_* = \tilde{c}/C_{perf}$, after 20 minutes in a cutplane at $z = 1.5$ m.

In Figure 3.3, the nondimensional concentration values, $\tilde{c}_*$, have been normalized (uniformly) by

$$C_{\text{perf}} = m_a/V,$$

in which m_a is the total mass of CB agent being released, and V is the total volume of the computational domain. C_{perf} thus represents the result of a simple "perfect mixing" model, which assumes that all of the agent would be dispersed uniformly in the room instantaneously.

As seen in Figure 3.3, the concentration fields are generally within the same order of magnitude, indicating that the height of the source has little effect on the subsequent dispersion. The two concentration fields shown in the figure have the largest overall statistical difference between two different release heights (for the same horizontal location) among all the simulated releases. This is seen from Table 3.1, to be discussed shortly, in which statistical comparisons between other source heights are also reported. As seen, the differences in concentration resulting from changes in source height appear to be minor for the cases that have been simulated.

Given these results, only releases at height 1.5 m are considered when comparing dispersion from different horizontal locations. Three different horizontal locations (A, B, and C) have been used in the simulations, cf. Figure 1.1b. Figure 3.4 shows the nondimensionalized CB agent concentration fields from these three locations 20 min after the releases have ceased.

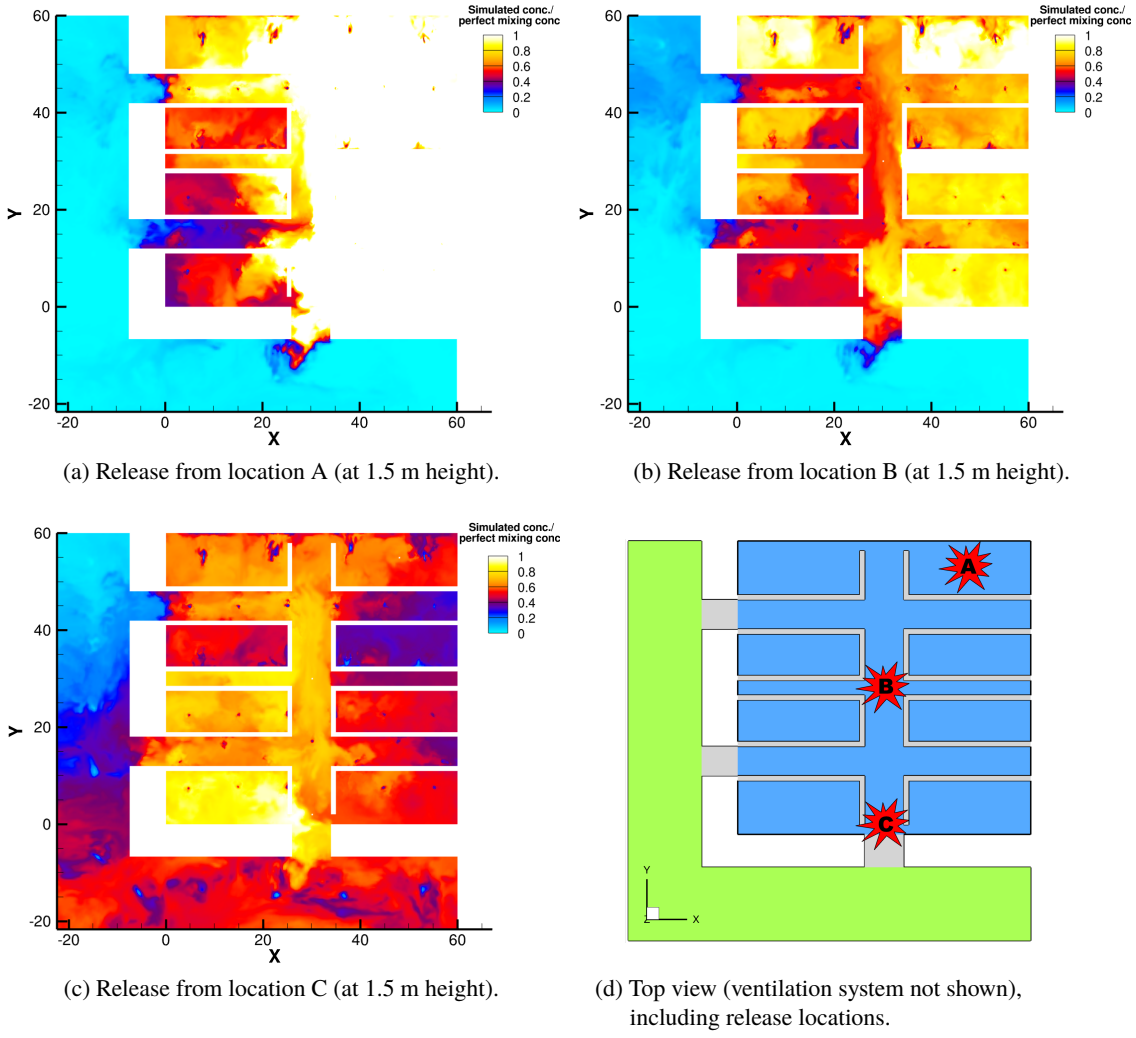

(a) Release from location A (at 1.5 m height).

(b) Release from location B (at 1.5 m height).

(c) Release from location C (at 1.5 m height).

(d) Top view (ventilation system not shown), including release locations.

Figure 3.4 Countours of nondimensional concentration, $\tilde{c}_* = \tilde{c}/C_{perf}$, after 20 minutes in a cutplane at $z = 1.5$ m. Contour levels are capped at $\tilde{c}_* = 1$ for clear visualization, even though $\max(\tilde{c}_*) = 2.2$. The bottom-right figure is identical to Figure 1.1b, included here for ease of access.

Clearly, the horizontal position of the source is significantly more important than the height; there are relatively large qualitative differences between the concentration fields from source A, B, and C. This is corroborated from the statistical differences summarized in Table 3.2.

Figure 3.4 also indicates that it takes time for the CB agent to disperse out into the corridor outside of the conference hall, regardless of source position. As expected, the source near one of the exits (at location C) appears to result in the highest concentration levels in the corridor.

A more quantitative measure of the effect of source position can be carried out via statistical analysis. The Mean Relative Bias (MRB) (Duijm et al., 1996) is a measure of systematic positive or negative deviations in one field, C_a, compared to another field, C_b. The metric is defined by

$$\mathrm{MRB} = \overline{2\left(\frac{C_a - C_b}{C_a + C_b}\right)},$$

where the overbar denotes an average over the entire data set. Note that negative MRB values imply a negative systematic deviation in C_a relative to C_b, and that positive MRB values consequently imply a positive systematic deviation in C_a.

A closely related metric is the Mean Relative Square Error (MRSE) (Duijm et al., 1996), which is related to the variance of the difference between the compared data sets and is thus a measure of the scatter in the comparison of the two sets of values. The MRSE is defined by

$$\mathrm{MRSE} = \overline{4\left(\frac{C_a - C_b}{C_a + C_b}\right)^2}.$$

When considered together, the MRB/MRSE pair gives information both on the level of positive and negative systematic deviations and to what extent the deviations are consistent within the compared data sets. This can be shown by noting that

$$MRSE \geq MRB^2,$$

with equality in the case of fully systematic discrepancies. If MRB vs. MRSE is plotted there will be an "ideal" trend curve with the shape of a parabola. The extent to which two data sets are consistent can be assessed by the distance of their MRB/MRSE pair to the trend curve.

Table 3.1 and 3.2 list the MRB and MRSE metrics for comparisons of the dispersion from sources of different heights and horizontal locations. The MRB/MRSE calculations presented here are based on data in a uniform grid spanning the entire z = 1.5 m plane. The tabular data is summarized in Figure 3.5, in which the aforementioned trend curve is also included. As discussed, the figure clearly shows that source height is rather insignificant, whereas the horizontal location of the source has a large effect on the dispersion pattern.

As seen from the MRB/MRSE data, the maximum systematic deviation between different release heights is MRB = -0.0645, corresponding to a -6.25 % deviation in concentration from the ground-level release at location B relative to the release at 1.5 m. For the comparison between different horizontal locations, on the other hand, the maximum deviation is 15.3 % (MRB = 0.1651). Moreover, the scatter in the data is generally below MRSE = 0.02 for the release height comparisons, whereas it lies between 0.2 and 0.35 for the horizontal changes in location. The latter fact indicates that *the dispersion pattern varies considerably when changing the horizontal location*, even if the average concentration in the room does not.

The differences between two data sets can also be evaluated from scatter plots in which the data values from each data set is plotted against each of the two respective axes. Thus, the straight line

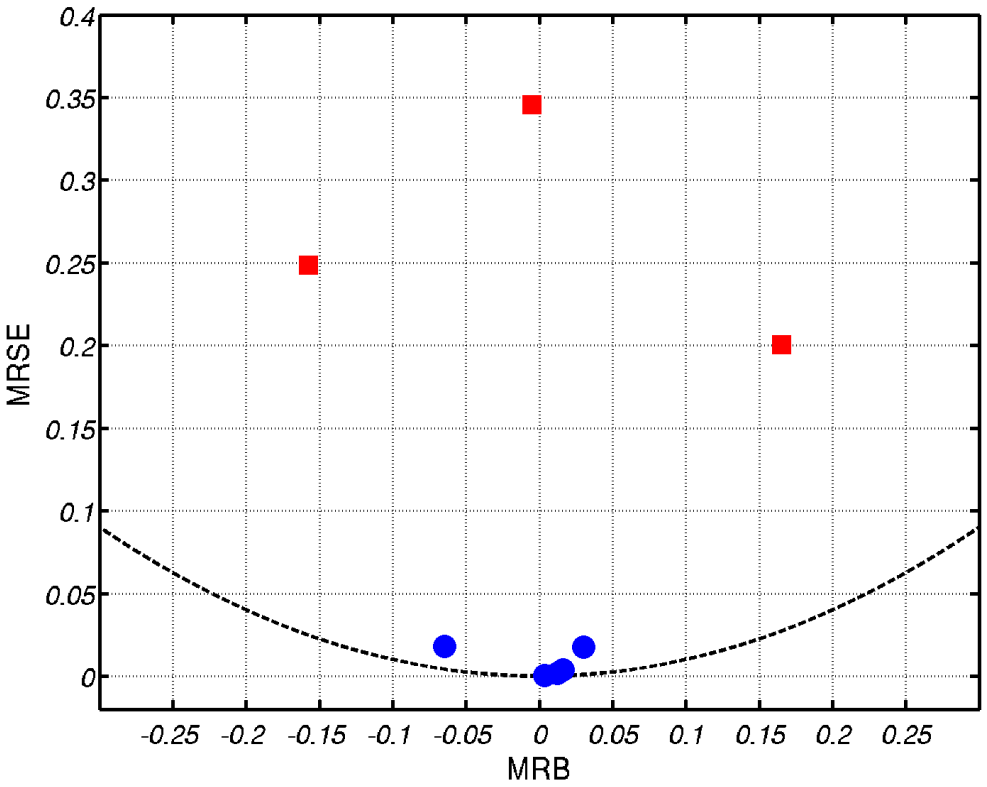

Figure 3.5 Statistical measures of differences between data sets, MRB vs. MRSE, for the present simulations 20 minutes after the releases have stopped. Blue circles correspond to comparison of variation in source height, whereas red squares represent the comparisons between different source locations. The black line is the ideal trend curve, MRSE = MRB^2.

Table 3.1 Statistical measures of differences between releases at different heights 20 minutes after the releases have stopped; Mean Relative Bias (MRB) and Mean Relative Square Error (MRSE).

Location	Heights compared [m]	MRB	MRSE
A	0 vs. 1.5	0.0038	0.0004
A	0 vs. 8.8	0.0164	0.0037
A	1.5 vs. 8.8	0.0126	0.0018
B	0 vs. 1.5	-0.0645	0.0178
C	0 vs. 1.5	0.0300	0.0175

Table 3.2 Statistical measures of differences between releases at different (horizontal) locations 20 minutes after the releases have stopped; Mean Relative Bias (MRB) and Mean Relative Square Error (MRSE).

Location	MRB	MRSE
A vs. B	0.1651	0.2006
A vs. C	-0.0049	0.3458
B vs. C	-0.1576	0.2488

$x = y$ represents the "ideal" trend line; if all points fall on this line, the two data sets are identical. Points further from the line indicate larger discrepancies between the data sets. Figure 3.6 shows such scatter plots for all the comparisons listed in Table 3.1 and 3.2.

The scatter plots in Figure 3.6 correspond well to the results already discussed. For example, the top left plot, comparing release heights 0 and 1.5 m at location A, indicates a strong correlation in the data. On the other hand, the bottom-left plot shows how the data correlation between dispersion from location A and dispersion from location C is virtually non-existent and with a lot of scatter. Overall, the scatter plots also support the notion that the source height is less important than the horizontal location of the source.

There are two main interrelated aspects of this particular flow that likely cause the exact source height to be insignificant: Firstly, the geometry is rather homogeneous in the vertical direction; only low dividing walls, some variation in the height of the ceiling, and very small ventilation pipes at select locations provide vertical inhomogeneities.

Additionally, the turbulent mixing is very efficient (as discussed in Section 3.1), which diffuses the concentration field and thus reduces differences in small-scale features between dispersion from different source heights.

Note that even though the height of the source is a near-negligible parameter for the present case, this is not a general conclusion; other geometries or ventilation systems might cause the height to be a critical parameter.

Due to the strongly horizontally inhomogeneous geometry, all the walls in particular, the significant effect of horizontal source location is no surprise. As time passes, however, the differences between the concentration fields from the different sources also decrease due to turbulent mixing.

3.2.1.1 Health effects

In this section, the health effects of different horizontal source locations will be evaluated by means of dosage calculations combined with the LD50 criteria, cf. Section 2.4.4. Differences in source

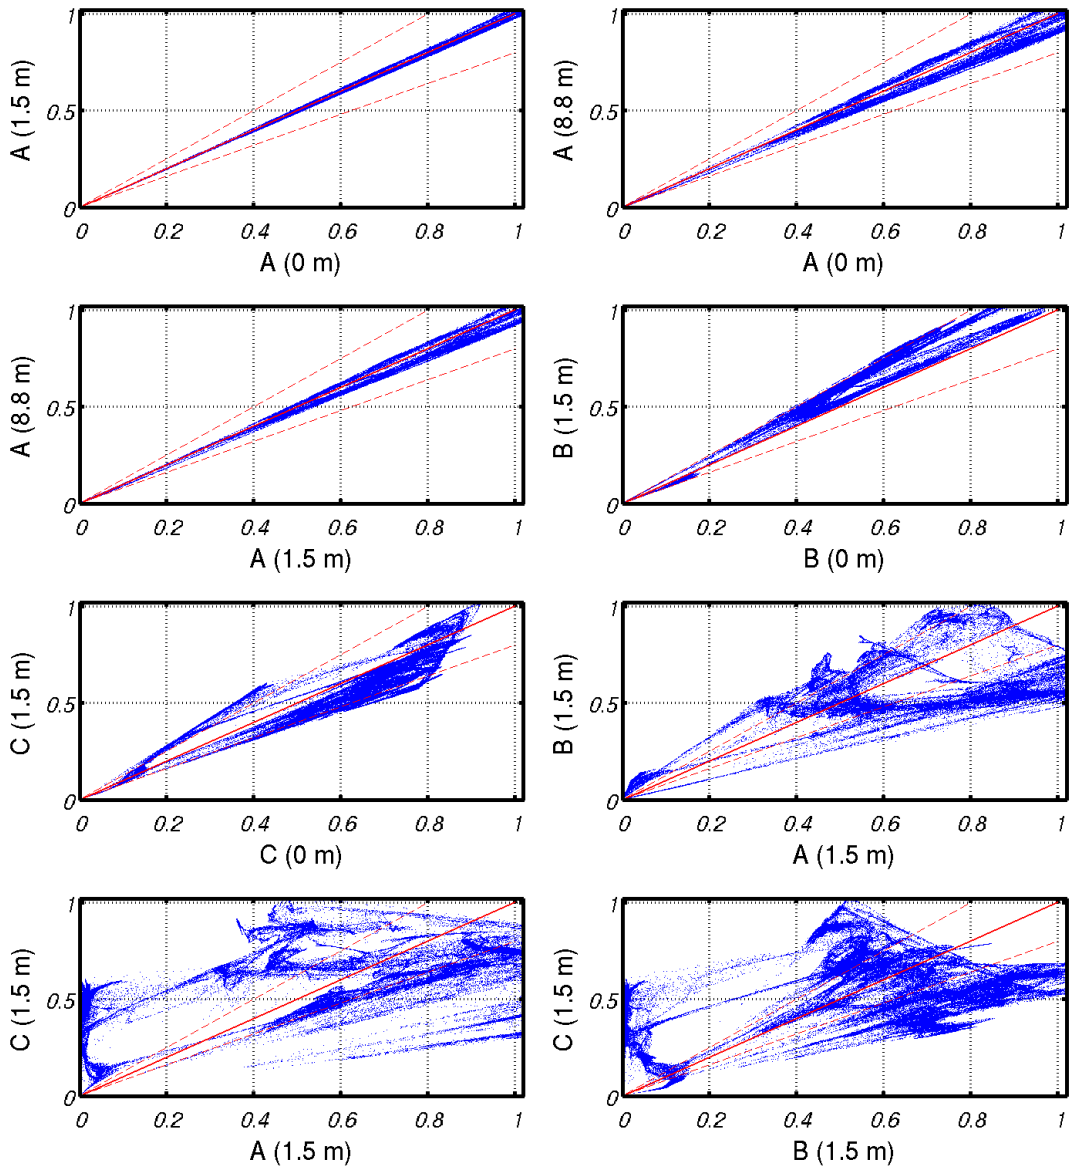

Figure 3.6 Scatter plots comparing $\tilde{c}_* = \tilde{c}/C_{perf}$ for two different release locations, as noted on the horizontal and vertical axes of each respective subplot. The scattered points represent grid data in a cutplane at height $z = 1.5$, taken 20 minutes after the releases have stopped. The solid red line indicates the "ideal" trend line ($y = x$), on which all points would fall if the data sets were identical. The two dashed red lines indicate the 20 % error margins.

height will not be discussed, in light of results in previous sections. Further health considerations for one single release is discussed a bit more extensively in Section 3.2.3.

As discussed in Section 2.4.4 (and its subsections), several parameters aside from the CB agent concentration itself will affect the mortality rate in an exposed population; properties of the exposed population (overall health, body mass, breathing rate, etc.) as well as the exposure time can be critical. In the following, it has been assumed that the average weight of the population is 70 kg, that the average breating rate is 10 liters/min and that most individuals are of sound health.

Figure 3.7 shows the time evolution of the average dosage accumulated for people at around head-level height when exposed to the dispersed CB agent for 5 and 30 minutes, respectively. The dosages are given in terms of LD50-threshold multiples, which implies that values above one correspond to a 50 % human casualty rate or higher. Only values significantly below one can be considered safe.

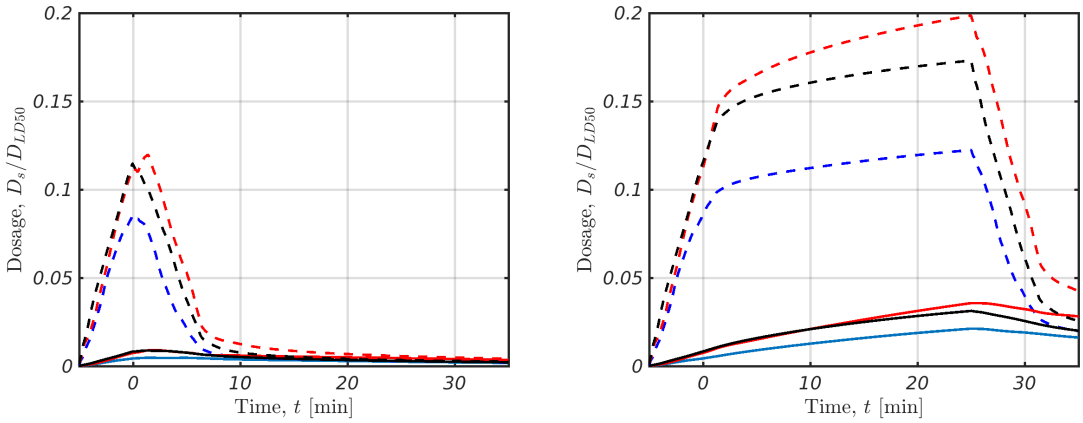

(a) 5-minute exposure period.

(b) 30-minute exposure period.

Figure 3.7 Accumulated dosages in terms of critical LD50-threshold multiples as a function of time after the release has stopped. Dosages are based on releases at 1.5 m height from locations A (blue lines), B (red lines) and C (black lines). Spatial averages (solid lines) and spatial averages plus two standard deviations (dashed lines) are shown. For consistency with the time reference used in previous figures, $t = 0$ corresponds to the time at which the CB release is stopped.

As expected, longer exposure corresponds to greater health risks. On average, the health risk is rather low for both exposure periods (for the amount presently released). However, there is large variation within the domain, as indicated by the large standard deviations seen in Figure 3.7. This is, in large part, due to the local nature of the location of the release, near which accumulated dosages will be much higher. Obviously, people close to the release location is exposed to significantly higher dosages.

It can be seen, particularly from the 30-minute exposure, that location B is the most dangerous release location in terms of health outcome. On average, it yields roughly twice the accumulated dosage compared to release from location A. This is not unexpected; as location B is in the middle of the conference hall, the released CB agent has more room to spread before being removed through the ventilation system. On the other hand, location A is situated near the corner of the room. A larger portion of the released agent will thus be transported upwards instead of horizontally outwards.

Now, consider a more potent source, either due to a more toxic CB agent or due to a larger amount of dispersed agent. Since the dosage integral is a linear operator, the previous results can simply be multiplied by an appropriate factor. For example, if the amount of CB agent was multiplied by a hundred, the LD50-scaled dosages would increase correspondingly, as shown in Figure 3.8. This would obviously have dramatic consequences for the outcome of the scenario in terms of human casualties (as shown in Section 3.2.4).

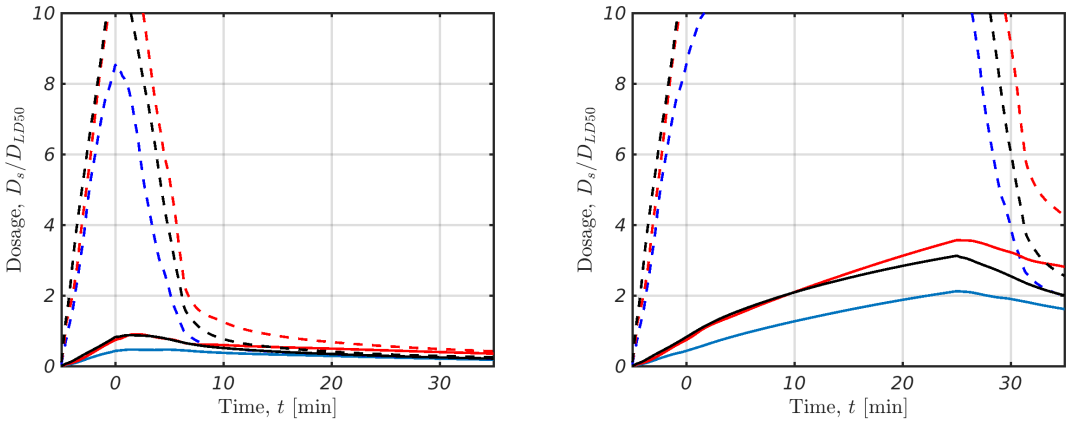

(a) 5-minute exposure period.

(b) 30-minute exposure period.

Figure 3.8 Accumulated dosages in terms of critical LD50-threshold multiples as a function of time after the release has stopped. Dosages are based on releases at 1.5 m height from locations A (blue lines), B (red lines) and C (black lines), but with 100 times more CB agent released than in Figure 3.7. Spatial averages (solid lines) and spatial averages plus two standard deviations (dashed lines) are shown. For consistency with the time reference used in previous figures, $t = 0$ corresponds to the time at which the CB release is stopped.

From Figure 3.8, it can be seen that even brief (5 minutes) exposures can be lethal, particularly in parts of the domain where concentrations are higher than average. For longer exposures, casualty rates will be severe, regardless of source location; it takes only 10 minutes before the critical LD50 threshold is passed, even for the CB agent released from location A, which yields the lowest risk among the three locations.

3.2.2 Ensemble averaging

The results discussed in previous sections represent individual release and dispersion events, i.e. single realizations. In order to assess the statistical significance of such individual simulations, the ensemble average should be estimated and considered, as discussed in Section 2.4.3.1.

In order to estimate the ensemble average, 40 different releases and subsequent dispersion processes have been simulated. Then, all the dispersion time lines have been averaged in order to produce an estimate of the (statistically transient) ensemble average. Several other ensemble-statistical estimates can be computed as well, such as the ensemble standard deviation and the ensemble maximum.

Figure 3.9 compares one of the 40 realizations to the ensemble estimate average. Qualitatively, it appears that a single realization represents a “typical” release and dispersion event quite well; there are no striking differences in either dispersion area or typical concentration values.

The ensemble average is more chaotic than might be expected. Generally, this could mean that

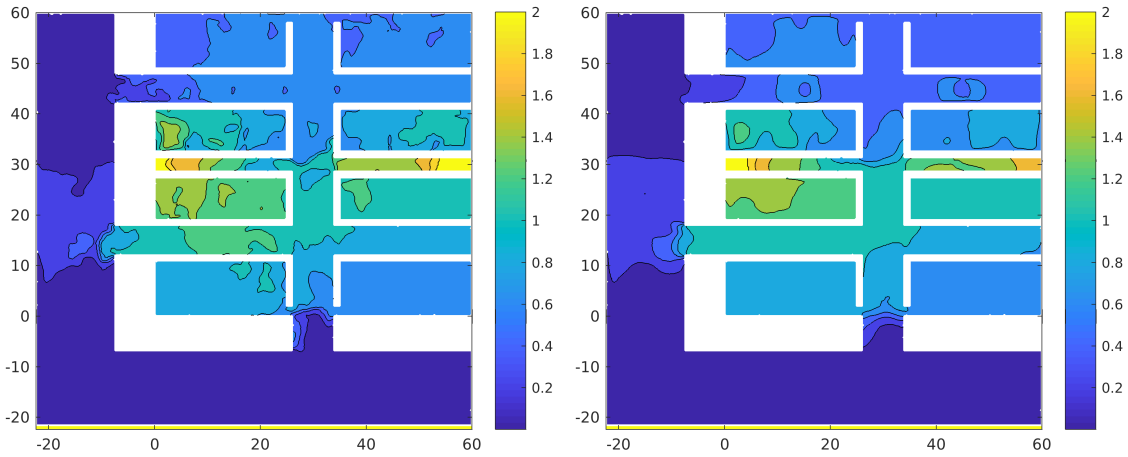

(a) Data from a single realization; $\tilde{c}_* = \tilde{c}/C_{\text{perf}}$.

(b) Ensemble-averaged data; $C_* = C/C_{\text{perf}}$.

Figure 3.9 Countours of nondimensional concentration in a cutplane at $z = 1.5$, 20 minutes after the release of CB agent has stopped.

the statistics are not converged, meaning that more realizations are required. Presently, however, the mean flow field is in fact likely to be rather chaotic, due to the complexity of the geometry and the large number of ventilation inlets. Indeed, the evolution of the ensemble average with increasing number of included realizations (cf. Appendix C) shows that an acceptable average is obtained with roughly 20 realizations for most points of interest.

The ensemble standard deviation sheds more light on the variability of different realizations; Figure 3.10a shows that the *relative* standard deviation, i.e. the standard deviation relative to the mean concentration is quite large in areas where the concentration values are low; this is not surprising and should be of little concern.

In the conference hall, the relative standard deviation is less than 0.3, usually much smaller. This also corroborates the notion that a single realization is relatively similar to the ensemble average; for example, at the time shown in Figure 3.9, the mean concentration is $C_* \approx 1$ in the area around location B, and the corresponding relative standard deviation is generally $\lesssim 0.2$. This implies that, statistically, about 95 % of all realizations can be expected to have concentration values in the range $0.6 \leq C_* \leq 1.4$, assuming a normal distribution[6].

Another way to use the ensemble statistics is to consider the pointwise maximum value of all the realizations. This will provide a more "worst-case" estimate than a single realization or the average, as it shows what maximum concentration was obtained from 40 simulated releases at a given position and time. The ensemble maximum is visualized in Figure 3.10b. As seen, the ensemble maximum is not strikingly different from the mean in the present case. The maximum values near location B are consistent with the 95 % confidence interval given by the standard deviation, discussed in the previous paragraph.

The results shown in the above indicate that a single realization is an acceptable representation of a typical ("average") CB event for the scenario discussed here. Nevertheless, in the following section, ensemble statistics will be used for all figures and subsequent discussion.

[6]The assumption of normally distributed values in the set of realizations may be incorrect here; the assumption is used simply to illustrate a point.

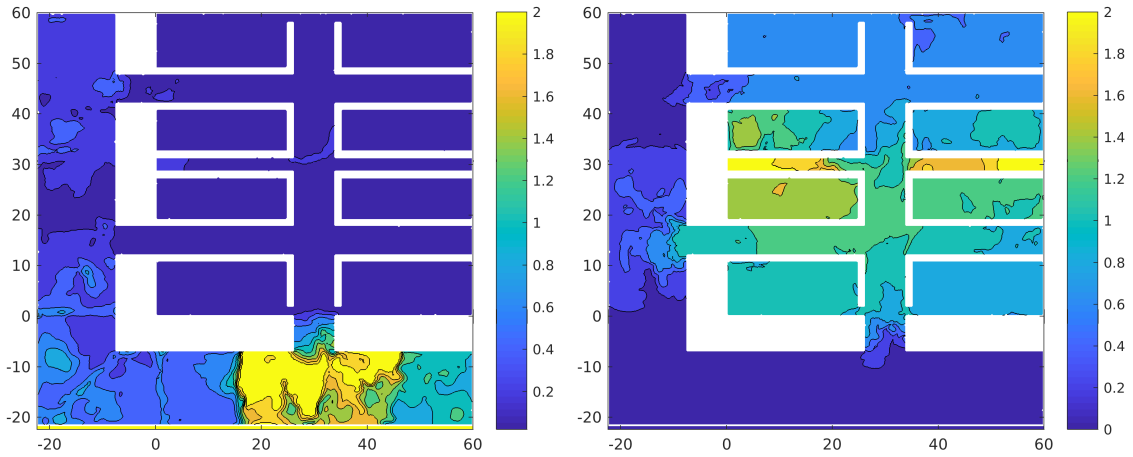

(a) Relative ensemble standard deviation of the concentration field.

(b) Ensemble-maximum nondimensional concentration.

Figure 3.10 Contours in a cutplane at $z = 1.5$, 20 minutes after the release of CB agent has stopped.

3.2.3 Temporal evolution of dispersion from location B

The evolution in time of the dispersion of CB agent released from location B (at 1.5 m) will now be considered. Figure 3.11 shows the ensemble-averaged normalized concentration field, $C_* = C/C_{\text{perf}}$ at 0, 15, 30 and 45 minutes after the release of agent has stopped.

From Figure 3.11, it can be seen that as the release stops (i.e. 5 minutes after it has started), the center of the hall is already filled with high concentrations of CB agent, whereas the corners of the room and the corridor outside is still unaffected. After 15 more minutes, however, the CB agent has dispersed significantly, filling most of the conference hall. It also begins to flow out into the corridor. At 30 and 45 minutes, the CB agent has been thoroughly mixed; the concentration is more spatially homogeneous and fills a larger volume, and the pointwise concentrations are thus lower.

It should be noted that the decrease in concentration in time is not only because of the dispersion, which "spreads" the CB agent over a larger volume. Additionally, polluted air is constantly being ejected from the domain through the ventilation outlets, while new, clean air is introduced trough the inlets. From Figure 3.12, it can be seen that 30 minutes after the release has stopped, only about half of the agent remains in the domain. After an hour, only 20 % of the agent remains.

The temporal evolution of the cloud of CB agent can also be summarized in a single figure by means of so-called "arrival time", i.e. the time at which a non-negligible concentration of agent reaches a location. Figure 3.13 shows that the dispersion within the conference hall is rather quick, whereas it takes considerably longer for the CB agent to reach the corridor outside. This is consistent with the results shown in Figure 3.11.

Figure 3.14 illustrates the accumulated dosage to which people at different locations in the domain will be exposed if they remain in the same place during the dispersion. The dosages are shown in terms of LD50 multiples, for which values > 1 indicate at least a 50 % statistical mortality rate (cf. Section 2.4.4).

The dosages are generally low; in most of the domain, the accumulated dosage does not exceed 10 % of the LD50 threshold for the type and amount of agent used presently. Unsurprisingly, close to the source there is considerably higher risk of injury. It can also be seen from Figure 3.14 that

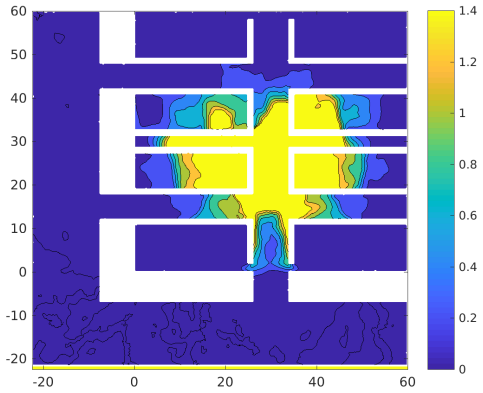

(a) Time $t = 0$ min after release has stopped.

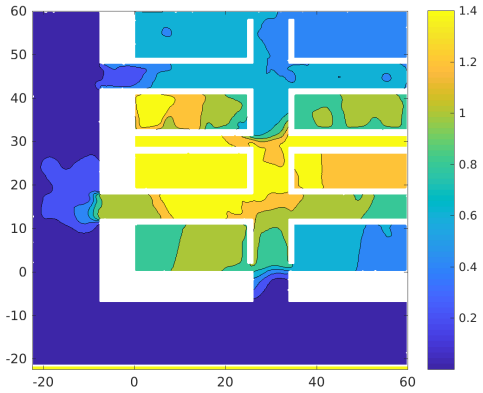

(b) Time $t = 15$ min after release has stopped.

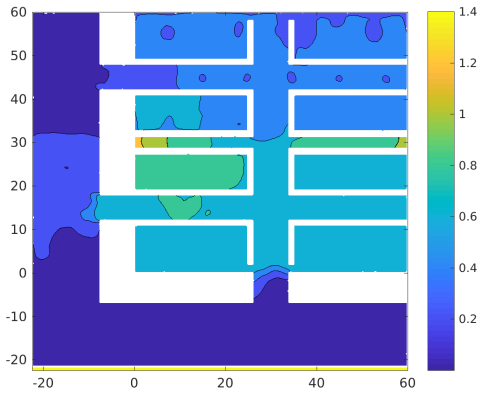

(c) Time $t = 30$ min after release has stopped.

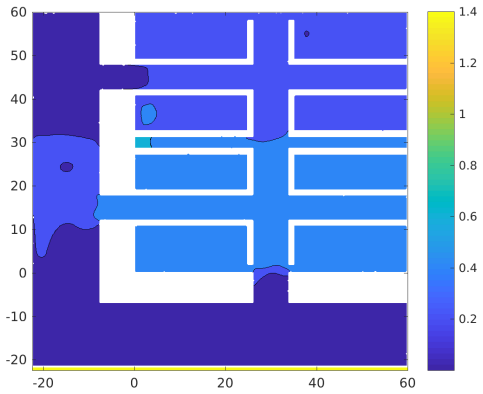

(d) Time $t = 45$ min after release has stopped.

Figure 3.11 Countours of ensemble-averaged nondimensional concentration, $C_* = C/C_{perf}$ in a cut-plane at $z = 1.5$ m.

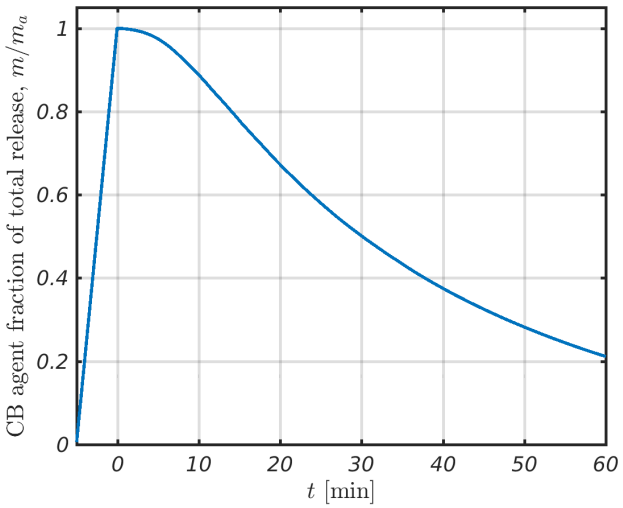

Figure 3.12 *Temporal evolution of the total amount of CB agent in the domain relative to the total amount released. For consistency with the time reference used in previous figures, $t = 0$ corresponds to the time at which the CB release is stopped.*

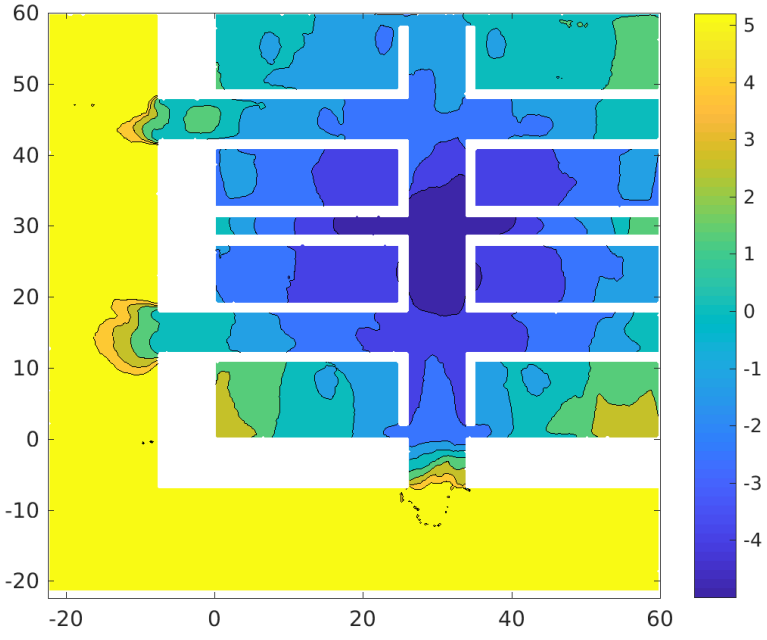

Figure 3.13 *Contours of arrival time for the cloud of CB agent, i.e. the earliest time at which $C_* > C_{thr}$, where $C_{thr} \sim 10^{-2}$ is a threshold value. For consistency with the time reference used in previous figures, $t = 0$ corresponds to the time at which the CB release is stopped. (Negative values thus correspond to minutes* before *the release is stopped.)*

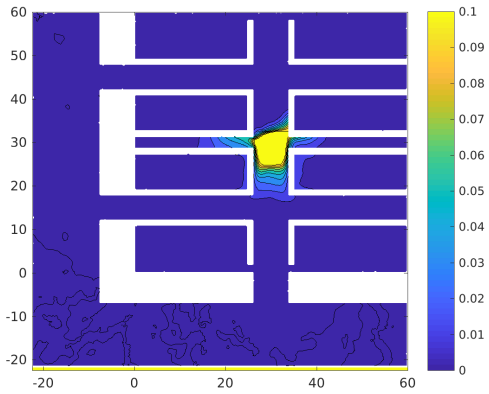

(a) Time $t = 0$ min after release has stopped (i.e. 5-minute exposure).

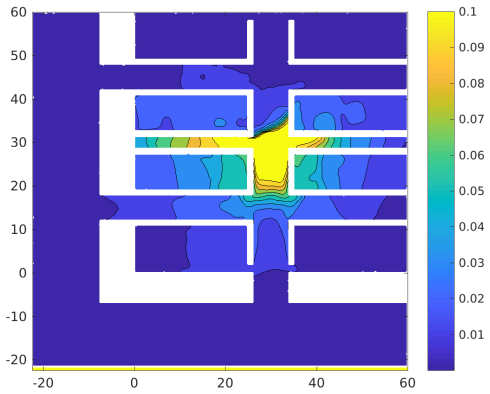

(b) Time $t = 15$ min after release has stopped (i.e. 20-minute exposure).

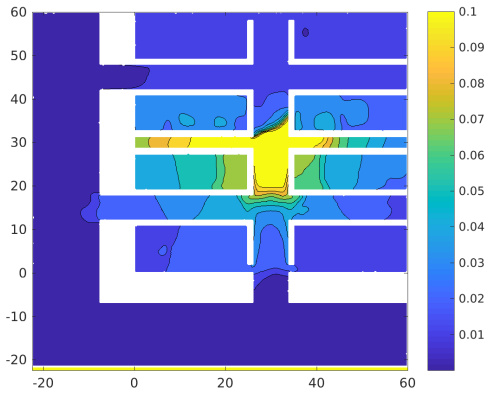

(c) Time $t = 30$ min after release has stopped (i.e. 35-minute exposure).

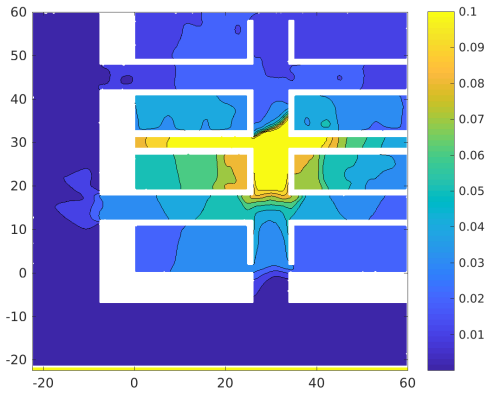

(d) Time $t = 45$ min after release has stopped (i.e. 50-minute exposure).

Figure 3.14 Countours of LD50-normalized accumulated dosage, $D_s/D_{LD50}m_p$ (cf. Section 2.4.4), in a cutplane at $z = 1.5$ m.

Table 3.3 Percentage sizes of a population in the conference hall exposed to different accumulated dosage intervals (relative to the LD50 threshold) of CB agent, 30 minutes after the release. The "original amount" corresponds to the amount of agent used in the original simulations discussed throughout this report.

D_s/D_{LD50} interval	Original amount	$100 \times$ orig. amount
< 0.5	100 %	10.8 %
$[0.5, 1)$	0 %	11.6 %
$[1, 2)$	0 %	22.4 %
≥ 2	0 %	55.2 %

after the inital release and subsequent 15 minutes, not much changes with respect to the pointwise accumulated dosage. This is consistent with the previous results which have shown that the cloud is vey diluted late in the dispersion process, and much of the CB agent has left the domain by means of the ventilation outlets.

Note, however, that for a more potent CB agent (or, equivalently, a larger amount of the present agent) the risk would be much higher than exemplified here. For example, if the amount released was multiplied by a hundred, the dosages would increase correspondingly. This would imply a dosage of more than one LD50-threshold for a large portion of the conference hall after only 15 minutes. Also keep in mind that the LD50 threshold implies a 50 % mortality rate, meaning that dosages should be significantly below this level to ensure no loss of life.

3.2.4 A practical example: Estimated casualty rates

The results discussed in the previous section, particularly the accumulated dosages, can be used to estimate casualty rates.

There are several choices involved in estimating casualties; if desired, complex models for human behavour can be used to make assumptions about walking paths, distribution of people etc. Similarly, assumptions about the population can be arbitrarily complex with regards to health level, body mass, respiration rates etc. In either case, the underlying concentration field is used to compute the dosages people are exposed to, cf. Section 2.4.4.

Presently, the following simple assumptions have been applied to exemplify how casualty rates can be estimated:

- The population has an average body mass of 70 kg.
- The population has an average respiration rate of 10 l/min.
- The population is generally healthy.
- The deposition efficiency of the CB agent in the airways is 100 %.
- The entire population is in the conference hall (i.e. the corridors are empty), and no one exits the hall.
- The concentration of people (i.e. their distribution in the conference hall) is uniform.

By applying these assumptions, in conjunction with the concentration field and dosage integrals, the fractions of the population exposed to different accumulated dosages can be computed. Table 3.3 give the sizes of such population fractions after 30 minutes for two different released amounts of CB agent.

Table 3.3 clearly reveals the significant differences in health effects from a small to a large

release of a CB agent. In the former case, no one is exposed to dosages higher than half the LD50 threshold, whereas for the large release, half the population are exposed to dosages more than twice the threshold.

It is not straight-forward to predict the death toll of a population purely from the dosages given in Table 3.3, for reasons given in Section 2.4.4. However, coarse estimates can be calculated, given crude assumptions about dosage response. Such assumptions are usually highly uncertain and must be regarded with great care.

For example, it could be assumed that if the dosage exceeds the LD50 threshold, regardless of how much, 50 % of the population dies. If the dosage is below the LD50 threshold, it is assumed that no one dies. This assumption is lacking for several reasons; for example, note that this assumption can never yield a casualty rate above 50 %, which is unrealistic for many situations.

Another assumption is that of a linear response to the accumulated dosage. For such an assumption, the casualty rate is directly related to the dosage by means of a linear function. That is, one could assume that for a dosage of one LD50 threshold, half the population dies, whereas for 0.5 and 2 times the LD50 threshold, a quarter and all of the population dies, respectively.

Given the latter of these two assumptions, a crude estimate of the number of casualties can be found. For example, if there are 500 people in the conference hall during the release and dispersion of the CB agent, the complete dosage data set (which is merely summarized in Table 3.3) implies that *for the small (original) release, approximately 7 people (1.4 %) will die*, whereas *for the larger release, 358 people (71.6 %) will die*.

4 Concluding remarks

4.1 Concluding remarks

It has been shown how CFD can be used to numerically simulate a CB release and dispersion event and thus predict the event's outcome in terms of health risks and potential casualty rates. The general CFD framework has been summarized, and the ability to easily assess effects of changing parameters, such as the location of the release or the amount of CB agent released, has also been exemplified. Moreover, a brief discussion of health-effect models and dosage calculations based on CFD data has been given.

In the particular scenario discussed in this report, it is seen that the height of the release location is relatively insignificant, while the horizontal location can have a larger effect on the dispersion of the CB agent. For all locations, the dispersion progresses rather quickly within the conference hall, but it takes more time ($\gtrsim$ 15 min) for the CB agent to reach the outside corridor. The flow field within the conference hall is complex and highly turbulent, consistent with the complexity of the geometry and the large number of ventilation inlets. This leads to efficient mixing of the released agent.

For the releases originally simulated in the present study, the health risks are generally low; it can be estimated that the casualty rate will be less than 2 %, given that the area is evacuated wihin 30 minutes. However, if the amount released is increased, the outcome quickly becomes worse; for example, if the released amount of CB agent is multiplied by 100, a resulting casualty rate of more than 70 % can be crudely estimated.

Expectedly, the area close to the release location is by far the more dangerous during and shortly following the release. It takes about 30 minutes for the CB agent to be well-mixed in the conference hall. At this time, 50 % of the agent has also been ejected from the domain through the ventilation outlets.

For a specific release location and type, a single simulated realization – i.e. simulation of one CB event – seems to be an acceptable representation of the ensemble-average, i.e. a "typical" such release, for this particular scenario. Hence, the concluding remarks given in the above can be expected to be statistically relevant.

4.2 Sources of error

The error sources in this study can be divided into two types.

There are potential sources of error pertaining to the scenario parameters, such as the following.

- The geometry of the conference hall and corridor might not be up to date. The geometry is based on a configuration from earlier (publicly available) studies, and it is a fact that the conference hall often changes between conferences. In particular, the number and positions of the dividing walls are not constant.
- The ventilation system parameters used here assume maximum ventilation. Depending on the weather and number of visitors, the ventilation system can perform quite differently from what is assumed here.
- The source description for the release, and the behavior of the CB agent, have been somewhat simplified in the present study. For the case presently under consideration, the simplifications are valid with negligible error. Generally, however, the near-source dispersion might differ

significantly for gases very much denser or lighter than air or for very large spray velocitites. Finally, second-order effects like far-field droplet agglomeration and surface deposition have been neglected, which may lead to lower concentration values than what has been simulated here.

- Input to the health effect models are based in large part on empirical data for the given CB agent (if known). Improving the range and quality of such empirical data can significantly reduce the uncertainties related to health effects in this study. Furthermore, knowledge about the typical population in a given scenario (body mass, health, respiration, etc.) also increases the certainty of the conclusions.

These uncertainties are difficult to quantify. Nevertheless, the reader should keep them in mind and be aware that the conclusions given in this report are not general, but depends on specific scenario information.

The other kind of potential errors relate to the numerical simulation itself. In particular, the choice of computational mesh represents a major uncertainty of this study. Based on unsystematic test simulations on different meshes and inspection of the results as a whole, the mesh utilized in the present simulations appears to be adequate for the present use. Thus, the general conclusions in this report are most likely valid.

However, as discussed in Section 2.4.1, a systematic mesh refinement study should ideally be performed to quanitify the uncertainty related to the chosen mesh resolution. Such a study has been omitted at present time, as the main objective has been to demonstrate fluid dynamical simulations as a means of assessing risk, and to develop the appropriate data processing routines for such a use of CFD.

4.3 Future work

This report merely gives an idea of the possibilities inherent in CFD simulation of CB dispersion events. Depending on the specific scenario and what kind of assessment is desired, a range of more detailed analyses can be performed.

In particular, much more extensive parameter studies can be performed, in which different amounts of agent, different behavior models for people (such as various walking patterns and evacuation routines) and different release locations can be investigated in more detail. It could also be of interest to study the differences in scenario outcome for different room configurations, i.e. changes in how the dividing walls are set up. Similarly, different ventilation-system settings could be studied.

Lastly, if a more extensive parameter study is performed, it should be preceded by a systematic mesh study to quantify the uncertainties inherent in the choice of mesh resolution.

References

Aalbergsjø, S. and Vik, T. Indoor dispersion of sarin by evaporation from liquid pools. Tech. rep., Tech. Rep. 2016/01242, Forsvarets Forskningsinstitutt, 2016.

ANSYS. *ANSYS Fluent 18.2 User's Guide*. ANSYS, Inc., Canonsburg, PA, USA, 2017.

Bruckner, J. V., Keys, D. A., and Fisher, J. W. The acute exposure guideline level (aegl) program: applications of physiologically based pharmacokinetic modeling. *Journal of Toxicology and Environmental Health, Part A*, 67 (8-10):621–634, 2004.

Coirier, W. J., Fricker, D. M., Furmanczyk, M., and Kim, S. A computational fluid dynamics approach for urban area transport and dispersion modeling. *Environ. Fluid Dyn.*, 5:443–479, 2005.

Duijm, N. J., Ott, S., and Nielsen, M. An evaluation of validation procedures and test parameters for dense gas dispersion models. *Journal of loss prevention in the process industries*, 9 (5):323–338, 1996.

Endregard, M., Burkhart, S., Burman, J., Gentilhomme, O., Robins, A., Wingstedt, E., Petterson Reif, B., Persson, L., Brännstrøm, N., Parmhed, O., , Björnham, O., Leroy, G., Eriksson, D., Vik, T., Tørnes, J., and Issartel, J.-P. Moditic – modeling the dispersion of toxic industrial chemicals in urban environments. Tech. rep., Tech. Rep. 2016/01478, Forsvarets Forskningsinstitutt, 2016.

Endregard, M., Reif, B. A. P., Vik, T., and Busmundrud, O. Consequence assessment of indoor dispersion of sarin—a hypothetical scenario. *Journal of hazardous materials*, 176 (1):381–388, 2010.

Evensen, F. and Olsen, B. Eksplosjon i ammoniakktank. *Tidsskriftet Brannmannen*, 2002. URL `http://brannmannen.no/cbrne/eksplosjon-i-ammoniakktank/`.

Fossum, H. and Petterson Reif, B. Computational fluid dynamics simulations of flow and dispersion in the aerosol test chamber. Tech. rep., Tech. Rep. 2012/01030, Forsvarets Forskningsinstitutt, 2012.

Fossum, H., Reif, B., Tutkun, M., and Gjesdal, T. On the use of computational fluid dynamics to investigate aerosol dispersion in an industrial environment: A case study. *Boundary-Layer Meteorology*, 144:21–40, 2012.

Germano, M., Piomelli, U., Moin, P., and Cabot, W. A dynamic subgrid-scale eddy viscosity model. *Physics of Fluids A: Fluid Dynamics*, 3:1760, 1991.

Gill, D. M. Bacterial toxins: a table of lethal amounts. *Microbiological reviews*, 46 (1):86, 1982.

Gjesdal, T., Helgeland, A., Vik, T., Petterson Reif, B. A., and Endregard, M. Simulering av utslipp og spredning fra et lagringsanlegg for svoveldioksid. Tech. rep., Tech. Rep. 2013/02880, Forsvarets Forskningsinstitutt, 2013.

Gresho, P. M. Incompressible fluid dynamics: Some fundamental formulation issues. *Annual Review of Fluid Mechanics*, 23 (1):413–453, 1991. Haf:003-18.

Ham, F. and Iaccarino, G. Energy conservation in collocated discretization schemes on unstructured meshes. *Annual Research Briefs*, 2004:3–14, 2004.

Ham, F., Mattsson, K., and Iaccarino, G. Accurate and stable finite volume operators for unstructured flow solvers. Tech. rep., Center for Turbulence Research, 2006. Haf:001-03.

Keating, A., Piomelli, U., Balaras, E., and Kaltenbach, H.-J. A priori and a posteriori tests of inflow conditions for large-eddy simulation. *Physics of Fluids*, 16:4696, 2004.

Lateb, M., Meroney, R., Yataghene, M., Fellouah, H., Saleh, F., and Boufadel, M. On the use of numerical modelling for near-field pollutant dispersion in urban environments- a review. *Environmental Pollution*, 208:271–283, 2016.

Lien, F.-S. and Yee, E. Numerical modelling of the turbulent flow developing within and over a 3D building array, Part I: A high-resolution Reynolds-Averaged Navier-Stokes approach. *Boundary-Layer Meteorology*, 112:427–466, 2004.

Lien, F.-S., Yee, E., Ji, H., Keats, A., and Hsieh, K. J. Progress and challenges in the development of physically-based numerical models for prediction of flow and contaminant dispersion in the urban environment. *Int. J. of Comput. Fluid Dyn.*, 20:323–337, 2006.

Lilly, D. K. The representation of small scale turbulence in numerical simulation experiments. In *Proceedings of the IBM Scientific Computing Symposium on Environmental Sciences*, vol. Form no. 320-1951, pp. 195–210 (IBM, 1967).

Lilly, D. K. A proposed modification of the Germano subgrid-scale closure method. *Physics of Fluids A: Fluid Dynamics*, 4 (3):633–635, 1992.

Mahesh, K., Constantinescu, G., Apte, S., Iaccarino, G., Ham, F., and Moin, P. Progress toward large-eddy simulation of turbulent reacting and non-reacting flows in complex geometries. *Annual Research Briefs*, pp. 115–142, 2002.

Mahesh, K., Constantinescu, G., and Moin, P. A numerical method for large-eddy simulation in complex geometries. *Journal of Computational Physics*, 197:215–240, 2004. Haf:001-01.

Osnes, A., Vartdal, M., and Petterson Reif, B. Numerical simulation of particle jet formation induced by shock wave acceleration in a hele-shaw cell. *Submitted: Shock Waves*, N/A:N/A, 2017.

Pita, R. and Domingo, J. The use of chemical weapons in the syrian conflict. *Toxics*, 2 (3):391–402, 2014.

Pope, S. B. *Turbulent Flows* (Cambridge University Press, Cambridge, United Kingdom, 2000).

Randhawa, M. A. Calculation of ld50 values from the method of miller and tainter, 1944. *J Ayub Med Coll Abbottabad*, 21 (3):184–5, 2009.

Sagaut, P. *Large Eddy Simulation for Incompressible Flows: An Introduction*. Scientific Computation (Springer, 2006), 3rd ed. ISBN 9783540263449. URL `http://books.google.no/books?id=ODYiH6RNyoQC`.

Santiago, J. L., Martilli, A., and Martin, F. Cfd simulation of airflow over a regular array of cubes. part i: Three-dimensional simulation of the flow and validation with wind-tunnel experiments. *Boundary-Layer Meteorology*, 122:609–634, 2007.

Smagorinsky, J. General circulation experiments with the primitive equations: I. The basic experiment. *Monthly Weather Review*, 91 (3):99–164, 1963.

Tu, A. T. Toxicological and chemical aspects of sarin terrorism in japan in 1994 and 1995. *Toxin reviews*, 26 (3):231–274, 2007.

Versteeg, H. K. and Malalasekera, W. *An Introduction to Computational Fluid Dynamics: The Finite Volume Method* (Pearson Education, 1995).

Vik, T., Tørnes, J., and Petterson Reif, B. Simulations of the release and dispersion of chlorine and comparison with the jack rabbit field trials. Tech. rep., Tech. Rep. 2015/01474, Forsvarets Forskningsinstitutt, 2015.

Weil, C. S. Economical ld50 and slope determinations. *Drug and chemical toxicology*, 6 (6):595–603, 1983.

Wenck, M. A., Van Sickle, D., Drociuk, D., Belflower, A., Youngblood, C., Whisnant, M. D., Taylor, R., Rudnick, V., and Gibson, J. J. Rapid assessment of exposure to chlorine released from a train derailment and resulting health impact. *Public health reports*, 122 (6):784–792, 2007.

Wingstedt, E. M., Fossum, H. E., and Reif, B. A. P. Simulation of bomb residue deposition following the oslo bombing july 22, 2011. Tech. rep., Tech. Rep. 2012/01836, Forsvarets Forskningsinstitutt, 2012a.

Wingstedt, E. M., Reif, B. A. P., and Fossum, H. E. On a small-scale indoor aerosol release. Tech. rep., Tech. Rep. 2012/00268, Forsvarets Forskningsinstitutt, 2012b.

Wingstedt, E. M. M., Osnes, A. N., Åkervik, E., Eriksson, D., and Reif, B. P. Large-eddy simulation of dense gas dispersion over a simplified urban area. *Atmospheric Environment*, 152:605–616, 2017.

Wu, X. Inflow turbulence generation methods. *Annual Review of Fluid Mechanics*, 49:23–49, 2017.

Wyngaard, J. *Turbulence in the Atmosphere* (Cambridge Univ Pr, 2010).

A Subgrid-viscosity modeling

The most common class of subgrid-stress models are subgrid-viscosity models, akin to the Boussinesq eddy-viscosity hypothesis originally used within the RANS modeling framework. Lilly (1967) assumed that the subgrid stresses were proportional to the resolved strain rate and suggested that Smagorinsky's 1963 eddy-viscosity form be used. In present-day terminology, that implies that

$$\tau_{ij}^{\text{Lilly}} = -2\nu_s \tilde{s}_{ij}^r, \tag{A.1}$$

where the subgrid viscosity $\nu_s(\boldsymbol{x}, t) = (C_d L_\Delta)^2 \sqrt{2\tilde{s}_{ij}^r \tilde{s}_{ij}^r}$, in which L_Δ is the LES filter width, C_d is the Smagorinsky coefficient, and $\tilde{s}_{ij}^r = (1/2)(\partial_j \tilde{u}_i^r + \partial_i \tilde{u}_j^r)$ is the resolved strain rate.

In practice, L_Δ is implicitly determined by the local computational cell size in most FVM solvers (see also Section 2.3.1).

In FVM-based LES solvers, the scalar subgrid flux, f_i, is frequently modeled similarly to the momentum subgrid stresses, i.e. by a subgrid-diffusivitiy model of the form $f_i = \gamma_s \partial_i \tilde{\theta}$. Here, γ_s is the subgrid diffusivity, for which various models exist, such as the very common relation $\gamma_s = \nu_s / \text{Pr}_s$, where the subgrid turbulent Prandtl number lies in the range $0.1 \leq \text{Pr}_s \leq 1$ in the literature, with $\text{Pr}_s = 0.6$ being typical (Sagaut, 2006, p. 463).

The original Smagorinsky model suffers from the fact that the Smagorinsky coefficient inherently depends on the flow regime, from zero in laminar flow to about $C_d \approx 0.15$ in high-Reynolds number turbulence (Pope, 2000, p. 619). The *dynamic* Smagorinsky model resolves this issue by determining a suitable local value of the coefficient. The model, proposed by Germano et al. (1991), can be outlined as follows (for convenience, the tilde symbols denoting full instantaneous fields are dropped for the time being):

Given a subgrid-viscosity model,

$$\tau_{ij} = 2C_d L_\Delta^2 |s^r| s_{ij}^r,$$

in which $|s^r| = \sqrt{2 s_{ij}^r s_{ij}^r} = 2C_d \alpha_{ij}$, and the superscript r denotes resolved (filtered) variables as usual, with the associated filter width L_Δ, a second model can be defined as

$$\tau_{ij}^* = 2C_d L_{\Delta*}^2 |s^{r*}| s_{ij}^{r*} = 2C_d \beta_{ij}.$$

Here, the superscript $r*$ denotes a second LES filter, with an associated filter size $L_{\Delta*}$. The second filter is usually referred to as a *test filter*, and its filter size is often around twice the size of the ordinary filter.

Now, letting $L_{ij} = u_i^{r*} u_j^{r*} - (u_i^r u_j^r)^{r*}$, the properties of the LES filter implies that

$$-L_{ij} = \tau_{ij}^* - (\tau_{ij})^{r*},$$

referred to as *Germano's identity*. Furthermore,

$$\begin{aligned} -L_{ij} &= 2(C_d \beta_{ij} - (C_d \alpha_{ij})^{r*}) \\ L_{ij} &\approx 2C_d(\alpha_{ij}^{r*} - \beta_{ij}) \\ &= 2C_d L_\Delta^2 M_{ij} \\ C_d &= \frac{L_{ij}}{2L_\Delta^2 M_{ij}}, \qquad \text{(no summation on } i, j) \end{aligned}$$

where $M_{ij} = (L_{\Delta*}/L_{\Delta})^2 |s^{r*}| s^{r*}_{ij} - (|s^r| s^r_{ij})^{r*}$. To improve the robustness of Germano's model, Lilly (1992) introduced a normalization and local spatial averaging procedure according to

$$C_d = \frac{1}{2L_{\Delta}^2} \frac{\langle L_{ij} M_{ij} \rangle}{\langle M_{rs} M_{rs} \rangle},$$

which is the basis for the dynamic Smagorinsky models used in the present study.

B Autocorrelation of the velocity signal

The autocorrelation function of a finite time series for a statistically steady velocity component, u_i, can be written

$$R_u(t) = \frac{1}{T\langle u_\alpha u_\alpha \rangle} \int_0^T u_\alpha(\tau) u_\alpha(\tau + t)\, \mathrm{d}\tau,$$

in which u_α implies use of velocity component α with *no summation on double indices*, t is the signal delay, and T is the chosen time range over which to average the series.

The temporal autocorrelation function measures the average correlation of a data signal with itself for varying time delays. A value of 1 (or -1) implies perfect correlation (or anticorrelation), whereas a value of zero implies no correlation at all. The correlation time of the signal is approximately the time at which the temporal autocorrelation approaches zero (within some threshold).

In Figure B.1, the temporal autocorrelation functions for components of the (fluctuating) velocity field, u_i, is plotted for two different locations in the domain. As seen, the evolution of the autocorrelation function with increasing time delay suggests that the correlation time estimate from Section 2.4.3.1 is appropriate; at (30, 10, 1.5), the autocorrelation decreases rapidly and, within 5 s, fluctuates around zero. At (30, 30, 3.3), the correlation time is longer; within roughly 30 s, the x and z components seem to approach zero, whereas the y component is slower.

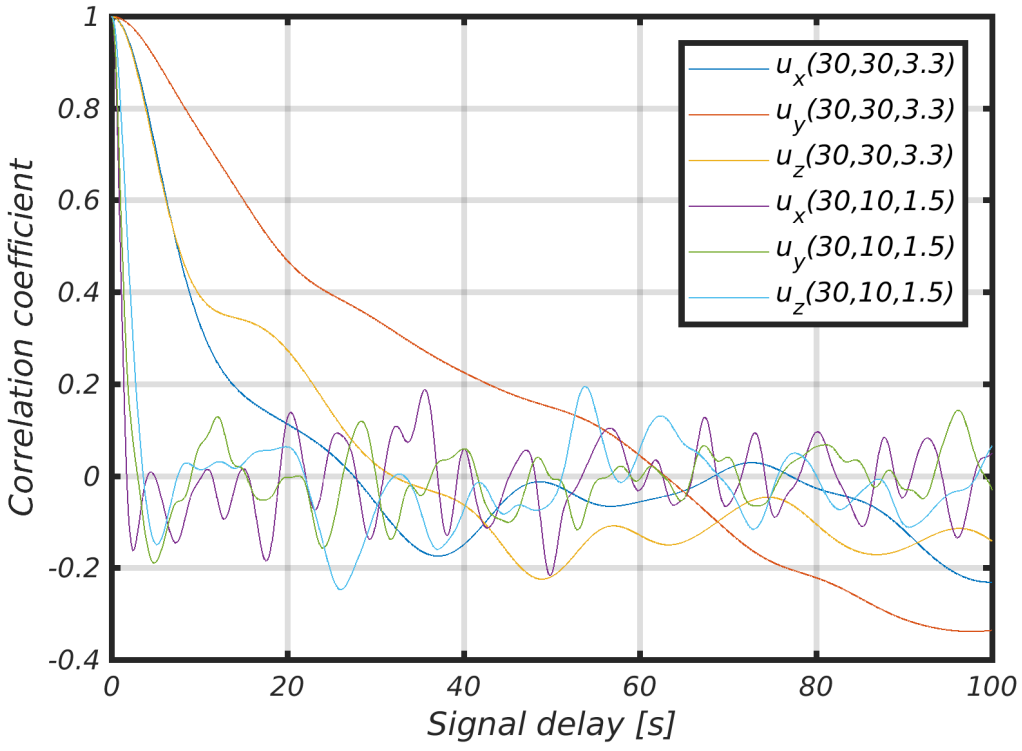

Figure B.1 Temporal autocorrelation functions for the (fluctuating) velocity field, u_i, at two different locations.

Only two locations are shown in Figure B.1, for clarity, but autocorrelation data at other locations corroborate the general picture discussed above: In the $z = 1.5$ plane (and below), the correlation

time generally seems rather short ($\lesssim$ 10 s), whereas in the $z = 3.3$ plane, the correlation time is longer ($\sim 60s$). For the ensemble-averaged dispersion realizations, which are released at $z = 1.5$, 60 s should thus be an acceptable time-separation between successive releases in order to obtain reasonably uncorrelated realizations.

C Statistical convergence of ensemble averaging

Figure C.1 shows how the approximately ensemble-averaged concentration field, $\langle \tilde{c} \rangle_R$, and the difference between successive approximations vary with respect to number of realizations used. Data from a uniform grid of 170 points in the $z = 1.5$ plane is shown in the figure. In Figure C.1b, an additional threshold of $C_* = 10^{-5}$ is used, below which data is discarded because the mean value approaches the computational tolerance of the numerical solver. This leads to a removal of 45 points, i.e. 26 %, from the data.

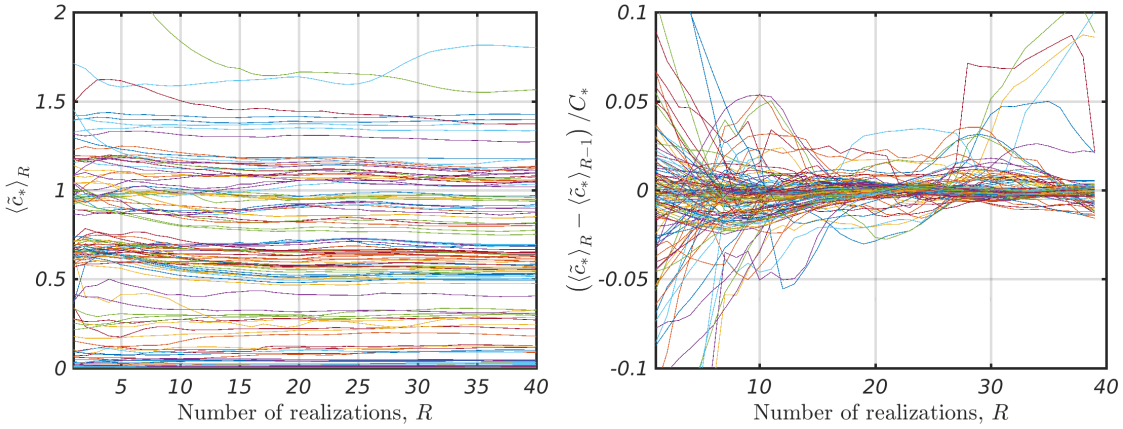

(a) Successive ensemble-mean approximations with increasing included number of realizations.

(b) Evolution of the difference between two successive ensemble-mean approximations as the included number of realizations increase.

Figure C.1 Statistical convergence of the approximately ensemble-averaged concentration field, $\langle \tilde{c} \rangle_R$, with increasing number of realizations, R.

As seen from C.1a, the ensemble-average for each data point generally appears to reach an approximately steady state at roughly 20 included realizations. This is confirmed by Figure C.1b, which shows that for most points (> 95 %), the "realization-derivative" of the ensemble-average is within 1.5 % when 40 realizations are used. These results are consistent with the estimate provided in Section 2.4.3.1.

About FFI

The Norwegian Defence Research Establishment (FFI) was founded 11th of April 1946. It is organised as an administrative agency subordinate to the Ministry of Defence.

FFI's MISSION

FFI is the prime institution responsible for defence related research in Norway. Its principal mission is to carry out research and development to meet the requirements of the Armed Forces. FFI has the role of chief adviser to the political and military leadership. In particular, the institute shall focus on aspects of the development in science and technology that can influence our security policy or defence planning.

FFI's VISION

FFI turns knowledge and ideas into an efficient defence.

FFI's CHARACTERISTICS

Creative, daring, broad-minded and responsible.

Om FFI

Forsvarets forskningsinstitutt ble etablert 11. april 1946. Instituttet er organisert som et forvaltningsorgan med særskilte fullmakter underlagt Forsvarsdepartementet.

FFIs FORMÅL

Forsvarets forskningsinstitutt er Forsvarets sentrale forskningsinstitusjon og har som formål å drive forskning og utvikling for Forsvarets behov. Videre er FFI rådgiver overfor Forsvarets strategiske ledelse. Spesielt skal instituttet følge opp trekk ved vitenskapelig og militærteknisk utvikling som kan påvirke forutsetningene for sikkerhetspolitikken eller forsvarsplanleggingen.

FFIs VISJON

FFI gjør kunnskap og ideer til et effektivt forsvar.

FFIs VERDIER

Skapende, drivende, vidsynt og ansvarlig.

FFI's organisation

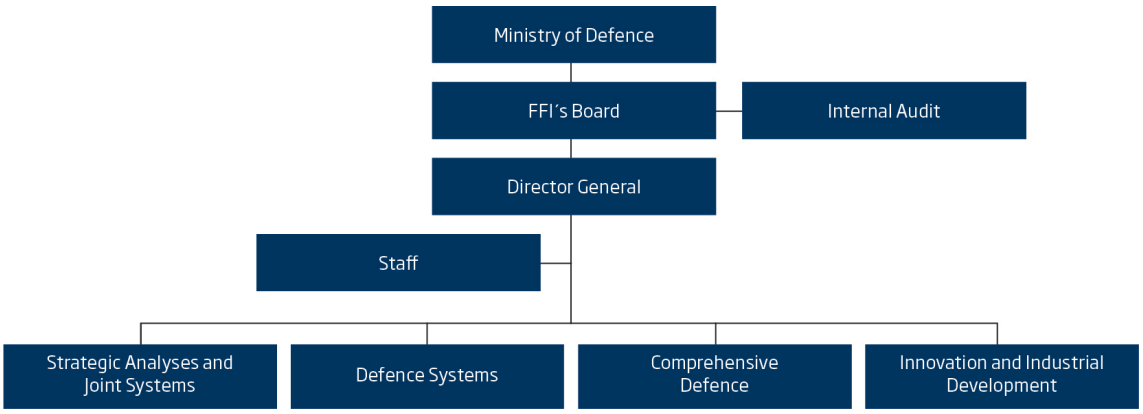

Forsvarets forskningsinstitutt
Postboks 25
2027 Kjeller

Besøksadresse:
Instituttveien 20
2007 Kjeller

Telefon: 63 80 70 00
Telefaks: 63 80 71 15
Epost: ffi@ffi.no

Norwegian Defence Research Establishment (FFI)
P.O. Box 25
NO-2027 Kjeller

Office address:
Instituttveien 20
N-2007 Kjeller

Telephone: +47 63 80 70 00
Telefax: +47 63 80 71 15
Email: ffi@ffi.no

FFI Forsvarets forskningsinstitutt
Norwegian Defence Research Establishment