Analysis of adaptive mesh refinement in a turbulent buoyant helium plume

Georgios Maragkos* | Elena Funk | Bart Merci

Department of Structural Engineering and Building Materials, Ghent University, Belgium

Correspondence
*Georgios Maragkos, Email: Georgios.Maragkos@UGent.be

Summary

The study evaluates OpenFOAM's adaptive mesh refinement (AMR) capabilities and accuracy when applied to numerical simulations of turbulent buoyant flows. To this purpose, large eddy simulations of Sandia's turbulent helium plume test case are considered, a scenario with similar characteristics (in terms of air entrainment and vortex shedding) to those encountered in large-scale fires. Comparison is made of the relative accuracy (in terms of both first and second order statistics), the predicted puffing frequencies and the computing times of numerical simulations using AMR and static meshes. Additionally, a sensitivity study on different AMR parameters is conducted and an evaluation of the performance of the dynamic Smagorinsky model when combined with AMR is made. Overall, the predictions of AMR are very satisfactory, in terms of both accuracy and computing times, compared to the use of static meshes of different sizes. Nevertheless, it is observed that a careful choice of a static mesh can result in equally accurate predictions with comparable, or even smaller, computing times than AMR for the case at hand. Additionally, the use of AMR slightly modified the entrainment characteristics in the near-field region of the plume and (significantly) altered some of the, dynamically determined, turbulence model parameters.

KEYWORDS:
AMR, FireFOAM, LES, turbulence, CFD

1 | INTRODUCTION

Nowadays, most real-life fire scenarios (e.g., involving industrial and/or high rise buildings, wildland fires) pose extreme challenges for Computational Fluid Dynamics (CFD), not only due to the different physical processes that need to be modelled (e.g., turbulence, combustion, radiation, soot, heat transfer, pyrolysis and flame spread), but also due to the wide range of length and times scales involved [1]. Both aspects directly influence the computational cost and accuracy of the numerical simulations (e.g., advanced physical modelling will be more accurate but will require higher computing times; simulations with coarse grid sizes will run faster but will lack in accuracy). A priori determination of the main regions of interest is required, so that a finer grid is employed in these areas, allowing for accurate simulations. Even though widely used in fire simulations, the use of static meshes, fixing the grid resolution from the beginning of a numerical simulation, can exhibit drawbacks depending on the case at hand. For example, a priori determination of the main regions of interest is not always straightforward (e.g., wildland or high-rise building fire simulations with variable wind conditions). Additionally, there are cases (e.g., involving flame spread)

where the a priori grid refined regions are not required throughout the whole duration of the simulation and will only lead to an (avoidable) increased computational cost. The use of adaptive mesh refinement (AMR), varying the grid resolution based on pre-defined refinement criteria, poses a promising alternative strategy for grid generation that could reduce the computing times in fire scenarios significantly.

The use of adaptive mesh refinement has been explored in the literature in the past for a wide variety of applications, including turbulent non-premixed combustion[2], liquid sprays/jets[3,4] and others[5,6,7,8,9,10,11]. Nevertheless, despite the high potential of reduction in computing times, only few researchers have investigated the use of AMR in fire applications[12,13,14]. Within this context, the present work is rather novel and aims to further pursue the application of AMR in scenarios involving turbulent buoyant flows. In order to avoid complexities and uncertainties related to combustion and radiation modelling, the consideration of a non-reacting buoyant plume is considered in this work. Overall, the gain from successfully applying and validating AMR in buoyant flows will be significant as AMR can potentially be useful for a variety of other fire-related applications, such as wildland fires, pollutant dispersion, thermal comfort in buildings, smoke transport. Additionally, it could be further applied to cases involving finite rate chemistry in order to significantly reduce the computational times associated with the use of more advanced combustion models.

The main objective of the paper is the application of AMR to a turbulent buoyant flow with similar characteristics, in terms of air entrainment and vortex shedding, as those encountered in large-scale fires. An evaluation of different refinement criteria will be performed, as well as an investigation in order to determine which AMR parameters are more important to consider for local grid resolution refinement without affecting the predictions of the numerical simulations (e.g., not altering the entrainment characteristics). These aspects have not been extensively studied in the context of buoyancy-driven flows up to now in the literature. For this reason, Large Eddy Simulations (LES) are considered, focusing on the near-field region of a turbulent buoyant helium plume[15]. The work compares the numerical predictions of AMR to results obtained with static meshes of different sizes and explores the following research aspects:

- Comparison of the relative accuracy in terms of first and second order statistics, as well as computing times of numerical simulations, using AMR and static meshes.

- Sensitivity study on different AMR related parameters (e.g., chosen refinement fields, frequency of grid refinement).

- Evaluation of the performance of the dynamic Smagorinsky turbulence model with the use of AMR (i.e., examination whether air entrainment and the turbulence model parameters are affected by the adaptive mesh refinement).

2 | NUMERICAL MODELLING

The CFD code FireFOAM 2.2.x, originally developed by FM Global, is employed for the numerical modelling, coupled with the available AMR functionality of the OpenFOAM platform[16]. FireFOAM is an LES solver that utilises the finite volume method on arbitrarily unstructured meshes and is capable of modelling complex fire applications (e.g., involving flame spread and fire suppression). A summary of the governing equations and the different sub-models employed, as well as an overview and the required input parameters of the AMR functionality used in the study, is presented below. It is worth noting that the code has been extensively validated in the past by the authors in a wide range of fire-related test cases (e.g.,[17,18,19]).

2.1 | Governing equations

The code solves the filtered Navier-Stokes equations along with transport equations for species mass fractions, using Favre-filtered quantities. An overview of the main equations is presented below.

$$\frac{\partial \overline{\rho}}{\partial t} + \nabla \cdot (\overline{\rho \widetilde{u}}) = 0 \tag{1}$$

$$\frac{\partial (\overline{\rho \widetilde{u}})}{\partial t} + \nabla \cdot (\overline{\rho \widetilde{u} \widetilde{u}}) = -\nabla \overline{p} + \nabla \cdot \left[\mu_{eff} \left(\nabla \widetilde{u} + (\nabla \widetilde{u})^T - \frac{2}{3} (\nabla \cdot \widetilde{u}) I \right) \right] + \overline{\rho} g \tag{2}$$

$$\frac{\partial (\overline{\rho} \widetilde{Y}_k)}{\partial t} + \nabla \cdot (\overline{\rho \widetilde{u}} \widetilde{Y}_k) = \nabla \cdot \left[\left(\overline{\rho} D_k + \frac{\mu_{sgs}}{Sc_t} \right) \nabla \widetilde{Y}_k \right], \quad (k = 1, ..., N_s - 1) \tag{3}$$

where $\overline{\rho}$ is the density, $\widetilde{u}$ is the velocity vector, $\overline{p}$ is the pressure, $\mu_{eff} = \mu + \mu_{sgs}$ is the effective dynamic viscosity, μ is the molecular viscosity, μ_{sgs} is the sub-grid scale viscosity, I is the identity tensor, g is the gravitational acceleration vector, $\widetilde{Y}_k$ is the species mass fraction, D_k is the species mass diffusivity and Sc_t is the turbulent Schmidt number.

2.2 | Sub-models

The dynamic Smagorinsky model[20] is used to model turbulence, calculating the sub-grid scale viscosity as:

$$\mu_{sgs} = \overline{\rho}(c_s \Delta)^2 |\widetilde{S}| \tag{4}$$

where Δ is the filter width (taken as the cube root of the cell volume) and $\widetilde{S}$ is the strain rate. A dynamic procedure is employed for determining c_s^2 as:

$$c_s^2 = \frac{\frac{1}{2}\langle L_{ij} M_{ij} \rangle}{\langle M_{ij} M_{ij} \rangle} \tag{5}$$

where the brackets denote averages which are computed as local averages of their face values. The Leonard term is defined as $L_{ij} = \widehat{\widetilde{u}_i \widetilde{u}_j} - (\widehat{\overline{\rho u_i}}\,\widehat{\overline{\rho u_j}})/\widehat{\overline{\rho}}$, $\beta_{ij} = -\widehat{\Delta}^2 \widehat{\overline{\rho}}|\widecheck{S}|(\widecheck{S}_{ij} - \delta_{ij}\widecheck{S}_{kk}/3)$, $\alpha_{ij} = -\Delta^2 \overline{\rho}|\widetilde{S}|(\widetilde{S}_{ij} - \delta_{ij}\widetilde{S}_{kk}/3)$ and $M_{ij} = \beta_{ij} - \widehat{\alpha_{ij}}$. The hat denotes the application of a test filter while $\widecheck{f} = \widehat{\overline{\rho}f}/\widehat{\overline{\rho}}$. Any negative values of the Smagorinsky parameter are clipped to zero to avoid potential numerical instabilities while no upper bound is defined. A simple top-hat filter is used as a test filter while the ratio of the test filter to the LES filter is set to $\sqrt{(6)}$[21].

The species mass diffusivity is calculated as:

$$D_k = \frac{\mu_{sgs}}{\overline{\rho} Sc_k} \tag{6}$$

where Sc_k is the molecular Schmidt number of species k.

The turbulent Schmidt number is computed dynamically as[23]:

$$Sc_t = \frac{c_s^2 \langle M_j M_j \rangle}{\langle L_j M_j \rangle} \tag{7}$$

where $M_j = -\widehat{\Delta}^2 \widehat{\overline{\rho}}|\widecheck{S}|\frac{\partial \widecheck{Y}}{\partial x_j} + \widehat{\Delta^2 \overline{\rho}|\widetilde{S}|\frac{\partial \widetilde{Y}}{\partial x_j}}$ and $L_j = \left(\frac{\widehat{\overline{\rho u_j}\,\overline{\rho Y}}}{\widehat{\overline{\rho}}} - \frac{\widehat{\overline{\rho u_j}}\,\widehat{\overline{\rho Y}}}{\widehat{\overline{\rho}}} \right)$. The Sc_t values are clipped between 0.05 and 1.0 to ensure numerical stability of the code.

The sub-grid kinetic energy is estimated as[22]:

$$k_{sgs} = c_I \Delta^2 |\widetilde{S}|^2 \tag{8}$$

with the model parameter c_I computed dynamically as:

$$c_I = \frac{\langle \frac{1}{2} L_{kk} M_{kk} \rangle}{\langle M_{kk} M_{kk} \rangle} \tag{9}$$

where the brackets denote averages, computed as local averages of their face values, with $L_{kk} = \widehat{\widetilde{u}_k \widetilde{u}_k} - \widehat{\overline{\rho u_k}}\,\widehat{\overline{\rho u_k}}/\widehat{\overline{\rho}}$ and $M_{kk} = \widehat{\Delta}^2 \widehat{\overline{\rho}}|\widecheck{S}|^2 - \widehat{\Delta^2 \overline{\rho}|\widetilde{S}|^2}$. Even though not explicitly used in the present study, the sub-grid scale kinetic energy, k_{sgs}, is a required input in the combustion models typically employed in fire modelling. Therefore, the predicted values of the model parameter c_I, involved for estimating k_{sgs}, are considered in the present study.

2.3 | AMR functionality

With the use of adaptive mesh refinement, the grid is dynamically changed within the computational domain (e.g., regions of large gradients will typically require fine grid sizes). AMR allows the initiation of numerical simulations on a relatively coarse grid, which is subsequently refined and un-refined, based on pre-defined criteria, depending on the spatial location within the domain. Therefore, relatively fine grid sizes are only used in the main regions of interest (e.g., regions with high strain rate, vorticity or large gradients) and not throughout the entire computational domain. The mesh refinement is performed by adding nodes along the mid-points of the computational cells (i.e., a cell is split into 8 cells) within regions determined from criteria specified by the user. However, the mesh adaptation does create an overhead time, which is not there with a static mesh.

Nevertheless, the benefit of AMR is that it can (potentially) lead to a reduction in the total number of cells required compared to the use of static meshes.

A general overview of the most important AMR parameters used in the OpenFOAM platform[16] is as follows:

- refineInterval (RI): Defines the frequency (i.e., number of time steps) for performing the mesh refinement.

- maxRefinement (MR): The levels of grid refinement (i.e., different grid sizes) used for AMR.

- nBufferLayers: The number of cells between each grid refinement level.

- maxCells: The (approximate) maximum number of cells allowed after grid refinement.

- field: The variable upon which the grid refinement is based.

- lowerRefineLevel (LRL): The minimum value of the variable used for grid refinement.

- upperRefineLevel (URL): The maximum value of the variable used for grid refinement.

- unrefineLevel: The variable value below which the grid is unrefined back to the base mesh size (set to the lower-RefineLevel).

The process of mesh refinement based on the selected refinement criterion (i.e., field) depends on the above-mentioned user-defined parameters (i.e., maxRefinement, nBufferLayers, lowerRefineLevel, upperRefineLevel, maxCells). For each refineInterval, the refinement criteria are examined first. Buffer layers protect the newly refined cells from coarsening and define the number of cells between refinement levels. Subsequently, the cell (un)refinement based on the user-defined criteria is performed. Finally, the fields are mapped onto the newly generated mesh and a flux correction procedure is performed to ensure mass conservation. The process of grid refinement follows the steps depicted in Figure 1. The reader is referred to other publications[8,10,12,16] for a more detailed overview regarding the implementation of the AMR functionality in OpenFOAM.
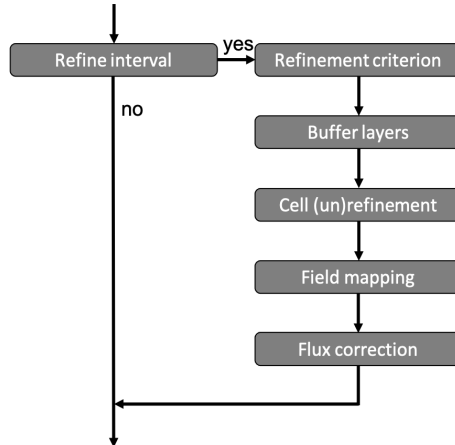

FIGURE 1 Summary of AMR refinement algorithm used in the simulations[8].

3 | EXPERIMENTAL TEST CASE

The experimental case is Sandia's helium plume experiments, reported by O'Hern et al.[15]. This set of experiments is one of the target cases of the MaCFP workshop and a detailed description of the experimental case, setup and boundary conditions can be found on https://github.com/MaCFP so only a brief overview is given here. The experimental facility consists of a 6.1 m cubical enclosure with a 2.4 m in diameter chimney, located on top of the chamber. The plume source is 1 m in diameter (surrounded by a 0.51 m wide floor), where helium was issued from a diffuser at an average velocity of 0.325 m/s. The inlet

mixture contained 96.4% helium (He), 1.7% acetone (C_3H_6O) and 1.9% oxygen (O_2) by volume, with a molecular weight of 5.45 g/mol. The ambient temperature and pressure during the experiments were 285 K and 80900 Pa, respectively. The average mixture Reynolds number was 3200 ± 0.6% while the average Richardson number was 76 ± 6.5%. The conditions at the exit of the diffuser place the flow in the transitional regime (i.e., $2000 < Re_D < 4000$[25]) for internal flows inside pipes. The scenario involves buoyancy-generated turbulence, due to Taylor-Rayleigh instabilities generated at the base of the plume, as previously described by Maragkos et al.[24]. The reported experimental uncertainties on the measured velocities and turbulent statistics are in the order of 20% and 30%, respectively, while for the mean and rms mass fractions are in the order of 23% and 21%, respectively.

4 | NUMERICAL SETUP

A cylindrical computational domain of 4 m x 4 m (diameter x height) is used to model the case. All simulations use a structured mesh with hexahedral cells. Apart from the use of a dynamic mesh (named 'AMR'), three static meshes (named 'Static - Small', 'Static - Medium' and 'Static - Large') are considered, with a local mesh refinement strategy based on the location inside the computational domain. All the different approaches use grid sizes that range from 1.5 cm to 6 cm (the 6 cm size being the base mesh). The 'Static - Small' mesh refines the grid size to 3 cm in a region of 1.5 m x 1.5 m (diameter x height) and to 1.5 cm in a region of 1 m x 1 m (diameter x height) above the fuel source. The 'Static - Medium' mesh refines the grid size to 3 cm in a region of 2 m x 2 m (diameter x height) and to 1.5 cm in a region of 1.5 m x 1.5 m (diameter x height) above the fuel source. The 'Static - Large' mesh refines the grid size to 3 cm in a region of 3 m x 3 m (diameter x height) and to 1.5 cm in a region of 2 m x 2 m (diameter x height) above the fuel source. The total number of cells is 267684, 525235 and 946201 for 'Static - Small', 'Static - Medium' and 'Static - Large' meshes, respectively. Given that the focus of the present study is on the near-field region of the turbulent buoyant helium plume (experimental measurements of radial profiles are only available up to a height of 0.6 m), AMR was only employed up to a height of 1 m above the fuel source. This aspect minimised the total number of cells required for AMR and resulted in an average number of cells with AMR in the order of 230000, comparable to the 'Static - Small' case. It is worth noting that the finest grid resolution of 1.5 cm chosen here, is based on the grid sensitivity studies previously conducted on the same test case by the authors[19,24]. An overview of the local grid refinement strategy considered in the paper, as well as a sample of the mesh used with AMR, is presented in Figure 2.

The corresponding mass flow rate (i.e., 0.047 kg/s based on an inlet velocity of 0.325 m/s and a mixture density of 0.186 kg/m^3) is applied at the fuel inlet with the experimentally reported mixture compositions (i.e., 96.4% He, 1.7% C_3H_6O and 1.9% O_2 by volume). The molecular Schmidt number is set to 0.22 for helium and to 0.7 for the other species. The 0.51 m wide ground plane, surrounding the fuel inlet, is included in the simulations, in order to obtain an as realistic as possible entrainment in the near-field region. The ambient temperature and pressure are set to the experimental conditions. A PIMPLE algorithm considering 3 outer loops is used for the pressure-velocity coupling with a Rhie-Chow interpolation to avoid odd-even decoupling. A mixed boundary condition is assigned for velocity at the open boundaries of the computational domain, setting a zero gradient for any outward flow and calculating the inlet velocity from pressure. A Dirichlet boundary condition is assigned for pressure at the open boundaries of the computational domain, fixing the total pressure (when the velocity changes, the dynamic pressure is adjusted accordingly). At the outlet (i.e., top boundary), a Neumann (i.e., zero gradient) boundary condition is assigned to all variables with no reverse flow allowed in order to avoid numerical instabilities. A no-slip boundary condition is applied for velocity at solid surfaces while a zero gradient boundary condition is used for the chemical species. The equations are advanced in time using a first order implicit Euler scheme. For the convective terms a second order central difference scheme (i.e., Gauss linear) is employed, for scalar transport a TVD scheme (i.e., limitedLinear01 0.5) is used, while for the diffusive terms a central difference scheme is applied. LimitedLinear is a first/second order, linear scheme that limits towards upwind in regions of rapidly changing gradients. The keyword 01 enforces bounding of scalars between 0 and 1, while the parameter 0.5 implies blending of 50% linear and 50% upwind schemes. More details regarding the discretization schemes used in the study can be found in[26]. The simulations are run for 35 s, with a variable time step, corresponding to a maximum Courant number of 0.9. Results are presented as averaged values over the last 30 s. All numerical simulations are run in serial mode which eliminates any possible uncertainties related to load balancing, which could potentially affect the computing times with AMR (see Table 2 later on), enabling a fair comparison between the serial and parallel simulations presented in the paper. Load balancing is a topic worth studying by itself (see e.g.,[8,27]) and is considered beyond the scope of the present paper, where several modelling parameters are investigated within the context of AMR. Inclusion of load balancing would only add to the complexity of the problem here, without significant added value to the conclusions.
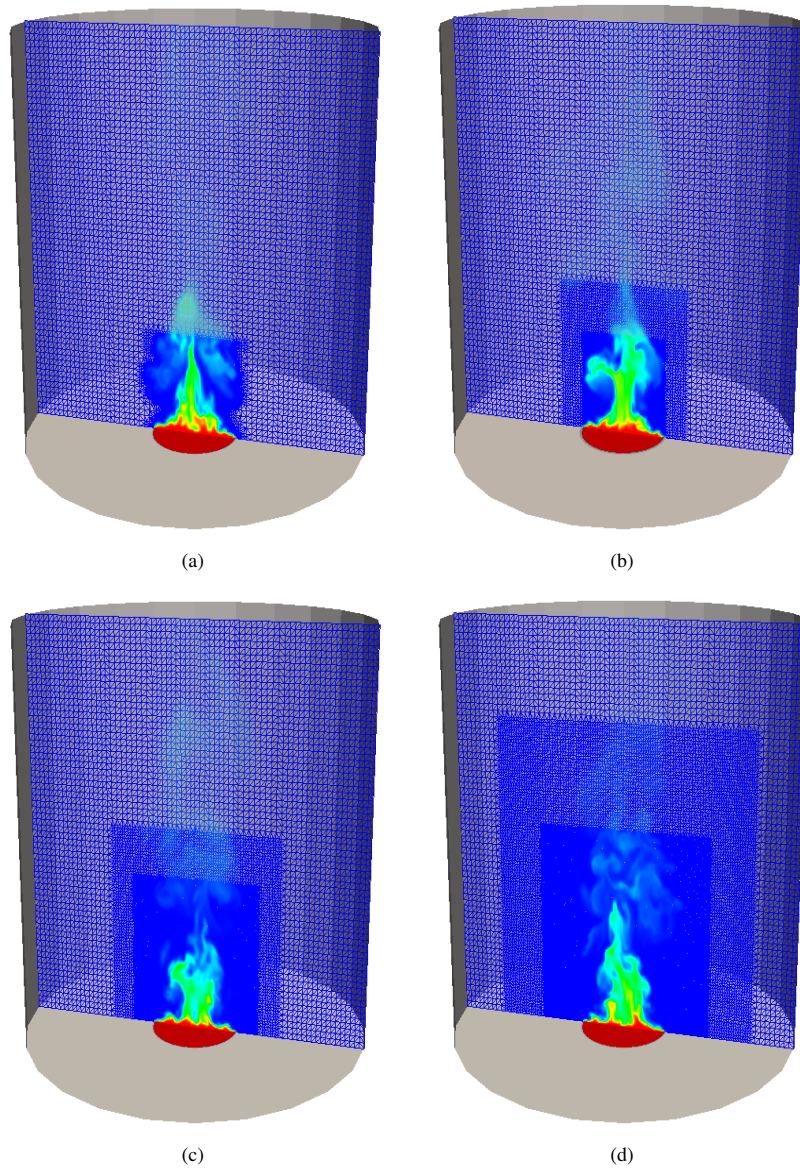
FIGURE 2 Computational domain and local grid refinement (i.e., from 6 cm to 1.5 cm) with (a) AMR, (b) Static - Small, (c) Static - Medium and (d) Static - Large.

5 | RESULTS

Initially, a sensitivity study on several AMR parameters is presented and their appropriate values are established. Subsequently, a comparison between the predictions obtained with the selected AMR settings and the results obtained with three different static meshes is made. More specifically, first and second order statistics at different axial locations are presented, as well as the resulting puffing frequencies and the predicted turbulence model parameters in the near-field region of the turbulent buoyant plume. For completeness of the study, selected validation results of the Fire Dynamics Simulator (FDS) (version 6)[28] on a similar grid size (i.e., 1.5 cm), the most widely used software for fire applications, have also been included in order to assess the accuracy of the simulations.

5.1 | AMR parameter sensitivity

A sensitivity study is presented in this section on: the choice of the refineInterval value (i.e., 1, 10 and 100); the chosen field (i.e., He mass fraction, vorticity magnitude, normalised vorticity magnitude, strain rate magnitude and normalised strain rate magnitude); the lower refinement level value and the maxRefinement value (i.e., 2 and 3) used for grid refinement with AMR. Normalisation of the vorticity and strain rate fields is conducted on the basis of their maximum values during each time step. In general, a refineInterval value of 1 should be the starting point of AMR studies because this choice will guarantee that the grid refinement will cover all regions of interest, in both the near- and far-field of the plumes, and at all times during the simulation. Going to higher refineInterval values will not only depend on the Reynolds number of the problem (i.e., the convective time scales involved), but also potentially on the configuration of the problem, as well as the exact regions of interest that need to be refined. The species mass fraction[31] or vorticity[32] are among the typical choices of fields used for grid refinement in the literature and hence are considered here. Moreover, the different refineInterval values considered, cover a wide range of convective time scales, which can be employed to different applications (e.g., fast moving flows will require small refineInterval values, while slowly moving flows could potentially use higher refineInterval values). An overview of the different AMR settings used in each numerical simulation is presented in Table 1.

TABLE 1 Overview of the numerical simulations using AMR. RI: refineInterval, LRL: lowerRefineLevel, URL: upperRefineLevel, MR: maxRefinement.

Case name	field	LRL	URL	RI	MR
He	Y_{He}	0.001	1.0	1	2
He-RI=10	Y_{He}	0.001	1.0	10	2
He-RI=100	Y_{He}	0.001	1.0	100	2
He-MR=3	Y_{He}	0.001	1.0	1	3
He-LRL=0.01	Y_{He}	0.01	1.0	1	2
Vorticity	Vorticity	1 s^{-1}	max	1	2
normVorticity	Normalised vorticity	0.05	1	1	2
StrainRate	Strain rate	1 s^{-1}	max	1	2
normStrainRate	Normalised strain rate	0.05	1	1	2

A first illustration of the local grid refinement with different fields for AMR, is presented in Figure 3. It is illustrated that, when using the helium mass fraction and the normalised vorticity or strain rate fields (i.e., Figures 3(a), 3(b) and 3(d), respectively), the grid refinement follows the plume structure closely, while a much wider region is refined with the use of either the (un-normalised) vorticity or strain rate fields (i.e., Figures 3(c) and 3(e), respectively). Obviously, the choice of the lowerRefineLevel value is relevant in this case: low values will refine a large area, while high values will refine a smaller area. An a priori choice for the lowerRefineLevel value when using vorticity or strain rate, in terms of absolute value, did not yield a generalised criterion that would refine an area closely following the plume shape at all time steps (in contrast to the use of the helium mass fraction or the normalised vorticity / strain rate). This choice will inevitably be scenario-specific and will depend on the turbulence levels of the (buoyant / fire) plume. On the other hand, a careful choice of the lowerRefineLevel value, based on the normalised value of the vorticity / strain rate fields, can result in a grid refinement region that resembles the one generated with the helium mass fraction. The value of 0.05 chosen here was deemed satisfactory for the scenarios studied in the paper. A sensitivity study, involving more scenarios, is recommended. Yet, the absolute value is not expected to differ significantly from this value: it can be interpreted as being similar to the criterion for determining the flame height based on the 95% value of the integrated heat release rate. The lowerRefineLevel value with the helium mass fraction as refinement field was set to a small number (i.e., 10^{-3}). However, a sensitivity study on this choice has been performed.

Selected results from the numerical simulations using different AMR settings, are presented in Figure 4 at the location closest to the fuel source (i.e., $y = 0.2$ m), as well as on the centreline. Given that the local grid refinement as generated based on the different fields, as reported in Figure 3, was similar, it is not surprising to observe only minimal differences between the different simulations at height $y = 0.2$ m, all agreeing well with the experimental data. Regardless of the refinement field choice, the near-field region of the plume is, almost always, refined to the finest grid size allowed (i.e., 1.5 cm). For this reason, and regardless of
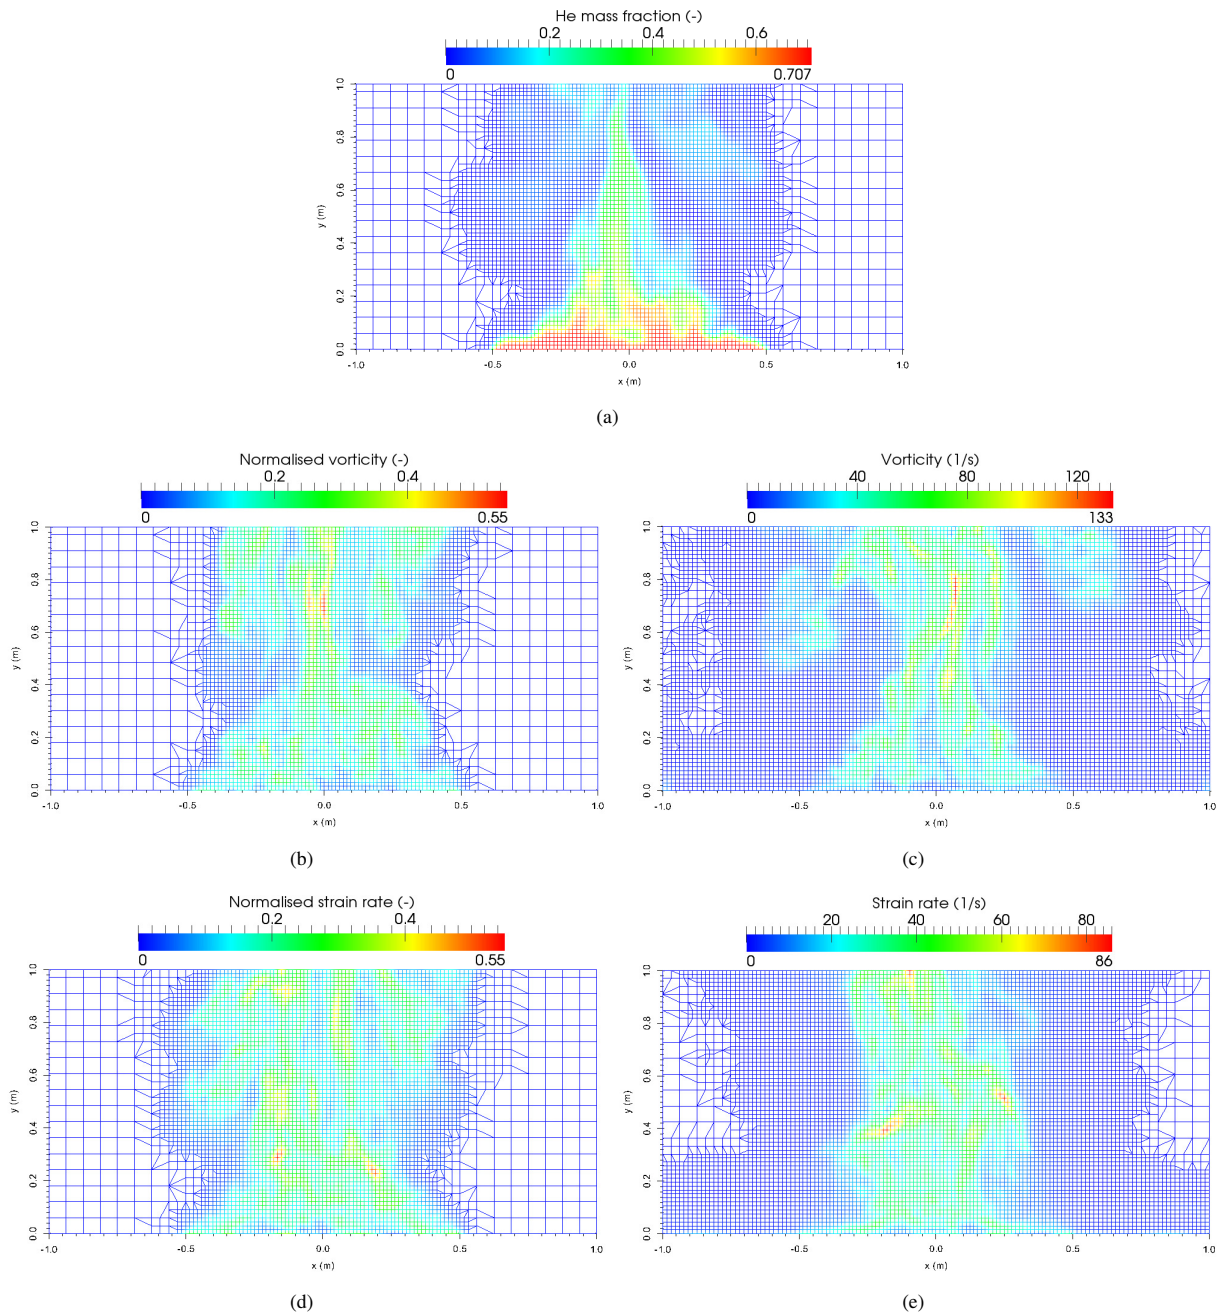
FIGURE 3 Mesh refinement in the near-field region of the plume using as refinement field the (a) He mass fraction, (b) normalised vorticity magnitude, (c) vorticity magnitude, (d) normalised strain rate magnitude and (e) strain rate magnitude.

the advective time scale of the scenario, it is also not surprising that the choice of the refineInterval value does not significantly affect the numerical predictions. Allowing for AMR in the entire computational domain would make the choice of refineInterval more relevant and more important if the numerical predictions further downstream would be of interest. Nevertheless, some differences at axial distances further downstream (i.e., $y > 0.6$ m) are evident in Figure 4(d) which are, most likely, caused by differences in the local grid refinement at radial locations due to the use of different refinement fields, as illustrated in Figure 3. What is interesting to note, however, is that the choice of maxRefinement is important, as shown in Figure 5(a), with refinement level 1 from 10 cm to 5 cm, refinement level 2 from 5 cm to 2.5 cm and refinement level 3 from 2.5 cm to 1.25 cm. In this case, a choice of 3 seems to negatively affect the numerical predictions. More specifically, the use of a very coarse base mesh (i.e., 10 cm) affects the air entrainment into the plume in the near-field region (i.e., lower radial velocities in Figure 4(b)) and makes

the plume more laminar. This results in stronger buoyancy forces, resulting in higher mass fractions (i.e., Figure 4(c)) and axial velocities (i.e., Figure 4(d)) around the centerline, due to lack of turbulent diffusion in the radial direction. It is worth noting that the coarseness of the base mesh (i.e., 10 cm in this case) is more important, in terms of affecting the air entrainment, and not so much the actual number of maxRefinement (i.e., 3 in this case). For example, a maxRefinement of 3 starting from a base mesh of 6 cm would not significantly alter the entrainment characteristics of the helium plume. Within this context, it is difficult to give a general guideline regarding the choice of maxRefinement to be applicable to a wide range of scenarios. Increasing the lowerRefineLevel value (i.e., 0.01) for the helium mass fraction, shown in Figure 5(b), made the refined mesh follow the plume structure more closely (i.e., effectively also using fewer cells) without apparent differences in the numerical predictions, compared to when a lower value (i.e., 0.001) was used. Overall, similar observations were made for the rms quantities at the same locations (the results are omitted to avoid repetition).
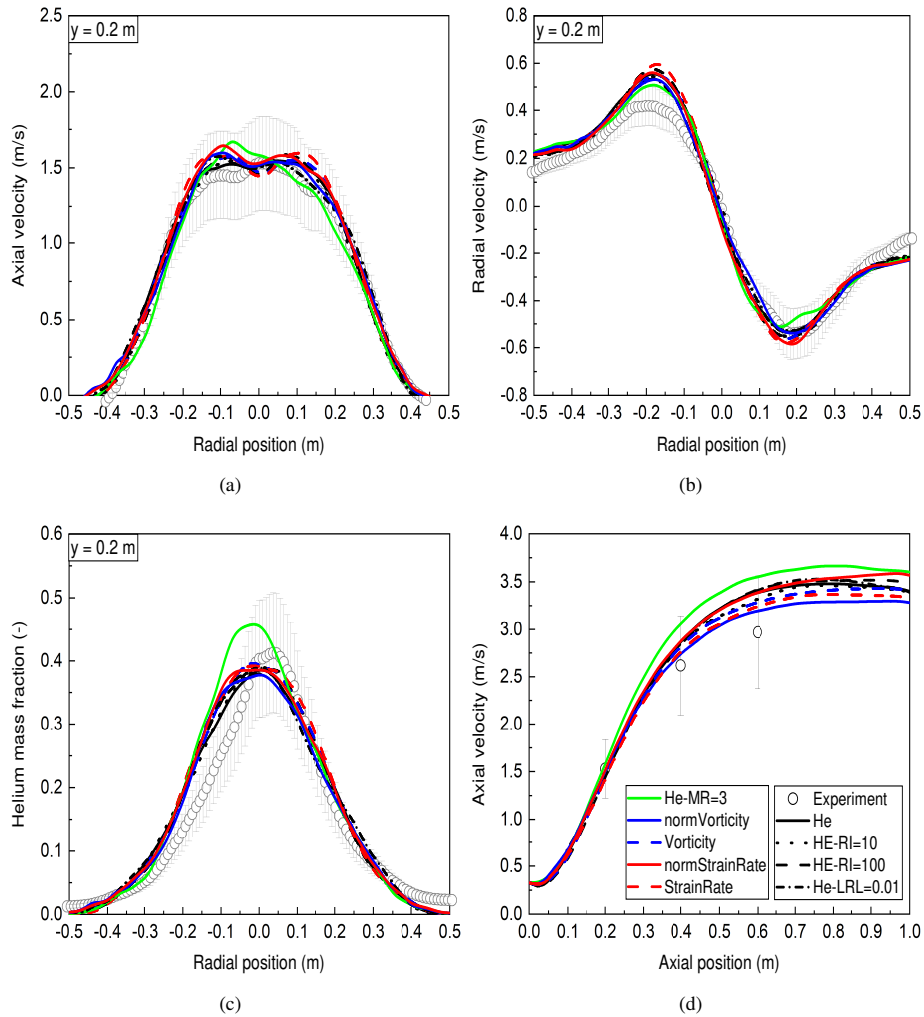

FIGURE 4 Sensitivity study with different AMR settings for the average (a) axial velocity, (b) radial velocity, (c) He mass fraction at height $y = 0.2$ m and (d) axial velocity on the centreline.

Based on the sensitivity study, the selected AMR settings used in the results presented below (named 'AMR') are the following: AMR is performed during every time step, considering the He mass fraction as refinement field with lower and upper threshold limits for grid refinement set to 10^{-3} and 1.0, respectively. The maxRefinement value is set to 2 (i.e., grid refinement from 6 cm to 3 cm (level 1) and from 3 cm to 1.5 cm (level 2)). A nBufferLayers value of 3, to ensure a smooth transition between the different refinement levels, is used while maxCells is set to 400000 cells.
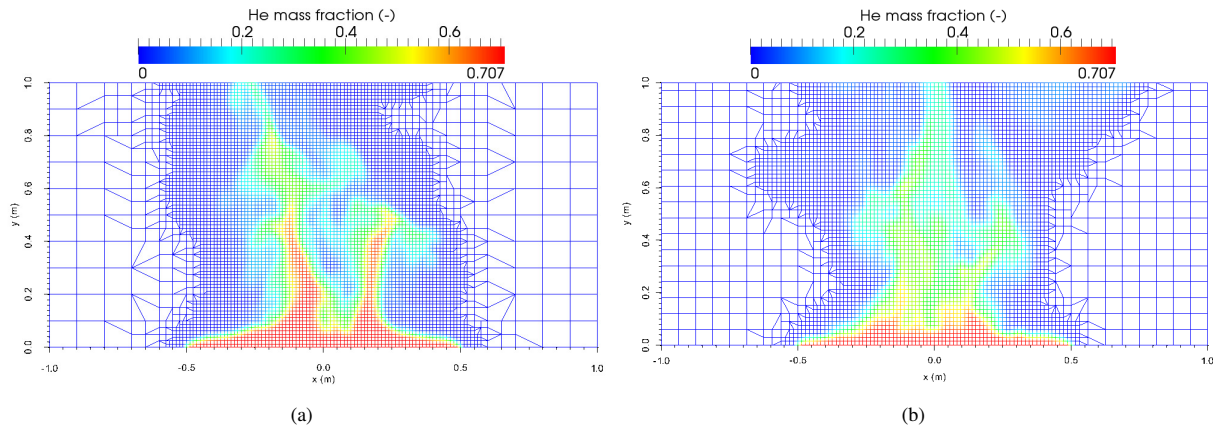
FIGURE 5 Mesh refinement in the near-field region of the plume for case (a) 'He-MR=3' and (b) 'He-LRL=0.01'.

5.2 | First / second order statistics

The first and second order statistics, in terms of axial and radial velocities as well as helium mass fractions, are presented in Figures 6-8 at different heights (i.e., $y = 0.2, 0.4, 0.6$ m) above the fuel source. Overall, the predicted mean and rms velocities for all cases are very close to each other and agree well with the experimental data, remaining within the experimental uncertainty at all locations examined. Compared to the static cases, slightly higher axial velocities are predicted with AMR at most locations, accompanied by slightly narrower profiles (see Figure 6) and lower radial velocities close to the fuel source (see Figure 7). This suggests a small influence on the air entrainment close to the fuel source (i.e., $y = 0.2$ m): AMR seems to make the plume less turbulent, affecting turbulent diffusion in the radial direction, and resulting in stronger buoyancy forces (seen as an increase in the axial velocities). This is further confirmed by the predicted helium mass fractions: the AMR profiles are narrower and the AMR values are approximately 20% higher than with the static meshes at all heights close to the centreline (see Figure 8). No apparent differences are evident in the rms velocities between the different cases. The higher rms helium mass fractions obtained with AMR, compared to the static meshes, stem from the resulting higher mean values as well.

When it comes to the predictions with AMR, the number of cells between each grid refinement level (i.e., nBufferLayers), the minimum value of the variable used for grid refinement (i.e., lowerRefineLevel) as well as the different levels of grid refinement (i.e., maxRefinement) are all parameters that can have a combined effect on the air entrainment in the plume, potentially making it less turbulent. For example, small nBufferLayers values (e.g., 1) might not result in a very smooth transition of the flow from one grid refinement region to another and the flow can experience a deceleration (hence affecting air entrainment) due to the added viscosity encountered in the coarser grid refinement region. This is particularly the case when large maxRefinement values are used (see the sensitivity study presented above). The lowerRefineLevel value can also be important with respect to how wide the grid refinement region will be with AMR (i.e., a small value will refine the grid in a much wider region, compared to a higher value). The chosen values in our work for the above-mentioned parameters (i.e., lowerRefineLevel = 0.001, maxRefinement = 2 and nBufferLayers = 3), enable the AMR to occur in a region that closely followed the plume structure, without excessively refining the grid in the radial direction. Using a lower value for the lowerRefineLevel parameter (e.g., 0.0001) would reduce/eliminate any effect of AMR on the plume air entrainment, but would come at the cost of excessively refining the grid in the radial direction (hence eliminating any gains associated with the use of AMR to some extent).

5.3 | Puffing frequency

The experimental puffing frequency obtained by O'Hern et al.[15] was 1.37 Hz / 1.45 Hz ± 0.1 Hz depending on the averaging procedure (10 test average / 4 test average, respectively). One experimental correlation[29], $f = 0.8 Ri^{0.38} U / D$, leads to a puffing frequency of 1.34 Hz for the case at hand, while another correlation[30], $f = 1.5/\sqrt{D}$, yields a puffing frequency of 1.5 Hz (independent of flow conditions). The predicted puffing frequencies, calculated by performing a Fast Fourier Transformation (FFT) on the time signal of the axial velocity on the centreline at a height of 0.5 m, is presented in Figure 9. The predicted values from the numerical simulations (i.e., 1.46-1.49 Hz) are well within the expected values based on the experimental data and the
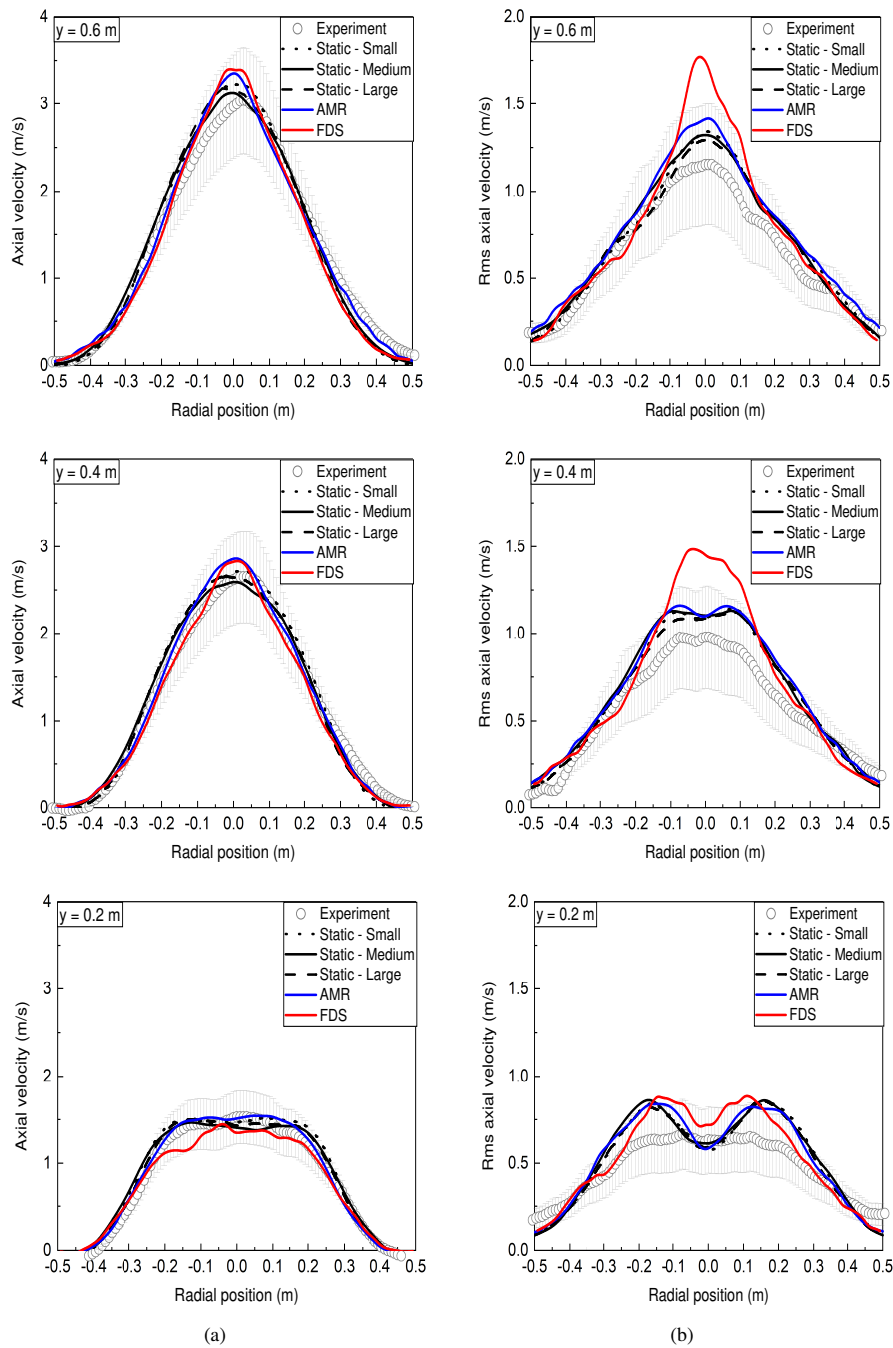
FIGURE 6 Radial profiles of (a) mean and (b) rms axial velocity at different heights.

correlations reported. Additionally, the decay of the energy spectrum has a slope that is very close to -5/3 (characteristic of a turbulent flow) for both the AMR and the static mesh cases. It is interesting to note that the predicted puffing frequency is not affected by the use of AMR, nor by the size of the static refinement region. This finding is not extremely surprising: the position where the time signal of the axial velocity is analysed (i.e., x, y, z = 0 m, 0.5 m, 0 m) is only half a diameter downstream of the fuel inlet, in a region where there will always be grid refinement to 1.5 cm (i.e., finest grid size considered in the simulations) when AMR is used. Nevertheless, it is important to note that this insensitivity is confirmed by the results of the numerical simulations: the use of AMR does not significantly affect the puffing behaviour of the buoyant plume. Finally, it is worth noting
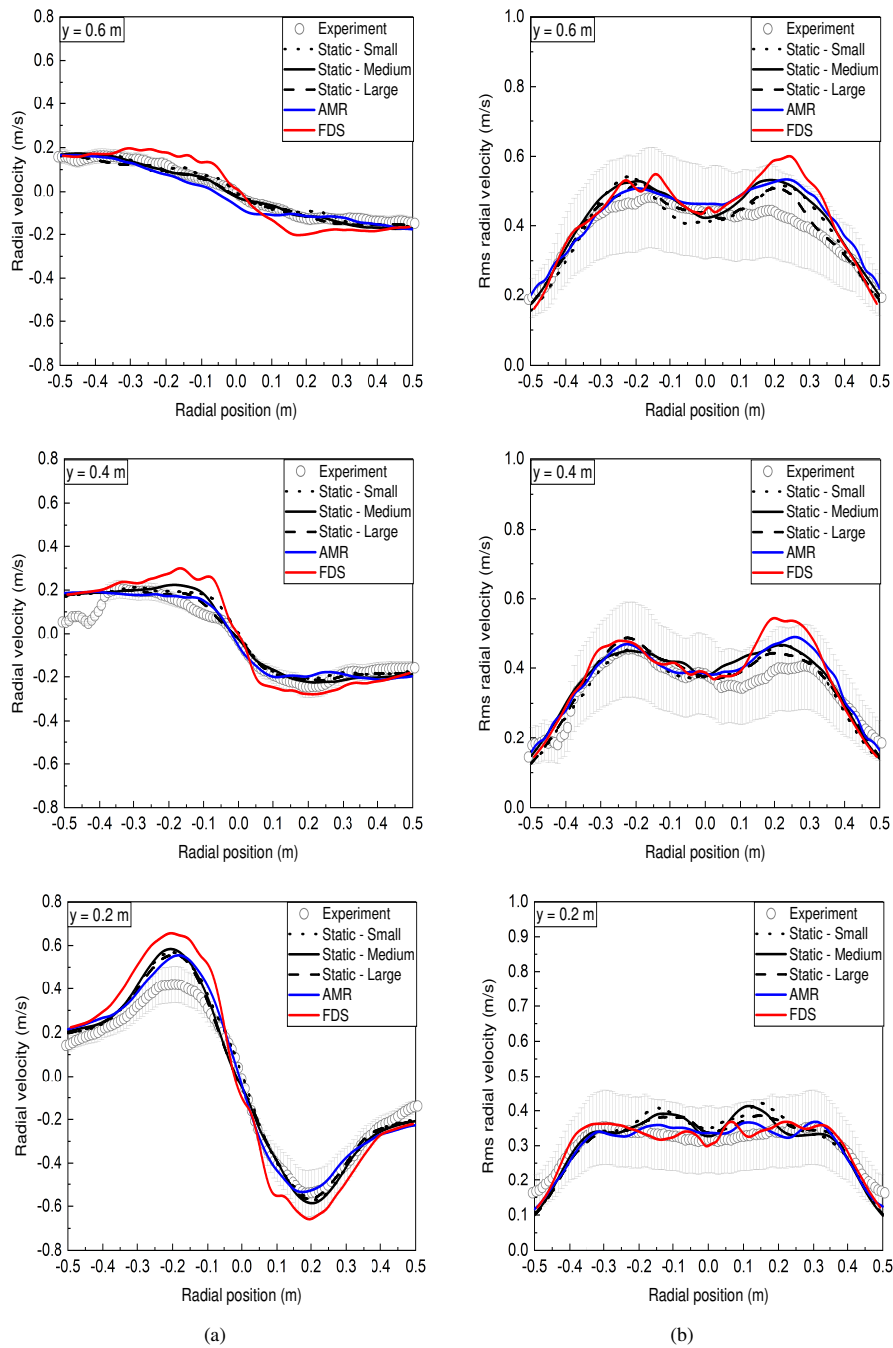
FIGURE 7 Radial profiles of (a) mean and (b) rms radial velocity at different heights.

that a slightly higher amplitude value is observed with the AMR case, compared to the static mesh cases, which is consistent with the higher axial velocities previously seen with AMR in Figure 6.

5.4 | Turbulence model parameters

The spatial variation of the turbulence model parameters (c_s, c_I, Sc_t) on the centreline is presented in Figure 10 for AMR and a static (i.e., 'Static-Small') mesh, up to a height of 1 m above the fuel source. This is the height for the local grid refinement strategy in the static meshes, as well as the AMR cases (where it was decided to refine the grid size to 1.5 cm only up to a
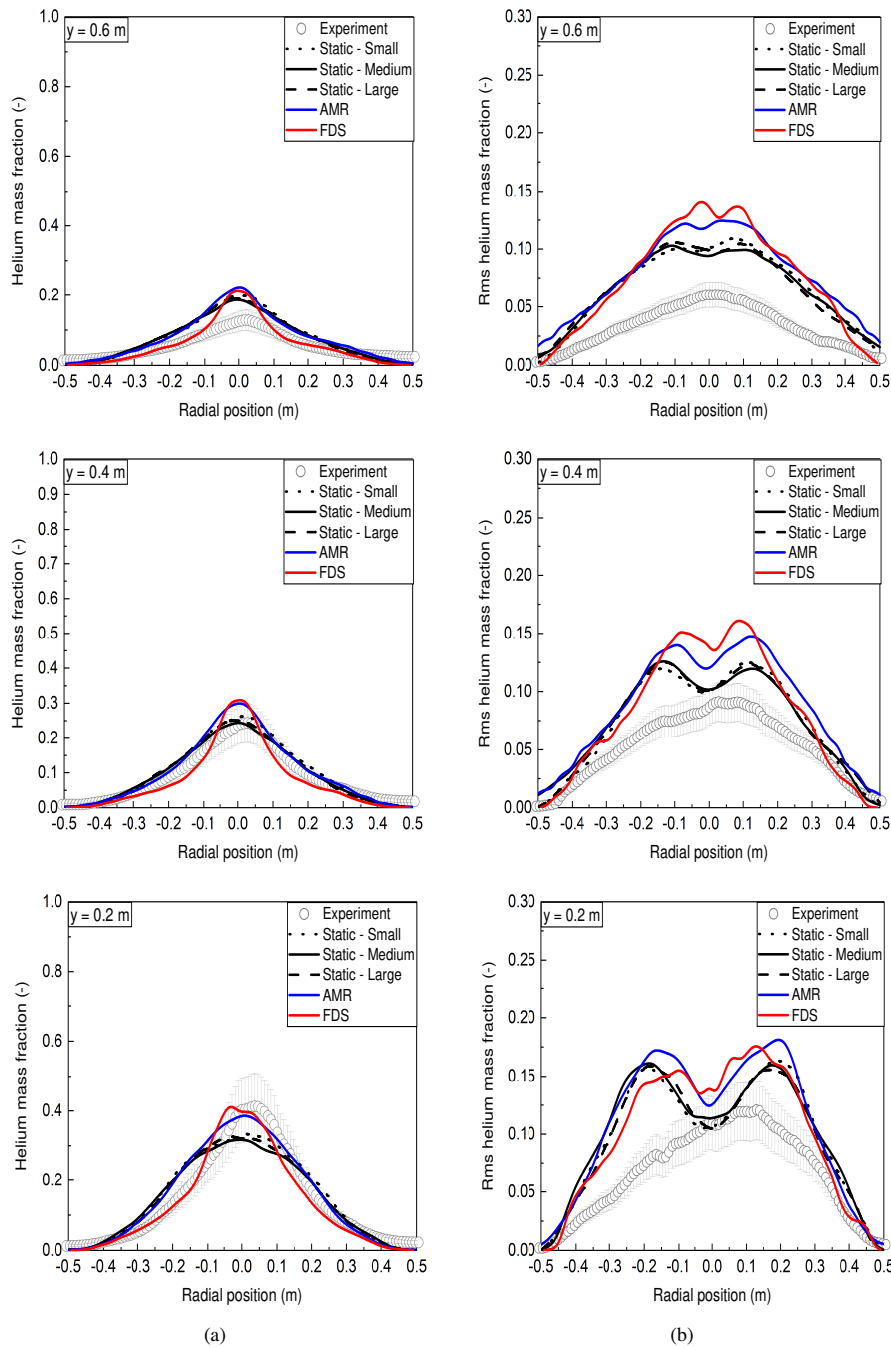
FIGURE 8 Radial profiles of (a) mean and (b) rms helium mass fraction at different heights.

height of $y = 1$ m). Results for the 'Static Small' case are presented because the total number of cells in this case, as well as the refinement regions close to the fuel source, are comparable to the AMR case.

There are noticeable differences in the predicted turbulence model parameters in the near-field region of the buoyant plume ($y < 0.4$ m) which diminish further downstream ($y > 0.4$ m). More specifically, higher c_s and Sc_t values are obtained with AMR directly above the fuel inlet ($y < 0.05$ m), compared to the static case, while the opposite is true further downstream (0.05 m $< D < 0.4$ m). The c_I values are also higher with AMR in this region. These results suggest that the use of AMR slightly affects the air entrainment in the near field-region of the plume. The (approximately 40%) lower c_s values close to the fuel source (0.05 m $< y < 0.4$ m) with AMR are in line with the less turbulent plume structure previously reported, when examining the results
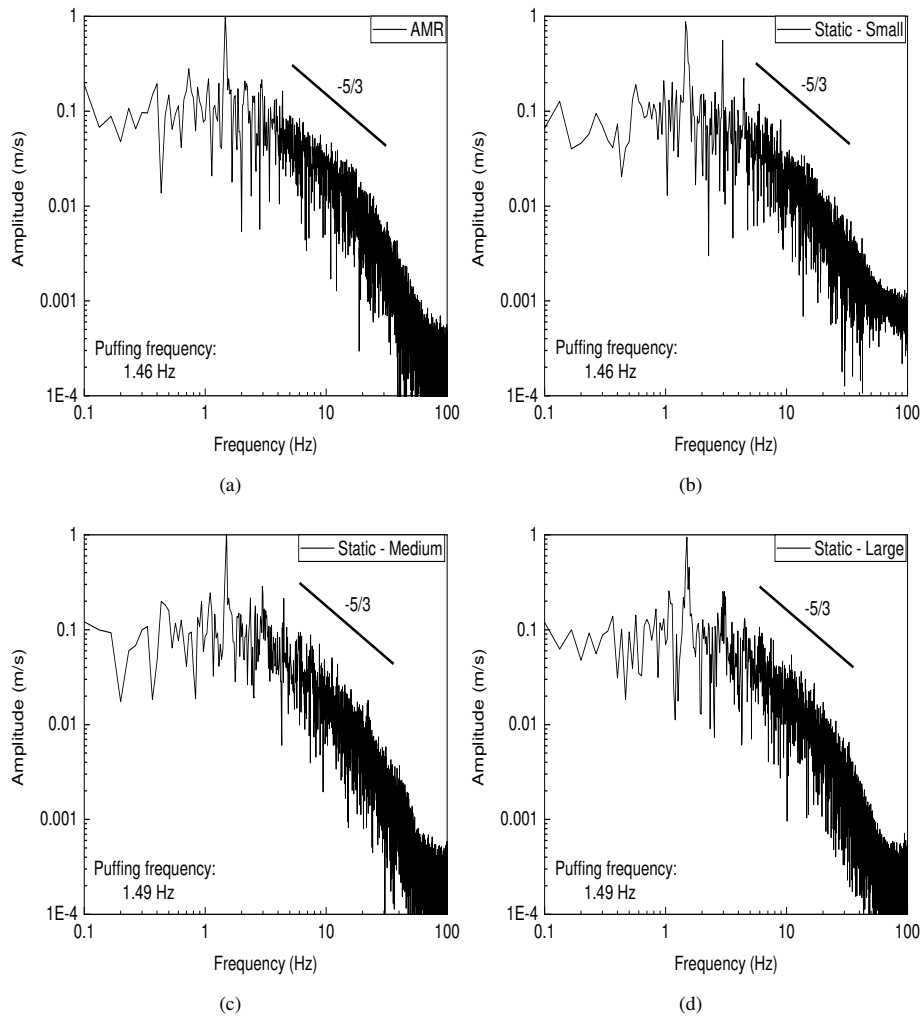
FIGURE 9 Predicted puffing frequency with (a) AMR, (b) Static - Small, (c) Static - Medium and (d) Static - Large.

for the velocities at different heights (i.e., higher axial velocities due to strong buoyancy forces in Figure 6 and lower radial velocities at height $y = 0.2$ m in Figure 7). Significant differences, up to 100%, are also evident for the predicted c_I values at the same axial locations. This result will inevitably affect the predicted sub-grid scale kinetic energies, k_{sgs}, and can potentially affect the turbulent mixing time scale (as typically used in combustion models with infinitely fast chemistry) in this region when reacting flows are considered. Differences in the order of 50% are also evident for the turbulent Schmidt number, Sc_t, but given that most turbulent (buoyant / fire) plume scenarios are dominated by advection, the influence of this parameter is expected to be less significant than that of c_s and c_I, because it will only slightly affect the turbulent (species / thermal) diffusion in the radial direction. For completeness, it also is reported that the choice of either the vorticity or the strain rate as refinement field in the AMR simulations, yielded similar profiles for the predicted turbulence model parameters. The predictions on the 3 different static meshes were also similar (results are not presented to avoid repetition).

5.5 | Computational time

In order to compare the computational cost of the different numerical simulations, a CPU cost parameter is defined as:

$$\text{CPU cost} = \frac{\text{Number of cores } \times \text{Wall clock time}}{\text{Simulation time } \times \text{Total cells}} \tag{10}$$

where the number of cores is 1 (i.e., serial mode) and the simulation time is 35 s for all cases.
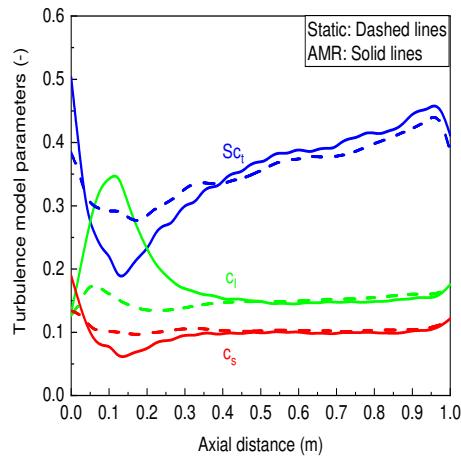
FIGURE 10 Predicted turbulence model parameters on the centreline up to $y = 1$ m.

A selection of the most representative results, using AMR and static meshes, is presented in Table 2. It is observed that when an AMR and a static case use comparable numbers of cells (i.e., AMR vs Static-Small), then the static case is faster by approximately 30% if AMR is employed every time step. Increasing the refineInterval value to 10 or 100 (i.e., decreasing the frequency of applying AMR), can result in comparable (i.e., AMR-RI=10 vs Static-Small) and even slightly shorter (i.e., AMR-RI=100 vs Static-Small) clock times, without significantly affecting the accuracy of the numerical simulations (in the near-field region of the plume, which was the focus of the study). Increasing the size of the regions where static grid refinement is performed (i.e., Static-Medium and Static-Large cases) also increases the CPU cost, for obvious reasons. However, the cost was still lower than the use of AMR (but higher than the cases with RI=10 or 100). Overall, it can be concluded that employing AMR during every time step will result in a relatively high CPU cost for the numerical simulation and will, most definitely, be slower than a static simulation using the same number of cells. On the other hand, an AMR simulation can be potentially faster than a simulation using a static mesh depending on the scenario examined. The observations made here are not strictly limited to the case at hand. Rather, they could be extended to (buoyant / fire) plume simulations in general. Within this context, it is important to note that a careful selection of a static mesh can be equally fast or even faster than the use of AMR. However, AMR has great potential if the region with desired grid refinement changes in time since the static mesh would not be capable of following the motion of the region of interest.

TABLE 2 Comparison of the wall clock times, total number of cells and CPU cost for the different numerical simulations.

Case	Wall clock time (s)	Total cells (-)	CPU cost (-)	CPU cost (%)
AMR	299339	230066	0.0372	145
AMR-RI=10	219898	232765	0.0270	105
AMR-RI=100	206224	236933	0.0248	97
Static-Small	239900	267684	0.0256	100
Static-Medium	547606	525235	0.0298	116
Static-Large	1110541	946201	0.0335	131

6 | CONCLUSIONS

Focusing on large eddy simulations, the study aimed at evaluating evaluating OpenFOAM's adaptive mesh refinement (AMR) capabilities and accuracy in the context of a turbulent buoyant helium plume scenario. Sandia's turbulent helium plume test case is a scenario with similar characteristics, in terms of air entrainment and vortex shedding, with those encountered in large-scale

fires. The simulation results using AMR were compared to simulations using static meshes of different sizes and to available experimental data (i.e., involving first and second order statistics and puffing frequencies). A comparison of the required computational times for the different simulations was also reported. Additionally, a sensitivity study on different AMR-related parameters was conducted and the performance of the dynamic Smagorinsky model, when combined with AMR, was evaluated.

Overall, the AMR predictions were comparable, in terms of accuracy, to the predictions on static meshes of different sizes. However, a slight influence of AMR on the entrainment in the near-field region of the plume was observed, pointing at reduced turbulence: lower radial velocities, resulting in higher axial velocities along the plume centreline due to stronger buoyancy forces, and therefore also higher helium mass fractions. The use of AMR also (significantly) altered some of the, dynamically determined, turbulence model parameters. More specifically, differences of up to 40%, 100% and 50% were evident for the c_s, c_I and Sc_t parameters, respectively. No influence in the predicted puffing frequencies of the plume was observed, all being close to the experimental values (Experiment: 1.37-1.45 Hz, AMR: 1.46 Hz, Static: 1.46-1.49 Hz). The use of a comparable number of cells for both the AMR and the static mesh cases made the latter (significantly) faster, due to an added overhead time needed for re-meshing with AMR. On the other hand, an AMR simulation could be (significantly) faster than a simulation on a static mesh if the size of the static mesh was not carefully selected, or if the region where mesh refinement is required, changes in time.

Overall, the use of AMR showed great potential, both in terms of accuracy and computing times, for the turbulent buoyant helium plume case. Nevertheless, a careful selection of the AMR related parameters (i.e., refinement field, threshold values for refinement and frequency of grid refinement) is necessary to fully benefit from the use of AMR. The influence of AMR in the entrainment characteristics in the near-field region of the plume can potentially lead to differences, compared to static meshes, when reacting flows are considered. Differences in terms of predicted sub-grid scale quantities (e.g., sgs kinetic energy, k_{sgs}) can potentially affect the turbulent mixing time scale, typically used by combustion models with infinitely fast chemistry. These aspects require attention and further investigation in the future if extensive use of AMR in fire related scenarios is to be made.

Author contributions

Conceptualization: Georgios Maragkos; Methodology: Georgios Maragkos; Writing - original draft preparation: Georgios Maragkos; Writing - review and editing: Elena Funk, Bart Merci; Funding acquisition: Bart Merci.

Financial disclosure

Conflict of interest

The authors declare no potential conflict of interests.

References

1. S.R. Tieszen, On the fluid mechanics of fires, Annu. Rev. Fluid Mech. 33 (2001) 67-92. https://doi.org/10.1146/annurev.fluid.33.1.67

2. X. Gao, C.P.T. Groth, Parallel Adaptive Mesh Refinement Scheme for Three-Dimensional Turbulent Non-Premixed Combustion, 46th AIAA Aerospace Sciences Meeting and Exhibit, 7-10 January 2008, Reno, Nevada. https://doi.org/10.2514/6.2008-1017

3. G. Hindi, E.E. Paladino, A.A.M. de Oliviera, Effect of mesh refinement and model parameters on LES simulation of diesel sprays, Int. J. Heat Fluid Flow 71 (2018) 246-259. https://doi.org/10.1016/j.ijheatfluidflow.2018.04.001

4. N. Liu, Z. Wang, M. Sun, R. Deiterding, H. Wang, Simulation of liquid jet primary breakup in a supersonic crossflow under Adaptive Mesh Refinement framework, Aerosp. Sci. Technol. 91 (2019) 456-473. https://doi.org/10.1016/j.ast.2019.05.017

5. J.J. Cooke, L.M. Armstrong, K.H. Luo, S. Gu, Adaptive mesh refinement of gas–liquid flow on an inclined plane, Comput. Chem. Eng. 60 (2014) 297-306. https://doi.org/10.1016/j.compchemeng.2013.09.007

6. E.V. Avdeev, V.A. Fursov, V.A Ovchinnikov, An adaptive mesh refinement in the finite volume method, Proceedings of Information Technology and Nanotechnology (ITNT-2015), CEUR Workshop Proceedings, 234-241, 2015.

7. M. Nuernberg, L. Tao, Three dimensional tidal turbine array simulations using OpenFOAM with dynamic mesh, Ocean Eng. 147 (2018) 629-646. https://doi.org/10.1016/j.oceaneng.2017.10.053

8. D. Rettenmaier, D. Deising, Y. Ouedraogo, E. Gjonaj, H. De Gersem, D. Bothe, C. Tropea, H. Marschall, Load balanced 2D and 3D adaptive mesh refinement in OpenFOAM, SoftwareX 10 (2019) 100317. https://doi.org/10.1016/j.softx.2019.100317

9. Z. Wang, L. Li, H. Cheng, B. Ji, Numerical investigation of unsteady cloud cavitating flow around the Clark-Y hydrofoil with adaptive mesh refinement using OpenFOAM, Ocean Eng. 206 (2020) 107349. https://doi.org/10.1016/j.oceaneng.2020.107349

10. Lm. Li, Dq. Hu, Yc. Liu, Bt. Wang, C. Shi, Jj. Shi, C. Xu, Large eddy simulation of cavitating flows with dynamic adaptive mesh refinement using OpenFOAM, J. Hydrodyn. 32 (2020) 398-409. https://doi.org/10.1007/s42241-019-0041-1

11. T. Dbouk, A computational framework with an adaptive mesh refinement technique for concentrated suspension flows, Particul. Sci. Technol. 38 (2020) 782-791. https://doi.org/10.1080/02726351.2019.1624663

12. C. Lapointe, N.T. Wimer, J.F. Glusman, A.S. Makowiecki, J.W. Daily, G.B. Rieker, P.E. Hamlington, Efficient simulation of turbulent diffusion flames in OpenFOAM using adaptive mesh refinement, Fire Saf. J. 111 (2020) 102934. https://doi.org/10.1016/j.firesaf.2019.102934

13. N.T. Wimer, M.S. Day, C. Lapointe, M.A. Meehan, A.S. Makowiecki, J.F. Glusman, J.W. Daily, G.B. Rieker, P.E. Hamlington, Numerical simulations of buoyancy-driven flows using adaptive mesh refinement: structure and dynamics of a large-scale helium plume, Theor. Comput. Fluid Dyn. 35 (2021) 61-91. 10.1007/s00162-020-00548-6

14. C. Lapointe, N.T. Wimer, J.F. Glusman, S. Simons-Wellin, J.F. Glusman, G.B. Rieker, P.E. Hamlington, Efficient Simulations of Propagating Flames and Fire Suppression Optimization Using Adaptive Mesh Refinement, Fluids 6 (2021) 323.

15. T.J. O'Hern, E.J. Weckman, A.L. Gerhart, S.R. Tieszen, R.W. Schefer, Experimental study of a turbulent buoyant helium plume, J. Fluid Mech. 544 (2005) 143–171. https://doi.org/10.1017/S0022112005006567

16. H. Jasak, A.D. Gosman, Automatic resolution control for the finite-volume method, Part 2: adaptive mesh refining and coarsening, Numer. Heat Tr. B-Fund 38 (2000) 257-271. 257-271. https://doi.org/10.1080/10407790050192762

17. G. Maragkos, T. Beji, B. Merci, Advances in modelling in CFD simulations of turbulent gaseous pool fires, Combust. Flame 181 (2017) 22-38. https://doi.org/10.1016/j.combustflame.2017.03.012

18. G. Maragkos, T. Beji, B. Merci, Towards predictive simulations of gaseous pool fires, Proc. Comb. Inst. 37 (2019) 3927-3934. https://doi.org/10.1016/j.proci.2018.05.162

19. G. Maragkos, B. Merci, On the use of dynamic turbulence modelling in fire applications, Combust. Flame 26 (2020) 9-23. https://doi.org/10.1016/j.combustflame.2020.02.012

20. P. Moin, K. Squires, W.H. Cabot, S. Lee, A dynamic subgrid-scale model for compressible turbulence and scalar transport, Phys. Fluids A 3 (1991) 2746-2757. https://doi.org/10.1063/1.858164

21. T.S. Lund, On the use of discrete filters for large eddy simulation, Annual Research Briefs, Center for Turbulence Research, NASA Ames/Stanford University, 83-95, 1997.

22. C. Fureby, G. Tabor, Mathematical and Physical Constraints on Large-Eddy Simulations, Theoret. Comput. Fluid Dynamics 9 (1997) 85-102. https://doi.org/10.1007/s001620050034

23. M.P. Martin, U. Piomelli, G.V. Candler, Sub-grid-Scale Models for Compressible Large-Eddy Simulations, Theoret. Comput. Fluid Dynamics 13 (2000) 361-376. https://doi.org/10.1007/PL00020896

24. G. Maragkos, P. Rauwoens, Y Wang, B. Merci, Large Eddy simulations of the flow in the near-field region of a turbulent Buoyant Helium plume, Flow Turb. Combust. 90 (2013) 511-543. https://doi.org/10.1007/s10494-012-9437-5

25. J.P. Holman, Heat Transfer, Tenth edition, McGraw-Hill, 2010.

26. G. Maragkos, S. Verma, A. Trouvé, B. Merci, Evaluation of OpenFOAM's discretization schemes used for the convective terms in the context of fire simulations, Comput. Fluids 232 (2022) 105208.

27. W.F. Mitchell, A refinement-tree based partitioning method for dynamic load balancing with adaptively refined grids, Journal of Parallel and Distributed Computing 67 (2007) 417-429. https://doi.org/10.1016/j.jpdc.2006.11.003

28. K. McGrattan, S. Hostikka, J. Floyd, R. McDermott, M. Vanella, Fire Dynamics Simulator Technical Reference Guide Volume 3: Validation, NIST Special Publication 1018-3, Sixth Edition (2021).

29. B.M. Cetegen, K.D. Kasper, Experiments on the oscillatory behaviour of buoyant plumes of helium and helium-air mixtures, Phys. Fluids 8 (1996) 2974–2984. https://doi.org/10.1063/1.869075

30. B.M. Cetegen, T.A. Ahmed, Experiments on the periodic instability of buoyant plumes and pool fires, Combust. Flame 93 (1993) 157–184. https://doi.org/10.1016/0010-2180(93)90090-P

31. M. Emmett, E. Motheau, W. Zhang, M. Minion, J.B. Bell, A fourth-order adaptive mesh refinement algorithm for the multicomponent, reacting compressible Navier–Stokes equations, Combust. Theory Model 23 (2019) 592-625. https://doi.org/10.1080/13647830.2019.1566574

32. D. Hartmann, M. Meinke, W. Schröder, A level-set based adaptive-grid method for premixed combustion, Combust. Flame 158 (2011) 1318-1339. http://dx.doi.org/10.1016/j.combustflame.2010.11.007