


MATHEMATICAL MODELS OF ATMOSPHERIC CONVECTION: A
LAGRANGIAN PERSPECTIVE

Samuel Joseph Wallace

A Thesis Submitted for the Degree of PhD
at the
University of St Andrews



2023

Full metadata for this thesis is available in
St Andrews Research Repository
at:

<http://research-repository.st-andrews.ac.uk/>

Identifiers to use to cite or link to this thesis:

DOI: <https://doi.org/10.17630/sta/485>
<http://hdl.handle.net/10023/27696>


This item is protected by original copyright

This item is licensed under a
Creative Commons License

<https://creativecommons.org/licenses/by/4.0>

Mathematical models of atmospheric convection: a
Lagrangian perspective

Samuel Joseph Wallace



University of
St Andrews

This thesis is submitted in partial fulfilment for the degree of

Doctor of Philosophy (PhD)

at the University of St Andrews

December 2022

General acknowledgements

First and foremost, I would like to thank my primary supervisor, David G. Dritschel for his continued support and patience throughout the past five years. It is no understatement to say that, without your guidance and encouragement, the work in this thesis would not have been possible and I am grateful beyond words.

I would also like to thank my secondary supervisors Dr Richard Scott (University of St Andrews) and Dr Alison Stirling (UK Met Office).

For your discussions and feedback over the years, I want to thank the other members of the Vortex group in St Andrews along with Dr Steven Boeing from the University of Leeds, whose input has been invaluable. I would also like to thank the administrative and support staff of the Mathematical Institute in St Andrews.

To my friends and academic family, who have made the last ten years living in St Andrews so special, I extend my heartfelt gratitude. This journey would not have been the same without all of you. It would be remiss not to also thank the members of the St Andrews Typhoons ice hockey team, both past and present, for being a constant source of comfort and camaraderie since the club's founding.

Finally, I would like to thank my family, Ricky, Jean, Laura, and Richard, who have all been nothing but supportive throughout this journey and had faith in me even when I had little faith in myself. I cannot thank you enough. This work would not have been possible without your support, despite the confused expressions throughout my attempts to explain it. It has not been an easy journey, but I have been incredibly lucky to have my family behind me every step of the way.

Funding

This work was supported by the UK Engineering and Physical Sciences Research Council (EPSRC), Award reference: 1950021.

Candidate's declaration

I, Samuel Joseph Wallace, do hereby certify that this thesis, submitted for the degree of PhD, which is approximately 48,000 words in length, has been written by me, and that it is the record of work carried out by me, or principally by myself in collaboration with others as acknowledged, and that it has not been submitted in any previous application for any degree. I confirm that any appendices included in my thesis contain only material permitted by the 'Assessment of Postgraduate Research Students' policy.

I was admitted as a research student at the University of St Andrews in September 2017.

I received funding from an organisation or institution and have acknowledged the funder(s) in the full text of my thesis.

23 May 2023

Date Signature of candidate

Supervisor's declaration

I hereby certify that the candidate has fulfilled the conditions of the Resolution and Regulations appropriate for the degree of PhD in the University of St Andrews and that the candidate is qualified to submit this thesis in application for that degree. I confirm that any appendices included in the thesis contain only material permitted by the 'Assessment of Postgraduate Research Students' policy.

23 May 2023

Date Signature of supervisor

Permission for publication

In submitting this thesis to the University of St Andrews we understand that we are giving permission for it to be made available for use in accordance with the regulations of the University Library for the time being in force, subject to any copyright vested in the work not being affected thereby. We also understand, unless exempt by an award of an embargo as requested below, that the title and the abstract will be published, and that a copy of the work may be made and supplied to any bona fide library or research worker, that this thesis will be electronically accessible for personal or research use and that the library has the right to migrate this thesis into new electronic forms as required to ensure continued access to the thesis.

I, Samuel Joseph Wallace, have obtained, or am in the process of obtaining, third-party copyright permissions that are required or have requested the appropriate embargo below.

The following is an agreed request by candidate and supervisor regarding the publication of this thesis:

Printed copy

No embargo on print copy.

Electronic copy

No embargo on electronic copy.

23 May 2023

Date Signature of candidate

23 May 2023

Date Signature of supervisor

Underpinning Research Data or Digital Outputs

Candidate's declaration

I, Samuel Joseph Wallace, understand that by declaring that I have original research data or digital outputs, I should make every effort in meeting the University's and research funders' requirements on the deposit and sharing of research data or research digital outputs.

23 May 2023

Date Signature of candidate

Permission for publication of underpinning research data or digital outputs

We understand that for any original research data or digital outputs which are deposited, we are giving permission for them to be made available for use in accordance with the requirements of the University and research funders, for the time being in force.

We also understand that the title and the description will be published, and that the underpinning research data or digital outputs will be electronically accessible for use in accordance with the license specified at the point of deposit, unless exempt by award of an embargo as requested below.

The following is an agreed request by candidate and supervisor regarding the publication of underpinning research data or digital outputs:

No embargo on underpinning research data or digital outputs.

23 May 2023

Date Signature of candidate

23 May 2023

Date Signature of supervisor

Abstract

This thesis presents extensions and enhancements to an existing Lagrangian model for atmospheric convection. The Moist Parcel-In-Cell (MPIC) method, developed by Dritschel et al. (2018), is a novel approach which avoids some shortcomings of conventional large-eddy simulation models, particularly concerning the representation of sub-grid turbulence. While the method is in a relatively early stage, we provide case studies to show the model's potential.

The first case study simulates the ascent of a rising thermal, subject to constant vertical wind shear. Using MPIC, we find that low to intermediate shear appears to promote cloud growth, but high shear tears the thermal apart. The air is partitioned into cloud-air and dry-air components based on the liquid water content and the evolution of the enstrophy associated with each component suggests that the increasing shear is causing more dry air to become turbulent and influence the growth of the cloud.

Our second study analyses the potential energy in MPIC and its sensitivity to numerical parameters. The results show an abnormal growth in the total energy at early times, which is attributed to a failure to enforce incompressibility in the method. The energy evolutions appear to converge with resolution and the numerical mixing parameters in the model for early to intermediate times, although the turbulent flow leads to greater discrepancies in the late-stage energy evolution.

Finally, we present a simulation of the shear-free atmospheric boundary layer, modelled through the implementation of surface fluxes of parcel attributes in MPIC. Our results show evidence of the two-layer entrainment zone structure observed in previous studies. Comparisons of the vertical enstrophy distribution and vorticity field further support this. We also compute entrainment rate parameters that suggest that MPIC underestimates entrainment compared to the zero-order model and overestimates entrainment when using the global buoyancy increment across the entire entrainment zone. Nonetheless, the entrainment rate in

MPIC compares favourably to results in the literature when using a local buoyancy increment computed at the height of minimum buoyancy flux.

Contents

1	Introduction and review of relevant literature	4
1.1	Introduction	4
1.2	Representation of convection in global models	7
1.2.1	Convective parameterisation schemes	7
1.2.2	Basics of convective parameterisation schemes	8
1.2.3	Modern adaptations and challenges	12
1.3	Explicit modelling of convection: Large-eddy simulation	15
1.3.1	Introduction	15
1.3.2	Basic model formulation	16
1.3.3	Representation of sub-grid turbulence	20
1.3.4	Representation of cloud microphysics	24
1.3.5	Radiative and surface processes	25
1.3.6	Applications of LES	26
1.3.7	Limitations of the LES models	28
1.4	A Lagrangian perspective	30
1.4.1	Lagrangian and semi-Lagrangian models	30
1.4.2	The Moist Parcel-In-Cell method	34
1.5	Overview of thesis structure	39
2	Overview of the Moist Parcel-In-Cell method	43
2.1	Governing Equations	43
2.2	Parcel representation	45
2.3	Vorticity inversion	47

2.4	Parcel splitting and mixing	50
2.5	Non-dimensionalisation	51
3	Response of a rising thermal to vertical wind shear in a Lagrangian framework	52
3.1	Introduction	52
3.2	Physical Model	55
3.3	Results	59
3.3.1	Flow evolution	59
3.3.2	Cloud volume and depth	64
3.3.3	Vertical motion	67
3.3.4	Enstrophy	70
3.4	Conclusions	78
4	Energy budget analysis of a rising thermal in the Moist Parcel-In-Cell method	81
4.1	Introduction	81
4.2	Parcel Energetics in MPIC	88
4.2.1	Derivation of Energy Components	88
4.2.2	Physical model	91
4.2.3	Comparison to <i>CAPE</i>	92
4.3	Results	98
4.3.1	Thermal Energetics	99
4.3.2	Resolution Sensitivity and grid vs parcel comparison	101
4.3.3	Sensitivity to parcel mixing	104
4.4	Discussion of Results	110
4.4.1	Incompressibility	113
4.5	Conclusions	115
5	A Lagrangian study of the convective boundary layer growing into a linearly stratified atmosphere	118
5.1	Introduction	118

5.2	Physical model	124
5.2.1	Surface fluxes of heat	124
5.2.2	Volume conservation	125
5.2.3	Prevention of spurious vorticity generation	129
5.2.4	Experimental setup	132
5.3	Results	135
5.3.1	Energy budget and estimation of Reynolds number	135
5.3.2	Evolution of the buoyancy field	140
5.3.3	Evolution of the vorticity field	146
5.3.4	Vertical distribution of thermodynamic properties	149
5.3.5	Vertical distribution of kinetic energy	155
5.3.6	Vertical distribution of enstrophy and related quantities	157
5.3.7	Measures of boundary layer depth	162
5.3.8	Estimates of sublayer thickness in the entrainment zone	167
5.3.9	Entrainment rate parameters	169
5.3.10	Sensitivity to numerical mixing	174
5.4	Conclusions	183
6	Summary and Outlook	187

Chapter 1

Introduction and review of relevant literature

1.1 Introduction

The study of atmospheric convection has long been of substantial importance to our understanding of weather and climate (Holloway et al. 2014; Bony et al. 2015). Convective phenomena can range from shallow cumulus clouds and stratocumulus atop the planetary boundary layer to the global scale circulations of Hadley and Walker cells (Hadley 1735; Walker 1923) and the Madden-Julian oscillation (Madden and Julian 1971). Deep atmospheric convection can give rise to intense thunderstorms, causing significant local disruption. If such thunderstorms organise, mesoscale convective systems may develop (Houze 2004), which have been observed to contribute strongly to the energetics of the monsoon and heavy rainfall events in South Asia (Dube et al. 2014; Singh et al. 2018; Singh et al. 2022). The coupling of convective processes to large-scale circulations can also give rise to extreme weather events. For example, Mekonnen et al. (2006) presented evidence that enhanced convective activity over the Darfur region typically precedes the presence of the African easterly waves. Further work in Mekonnen and Rossow (2011) also identified the interactions between upper level waves and




Figure 1.1: Images showing: (a) fair weather cumulus clouds over Florida, (b) a cumulonimbus cloud with anvil and (c) a mesoscale convective system forming over a landmass in the tropics. Image sources NOAA Library: (a) and (b); NASA Library (c).

convective systems over the Ethiopian Highlands as a source of the lower level waves. These lower level waves were found to, under favourable conditions, serve as the ‘seedlings’ for tropical cyclones which can cause considerable damage and displacement in the United States and elsewhere.

These phenomena all serve to illustrate that the study of atmospheric convection spans a vast range of spatial scales: from the order of meters at cloud edges and in high liquid water cores (Austin et al. 1985; Grabowski and Clark 1993; Blyth et al. 2005; Blyth et al. 2015; Seifert et al. 2015) to the global scales of tens, or even hundreds, of kilometres. As such, the representation of convection in global climate models has presented an immense challenge to the atmospheric community (Randall et al. 2003; Bony et al. 2015; Rio et al. 2019; Lin et al. 2022c), effectively forming a bottleneck in the development of more sophisticated models for climate and weather prediction.

This problem is compounded further by the complexity of the microphysics at play within convective clouds. The phase changes of water contribute significantly to the energy budget within moist convection and can enhance or hinder the associated vertical motions. Attempts to model cloud water and other hydrometeor species (rain, graupel and hail, for example) must contend with the numerous states of water within clouds and consider the distribution of droplet sizes and collisions for precipitation (Kessler 1969, Pruppacher et al. 1998, Swann 1998).

Other sources of feedback to be considered are the effects of radiative heating (e.g. Slingo 1989) and the impact of aerosol species and their transport (Gidel 1983). Adding further layers to the issue are the interactions with the boundary layer and ocean or land surface couplings (Wu et al. 1998; Patton et al. 2005; Sullivan et al. 2014). In short, the modelling of convection is an increasingly complex problem and representing it explicitly on a global scale is beyond the present limitations of computational power.

Many of these feedback processes are interlinked and make clouds one of the most sensitive factors in climate simulations (Sanderson et al. 2010; Sherwood et al. 2014). For example, high concentrations of cloud ice will lead to a greater reflection of solar radiation at the cloud top, resulting in a cooling effect within the Earth's atmosphere. On the other hand, as clouds grow they produce a blanketing effect that may trap heat, resulting in a warming effect on the atmosphere. The concentrations of chemical species will also impact the cooling and warming effects of clouds. Similarly, precipitation cools the lower atmosphere, and the effects of entrainment and detrainment can result in upwards or downwards transport of heat and moisture. To further stress this, a study by Klein et al. (2018) found that low-cloud feedback is currently one of the largest sources of uncertainty in climate models. With regard to weather forecasting, the scales involved in cloud processes are often orders of magnitude smaller than the scales of regional models typically used. Consequently, these models often rely on parameterisation of the unresolved cloud processes, which themselves introduce considerable uncertainty to the problem, for example in the prediction of precipitation levels (Slingo and Palmer 2011).

Another prominent feature of atmospheric convection that is difficult to represent on the global scale is turbulence. Turbulence is what drives the mixing of clouds and the surrounding environmental air via entrainment and has significant implications for the life cycle of clouds (Blyth 1993). Turbulent mixing occurs on scales far smaller than the typical grid resolution of global models. The process of mixing can lead to shells of partially mixed air forming around rising plumes of

cloud air (Bretherton and Smolarkiewicz 1989; Hannah 2017) and impacts entrainment rate in the lower troposphere (Becker et al. 2018) due to the dilution of the air entrained into rising thermals. The turbulent motions of air within the cloud also influence the microphysical processes that dictate transitions between different species of hydrometeor, as well as the precipitation intensity. For example, the frequency of collisions between suspended water droplets increases as turbulent motions become more intense, causing them to grow more readily (Grover and Pruppacher 1985; Falkovich et al. 2002; Pinsky and Khain 2002). Similarly, in studies of the atmospheric boundary layer, turbulence drives the entrainment of dry air from above, feeding the growth of the layer. The cascading nature of turbulence means that even simulations that explicitly resolve convective motions have difficulty resolving the small scales required. These simulations must still account for the influence of turbulence, however, and typically rely on sub-grid scale models to represent unresolved motions.

The coupling between the small and large scale processes is a recurring theme in the study of atmospheric convection and presents a fundamental challenge of the field. Because the processes involved are interconnected across scales, both global scale representations that simplify many key processes and techniques designed to explicitly model convection at smaller scales are equally important in expanding our understanding of convection and its impact on weather and climate. As such, the advancement of our representations of convection is crucial to improving both weather forecasting capabilities (Slingo and Palmer 2011) and our ability to predict the response of the Earth’s climate (Bony et al. 2015) to factors such as rising sea surface temperatures and increases in aerosol concentration.

1.2 Representation of convection in global models

1.2.1 Convective parameterisation schemes

Due to computational limitations, as previously noted, most global circulation models do not explicitly model convection. Instead, the most common approach

is to use a convective parameterisation scheme. The problem of convective parameterisation is best summarised in Arakawa (2004) as “the problem of formulating the statistical effects of moist convection to obtain a closed system for predicting weather and climate”. Any system of large-scale equations for heat and moisture will always have undetermined parameters associated with the effects of unresolved convection. A convective parameterisation scheme aims to make assumptions about the statistical properties of convective motions on the sub-grid scale that can close the system.

Thus, these schemes are very much defined by the assumptions made about convective motions to close the set of equations, the primary closure assumption. The earliest convective parameterisation schemes were developed by Arakawa and Schubert (1974), Kuo (1974) and Manabe et al. (1965), who used closure schemes based on mass-flux quasi-equilibrium, moisture convergence and energetic or state based quasi-equilibrium respectively. The modern parameterisation schemes used today in weather forecasting primarily originate from the successors to these schemes developed in the 80s and 90s (Tiedtke 1993; Kain and Fritsch 1990; Grell 1993).

Because global models cannot resolve the effects of individual clouds, these schemes instead attempt to simulate the collective influence of an ensemble of clouds coupled to the large-scale circulations (Arakawa and Schubert 1974; Arakawa 2004). The two main approaches to this problem are bulk methods, which model the ensemble as a single entraining-detraining plume with the properties representative of the ensemble as in Yanai et al. (1973), and spectral methods, which characterises plumes in the ensemble based on properties such as updraft velocity or cloud top (Arakawa and Schubert 1974; Tiedtke 1993).

1.2.2 Basics of convective parameterisation schemes

A common approach in the parameterisation of convection is to assume that the mass-flux of a convective ‘plume’ (usually referring to a cloud or convective updraft) dominates the upwards transport in the system, as in Arakawa and Schubert 1974. The basic principle of the mass flux approach is that the i^{th} plume covers




Figure 1.2: A schematic from Arakawa and Schubert (1974) showing an ensemble of clouds penetrating through a layer unit surface area at some point below the cloud top. the i^{th} cloud occupies an area σ_i and has associated mass flux M_i . Clouds that lose buoyancy detrain air into the environment. Reproduced from Figure 1 in ‘Interaction of a Cumulus Cloud Ensemble with the Large-Scale Environment, Part I’ by Arakawa and Schubert (1974) with permission. **Published 1974 by the American Meteorological Society**

a horizontal area $\sigma_i(z, t)$ of the grid cell (see Figure 1.2). As such, the mass flux, M_i through the area σ_i can be represented by the equation

$$M_i = \int_{\sigma_i} \rho w d\sigma = \rho \sigma_i w_i, \quad (1.1)$$

where ρ is the density of the air and w_i is the vertical velocity of air in the i^{th} plume. Figure 1.3 shows a schematic of the mass continuity for a single plume. For a property χ , we can determine the mass-flux representation of its vertical flux from the expression,

$$\langle \rho' \chi' w' \rangle = \sum_i M_i (\chi_i - \tilde{\chi}) \quad (1.2)$$

where angled brackets denote horizontal averaging, a prime denotes the deviation from horizontal plane averages and the tilde denotes the value in the non-cloudy environment (Plant 2010). These vertical fluxes of properties are used to quantify the effects of plumes on their environment under the mass flux approximation.




Figure 1.3: Schematic showing the mass continuity for a thin layer of thickness Δz in the i^{th} cloud. M_i is the mass flux of the rising plume and $(\partial M_i / \partial z) \Delta z$ is the mass flux into the layer from the environment. Reproduced from Figure 2 in ‘Interaction of a Cumulus Cloud Ensemble with the Large-Scale Environment, Part I’ by Arakawa and Schubert (1974) with permission. **Published 1974 by the American Meteorological Society.**

This approximation assumes that plumes are rising much faster than the vertical motions of ambient air, i.e., $w \ll w_i$, and that the fractional area covered by a single plume is much smaller than the grid cell, i.e. $\sigma_i \ll 1$.

To develop a convective parameterisation scheme using the mass-flux approach, the plume budget equations for momentum, heat, moisture and liquid water are temporally averaged over the lifetime of the plume. The budget equations are then integrated from the cloud base, to the top of the detrainment layer at the top of the plume where air is detrained into the surrounding environment (Yanai et al. 1973; Arakawa and Schubert 1974; Plant 2010). It is here that the spectral and bulk methods diverge from one another, as the bulk approach of Yanai et al. (1973) made use of a mass-flux weighted average to produce the properties of the single bulk plume:

$$\langle \rho' \chi' w' \rangle = \sum_i M (\chi_B - \tilde{\chi}) \quad (1.3)$$

where

$$M = \sum_i M_i \quad (1.4)$$

and the bulk value of the chosen parameter, χ is given by

$$\chi_B = \frac{\sum_i M_i \chi_i}{M}. \quad (1.5)$$

Quantities denoted with a prime in these equations represent a deviation from the horizontal averages and quantities denoted with a tilde represent the values in a cloud-free environment. For the plume equations to take the same form in terms of bulk variables, some assumptions must be made. The first is in regard to the detrainment rate. When computing the bulk detrainment rate of χ , it is found that it does not add up to the sum of the plume detrainment rates, D_i , across the ensemble and is, in fact, different for each variable. This is what prompted the assumption used in Arakawa and Schubert (1974) that detrainment only occurs at the top of the plume, in the detrainment layer. The assumption is that D_i is assumed to be zero elsewhere and the entrainment rate of the plume, E_i , is assumed to be zero in the detrainment layer. Arakawa and Schubert (1974) provided two sets of equations for each plume’s entraining and detraining regions. In the bulk approach of Yanai et al. (1973), however, they relied on the assumption that the liquid water detrained from each plume is equal to the bulk value, i.e. $l_i = l_B$. Yanai et al. (1973) go on to describe this as “a gross assumption but needed to close the set of equations”. It is also assumed that the plumes within the ensemble are non-interacting with one another.

The closure assumption in Arakawa and Schubert (1974) computes a cloud work function, $A(\lambda)$, for each plume type (sorted by some parameter λ) that closes the system of equations, decomposed into cloud and large-scale contributions. The closure comes from the quasi-equilibrium assumption that the work function’s net rate of change should be approximately zero, $dA/dt \approx 0$. The bulk closure assumption of Yanai et al. (1973) is based on similar principles in terms of defining a work function, but differs from the quasi-equilibrium requirement of Arakawa and Schubert (1974). Instead, closure requires that the timescale associated with the evolution of the large-scale contribution greatly exceeds that of the cloud contribution. Alternative closure assumptions may use the convective available

potential energy (*CAPE*) to close the system of equations (Plant 2010).

While more computationally intensive, the spectral approach tends to be favoured in modern global models of the atmosphere. Because the distribution of entrainment rates, E_i , in an ensemble can be studied through radar observations of individual clouds, this can be used to provide a more realistic representation of plume properties within the parameterisation scheme. Both parameterisations are typically evaluated against observational data and results from convection resolving models at smaller scales (i.e. Siebesma and Cuijpers 1995; Angevine et al. 2018; Zhang et al. 2021). That said, comparisons to large-eddy simulations suggest that bulk representations of convection tend to underestimate significantly the entrainment rates associated with convective clouds (Roms 2010). On the other hand, some studies suggest this may help to account for the dilution of air around plumes (Hannah 2017). Knight et al. 2007 also found that the entrainment parameter in convection schemes is one of the most sensitive aspects of global circulation models.

1.2.3 Modern adaptations and challenges

Modern convective schemes have faced a range of challenges over the years and are the result of constant adaptation to observational data. A common refinement made in more modern schemes is the implementation of stochastic processes to enhance the triggering and development of convection. Examples include stochastic perturbations of the parameters (Grell and Freitas 2014) or stochastic trigger conditions (D’Andrea et al. 2014) and even stochastic distributions of cloud types (Goswami et al. 2017). Stochastic processes have led to improvements in the intraseasonal variability within climate models.

Another relatively new advance in these schemes is the introduction of scale-awareness, such as in Han et al. (2017). In this scheme, the mass flux associated with clouds decreases as the resolution is increased, resulting in an improved prediction of summer precipitation in the U.S. when compared to observational data. These schemes are particularly useful as resolutions approach the ‘gray-zone’ in the region of 5 km to 25 km and models are required to differentiate between resolved

and parameterised convection (Kwon and Hong 2017).

Despite the success of modern techniques, there are still numerous difficulties facing convective parameterisation schemes. For example, a common problem faced in global circulation models is that convection within the model is too readily triggered, leading to precipitation biases, as in Huang et al. (2018), although this was similarly observed in high resolution simulations too. Trigger conditions are a vital part of these schemes, and are based upon a range of factors. Possible trigger conditions include the moisture convergence within a grid cell or the updraft velocities of the strongest thermals and whether they can overcome the convective inhibition of the environment. A consequence of models too readily enabling convection is a persistent low level of precipitation, very different to the heavy rain events observed in some regions. This bias can influence the large-scale circulation features and even disrupt the formation of stratocumulus in some regions. Other tests of scale-aware schemes over Taiwan have found a warm bias (Lin et al. 2022a). Furthermore, comparisons with course-grid convection resolving simulations have also found disagreement in the mass flux profiles of some parameterisation schemes (Zhang et al. 2021).

Another weakness of most convective parameterisation schemes is that they were primarily developed to capture the effects of deep atmospheric convection, with lengthscales extending beyond 5 km and upwards towards the tropopause. As such, the representation of shallow convection associated with boundary layer circulations and stratocumulus clouds is difficult to accommodate in models using traditional schemes (Arakawa 2004) and remains a major source of uncertainty in climate models (Sanderson et al. 2010; Sherwood et al. 2014). There has been a push to generate unified schemes that include both shallow and deep convection (Hohenegger and Bretherton 2011). However, some argue that due to the difference in driving forces between shallow and deep convection (i.e. deep organised convection is heavily influenced by wind shear, while shallow convection typically arises from surface fluxes) that different mass flux schemes are required for each (Rio et al. 2013).

Furthermore, in studies of shallow convection, the atmospheric boundary layer is topped by an inversion layer called the entrainment zone. The entrainment zone is a transition region between the turbulent boundary layer and the free atmosphere with structural features that the grid scales of global climate models cannot resolve (Fedorovich et al. 2004; Garcia and Mellado 2014). The entrainment rate into the boundary layer can be estimated in parameterisation schemes and use representations such as the zero-order model of Zilitinkevich (1991), based on the buoyancy jump across the entrainment zone. Higher resolution studies into the vertical structure, however, have suggested that such representations may not be appropriate (Sorbjan 1999; Garcia and Mellado 2014).

These difficulties have led to some researchers moving away from the traditional parametrisations of convection and towards alternative methods that depend more heavily on explicit convection modelling via large-eddy simulation. For example, Sansom (2021) demonstrates the use of machine learning techniques in improving the ability of models to simulate shallow stratocumulus convection. In particular, they used the Met Office NERC model (Brown et al. 2015) to train a neural network to improve how convection is represented at global scales using statistical emulation.

Other researchers have moved towards the growing accessibility of embedded convection resolving models or even superparameterisation frameworks (Randall et al. 2003; Grabowski 2016). These approaches are especially promising for future studies of the atmosphere. Superparameterisations also represent a natural progression from many of the scale-aware schemes currently in use. The computational requirements to run such simulations on climate timescales are still beyond current capabilities. Nonetheless, such models are likely to see increasing use within the weather forecasting community in the coming years.

Superparametrisations work by coupling a large-scale model to an explicit model at convective resolutions. Within each grid cell, convection is simulated, independent of the surrounding grid cells, but subject to forcings from the large-scale. Typically, this is accomplished with two-dimensional convection models, but recent

advances are making the possibility of fully three-dimensional embedded models increasingly achievable. The bulk properties of the convection simulation in each cell along with vertical distributions of prognostic variables and fluxes etc. are shared back into the main model and their effects on the large-scale flow computed. An example is given in Jansson et al. (2019) where they demonstrated the application of a regional-scale superparameterisation model using a fully three-dimensional convection model within each grid cell. While there remained some issues with load imbalance requiring a modular approach to the problem, the results demonstrated the growing potential of the approach. It is likely that increasing demand will be placed on computationally efficient versions of existing convection resolving models and large-eddy simulations or equivalent techniques that can be embedded within a superparameterisation model.

That is not to say that conventional convective parameterisation schemes are likely to go away anytime soon, far from it. As previously noted, the timescales involved in climate simulations (typically over 100 years) mean it may be decades before even a superparameterised model can be applied, let alone a global convection resolving simulation. As such, we believe that the development of new techniques should continue to supplement the existing representations of convection on the global scale. Simultaneously, more efficient or creative ways to incorporate explicit convection in global circulation models should also be considered as part of the push towards global convection resolving models.

1.3 Explicit modelling of convection: Large-eddy simulation

1.3.1 Introduction

First explored in the 1970s by Lilly (1966) and Deardorff (1970), the large-eddy simulation (LES) technique has become a workhorse of the atmospheric community. The distinction between LES and cloud resolving models (CRM) is primarily one of scales, with CRMs operating at larger scales and thus requiring a different sub-grid formulation. The method was employed to precisely resolve the convec-

tive motions of clouds on much smaller scales than in global circulation models, with the first simulations of shallow clouds performed by Sommeria (1976). Later work by Deardorff (1980) produced early simulations of the stratocumulus topped boundary layer, and since then the technique has developed rapidly, with increasingly sophisticated models being developed (Brown et al. 2015; van Heerwaarden et al. 2017). LES models have granted us a greater insight into the small-scale mixing processes occurring within clouds (Grabowski and Clark 1993; Heus et al. 2008) and see frequent use as a feedback method for the refinement of convective parameterisation schemes described in Section 1.2 (i.e. Siebesma and Cuijpers 1995). Advances in computational power (Khairoutdinov et al. 2009) have also expanded the sizes of domains that can be feasibly modelled in an LES, bringing us even closer to resolving convection on global scales. At the very least, we are seeing notable advances towards superparameterisation and embedded schemes, such as those discussed in Randall et al. (2003) or Grabowski (2016).

1.3.2 Basic model formulation

Large-eddy simulations are non-hydrostatic models that numerically resolve the Navier-Stokes equations on an Eulerian grid. LES models occupy an intermediate space between methods that attempt to fully resolve turbulent motions via direct numerical simulation (DNS) and those that rely on averaging of the Navier-Stokes equations while representing turbulence exclusively through a turbulence model (i.e. the Reynolds averaged Navier-Stokes equations, RANS). As such, they are often viewed as an appealing alternative to DNS due to the reduced computational expense, while offering a greater resolution of turbulent eddies than RANS models. The crux of the LES method is the application of a spatial filter to the Navier-Stokes equations, removing the smallest scale motions from the model (Sagaut and Lee 2002). A turbulence model is then employed to represent the turbulent motions below the filter scale and the resulting stresses before computing their effects on the resolved flow. This filtering means that LES models still resolve the largest eddies in a turbulent flow, unlike RANS models, which represent all turbulent eddies in the flow via turbulence modelling. As such, the method is

highly advantageous in cases where the averaged equations are not appropriate, while being less computationally expensive than DNS models. The advantages of LES models make them a natural choice for highly turbulent phenomena in the atmosphere such as clouds and boundary layer interfaces where DNS may be too expensive.

The filtering in an LES model is represented, in physical space, as a convolution product. The resolved part of a variable, ϕ , is defined as

$$\bar{\phi}(\mathbf{x}, t) = \frac{1}{\Delta} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} G\left(\frac{\mathbf{x} - \mathbf{r}}{\Delta}, t - \tau\right) \phi(\mathbf{r}, t) d\tau d\mathbf{r}, \quad (1.6)$$

where G is the convolution kernel which determines the shape of the filter, such as a Gaussian filter, or a sharp cutoff and Δ is a spatial representation of the filter width. In this equation and those that follow in this section, any variable $\phi(\mathbf{x}, t)$ is expressed as a decomposition between their resolved values after application of the spectral filter, $\bar{\phi}$, and the sub-grid deviations from these values, ϕ' . As such, any property ϕ can be written as $\phi = \bar{\phi} + \phi'$. It follows that the terms given by $\overline{u'_j \phi'}$, therefore, can be interpreted as the turbulent flux of ϕ . Terms denoted by a hat, such as $\hat{\phi}$, correspond to the space-time variable defined in Fourier (spectral) space. The kernel of the filter is typically associated with a spatial width, Δ which is determined by the LES model. The most common approach is to set Δ to be equivalent to the grid spacing of the model and choose the kernel, G a sharp cutoff for the kernel in physical space, such that

$$G(\mathbf{x} - \mathbf{r}) = \frac{\sin(k_c(\mathbf{x} - \mathbf{r}))}{k_c(\mathbf{x} - \mathbf{r})}, \quad (1.7)$$

where $k_c = \pi/\Delta$ is the cutoff wavenumber of the filter. When converted to spectral space, via the use of Fourier transforms, this becomes a Heaviside step function with respect to the wavenumber, k . The spectral-space form of the kernel has a value of 1 for wavenumbers of magnitude $\|k\| < k_c$ and a value of 0 otherwise. In spectral space, the filtering operation is performed by multiplying the spectral form of the kernel, \hat{G} by the variable to be filtered in spectral space $\hat{\phi}$,

$$\overline{\hat{\phi}} = \hat{G}\hat{\phi}. \quad (1.8)$$

Applying the sharp cutoff filter defined previously in this fashion serves to completely isolate the sub-grid component of ϕ which can then be removed from the governing equations to produce the filtered versions of the Navier-Stokes equations used in large-eddy simulations.

Large-eddy simulations of convective clouds require prognostic equations for the velocity components of the wind $\mathbf{u} = (u, v, w)$ with w being the vertical component. A temperature variable is also required, which is often the potential temperature, θ , defined as the temperature a parcel of air would attain if raised to a reference pressure. Another commonly used temperature variable is the liquid water potential temperature, θ_l , which is defined as the potential temperature that a parcel of air would attain if all liquid water content in the parcel were evaporated. This is popular in studies of cumulus convection as it incorporates the effects of moisture into the primary thermodynamic variable in the equations. Also required are at least two water variables, for example water vapour and cloud mixing ratios h_v and h_c . The number of water variables used will vary depending on the microphysical scheme.

Below, we present a minimal set of filtered equations for use in an LES, as detailed in Guichard and Couvreur (2017). The notation used is chosen to be consistent with existing literature on LES models, such as in Heus et al. (2010) and van Heerwaarden et al. (2017). The anelastic approximation of Ogura and Phillips (1962) is used, which approximates the density, ρ as only being a function of the vertical coordinate, i.e. $\rho(z) = \rho_r$. In many cases, the Boussinesq approximation is sufficient, particularly in the simulation of shallow cumulus and boundary layer studies. We use θ as the temperature variable and (u, v, w) for the wind with w being the vertical component. The governing equations are as follows. The equation of state is

$$\frac{T}{\theta} = \left(\frac{P}{P_0} \right)^{\frac{R_d}{C_p}}, \quad (1.9)$$

where T is the temperature, P_0 is a reference pressure, R_d is the dry air gas constant and C_p is the specific heat coefficient.

The filtered governing equations in an LES model are obtained by applying the spatial filter to the Navier-Stokes equations and any prognostic equations in order to remove the sub-grid terms ϕ' . The filtered continuity equation is,

$$\rho_r \frac{\partial \bar{u}_j}{\partial x_j} + \bar{w} \frac{\partial \rho_r}{\partial z} = 0. \quad (1.10)$$

The subscript i refers to the i^{th} component of the corresponding vector quantity, for example $u_3 = w$ is the vertical component of the velocity.

The filtered dynamic equation is

$$\frac{\partial \bar{u}_i}{\partial t} = -\frac{1}{\rho_r} \frac{\partial (\rho_r \bar{u}_i \bar{u}_j)}{\partial x_j} - \frac{1}{\rho_r} \frac{\partial P}{\partial x_i} + B \delta_{i,3} - \frac{1}{\rho_r} \frac{\partial (\rho_r \overline{u'_i u'_j})}{\partial x_j} - 2\epsilon_{i,j,k} \Omega_j \bar{u}_k \quad (1.11)$$

where the buoyancy force is given by $B = (g/\theta_r) (\overline{\theta_{vl}} - \theta_r)$ with g being the gravitational acceleration. θ_{vl} is the virtual potential temperature, which is defined as the potential temperature of dry air that would have the same buoyancy as moist air at a reference pressure, and θ_r is a reference potential temperature. The fourth term on the right-hand side corresponds to the force density on the fluid due to turbulent fluctuations and is given by the divergence of the Reynolds stress term, $\rho_r \overline{u'_i u'_j}$. Ω_j is the j^{th} component of the Earth's angular velocity, such that $2\epsilon_{i,j,k} \Omega_j \bar{u}_k = F_{cor}$ represents the effects of the coriolis force.

The thermodynamic equation is

$$\frac{\partial \bar{\theta}}{\partial t} = -\frac{1}{\rho_r} \frac{\partial (\rho_r \bar{u}_j \bar{\theta})}{\partial x_j} - \frac{1}{\rho_r} \frac{\partial (\rho_r \overline{u'_j \theta'})}{\partial x_j} + Q_{\text{rad}} + Q_{m\Phi} + \left(\frac{\partial \bar{\theta}}{\partial t} \right)_{\text{LS}}, \quad (1.12)$$

where Q_{rad} corresponds to heating due to radiative processes and $Q_{m\Phi}$ to heating from microphysical processes (i.e. phase changes of water). Terms denoted LS correspond to the influence of large-scale advection on a scalar quantity.

We also require prognostic equations for other scalar properties of the flow, such

as water vapour and hydrometeor species,

$$\frac{\partial \overline{\phi_x}}{\partial x_j} = -\frac{1}{\rho_r} \frac{\partial (\rho_r \overline{u_j \phi_x})}{\partial x_j} - \frac{1}{\rho_r} \frac{\partial (\rho_r \overline{u'_j \phi'_x})}{\partial x_j} + S_x + \left(\frac{\partial \overline{\phi_x}}{\partial t} \right)_{\text{LS}}. \quad (1.13)$$

Here S_x is a term used to group together all the sources and sinks of property ϕ_x , for example microphysical processes that influence the population of hydrometeor and water vapour mixing ratios (i.e. h_c and h_v).

The filtered governing equations must then be solved numerically over the whole domain, typically through finite-difference schemes or spectral methods. It is worth noting that choosing an appropriate numerical method depends largely on the set of equations and any simplifications. Other important numerical choices include the grid spacing, which must be chosen to satisfy the CFL condition, but using too fine a grid will dramatically increase computational expense. LES models tend to use a near isotropic grid spacing (i.e. $\Delta x = \Delta y = \Delta z$) but staggered grids can also be used, where different sets of variables are defined at different points on the grid cells. An example of this is shown in Mesinger and Arakawa (1976), where scalar variables such as potential temperature are defined at the centre of the grid cell while vector quantities such as wind speed are defined at the edges. A staggered grid can serve to reduce the distance between points on the finite difference stencils used, allowing a higher effective resolution.

1.3.3 Representation of sub-grid turbulence

As previously stated, large-eddy simulations operate by applying a low-pass spatial filter to the Navier-Stokes equations, effectively removing any motions smaller than the filter size, typically chosen in accordance with the model's grid spacing. However, this then means that the effects of the sub-grid scale motions must be represented in the model via parametrisation. In essence, this is very similar to how global-circulation models represent convection in that a closure assumption is required. The turbulent fluxes of variables are written as functions of their local gradients, i.e.

$$\overline{u'_j \phi'} = -K \frac{\partial \overline{\phi}}{\partial x_j} \quad (1.14)$$

where K is an eddy-diffusivity coefficient and describes the diffusion of ϕ due to turbulent eddies below the grid scale. Through this, the sub-grid model can model the turbulent motions in terms of the resolved flow. The filtering option in LES is designed to resolve fluid motions down to the inertial subrange, and the sub-grid model takes over in the viscous subrange (see Figure 1.4).

The transfer of energy across scales is well described in Vincent and Meneguzzi (1994) but the basic principle is that the largest eddies (with the smallest wavenumbers, k) are responsible for removing kinetic energy from the resolved flow. These eddies are typically found in the regions where the strain is particularly strong (i.e. high shear regions) and as such rapidly become unstable, breaking down into vortex tubes that become stretched and give way to even smaller-scale structures. This effect gives rise to the energy cascade illustrated in Figure 1.4. The exact nature of the boundary between the resolved and sub-grid scales in an LES will vary depending on the type of filter chosen (characterised by the convolution kernel, G), although a sharp cut-off is commonly used. At the end of this transfer of energy, the molecular diffusion takes over and kinetic energy is dissipated.

In the inertial subrange, the energy can be related to the wavenumber by the expression (Onsager 1949):

$$E(k) = \alpha \epsilon^{\frac{2}{3}} k^{-\frac{5}{3}} \quad (1.15)$$

where α is the Kolmogorov constant and ϵ is the viscous dissipation rate of turbulence kinetic energy. The universality of this constant has been discussed in Sreenivasan (1995).

The fundamental idea behind implementing a sub-grid scheme is that the transfer of energy across scales from resolved to sub-grid occurs via processes similar to molecular diffusion in the viscous range. As such, it was proposed to represent this transfer similarly, using a sub-grid eddy-viscosity, ν_{SGS} analogous to the




Figure 1.4: Schematic drawn showing the forward cascade of energy through the inertial subrange down to the Kolmogorov lengthscale, L_k , for isotropic turbulence (Kolmogorov 1941). The wavenumber, k , is represented in terms of the associated lengthscale, for, example, $k = 1/(L_k)$ is the wavenumber associated with the Kolmogorov lengthscale. The spirals are added to illustrate the decreasing size of the eddies associated across the subrange, and we mark the ranges in an LES that are resolved by the filtered equations and parameterised by the sub-grid model.

molecular viscosity. The main difference is that ν_{SGS} is effective at intermediate scales smaller than the resolved flow, but too large for molecular diffusion to be relevant. To define ν_{SGS} , characteristic lengthscales and timescales are required, i.e.

$$\nu_{SGS} = \frac{l_0^2}{t_0} \quad (1.16)$$

which is obtained by dimensional analysis, since the same units as diffusivity are required.

One of the more widely used sub-grid models is the Smagorinsky-Lilly model (Smagorinsky 1963; Lilly 1962). This sub-grid model and adaptations of it are used heavily in many of the LES-based models in use today due to its simplicity and efficiency. Smagorinsky (1963) assumed that any shear or buoyancy generated by the turbulence kinetic energy (TKE) is balanced by molecular dissipation and that the flow is in spectral equilibrium. A Kolmogorov spectrum similar to the illustration in Figure 1.4 was also assumed. The lengthscale used is determined by the filter width, Δ , multiplied by the Smagorinsky constant, C_s . The timescale associated with this viscosity is computed from the resolved component of the local strain rate tensor, giving the expression

$$\nu_{SGS} = C_s^2 \Delta (2\bar{S}_{ij}\bar{S}_{ij})^{\frac{1}{2}} \quad (1.17)$$

where

$$\bar{S}_{ij} = \left(\frac{\partial \bar{u}_i}{\partial x_j} + \frac{\partial \bar{u}_j}{\partial x_i} \right) \quad (1.18)$$

is the resolved strain rate tensor. Various studies have tried to determine the value of C_s , and it may not be universally constant, with Canuto and Cheng (1997) suggesting the properties of the flow determine it. Values typically used lie in the range $0.1 \leq C_s \leq 0.2$.

The sub-grid model is intended to capture the cascade of scales in turbulent flows and consequent mixing of dynamical and microphysical properties. While

the Smagorinsky model represents a simple model of this, there have been multiple adaptations based on it, such as the dynamic Smagorinsky model of Germano et al. (1991).

1.3.4 Representation of cloud microphysics

Hydrometeors in LES models are typically separated into cloud water and precipitation such as rain, hail etc. However, many processes contribute to the concentration and size of each. For example, rain drops can grow in size through the collection of cloud water in their surroundings or collisions with other rain droplets. Cloud water freezes to form ice droplets which can aggregate onto snow and riming of cloud ice can produce graupel and hail. Most LES-type models adopt a bulk water approach to hydrometeors and precipitation (Kessler 1969; Lin et al. 1983) which fits naturally into the Eulerian framework, computing the fraction of each species in the grid box to be converted by each process and updating after each timestep. These schemes are typically classified by the number of moments described by the distribution function of particle sizes. These moments include the mass/mixing ratio of each water species and the number concentration. In a single-moment scheme, the mixing ratio of each species is used, and droplet size distributions are computed as monotonic mixing ratios. Double-moment schemes also exist, which represent the mixing ratio and number concentration of each hydrometeor species prognostically. A double moment microphysical scheme, implemented in the Met Office Large Eddy Model (LEM) was able to predict a size sorting effect due to the increased fall speed of larger particles. The predictions of the model were consistent with radar observations, according to Swann (1998) but single moment schemes were not able to reproduce these results. The LEM microphysics scheme involved 34 distinct microphysical processes between 6 hydrometeor species.

A bulk-water approach is not the only option for modelling precipitation. Recently, Lagrangian ‘super-droplet’ schemes, such as the one described in Shima et al. (2009) have been developed. These models represent large quantities of rain

droplets by a single ‘super-droplet’, carrying size distributions and other information. The super-droplets are then advected in a Lagrangian manner according to their terminal velocity and the local wind speeds. Adapting such schemes to include other hydrometeor species may prove difficult and even then, it adds significant computational expense when compared to the conventional bulk water approach due to the large number of droplets required. As such, the method may be less feasible over larger domains. Needless to say, a fully Lagrangian scheme modelling individual hydrometeors is well beyond the limitations of current computational power.

1.3.5 Radiative and surface processes

A radiative scheme needs to account for the range of wavelengths present in solar radiation, absorption and emission lines of water molecules and any aerosols present. In the Met Office LEM, for example, the representation of radiative processes used is described in Edwards and Slingo (1996). The basic premise is that two spectral files containing gaseous absorption, impact of aerosols and Rayleigh scattering are read in, corresponding to solar and terrestrial radiation. The code uses a two-stream radiation scheme and splits the absorption lines in each type of radiation into 6 spectral bands for solar radiation and 8 for terrestrial radiation. It also accounts for the influence of precipitating hydrometeors and cloud ice when determining the volume extinction coefficients, single scatter albedo and asymmetry. Most conventional models also make use of the Eddington approximation for plane parallel atmosphere, which assumes that the spectral radiance is a linear function of the cosine of the angle of incidence, commonly used in stellar physics. Some LES models use similar representations of radiative transfer to those used in large-scale parameterisation schemes (i.e. Slingo 1989 and Ebert and Curry 1992).

The effects of surface processes also contribute, such as moisture and heat fluxes arising from the evaporation of seawater (or lakes, ponds in urban areas etc.) or other sources. Surface effects can be represented via the coupling of land or ocean surface models to the LES (Patton et al. 2005) to generate surface boundary

conditions. Alternatively, the model can prescribe surface fluxes to investigate the atmospheric response to given boundary conditions. Coupling of these models is often valuable for analysing the feedback loops between clouds and the larger environment, and more detailed land surface models can account for the effects of vegetation and soil levels. Ocean-atmosphere models have also been used to investigate interactions at the mesoscale (Wu et al. 1998) and the sensitivity of the boundary layer waves in the ocean (Sullivan et al. 2014).

1.3.6 Applications of LES

Since its introduction, LES models have proven to be a potent tool in developing our understanding of the atmosphere and is widely used throughout the atmospheric community. However, their application only started to take off in the 1990s, nearly 25 years after their initial development. Figure 1 in Stoll et al. (2020) shows the increase in articles mentioning LES and illustrates that the number of references to LES began to grow around this time. This was mostly likely due to advances in computing, allowing for the higher resolution studies that would demonstrate the method's strengths. While initially applied to studies of shallow convection similar to those seen in Sommeria (1976), LES models now see use in engineering, combustion (Pitsch 2006) and acoustics (Wagner et al. 2007). Simultaneously, advances in observational techniques allowed for better comparisons to real world convection on LES scales.

One of the biggest uses of LES models has been in the study of clouds interacting with their environments, particularly in cases of shallow convection. The use of passive tracers in Zhao and Austin (2005) allowed for a study of the life cycle of six cumulus clouds formed in a sheared environment. Their study revealed that a vortical structure in the ascending cloud top modulates entrainment through the top and sides. Modulating entrainment in this way produces a mushroom-like core as air is entrained into the cloud from below and rises up through the centre, mixed by turbulent eddies within. .

In other studies, Heus et al. (2008) used Lagrangian tracers to analyse the

mixing between clouds and their surroundings in the subsiding shells found around convective clouds. Later works by Hannah (2017) and Becker et al. (2018) also found these to play an important role in interactions between convective plumes and the large-scale environment, as the air around the cloud becomes diluted, resulting in the entrainment of less environmental air.

LES models are not exclusively limited to studies of shallow convection, however. Work by Böing et al. (2012) used the DALES model of Heus et al. (2010) to study the role of cold pools in promoting the transition between shallow and deep convection through Lagrangian tracers. Their study found that the presence of cold pools in the sub-cloud layer, formed by evaporative cooling due to rainfall, enhanced the intensity of the convection. It was also found that these cold pools promote the formation of deeper clouds with more intense precipitation, which in turn form larger cold pools in a positive feedback loop. Their results also suggested that convective organisation within the sub-cloud layer was important to the triggering of deep convection.

Creative extensions of LES models have also been designed over the years, for example in the simulation of a tropical cyclone in Rotunno et al. (2009). In their study, a nested grid was used to provide increased resolution in the central regions of the circulation where the need for high resolution was greatest, avoiding the computational expense of maintaining the fine grid spacing over a large domain. This method was used again in Ren et al. (2020) to study the changes in wind speed distribution to sea surface temperature changes. Ren et al. (2022) used the same model to determine the radius of maximum wind in a hurricane and related this to an inflection point in the angular momentum contours. Other examples of this approach can be found in Moeng et al. (2007) and Mirocha et al. (2014), where an LES model was embedded in a mesoscale simulation.

Meanwhile, advances in computational power are leading to applications of LES models over larger domains, such as those carried out in Heinze et al. (2017) and Stevens et al. (2020). Similarly, the development and application of superparameterisation schemes (Lin et al. 2022b) is a rapidly developing field. While global

LES modelling may be a distant dream, great advances have been made in integrating them into the development and refinement of convection schemes in global models such as in Angevine et al. (2018). The growing use of machine learning techniques in this field too (Mooers et al. 2021; Sansom 2021) mean that LES and similar models will play a major role in the global representation of convection for a long time

These examples all serve to illustrate just how wide the range of applications for LES models is. Trying to recount every significant result obtained from LES studies would far exceed the scope of this review, so we have chosen a range of examples to emphasise its effectiveness. Despite the limits of computational power, LES results have influenced our understanding of convection across a vast range of scales and are likely to remain one of the pillars of the field for the foreseeable future as new developments continue.

1.3.7 Limitations of the LES models

Despite their wide range of successes, LES models have some shortcomings of their own. These, in some ways, parallel the problems of global circulation models in representing convection. Despite operating at finer scales than global or regional models, LES models are often limited by their representation of sub-grid processes. For example, in a mirror to the issues of global models, there is a ‘gray zone’ close to the grid spacing in which work-sharing can occur as the model must differentiate between resolved and sub-grid processes. This can have consequences for the computational efficiency.

Studies of cumulus clouds by Matheou et al. (2011) and Pressel et al. (2015) have shown that LES models are susceptible to the choice of numerical methods (Stevens et al. 1999). Increasing numerical dissipation results in lower precipitation and reduced cloud formation, displaying a sensitivity to the resolution of the model. Similarly, different numerical schemes found variations in precipitation. These results have been similarly observed in studies of stratocumulus clouds (Stevens et al. 2005; Pressel et al. 2015). Likewise, Takemi and Rotunno (2003) found that

varying the eddy-viscosity coefficient offered increased resolution of convective cells within a mesoscale system. However, making the constant too large caused the cells to become smoothed, losing many features of the fine-scale structure.

One area in which these difficulties become particularly relevant is in the convective boundary layer. Sullivan and Patton (2011) investigated the effects of varying grid spacing on the properties of the boundary layer. They found that, for low-order moment statistics such as fluxes, they become independent of the grid spacing once there is sufficient scale separation between the boundary layer height and the cut-off scale for the sub-grid scheme (i.e. when $z_i \gg C_s \Delta$). Not only that, but they found that the bulk entrainment rate of the boundary layer was a highly sensitive measure of convergence of the LES, increasing as the grid becomes coarser. However, the entrainment rate still became independent of the grid spacing at sufficient resolution. This is likely due to the failure to resolve temperature gradients in the entrainment zone adequately. Furthermore, in the entrainment zone, the temperature variance was found to increase as the grid became finer.

As a result of the sub-grid models and the highly turbulent nature of the boundary layer causing difficulties in LES studies, subsequent studies have opted to use direct numerical simulation (DNS) (Garcia and Mellado 2014; Mellado et al. 2017; Haghshenas and Mellado 2019). These studies have offered considerable insight into the vertical structure of the boundary layer, but are limited to relatively low Reynolds numbers, orders of magnitude lower than typical atmospheric values.

These shortcomings do not diminish the great successes of LES and similar models, it should be stressed. The weaknesses of LES models are far outweighed by the numerous improvements they have provided to our understanding of atmospheric convection. Advances in computing and new approaches to the representation of sub-grid turbulence are likely to see LES models improve in the future, and they remain promising candidates to be embedded in large-scale superparameterisation models. However, the sub-grid representation leaves somewhat undesirable numerical sensitivities that may present an opening for alternative methods of modelling turbulence. Such techniques would be particularly beneficial in the convective

boundary layer, where DNS studies are computationally expensive and struggle to approach atmospheric Reynolds numbers. Nonetheless, LES models will likely remain a staple of research into atmospheric convection far into the foreseeable future. Any alternative methods developed are more likely to supplement the insights gained from LES studies rather than aim to supplant them.

1.4 A Lagrangian perspective

1.4.1 Lagrangian and semi-Lagrangian models

Lagrangian methods within fluid dynamics involve the simulation of fluid motions following particles of fluid, themselves moving with the flow. The discretisation of the fluid is based on the definition of these particles. This is in contrast to the traditional Eulerian method used within many atmospheric models, in which the fluid is discretised onto a grid (Cartesian or otherwise). The absence of this grid in Lagrangian methods means they naturally lend themselves to seeking alternative representations of turbulence in comparison to the eddy-viscosity schemes used in LES models. Similarly, the representation of the fluid in terms of particles itself comes with many advantages.

Models developed in a purely Lagrangian framework have seen considerable success in other fields, for example, smooth particle hydrodynamics (SPH) (Lucy 1977; Monaghan 1992). This method forgoes an underlying grid and uses a smoothing function to model the distributions of mass and other properties around the Lagrangian particles. The method can naturally incorporate an adaptive resolution, as the lengthscale associated with the smoothing function can be adjusted in more active regions of the flow to provide a greater spatial resolution (Gingold and Monaghan 1977; Rosswog 2009). As such, it lends itself well to the vast range of scales present in astrophysical flows, ranging from studies of star forming regions such as Smilgys and Bonnell (2017) to the evolution and structure of black hole accretion disks (Nealon et al. 2015). Moreover, the absence of an underlying grid allows the method to easily accommodate complex boundary dynamics, such as those found

in free-surface flows in engineering (Violeau and Rogers 2016). SPH has even been used to simulate multiphase flows (Chen et al. 2015) such as at the air-sea interface or oil-water interactions. Such features may also support a more natural implementation of the Lagrangian precipitation scheme proposed by Shima et al. (2009) to develop a fully Lagrangian representation of cumulus convection.

The representations of turbulence within SPH are still a developing area of research, and most studies are limited to two-dimensional turbulence within engineering flows (Monaghan 2017; Monaghan and Mériaux 2018) rather than the study of atmospheric phenomena. Three-dimensional turbulence within SPH, in general, is very difficult to simulate accurately. Even in idealised channel flows such as in Mayrhofer et al. (2015), the model failed to reproduce the statistical properties observed in DNS of the Taylor-Green vortex flow (Taylor and Green 1937).

However, one promising avenue of research comes from the idea that the smoothing operation within SPH is similar in framework to the filtering carried out in LES, despite the difference in model formulation. The proposal of a Lagrangian SPH-LES model in Di Mascio et al. (2017) outlines the potential framework required for a meshless SPH model that incorporates the turbulence representation of LES. It should be noted, however, that these models would rely on Lagrangian adaptations of the turbulence schemes commonly used within LES models (i.e. Smagorinsky 1963 and Lilly 1966). Therefore, it is likely they may encounter similar limitations and sensitivities. It remains to be seen if the Lagrangian formulation and built-in adaptive resolution are sufficient for the study of atmospheric flows while avoiding the drawbacks of an Eulerian approach. Nonetheless, combining these techniques may yet lead to exciting developments as the methods are refined.

Such models also come with unique challenges, however. One of the common issues found in mesh-free methods is incompressibility, as the clustering of particles can result in ‘holes’ elsewhere in the flow and consequently produce inaccuracies in the density and pressure fields of the flow (Swegle et al. 1995; Monaghan 2000). In hybrid approaches to the problem, with an underlying grid, there remains the

option of redistribution of particles periodically to ensure the flow remains incompressible, such as in the combined Lagrangian advection method of Dritschel and Fontane (2010). This involves compromising the Lagrangian integrity of the particles occasionally, but in turn allows for the quasi-conservative redistribution of particles and their properties to prevent the numerical instabilities incurred by insufficient particle concentrations. Other issues arise at the boundaries in SPH models. While well-suited to free-surface flows and systems with displaced boundaries, the implementation of wall boundary conditions or even inlet and outlet conditions has proven challenging to simulate in SPH. In this case, the absence of an underlying grid in any form makes these boundaries particularly problematic to represent. The development of such boundary conditions for SPH models is considered one of the ‘grand challenges’ of the SPH research and engineering international community (SPHERIC), according to Vacondio et al. (2021).

While the development of combined SPH-LES schemes is promising, it seems unlikely that a fully Lagrangian model will be able to accurately simulate highly turbulent atmospheric flows such as those found in convective clouds and the boundary layer. This is likely to remain the case, at least for the near future. However, it is not impossible, and the SPH method would bring considerable advantages due to its adaptive resolution and the ability to incorporate complex boundary dynamics. Instead, it may be more beneficial to pursue the development of hybrid models, in which an underlying grid supports the Lagrangian dynamical core in some capacity.

Even within LES, semi-Lagrangian methods have already seen use through the introduction of tracers to study the mixing of clouds and their environment (Zhao and Austin 2005; Heus et al. 2008; Böing et al. 2012). Some LES models have even adopted semi-Lagrangian schemes for advection. These schemes work by modelling the advection of Lagrangian particles according to the gridded velocity fields at each timestep. The particle properties are then interpolated back onto the grid to reconstruct the gridded fields afterwards (see Zerroukat et al. 2002 and Kaas 2008 for examples). This process is essential to ensure the conservation of mass

in such schemes because neither the continuity nor transport equations appear in flux form within semi-Lagrangian advection schemes. As such, they do not explicitly conserve mass. Introducing a mass-conservation scheme to these models also incurs additional computational expense that can make them undesirable.

Similar tracer advection methods have also been used in climate scale studies of atmospheric chemistry and transport in the stratosphere (Konopka et al. 2004; Stenke et al. 2008; Stenke et al. 2009). These methods are relatively effective when the properties carried by the Lagrangian tracers are not directly coupled to the dynamics of the flow, such as in passive transport studies. The key advantage of these methods is that passive tracers in a Lagrangian framework do not require additional prognostic equations as they would in a grid-based method, because the Lagrangian particles carry tracer concentrations as they are advected. However, the use of such schemes introduces inconsistencies between the Lagrangian and Eulerian descriptions of the local air density which, in turn, produced discrepancies in tracer concentrations and other properties. Grewe et al. (2014) attributed these inconsistencies to an incompatibility between the tracers' advection and the model's dynamical core, as column-based chemical processes carried out using the grid often led to errors in the associated Lagrangian quantities. The proposed solution in Grewe et al. 2014 was that models with a Lagrangian dynamical core should be employed. In their work, they explore two candidates: the finite mass method (FMM) of Gauger et al. (2000) and the Hamiltonian particle-mesh (HPM) method of Frank et al. (2002).

The FMM is another mesh-free method, and the primary feature of this model is that, rather than discretising the spatial domain, the system's mass is discretised into a series of overlapping packets. These packets each carry their own mass distributions, often described through the use of a smooth base function (Klinger et al. 2005), similar to the SPH method using a smoothing kernel for particle properties. The method requires prognostic equations for both the centre of each mass packet and the associated deformation matrix describing the mass distribution. The fluid density at a given point in space and time can then be obtained

by summing the corresponding densities of any overlapping mass packets at that point. Grewe et al. (2014) propose a reshaping of the mass packets to form an underlying auxiliary grid used for chemical processes. Applications of this method are somewhat limited. One such scheme has been developed by Dong et al. (2015) and was tested in three dimensions using the flow generated by an LES model.

The HPM method is a hybrid method which uses both particles and an underlying longitude-latitude grid that initialises many more particles than the FMM (approximately an order of magnitude greater). Whereas in the FMM, the volume of each particle is defined by its deformation matrix, HPM particles do not carry volume. The prognostic variables for this method are the particle centres and the pressure gradient derived from the overall particle mass distribution. Also, unlike the FMM, there is no additional step required to construct an auxiliary grid, as the model already makes use of one. However, studies using the HPM have mainly been limited to two-dimensional models using the shallow water equations (e.g. Frank and Reich 2004 and Bresch et al. 2020).

While these semi-Lagrangian schemes have been primarily applied to transport studies in the larger atmosphere, their application to cumulus convection is relatively limited. Still bound to the Eulerian nature of LES-based models, they are subject to the same sub-grid schemes representing turbulent mixing and thus many of the resulting sensitivities. This presents a niche for models of a more Lagrangian nature, somewhere intermediate to these methods without altogether abandoning the Eulerian grid as in SPH. Furthermore, such a model would be well-suited to the transport of tracer concentrations due to its Lagrangian nature, and able to adopt a more explicit representation of sub-grid turbulence with Lagrangian particles providing a similar adaptive resolution.

1.4.2 The Moist Parcel-In-Cell method

The Moist Parcel-In-Cell (MPIC) method of Dritschel et al. (2018) represents an attempt to create a model for studying atmospheric convection that is almost entirely Lagrangian in nature. Although it currently exists in a highly idealised

setting, it has shown promising comparisons to an existing LES model in the Met Office NERC model (MONC, based on the LEM in Brown et al. 2015), documented in Böing et al. (2019). The method is worth exploring, as it occupies a position intermediate between fully mesh-free techniques like SPH and the traditional Eulerian approach to convection. As previously discussed, this framework allows the model to explore more physical representations of turbulence and sub-grid mixing. Simultaneously, the formulation of the model makes adding tracers and chemical species relatively straightforward.

The model is similar to the existing particle-in-cell methods commonly used in plasma physics (Harlow 1962; Brackbill 1991; Markidis and Lapenta 2011), but draws inspiration more heavily from the vortex-in-cell method of Christiansen and Zabusky (1973). Similar methods have even been used previously in convection studies, albeit in two dimensions (Gadian 1991). The basic principle of the model is that all physical quantities are evolved on the Lagrangian parcels, using an underlying Cartesian grid to determine the physical distributions of properties. The particles carry vorticity and are used to determine the system’s dynamical evolution while the grid is used to compute the velocity field for advection. This is accomplished through interpolating the parcel vorticity to produce a gridded vorticity field, which is then inverted to generate the velocity field (more information on this is given in Chapter 2, where we present an overview of the method).

Turbulence and mixing in MPIC differ from LES models and are represented through the splitting and removal of parcels. In particular, the integrated vorticity stretch of a parcel is calculated at each dynamical timestep. Once it exceeds a specified threshold, the parcel splits into two identical parcels, with the original parcel’s properties shared between them. This mirrors the adaptive resolution effect seen in SPH by increasing the number of parcels, and simultaneously decreasing their size, in highly turbulent flow regions. Of course, this cannot continue indefinitely, as the computational expense incurred would be too great. Therefore, parcels are removed from the model once they fall below a specified volume. Their properties are then redistributed onto the grid and subsequently interpolated onto the sur-

rounding parcels to ensure conservation. This process presents a more physical, explicit representation of the sub-grid turbulence, although other model adaptations have modified the mechanism determining when splitting occurs (Frey et al. 2022).

The advantage of this sub-grid representation is that the numerical diffusion is considerably weaker than in an LES-type model of similar grid resolution. The parcel formulation of MPIC allows it to carry more information about the sub-grid scales, meaning that coarser grid resolutions can be used, and thus can be more efficient computationally. At the same grid resolution, the parcel formulation of MPIC allows for a higher effective resolution in active regions of the flow.

Initial comparisons of the method with MONC in Böing et al. (2019) show very good agreement between the two methods, despite the vast differences in their formulation. As such, any subsequent work on modelling known case studies in MPIC would do well to compare to existing LES and DNS models. The visual appearance of the clouds formed in Böing et al. (2019) are very similar, although due to the highly turbulent nature of the flow, the smallest scales differ (see the comparison of cloud edges from Böing et al. 2019 in Figure 1.5). The bulk properties of the flow too are found to be in strong agreement with one another, which is a mark in the method's favour.

At smaller scales, there is more considerable deviation from the results observed in MONC, in particular that the effects of mixing in MPIC are weaker, especially at low resolution. In an LES model, the turbulence scheme will rapidly diffuse thermodynamic properties such as liquid water once they enter the sub-grid regime. This is not the case in the MPIC simulations carried out in Böing et al. (2019), best shown in the liquid water probability distribution functions where there are sharp peaks in the distribution at higher wavenumbers. This represents an undesirable resolution sensitivity within MPIC, although the distributions do appear to converge with increasing resolution.

The representation of the vorticity field within MPIC is also particularly useful, as studies of the vorticity dynamics in cloud mixing are relatively limited. In stud-

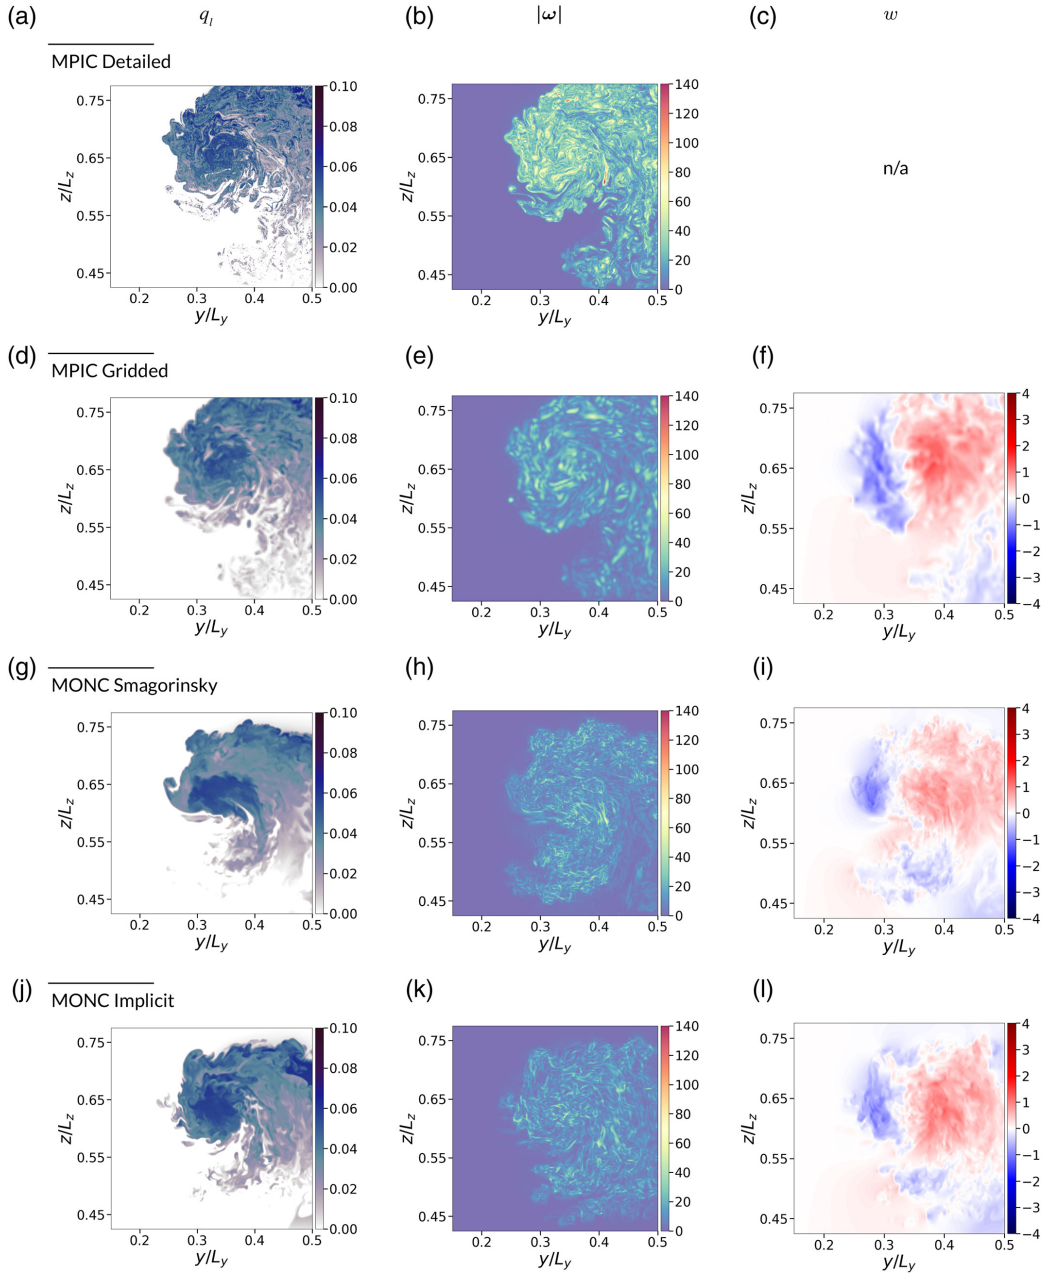


Figure 1.5: Detailed view of a cloud edge showing the liquid water specific humidity (a, d, g, j), vorticity (b, e, h, k) and vertical velocity (c, f, i, l) fields for clouds formed in MPIC using high resolution interpolation of the parcel field (a, b, c) and the grid resolution (d, e, f). These are compared to MONC simulations using Smagorinsky (g, h, i) and implicit (j, k, l) sub-grid models. Reproduced from Figure 2 of ‘Comparison of the Moist Parcel-In-Cell (MPIC) model with large-eddy simulation for an idealized cloud’ by Böing et al. (2019) in QJRMS under the Creative Commons license (CC-BY 4.0).

ies of the boundary layer, the spatial distribution of enstrophy has been proposed as a potential conditioning parameter to distinguish between the turbulent and nonturbulent regions of the atmosphere (Fodor and Mellado 2020). This may have some relevance too to the two-layer structure observed in DNS simulations of the entrainment zone, where the upper sublayer is believed to represent a transition layer between two regions. A model like MPIC is therefore uniquely positioned to study the vorticity in this region, and perhaps offer further insights into the structure of the boundary layer.

Furthermore, the addition of Lagrangian tracers representing aerosols, chemical species or other parcel attributes fits naturally into the MPIC framework. Tracer concentrations would be carried as materially conserved attributes on each parcel, subject to the same splitting and mixing operations and, as such, do not require additional prognostic equations to be solved as they would in an Eulerian method. This also makes the addition of tracers and attributes within MPIC far less expensive computationally in comparison to LES models, for example, and may lend itself well to studying distributions of different chemical or hydrometeor species in future studies. One straightforward example of tracer usage in MPIC was shown in Böing et al. (2019), tracking the origin height of parcels to compute their vertical displacement (Fig 1.6). This could be useful in studies of the boundary layer or cloud development to identify the origin of entrained air. The main advantage of MPIC, over the handling of tracers in LES models, is that there is no need for additional prognostic equations associated with each attribute. Not having to solve these additional prognostic equations reduces the cost of potential tracer implementations in MPIC. Likewise, many of the consequences for tracer concentrations resulting from the inconsistency between Lagrangian parcels and an Eulerian dynamical core could be avoided. Another benefit to the use of tracers in MPIC is that, because it keeps the tracers monotone, it does not create false extrema or spurious negative values. However, there are still concerns regarding incompressibility within Lagrangian models.

Despite its weaknesses, the MPIC method remains a promising candidate for




Figure 1.6: Cross-section through the centre of the domain showing the vertical displacement of parcels for an MPIC simulation at a resolution of 256^3 , shown at (a) $t=4$, (b) $t=6$ and (c) $t=8$. Reproduced from Figure 10 of ‘Comparison of the Moist Parcel-In-Cell (MPIC) model with large-eddy simulation for an idealized cloud’ by Böing et al. (2019) in QJRMS under the Creative Commons license (CC-BY 4.0).

a Lagrangian model of atmospheric convection. The model retains many of the strengths of fully mesh-free models such as SPH, while avoiding the difficulties of representing boundary conditions due to the underlying grid. The explicit sub-grid representation of turbulence makes it a very appealing method in highly turbulent flows where the mixing of air is important, such as in the growing boundary layer or cloud edges. The model may also, in later stages of development, be suitable as an embedded scheme within superparameterisation models and consequently the favourable comparisons of bulk flow properties to an existing LES model in Böing et al. (2019) are particularly promising.

1.5 Overview of thesis structure

In this thesis, we aim to further develop the MPIC method and showcase its application beyond the initial work shown in Dritschel et al. (2018) and Böing et al. 2019. These changes are made with the goal of applying the MPIC method to atmospheric problems and case studies that were not previously possible with the method and have not previously been modelled with Lagrangian methods, such as the atmospheric boundary layer, as is done in Chapter 5. Such problems, in the future, may benefit greatly from the Lagrangian nature of the method through the

computationally efficient implementation of tracer populations to study transport problems, for example in atmospheric chemistry. Analysis of the energy budget in MPIC may also lend to its use in guiding parameterisation schemes in larger scale models. It is hoped that the work in this thesis will provide a stepping stone towards a more realistic Lagrangian model of convection that may be used alongside existing methods to further our understanding of cloud development and boundary layer convection.

To this end, we demonstrate the application of MPIC to additional case studies corresponding to sheared thermals and the dry atmospheric boundary layer that may be of interest to the atmospheric community. While still in a relatively idealised state, the MPIC method provides valuable insight into these case studies. Moreover, as additional layers of complexity are added to MPIC’s microphysical representation, or through the addition of Lagrangian diagnostics these studies can be further extended. One of the key advantages of MPIC is the natural implementation of additional parcel attributes without the need to solve further prognostic equations and as greater microphysical and chemical complexity is added to the model, such an advantage only becomes more pronounced. As such, the work documented here primarily presents some early steps in moving MPIC towards fully realistic modelling of convection, while demonstrating the method’s potential, even in its current state. Many of the diagnostics used in later chapters, such as the vertical distributions of thermodynamic properties and energies associated with thermals, may even be relevant for coupling to larger-scale circulations as a potential embedded convection model.

An overview of the method and numerical parameters used in this thesis are presented in Chapter 2. The majority of this chapter is a recap of the methodology outline in the original papers on the method. We do, however, present a modification to the vertical solver within MPIC, moving from a semi-spectral method to a tri-diagonal solver. This change is necessary to facilitate the implementation of surface fluxes within the method, as it was not possible using the model presented in Dritschel et al. 2018. This modification is used in all subsequent works in

this thesis, and is particularly relevant to later chapters studying boundary layer growth.

Chapter 3 studies a modification of the original MPIC test case. We have smoothed the thermal edge to limit the sharp buoyancy gradients mentioned in Böing et al. (2019) that gave rise to more intense vorticity than would normally occur. We study the ascent of the new thermal, subject to a constant vertical wind shear of different strengths. While vertical shear is thought to be important for the organisation of deep convection, the effects of wind on individual thermals has not been studied in particular depth. Furthermore, our analysis studies the enstrophy of the air, partitioned into cloud and dry air components by liquid water content to assess the effects of entrainment into the rising thermal on the cloud properties.

Chapter 4 presents an analysis of the rising thermal in the absence of vertical shear, similar to the previous test case. The analysis presented can easily be applied to sheared cases too, but this chapter aims to assess how well energy is conserved within MPIC and the effects of numerical choices within the method. As such, it benefits from studying a simplified test case. However, we now derive a measurement of the potential energy associated with the thermal that is available for convection and briefly assess its sensitivity to numerical parameters within MPIC. Furthermore, the potential energy derived here is analogous to the convective available potential energy used in convective parametrisations and some closure assumptions. We discuss how this energy varies with height for parcels of different origins within the test case. Showcases of possible modifications to the background stratification that could subject parts of the thermal to convective inhibition are also given, which may be useful in studies of convection driven by surface fluxes.

Finally, Chapter 5 presents a study of the dry convective boundary layer growing into a linearly stratified atmosphere in the absence of vertical shear. This is, to the author's knowledge, the first study of its kind to use a Lagrangian method like MPIC and requires substantial changes to the method to include the effects

of surface fluxes of buoyancy. The flux implementation is highly advantageous within the MPIC framework, as the same method can be applied to fluxes of any parcel attribute, normally requiring additional prognostic equations in an LES or DNS model. As such, future works using MPIC would easily be able to implement surface fluxes of, for example, aerosol species, moisture, or even vorticity. We include analysis of the enstrophy field and entrainment rate parameters of the boundary layer in comparison to existing LES and DNS studies of the entrainment zone. In particular, we compare our results to the two-layer vertical structure of the entrainment zone observed in Garcia and Mellado 2014 and assess MPIC's ability to model the boundary layer, despite the differences in framework.

Chapter 6 presents a summary of the key findings in these studies and a discussion of future extensions and applications for the method. We provide an assessment of the potential of MPIC in its present state as an alternative or supplement to conventional convection models on both the small and large scale.

Chapter 2

Overview of the Moist Parcel-In-Cell method

This chapter largely presents a recap of the work of Dritschel et al. (2018) and Böing et al. (2019). While not original work, details of the method are essential to the understanding of subsequent chapters, and as such are presented here. Full details of the MPIC method are given in Dritschel et al. (2018). We also detail the modifications made to the vertical solver in section 2.3 which is used for all subsequent studies and has since been integrated into later versions of MPIC and its parallel version PMPIC. Chapter specific changes to the model will be discussed in the corresponding sections discussing the experimental setup for each chapter.

2.1 Governing Equations

The governing equations are given below. The notation used in this chapter differs from the description of LES models in Section 1.3 as it was chosen to be consistent with previous papers on the development of the MPIC method (Dritschel et al. 2018; Böing et al. 2019)

$$\frac{D\mathbf{u}}{Dt} = -\frac{\nabla p}{\rho_0} + b\hat{\mathbf{e}}_z, \quad (2.1)$$

$$\frac{Db_l}{Dt} = 0, \quad (2.2)$$

$$\frac{Dq}{Dt} = 0, \quad (2.3)$$

$$\nabla \cdot \mathbf{u} = 0, \quad (2.4)$$

where \mathbf{u} is the velocity, p is the non-hydrostatic pressure, b_l is the liquid water buoyancy and q is the specific humidity. $D/Dt = \partial/\partial t + \mathbf{u} \cdot \nabla$ is the material derivative. ρ_0 is a constant reference density, equivalent to the ρ_r in Section 1.3. At present, the density in these equations differs from the LES example given in section 1.3 because we use the incompressible Boussinesq approximation to simplify the equation set and filter out sound waves. The approximation assumes that fluid density variations are small in comparison to the average fluid density in the domain (Durran and Arakawa 2007) and as such the density profile is approximated to be constant with respect to the vertical coordinate, z . This causes the pressure term to vanish in a vorticity formulation (see equation 2.7).

The total buoyancy is approximated by

$$b = b_l + \frac{gL}{c_p\theta_{l0}}q_l, \quad (2.5)$$

where

$$q_l = \max(0, q - q_0e^{-\lambda z}), \quad (2.6)$$

is the liquid water content, q_0 is the surface saturation humidity and λ is the inverse condensation scale height. The liquid water buoyancy is given by $b_l = g(\theta_l - \theta_{l0})/\theta_{l0}$ with θ_l being the liquid water potential temperature and θ_{l0} a reference value. This equation represents a simple, height-based representation of condensation. Once a parcel of air reaches the height at which the parcel specific

humidity exceeds that of the background profile, the excess humidity is converted into liquid water.

The prognostic equation for vorticity, $\boldsymbol{\omega} = (\xi, \eta, \zeta)$, is obtained by taking the curl of the momentum equation (2.1),

$$\frac{D\boldsymbol{\omega}}{Dt} = \boldsymbol{\omega} \cdot \nabla \mathbf{u} + \left(\frac{\partial b}{\partial y}, -\frac{\partial b}{\partial x}, 0 \right). \quad (2.7)$$

The solenoidal condition ($\nabla \cdot \boldsymbol{\omega} = 0$) is also enforced on the vorticity field (taking the curl of equation 2.4). We note that the equivalent form of the vorticity equation discussed in Dritschel et al. (2018), as proposed by Cottet and Koumoutsakos (2000), is no longer used in MPIC in favour of numerically solving the prognostic equation for vorticity on the parcels.

2.2 Parcel representation

The core feature of MPIC is its representation of the fluid as a collection of parcels, in many ways similar to the particle representation used in smoothed particle hydrodynamics or the hamiltonian particle mesh method. As was stated in Chapter 1, MPIC originates from the vortex-in-cell method of Christiansen and Zabusky (1973) (see also Cottet and Koumoutsakos 2000). All fields within MPIC are evolved using the parcels, which move freely within the domain. In order to reconstruct the fields for visualisation and inversion, an underlying Cartesian grid is used. The inversion of the vorticity field on this grid recovers the velocity field that is then re-interpolated back to the parcels for advection.

Parcels in MPIC carry vorticity, $\boldsymbol{\omega}$, liquid water buoyancy b_l , specific humidity q and parcel volume V . Except for vorticity, these are materially conserved quantities that lend themselves well to the parcel representation of the model, since no further prognostic equations are required to describe their evolution. This feature of the model will be especially beneficial in studies involving tracer concentrations or the addition of aerosol species, since the additional computational expense is

reduced in comparison to grid-based models. Parcels also allow for a relatively straightforward representation of processes such as condensation and evaporation, although it is unclear how effective this will be in the case of a full microphysical scheme. Finally, in order to accurately model the flow with MPIC, the entire domain is filled with parcels. In the simulations in this thesis, MPIC is initialised with 8 parcels per grid cell in regions not occupied by an initial thermal, and 64 parcels per grid cell in cells where an initial thermal is present.

In MPIC, we transfer the parcel properties to the underlying grid by means of trilinear interpolation in which the grid box is broken into smaller boxes which are used as weights for interpolation. We denote gridded quantities by a bar, i.e. \bar{q} while parcel quantities have the subscript i . The gridded value of specific humidity $\bar{q}(\bar{\mathbf{x}}, t)$ is obtained at a grid point $\bar{\mathbf{x}}$ from the equation,

$$\bar{q}(\bar{\mathbf{x}}, t) = \bar{V}^{-1} \sum_{i \in \mathcal{P}(\bar{\mathbf{x}})} \phi(\mathbf{x}_i - \bar{\mathbf{x}}) q_i V_i, \quad (2.8)$$

where

$$\bar{V}(\bar{\mathbf{x}}, t) = \sum_{i \in \mathcal{P}(\bar{\mathbf{x}})} \phi(\mathbf{x}_i - \bar{\mathbf{x}}) V_i, \quad (2.9)$$

and the trilinear weights are given by

$$\phi(\mathbf{x}_i - \bar{\mathbf{x}}) = \left(1 - \frac{|x_i - \bar{x}|}{\Delta x}\right) \left(1 - \frac{|y_i - \bar{y}|}{\Delta y}\right) \left(1 - \frac{|z_i - \bar{z}|}{\Delta z}\right). \quad (2.10)$$

Here \mathcal{P} denotes the set of all parcels within the 8 grid boxes surround $\bar{\mathbf{x}}$. We can then do the reverse operation to interpolate the velocity field from the grid back to the parcel positions,

$$\mathbf{u}(\mathbf{x}_i, t) = \sum_{i \in \mathcal{G}_i(\bar{\mathbf{x}})} \phi(\mathbf{x}_i - \bar{\mathbf{x}}) \bar{\mathbf{u}}(\bar{\mathbf{x}}, t), \quad (2.11)$$

where \mathcal{G}_i is the set of all 8 grid points at the corners of the grid point containing

the i^{th} parcel.

2.3 Vorticity inversion

To obtain the velocity field needed for parcel advection, we must invert the vorticity field obtained from equation 2.7. In the horizontal, a spectral solver is used with periodic boundary conditions, further details of which are included in Dritschel et al. (2018). Previous versions of MPIC required that the horizontal vorticity components, ξ and η must be zero on the lower boundary to allow their representation in the inversion problem as a sine series in the vertical coordinate. This was done to allow for a simple and efficient spectral solution in the vertical. Consequently, the buoyancy at the lower surface could not vary, as the horizontal vorticity components are forced by buoyancy gradients. In order to implement surface heating for later studies on realistic convection and boundary layer turbulence, it would be beneficial to allow for surface heating of the domain, represented by a flux of buoyancy at the lower surface. To this end, a tridiagonal solver is implemented, with fourth-order compact differencing (Esfahanian et al. 2005) supplemented by one-sided compact differences at the boundaries (Ghader et al. 2011), resulting in third-order accuracy. This method also allows for the eventual addition of variable fluxes of other parcel properties, such as moisture, at the lower boundary of the simulation. This marks the main deviation from the method described in Dritschel et al. (2018) and is used in all subsequent works within this thesis.

The vertical extent of the domain in MPIC is between the lower boundary at the surface $z = 0$ and a rigid boundary at the top of the domain $z = L_z$. In order for the buoyancy to vary at the lower surface, it must be possible to have non-zero horizontal vorticity components at the lower boundary, since these are forced by the horizontal derivatives of buoyancy.

$$\frac{D\xi}{Dt} = \boldsymbol{\omega} \cdot \nabla u + b_y, \quad \frac{D\eta}{Dt} = \boldsymbol{\omega} \cdot \nabla v - b_x. \quad (2.12)$$

Our inversion problem is,

$$\nabla^2 \mathbf{A} = \boldsymbol{\omega}, \quad \mathbf{u} = -\nabla \times \mathbf{A}, \quad (2.13)$$

where \mathbf{A} is a vector potential and $\nabla \cdot \mathbf{A} = 0$ has been imposed without loss of generality. We then convert to semi-spectral space by taking Fourier transforms in the x and y directions (with wavenumbers k and l respectively). The velocity components are then given by,

$$\hat{u} = \hat{B}_z - il\hat{C}, \quad \hat{v} = ik\hat{C} - \hat{A}_z, \quad \hat{w} = il\hat{A} - ik\hat{B}, \quad (2.14)$$

where the subscript z denotes differentiation with respect to the vertical coordinate and \hat{A} , \hat{B} and \hat{C} are the components of $\hat{\mathbf{A}}$, the vector potential in semi-spectral space. We can use the non-divergence of the vector potential to define the boundary conditions,

$$\hat{A} = \hat{B} = \hat{C}_z = 0, \quad (2.15)$$

on the boundaries $z = 0$ and L_z , since we must have $\hat{w} = 0$.

Thus, the interior equations satisfied by \hat{A} , \hat{B} and \hat{C} are

$$\begin{aligned} \hat{A}_{zz} - (k^2 + l^2)\hat{A} &= \hat{\xi}, \\ \hat{B}_{zz} - (k^2 + l^2)\hat{B} &= \hat{\eta}, \\ \hat{C}_{zz} - (k^2 + l^2)\hat{C} &= \hat{\zeta}, \end{aligned} \quad (2.16)$$

where $\hat{\xi}$, $\hat{\eta}$ and $\hat{\zeta}$ are the x , y and z vorticity components (semi-spectral) respectively.

This problem can be solved to second order accuracy using centred differencing to discretize into the tridiagonal problems for the vector potential components. However, we opt to use the fourth-order compact differencing of Esfahanian et al. (2005) which gives us the tridiagonal problem,

$$\begin{aligned} \frac{\hat{A}_{j+1} - 2\hat{A}_j + \hat{A}_{j-1}}{\Delta z^2} - \frac{K^2}{12}\hat{A}_{j+1} - \frac{5K^2}{6}\hat{A}_j - \frac{K^2}{12}\hat{A}_{j-1} \\ = \frac{1}{12}\hat{\xi}_{j+1} + \frac{5}{6}\hat{\xi}_j + \frac{1}{12}\hat{\xi}_{j-1}, \end{aligned} \quad (2.17)$$

$$\begin{aligned} \frac{\hat{B}_{j+1} - 2\hat{B}_j + \hat{B}_{j-1}}{\Delta z^2} - \frac{K^2}{12}\hat{B}_{j+1} - \frac{5K^2}{6}\hat{B}_j - \frac{K^2}{12}\hat{B}_{j-1} \\ = \frac{1}{12}\hat{\xi}_{j+1} + \frac{5}{6}\hat{\eta}_j + \frac{1}{12}\hat{\eta}_{j-1}, \end{aligned} \quad (2.18)$$

$$\begin{aligned} \frac{\hat{C}_{j+1} - 2\hat{C}_j + \hat{C}_{j-1}}{\Delta z^2} - \frac{K^2}{12}\hat{C}_{j+1} - \frac{5K^2}{6}\hat{C}_j - \frac{K^2}{12}\hat{C}_{j-1} \\ = \frac{1}{12}\hat{\zeta}_{j+1} + \frac{5}{6}\hat{\zeta}_j + \frac{1}{12}\hat{\zeta}_{j-1}, \end{aligned} \quad (2.19)$$

$$j = 1, \dots, n_z - 1, ,$$

where $K = \sqrt{k^2 + l^2}$, to be solved for \hat{A} , \hat{B} and \hat{C} .

Because the boundary conditions for the vertical component of the vector potential, C , are different to the two horizontal component, we need to assign a value to $\hat{\zeta}_{-1}$ across the lower boundary and similarly across the upper boundary in order to define the source terms on the right-hand side of 2.19. To do this, we use the semi-spectral version of the solenoidal condition on vorticity,

$$ik\hat{\xi} + il\hat{\eta} + \hat{\zeta} = 0, \quad (2.20)$$

which can be used to obtain $\hat{\zeta}$ in terms of $\hat{\xi}$ and $\hat{\eta}$. Applying this at the boundaries allows us to obtain expressions for the required terms,

$$\hat{\zeta}_{-1} = \hat{\zeta}_1 + 2\Delta z \left(ik\hat{\xi}_0 + il\hat{\eta}_0 \right) \quad (2.21)$$

$$\hat{\zeta}_{n_1} = \hat{\zeta}_{n_z-1} - 2\Delta z \left(ik\hat{\xi}_{n_z} + il\hat{\eta}_{n_z} \right). \quad (2.22)$$

This second order approximation of the source terms, leads to a fourth-order ac-

curate representation of \hat{C} .

To obtain the vertical derivatives, \hat{A}' and \hat{B}' , we must solve another tridiagonal problem,

$$\frac{1}{6}\hat{A}_{z,j+1} + \frac{2}{3}\hat{A}_{z,j} + \frac{1}{6}\hat{A}_{z,j-1} = \frac{\hat{A}_{j+1} - \hat{A}_{j-1}}{2\Delta z}, \quad j = 0, \dots, n_z. \quad (2.23)$$

These are supplemented by one-sided compact differences at the boundaries, as in Ghader et al. (2011), due to the finite difference stencils extending beyond the domain. However, this actually results in a loss of accuracy, and the scheme only achieves overall third-order accuracy.

2.4 Parcel splitting and mixing

In MPIC, mixing is represented by the splitting and removal of parcels from the model. The splitting of parcels is controlled by the integrated vorticity stretching, estimated by

$$\gamma_i(t) = \int_{t_0}^t \left(|\boldsymbol{\omega}_i \cdot \frac{d\boldsymbol{\omega}_i}{dt}| \right)^{\frac{1}{3}} dt. \quad (2.24)$$

The parcel is split when this value exceeds $\gamma_{\max} = 4$, although tests in Dritschel et al. (2018) show that the model is not sensitive to the choice of γ_{\max} . Upon splitting, the parcel is replaced by two equal sized parcels separated by a distance determined by the original parcel volume. A minimum volume is then defined to decide at which point parcels should be removed from the model. When a parcel is removed, its attributes are redistributed onto the grid corners and interpolated back onto neighbouring parcels. This procedure is covered in more detail in Dritschel et al. (2018). Currently, the minimum volume for splitting is set to $\Delta V/6^3$ where ΔV is the grid cell volume, although this choice is analysed further in Chapter 3.

2.5 Non-dimensionalisation

The equations used in MPIC are non-dimensional, in particular, we non-dimensionalise time using the characteristic squared buoyancy frequency $g\lambda\Delta\theta_{l0}/\theta_{l0} = 1$ where $\Delta\theta_{l0}/\theta_{l0} = 0.01$ is a characteristic fractional variation of the liquid water potential temperature. The lengthscale, λ is non-dimensionalised using $1/\lambda = 1$ where $\lambda = 2000$ m is the scale height. Combining these gives a non-dimensional gravity of $g = 1/(\Delta\theta_{l0}/\theta_{l0}\lambda)$, such that $g = 1/(0.01 * 1) = 100$. The specific humidity is then scaled by its saturation value at ground level q_0 and the dimensionless equation for buoyancy is

$$b = b_l + b_m \max(0, \bar{q} - e^{-z}), \quad (2.25)$$

where $\bar{q} = q/q_0$ and

$$b_m = \frac{gLq_0}{c_p\Theta_{l0}}. \quad (2.26)$$

We use values of $L/c_p = 2500$ K, $q_0 = 0.015$, and $\theta_{l0} = 300$ K, giving $b_m = 12.5$. As previously stated, we make use of the incompressible Boussinesq approximation, which is appropriate for shallow convection. In such a case it can be argued that the Boussinesq approximation is not valid as this enters the realm of deep convection. The approximation is used regardless for the purpose of keeping the dynamics relatively straightforward. It may be reasonable to reduce the scale of the case-defining heights to more typical values for shallow cumulus convection and adopt an anelastic approach in which the density is a function of height (e.g. Ogura and Phillips 1962) in future work.

Chapter 3

Response of a rising thermal to vertical wind shear in a Lagrangian framework

3.1 Introduction

Clouds and moist convection in the atmosphere can significantly influence the large-scale environment (Arakawa and Schubert (1974), Mekonnen et al. (2006)) and play a large role in the global regulation of heat and moisture. In particular, the large-scale organisation of convection can lead to the formation of mesoscale convective systems (MCSs) (Houze 2004) such as squall lines and thunderstorms which bring intense precipitation (Parker and Johnson 2000) and can cause severe disruption. As such, it is important to understand the processes that influence the development of such systems and their effects on convective systems of all scales.

In particular, vertical wind shear has a large impact on the formation and development of MCSs, especially squall lines (Rotunno et al. 1988, Takemi 2006), and individual clouds. Several studies have shown how the structure of a MCS can change drastically in response to strong wind shear at different levels (e.g Chen et al. 2015). For example, strong wind shear in the lower troposphere (from 0 km to 5 km) leads to interactions with cold pools in the boundary layer and causes a preference to develop into lines of clouds (Rotunno et al. 1988, Moncrieff 1992)

while mid-level (5 km to 10 km) shear appears to encourage rotation and favours development of supercells (Coniglio et al. 2006) such as those in which tornadoes form. Strong shear in levels above this point seems to cause a flattening of cloud tops (Sathiyamoorthy et al. 2004) and can inhibit convection through increased exposure to radiation. Vertical wind shear can also affect the entrainment of dry air and the turbulent mixing, which is important when considering the spread of aerosol particles in the atmosphere (Wang and Prinn 1998, Yamaguchi et al. 2019).

Most work is focused on the role of wind shear in convective organisation and as such make use of cloud resolving models (CRM), grid-based models with a resolution typically of the scale 1 km in comparison to the 100 m scale typically seen in large-eddy simulation (Sommeria 1976, Deardorff 1980). These models are limited by computational power and do not resolve motions on the smallest scales, instead using sub-grid parametrisations involving eddy-closure assumptions (Lilly 1962, Smagorinsky 1963). As such, there is little discussion of the impacts of vertical wind shear on the smallest scales or for individual clouds, where thermodynamic properties can vary on lengthscales of a few metres (Austin et al. 1985). Lagrangian approaches, such as the Hamiltonian particle method (HPM, Klingler et al. 2005) and finite mass method (FMM, Frank et al. 2002) have seen some success in studies of atmospheric chemistry/transport (Grewe et al. 2014) when compared to Eulerian methods and may be able to afford better resolution, but are underused due to the need for tuning of various numerical parameters.

While the impacts of shear on convective organisation are well documented in the literature, previous studies have overwhelmingly focused on advanced stages of deep convective systems, such as supercells (Weisman 1992; Weisman and Rotunno 2000; Warren et al. 2017). Comparatively, studies of isolated thermals in the literature are limited. Despite this, they offer considerable insight into the role of wind shear in the initiation and development of early-stage convection, such as shallow cumulus clouds. Some studies into sheared thermals have suggested that sheared up and downdrafts are weaker than their shear-free counterparts (Peters 2016), despite the fact that shear is known to enhance convection in organised

convective systems. A more recent study in Peters et al. (2019) also found that the deepening of convection in sheared environments was actively inhibited by vertical wind shear.

Moreover, the study of individual thermals and plumes is relevant to the development and refinement of convective parameterisation schemes. Because such schemes aim to represent an ensemble of clouds and plumes through their statistical properties, they depend heavily on our own understanding of individual clouds. A core assumption of the earlier schemes (Arakawa and Schubert 1974; Yanai et al. 1973) are that the plumes are non-interacting with one another. These schemes also have difficulties with representing shallow convection (Arakawa 2004; Sander-son et al. 2010). As such, the application of higher resolution cloud simulations, like LES and MPIC, to the study of individual clouds and their response to vertical wind shear may yield results of use to global models that make use of modern parameterisation schemes. In particular, some studies on isolated thermals aim to assess the role of vertical wind shear in the initiation of convection, feeding into the trigger functions and closures used in convective parameterisation schemes (i.e. Pan and Randall 1998). As such, further analysis of sheared thermals using novel approaches and analysis may offer further insight to be used in the refinement of the representation of convection on the global scale. For example, a particle based method may be beneficial in identifying exactly how vertical shear inhibits convective motions, potentially through the use of tracers, or even in sheared transport studies.

To this end, we employ the Moist Parcel-In-Cell (MPIC) method described in Dritschel et al. 2018, a Lagrangian, vorticity dynamics based model for studying moist convection that uses an underlying grid. The MPIC method is still in an early stage, but an initial comparison of the model to the Met Office NERC model (MONC) detailed in Böing et al. (2019) has shown promising results. In particular, MPIC is able to resolve cloud features on much smaller scales than a MONC simulation of the same grid resolution due to its Lagrangian nature. This is especially helpful in studies of turbulence and mixing/entrainment. The model also

uses a more physical representation of mixing than the parametrisation schemes adopted in LES/CRM simulations through the splitting and removal of parcels. These features make it a strong choice to analyse the influence of vertical wind shear on the development of an individual cloud and to look at properties such as enstrophy which can be used as a measure of turbulence.

We set up idealised simulations of the ascent of a spherical thermal, initially at rest, in environments with a linear vertical shear profile of varying strength and compute bulk properties, such as enstrophy and liquid water volume, to gain some insight into the effects of shear strength on the development and structure of the resulting cloud. We separate the enstrophy and kinetic energy into cloud and dry air components to demonstrate the increasing degree of entrainment. In Section 3.2 we provide an overview of the test case. Section 3.3 contains the main results and discusses their implications. Finally, Section 3.4 summarises our results and considers possibilities for refinements and extensions.

3.2 Physical Model

We use the MPIC model as described in Chapter 2 and our test case is similar to that described in Dritschel et al. (2018). The initial setup of the thermal is shown later in Figure 3.1. The lower region of the environment is a mixed layer, with neutral liquid water buoyancy, $b_l = 0$ and a uniform specific humidity $\tilde{q}_{env} = \mu\tilde{q}_{th}$, where $\mu = 0.9$ is a constant and \tilde{q}_{th} is the specific humidity inside the thermal. Above this is a stratified region, beginning at the level of dry neutral buoyancy, $z_b = 2.38$ (corresponding to the humidity of parcels located in the mixed layer, $\tilde{q}_b = 0.8$), and the dimensionless Brunt-Väisälä frequency in this region is $N = 0.97$, such that there is no discontinuity in the buoyancy field at this point. We also specify the height of the lifting condensation level, $z_c = 2.5$, which is used to determine the specific humidity of parcels located within the thermal, $\tilde{q}_{th} = e^{z_c}$. Finally, we specify the level of moist neutral stratification, $z_m = 5.0$, and the level of dry neutral stratification, $z_d = 4.0$. We use periodic horizontal boundary conditions with a rigid, free-slip boundary to represent the surface and

upper boundary of the domain.

We can calculate the buoyancy inside the thermal, from

$$b_{l,th} = N^2 (z_d - z_b). \quad (3.1)$$

We also introduce an asymmetry in buoyancy within the thermal such that

$$b_l = b_{l,th} \left(1 + \frac{e_1 x' y' + e_2 x' z' + e_3 y' z'}{R^2} \right), \quad (3.2)$$

where x' , y' and z' represent the displacement from the centre of the thermal and $e_1 = 0.3$, $e_2 = -0.4$ and $e_3 = 0.5$ are dimensionless constants. These are chosen to ensure that the mean value of b_l is unchanged. This asymmetry is carried over from the initial test case in Dritschel et al. (2018) and is used to break the symmetry of the simulation. While the shear in our test cases may accomplish this, a small asymmetry can still lead to more interesting dynamics, and as such, we have kept it in our study.

Where this initial setup differs is that we have now opted to smooth the edge of the initial thermal in an attempt to make the model more realistic. The sharp edges in previous MPIC simulations lead to strong vorticity and resolution dependence so we have altered the properties accordingly. The radius of the thermal used in this work is now $R = 0.88$ and we allow buoyancy and humidity to transition to the background levels in the mixed layer over the outer 20% of the radius.

To accomplish this, we introduce a function of the radial position within the thermal that tells us how far between the inner edge and outer edge of the smoothed region each parcel, i , is during initialisation and apply this to each parcel within the smooth region:

$$d_i = \frac{\sqrt{x_i'^2 + y_i'^2 + z_i'^2} - 0.8R}{R - 0.8R} = \frac{r_i - 0.8R}{R - 0.8R} \quad (3.3)$$

Note that $d_i = 1$ when $r_i = 1$ and $d_i = 0$ when $r_i = 0.8R$ and the first derivative is zero at both. With such a function, we can represent the buoyancy of a parcel,




Figure 3.1: x - z cross-sections of the initial buoyancy (a) and specific humidity (b) fields for the test case.

b_i in this region by

$$b_i = b_{l,th} (1 - F_{sm}(d_i)) \quad (3.4)$$

where $b_{l,th}$ has the same meaning as before and F_{sm} is a smoothing function used to smooth the sharp edges of the thermal. This will give a parcel at $r/R = 1$ a buoyancy of zero, making a continuous transition to the environmental value. The humidity for a parcel can be calculated from

$$\tilde{q}_i = q_{th} (1 + (\mu - 1) F_{sm}(d_i)), \quad (3.5)$$

such that the humidity of a parcel at the edge of the thermal is $q_i = \mu \tilde{q}_{th}$.

The smooth step function used for this purpose is

$$F_{sm}(d) = 6d^5 - 15d^4 + 10d^3. \quad (3.6)$$

The buoyancy and specific humidity of the initial thermal at rest are shown in Figure 3.1.

The final and most important change to the initial conditions is the implementation of a vertical wind profile across the entire domain. In MPIC, this is done

relatively simply by applying an initial distribution of the horizontal vorticity components. Because we use free-slip boundary conditions at the vertical boundaries, this vorticity will persist at all times in the simulation, resulting in a constant in time background wind profile. A wind shear profile, can be obtained by applying a background vorticity profile to the domain, i.e.

$$S_x(z) = \frac{d\langle u \rangle}{dz} = \langle \eta \rangle, \quad S_y(z) = \frac{d\langle v \rangle}{dz} = -\langle \xi \rangle. \quad (3.7)$$

Angled brackets refer to the horizontal averages of quantities. In this work, we take the simplest possible case, a horizontal wind in the x -direction that increases linearly with height and is zero at the bottom of the domain. This corresponds to adding a constant value to $\langle \eta \rangle$, since the vorticity is equivalent to the vertical gradient of horizontal velocity, which determines the strength of the vertical shear profile. It may be feasible in future works to apply more complex shear profiles.

To investigate the effects of increasing shear, we use a constant value for the shear parameter S_x throughout the domain and vary this from 0.0 to 1.0 in intervals of 0.1. The upper limit here is roughly equivalent to the non-dimensional squared buoyancy frequency (N^2) in the environment. Going beyond this value would likely lead to overturning and in runs where this was attempted, the vorticity values became excessively large early in the run. This is likely because the buoyancy frequency determines the stratification and stability of the atmosphere, so as the vertical shear becomes stronger, the environment becomes susceptible to Kelvin-Helmholtz instabilities. While these instabilities can occur within clouds in real convection, if they affect the entire domain due to the shear being too strong, the entire environment becomes unstable, leading to the observed vorticity blow-up. It is reasonable to assume that such cases are generally not common in the atmosphere and do not offer much insight into cloud development. The simulations are run out to a non-dimensional time of 20 (roughly double the duration of those presented in Dritschel et al. (2018)) to study the development of the system beyond the initial ascent of the thermal.

3.3 Results

Here we present the various diagnostics and bulk properties computed to analyse the impact of vertical wind shear. Unless stated otherwise, these properties were computed using the gridded values of the various attributes. For any diagnostic which depends linearly on a prognostic variable, the values obtained from the parcel and gridded based variables are identical by design.

3.3.1 Flow evolution

In this section, we present visual representation of the flow evolution for a selection of shear parameters ($S = 0.3, 0.5, 0.7, 1.0$). This is shown in Figures 3.2-3.5.

In these figures, the actual 'cloud' that is formed by the rising thermal is visualised by the liquid water content shown in the third column of each figure. As such, this is immediately useful to qualitatively assess the effects of shear on the shape of the cloud. The first snapshots, taken at $t = 5$ can highlight the impact of shear on the initial ascent of the thermal. In the lower shear cases, it is immediately apparent that the thermal penetrates the stratified region of the domain to a greater vertical extent, and is less stretched horizontally. The cloud at the early stages shows a 'mushroom'-like shape not dissimilar to the shear-free test case in Dritschel et al. (2018) and Böing et al. (2019), albeit stretched in the direction of the horizontal wind. In the case of stronger vertical shear, this is not present, and even at early times, the cloud starts to fragment. This is in line with the results obtained in other studies of sheared thermals (Peters et al. 2019) which also found that the vertical extent and deepening of convection is inhibited by vertical shear.

At later stages in the cloud's evolution, it becomes spread across the upper region of the domain. In the case of weaker vertical shear, the cloud area remains mostly contiguous. In the higher shear cases, however, the cloud water is scattered across the domain in various fragments, and it is debatable whether the images really depict a cloud at all. At a glance, there also appears to be less cloud volume in the domain as the shear strength is increased, and it occupies a lesser vertical

extent than low shear cases.

The buoyancy field in these figures can, to some extent, be used to track the displacement of air in the system as the thermal rises. Dark red regions correspond to cooler air from the mixed layer that becomes entrained into the thermal as it rises through the domain. As the thermal rises through the domain, air from the mixed layer is pulled into the space behind the rising thermal. Similarly, air above the thermal is pushed out of the way during its ascent and forced downwards. Consequently, air from below is pulled up into the stratified region and vice versa. To distinguish between air entrained from below and the air associated with the thermal, we can compare the buoyancy cross-sections to the corresponding liquid water cross-sections, as buoyant air from the thermal that rises above the condensation level will form cloud water. As such, less buoyant air in the upper regions of the domain that do not overlap with the cloud itself is likely to have been entrained from the mixed layer. In the highest shear case (Figure 3.5) we can clearly see wisps of air from the mixed layer being pulled up above the cloud. This is particularly evident at $t = 10$ on the right of figures 3.5 (e) and (g) and is present at later times too but appears to be mixed with the upper air. We can see that more air is being caught up within the cloud itself as the shear is increased. From these cross-sections alone, however, it is difficult to assert that vertical shear is enhancing the entrainment of air into the stratified region of the domain. Many of these features are also visible in the specific humidity cross-sections shown in the second column.

Finally, the vorticity field paints an interesting picture and can show how turbulence spreads through the domain as a result of the cloud's ascent. At early times in all simulations, the regions of high vorticity are primarily located at the edges of the thermal. However, we can also observe that the vorticity field appears to trace out the path of the thermal during its ascent, leaving a 'trail' behind it where the environmental air has been disrupted. This is particularly evident in the low or mid-shear cases. As the system is allowed to evolve further, vorticity rapidly spreads throughout the domain as environmental air becomes caught up in the




Figure 3.2: x-z cross sections through the centre of the domain for shear parameter $S_x = 0.3$ showing the buoyancy, b , (a,e,i,m), specific humidity, q , (b,f,j,n), liquid water content, q_l , (c,g,k,o) and vorticity magnitude, $|\omega|$, (d,h,l,p). The cross-sections are taken at $t = 5$ (a-d), $t = 10$ (e-h), $t = 15$ (i-l) and $t = 20$ (m-p).




Figure 3.3: x - z cross sections through the centre of the domain for shear parameter $S_x = 0.5$ showing the buoyancy, b , (a,e,i,m), specific humidity, q , (b,f,j,n), liquid water content, q_l , (c,g,k,o) and vorticity magnitude, $|\omega|$, (d,h,l,p). The cross-sections are taken at $t = 5$ (a-d), $t = 10$ (e-h), $t = 15$ (i-l) and $t = 20$ (m-p).

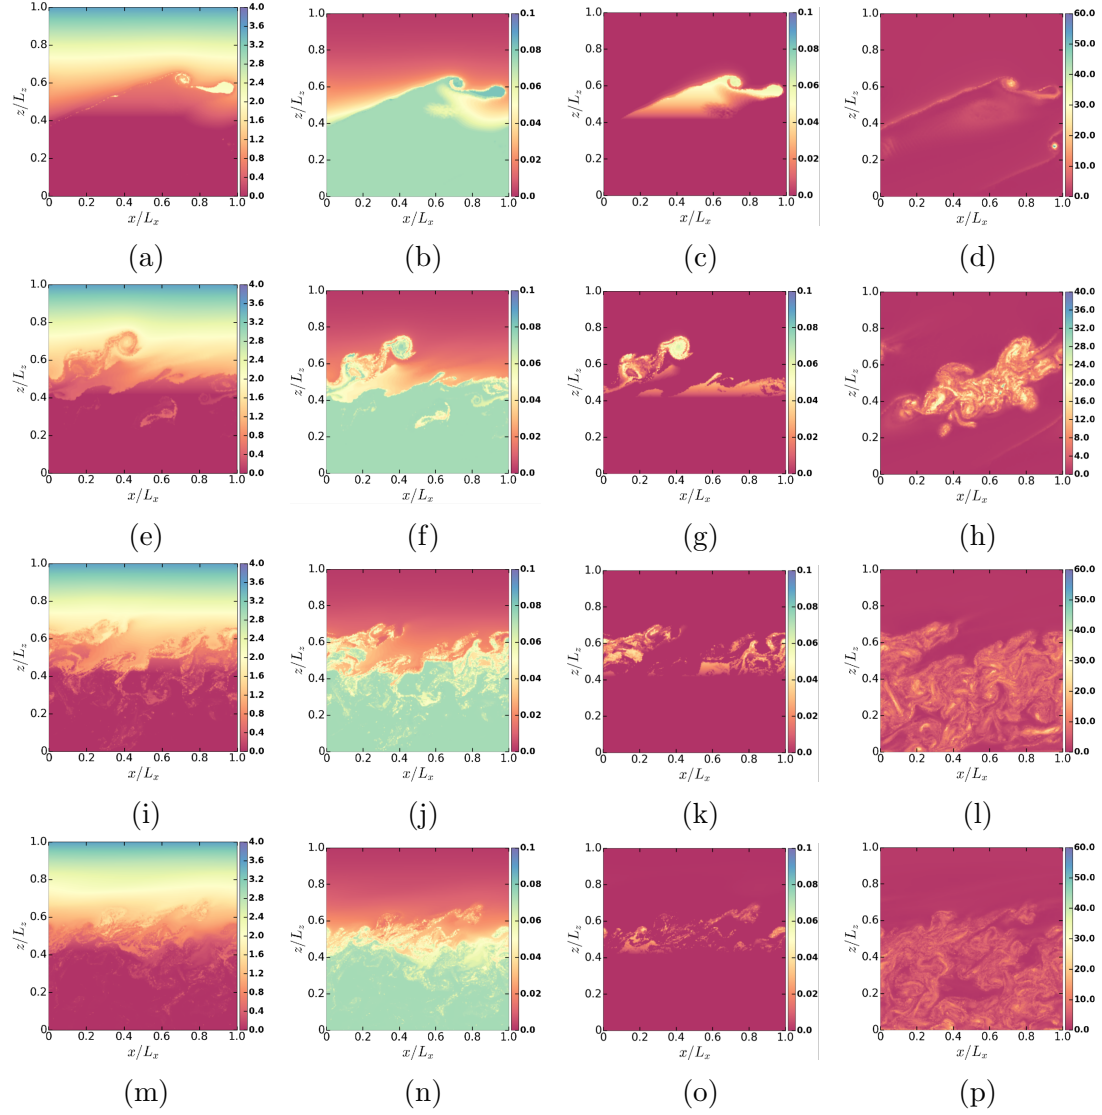


Figure 3.4: x - z cross sections through the centre of the domain for shear parameter $S_x = 0.7$ showing the buoyancy, b , (a,e,i,m), specific humidity, q , (b,f,j,n), liquid water content, q_l , (c,g,k,o) and vorticity magnitude, $|\omega|$, (d,h,l,p). The cross-sections are taken at $t = 5$ (a-d), $t = 10$ (e-h), $t = 15$ (i-l) and $t = 20$ (m-p).

turbulent eddies forming around the thermal's path. At the latest stages of these simulations, large regions of the mixed layer have become turbulent. This process happens much more quickly in the case of strong vertical shear and vorticity rapidly covers almost the entirety of the lower domain.

These cross-sections offer a qualitative insight into the ascent of the thermal and evolution of the system in response to vertical shear. However, a quantitative analysis would be greatly beneficial to quantify the impact of increasing shear strength on an individual clouds development.

3.3.2 Cloud volume and depth

The most obvious property to study when determining how shear influences the development of cumulus clouds is the liquid water volume – the total amount of cloud produced in each simulation. This is relatively easily computed by summing the liquid water fraction across all grid cells/parcels and multiplying by the volume occupied by each. We use the gridded fields here as they are saved more frequently due to taking up less space and therefore offer a greater temporal resolution. The time evolution of the liquid water volume is shown in Figure 3.6.

The first thing of note is that the point at which the cloud volume reaches its peak value seems to occur at a later time as the vertical shear strength increases, which suggests that the shear may be hindering the vertical motion of the thermal. This is in agreement with the results of Peters (2016) and Peters et al. (2019), as well as the observed cloud cross-sections in Section 3.3.1. It is also interesting that the initial peak is actually higher in the low-mid shear runs, with the highest value occurring for a shear parameter of $S_x = 0.4$. As the vertical shear gets stronger past this point ($S_x = 0.7 - 1.0$) the initial growth of the cloud becomes very slow, although there is a greater accumulation of liquid water than the low shear cases. This is likely a result of shear slowing the vertical motions of cloud parcels in the simulation, meaning not only that they rise more slowly, but also do not descend as rapidly afterwards. If we focus on the initial growth though, it suggests that some degree of wind shear may be favourable for the growth of larger clouds. This




Figure 3.5: x - z cross sections through the centre of the domain for shear parameter $S_x = 1.0$ showing the buoyancy, b , (a,e,i,m), specific humidity, q , (b,f,j,n), liquid water content, q_l , (c,g,k,o) and vorticity magnitude, $|\omega|$, (d,h,l,p). The cross-sections are taken at $t = 5$ (a-d), $t = 10$ (e-h), $t = 15$ (i-l) and $t = 20$ (m-p).

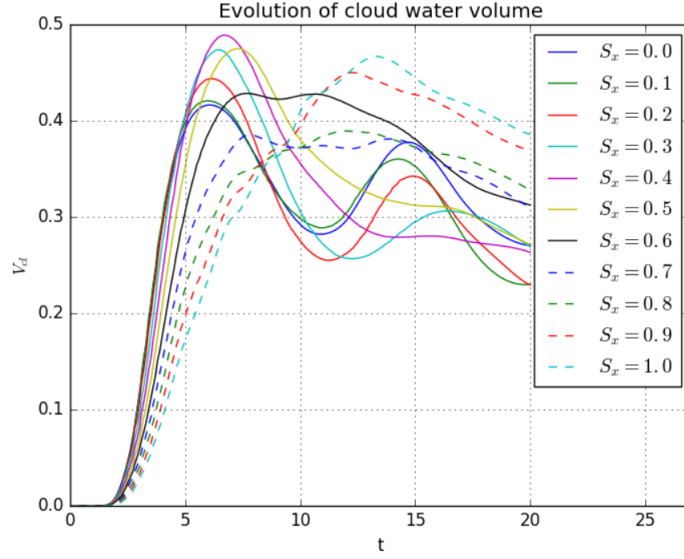


Figure 3.6: Time evolution of total liquid water volume for shear values ranging from $S_x = 0.0$ to $S_x = 1.0$

result is not immediately obvious from the cross-sections shown previously. In particular, shear in the lower levels may have the greatest influence on the early development of an individual cumulus cloud. Separating the shear profile into lower, mid and upper level components would be an interesting exercise to see at which level shear has the greatest impact.

We can look further into the vertical extent of the clouds by calculating a liquid water weighted mean height and standard deviation with serves as a ‘half-depth’. These are shown in Figure 3.7. As expected, the cases with little to no vertical shear generally have a higher mean height and this decreases as the vertical shear gets stronger. It is also worth noting that the mean cloud height in every case seems to oscillate slightly while slowly settling to a steady value, suggesting it evolves towards an equilibrium state. Higher shear values seem to result in a larger half-depth, meaning the cloud is more spread out vertically in the case of high wind shear. We also see that the initial growth slows down as shear increases, similar to the cloud volume. At first glance, this could be interpreted as the highest shear cases resulting in larger, deeper clouds being built up over a longer period

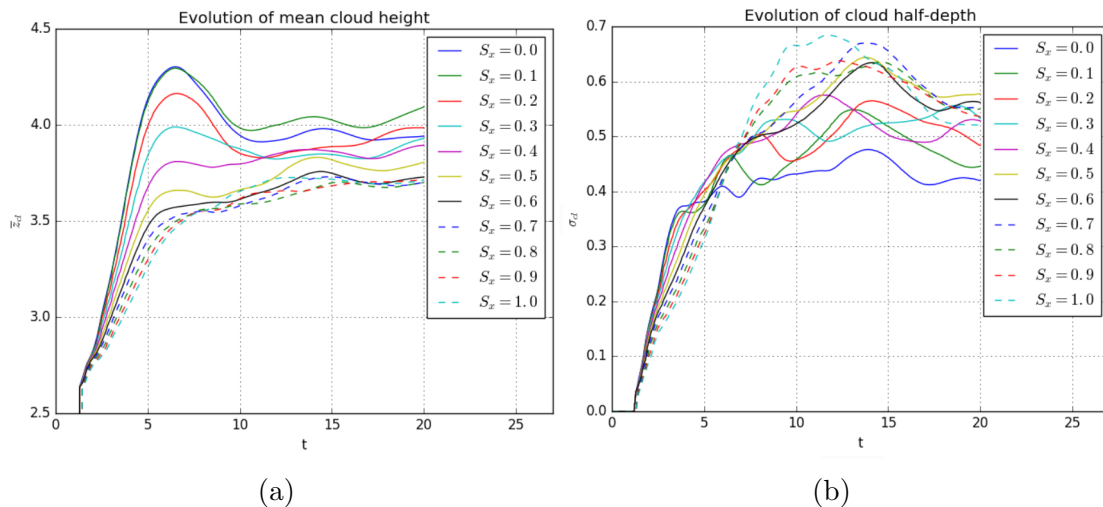


Figure 3.7: Time evolution of mean cloud height (a) and half-depth (standard deviation, (b)) for all values of shear parameter S_x .

of time. However, the results in Section 3.3.1 clearly show that the clouds in these cases are being pulled apart by the vertical shear and the convection at work is much less intense, resulting in very little vertical motion of cloud water.

3.3.3 Vertical motion

To investigate this further, we calculate the maximum updraft velocity for each simulation to gain some insight into the influence of shear on vertical motions within the domain. This is done in a simple manner, simply finding the maximum value of w at each save time in the simulation. We also compute the total kinetic energy associated with vertical motions of the system integrated over the entire domain, $K_{v,tot} = \frac{1}{2} \int_V w^2 dV$.

Results are shown in figure 3.8. They show the strong impact of shear strength on the initial ascent of the thermal, with the peak updraft velocity in this stage decreasing steadily as S_x increases. The vertical kinetic energy (KE) shows a similar trend at first, but at later times in higher shear cases we can see that the peak vertical kinetic energy grows with increasing shear. It almost appears as if there are two regimes visible in the kinetic energy evolutions, with a transition

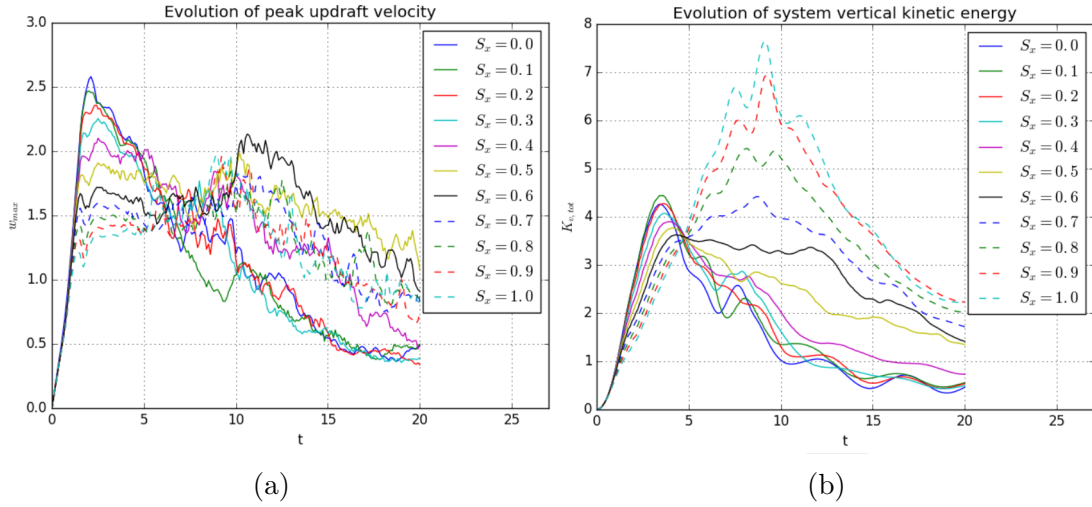


Figure 3.8: Time evolution of peak updraft velocity (a) and vertical kinetic energy (b) for shear values ranging from 0.0 to 1.0

occurring at shear strengths of around $S_x = 0.6$ or $S_x = 0.7$. In the weak-shear regime, we see that the vertical KE peaks rapidly during the initial ascent of the thermal, before also decaying rapidly at later times. In the high-shear regime, the vertical KE continues to grow, peaking at much later times and decaying slowly afterwards.

Many of the vertical KE evolutions appear to intersect at a time of approximately $t = 4.3$. Here the models transition from the early ascent phase where shear is inhibiting the vertical motions to later stages where vertical motions are enhanced by the wind shear, as evidenced by the vertical KE evolution after this point.

At a glance, these results tell us that the vertical motions at later times are more intense in the higher shear cases. However, it is possible that a large component of this energy is due to dry air which has been entrained into the cloud during its ascent. To investigate this, we set a liquid water threshold which defines which grid cells contain significant cloud water and separate the vertical kinetic energy into cloud and dry air components. This threshold is set at $h_{l,threshold} = 0.01$ in line with the value used in Böing et al. (2019) when calculating fractal dimension.

The time evolution of these energy components are shown in Figure 3.9 for all

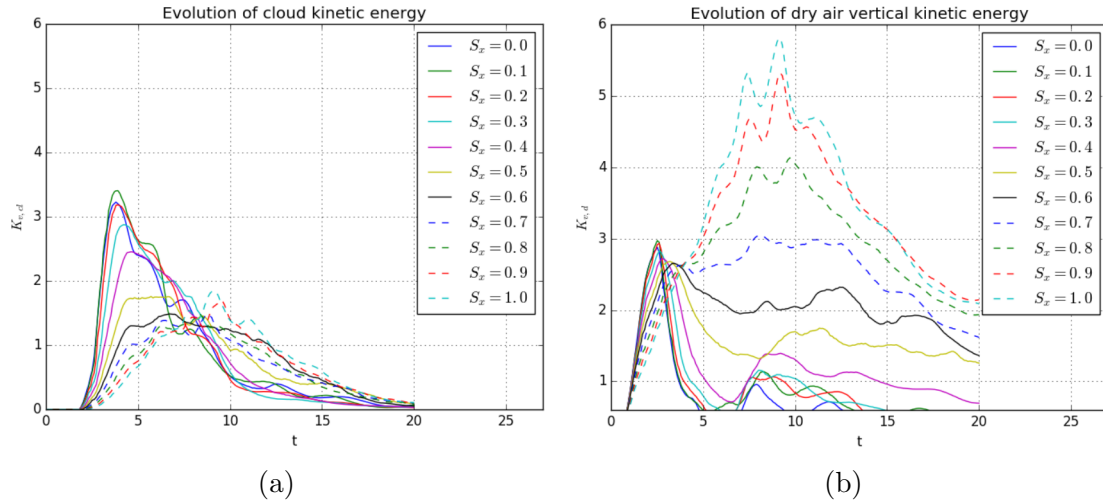


Figure 3.9: Time evolution of vertical kinetic energy for cloud air (a) and dry air (b) for shear values ranging from 0.0 to 1.0

runs. In the figure we see that, for low-mid shear values, the same initial increase occurs for dry air prior to the onset of condensation with a noticeable decline once condensation starts to occur. In the higher shear simulations, the dry air vertical KE is higher for the majority of the simulation. Also interesting is that the general shape of the cloud and dry air vertical KE curves are actually similar in the highest shear cases, even though the cloud values are significantly lower. Evidently, the vertical motions in high shear cases are dominated by the dry air which is disrupted by the thermal during its ascent. This serves to confirm that the vertical motions of cloud water are relatively weak for high shear, leading to the observed accumulation of liquid water at later times.

The differences in vertical KE between weak-shear and high-shear regimes noted in Figure 3.8 are equally apparent in the dry air vertical KE evolution. The transition here is also seen to occur at shear strengths in the $S_x = 0.6$ to $S_x = 0.7$ range. This implies that there is some link between the difference in the vertical KE evolution between the two regimes and the entrainment of dry air into the cloud at low levels, since the moist component of the kinetic energy is relatively low. However, work in Peters et al. (2019) suggests that the role of entrainment in inhibiting vertical motions is not as significant when compared to dynamic pressure

acceleration. They do also note, though, that entrainment may play a greater role in the dilution of buoyancy in the middle troposphere, which may be more comparable to the domain sizes considered here. An experiment with only low-level vertical shear would be an interesting means to investigate this further and analysis of dilution in our case studies would also present an interesting avenue of research. This is presently beyond the scope of this study, but developing methods to better quantify entrainment within particle methods would be highly beneficial.

3.3.4 Enstrophy

While the study of vertical kinetic energy has given a clear indication of significant entrainment in high vertical shear simulations, the total enstrophy of the system is generally a stronger indicator of turbulence. Hence, the enstrophy field should be expected to show a much more pronounced difference between dry air and cloud air in the case of significant entrainment. To investigate this, we integrate the squared vorticity magnitude over the entire domain to obtain a measure of total enstrophy,

$$\mathcal{E} = \int_V |\boldsymbol{\omega}|^2 dV. \quad (3.8)$$

We subtract the initial vorticity profile from $\boldsymbol{\omega}$ to get the vorticity field relative to background wind. As above, we separate this into cloud and dry air components, but this time we compute the ratio between the two terms. This will provide a measure of the turbulence present in the cloud itself compared to the environmental air, which can be used to illustrate the extent of entrainment due to wind shear. This is shown in Figure 3.10. We also show the values obtained for cloud and dry air separately in Figure 3.11.

This paints a very clear picture. In the case with no shear, a sharp peak occurs during the initial ascent of the thermal after the onset of condensation before decaying at later times. We note that this value remains above 1.0 until around $t = 15$ suggesting that the majority of the turbulence is due to the cloud itself. As the vertical shear is increased, however, this peak becomes less and less pronounced, suggesting that more dry air is becoming turbulent as a result of interactions with

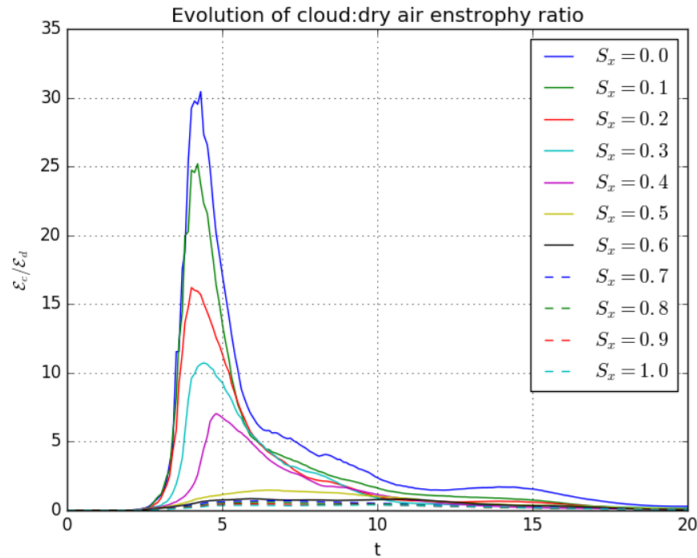


Figure 3.10: Time evolution of ratio between cloud and dry air enstrophy for shear values ranging from 0.0 to 1.0.

the rising thermal. And for the strongest shear values ($S_x \geq 0.7$), this ratio does not exceed 1.0 at any point in the simulation. The implication of this is that significantly more dry air is disturbed and entrained into the cloud during its ascent as the strength of the vertical shear increases. Physically, the entrainment of dry air is causing the cloud to spread out horizontally as dry air enters the thermal.

Looking at the values of enstrophy for cloud and dry air individually, we see that while there is some variation in cloud enstrophy with vertical shear strength, this does not account for the low ratios observed in higher shear cases (in fact some high shear cases actually show greater values for cloud enstrophy). Looking at the dry air kinetic energy confirms this: the observed decay in enstrophy ratio is largely due to a considerable growth in the dry air enstrophy with increasing vertical shear strength – significantly more dry air is entrained by the initial thermal during its ascent. This causes the thermal to rapidly spread out and disrupt more of the environmental air, causing it to rapidly become turbulent. This correlates with the observed effects on liquid water volume: if entrainment and wind shear are pulling

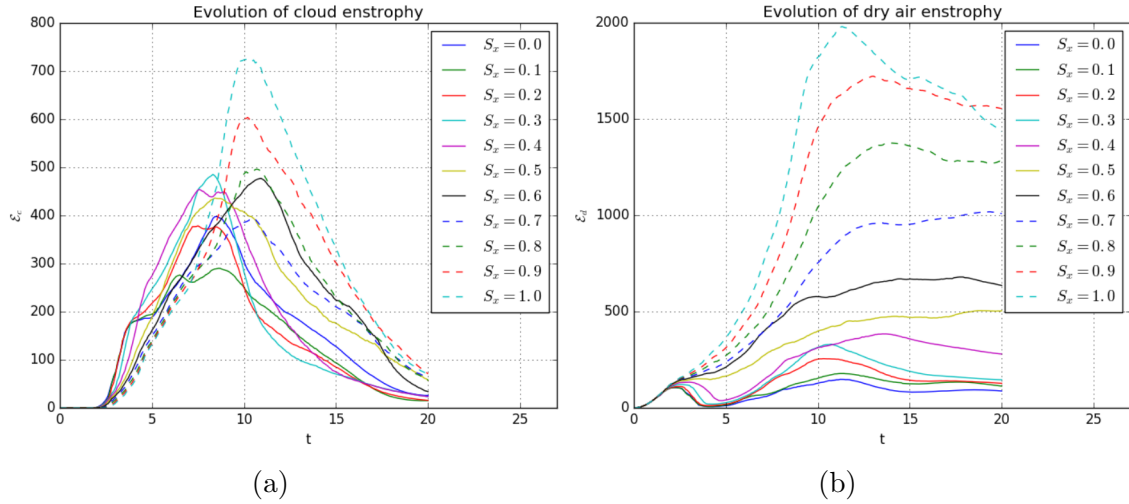


Figure 3.11: Time evolution of enstrophy for cloud air (a) and dry air (b) for shear values ranging from 0.0 to 1.0

the thermal apart as it rises initially, a smaller portion of the thermal will rise above the condensation level initially while the rest will lag behind. Disentangling the effects of shear and entrainment in MPIC is somewhat difficult, as we need to develop a means to quantify entrainment within the Lagrangian framework. This is significantly easier in grid-based Eulerian methods. Because entrainment and wind shear inhibit the vertical motion of cloud water, the cloud above the condensation level takes longer to move downwards than in low shear cases and the rest catches up, resulting in the slow accumulation of liquid water observed in Figure 3.6. The increasing half-depth with shear strength actually suggests that the cloud itself may be pulled apart by the vertical shear, becoming more fragmented and scattered across the upper region of the domain. It is also possible that a significant portion of the dry air enstrophy may be contained in the lower region of the domain (mixed layer) as the flow becomes turbulent during the thermal's ascent.

To further investigate this, we look at the x - z cross-sections of liquid water content and vorticity magnitude through the centre of the domain, as well as vertical profiles of liquid water volume and total enstrophy for each of the shear values. We only present images for $S_x = 0.3, 0.5, 0.7, 1.0$ to provide a representative range of low, mid and high shear examples. We attempt to capture the same stage




Figure 3.12: x - z cross-sections through the centre of the domain for liquid water (a, c, e, g) and vorticity magnitude (b, d, f, h) for simulations with shear values $S_x = 0.3$ (a, b), 0.5 (c, d), 0.7 (e, f) and 1.0 (g, h) at the time of peak cloud enstrophy.

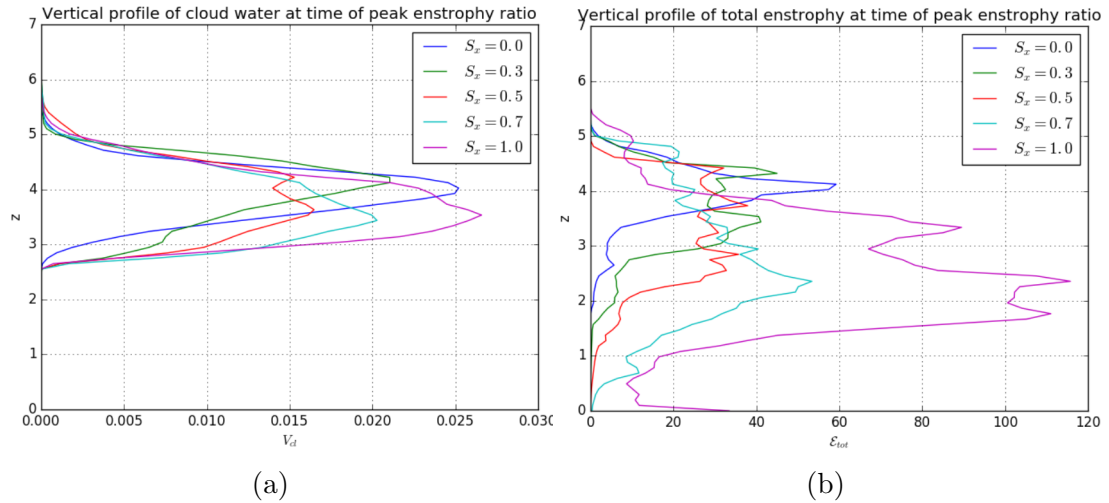


Figure 3.13: Vertical profiles of liquid water volume (a) and total enstrophy (b) for simulations with shear values $S_x = 0.0, 0.3, 0.5, 0.7, 1.0$ at the time of peak cloud enstrophy.

in cloud evolution for all of these, so rather than using the same fixed time, we identify the time at which the total cloud enstrophy is at its peak (when the cloud itself is at its most turbulent). For the aforementioned shear strengths, these times are $t_p(S_x = 0.0) = 8.5$, $t_p(S_x = 0.3) = 8.3$, $t_p(S_x = 0.5) = 8.4$, $t_p(S_x = 0.7) = 10.5$, $t_p(S_x = 1.0) = 10.1$.

The cross-sections are shown in Figure 3.8 and show roughly what we would expect to observe in this case. The clouds formed in low to intermediate shear simulations still appear largely concentrated whereas the higher shear simulations show significant fragmentation and spread across the upper region of the domain. This is consistent with the results noted in previous sections, as the greater vertical spread of the cloud in high shear cases was linked to significantly more vertical spread. Whether the highest shear cases still constitute an individual cloud is worth studying in further detail, perhaps through the use of liquid water probability distribution functions. At present, the visual difference still shows a clear structural deviation from the low shear cases.

Looking next at vorticity we can see that the low shear cases have very little

vorticity in the mixed layer below the cloud at this point in the simulation, compared to the high shear cases. This is somewhat expected, as the enstrophy ratio shows a lesser degree of entrainment in these simulations. In the high shear cross-sections we see significant turbulence in all regions except the highest points in the domain, where the cloud water has not reached. This is particularly revealing as it shows how much air in the mixed layer has been entrained by the cloud during its initial ascent. The cross-section for $S_x = 0.3$ shows clear eddies forming around the thermal as it ascends and displaces air in the stratified region. In the $S_x = 0.5$ run these eddies are still apparent, but there are now eddies forming in the lower levels of the domain, formed as the thermal entrained air during its ascent. As the shear is increased further, the disturbance in the low levels becomes increasingly pronounced and widespread and in the $S_x = 1.0$ cross-section we see that most of the domain below the cloud top is filled with turbulent eddies.

Mid-level shear causes the vortices associated with the cloud to become stretched horizontally compared to the no shear case. These observations are in line with Coniglio et al. (2006) who showed that strong mid level shear favoured rotation. The effects of high level shear in the simulations can also be observed through limiting the vertical extent of the liquid water and cloud vorticity. This is consistent with previous observations of the kinetic energy in Figures 3.9 and 3.10, as the late time kinetic energy of the system is largely dominated by the motions of dry air. It is worth noting that these are only cross-sections through the centre of the domain and still only provide a limited view of the full cloud structure. Further analysis of the vertical structure could be obtained through studying profiles of the horizontally averaged buoyancy field and its vertical flux and the evolution of both. $x - y$ cross-sections through key heights in the domain may also offer better insight into the horizontal distribution of liquid water. Alternatively, it may be worthwhile computing the difference between the buoyancy and the background profile or horizontal average to offer a better representation of the origin of air in the stratified region and identify entrainment. It may also be feasible to map key heights based on properties of the system, such as heights of peak enstrophy (or the cloud/dry air enstrophy ratio), cloud water and updraft velocity etc.

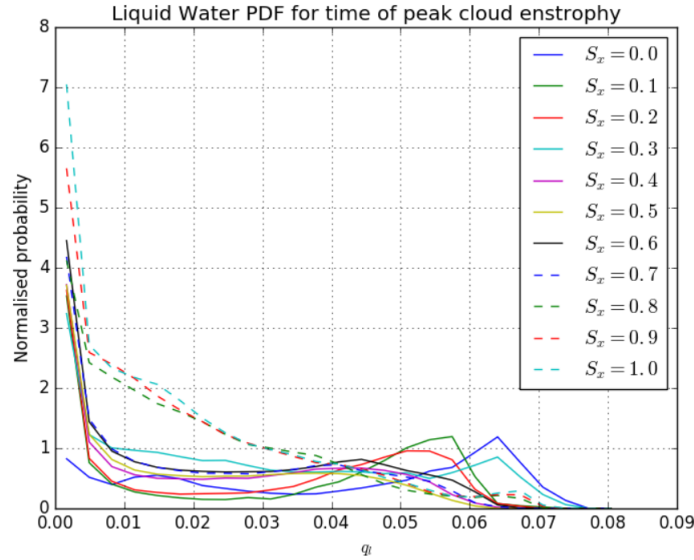


Figure 3.14: Liquid water probability distribution functions for shear values ranging from 0.0 to 1.0 calculated as in Böing et al. (2019). These use the parcel data at the save interval closest to the time of peak enstrophy for each simulation.

The vertical profiles of liquid water and total enstrophy in Figure 3.13 mostly serve as a companion to the above results, demonstrating the vertical structure of the resulting clouds and turbulence in the lower region of the domain. It reaffirms that the cloud is less spread out vertically in lower shear cases, but the peak of the cloud water distribution is located at a higher altitude. It appears some weak shear can result in more cloud water being concentrated at the base of the cloud, causing its mean height to drop as observed. Higher shear cases cause this effect to become even more dramatic, resulting in a cloud within which most of the water volume is concentrated in the cloud base, with very little reaching the upper regions. In this case, using the standard deviation of cloud height as a measure of the cloud depth may not be an accurate measure of how deep the convection is. While there is still a portion of cloud water at upper levels in the highest shear cases, it has very likely been torn apart by the shear in these levels, leaving small fragments in the upper region. In simulations with more sophisticated microphysics these would likely be removed by evaporation.

We also calculate the liquid water probability distribution functions using the same algorithm from Böing et al. (2019). These are shown in Figure 3.14. Because saving parcel data requires a lot of storage, the parcel save intervals are significantly less frequent than the gridded data ($\Delta t_{save,p} = 2.0$) so we used the parcel data closest to the time at which peak cloud enstrophy occurs for each simulation. The gridded field data is saved at intervals of ($\Delta t_{save,g} = 0.1$) and so offers a much greater temporal resolution. Reducing the discrepancy between these would also be beneficial, although saving parcel data frequently increases storage demands. It may be more reasonable to look at the same point in all simulations, but for consistency with the other comparisons drawn above, we attempt to stay as close to the time of peak cloud enstrophy as possible. Regardless, these still reflect the effects of vertical shear on the structure of the cloud.

In the no shear case, we see a peak in liquid water fraction just below $h_l = 0.07$, suggesting that the majority of the cloud volume is contained in a smaller number of parcels with higher liquid water content. As the shear value is increased, the effects on the curve are very clear: the concentration of parcels with high liquid water are reduced, causing this peak to flatten (and drift to the left); and the concentration of parcels with small liquid water fractions (or none at all) increases as the vertical shear is increased. The $S_z = 0.3$ case seems to behave differently, with the high liquid water peak occurring nearer the shear-free case. The lower liquid water bins are increasingly filled as the shear strength is increased. Furthermore, the highest shear cases ($S_x = 0.8, 0.9, 1.0$) appear to show an almost identical liquid water distribution and vary dramatically compared to the other cases. These results reaffirm that, especially in cases with higher vertical shear, the cloud is increasingly pulled apart, as it becomes spread out over a greater number of parcels. The liquid water distributions and their dependence on shear are particularly interesting when we also consider that the higher shear cases have a greater total cloud volume at the time of peak cloud enstrophy.

It is also worth noting that the leftmost values in Figure 3.14 actually occur below the threshold value of liquid water fraction used in the previous results,

although it is unclear whether the threshold should be adjusted to include these. A final point of interest is the behaviour of the curve for shear parameters of $S_x = 0.7$ and $S_x = 0.8$ where there appears to be a transition in the behaviour of the curve. The highest shear cases ($S_x = 0.8, 0.9, 1.0$) appear to show an almost identical liquid water distribution are vastly different compared to the other cases, further emphasising the division between low and high shear regimes.

At a glance though, our results do seem to suggest that low-mid strength vertical wind shear (particularly in the lower/mid levels) can actually produce a greater volume of liquid water and deeper clouds without significant fragmentation. A more detailed analysis of the structure may be necessary, through vertical profiles of heat and moisture fluxes, or studying buoyancy anomaly compared to the environment or other means as mentioned previously. Nonetheless, our results at higher shears show clear evidence that the vertical convective motions are being strongly inhibited, as seen in the vertical kinetic energies and cross sections of cloud water. The inhibition of convection due to vertical shear is documented in existing studies (Peters et al. 2019) and reproducing this effect, despite the differences in formulation between MPIC and conventional approaches like LES gives some validity to the method.

3.4 Conclusions

We have presented a simple model of cumulus convection subject to vertical wind shear, using a Lagrangian, particle-based method with the aim of providing some insight into the effects of different shear strengths on the development of an individual cloud in an idealised setting. The choice of a linear shear layer with a depth of 12 km may not be an accurate representation of typical atmospheric conditions, but it provides a starting point from which to study the effects of more complex shear profiles. One possible extension would be to investigate the effects of strong wind shear concentrated in different levels of the domain. Regardless, the results show that vertical wind shear can significantly influence the evolution of a shallow cumulus cloud, even in the absence of sophisticated microphysics. Improvement

to the microphysical scheme and potentially implementing precipitation may lead to increased organisation of convection in the system. The vertical shear causes updrafts and the resulting clouds to become slanted, and, as such, the highest region of the cloud, where condensation occurs more readily, will be located on one side of the cloud. The evaporation of rainwater on this side of the cloud will produce a cooling effect on the environment and may enhance the organisation of convection. In the absence of precipitation and evaporation, the entrainment of dry air is a very important factor in the growth and structure of a rising thermal. Some low-mid shear (particularly in the lower levels) will delay the initial ascent, but leads to a greater overall liquid water volume, with a larger fraction of this being at the cloud base, resulting in clouds with a much greater vertical extent. In the highest shear cases, liquid water slowly accumulates in the upper region of the domain, but is limited vertically.

The impacts of vertical shear even more apparent when looking at the kinetic energy and enstrophy of the cloud and dry air components of the domain. At lower shears, the enstrophy ratio between cloud air and dry air shows a very clear peak which becomes less and less pronounced as the shear is strengthened. We observe that the initial ascent of the thermal is hindered greatly in the highest shear cases, while the majority of the total system enstrophy is concentrated in the dry air being entrained. Looking at the structure of the cloud suggests that the shear and entrainment tear it apart in cases where the vertical shear is sufficiently strong. More detailed analysis is required here. Studies of fractal dimension and liquid water probability distribution functions or other measures of parcel dispersion may offer more insight into the effects of shear on individual cloud structure and convective initiation. This is of particular interest to the atmospheric community due to the importance of trigger function in the closure assumptions for parameterisation schemes.

Future work on this topic should endeavour to study the potential energy of the system to see how much of it is being used to move dry air, especially in cases where there is significant entrainment and to look at more physically realistic

shear profiles. An especially important step forward from the work presented here is to develop a means to quantify entrainment within MPIC to disentangle its influence on vertical motions from the wind shear itself. One potential means to do this is through the addition of a tracer population in the environmental air, and would take advantage of the natural addition of parcel attributes in MPIC. This would allow us to follow parcels of air from the mixed layer and their vertical displacement as a result of the thermal's ascent. Implementing more complex microphysics in the model, in particular precipitation and evaporation, would also be particularly interesting to investigate whether convective organisation occurs on the scales studied in these models. An anelastic framework would also be helpful here, given the vertical extent of the domain brings the validity of using the Boussinesq approximation into question. This framework would allow for a study of deep convection subject to vertical shear. One could even go further and add in horizontal wind shear, although it would need to be made consistent with the boundary conditions (either by choosing a periodic profile or modifying the horizontal boundaries) to study the effects of rotation.

The list of refinements and extensions that could be made to produce a more realistic setting is extensive. However, it should not detract from the benefits of the idealised test case presented. Even in a largely simplified setting, we have been able to demonstrate the effects of vertical wind shear on a developing cloud and the subsequent consequences for its bulk properties and structure. Our results at high shears are consistent with the limited existing literature on isolated thermals, although the increased cloud content at low shear is interesting. Further analysis of these test cases in a Lagrangian framework, incorporating tracers, for example, may offer a greater insight into potential sensitivities of convective systems to shear and entrainment, particularly in the lower troposphere, as was suggested in Peters et al. (2019).

Chapter 4

Energy budget analysis of a rising thermal in the Moist Parcel-In-Cell method

4.1 Introduction

Many factors contribute to the intensity of convection observed in the atmosphere and need to be accounted for in large-scale circulation models. Due to the turbulent nature of atmospheric convection and complex microphysics, the processes that primarily drive moist convection are highly nonlinear. For example, the potential energy released by condensation in clouds can often allow them to overcome the stratification of the atmosphere. This can cause deep convection, extending as far as the tropopause, often accompanied by intense rainfall and thunderstorms. If external conditions are favourable, deep convective cells can even organise and give rise to mesoscale convective systems (Houze 1997; Houze 2004). Other cloud microphysical processes, such as freezing and riming, also contribute to the energy budget of these systems (Zipser 1994; Zipser and Lutz 1994). Meanwhile, turbulence can lead to a dissipation of kinetic energy via molecular diffusion at the smallest scales (Kolmogorov 1941; Onsager 1949; Smagorinsky 1963). As such, understanding the energy transfers occurring within clouds is important for developing representations of convection at all scales.

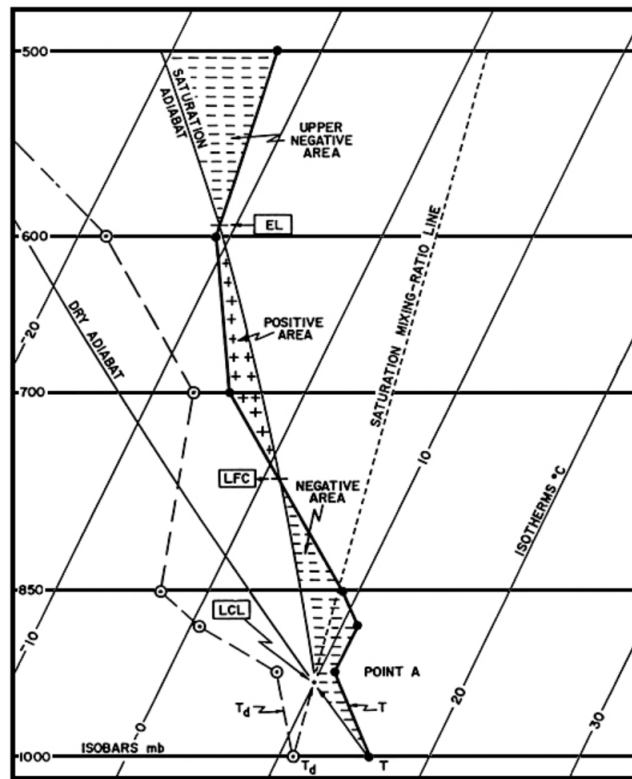


Figure 4.1: Skew-T plot showing key features. The positive area denotes a region of convective available potential energy (*CAPE*) and the lower negative area represents a region of convective inhibition (*CIN*). Isotherms are shown at 45 degree angles, hence the name skew-T. The lifting condensation level, LFC and equilibrium level are labelled and the moist and dry adiabats are shown alongside the sounding data (dark line with data points marked). Image source: NOAA - http://laps.noaa.gov/szoke/Korea/introconvective_stuff.html

The triggering of deep convection depends heavily on the energy available in the environment (Moncrieff and Miller 1976). This potential energy typically consists of two components, split between the gravitational potential energy due to the buoyancy of the air and a component due to latent heat released from the condensation of moisture. Studying the potential energy of a given environment can allow us to assess whether it is favourable for deep convection or will resist convective motions due to vertical stability. The data used to study the susceptibility of the atmosphere to convection comes from sounding studies and is often shown using a skew-T plot tracking the temperature of the background air with height. Skew-T plots show the air temperature variation with height and how it compares to the temperature of a dry or saturated parcel being raised adiabatically through the environment. These lines on the skew-T plot are called adiabats.

In Figure 4.1, there are two regions of particular interest. These are the areas between the temperature curve and the saturated adiabat (which describes the adiabatic cooling/heating of a saturated air parcel on a graph of temperature and pressure). Saturated refers to parcels of air in which water can exist both as a vapor and a liquid. We note that the line representing the air temperature intersects the saturated air adiabat at two particular points on the graph. The lowest of these represents the point at which ambient air temperature begins to decrease more rapidly than the moist adiabatic lapse rate of the parcel in question. A parcel raised to this point will then be able to rise through its surroundings further due to its higher potential temperature and lower density. This is known as the level of free convection (LFC). The second point of intersection represents the point at which the air parcel starts to become denser than the surrounding air again. The vertical structure of the environment inhibits convection above this point. This intersection is the parcel's equilibrium level.

The convective available potential energy (*CAPE*) is defined as the work done by the buoyancy force in raising a parcel from the free convection level to its equilibrium height. *CAPE* is the energy available to the parcel once it ascends into the region where free convection can occur and is a property of the environment

rather than a property of the parcel itself. Figure 4.1 illustrates this energy as the area between the temperature line and the saturated adiabat between the LFC and the lifting condensation level. It is important to note the sign here, as the *CAPE* must be positive for free convection. A negative value would actively hinder convection. Non-zero values of the *CAPE* are typically associated with the formation of storm systems (Craven and Brooks 2004), but this alone is not a sufficient condition for triggering deep convection. We can calculate the *CAPE* by integrating the difference between the temperature curve and the moist adiabat, and this gives the expression

$$CAPE = \int_{z_f}^{z_n} g \left(\frac{\theta_{v,parcel} - \theta_{v,env}}{\theta_{v,env}} \right) dz \quad (4.1)$$

where $\theta_{v,parcel}$ is the virtual temperature of the parcel of air, $\theta_{v,env}$ is the virtual temperature of the environmental background air, z_n is the level of neutral buoyancy of the parcel, and z_f is the level of free convection. The virtual temperature of an air parcel implicitly contains the effects of moisture, by representing a moist air parcel as an equivalent theoretical parcel of dry air.

The second region of interest is below the LFC, where the air parcel is denser than its surroundings. This typically occurs when a parcel of air is under a region of warmer air and thus has a lower density than its surroundings, resulting in a negative buoyancy. This difference in density actively suppresses convective motions and can prevent the parcel from reaching the LFC to begin free convection under the buoyancy force. The negative area between the temperature curve and saturated adiabat is known as convective inhibition (*CIN*).

The *CAPE* is defined as the work done against the gravitational/buoyancy force in lifting a parcel from the bottom of a system to the level of free convection. As such, *CIN* can be calculated similarly to *CAPE*:

$$CIN = \int_{z_{bot}}^{z_f} g \left(\frac{\theta_{v,parcel} - \theta_{v,env}}{\theta_{v,env}} \right) dz \quad (4.2)$$

here z_{bot} is the lower boundary of the domain (surface level). These two quantities work in opposition to one another. For deep convection to occur, the *CIN* of the environment must be overcome either through mechanical forcing such as orographically forced clouds or surface heating, giving the air sufficient energy to rise to the lifting condensation level. Once air rises to the condensation level, the phase change of water vapour releases additional heat. Deep convection can occur if this additional energy release is sufficient to allow the parcel to ascend above the LFC. *CAPE* can, in some cases, be a much more significant factor than *CIN* in determining the intensity and distribution of precipitation (Myoung 2010).

Consequently, the difference in altitude between the LFC and lifting condensation level becomes a critical factor in the formation of storm systems. As the gap between these two levels narrows, the *CAPE* of the environment decreases and deep convection occurs much more readily. One consequence of this occurs in the formation of tornadoes. Sounding data from significant tornadoes suggests that they primarily form where there is a relatively high condensation level (≥ 1300 m) and the LFC height is comparatively low (≤ 2000 m) (Rasmussen and Blanchard 1998; Thompson et al. 2003; Craven and Brooks 2004). This means that the latent heat release due to condensation allows the air to easily overcome the shallow layer of *CIN* and reach the LFC to access the *CAPE* above. Conversely, a large difference between the lifting condensation level and LFC results in a large region of *CIN* that can prevent deep convection.

The environmental *CIN* strongly influences the triggering of deep convection (Emanuel et al. 1994; Mapes 2000). When representing convection in global scale models via parameterisation schemes, the *CIN* sometimes appears as part of the trigger condition for convection in a grid cell. For example, Rochetin et al. (2014) used a stochastic distribution of thermals with varying updraft speeds. The triggering condition in their scheme depends on the updraft speed of the strongest thermal in the distribution and if it is sufficient to overcome the *CIN* of the environment. Other effects can also encourage environmental air to push through the inhibition layer, such as orographic forcing due to terrain. Studies have suggested

that when convection is triggered this way, where the *CIN* is reduced through external means, the resulting convection is more intense than if no *CIN* were initially present. The build-up of heat in the lower layer likely enhances convection once the air can rise to the LFC. This is why intense rainfall often occurs near mountain ranges where the air blows inland from the sea, and is forced upwards over the terrain.

Similarly, the environmental *CAPE* is of utmost importance in global weather prediction and climate models. Cumulonimbus clouds and thunderstorms transfer *CAPE* into the upper atmosphere, regulating the temperature balance between the upper and lower troposphere, and often occur on rapid timescales (Raymond and Herman 2011). Since *CAPE* indicates an instability to convection in the environment, its build-up is often indicative of severe weather events. Therefore, climate models need to pay particular attention to how the *CAPE* (and *CIN*) respond to rising sea surface temperatures and other factors (Chen et al. 2020). Similarly, the vertical distribution of *CAPE* in the atmosphere is also important since low-level *CAPE* can encourage convection.

CAPE and *CIN* are widely used in global models to diagnose whether domain regions are unstable to convection. However, studies on the scale of individual clouds are important too. Analysis of the energetics of clouds during their formation and development can help improve our understanding of cloud processes. Using smaller-scale models such as LES to feed into large-scale global models to develop more sophisticated parametrisations schemes is not uncommon in the field (i.e. Siebesma and Cuijpers 1995; Angevine et al. 2018 and Zhang et al. 2021). As such, it is beneficial for any convection resolving model or large-eddy simulation to estimate the energy available for convection in the system. This is especially relevant if such a model is to be considered for use in an embedded or superparameterisation scheme. Analysis of the energy transfers during cloud formation and development can also be beneficial in evaluating microphysical schemes and their implications for cloud development in more unpredictable regions (Rybka et al. 2021).

Furthermore, some convective schemes may also make use of *CAPE* as part of their closure assumption (Nordeng 1994; Neggers et al. 2004). These closure assumptions differ from more conventional quasi-equilibrium approximations but are not typically favoured in current representations of convection. Again, evaluating these assumptions via LES studies is a crucial component of refining convection schemes in global models. Being able to reliably observe and model the evolution of potential energy within a convection model is of great relevance to our understanding the triggering of deep convective systems. Because *CAPE* is directly related to deep convection, parameterisation schemes attempting to model the transition between shallow and deep convection often use explicit simulations or observations of low-*CAPE* environments to evaluate their performance (Yano et al. 2013; Baba 2020). Analysis of the energy released by clouds during convection is of interest when considering their representation on all scales.

To this end, we aim to derive quantities within a Lagrangian convection model, the Moist Parcel-In-Cell model of Dritschel et al. (2018), that will allow us to analyse the potential energy available to the system. These are intended to tell us about the transfer between potential and kinetic energy within the model and the energy released by condensation, which previous MPIC studies have not yet investigated. In future works, the energetics within MPIC may be able to feed into larger-scale models and parameterisation schemes, similarly to how LES models are often used in the atmospheric community. The Lagrangian formulation of MPIC allows these quantities to be calculated for individual parcels instead of on grid cells. While *CAPE* is an integrated quantity and is fixed for a given environment, the potential energies calculated in this chapter vary for parcels of air as they rise through the domain. As such, the parcel energetics may offer greater insight into the distribution of potential energy within the cloud during development in future works. For example, we could map the available potential energy within a cloud, or implement a tracer population to assess how the energy of entrained air varies during convection. Calculating the energy on the parcels also produces a more accurate estimation of the energies present in the system since no additional interpolation steps are required.

We apply this energy budget analysis to a variation of the standard MPIC test case (see Chapter 3) and study the balance of potential and kinetic energy. We also assess how the parcel energetics respond to varying numerical parameters used in MPIC, particularly the grid resolution and the parcel splitting threshold. Section 4.2 describes the derivation of the potential energy in MPIC and presents a comparison to environmental *CAPE* and section 4.3 presents our results. In section 4.4 we discuss the broader implications of this study and briefly discuss the impact of a new scheme to ensure incompressibility in MPIC. Finally, in section 4.5, we summarise our findings and discuss possible extensions of this work.

4.2 Parcel Energetics in MPIC

4.2.1 Derivation of Energy Components

We start from the governing equations in MPIC as presented in Chapter 2 (Dritschel et al. 2018; Böing et al. 2019), which are as follows:

$$\frac{D\mathbf{u}}{Dt} = -\frac{\nabla p}{\rho_0} + b\hat{\mathbf{e}}_z, \quad (4.3)$$

$$\frac{Db_l}{Dt} = 0, \quad (4.4)$$

$$\frac{Dq}{Dt} = 0, \quad (4.5)$$

$$\nabla \cdot \mathbf{u} = 0, \quad (4.6)$$

with the equation of state being given by,

$$b = b_l + \frac{gL}{c_p\theta_{l0}}q_l, \quad (4.7)$$

where

$$q_l = \max(0, q - q_0e^{-\lambda z}), \quad (4.8)$$

is the liquid water content. More details of the vorticity dynamics approach used to solve these equations for parcels are provided in Chapter 2 on the MPIC method.

We intend to derive an expression for the system's total energy, divided into kinetic and potential energy components. To this end, we follow the work of Vallis (2017), who describes the energetics for a Boussinesq system with a general equation of state. We use this to derive expressions for an air parcel's kinetic and potential energy within MPIC.

First, we look at the kinetic energy distribution, which is relatively straightforward to calculate. We first define the kinetic energy density, K , as

$$K = \frac{1}{2} |\mathbf{u}|^2. \quad (4.9)$$

The evolution of this component, as per Vallis (2017), is given by

$$\frac{DK}{Dt} = bw - \nabla \cdot (\phi \mathbf{u}), \quad (4.10)$$

where w is the vertical component of the parcel velocity, b is the total parcel buoyancy and $\phi = p/\rho_0$. Note that p here is the non-hydrostatic pressure, equivalent to δp in Vallis (2017)).

Next, we must define the potential energy from the equation of state. In Vallis (2017), it is shown that in order to conserve energy, the equation of state must be a function of the vertical coordinate, i.e. of the form $b = b(b_l, q, z)$, where b_l and q are state variables. In MPIC, we use the liquid water buoyancy b_l and specific humidity q as state variables. The procedure is essentially the same. We define the potential energy density, Π , at vertical position z as

$$\Pi(b_l, q, z) = - \int_a^z b dz', \quad (4.11)$$

where a is a constant reference height that may be chosen freely. The choice of a is arbitrary but determines where the zero-point in available potential energy occurs within the system. We could, for example, use the condensation level, defined as

$$z_c = -\frac{1}{\lambda} \ln \left(\frac{q}{q_0} \right), \quad (4.12)$$

although the specific humidity determines this and thus it is different for each parcel. This would also lead to a negative available potential energy should a parcel rise above z_c . Another option is to set the reference height at the top or bottom of our domain. While the choice of domain height does not have any physical significance in the MPIC test case, it does represent the highest point a parcel can reach within the simulation. The boundary is rigid, but later developments may seek to add a ‘sponge layer’ or damping effect close to the boundary. For now, this seems a sensible point to set the zero of available potential energy.

First, we consider the potential energy density, with $a > z_c$, such that the reference height is above the condensation level. So our integration is now:

$$\Pi(b_l, q, z) = - \int_a^z b dz' = \int_z^a b dz'. \quad (4.13)$$

All parcels will have a component due to moisture in their available potential energy, which does not contribute when integrating from the lower boundary for a parcel below z_c . This is not strictly true for drier parcels in the upper regions of the domain. Sometimes, the condensation height can exceed the domain height and will approach infinity as the specific humidity approaches zero, due to the dependence on $\ln q/q_0$ in equation 4.12. In the case of such parcels, the moist component of their potential energy will be unavailable at any point in the domain, so we can neglect this term. We do this by setting their condensation height z_c equal to the upper boundary of the domain since the potential energy density must be continuous across z_c . This will cause any condensation terms to vanish in the integral as they will only contribute for heights in the range $z_c < z < z_{top}$. Carrying out the above integral, we obtain two expressions for parcels above and below their condensation level,

$$\Pi(b_l, q, z) = b_l(a - z) + \frac{gL}{c_p\theta_{l0}} \left[q(a - z) + \frac{q_0}{\lambda} (e^{-a} - e^{-z}) \right], \quad q > q_0 e^{-\lambda z}, \quad (4.14)$$

and

$$\Pi(b_l, q, z) = b_l(a - z) + \frac{gL}{c_p\theta_{l0}} \left[q(a - z_c) + \frac{q_0}{\lambda} (e^{-a} - e^{-z_c}) \right], \quad q < q_0 e^{-\lambda z}, \quad (4.15)$$

where only the dry component varies with z . For dry parcels where z_c has been reset to z_{top} , the moist component (second term on the right-hand side) vanishes for the choice $a = z_{top}$. For the rest of this chapter, we will set $a = z_{top}$, thus assigning the zero in potential energy to the domain top.

Taking the material derivative, as in Vallis (2017), we obtain

$$\frac{D\Pi}{Dt} = -bw, \quad (4.16)$$

which can be combined with the expression for kinetic energy density above,

$$\frac{\partial}{\partial t} (K + \Pi) + \nabla \cdot [\mathbf{u} (K + \Pi + \phi)] = 0. \quad (4.17)$$

The second term in this expression is a flux term and vanishes when integrated over the entire domain. Because we assume density variations in the flow are small (Boussinesq approximation), all parcels have approximately the same density, ρ_0 . Hence, to obtain the system's total energy, we compute a volume-weighted sum of K and Π over all parcels. This value should remain approximately constant in the absence of any external fluxes of heat or moisture, ignoring the effects of numerical dissipation and parcel mixing at later times.

4.2.2 Physical model

We will use the expressions derived above to analyse the behaviour of kinetic and potential energy components for the chosen test case. This should allow us to

study the evolution of the energies involved in cloud development and also provide an opportunity to assess how effectively total energy is conserved in the model. We can also study the enstrophy of the system and the mixing of parcels to observe the effects these have on energy conservation. Our test case is identical to that shown in Chapter 3, where the response of the rising thermal to vertical wind shear was analysed, albeit with the shear removed.

4.2.3 Comparison to *CAPE*

As seen above, the choice of reference height determines the vertical position at which the potential energy of a given air parcel will be zero. In the later sections showing the energetics, we primarily look at the change in the potential energy since this should be similar to the change in *CAPE* of a parcel with height. There are several differences between the potential energy derived above and *CAPE*. In particular, *CAPE* is an integrated quantity and thus fixed for a given environment and only includes the effects of water vapour in the environment. Liquid water already present does not contribute to the *CAPE*. What we have computed instead is effectively a profile version of *CAPE* that can be computed on a given parcel of air in the simulation and varies as the system evolves.

We take a set of three test parcels using typical values observed in the standard MPIC test case for the mixed layer ($b_{l,ml} = 0.0$, $q_{ml} = 0.074$), the centre of the thermal ($b_{l,th} = 1.52$, $q_{th} = 0.082$), and near the top of the domain ($b_{l,top} = 3.78$, $q_{top} = 0.0001$). We plot the evolution of the available potential energy of each parcel as they rise through the domain. Initially, we set a to be at the domain's top.

In Figure 4.2, we can see that the potential energy of each parcel reaches a value of 0 at the top of the domain due to the choice of a . Each initially starts as a straight line (with slope b_l for the respective parcel type) before curving due to the effects of condensation once the parcel crosses its respective lifting condensation level, except for the upper layer parcel for which z_c occurs at the top of the domain. In the case of mixed-layer parcels, the available potential energy

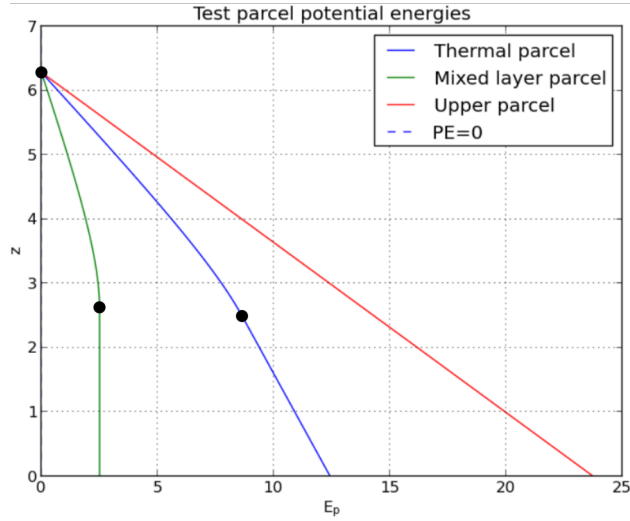


Figure 4.2: Plot of the variation in potential energy with vertical position for three test parcels using typical values of specific humidity and buoyancy from the thermal, mixed layer and upper region of the domain. Green: $b_{l,ml} = 0.0, q_{ml} = 0.074$; blue: $b_{l,th} = 1.52, q_{th} = 0.082$; red: $b_{l,top} = 3.78, q_{top} = 0.0001$. Black markers are used to denote the approximate location of the condensation height for each test parcel.

does not change until condensation occurs since the mixed layer in the test case has zero buoyancy. The condensation height for parcels in the mixed layer and the thermal are relatively similar.

In reality, the potential energy of a parcel is zero at its associated equilibrium level (if all parcels in a given environment are at their equilibrium level, the *CAPE* is also zero). It thus makes more sense for us to use this as our upper limit of integration in the above derivation, such that we have a region of positive potential below, where the parcel rises freely, and a region above, where the negative potential energy inhibits convection. This follows the structure of *CAPE* and *CIN* in real-world observations. To do so, we may construct something comparable to the skew-T (Figure 4.1) shown in Section 4.1 for the environment in MPIC. Figure 4.3 shows a similar plot to the skew-T, however we have replaced temperature with buoyancy as the thermodynamic variable and the vertical coordinate, z , is used in place of the pressure.

In MPIC, the potential temperature of parcels is included in the buoyancy, which means that we cannot directly plot the temperature. The buoyancy of a parcel at any given point is related to its virtual potential temperature by $b = g\theta_v/\theta_0$, which is the same for all parcels. Hence, we can compare the parcel's total buoyancy and the environment's vertical structure to determine where the LFC and equilibrium level will lie for the parcel in question. We use the initial background buoyancy profile described in the standard MPIC test case, Dritschel et al. (2018) and Chapter 3 of this thesis, as our reference profile. Here we have $b_{l,bg} = 0.0$ below z_b where z_b is the top of the mixed layer and $b_{l,bg} = N^2(z - z_b)$ in the stratified region above, where $N^2 = 0.97048$ is the squared buoyancy frequency.

We can then take the same test parcels and use equation 4.7 to compute their total buoyancy as they rise through the domain. This allows us to identify the regions in which they are more or less buoyant than their surroundings, along with the points of intersection (LFC and equilibrium level) to serve as the limits of integration for analogues to *CAPE* and *CIN*.

Mixed layer parcels do not ever cross the line corresponding to the buoyancy of the reference profile, although the two lines are identical below the top of the mixed layer. This is by design, as the background for the MPIC test case was intended to be stable to convection so the study could focus on the thermal's ascent. If the environment were unstable, we would see dry air being lifted by the thermal above the free convection level (which may make for an interesting study in itself). Consequently, this test case has no LFC for the mixed-layer parcels. If they rise above the mixed layer, they will always be less buoyant than the surrounding air and forced downwards under gravity. The equilibrium level of the mixed layer parcel is z_b at the top of the mixed layer, but it technically occupies a range between the surface and this point since the buoyancy in the mixed layer is constant.

If moved down from the upper layer, the upper-level parcel will almost always be more buoyant than its surroundings and be forced upwards towards the equilibrium level (approximately $z_n = 2\pi$). There is little else to say here, but if the domain

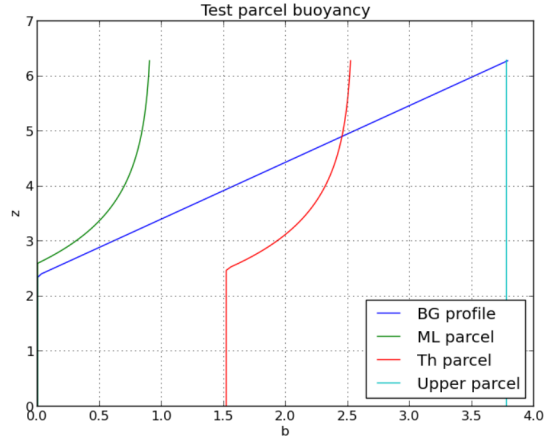


Figure 4.3: Plot showing the variation in buoyancy for the three test parcels. Green: $b_{l,ml} = 0.0, q_{ml} = 0.074$; red: $b_{l,th} = 1.52, q_{th} = 0.082$; cyan: $b_{l,top} = 3.78, q_{top} = 0.0001$. Also shown is the background environmental profile (blue). The points at which the buoyancy curves intersect with the environmental buoyancy profile indicates either the lifting condensation level (if approaching from the left, i.e. lower buoyancy) or equilibrium level of the parcel (if approaching from the right, i.e. higher buoyancy).

were extended above this point, the potential energy would become negative when raising the parcel further. This is the same for any environmental air parcels displaced during the ascent of the thermal.

In contrast to this, the parcels in the thermal studied in this chapter are always more buoyant than the environmental air in the mixed layer until a point at an altitude of around $z = 5.0$. At this point, they then become less buoyant than the surrounding air, suggesting this is the level of neutral buoyancy. Again, this is by design as the initial MPIC tests specify the levels of dry and moist neutral buoyancy for the initial thermal in setting up the test case ($z_d = 4.0$ and $z_m = 5.0$) and the buoyancy frequency N^2 is determined using these values. As such, the level of free convection for parcels in the initial thermal is effectively the lower surface of the domain. In physical terms, the thermal has already been heated sufficiently to overcome the *CIN* and thus experiences free convection until it reaches the equilibrium level. However, we can discuss hypothetical test problems that could

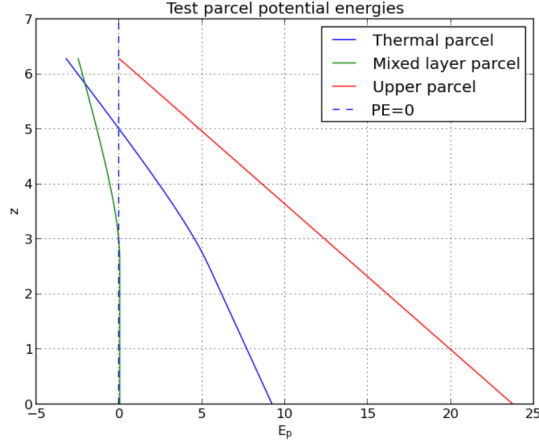


Figure 4.4: Plot showing the variation of potential energy with height for the three test parcels, Green: $b_{l,ml} = 0.0, q_{ml} = 0.074$; blue $b_{l,th} = 1.52, q_{th} = 0.082$; red: $b_{l,top} = 3.78, q_{top} = 0.0001$. The blue dashed line indicates the x -axis ($PE = 0$) This time the equilibrium levels used in potential energy calculation are determined from identifying the point at which they match the background buoyancy profile as detailed in this section. Note that the overall shape of the curves remain unchanged, but the x -intercept (zero potential energy) occurs at the equilibrium level.

be studied in future works with MPIC in which CIN must be overcome.

In this case, to estimate the equilibrium level of an air parcel, we could also solve the equation $b_{tot}(z) = b_{l,bg}(z)$ for z to obtain the points at which the parcel buoyancy matches the reference profile (background) buoyancy. So we solve

$$b = b_l + \frac{gL}{c_p \theta_{l0}} \max(0, q - q_0 e^{-\lambda z}) = N^2 (z - z_b), \quad (4.18)$$

for each parcel at a given time. We revisit our test parcel energies and use the values of z_n obtained by solving 4.18 to set the zero points for available potential energy. The resulting curves are shown in Figure 4.4.

Another possibility would be to adjust the structure of the mixed layer slightly. While the region should maintain an approximately constant potential temperature (included in the buoyancy in MPIC), the moisture content of the air in the mixed




Figure 4.5: Variation of total buoyancy for a mixed layer parcel, $b_{l,ml} = 0.0$, $q_{ml} = 0.07$ (green) and the background buoyancy profile described in Section 4.2.3 with $z_b = 2.0$, $N^2 = 0.27$ (blue) to illustrate where the LFC and equilibrium level for a mixed layer parcel would lie in such a setup.

layer is typically non-homogenous, and it may be worth considering a height-dependent profile in this region. Adjusting the humidity of the mixed layer would affect the stability of the environment to convection. With some changes to the vertical structure, we could design a vertical structure that more closely mimics natural convection, in which a region of *CIN* lies below the LFC with a region of *CAPE* above it. In Figure 4.5, we have lowered the top of the mixed layer to $z = 2.0$ and adjusted the strength of the stratification in the upper region of the domain, reducing the buoyancy frequency to approximately $N^2 = d\bar{b}/dz = 0.27$. This adjustment makes the domain less stable to convection. We then plot the vertical buoyancy profile of a test parcel in the mixed layer. While these may be arbitrary adjustments, they illustrate that we can alter the vertical structure of the MPIC setup to create an environment that is conditionally unstable to convection without the need to generate an initial thermal.

Through this, we can identify regions of *CIN* and *CAPE* for such an environmental profile and the level of free convection ($z = 2.9$) and equilibrium level ($z = 5.1$). There now exists a region of *CIN* that will push the parcel back down into the mixed layer, even if condensation can occur. In this case, we expect to see

shallow cumulus clouds forming below the LFC due to the position of the condensation level. If these parcels receive enough energy from external sources to ascend above the LFC, they will rise freely to the equilibrium level, which becomes the nominal cloud top.

While not used for the isolated thermal simulations in the rest of this chapter, the implications of such a setup are more interesting if we allow for a source of heat and moisture at the lower surface of the domain. This could raise the buoyancy of the parcels in the mixed layer to the point where they can reach the LFC and initiate deeper convection. By adding a simple mechanism for cooling and drying, we can observe the environment's evolution and investigate how strong such a heat source would need to be to allow for deep convection in MPIC. Moreover, external factors, such as vertical wind shear, could be investigated in terms of whether they reduce or increase the required heating to initiate deep convection in such a setup. This would open up many possibilities for exploring the evolution of convective systems without an initial thermal in MPIC, moving away from the more idealised test problems shown in this chapter and chapter 3. That said, it may also be worthwhile to study what would happen in the thermal test case if the environment were adjusted similarly.

4.3 Results

The simulations in this chapter are allowed to evolve to a non-dimensional time of $t = 20$ and the energy components and gridded properties are saved at intervals of $\Delta t_{save} = 0.1$ with a maximum timestep of $\Delta t = 0.05$. Unless stated otherwise, all results use a grid resolution of 64^3 . These settings are the same as was used in chapter 3. Quantities calculated in the subsequent analysis are calculated using the values on parcels and not the gridded fields unless specified. Energies calculated refer to the change in that component of the total energy relative to its initial value at $t = 0.0$, unless stated otherwise.

4.3.1 Thermal Energetics

Here we present an analysis of the energetics for the thermal test case detailed in 3.2, in the absence of wind shear. We have separated the total available potential energy into two components, corresponding to the dry potential energy and the latent heat released by condensation. The energy components in subsequent analysis are defined as follows. E_T is the total energy, E_K is the kinetic energy, E_P is the total potential energy including moisture, E_D is the potential energy not including moisture and E_L is the latent heating component of potential energy. Note that E_k and E_P are the volume integrals of the energy densities defined in section 4.2.1, K and Π , respectively. All subsequent analysis referring to these components will refer to the differences between their values at a given time and their value at $t = 0$, unless explicitly stated otherwise. We also added a curve to represent the sum of the change in kinetic and potential energy to visualise the extent to which the model conserves the total energy.

In Figure 4.6 we can see that, at the earliest times, the total system energy is conserved, as the change in total potential energy mirrors the kinetic energy evolution (note $E_k(t = 0) = 0$ here since the system is initially at rest). This plot also allows us to see when condensation begins to occur in the model, denoted by the latent PE component (difference between total PE and dry component) becoming non-zero. At later times, the total energy starts to decay, with the loss in energy eventually becoming comparable to the kinetic energy.

A potentially concerning feature of these results is that the total energy does not decrease monotonically. There are times at which the total energy rises. In particular, during the early stages of the thermal's ascent, the total energy appears to grow substantially until around $t = 5$ to $t = 6$, after which it begins to level out and decay. This is somewhat contradictory to expectations in a simulation of three-dimensional turbulence. Choosing appropriate boundary conditions and integrating equation 4.17 over the domain would also offer some insight into global conservation of energy in MPIC and this is worth investigating in future works. A likely cause of this is that MPIC, in its current state, does not enforce incompress-

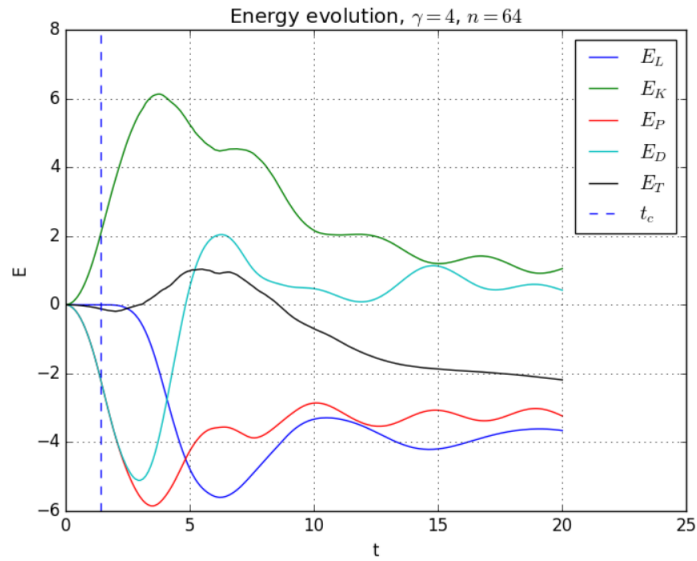


Figure 4.6: Plot showing the time evolution of the difference in each energy component and their initial values. E_T is the total energy, E_P is the total potential energy including moisture, E_D is the potential energy not including moisture and E_L is the latent heating component of potential energy. γ is the current splitting threshold for parcels and n denotes the resolution, 64^3 . The dashed line, t_c , shows the point at which condensation begins to occur in the simulation. Note that these energy values here are the change in each component from their initial values, as stated previously.

ibility. A fix for this has since been developed in two dimensions in an alternative model version (Frey et al. 2022). However, the work in this chapter was carried out long before its three-dimensional implementation. As such, we acknowledge this as a significant limitation of the study but will briefly discuss the impacts of the incompressibility corrections in a later section.

4.3.2 Resolution Sensitivity and grid vs parcel comparison

One of the first things we want to test in studying the energetics of cloud development is the effect of numerical diffusion within the model, particularly at early times before the flow has become highly turbulent. To this end, we vary the resolution of the underlying grid in MPIC and run the same test case at grid resolutions of 32^3 , 48^3 , 64^3 and 96^3 in order to assess the effects of resolution on energy conservation. It is worth adding that the effective resolution achieved by the parcels in MPIC is greater than that of the grid used for inversion since the minimum parcel volume is $1/6^3$ of the grid cell volume (Dritschel et al. 2018).

In the early stages of the thermal's ascent in Figure 4.6, we see that the system's total energy is converging towards the horizontal axis. As the system evolves, however, we observe a significant decay in total energy. At later stages in the development of the cloud, the flow has become highly turbulent, so we expect the total energy to decrease as a result. Moreover, as the flow becomes increasingly turbulent, the small-scale variations that may occur due to resolution have an increasing effect on the flow. It is also worth noting that the vorticity field in MPIC exhibits some sensitivity to resolution (Dritschel et al. 2018) and as such higher resolutions may result in increased turbulence at the smallest scales, although the effect on the decay of total energy is not totally clear.

We can also decompose the energy into each component (kinetic, dry potential, latent heat) to study their sensitivity to changes in grid resolution, and we show this in Figure 4.7. Starting with the system's kinetic energy, we notice that the peak values of kinetic energy increase as the resolution is improved. At later times, we note some degree of convergence with resolution. What stands out

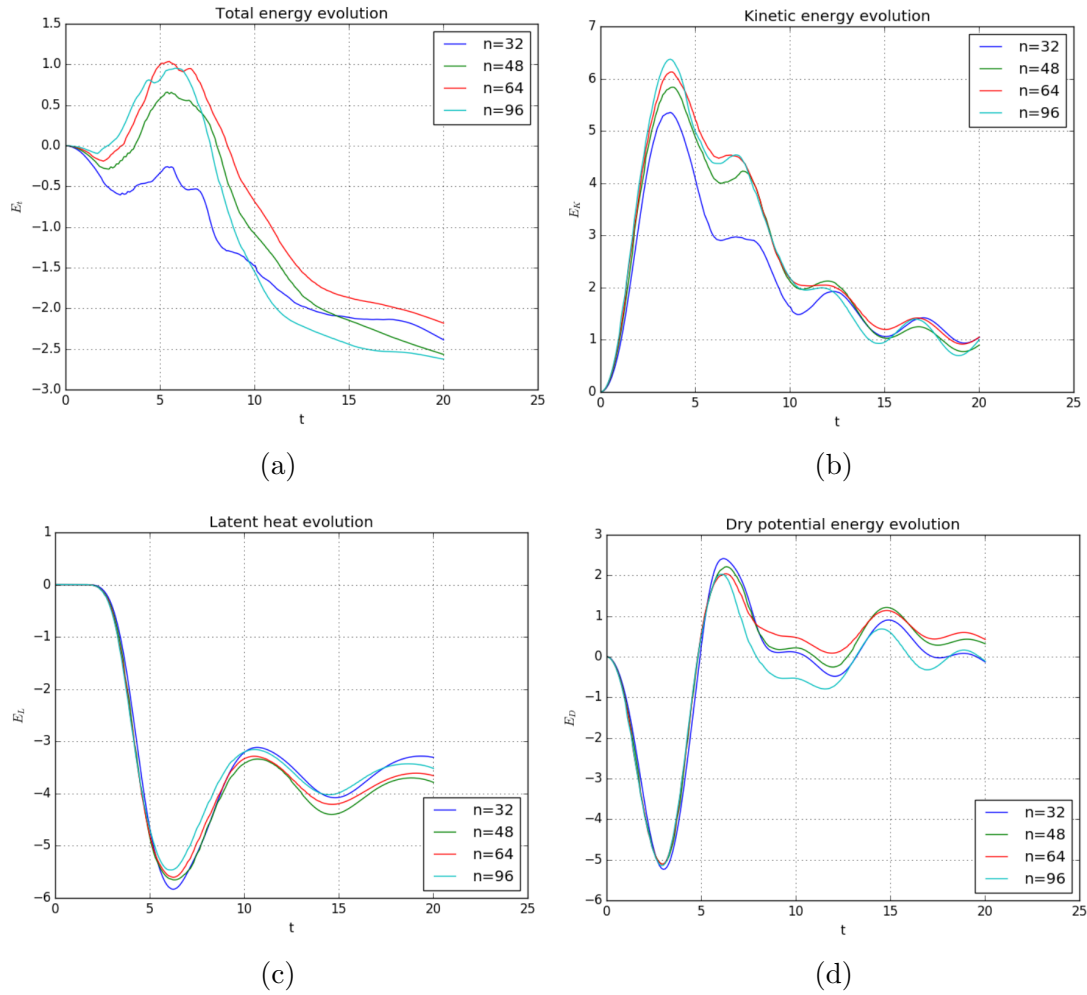


Figure 4.7: Evolution of the difference in total energy (a), kinetic energy (b), latent heat potential energy (c) and dry potential energy (d) for grid resolutions of $n = 32, 48, 64, 96$.

is that, even though the kinetic energy grows more rapidly during the ascent of the thermal, it also decays more rapidly at later times, eventually seeing the 96^3 and 64^3 runs decrease to comparable and occasionally smaller values than the run on a 32^3 grid. There are a few possible causes for this behaviour, the first being the resolution dependence of the vorticity field in MPIC mentioned above. The fact that kinetic energy increases and decreases more rapidly appears to be consistent with the idea that there is more turbulence at smaller scales due to the higher vorticity values on the parcels. It is also possible that a grid resolution of 32^3 is too coarse to fully capture the behaviour in regions where the lengthscales involved become very small, for example, at cloud edges where there is mixing between the cloud and surrounding air. Another possibility is related to the previously mentioned concerns regarding incompressibility. Since the number of parcels present is proportional to the number of grid cells used, we would expect higher resolutions to show variations in the density field that would influence the calculation of the energy components.

The two potential energy components also appear to exhibit a convergence with resolution at early to intermediate times (in fact, they are nearly indistinguishable until after $t = 5$). Towards the end, we see more significant variation in the latent heating component. Whether this is a consequence of turbulence and increased mixing in higher resolution simulations is currently unclear. That aside, these results still show that the derived expressions capture the system's behaviour better as the resolution is increased prior to extensive mixing and strong turbulence.

At this stage, it is instructive to compare the values of the kinetic energy of the fluid when integrating over the grid cells and parcels, respectively. This is shown in Figure 4.8 for reference, although it demonstrates that there is not a substantial difference between the two. Since the velocity field in MPIC is computed on the grid initially and then interpolated to the parcel positions, it is perhaps unsurprising that there is no substantial difference. This result is included for completeness.

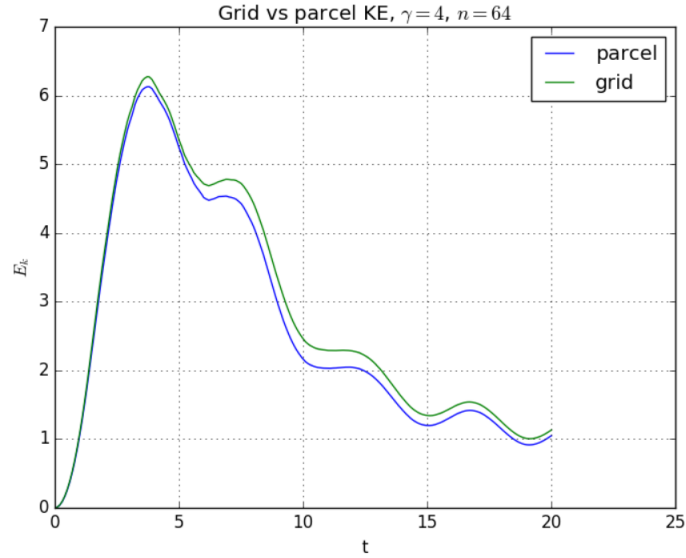


Figure 4.8: Comparison of total kinetic energy calculated using gridded and parcel fields for the simulation run at 64^3 resolution to show that deviations between the two are small.

4.3.3 Sensitivity to parcel mixing

Because turbulence is strongly related to the diffusion of kinetic energy, it makes sense to investigate how the MPIC representation of small-scale mixing affects the energetics. Recall that, in conventional LES models, a sub-grid turbulence model is used to parameterise the dissipation of energy down to the viscous sub-range where molecular diffusion takes over (Kolmogorov 1941, Smagorinsky 1963). MPIC differs substantially from these representations because it does not use an eddy-viscosity scheme and has no explicit molecular diffusion. Mixing in MPIC is modelled through the splitting and removal of parcels as described in Chapter 2. Hence, it is worth studying exactly how the process impacts the energy evolution of the system. The splitting and removal of parcels also dissipate the high vorticity values found in turbulent flow regions.

The primary parameter controlling the mixing effects in MPIC is the splitting threshold γ_{\max} . As previously described, parcels in MPIC are split when their integrated vorticity stretching,




Figure 4.9: Plot showing evolution of the total energy difference and each of its components for a simulation run at 64^3 resolution with parcel splitting turned off ($\gamma = \infty$) to illustrate the impact of parcel mixing on system energetics.

$$\gamma_i(t) = \int_{t_0}^t \left(|\boldsymbol{\omega}_i \cdot \frac{d\boldsymbol{\omega}_i}{dt}| \right)^{\frac{1}{3}} dt \quad (4.19)$$

exceeds this threshold. In most of the work in this thesis, we use the same value as Dritschel et al. (2018) and Böing et al. (2019), which is $\gamma_{\max} = 4.0$. During the initial development and testing of MPIC, this parameter was noted to have a minimal effect on the flow evolution.

We can study the effects on energy by making changes to the splitting threshold, γ_{\max} , and investigating how the total energy responds. The obvious first step is to switch the splitting of parcels off altogether and allow the system to evolve without mixing. The resulting energy evolution is shown in Figure 4.9. It is interesting to note that there is still some growth in the total energy at early times, although it is significantly reduced compared to the initial results for $\gamma_{\max} = 4.0$. There appears to be a point where the total energy is balanced again.

In contrast to the initial results in Section 4.3.1, we observe that at no point does

the system's total energy exceed the initially available energy. We still observe a brief period of decay during the initial ascent as in the previous case, but the final values of the total energy are similar to the $\gamma_{\max} = 4.0$ simulation. The absence of the initial growth suggests some link between the diffusion of energy and the rates at which splitting and removal occur, suggesting that the process of splitting parcels is part of the error in total energy. Consequently, the method is sensitive to the numerical choices surrounding the splitting and mixing of parcels and this should be considered if it is used in future studies. The results also hint that this is not the only factor contributing to the abnormal growth in total energy observed early in the simulation since we still do not observe monotonic decay.

To investigate further, we will study a range of values for the splitting threshold between $\gamma_{\max} = 4.0$, and $\gamma_{\max} = 12.0$ and observe the impact this has on the behaviour of the total energy curve. The features of interest in each curve are the rise in the total energy at early times and the dissipation due to turbulence at later times. Figure 4.10 shows how the total energy curve of the simulation changes as the splitting parameter is adjusted between $\gamma_{\max} = 4.0$ and $\gamma_{\max} = 12.0$, but also includes the run with splitting disabled for comparison. Studying these results suggests that using a slightly higher threshold for splitting may help to conserve energy in the system. While Dritschel et al. (2018) explains that the system's overall behaviour is generally insensitive to the choice of this parameter, the work covered here reveals that it impacts the overall evolution of energy in the system more substantially than anticipated. The convergence as the threshold increases is another noteworthy feature of the results. This is expected at exceptionally high splitting thresholds (since close to no splitting and mixing would occur). We see good agreement with the case with no splitting, even at $\gamma_{\max} = 10.0$. As with resolution, late stages in the simulation become more complicated due to the highly turbulent nature of the flow.

Finally, we seek to determine which components of the system energy are most susceptible to the choice of γ_{\max} , so we decompose the total energy into kinetic, dry potential and latent heating components and plot these on the same axes.

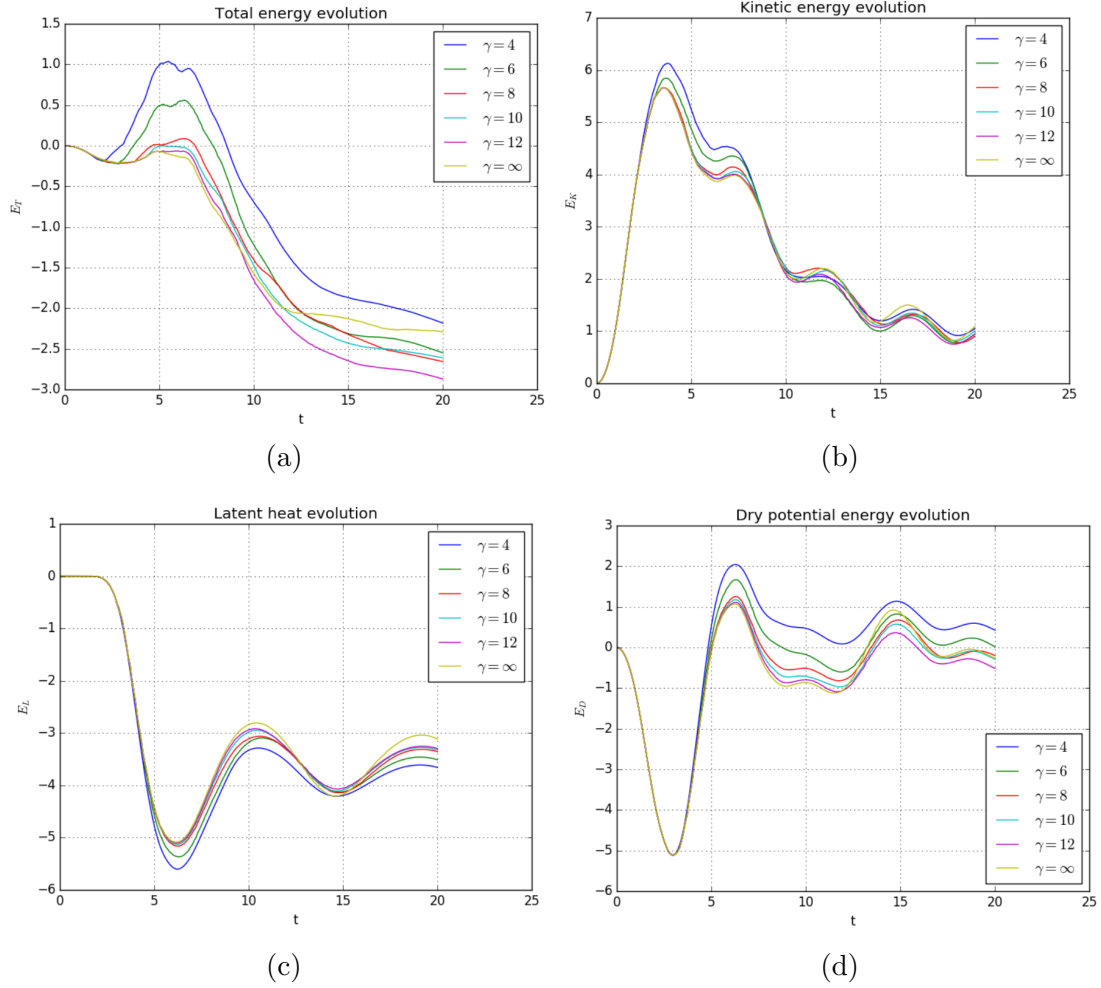


Figure 4.10: Evolution of the difference in total energy (a), kinetic energy (b), latent heat potential energy (c) and dry potential energy (d) for parcel splitting thresholds of $\gamma = 4, 6, 8, 10, 12, \infty$.

In Figure 4.10, we can see that, before $t = 5$, when the total energy starts to increase, the kinetic energy in runs with a lower splitting threshold is higher and decreases as γ_{\max} is increased. At this point, the difference between the values of kinetic energy is roughly $\Delta K = 0.5$ between the highest and lowest splitting thresholds. Hence, splitting seems to cause a considerable increase in the system's kinetic energy at early times, although as the run continues, this difference becomes less significant. The kinetic energy variation is only especially significant between the lowest values, $\gamma_{\max} = 4.0$ and $\gamma_{\max} = 6.0$. The kinetic energy curves for higher splitting thresholds remain very close to each other.

The potential energy shows a much larger sensitivity to the choice of γ_{\max} at times just beyond this point. In particular, the dry component of potential energy shows a substantial decrease $\Delta E_{p,dry} = 1.0$ between the highest and lowest values of γ_{\max} at approximately $t = 6.5$ which coincides with the peak of the rising total energy at early times. Unlike the kinetic energy, this difference remains approximately constant at later times. Also of note, is that the dry potential energy appears to be the most sensitive to variations in the splitting parameter in the later stages of the cloud's evolution.

Finally, we look at the potential energy component due to the latent heat of condensation. At early-mid times (around $t = 6$ to $t = 7$), we see a similar degree of variation to the kinetic energy, albeit negative, since the potential energy initially decreases in time. This coincides with the minimum value of this energy component. Shortly after this point, the values of $E_{p,lat}$ appear to converge again until around $t = 9$ to $t = 10$, when the impact of changing the splitting parameter has a substantial effect. For example, by the end of the simulation at $t = 20$, the difference between the highest and lowest values of γ_{\max} is approximately $\Delta E_{p,lat} = 1.2$. This is comparable to the variation in dry potential energy at early-mid times. It is also interesting to note that at early times this energy component appears to increase with increasing γ_{\max} whereas, after $t = 9$, things become inverted. By the end of the simulation, $E_{p,lat}$ decreases as γ_{\max} is increased.

These results tell us that the numerical choices made in the initial development




Figure 4.11: Plot showing the evolution of the total system enstrophy for simulations of a rising thermal carried out with parcel splitting thresholds of $\gamma = 4, 6, 8, 10, 12, \infty$.

of MPIC likely have more substantial effects than initially anticipated. All energy components are sensitive to the readiness with which mixing is allowed to occur within the model. This sensitivity must be taken into consideration when the model is applied to future case studies. The kinetic and dry potential energy seem to be most susceptible to this. Also of note is that the rate of decay of energy components appears to increase as the splitting threshold is increased. It would seem that the numerical diffusion of energy in MPIC occurs more readily when fewer parcels are being split and removed, although it becomes complicated at later times. The plot of the total system enstrophy shown in Figure 4.11 reflects this. We observe that the reduction of parcel splitting results in a much more significant build-up of enstrophy within the system. This is very likely to be a source of excess diffusion. It also appears that the increased diffusion resulting from higher values of γ_{\max} is enough to offset the observed energy growth at early times. While we attribute this primarily to a violation of incompressibility in the model, we must also consider the role of parcels splitting excessively. These results suggest that

using the vorticity stretching of parcels to dictate mixing may impact the system's energy budget. As such, it may be worth considering alternative mechanisms to determine if a parcel should split such as in Frey et al. (2022)).

4.4 Discussion of Results

Here we seek to evaluate the above results and place them in a broader context.

One of the most prominent criticisms of this work in its current state is the highly idealised nature of the test case. At this stage, a key feature missing from the model is a more sophisticated microphysics scheme, including more realistic condensation and precipitation. The ability to simulate the effects of precipitation will play a significant role in the energy budget of a developing cloud. There are similar concerns regarding the validity of the Boussinesq approximation in the thermal test case. In these simulations, the scale height is around 2000 m, meaning the domain is approximately 12 km in height. It would be more reasonable to consider an anelastic approximation in which the pressure and density are functions of the vertical coordinate (Ogura and Phillips 1962).

We must also consider the vertical structure of the environment and initial conditions since we start with an already formed thermal. In this work, the initial thermal has already overcome the environmental *CIN* to enable convection. In reality, it would make sense to position a heat source on the lower surface and allow the convection to develop more naturally. Such a setup would bear a greater resemblance to convection over the land or ocean surface.

While the idealised nature of the test case does not allow for more realistic simulations in this work, it still provides information about the exchange of energy occurring during the development of a cloud and allows us to study the structure of potential and kinetic energy within. The test case described are also helpful in analysing the effects of mixing and turbulence on the energy balance in MPIC.

We do not achieve energy conservation in the thermal test case, and several contributing factors remain to be investigated. While they do substantially influ-

ence the energetics, the parcel removal and splitting are not the sole causes of the observed increases in the total energy at early times. Further work into the possibility of numerical errors are also worth considering, there may be more numerical choices made in the development of the MPIC method that also contribute to deviations from energy conservation. Test cases which aim to study the evolution of a system with minimal turbulence may be more capable of isolating the sources of variation in total energy more reliably than the standard rising thermal test case used in this work.

The algorithm used for parcel mixing in MPIC may contribute to anomalies observed in the evolution of the total energy. Implementing a more physically intuitive algorithm for splitting and removing parcels may be worthwhile. This could be done, for example, by representing them as ellipsoids. Such a method has recently been developed and tested in two dimensions (Frey et al. 2022) and may allow for better conservation of total energy if implemented.

We can also estimate the total initial available potential energy of the thermal described within our test case using the expressions derived in 4.2.1. In doing so, we elect to set the level of neutral buoyancy within the thermal to the moist neutral buoyancy level specified in the original MPIC test case ($z_m = 5.0$). This provides an approximation of the total potential energy available to the system for convection. The thermal parcels all have a similar equilibrium level, with some small variations at the edges, so any errors from this approach will be minimal. This also serves as a sanity check to ensure we do not significantly overestimate the kinetic energy. Since the plots above all use the change in potential energy, the choice of the upper limit of integration does not change their behaviour in any meaningful way.

Computing the available potential energy by summing over all parcels at $t = 0$ and applying the volume weighting gives an initial estimate of $E_{tot} = 16.6$. We also assume that, since background parcels are already initialised at their equilibrium level, they do not contribute to the initial potential energy. Comparing this to the peak values of kinetic and change in potential energy, we can estimate the

percentage of potential energy converted into kinetic energy. When kinetic energy reaches its peak, approximately 34% of the initially available potential energy has been converted to kinetic energy. Another benefit of knowing the initial potential energy is that we can analyse how significant the loss in total energy is relative to the energy available at the start of the simulation. Using this calculation, we can infer that, by the time the system has evolved out to $t = 20$, approximately 16% of the initially available energy has been lost. Comparatively, the abnormal growth in total energy at early times is only 6% of the initially available potential energy so may not be an especially serious concern. The loss in energy is not unreasonable for a highly turbulent flow, such as this. Kolmogorov theory predicts a strong dissipation due to the forward cascade of energy (Figure 1.4). Because of the high vorticity values attained in these simulations, the lengthscales associated with the flow rapidly reach the viscous range where diffusion occurs. The effects of resolution on the dissipation are limited, because, even as the dissipation scales become smaller, the flow still reaches them quickly in turbulent regions. Moreover, the results observed for varying the splitting threshold agree with the predictions of Kolmogorov theory. As the splitting threshold is increased, the system enstrophy grows substantially at intermediate to late times (Figure 4.11). Enstrophy reflects turbulence and, as such, is associated with the dissipation of kinetic energy. Hence, while reducing the splitting of parcels might remove some numerical diffusion due to the splitting/removal operations, the subsequent increase in the system enstrophy actually results in increased dissipation of energy. This is why increasing the splitting threshold appears to counteract the erroneous growth in energy at early times.

These results are only for a 64^3 grid, and enhancing the resolution will likely limit energy diffusion due to such effects. We also note that the original MPIC test cases only ran as far as $t = 10$. At the same point in this test case, diffusion has removed approximately 8% of the initial energy. While these losses are somewhat significant, we again emphasise that the development of a cumulus cloud, as depicted here, is a highly turbulent flow, and the enstrophy evolution of the system reinforces this. As such, we do not believe such losses to be cause for significant

concern.

The Lagrangian nature of MPIC will also be beneficial to studying energy, as being able to map regions of high and low *CAPE* in real-world environments plays a role in predicting extreme weather. Since we calculate the available potential energy of parcels, this can help us study how the distribution of *CAPE* and *CIN* evolve within a cloud as it entrains air from the surrounding environment. Another consequence is that, as MPIC is developed further, the same methods can be used to analyse the energetics of clouds developing in more complex and realistic physical settings. Flows driven by surface fluxes or where there is a vertical wind shear present (this is a common factor in the formation of some convective systems such as squall lines) are just a few examples. Here we have simply laid the groundwork for studying cloud energetics in MPIC. Another goal, eventually, is to implement a precipitation scheme within the model, which will introduce further complexity to the energy budget analysis. Depending on the scheme chosen, the above equations will likely require refinement to accommodate more sophisticated microphysics as MPIC develops further. However, the basic principles behind the derived energies remain valid as long as the system is still Boussinesq and the state variable is a function of the vertical coordinate. As stated before, the Boussinesq approximation may not be valid, given the domain depth in MPIC. So this may need to be re-evaluated later if an anelastic version is considered. In such a case, it may also prove helpful to investigate whether the anomalous behaviour of the total energy is due to volume/density fluctuations being mishandled due to the use of the Boussinesq approximation in these results.

4.4.1 Incompressibility

Further work on enforcing incompressibility within MPIC has been carried out in Frey et al. (2022) during the development of the elliptical parcel-in-cell model. A three-dimensional extension of these corrections is used in a later section of this thesis. This was not implemented when recording the initial results in this chapter. As such, we can confidently say that the anomalous increases in total




Figure 4.12: Sample energy difference evolution from a simulation using the three-dimensional extension of the volume correction applied in EPIC (Frey et al. 2022).

energy observed in the higher resolution simulations appear to be primarily a result of the non-conservation of volume. When the same test cases are run with these corrections applied, the energy decays slowly during the initial ascent and more rapidly at later stages when the flow becomes highly turbulent. Energy conservation within parcel-in-cell models has always been a problematic aspect of the method, as they do not tend to conserve energy naturally (Markidis and Lapenta 2011). As such, we expect to see some loss in energy within MPIC, especially as the flow becomes increasingly turbulent.

In Figure 4.12 we show the energy evolution of a short simulation run out to a time of $t = 10$ to demonstrate the impact of enforcing incompressibility in MPIC. This shows that applying the correction removes the spurious growth in total energy observed in previous simulations' early stages. One consequence of this appears to be more substantial energy loss from the system at intermediate to late stages. A possible consequence of this is that the correction contributes to the numerical dissipation of energy in MPIC. Looking at previous results, we see that the overall energy loss between $t = 5$ and $t = 10$ is comparable to that observed in

previous simulations at similar resolutions. In the previous simulations, the initial growth in total energy offsets this. The impacts on numerical diffusion are likely relatively low, although this may be worth investigating further.

4.5 Conclusions

In this chapter, we have presented a means to analyse the energetics of clouds modelled using the current version of MPIC. Doing so has shed light on the nature of energy conservation within the method and provided us with tools to study how kinetic and potential energy evolve within more complex systems. The potential energy derived and presented in these results can be considered a profile version of (*CAPE*), often used in weather prediction and studying convective systems in the atmosphere. Being able to derive this for the initial conditions can also let us identify the convective inhibition (*CIN*) in settings where it is relevant and assess whether a given setting is unstable to convection. In the case of an initial thermal, this is less relevant since the thermal has already overcome the environmental *CIN* during its formation, and it rises freely under buoyancy. This may become more important for settings where a heat source is present in the lower domain, depending on the vertical structure of the environment. The available potential energy of the initial thermal can still be used to determine how ‘efficiently’ potential energy is converted to kinetic energy within the system and giving a yardstick for measuring any energy loss.

The results suggest that at very early times, barring numerical diffusion, MPIC conserves energy during the initial stages of cloud development relatively well, and the total energy curve at early times converges towards zero. As the flow rapidly becomes turbulent and the splitting and removal of parcels begin to take place, the behaviour of the energy becomes somewhat less stable. By varying the threshold values used for parcel splitting, we were able to identify that the splitting and removal of parcels in MPIC seems to cause a slight increase in the system’s total energy. It is unclear if this effect results from how parcel properties are redistributed after removal, or if the energy gain is a dynamical effect of the mixing

process smoothing out turbulence in the flow. We note that the potential energy components appear to be more sensitive to variations of the splitting threshold. This would suggest that the redistribution of parcel properties may contribute to this more than the dynamical effects of mixing. The decay in total observed energy is primarily attributed to turbulence and numerical dissipation in the model. Even though the losses at later times become quite large (16% of *CAPE* at late times), they are not unreasonable in a highly turbulent flow, and so we do not believe them to be a significant concern.

Further work done after this initial study using a correction to the volume field to preserve incompressibility has offered further insight into energy conservation within MPIC. Moreover, a three-dimensional adaptation of the scheme used to preserve incompressibility in Frey et al. (2022) has been implemented in later work in this thesis. A quick comparison in this chapter shows the impact this has on energy conservation. The results are promising, as the spurious increase in total energy is almost completely eliminated.

In the future, the analysis in this chapter may be able to take advantage of the Lagrangian nature of MPIC to analyse the distributions of potential and kinetic energy and their overall evolution in a range of test cases. The change in potential energy for a given parcel in MPIC is related to its vertical displacement, for example. This could be used to track the origin of environmental air in studies of flux driven convection. The presence of different chemical species in MPIC is also easily implemented through the addition of parcel attributes and their effects on potential energy could be represented in the equation for total energy density (4.17), for example some aerosols serve as condensation nuclei that can encourage condensation and reduce the condensation level of a given parcel. Such a modification may be beneficial in studies of atmospheric chemistry and transport in the future. At present, however, the work presented here has detailed a study into the properties of interest for the simplified test case. It will serve as a basis for analysing energetics presented in future works. Assessing the validity and significance of the results, even in a simple, idealised setting, has been very helpful in

moving forward with confidence.

Chapter 5

A Lagrangian study of the convective boundary layer growing into a linearly stratified atmosphere

5.1 Introduction

One area of atmospheric convection in which a Lagrangian model such as MPIC may be greatly beneficial is in the study of convective boundary layers (CBLs): turbulent layers at the base of the troposphere in which air entrained from above is mixed by convective motions. At the top of the CBL lies the entrainment zone (or the inversion layer), the region in which air is entrained downwards to feed the growth of the CBL. The primary driving force behind CBL growth is the flux of heat from the surface, although they are also strongly influenced by vertical wind shear and, in some cases, wind shear can form boundary layers by itself (Stull 1988). Both sheared and shear-free boundary layers occur regularly in the atmosphere and influence not only the surface weather, but often larger scale convective motions. Purely shear-driven boundary layers are less common, and mostly arise within the atmosphere due to surface friction, although they are common in oceanic cases where the surface wind can give rise to a shear stress at

the ocean-atmosphere boundary. In this chapter, we will primarily look to analyse the shear-free boundary layer.

Boundary layers are particularly important in larger-scale studies of atmospheric circulations, as they control the transfer of heat, moisture and momentum from the surface into the upper troposphere. Likewise, the thermodynamic properties of air entrained into the boundary layer can strongly influence its development, leading to potential drying or moistening of the layer which influences cloud development (Deardorff 1974; Mahrt 1991). As such, entrainment is a very important factor in determining the vertical distributions of heat, moisture, and momentum (Sorbjan 2005; Mellado et al. 2017) in the CBL. Because of the crucial role of entrainment in the feedback loop between CBLs and the rest of the atmosphere, studies into the behaviour of entrainment parameters and how they respond to variations in stratification or wind shear are of particular interest (Moeng and Sullivan 1994; Kim et al. 2003, Fedorovich et al. 2004; Conzemius and Fedorovich 2006a; Liu et al. 2018).

Unfortunately, the scales associated with the convective motions that drive entrainment into CBLs are much smaller than those we are able to resolve in atmospheric models. Turbulent motions within clouds, for example, occur on length-scales of a few meters (Austin et al. 1985; Blyth et al. 2005). As such, it is necessary to derive boundary layer parametrisations that encompass the effects of entrainment (Lilly 1968; Betts 1974; Deardorff 1979; Zilitinkevich 1991; Fedorovich 1995; Fedorovich et al. 2004; Conzemius and Fedorovich 2006b; Pino et al. 2006; Garcia and Mellado 2014) in order to represent the effects of the CBL. These models typically make use of bulk properties such as the entrainment rate, defined as the rate of change of the CBL depth, z_i , in time, dz_i/dt . There is some debate as to exactly how the CBL depth is defined, as well as the vertical extent of the entrainment zone above it (Garcia and Mellado 2014 provides a range of potential definitions for z_i). The most common definition used in the literature is the point at which the minimum in kinematic heat flux occurs.

Of the models used to parameterise the CBL, the most commonly used is the




Figure 5.1: Schematic of the vertical structure of the CBL in zero-order model of Zilitinkevich (1991). Vertical profiles of buoyancy, b (a), and vertical buoyancy flux B (b), are shown. Thick black lines indicate zero-order model predictions of the quantities. Dashed horizontal lines indicate the upper, $z_{i,u}$, and lower, $z_{i,l}$, boundaries of the entrainment zone, as well as the zero-order model discontinuity line at z_i . B_0 is the buoyancy flux at the lower surface and Δb is the buoyancy increments across the whole entrainment zone and δb is the local buoyancy increment at the discontinuity. δ represents the thickness of the entrainment zone (heights are not necessarily drawn to scale).

zero-order bulk model proposed by Lilly (1968), which was further developed in Zilitinkevich (1991) and Fedorovich (1995) and Fedorovich et al. (2004). The simplicity of this model makes it highly appealing for parameterisation schemes. The model approximates the entrainment zone atop the CBL as an infinitesimally thin layer across which the properties of potential temperature, θ , and horizontal velocity, u , are discontinuous. The potential temperature within the CBL is constant and above the defined top, z_i , it scales linearly with height according to $d\theta/dz = N^2\theta_0/g$ where N is the Brunt Väisälä frequency and θ_0 is the potential temperature at the lower surface. The buoyancy, b , is related to the potential temperature by $b = (g/\theta_0)(\theta - \theta_0)$, and vertical buoyancy flux, B , is related to the heat flux by, $B = (g/\theta_0)Q$. The horizontal velocity component is constant above and below the discontinuity line, with a jump at this point. Finally, the heat flux decays uniformly as the vertical coordinate increases from its value at the surface, Q_s , to $Q_i = \Delta\theta(dz_i/dt)$ at the top of the CBL, where Q_i is the entrainment heat flux and $\Delta\theta$ is the jump in potential temperature across the entrainment zone continuity. The corresponding buoyancy increments across the discontinuity are

defined as $\Delta b = (g/\theta_0) \Delta\theta$. dz_i/dt is the rate of growth of the CBL. Example profiles of the buoyancy and flux are shown in Figure 5.1.

These properties can be compared with the output of LES studies as in Fedorovich et al. (2004) where the entrainment ratio is directly obtained from the LES profiles. The value of the entrainment ratio, A , in the equilibrium entrainment regime is typically around $A = 0.2$, in agreement with observational and experimental data (Stull 1976; Deardorff 1974; Deardorff 1980; Fedorovich et al. 2004). Higher-order bulk models of CBL growth have also been developed (Betts 1974; Vanzanten et al. 1999) as well as general structure models (Fedorovich and Mironov 1995; based on Deardorff 1979), although Fedorovich et al. (2004) finds that these models are less effective in modelling the entrainment zone, likely due to an incomplete understanding of its structure. As such, higher-order models are generally unused in favour of the zero-order model, which remains applicable and benefits further from its simplicity and lack of computational expense as a parameterisation scheme.

Because the zero-order model assumes the entrainment zone is infinitesimally thin, however, it struggles to predict entrainment rate parameters in boundary layers with much thicker entrainment zones (as shown in Sullivan et al. 1998 ; Fedorovich et al. 2004). Work on the vertical structure of the entrainment zone in higher-order bulk models is relatively limited and has led to difficulties in defining the entrainment rate parameters. Garcia and Mellado (2014) suggested that a structure consisting of two overlapping sublayers may explain some discrepancies seen in other models. Their study used DNS at moderate Reynolds numbers to probe the vertical structure of the entrainment zone, identifying self-similar behaviour in the buoyancy profile.

The majority of studies on boundary layer entrainment are carried out using LES models and similar methods and results are compared to the zero-order model when calculating entrainment parameters (Fedorovich et al. 2004; Liu et al. 2018). LES studies, however, are limited by the parametrisations of sub-grid turbulence. Convective entrainment is highly turbulent and as such, the use of sub-grid turbulence

models may not be best suited to accurately probe the entrainment parameters and structure of the entrainment zone (Sullivan and Patton 2011). In previous works, this has motivated the use of DNS to fully resolve the turbulent flow of, shear-free CBL growing into a linearly stratified atmosphere (Garcia and Mellado 2014; Mellado et al. 2017). These studies yielded considerable insight into the vertical structure of the entrainment zone. DNS does have one distinct weakness in that, due to the computational expense of modelling turbulence, it struggles to model the high Reynolds numbers associated with atmospheric flows. Garcia and Mellado (2014) cite Reynolds number similarity – some flow properties become independent of the effective Reynolds number once it exceeds a certain value (Moin and Mahesh 1998; Dimotakis 2000) as partial justification for the use of DNS in their study. DNS has been proven to be a powerful tool in the study of atmospheric turbulence, in spite of the Reynolds number limitations.

We propose the MPIC method of Dritschel et al. (2018) as an alternative approach to modelling the growth of CBLs and probing entrainment zone structure. As a semi-Lagrangian model with an explicit sub-grid representation of mixing, this model avoids the parametrisations used by LES models and eliminates a degree of uncertainty at the finer scales. The model compares favourably with existing convection models (Böing et al. 2019) in a simple test case based on a rising thermal, and the parcel representation allows it to model the flow well below the resolution of its underlying grid. The particle-based nature of the model means that there is no explicit kinematic viscosity directly introduced into the model outside of limited numerical diffusion that occurs during the mixing process. As a result, a model like MPIC may have the potential to enable study CBL growth at higher effective Reynolds numbers than existing DNS studies, while offering a greater degree of computational efficiency and avoiding LES parametrisations of turbulence. At the very least, Lagrangian methods may offer an alternative to conventional atmospheric models that can attain similar Reynolds numbers and results typical of atmospheric turbulence at lower resolutions than is required within LES and DNS simulations.

In this chapter, we present the results of a simulation in MPIC modelling the growth of a convective boundary layer into a linearly stratified atmosphere, in the absence of wind shear or moisture effects. To the author’s knowledge, this is the first time a particle-based model has been used to analyse the atmospheric CBL and entrainment zone. The chapter is organised as follows. In section 5.2 we present the modifications made to the base MPIC method described in chapter 2 to allow for modelling CBL development, in particular the modifications to allow for a flux of heat to be applied at the lower surface of the system while preserving volume conservation and avoiding spurious growth of vorticity at the surface. In section 5.3 we present the results of these simulations. In section 5.3.1 we present a brief study of the energy budget in the simulation and determine an estimate for the buoyancy Reynolds number of the flow. Sections 5.3.2 and 5.3.3 show the evolution of the boundary layer through cross-sections of the buoyancy and vorticity fields, respectively. Section 5.3.4 focuses on the vertical distributions of thermodynamic properties of the boundary layer and entrainment zone. Section 5.3.5 discusses the vertical distributions of kinetic energy in the system. In section 5.3.6 we present analysis based on the vorticity dynamics based approach in MPIC in which enstrophy is used to trace turbulence within the domain, similarly to work done in Fodor and Mellado (2020), Borrell and Jiménez (2016) and references therein, and we relate these to the vertical structure of the entrainment zone noted in previous sections. In section 5.3.7 we use the vertical distributions of these properties to derive a series of measures of the boundary layer height and relate these to the structure of the entrainment zone. A brief discussion of the estimated thickness of the two entrainment zone sublayers is presented in section 5.3.8. Finally, in section 5.3.9, we compute the mean entrainment rate and other entrainment parameters of the flow, and compare these results to the predictions of zero-order models and existing computational studies to assess the performance of MPIC in relation to parameters important to large-scale models and convective parametrisations. Section 5.4 presents a summary of our findings and comments on potential extensions to sheared boundary layers, including the effects of moisture and cloud development.

5.2 Physical model

5.2.1 Surface fluxes of heat

The implementation of surface fluxes within MPIC is relatively straightforward, with the changes made to the solver in Chapter 2. A point of note in comparison to existing models is that MPIC does not directly calculate the potential temperature of a parcel. As such, we use the associated buoyancy as a thermodynamic variable in our results and the vertical fluxes of heat are obtained by calculating the kinematic flux of buoyancy through the domain. Similarly, the transfer of heat from the lower surface into the domain is modelled by incrementing the buoyancy of parcels accordingly, yielding a relatively simple model for a surface heating driven boundary layer.

First we prescribe a surface profile for the buoyancy and humidity flux on a two-dimensional grid matching the resolution of the grid used for inversion. We then identify all parcels in the simulation that are within the lowest grid layer, $z_i < z_{g1}$ where z_{g1} is the top of the lowest grid cells. The buoyancy flux is then interpolated onto the parcels using bi-linear interpolation, with a linear decay to 0 at the top of the first grid layer. For example, the buoyancy increment received by a parcel, at position \mathbf{x}_i is given by

$$b_{i,inc} = \sum_{\bar{\mathbf{x}} \in G_i} \phi(\mathbf{x}_i - \bar{\mathbf{x}}) b_{inc}(\bar{\mathbf{x}}) (z_{g1} - z_i) \Delta t, \quad (5.1)$$

where $\bar{\mathbf{x}} \in G_i$ are the four corners of the grid box containing the parcel on the lower boundary of the domain and Δt is the timestep. $b_{inc}(\bar{\mathbf{x}})$ is the buoyancy increment on the corner of the 2D grid on the lower boundary (from the prescribed profile). The $\phi(\mathbf{x}_i - \bar{\mathbf{x}})$ represent the bilinear weights used to interpolate from the prescribed surface flux profile to the horizontal parcel locations. These expressions are then weighted by the volume fraction carried by each parcel in the grid layer, \hat{V}_i .

Because the buoyancy increment is applied to each grid cell at the lower bound-

ary, we apply a global scaling factor to eliminate the effects of resolution and convert it to the input buoyancy flux B_0 . To ensure that the input flux is independent of the horizontal resolution, we must scale it according to the horizontal grid spacing. More specifically, we multiply the input flux by a factor of $\Delta x \Delta y$. This effectively converts the flux from a buoyancy input per unit area to the total buoyancy applied to a grid cell of area $\Delta x \Delta y$ which can then be interpolated onto the parcels within. This does, unfortunately, include some degree of resolution dependence the surface flux scheme but it remains a simple and effective means to incorporate surface heating in MPIC, since the region in which heat enters the system shrinks as resolution is increased. It may be more physical to define a physical depth at the base of the domain and apply the heating within and investigation into this is a topic for future works.

This represents the simplest possible implementation of surface heating within MPIC, but it is not the only possibility, as an alternative approach would involve the injection of new parcels into the flow over time. We have opted for the simplest approach initially to showcase the ability of the model to simulate boundary layer development, but a more detailed evaluation is left for future works.

5.2.2 Volume conservation

One issue that initially arose from the implementation of surface heating was a depletion of parcel concentrations in the lowest regions of the domain. This is not an uncommon issue within particle based models, so it is unsurprising that we might encounter a similar effect here. The issue arises in MPIC flux simulations because the surface heating forces parcels away from the lower boundary of the domain, and the parcels are not sufficiently replaced by the constant forcing. In order to interpolate properties correctly, MPIC requires a sufficient number of parcels in the nearby grid cells for interpolation. Consequentially, properties of the flow in the lower region become unstable. This depletion of parcels at the lower surface was found to impact the rate at which the surface flux was added to the domain, because this is directly related to the volume fraction of the grid

cells in the surface layer of the domain. This resulted in the input flux gradually declining over time as flux driven simulations were allowed to continue. In fact, the parcel depletion effect has been found in previous simulations of rising thermals in MPIC in general, although not to the extent observed in flux-driven flows

This issue was also encountered in the latest version of MPIC, the Elliptical parcel-in-cell method (EPIC) (Frey et al. 2022) which replaces the vorticity-based parcel splitting with elliptical parcels for which a more physical measure of parcel stretch can be computed. In this model, Frey et al. (2022) propose two corrections used to enforce volume conservation, the first of which ensures that the flow is divergence free, and the second is based on spatial gradients of the volume field. A two-dimensional version of these corrections and their usage is described in detail in Frey et al. (2022). The basic premise of the second correction is that parcels are shifted along the gradients of density within a given grid cell. In one dimension, the displacement of a parcel, $\delta\tilde{x}$ in this correction is given in Frey et al. (2022) by,

$$\delta\tilde{x} = C\tilde{x}(1 - \tilde{x}), \quad (5.2)$$

where \tilde{x} here describes the relative position of the parcel within the grid cell, ranging from 0 at one edge, to 1 at the opposite edge in the x -direction. C is a prefactor which is dependent on the change in gridded volume fraction across the grid cell. The magnitude of C determines the maximum extent to which parcels are compressed at the edges of the cell. To obtain the actual shift in parcel position, we simply multiply $\delta\tilde{x}$ by the grid spacing Δx . In one dimension, C is given, in the i^{th} grid cell by,

$$C = -\beta \frac{\bar{V}_{i+1} - \bar{V}_i}{V_{cell}}, \quad (5.3)$$

where \bar{V} represents the volume of each grid cell and V_{cell} is the desired volume to be occupied by a grid cell in the domain (i.e. the correct grid cell volume). The coefficient, β , is chosen to be 1.8, based on numerical tests (Frey et al. 2022).

To extend this to two dimensions, this correction is applied in each dimension

independently. However, to obtain the corresponding density gradient in the x -direction for a parcel also depend on its relative position along the vertical axis, \tilde{y} . To obtain the appropriate prefactor, C_x for the shifting along the x -axis, we linearly interpolate the volume gradients between the edges of the grid cell at $\tilde{y} = 0$ and $\tilde{y} = 1$ to the particle's vertical position. The process is repeated for the vertical parcel shifting, this time interpolating to the parcel's horizontal position within the cell.

To obtain a three-dimensional correction along a given axis, the process is identical, however now we must linearly interpolate the volume gradient along both of the other axes. For example, to obtain $\delta\tilde{x}$ we must obtain C_x by computing the volume gradient in the x -direction, interpolated between $\tilde{y} = 0$ and $\tilde{y} = 1$ as well as between $\tilde{z} = 0$ and $\tilde{z} = 1$. For example, when shifting a parcel according to density gradients in the x -direction, we have,

$$C_x = -\frac{\beta}{V_{cell}} \left[(1 - \tilde{z}) \left[(1 - \tilde{y}) (\bar{V}_{iy,ix+1,iz} - \bar{V}_{iy,ix,iz}) + \tilde{y} (\bar{V}_{iy+1,ix+1,iz} - \bar{V}_{iy+1,ix,iz}) \right] + \tilde{z} \left[(1 - \tilde{y}) (\bar{V}_{iy,ix+1,iz+1} - \bar{V}_{iy,ix,iz+1}) + \tilde{y} (\bar{V}_{iy+1,ix+1,iz+1} - \bar{V}_{iy+1,ix,iz+1}) \right] \right], \quad (5.4)$$

where ix, iy, iz refer to the x, y and z indices of the grid cell in which the parcel being shifted is located. Similarly, when shifting parcels according to density gradients in the y -direction, we use,

$$C_y = -\frac{\beta}{V_{cell}} \left[(1 - \tilde{z}) \left[(1 - \tilde{x}) (\bar{V}_{iy+1,ix,iz} - \bar{V}_{iy,ix,iz}) + \tilde{x} (\bar{V}_{iy+1,ix+1,iz} - \bar{V}_{iy,ix+1,iz}) \right] + \tilde{z} \left[(1 - \tilde{x}) (\bar{V}_{iy+1,ix,iz+1} - \bar{V}_{iy,ix,iz+1}) + \tilde{x} (\bar{V}_{iy+1,ix+1,iz+1} - \bar{V}_{iy,ix+1,iz+1}) \right] \right]. \quad (5.5)$$




Figure 5.2: Comparison of the observed RMS errors in gridded volume fraction (a) and maximum fractional volume error (b) for multiple applications of the volume corrections presented in Frey et al. (2022) for the simple rising thermal test case. 0x, 2x denote the number of iterations of the volume corrections were applied to the simulation. The colour of the line indicates the number of times that the divergence and gradient based corrections were applied.

Finally, the prefactor for shifting parcels in the vertical is given by,

$$\begin{aligned}
 C_z = -\frac{\beta}{V_{cell}} & \left[(1 - \tilde{x}) \left[(1 - \tilde{y}) (\bar{V}_{iy,ix,iz+1} - \bar{V}_{iy,ix,iz}) + \tilde{y} (\bar{V}_{iy+1,ix,iz+1} - \bar{V}_{iy+1,ix,iz}) \right] \right. \\
 & \left. + \tilde{x} \left[(1 - \tilde{y}) (\bar{V}_{iy,ix+1,iz+1} - \bar{V}_{iy,ix+1,iz}) + \tilde{y} (\bar{V}_{iy+1,ix+1,iz+1} - \bar{V}_{iy+1,ix+1,iz}) \right] \right].
 \end{aligned}
 \tag{5.6}$$

The results in figure 5.2 showcase the effects of applying each correction multiple times. These results are obtained for the standard MPIC test case described in previous sections of a rising thermal initially at rest and show the RMS errors in the gridded volume fraction. These results show that just applying the correction twice is sufficient to reduce the fractional RMS volume errors to well below 10^{-3} and the maximum volume fraction errors to a few percent. We also note that the observed errors appear to show convergence on the log scales used as the

corrections are applied multiple times. Beyond a certain point, we see diminishing returns for the computational cost incurred. The largest errors are observed at early times in the rising thermal test case in runs with fewer applications of the corrections, and this is likely because the strongest vertical motions are observed during the initial ascent. This means that parcels within the thermal are displaced rather rapidly, causing depletion to occur more strongly than can be counteracted. As such, we elect to apply the correction twice to balance between the additional computational cost and accuracy of the volume field.

5.2.3 Prevention of spurious vorticity generation

Another issue that emerges as a consequence of the surface fluxes is the generation of vorticity, ($\boldsymbol{\omega} = (\xi, \eta, \zeta)$), at the lower boundary. More specifically, this is a consequence of the free-slip boundary conditions used in MPIC. These, combined with the horizontal buoyancy gradients arising from a surface flux profile, can lead to vortices becoming trapped at the lower (and occasionally upper) boundaries. These vortices grow in strength and lead to the formation of a spurious shear in the domain that can eventually grow to dominate the motions of the system.

In conventional simulations of convective boundary layers, a no-slip boundary condition is enforced (see any of the previous studies cited in this chapter). Whether a free-slip boundary can be considered realistic within studies of atmospheric convection is unclear. We do not attempt to force no-slip boundaries in MPIC to minimise the effects of additional viscosity and maintain higher Reynolds numbers across the domain. We do, however, acknowledge that some degree of diffusion is required in order to offset the vorticity growth at the vertical boundaries. Hence, we aim to apply a damping effect on the vorticity of parcels very close to the boundaries.

We start from the vorticity equation presented in Chapter 2

$$\frac{D\boldsymbol{\omega}}{Dt} = \boldsymbol{\omega} \cdot \nabla \mathbf{u} + (b_y, -b_x, 0), \quad (5.7)$$

which can be split into the individual components

$$\begin{aligned}\frac{D\xi}{Dt} &= \xi u_x + \eta u_y + \zeta (\eta + w_x) + b_y, \\ \frac{D\eta}{Dt} &= \xi v_x + \eta v_y + \zeta (\eta + w_y) - b_x, \\ \frac{D\zeta}{Dt} &= \xi w_x + \eta w_y - \zeta (u_x + v_y).\end{aligned}\tag{5.8}$$

The simplest course of action would be to introduce an additional term to each component consisting of a dimensionless damping rate, α , multiplied by the corresponding vorticity component such that, for example, $D\xi/Dt \rightarrow D\xi/Dt - \alpha\xi$. This linear damping was then applied to all parcels located within the lowest grid cell of the domain, the same region as those receiving increments of buoyancy from the lower surface. This method left much to be desired, as it is heavily dependent on the choice of α . For example, if the system is overdamped we generate an instability at the boundary and underdamping fails to prevent the excessive growth of vorticity. When employing a nonlinear damping of the form, $D\xi/Dt \rightarrow D\xi/Dt - \alpha\xi^2$ we found similar issues.

Instead, we look to a more sophisticated means of curbing the growth of vorticity at the vertical boundaries. We introduce a diffusion term to the vorticity equation that is dependent on the spatial gradients of the vorticity components and controlled by a nonlinear function of the vertical coordinate, chosen such that it disappears above the damping region. Using the x -component of vorticity again as an example, the tendency now takes the form

$$\frac{D\xi}{Dt} = \xi u_x + \eta u_y + \zeta (\eta + w_x) + b_y + \nabla \cdot (\kappa(z) \nabla \xi),\tag{5.9}$$

where $\kappa(z)$ is a height-dependent eddy viscosity which defines the form of the diffusion profile. More specifically, we choose

$$\kappa(z) = \alpha (z_2 - z)^2,\tag{5.10}$$

where α is a dimensionless damping rate. z_2 is the vertical coordinate associated

with the second grid layer such that the diffusion is spread across the two grid layers closest to each vertical boundary. Computing the diffusion term requires the spatial derivatives of ξ and so we convert to spectral space in the horizontal, obtaining

$$\hat{\xi}_t + (k^2 + l^2) \kappa(z) \hat{\xi} = \hat{S}_\xi + \frac{d}{dz} \left(\kappa(z) \frac{d\hat{\xi}_h}{dz} \right), \quad (5.11)$$

where \hat{S}_ξ is the usual tendency of ξ in spectral space and k and l are the horizontal wavenumbers as in Chapter 2. We then evaluate the vertical derivatives using second order finite differences. At z_0 we obtain

$$\frac{d}{dz} \left(\kappa(z) \frac{d\hat{\xi}_h}{dz} \right) = \left[\kappa \left(z_{\frac{1}{2}} \right) \frac{d\hat{\xi}}{dz} \left(z_{\frac{1}{2}} \right) - \kappa \left(z_{-\frac{1}{2}} \right) \frac{d\hat{\xi}}{dz} \left(z_{-\frac{1}{2}} \right) \right] / \Delta z. \quad (5.12)$$

The vorticity components are chosen to be symmetrical across the boundary between, $n_z = -1, \dots, 1$ such that the second term in the numerator is the product of an odd and an even function. This results in a centred difference at $z = z_0$,

$$\frac{d}{dz} \left(\kappa(z) \frac{d\hat{\xi}_h}{dz} \right) \approx \frac{2\kappa \left(z_{\frac{1}{2}} \right) - \left(\hat{\xi}_1 - \hat{\xi}_0 \right)}{(\Delta z)^2}, \quad (5.13)$$

and similarly, we obtain the finite differencing at the base of the second grid cell, $z = z_1$,

$$\frac{d}{dz} \left(\kappa(z) \frac{d\hat{\xi}_h}{dz} \right) \approx \frac{\kappa \left(z_{\frac{3}{2}} \right) \left(\hat{\xi}_2 - \hat{\xi}_1 \right) - \kappa \left(z_{\frac{1}{2}} \right) \left(\hat{\xi}_1 - \hat{\xi}_0 \right)}{(\Delta z)^2}. \quad (5.14)$$

The process is identical for the upper boundaries between $z = z_{n_z}$ and $z = z_{n_z-2}$. Initial tests suggest that in the absence of this damping, vorticity can become trapped at the top of the domain and accumulate. To avert this, we implement the same damping layer at the upper surface of the domain. It may be preferable to introduce a proper sponge layer in this region that relaxes all properties towards the initial state near the upper surface instead. This is a somewhat more involved procedure within MPIC however and would require tuning of further numerical parameters, so we have opted to simply damp the vorticity. Our testing suggests




Figure 5.3: Schematic diagram showing the experimental setup for the boundary layer simulation. The buoyancy field at time $t = 0$ is shown alongside a graph showing the linear variation of buoyancy with height, $\langle b_0 \rangle (z) = N^2 z$. B_0 represents the strength of the buoyancy flux at the lower surface and L_x and L_z are the dimensions of the domain of the horizontal and vertical directions.

that a value of $\alpha = 0.05$ is sufficient to prevent the build up of vorticity in CBL simulations using MPIC without overdamping the system.

This method is perhaps not the ideal means with which to address the issue, but it does provide a relatively simple solution to the problem of spurious vorticity growth. Damping the vorticity in this way leaves the bulk of the MPIC method unaltered and avoids having to directly enforce no-slip boundaries, which would require considerable modifications to the vertical solver.

One concern with this method is that it introduces resolution dependence, since vorticity is only damped within the lowest two grid cells. It may be worth revisiting the damping and defining a physical extent of the diffusive region which does not change with vertical resolution. On the other hand, the resolution dependence may be beneficial as the diffusion takes effect only within a smaller region of the domain retaining high Reynolds numbers in all but the lowest and uppermost grid cells. It may be necessary to adjust the damping rate in such cases, however.

5.2.4 Experimental setup

Here, we present the parameters used for the simulation of a dry convective boundary layer in MPIC. Quantities used are non-dimensionalised according to the values

in Chapter 2 unless stated otherwise. The majority of the numerical parameters within MPIC are the same as in previous chapters, with the exception of the splitting threshold for parcels which, for this simulation, was set to $\gamma_{max} = 8.0$. This value was chosen because Chapter 4 has shown that reduced splitting of parcels offers greater conservation of energy in the simulation. The minimum parcel volume was kept as $1/6^3$ of the grid cell volume and the model was initialised with 4 parcels within each grid cell.

Compared to previous simulations of rising thermals in MPIC, we have adopted a much wider horizontal domain. This allows for much better sampling of properties used to generate vertical profiles. In an attempt to emulate existing works such as Garcia and Mellado (2014) and Mellado et al. (2017), we expand the horizontal extent of the domain to be four times larger than the vertical, while keeping the grid spacing the same. Figure 5.3 shows a schematic of the initial setup, including a cross-section through the domain along with the vertical profile of the horizontally averaged initial buoyancy, $\langle b_0 \rangle$. The domain size is $L_x \times L_y \times L_z = 8\pi \times 8\pi \times 2\pi$ at a resolution of $128 \times 128 \times 32$. Our reasoning for keeping the horizontal and vertical grid spacing the same is due to the homogeneous nature of turbulence in shear-free boundary layers. The upper and lower boundaries of our domain are free-slip, stress-free boundaries as in previous MPIC simulations and we use periodic boundary conditions in the horizontal directions. As discussed previously, we apply a surface flux of buoyancy, with strength B_0 to all parcels in the lowest grid cell of the domain. We also use the vorticity damping at the upper and lower vertical boundaries outlined in Section 5.2.3, with a dimensionless damping rate of $\alpha = 0.05$.

The vertical structure of the initial environment is very straightforward, being a dry linearly stratified atmosphere with $db/dz = N^2$ where N is the Brunt Väisälä frequency. This surface is heated from below by a uniform buoyancy flux of strength B_0 in order to drive the growth of the boundary layer. In order to break symmetry, we add a random noise to each grid cell in the prescribed buoyancy profile, generating a random perturbation between -1% and $+1\%$. An analytic

perturbation to the flux profile was also tested in early stages, using a superposition of sine and cosine waves but this was removed in favour of the random noise. We also note that the choice of initial conditions has little effect on the long term evolution of the system, at late times the bulk properties of the flow are near identical in both cases. The random noise was more effective in breaking down symmetry at early times, despite being dependent on the horizontal resolution.

The most important parameters in determining the evolution of the system are the surface buoyancy flux, B_0 , and the background stratification, N^2 . We use these to define an Ozmidov lengthscale, L_0 , which can be interpreted as the integral lengthscale of turbulence within the system (Ozmidov 1965) because the viscous dissipation rate in a shear-free boundary layer is of similar order of magnitude to the surface buoyancy flux (Fedorovich et al. 2004). This Ozmidov lengthscale can also be interpreted as the minimum thickness of the CBL (see Appendix A in Garcia and Mellado 2014). L_0 incorporates the effects of the stratification and surface heating and so can be useful in further analysis, or even used to represent the development stage of the CBL in place of non-dimensional time (Garcia and Mellado 2014; Haghshenas and Mellado 2019). L_0 is given by the equation,

$$L_0 = \left(\frac{B_0}{N^3} \right)^{\frac{1}{2}}. \quad (5.15)$$

In this simulation of a convective boundary layer in MPIC we use dimensionless values of $B_0 = 0.0044$, $N = 0.97$ and obtain $L_0 = 0.07$ corresponding to physical values of $L_0 = 140$ m, $B_0 = 0.007 \text{ m}^2 \text{ s}^{-3}$ and $N = 0.007 \text{ s}^{-1}$. We note that the values of B_0 and N are comparable to typical atmospheric values (Fedorovich et al. 2004). The value of L_0 also lies well within the range of values studied in Garcia and Mellado (2014). Through the course of the simulation we see the CBL grows to depths of $z_{enc}(t)/L_0 = 10$ by a dimensionless time of $t = 50$ where $z_{enc}(t)$ is the encroachment height, representing the upper boundary of the mixed layer at time t . The computation of the encroachment height (Lilly 1968; Carson and Smith 1975) is discussed further in section 5.3.7. The run was allowed to continue to a dimensionless time of $t = 400$, encompassing stages of development studied

in Garcia and Mellado (2014) in which the boundary layer grows to a depth of $z_{enc}/L_0 = 26$.

5.3 Results

5.3.1 Energy budget and estimation of Reynolds number

While studying the development of the boundary layer, it is important to understand how accurately the flow resembles real turbulence. Hence, we aim to estimate the buoyancy Reynolds number of the flow in MPIC. This is typically obtained from the expression

$$Re_0 = \frac{B_0}{\nu N^2}, \quad (5.16)$$

where all symbols have their previously stated meanings, and ν is the kinematic viscosity of the fluid. Within the semi-Lagrangian model MPIC, this becomes somewhat more complicated to calculate, because the model does not include an explicit representation of viscosity. This is due to how the mixing process is handled within MPIC. More details on this can be found in Chapter 2. Because MPIC simulates mixing through the splitting and removal of parcels based on their volume and the vorticity stretching integral, the only dissipation that occurs is numerical. Thus, we must look for an alternative means to estimate the effective buoyancy Reynolds number.

To do this, we must consider what it is that limits the resolution of turbulence within MPIC. In particular, we must identify the smallest lengthscale within the system that energy can cascade down to before it is subjected to numerical diffusion. We may instead represent the kinematic viscosity of the system in terms of the Kolmogorov scale given by the equation

$$l_k = \left(\frac{\nu^3}{\epsilon_D} \right)^{\frac{1}{4}}, \quad (5.17)$$

where ϵ_D is used to represent the average rate of dissipation of turbulence kinetic energy and l_k is the Kolmogorov scale at which viscous forces become dominant




Figure 5.4: Evolution of the volume-averaged total energy (black line) and the energy generated by the surface flux (blue line).

and kinetic energy is dissipated.

This term can be rearranged to give an estimate of the kinematic viscosity due to numerical mixing, ν , within MPIC but we still need to determine the lengthscale l_k . A reasonable and simple way to estimate this scale is to use the volume threshold that determines when parcels are removed from the model and their properties redistributed to the grid. In the version of MPIC used in this study, parcels are removed once their volume becomes smaller than $(1/6^3)\Delta V$ where ΔV is the grid cell volume. So we can make a crude estimation that the minimum lengthscale is approximately six times smaller than the grid spacing. This would give us a value of $l_k \approx 0.03$ which is equivalent to about 60 m, several orders of magnitude larger than typical atmospheric values. Given the domain height is the same as in previous simulations at 12 km, that means this lengthscale is about 0.5% of the domain height. Also of note is the relatively low resolution at which this study has been carried out. It is likely that focusing on a shallower domain and increasing the grid resolution could reduce this value considerably.

The other term needed to estimate the Reynolds number is ϵ_D , the dissipation rate of kinetic energy density, represented as a volume averaged quantity. There is no simple way to estimate this within MPIC a priori, and so we must study the evolution of the total system energy across the duration of the simulation. We calculate the potential energy of the parcels as described in Chapter 4 and compare this to the energy added to the system by the surface buoyancy flux. The comparison of these is shown in Figure 5.4, and we can use this to estimate that by the end time of $t = 400$, the volume averaged total energy of the system is approximately $\Delta E = 0.1$ less than the energy density that entered through the lower surface. Because the rate of change of total energy is approximately linear in Figure 5.4, we can estimate the value of ϵ_D by simply dividing by the run time of the simulation, giving us $\epsilon_D = 0.00025$ in non-dimensional units.

We can then substitute the rearranged expression for the kinematic viscosity in terms of the lengthscale l_k and dissipation rate ϵ_D into the expression for the buoyancy Reynolds number, Re_0 giving,

$$Re_0 = \frac{B_0}{l_k^{\frac{4}{3}} \epsilon_D^{\frac{1}{3}} N^2}. \quad (5.18)$$

Substituting in the values obtained above gives an estimated value of $Re_0 = 8.5$. In comparison to typical atmospheric cases this is incredibly small, and even compared to work done in Garcia and Mellado (2014) it is still an order of magnitude lower than the value obtained in their second test case, ($Re_0 = 117$). It is important to note that there is a massive difference in resolution between the simulations carried out here and those in other studies. For example, Garcia and Mellado (2014) use a grid resolution of $5120 \times 5120 \times 840$ in the second test case which is more than ten times the grid resolution used here in MPIC in each dimension. Another example can be found in Liu et al. (2018) who used $400 \times 400 \times 400$ grid points at a resolution of $25 \text{ m} \times 25 \text{ m} \times 25 \text{ m}$ compared to a homogeneous grid spacing of $\Delta x = \Delta y = \Delta z = 0.196 \approx 400 \text{ m}$ in the MPIC study presented here. Given that the radius of the smallest parcels used to estimate l_k is directly proportional to the grid spacing, we can see that this Reynolds number will have a $n^{4/3}$ scaling

with the grid resolution. This would likely yield comparable, if not higher, values of Re_0 than in DNS and LES in a study of comparable grid resolution. Chapter 4 also shows that the conservation of energy in MPIC can be improved with increasing resolution and tuning of the parcel splitting parameters, which will also impact the value of Re_0 . Because Re_0 only depends on $\epsilon_D^{1/3}$, the effect of this will be less pronounced.

It is also worth noting that the only source of viscosity within the model beyond the numerical mixing is the vorticity damping in effect at the upper and lower surfaces due to their free-slip nature. If this is indeed responsible for the dissipation of energy observed in Figure 5.4 then we should note that the behaviour of the remainder of the domain, in particular the entrainment zone at later times, may have a higher local Reynolds number than the estimates provided here. To examine this further, we would need to be able to map the dissipation of energy within the system.

In relation to the characteristic lengthscale L_0 we find that the estimated Kolmogorov scale is slightly less than half the value of L_0 . One could then also argue that the buoyancy Reynolds number can be determined by using the ratio of these two lengthscales, i.e.

$$Re_0 = \left(\frac{L_0}{l_k} \right)^{\frac{4}{3}}. \quad (5.19)$$

This yields an estimated value of $Re_0 = 3.1$ which is even lower still. This does not, however, account for the dissipation of energy and effective viscosity introduced by the numerical method in MPIC, so we believe the initial estimate to be more reliable. It would not be correct to dismiss this outright either, so we present it here regardless as an alternative estimation.

Our ability to reliably estimate Re_0 here is also somewhat limited by the fact that we can only assume that the Kolmogorov scale is the size of the smallest parcels, and that we must estimate the dissipation rate from the loss in total energy. This remains a crude estimation used to provide some comparison to other studies of convective boundary layers in the atmosphere. A comparison of the kinetic energy




Figure 5.5: Evolution of the change in kinetic energy as a fraction of the total energy in the system. This is not the same as the total energy entering the system due to the surface flux (see Figure 5.4).

of the system as a fraction of the total energy is shown in Figure 5.5, and we see that the overwhelming majority of the energy input is used to heat the domain and stored as potential energy. Kinetic energy constitutes a fraction of 0.20 – 0.25% of the total energy available, much lower even than the fraction lost to dissipation. After the initial spike in kinetic energy, this fraction actually becomes close to steady in time.

The implications of this are unclear, but may well have an impact on the reliability of our estimation of the dissipation rate. The considerable difference in magnitude of the kinetic and potential energy may well suggest that the dissipation rate is much higher than expected, if it is the result of kinetic energy being rapidly dissipated into heat via turbulent mixing. This may lend some credibility to the estimated value based on the ratio of the characteristic lengthscale L_0 and the Kolmogorov scale, l_k , which has a lower value. It is unlikely, however, that a factor of two in Re_0 would result in dissipation capable of accounting for the several orders of magnitude between kinetic and potential energy. On the other

hand, if the total kinetic energy is so much lower than even the total energy loss, it may be possible that energy other than kinetic energy is being diffused. This would suggest that our value of ϵ_D above is too high, since it describes only the dissipation of turbulent kinetic energy. It is unclear which of these is likely to be the case, but in later sections we endeavour to keep the estimated effective Reynolds number in mind as we describe the evolution of the boundary layer and its structure.

5.3.2 Evolution of the buoyancy field

Here we illustrate the evolution of the system through cross-sections of the buoyancy and vorticity field taken through the centre of the domain at different stages in the boundary layer's development. This provides a useful visualisation of how the domain evolves in time. Figure 5.6 shows the early thermals rising into the stratified atmosphere after the surface heating through the buoyancy field alongside its variation relative to the initial background profile $b_0(z)$ and the horizontal average, $\langle b \rangle$. We also include the vorticity field. The analysis presented is largely a qualitative assessment of these properties and the insight they may give us regarding the structure of the boundary layer and entrainment zone.

Prior to around $t = 35 - 40$ the heat generated by the surface flux accumulates within the lower region of the domain before the warm air eventually organises into small thermals. These thermals rise upwards a short distance before dispersing and entraining air from above into the lower region of the domain. This process is what feeds the growth of the boundary layer. The distribution of these thermals is likely determined by the small random perturbations applied to the near homogeneous surface heat flux, and is much less uniform than if we had used an analytic function as was originally tested. Beyond this initial growth phase, we find that the effects of the form taken by the small spatial variations of the surface flux have a limited impact on the overall development of the boundary layer at later times. The vorticity field shown provides a much clearer visualisation of the influence of the thermals on the ambient air. In particular, the vorticity does not necessarily




Figure 5.6: x - z cross-sections through the centre of the domain, taken at time $t = 35$ showing (a) the buoyancy field, b ; (b) the buoyancy anomaly relative to the initial profile, $b - b_0(z)$; (c) the buoyancy anomaly relative to the horizontal average, $b - \langle b \rangle$; (d) the vorticity field magnitude, $|\omega|$, during the rise of the initial thermals resulting from the surface heating.

overlap with the regions in which the buoyancy field shows the dispersing thermals but rather it is concentrated at the edges of the dispersing thermals and along their upward trajectory as air is pulled into the region below them and can hence show us where the strongest mixing is likely taking place.

Because we keep many of the parameters the same as in previous MPIC simulations of individual thermals (see Dritschel et al. 2018 and Böing et al. 2019 or Chapters 3 and 4) we can actually compare the time taken for these thermals to form, rise and disperse to the existing simulations. The original MPIC case study was run out to a non-dimensional time of $t = 10.0$ but in this simulation the early thermals are not seen until a time of around $t = 35$. Once formed, they reach the peak of their ascent quite rapidly, although they are each individually smaller and less buoyant and hence only rise a short distance vertically. This does, however, represent the first time within this framework in which we have simulated naturally occurring convection without the need to start from an initial thermal. This opens the way for studies of surface flux driven cloud fields in a Lagrangian framework.

Once the initial thermals have ascended and entrained air downwards, it begins to mix with the air in the lower region of the domain, and we start to see the formation of the boundary layer. The further evolution of the buoyancy field is shown in Figure 5.7. We also include the time evolution of the buoyancy anomaly relative to the initial background profile and horizontally averaged buoyancy in Figures 5.8 and 5.9. These are included to better illustrate the plumes and thermals forming in the mixed layer. The $b - b_0(z)$ shows the extent of the boundary layer particularly well, while the $b - \langle b \rangle$ profile highlights the plumes in the boundary layer and the air displaced by convective updrafts. These results show that the mixed layer at the base of the domain appears to grow rapidly in time and as the layer becomes more strongly mixed we no longer observe clear thermals but rather wisps and plumes rising from the surface and pushing air up into the stratified region. The entrainment zone is most clearly visible within the early-mid times of the cross-sections (Figures 5.7, 5.8, 5.9) where pockets of cool air from the




Figure 5.7: x - z cross-sections through the centre of the domain taken at times (a) $t = 100$, (b) $t = 200$, (c) $t = 300$ and (d) $t = 400$ showing the buoyancy field, b .




Figure 5.8: x - z cross-sections through the centre of the domain taken at times (a) $t = 100$, (b) $t = 200$, (c) $t = 300$ and (d) $t = 400$ showing the buoyancy anomaly relative to the initial background profile, $b - b_0(z)$.




Figure 5.9: x - z cross-sections through the centre of the domain taken at times (a) $t = 100$, (b) $t = 200$, (c) $t = 300$ and (d) $t = 400$ showing the buoyancy anomaly relative to the horizontally averaged buoyancy, $b - \langle b \rangle$.

stratified region are seen being pulled down into the boundary layer. By the end of the simulation, it appears that the mixed layer has grown to nearly a third of the domain height.

Another feature of note in the buoyancy field is the appearance of perturbations in the upper regions of the domain, far above the boundary layer and entrainment zone. The cross-sections of the buoyancy anomaly relative to the horizontal average ($b - \langle b \rangle$) highlight these perturbations more clearly. It is, at present exactly what the source of these is, but they may arise to numerical effects.

5.3.3 Evolution of the vorticity field

While the buoyancy field of convective boundary layers growing into stratified atmospheres has been studied in previous works (Fedorovich and Mironov 1995; Sorbjan 1999; Fedorovich et al. 2004; Garcia and Mellado 2014; Liu et al. 2018), studies pertaining to the vorticity field are somewhat rarer. Typically, it is used to distinguish between the turbulent regions of the boundary layer and the non-turbulent region above (Bisset et al. 2002; Mellado et al. 2009; da Silva et al. 2014b; Fodor and Mellado 2020) as well as in studies of the interface between such regions (da Silva et al. 2014a). These studies typically use DNS to resolve the enstrophy field. This, to the author's knowledge, is the first time a Lagrangian model has been used in this context and due to this novelty, the comparison to existing studies will be beneficial to improving the method in the future. As a model formulated from the vorticity equation, the study of enstrophy in MPIC seems a natural application of the method. Because enstrophy is directly related to kinetic energy and often traces regions of dissipation in the system, we can use it to identify the regions where turbulent mixing is occurring most strongly.

The evolution of the vorticity field, presented in Figure 5.10 paints an interesting picture. At $t = 100$, we see that the boundary layer is still in a relatively early stage of its development, and the vortices found in the lower region of the domain are fairly small and spread around the thermals and plumes observed in the buoyancy field at this time.




Figure 5.10: x - z cross-sections through the centre of the domain taken at times (a) $t = 100$, (b) $t = 200$, (c) $t = 300$ and (d) $t = 400$ showing the magnitude of the vorticity field, $|\boldsymbol{\omega}|$.

As the boundary layer grows at intermediate times, $t = 200$ to $t = 300$, the vorticity field begins to change. The many small pockets of vorticity seem to have dispersed or merged, giving way to larger, stronger regions of vorticity that are concentrated primarily just below the upper extent of the layer. These vortices extend into the entrainment zone itself, where the cool air is being pulled from. This implies that the majority of the mixing is found just below the base of the entrainment zone, which seems reasonable given that this is where the air rising through the mixed layer is likely to meet the pockets of air entrained from above. There are particularly prominent vortices found near the centre of the domain at both $t = 200$ and $t = 300$ that illustrate this point clearly. If we compare the same snapshots of the buoyancy fields, we can see that these vortices roughly align with particularly strong updrafts and the resulting entrainment of air from above.

In the upper region of the domain, we see vorticity aligning with the buoyancy perturbations observed in Figures 5.7-5.9. In this figure their distribution is much more clear, in particular at $t = 200$. The intensity of the vorticity appears to decay by the end of the simulation. This is likely caused, in part, by the damping at the upper and lower surfaces.

Directly above the vortices driving the entrainment, there appears to be a transition region. Across this layer, the vorticity field changes from being dominated by updrafts and the mixing of entrained air to being largely dominated by small regions of weak vorticity scattered across the upper half of the domain. The vorticity falls off rapidly, giving rise to a region of depleted vorticity between the boundary layer and the stratified region aloft. This is particularly evident in the cross-section for $t = 200$. These findings are consistent with the work of Garcia and Mellado (2014), who suggested that the entrainment zone can be represented by two overlapping sublayers. The upper sublayer in this description of the entrainment zone is characterised by turbulent thermals penetrating into a smooth environment, and it is described as a transition region between turbulent and non-turbulent air. Further analysis of the enstrophy field in Fodor and Mellado (2020) found that enstrophy drops off dramatically across the entrainment zone, which

is reflected in our own results. The scale separation between the free atmosphere above and the boundary layer below was used to determine suitable thresholds for the conditioning of air into turbulent and nonturbulent regimes. This is illustrated in the MPIC vorticity field by the large vortices failing to rise beyond the upper limits of the entrainment zone. Section 5.3.5 and 5.3.6 will study this further through the vertical distributions of enstrophy and kinetic energy.

5.3.4 Vertical distribution of thermodynamic properties

Qualitative analysis of the observed fields can only provide so much information. As discussed in Section 5.1, many bulk models used for representing the convective boundary layer (Fedorovich et al. 2004 provide an overview of both zero-order and general structure models) make use of the vertical distributions of thermodynamic properties of the boundary layer. This approach can offer some insight into the growth of the layer and provides a means of validating our results against other studies.

One of the most commonly used quantities in defining the vertical extent of the convective boundary layer is the vertical flux of buoyancy since this can identify where the entrainment of buoyant air from above is occurring. This is given by

$$B = \langle b'w' \rangle - \kappa_m \frac{\partial b}{\partial z}, \quad (5.20)$$

where angled brackets denote the horizontal averages of the quantities within, and primes denote variations from the horizontal average. κ_m is the molecular diffusivity. In MPIC, we do not model molecular diffusion explicitly, so κ_m is zero for the purpose of this study.

The vertical profiles for the horizontally averaged vertical buoyancy flux, $B = \langle b'w' \rangle$, are shown in Figure 5.11. The formulation of the MPIC model requires that the vertical velocity at the lower surface be zero, hence the sudden decay in the lowest grid cell in these profiles. In reality the flux at this level will be $B = B_0 = 0.0044$. The region in which the buoyancy flux decays at the base of

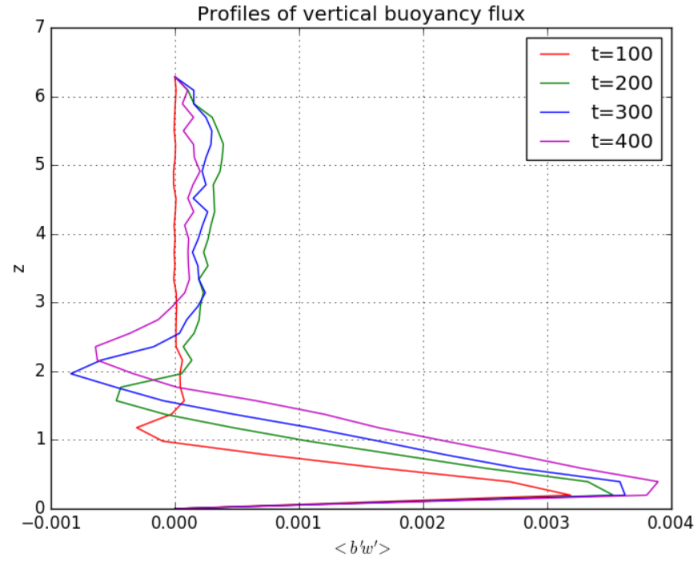


Figure 5.11: Vertical distributions of the horizontally averaged vertical buoyancy flux $B = \langle b'w' \rangle$ for $t = 100$ (red), $t = 200$ (green), $t = 300$ (blue), $t = 400$ (magenta).

the domain also appears to overlap with the vorticity damping layer, which may influence the vertical velocity to some extent.

The most important feature of these profiles is the region in which the buoyancy flux becomes negative, the entrainment zone. A negative buoyancy flux indicates that air more buoyant than the horizontal average is being entrained downwards (positive b' , negative w') or that air less buoyant than the horizontal average is detraining upwards (negative b' , positive w'). Typically, the point at which the buoyancy flux first becomes negative (the zero-crossing point) is regarded as the base of the entrainment zone (Fedorovich and Mironov 1995; Sorbjan 1999; Fedorovich et al. 2004), and the negative buoyancy flux observed above is the entrainment buoyancy flux (Lilly 1968).

The negative peak marking the approximate centre of this region moves upwards in time as the boundary layer grows. It extends more in the vertical direction, resulting in an entrainment zone that spans a comparable depth to the boundary layer below. It may be that this results from the lack of a sponge layer in the

upper region of the domain. It may also be a consequence of us neglecting the molecular diffusion term in equation 5.20. Comparing this to the vertical cross-section shown in Figure 2 and the vertical profiles across the entrainment shown in Figure 4 in Garcia and Mellado (2014), it would appear that MPIC has produced a much thicker entrainment zone. In a low Reynolds number flow, we would expect the region of negative buoyancy flux to be narrower because it is harder for air rising from thermals and plumes to penetrate so far into the stratified region. Likewise, the turbulent motions that drive entrainment would be damped more strongly. This is also of note because the zero-order bulk model assumes that the entrainment zone is infinitesimally thin, as discussed in Section 1. This model becomes less reliable for boundary layers with a relatively thick entrainment zone (Sullivan et al. 1998; Fedorovich et al. 2004) and this should be kept in mind in subsequent comparisons.

One more feature of note in these profiles of the vertical buoyancy flux is the region above the entrainment zone, where the buoyancy flux becomes positive (or zero) again. This seems to suggest that there is upward transport of heat through the stratified region above the boundary layer, it may be, in some way, linked to the disturbances observed in the vorticity in this region, although the origin is not immediately clear. We can see at $t = 100$ that the buoyancy flux above the entrainment zone is effectively zero.

The region below the entrainment zone is the mixed layer, characterised by rising updrafts of buoyant air heated by the surface flux. The buoyancy in the mixed layer is close to constant. Consequently, once air has been heated to a certain point (see Figure 5.12), it can rise freely through this region, typically overshooting into the layer above due to its momentum. Hence, we typically observe a strong positive flux of buoyancy that decays with height. The zero-order bulk model approximates this as a linear decrease in the buoyancy flux (Zilitinkevich 1991; Fedorovich et al. 2004), and the results obtained in MPIC seem to exhibit similar behaviour in this region.

The vertical distribution of buoyancy shown in Figure 5.12 illustrates how the




Figure 5.12: Vertical distributions of the horizontally averaged buoyancy, b , for $t = 100$ (red), $t = 200$ (green), $t = 300$ (blue), $t = 400$ (magenta).




Figure 5.13: Vertical distributions of the horizontally averaged vertical buoyancy gradient, $d\langle b \rangle/dz$ for $t = 100$ (red), $t = 200$ (green), $t = 300$ (blue), $t = 400$ (magenta).

surface flux influences the vertical structure of the environment. Close to the lower surface, we see a small region where the air is unstable to convection (buoyancy decreases with height) due to the heating. Above this lies a mixed layer where the buoyancy is somewhat close to constant or grows slowly before curving to match the background stratification again as the profile crosses through the entrainment zone above.

There are some small variations in the horizontally averaged buoyancy at the very top of the domain at later times. The buoyancy at the upper surface appears to decay in time. It is unclear exactly what the cause of these is, but it may be a numerical effect. Alternatively, the decrease in buoyancy could be due to unforeseen dissipation caused by the vorticity damping at the upper surface. That said, it does not appear to be a significant result concerning the development of the boundary layer as a whole but is worth drawing attention to.

Computing the vertical gradient of the buoyancy field also offers some insight into the vertical structure of the boundary layer and entrainment zone. In these plots, the two-layer structure of the entrainment zone discussed by Garcia and Mellado (2014) is apparent. The main argument in favour of this structure is that the buoyancy gradient scales differently with time at different heights relative to the point at which its maximum occurs.

Garcia and Mellado (2014) observe that in the upper sublayer, the maximum value of the buoyancy gradient grows with time. This sublayer is centred on the maximum gradient's height. The lower sublayer is defined as the region between the height of the maximum buoyancy gradient, shown in Figure 5.13, and the zero-crossing point of the buoyancy flux, visible in Figure 5.11. The length and buoyancy scales evolve differently in each of these layers. Our results in Figure 5.13 partially support this idea because the magnitude of the observed peak in the buoyancy gradient grows with time. The thickness of the region where the buoyancy gradient exceeds the background stratification before returning to the background level also increases in the later stages of the growth of the boundary layer.

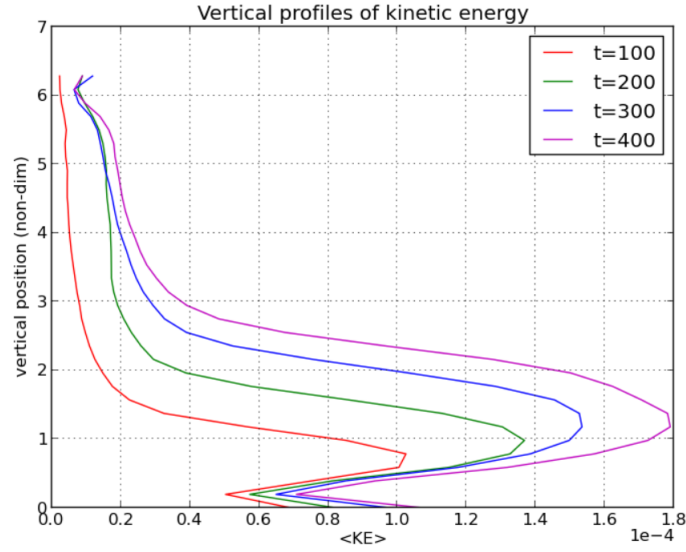


Figure 5.14: Vertical distributions of the horizontally averaged kinetic energy, $\langle KE \rangle$, for $t = 100$ (red), $t = 200$ (green), $t = 300$ (blue), $t = 400$ (magenta).

The region below this shows the vertical buoyancy gradient at each stage in the boundary layer's development, all evolving parallel to one another. The curves effectively shift upwards as the system evolves, converging towards zero at almost the same height regardless of the stage in the simulation. Below this is the unstable region where the buoyancy gradient is negative. The region in which the different evolution stages diverge is the mixed layer, which grows in time. As the system evolves, its buoyancy gradient becomes closer to zero. It follows that the first sublayer of the entrainment zone likely begins at the top of this mixed layer, where the vertical profiles of buoyancy gradient are parallel to one another. The upper sublayer represents a transition region between this layer and the stratified atmosphere above the entrainment zone. The existence of these sublayers is not immediately apparent from the profiles alone. Analysis of kinetic energy and vorticity-related properties may further support this view of the entrainment zone.

5.3.5 Vertical distribution of kinetic energy

To better understand the structure of the boundary layer and entrainment zone, it is beneficial to study the kinetic energy associated with the air motions in the system. The evolution of kinetic energy distribution in Figure 5.14 shows that the kinetic energy within the boundary layer grows over time. A more expansive, well-mixed layer makes it much easier for the warm air in thermals and plumes to rise through, accelerating upwards towards the stratified upper region of the domain. We would expect that the point at which the kinetic begins to decay is the base of the entrainment zone. Hence, the kinetic energy profiles show a smooth rise and fall with height.

At the lower surface, the kinetic energy is non-zero due to the free-slip boundary conditions used, but it decreases with height immediately above this. This is a consequence of the vorticity damping described in Section 5.2, which prevents extreme vorticity growth at the lower surface. Failing to damp the vorticity would give rise to intense horizontal motions and consequent shear layers dominating the system's dynamics.

To better understand the picture within the boundary and entrainment zone, it can be helpful to decompose the kinetic energy into components associated with the horizontal and vertical motions of the flow. Examining the lower surface first in Figure 5.15, we can see that the damping of vorticity here is working as intended. Only the horizontal motions decay with height in the lowest grid cells, and while there is some evidence of a potential shear layer, it does not grow to the extent that it dominates the motions across the whole domain. The kinetic energy associated with vertical motions appears largely unaffected and decays smoothly to zero at the lower surface.

There is a small region between $z = 1.85$ and $z = 2.75$ within which the horizontal kinetic energy exceeds the vertical component. As the stratified environment begins to inhibit the vertical motion of rising air, it decelerates, and the vertical component decreases. At the same time, more air from below is pushed up underneath it, and it is forced to spread out horizontally once it can rise no further.




Figure 5.15: Vertical distributions of the horizontally averaged kinetic energy (black line) along with its vertical (blue line) and horizontal (red line) components for $t = 300$.

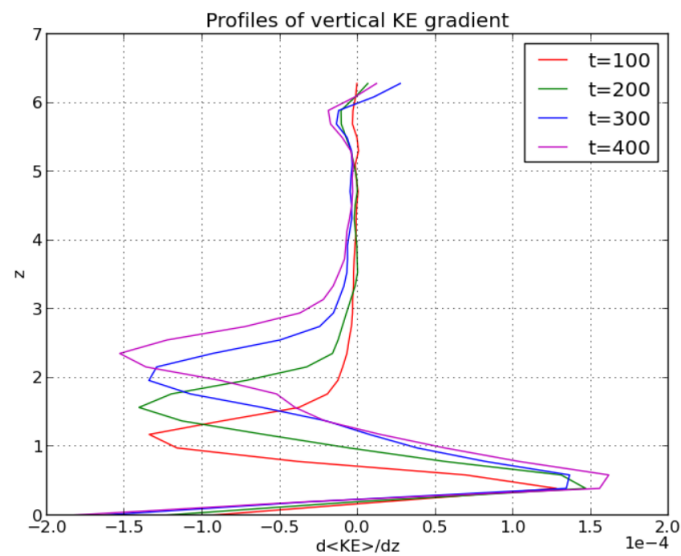


Figure 5.16: Vertical distributions of the horizontally averaged kinetic energy gradient for $t = 100$ (red), $t = 200$ (green), $t = 300$ (blue), $t = 400$ (magenta).

The negative buoyancy flux in Figure 5.11 also overlaps with this region of the flow, although the zero-crossing point of $\langle b'w' \rangle$ lies below the intersection of the kinetic energy components at $z = 1.5$. The base of this region is better aligned with the minimum in the buoyancy flux profile at the corresponding time. This may suggest that this layer is linked to the driving of entrainment, as the rising air begins to move out horizontally and forces ambient air at the same height downwards. Relating this to the two sublayer structure in Garcia and Mellado (2014) and Haghshenas and Mellado (2019), it may also be possible that this is part of the transitioning upper sublayer as it encompasses the peak in buoyancy gradient found at $t = 300$ at $z = 2.5$.

To further investigate the structure of this region, the behaviour of the vertical gradient of the kinetic energy is shown in Figure 5.16. The behaviour of this gradient is interesting, as it appears to oscillate between positive and negative peaks before dropping to zero in the upper half of the domain. The height at which the minimum occurs rises in time as the boundary layer grows, and the negative regions are similar in shape, effectively being vertical translations of one another. The thickness of the negative kinetic energy gradient regions suggests that the decline occurs gradually over a large vertical range, as observed in the smooth profiles of kinetic energy in Figure 5.14 and Figure 5.15. This gradient's zero-crossing point consistently occurs well below the zero-crossing point of buoyancy flux.

5.3.6 Vertical distribution of enstrophy and related quantities

Analysis of the enstrophy distribution across the boundary layer has been done in other studies of the entrainment zone and in studies of the interface between turbulent and nonturbulent regions (da Silva et al. 2014b; da Silva et al. 2014a; Fodor and Mellado 2020). Comparing results obtained in MPIC to the results of such studies using DNS is essential to assess the model's performance. At the same time, evidence of the two-layer entrainment zone structure proposed in Garcia and Mellado (2014) and Haghshenas and Mellado (2019) found within MPIC provides

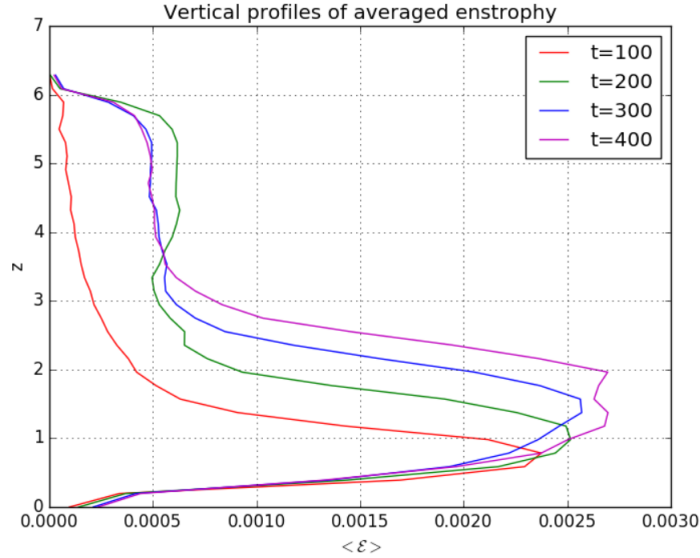


Figure 5.17: Vertical distributions of the horizontally averaged enstrophy for $t = 100$ (red), $t = 200$ (green), $t = 300$ (blue), $t = 400$ (magenta).

validation despite the different modelling approach used in our study. We use the definition of enstrophy from Pope (2000) as $\mathcal{E} = \boldsymbol{\omega} \cdot \boldsymbol{\omega}$.

Studies of enstrophy in the boundary layer have recorded that it drops dramatically across the entrainment zone. This is observed in the vertical profile of the horizontally averaged enstrophy in Figure 5.17. These plots show a region of high enstrophy located within the mixed layer that falls off rapidly with height above. As the boundary layer grows, this region becomes considerably broader vertically. Despite this, the peak enstrophy contained within does not grow considerably throughout the simulation.

Compared to the enstrophy PDF shown in Figure 1 of Fodor and Mellado (2020), we see that many of the same features appear in both. An equivalent stage in the boundary layer development to this figure would be around $t = 300$ in the MPIC simulation. In particular, the enstrophy decays by around a factor of 5 across the entrainment zone, a much smaller drop than occurs within their study. This difference of scales between the enstrophy in the turbulent and nonturbulent regions validates the estimates of the Reynolds number earlier in this work, as it is




Figure 5.18: Vertical distributions of the horizontally averaged vertical enstrophy flux for $t = 100$ (red), $t = 200$ (green), $t = 300$ (blue), $t = 400$ (magenta).

thought to be directly proportional to the Reynolds number of the system (Fodor and Mellado 2020). In this study, we see that the factor of 5 difference in the horizontally averaged enstrophy of the boundary layer and the free atmosphere is comparable to the estimates of $Re_0 = 8.5$ or $Re_0 = 3.1$ obtained in Section 5.3.1.

Another point of note is that in Fodor and Mellado (2020), the enstrophy within the boundary layer is approximately constant. It grows at the lower surface in the simulation with a Froude number of 20. This highlights a core difference between the two models, as the MPIC simulation uses a free-slip boundary condition that diffuses vorticity in the lower regions of the domain. Meanwhile, the typical DNS and LES simulations use a no-slip boundary at the lower surface and a free-slip (stress-free) upper boundary. Within the MPIC framework, implementing no-slip boundary would require substantial changes to the numerical solver, possibly through the implementation of virtual parcels or otherwise as is done in Macià et al. (2011) for SPH. The implementation of no-slip boundaries in particle-based methods is less straightforward than for their gridded counterparts. This may however be worth attempting in future work to allow for better comparisons to

existing LES studies. The damping in MPIC causes the enstrophy in the mixed layer to decay towards zero at both surfaces.

The lower vertical resolution of the MPIC study exacerbates these problems. With only 32 vertical grid cells, the vorticity diffusion layer at the lower surface extends through approximately 6% of the domain height. As the resolution increases, the size of the diffusion layer will decrease, and its impact on the evolution of the boundary layer will be reduced. Whether the diffusion strength is too great is also worth investigating, although the fact that we achieve comparable results despite the model differences is noteworthy.

Early studies with MPIC revealed that the vorticity field in the model displays a sensitivity to resolution (Dritschel et al. 2018, Böing et al. 2019). The small-scale buoyancy gradients generate more intense, localised vorticity, so studies at higher resolutions may yield a much greater scale separation than observed here. Local buoyancy gradients within the boundary layer are likely to be much stronger than in the stratified region above the entrainment zone, which may also impact the enstrophy scales, but this remains to be investigated.

It also appears that MPIC overestimates the vorticity in the upper region of the domain. This may be due to the absence of a sponge layer in MPIC, as the only relaxation in the domain is the damping on vorticity at the upper and lower boundaries. Figure 5.18 shows the horizontally averaged vertical flux of enstrophy and indicates that enstrophy is not being carried upwards above the entrainment zone by fluid motions. This flux is calculated in the same manner as the vertical buoyancy flux,

$$\mathcal{F} = \langle \mathcal{E}' w' \rangle, \quad (5.21)$$

where \mathcal{F} is used to denote the vertical flux.

The profiles show a similarly broad region within the mixed layer as the original profiles of enstrophy. Across the entrainment zone, they decay rapidly to zero and remain there in the free atmosphere above.

One of the key features shown in Figure 1 of Fodor and Mellado (2020) is the




Figure 5.19: Vertical distributions of the horizontally averaged vertical enstrophy gradient for $t = 100$ (red), $t = 200$ (green), $t = 300$ (blue), $t = 400$ (magenta).

saddle point in the enstrophy PDF (denoted by a star in their work). This saddle point is used to determine the enstrophy threshold for the conditioning in their study. We can aim to find something similar by studying the vertical gradient of the enstrophy profile within MPIC (Figure 5.19). The minimum of this profile is an approximate equivalent to the saddle point and shows the point within the entrainment zone where the decay of enstrophy with height is most substantial. As expected, the enstrophy gradient is negative throughout the entrainment zone, and the minimum rises through the domain as the boundary layer grows.

In relation to the two-layer structure, the point of minimal enstrophy gradient may be relevant to the location of the upper sublayer and the transition between turbulent and nonturbulent air. We know that enstrophy experiences a rapid decline with height, so the point at which this is the strongest may denote the base of the upper sublayer, centred on the height of the maximum buoyancy gradient. Interestingly, as the boundary layer grows between $t = 100$ and, $t = 300$ the minimum value of the enstrophy gradient actually decreases (although it grows slightly between $t = 300$ and $t = 400$). In section 5.3.7, we will study in more

detail how the height at which this minimum occurs compares to other measures of the entrainment zone height.

Below, a region of near constant enstrophy forms towards later times, with a layer of high enstrophy gradient at the base due to the diffusion layer. The profiles of the enstrophy gradient also have a similar shape to those of buoyancy flux (and, naturally, the kinetic energy gradient), although it should be noted that the enstrophy gradient does not grow linearly across the mixed layer.

5.3.7 Measures of boundary layer depth

The vertical profiles shown previously have been interpolated quadratically in post-processing to a higher resolution to give a more smooth variation of the various properties used to measure the heights used in the following subsections. This has a minimal effect on the actual results, but since less than half of the domain is occupied by the boundary layer and entrainment zone, using fewer than 16 grid points makes it much harder to distinguish between different heights.

When discussing measures of boundary layer depth, it seems only fair to start from the encroachment height (Lilly 1968; Carson and Smith 1975) to estimate where the base of the entrainment zone should lie. This can be obtained from an expression similar to equation 5 of Garcia and Mellado (2014) where the boundary layer depth is quantified in terms of the total energy input into the system,

$$z_{enc}(t) = \left[\frac{2B_0}{N^2} (t - t_0) \right]^{\frac{1}{2}}, \quad (5.22)$$

where the symbols have their usual meaning. t_0 is a virtual origin time chosen to be after the initial transient. The total buoyancy entering the system due to the surface flux between can also be written as,

$$B_0(t - t_0) = \int_0^{z_\infty} \langle b \rangle(z, t) - b_0(z) dz, \quad (5.23)$$

in the absence of dissipation. z_∞ is chosen to be far into the stratified region (i.e.




Figure 5.20: Time evolution of the right-hand side of equation 5.23.

the top of the domain. Choosing $t_0 = 0$ means that this equation allows us to represent the boundary layer depth in terms of the total buoyancy entering the system. Combining this with equation 5.22 gives us a modified version of equation 7 in Van Heerwaarden and Mellado (2016),

$$z_{enc}(t) = \left[\frac{2}{N^2} \int_0^{z_\infty} \langle b \rangle(z, t) - b_0(z) dz \right]^{\frac{1}{2}}. \quad (5.24)$$

The reason we use the integrated buoyancy difference throughout the domain is because this gives a more accurate representation. Computing $B_0 t$ gives the total energy input from the surface flux, but integrating the buoyancy difference over the domain also accounts for energy lost to dissipation. Figure 5.20 shows the time evolution of the right-hand side of equation 5.23, which determines the encroachment height in our simulations. The graph follows an almost straight line, except for a minor blip at the beginning ($t = 35$ to 40). This is around the time when the initial thermals take shape, as seen in Figure 5.6, and is either the result of the initial transient developing into a convective instability or the vorticity diffusion layer inhibiting the buoyancy growth beyond the lower two grid

cells. It is easy to estimate z_{enc} . Using equation 5.24, we no longer need to specify the virtual origin time, t_0 . This method also gives a more realistic impression of what is happening within the system since it is based on the buoyancy field as opposed to just the input flux and effective Reynolds number.

In order to evaluate the growth of the boundary layer compared with the encroachment height, we need to define some measure of depth. Thankfully, there are multiple measures used throughout the literature, primarily based on the vertical profiles of buoyancy and buoyancy flux (Fedorovich et al. 2004, Garcia and Mellado 2014). So we use some of these while adding two additional measures based on the vertical distribution of enstrophy and its gradient. The different measures used are as follows:

1. z_{if} : The conventional measure of entrainment zone height used in the zero-order model. This is where the minimum value of $\langle b'w' \rangle$ occurs.
2. z_i : The zero-crossing point of $\langle b'w' \rangle$. This typically denotes the base of the entrainment zone.
3. z_{ig} : The height at which the maximum value of vertical buoyancy gradient, $d\langle b \rangle/dz$, occurs. The upper entrainment zone sublayer is centred on this height according to Garcia and Mellado (2014).
4. z_{iE} : The height at which enstrophy reaches its maximum value.
5. $z_{iE_{gm}}$: The height at which the minimum value of the enstrophy gradient occurs. This represents an approximation to the saddle point discussed in Fodor and Mellado (2020).

These are shown along with the encroachment height in Figure 5.21. It is immediately apparent that the encroachment height and the zero-crossing height of the buoyancy flux are very closely aligned. This is particularly notable as it reaffirms that the model has accurately estimated the growth of the boundary layer itself (and the base of the entrainment zone) despite the differences in its formulation and limited spatial resolution. The other heights appear to follow a very similar time-dependence to z_{enc} , which is also reassuring, and we visualise this by




Figure 5.21: Time evolution of the various heights relevant to the entrainment zone and boundary layer structure.

normalising each height with respect to z_{enc} in Figure 5.22.

We disregard results from particularly early times (pre $t = 50$) because these are primarily a result of the initial heating of the lower region of the domain. The large value of the buoyancy gradient-based height at this time is likely an anomaly due to the initial setup. This issue is resolved by the time the initial thermals form. The substantial spike in z_{if} likely captures a significant overshoot by the thermals causing significant initial entrainment.

The height of maximum enstrophy remains consistently below the encroachment height, which is to be expected given that the decay of enstrophy occurs in the entrainment zone. Meanwhile, the minimum enstrophy gradient height $z_{iE_{gm}}$ starts below the height of minimum buoyancy flux z_{if} before rising above it by $t = 150$ and remaining only slightly above until the end of the simulation. This is consistent with observations in Fodor and Mellado (2020) in which the saddle point was found to lie between z_{if} and z_{ig} .

From Figure 5.22, we see that most of the heights, except for z_{iE} , seem to become




Figure 5.22: Time evolution of the various heights relevant to the entrainment zone and boundary layer structure, normalised by the encroachment height z_{enc} .

nearly constant with respect to z_{enc} as the system evolves. The fact that this is also true for z_{iEgm} lends some strength to the argument that it is representative of one of the sublayers, possibly the base of the upper sublayer in Garcia and Mellado (2014). This suggests that each height can be written similarly to Garcia and Mellado (2014) as a constant multiplying the encroachment height. For example, the buoyancy flux minimum height is given by

$$z_{if} = C_{if} z_{enc}, \quad (5.25)$$

where C_{if} is the approximate ratio between z_{if} and z_{enc} at later times in the simulation. The fact that all the previously listed heights associated with the entrainment zone tend towards the $t^{1/2}$ scaling of z_{enc} indicates that the boundary layer is within the equilibrium entrainment regime. The ratio (z_{enc}/L_0) exceeding an order of magnitude is shown to be an indicator of this in Fedorovich et al. (2004) and Garcia and Mellado (2014) and we cross the threshold of $z_{enc}/L_0 = 10$ by a time of approximately $t = 75$.

There are some discrepancies with the results shown in Figure 1 of Garcia and Mellado (2014). The ratios between each relevant height and z_{enc} are considerably larger in MPIC. This suggests that the entrainment zone produced in this work is actually considerably thicker than in LES and DNS studies (Fedorovich et al. 2004; Garcia and Mellado 2014), even though much of the overall structure remains similar. This may be due to the lack of a sponge layer in the MPIC simulation, which may inhibit the vertical growth somewhat depending on where it is located.

5.3.8 Estimates of sublayer thickness in the entrainment zone

We can also attempt to estimate measures of thickness for the two overlapping sublayers that comprise the entrainment zone using the heights we have already obtained. We can use the buoyancy gradient thickness definition (equation 21 in Garcia and Mellado 2014, see also Pope 2000):

$$\delta_b = \frac{\langle b \rangle(z_{ig}) - b(z_{ig})}{(\partial \langle b \rangle / \partial z)(z_{ig}) - N^2}. \quad (5.26)$$

We also include the thickness of the lower sublayer, simply given by $z_{ig} - z_{enc}$, along with two more measures based on the enstrophy gradient. The first is simply the difference between the enstrophy gradient-based height and the buoyancy gradient-based height, $z_{ig} - z_{iE_{gm}}$, chosen for the purpose of comparison to δ_b to investigate whether $z_{iE_{gm}}$ could potentially be considered the base of the upper sublayer. The second is the vertical distance between $z_{iE_{gm}}$ and the point at which the negative enstrophy gradient has returned to half of its minimum value. This may appear an arbitrary measure, but given that the upper sublayer is considered a transition region between the turbulent and nonturbulent fluid, it is worth examining a point at which the decay of enstrophy with height has slowed significantly.

It is expected that the thickness of the first sublayer will grow proportionally to z_{enc} and that the thickness of the upper sublayer will decrease with respect to z_{enc} and this is observed in Figure 5.23. However, what is less expected is how closely the three, independently calculated, measures of the upper sublayer thickness align




Figure 5.23: Time evolution of the various measures of sublayer thickness, including the buoyancy gradient thickness definition in Equation 5.26 (black); the lower sublayer thickness (green); the enstrophy gradient based thickness (red); and the difference between z_{ig} and $z_{iE_{gm}}$.

with each other. All three maintain a near constant average between 0.35 and 0.4 with some fluctuations, although we do note that the buoyancy gradient definition appears to be growing very slowly. If the upper sublayer is believed to be centred on z_{ig} then this would mean that the enstrophy gradient minimum occurs roughly a distance of δ_b below it and would, in this case at least, represent the base of the upper sublayer. Simultaneously, after reaching its minimum value at $z_{iE_{gm}}$, the enstrophy gradient has returned to half of its minimum value at the centre of the upper sublayer.

These results further support the picture of the upper entrainment zone sublayer as a transition region. It does also bear keeping in mind that MPIC seems to overestimate the vertical extent of the entrainment zone, so further studies may wish to examine enstrophy and its vertical gradient for more typical model formulations (LES and DNS).




Figure 5.24: Time evolution of the mean entrainment rate, E , at height z_{if} (blue markers) on a log-log scale. The red line denotes the value of equation 5.30 with an entrainment ratio of $A = 0.1$.

5.3.9 Entrainment rate parameters

Finally, we attempt to derive entrainment rate parameters from the MPIC simulation. These parameters are useful in representing convection in larger scale models and will provide a useful comparison between the MPIC simulation and existing parametrisations of the entrainment zone, such as the zero-order model.

We can start by simply calculating the entrainment rate at a specific point in the entrainment zone. This is typically done at the height of minimum buoyancy flux, so we will use z_{if} where we have assumed that beyond $t = 75 - 100$ that $z_{if} \approx 1.25z_{enc}$ (i.e. it is constant with respect to z_{enc} , as in equation 5.25). This removes many of the small scale variations and makes the comparison to predicted entrainment rates much easier.

The mean entrainment rate at the height z_{if} is given by

$$E = \frac{w_{ef}}{w_{*f}} \quad (5.27)$$

where

$$w_e = \frac{dz_{if}}{dt}, \quad (5.28)$$

is the rate of growth of the measure of entrainment zone height and

$$w_{*f} = (B_0 z_{if})^{\frac{1}{3}}, \quad (5.29)$$

is the convective velocity at height z_{if} . A power law in the zero-order model for the mean entrainment rate is given in Fedorovich et al. (2004), for the equilibrium entrainment regime, as

$$E = z_i^{-\frac{1}{3}} \frac{dz_i}{dt} = \frac{1}{2} [2(1+A)]^{\frac{1}{3}} t^{-\frac{2}{3}}, \quad (5.30)$$

where A is the entrainment ratio. The mean entrainment rate for the MPIC simulation is shown in Figure 5.24. We compare this to the predicted entrainment rate from equation 5.30 for an entrainment ratio of $A = 0.1$, denoted by the red line. At late times, we see a very strong agreement between the observed entrainment rate in MPIC and the predicted value, although we must note that the typical value of A used in the zero-order models is $A = 0.2$ and is supported by LES studies carried out in Fedorovich et al. (2004), which suggests that MPIC is underestimating the entrainment. It is possible this may change at higher resolutions, since MPIC at the current resolution is operating at lower Reynolds numbers. The numerical mixing of parcels may also play a role.

We can investigate this further by calculating the entrainment ratio independently. A can be obtained from the expression

$$A = \frac{\Delta b_i}{B_0} \frac{dz_i}{dt}, \quad (5.31)$$

for a given entrainment zone height z_i , which we will represent by z_{if} . The term Δb_i refers to the buoyancy increment across the entrainment zone, shown in Figure 5.1. In the case of the zero-order model, this refers to the jump in buoyancy between z_{il} and z_{iu} (Fedorovich et al. 2004). This becomes more difficult to de-




Figure 5.25: Time evolution of total buoyancy increment (blue) and the local buoyancy increment at z_{if} (red).

termine in our simulation. We compute the buoyancy increment across the entire entrainment zone in MPIC, and this is shown in Figure 5.25. We measure the buoyancy increment between z_i and $z_{ig} \approx 1.5z_{enc}$, giving us

$$\Delta b_{tot} = N^2 (z_{ig} - z_i), \quad (5.32)$$

as a substitute for this value in MPIC. We also consider the local buoyancy increment at the height of the minimum buoyancy flux as an alternative. The difference between this increment and the total buoyancy increment is shown in Figure 5.1 and in Figure 10 of Garcia and Mellado 2014. δb is used to denote the local buoyancy increment in these examples, so we use δb_{if} to include the height at which it is defined. The main reason to consider this is that many of the zero-order and general structure model approximations are not valid for entrainment zones of considerable thickness (e.g. Sullivan et al. 1998). We have already observed that the boundary layer in MPIC has created an especially thick entrainment zone, and as such, the buoyancy increment across it, Δb_{tot} , is likely to be much larger than is expected. In Figure 5.25, we show a comparison between this and the local




Figure 5.26: Time evolution of the entrainment ratio, A , from equation 5.31, calculated using total buoyancy increment (blue) and the local buoyancy increment at z_{if} (red).

buoyancy increment at the height of minimum buoyancy flux, given by,

$$\delta b_{if} = \langle b \rangle (z_{if}) - b_0 (z_{if}). \quad (5.33)$$

As expected, the total buoyancy increment greatly exceeds the local increment. Because we are only interested in the equilibrium entrainment regime, we may ignore the exceedingly large values at early times, as this is likely measuring the increment across most of the domain until the boundary layer is clearly formed.

The resulting estimates of the entrainment ratio are shown in Figure 5.26. The ratio using the total buoyancy increment starts very large (probably due to the high initial value of z_{ig}) and decays towards a value of around 0.5 at the end of the run. Meanwhile, the entrainment ratio calculated using δb_{if} has a nearly constant value just above $A = 0.1$ from $t = 50$ onwards, rising slightly to around $A = 0.12$ by the end of the simulation.

Very clearly, the entrainment ratio predicted using the lower buoyancy increment

is much closer to the observed entrainment rate in Figure 5.24. This further suggests that MPIC underestimates entrainment compared to standard zero-order model predictions. Meanwhile, using the total buoyancy increment leads to an overestimation of entrainment. We can also compare our results in MPIC to lab-based experiments carried out in Jonker and Jiménez (2014) using a saline water tank. Figure 10 in Jonker and Jiménez (2014) shows that at the lowest Reynolds numbers in their experiments ($Re \approx 50 - 100$) they observe entrainment ratios in the range of $A = 0.8 - 1.3$. These values are very close to the results obtained at MPIC and lend some credibility to the notion that MPIC is underestimating entrainment due to a low effective Reynolds number. As such, future work should prioritise higher resolution simulations which can attain greater values of Re_0 to verify this.

With these increments calculated, we can also attempt to verify the Richardson number power law in MPIC. This power law directly relates the mean entrainment rate, and thus the growth rate of a boundary layer to the convective Richardson number. Several computational studies (Sullivan et al. 1998; Fedorovich et al. 2004; Garcia and Mellado 2014), observational data from radar aircraft (Träumner et al. 2011) and even some lab based studies in water tanks (i.e. Jonker and Jiménez 2014) have verified the existence of a power law-like relationship between the Richardson number and entrainment rate. The entrainment rate is related to a Richardson number by the expression

$$E = ARi_*^{-n}, \quad (5.34)$$

where $1 \leq n \leq 2$ and

$$Ri_* = \frac{z_{enc} \Delta b_{tot}}{w_*^2}, \quad (5.35)$$

is the convective Richardson number using the total buoyancy increment. A Richardson number for the local increment is also given by equation 60 in Garcia and Mellado (2014)

$$Ri_f = \frac{z_{enc} \delta b_{if}}{w_*^2}. \quad (5.36)$$

The value of n is typically understood to be close to $n = 1$ (Fedorovich et al. 2004;) and so in Figure 5.27, we show plots of the relationship between E and Ri for Richardson numbers calculated using the total and local buoyancy increments alongside one another similar to Figure 11 in Garcia and Mellado (2014).

From this we can see that, for intermediate values of the Richardson number between approximately $Ri = 0.2$ and $Ri = 35$, the MPIC simulation recovers the power law relatively well (with some outliers). However, entrainment rates predicted by the power law using the total buoyancy increment are estimated using an entrainment ratio of $A = 0.55$. This is more than double the typical value found in zero-order model studies using LES ($A = 0.2$). Meanwhile, the power law predictions of entrainment rate using the local buoyancy increment use $A = 0.1$. Therefore, we must conclude that while the MPIC simulation of the shear-free boundary layer is effective in recovering many of the properties of entrainment, it still does not reliably reproduce the results predicted by common representations of the entrainment zone. The extended entrainment zone in MPIC may, at least in part, explain why the results in Figures 5.26 and 5.27 predict a much larger entrainment ratio when using the total buoyancy increment.

On the other hand, the results using the local buoyancy increment are very similar to those obtained in Figure 10 of Garcia and Mellado (2014), which also compared the entrainment rate to $0.10Ri_f^{-1}$. This is reassuring as many of the results obtained using MPIC have been aligned with the predictions in their work, particularly relating to the two sublayer structure of the entrainment zone. It is also promising that the entrainment rate—Richardson number relationship recovered is so similar to their results, despite the difference in numerical method and spatial resolution of the two models.

5.3.10 Sensitivity to numerical mixing

Studying the vertical profiles of buoyancy and buoyancy flux near the boundaries in previous sections has revealed some unusual behaviour (see Figure 5.11 and 5.12). For example, the vertical flux of buoyancy at the surface appears to fall off rapidly.




Figure 5.27: Relationship between entrainment rate, E , and Richardson number, Ri , computed for both the local buoyancy increment (red markers) and total buoyancy increment (blue markers). The blue line corresponds to $E = 0.1 Ri_f^{-1}$ while the green line corresponds to $E = 0.55 Ri_*^{-1}$ and shows how the data compares to the predictions of the power law using rough estimates of A for each buoyancy increment.




Figure 5.28: Vertical distributions of the horizontally averaged buoyancy flux, $\langle b'w' \rangle$ for the simulation with $\gamma_{max} = 4.0$ at $t = 100$ (red), $t = 200$ (green), $t = 300$ (blue), $t = 400$ (magenta).

Similarly, the buoyancy at in the top grid cell appears to change as the boundary layer has been allowed to grow too. In this section, we aim to investigate if and how these effects are linked to numerical choices made during the development of the model. To this end, we have carried out a second experiment with a different value of the parcel splitting threshold, γ_{max} , as was used to investigate energy conservation in Chapter 4. In the previous results, we have used a value of $\gamma_{max} = 8, 0$, double what was used in the original works on MPIC. This was chosen as we found better conservation of energy in Chapter 4 when the threshold was increased.

Here, we use a value of $\gamma_{max} = 4.0$ to allow mixing to occur more readily within the simulation and examine how the results obtained respond. Doing so is useful to offer further insight into the sensitivity of results obtained using MPIC in the future to the choice of numerical parameters related to mixing. Moreover, results of these investigations can also serve to provide guidelines for tuning these parameters in future studies using MPIC.




Figure 5.29: Vertical distributions of the horizontally averaged buoyancy, $\langle b \rangle$, for the simulation with $\gamma_{max} = 4.0$ at $t = 100$ (red), $t = 200$ (green), $t = 300$ (blue), $t = 400$ (magenta).

Figure 5.28 shows the buoyancy flux profiles for this test case in at $t = 100, 200, 300, 400$ as was done for the previous test case. Overall, these results are largely very similar to the observations. Although the flux profile at the later stages in the simulation does start to deviate from the previous simulation, in particular, at $t = 300$ and 400 in the centre of the entrainment zone. This makes some sense as the flow at this stage of the simulation has become highly turbulent and thus the role of numerical mixing is likely to become increasingly significant. Whether this significantly impacts the growth of the boundary layer and the entrainment rate parameters will be investigated further later in this section.

Due to the requirement in MPIC that the vertical velocity at the lower surface must be zero, we see $\langle b'w' \rangle$ decays to zero at this point as in the previous case, but if this were the sole reason we would expect a sharp decay to zero between the first grid cell and the lower surface. However, as noted previously, the decay in buoyancy flux actually starts above this point, This is very similar to what is observed in the first test case with $\gamma_{max} = 8.0$ and the increase in parcel




Figure 5.30: Vertical distributions of the horizontally averaged buoyancy gradient, $\langle db/dz \rangle$, for the simulation with $\gamma_{max} = 4.0$ at $t = 100$ (red), $t = 200$ (green), $t = 300$ (blue), $t = 400$ (magenta).

splitting/mixing in this simulation appears to have very little influence on this. The peak in buoyancy flux seems to start at a height of approximately $z = 0.4$. This aligns with the top of the second grid cell in the model, since $2\pi/32 = 0.19$, which is also the top of the vorticity damping layer. With parcel splitting ruled out as a cause of this error, we must consider that the damping layer is somehow influencing the buoyancy flux in the lowest two grid cells of the domain. Further investigations into this effect should aim to increase the resolution and study whether the region in which buoyancy flux decays shrinks with the damping layer. If this is the case, then it may be worthwhile to study exactly how strongly the damping method described in this work impacts the buoyancy fluxes and consequently the growth of the boundary layer as a whole.

Above the entrainment zone, the buoyancy flux behaves similarly to what was observed in the $\gamma_{max} = 8.0$ test case too. This suggests that the positive buoyancy flux above the boundary layer are likely not a consequence of the splitting procedure, but a further investigation may wish to completely disable parcel splitting to




Figure 5.31: Time evolution of the various heights relevant to the entrainment zone and boundary layer structure for the simulation with $\gamma_{max} = 4.0$.

see if this has a more dramatic effect. Variations in the behaviour of the damping layer at the upper boundary may also be worth investigating. Similarly, resolution studies would also be crucial.

While on the topic of the upper boundary, we show the horizontally averaged buoyancy profile in Figure 5.29. These are near identical to those observed in the previous simulation, including the unusual changes at the top of the simulation. The slope of the buoyancy profile at the upper surface seems to be decaying in time, although the region in which this occurs also seems to align with the vorticity damping layer, this time at the top of the domain. Because of the similarity of these profiles to the previous simulation, it appears highly unlikely that the horizontally averaged buoyancy profile is especially sensitive to the choice of splitting threshold. However, a study at higher (or even lower) resolution would likely reveal whether this behaviour is directly related to the vorticity damping.

We show a profile of the buoyancy gradient in Figure 5.30 as a supplement, as this clearly shows the region in which the changes start to occur. This highlights that




Figure 5.32: Time evolution of the mean entrainment rate, E , at height z_{if} (blue markers) on a log-log scale for the simulation using $\gamma_{max} = 4.0$. The red line denotes the value of equation 5.25 with an entrainment ratio of $A = 0.1$.

the region in which buoyancy gradient decays extends well beyond the damping layer. Compared to the previous simulation, the decay at later times appears to be slightly smoother, although the overall effect is largely the same. As before, it may be worth revisiting this at higher resolution and with splitting completely switched off to see how this impacts the vertical structure of buoyancy. The sensitivity to splitting and mixing may also become more apparent at higher resolutions too.

Figure 5.31 shows the evolution of the various heights associated with the entrainment zone used in previous sections for the simulation with $\gamma_{max} = 4.0$. The overall results are highly similar to the previous case, and we still see a strong agreement between the encroachment height, z_{enc} and the height of zero buoyancy flux, z_i . This figure suggests that the overall growth of the boundary layer and its bulk properties are not particularly sensitive to the choice of splitting threshold, although there may be some deviation in particularly turbulent regions of the flow where splitting occurs more frequently.

To take this a step further, we compute the mean entrainment rate and compare

this to the predictions made in section 5.3.9 (equation 5.30). The results are shown in Figure 5.32 and compared to the same estimated entrainment ratio of $A = 0.1$ and we see that the resulting entrainment rate is shows a near identical agreement, as in the previous section. This implies that the mean entrainment rate and thus the rate of growth of the boundary layer are largely insensitive to the parcel splitting threshold.

We also attempt to recover the Richardson number power law as in the previous simulation for these new results. While the mean entrainment rate in Figure 5.32 shows a strong agreement, Figure 5.33 shows that the entrainment ratios, A , and power scaling of the power law, n , require some slight tuning to produce better agreement with the results in this simulation. This is likely due to the differences in the minima of the vertical buoyancy flux observed in Figure 5.28. Here we find that the relationship for the local buoyancy increment is closer to $E = 0.1 Ri_f^{1.05}$ and similarly, for the local buoyancy increment, $E = 0.6 Ri^{1.05}$. This still lies somewhat within the predictions of the power law, as n is typically expected to be between 1 and 2. However, the predicted entrainment ratio and scaling are still larger, revealing a sensitivity to the parcel mixing.

Further work into the numerical sensitivities of boundary layer simulations in MPIC should aim to investigate the effects of increasing resolution and potentially turning off the parcel splitting entirely. The depth of the damping layer for vorticity may also be worth investigating, as this presently depends on the resolution of the test case. Avenues for exploration here could be extending the layer to higher grid cells and studying if the anomalous behaviour in the buoyancy and flux profiles change in response, or simply increasing resolution. Alternatively, defining the damping layer in terms of a fixed, physical depth instead of being resolution-dependent as it is currently may also be worth studying. The results obtained in this section suggest that there are still some numerical sensitivities within MPIC when it is applied to boundary layer studies that would be beneficial to understand in greater detail if the method is to be applied to more ambitious test cases. such as cloud-topped boundary layers or sheared cases.




Figure 5.33: Relationship between entrainment rate, E , and Richardson number, Ri , computed for both the local buoyancy increment (red markers) and total buoyancy increment (blue markers) for the simulation using $\gamma_{max} = 4.0$. The blue line corresponds to $E = 0.1 Ri_f^{-1.05}$ while the green line corresponds to $E = 0.6 Ri_*^{-1.05}$ and shows how the data compares to the predictions of the power law using rough estimates of A for each buoyancy increment. Note that we use $n = 1.05$ instead of $n = 1$ here unlike in the analysis of the previous simulation.

5.4 Conclusions

The results in this chapter demonstrate the ability of MPIC to model the growth of a convective boundary layer into a linearly stratified atmosphere. For the MPIC method, this marks a major step forward as it has previously only been used in the study of a single rising thermal. The implementation of surface fluxes of heat can easily be applied to the moisture field and could allow for the study of cloud fields resulting from these fluxes alone. A good next step may be to use MPIC to study the moisture distribution across a boundary layer or even cloud field as it evolves in time, similarly to Mellado et al. (2017). While the effectiveness of the free-slip boundary conditions with a vorticity diffusion layer is yet to be further validated, it likely provides a useful stepping stone toward more realistic or conventional boundary conditions in future studies.

The model is not without its shortcomings, being different in the numerical method to conventional LES and DNS models, and we see this through the greater vertical extent of the entrainment zone formed. Similarly, the low grid resolution of this study means that the results obtained are only at low Reynolds numbers in the range $Re_0 = 3.1 - 8.5$. Consequently, we do not attain the same separation of scales in the enstrophy field observed in other studies of turbulent and nonturbulent flow regions. We observe a much higher averaged enstrophy than occurs here in other studies (Fodor and Mellado 2020), but it is possible that higher resolution simulations in MPIC would improve this issue, given that the estimates of Re_0 are largely dependent on the minimum parcel size, which in turn depends on the grid spacing. It is also possible that these anomalies arise from other numerical effects, but this warrants further investigation into the effects of resolution and numerical mixing. Similarly, MPIC experiences some resolution dependence in the vorticity field, which may also see the difference in enstrophy between the two regions change significantly at higher resolutions.

We were able to provide a range of estimates of the depth of the entrainment zone and estimates of the thickness of the upper sublayer. These results illustrate that the lower boundary of the entrainment zone, defined as the zero-crossing height

of the buoyancy flux profile, shows excellent agreement with the encroachment height predicted by equation 5.24. This indicates that MPIC accurately models the growth of the boundary layer base into the atmosphere. The other measures of entrainment zone depth, however, occur at greater heights relative to this, meaning that MPIC is producing a thicker entrainment zone than is observed in previous works.

The vertical gradient of enstrophy is also studied to provide an approximation to the saddle point used in studies that condition air into turbulent and nonturbulent air based on an enstrophy threshold. The height at which the vertical enstrophy gradient is at its minimum value, $z_{iE_{gm}}$, is found to lie between the heights of minimum buoyancy flux and maximum buoyancy gradient. This grows only slightly with respect to z_{enc} in the equilibrium regime, being nearly constant. We propose that this point may in some way be linked to the upper sublayer proposed in previous studies. In Figure 5.23, we do indeed find that the separation between this point and the centre of the upper sublayer located at z_{ig} is close to the buoyancy gradient thickness of the upper sublayer obtained from equation 21. The significance of this is further illustrated when including the distance from $z_{iE_{gm}}$ at which the enstrophy gradient has returned to half of its minimum value, as this also agrees with the aforementioned measures of sublayer thickness. These measures are close to constant in time and thus decay with respect to the thickness of the lower sublayer, which is directly proportional to z_{enc} , again in agreement with previous studies (Borrell and Jiménez 2016; Fodor and Mellado 2020).

The importance of these results is that even with a very different model and boundary conditions to other boundary layer simulations, we are still able to observe evidence of this structure in the entrainment zone with MPIC. And while there are some differences in the size of the entrainment zone, much of its structure remains, and it appears that MPIC accurately models the growth of the regions between z_{enc} and z_{if} , the lower sublayer.

Computing entrainment parameters for this simulation reproduces some well-known power laws associated with the equilibrium entrainment regime (equations

5.25 and 5.29). And studying the time evolution of the entrainment rate alone suggests that MPIC may be underestimating entrainment, with a value of $A \approx 0.10$ being found in Figure 5.24. We then directly computed the entrainment ratio from the global and local buoyancy increments, and found values in the ranges $A = 0.10 - 0.12$ and $A = 0.5 - 0.6$, respectively. This would suggest that MPIC's thicker entrainment zone results in a much greater entrainment rate across the whole layer, whereas the local entrainment ratio at the height of minimum buoyancy flux is closer to estimated values in other works (Garcia and Mellado 2014). We then computed the relationship between the Richardson number and entrainment rate, using two different Richardson numbers based on the different buoyancy increments. The power law in equation 5.29 was very closely recovered in both cases for $A = 0.10$ for the local buoyancy increment and $A = 0.55$ for the global increment. The latter value represents a significant overestimation of A in comparison to previous works, but we note that the local increment relationship is very close to that obtained in Figure 10 of Garcia and Mellado (2014). This suggests that, at least around height z_{if} , MPIC is producing similar estimates of entrainment to DNS.

The discrepancies between MPIC and other studies seem to lie in the upper regions of the entrainment zone, where it is much less compact. Nonetheless, this chapter shows that a semi-Lagrangian approach to modelling convection in a vorticity dynamics framework can still reproduce many of the features of convective boundary layer growth. Subsequent work with MPIC should strive to perform studies at higher resolutions and seek to probe the impacts of wind shear on the boundary layer growth to see if any of the established relationships between shear strength and entrainment zone structure are maintained in this framework. Alternatively, it may be worth studying the distributions of moisture starting from a dry domain subject only to surface fluxes and assessing the point at which cloud formation can begin to occur. By altering the environmental structure as discussed in Chapter 4, for example, we could study the formation of flux driven convective clouds in the presence of convective inhibition to assess how overcoming this through surface heating influences the intensity of convection. Furthermore,

studies into boundary layer entrainment could take advantage of the Lagrangian nature of MPIC by implementing tracer populations to track the origin of air in the entrainment zone with minimal additional computational expense. This could similarly be applied to chemical transport problems by adding surface fluxes of chemical populations. This would offer some insight in the vertical transport of aerosols and chemical species into the upper atmosphere and how this responds to wind shear and changes in surface heating profile. Refinements to resolution and numerical effects in MPIC will also allow further investigation into the transition regime of the entrainment zone and its structure.

This work showcases the beginning of many possible applications of the model now that it no longer needs to be initialised with a thermal already present. Future enhancements and extensions to the method will open the way for many more applications of Lagrangian models to studies of convection. This represents a major step forward for the MPIC method that can be built upon in future works.

Chapter 6

Summary and Outlook

In this chapter, we summarise the results of the work carried out for this thesis and the outlook for Lagrangian modelling of atmospheric convection, in particular, using the MPIC model.

Developing our understanding of convective processes in the atmosphere is essential to improving weather forecasting and climate prediction in the future. On the global scale, these processes are represented by parameterisation schemes and are typically refined through explicit convection simulations via LES or DNS. While conventional Eulerian models like LES have proven very effective in modelling atmospheric convection, they are not without their shortcomings. To this end, we have endeavoured to demonstrate the promise of a Lagrangian approach to the problem by applying the MPIC model to case studies of isolated thermals and the boundary layer. In this final chapter, we will start by revisiting the key findings from studies of convection in MPIC, despite its currently idealised state. Then we will present ideas for additional extensions and enhancements of the model, particularly concerning the surface flux implementation described in Chapter 5.

Chapter 3 presents the simple test case of a rising thermal subject to constant vertical wind shear. Wind shear at different altitudes is known to impact the evolution of convective clouds and even contribute to the organisation of convective

systems, such as squall lines (Rotunno et al. 1988; Takemi 2006). Many studies primarily focus on the larger-scale impacts of shear. Our analysis of its effects on an isolated thermal may be relevant to some convective schemes that seek to represent convection through the statistical properties of an ensemble of plumes (Arakawa and Schubert 1974). The simple approach of using a constant shear at all heights makes it easier to assess the effects of the shear strength within the model. Due to the free-slip boundary conditions within MPIC, the implementation of the wind shear profile is relatively straightforward, requiring only an initial vorticity component applied throughout the domain.

The effects on the evolution of the cloud were significant. We saw from studying the volume of cloud water in each simulation that low to intermediate shear strengths led to a larger (in terms of liquid water volume) cloud forming during the initial ascent. This suggests that wind shear can enhance cloud development even for an isolated thermal. A likely cause of this enhancement was the entrainment of ambient air from the environment. Since the wind shear stretches the initial thermal horizontally, more air is pulled into its path during the ascent. A curious feature of the liquid water volume evolution is that, for strong shear, the initial ascent of the cloud is strongly inhibited, but the overall cloud volume formed at later stages is considerably greater than for low-shear simulations.

Further diagnostics taken include the liquid water-weighted mean cloud height and its standard deviation. These offered insights into the vertical distribution of cloud water subject to a given shear strength. In the high-shear simulations, we saw that the estimated cloud height was lower than in low-shear cases, but they exhibited a larger vertical dispersion of liquid water. This was also likely a consequence of increased dry air entrainment into the rising thermal due to the wind shear.

We confirmed the role of entrainment by partitioning the air in the domain into cloud air or dry air based on the liquid water fraction threshold used in previous works (Böing et al. 2019). We then computed the total enstrophy ratio between the cloud and dry air components. In low-shear simulations, we see a

pronounced spike during the initial ascent of the thermal, which smooths out as dry air enters the thermal. The extremity of this spike decreased significantly as the shear strength increased, reaffirming that more entrainment is occurring at lower levels. In the high-shear cases (above $S_x = 0.7$), we saw that the ratio has particularly low values, and most of the system enstrophy is contained in dry air. Comparing this to cross-sections through the domain around the time of the peak cloud enstrophy reveals large vortical structures in the lower levels for high-shear cases. This analysis demonstrates that enstrophy can be used to identify the entrainment of air into a cloud. Moreover, we identified the visible effects of excessive entrainment on the cloud through the liquid water cross-sections. At higher values of shear, the turbulent entrainment of dry air at low levels appeared to tear the initial thermal apart, resulting in a fragmented cloud formation.

Analysis of the vertical motions of cloud and dry air also exhibited apparent differences for high and low shear cases. As the strength of the shear increases, the vertical motions of dry air start to become dominant, as shown in Figures 3.5 and 3.6. The growth of the kinetic energy associated with vertical motions was suppressed during the initial ascent of the thermal. Instead, there was a late period of kinetic energy growth, primarily associated with dry air motions in the domain's lower region.

While not groundbreaking, these results showcase MPIC as a tool for analysing convection and its response to external forcing. The Lagrangian model enables a straightforward partitioning of parcels based on thresholds associated with any attributes. For example, one could divide the parcels based on their enstrophy, as is sometimes done in studies of the convective boundary layer (Fodor and Mellado 2020). Moreover, varying the initial thermal properties may change the response to the vertical wind shear. This may be useful when examining statistical properties within an ensemble of clouds. However, we note that there are limitations of the simplified state of MPIC at the time of writing. A more sophisticated microphysics representation within the model or even more complex shear profiles with differing strengths at low and high altitudes would both be beneficial modifications to this

setup for further studies. The microphysics in particular represent one of the most important targets to further the application of MPIC studies of the atmosphere. Similarly, initiating multiple thermals or applying surface heating to the domain may be worth initiating. Nonetheless, the results still offer insight into how vertical wind shear can impact the growth of an individual cloud.

In Chapter 4, we investigated the potential energy of a rising thermal in MPIC. Particle-in-cell methods typically do not conserve energy (Markidis and Lapenta 2011), so assessing the energy transfers within MPIC would offer some insight into the model's performance. We also used this analysis to assess energy dissipation due to turbulence. The distribution of potential energy is a significant quantity and tells us much about whether an environment is favourable for the triggering of deep convection. This is typically represented through the convective available potential energy (*CAPE*) and convective inhibition (*CIN*). In MPIC, the derived potential energy is a non-static version of the *CAPE* and provides a quantification of the energy available for a parcel of air within the cloud. To further this comparison, we demonstrated the variation of potential energy with height for a selection of test parcels from within the MPIC test case. This allowed us to identify modifications to the vertical structure of the environment that would create regions of *CIN* and *CAPE* for a given parcel of air. These examples showed that the potential energy derived can be used in MPIC to predict whether the environment is stable to convection or if the vertical motions will be resisted. In future studies where mechanical forcing, such as surface heating, is included, these properties may be beneficial.

We studied how numerical settings within MPIC affect the evolution of the kinetic and potential energy in MPIC, in particular, the resolution and parcel splitting threshold. For both of these, we observe that the total energy converges at early times in the simulation, but this is primarily lost at later stages when the flow has become fully turbulent. Furthermore, this revealed an abnormal growth in the total energy at early to intermediate times. The cause of this growth was later identified as a lack of incompressibility in the model, a common problem in fully

Lagrangian methods. A three-dimensional adaptation of the volume conservation correction detailed in Frey et al. (2022) solved this issue.

Finally, Chapter 5 presents a study of the dry convective boundary layer in MPIC. To the author’s knowledge, this represents the first time such a model has been applied to study the boundary layer. The test case for this setup was primarily based on the works of Garcia and Mellado (2014) and Liu et al. (2018). Studies of the boundary layer typically employ LES or DNS, although DNS appears to be preferred recently (Garcia and Mellado 2014; Mellado et al. 2017; Fodor and Mellado 2020). DNS suffers from computational limitations on how high the Reynolds number of a simulation can become. As such, applying a Lagrangian model like MPIC provides an alternative means to analyse the properties of the boundary layer and entrainment zone.

Furthermore, the vorticity dynamics formulation in MPIC lends itself naturally to studies of enstrophy in boundary layer turbulence. We apply a uniform heat flux, with minor random variations, to the lower boundary of a linearly stratified atmosphere and analyse the growth of the resulting boundary layer. The results serve to further our understanding of the entrainment zone and assess the performance of MPIC in such a setting. While our estimates of the Reynolds number in these simulations fall short of even DNS values, we stress that these were carried out at much lower resolutions than other studies in the field. Studies of MPIC carried out at comparable grid resolutions are likely to yield much higher Reynolds numbers than are typically obtained in DNS. While these may prove to be computationally expensive, similarly high resolution studies in DNS and LES models are likely to require similar resources. A higher resolution study of the boundary layer in MPIC, likely using a massively parallel computer, would be a beneficial next step in the comparison to existing studies. A discussion of the computational expense of MPIC in comparison to existing Eulerian methods is provided in Section 2 of Böing et al. (2019) and discusses the scalability of the base MPIC code on multiple cores. For reference the simulations carried out in this work were run using 6 cores with runtimes varying significantly depending on the simulation in

question, with the boundary layer simulations taking 3-4 days depending on the parcel splitting threshold used. A new highly parallel version, called PMPIC has been developed since then, although the developments documented in this thesis are not yet implemented in PMPIC.

This case study required substantial modifications to the MPIC method to incorporate the buoyancy's surface flux. The modifications to the vertical solver for inversion were discussed in Section 2.3. These solver changes were required to allow the buoyancy gradients to vary at the lower surface for the flux implementation. Surface heating in MPIC is simulated by adding increments of buoyancy to the parcels in the lowest grid cell of the domain. These are weighted by the vertical position of the parcels within each cell. In conjunction with the three-dimensional volume correction, it ensures the correct flux is applied to the domain at each timestep. While this may seem like a somewhat simple change to the model, it is a potent tool for future studies. The same flux method can be applied to introduce a surface flux of any parcel attribute, and we will discuss some possible extensions later in the chapter. For now, it is enough that this method brings the MPIC model even closer to more realistic studies of convection, such as the boundary layer study in Chapter 5.

The growth of the boundary layer in MPIC appears to agree relatively well with predictions made by other studies. Defining the base of the boundary layer as the zero-crossing height of the buoyancy flux profile and comparing to the encroachment height (Fedorovich et al. 2004; Garcia and Mellado 2014), we saw that the two measures are almost identical. This reaffirms the surface flux implementation within MPIC, as the encroachment height defines the base of the boundary layer based on information about the heat flux into the system. Another important detail about this measure is that it confirmed that our study captured the boundary layer in the equilibrium entrainment regime at intermediate to late times. A comparison of the encroachment height to other measures of the boundary layer depth also found that at late times, the heights defined by buoyancy gradients and fluxes were approximately constant with respect to the base of the boundary layer.

One concern is that the entrainment zone observed within MPIC was substantially thicker than in other studies and contrasted strongly with the zero-order model of Zilitinkevich (1991), based on an infinitesimally thin layer. This may result from not implementing a sponge layer in the upper domain to suppress deviations from the background state or be a sign that MPIC overestimated the entrainment’s vertical influence. This does not detract from the overall significance of our results, as the overall entrainment zone structure was mostly unchanged.

Our study also analyses the vorticity field and vertical distribution of enstrophy in the boundary layer, which is not commonly studied in LES studies. x - z cross-section through the centre of the domain revealed large eddies within the boundary layer itself, topped by a region of depleted vorticity. This depleted region is consistent with the description of the entrainment zone in Garcia and Mellado (2014), who observed that the entrainment zone is comprised of two overlapping sublayers. The upper sublayer is described as a transition region between the turbulent entrainment zone and the nonturbulent stratified region aloft. As such, we believe that the vorticity depletion above the boundary layer is aligned with this upper sublayer as the flow transitions nonturbulent region above. The vertical profile of the horizontally averaged enstrophy and its gradient reflect this. The enstrophy profile agrees with the work of Fodor and Mellado (2020), who observed a sharp decay of enstrophy at the top of the entrainment zone. They used the saddle point of this distribution to partition air into turbulent and nonturbulent components. This is comparable to the minimum of the vertical enstrophy gradient in our boundary layer study. We found that the vertical position of this minimum, at later stages of the boundary layer’s evolution, became close to constant with respect to the encroachment height. Its position above the height of minimum buoyancy flux was also consistent with the observations of Fodor and Mellado (2020).

We propose that the minimum of the enstrophy gradient may represent the base of the upper sublayer in the entrainment zone. The measures of sublayer thickness computed in section 5.3.8 support this. The buoyancy gradient thickness definition

used in Garcia and Mellado (2014) was compared to the difference between the maximum buoyancy gradient height and the minimum enstrophy gradient height. Since the upper sublayer is believed to be centred on the buoyancy gradient maximum, the fact that both quantities agreed seems to support this claim. Moreover, we computed the height at which the enstrophy gradient decayed to half its minimum value and found that this also aligned with the other measures of the upper sublayer thickness. These results and the definition of the upper sublayer as a transition region seem to support the notion that its base is aligned with the minimum vertical enstrophy gradient. Further studies may wish to investigate this relationship further. Notably, despite the different models used, our study appears to support the two-layer structure observed in Garcia and Mellado (2014) and subsequent works.

Finally, we computed parameters associated with the entrainment rate. These were compared to the predictions and results of Fedorovich et al. (2004), Liu et al. (2018) and Garcia and Mellado (2014). Initial estimates of the mean entrainment rate in MPIC seemed to show a good agreement with the $t^{-2/3}$ scaling predicted by LES results for the shear-free boundary layer. However, the models differed when comparing entrainment ratios, with MPIC underestimating entrainment ($A = 0.1$ in MPIC compared to the generally accepted value of $A \approx 0.2$). The underestimation of the entrainment ratio was mirrored in comparisons of the entrainment rate-Richardson number power law to Garcia and Mellado (2014). The power law was recovered for the entrainment rate computed using the buoyancy increment at the minimum flux height. We find that it aligned well with the DNS study's entrainment ratio of $A = 0.1$. When using the buoyancy increment across the whole entrainment zone, we found that MPIC significantly overestimated the entrainment ratio. A likely source of this is the thicker entrainment zone observed in the MPIC simulation. Overall, this initial simulation of the boundary layer in MPIC showed good initial comparisons to existing studies using DNS and LES. The fact that we obtained good agreement with existing results despite the difference in model formulation and the lower boundary conditions is significant, and illustrate the potential of a Lagrangian approach.

The results of this thesis demonstrate the application of a Lagrangian method for modelling convection to studies of both isolated thermals and the convective boundary layer and the influence of wind shear. Simultaneously, we have developed modifications and diagnostics within the model that have allowed us to gain new insights into convective processes from a new perspective while presenting comparisons to existing studies, as in Chapter 5. These results showcase the MPIC method’s potential as a possible alternative to conventional Eulerian models like LES in studies of atmospheric convection. The modifications have opened up many avenues of exploration that may further develop our understanding of convective processes. Meanwhile, many of the diagnostics obtained are relevant to convective parameterisation schemes, such as the available potential energy, vertical fluxes and vertical distributions of thermodynamic properties. These diagnostics are also particularly relevant to the potential implementation of MPIC within a superparameterisation scheme since many are used in coupling to large-scale circulations.

The surface flux implementation, in particular, represents a significant model enhancement. While the analysis of pre-initiated thermals may be of use in studies of cloud mixing and the response to external forcing like wind shear, surface fluxes open the door for more realistic studies of convection. In Chapter 5, we focused on the growth of a dry convective boundary layer to better understand the structure of the entrainment zone and for comparison to existing studies. One straightforward extension to the existing study would be a simulation of the sheared boundary layer and comparison to the results in Liu et al. (2018). Entrainment relationships in the sheared boundary layer differ from the shear-free case significantly, and this would provide another opportunity to test the model. Another straightforward extension would be adding a surface flux of moisture and studying how the moist convective boundary layer develops. Mellado et al. (2017) studied moisture statistics in the entrainment zone using a DNS simulation, which may provide a point of reference for similar studies using MPIC. As previously stated, the implementation of surface fluxes for other parcel attributes can be handled exactly as the buoyancy fluxes, so this would be a relatively straightforward addition. Moreover, such an approach

would open the door to simulating the stratocumulus-topped boundary layer and the triggering of deep convection, both significant areas of uncertainty in convective parameterisation schemes.

Another extension of this surface scheme lies in representing aerosols and tracers in MPIC. The parcel approach of MPIC makes implementing aerosol and tracer concentrations straightforward and computationally inexpensive. This is because parcel attributes in MPIC are evolved on the parcels and do not require additional prognostic equations to be solved. Surface fluxes of these concentrations could be added in the same way as buoyancy. Tracer populations fit naturally into the MPIC framework, since they simply become additional attributes carried by the parcels. This represents one of the core advantages of the method in potential future studies of convection, since it becomes straightforward to track properties such as the origin height of air in the boundary layer to see the vertical extent of entrainment, or the transport and mixing of hydrometeor and chemical species. Such studies are relevant to climate simulations due to the influence of aerosols on radiative transfer and cloud parametrisations through the resulting changes in heating and condensation nuclei.

Further applications of MPIC may wish to move away from the currently simplified microphysics representation. In this thesis, we have noted the limitations of such an approach. A complete microphysics scheme within MPIC is a logical next step for the method but raises questions about how exactly to incorporate this. Simple steps would involve changing the saturation height-based approach to condensation and replacing it with a more physical process. If MPIC is to become an appealing alternative to LES models it must seek to represent the effects of evaporation and precipitation realistically. Precipitation presents the most significant challenge. A simple approach might be to interpolate liquid water onto the grid and adopt a bulk water approach (Kessler 1969), as is typically done in LES models. This method is appealing in the simplicity of its implementation and the easy comparison to existing LES results. However, one could argue that such an approach would compromise some of the novelty of the Lagrangian approach in

MPIC. As such, it may be beneficial to look to the Lagrangian ‘super-droplet’ scheme proposed by Shima et al. (2009), which may fit more naturally into the MPIC framework. Precipitation would significantly enhance the existing MPIC model regardless of the scheme used and allow for comparisons to observational data and parameterisation schemes. It would, for example, be possible to recreate the study of Böing et al. (2012) on the influence of cold pools on convective organisation if precipitation and evaporation were added to MPIC simulations. This would also provide an interesting comparison of the method to existing results obtained from LESs in terms of more physically relevant diagnostics, like rainfall distributions and intensity.

Another limitation of MPIC’s current state, which we have previously acknowledged in this thesis, is the use of the Boussinesq approximation. The scale height used in all studies in this work has remained at 2000 m as in the original test cases (Dritschel et al. 2018; Böing et al. 2019). This results in the estimated height of the domain being 12 km, extending far above the range for which the Boussinesq approximation is valid. It is worth considering whether an anelastic approximation (Ogura and Phillips 1962) is more appropriate for future works. An anelastic model would also allow MPIC to study deep convection more accurately. Implementing a height-dependent density profile is somewhat complex and would depend heavily on the vertical density profile’s exact form. In particular, the form of the density profile would alter the tridiagonal problem that needs to be solved in the vertical, and doing this for a generalised form $\rho(z)$ would likely require numerical differentiation of the density profile. This would also introduce additional mixed derivatives into the tridiagonal problem in the vertical solver. In this case, solving the system of equations is non-trivial, and the compact differencing must be modified accordingly. A similar inversion problem was encountered in Section 3 of Mohebalhojeh and Dritschel (2007), and solving the inversion problem in MPIC may be possible by adapting the solution presented therein. The exact nature of the tridiagonal problem will still depend on the form of the vertical density profile. As such, the generalisation of the solver to any prescribed profile may be more complicated. However, for a prescribed, analytic density profile, the tridiagonal

problem becomes more straightforward. This may be a worthwhile endeavour, and a simple test case for an exponential profile (i.e. $\rho(z) = \rho_0 e^{-z}$) is certainly feasible. In log-pressure coordinates, this density profile is accurate for an isothermal atmosphere, where the log-pressure height is approximately equal to the physical height. Within the troposphere this appears to be a reasonable approximation. As such, it would represent a natural step towards developing a fully anelastic version of MPIC.

One final aspect of MPIC that could be improved is the representation of mixing. The integrated vorticity stretching of parcels determines when they are split (equation 2.19). It may instead be beneficial to consider a more physical approach for this. Frey et al. (2022) present a two-dimensional implementation of an alternative called the Elliptical parcel-in-cell method (EPIC). In EPIC, parcels are modelled as ellipsoids rather than point parcels, and the deformation of the ellipsoids determines whether parcels split. The threshold for splitting is now defined based on a critical aspect ratio between the major and minor axes of a given parcel. The removal of parcels is replaced by merging smaller parcels into their nearest neighbours rather than transferring residuals onto the grid in the version of MPIC used here. This would likely reduce the effects of numerical diffusion associated with the mixing within MPIC. Initial tests of the two-dimensional version of EPIC show that the new mixing scheme performs well and shows excellent conservation with little extra computational expense. A fully three-dimensional version of this approach to mixing is currently in development and would significantly enhance the representation of turbulence. Even within the work in this thesis, we have drawn inspiration from the development of EPIC. For example, a three-dimensional adaptation of the corrections used to ensure incompressibility in EPIC was applied in Chapter 5.

While many possible refinements to the method remain, our results show that MPIC presents a promising alternative to existing convection models. Moreover, we have shown that a Lagrangian approach can also enhance our understanding of the phenomena associated with convective processes. Lagrangian models offer a

new perspective to supplement and go beyond the knowledge obtained from Eulerian models. While limited in scope, we have demonstrated some early steps and applications towards a more complete Lagrangian model of atmospheric convection.

Bibliography

- Angevine, W. M. et al. (2018). “Shallow Cumulus in WRF Parameterizations Evaluated against LASSO Large-Eddy Simulations”. *Monthly Weather Review* 146.12, pp. 4303–4322.
- Arakawa, A. and W. H. Schubert (1974). “Interaction of a Cumulus Cloud Ensemble with the Large-Scale Environment, Part I.” *Journal of Atmospheric Sciences* 31.3, pp. 674–701.
- Arakawa, A. (2004). “The cumulus parameterization problem: Past, present, and future”. *Journal of Climate* 17.13, pp. 2493–2525.
- Arakawa, A. and W. H. Schubert (1974). “Interaction of a Cumulus Cloud Ensemble with the Large-Scale Environment, Part I” . *Journal of Atmospheric Sciences* 31.3, pp. 674–701.
- Austin, P. H. et al. (1985). “Small-Scale Variability in Warm Continental Cumulus Clouds.” *Journal of Atmospheric Sciences* 42, pp. 1123–1138.
- Baba, Y. (2020). “Shallow convective closure in a spectral cumulus parameterization”. *Atmospheric Research* 233, p. 104707.
- Becker, T. et al. (2018). “Estimating Bulk Entrainment With Unaggregated and Aggregated Convection”. *Geophysical Research Letters* 45.1, pp. 455–462.
- Betts, A. (1974). “Reply to comment on the paper ‘Non-precipitating cumulus convection and its parameterization’”. *Quarterly Journal of the Royal Meteorological Society* 100.425, pp. 469–471.
- Bisset, D. K. et al. (2002). “The turbulent/non-turbulent interface bounding a far wake”. *Journal of Fluid Mechanics* 451, pp. 383–410.

- Blyth, A. M. et al. (2005). “A study of thermals in cumulus clouds”. *Quarterly Journal of the Royal Meteorological Society* 131, pp. 1171–1190.
- Blyth, A. M. (1993). “Entrainment in Cumulus Clouds”. *Journal of Applied Meteorology and Climatology* 32.4, pp. 626–641.
- Blyth, A. M. et al. (2015). “High-resolution observations of precipitation from cumulonimbus clouds”. *Meteorological Applications* 22.1, pp. 75–89.
- Bony, S. et al. (2015). “Clouds, circulation and climate sensitivity”. *Nature Geoscience* 8, pp. 261–268.
- Borrell, G. and J. Jiménez (2016). “Properties of the turbulent/non-turbulent interface in boundary layers”. *Journal of Fluid Mechanics* 801, pp. 554–596.
- Brackbill, J. (1991). “FLIP MHD: A particle-in-cell method for magnetohydrodynamics”. *Journal of Computational Physics* 96.1, pp. 163–192.
- Bresch, D. et al. (2020). “Augmented skew-symmetric system for shallow-water system with surface tension allowing large gradient of density”. *Journal of Computational Physics* 419, p. 109670.
- Bretherton, C. S. and P. K. Smolarkiewicz (1989). “Gravity Waves, Compensating Subsidence and Detrainment around Cumulus Clouds”. *Journal of Atmospheric Sciences* 46.6, pp. 740–759.
- Brown, N. et al. (2015). “A Highly Scalable Met Office NERC Cloud Model”. *Proceedings of the 3rd International Conference on Exascale Applications and Software*. EASC '15. Edinburgh, UK: University of Edinburgh, pp. 132–137. ISBN: 978-0-9926615-1-9.
- Böing, S. J. et al. (2019). “Comparison of the Moist Parcel-in-Cell (MPIC) model with large-eddy simulation for an idealized cloud”. *Quarterly Journal of the Royal Meteorological Society*.
- Böing, S. J. et al. (2012). “Influence of the Subcloud Layer on the Development of a Deep Convective Ensemble”. *Journal of the Atmospheric Sciences* 69.9, pp. 2682–2698.
- Canuto, V. M. and Y. Cheng (1997). “Determination of the Smagorinsky–Lilly constant CS”. *Physics of Fluids* 9.5, pp. 1368–1378.

- Carson, D. and F. Smith (1975). “Thermodynamic model for the development of a convectively unstable boundary layer”. *Advances in Geophysics*. Vol. 18. Elsevier, pp. 111–124.
- Chen, J. et al. (2020). “Changes in Convective Available Potential Energy and Convective Inhibition under Global Warming”. *Journal of Climate* 33.6, pp. 2025–2050.
- Chen, Q. et al. (2015). “Roles of wind shear at different vertical levels: Cloud system organization and properties”. *Journal of Geophysical Research: Atmospheres* 120.13, pp. 6551–6574.
- Chen, Z. et al. (2015). “An SPH model for multiphase flows with complex interfaces and large density differences”. *Journal of Computational Physics* 283, pp. 169–188.
- Christiansen, J. P. and N. J. Zabusky (1973). “Instability, coalescence and fission of finite-area vortex structures”. *Journal of Fluid Mechanics* 61, pp. 219–243.
- Coniglio, M. C. et al. (2006). “Effects of Upper-Level Shear on the Structure and Maintenance of Strong Quasi-Linear Mesoscale Convective Systems.” *Journal of Atmospheric Sciences* 63.4, pp. 1231–1252.
- Conzemius, R. J. and E. Fedorovich (2006a). “Dynamics of sheared convective boundary layer entrainment. Part I: Methodological background and large-eddy simulations”. *Journal of the Atmospheric Sciences* 63.4, pp. 1151–1178.
- (2006b). “Dynamics of sheared convective boundary layer entrainment. Part II: Evaluation of bulk model predictions of entrainment flux”. *Journal of the Atmospheric Sciences* 63.4, pp. 1179–1199.
- Cottet, G.-H. and P. D. Koumoutsakos (2000). *Vortex Methods*, p. 328.
- Craven, J. and H. Brooks (2004). “Baseline climatology of sounding derived parameters associated with deep, moist convection”. *Natl. Wea. Dig.* 28, pp. 13–24.
- Da Silva, C. B. et al. (2014a). “Characteristics of the turbulent/nonturbulent interface in boundary layers, jets and shear-free turbulence”. *Journal of Physics: Conference Series*. Vol. 506. 1. IOP Publishing, p. 012015.

- Da Silva, C. B. et al. (2014b). “Interfacial layers between regions of different turbulence intensity”. *Annu. Rev. Fluid Mech* 46.1, pp. 567–590.
- Deardorff, J. W. (1980). “Stratocumulus-capped mixed layers derived from a three-dimensional model”. *Boundary-Layer Meteorology* 18, pp. 495–527.
- Deardorff, J. W. (1970). “A numerical study of three-dimensional turbulent channel flow at large Reynolds numbers”. *Journal of Fluid Mechanics* 41.2, 453–480. DOI: 10.1017/S0022112070000691.
- Deardorff, J. W. (1974). “Three-dimensional numerical study of the height and mean structure of a heated planetary boundary layer”. *Boundary-Layer Meteorology* 7.1, pp. 81–106.
- Deardorff, J. (1979). “Prediction of convective mixed-layer entrainment for realistic capping inversion structure”. *Journal of Atmospheric Sciences* 36.3, pp. 424–436.
- Di Mascio, A. et al. (2017). “Smoothed particle hydrodynamics method from a large eddy simulation perspective”. *Physics of Fluids* 29.3, p. 035102.
- Dimotakis, P. E. (2000). “The mixing transition in turbulent flows”. *Journal of Fluid Mechanics* 409, pp. 69–98.
- Dong, L. et al. (2015). “Lagrangian advection scheme with shape matrix (LASM) v0.2: interparcel mixing, physics–dynamics coupling and 3-D extension”. *Geoscientific Model Development* 8, pp. 2675–2686.
- Dritschel, D. et al. (2018). “The moist parcel-in-cell (MPIC) method”. *Quarterly Journal of the Royal Meteorological Society*.
- Dritschel, D. G. and J. Fontane (2010). “The combined Lagrangian advection method”. *Journal of Computational Physics* 229.14, pp. 5408–5417. ISSN: 0021-9991.
- Dube, A. et al. (2014). “Forecasting the heavy rainfall during Himalayan flooding—June 2013”. *Weather and Climate Extremes* 4, pp. 22–34.
- Durrán, D. R. and A. Arakawa (2007). “Generalizing the Boussinesq approximation to stratified compressible flow”. *Comptes Rendus Mécanique* 335.9-10, pp. 655–664.

- D’Andrea, F. et al. (2014). “Triggering Deep Convection with a Probabilistic Plume Model”. *Journal of the Atmospheric Sciences* 71.11, pp. 3881–3901.
- Ebert, E. E. and J. A. Curry (1992). “A parameterization of ice cloud optical properties for climate models”. *Journal of Geophysical Research: Atmospheres* 97.D4, pp. 3831–3836.
- Edwards, J. M. and A. Slingo (1996). “Studies with a flexible new radiation code. I: Choosing a configuration for a large-scale model”. *Quarterly Journal of the Royal Meteorological Society* 122, pp. 689–719.
- Emanuel, K. A. et al. (1994). “On large-scale circulations in convecting atmospheres”. *Quarterly Journal of the Royal Meteorological Society* 120.519, pp. 1111–1143.
- Esfahanian, V. et al. (2005). “On the use of the super compact scheme for spatial differencing in numerical models of the atmosphere”. *Quarterly Journal of the Royal Meteorological Society* 131, pp. 2109–2129.
- Falkovich, G. et al. (2002). “Acceleration of rain initiation by cloud turbulence”. *Nature* 419.6903, pp. 151–154.
- Fedorovich, E. and D. Mironov (1995). “A model for a shear-free convective boundary layer with parameterized capping inversion structure”. *Journal of Atmospheric Sciences* 52.1, pp. 83–96.
- Fedorovich, E. (1995). “Modeling the atmospheric convective boundary layer within a zero-order jump approach: An extended theoretical framework”. *Journal of Applied Meteorology and Climatology* 34.9, pp. 1916–1928.
- Fedorovich, E. et al. (2004). “Convective Entrainment into a Shear-Free, Linearly Stratified Atmosphere: Bulk Models Reevaluated through Large Eddy Simulations”. *Journal of the Atmospheric Sciences* 61.3, pp. 281–295.
- Fodor, K. and J. P. Mellado (2020). “New Insights into Wind Shear Effects on Entrainment in Convective Boundary Layers Using Conditional Analysis”. *Journal of the Atmospheric Sciences* 77.9, pp. 3227–3248.
- Frank, J. and S. Reich (2004). “The Hamiltonian particle-mesh method for the spherical shallow water equations”. *Atmospheric Science Letters* 5.5, pp. 89–95.

- Frank, J et al. (2002). “The Hamiltonian particle-mesh method, Meshfree Methods for Partial Differential Equations”. *Lecture Notes in Computational Science and Engineering* 26, pp. 131–142.
- Frey, M. et al. (2022). “EPIC: The Elliptical Parcel-In-Cell method”. *Journal of Computational Physics: X* 14, p. 100109.
- Gadian, A. (1991). “Lagrangian Modelling of an Atmospheric Plume and an Analysis of the Comparative Advantages of Lagrangian and Eulerian Methods”. *Computer Modelling in Environmental Sciences*. Ed. by D. Farmer and M. Rycroft. Oxford University Press, pp. 255–273.
- Garcia, J. R. and J. P. Mellado (2014). “The Two-Layer Structure of the Entrainment Zone in the Convective Boundary Layer”. *Journal of the Atmospheric Sciences* 71.6, pp. 1935 –1955.
- Gauger, C. et al. (2000). “The Finite Mass Method”. *SIAM Journal on Numerical Analysis* 37.6, pp. 1768–1799.
- Germano, M. et al. (1991). “A dynamic subgrid-scale eddy viscosity model”. *Physics of Fluids A: Fluid Dynamics* 3.7, pp. 1760–1765.
- Ghader, S. et al. (2011). “High-order compact scheme for Boussinesq equations: implementation and numerical boundary condition issue”. *International Journal for Numerical Methods in Fluids* 69.3, pp. 590–605.
- Gidel, L. T. (1983). “Cumulus cloud transport of transient tracers”. *Journal of Geophysical Research: Oceans* 88.C11, pp. 6587–6599.
- Gingold, R. A. and J. J. Monaghan (1977). “Smoothed particle hydrodynamics: theory and application to non-spherical stars”. *Monthly notices of the royal astronomical society* 181.3, pp. 375–389.
- Goswami, B. B. et al. (2017). “Improving synoptic and intraseasonal variability in CFSv2 via stochastic representation of organized convection”. *Geophysical Research Letters* 44.2, pp. 1104–1113.
- Grabowski, W. W. and T. L. Clark (1993). “Cloud-Environment Interface Instability: Part II: Extension to Three Spatial Dimensions.” *Journal of Atmospheric Sciences* 50, pp. 555–573.

- Grabowski, W. W. (2016). “Towards Global Large Eddy Simulation: Super-Parameterization Revisited”. *Journal of the Meteorological Society of Japan. Ser. II* 94.4, pp. 327–344.
- Grell, G. A. and S. R. Freitas (2014). “A scale and aerosol aware stochastic convective parameterization for weather and air quality modeling”. *Atmospheric Chemistry and Physics* 14.10, pp. 5233–5250.
- Grell, G. A. (1993). “Prognostic Evaluation of Assumptions Used by Cumulus Parameterizations”. *Monthly Weather Review* 121.3, pp. 764–787.
- Grewe, V. et al. (2014). “On the theory of mass conserving transformations for Lagrangian methods in 3D atmosphere-chemistry models”. *Meteorologische Zeitschrift* 23.4, pp. 441–447.
- Grover, S. N. and H. R. Pruppacher (1985). “The Effect of Vertical Turbulent Fluctuations in the Atmosphere on the Collection of Aerosol Particles by Cloud Drops”. *Journal of Atmospheric Sciences* 42.21, pp. 2305–2318.
- Guichard, F. and F. Couvreux (2017). “A short review of numerical cloud-resolving models”. *Tellus A: Dynamic Meteorology and Oceanography* 69.1, p. 1373578.
- Hadley, G. (1735). “VI. Concerning the cause of the general trade-winds”. *Philosophical Transactions of the Royal Society of London* 39.437, pp. 58–62.
- Haghshenas, A. and J. P. Mellado (2019). “Characterization of wind-shear effects on entrainment in a convective boundary layer”. *Journal of Fluid Mechanics* 858, 145–183.
- Han, J. et al. (2017). “Updates in the NCEP GFS Cumulus Convection Schemes with Scale and Aerosol Awareness”. *Weather and Forecasting* 32.5, pp. 2005–2017.
- Hannah, W. M. (2017). “Entrainment versus Dilution in Tropical Deep Convection”. *Journal of the Atmospheric Sciences* 74.11, pp. 3725–3747.
- Harlow, F. H. (1962). *The particle-in-cell method for numerical solution of problems in fluid dynamics*. Tech. rep. Los Alamos National Lab.(LANL), Los Alamos, NM (United States).

- Heinze, R. et al. (2017). “Large-eddy simulations over Germany using ICON: a comprehensive evaluation”. *Quarterly Journal of the Royal Meteorological Society* 143.702, pp. 69–100.
- Heus, T. et al. (2010). “Formulation of the Dutch Atmospheric Large-Eddy Simulation (DALES) and overview of its applications”. *Geoscientific Model Development* 3.2, pp. 415–444.
- Heus, T. et al. (2008). “Mixing in Shallow Cumulus Clouds Studied by Lagrangian Particle Tracking”. *Journal of the Atmospheric Sciences* 65.8, pp. 2581–2597.
- Hohenegger, C. and C. S. Bretherton (2011). “Simulating deep convection with a shallow convection scheme”. *Atmospheric Chemistry and Physics* 11.20, pp. 10389–10406.
- Holloway, C. E. et al. (2014). “Understanding and representing atmospheric convection across scales: recommendations from the meeting held at Dartington Hall, Devon, UK, 28-30 January 2013”. *Atmospheric Science Letters* 15.4, pp. 348–353.
- Houze, R. A. (2004). “Mesoscale convective systems”. *Reviews of Geophysics* 42, RG4003, RG4003.
- Houze, R. A. (1997). “Stratiform precipitation in regions of convection: A meteorological paradox?” *Bulletin of the American Meteorological Society* 78.10, pp. 2179–2196.
- Huang, D. et al. (2018). “Uncertainty of global summer precipitation in the CMIP5 models: a comparison between high-resolution and low-resolution models”. *Theoretical and Applied Climatology* 132.1.
- Jansson, F. et al. (2019). “Regional Superparameterization in a Global Circulation Model Using Large Eddy Simulations”. *Journal of Advances in Modeling Earth Systems* 11.9, pp. 2958–2979.
- Jonker, H. J. and M. A. Jiménez (2014). “Laboratory experiments on convective entrainment using a saline water tank”. *Boundary-layer meteorology* 151.3, pp. 479–500.

- Kaas, E. (2008). “A simple and efficient locally mass conserving semi-Lagrangian transport scheme”. *Tellus A: Dynamic Meteorology and Oceanography* 60.2, pp. 305–320.
- Kain, J. S. and J. M. Fritsch (1990). “A One-Dimensional Entraining/Detraining Plume Model and Its Application in Convective Parameterization”. *Journal of Atmospheric Sciences* 47.23, pp. 2784 –2802.
- Kessler, E. (1969). “On the Distribution and Continuity of Water Substance in Atmospheric Circulations”. *On the Distribution and Continuity of Water Substance in Atmospheric Circulations*. American Meteorological Society, pp. 1–84.
- Khairoutdinov, M. F. et al. (2009). “Large-Eddy Simulation of Maritime Deep Tropical Convection”. *Journal of Advances in Modeling Earth Systems* 1, 15, p. 15.
- Kim, S.-W. et al. (2003). “Entrainment processes in the convective boundary layer with varying wind shear”. *Boundary-layer meteorology* 108.2, pp. 221–245.
- Klein, S. A. et al. (2018). “Low-Cloud Feedbacks from Cloud-Controlling Factors: A Review”. *Shallow Clouds, Water Vapor, Circulation, and Climate Sensitivity*. Ed. by R. Pincus et al. Springer International Publishing, pp. 135–157.
- Klingler, M. et al. (2005). “The Finite Mass Method on Domains with Boundary”. *SIAM Journal on Scientific Computing* 26.5, pp. 1744–1759.
- Knight, C. G. et al. (2007). “Association of parameter, software, and hardware variation with large-scale behavior across 57,000 climate models”. *Proceedings of the National Academy of Sciences* 104.30, pp. 12259–12264.
- Kolmogorov, A. N. (1941). “The local structure of turbulence in incompressible viscous fluid for very large Reynolds numbers”. *C. R. Acad. Sci. URSS* 30, pp. 301–305.
- Konopka, P. et al. (2004). “Mixing and ozone loss in the 1999–2000 Arctic vortex: Simulations with the three-dimensional Chemical Lagrangian Model of the Stratosphere (CLaMS)”. *Journal of Geophysical Research: Atmospheres* 109.D2.
- Kuo, H. L. (1974). “Further Studies of the Parameterization of the Influence of Cumulus Convection on Large-Scale Flow”. *Journal of Atmospheric Sciences* 31.5, pp. 1232 –1240.

- Kwon, Y. C. and S.-Y. Hong (2017). “A Mass-Flux Cumulus Parameterization Scheme across Gray-Zone Resolutions”. *Monthly Weather Review* 145.2, pp. 583–598.
- Lilly, D. K. (1962). “On the numerical simulation of buoyant convection”. *Tellus* 14, p. 148.
- Lilly, D. K. (1966). “The representation of small-scale turbulence in numerical simulation experiments”. Ed. by H. Goldstine. Proc. IBM Scientific Computing symp. on Environmental Sciences, 195–210.
- Lilly, D. K. (1968). “Models of cloud-topped mixed layers under a strong inversion”. *Quarterly Journal of the Royal Meteorological Society* 94, pp. 292–309.
- Lin, C.-H. et al. (2022a). “The Impact of Scale-Aware Parameterization on the Next-Generation Global Prediction System in Taiwan for Front Predictions”. *Atmosphere* 13.7.
- Lin, G. et al. (2022b). “Mesoscale convective systems in a superparameterized E3SM simulation at high resolution”. *Journal of Advances in Modeling Earth Systems* 14.1, e2021MS002660.
- Lin, J. et al. (2022c). “Atmospheric Convection”. *Atmosphere-Ocean* 60.3-4, pp. 422–476.
- Lin, Y.-L. et al. (1983). “Bulk Parameterization of the Snow Field in a Cloud Model”. *Journal of Applied Meteorology and Climatology* 22.6, pp. 1065–1092.
- Liu, C. et al. (2018). “Revisiting entrainment relationships for shear-free and sheared convective boundary layers through large-eddy simulations”. *Quarterly Journal of the Royal Meteorological Society* 144.716, pp. 2182–2195.
- Lucy, L. B. (1977). “A numerical approach to the testing of the fission hypothesis”. *The astronomical journal* 82, pp. 1013–1024.
- Macià, F. et al. (2011). “Theoretical Analysis of the No-Slip Boundary Condition Enforcement in SPH Methods”. *Progress of Theoretical Physics - PROGRESS THEOR PHYS KYOTO* 125, pp. 1091–1121.
- Madden, R. A. and P. R. Julian (1971). “Detection of a 40–50 Day Oscillation in the Zonal Wind in the Tropical Pacific”. *Journal of Atmospheric Sciences* 28.5, pp. 702–708.

- Mahrt, L (1991). “Boundary-layer moisture regimes”. *Quarterly Journal of the Royal Meteorological Society* 117.497, pp. 151–176.
- Manabe, S. et al. (1965). “Simulated Climatology of a General Circulation Model with a Hydrologic Cycle”. *Monthly Weather Review* 93.12, pp. 769–798.
- Mapes, B. E. (2000). “Convective inhibition, subgrid-scale triggering energy, and stratiform instability in a toy tropical wave model”. *Journal of the Atmospheric Sciences* 57.10, pp. 1515–1535.
- Markidis, S. and G. Lapenta (2011). “The energy conserving particle-in-cell method”. *Journal of Computational Physics* 230.18, pp. 7037–7052.
- Matheou, G. et al. (2011). “On the Fidelity of Large-Eddy Simulation of Shallow Precipitating Cumulus Convection”. *Monthly Weather Review* 139.9, pp. 2918–2939.
- Mayrhofer, A. et al. (2015). “DNS and LES of 3-D wall-bounded turbulence using Smoothed Particle Hydrodynamics”. *Computers & Fluids* 115, pp. 86–97.
- Mekonnen, A. and W. B. Rossow (2011). “The Interaction Between Deep Convection and Easterly Waves over Tropical North Africa: A Weather State Perspective”. *Journal of Climate* 24, pp. 4276–4294.
- Mekonnen, A. et al. (2006). “Analysis of Convection and Its Association with African Easterly Waves”. *Journal of Climate* 19, p. 5405.
- Mellado, J. P. et al. (2009). “Buoyancy reversal in cloud-top mixing layers”. *Quarterly Journal of the Royal Meteorological Society: A journal of the atmospheric sciences, applied meteorology and physical oceanography* 135.641, pp. 963–978.
- Mellado, J. P. et al. (2017). “Moisture statistics in free convective boundary layers growing into linearly stratified atmospheres”. *Quarterly Journal of the Royal Meteorological Society* 143.707, pp. 2403–2419.
- Mesinger, F. and A. Arakawa (1976). *Numerical methods used in atmospheric models*. Tech. rep.
- Mirocha, J. et al. (2014). “Resolved turbulence characteristics in large-eddy simulations nested within mesoscale simulations using the Weather Research and Forecasting Model”. *Monthly Weather Review* 142.2, pp. 806–831.

- Moeng, C. et al. (2007). “Examining two-way grid nesting for large eddy simulation of the PBL using the WRF model”. *Monthly weather review* 135.6, pp. 2295–2311.
- Moeng, C.-H. and P. P. Sullivan (1994). “A comparison of shear-and buoyancy-driven planetary boundary layer flows”. *Journal of Atmospheric Sciences* 51.7, pp. 999–1022.
- Mohebalhojeh, A. R. and D. G. Dritschel (2007). “Assessing the Numerical Accuracy of Complex Spherical Shallow-Water Flows”. *Monthly Weather Review* 135.11.
- Moin, P. and K. Mahesh (1998). “Direct numerical simulation: a tool in turbulence research”. *Annual review of fluid mechanics* 30.1, pp. 539–578.
- Monaghan, J. J. (1992). “Smoothed particle hydrodynamics”. *Annual Review of Astronomy and Astrophysics* 30, pp. 543–574.
- Monaghan, J. (2000). “SPH without a Tensile Instability”. *Journal of Computational Physics* 159.2, pp. 290–311.
- (2017). “SPH- simulation of 2D turbulence driven by a moving cylinder”. *European Journal of Mechanics - B/Fluids* 65, pp. 486–493.
- Monaghan, J. and C. A. Mériaux (2018). “An SPH study of driven turbulence near a free surface in a tank under gravity”. *European Journal of Mechanics - B/Fluids* 68, pp. 201–210.
- Moncrieff, M. W. (1992). “Organized Convective Systems: Archetypal Dynamical Models, Mass and Momentum Flux Theory, and Parametrization”. *Quarterly Journal of the Royal Meteorological Society* 118.507, pp. 819–850.
- Moncrieff, M. W. and M. J. Miller (1976). “The dynamics and simulation of tropical cumulonimbus and squall lines”. *Quarterly Journal of the Royal Meteorological Society* 102.432, pp. 373–394.
- Mooers, G. et al. (2021). “Assessing the Potential of Deep Learning for Emulating Cloud Superparameterization in Climate Models With Real-Geography Boundary Conditions”. *Journal of Advances in Modeling Earth Systems* 13.5.

- Myoung, B. (2010). “The Convective Instability Pathway to Warm Season Drought in Texas. Part I: The Role of Convective Inhibition and Its Modulation by Soil Moisture”. *Journal of Climate* 23.
- Nealon, R. et al. (2015). “On the Bardeen–Petterson effect in black hole accretion discs”. *Monthly Notices of the Royal Astronomical Society* 448.2, pp. 1526–1540.
- Neggers, R. et al. (2004). “An Evaluation of Mass Flux Closures for Diurnal Cycles of Shallow Cumulus”. *Monthly Weather Review* 132 (2004) 11 132.
- Nordeng, T. E. (1994). “Extended versions of the convective parametrization scheme at ECMWF and their impact on the mean and transient activity of the model in the tropics.” *Research Department Technical Memorandum* 206, pp. 1–41.
- Ogura, Y. and N. A. Phillips (1962). “Scale Analysis of Deep and Shallow Convection in the Atmosphere.” *Journal of Atmospheric Sciences* 19, pp. 173–179.
- Onsager, L. (1949). “The Distribution of Energy in Turbulence. Abstract”. *Phys. Rev* 68.11, p. 286.
- Ozmidov, R. (1965). “On the turbulent exchange in a stably stratified ocean”. *Izv. Acad. Sci. USSR, Atmos. Oceanic Phys* 1.8, pp. 853–860.
- Pan, D.-M. and D. D. Randall (1998). “A cumulus parameterization with a prognostic closure”. *Quarterly Journal of the Royal Meteorological Society* 124.547, pp. 949–981.
- Parker, M. D. and R. H. Johnson (2000). “Organizational Modes of Midlatitude Mesoscale Convective Systems”. *Monthly Weather Review* 128.10, p. 3413.
- Patton, E. G. et al. (2005). “The Influence of Idealized Heterogeneity on Wet and Dry Planetary Boundary Layers Coupled to the Land Surface.” *Journal of Atmospheric Sciences* 62, pp. 2078–2097.
- Peters, J. M. (2016). “The impact of effective buoyancy and dynamic pressure forcing on vertical velocities within two-dimensional updrafts”. *Journal of the Atmospheric Sciences* 73.11, pp. 4531–4551.
- Peters, J. M. et al. (2019). “The Influence of Vertical Wind Shear on Moist Thermals”. *Journal of the Atmospheric Sciences* 76.6, pp. 1645 –1659.

- Pinkus, M. B. and A. P. Khain (2002). “Effects of in-cloud nucleation and turbulence on droplet spectrum formation in cumulus clouds”. *Quarterly Journal of the Royal Meteorological Society* 128.580, pp. 501–533.
- Pino, D. et al. (2006). “Representing sheared convective boundary layer by zeroth- and first-order-jump mixed-layer models: Large-eddy simulation verification”. *Journal of Applied Meteorology and Climatology* 45.9, pp. 1224–1243.
- Pitsch, H. (2006). “Large-Eddy Simulation of Turbulent Combustion”. *Annual Review of Fluid Mechanics* 38.1, pp. 453–482.
- Plant, R. S. (2010). “A review of the theoretical basis for bulk mass flux convective parameterization”. *Atmospheric Chemistry and Physics* 10.8, pp. 3529–3544.
- Pope, S. B. (2000). *Turbulent flows*. Cambridge university press.
- Pressel, K. G. et al. (2015). “Large-eddy simulation in an anelastic framework with closed water and entropy balances”. *Journal of Advances in Modeling Earth Systems* 7, pp. 1425–1456.
- Pruppacher, H. et al. (1998). “Microphysics of Clouds and Precipitation”. *Aerosol Science and Technology* 28, pp. 381–382.
- Randall, D. et al. (2003). “Breaking the Cloud Parameterization Deadlock”. *Bulletin of the American Meteorological Society* 84.11, pp. 1547–1564.
- Rasmussen, E. N. and D. O. Blanchard (1998). “A Baseline Climatology of Sounding-Derived Supercell and Tornado Forecast Parameters”. *Weather and Forecasting* 13.4, pp. 1148–1164.
- Raymond, D. J. and M. J. Herman (2011). “Convective quasi-equilibrium reconsidered”. *Journal of Advances in Modeling Earth Systems* 3.3.
- Ren, H. et al. (2020). “Large-Eddy Simulation of Idealized Hurricanes at Different Sea Surface Temperatures”. *Journal of Advances in Modeling Earth Systems* 12.9, e2020MS002057.
- (2022). “The size characteristics and physical explanation for the radius of maximum wind of hurricanes”. *Atmospheric Research* 277, p. 106313.
- Rio, C. et al. (2013). “Control of deep convection by sub-cloud lifting processes: the ALP closure in the LMDZ5B general circulation model”. *Climate Dynamics* 40.9, pp. 2271–2292.

- Rio, C. et al. (2019). “Ongoing Breakthroughs in Convective Parameterization”. *Current Climate Change Reports* 5.2, pp. 95–111.
- Rochetin, N. et al. (2014). “Deep convection triggering by boundary layer thermals. Part II: Stochastic triggering parameterization for the LMDZ GCM”. *Journal of the Atmospheric Sciences* 71.2, pp. 515–538.
- Romps, D. (2010). “A Direct Measure of Entrainment”. *Journal of The Atmospheric Sciences - J ATMOS SCI* 67, pp. 1908–1927.
- Rosswog, S. (2009). “Astrophysical smooth particle hydrodynamics”. *New Astronomy Reviews* 53.4, pp. 78–104. ISSN: 1387-6473.
- Rotunno, R. et al. (1988). “A Theory for Strong, Long-Lived Squall Lines.” *Journal of Atmospheric Sciences* 45.3, pp. 463–485.
- Rotunno, R. et al. (2009). “Large-Eddy Simulation of an Idealized Tropical Cyclone”. *Bulletin of the American Meteorological Society* 90.12, pp. 1783–1788.
- Rybka, H. et al. (2021). “The behavior of high-CAPE (convective available potential energy) summer convection in large-domain large-eddy simulations with ICON”. *Atmospheric Chemistry and Physics* 21.6, pp. 4285–4318.
- Sagaut, P. and Y.-T. Lee (2002). “Large Eddy Simulation for Incompressible Flows: An Introduction. Scientific Computation Series”. *Applied Mechanics Reviews* 55, pp. 115–.
- Sanderson, B. M. et al. (2010). “Climate feedbacks determined using radiative kernels in a multi-thousand member ensemble of AOGCMs”. *Climate Dynamics* 35.7, pp. 1219–1236.
- Sansom, R. (2021). “Demystifying the complex nature of stratocumulus clouds with machine learning”. *Weather* 76.6, pp. 185–191.
- Sathiyamoorthy, V. et al. (2004). “Influence of the Upper-Tropospheric Wind Shear upon Cloud Radiative Forcing in the Asian Monsoon Region.” *Journal of Climate* 17.14, pp. 2725–2735.
- Seifert, A. et al. (2015). “Large-eddy simulation of the transient and near-equilibrium behavior of precipitating shallow convection”. *Journal of Advances in Modeling Earth Systems* 7, pp. 1918–1937.

- Sherwood, S. et al. (2014). “Spread in model climate sensitivity traced to atmospheric convective mixing”. *Nature* 505, pp. 37–42.
- Shima, S. et al. (2009). “The super-droplet method for the numerical simulation of clouds and precipitation: a particle-based and probabilistic microphysics model coupled with a non-hydrostatic model”. *Quarterly Journal of the Royal Meteorological Society* 135, pp. 1307–1320.
- Siebesma, A. P. and J. W. M. Cuijpers (1995). “Evaluation of Parametric Assumptions for Shallow Cumulus Convection”. *Journal of the Atmospheric Sciences* 52.6, pp. 650–666.
- Singh, K. et al. (2018). “Prediction of heavy rainfall over Chennai Metropolitan City, Tamil Nadu, India: Impact of microphysical parameterization schemes”. *Atmospheric Research* 202, pp. 219–234.
- Singh, K. S. et al. (2022). “Impact of Time Step Size on Different Cumulus Parameterization Schemes in the Numerical Simulation of a Heavy Rainfall Event Over Tamil Nadu, India”. *Pure and Applied Geophysics* 179.1, pp. 399–423.
- Slingo, A. (1989). “A GCM Parameterization for the Shortwave Radiative Properties of Water Clouds”. *Journal of Atmospheric Sciences* 46.10, pp. 1419–1427.
- Slingo, J. and T. Palmer (2011). “Uncertainty in weather and climate prediction”. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 369.1956, pp. 4751–4767.
- Smagorinsky, J. (1963). “General Circulation Experiments with the Primitive Equations”. *Monthly Weather Review* 91, p. 99.
- Smilgys, R. and I. A. Bonnell (2017). “Formation of stellar clusters”. *Monthly Notices of the Royal Astronomical Society* 472.4, pp. 4982–4991.
- Sommeria, G. (1976). “Three-Dimensional Simulation of Turbulent Processes in an Undisturbed Trade Wind Boundary Layer.” *Journal of Atmospheric Sciences* 33, pp. 216–241.
- Sorbjan, Z. (1999). “Similarity of Scalar Fields in the Convective Boundary Layer”. *Journal of the Atmospheric Sciences* 56.13, pp. 2212–2221.

- Sorbjan, Z. (2005). “Statistics of scalar fields in the atmospheric boundary layer based on large-eddy simulations. Part 1: Free convection”. *Boundary-layer meteorology* 116.3, pp. 467–486.
- Sreenivasan, K. R. (1995). “On the universality of the Kolmogorov constant”. *Physics of Fluids* 7.11, pp. 2778–2784.
- Stenke, A et al. (2009). “Implications of Lagrangian transport for coupled chemistry-climate simulations”. *Atmos. Chem. Phys. Discuss* 8.18, pp. 727–18.
- Stenke, A. et al. (2008). “Lagrangian transport of water vapor and cloud water in the ECHAM4 GCM and its impact on the cold bias”. *Climate dynamics* 31.5, pp. 491–506.
- Stevens, B. et al. (1999). “Large-Eddy Simulations of Radiatively Driven Convection: Sensitivities to the Representation of Small Scales.” *Journal of Atmospheric Sciences* 56, pp. 3963–3984.
- Stevens, B. et al. (2005). “Evaluation of Large-Eddy Simulations via Observations of Nocturnal Marine Stratocumulus”. *Monthly Weather Review* 133, p. 1443.
- Stevens, B. et al. (2020). “The Added Value of Large-eddy and Storm-resolving Models for Simulating Clouds and Precipitation”. *Journal of the Meteorological Society of Japan. Ser. II* 98.2, pp. 395–435.
- Stoll, R. et al. (2020). “Large-Eddy Simulation of the Atmospheric Boundary Layer”. *Boundary-Layer Meteorology* 177.2, pp. 541–581.
- Stull, R. B. (1976). “The energetics of entrainment across a density interface”. *Journal of Atmospheric Sciences* 33.7, pp. 1260–1267.
- (1988). *An introduction to boundary layer meteorology*. Vol. 13. Springer Science & Business Media.
- Sullivan, P. P. et al. (2014). “Large-Eddy Simulation of Marine Atmospheric Boundary Layers above a Spectrum of Moving Waves”. *Journal of Atmospheric Sciences* 71, pp. 4001–4027.
- Sullivan, P. P. and E. G. Patton (2011). “The Effect of Mesh Resolution on Convective Boundary Layer Statistics and Structures Generated by Large-Eddy Simulation”. *Journal of the Atmospheric Sciences* 68.10, pp. 2395–2415.

- Sullivan, P. P. et al. (1998). “Structure of the entrainment zone capping the convective atmospheric boundary layer”. *Journal of the Atmospheric Sciences* 55.19, pp. 3042–3064.
- Swann, H. (1998). “Sensitivity to the representation of precipitating ice in CRM simulations of deep convection”. *Atmospheric Research* 47, pp. 415–435.
- Swegle, J. et al. (1995). “Smoothed Particle Hydrodynamics Stability Analysis”. *Journal of Computational Physics* 116.1, pp. 123–134.
- Takemi, T. (2006). “Impacts of moisture profile on the evolution and organization of midlatitude squall lines under various shear conditions”. *Atmospheric Research* 82.1-2, pp. 37–54.
- Takemi, T. and R. Rotunno (2003). “The Effects of Subgrid Model Mixing and Numerical Filtering in Simulations of Mesoscale Cloud Systems”. *Monthly Weather Review* 131.9, pp. 2085–2101.
- Taylor, G. I. and A. E. Green (1937). “Mechanism of the production of small eddies from large ones”. *Proceedings of the Royal Society of London. Series A - Mathematical and Physical Sciences* 158.895, pp. 499–521.
- Thompson, R. L. et al. (2003). “Close Proximity Soundings within Supercell Environments Obtained from the Rapid Update Cycle”. *Weather and Forecasting* 18.6, pp. 1243–1261.
- Tiedtke, M. (1993). “Representation of Clouds in Large-Scale Models”. *Monthly Weather Review* 121.11, pp. 3040–3061.
- Träumner, K et al. (2011). “Convective boundary-layer entrainment: Short review and progress using Doppler lidar”. *Boundary-layer meteorology* 141, pp. 369–391.
- Vacondio, R. et al. (2021). “Grand challenges for Smoothed Particle Hydrodynamics numerical schemes”. *Computational Particle Mechanics* 8.3, pp. 575–588.
- Vallis, G. K. (2017). *Atmospheric and Oceanic Fluid Dynamics: Fundamentals and Large-Scale Circulation*. 2nd ed. Cambridge University Press.
- Van Heerwaarden, C. C. et al. (2017). “MicroHH 1.0: a computational fluid dynamics code for direct numerical simulation and large-eddy simulation of atmospheric boundary layer flows”. *Geoscientific Model Development* 10.8, pp. 3145–3165.

- Van Heerwaarden, C. C. and J. P. Mellado (2016). “Growth and decay of a convective boundary layer over a surface with a constant temperature”. *Journal of the Atmospheric Sciences* 73.5, pp. 2165–2177.
- Vanzanten, M. C. et al. (1999). “Entrainment parameterization in convective boundary layers”. *Journal of the Atmospheric Sciences* 56.6, pp. 813–828.
- Vincent, A. and M. Meneguzzi (1994). “The dynamics of vorticity tubes in homogeneous turbulence”. *Journal of Fluid Mechanics* 258, 245–254.
- Violeau, D. and B. Rogers (2016). “Smoothed particle hydrodynamics (SPH) for free-surface flows: past, present and future”. *Journal of Hydraulic Research* 54, pp. 1–26.
- Wagner, C. et al. (2007). *Large-eddy simulation for acoustics*. Vol. 20. Cambridge University Press.
- Walker, G. T. (1923). “Correlation in seasonal variation of weather. VIII : A preliminary study of world weather”. *Memoirs of India Meteorological Department* 24, pp. 75–131.
- Wang, C. and R. G. Prinn (1998). “Impact of the horizontal wind profile on the convective transport of chemical species”. *Journal of Geophysical Research: Atmospheres* 103.D17, pp. 22,063–22,071.
- Warren, R. A. et al. (2017). “Impact of variations in upper-level shear on simulated supercells”. *Monthly Weather Review* 145.7, pp. 2659–2681.
- Weisman, M. L. (1992). “The role of convectively generated rear-inflow jets in the evolution of long-lived mesoconvective systems”. *Journal of Atmospheric Sciences* 49.19, pp. 1826–1847.
- Weisman, M. L. and R. Rotunno (2000). “The use of vertical wind shear versus helicity in interpreting supercell dynamics”. *Journal of the Atmospheric Sciences* 57.9, pp. 1452–1472.
- Wu, X. et al. (1998). “Long-Term Behavior of Cloud Systems in TOGA COARE and Their Interactions with Radiative and Surface Processes. Part I: Two-Dimensional Modeling Study.” *Journal of Atmospheric Sciences* 55, pp. 2693–2714.

- Yamaguchi, T. et al. (2019). “Aerosol-Cloud Interactions in Trade Wind Cumulus Clouds and the Role of Vertical Wind Shear”. *Journal of Geophysical Research: Atmospheres* 124.22, pp. 12244–12261.
- Yanai, M. et al. (1973). “Determination of Bulk Properties of Tropical Cloud Clusters from Large-Scale Heat and Moisture Budgets”. *Journal of Atmospheric Sciences* 30.4, pp. 611–627.
- Yano, J.-I. et al. (2013). “Phenomenology of convection-parameterization closure”. *Atmospheric Chemistry and Physics* 13.8, pp. 4111–4131.
- Zerroukat, M. et al. (2002). “SLICE: A Semi-Lagrangian Inherently Conserving and Efficient scheme for transport problems”. *Quarterly Journal of the Royal Meteorological Society: A journal of the atmospheric sciences, applied meteorology and physical oceanography* 128.586, pp. 2801–2820.
- Zhang, X. et al. (2021). “Evaluation and Comparison of Two Deep Convection Parameterization Schemes at Convection-Permitting Resolution”. *Monthly Weather Review* 149.10, pp. 3419–3432.
- Zhao, M. and P. H. Austin (2005). “Life Cycle of Numerically Simulated Shallow Cumulus Clouds. Part II: Mixing Dynamics”. *Journal of the Atmospheric Sciences* 62.5, pp. 1291–1310.
- Zilitinkevich, S. S. (1991). *Turbulent penetrative convection*. Avebury Technical.
- Zipser, E. J. (1994). “Deep cumulonimbus cloud systems in the tropics with and without lightning”. *Monthly weather review* 122.8, pp. 1837–1851.
- Zipser, E. J. and K. R. Lutz (1994). “The vertical profile of radar reflectivity of convective cells: A strong indicator of storm intensity and lightning probability?” *Monthly Weather Review* 122.8, pp. 1751–1759.