

POLITECNICO DI TORINO

Master's Degree in Physics of Complex Systems



**Politecnico
di Torino**

Master's Degree Thesis

Relative dispersion and Lagrangian correlations in inhomogeneous turbulence: experimental and numerical study

Supervisor

Prof. Daniela TORDELLA

Candidate

Niccolò GALLINO

a.a. 2023-24

Abstract

This work explores two aspects of the motion of particles inside a turbulent fluid flow: relative dispersion and Lagrangian correlations. These are studied experimentally, by means of custom-made miniaturized radiosondes launched in clusters to obtain data about the turbulence in the atmospheric boundary layer, and numerically, by means of Direct Numerical Simulations of a cloud's border region, modeled as a turbulent mixing layer in which the advected Lagrangian particles are water droplets.

Relative dispersion is one of the open problems of fluid turbulence, concerning how fast particles are spread apart by the fluid flow. The relevant theoretical framework is Kolmogorov's and Obukhov's K41 theory, but this only holds for the case of stationary, homogeneous, isotropic turbulence. However, real flows rarely hold these properties, and thus studies of dispersion far from these ideal conditions are necessary both for theoretical insight into real turbulence and for correct modelling of real-world phenomena and forecasting. Our experimental results are compared with those from K41, in terms of the exponents and coefficients found for the dispersion laws, and deviations are found that can be explained in terms of the inhomogeneities inherently contained in atmospheric turbulence. On the numerical side, preliminary studies are conducted on the currently available simulations, that lack the necessary features to allow for definitive investigation, in order to pave the way for newer upcoming simulations.

Lagrangian correlations are a far less explored branch in the study of turbulence, and while some specific studies have been conducted there is no solid and widely recognized theoretical framework to interpret them. We thus present novel results about these quantities both for velocity components, which are what most studies investigate, and other relevant quantities for the cases of the cloud border simulations and the radiosondes released in the atmosphere. Symmetrically with respect to the discussion on dispersion, results on Lagrangian correlations see a stronger contribution from the numerical simulations.

Acknowledgements

First and foremost I would like to thank Prof. Daniela Tordella for the (sometimes literal) journey that this thesis entailed. I think it is fair to say that my experience as a Master thesis writer was unlike that of most, and for this I am deeply thankful. Never would I have thought to have half the opportunities she offered me.

In second I need to thank Shahbozbek Abdunabiev, PhD student in Prof. Tordella's group, for being a guide in a great many occasions. The amount of things that I have learned from him can hardly be overstated.

In third I thank all people I have met while carrying out my work: everyone at the Chilbolton Observatory who welcomed us and helped us in our experiments, as well as all members of the MET Office and NCAS who continuously collaborated with us. With the same breath I thank all of the staff of the Universitat Politècnica de València that made my participation in ETC18 a memorable experience. Of the above I thanked none by name, to counteract my forgetfulness, but I will break this rule for two people: Bianca Tenti and Prof. Ladislau Matekovits, precious companions during our travels in England.

Finally, I thank all people who were rooting for me, friends and family, out of which two other names must be specially called out: mom and dad.

“ ‘It could be worse.’
‘How?’
‘It could be raining.’ ”

Young Frankenstein

Table of Contents

List of Figures	VII
Introduction	1
1 Turbulent flow and particles	7
1.1 General features of turbulence	8
1.2 The Kolmogorov - Obukhov scaling theory	10
1.2.1 Choice of parameters and similarity hypotheses	11
1.2.2 Structure functions	13
1.2.3 Energy spectrum	14
1.2.4 Turbulent diffusion coefficient	15
1.3 Turbulent relative dispersion	16
1.3.1 Richardson formalism: distance-neighbor-graph	17
1.3.2 Modern approach: scaling theory	20
1.3.3 Refinements, criticism and recognized results	21
1.4 Lagrangian correlations	23
2 Simulation and phenomenology of a cloud border region	27
2.1 Cloud physics in a nutshell	28
2.1.1 Parcel theory for atmospheric supersaturation	29
2.1.2 Water vapor at equilibrium with droplets	31
2.1.3 Growth of water droplets by diffusion	34
2.1.4 Collision and coalescence	36
2.2 Simulation of a cloud-clear air interface	37
2.2.1 Fields simulation: velocity, temperature, water vapor	37
2.2.2 Droplet simulation: Lagrangian dynamics and growth	41
2.3 Results so far: cloud top boundary phenomenology	42
2.3.1 Acceleration of droplet population dynamics	43
2.3.2 Time and space-dependent collision kernel	44
2.3.3 Microphysical timescales	46
2.3.4 Supersaturation balance	47

3	Experimental setup: radiosonde cluster	55
3.1	Idea and main requirements	56
3.2	The radiosondes	58
3.3	Balloons	61
3.4	Lab tests	64
3.4.1	Temperature and humidity tests	64
3.4.2	Position	66
3.5	In-field tests	66
3.5.1	Sondes vs. other instruments	66
3.5.2	Cluster tests	70
4	Results on relative dispersion	77
4.1	Experimental results: Launch 1	79
4.1.1	Environmental characterisation	80
4.1.2	Dispersion results	83
4.2	Experimental results: Launch 2	90
4.2.1	Environmental characterisation	91
4.2.2	Dispersion results	91
4.2.3	Discussion of the results	97
4.3	Numerical results: preparatory analysis	99
4.3.1	Visualizations	101
4.3.2	Mean square separation	103
5	Results on Lagrangian correlations	107
5.1	Numerical results	109
5.1.1	Monodisperse distribution, unstable stratification	111
5.1.2	Polydisperse distribution, unstable stratification	113
5.1.3	Polydisperse distribution, stable stratification	116
5.1.4	Discussion of results	121
5.2	Experimental results	122
5.2.1	Launch 1	122
5.2.2	Launch 2	129
	Conclusions	135
	Bibliography	141

List of Figures

1.1	Energy spectrum of a fluid in turbulent motion in logarithmic space.	16
1.2	Distance-neighbour graph and its evolution in time (not normalised, hence the high values). Values obtained in the context of the simulation of a cloud top border, see Chapter 2.	19
2.1	Köhler curve (solid black) for droplets formed on hydrosoluble aerosol. The Kelvin curve plotted in dashed black allows us to appreciate the influence of the Raoult effect: in the absence of aerosol, small droplets require very high values of saturation to form and be stable.	34
2.2	(a) Geometry of the computational domain. (b) Visualization of the vapor concentration in the cloud portion of the domain. Plots for the kinetic energy, vapor density and temperature. [30]	38
2.3	Spatial and temporal evolution for kinetic energy and dissipation rate in the computational domain [30]. (d) includes plots for the Kolmogorov length η [31].	39
2.4	Temporal evolution of droplet populations with a monodisperse initial distribution. Top plots represent number density, bottom plots the mass density [31].	44
2.5	Temporal evolution of droplet populations with a polydisperse initial distribution. Top plots represent number density, bottom plots the mass density [31].	45
2.6	Time and space-dependent values for the collision kernel in the monodisperse initial population case. On the left are the values for the interface region, on the right for the cloud bulk. Time evolution can be observed in three steps from top to bottom [31].	49
2.7	Time and space-dependent values for the collision kernel in the polydisperse initial population case. On the left are the values for the interface region, on the right for the cloud bulk. Time evolution can be observed in three steps from top to bottom [31].	50

2.8	Pearson correlation coefficient between standard deviation (σ), skewness (S) and kurtosis (K) of the spatial derivative in the anisotropic x_3 direction of the corresponding velocity component and the number of collisions throughout the transient [31].	51
2.9	Plots for the spatial variation of microphysical timescales inside cloud top border simulations at times close to the starting and ending points, along with supersaturation plots [30].	52
2.10	Excess in the balance of supersaturation evolution terms, alongside covariances of the supersaturation S and velocity spatial derivatives [30].	53
2.11	Correlation coefficients (blue) between the turbulent supersaturation production term and its proposed forms in terms of covariances. Multiplicative constant (red) that would be required for the proposed term to account for numerical values [30].	54
3.1	Experimental setup for radiosonde usage [35].	56
3.2	(a) Complete setup of one radiosonde with its MaterBi balloon. (b) Radiosonde and connected battery [35].	59
3.3	Block diagram of the radiosonde components [36].	59
3.4	Current design of MaterBi “petals” for production of biodegradable balloons.	63
3.5	Plot that links balloon radius to the attainable altitude for different values of the payload m_r , considering MaterBi as the balloon material [35].	63
3.6	Temperature test of implemented sensors inside the climatic chamber. Humidity is kept constant at RH = 30% while temperature is increased from -5 °C to 10 °C. [36].	65
3.7	Comparison between a smartphone GPS signal and our GNSS receiver, test carried out by simultaneous movement of the two devices on Politecnico grounds. [36]	67
3.8	Number of packets received during the 2021 Levaldigi dual launch as a function of (a) time and (b) altitude. [36]	68
3.9	Trajectory of the 2023 Vaisala-COMLETE dual launch in Chilbolton, UK.	68
3.10	Comparison of position readings for the 2021 Levaldigi dual launch. Our radiosonde is marked as COMPLETE. The inset plot in (b) shows the difference between the two readings for a selected time range. [35]	69

3.11	(a) Comparison of wind speed readings for the 2021 Levaldigi dual launch. Our radiosonde is marked as COMPLETE, and the raw data is indicated by the blue dots, while the red line is given by interpolation (some of the over/under-shoot is an artifact of the interpolation, whereas some is an actual feature of the dataset. (b) Power spectrum from the velocity readings of the probes. [35] . . .	70
3.12	Temperature, Pressure and Humidity comparison for the Levaldigi Vaisala-COMPLETE dual launch. [35]	71
3.13	Comparison between two dual launches with different configurations of attachment of the Vaisala and COMPLETE sondes. The larger error experienced in the Chilbolton launch by our sensors may be due to heating from the other sonde. [35]	72
3.14	Left: Comparison of readings from multiple sondes and a Vaisala weather station at INRiM facilities. Right: Readings from a cluster of radiosondes in a tethered launch. [35]	73
3.15	Comparison of readings from radiosonde cluster and reference sensors provided by INRiM. In yellow and purple (Ref USH 1 and Ref USH 2) reference sensors without shielding from solar radiation, and in green (Ref SH) sensor with shielding. It is clear that solar radiation has a considerable effect on the readings, and the sondes are in good agreement with the unshielded sensors. [35]	74
3.16	Trajectories of all sondes during the 2022 OAVdA cluster launch, with the color map indicating altitude. [35]	75
3.17	Temperature and humidity readings of all sondes for the 2022 OAVdA cluster launch. [35]	75
3.18	Energy spectra for 3 probes as computed from the OAVdA launch dataset. Dotted lines are power laws plotted for reference. [35] . . .	76
4.1	Map of the July 5th launch at Chilbolton observatory with cluster trajectories.	79
4.2	Spreading of the sondes in the three directions North, East and up. Launch 1.	80
4.3	Wind profiling by the Chilbolton radar instrumentation on July 5th, day of our first launch. Launch time was $\sim 13:52$. Top plot shows updraft, bottom plot shows horizontal wind speed and direction, the direction being identified by the orientation of the L shapes. . . .	81
4.4	Measurements from the Vaisala sonde launched during the July 5th IOP. Sonde launched at 14:00, approximately 8 minutes after our cluster launch.	82
4.5	Fraction of area above the Chilbolton observatory occupied by clouds as modeled by the ECMWF IFS forecast [38].	83

4.6	Distance-neighbor graph time evolution for the experimental dataset of launch 1. Not normalized. Size of intervals $h = 100\text{m}$ (refer to Chapter 1).	85
4.7	Interpolated distance-neighbor graph color-map visualizations. The vertical axis indicates the distance intervals, the horizontal axis is time, and the color indicates the number of sondes as given by interpolation.	86
4.8	Logarithmic plot of the mean square separation distance in time, launch 1. 21 sonde pairs.	87
4.9	Logarithmic plot of the mean square separation distance in time, launch 1, power-law regime only. Cut-off time for the first, spurious regime set at 470 seconds to eliminate all sources of uncertainty. 21 sonde pairs.	88
4.10	Power-law fitting of the mean separation distance in time. Exponent, coefficient and Relative Root Mean Square Error are reported under the legend.	89
4.11	Map of the July 6th launch at Chilbolton observatory with cluster trajectories.	90
4.12	Spreading of the sondes in the three directions North, East and up. Chilbolton, Launch 2.	91
4.13	Wind profiling by the Chilbolton radar instrumentation on July 6th (Launch 2). Launch time was $\sim 13:45$. Top: color map indicates updraft. Bottom: color map indicates horizontal wind speed, L shapes indicate wind direction in compass coordinates, i.e. upwards being north and right being east. [39]	92
4.14	Measurements from the Vaisala sonde launched on July 6th at 11:00. Launch time was 13:30, but no closer launches were available due to July 6th not being included in an IOP.	93
4.15	Fraction of area above the Chilbolton observatory occupied by clouds as modeled by the ECMWF IFS forecast [39].	93
4.16	Distance-neighbor graph time evolution for the experimental dataset of launch 2. Not normalized. Size of intervals $h = 100\text{m}$ (refer to Chapter 1).	94
4.17	Interpolated distance-neighbor graph color-map visualizations for launch 2. The vertical axis indicates the distance intervals, the horizontal axis is time, and the color indicates the number of sondes as given by interpolation.	95
4.18	Logarithmic plot of the mean square separation distance in time, launch 2. 45 sonde pairs. A straight line would be expected, since we expect a power-law behavior, but instead we find a more exponential-like curve.	96

4.19	Power-law fitting of the mean separation distance in time for Chilbolton Launch 2. Exponent, coefficient and Relative Root Mean Square Error are reported under the legend.	97
4.20	Exponential fitting of the mean separation distance in time for Chilbolton Launch 2. Exponent coefficient, multiplicative coefficient and Relative Root Mean Square Error are reported under the legend.	98
4.21	Plots of the droplet positions in time in the cloud top boundary simulations. Plots given for the currently available combinations of initial droplet populations and thermal stratification: monodisperse unstable (right), polydisperse unstable (middle), polydisperse stable (right).	102
4.22	Q plots time evolution for the available combinations of stability conditions and initial droplet populations. Left plots are for droplets starting in the cloud bulk, right plots are for droplets starting in the mixing layer at the cloud border.	104
4.23	Mean square separation for the available combinations of stability conditions and initial droplet populations. Plots for droplets starting both in the cloud bulk (blue) and in the mixing region (red).	105
5.1	Distribution along homogeneous axes of the particles used for Lagrangian correlations computation. Five cubic volumes with a side of 2 cm were selected, and then droplets starting out inside these volumes were selected for inhomogeneous axis z values corresponding to the center of the cloud region and the center of the mixing region.	110
5.2	Visualization of the trajectories of 30 randomly marked particles starting out in the five initial volumes in the case of monodisperse population, unstable stratification.	112
5.3	Lagrangian autocorrelations for velocity components in the monodisperse population, unstable stratification case.	113
5.4	Lagrangian autocorrelations for radius and growth rate in the monodisperse population, unstable stratification case.	114
5.5	Visualization of the trajectories of 30 randomly marked particles starting out in the five initial volumes in the case of polydisperse population, unstable stratification.	115
5.6	Lagrangian autocorrelations for velocity components in the polydisperse population, unstable stratification case.	116
5.7	Lagrangian autocorrelations for radius and growth rate in the polydisperse population, unstable stratification case.	117
5.8	Visualization of the trajectories of 30 randomly marked particles starting out in the five initial volumes in the case of polydisperse population, stable stratification.	118

5.9	Lagrangian autocorrelations for velocity components in the polydisperse population, stable stratification case.	119
5.10	Lagrangian autocorrelations for radius and growth rate in the polydisperse population, stable stratification case.	119
5.11	Velocities of the sondes along the three directions East, North and Up in the first Launch carried out in Chilbolton.	123
5.12	Velocity components autocorrelations from Launch 1 in Chilbolton.	124
5.13	Temperature, humidity and pressure measurements from all sondes for Launch 1 in Chilbolton.	125
5.14	Autocorrelations for temperature, relative humidity and pressure obtained from the dataset gathered during Launch 1 in Chilbolton1.	126
5.15	Traslational autocorrelations for velocity components for Launch 1 in Chilbolton.	127
5.16	Traslational autocorrelations for temperature, humidity and pressure for Launch 1 in Chilbolton.	128
5.17	Velocities of the sondes along the three directions East, North and Up in the second Launch carried out in Chilbolton.	129
5.18	Velocity component autocorrelations from Launch 2 in Chilbolton. .	130
5.19	Temperature, humidity and pressure measurements for all sondes for Launch 2 in Chilbolton.	131
5.20	Autocorrelations for temperature, humidity and pressure from Launch 2 in Chilbolton.	132
5.21	Traslational autocorrelations for velocity components for Launch 2 in Chilbolton.	133
5.22	Traslational autocorrelations for temperature, humidity and pressure for Launch 2 in Chilbolton.	134

Introduction

I have always been interested in atmospheric phenomena. Their scale and their beauty appealed to me. Still, it was not a fascination I consciously realised: I just thought it was cool when it rained. Then, towards the end of my Bachelor's, I was considering enrolling in the Physics of Complex Systems Master's program, and as one does I was trying to imagine what exactly I might end up studying if I did decide to take that path, since it is an area of physics less easy to conceptualize than, say, particle physics. One day, all of a sudden and quite by chance, possibly stirred by the vision of some distant cumulus formation, a thought popped up in my head: "Hey, clouds are probably a pretty complex system. It would be fun to study those."

I could not know at the time, but as a matter of fact I hit the nail quite square on its head with that naive realisation. When I approached Prof. Tordella to have more information about her research and the activities of her group I still had little idea of what the study of clouds actually entailed, as there were no courses in my curriculum that went any deeper than a couple paragraphs into it. Specifically, although *a posteriori* it may seem obvious, I was not fully ready for how tight the bond between the study of clouds and fluid dynamics really is. That a bond should exist was clear enough, clouds being made up of fluids moving through a fluids (with the exception of ice in high clouds, but those still arise from the freezing of water vapor mixed with air, so the point remains), but one aspect of it could not but escape me: it is not only with fluid mechanics in general that cloud dynamics are coupled, but specifically with turbulent fluid motion. There is a popular science article by Bodenschatz *et al.* [1] whose title does a rather complete job of summing up what research both old and new has discovered about this link: "*Can we understand clouds without turbulence?*". The implied answer, obviously, is "No". For this reason I unexpectedly had to dive into the wild world of turbulence theory, which I already knew to be among the paramount examples of complexity in physical systems. Luckily, I quickly found that turbulence interested me just as much as clouds did, and therefore working on their relationship proved to be an extremely satisfying experience. However, the theory of turbulence comes in a great many flavors (in fact virtually all flavors of scientific endeavors), and as is often the

case with very fundamental problems the more theoretical parts of the research in this field have become difficult to reconcile with applications. While obviously important, the sort of idealised version of turbulence that is used to investigate it at the fundamental level proves only partially valid when the need is to understand the impact of real-world turbulence on various phenomena. This is mostly due to the fact that the best framework for the theoretical studies, namely homogeneous, isotropic turbulence (HIT), is but a model for real turbulence, which in the great majority of cases cannot display neither homogeneity nor isotropy. This issue will be relayed more thoroughly in Chapter 1, but is noted here in passing to provide a justification for the fact that when needing to investigate the relationship between real turbulence and the phenomena taking place inside it the purely theoretical tools are very seldom used. Rather, they provide a measure for comparison with the results that are obtained, and should be employed as a foundation on which to build more concrete models.

With pure theory having to be put aside, two methods of investigation remain available: numerical simulation and experimental studies. Most of the past works by Prof. Tordella's group were numerical in nature, so when I joined I was prepared for some type of computational effort. More recently though, a new avenue had been opened thanks to an H2020 project by the name of COMPLETE. The group started developing an innovative experimental setup for atmospheric measurements aimed at cloud research, based around the deployment of miniaturized radiosonde clusters. By its nature in-field experimental work requires a lot of manpower, and therefore despite not being my main source of interest at that time I happily took part in the organisation of the first free-flight launch of the sondes. It was so that the experimental type of investigation started to grow on me, and I slowly appreciated more and more how complementary computational and experimental research are. Obviously enough both have their strengths and their weaknesses, but at the beginning I had not quite realised that in many instances what proves hard in one case is easy in the other. For a while then I just went with the flow and participated in the development of both the experimental and the numerical sides of the work. Because of their intersection, though, it felt unfair to only include one of the two in the final account of my work, namely this thesis. It was thus decided to include both, in an effort to highlight the importance of each approach.

At this point, I should start to be more precise about both the simulations that were employed to obtain the numerical results and the aforementioned experimental setup. They will each be presented in full detail in their respective chapters, but a shorter account may be given here for the reader that might still be deciding whether those are worth reading. First and foremost I should specify that I did not run any new simulations myself (this might hopefully change soon, as told in the Conclusions chapter), as computational fluid dynamics is notoriously extremely expensive in terms of computational resources and performing new runs requires

a somewhat lengthy process of approval from the institutions that may provide such resources. Rather, I performed novel post-processing analysis on the datasets obtained from the most recent simulations run by the group shortly before my arrival. These simulations are Direct Numerical Simulations of the fluid flow and water droplets dynamics at the border of a cloud. This has been identified as a region worthy of interest owing precisely to the inhomogeneity of the turbulence therein, due to the fact that clouds, broadly speaking, are more turbulent than the clear air surrounding them. The experimental setup, instead, consists of the radiosondes, attached to fixed-size helium balloons for neutral buoyancy at a predefined altitude, in communication with one or more ground stations for data reception. The sondes are to be released in clusters, which is a crucial aspect of the methodology: indeed, many measurements all over the world are performed by means of single sondes, which are unable to provide same-time, different-space point information, necessary to understand the structure and statistical behavior of the turbulent flow.

With this rough explanation of the methods in place we can now discuss exactly which aspects of turbulence have been investigated. In doing this we should keep in mind our ultimate goal: understand how turbulence impacts the behavior of clouds. This is because, at the basic level of modelling, clouds can be seen as systems that are *advected by* the turbulence, and in more than one sense. At the microscopic level this should be quite intuitive: water droplets or ice crystals are carried around by the air flow. At the level of intermediate scales, though, this is also true: this interpretation is given by what is called “parcel theory”, which basically treats small portions of air as if they were particles moving in a fluid. This has its applications, as exposed in Chapter 2, but for now it simply serves the purpose of illustrating the fact that we are interested in how turbulence transports whatever the fluid may contain. The adjective for this type of studies is “Lagrangian”, which refers to the Lagrangian approach to fluid mechanics, meaning an approach focusing on describing the behavior of the fluid particles, contrasted by the Eulerian approach that treats the fluid as a field (and is in fact the most commonly used). In this thesis, two problems in Lagrangian turbulence are included, reflective of the direction in which we are developing the research. One is often considered to be among the central issues in the field, as well as of turbulence at large; the other instead has so far not been looked into as deeply by the community, and is only recently surfacing in a few newer studies. This sort of juxtaposition was not necessarily intended, but it provides for a nice parallel.

The first one is the problem of relative dispersion. The question as its core is, simply put: “How are particles spread apart by turbulent flow?”. We will be more precise in the following, but this will suffice for now. The reason why it is held in such high consideration is that there is a particularly nice result that answers the question for the case of HIT, called the Richardson-Obukhov law. Despite

its simplicity and elegance, this result has never received definitive confirmation even in its simple framework, which explains its status as a core issue. The second one, instead, is the topic of Lagrangian correlations. What this means is the study of the time and space dependence of correlation functions of various quantities as they evolve over Lagrangian trajectories of particles. Their importance, other than fundamental theoretical investigation, lies mostly in the modelling of various Lagrangian aspects of turbulence, but their research has long been hindered by the inherent difficulties of carrying out Lagrangian measurements. Our environments, both computational and experimental, provide instead a convenient context in this sense.

Some words should be spent on the experiments that allowed us to obtain the datasets that lead to the results presented here. They were performed during our participation in the Wessex Convection (WESCON) 2023 experimental campaign, organised by the British Met Office. For 5 days we were kindly taken care of and aided by the folks at the Chilbolton Observatory, where we carried out two separate cluster launches with 7 and 10 sondes respectively. A special thanks must go to Darcy Ladd, Jeremy Price and Chris Walden, who have been our main points of contact with the British research apparatus and whose contribution has been (and hopefully will continue to be) crucial in the success of our experiments and their analysis. Of course they were not alone, and my thanks are extended to all other people that helped us both during our permanence in England and afterwards.

One last piece of information before coming to the content of the thesis: the results about relative dispersion, both the experimental ones and the preliminary computational work, were presented in two separate sections both by me and Shahbozbek Abdunabiev at the 18th European Turbulence Conference on September 5th 2023, in Valencia (which incidentally also turned out to be the last European Turbulence Conference, since starting in 2024 it will be merged with the European Fluid Mechanics Conference to give rise to the European Fluid Dynamics Conference). My thanks in this case goes to the lead organizer Prof. Sergio Hoyas, as well as to all the people at the Universitat Politècnica de València that made this event an extremely fun experience for all of us to attend.

We can finally come to the matter at hand and illustrate the structure of my work:

Chapter 1 contains a brief account of the theoretical framework that is necessary to understand the subsequent results and appreciate their relevance. Some space is given to the Kolmogorov-Obukhov theory of turbulence, which allows us to then step comfortably into relative dispersion. Finally the concept of Lagrangian correlation is introduced.

Chapter 2 starts with a very condensed summary of the principal aspects of cloud physics, focusing on the ones that are necessary to discuss the numerical simulations of cloud borders, which are presented right after. Along with their

general aspects some results that have been achieved so far thanks to those same simulations are also illustrated, to further justify our interest in this special region of a cloud's body.

Chapter 3 is the last preparatory chapter, and it gives a complete overview of the experimental setup. All of the components of the sondes are listed and the rationale behind their design explained, along with an account of many tests that have been conducted to ascertain their viability as measurement tools.

Chapter 4 relates the results in the matter of relative dispersion. Most of the chapter is taken up by the experimental results which are presented for the two Chilbolton launches separately. The numerical results follow, which for reasons that will be discussed at the beginning of the appropriate section mostly consist in a preliminary analysis, and serve as a stepping stone for upcoming studies.

Lastly, Chapter 5 relates the results for Lagrangian correlations. In this case numerical results come first, the simulations serving as an ideal ground for this type of analysis. Then the experimental results close the discussion. Importantly, Lagrangian correlations are most often computed for the velocity components of particles, whereas I present results for other characterising quantities as well.

While many results are shown, much work remains to be done on both fronts discussed, and many more results are on their way as I write these lines. Unfortunately, time is a deciding factor in how much can fit into a Master's Degree thesis work, and as a matter of fact one can imagine that if I wanted to include all possible future developments stemming from the illustrated results this thesis would probably never be concluded. For this reason, in the Conclusions I did my best to give an overlook of the direction the research here presented is moving towards in the near future, so that the interested reader may be able to be on the lookout for future updates.

Chapter 1

Turbulent flow and particles



The aim of this chapter is to provide the necessary background to contextualise the subsequent results. Naturally, the immense topic of turbulence has to be shrunk to a very serious degree in order to fit in an introductory chapter. We will therefore limit ourselves to the very backbone of the theory, meaning the elements that it fundamentally relies on combined with the specific topics of our interest.

We will begin with an overview of what turbulence actually is and its main characteristics. Then we will skip a chapter that is often found at the very beginning of turbulence textbooks: the semiempirical theories. While obviously extremely important in many areas of science and engineering, they are beyond the scope of this work, in that they often deal with specific practical problems rather than with

some underlying physical content (this is of course a broad, and thus only partially correct, statement: the two things are clearly related). We will instead discuss the theory developed by Kolmogorov, and in parallel by Obukhov, that focuses on the fundamental structure of turbulence, and by some simple assumptions is able to derive some impressively deep results. From this we will move on to the topic of turbulent relative dispersion, real core of our discussion, and illustrate some of its key aspects as well as shortcomings. Finally we will illustrate the topic of Lagrangian correlations, which remains so far somewhat elusive when compared to the body of work discussing other topics in turbulence.

1.1 General features of turbulence

Turbulent flow exists in opposition to laminar flow. The latter is characterised by order and smoothness, whereas the former is characterised by disorder, in the form of very strong fluctuations in the flow velocity (and, by extensions, of many other measurable quantities). The nondeterministic and nonlinear nature of these fluctuations has basically spawned the whole field of research: why do they arise? Can we say anything at all about the resulting flow? How do we deal with them?

As said, the last question is mostly the concern of the semiempirical theories: the problem is far too complicated and far too pervasive to wait for a full solution before it is dealt with in applications, and so many methods have come up to deal with it even without knowing everything that lies beneath. The first question, while interesting even at a fundamental level, obviously has many and lengthy answers, that are provided by the studies in dynamic instability of the flow. We will mostly focus on the question in the middle: what does turbulence look like, and what can we know about it once it's there? This implicitly states that we will position ourselves far from the moment where the flow first becomes unstable, in a regime commonly referred to as *fully developed turbulence*.

To keep the discussion tidy, let us first state the problem a bit more formally, limiting ourselves to the case of incompressible fluids (as we will do throughout this thesis). Given a set of boundary conditions, the problem of turbulence amounts to finding the probability distribution $P(d\omega)$ in the phase space of turbulent flow Ω , in which each point represents an instance of the solenoidal velocity field $\vec{u}(\vec{x}, t)$ (meaning vector fields that satisfy the equations of fluid mechanics) [2]. This is obviously a huge task, and indeed it remains unsolved. Because of its complexity, much work has went into solving small parts of it: rather than solving the full distribution, the focus is usually on finding its first few moments.

With all this in mind we can start to take a look at the features of turbulent flow. The most obvious one for any onlooker is the presence, amidst the aforementioned disorder, of spiraling structures that can be clearly identified. These are called

eddies, and are indeed at the core of turbulent motion. These eddies appear to span a very wide range of scales: while looking at a certain flow the largest ones will typically be of the size of the domain in which the flow is confined, and these will contain others that are smaller and smaller down to a scale that is much smaller than the size of the domain. The largest vortices arise due to instabilities in the mean flow, and they generate the smaller ones by loss of stability in their own flow, and so on to the smallest scale, which by definition must then be stable. The parameter that describes this stability (or lack thereof) is the well known adimensional quantity $\text{Re} = \frac{UL}{\nu}$, called the Reynolds number. Here U is the typical scale of flow velocity for the flow as seen at the spatial scale L , and ν is the kinematic viscosity of the fluid. Thus the Reynolds number expresses the balance between the inertial forces in the fluid and its viscosity. Roughly speaking, the higher the Reynolds number, the more unstable the flow and so the more developed the turbulence, although the exact value at which the transition starts to occur depends on the exact flow we are considering. In the context of the hierarchy of eddies, then, it must be true that at large scales the viscosity plays a very marginal role, so much so that it can be neglected. The Reynolds number then decreases by descending in the length scale of the vortices, down to the smallest scale that, by being stable, indicates that the viscosity has a strong enough effect as to prevent the further breakdown of the flow. Because viscosity is what determines energy dissipation in a fluid this leaves us with the following picture: energy is injected in the largest scales of turbulence by the mean flow; this induces an *energy cascade* that runs from the largest scales through the smaller ones, down to the point in which viscosity becomes prevalent again and the energy is dissipated. At this point it would be unfair, especially in a thesis partially focusing on relative dispersion, not to include the celebrated poem by Richardson, that elegantly encapsulates what we have described so far:

*Big whirls have little whirls,
Which feed on their velocity;
And little whirls have lesser whirls,
And so on to viscosity
(in the molecular sense)*

This simple intuition has a very powerful consequence: while the large scales will be affected by the boundary conditions, and the small scales by the viscosity of the specific fluid, we expect (and indeed observe, at least in some ways) the intermediate scales to present a universal behavior. This leads to a division of all possible scales into different ranges. Calling L the spatial scale of the domain in which the flow is embedded and η the scale of the smallest eddies we will have the following partition: scales $\ell \sim L$ will belong to the *energy range*; scales such

that $\eta \ll \ell \ll L$ will belong to the *inertial range*; scales $\ell \sim \eta$ will belong to the *dissipation range*.

The last concept we introduce before delving into Kolmogorov's theory of turbulence is needed as a bridge towards it, and it is the idea of considering a field of *homogeneous* and *isotropic* turbulence. This framework was first introduced by G.I. Taylor, and has proven extremely useful in the process of making the study of turbulence theoretically approachable (and constitutes the de facto standard in the study of turbulence by statistical physics means). It amounts to assuming that the probability distributions of fluid mechanical quantities are invariant under orthogonal transformations of a system of three-dimensional coordinates [2]. Simply stated, the statistical landscape for our flow will be the same irrespective of the specific point in the domain or the specific direction we look in. We can express this in terms of the velocity correlation functions, which are the main quantity used to characterise the structure of turbulence, as [3]

$$\langle u_\alpha(x)u_\beta(x') \rangle = f_{\alpha\beta}(|x - x'|) = f_{\beta\alpha}(|x - x'|) \quad (1.1)$$

where $f_{\alpha\beta}$ is just some generic function that depends on direction, as specified by the order of the velocity component indices. Clearly this is an assumption that can only ever hope to be valid at sufficiently small scales $\ell \ll L$, as boundary conditions inherently introduce inhomogeneities. Even then it is obvious that most real flows (as we will see later) are quite far from being homogeneous and isotropic, but the approximation is still a highly regarded mathematical tool for the investigation of some fundamental, basic statistical properties of turbulence. This is partly due to the fact that, for sufficiently high Reynolds numbers, many types of flow can be considered to be *locally* homogeneous and isotropic, thus allowing us to disregard the boundary conditions altogether. This is, once again, only valid until sufficiently small scales are considered. If a system presents inhomogeneities of any kind at scales of order l_{in} , the assumption will only hold for $\ell \ll l_{in}$, irrespective of the fact that we may be far from the walls of the domain.

1.2 The Kolmogorov - Obukhov scaling theory

Often referred to simply as K41 theory, the Kolmogorov-Obukhov theory is the framework that has given to the research community some of the most dearly held and most highly regarded results about turbulence. In its simplicity, it was able to discover some laws that now constitutes some of the reference points for turbulence theory. This nonetheless, it has had and, up to a point, still has its detractors, and it can be said that consensus over its validity has not quite been reached. As fascinating as it is, we must keep in mind that turbulence flow structure, which is really what most of the results are about, is not quite the focus of this work. We

will therefore try to abstain from listing all points of controversy, and to stick to the results that are useful for our discussion.

Because of its importance, this piece of theory is treated more or less deeply in most books about turbulence. Aside for the two original papers by Kolmogorov [4] [5] and one by Obukhov [6], our discussion will loosely draw from the second volume of the hallmark book by Monin and Yaglom [7], two books by David McComb [8] [3] with more of a theoretical focus, and the well-known Landau and Lifshitz volume on fluid mechanics[9].

1.2.1 Choice of parameters and similarity hypotheses

As stated above, the context in which K41 is developed is homogeneous isotropic turbulence. Since it was still a relatively new concept, Kolmogorov originally formulated it in terms of local homogeneity and isotropy, as well as implicitly introducing the additional constraint of stationarity. The assumption then is that we are looking at scales (Kolmogorov refers to the “observation scale” for the phenomenon) much smaller than the scale of the mean motion, and we’re doing so for times much smaller than the characteristic mean motion time L/U . We can then assume that the scales we are looking at have “forgotten” the effect of the mean motion and satisfy our hypotheses. While these considerations might be important for certain types of experimental work, from the theoretical point of view it has become much more common to employ homogeneous isotropic (stationary) turbulence purely as a mathematical mean to study the fundamental properties of the turbulent field. Thus we can simply imagine an infinite field of stationary turbulence characterised by the same statistical properties as those of real turbulence far from the length scales of the domain walls, in a way forgetting about the boundary conditions. The two approaches are largely equivalent for the present discussion.

In the described landscape we would like to isolate those parameters that the probability distribution of the velocity field depends on. Specifically we now focus on velocity differences:

$$\vec{v}(\vec{r}, \tau) = \vec{u}(\vec{x}_0 + \vec{r}, t_0 + \tau) - \vec{u}(x_0, t_0) \quad (1.2)$$

with respect to some reference point and time x_0, t_0 (the previous paragraph amounts to saying $r \ll L, \tau \ll L/U$). Because we are discarding boundary conditions, the only way in which the larger scales can influence the smaller ones is by means of the energy dissipation rate $\bar{\epsilon}$. This is the energy that flows from the top down in the “eddy cascade” from the previous section, and it flows from the energetic scales all the way down to the smallest possible ones, where the dissipation into heat actually happens. In addition to this, we should consider the properties

of the specific fluid that is set into motion, which in the case of incompressible fluids amount to the density ρ and the kinematic viscosity ν . However, since the velocities do not depend on the chosen units of mass, the distributions for $\vec{v}(\vec{r}, \tau)$ cannot depend on ρ . This leads to the formulation of the **first Kolmogorov similarity hypothesis**:

For homogeneous, isotropic turbulence at sufficiently high Reynolds numbers, the probability distributions for the relative velocities $\vec{v}(\vec{r}, \tau)$ are uniquely defined by the values of $\bar{\varepsilon}$ and ν .

This simple but sophisticated assumption allows us, just on dimensional grounds, to determine many of the quantities that may interest us when dealing with turbulent flow. First and foremost, we can give an expression for the length scale of the smallest eddies, at which viscosity stabilises the flow and takes over to dissipate the energy. The unique combination of $\bar{\varepsilon}$ and ν that has the dimensions of length is

$$\eta = \left(\frac{\nu^3}{\bar{\varepsilon}} \right)^{\frac{1}{4}} \quad (1.3)$$

and this is related to a velocity scale and a time scale:

$$v_\eta = (\nu \bar{\varepsilon})^{\frac{1}{4}} \quad \tau_\eta = \left(\frac{\nu}{\bar{\varepsilon}} \right)^{\frac{1}{2}} \quad (1.4)$$

We can use these quantities to build a dimensionless distribution from the relative velocities:

$$\vec{w}(\vec{\xi}, s) = \frac{\vec{v}(\vec{\xi}, s \tau_\eta)}{v_\eta} \quad (1.5)$$

and we expect these to be universal. The first hypothesis can then be reformulated in a second version:

For the case of homogeneous, isotropic turbulence at high enough Reynolds numbers, the probability distributions for the adimensional random field $\vec{w}(\vec{\xi}, s)$ are the same for all flows.

It is easy to show the dependence of the Kolmogorov length scale on the Reynolds number of the flow:

$$\eta \sim L \cdot \text{Re}^{-\frac{3}{4}} \quad (1.6)$$

Intuitively enough, the larger the Reynolds number, meaning the larger the dominance of the inertial forces over viscosity, the smaller sized eddies we will find

in the flow. This is the basis a second hypothesis: let us consider a situation in which the flow possesses a large enough inertial range, meaning a sufficiently wide range of scales in between the energetic, mean flow scales and η . On a strictly theoretical level, this amounts to adding a $\eta \rightarrow 0$ limit to the already considered $L \rightarrow \infty$ limit. For any scale ℓ in this range, we expect the corresponding Reynolds number $\text{Re}_\ell = \ell v_\ell / \nu$ to be very large, which is to say that the viscosity should be negligible. Indeed, once again on theoretical terms, an $\eta \rightarrow 0$ limit can be obtained from a $\text{Re} \rightarrow \infty$ limit, according to the behavior shown above. We then arrive to the **second Kolmogorov similarity hypothesis**:

For the case of homogeneous, isotropic turbulence, the probability distributions for the relative velocities $\vec{v}(\vec{r}, \tau)$ in the inertial subrange of scales are uniquely defined by the value of $\bar{\epsilon}$.

With these two statements in mind, we can now move on to deriving equations for other important quantities.

1.2.2 Structure functions

Structure functions are a particular kind of two-point correlation function for the flow field. Specifically, they are correlation functions for the (powers of the) velocity differences between space points:

$$S_n(r) = \langle [u(\vec{x} + r\hat{x}) - u(\vec{x})]^n \rangle \quad (1.7)$$

where \hat{x} is the modulus 1 vector in the direction of the displacement between the two points for which S_n is being computed and the velocities u at the two points are taken along this direction. Indeed this is actually the definition of the *longitudinal* structure functions, and we could also define an analogous quantity in the transverse direction, but their discussion is a bit outside our scope. Because we are restricting our attention to homogeneous isotropic turbulence, the above equation reduces to

$$S_n(r) = \langle [u(r) - u(0)]^n \rangle \quad (1.8)$$

Using the second similarity hypothesis, in the inertial range we expect the structure functions to depend only on $\bar{\epsilon}$ in addition to r . Dimensional analysis then brings us straightforwardly to the second great result of K41, concerning the form of the second order structure functions:

$$S_2(r) = C_2 \bar{\epsilon}^{2/3} r^{2/3} \quad (1.9)$$

This is usually called the “**two-thirds law**”. C_2 was thought by Kolmogorov to be a universal constant, and since then there has been rather solid experimental

verification of this [8]. It is also to be noted that, due to the hypotheses we are working under, “constant” can simply mean “independent of $\bar{\varepsilon}$, ν and r ”.

A second result can be obtained for structure functions, linking those of second and third order. The starting point is the Kármán-Howart equation, which for structure functions takes the form:

$$-\frac{2}{3}\bar{\varepsilon} - \frac{1}{2}\frac{\partial S_2}{\partial t} = \frac{1}{6r^4}\frac{\partial}{\partial r}(r^4 S_3) - \frac{\nu}{r^4}\frac{\partial}{\partial r}\left(r^4\frac{\partial S_2}{\partial r}\right) \quad (1.10)$$

Thanks to stationarity we can (exactly) neglect the time derivative:

$$-\frac{2}{3}\bar{\varepsilon} = \frac{1}{6r^4}\frac{\partial}{\partial r}(r^4 S_3) - \frac{\nu}{r^4}\frac{\partial}{\partial r}\left(r^4\frac{\partial S_2}{\partial r}\right)$$

We can then manipulate the equation further as:

$$-\frac{2}{3}\bar{\varepsilon}r^4 = \frac{1}{6}\frac{\partial}{\partial r}(r^4 S_3) - \nu\frac{\partial}{\partial r}\left(r^4\frac{\partial S_2}{\partial r}\right) \quad (1.11)$$

$$-\int_0^r \frac{2}{3}\bar{\varepsilon}r'^4 dr' = \int_0^r \frac{1}{6}\frac{\partial}{\partial r'}(r'^4 S_3)dr' - \int_0^r \nu\frac{\partial}{\partial r'}\left(r'^4\frac{\partial S_2}{\partial r'}\right)dr' \quad (1.12)$$

$$-\frac{2\bar{\varepsilon}}{3}\frac{r^5}{5} = \frac{1}{6}r^4 S_3(r) - \nu r^4\frac{\partial S_2}{\partial r} \quad (1.13)$$

$$S_3(r) = -\frac{4}{5}\bar{\varepsilon}r + 6\nu\frac{\partial S_2}{\partial r} \quad (1.14)$$

This result is still exact under our hypotheses, and indeed gives the behavior of third order structure functions in terms of the derivative of those of second order. Then, in the inertial range, we can assume the viscosity term to be negligible, and we find

$$S_3(r) \simeq -\frac{4}{5}\bar{\varepsilon}r \quad (1.15)$$

Which is known as the **four-fifths law**. Again, experimental verification so far has strongly leaned in favor of this result.

One point to be made about this result is that, since the KHE is an equation for the conservation of energy, we would expect a term describing the forcing that keeps the turbulence active to appear. Because the equation is local in the scale r though we can assume the forcing to act at the large scales and to be focusing on the small ones. This is also coherent with our infinite domain assumption.

1.2.3 Energy spectrum

What is probably the most recognizable result of the K41 theory is the behavior of the energy distribution in Fourier space. The quantity under study is the density

function $E(k)$ defined by

$$E = \int_0^{\infty} E(k) dk \quad (1.16)$$

By using again the first Kolmogorov hypothesis we can easily use dimensional arguments to find a scaling form for it:

$$E(k) = \nu^{5/4} \bar{\varepsilon}^{1/4} f(k\eta) \quad (1.17)$$

Where f is some (adimensional) universal function. Then applying the second hypothesis, meaning limiting the discussion to the inertial range, we know that the dependence on the viscosity should vanish, implying that the spectrum would take the form:

$$E(k) = C_1 \bar{\varepsilon}^{2/3} k^{-5/3} \quad (1.18)$$

Known as the “**five-thirds law**”. As said, this is by far the most highly regarded result of K41, as it has been experimentally and numerically verified with good agreement in a number of different contexts. This is not to say that it does not have its good amount of criticism, much like the whole of K41, but the discussion of such criticism would be too much of a digression. Suffice it to say that most of the research community holds this result quite dearly and, should one find very large deviations from it in the course of their work, some eyebrows would definitely raise, whether for good or for bad.

Because it has been looked at abundantly, the general form of the energy spectrum across wavenumber ranges has itself become a very recognizable feature of turbulent flow. Qualitatively speaking, it goes as shown in Fig.1.1. A first region is seen at the low k (meaning large l) scales, where usually the forcing takes place. This is the energy range from earlier. Then we enter the inertial range and observe the five-thirds law. Finally, at the very large k 's where viscosity comes back into the picture and we enter the dissipation range, the energy dips further until the smallest eddies are reached.

1.2.4 Turbulent diffusion coefficient

This last paragraph serves as a convenient segue towards the next section. At the very end of his 1941 paper [6], Obukhov shortly takes notice of the fact that the dimensional theory developed so far by Kolmogorov and him can be applied to the issue of the turbulent dispersion of particles. An effective turbulent dispersion coefficient F would have dimensions $[F] = [L]^2[T]^{-1}$. In the inertial range we expect this coefficient to only be dependent on the scale over which the polluting particles




Figure 1.1: Energy spectrum of a fluid in turbulent motion in logarithmic space.

are distributed, ℓ , and, as should by now be clear, the energy dissipation rate $\bar{\varepsilon}$. The only possible combination of these quantities that yields proper dimensions is

$$F(\ell) \sim \bar{\varepsilon}^{1/3} \ell^{4/3} \quad (1.19)$$

In the upcoming section we will study this in much more detail and understand what prompted Obukhov to add this remark at all, starting with the pioneering and long wrongfully neglected work of Richardson.

1.3 Turbulent relative dispersion

With the previous result we can start to venture into one of the main topics of this thesis work. The core question underlying the field of turbulent dispersion is: how are particles advected by the turbulent flow? Given, as we have explained, that the nature of turbulence is stochastic and that we study and measure it as such, what can we say about the stochastic process that is the resulting motion of particles moving with the fluid? If we then focus on the reciprocal motion of particles, rather than their absolute displacement with respect to some fixed frame of reference, we can say that we are dealing with “relative” dispersion. This means that what interests us is fundamentally the distance in between the particles at different points in time. It should be mentioned that the earliest work in turbulent dispersion was carried out by G.I. Taylor, in his theory that is often studied in

basic Fluid Mechanics courses, but his perspective was somewhat different. As important as it was and is, we will bypass its discussion as it does not strictly relate to how the field of relative dispersion evolved and is treated today.

At this point I believe it necessary to be specific about the distinction in the two approaches usually seen in opposition in the context of fluid dynamics. The most commonly adopted in most fields is the Eulerian one, which is also the one we implicitly adopted so far: the fluid is treated as a field by specifying the velocity at each point in the domain. The second one, which is actually conceptually more immediate, is the Lagrangian one: we look at the positions and velocities of specific particles in the fluid by describing their trajectories. The Eulerian approach came second chronologically, but it allowed for a much more comprehensive and compact treatment of fluid dynamics. Still, for some areas, the Lagrangian view (they are also sometimes called frames of reference, although they don't quite match the proper definition) is very useful, if not necessary. Clearly, the current issue of particle displacements is one of them. Importantly, the definition of “particle” here is quite loose: we might be tracking a single molecule of fluid (more in theory than in practice), or a small volume of fluid that we choose to regard as a small indivisible unit for the purpose of our analysis (sometimes also called a parcel), or still we might be looking at particles that are suspended in the fluid but do not actually belong to it and are being transported by its motion. As far as experiments go, this last one is by far the most likely scenario.

Let us then start to unravel the issue at hand by introducing the field as it was introduced historically by Lewis Richardson in his 1926 paper [10]. Although it was in many ways preliminary to the more rigorous work that came afterwards, it contains some bright intuitions about the nature of the phenomenon in addition to introducing the formalism apt to study the problem. The discussion in this section is mostly based on the two famous reviews by Sawford [11] and Salazar and Collins [12], as well as the material contained in the already mentioned second volume of the book by Monin and Yaglom [7]. More specific papers will be cited as the need arise.

1.3.1 Richardson formalism: distance-neighbor-graph

The original question Richardson was investigating when he started formalising this issue was the spreading of clouds (either of water droplets or contaminants) about their center of mass in the atmosphere. The commonplace empirical observation is that a puff starting out in a certain contained volume and let loose in the atmosphere will quickly expand, as the particles that make it up are spread apart by the flow. This happens at a much faster rate than would be justified by molecular diffusion. The first intuition Richardson had was of attributing this to the turbulent eddies that make up the flow of atmospheric air: as particles spread apart, even

just by molecular diffusion if they start out extremely close, they enter different eddies and are carried away by them. According to the concept of eddie hierarchy we have introduced, their separation will increase at a faster and faster rate as they grow apart because it will be influenced by eddies of larger and larger scale. We then need a framework to study this phenomenon to understand how fast they are spread by this mechanism.

The quantity Richardson first defined is intended to be “a statistic for clusters”. Let us assume that we marked some number of particles N that we keep track of, and let us label them by capital letters. Now let us also take some arbitrary length measure h . The choice of this parameter is effectively of some influence in the analysis, but for now let us assume it is chosen adequately. We define $A_{n,n+1}$ as the number of particles whose distance ℓ from A fall in the range $nh \leq \ell \leq (n+1)h$. We then define $Q_{n,n+1}$ as the mean of this quantity over all marked particles:

$$Q_{n,n+1} = \frac{1}{N}(A_{n,n+1} + B_{n,n+1} + C_{n,n+1} + \dots) \quad (1.20)$$

Q , meant as the ordered set of $Q_{n,n+1}$ quantities (one can think of an array in computational terms), is named **Distance-Neighbor Graph**, and lies at the core of Richardson’s analysis. A plot of an example of this quantity computed from one of the simulations described in the next chapter is shown in Fig. 1.2. Richardson originally explained the derivation with particles placed on an oriented line, which is why his plots are for both positive and negative ℓ values, but in the 3D treatment ℓ just has the meaning of a distance and is always positive. The plots are mirrored for visual convenience (it is also true that even in the 1D oriented line case the plots are always symmetric).

Richardson then proceeds to take the continuum limit of $Q(\ell, t)$. He gives a definition in terms of scalar field of concentration, but to keep the discussion lean we will just call it for what it is: a probability density function of the separation distances between the particles in the marked cluster. We will follow along his work and call this continuous version $q(\ell, t)$ in lower case. The next step is to look for an equation that may describe the evolution of this quantity. Coherently with the topic at hand, and by noting that the space integral of q is a conserved quantity, we can expect this equation to be a non-Fickian diffusion equation like

$$\frac{\partial q}{\partial t} = \frac{\partial}{\partial \ell} \left(F(\ell) \frac{\partial q}{\partial \ell} \right) \quad (1.21)$$

The key feature of this equation is the dispersion coefficient $F(\ell)$, which carries a dependency on the separation distance. Specifically, according to the deductions above, it should be an increasing function of ℓ . This is where Richardson’s genius really struck: he gathered what data there was available at the time on atmospheric




Figure 1.2: Distance-neighbour graph and its evolution in time (not normalised, hence the high values). Values obtained in the context of the simulation of a cloud top border, see Chapter 2.

dispersion at a wide range of scales (as the reader can imagine, it was not much) and tried to interpolate it to get a power-law scaling of $F(\ell)$. What he got was

$$F \sim \ell^{4/3} \quad (1.22)$$

We should immediately recognize the same length scaling as derived thanks to similarity theory in the previous section. Obviously the chronological order of these results is reversed: Richardson achieved it based on observations in 1926, some fifteen years before the development of the Kolmogorov-Obukhov theory, and therefore did not yet have the tools to interpret it on a theoretical basis. In his paper cited above at the end of the section, Obukhov was addressing Richardson’s empirical law and confirming it on the grounds of the recent theory. Considered how Richardson achieved his result, this was really quite remarkable.

We should now have all the tools in place to move on to the modern field of relative dispersion. Although the core questions remained the same (as, in fact, they have not fully been answered almost 100 years later), a lot was added to it and effort was put into trying to test both the original results and the outcomes of these additions.

1.3.2 Modern approach: scaling theory

It has become common to refer to relative dispersion as “two-particle dispersion”. This highlights the fact that we are focusing on statistics of the separation in between particle pairs. Indeed, the consideration of statistics based on more than two particles introduces a whole new array of problems [13] which will not be dealt with in our cases.

With the above diffusion-like equation in mind, we can recover Obukhov’s result for the dispersion coefficient

$$F(\ell) = k_0 \bar{\varepsilon}^{1/3} \ell^{4/3} \quad (1.23)$$

where we have now also included the dimensionless proportionality constant k_0 , and use it to get a solution for a point source. This comes in the form

$$q(\ell, t) = A \left(\frac{1}{\pi \langle \ell^2 \rangle} \right)^{3/2} \exp \left[- \left(\frac{B \ell^2}{\langle \ell^2 \rangle} \right)^{1/3} \right] \quad (1.24)$$

where A and B are just numerical factors. This solution implies a behavior for the mean square separation like

$$\langle \ell^2 \rangle = g \bar{\varepsilon} t^3 \quad (1.25)$$

where g is called the Richardson constant. This last result is among the most relevant ones for the investigation of relative dispersion, and is called Richardson-Obukhov law. It is obviously related to the law for $F(\ell)$ by the fact that $F(\ell) = d\langle \ell^2 \rangle / dt$, so looking at the two quantities is broadly speaking equivalent.

This is approximately where Richardson’s original theory runs its course, and we can start adding to it on one side and understanding its limitations on the other. After the contributions by Obukhov it was Batchelor [14] who famously developed the field through a more detailed application of K41, and he did so in two ways. Specifically, he noted that at short times the dispersion should depend on the initial separation between the particles through the flow structure functions. This leads to the combined result

$$\langle (\ell - \ell_0)^2 \rangle \simeq \begin{cases} \frac{11}{3} C_2 (\bar{\varepsilon} \ell_0)^{2/3} t^2 & \text{for } t \ll t_B \\ g \bar{\varepsilon} t^3 & \text{for } t_B \ll t \ll T_L \end{cases} \quad (1.26)$$

where T_L is the Lagrangian integral time scale (see next section) and t_B is a newly introduced time scale before which the initial separation plays a role, named the Batchelor time. By dimensional arguments it can be seen that $t_B = \ell_0^{2/3} \bar{\varepsilon}^{-1/3}$.

Batchelor’s second contribution [15] was analysing relative dispersion in the dissipation subrange. We have so far put a lot of stress on the fact that we wanted to focus on the supposedly self-similar inertial range, which is supposed to give universal behaviors, but we will include the results for the dissipation subrange for reasons that will be clear in following chapters. Simply put, the mean square separation’s evolution is described by

$$\frac{d\langle \ell^2 \rangle}{dt} = \langle \vec{\ell}(t) \cdot \vec{w}(t) \rangle \quad (1.27)$$

where $\vec{w}(t)$ is the relative velocity between two particles. If the separation between them is in the dissipation subrange scales we can expand the velocity in Taylor series and stop at the first order (on account of them being close, in terms of turbulent fluctuations), getting $\vec{w}(t) = \vec{\ell}(t) \cdot \nabla \vec{u}$. As a consequence we arrive, for the dissipation range, at the very different law

$$\langle \ell^2 \rangle = \ell_0^2 \exp(\xi t) \quad (1.28)$$

This completes the classical theoretical ground for turbulent relative dispersion. While it might appear simple on the surface, it is actually rich in nuance and more or less hidden problems, which is why to this day it remains a very active field of research. As a last subsection we will discuss some of these more faceted aspects.

1.3.3 Refinements, criticism and recognized results

The first and most obvious specification to be made is that, while a lot of literature on relative dispersion focuses on the laws that describe the behavior of either the mean square separation or the dispersion coefficient, the only quantity that fully describes the dispersion process is $q(\ell, t)$. Indeed, $\langle \ell^2 \rangle$ is just a moment of the distribution, and several distributions can result in the same behavior of $\langle \ell^2 \rangle$. In general, if we take the diffusion equation for $q(\ell, t)$ as a given, the solution will be determined by the expression for K that we use. We can extend K to also be a function of time by writing $K(\ell, t) = k_0 \bar{\varepsilon}^a t^b \ell^c$ and note that any combination of a , b and c satisfying the system

$$\begin{cases} 2b + 3c = 4 \\ a + \frac{c}{2} = 1 \end{cases} \quad (1.29)$$

will satisfy the requirements of K41 and also implies $\langle \ell^2 \rangle = g \bar{\varepsilon} t^3$. Conversely, the distributions solving the resulting diffusion equation will take on rather different forms.

In second, we have to address the assumption we just made: that the process is governed by a diffusion equation. Indeed, implicit in the derivation of diffusion equations is the fact that the underlying mechanism that induces the random motion, in this case the turbulent flow, should be characterised by short (ideally point-like, think of the derivation of the diffusion equation in terms of random walks) time correlations [16]. This is not the case for fluid turbulence, and as such it may very well be that a diffusion equation is not best suited to describe the process. Specifically, it might be necessary to take into account non-local effects. Various avenues have been tried over the last decades, almost all of them sharing the $\langle \ell^2 \rangle \sim t^3$ result but leading to different forms of $q(\ell, t)$. Some of the examples provided in the sources already cited are somewhat dated, so a more recent treatment of non-locality can be found for example in the work of Malik and Hussain [17].

So how did the theory stand up to tests so far? As said, experimental verification of relative dispersion is rather complicated. First of all, turbulence in the outside world will never be isotropic or homogeneous. We are bound to have to deal with the effects of thermal stratification, energetic inhomogeneities, geometric anisotropies and all the like. On the flip side, laboratory experiments are always limited to the relatively low Reynolds numbers achievable. Still, all this has not stopped many groups (including us, since experimental investigation is a large portion of the work described in the later chapters) from trying to tackle the task. Similar considerations can be made for numerical studies: while it is easy to generate exactly homogeneous isotropic turbulence in this case, the Reynolds number achievable are generally still lower than in laboratory experiments, and it should be remembered that one of the hypotheses underlying K41 is basically of $\text{Re} \rightarrow \infty$. The challenge then is resolving the inertial range of scales (see the appropriate section in Chapter 4). All this results in the fact that, although many groups are working on the topic, consensus has been reached over very few aspects, if any.

Among the ones that at this point in time seem to be more or less proven is the value of g , the Richardson constant. Although obviously there is no unique exact value found in laboratory experiments, many of them have obtained values very close to 0.5 (Ott and Mann [18], Berg *et al.* [19]). Similar values have been found by early direct numerical simulations (Boffetta and Sokolov [16], Biferale *et al.* [20], Sawford *et al.* [21]), with some studies presenting slightly larger values.

As far as in-field experiments are concerned, probably the most notable one was the EOLE experiment carried out in 1971 reaching into early 1972. 480 large balloons were tracked for a long period of time and this led to an identification first of an exponential regime and then to a R-O one. It should be noted, for reasons that become clear in the following chapter, that these large balloons are hardly ideal tools for investigation of turbulent flows, owing to their large inertia.

They can still be considered valuable for extremely large-scale experiments such as this one.

This concludes our short introduction to the topic of turbulent relative dispersion. The key takeaway is this: while theory provides us with some important results, care should be taken when applying it to the real world and its far-from-ideal conditions. We will deal with this aspect again later, when I present our experimental results. For now we can move on to the final theoretical topic, to lay the groundwork for the second set of results.

1.4 Lagrangian correlations

As a last section in this chapter we briefly introduce the second quantity that we will investigate in the later chapters, namely Lagrangian correlations. Their treatment in the literature is once again quite differentiated, possibly even more so than the preceding topics. I will base my discussion mostly on two papers, by James M. Wallace [22] and He *et al.* [23], which are the only reviews I managed to find that offer a relatively complete overview of the various aspects of this field, whereas most articles are rather specialised, often adopting slightly different definitions or parameters for observations to suit their purposes. The abundance of material is mostly related to the fact that correlations, although we have not discussed them in these terms before, are arguably the most important part of turbulence modeling for numerical applications. Indeed, to account for its statistical nature, virtually all techniques used to deal with turbulence at the practical level (in fields such as engineering) rely on an exact solution of mean flows in parallel to some approximations for the statistical terms (namely so-called Reynolds-Averaged Navier-Stokes models) or on an exact solution of large vortices with some approximations accounting for the smaller ones at a probabilistic level (called LES, for Large Eddy Simulations). These approximations are called “closures”, and being statistical models they depend on some parameters, which generally speaking include velocity correlations. In the field of “particle-laden” turbulence these correlations need to be Lagrangian in nature, to properly mirror the motion of the particles.

Historically speaking, Lagrangian correlations started to be looked at much more closely thanks to the development by Robert Kraichnan of his direct interaction approximation (DIA). This is a theoretical approach that relies on some rather strong assumptions to attempt to derive a closure for Navier-Stokes equations describing fully developed, homogeneous, isotropic turbulence. This theory had the macroscopic issue of yielding an energy spectrum power law of type $E(k) \sim k^{-3/2}$. At the time of Kraichnan’s first paper, i.e. 1959, Kolmogorov’s 5/3 prediction had not yet received much confirmation, so the DIA theory hung on for a while, but

Kraichnan was forced to revise it when experimental results started popping up that displayed the K41 spectrum. He then modified DIA by focusing on Lagrangian velocities, which he defined in a sort of hybrid fashion as $v(\vec{x}, s; t)$: this means that v is the velocity at time t of a particle that was in position \vec{x} at time s . This updated version was dubbed Lagrangian History Direct Interaction Approximation (LHDIA), and it managed to display a 5/3 spectrum. In spite of this, many issues remained unaddressed, and the very assumptions at the core of the theory are still a matter of debate today. Some further improvements were developed, but the results involve extremely complicated equations that need serious simplification to be tackled. Because of all this, approaching the subject from the theoretical standpoint remains a challenge that is oftentimes excessive with respect to the results.

The experimental approach has for a long time remained almost as problematic, as Lagrangian quantities are inherently hard to measure. Only recently has the situation improved thanks to the development of specialised tools such as 3D particle tracking. One notable study that takes advantage of this technology is the one by Oliveira *et al.* [24]: in a series of papers describing their experiment they investigated, among other aspects, Lagrangian correlations, at the considerable Reynolds number of 10300. They found good overlap with numerical simulations. These obviously still remain the preferential tool for this kind of analysis, as they most promptly provide access to Lagrangian-type datasets. It is important to stress, though, that there does not appear to be any kind of consensus in the matter of Lagrangian correlations, meaning some important result having at least enough of foundation and recognition to serve as a reference for new studies (such as may be the Richardson-Obukhov law in the case of dispersion). The functional form itself is not certain: many papers report exponential evolution in time, but many do not.

Because of all of the above, we will try to be as practical as possible and focus on the very basics. The quantity at the core of our analysis will be a normalised version of the correlation function, which is the Pearson correlation coefficient:

$$R_L(\vec{x}, \tau) = \frac{\langle V_i(\vec{x}, t_0) V_j(\vec{x} + \vec{r}(t_0 + \tau), t_0 + \tau) \rangle}{\sqrt{\langle V_i^2(\vec{x}, t_0) \rangle} \sqrt{\langle V_j^2(\vec{x} + \vec{r}(t_0 + \tau), t_0 + \tau) \rangle}} \quad (1.30)$$

where V_α are velocity components, and $\vec{r}(t_0 + \tau)$ is the displacement of the particle after a time interval τ . As usual, in the case of stationary, homogeneous, isotropic turbulence the correlation coefficient is independent of the initial position and the initial time, so that effectively the only controlling parameter is τ . We kept all dependencies in the formula since in our cases, instead, they will be relevant. We also stress, at odds with some definitions found in the literature (including [22]), that the velocities $V(\vec{x}, t)$ are the Lagrangian velocities of the particles forming our ensemble, and not Eulerian velocities at the particle's location. This is due to

the fact that, while the interest of many authors lies in the particles that make up the fluid, and as such Eulerian and Lagrangian velocities coincide at a certain point in space, we will deal with massive particles suspended in the fluid, with their own inertia, that in general might have a different velocity than the surrounding fluid. Moreover, the definition given here only refers to velocity components, as indeed these are the correlation functions that most of the theory focuses on, but in principle the correlation coefficient can be defined for any quantity that a Lagrangian particle moving with the flow might be characterised by. In the related chapter we will give a more general definition that reflects this fact.

Once we have the correlation coefficient we can define another quantity which is usually employed to actually parameterise the correlations, the Lagrangian integral time scale:


$$T_L = \int_0^\infty R_L(\vec{x}, \tau) d\tau \quad (1.31)$$

This is a measure of how quickly correlations decay in time, and it is again important to underline that, while often this is considered a global quantity of the flow, in our case it keeps a space dependence. Its relationship with the Eulerian integral time scale is another common topic of research. The general consensus so far seems to be that the Lagrangian scale is larger than the Eulerian one, meaning that correlations decay more slowly while following a particle, but the exact relationship is far from a solved problem. Again, in Chapter 5 we will discuss this quantity further, to attempt to make its computation viable in our simulations.

This brief last section concludes the first part of the theoretical introduction necessary for the discussion of the results achieved in the various aspects of turbulence described above. We can now move to a somewhat more specific description of the context in which these results were obtained.

Chapter 2

Simulation and phenomenology of a cloud border region



This chapter deals with the various results that have been obtained by the group of Prof. Tordella in the context of the direct numerical simulation of clouds at the small scale. These are the subject of a few recent (as well as less recent, as far as the DNS code is concerned) papers, all published in the years shortly preceding the start of my thesis work.

The reason why the group chose to focus on the region of the cloud bordering the clear air around it is that, despite the fact that the phenomena taking place in the bulk of clouds are themselves far from trivial, this specific region presents some interesting peculiarities. Specifically, it is a known fact that the turbulence inside clouds is more intense than that around them. The exact reason is still a matter of speculation: the most immediate phenomenon that surely contributes to this is the release of latent heat that occurs during the condensation of water vapor, but a variety of others are still under investigation. Whatever the reasons may be, the implication is that the border between a cloud and the clear air is to be modelled as a turbulent mixing layer. A lengthy discussion about mixing layers would require more than can be afforded here in terms of both time and space, so for the moment we will just say that they are interesting to different areas of fluid dynamics for different reasons, but to us there is one that stands out from the others: due to the nature of the mixing layer, the turbulence inside it is by definition inhomogeneous and anisotropic. These qualities depart from the “idealised” description that is often used for turbulence, as explained in Chapter 1, and allow us to study the effect of intermittency both on the flow itself and on the various physical properties and phenomena that might occur inside it.

The chapter will have the following structure: in the first section I will attempt to give a very summarised account of cloud physics, which is of course as broad a subject as one can imagine, but I will make an effort to limit the discussion to those part of the theory that should provide the reader with the tools necessary to appreciate the modelling choices and the results that will follow; in the second section I will describe the simulations that have been carried out of the cloud border region, which are the ones that were used both for the results presented in the past papers and for the ones presented in the current work; finally, in the third section, I will summarise the past results, so as to show the interesting phenomena that have been analysed so far and hopefully justify our interest in pursuing the matter further.

2.1 Cloud physics in a nutshell

As said, many books can be and have been written on the topic of cloud physics. This is not only due to the variety of perspectives that can be used to approach many broad subjects, but also to the intrinsic complexity that all atmospheric phenomena possess, clouds and precipitation being perhaps the most prominent example of this. Here I will try to shrink the whole subject to only those part that will allow for a good understanding of the relevance of our choices and results, as well as to the parts that are generally most widely accepted as solid and verified (which cannot be said for all of them). It is inevitable that the matter will suffer

a considerable simplification, but I trust that my selection will provide a good enough picture to carry on smoothly with the discussion.

Most of the topics I will cover are basic enough that their presentation has been sort of “standardised” over the years, so they can be found in very similar fashions in various sources. I used as reference mostly three books: the one by Pruppacher and Klett [25], which has become a bit of a standard reference for cloud physics; the more concise book by Rogers and Yau [26]; the more recent one by Wang [27].

At the most basic level, clouds are suspensions of water particles in air. These particles can be in the form of liquid water or of ice, depending partially on the temperature at which the clouds form but also on some other parameters. While many of their features are similar the discussion is generally separated for the two types of clouds, and we will stick to the type containing only liquid water droplets. These are generally referred to as “warm clouds”. They form in the atmosphere thanks to some portions of warm, moist air raising to altitudes at which the local temperature, being lower than the one the air started at, brings the water vapor inside the air beyond the point of saturation and thus induces condensation. We will thus begin by looking at the model describing this process, often called “parcel theory”.

2.1.1 Parcel theory for atmospheric supersaturation

Let us start by considering a parcel of air containing no moisture at all, which we will call “dry”. This dry air parcel, at a certain point, starts to rise up. It may be purely due to buoyancy forces related to its density (or equivalently temperature) or to some slow updraft, but once the movement is started we assume buoyancy to be the only remaining force on the parcel. We will assume two other things: that the ascent is faster than the typical timescale of heat exchange between the parcel and the ambient air surrounding it, and that no ambient air will mix with the parcel. Thus our parcel is, in first approximation, a closed system, and what we are describing will have the characteristics of an adiabatic process. What will happen to it while it rises is that it will encounter lower and lower ambient pressures, and so it will expand. Intuitively we know that this will cause the temperature in the parcel to fall. At the same time, along with the pressure, the temperature also decreases while moving to higher altitudes in the atmosphere, and so the buoyancy of the air parcel will change as it goes up. The question then becomes: will the air parcel’s temperature lower fast enough to, at some point, stop its ascent by encountering levels of the atmosphere warmer than itself? Or will it lower more slowly than the ambient temperature, thus causing the parcel to rise indefinitely?

Let us start by using the first law of thermodynamics. Since we are discussing an adiabatic process, it reads

$$dQ = 0 = C_v dT + PdV \quad (2.1)$$

Where dQ would be the heat exchange, which in our case is null, C_v is the heat capacity of air at constant volume, and T , P and V are temperature, pressure and volume. This shows quantitatively that because dV is positive, dT must be negative. Since we are not interested in the extensivity of the quantities at play, let us divide everything by the parcel mass m . This will turn the equation into the similar form

$$dQ = 0 = c_v dT + Pdv \quad (2.2)$$

Where we have defined the specific heat $c_v = C_v/m$ and the specific volume $v = V/m$. Next we should keep in mind that, for the purpose of this simple model, the air in the atmosphere can be considered to be in hydrostatic equilibrium, meaning that its pressure will satisfy the equation

$$dP = -\rho_a g dz \quad (2.3)$$

ρ_a being the air density, g the gravitational acceleration and dz the displacement along the vertical direction. How can we relate the two laws above? By supposing (with good reason) that the air behave as an ideal gas, we can use the ideal gas law in the form:

$$Pv = R_a T \quad (2.4)$$

Where R_a is the dry air gas constant. By differentiating this expression we get

$$Pdv + vdP = R_a dT \quad (2.5)$$

and plugging this into 2.2

$$c_v dT + R_a dT - vdP = 0 \quad (2.6)$$

$$c_p dT - vdP = 0 \quad (2.7)$$

Where c_p is now the specific heat at constant pressure. Substituting the pressure differential from the hydrostatic equilibrium equation we finally get a form for the temperature variation with height of the air parcel

$$\Gamma_d = -\frac{dT}{dz} = \frac{g}{c_p} \quad (2.8)$$

which we call "dry adiabatic lapse rate". It is important to keep in mind that this is the temperature variation with altitude *of the air parcel*, which we are going to compare with a similar quantity defined for the atmosphere around it:

$$\gamma = -\frac{\partial T}{\partial z} \quad (2.9)$$

As said above, the parcel will keep rising as long as it is lighter than the air around it, and because the pressure of the parcel is always assumed to be the same as the environmental pressure this amounts to saying that it will rise as long as it is hotter than the air around it. This is what is meant when talking about “atmospheric stability”: the atmosphere is said to be stable if the ascent of a warm air parcel will be damped and stopped by the environmental temperature gradient, and unstable in the opposite case. In terms of the lapse rate we just derived, the stability conditions read:

$$\gamma < \Gamma_d \quad \text{STABLE} \quad (2.10)$$

$$\gamma = \Gamma_d \quad \text{NEUTRAL} \quad (2.11)$$

$$\gamma > \Gamma_d \quad \text{UNSTABLE} \quad (2.12)$$

At this point we have to turn back to the beginning and ask: why bother with a parcel of dry air if what we are trying to describe is the formation of clouds, which necessarily requires water vapor? As it turns out, had we considered a moist parcel of air, meaning one containing some amount of water vapor, the discussion would have been rather similar if somewhat more convoluted, to account for the necessary corrections. The main one that has to be considered is the effect of latent heat: once the water vapor in the parcel reaches (and goes beyond) saturation, it will start to condensate, and this causes the release of heat. Due to this effect the temperature drop in the moist parcel is actually slower than in the dry parcel. Still, at the end we obtain a quantity analogous to Γ_d , called “pseudoadiabatic lapse rate” with symbol Γ_s , that incorporates all the necessary corrections and can be substituted in the stability condition. Importantly, according to the explanation above, we always have $\Gamma_s < \Gamma_d$.

We can now turn to the analyses of the process just mentioned: the condensation of water vapor, once saturation is reached (or rather *supersaturation*), to form water droplets. Once again, we will focus on water droplets despite the fact that water vapor sometimes aggregates directly to ice particles, in which case we would talk about deposition rather than condensation.

2.1.2 Water vapor at equilibrium with droplets

The process of condensation starts on a few aggregates of water molecules known as nuclei, which gives to their formation the name of *nucleation*, and have to start this subsection by drawing an important distinction. Water vapor can

condensate in two different ways: *homogeneous nucleation* happens by simple collision of water molecules which may find it energetically convenient to stick together, and progressively form groups that will serve as the condensation nuclei; *heterogeneous nucleation*, instead, happens in the presence of small particles of other substances, known as *aerosols*, onto which water molecules latch and that play the role of kickstarting the droplets. In heterogeneous nucleation the term "nuclei" is thus usually referred to the aerosol mixed in with the water vapor. Now, while homogeneous nucleation can and does occur, in reality it requires saturations of several hundred percents to give rise to appreciable droplets. What takes place in the atmosphere is, virtually always, heterogeneous nucleation. The types of aerosols that can serve as cloud condensation nuclei (often abbreviated CCN) are quite varied and still a matter of abundant research, but for what concerns us it should be sufficient to make the distinction between soluble and insoluble aerosols: while the first ones greatly facilitate condensation, the second ones can even impair it.

Let us now take a closer look to the equilibrium between liquid water and water vapor at saturation. If we have a simple closed container of water it is described by the well known Clausius-Clapeyron equation:

$$\frac{de_{sat}}{dT} = \frac{L_e}{T(v_V - v_W)} \quad (2.13)$$

where e_{sat} is the saturation vapor pressure of water, L_e is the condensation-evaporation latent heat, v_V and v_W are the molar volumes of the vapor and the liquid water and T is of course the temperature. In the case of cloud droplets there are two aspects that this equation fails to capture. The first one is the effect of curvature at the interface between the two phases, known as the Kelvin effect. In droplets surface tension induces a pressure difference between the liquid and the vapor that satisfies

$$P_W - P_V = \frac{2\sigma}{R} \quad (2.14)$$

where σ is the surface tension and R is the droplet radius. Secondly, as mentioned above, we are interested in the case in which water condensates on soluble aerosol, as that is the most likely (if not almost only) mechanism that leads to cloud formation. This implies that the water will tend to dissolve the aerosol particle, and this impacts the saturation vapor pressure (which is known as the Raoult effect). Naively put, the saturation pressure is lower for a solution than for pure water due to the solute occupying some of the surface of the liquid phase. The influence of the solute is described at the thermodynamic level by the *activity* parameter a_k , which specifically acts on the chemical potential of the k -th component of a system. Both

of these effects can be incorporated in the derivation of the Clausius-Clapeyron equation, which we will skip as it becomes somewhat cumbersome although not especially difficult, and what is obtained is called the *generalised Clausius-Clapeyron equation* which reads:

$$-\frac{(h_{V,0} - h_{W,0})}{T^2}dT + \frac{(v_{V,0} - v_{W,0})}{T}dP_V - \frac{2v_{W,0}}{T}d\left(\frac{\sigma}{R}\right) + R_g d \ln\left(\frac{a_V}{a_W}\right) = 0 \quad (2.15)$$

where the only new quantities are $h_{V,0}$ and $h_{W,0}$ which are molar enthalpies. The 0 subscript refers to the case of pure substances. It is easy to see that the first two terms are the ones in the standard Clausius-Clapeyron equation, and the ones following bring the contribution of the Kelvin and Raoult effects respectively.

Once we have this, let us analyse the equilibrium between the two phases by supposing it was reached without changing to pressure or water vapor, i.e. imposing $dP_V = dT = 0$. This leaves us with the simpler expression

$$-\frac{2v_{W,0}}{T}d\left(\frac{\sigma}{R}\right) + R_g d \ln\left(\frac{a_V}{a_W}\right) = 0 \quad (2.16)$$

Because in general the amount of water vapor mixed with the air will be small with respect to the total gaseous mass we can approximate the activity in terms of partial pressure $a_V \simeq e_{sat,W}/P$. Then upon integration we get

$$-\frac{2v_{W,0}}{T} \int_{\infty}^a d\left(\frac{\sigma}{R}\right) + R_g \int_{e_{sat,\infty}}^{e_{sat,R}} d \ln\left(\frac{e_{sat,0}}{P}\right) - R_g \int_{e_1}^{a_W} d \ln a_W = 0 \quad (2.17)$$

$$-\frac{2v_{W,0}}{T} \left(\frac{\sigma}{R}\right) + R_g \ln\left(\frac{e_{sat,R}}{e_{sat,\infty}}\right) - R_g \ln a_W = 0 \quad (2.18)$$

$$\frac{e_{sat,R}}{e_{sat,\infty}} = a_W \exp\left(\frac{2v_W \sigma}{R_g T R}\right) \quad (2.19)$$

The final expression is known as the Köhler equation, and it is the fundamental result for the equilibrium of water vapor in atmospheric clouds. The fraction at the left-hand side is usually called *saturation ratio* with symbol S_{VW} (or sometimes just S , but we will later use S simply for the supersaturation). The plot of the Köhler equation with respect to the droplet radius is known as the Köhler curve, which generally differs for the amount and type of solute but whose general behavior is given in Fig. 2.1. As can be seen the curve displays a maximum, and is to be interpreted in the following way: up until reaching the argument of the maximum, known as *critical radius* R_C , the radius of the droplets increases with the saturation ratio (a notable feature is that, thanks to the Raoult effect, small droplets are in equilibrium with subsaturated air); once the critical radius is overcome, droplets of

larger radius will be at equilibrium with lower amounts of supersaturation. This means that past the critical radius water vapor will keep flowing to the droplets, depleting the air until the saturation point. For this reason the droplets that reach the critical radius are usually described as *activated*. We will now see how this growth occurs, and then explain why, past a certain point, the mechanism of coalescence tends to take over.




Figure 2.1: Köhler curve (solid black) for droplets formed on hydrosoluble aerosol. The Kelvin curve plotted in dashed black allows us to appreciate the influence of the Raoult effect: in the absence of aerosol, small droplets require very high values of saturation to form and be stable.

2.1.3 Growth of water droplets by diffusion

In the first phase of their growth droplets acquire new water molecules by simple diffusion. All this means is that the vapor molecules, floating around in the air parcel under discussion, sometimes will encounter a droplet and remain attached to it. We will go through a very simple (yet effective) description of this process in an exact way, but we will keep it relatively short as in reality, as we will see later, the numerical simulation of cloud droplets relies on an approximation that allows to compute the size variations more readily.

Let us start by considering a single, static droplet in a field of water vapor, and by calling $n(r)$ the concentration of water vapor molecules at a distance r from the center of the droplet. This quantity must satisfy a diffusion equation

$$\frac{\partial n}{\partial t} = D\nabla^2 n \quad (2.20)$$

D being the diffusion coefficient. If we imagine to be considering an infinitesimal increase (or decrease) in size of the droplet we can approximate the process as being stationary, meaning that we ignore the time variation of the concentration. This easily leads to a solution by imposing boundary conditions $n(r \rightarrow \infty) = n_0$, corresponding to the ambient concentration, and $n(r \rightarrow R) = n_R$ the concentration at the surface of the droplet:

$$n(r) = n_0 - \frac{R}{r}(n_0 - n_R) \quad (2.21)$$

Because the flux of molecules to the droplet surface is $D(\partial n/\partial r)_{r=R}$ we can easily write the rate of mass increase of the droplet as

$$\frac{dm}{dt} = 4\pi R^2 D \left(\frac{\partial n}{\partial r} \right)_{r=R} m_0 \quad (2.22)$$

where m_0 is the mass of a single water molecule. By using the above solution for $n(r)$ we finally get

$$\frac{dm}{dt} = 4\pi R^2 D (n_0 - n_R) m_0 \quad (2.23)$$

Of course, to complete the picture, we would have to couple this equation with one accounting for the influence of condensation (or evaporation) on the temperature through latent heat release (or absorption), which in turn influences the saturation. This doesn't lead us much further in terms of understanding the process though, and as said usually approximations are employed to account for all the various effects. One additional fact worth mentioning is that this short discussion has focused on a single droplet, by in real clouds the droplet population competing for the same limited supply of water vapor introduces some serious and non-negligible limitations in the growth process. In fact, we need to account for the variation of the saturation ratio by means of an equation of type

$$\frac{dS_{VW}}{dt} = P + C \quad (2.24)$$

where P is a term accounting for saturation production due to updraft and C accounts for depletion due to condensation (generally $C < 0$). The derivation of the two terms is rather involved and will be skipped, but we will see later that

the form of this equation will be analysed in the turbulent setting and will yield some interesting results. We can then proceed by supposing that some number of droplets have reached a stable size by diffusional growth and have started (although truthfully they never stopped at all) to move around the newly formed cloud.

2.1.4 Collision and coalescence

Very simply put, coalescence happens when two droplets collide and, as a result, merge. Until at least some drops have reached a considerable radius (of the order of around 10-20 μm), the collisions between drops are rather negligible and the mechanism of coalescence does not impact the droplet *spectrum* (which is the distribution of droplets as a function of size, usually parametrised by radii). After this point though it starts to become prevalent and quickly overcomes diffusional growth. There are a number of ways of treating coalescence, mostly of a statistical nature, for obvious reasons. In the case of our simulations, though, this mechanism is incorporated in the droplet population in a very direct way, without the use of statistical models. For this reason we will not delve on the modelling of coalescence, as interesting as it is, but will just discuss two parameters that are important to consider to understand the choices made in the simulations.

The first one is called *collision efficiency*, and is a product of the fact that water droplets are moving inside a fluid field rather than in empty space. When we think of gas particles, for examples, the collisions between them occur with certainty if their trajectories intersect, since there is nothing in between to prevent it. Clouds droplets, during their motion, change the flow field around them, and because of this may actually not come into contact even if their trajectories seemed to be intersecting before they got close. The governing parameter to determine the probability of collision between two drops is the ratio of the radii, and various efforts have been done to determine this dependency. The field was basically started thanks to the superposition method introduced by Irving Langmuir and used in the calculations of two papers often cited for dataset references, by Lin and Lee[28] and Schlamp et al[29].

The second parameter is the *coalescence efficiency*, which measures the probability of two drops actually merging together after colliding. In general, other than coalescing, they might simply bounce out of the collision or partially merge temporarily and then separate again. For cloud droplets, though, these two events seem to constitute a rather minute fraction of the outcomes of collisions, such that the coalescence efficiency is often approximated to 1. Sometimes the two parameters are defined as a single one called *collection efficiency*, which is simply their product.

One important fact to highlight is that most classic studies on the collection efficiency do not actually account for turbulence effects. All they consider is the

buoyancy effect that accounts for the different fall speeds of droplets of different radii. It is only recently that the influence of turbulence has started to be considered, which is the context in which our research fits.

2.2 Simulation of a cloud-clear air interface

With the previous section I hope to have cleared the way for the more specific topics to come. This one will be devoted to the description of the domain that the numerical experiments of this work and the recent papers mentioned at the start of the chapter have been carried out into.

The region that is simulated is the border between a cloud and the surrounding droplet-free air, generally called clear air. The reasons for the interest in this specific portion of a cloud's body are mostly contained in its turbulent characteristics. Because of the difference in the flow of the two region (more on this shortly), it is a turbulent mixing layer. Like many other topics touched upon so far, mixing layer theory is a whole large branch in and of itself, full of theoretical aspects and practical applications. For our purposes, the interest lies in the fact that it is an instance of anisotropic turbulence. This anisotropy makes it so that cloud droplets experience very different phenomena with respect to what they might experience in the bulk of the cloud: mixing with dry air, intermittency of the flow and relevant scalar fields, stratification effects among others. A characterization of the anisotropy in our specific domain was given by Tordella and Iovieno using an earlier version of the same code used for the present simulation. The important metrics to analyse when looking at anisotropy at intermittency are the skewness and the kurtosis of the distribution of whatever physical quantity we are probing, which is why they will appear in a number of subsequent plots. Let us now look deeper into the two components that the simulation is comprised of: the Eulerian fields and the Lagrangian droplets.

2.2.1 Fields simulation: velocity, temperature, water vapor

A visualization of the computational domain can be seen in Fig. 2.2 A portion of the cloud interface region is modeled as a parallelepiped made of two stacked cubes with a size of $512 \times 512 \times 512$ points, amounting to $0.512 \times 0.512 \times 0.512$ m. This means the resolution is of 1 mm, which in turn means that the smallest possible scale to be resolved will be of about 0.5 mm. One cube contains the cloud part, and the other the clear air. As the title introduces, the Eulerian part of the simulation works by simultaneous, direct solution of equations for five fields: three velocity components, temperature and water vapor concentration. A big part of the modeling consists of how these fields are initialized. The reason why this juxtaposition gives rise to a mixing layer is that clouds have been found to usually contain more intense




Figure 2.2: (a) Geometry of the computational domain. (b) Visualization of the vapor concentration in the cloud portion of the domain. Plots for the kinetic energy, vapor density and temperature. [30]

turbulence than the surrounding air (as already mentioned in the previous section). To introduce this aspect in the simulation, the same velocity field is initialized for both cloud and clear-air cubes, but the values for the cloud part are then multiplied by a constant. This has the important outcome of limiting the inhomogeneity of the flow to just the kinetic energy content while leaving the turbulent integral scale unaltered. Similarly, the temperature in the cloud is initialised to be higher, and the same for the water vapor concentration (which is above saturation in the cloud). All values are chosen to reflect the best available measurements for these quantities in real clouds. The two cubes are initially generated separately as homogeneous regions, then a narrow hyperbolic tangent type of function is superimposed to the values so as to smooth out the border in between the two. It is important to underline that the turbulence will be freely decaying, meaning that no forcing will be applied to the flow (this is one more way in which this situation differs from the stationary turbulence from Chapter 1). Fig. 2.3 gives an overview of the initial values for some of the quantities just listed, as well as their temporal evolution.

The DNS code used is an iteration of a code that has been in development since 2001, when it was first published [32]. The full equations it solves are incompressible Navier-Stokes in the Boussinesq approximation for buoyancy for the flow, plus two




Figure 2.3: Spatial and temporal evolution for kinetic energy and dissipation rate in the computational domain [30]. (d) includes plots for the Kolmogorov length η [31].

for the scalar fields (Einstein notation is adopted):

$$\frac{\partial u_j}{\partial x_j} = 0 \quad (2.25)$$

$$\frac{\partial u_i}{\partial t} + u_j \frac{\partial u_i}{\partial x_j} = -\frac{1}{\rho_0} \frac{\partial p}{\partial x_i} + \nu \frac{\partial^2 u_i}{\partial x_j^2} - \mathcal{B} \delta_{zi} \quad (2.26)$$

$$\frac{\partial T}{\partial t} + u_j \frac{\partial T}{\partial x_j} = \kappa \frac{\partial^2 T}{\partial x_j^2} + \frac{L}{c_p} C_d \quad (2.27)$$

$$\frac{\partial q_v}{\partial t} + u_j \frac{\partial q_v}{\partial x_j} = D \frac{\partial^2 q_v}{\partial x_j^2} - C_d \quad (2.28)$$

where u_i are the velocity components, ρ_0 is known as Boussinesq density and it is the mean air density at an altitude of 1 km, p is the pressure field, ν is the kinematic viscosity of air, T is temperature, κ the thermal diffusion coefficient, L the latent heat of condensation-evaporation, c_p the specific heat air, ρ_v the vapor concentration and finally D the molecular diffusion coefficient. As can be seen, the scalar fields are transported passively. It is important to remember that, for numerical reasons, periodic boundary conditions are used, which should make us wary about the values of the fields near borders, especially the top and bottom ones. \mathcal{B} and C_d are the buoyancy and condensation terms, respectively. \mathcal{B} has the typical Boussinesq form

$$\mathcal{B} = g \left(\frac{T - \langle T \rangle_z}{\langle T \rangle_z} + \left(\frac{m_a}{m_v} - 1 \right) \frac{\rho_v - \rho_0}{\rho_0} \right) \quad (2.29)$$

g being gravitational acceleration, m_a and m_v are the molar masses for air and vapor and the $\langle \cdot \rangle_z$ averages are taken over planes perpendicular to the z axis (which is the direction of anisotropy). The condensation term is somewhat delicate in that it is actually properly defined in the Lagrangian frame of the droplets, since condensation occurs at their position. To incorporate it into the Eulerian equations we have to transform the Lagrangian information about droplet growth into a field. This is done by means of the expression

$$C_d(x_n, t) = \frac{1}{\Delta x^3} \frac{dM_w(x_n, t)}{dt} = \frac{4\pi\rho_w}{\Delta x^3} \sum_{j=1}^{N_n} R_j^2(t) \frac{dR_j(t)}{dt} \quad (2.30)$$

x_n refers to the n -th cell in the computational domain, and correspondingly $M_w(x_n, t)$ is the total amount of liquid water in that cell at time t . Δx^3 is the cell volume. j is a counter that runs over all droplets inside the cell being considered, the number of which is N_n . R_j is the radius of the j -th droplet. With this, we can move to the other reference frame and introduce the droplet population.

2.2.2 Droplet simulation: Lagrangian dynamics and growth

Simulations were run with two different types of droplet populations: monodisperse, meaning that all droplets start with the same radius, and polydisperse, meaning that the population is initialised to a wider spectrum of radii. As mentioned earlier, there is no definitive consensus on the droplet size distribution in real atmospheric clouds, and therefore the two extremes were chosen for inspection: the monodisperse distribution being delta-peaked on a given value, and the polydisperse one being uniform in mass (meaning that each radius interval contains the same water mass, implying in turn that small droplets are much more numerous). The monodisperse distribution may seem very unlikely to correspond to a real instance, but it is worth studying for at least two reasons: first, it has been the one used in many past studies from a variety of authors, and as such it allows for result comparison; secondly, it is not impossible for some mechanisms operating in the early phases of cloud formation to have a uniforming effect on droplet radii, thus bringing the spectrum closer to a monodisperse one. The total liquid water content was also chosen to correspond to real cloud measurements, and then distributed on the two droplets populations. This results in 8×10^6 droplets in the monodisperse simulations, and a number of comparable magnitude in the polydisperse ones. Each droplet is given an identification number, three position components, three velocity components and a radius value

The motion of each i -th droplet is given by simple dynamic equations:

$$\frac{d\vec{X}_i}{dt} = \vec{V}_i \quad (2.31)$$

$$\frac{d\vec{V}_i}{dt} = \frac{\vec{u}(\vec{X}_i, t) - \vec{V}_i}{\tau_i} + \left(1 - \frac{\rho_0}{\rho_w}\right) \vec{g} \quad (2.32)$$

where obviously \vec{X}_i and \vec{V}_i are the position and velocity, and τ_i is a timescale that allows to keep into account drag, defined as

$$\tau_i = \frac{2 \rho_w R_i^2}{9 \rho_0 \nu} \quad (2.33)$$

We note that the transport of the droplets is also passive, meaning that they do not modify the flow field at all.

As anticipated, the diffusional growth of droplets is regulated by Köhler's theory. As it was introduced theoretically in the dedicated subsection, though, it is rather cumbersome to implement numerically. It has then become a standard to use an approximated version of the theory, that still deals with all the same effects. It is summed up by the equation for the radius growth rate:

$$R_i \frac{dR_i}{dt} = K_s \left(S - \frac{A}{R_i} + \frac{Br_d^3}{R_i^3} \right) \quad (2.34)$$

Where S is the supersaturation, whose value is derived from the vapor mixing ratio, r_d is the dry particle radius that describes the amount of solute in the droplet, and A and B are the terms that account for curvature and solute effects, respectively. There is some literature available to choose these parameters appropriately, which was employed in our case to choose the values $A = 1.15 \cdot 10^{-7}$ and $B = 0.7 \cdot 10^{-18}$.

The last mechanism to cover is the coalescence of particles. The implementation is quite straightforward: whenever the distance between two particles becomes smaller than the sum of their radii the particles coalesce, meaning that a new drop is born with a volume equal to the sum of the volumes. In view of the preceding discussion it should be clear what this implies in terms of stochastic process parameters: indirectly, both the collision efficiency and the coalescence efficiency are set to one. This constitutes one way in which the modelling slightly simplifies reality, in that we know that such a value is not quite accurate. The choice is justified by the fact that the code is already rather complex and running time is a major concern, so each new calculation needs to be evaluated in terms of what it adds to the results. In this case, the intent of the simulations is to capture behaviors of turbulent flows and droplet populations, and specifically to compare these behaviors in the bulk of the cloud and in the mixing region. It is apparent that if we had considered a lower coalescence efficiency, put keeping it the same in both regions, these behavior would have changed as far as absolute values are concerned, but not in terms of their relationship. A practical example: we will see shortly that in the mixing region droplets are more prone to collision, and by extension to coalescence. If the collision or the coalescence efficiency were lower than one, we would find a lower number of collisions overall, but they would still be more in the mixing region proportionally. One effect that may be investigated in the future is the possibility for these parameters to take different values in the two regions, but it was beyond the scope of these simulations.

With all of this said, the context for the achieved results should be clear, and we can move on to looking at what has been seen from these simulations so far.

2.3 Results so far: cloud top boundary phenomenology

The title of the section gives one first specification: the results obtained in terms of water droplet dynamics were obtained by looking at simulations in which the cloud portion is the one on the bottom, corresponding to what would be the top region of a cloud. This is important to mention, and to differentiate with respect

to the bottom of the cloud, because of stratification effects. Like some other topics touched on so far, a discussion of fluid stratification in the atmosphere would lead us too far astray, and will therefore be put aside, but suffice it to say that stratified turbulence is a whole field of research in and of itself and the influence of density-temperature-buoyancy interplay is extremely relevant to atmospheric phenomena. A study of these effects in a region analogous to our was carried out by Gallana *et al.*[33] by using the same code illustrated above but excluding particles, which allowed them to achieve a larger domain size of $1024 \times 1024 \times 2048$.

The results I am going to illustrate were the subjects of two very recent papers, by Golshan *et al.* (2021) [31] and Fossà *et al.* (2022) [30]. For convenience we will proceed in chronological order, starting with a discussion on the dynamics of the droplet size spectrum.

2.3.1 Acceleration of droplet population dynamics

The first effect we are going to look at is the differentiated dynamics that the droplet populations go through in the mixing region with respect to the cloud bulk. As a reminder, the monodisperse distribution is initialised with all droplets having a radius of $15 \mu m$, whereas the polydisperse distribution is uniform over mass and results in a radius distribution spanning from $0.6 \mu m$ to $30 \mu m$.

The temporal evolution for the two populations is displayed in figures 2.4 and 2.5. We can immediately identify some key features.

Let us start by considering the monodisperse population. In both regions the initial delta-like peak at $15 \mu m$ mostly spreads towards lower values, which is the sign of evaporation. This is very clearly happening at a faster rate in the mixing region, which just by inspection can be seen to be much larger at the final times. In addition to this spreading peak there is a smaller peak at much larger R values, which is due to droplet coalescence. It is apparent that in the mixing region also the coalescence-related part appears much wider than in the cloud bulk, thanks to a wider variety of droplet sizes being able to collide and generate new droplets.

For the polydisperse population, instead, the opposite effect seems to take place. While in the cloud region the original distribution shows very little variation, in the mixing a shrinking is observed. The effect is obviously not drastic, but we should remember that we are only observing the system for what amounts to few seconds. If a prolonged observation were to confirm this trend as continuing similarly it could well lead to a considerable peaking of the distribution. Additionally, in the mixing condensation is severely limited, so that coalescence is by far the main growth mechanism.

So it looks like the two initial populations tend to evolve rather differently, and the mixing at the interface region exacerbates these behaviors. We would then like to have a measure of the relative evolution for the two populations in

Monodisperse initial distribution




Figure 2.4: Temporal evolution of droplet populations with a monodisperse initial distribution. Top plots represent number density, bottom plots the mass density [31].

both regions. Since one distribution is spreading while the other is shrinking (in terms of "width", which in quantitative terms means standard deviation), by interpolating the evolution for the standard deviation and equating the resulting laws we can deduce how long it would take for the two distributions to achieve the same standard deviation value. The result is that in the cloud region this would take about $100\tau_0$ (τ_0 is the eddy turnover time, a turbulent time-scale often employed in simulations to describe quantities in terms of simulation parameters rather than absolute dimensions). In the mixing region a much smaller value is found: $18\tau_0$, meaning a more than five-fold speed-up. Notably, this is in spite of the fact that the turbulent energy in the mixing region is much smaller (something that will be important in the next subsection as well).

2.3.2 Time and space-dependent collision kernel

I purposefully skipped a discussion of stochastic models that describe the collision coalescence process, because it would have been excessively time-consuming to go through and it would not have been justified by sufficient pertinence with the subject at hand. To explain the next result, though, I will have to go back on my word at least partially to introduce the quantity in the title, the collision

Polydisperse initial distribution




Figure 2.5: Temporal evolution of droplet populations with a polydisperse initial distribution. Top plots represent number density, bottom plots the mass density [31].

kernel. Suppose we are trying to describe droplet growth in terms of the stochastic evolution of a size distribution function $f(R)$, it will satisfy an equation of the Boltzmann type, with terms accounting for all various sources and sinks of droplets of radius R (see for example [34]). It will then contain a term describing the coalescence of particles of radii r_1 and r_2 to form one of radius r , which will look something like

$$\int_0^r \int_0^r \Gamma(r_1, r_2) f(r_1) f(r_2) dr_1 dr_2 \quad (2.35)$$

with the constraint that the volume of the two colliding droplets should sum up to the volume of the droplet of radius r . $\Gamma(r_1, r_2)$ is the collision kernel, which is a parameter that describes (very roughly speaking) the probability for two droplets of radii r_1 and r_2 to collide and coalesce (indeed, it is also called collection kernel, coherently with the terminology already introduced). In most well known literature the collision kernel is defined by a constant value, because those studies are generally carried out for stationary, homogeneous isotropic turbulence. We can then try to extend the concept to our case of time-decaying and inhomogeneous turbulence by computing the collision kernel as

$$\Gamma_{FS}(r_1, r_2, t, V) = \frac{N_{coll}}{n_1 n_2} \frac{V}{t_2 - t_1} \quad (2.36)$$

where FS indicates the flow structure, meaning either homogeneous or inhomogeneous, t is a time label that indicates the time interval $[t_1, t_2]$, V is a certain volume selected for counting, N_{coll} is the number of collisions between droplets of radii r_1 and r_2 , and n_1 and n_2 are the number of droplets having radii r_1 and r_2 (clearly, because of the small probability of two droplets having exactly the same radius, some sort of binning has to be employed). The result of the analyses can be appreciated in figures 2.6 and 2.7. The two axes in the plots contain the radii of the colliding particles, whereas the color mapping shows the value of the collision kernel. The simulation time was divided in three parts, stacked vertically in the figures, and the left and right plots show values for the mixing and cloud regions, respectively.

The difference between the two initial populations is quite plain to see: due to the fact discussed above, namely that similarly sized drops are less likely to collide, there are many parts of the plots of the monodisperse distribution left white, meaning that no collisions happened at all. The plots for the polydisperse distribution are much richer in that sense. In both cases, though, it is clear that the values in the interface region are much higher than in the cloud bulk, for the same radius pairs. As mentioned above, this is in contrast to the lower turbulence intensity in this region, so the result is somewhat counter-intuitive. As a quantitative measure, in the monodisperse case the number density of collisions decreases by 76% throughout the transient, alongside a decay in kinetic energy by 92%. In comparison, in the interface region the collision number density only decreases by 50%, with a similar energetic decay of 86%. It is then reasonable to imagine that there may be some correlation between the anisotropic properties of the flow and the number of collisions. In this regard, one more result can be given from the simulations: Fig. 2.8 shows the Pearson correlation coefficient between moments of the velocity spatial derivative in the anisotropic direction (which, as we said, are a measure of the anisotropy) and the number of collisions for both cloud bulk and interface regions. While the values oscillate around 0 for the bulk region, in the interface even quite high values are reached. This aspect will of course need further development, as it is quite hard to characterise, but the perspective seems promising.

2.3.3 Microphysical timescales

The next result that I will quickly illustrate relates to the parametrisation of cloud phenomena. This is one of, if not the, major concerns of cloud physics as far as weather forecasting is concerned. Indeed, we are very far from the day in which

clouds will be inserted in numerical models at full resolution at the particle level (if such a day should ever come), and they are therefore treated on a stochastic basis by means of parameters. Among the interesting parameters to choose and use are time-scales. They are already widely used in turbulence studies, and we have encountered some of them so far. Analogously, they can be defined for different processes that take place in cloud droplet populations. We will define a few of them and then take a look at their behavior in the simulations described.

First we introduce the phase relaxation time, τ_{phase} . This is the time scale that describes the evolution of the supersaturation field and is defined as $\tau_{phase} = (4\pi\kappa_v N \bar{R})^{-1}$, with N the total number of droplets and \bar{R} the average radius. Second is the evaporation time scale, defined as the time it would take for a droplet of radius R to completely evaporate in a subsaturated environment at $S_0 < 0$. Numerically it is $\tau_{evap} = -R^2/K_s S_0$. Note that this timescale is only defined in subsaturated regions. We can therefore define a similar one for the process of condensation, defined in regions above the saturation point. We define a τ_{cond} as the time it would take for a droplet of radius R to achieve a radius of $2R$, i.e. $\tau_{cond} = 3R^2/2K_s S_0$. Finally, we want to define a time that inherently takes into account the fact that in our region both saturation and integral radius (defined as $N\bar{R}$) fluctuate, and therefore all processes occur concurrently. We call this time scale τ_{react} , and we define it as the shortest time before a droplet has either completely evaporated or the parcel in which it is located has become supersaturated.

In Fig. 2.9 the behavior of the time scales, averaged along planes perpendicular to the anisotropic direction, is displayed. The interesting feature is that there is a point, inside the mixing layer, in which the timescales seem to converge to very similar values. This gives hope that the various processes may be parametrised by a single time scale, thus reducing the modelling complexity.

2.3.4 Supersaturation balance

The fourth and last result I want to discuss relates to the supersaturation balance equation, which was introduced in section 2.1.3. There it was written in terms of saturation ratio, but the two quantities are easily related since supersaturation is just the excess vapor concentration beyond the saturation point. The equation that was introduced was

$$\frac{dS}{dt} = P + C \tag{2.37}$$

As explained term P accounts for supersaturation production thanks to updraft, and therefore in our case is null. Term C instead accounts for supersaturation absorption by the condensation process, and in our case is written as $C = -S/\tau_{phase}$. Since $P = 0$, we should expect that $dS/dt - C \simeq 0$. By inspection of Fig. 2.10 it is

evident that, most noticeably in the mixing region, the two terms do not actually balance out, and some additional effect is generating supersaturation fluctuations. This is most promptly identified in the turbulent, intermittent nature of the fluid motion. The balance equation then assumes the new form, which is given in terms of the usual planar averages to reflect the analysis along the direction of anisotropy:

$$\frac{\overline{dS}}{dt} + \frac{\overline{S}}{\tau_{phase}} = \overline{P}_t \quad (2.38)$$

The form of P_t is obviously hard to pin down exactly. By observing that the new term is most active in the most highly intermittent region, though, we can deduce that it will be somehow related to the quantities that characterise this intermittency. It is reasonable to put forward the idea that, for example, we might have $\overline{P}_t \sim \text{cov}_{S, \partial u_i / \partial x_i}$. Two velocity derivatives are plotted alongside the production term in fig. 2.10. To validate this hypothesis, like in the previous subsection, the most immediate tool is given by the Pearson correlation coefficient. We generally need to be wary of low correlation coefficient values, as the Pearson coefficient is only able to capture linear correlations, but high values are usually a good indicator. Indeed, as can be seen in Fig. 2.11, the values are quite high, meaning that some sort of linear correlation between the production term and the covariance described can be expected. In red in the figure is plotted the factor that the covariance would need to be multiplied by to obtain the computed production term, K .

This concludes the discussion about the latest results obtained for the simulations. We can now move on and take a look at the experimental tools that complete the picture of our most recent studies and observations.

Monodisperse initial distribution




Figure 2.6: Time and space-dependent values for the collision kernel in the monodisperse initial population case. On the left are the values for the interface region, on the right for the cloud bulk. Time evolution can be observed in three steps from top to bottom [31].

Polydisperse initial distribution





Figure 2.7: Time and space-dependent values for the collision kernel in the polydisperse initial population case. On the left are the values for the interface region, on the right for the cloud bulk. Time evolution can be observed in three steps from top to bottom [31].

(a) Pearson's correlation index between air flow intermittency and collision count inside cloud



(b) Pearson's correlation index between air flow intermittency and collision count inside cloud interface




Figure 2.8: Pearson correlation coefficient between standard deviation (σ), skewness (S) and kurtosis (K) of the spatial derivative in the anisotropic x_3 direction of the corresponding velocity component and the number of collisions throughout the transient [31].




Figure 2.9: Plots for the spatial variation of microphysical timescales inside cloud top border simulations at times close to the starting and ending points, along with supersaturation plots [30].




Figure 2.10: Excess in the balance of supersaturation evolution terms, alongside covariances of the supersaturation S and velocity spatial derivatives [30].





Figure 2.11: Correlation coefficients (blue) between the turbulent supersaturation production term and its proposed forms in terms of covariances. Multiplicative constant (red) that would be required for the proposed term to account for numerical values [30].

Chapter 3

Experimental setup: radiosonde cluster



In this chapter I will present the tools that were used in the experimental part of the work we conducted (and are currently conducting). (As presented in the introduction,) These consist in a cluster of free-floating transceivers carried by helium balloons, one or more ground stations for data reception, and a post-processing machine for data elaboration and storage (see Fig. 3.1). Because the focus of this thesis is not the project and development of these tools themselves,




Figure 3.1: Experimental setup for radiosonde usage [35].

but rather their application to the investigation of the physics of turbulence and the attained results, I will mostly try to limit the scope of the discussion to the relevant information: what the radiosondes are built for, what they are capable of, and the tests they went through to ascertain their usability. The chapter will then be structured as follows: I will start by describing the sondes and their components; then I will show the first tests that were conducted in lab to validate the sensors in very controlled conditions; then we will move to the field by taking a look at similar tests conducted out in the open with the help of other industry-established sondes commonly used in meteorological telemetry; after this we will look at the comparison of some measurements taken by the sondes in cluster, to check them for consistency; finally we will discuss the balloons, which are a crucial part of the design to ensure that the displacement measurements are valid.

3.1 Idea and main requirements

Atmospheric measurements are continuously carried out in most of the world, in the great majority of cases for the purpose of weather forecasting. A great variety of tools are employed for these measurements: weather ground stations with sensors

for virtually all relevant quantities, radars and lidars for remote measurement of wind velocities and some cloud and hydrometeor properties (among other possible applications), satellites for large scale imaging and, most relevantly to us, weather balloons. Generally speaking, these are very large balloons with very heavy electronic circuits attached to them that are built to be as robust as possible and to provide reliable, consistent measurements for atmospheric profiling. What these are not built for, however, is the measurement of turbulence properties. Indeed, their size and weight makes them unable to adequately follow the air flow in its characteristic turbulent behavior, and as such they cannot be used for the kind of research our group has devised. Additionally, because of their rather high cost, the launching of large clusters of them would have prohibitively high prices, which just goes to underline the need for a new tool.

Indeed, the main characteristic of the tool to be developed was to be launched in clusters. As I hopefully made clear in Chapter 1, when looking at turbulence the interest lies in the fluctuations of physical quantities. These fluctuations are both in time and space, and as such there is a need to measure the same quantities at several different points in the flow. This is commonly done in lab settings with fixed point sensor arrays, and to a lesser degree it can be done in the atmosphere. When the object of investigation becomes the Lagrangian properties of turbulence, though, the measurement has to be performed by passively following the flow. With all this in mind, it should be intuitive to understand which are the core requirements that were set before development. First and foremost, the sondes should be small and light. The two clearly go together: as will be explained in detail later, the size of the balloon (which is obviously much larger than the sonde) depends on the desired floating altitude and on the payload, i.e. the radiosonde weight. The circuit which makes up the device has thus been reduced to the very bare bones in order to diminish the weight as much as possible, and more optimization is on its way in future prototypes.

Regarding the sensors, there is a great variety of potentially interesting quantities. Because the project is still moving its first steps, though, the current versions of the sondes are limited to measuring position, velocity, acceleration and then temperature, pressure and humidity. This allows to keep the weight in check, and to make operation as smooth as possible. As far as position is concerned, the ideal would be to have an uncertainty close to the size of the balloons. With current technology, though, this is virtually impossible, and we have to accept the accuracy of GPS sensors, of the order of few meters (more on this later). For temperature, the focus at the time of development was the study of warm clouds, so a temperature range of 0°C to $\sim 30^{\circ}$ should be sufficient. Luckily, most sensors provide a much wider range, so that poses no problems. For pressure, the range of interest is about 400 mbar to 1100 mbar. Humidity is obviously the most problematic quantity: the required range is 0% to 100%, but ideally even beyond this, as supersaturation

is one of the key players in cloud formation. Unfortunately, there are no sensors that can detect supersaturation. The next best thing is to try estimating it by combining the readings of different quantities, which is among the future todos, when the launches will actually target clouds.

One unfortunate factor that also impacts the design is, as mentioned above, cost. The reason is twofold: on one hand, just realistically, experiments could not be conducted if the cost were too high, so it is necessary to keep it contained no matter what; on the other hand though, the prospective applications of the new technology and methodology could potentially make it palatable for industries. Indeed, while to us they are experimental tools to probe atmospheric turbulence, they could be used in many other contexts such as environmental monitoring in cities, pollution control, weather prediction and the like. This obviously requires them to be comparatively affordable, both to make them fit in the budget of institutions and companies usually carrying out these tasks and because they would be required in large amounts.

Lastly, environmental concerns are to be taken into account. At a time in which climate consequences of human behavior are becoming more and more evident, it was decided to try and reduce the impact of the sondes as much as possible, since they are not meant to be recovered after use. We thus looked for biodegradable materials to build them out of. The main issue is with the balloons, due to their size, but in the future there is an intention to look for valid alternatives for the circuit board as well. The solution adopted for the balloons will be detailed in the appropriate section.

3.2 The radiosondes

This section is devoted to the description of the radiosonde component and their characteristics. As can be seen in 3.2, in their current form the sondes are made up of nothing other than the electronic circuit connecting the components and the board on which they are mounted. There is good reason for this: any additional casing or non-functional feature would increase the weight and become counter productive.

Control unit This is the heart of the circuit, where all the data processing happens and where the signals from the various sensors are managed. The current versions of the probes mount the ATmega328 microcontroller from Microchip, with dimensions 9 mm x 9 mm x 1 mm at a weight of 70 mg. The required supply voltage range is 1.8 - 5.5 V, the temperature range is -40 - +85 °C and it also features a power-saving mode.




Figure 3.2: (a) Complete setup of one radiosonde with its MaterBi balloon. (b) Radiosonde and connected battery [35].




Figure 3.3: Block diagram of the radiosonde components [36].

Communication system For sonde - ground communications the LoRa technology was adopted. This is a modulation technique which is most commonly implemented in the framework of the LORAWAN protocol. In our case, instead, it was used to create a private network suiting the purposes of our system. The on-board component implementing the technology is the RFM95 module from HopeRF. One of the most prominent features of LoRa is its low power requirement, essential for us to keep battery size small. The chip requires a voltage range of 1.8 - 3.7 V, it has a working temperature range of -20 - +70 °C and features, aside for sleep and idle modes, three different transmit modes at +7 dBm, +13 dBm and +20 dBm output powers.

Antennas The boards mount two different antennas, one outputting the collected data to the ground stations and the other for reception of tracking data from satellites. The first one is 5 mm x 3 mm x 0.5 mm and the second one is 3.2 mm x 1.6 mm x 0.5 mm, and both are linearly polarized. Shielding was implemented to ensure minimal interference, and an L-type matching network was adopted for both to maximise power delivery. Additionally, both antennas were tuned to the appropriate working frequencies of 868 MHz and 1575 MHz respectively. The matching yielded very good results, which are summarised in the table below.

Frequency	Initial S_{11}	Final S_{11}
868 MHz	-0.56 dB	-21.09
1575 MHz	-1.22	-23.09

(Briefly speaking, S_{11} is the component of the scattering matrix for the circuit - antenna system that describes how much power the antenna reflects back at the circuit. Simply put, the value should be as low as possible to ensure most of the power is delivered to the antenna.)

Sensors For the measurement of physical quantity it is nowadays commonplace to adopt several-in-1 integrated solutions, which are particularly convenient and offer easy data manipulation off the shelf. For this project the BME280 sensor from Bosch was adopted. It offers temperature, pressure and relative humidity readings and, once again, keeps power consumption basically as low as possible. The size is of 2.5 mm x 2.5 mm x 0.93 mm, the required voltage is 1.2 - 3.6 V. The ranges it should be able to capture are 0% to 100% RH for relative humidity, 300 hPa to 1100 hPa for pressure, and -40 °c to +85 °C for temperature, with respective uncertainties of $\pm 3\%$ RH, ± 1 hPa and ± 1 °C.

Tracking There are two components that have the purpose of giving information about positioning. One is an inertial measurement unit (IMU), meaning a sensor that provides relative displacement and rotation information, and the other is a GNSS (Global Satellite Navigation System) receiver similar to those found in smartphones and navigation systems. The selected IMU is the LSM9DS1 by STMicroelectronics, dimensions of 3.5 mm x 3 mm x 1 mm, supply voltage range of 1.9 - 3.6 V. This sensor provides readings on acceleration, angular rate and also magnetic field. The GNSS receiver is the ZOE-M8B by u-blox. The size is of 4.5 mm x 4.5 mm x 1 mm and the required voltage range is 1.71 - 1.89 V. Once again, one of the reasons to choose this particular sensor was the availability of a power saving mode (Super-E) that allows to optimise battery consumption. The input the GNSS is also aided by a low noise amplifier to optimise performance. The accuracy provided by the GNSS is of 2.5 m in the standard operational mode (called continuous mode) and of 3.5 m and 4 m in the two power saving modes (Super-E and Super-E power save).

Power supply It should be clear by now that battery capacity is one of the key issues to be dealt with in the design. A balance needs to be struck between the weight and the capacity, since we would like to acquire signals for as long a time as possible. The batteries used for the launches described in this work weigh 8 grams (more than the sonde itself), and has a nominal capacity of 125 mAh. While future prototypes of the sonde are going to shave more weight off of it (the one being developed at the moment weighs just 3 grams), the battery weight remains problematic as the power density is roughly fixed.

3.3 Balloons

The balloons are, as of yet and maybe surprisingly, the most problematic part of the system. This is due to the fact that, despite their apparent simplicity, many issues concur to making their production particularly delicate. Chiefly, the balloon's volume needs to remain constant, so as to make sure that the base altitude at which they float does not vary and their motion is dictated only by the air flow rather than by buoyancy effects. This rules out all expanding materials that are most commonly used for weather balloons such as latex. Secondly, as mentioned before, there is a desire to keep the whole project as environmentally friendly as possible, and while some sources of pollution are hard to replace (such as the batteries) the balloons constitute a large portion of material that can be replaced by some sustainable option. This section explains the steps that went into the design and that are currently taken for production of the balloons used in some of the launches. It is important to note, though, that due to the aforementioned issues production

remains problematic, forcing us for the time being to also rely on some off-the-shelf Mylar balloons that, while not biodegradable, have excellent performance in terms of helium retention and durability.

Material An in-depth study of the material to make the balloons out of has been conducted at the early stages of the project [37]. Specifically, after some initial research to pick possible candidates, two relatively available materials were compared: MaterBi and PLA. Ultimately MaterBi performed best, especially as pertains to durability, as PLA reacted poorly to pressure increments, which during an atmospheric flight are to be expected. In the same study some coatings to be applied on top of the MaterBi are also tested. This is due to the fact that biodegradable materials, unsurprisingly, do not hold up well to moisture, and since the final goal is to use the sondes for measurements in cloud a way must be found to improve performance under this aspect. It was found that a spray coating of SiO₂ nanoparticles helped greatly with both hydrophobicity and helium impermeability.

Design A number of Master and PhD students have, by now, contributed to the design of the balloons, which goes hand in hand with their production. The current version consists of 4 sheets of MaterBi cut in the shape as displayed in Fig.3.4. On the left side is the neck in which helium is injected to fill the balloon. The wedge protruding on the right side is my own personal contribution: before this part was added the four sheets would be attached to a circular additional piece of MaterBi, serving as a sort of “hat” at the top. The cutting of a differently shaped piece, though, as well as the attachment, was rather encumbering during production. I thus got rid of it and adopted these extra slices of material that, by being layered on top of each other, allow to close the top of the balloon without the need for a cap.

The most delicate aspect of design is, of course, the sizing. There is a convenient formula to be used in this case:

$$V_b = \frac{m_r + m_b}{\rho_a - \rho_g} = \frac{m_r + m_b}{\rho_a(1 - M_g/M_a)}, \quad (3.1)$$

Here V_b is the volume required for the balloon to lift to a given altitude a payload of mass m_r , while having itself a mass of m_b , when filled with a gas of molar mass M_g and surrounded by air of molar mass M_a and density ρ_a at the target altitude (equivalently, ρ_g is the density of the gas inside the balloon at the same altitude). Obviously enough, for circular balloons this can also be expressed in terms of radius, and Fig. 3.5 sums up the requirement in a plot for various payloads.

For our case, in which the probe attached to a battery weighs 17.5 grams and we want to keep the altitude between 1000 m and 2000 m, it is evident that the radius should be around ~ 20 cm.




Figure 3.4: Current design of MaterBi “petals” for production of biodegradable balloons.




Figure 3.5: Plot that links balloon radius to the attainable altitude for different values of the payload m_r , considering MaterBi as the balloon material [35].

Production By far the most problematic aspect as far as the balloons are concerned is the actual production. Because they are completely customised research tools there is, as of yet, no automatised production method, and they have been entirely crafted by hand. This involves procuring MaterBi in the form of a

large double-layered sheet, tracing the petal shape from a cut-out, cutting it, and then welding it all together by using a plastic welder, meaning a small heated wheel to be run across the edge of the material to fuse the petals. The temperature of the wheel, though, is not originally intended for use with biodegradable plastics, and therefore is much too high for use directly on MaterBi. Various solutions have been adopted to circumvent this problem, the latest one being the interposition of a layer of Teflon between the wheel and the balloons during welding. Still, the operation puts a lot of stress on the MaterBi sheets, and tears are almost unavoidable. These can usually be fixed with the help of tape and the like, but other times can be hard to identify and present a significant issues. Additionally, as is often true for manual work, it can be hard to keep the production at a sufficient level of consistency and exactness. For our application an apparently tiny difference in the balloon radius can make for a rather noticeable difference in floating altitude. All this, together with the considerable time it takes to go through the manual process, has led us to explore the possible avenues of automation. These remain a challenge for the future, and have two major problems to solve: the first one is the fact that the current structure of the balloons is not particularly suitable for industrial automatic production, but at the same time it is not trivial to modify it in a way that keeps the desired characteristics; secondly, once again, cost becomes very much a factor, since this type of highly customised work necessitates the construction of specially manufactured tools that are far from cheap.

3.4 Lab tests

Once the components were chosen, and the first sondes assembled, some tests were carried out under heavily controlled conditions to determine their accuracy. We will briefly relate them to show that the sondes generally yield good readings and the components work well within the overall systems.

3.4.1 Temperature and humidity tests

The PHT sensor was tested at the Applied Thermodynamics Laboratory in the INRiM (Italian National Metrology Institute) facilities. They provided access to a climatic chamber (model Kambic KK190 CHLT) that offers precise control over the quantities of interest. The probes were tested against high-precision tools calibrated at INRiM: Pt100 sensors paired with a FLUKE 1594 thermometer for temperature and a Delta Ohm sensor for relative humidity. To get a sense of the average performance of the probes three of them were used at the same time.

Fig. 3.6 sums up the main temperature test. While keeping the humidity constant at $RH = 30\%$, the initial temperature is set at $T = +24\text{ }^\circ\text{C}$. It is then dropped to $-5\text{ }^\circ\text{C}$ and kept constant for an hour. After this time, increments of 5

$^{\circ}\text{C}$ and 10°C are applied until reaching $T = 10^{\circ}\text{C}$. Each step is against held for an hour, which is necessary for the system to stabilise at the desired temperature (some overshoot can indeed be seen when the temperature is first raised). As can be seen, the probe sensors follow the variations quite truthfully.




Figure 3.6: Temperature test of implemented sensors inside the climatic chamber. Humidity is kept constant at $\text{RH} = 30\%$ while temperature is increased from -5°C to 10°C . [36]

A similar procedure was followed for relative humidity measurements. The temperature was kept constant at $T = 30^{\circ}\text{C}$, and the humidity was increasingly set to 10% RH, 20% RH, 40% RH and 60% RH, holding each value for about 30 minutes. The quantitative results for bot tests are reported in the two tables 3.1 and 3.2.

Temp. set point [$^{\circ}\text{C}$]	Ref. sensor T [$^{\circ}\text{C}$]	Sonde 1		Sonde 1		Sonde 1	
		Mean [$^{\circ}\text{C}$]	Mean error [$^{\circ}\text{C}$]	Mean [$^{\circ}\text{C}$]	Mean error [$^{\circ}\text{C}$]	Mean [$^{\circ}\text{C}$]	Mean error [$^{\circ}\text{C}$]
-5	-5.063	-5.31	0.25	-5.30	0.24	-5.25	0.18
0	0.002	-0.25	0.25	-0.23	0.23	-0.17	0.18
10	9.878	9.82	0.065	9.75	0.13	9.74	0.13

Table 3.1: Results of temperature comparative test of the implemented sensors carried out inside INRiM climate chamber, see Figure 3.6. [36]

The outcome of the tests is roughly within the expectation according to the

Humidity set point [%RH]	Ref. sensor RH [%RH]	Sonde 1		Sonde 1		Sonde 1	
		Mean [%RH]	Mean error [%RH]	Mean [%RH]	Mean error [%RH]	Mean [%RH]	Mean error [%RH]
10	10.50	13.12	2.62	14.74	4.24	14.16	3.66
20	19.75	19.85	0.09	21.35	1.60	21.09	1.34
40	37.68	35.31	2.37	35.64	2.04	36.06	1.62
60	59.70	56.13	3.57	54.53	5.17	55.69	4.01

Table 3.2: Results of humidity comparative test of the implemented sensors carried out inside INRiM climate chamber. [36]

data provided by manufacturers. In most cases, the temperature accuracy is within ± 1 °C and for humidity within $\pm 3\%$ RH. The anomalies in the humidity readings will be confirmed by the other tests, which show that this sensor can sometimes suffer from some inaccuracy.

3.4.2 Position

A very simple test was also conducted to make sure the GNSS receiver would be suitable for our application. Its positioning data was compared to the one of a smartphone by walking around our university campus while acquiring data for both. The test showed good agreement between the two. Of course, this comparison is less strict than the previous one, as the smartphone chip is not necessarily more accurate than the sonde one, but it confirmed that the GNSS component was well integrated in the sonde design and is usable for our experiment. Fig. 3.7 shows the overlap between the two readings for latitude and longitude.

3.5 In-field tests

This section is devoted to the various tests that were conducted directly on the field. It is divided in two subsections: the first one will deal with tests in which some reference instrumentation is used to compare the readings from a sonde; the second one will show some readings from a cluster of sondes, thus allowing to check for consistency among them.

3.5.1 Sondes vs. other instruments

On three occasions we have had the precious opportunity of attaching one of our boards to weather balloons used by national institutions performing routine




Figure 3.7: Comparison between a smartphone GPS signal and our GNSS receiver, test carried out by simultaneous movement of the two devices on Politecnico grounds. [36]

atmospheric profiling through dedicated probes. The first two were at Levaldigi Airport, Cuneo, Italy (October 2020 and June 2021), whereas the (at the moment of writing) latest one was in Chilbolton Observatory, UK (July 2023), while taking part in the Wessex Convection Campaign (more or it later). During the first launch in Levaldigi some interference was observed, likely due to attaching our sonde directly to the other one. While some results from that launch are available, I have excluded them from the present discussion for the sake of brevity, believing that the results from the more successful launches will prove satisfactory. The Levaldigi launch was paired with a Vaisala RS41-SG sonde and the Chilbolton one was with a Vaisala RS41-SGP, two very similar models.

The first result pertains to the efficacy of data transmission from the sondes to the ground stations. The data travels in small packets that the sondes send out at regular intervals. Not all these packets reach the stations, and the amount of them that can be successfully received is obviously a function of distance but much more so of the presence of obstacles. The presence of several ground stations is itself already a measure to combat packet loss through oversampling, which proved quite valuable in the Chilbolton cluster launches. Interpolation is then used in post-processing to compensate for missing packets. Fig. 3.8 shows the number of packets received as a function of time and altitude for the 2021 dual launch.

During the Chilbolton dual launch, with clear weather and an unobstructed field of operation, we were able to keep receiving packets up to a distance of more than 34 km during the Chilbolton dual launch (Fig. 3.9). This greatly exceeds




Figure 3.8: Number of packets received during the 2021 Levaldigi dual launch as a function of (a) time and (b) altitude. [36]

expectations when using the LoRa technology, and is very promising for future cluster launches. The true challenge will be to make signal transmission robust against more adverse weather and liquid water content of clouds.




Figure 3.9: Trajectory of the 2023 Vaisala-COMLETE dual launch in Chilbolton, UK.

Moving on to position and velocity measurements, Fig. 3.10 shows the measured position from the Vaisala and our sonde for the 2021 launch. The reason for the jagged behavior of the Vaisala trajectory is the lower resolution of its latitude

and longitude readings, which also impacts the error computation. For longitude, latitude and altitude measurements the RMSE was found to be $2.97\text{e-}4$ degrees, $3.7\text{e-}4$ degrees, and 22.3 m.




Figure 3.10: Comparison of position readings for the 2021 Levaldigi dual launch. Our radiosonde is marked as COMPLETE. The inset plot in (b) shows the difference between the two readings for a selected time range. [35]

Fig. 3.11 provides instead a comparison of wind speed readings. Again, the agreement is reasonable, although some overshoot and undershoot can be detected at some points. I chose to include also Fig. mostly because it is the first occasion to see the five-thirds law in action in our experimental context. Of course, due to the noisy nature of the dataset, it is debatable whether a $5/3$ power law is truly the best fit, but as explained in Chapter 1 we should remember that, since the range of turbulence scales visited by the sondes is far from isotropic and homogeneous, we have in principle no reason to believe it would be.

The results of the comparison for temperature, pressure and humidity measurements are shown in Fig. 3.12, once again in reference to the Levaldigi launch. By inspection, while the pressure seems to pose no huge problems for our sensor, there are some unwanted features in the temperature and humidity plots.

We see that, after an altitude of about 4000 m, the detected temperature for our probe starts to get considerably higher than the Vaisala readings. This problem is further explored in Fig. 3.13. As the panels on the left highlight, there is an intrinsic offset that needs to be accounted for, since our temperature readings show a constant temperature bias of a few degrees. However, even after correcting this bias, in both launches there is a linear increase in the measured temperature with respect to the reference. The current hypothesis is that this effect is due to solar irradiance. Indeed, this is a known factor to be accounted for in atmospheric measurement, and Vaisala corrects for it in its readings, whereas we have not yet had the opportunity to put in place a counterbalance for this phenomenon.




Figure 3.11: (a) Comparison of wind speed readings for the 2021 Levaldigi dual launch. Our radiosonde is marked as COMPLETE, and the raw data is indicated by the blue dots, while the red line is given by interpolation (some of the over/under-shoot is an artifact of the interpolation, whereas some is an actual feature of the dataset). (b) Power spectrum from the velocity readings of the probes. [35]

As concerns the humidity readings, we can see that they seem to lag behind the reference ones, being unable to cover the full extent of the variations. This is a known issue with these humidity sensors, which suffer from a slow response time. Additionally, the humidity measurement is carried out through a small vent hole in the PHT sensor, that needs proper air flow through it to perform correctly. During flight it is hard to ensure such a condition, even though the positioning of the hole was carefully picked during the circuit design. It is possible to somewhat correct these issues through a posteriori compensation, but the employment of different sensors is also being considered for future prototypes.

3.5.2 Cluster tests

This last portion of the chapter is dedicated to a brief illustration of some early measurements of the quantities listed so far during collection of data from several sondes. Aside for informal, desktop tests, this has been done three times before the two Chilbolton launches: twice by keeping the sondes tethered to the ground to avoid losing them, and once in free flight. The first tethered tests were carried out at one of the INRiM facilities, and the second one took place in the same spot that later hosted the first launch: the Astronomical Observatory of Valle d’Aosta (OAVdA), located in Saint Barthelemy, Italy.

Fig. 3.14 sums up the results of the first tethered launch at INRiM. Five probes were turned on and their readings compared to those of a Vaisala WXT510 weather station, before releasing them above the facility. The panels on the left show the





Figure 3.12: Temperature, Pressure and Humidity comparison for the Levaldigi Vaisala-COMPLETE dual launch. [35]

performed check, which yielded average RMSDs of 0.065 hPa for pressure, 3.57% RH for humidity and 1.21 °C for temperature. These are in good agreement (much lower for pressure, slightly higher for humidity and temperature) with producer specifications. The panels on the right show the quantities as measured by the sondes during tethered flight. The sondes were kept relatively close, so we expect trends to show a similar behavior, but the differences shown are also to be expected.

To avoid repetitions I will skip the discussion of the second tethered set of tests at OAVdA, as it was mostly a preliminary step to prepare for the first free launch. We can then move on directly to the free launch itself. Before release the sondes were checked against portable INRiM instrumentation, which gave us an opportunity to confirm the aforementioned effect of solar radiation on our temperature readings. It is visible in Fig. 3.15: our probes, attached to a fence and ready for launching, are shown alongside three Pt100 thermometers. Two of them had no protection against solar irradiance (USH 1 and USH 2), just like our sondes,

Chilbolton, UK. July 6th, 2023. COMPLETE probe was directly fastened to Vaisala RS-41 SGP probe with a scotch tape, **possible heating from RS41-SGP**.



Levaldigi Airport, Cuneo, Italy. June 9th, 2021. COMPLETE probe attached with a 80-cm long wire to RS41-SG. **No heating from RS41-SG**.




Figure 3.13: Comparison between two dual launches with different configurations of attachment of the Vaisala and COMPLETE sondes. The larger error experienced in the Chilbolton launch by our sensors may be due to heating from the other sonde. [35]

whereas a third one was provided with a solar shield (SH). The effect is evident: the green connected triangles positioned a few degrees below all other lines are the readings from the shielded sensor, whereas the readings from our probes generally agree well with the unprotected ones. The average difference, as computed between shielded and unshielded reference sensors, was of 1.28 °C. The gray area in the left half distinguishes the moment in which the sondes were tied to the fence from the moment we picked them up, preparing for launch, but kept the checks going.

After these initial tests the sondes were finally released. Their position is tracked in Fig 3.16, whereas the readings for temperature and humidity in time are shown in Fig. 3.17. We were able to receive data for about 35 minutes, although many of the sondes stopped earlier. As the Chilbolton launches will show, the environment in which the launches are performed influences the reception quite strongly. In this first case the sondes were quickly carried above some of the lower mountain




Figure 3.14: Left: Comparison of readings from multiple sondes and a Vaisala weather station at INRiM facilities. Right: Readings from a cluster of radiosondes in a tethered launch. [35]

crests surrounding the observatory by updraft currents, subsequently disappearing behind the rocky landscape and thus, in some instances, losing signal.

The last interesting result I want to include pertains, once again, to the measured power spectrum. Fig. 3.18 shows the spectrum as measured by three of our sondes. Once again, the five-thirds rule seems like a very reasonable estimation, although everything said above regarding the spectrum computed from the Vaisala-COMplete dual launch also applies here.




Figure 3.15: Comparison of readings from radiosonde cluster and reference sensors provided by INRiM. In yellow and purple (Ref USH 1 and Ref USH 2) reference sensors without shielding from solar radiation, and in green (Ref SH) sensor with shielding. It is clear that solar radiation has a considerable effect on the readings, and the sondes are in good agreement with the unshielded sensors. [35]




Figure 3.16: Trajectories of all sondes during the 2022 OAVdA cluster launch, with the color map indicating altitude. [35]




Figure 3.17: Temperature and humidity readings of all sondes for the 2022 OAVdA cluster launch. [35]




Figure 3.18: Energy spectra for 3 probes as computed from the OAVdA launch dataset. Dotted lines are power laws plotted for reference. [35]

Chapter 4

Results on relative dispersion



This chapter gathers the results we have achieved in the investigation of relative dispersion. On the experimental front these have been achieved by means of the radiosondes described in Chapter 3. On the numerical side, as we will see, the analyses is mostly preliminary and serves as a preparation for further work to be carried out on longer simulations that are being prepared as I am writing these lines. The common goal for both methods is the study of dispersion in a

context far from the homogeneous isotropic turbulence in which it was described in Chapter 1. There are two main reasons for this. Because it is largely unexplored territory, both agreement and disagreement with the theoretical, K41-based theory can be ground for some interesting development: in case of agreement, which is obviously hardly expected, it would allow either for the relaxation of some K41 hypotheses or for a simplification of the treatment of inhomogeneous flows; in case of disagreement, which a priori is much more likely, it would allow for some insight into the behavior of particles in realistic flows, with implications for modelling and technological applications. Indeed, as we have somewhat touched on at a few points in the previous chapters, most of the real-world forecasting is based on dispersion models rather than a rigorous application of any theory or equation solving, so this prospective is just as important as theoretical advances. We will start by describing the experimental results, and then analyse the situation for the numerical simulations.

As already mentioned, the experiments were carried out in the context of the Wessex Convection (WESCON) campaign organised by the British Met Office. This was a large scale experimental campaign that took place in the summer months of 2023, and involved a large number of both British and foreign researchers with the main goal of improving atmospheric forecasting models by means of in-field measurements. From the 3rd to the 7th of July 2023 we were hosted by the Chilbolton Observatory, one of the facilities taking part in the campaign, and there we carried out two launches of radiosonde clusters, as well as a dual launch paired with a Vaisala radiosonde as already described in Chapter 3. It should be kept in mind that, while very valuable datasets were extracted, these are in some respect still preliminary launches, as some of the experimental steps that make up the launch procedure still need to be properly adjusted. For one, the sondes in their current configuration are not suitable for in-cloud measurements, as they are still sensitive to humidity and most of all to the presence of liquid water. While some ways to work around this limitation have been studied and mentioned in Chapter 3, they were not yet in place at the time of these launches. As such, the study involves relative dispersion in the atmospheric boundary layer (ABL) rather than in clouds. Secondly, the main source of uncertainty that we seek to overcome in the future is the composition of the balloon fleet and its inflation. Indeed, the logistical complications of carrying out an experiment far away from our main research facility are many and sometimes not fully in our control. They particularly struck in one instance: we came to Chilbolton bringing with us a custom built electronic valve that, by connecting to the helium tanks kindly provided by the Observatory, would have allowed us to a good degree of certainty to inject the exact same amount of helium in each balloon, thus guaranteeing that they have the same buoyancy. Unfortunately, due to a communication error, the fitting on the tank

and on the valve was actually different, and we were therefore unable to employ it. The balloons were thus filled by hand, and although they are non-expandable some slight margin of uncertainty in their buoyancy is to be expected, although hard to quantify. Hopefully the next launch, which will make use of the valve, will provide some further insight.

Despite these setbacks, the results are rich and present many interesting features. Let us now begin with a description of launch 1 and its outcomes.

4.1 Experimental results: Launch 1




Figure 4.1: Map of the July 5th launch at Chilbolton observatory with cluster trajectories.

Launch 1 was performed on July 5th, 2023. The actual launch took place at 13:52, but preparations started in the morning: each sonde is tested for operation, the balloons are inflated and tested for leakage, and then the sondes are secured to them. We selected the 5th as the first day of launch since it was an Intense Observation Period (IOP) of the Wescon campaign. This meant that each hour, starting in the early morning, a Vaisala radiosonde was launched for vertical atmospheric profiling. This allowed us to enact our plan for data packet oversampling: one of

the receiver stations would remain fixed at the Observatory, while a second one would be displaced at a point about 10 km away in the direction of the wind at the expected cruising altitude for the sondes (measured by the Vaisala launch). This would allow us to gather more data for some period of time, while both receivers were active, and then follow the sondes for longer when they would be too far for the receiver in Chilbolton. However, this plan actually evolved shortly before launch. I was part of the team that moved to the second location, and we were unable to receive any of the test packets sent before launching. Lacking confidence in our estimation of the sondes' trajectory, we switched to our second option: we went back to the observatory and left on car after the launch, following the sondes. The secondary fixed station thus changed to a mobile station. For a short time we were able to follow the sondes by sight, and after that we relied on the number of packets we were receiving (checked by inspection on a computer monitor) and on wind information provided by the Vaisala launches to pinpoint the right direction. This indeed turned out to be effective in prolonging the lifetime of the experiment. In fact, the spatial extension of the data we were able to receive is remarkable for a system utilizing LoRa technology, as can be seen in Fig. 4.1. The overall maximum distance from the stationary ground station was in excess of 44 km. Fig. 4.2 also visualizes the trajectory and spreading of the sonde cluster along the North, East and up directions.




Figure 4.2: Spreading of the sondes in the three directions North, East and up. Launch 1.

Before coming to the actual dispersion results, I will spend some words to illustrate some environmental features that played a role in the outcome of the launch. Because we are working in the field, it is important to check for what the experimental ground actually looks like.

4.1.1 Environmental characterisation

Among the advantages of carrying out our launches in the context of a larger experimental campaign is the wide availability of data from other sources about

a wide variety of different quantities. In fact, much of this data still has to be analysed by us after recently gaining access to the Jasmin repository, where most of it is stored in its raw form, but some characterization is already possible thanks to figures provided by Chris Walden, from NCAS, and to processed data found on the Cloudnet website, an open database that gathers a subset of the measurements performed daily by instruments spread in various sites around Europe and the British islands. Among them, the Chilbolton observatory is home to a lot of instrumentation for atmospheric monitoring, the most prominent of which is a 25 m-diameter antenna, and it continually provides atmospheric data for the purpose of weather forecasting (as well as other scientific endeavors).

05-July-2023 Chilbolton Radar Wind Profiler




Figure 4.3: Wind profiling by the Chilbolton radar instrumentation on July 5th, day of our first launch. Launch time was $\sim 13:52$. Top plot shows updraft, bottom plot shows horizontal wind speed and direction, the direction being identified by the orientation of the L shapes.

Figure 4.3 shows the updraft (top) and horizontal wind velocity (bottom) above Chilbolton throughout the whole day, as measured by the radar instrumentation. The color mapping identifies wind intensity in both plots, whereas the L-shaped marks on the bottom plot identify the wind direction and are to be read as if seen from a top view. On this specific day the wind was blowing to the east at a moderate speed, and inspection of the wind direction at the time of launch at an altitude around the 1000 m mark matches with the direction of the sondes' path as displayed in the map above. Ideally we would have liked wind at the ground level to be completely absent, as the moment of launch is particularly delicate, but we quickly found out from real-time forecasts that we would not have such a condition at any point in the day. This did indeed cause a couple of sondes to crash shortly after launch.

05-July-2023 14:00 Radiosonde launch from Chilbolton




Figure 4.4: Measurements from the Vaisala sonde launched during the July 5th IOP. Sonde launched at 14:00, approximately 8 minutes after our cluster launch.

Fig. 4.4 shows the quantities measured by one of the Vaisala sondes launched

for the IOP. The 14:00 sonde was selected as the closest to the time of our launch. We have readings for temperature, humidity and wind speed. This last one is still the most interesting to us for now, and it displays the characteristic boundary layer behavior. The humidity reading is related to the next figure: while it is overall rather high up until about 2500 m, the spikes toward the 100% value may indicate the presence of supersaturation (which as we already know is not directly measurable). Fig. 4.5 shows that, despite the possible presence of supersaturation or values really close to it, only some mild condensation must have been taking place at the time of our launch. Indeed, the figure plots the values of “cloud fraction” against height and time above the Chilbolton Observatory area. Rather than the measurement by a single piece of instrumentation, this is a model-based quantity that compiles data from various instruments and makes some inference about how cloudy the area is. The modeling software is the Integrated Forecasting System (IFS) by the European Centre for Medium-Range Weather Forecasts (ECMWF). The current iteration of this model was implemented shortly before our experiments, on July 23rd, 2023.




Figure 4.5: Fraction of area above the Chilbolton observatory occupied by clouds as modeled by the ECMWF IFS forecast [38].

So, despite the high measured humidity, the sky was quite clear at the time of launch, with at most some partial condensation around 3000 m, much higher than our sondes reached. The temperature plot shows no particular irregularities, with an overall standard linear behavior.

4.1.2 Dispersion results

Both launches were planned to be performed with 10 radiosondes. Unfortunately, as mentioned above, during in-field experiments the teams’ plans are bound to clash with the plans of the environment, as careful as one may be. Indeed, at the

moment of launch as throughout the whole day, gusts of wind were blowing at ground level. While not especially strong (as shown by the plots above), these were strong enough as to disrupt the departure of a few sondes. Out of the 10 released, 3 did not make it very far: two were pushed to the ground by wind at the time of release, and one crashed into a line of trees some hundreds of meters away from the observatory. This left us with 7 active sondes, meaning 21 pairs to form our two-particle dispersion ensemble. Clearly this is an intermediate value: it is enough to get some significant averaging, but it also is not the tens of thousands of pairs that can be tracked in numerical simulation. This leads us to the next point: launch 1 took place at 13:52, and the sonde that was followed the longest was sonde 7, that stopped responding at 15:15, so almost one and a half hours past launch time. Most stopped responding after a little more then one hour, but sonde 5 stopped the earliest at 14:36. This means that a choice needs to be made: either we discard sonde 5 from the analysis to keep the longer time frame while reducing our ensemble size, or we include it at the price of some minutes. At this ensemble size, though, discarding it would rob us of 6 pairs out of 21, meaning almost 30% of the total. It was then decided to keep it, and the resulting time frame is of 41:30 minutes or 2490 seconds . One additional important remark is that at the very beginning of the dataset very scattered values of the position are found, most likely due to the initialization of the GNSS sensor, and the analyses is therefore started a few minutes after the launch itself. The initial value for the mean distance is thus of 322 m (which we can compare to the maximum relative distance we measured at the end of the flight, 2819 m).

The next point I want to discuss is in fact the dataset itself. Since this thesis includes both results for numerical simulation and for experiments, it should be remarked that the datasets emerging from these two types of investigation differ greatly. While often massive and somewhat cumbersome, numerical simulation datasets are extremely tidy: a fixed number of time instances are available, at which all flow and particle quantities are equally available. Experimental datasets, instead, are all but tidy. As said in Chapter 3, each sonde communicates with the receiver through packets. The current receivers are all single-channel, meaning they can only receive one packet at a time. This in and of itself introduces a source of uncertainty when considering pair-related quantities: in the raw dataset the quantities are measured at slightly different times in between sondes. In addition, many packets just go missing, which can happen for a number of reasons (refer to Fig. 3.8). This means that some method needs to be put in place to clean the dataset and have synchronous values that can be compared. In our case this consists of linear interpolation: from the raw dataset, consisting single-sonde data packets all at different times, we obtain a dataset that has one set of values for each sonde at 10 second intervals.

We can finally start illustrating the results. As a first visualization, we can take

a look at the separation distance distribution Q as the one shown in Chapter 1 in Fig. 1.2. This is shown in Fig. 4.6, with time evolution at intervals of 6:40 min. For this visualization $h = 100\text{m}$ was selected as the most representative value, but with the current sample size it is obvious that the figure is quite crude. A better visualization can be once again recovered by means of interpolation, combined with color mapping. These are illustrated in Fig. 4.7 for different h values. Here the dispersion pattern can clearly be observed, and the influence of the choice of h is also much clearer: it should be chosen as small as possible, but large enough for each ℓ interval to contain a significant number of particles (which again, in our case is only appreciable thanks to interpolation).




Figure 4.6: Distance-neighbor graph time evolution for the experimental dataset of launch 1. Not normalized. Size of intervals $h = 100\text{m}$ (refer to Chapter 1).

Coming to the more quantitative results, as is the praxis when looking for a power-law behavior, we can plot $\langle \ell^2 \rangle$ on a logarithmic plot against time and expect to see a straight line, whose angular coefficient will indicate the exponent of the power law and the y -axis intercept the multiplicative constant. This is shown in Fig. 4.8. It is clear that two different regimes appear, the first one featuring some unexpected behavior. Not only does the average separation seem to remain almost constant for quite some time, at a certain point it even seems to diminish by a considerable amount. At this stage it is difficult to characterise this spurious behavior. As can be seen in the map in Fig. 4.1 the sondes, while spreading apart, still remain trapped in some kind of coherent flow structure that leads them all in the same direction virtually throughout the flight. Thus, it is not




Figure 4.7: Interpolated distance-neighbor graph color-map visualizations. The vertical axis indicates the distance intervals, the horizontal axis is time, and the color indicates the number of sondes as given by interpolation.

impossible that before spreading apart enough to start being separated by larger vortices they would be advected in such a way as to lower their average distance. In particular, two specific sondes could end up very close and lower the average significantly. At the same time, a physical explanation is not the only possible one. At the moment of launch the sondes started changing position rather abruptly, as described in the previous section, due to some wind being present at ground level. This might have made it necessary for the GNSS sensors to readjust for some time, influencing the dataset. Indeed the behavior regularises at about the 400 seconds mark, meaning after about 7 minutes, which is not an unreasonable time for the GNSS to experience problems. Based solely on a single launch it is not possible to say for certain what the real cause for this behavior is.

Regardless, we can just identify this as a different regime, produced by either physical or technical causes, and focus on the next one, which instead presents itself as an approximately straight line. A plot of the isolated second regime is




Figure 4.8: Logarithmic plot of the mean square separation distance in time, launch 1. 21 sonde pairs.

shown in Fig. 4.9. While the power law behavior is evident, a second noisy feature can be observed at the center of the plot, which is once again hard to characterise. In this case, though, a technical issue seems more likely, as the exact same regime as before the noisy drop is recovered right after it.

At this point the most basic method forward would be linear fitting on the logarithmic plot. However this procedure is somewhat outdated, as many non-linear fitting algorithms exist that yield much better results. I used a MatLab function performing non-linear regression by means of an iterative least squares method. The results are shown in picture Fig. 4.10. The emerging power law, with the fitted values, looks like

$$\langle \ell^2 \rangle = (5.21 \cdot 10^{-5}) t^{3.2819} \quad (4.1)$$

Let us now elaborate on it. First of all we would like to compare this to the R-O law from Chapter 1, which is $\langle \ell^2 \rangle = g\bar{\epsilon}t^3$. Clearly our exponent is not very far from the R-O one. The coefficient is slightly harder to analyse, since we need information about the average dissipation in the air flow. One point to stress once again here is that for the R-O law the context is the K41 theory, which is thoroughly based on the assumption of stationary, homogeneous, isotropic turbulence. As is evident from the wind speed plots above, this is very far from our conditions. As it relates to the issue of the coefficient specifically, in HIT $\bar{\epsilon}$ is a global quantity that characterises the whole flow, because despite the possible turbulent fluctuations in ϵ




Figure 4.9: Logarithmic plot of the mean square separation distance in time, launch 1, power-law regime only. Cut-off time for the first, spurious regime set at 470 seconds to eliminate all sources of uncertainty. 21 sonde pairs.

(i.e. the instantaneous energy dissipation, not averaged) the fluctuations themselves are homogeneous in the domain. Conversely, being related to the kinetic energy, ϵ will vary greatly during the sondes' flight, not only in the form of turbulent fluctuations but in terms of overall mean value as well. $\bar{\epsilon}$, then, carries a lot less meaning than in the HIT case, and the comparison is bound to be imperfect. In addition to all this, $\bar{\epsilon}$ is not an easy quantity to calculate in-field. We can estimate it by computing the energy spectrum from the sondes' velocities and inverting the K41 relation:

$$\bar{\epsilon} = \left(\frac{k^{5/3} E(k)}{C_1} \right) \quad (4.2)$$

For this launch we get a value of order 10^{-2} . The Richardson constant g , as explained in Chapter 1, is estimated to be of order 10^{-1} . Even accounting for the relatively low accuracy on the $\bar{\epsilon}$ estimation, we are still two orders of magnitude away from the factor we found by fitting. It is then evident that some other mechanism is at play, as could be expected from the extensive discussion about the limitations of K41.




Figure 4.10: Power-law fitting of the mean separation distance in time. Exponent, coefficient and Relative Root Mean Square Error are reported under the legend.

4.2 Experimental results: Launch 2

Cluster Launch 2, July 6, 2023, time interval: 13:30 - 15:28, joint dataset




Figure 4.11: Map of the July 6th launch at Chilbolton observatory with cluster trajectories.

Launch 2 took place the day after the first one, on July 6th, 2023. Launch time was 13:30: we wanted to keep the time of day as consistent as possible in between the two launches, but forecasting displayed a threat of bad weather later during the day, and since we were leaving on the 7th we could not afford to miss the launch. This time we directly adopted the moving secondary station strategy, but with an improvement: rather than relying on wind information from the latest Vaisala launch for the direction of the probes, I manually checked their coordinates as they came in through the user interface of the receiver and plugged it into my smartphone's maps application. This gave us real-time positions to work with and we were able to more accurately pursue them. Indeed, the range extended even further than in the first launch: the furthest sondes were followed for more than 53 km.

Let us once again take a look at what the atmospheric surroundings looked like during our launch, and then inspect the dispersion results.




Figure 4.12: Spreading of the sondes in the three directions North, East and up. Chilbolton, Launch 2.

4.2.1 Environmental characterisation

We will use the same plots as in the corresponding section for launch 1. We start with the wind profiling radar measurements in Fig. 4.13. By inspecting the top plot we can see that a spot of strong updraft can be identified at the low altitudes. However, this did not seem to particularly influence the sondes, as we actually measured a slower ascent with respect to the first launch. Wind intensity was slightly lower than during the first launch, whereas the direction this time was very different, at almost a 90° angle.

Fig. 4.14 shows measurements taken by a Vaisala sonde launch that was performed at 11 am. Unfortunately July 6th was not an IOP, so measurements closer to our launch were unavailable. With respect to the first launch a point of stronger thermal inversion can be identified, but it is at an altitude that our sondes did not reach, at around 2300 m. In the altitude range that we did reach the temperature profile is quite similar to the previous day, which is good for consistency with respect to stratification effects. The humidity profile, though rising more slowly, holds consistently high values in the range just above the 1500 m mark up to about 2250 m. Comparing this with the modelled cloud fraction in Fig. 4.15 we can confirm that in this range clouds were present, so humidity was likely higher than 100%. It is possible that our highest flying sonde entered the bottom portion of a cloud towards the end of its flight, but for the majority of the measurement they were all in clear air.

4.2.2 Dispersion results

We can now illustrate the results of the second launch. This time the launch was comprised of 12 sondes at the start, but two stopped shortly after their release, so




Figure 4.13: Wind profiling by the Chilbolton radar instrumentation on July 6th (Launch 2). Launch time was $\sim 13:45$. Top: color map indicates updraft. Bottom: color map indicates horizontal wind speed, L shapes indicate wind direction in compass coordinates, i.e. upwards being north and right being east. [39]

the analyses is carried out on 10. This gives 45 pairs, meaning more than double the pairs of the previous launch. The period of observation was also longer: as per the map the observation lasted almost two hours. Like for the previous launch, though, this is not true for all sondes, and the viable time frame for dispersion analysis had to be shrunk accordingly. The resulting interval was of 1:23:30 or 5010 seconds, again more than double that of the previous launch. In this second case as well the first part of the position dataset was completely scattered, and thus the analyses is carried out starting a few minutes after launch, so that the initial average distance is 278 m.

As above, we can start by visualizing Q with the slightly higher number of sondes. The plots in Figg. 4.16 and 4.17 show the standard plot and the color map, respectively. The biggest difference can be seen in the line plot, that at 10 sondes

06-July-2023 11:00 Radiosonde launch from Chilbolton




Figure 4.14: Measurements from the Vaisala sonde launched on July 6th at 11:00. Launch time was 13:30, but no closer launches were available due to July 6th not being included in an IOP.




Figure 4.15: Fraction of area above the Chilbolton observatory occupied by clouds as modeled by the ECMWF IFS forecast [39].

is starting to appear more regular, although we are far from the smoothness found in the simulations.




Figure 4.16: Distance-neighbor graph time evolution for the experimental dataset of launch 2. Not normalized. Size of intervals $h = 100\text{m}$ (refer to Chapter 1).

At this point we can take a look at the logarithmic plot, once again expecting a straight line. This is where surprise strikes: Fig 4.18 looks nothing like a straight line. Such a behavior on a log-plot is much more indicative of an exponential law. We can choose to disregard this fact and perform power law fitting all the same, and the results are not quite as terrible as the logarithmic plot would seem to indicate. They are shown in Fig. 4.19. First of all, while it is too compressed to be clearly seen in the logarithmic plot, we can identify a spike in the mean square distance that, much like the dip in the plot of launch 1, is hard to characterize but is likely due to an instrumental error. Putting that aside, we can see that the RRMSE is quite low. The exponent α , though, now departs more significantly from the K41 value and is approaching 4. At the same time, the coefficient m is getting really quite low at an order of magnitude of 10^{-7} . While, as said, we do not expect these values to exactly reflect K41, such a low coefficient makes us suspicious as to whether we would be justified in our power law fitting, in addition to the inspection of the log-plot.

We can then easily repeat the fitting procedure by using an exponential function instead, which is shown in Fig. 4.20. This results in a significantly better fit, as the lower RRMSE testifies. As is clear from the discussion of Chapter 1, though, an exponential behavior is something unexpected in this range. In spite of this, it seems




Figure 4.17: Interpolated distance-neighbor graph color-map visualizations for launch 2. The vertical axis indicates the distance intervals, the horizontal axis is time, and the color indicates the number of sondes as given by interpolation.

to be the best guess for the fitting of this launch’s results, and we therefore have to take into account the possibility that this exponential law might capture some features of the flow in which the sondes were embedded. Obviously the potential reasons for this behavior are numerous, and one single experiment displaying it is not enough to say anything conclusive. Still, we can try get a better idea of what could be happening, or at least of method of figuring out some possible avenues for investigation.

As the reader may remember, while in the inertial range we expect the mean square separation to follow a power law, there is another range in which an exponential behavior is predicted: the dissipation range. It would be clearly wrong to imagine that our measurements would fall in the dissipation range, but there is another possibility that should be tested in case such measurements were confirmed by further experiments. As opposed to distinguishing between dissipation range and inertial range we can draw the distinction between smooth and non-smooth




Figure 4.18: Logarithmic plot of the mean square separation distance in time, launch 2. 45 sonde pairs. A straight line would be expected, since we expect a power-law behavior, but instead we find a more exponential-like curve.

velocity fields (this is the version adopted in the already mentioned paper by Falkovich *et al.* [13]). Given the scaling of velocity differences with the distance as

$$|u(\vec{R}_1) - u(\vec{R}_2)| \sim |\vec{R}_1 - \vec{R}_2|^\delta \quad (4.3)$$

where \vec{R}_1 and \vec{R}_2 are two points in space, a smooth velocity field corresponds to the case $\delta = 1$, whereas a non-smooth velocity fields corresponds to all values $\delta < 1$. It is then true that the dissipation range of scales corresponds to a locally smooth velocity field, whereas the inertial range is characterised by non-smoothness, but the focus on this characteristic is more general. It could then be possible that the flow that advected the sondes during their flight might, at some level, display smooth characteristics. This, in general, is quite hard to check, as parallel Eulerian measurements are needed, and they should be performed at various scales. Indeed, it is clear that not all scales will be smooth, as Reynolds numbers in the atmosphere are incompatible with this hypothesis. If, however, some sort of “large-scale smoothness” could be detected, this would remain a possible avenue to explain the observed behavior. As pertains to the coefficient and exponent, C would be just the mean square separation at the start of the observation, and this is coherent with the actual value $278^2 = 7.72 \cdot 10^4$, where the error is simply due to the fitting.




Figure 4.19: Power-law fitting of the mean separation distance in time for Chilbolton Launch 2. Exponent, coefficient and Relative Root Mean Square Error are reported under the legend.

The coefficient at exponent, γ , obviously has units of inverse time and its reciprocal, which we can call τ_γ , is a time scale that should be characteristic of the process. In our case its value is $\tau_\gamma = 1055$ s, or 17.5 minutes.

We will now have a brief discussion of the results before moving to the numerical studies.

4.2.3 Discussion of the results

As mentioned at the beginning of the chapter, these were respectively the second and third free-flight sonde launches. Because of this it should be restated that it is early to say anything conclusive about the results. Despite this, we can still engage in some analysis to at least form some expectations about the possible outcome of future experiments. First of all, a comment on the numerical side is in order: since, as said, the analysis starts when the sondes are still some distance apart, we might have expected a fit of type $\langle \ell^2 \rangle = mt^\alpha + R$, with R the initial separation, to be more accurate. However, due the fact that the data spans several orders of magnitude, the addition of this third parameter introduces convergence issues in the regression algorithm. It was then decided to leave it out, as the predicted parameters produce




Figure 4.20: Exponential fitting of the mean separation distance in time for Chilbolton Launch 2. Exponent coefficient, multiplicative coefficient and Relative Root Mean Square Error are reported under the legend.

a rather satisfying agreement of the power law with the data. Secondly, in light of the previous paragraph, it should be noted that the two observed behaviors, meaning the power law from the first launch and exponential from the second, do not necessarily contradict each other. Indeed it is possible for the speculated smooth characteristics of the flow to only have been present in the second launch. More in general, even if the real motivation differed from the smoothness argument, it is possible that it influenced the second launch differently from the first one. In this regard, a great deal of information is needed to adequately characterize the flow. To a certain extent this will be possible to carry out thanks to the access we were granted to the Jasmin database, where a large number of the measurements performed during the Wescon campaign is stored. It is not fully clear what type of characterisation is appropriate, though, and more research is needed to understand what causes might underlie the observations. This, of course, is coupled to the need for more, larger and longer similar observations. The COMPLETE project is continuing development, and the next launches are already being scheduled. At the present scale, the Chilbolton launches already pushed the limits of our capabilities, greatly surpassing our expectations as far as observation time and distance are concerned. To upscale the next experiments there will be a need for more receiver

stations, with better tools to ensure the possibility of following the sondes as long as they keep transmitting. This links to the next necessary development: smaller sondes, which are already in the prototyping stage, and improved energetic efficiency. Indeed, it is estimated that in the current configurations the sondes are able to transmit data for a little more than 2 hours. As laid out in this chapter the second launch was already near this limit, and it is possible that even if we had been able to remain in the sondes' transmission range for longer the batteries would have simply run out. We will go into slightly deeper detail regarding future prospectives in the conclusion.

One last comment, more theoretical in nature, concerns the assumptions that the performed fitting implies. Indeed, when fitting to both power law and exponential, the basic assumption is that there is a single factor in the formula for $\langle \ell^2 \rangle$ that is time dependent. This, in general, may not be true. Let us take, as an example, the usual case of homogeneous isotropic turbulence, but let us relax the hypothesis of stationarity. In the absence of forcing, the kinetic energy will decay, and consequently so will $\bar{\varepsilon}(t)$, the usual control parameter, where now we specify that the averaging is performed either as an ensemble average or in space, the two being analogous thanks to HIT hypothesis, but the value will be time dependent. For the sake of simplification, let us postulate that $\langle \ell^2 \rangle$ still only depends on $\bar{\varepsilon}$ and t , meaning no new mechanism is introduced by the decay. At each time instant, the infinitesimal increment in $\langle \ell^2 \rangle$ will depend on the current value of $\bar{\varepsilon}(t)$, which would significantly alter the overall behavior even if an explicit t^3 were to be kept. This kind of behavior would not be captured by our fitting. This means that, if we are to accept the laws we have found by fitting the data, the only quantities that should be considered when trying to define the various factors in terms of physical meaning are quantities that remain constant in time. Additionally, because we are not referring to local, Eulerian measurements but rather to Lagrangian measurements that explore a large domain, the quantities should also be global, meaning constant in space, or they would inherit an implicit dependence on time. If this simplification is not found to be acceptable, the process would need to be described in terms of a differential equation keeping track of the time variation of all relevant quantities, which would likely result in a solution deprived of self-similarity.

This officially concludes our presentation of the experimental results as far as relative dispersion is concerned. We can now move on to describing a similar analysis carried out in the cloud numerical simulations, and its limitations.

4.3 Numerical results: preparatory analysis

The second set of results on relative dispersion actually constitutes a starting point for future studies to be carried out on simulations that are currently being set

up, both in terms of code updating and of resource gathering. Indeed, the ones that are currently at our disposal are not suited for this kind of analysis, for two separate reasons.

The first one relates to the intensity of the turbulence found in them. As was stated many times, the turbulent regime we would like to access is the inertial range, which means that we must have a broad spectrum of resolved scales in our simulation. The parameter that allows us to measure how many scales we are able to resolve is the Taylor scale Reynolds number, Re_λ , which is obviously related to the Taylor length scale λ . This length scale indicates the scale at which the transition from the energy-containing scales to the inertial range happens. If it is too small, and by extension if so is Re_λ , the inertial range is not adequately resolved, and the energetic scales basically give immediate way to the dissipation range. Conversely, the higher Re_λ , the more developed the inertial range. As touched upon shortly in first chapter, obtaining high Reynolds numbers is one of the main challenges currently limiting computational fluid dynamics. Raising it by very little greatly increases the computational cost of the simulations, and the highest currently achievable Re_λ is of the order of ~ 1000 . In their already cited paper, Malik and Hussein [17] argue that this is the bear minimum value at which one can hope to find real R-O scaling. At the same time, many of the seemingly promising results listed in Chapter 1 were obtained for lower Re_λ , but the point remains that there is a need for higher and higher values. To the already hard task of simulating high Re_λ flows, our simulations add many intrinsically heavy task: not only are we simulating the flow, but also an order of 10^7 particles with their growth dynamics, along with two scalar fields. This makes it all the more problematic for us to reach the desired range, and indeed in the past simulations Re_λ was set at values around 50, not high enough to resolve the inertial range. The chief strategy to increase this value is to augment the resolution of the computational domain: this allows for increasing the number of scales that are possible to resolve, thus allowing for the necessary increase in the kinetic energy .

On a second front, the current simulations are also somewhat too short to adequately perform dispersion analysis, and this also relates to the results in the next chapter. As far as physical time goes, the longest ones span about 2.7 s. The preferred indication in computational fluid dynamics though is the adimensional quantity t/τ_0 , where τ_0 is the large eddy turnover time, meaning the time scale related to the large scales of the turbulence. Our simulations span about 8 eddy turnover times. By comparison, we can estimate the eddy turnover time of our experiments to be in the vicinity of 50 s, meaning that the measurements span a range of 40 to 80 eddy turnover times. This has lead us to start organising new runs that will have the purpose of continuing these previous simulations, by using the current final configurations as initial conditions.

While these further simulations are under way, we can start looking at the current

ones under the light of relative dispersion, to form a better idea of how the analysis can be carried out. The reader can refer to Chapter 2 for a detailed description of their various aspects. Briefly stated, we compare the same quantities for the cases of monodisperse and polydisperse initial droplet population distribution, and for the cases of stable and unstable stratification. A monodisperse and stable-stratified simulation if as of now not available, and is in fact in the to-dos for the upcoming ones. In the first section we will look at some visualisations that are analogous to those shown for the experiments, and in the second one we will check first-hand what the outcome of the low Re_λ value is on the $\langle \ell^2 \rangle$ behavior.

4.3.1 Visualizations

To start, we highlight that the differences we would like to detect are in the behavior of the mean separation in the cloud bulk, modeled as an HIT region, and its behavior in the mixing region, localized in the middle of the domain and expanding in time. For this reason, our analysis starts by extracting a subset of the particles localised in a small volume in the desired regions. This is done by specifying the size of the volume, which will be a cube, and its location. The location of all particles at the starting time is then parsed, and the ones whose position is included in our volume are added to a hashset which will enable us to follow them throughout their motion. This obviously implies that the larger the initial volume, the larger the number of droplets in the hashset. Fig 4.21 shows the “marked” particles at their initial position and their evolution in time. We note that there is a stark difference in the behavior of the monodisperse population and that of the polydisperse one: in the polydisperse case the larger droplets sink considerably towards the bottom, due to their larger weight. This is useful information for assigning priority to the future simulations: while the monodisperse case may be more idealised, it might be more suited for an analysis of dispersion that highlights the effects of turbulence over those of buoyancy.

We can then easily compute relative distances and track them over time, which gives us $Q(\ell, t)$. As said in the first chapter and repeated in passing in this one, the choice of h (the length that makes up the various “bins” of Q , since it is a discretised distribution) plays a role in how significant the distribution will be: choose it too small and most bins will contain an equally low number of particles, thus flattening the distribution; choose it too big and the distribution will become very jagged, with sharp features that make it harder to interpret. Some trial and error allowed us to find that, as long as the starting volume is small, keeping h of the same order as its linear dimension is a good compromise. The Q plots in Fig. 4.22 start giving us an idea of the problems listed above: by visual inspection we can see that they can hardly be said to be expanding at an accelerated rate, as the R-O law would predict. At the same time some differences in between




Figure 4.21: Plots of the droplet positions in time in the cloud top boundary simulations. Plots given for the currently available combinations of initial droplet populations and thermal stratification: monodisperse unstable (left), polydisperse unstable (middle), polydisperse stable (right).

them can be detected, which perk up our interest in pursuing the matter further. Most prominently we see that the initially monodisperse populations expand at a slower rate, which is certainly due to the presence of very small droplets in the polydisperse cases which more readily follow the fluid flow thanks to the lower inertia. Secondly, although the effect is not extreme, the distributions relative to the mixing region do seem to spread slightly faster. This can also be observed in the previous Fig. 4.21, although there it is even more difficult to quantify.

4.3.2 Mean square separation

This last observation is confirmed by the computation of $\langle \ell^2 \rangle$, plotted in Fig. 4.23. As we can see, not much trace of turbulent contribution can be detected, and the behavior appears mostly linear. As said above, for now it is also computed for a small time window, in addition to the small Re_λ value, so that even if some superlinear behavior were present we would only be able to plot it for a short duration and thus possibly not identify it as such.

What these plots allow us to confirm though is that the mixing region seems to keep up with the cloud bulk in terms of transport, which is at odds with the fact that the kinetic energy contained in it is lower. In the monodisperse case the line relative to the mixing very clearly takes over after just a couple eddy turnover times, and in the stable polydisperse case we can start to see the same thing happening towards the end of the simulation. Another interesting observation is that, as we might expect on theoretical grounds, the stability configuration of the domain impacts transport: if we check the final separation values, in the unstable polydisperse case they are upwards of 180 cm^2 whereas in the stable case they are shy of 150 cm^2 .

This brings us to the end of the discussion of our results in the topic of relative dispersion. As I hope I made clear there is much room for future development, and some of the possible routes towards it I will outline in the Conclusions chapter. We can now move on to the second main topic, that will see a much more prevalent role of the numerical simulations.




Figure 4.22: Q plots time evolution for the available combinations of stability conditions and initial droplet populations. Left plots are for droplets starting in the cloud bulk, right plots are for droplets starting in the mixing layer at the cloud border.





Figure 4.23: Mean square separation for the available combinations of stability conditions and initial droplet populations. Plots for droplets starting both in the cloud bulk (blue) and in the mixing region (red).

Chapter 5

Results on Lagrangian correlations



In this last chapter we will analyse the behavior of Lagrangian correlations of a variety of quantities in both the numerical and experimental datasets that have been described. As outlined in Chapter 1, a far thinner framework is available for comparison in this case. This implies an obvious difficulty in the interpretation of the results, since there is no unique result comparable to the Richardson-Obukhov law of relative dispersion that one might try to stack against their own findings. If this is true for velocity correlations, it is even truer for the Lagrangian correlations of other quantities. Indeed, suppose that the Lagrangian trajectory of a particle being advected by some fluid flow can be characterised by properties $\{f_1, \dots, f_N\}$, some of which are properties of the particle (such as velocities and mass) and some

of which may be local properties of the fluid (such as humidity, or any other scalar concentration). Then, departing from most of the literature, which only focuses on velocities as reported in Section 1.4, we can define a more general Lagrangian correlation coefficient as:

$$R_L^{(f_i f_j)}(\vec{x}, \tau) = \frac{\langle f_i(\vec{x}, t_0) f_j(\vec{x} + \vec{r}(t_0 + \tau), t_0 + \tau) \rangle - \langle f_i(\vec{x}, t_0) \rangle \langle f_j(\vec{x} + \vec{r}(t_0 + \tau), t_0 + \tau) \rangle}{\sqrt{\langle f_i^2(\vec{x}, t_0) \rangle - \langle f_i(\vec{x}, t_0) \rangle^2} \sqrt{\langle f_j^2(\vec{x} + \vec{r}(t_0 + \tau), t_0 + \tau) \rangle - \langle f_j(\vec{x} + \vec{r}(t_0 + \tau), t_0 + \tau) \rangle^2}} \quad (5.1)$$

This is important to the cases we are analysing because we are not only interested in the turbulent flow, fully defined in terms of velocities, but also in its interplay with the phenomenology of the atmosphere at large and of clouds specifically, which are clearly dependent on many other quantities.

Clearly, while the correlation coefficients are a quantitative measure, their behavior in time can hardly serve as a parameter, both in terms of modelling and in terms of comparison. For this reason, the Lagrangian time scale was introduced (refer again to Section 1.4). The definition, though, is given with the assumption that the correlations decay to 0 rapidly enough, so that their integral between 0 and infinity will be convergent. As we will see, we will not find this to always be the case. Additionally we have to deal with the fact that our simulations are not long enough to make a some over the while time domain a good approximation of an integral to infinity. It is therefore more convenient, to allow for a proper comparison in between the various cases, to define a Lagrangian time scale that is cut short up to some n -th eddy turnover time τ_0 :

$$T_{f_i f_j}^{(n)}(\vec{x}) = \int_0^{n\tau_0} R_L^{(f_i f_j)}(\vec{x}, \tau) d\tau \quad (5.2)$$

where once again we highlight that in some of the cases relevant to us the space dependence will remain, specifically to contrast the cloud bulk and the mixing region. The n value that was selected to illustrate the present results is 4, as it is the largest number that is common to all simulations.

We will start by looking at the numerical simulations. Contrarily to the case of relative dispersion, these are a much more suitable environment with respect to the radiosonde launches in their current form. Indeed, a great importance is put on the initial value of f_i , and a noisy dataset can make this value much less significant than in the fully deterministic case of the simulations. We will see the effect of this in the section dedicated to the experiments, and we will discuss one possible avenue to circumvent the problem.

5.1 Numerical results

In the numerical simulations, which were described in detail in Chapter 2, the goal will be to compare the correlations inside the cloud bulk, modeled as a region of homogeneous and isotropic turbulence, with the correlations inside the mixing layer formed at the interface with the clear air, which instead displays anisotropy along the vertical axis. Results will be provided for three different kinds of cloud-clear air border region simulation, which consist of the different matching of two initial conditions: droplet population and thermal stability. The initial droplet population has been outlined in detail in Chapter 2, but we quickly recall that the two options consist in a monodisperse population, in which all droplets in the cloud portion of the domain are initialized to a radius of $15 \mu m$, and a polydisperse population, in which initial radii span the range $0.6 - 30 \mu m$. The stability instead relates to the direction of the temperature gradient, and while it is clearly related to the stability defined in the context of parcel theory from Section 2.1.1 it is somewhat more general: a fluid region with thermal inhomogeneities is stable when the thermal gradient points upwards, meaning that the colder and denser region is below the hotter one so that buoyancy effects will not induce any convection, and unstable in the opposite case. At the present moment simulations are available for the three cases of polydisperse and stable, polydisperse and unstable and monodisperse unstable. The monodisperse stable case is missing, and is already included in a simulation project currently submitted for approval at CINECA, the facility that has provided computational resources for the ones we do have. We should therefore be able to report on it shortly, but it will have to be excluded from this thesis. Each configuration is interesting for its own reasons, in addition to the comparisons we can draw in between them. The order of presentation is thus mostly arbitrary.

At this point a brief comment should be made on the datasets used for the computation of the results presented in the upcoming subsections. In the case of the preliminary dispersion studies droplets were extracted (tantamount to saying “marked”) from the total amount present in the simulations based on them starting out at the first time instance inside a selected volume. This was done for one volume at the center of the cloud region and one volume at the center of the mixing region, so that the evolution of the two volumes could be compared. As far as Lagrangian correlations studies are concerned, though, this setup is flawed: because they start out very close to one another there is an increased probability of the droplets behaving similarly in regards to the observed quantities, thus introducing spurious correlations that do not get filtered out by the ensemble averaging. Because of this I modified the extraction procedure so as to extract five volumes instead of one, distributed according to Figure 5.1 along the x and y axes (which are the isotropic axes in the simulations). These volumes were then used to mark droplets starting out at z (the axis along which the anisotropy is imposed) values corresponding

to the center of the cloud region and to the center of the mixing region. The volumes are cubical, with a side of 2 cm. This resulted in numbers of particles in between 2000 and 3000 in each dataset, meaning the five volumes. The exact number obviously depends on the simulation and z position, since the initial particle position is randomized homogeneously in the cloud region. Additionally, it should be kept in mind that some droplets completely evaporate before the end of the simulated time. To ensure consistency in the ensemble size only particles present at the final time were kept from the extracted datasets.




Figure 5.1: Distribution along homogeneous axes of the particles used for Lagrangian correlations computation. Five cubic volumes with a side of 2 cm were selected, and then droplets starting out inside these volumes were selected for inhomogeneous axis z values corresponding to the center of the cloud region and the center of the mixing region.

A second type of ensemble averaging was also put in place: because the behavior of the correlations inside a single simulation depends on the underlying velocity field, simulations with the same population and stability conditions but different initial velocity fields were used for averaging of the correlation coefficients, thus reducing the effect of the specific simulation on the overall values. Specifically, four simulations of the monodisperse unstable case were available, and three of the polydisperse unstable case. Unfortunately the polydisperse stable case was only used in one simulation, and the difference in behavior with respect to the other two cases will be noticeable. Ideally even more simulations could be used, but in our case of decaying turbulence it is especially expensive: because the energy does not remain constant, we cannot simply take an instant after the initial one and

used droplets marked at that time for averaging, as you could potentially do in the case of HIT, as one we would be comparing particles starting out in very different conditions.

In the numerical discussion we will include results on correlations for five quantities: velocity components, droplet radius and droplet growth rate. These are chosen because they are properties of the particles themselves, and as such are all saved in the same dataset when extracting the chosen droplets from the simulations. A future step will be to also extract temperature and water vapor density information along the Lagrangian trajectories, but this requires an algorithm that will fetch the quantities from their own binary file based on the Lagrangian trajectories. This is, once again, already in the works, but can unfortunately not make it into the present discussion.

5.1.1 Monodisperse distribution, unstable stratification

We will start out with the initial monodisperse droplet distribution case, with unstable thermal stratification conditions. Figure 5.2 is a visualisation of the trajectories of 30 randomly extracted particles from one of the datasets used in the averaging. Owing to the comparatively large size of the droplets we can see that gravity has a rather strong effect on the mean displacement, and most droplets tend to sink towards the bottom of the domain. Figure 5.3 is the first quantitative result, namely the velocity component autocorrelations. We can see that, despite the ensemble averaging procedures, some seemingly spurious behaviors remain, most evident in the autocorrelation of the velocity along y . Indeed we would expect the velocities along x and y to behave almost exactly in the same way, since there is no anisotropy that differentiates between the two. As said above, this is likely due to the initial conditions used for the simulations (one should remember that even homogeneous, isotropic turbulence is only such on average, since locally and instantaneously the flow structure may well be anisotropic). With these caveats in mind, thus allowing for some leeway, we can see that the decay appears approximately exponential. We can observe that there is an inversion in the behavior of the autocorrelation of w , the velocity along z : it is the fastest decaying function in the cloud region, whereas it is the slowest one in the mixing up until about 2 eddy turnover times, a time by which all three components have decayed to low values. Indeed, $T_{ww}^{(4)}$ is the one that is overall most affected by moving from the cloud region to the mixing layer: its variation is more than then twice the variation in either $T_{uu}^{(4)}$ (varying by about 0.7) or $T_{vv}^{(4)}$ (varying of just about 0.25). Interestingly, the changes in $T_{uu}^{(4)}$ and $T_{vv}^{(4)}$ go in opposite directions, with the first one increasing and the second one decreasing. Since there is nothing intrinsic to the flow that should cause this difference, this is a further hint of the dependence of the correlations on the initial conditions. An important note is that




Figure 5.2: Visualization of the trajectories of 30 randomly marked particles starting out in the five initial volumes in the case of monodisperse population, unstable stratification.

the behavior of the v autocorrelations is rather odd, in that at a certain point it becomes negative. While this probably amounts to one of the issues discussed above, having to do with the initial condition dependence of the correlations, it forces us to make a decision about the computation of $T_{vv}^{(4)}$: if we keep the negative values the Lagrangian time scale will actually appear to be shorter due to some negative values being added to it after about $1.2\tau_0$. For this reason it was decided to actually integrate (i.e. sum, as far as the actual computation is concerned) the absolute value of $R_L^{(vv)}$, so as to keep into account both positive and negative correlations. It is hard to implement any better solution at the moment, since the actual origin of the unusual behavior is not known. All the same is true, to a lesser degree, also for $R_L^{(ww)}$.

We can then introduce the quite novel results about the radius and growth rate




Figure 5.3: Lagrangian autocorrelations for velocity components in the monodisperse population, unstable stratification case.

autocorrelations, which are plotted in Figure 5.4. As we will learn in the next section, the monodisperse population provides a much more ideal environment as far as radius correlation computation is concerned, and partly also for the growth rate autocorrelations. Their behavior looks quite regular, and the differences between the cloud region and the mixing are plain to see: in the cloud we observe what looks like a simple exponential decay for both quantities, analogous to the velocity autocorrelations; in the mixing, instead, both functions decay much more slowly, and would likely require a rather large amount of eddy turnover times to approach 0, but they also display the rapid decay only after a brief initial interval in which the slope is much more gentle. This means that the mixing does not only affect the speed of the decay but also the functional shape of the correlations. The slower descent is obviously reflected in the values of $T_{RR}^{(4)}$ and $T_{GRGR}^{(4)}$, both approximately doubling. When looking at this result one aspect should be kept in mind: these correlations carry no information about the coalescence process, since the growth rate is only computed from the Köhler theory for the condensation/evaporation process.

5.1.2 Polydisperse distribution, unstable stratification

Moving to the polydisperse distribution, we again visualise the motion of 30 random particles in Figure 5.5. We can see that the downwards motion is less generalised, owing to the presence of smaller particles that are less affected by gravity. Figure




Figure 5.4: Lagrangian autocorrelations for radius and growth rate in the monodisperse population, unstable stratification case.

MU	Cloud	Mixing
$T_{uu}^{(4)}$	0.5298	0.6302
$T_{vv}^{(4)}$	0.6804	0.6373
$T_{ww}^{(4)}$	0.4409	0.7191

MU	Cloud	Mixing
$T_{RR}^{(4)}$	1.1061	2.1223
$T_{GRGR}^{(4)}$	0.5043	1.1669

Table 5.1: Lagrangian time scales for velocity components, radius and growth rate autocorrelations in the monodisperse population, unstable stratification case

5.6 shows the velocity component autocorrelations. The overall behavior is similar to the monodisperse case, again displaying a seemingly exponential decay. However, an effect can be easily identified that was not present before: some correlations seem to relax to values different from 0, specifically correlations of u and w in the cloud bulk and only of w in the mixing region. We will discuss this more in the dedicated section after all results have been presented, but it is an interesting feature that will need to be investigated further. By inspecting the Lagrangian Time scale values, we can see that the value of $T_{vv}^{(4)}$ decreases significantly in the mixing, due to the fact that it actually decays to zero rather than to a higher value. The increase in $T_{uu}^{(4)}$ in the mixing region is quite similar to the monodisperse case. $T_{ww}^{(4)}$ slightly decreases in the mixing, at odds with the monodisperse case in which it grew significantly.

The autocorrelations for radius and growth rate are, in this case, quite a bit




Figure 5.5: Visualization of the trajectories of 30 randomly marked particles starting out in the five initial volumes in the case of polydisperse population, unstable stratification.

more problematic. Indeed, we can see that the autocorrelations for the radius remain basically squeezed very near 1 both in the mixing and the cloud, even though we know for a fact that the radii are changing, and especially for the smaller droplets they change a great deal more than in the monodisperse case. The odd behavior of the correlation coefficient stems, in fact, from the large spread of the size distribution. Indeed, the current treatment of the correlations sees the various radii simply as different instances of a single random variable. The variations in time in radius of each single droplet are, then, very much smaller than the standard deviation of the size distribution itself, and this leads to the effect that radii seem to remain very correlated in time despite the fact that they may vary considerably, and certainly much more than in the monodisperse case. How to deal with this issue is not obvious. One possibility would be by binning of the radii into groups that




Figure 5.6: Lagrangian autocorrelations for velocity components in the polydisperse population, unstable stratification case.

have a similar starting value and check for an initial-radius-dependent correlation coefficient. Because the variations are small though, there are very few droplets in the dataset starting with radii that are close enough in value, and thus averaging would become problem because of sample sizes. Obviously the most direct way around this would be to increase the whole dataset size, but a rather considerable number of particles would have to be taken for this to work. More on this will come in the section dedicated to the discussion.

The autocorrelations for the growth rate, however, despite featuring a similar issue in the cloud region, shows a decay that is rather similar to the monodisperse case in the mixing, displaying the mild initial slope and the exponential decay afterwards. This stark difference is quite surprising, and will need some further analysis to be fully understood, but at a first level might be attributed to the fact that in the cloud, while fluctuations in the vapor concentration field are present, the local values always remain quite high. In the mixing, instead, the entrainment of dry air causes the droplets to come in contact with regions of very low concentration, which induces a much more rapid change in the growth rate

5.1.3 Polydisperse distribution, stable stratification

The last case we can deal with is that of polydisperse initial droplet distribution and stable thermal stratification. Because convection along the vertical axis is suppressed by the stratification effect, we expect upwards transport of droplets




Figure 5.7: Lagrangian autocorrelations for radius and growth rate in the polydisperse population, unstable stratification case.

PU	Cloud	Mixing
$T_{uu}^{(4)}$	0.7657	0.8524
$T_{vv}^{(4)}$	0.9174	0.4640
$T_{ww}^{(4)}$	0.9527	0.8887

PU	Cloud	Mixing
$T_{RR}^{(4)}$	4.0349	4.0364
$T_{GRGR}^{(4)}$	3.9392	1.2531

Table 5.2: Lagrangian time scales for velocity components, radius and growth rate autocorrelations in the polydisperse population, unstable stratification case

to be reduced. Indeed, in the visualisation in Figure 5.8, when compared to the unstable case, does show droplets that seem on average to be less prone to move towards the clear air (it should be stated again that, since these are just a few droplets and not the entire dataset, these visualizations should not be taken at face value and are only meant to help illustrate the domain). As said in the preamble, currently only one simulation is available for the stable configuration, and thus it was not possible to perform ensemble averaging over different initial flow fields. This is clearly reflected in the behavior of the velocity autocorrelations shown in Figure 5.9: while the general exponential trend does appear, their curves are a lot less smooth. This makes it even harder to draw conclusions about their decay, but still some features do seem relevant. By comparing the Lagrangian time scales in table 5.3 we see that the trends are the same as in the polydisperse unstable case, with only $T_{uu}^{(4)}$ increasing and the other two decreasing, but the amount by which




Figure 5.8: Visualization of the trajectories of 30 randomly marked particles starting out in the five initial volumes in the case of polydisperse population, stable stratification.

they vary is different: specifically, the decrease in $T_{ww}^{(4)}$ is much stronger in the stable case. This can again be interpreted by supposing that velocity components w are inhibited by the stratification, causing faster decorrelation.

The correlations for radius and growth rate display an overall behavior that is analogous to the stable case, where the discussion about the problematic nature of the polydisperse distribution is obviously still valid. The surprising feature is that in the mixing the growth rate autocorrelations decay even faster than in the unstable case. This is unexpected, in that the stable configuration should reduce the entrainment of dry air that was speculated to cause the rapid decorrelation.

This wraps up the presentation of the numerical results. The next section is dedicated to some discussion and possible interpretations, while the consideration of possible avenues for development will be given in the Conclusions chapter.




Figure 5.9: Lagrangian autocorrelations for velocity components in the polydisperse population, stable stratification case.




Figure 5.10: Lagrangian autocorrelations for radius and growth rate in the polydisperse population, stable stratification case.

PS	Cloud	Mixing
$T_{uu}^{(4)}$	1.0844	1.1157
$T_{vv}^{(4)}$	0.6832	0.4340
$T_{ww}^{(4)}$	1.1469	0.7491

PS	Cloud	Mixing
$T_{RR}^{(4)}$	4.0408	4.0350
$T_{GRGR}^{(4)}$	3.9410	0.7164

Table 5.3: Lagrangian time scales for velocity components, radius and growth rate autocorrelations in the polydisperse population, stable stratification case

5.1.4 Discussion of results

As stated above, the interpretation of the results given proves much harder than in the case of relative dispersion. That is itself far from trivial, since the environments in which they were studied differ greatly from the HIT in which the theoretical framework for it is developed, but at least such a theoretical framework exists. In the case of Lagrangian correlations the same cannot be said, and terms for comparison are rather lacking. Therefore, the results must be analysed chiefly in view of what we know about our domain, rather than on some general laws that may or may not be observed.

First and foremost we should address the fact that the correlations do not go smoothly to zero, and in some cases do not go to zero at all but rather seem to have some horizontal asymptote at higher values. In stationary homogeneous turbulence we would not expect this behavior even without any ensemble averaging, since the persistent randomness of the field would prevent quantities from remaining correlated (although the decay in a single simulation may still not be smooth). In this regard the most likely culprit we can identify in our case is that the turbulence is decaying. Indeed, in decaying turbulence the small vortices are the first to disappear, and by the end of the transient only much larger ones remain, the same ones that were previously advecting the small scales. This might explain the remaining non-zero correlations: as the turbulent field becomes smoother, meaning as the small scales of the random motion disappear, some correlation remains with the initial velocities in the form of the large vortices.

Secondly, one interesting observation concerns $T_{ww}^{(4)}$. The reason to focus on this quantity, as mentioned earlier, is that we expect the anisotropy in our domain to mainly affect quantities along the same axis along which the anisotropy is found. By checking the values of $T_{ww}^{(4)}$ in the mixing region we can see that they are not too dissimilar, despite the values in the cloud for the various setups differ by a lot (0.4409 in the monodisperse unstable case versus 1.1469 in the polydisperse stable case). This suggests that the mixing region and the anisotropy contained in it may have some converging effect on this time scale. Clearly more simulations are needed to be able to confirm this fact, and it should be remembered that the length of our current simulations made it necessary for us to use $T_{f_i f_j}^{(4)}$ as a time scale, whereas this effect might become more evident by using time scales for larger values of n (also reminding ourselves that the ideal definition of T_L requires integration to infinity).

We can now move on to the experimental results, which differ greatly from the numerical ones mostly, as said, because of the current experimental sample size.

5.2 Experimental results

All necessary information about the experimental setup and the atmospheric conditions in which the launches took place have already been given in the corresponding sections in Chapter 4, and will therefore not be repeated here. The reader can refer to those and to Chapter 3 for all technical aspects of the data collection process. The results will be given for autocorrelations of velocity components and then for temperature, relative humidity and pressure. Because of this, the temperature, humidity and pressure measurements have not been shown when illustrating the launches in Chapter 4, as I preferred to keep only information relevant to position measurements in that chapter, and will instead be presented here along with the velocity measurements. The velocity is computed by finite differences using the position measurements, and not by means of the direct velocity measurements as obtained through the IMU, as these have been found to be less reliable. They are given in the East, North, Up reference frame. Again we will separate the two launches for discussion, although in this case differences in between the two will not be quite so easy to spot.

5.2.1 Launch 1

For a start, Figure 5.11 displays the velocities of the sondes as computed from position information. Notice that the plots only extend to 25 eddy turnover times for improved legibility, as displaying them for the whole duration of the experiment (about $50 t/\tau_0$) would compress the lines too much, and without providing much more information: the velocity fluctuations show a similar behavior throughout the flight, and possible differences would not be noticeable by the naked eye anyway. Coherently with what we know about the sondes' trajectory (see map in Figure 4.1), the velocities in the East direction display the largest mean value.

We can then compute autocorrelations for velocity components, which are shown in Figure 5.12. This clearly shows a much messier behavior than in the numerical simulations: the plots show very strong oscillations, which do seem to be reducing in amplitude towards the final time instants but are still quite far from decaying to zero. These of course are of rather difficult interpretation, and while it is certainly true that they are probably caused to a large extent by the sample size we should remember that these are still the result of some averaging and thus may potentially include some physical insight. One possible observation is that the oscillations appear to span many different wavelengths but they are not simple noise. Therefore, while the overall decay seems to escape the analyses at the current sample size, it may be possible to create a link between these oscillations and the nature of the underlying flow.

Figure 5.13 shows the measurements for temperature, humidity and pressure




Figure 5.11: Velocities of the sondes along the three directions East, North and Up in the first Launch carried out in Chilbolton.

for each sonde. This serves as a visualization to make sure that no outliers are to be found in the data. This is obviously not by chance, as outlier removal procedures are in place to clean the dataset before usage for any type of analysis, but it is important to check the regularity of the measurements to confirm that the upcoming correlations do not owe their weird shape to any measurement errors. Indeed, the autocorrelations are plotted in Figure 5.14, and they also show some strongly oscillatory behavior. Additionally, all of them abundantly explore the region of negative correlations while seeming to start rising again towards the end of the measurement period. Somewhat surprisingly the pressure curve does seem quite smooth. These negative correlations would mean that initial higher values for each of the quantities imply smaller values during the middle stages of the flight, which is not a very intuitive result (especially considering that the sondes started out at very similar measured values for all quantities concerned).

All of this makes it very hard to make any definitive statement about the nature of the Lagrangian autocorrelations computed from the measurements, and forces us to think of possible ways to extract more information from the datasets, short of




Figure 5.12: Velocity components autocorrelations from Launch 1 in Chilbolton.

repeating the experiments with more sondes, which is obviously the only fully sound method. One possibility, that should be treated with an abundance of caution but that we still report here for completeness, is to employ a different kind of correlation coefficient, much more commonly used in the field of economics. This is effectively quite a different approach and deserves a moment for explanation. To distinguish this different approach from the previous (which we underline is the most appropriate one) we will refer to it as “traslational autocorrelations”, for reasons that will be clarified shortly.

Putting aside space dependence for the moment, the autocorrelation as computed until now is a function of time. This is because the initial time is implicitly assumed to be the initial time of observation, meaning the beginning of the simulation of the experiment. The observations that make up our ensemble are the values of any of the measured quantities picked from all the sondes at the starting time and then the values from all sondes at all subsequent times. From these correlation coefficients are computed so that the information we will obtain is the correlation of the values at any time with those at the initial time. Traslational autocorrelations, instead, are a function not of time but of a time lag τ , and the ensemble that concurs to the computation of correlations for a certain value of τ are all pairs of values of a certain quantity that are a distance τ apart in time. So, for example, for $\tau = 1$ we will pair values *from a single sonde’s measurement* at times 1 and 2, 2 and 3, 3 and 4 and so on, and use these pairs for the computation of a correlation coefficient. The information that this coefficient will convey than is not about the correlation of values at different times with the initial values, but rather about




Figure 5.13: Temperature, humidity and pressure measurements from all sondes for Launch 1 in Chilbolton.

the correlation of any value at a given time with a value measured at a distance τ earlier in time. As italicized above, this means that this correlation function can be computed from the dataset coming from a single sonde, and that for a long measurements period the sample size will be large enough to provide a relatively smooth behavior. In our case this can be further enhanced thanks to averaging over the results from the various sondes, this obtaining a curve that is, up to some degree, representative of the whole experiment. This point should be stressed: while the previously discussed autocorrelations reflect how the events occurring during flight impact the starting values (eventually departing from them), these translational autocorrelations average over points distant τ occurring over the whole measurement period, and as such they do not strictly relate to the initial values but rather show the average decay of correlations over the whole flight. This means that, in a way, we are looking at a global quantity, which does not possess the “resolution” to distinguish between different moments of the experiment since what matters is the time distance between two points and not the absolute time. In this sense we say that the time lag is “translated” over the whole dataset for averaging,




Figure 5.14: Autocorrelations for temperature, relative humidity and pressure obtained from the dataset gathered during Launch 1 in Chilbolton1.

hence the name.

Results for these translational autocorrelations are reported in figures 5.15 and 5.16 for velocity components and temperature, humidity and pressure, respectively. The colored lines are the autocorrelations coming out of each sonde's measurements, whereas the thick black line is the average. As expected, the behavior appears much smoother, especially for the averaged plots. Overall we seem to observe again an exponential type of decay, although the velocity along the North direction seems to enter a linear regime (it is the slowest decaying component, which is surprising considering that the mean motion is mostly towards the east). The decorrelations for temperature, humidity and pressure instead show surprising similarity amongst themselves, with the humidity decaying slightly faster. Here, an interesting observation is that these scalars show a qualitative similarity with the decay of the radius and droplet growth rate in the simulations, showing an initial milder slope before a steeper descent. Of course the effect is minor, and we should use care in drawing hurried comparisons, but it is still ground for further inspection.




Figure 5.15: Traslational autocorrelations for velocity components for Launch 1 in Chilbolton.




Figure 5.16: Traslational autocorrelations for temperature, humidity and pressure for Launch 1 in Chilbolton.

5.2.2 Launch 2

Many of the same considerations made for Launch 1 remain true for Launch 2, so the discussion will be shorter. Unfortunately the increased number of sondes (10, against 7 in Launch 1) is not sufficient to provide cleaner values for the standard autocorrelations, so the traslational autocorrelations are given again at the end. Figure 5.17 shows the computed velocities for all sondes, and Figure 5.18 the respective autocorrelations.




Figure 5.17: Velocities of the sondes along the three directions East, North and Up in the second Launch carried out in Chilbolton.

Owing to the longer duration of the experiment we can see somewhat more clearly that the amplitude of the oscillations seems to be reducing as time goes by, although they are still rather far from zero. One other observation is that the wavelength of the largest oscillations seems to be much shorter at earlier times, and to grow during the duration of the experiment. Again, this is a qualitative statement based upon visual inspection of the figures, but it is noticeable nonetheless.

Figure 5.19 shows the temperature, humidity and pressure trends for the second




Figure 5.18: Velocity component autocorrelations from Launch 2 in Chilbolton.

Launch. Temperatures seem to rise quite considerably towards later times, and relative humidity displays a more regular behavior with respect to Launch 1. The visualizations of the measurements prove particularly useful in second case because a weird feature can be spotted in the autocorrelations for pressure, shown in Figure 5.20. There is a sharp dip in this curve that rises again after just a few eddy turnover times, without any further strong fluctuations. This figure provides an interesting comparison with Launch 1, since both pressure and temperature show radically different behaviors, whereas relative humidity seems more consistent in between the two.

Exploring traslational autocorrelations for velocities we can see that they appear quite consistent with Launch 1 (notice that they seem to decay more quickly only because the plots contain more time lags τ , but for example the approach at zero is obtained for similar lags). This is interesting if we consider that the conditions for the two launches were quite dissimilar, as were the results for relative dispersion).

This similarity is not found in the temperature, humidity and pressure correlations, which instead decay more slowly in this second launch.

With this, the presentation of results draws to an end. All that is left are some conclusive remarks, mostly regarding future outlooks to move forward with the research presented here.




Figure 5.19: Temperature, humidity and pressure measurements for all sondes for Launch 2 in Chilbolton.




Figure 5.20: Autocorrelations for temperature, humidity and pressure from Launch 2 in Chilbolton.




Figure 5.21: Traslational autocorrelations for velocity components for Launch 2 in Chilbolton.




Figure 5.22: Traslational autocorrelations for temperature, humidity and pressure for Launch 2 in Chilbolton.

Conclusions

As written a few times throughout the text, drawing conclusions is somewhat of a challenge at the present time. Indeed, many different results have been presented, innovative in nature and displaying a variety of unique features. Because of this, more work will be necessary to make more definitive statements about the observed phenomena. This is not to underplay the relevance of what has been shown so far, but rather to highlight that the many interesting observations that were made are but a peak into the largely mysterious world that is inhomogeneous turbulence, and therefore they should be used as a launchpad for future investigation. Thus, I wish to devote this final chapter to the discussion of the possible avenues that may lead to more results, or to the confirmation and explanation of the existing ones. I will do so separately for the experimental and numerical parts, since the issues and possible improvements are tightly linked to the technology and the approach.

Experimental outlooks

The current experimental setup is far from basic, and already features some rather advanced characteristics. The main challenge, currently, is its upscaling in order to perform launches comprising more sondes: as seen, this improves the sample size for averaging quantities and therefore is among the main predictors in the accuracy of an experiment. It would be easy, then, to simply say that more money is needed to buy more sondes and carry out larger experiments and leave it at that. However, (true as that may be) many aspects of the setup need to be carefully considered when imagining a larger version of it. In Chilbolton, five people had an actively operational role in the collection of the dataset, and at the moment of launch a total of about a dozen helped in the deployment of the sondes. This was for 7 sondes in the first launch and 10 in the second. Additionally, we note that this was the available number of people, not the *optimal* number of people. Indeed, a larger number would have allowed for a more careful and precise handling of the sondes, inflating of the balloons and releasing of the floating systems. It is then apparent to see that a launch featuring 20, 30 or 100 sondes would not simply be a matter of purchasing the materials, but would consist of a complex and delicate operation that, if not planned down to the very details, would have no hope of success. The

planning would then have to go hand in hand with the education of the staff, that should be instructed about at least the basic procedures of preparation and deployment. After launch, the issue would then be the pursuit of the sondes, which as told in Chapter 4 was done by car for the WESCON experiments. While this was the most versatile option at the time, it can be greatly optimized, despite requiring some definitely non-trivial work. The most crucial improvement would be the implementation of an interface allowing for immediate real-time monitoring of the sondes' position, that would allow the driver to easily follow them along their flight. In conditions where even slight convection is present the sondes are much more likely to exit the area of operation of the receiving ground station due to the wind mean flow rather than due to spreading about the cluster's center of mass (as can be clearly seen by the maps in Fig. 4.1 and 4.11) therefore a single moving station would likely be enough to cover all the necessary distance, but with larger clusters this may stop being true once a few sondes drastically change direction due to being swept away by some different atmospheric structure (see for example Sonde 21 in Fig. 4.11). Depending on the wind conditions several vehicles may be necessary to avoid seeing a dramatic reduction in sample size at the later times of the experiment. One yet more ambitious possibility for improvement, which would be a whole project in and of itself, would be to mount the moving receiver not on a car but on a drone. In principle, the drone could be programmed to read the incoming position data and automatically move towards the cluster's center of mass, ensuring optimal positioning at all times. This is far easier said than done, but is worth keeping in mind should the opportunity to develop this idea ever arise.

In second we can discuss the technical improvements that can be made to the flying system, namely the sondes and the balloons. As said before a new, lighter prototype of the sondes is already in the works. The reduced weight is indeed among the main points of improvement, due to the lower inertia and most of all to the smaller balloons necessary to carry it. The second key issue is the sonde's lifespan: as mentioned in previous chapters, all adjustments to the experimental procedure allowing us to follow the sondes for many hours would be to no avail if the battery duration is not extended accordingly. Unfortunately this point is obviously in contrast with the former one: longer lifespan would imply larger batteries, which would weigh more. The impossibility of having smaller batteries with the same (or higher) energy storage is due to the power volume density of a material (or of a technology to be more precise) being a constant, which implies that a battery's capacity is unequivocally linked to its size. There are then two possibilities: one is to optimize the power usage from the components, although these have already been selected among the most efficient commercially available components; the other is to use a different battery technology than the currently used lithium-ion. Power storage research has been rampant over recent years, partly due to the rise of electric motor vehicles, and a few new technologies seem promising in delivering

much better performances than the current standard. Just a few days ago at the moment of writing this the Chinese company Betavolt has announced the first miniaturized nuclear-powered battery, with a claimed lifespan of over 50 years. This of course would be an absolute revolution, although we seem to still be far from consumer-level usage.

Regarding the balloons, their development continues, mostly from the standpoint of their production rather than of their actual properties. In this case it may be more appropriate to describe the issue in monetary terms: industrial prototyping and production *ex novo* is a costly endeavor, and the price is likely to be a significant obstacle even when what needs doing is exactly known.

Computational outlooks

I have mentioned in the relevant sections that, after doing abundant use of the simulations described in this thesis, new ones are now under way thanks to a recently submitted project that aims at obtaining computational resources from the CINECA institute. More than running completely new simulations, the idea behind this project is to extend the datasets stored so far (despite some new simulations being present in the plan, like the missing monodisperse population, stable stratification case). The straightforward method will be to plug the final time instances of the current datasets as initial conditions to the new simulations. The idea is to reach a number of eddy turnover times closer to the order we obtained in the experimental datasets, namely about $\sim 40 - 50$. This is important for the Lagrangian correlation studies because, as we saw, the current duration is not sufficient to observe the full decay of all correlations, which we expect in any case due to the decay of the turbulent kinetic energy. Concerning the dispersion studies, instead, although we have no bigger hope to observe the Richardson-Obukhov regime due to the low Reynold's number, the longer duration may help us in identifying any existing superlinear behavior, such as for example the Batchelor t^2 regime. Other than the issues discussed in this thesis, longer simulations will also allow to observe the development the droplet population under other lights, both old and new.

After the present simulations are extended in this way we may turn to the possibility of producing new ones. Specifically, as should by now be abundantly clear, the current ones were not originally intended to produce relative dispersion studies, and were conceived with very different studies in mind. To study relative dispersion we have to focus on the key controlling parameters, and adequately incorporate them in the simulations, all the while making sure not to alter them in such a way as to deprive them of their original goal (that is modelling a cloud, and not just some suspension of undefined particles inside some fluid). With this in mind, the paramount issue is to increase the system's Reynold's number. On one side,

this calls for nothing other than more computational resources: in DNS there is no way around the fact that equations need to be solved, and this requires processing power. On the other side though there are many ways to optimize the computation, both at the level of the specific simulation and at the level of the code used as a whole. A simulation made for the precise scope of studying relative dispersion could be run by switching the current droplet populations, that are extremely numerous and occupy the whole cloud region, for a much smaller population (obviously still large enough to ensure an appropriate sample size), to be localized in two small volumes inside the cloud bulk and the mixing regions. The droplet dynamics would then have a much smaller impact on the computational time, and the resources formerly allocated to their evolution could be allocated to enlarging the domain on one side and intensifying the turbulence on the other. As far as general code efficiency is concerned, instead, one path that seems very promising is to work on code hybridization, meaning the implementation of libraries that involve the usage of a GPU in addition to the standard CPU. This is not the right place to delve deep into the exact advantages this would bring, but simply put it would add one level of parallelization to the already adopted core-level parallelization, which can potentially result in a very significant speedup. We (meaning me and one or two other people in the group, depending on their availability) are currently enrolled in a Hackaton-type event with the specific purpose of developing GPU-parallelised code starting from a regular CPU oriented one, which will take place towards the end of may.

Hopefully I have done a good enough job of conveying the message I had in mind: much has been done, much needs to be done. The direction is quite clear, although one of the main lessons I've learnt during the course of my work so far is that the most interesting things are not known at the start. Instead, they pop up while trying to navigate unknown waters, revealing new paths to tread, should one be so inclined.

Bibliography

- [1] E. Bodenschatz, S. P. Malinowski, R. A. Shaw, and F. Stratmann. «Can We Understand Clouds Without Turbulence?» In: *Science* 327.5968 (2010), pp. 970–971. DOI: 10.1126/science.1185138 (cit. on p. 1).
- [2] A.S. Monin and A.M. Yaglom. *Statistical Fluid Mechanics: Mechanics of Turbulence Volume 1*. The MIT Press, 1971 (cit. on pp. 8, 10).
- [3] David W. McComb. *The Physics of Fluid Turbulence*. Oxford Science Publications, 1990 (cit. on pp. 10, 11).
- [4] A.N. Kolmogorov. «The local structure of turbulence in incompressible viscous fluid for very large Reynolds numbers». In: *C.R. Acad. Sci. URSS* 30 (1941), pp. 301–305 (cit. on p. 11).
- [5] A.N. Kolmogorov. «Dissipation of energy in locally isotropic turbulence». In: *C.R. Acad. Sci. URSS* 32 (1941), pp. 16–18 (cit. on p. 11).
- [6] A.M. Obukhov. «Spectral energy distribution in a turbulent flow». In: *Dokl. Akad. Nauk SSSR* 32.1 (1941) (cit. on pp. 11, 15).
- [7] A.S. Monin and A.M. Yaglom. *Statistical Fluid Mechanics: Mechanics of Turbulence Volume 2*. The MIT Press, 1975 (cit. on pp. 11, 17).
- [8] David W. McComb. *Homogenous Isotropic Turbulence*. Oxford Science Publications, 2014 (cit. on pp. 11, 14).
- [9] L.D. Landau and E.M. Lifshitz. *Fluid Mechanics*. Pergamon press, 1987 (cit. on p. 11).
- [10] Lewis Fry Richardson. «Atmospheric diffusion shown on a distance-neighbour graph». In: *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character* 110.756 (1926), pp. 709–737 (cit. on p. 17).
- [11] Brian Sawford. «TURBULENT RELATIVE DISPERSION». In: *Annual Review of Fluid Mechanics* 33.1 (2001), pp. 289–317. DOI: 10.1146/annurev.fluid.33.1.289 (cit. on p. 17).

- [12] Juan PLC Salazar and Lance R Collins. «Two-particle dispersion in isotropic turbulent flows». In: *Annual review of fluid mechanics* 41 (2009), pp. 405–432 (cit. on p. 17).
- [13] G. Falkovich, K. Gawedzki, and M. Vergassola. «Particles and fields in fluid turbulence». In: *Rev. Mod. Phys.* 73 (4 2001), pp. 913–975. DOI: 10.1103/RevModPhys.73.913 (cit. on pp. 20, 96).
- [14] G. K. Batchelor. «The application of the similarity theory of turbulence to atmospheric diffusion». In: *Quarterly Journal of the Royal Meteorological Society* 76.328 (1950), pp. 133–146. DOI: <https://doi.org/10.1002/qj.49707632804> (cit. on p. 20).
- [15] George Keith Batchelor and Geoffrey Ingram Taylor. «The effect of homogeneous turbulence on material lines and surfaces». In: *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences* 213.1114 (1952), pp. 349–366. DOI: 10.1098/rspa.1952.0130 (cit. on p. 21).
- [16] G. Boffetta and I. M. Sokolov. «Relative Dispersion in Fully Developed Turbulence: The Richardson’s Law and Intermittency Corrections». In: *Phys. Rev. Lett.* 88 (9 2002), p. 094501. DOI: 10.1103/PhysRevLett.88.094501 (cit. on p. 22).
- [17] Nadeem A Malik and Fazle Hussain. «New scaling laws predicting turbulent particle pair diffusion, overcoming the limitations of the prevalent Richardson–Obukhov theory». In: *Physics of Fluids* 33.3 (2021), p. 035135 (cit. on pp. 22, 100).
- [18] Søren Ott and Jakob Mann. In: *Journal of Fluid Mechanics* 422 (2000), pp. 207–223. DOI: 10.1017/S0022112000001658 (cit. on p. 22).
- [19] Jacob Berg, Beat Lüthi, Jakob Mann, and Søren Ott. «Backwards and forwards relative dispersion in turbulent flow: An experimental investigation». In: *Phys. Rev. E* 74 (1 2006), p. 016304. DOI: 10.1103/PhysRevE.74.016304 (cit. on p. 22).
- [20] L. Biferale, G. Boffetta, A. Celani, B. J. Devenish, A. Lanotte, and F. Toschi. «Lagrangian statistics of particle pairs in homogeneous isotropic turbulence». In: *Physics of Fluids* 17.11 (2005), p. 115101. ISSN: 1070-6631. DOI: 10.1063/1.2130742 (cit. on p. 22).
- [21] Brian L. Sawford, P. K. Yeung, and Jason F. Hackl. «Reynolds number dependence of relative dispersion statistics in isotropic turbulence». In: *Physics of Fluids* 20.6 (2008), pp. 065111–065111–13. DOI: 10.1063/1.2946442 (cit. on p. 22).

- [22] James M. Wallace. «Space-time correlations in turbulent flow: A review». In: *Theoretical and Applied Mechanics Letters* 4.2 (2014), p. 022003. ISSN: 2095-0349. DOI: <https://doi.org/10.1063/2.1402203>. URL: <https://www.sciencedirect.com/science/article/pii/S2095034915303081> (cit. on pp. 23, 24).
- [23] Guowei He, Guodong Jin, and Yue Yang. «Space-Time Correlations and Dynamic Coupling in Turbulent Flows». In: *Annual Review of Fluid Mechanics* 49.1 (2017), pp. 51–70. DOI: 10.1146/annurev-fluid-010816-060309 (cit. on p. 23).
- [24] J. L. G. Oliveira, C. W. M. van der Geld, and J. G. M. Kuerten. «Concentration and velocity statistics of inertial particles in upward and downward pipe flow». In: *Journal of Fluid Mechanics* 822 (2017), pp. 640–663. DOI: 10.1017/jfm.2017.289 (cit. on p. 24).
- [25] H.R. Pruppacher and J.D. Klett. *Microphysics of Clouds and Precipitation*. Atmospheric and Oceanographic Sciences Libraries, 1997 (cit. on p. 29).
- [26] R.R. Rogers and M.K. Yau. *A short course in cloud physics*. Pergamon Press, 1996 (cit. on p. 29).
- [27] Pao K. Wang. *Physics and Dynamics of Clouds and Precipitation*. Cambridge University Press, 2013 (cit. on p. 29).
- [28] CL Lin and SC Lee. «Collision efficiency of water drops in atmosphere». In: *JOURNAL OF THE ATMOSPHERIC SCIENCES* 32.7 (1975), pp. 1412–1418. ISSN: 0022-4928. DOI: 10.1175/1520-0469(1975)032<1412:CEOWDI>2.0.CO;2 (cit. on p. 36).
- [29] RJ SCHLAMP, SN GROVER, HR PRUPPACHER, and AE HAMIELEC. «NUMERICAL INVESTIGATION OF EFFECT OF ELECTRIC CHARGES AND VERTICAL EXTERNAL ELECTRIC-FIELDS ON COLLISION EFFICIENCY OF CLOUD DROPS». In: *JOURNAL OF THE ATMOSPHERIC SCIENCES* 33.9 (1976), pp. 1747–1755. ISSN: 0022-4928. DOI: 10.1175/1520-0469(1976)033<1747:ANIOTE>2.0.CO;2 (cit. on p. 36).
- [30] Ludovico Fossà, Shahbozbek Abdunabiev, Mina Golshan, and Daniela Tordella. «Microphysical timescales and local supersaturation balance at a warm cloud top boundary». In: *Physics of Fluids* 34.6 (2022), p. 067103. DOI: 10.1063/5.0090664 (cit. on pp. 38, 39, 43, 52–54).
- [31] Mina Golshan, Shahbozbek Abdunabiev, Mattia Tomatis, Federico Fraternali, Marco Vanni, and Daniela Tordella. «Intermittency acceleration of water droplet population dynamics inside the interfacial layer between cloudy and clear air environments». In: *International Journal of Multiphase Flow* 140 (2021), p. 103669. ISSN: 0301-9322. DOI: <https://doi.org/10.1016/j.ijmultiphaseflow.2021.103669> (cit. on pp. 39, 43–45, 49–51).

- [32] M. Iovieno, C. Cavazzoni, and D. Tordella. «A new technique for a parallel dealiased pseudospectral Navier-Stokes code». In: *Comput. Phys. Commun.* 141 (2001), pp. 365–374 (cit. on p. 38).
- [33] L. Gallana, S. Abdunabiev, M. Golshan, and D. Tordella. «Diffusion of turbulence following both stable and unstable step stratification perturbations». In: *Physics of Fluids* 34.6 (2022), p. 065122. DOI: 10.1063/5.0090042 (cit. on p. 43).
- [34] B. Devenish et al. «Droplet growth in warm turbulent clouds». In: *Quarterly Journal of the Royal Meteorological Society* 138.667 (2012), pp. 1401–1429. DOI: 10.1002/qj.1897 (cit. on p. 45).
- [35] Abdunabiev et al. «Validation and traceability of miniaturized multi-parameter cluster radiosondes used for atmospheric observations». In: *Measurement* 224 (2024). ISSN: 0263-2241. DOI: <https://doi.org/10.1016/j.measurement.2023.113879> (cit. on pp. 56, 59, 63, 69–76).
- [36] Paredes et al. «Innovative mini ultralight radioprobes to track Lagrangian turbulence fluctuations within warm clouds: electronic design». In: *Sensors* 21.4 (2021). ISSN: 1424-8220. DOI: 10.3390/s21041351 (cit. on pp. 59, 65–68).
- [37] Tessa C. Basso, Giovanni Perotto, Chiara Musacchio, Andrea Merlone, Athanasia Athanassiou, and Daniela Tordella. «Evaluation of Mater Bi and Polylactic Acid as materials for biodegradable innovative mini-radiosondes to track small scale fluctuations within clouds». In: *Materials Chemistry and Physics* 253 (2020), p. 123411. ISSN: 0254-0584. DOI: 10.1016/j.matchemphys.2020.123411 (cit. on p. 62).
- [38] Ewan O’Connor. *Model data from Chilbolton on 5 July 2023*. 2023. URL: <https://hdl.handle.net/21.12132/1.bb07bdeb9739418f> (cit. on p. 83).
- [39] Ewan O’Connor. *Model data from Chilbolton on 6 July 2023*. 2023. URL: <https://hdl.handle.net/21.12132/1.5088f7ded8124af7> (cit. on pp. 92, 93).