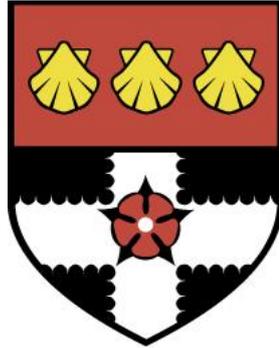


UNIVERSITY OF READING

Department of Meteorology



**Island Convection and its
Representation in Numerical
Weather Prediction Models**

MICHAEL CAVIN JOHNSTON

A thesis submitted for the degree of Doctor of Philosophy

December 2019

Declaration

I confirm that this is my own work and the use of all material from other sources has been properly and fully acknowledged.

Michael Cavin Johnston

Acknowledgements

I would like to start by thanking my supervisors - Chris Holloway and Bob Plant for taking on my project and their guidance, support, and advice over the course of the PhD. They were particularly speedy and helpful with feedback as this thesis came together near the end. In addition, I extend my thanks toward my monitoring committee - Andy Turner and Miguel Teixeira - who have provided essential big-picture perspective that has helped keep this PhD on track. I must also acknowledge the many people in the Department, the Met Office, and the Bermuda Weather Service who have provided essential feedback in research groups, aid in operating the idealised UM, and inspiration and data for using Bermuda as a case study island. These include many presentation opportunities and resulting good discussion in Mesoscale Group, Tropical Hour, the ParaCon Group - in addition to separate discussions with Peter Clark, Carol Halliwell, Kirsty Hanley, and Rachel Stratton.

Furthermore, I would be amiss if I didn't extend my thanks to my peer support groups. This includes my original 179 Harry-Pitt office, H el ene Bresson, Ying Ying Toh, Davi Mignac, Will Morrison, Tom Hall, and Tom Eldridge who helped guide me through the early days of the PhD and made me feel welcome on office getaways. While 179 Harry-Pitt was eventually split up, my peer support continued into 1U08 Meteorology where I joined Aga Faulkner, Liz Cooper, and later Kaja Milczewska. I feel we truly bonded as an office with coffee breaks taken as an office - even if just to look away for a few minutes and have a quick chat. My gratitude extends to the department's many other PhD candidates who have always been a group I can turn to - each person for a different reason.

Last, but not least, I have to thank my parents, my sister, and my uncles. My family have supported me through the PhD in countless ways from day one. In particular, financially and emotionally. This thesis truly has been a group effort with much help from behind the scenes. Without this army of support over these three-plus years, the technical difficulties which lurked around each and every corner would have won.

Abstract

Cloud Trails (CTs) are organised, thermally-forced bands of shallow convection which form downwind of small heated islands. Increased availability of high-frequency satellite imagery offers new opportunities to investigate CT occurrence. There are few examples of the environments associated with CT occurrence, with few links to the CT system in general, and none for Bermuda. Furthermore, Numerical weather prediction (NWP) systems are currently evolving toward higher resolution. As a result, CT-producing islands can now be partially represented by regional models. This motivates renewed exploration of controls on CT occurrence, the role of CTs for the evolution of the larger-scale, and how CT representation is degraded when poorly-resolved. We extended our understanding of CT and shallow convection with a two-pronged observational-numerical approach.

First, a five-year climatology of CT occurrence covering the warm season at Bermuda is constructed using a novel automated detection algorithm trained on manually classified, half-hourly visible satellite imagery and surface wind observations. A total of 16,400 images are classified with more than 5,000 CT images identified. CT are most common in July and the afternoon: coincident with peak solar heating and the peak strength of the large-scale subtropical ridge and associated settled summertime weather in Bermuda. Additionally, CT form more frequently when the low-level environment is warmer and more humid than climatology. Wind speed and direction are of further importance because of their control on the total heating of the flow as it crosses the island.

Next, the idealised UM is used to perform numerical experiments to examine the CT system more closely. A control experiment based on observations from Bermuda help us to extend the existing conceptual model of the CT system by expanding the role of the warm plume in stabilising the boundary layer in the wake of the island. Variants on the control experiment, are used to investigate changes in CT behaviour in response to their forcing and environment. We find that CT strength varies linearly with island heating as expected from scale analysis. Then we demonstrate that a CT circulation

of equivalent strength can form without clouds given a sufficiently dry environment. Finally, we show that the CT system evolves as it intensifies with decreasing wind speed. The CT becomes precipitating at lower wind speeds in our experiments, with more impact on the larger-scale environment and potentially for the public on or downwind of such islands.

The final part of this work introduces the problem of representing shallow convection at grid lengths where it is poorly resolved (i.e. the shallow convection “grey zone”). We show that shallow convection becomes too-strong and too-deep in the UM for grey zone resolutions. A crude representation of the CT system is possible for grid lengths of 1.6 km (similar to the resolution of the MetOffice’s UKV). But we show that when larger grid spacing is used, the CT becomes more intense and more strongly precipitating compared to a well-resolved CT. We then echo calls for the use of a shallow convection parametrisation in kilometre-scale models, and demonstrate some of the current challenges.

Contents

Declaration	i
Acknowledgements	iii
Abstract	v
Table of contents	vii
List of Figures	xi
List of Tables	xxix
1 Introduction and Motivation	1
1.1 Background and Motivation	2
1.1.1 Circulation Generation	2
1.1.2 Atmospheric Convection	6
1.1.3 Island Convection in Literature	9
1.1.4 Island Convection in Numerical Weather Prediction	13
1.2 Thesis Aims and Structure	17
2 Data Sources	21
2.1 Observations	21
2.1.1 Bermuda	21
2.1.2 In-Situ Observations	22
2.1.3 Satellite Data	26
2.1.4 Satellite Calibration	27
2.1.5 ERA-interim Reanalysis	28
3 Climatology of Cloud Trails at Bermuda	31
3.1 Introduction	31
3.1.1 Motivation	32
3.1.2 Background	32
3.2 Methods	35
3.2.1 Manual Classification Method	35
3.2.2 The Classification Algorithm	37
3.2.3 Algorithm Parameters	42
3.3 Results and Discussion	45
3.3.1 Algorithm Validation	45

3.3.2	CT Climatology	49
3.3.3	Environmental Characteristics	54
3.4	Algorithm For Barbados	61
3.5	Summary and Conclusions	63
4	Numerics of Shallow Convection: Resolution, Domain, and Idealisations	67
4.1	Introduction	67
4.2	The Met Office Unified Model	68
4.2.1	ENDGame Dynamical Core	69
4.2.2	Model Physics	71
4.2.3	Idealisations	80
4.3	Introduction to the BOMEX case	85
4.3.1	The BOMEX setup	86
4.3.2	The Equilibrium state	87
4.4	Evaluation of the Control BOMEX Case	92
4.4.1	Cloud Cover, Cloud Liquid, and TKE	92
4.4.2	Maintaining Equilibrium State	95
4.5	Sensitivity to grid and domain changes	97
4.5.1	BOMEX at Coarser Resolution	98
4.5.2	BOMEX with rigid model lid at 40 km	103
4.6	Method for initialising Cloud Trail simulations	108
4.6.1	Choosing Initial Conditions	109
4.6.2	Spin-up Simulations	112
4.7	Summary and Conclusions	119
5	Idealised Simulations of Cloud Trails	121
5.1	Introduction	121
5.2	The Cloud Trail System	122
5.2.1	Details of the Control Simulation	123
5.2.2	Results from the Control Simulation	126
5.3	Exploring the CT response	143
5.3.1	Island Surface Fluxes	146
5.3.2	The importance of Relative Humidity	152
5.3.3	Wind Speed	159
5.4	Summary and Conclusions	170
6	Cloud Trails in the Shallow Convective ‘Grey Zone’	173

6.1	Introduction	173
6.2	Methods for the Grey Zone	174
6.2.1	Choice of Initial Profiles	175
6.2.2	Design of the CT Experiments in the Grey Zone	175
6.3	Sensitivity of Spin-up Simulations to Grid Length	176
6.3.1	Results for the Grey Zone	181
6.4	The Shallow Convection Parametrisation Perspective	190
6.4.1	The Convection Scheme	190
6.4.2	Cloud Census and Population Statistics	194
6.4.3	Parametrised Results	204
6.5	Summary and Conclusions	208
7	Conclusions and Future Work	211
7.1	Summary of Work	211
7.1.1	Cloud Trail Climatology at Bermuda	212
7.1.2	Idealised Simulations: BOMEX and Shallow Convection	214
7.1.3	Idealised Simulations: The Well-Resolved CT System	215
7.1.4	Idealised Simulations: The Poorly-Resolved CT System	218
7.2	Future Work	220
	References	223

List of Figures

- 1.1 True colour imagery from the Aqua Moderate resolution Imaging Spectroradiometer (MODIS) on the 23rd May 2017 at approximately 1230 Local Time. The islands of Bermuda and the shallow waters of the lagoon to the north of the island are seen in cyan and green colours, the deeper ocean surrounding the island appears as dark blue, and the clouds appear white. A red box highlights the location of the cloud trail at this time. The red star marks the eastern end of Bermuda, near where the Bermuda Weather Service is located. [Inset] Webcam picture of the cloud trail at about the same time as the satellite imagery. Camera is looking toward the north-northeast. Courtesy of the Bermuda Weather Service. The clouds with darker grey cloud bases associated with the cloud trail are in the left half of this picture. 3
- 1.2 A schematic of the cloud trail circulation based on observations from the island of Nauru in [Matthews et al. \(2007\)](#). Schematic shows the island in black, streamlines of the full circulation. Areas where clouds preferentially form are indicated by small cloud shapes. [The source of this material is the COMET Website at <http://meted.ucar.edu/> of the University Corporation for Atmospheric Research \(UCAR\), sponsored in part through cooperative agreement\(s\) with the National Oceanic and Atmospheric Administration \(NOAA\), U.S. Department of Commerce \(DOC\). ©1997-2017 University Corporation for Atmospheric Research. All Rights Reserved.](#) 11
- 1.3 An example of many nearby and interacting cloud trails (along with other forms of island and non-island convection) as seen from Terra MODIS True Color imagery around Fiji on the 21st October 2019. . . . 16

- 2.1 A Digital Elevation Model of Bermuda based on data from [Sutherland et al. \(2013\)](#). Data are at a horizontal grid spacing of 1 arc-second (about 30 m). The coast, as defined by the 0 m above mean sea level contour, is outlined in blue. The maximum elevation from this data set is also included as a red triangle with the elevation marked. The location of the anemometers at runway 12 (RWY12) and runway 30 (RWY30) are marked and labelled as red crosses. The anemometers mark each end of the runway at the L. F. Wade International Airport. 23
- 3.1 GOES-13 visible satellite imagery showing example scenes. (a) Cloud Trail scene where clouds organise into a band downwind of Bermuda as indicated by a southwest-northeast oriented band of higher albedo. (b) Non-Trail scene in which there are few clouds and some of the higher albedo near and over the island might be shallow water and land showing up rather than cloud. (c) Obscured scene where the island (and much of the surroundings) is obscured from view by widespread cloud as indicated by high albedo throughout the scene. In each example, a wind barb is plotted showing 10 m wind direction and speed in kts at TXKF provided by the BWS. 36
- 3.2 Part of the sensitivity analysis for cloud mask thresholds for 0700 - 1115 Local Time on 18 Aug 2016 (a mostly cloudless morning). The top left is the mean morning albedo. Mean albedo over the island is higher than the surrounding water. The remaining panels are the morning cloud frequency for increasing cloud mask thresholds (cmt). The title of each panel states the cmt used. The 15 m water depth is given by a thin blue contour to highlight the extent of the shallow waters of the lagoon to the north of Bermuda which can be seen in the cloud frequency for the cmt = 0.06 and cmt = 0.09 panels. 38

3.3	Visible-channel satellite imagery as in Fig. 3.1. An example of a case of high solar zenith angle at sunrise. The left half of this image has not yet had sunrise and appears black. The right half of this image has just had sunrise and so the incoming solar radiation is weak. The albedo (the ratio of the reflected solar radiation to incident solar radiation) is artificially high.	39
3.4	Walk-through of the satellite image classification algorithm. a) The visible-channel satellite scene is screened for solar zenith angle less than 75° to ensure it is well-illuminated; b) a cloud mask threshold is applied to the imagery; c) the cloud fraction in the grey circle, excluding the island is computed to determine whether the scene is obscured or not; d) if the scene is not obscured, the maximum cloud fraction in the upwind and downwind quadrants are compared. If the difference between the two is greater than some threshold then the scene is cloud trail, otherwise non-trail.	40
3.5	Cumulative distribution function of cloud fractions for manually classified non-OB scenes compared to the inverse cumulative distribution function of cloud fractions for manually classified OB scenes. The cloud fraction is calculated over the grey region illustrated in Figure 3.4c. By taking the cloud fraction where these two distributions intersect as the threshold to distinguish between OB and non-OB, we maximise the number of obscured scenes which are correctly classified.	43
3.6	Cumulative distribution function of the maximum difference in cloud fraction between the upwind and downwind regions, δF , for manually classified NT scenes compared to the inverse cumulative distribution function of δF for manually classified CT scenes. The regions used to calculate δF are illustrated in Figure 3.4d. Only the scenes which are manually classified as non-OB are included here. By taking the δF where these two distributions intersect as the threshold for to separate non-OB scenes into NT and CT, we maximise the number of correctly classified CT scenes.	44

3.7	A comparison between i) the algorithm classifications for 2012 and ii) the manual classifications for 2012 of the percentage of scenes classified as a) CT, b) NT, and c) OB sorted by local time from left to right and month from top to bottom. A black ‘x’ is in the cell which has no classified scenes because of high solar zenith angles.	48
3.8	As in 3.7 but just for the algorithm applied to the whole period of record, 2012 - 2016.	51
3.9	Frequency of different contiguous CT periods per day for the warm season of 2012. Results from a) the algorithm classification, and b) the manual classification.	52
3.10	a) The climatological potential temperature (θ) profile from radiosondes at 0900 LT JJA, 2012-2016. b) The composite θ anomalies for CT (black dashdot line and gray shading), NT (purple dashed line and shading), and OB (dark red dotted line and shading). The shaded region represents the uncertainty about the mean anomaly: $\pm \frac{\sigma}{\sqrt{n}}$. All composites are of 0900 LT radiosondes for JJA in 2012-16. c) as in a) but for relative humidity (RH). The mean LCL for each classification and for the climatology are shown as horizontal line segments with corresponding line colours and styles. d) as in b) but for RH composites. In each panel the surface value is given by a dot with a range given by $\pm \frac{\sigma}{\sqrt{n}}$. LCL pressures are calculated using the radiosonde temperature and dewpoint at the lowest reported altitude and Eqs. 2.2a, 2.2b, and 2.2c from Chapter 2.	56
3.11	JJA 2012-2016 ERA-Interim reanalysis composites centered on Bermuda. Composites for a), b) cloud trails; c), d) non-trails; and e), f) obscured. (left) The mean sea level pressure in hPa (solid black contours), 1000-hPa temperature in °C (dashed grey contours), and 1000-hPa specific humidity in g kg ⁻¹ (brown to green-blue shading). (right) The 500-hPa vertical velocity in Pa s ⁻¹ , ascent (red) and subsidence (blue) are shown in dashed and solid contours, respectively.	60

- 3.12 Cloud frequency by classification. a), b), and c) are the cloud frequency for the CT, NT, and OB categories respectively for Bermuda. Similarly, d), e), and f) are the same respective categories but for Barbados. A blue star marks the centre-point of the island in question. Each scene is cropped to be $\pm 0.5^\circ$ away from the island in either direction. See text for more information. 62
- 4.1 Critical relative humidity (RH_{crit}) plotted against horizontal grid spacing (i.e. flight-leg distance from field campaigns). The mean is in black, and the 5th and 95th percentiles are in blue. The 5th percentile is described by two lines of best fit (equation 4.16a), one for grid spacing less than 20 km as a dashed green line, and one for grid spacing greater than 20 km as a solid green line, however, because there is so much variability for grid spacing above 20 km, a constant value is used. The best fit to the mean is in a red solid line, and the best fit for the 95th percentile is a green solid line. 77
- 4.2 Map of the BOMEX ship array during the first three observation periods 3-15 May, 24 May to 10 June, and 19 June to 2 July, 1969; as in Holland (1972) (©Copyright 2020 AMS). 85
- 4.3 The initial conditions and large-scale forcing profiles applied in the BOMEX simulations following the set-up described in Siebesma and Cuijpers (1995) and Siebesma et al. (2003). a) The initial θ (solid line) and q_T (dashed line) profiles. b) the initial u (solid line) and v (dashed line) profiles. c) the geostrophic forcing profiles for u_g (solid line) and v_g (dashed line). d) the subsidence profile . e) the radiative cooling profile (solid line) and the advective drying profile (dashed line). The grey shading between 500 and 1500 m represents the expected cloud layer. 88

- 4.4 Liquid water path shaded in greys at a) during the spin-up, 30 minutes into the simulation; b) at the start of the quasi-steady-state, 180 minutes into the simulation; and c) during a maximum in domain mean liquid water path, 320 minutes into the simulation (see Fig. 4.5b). A blue contour is drawn to mark 0 g m^{-2} liquid water path. 93
- 4.5 A time-series of macro-scale cloud properties for the first 6 hours of our BOMEX simulation. a) The total cloud cover as a fraction of the horizontal space with non-zero liquid water path; b) the horizontal mean liquid water path; and c) the horizontal mean, vertically integrated turbulence kinetic energy. The results from the UM simulation are in red. Estimated from Siebesma et al. (2003) Figure 2 are the mean (black line) and $\pm 2\sigma$ (grey shading) of the 10-model intercomparison. 94
- 4.6 Horizontal-mean profiles over the last 3 hours of a) potential temperature, θ , b) water vapour specific humidity, q_v , c) the horizontal wind components, u and v , d) the cloud liquid water, q_l . The red solid lines represent the mean profile during this time-frame. In a), b), and c) the dashed black line represents the initial conditions. Estimated from Siebesma et al. (2003) Figure 3d are the mean (black line) and $\pm 2\sigma$ (grey shading) of the 10-model intercomparison for the cloud liquid water. 96
- 4.7 Box plots of the horizontal-mean conditions over model hours three through six of the BOMEX grid length experiments summarised in Table 4.4. Presented are the total cloud cover in a), the liquid water path in b), and the vertically-integrated TKE in c). For each box, the middle line represents the median value, the top and bottom of the box represent the interquartile range, and the top and bottom of the whiskers represent the interdecile range. Small blue dots either side of the whiskers mark outliers. 100

4.8	Vertical cross sections of through a particularly strong convective plume in DX800. a) an x-z cross section, and b) a y-z cross section. Cloud liquid water, q_{cl} , is presented in grey shading with black contours of vertical velocity every 1 m s^{-1} , w , overlaid (solid where positive, dashed where negative, a heavy black line marks the zero vertical velocity).	101
4.9	A summary of the spin-up behaviour across the BOMEX grid length experiments. a) the time at which total cloud cover first increases to 0.05, b) the peak total cloud cover during the spin-up, c) the peak liquid water path during the spin-up, and d) the peak vertically-integrated TKE during the spin-up. In each panel dots are coloured by grid length.	102
4.10	Difference between the initial q_v profile and the horizontal-mean, 3-hour time-mean q_v profiles for model hour three through six for the BOMEX grid length experiments. The percentage change from the initial conditions is reported next to the corresponding legend entry.	103
4.11	An illustration of the spacing between vertical levels for the control experiment (solid grey lines) and DX100L140 (dotted black lines). Three layers in the lowest 3 km are shown.	104
4.12	As in Fig. 4.5 with the DX100L140 values added to the figure in blue. An experiment with interactive surface fluxes, DX100L140if (to be introduced below), is also included in purple.	106
4.13	As in Fig. 4.6 panels a and b, but only the lowest 600 m. The DX100L140 values are added to the figure in blue.	107

4.14 Skew-T log-P diagram of the observed temperature in a solid red curve and observed dew point in a solid blue curve. A parcel ascent (based on the mean properties of the lowest 500 m) is shown in a solid grey curve. The pressure of the LCL is indicated as a solid black line. Pressure, in hPa, is marked by dotted grey lines; temperature, in °C, is marked by dashed grey lines; mixing ratio, in g kg^{-1} , is marked by dashed purple lines; dry adiabats, in °C, are marked by thin green lines; and moist adiabats, in °C, are marked by thin blue lines. Winds barbs are used to note the wind speed and direction with height on the right, where a full barb indicates 10 kts, a half barb indicates 5 kts, and a circle indicates calm winds. 110

4.15 A comparison between the radiosonde and simplified initial condition profiles of a) the potential temperature, b) the relative humidity, and c) the wind speed. 111

4.16 Time series of a) The mass-weighted boundary layer mean u- and v-wind components, b) the boundary layer height, z_h , which is diagnosed as the height at which the local Richardson Number, Ri , increases to above a threshold critical Richardson Number, $Ri_c = 0.25$, and c) the surface energy fluxes for the ten-day spin-up simulation. The start of day six is marked on each plot with a vertical dotted grey line, the region to the right of this line (i.e. the last four simulated days) is considered to be in ‘equilibrium’. 115

4.17 Comparison of the simplified profiles used to initialise the spin-up simulations with the equilibrium profiles from the Control spin-up simulation. As in Fig. 4.15 116

- 4.18 An illustration of how the iterative end-point fit algorithm reduces a highly resolved example curve ($\pi\cos(x)$) with 70 points to a curve with 13 points and a maximum error of $\varepsilon_y = 0.5$. In each panel the well-resolved curve is marked in black. a) The first guess curve (blue dashed line), the tolerance envelope (blue shaded region), and point with greatest distance outside that envelope (red dot). If more than one point has the same distance, the first occurrence is considered. b) The red dot from a) becomes the end point of the line for this iteration (blue dot), but there is still a point outside the tolerance (red dot). c) the red dot from b) becomes the end of the line for this iteration (blue dot), but now no points lie outside the tolerance, and the next segment is initiated from the red point in c) to the end of the well-resolved curve. This is repeated until the whole curve is represented, as in d). 119
- 5.1 Schematic where “sea” grid points are shown in white and “land” grid points are shown in black for the experiments with 100 m horizontal grid spacing. The land points correspond to the 50 km² island which is centred at $x = 108$ km, $y = 16$ km. 125
- 5.2 Snapshots of the liquid water path in g m^{-2} in greyscale shading. A contour of $LWP = 0 \text{ g kg}^{-1}$ is outlined in blue, and the island coastline outlined in red. Snapshots are taken every two hours from 3 - 11 inclusive. Surface fluxes, H and E , over the island at the time of the snapshots are noted in each panel’s title. 127
- 5.3 The daytime (hour 06 - 18 model time) cloud frequency for the reference simulation. a) as a function of distance away from the leeward edge of the island and b) the spatial distribution of cloud frequency. The cloud frequency in a) is computed by taking the average in the y' direction across the CT region marked as a blue rectangle in b). In a) a grey horizontal line denotes the spatial mean daytime cloud frequency outside the CT region and represents the ambient and suppressed cloud frequency. The island is highlighted in red in both panels. 129

- 5.4 Snapshots of the potential temperature anomaly from the horizontal mean at each time, θ' , in K. θ' is taken at the lowest model level, 2 m. Snapshots are taken every two hours from hour 3 to 11 inclusive. The island sensible heat flux, H , and the island latent heat flux, E , at the time of the snapshot are noted in the title. 132
- 5.5 Schematic illustrating the relationship between the ambient flow, u , the thermally-forced flow perturbations v' , and the island geometry, given by island radius, r , combine for the triangular shape of the low-level warm plume. a) Example values for u , v' , and r . b) The triangular warm plume, red shaded region is determined by how long it takes v' to advect the edges of the warm plume inward where they meet in the middle after some time Δt . During Δt , the whole system has also moved downwind with u leading to the characteristic triangular warm plume shape in red shading. This process also controls the length of the warm plume. 133
- 5.6 Snapshot at hour 12 of a) the difference between θ of a surface parcel and θ of the environment along the CT region, and b) the anomaly of that difference along the CT region from the horizontal mean. In both panels, the bold line is the mean height of the LCL and the thin line is the mean boundary layer height, z_h , both are along the CT region. The capping influence of the warm plume is highlighted as positive anomalies in b) and z_h tracks the bottom of this capping warm layer where the difference between the LCL and z_h is largest. 134
- 5.7 Snapshots of the potential temperature anomaly, θ' at the model level nearest to the horizontal mean boundary layer height at each time a) $z_h = 622$ m at model hour 6, b) $z_h = 609$ m at model hour 9, and c) $z_h = 586$ m at model hour 12. The 'V'-shaped part of the warm plume is marked with two arrows with dashed lines that parallel the arms of the 'V', and the cool anomaly aligned with the cloud band is marked by an arrow with a dotted line. 135

5.8 Snapshots of the 10 m wind speed, U , for a) model hour 6 (at dawn), b) model hour 9, and c) model hour 12. The horizontal-mean wind direction vector at this level points from east-northeast to west-southwest and is given as a black arrow in c). 136

5.9 Midday (model hour 12) snapshots of a) the along-flow wind anomaly, s' , averaged across the CT region and b) the across-flow wind anomaly, n' , averaged along the CT region. In a), downwind is to the right, and in b), the perspective is centred on the flow with the right side of the plot to the right of the flow, and the left side of the plot to the left of the flow. In both panels, positive anomalies point to the right. 138

5.10 As in Fig 5.8 but for the 10 m across-flow component of the wind, n' . Here, positive denotes flow from the south-southeast, and negative from the north-northwest. 139

5.11 A snapshot of a) cumulative distribution function (CDF) of cloud top heights, z_{top} , at model hour 6 (in black) and 12 (in blue), b) the difference between the midday CDF and dawn CDF, and c) the midday cloud mask with a blue rectangle containing the CT region. Annotated are the mean midday cloud top heights, $\overline{z_{top}}$ for the CT region in blue, and elsewhere in black. 141

5.12 A schematic diagram of the main features of the Cloud Trail System as identified in this section. a) The across-flow system where the land-sea contrast is surface buoyancy fluxes generates a convergent low-level circulation at and downwind of the island similar to the traditional sea-breeze. Cloud forms in the turbulent warm plume that forms over and downwind of the island, and atop the ascending branch of the across-flow circulation. b) The along-flow system where the traditional sea-breeze circulation is advected downwind and transformed into the “heat island circulation”. The main cloud band is supported by the across-flow circulation which is advected downwind from the island, while, at least in the Control simulation, the cloud band and turbulent warm plume cloud feature are separated due to along-flow variations in static stability and dominant mechanisms for ascent. 142

5.13 Daytime cloud frequency for the eight heat flux experiments in which only the peak island surface sensible heat flux is varied (i.e. excluding H125E375 and H375E125). Here the Control experiment, H250E250, is shown in panel e). 148

5.14 Instantaneous potential temperature anomaly from the horizontal mean, θ' , at hour 12 (midday), on the lowest model level (2 m). In each panel the island coastline is marked in black. Here, the Control experiment, H250E250, is shown in panel e). 149

- 5.15 The circulation mass flux, M_c , against a) peak total island turbulent surface energy flux, Q_{net} , b) peak island surface sensible heat flux, H_0 , and c) peak island surface buoyancy flux. A line of best fit is drawn as a dashed black line and the equation for the line of best fit is reported in each panel's title. The symbols are coloured with a darker red for higher H_0 in each experiment, and are circular for experiments where only H_0 is changes, and triangular for experiments where the Bowen ratio, β , is changed but Q_{net} is kept constant. An example of the time-mean mass flux between 9 AM and 12 PM at the height LCL is shown in the shading in d). In d), a bold black circle represents the island coastline, and a dashed rectangle represents the area over which M_c is computed. 151
- 5.16 Initial profiles of a) RH, b) specific humidity, q_v , for BLm25 (red dashed line), FAm25 (blue dashed line), Control (short) (black solid line). The remainder of the profile above 4 km remains unchanged between the initial conditions. 153
- 5.17 As in Fig. 5.16 but for the conditions at dawn, after two model hours in these experiments. 155
- 5.18 The instantaneous across-flow wind component, n' , at 10 m (blue-white-red shading), area covered by daytime cloud frequency greater than 25 % (grey shading), and the instantaneous warm plume extent as defined by the 0.1 K θ' contour at 2 m. The n' and the warm plume extent are taken at model hour 12. Fields are presented for a) Short Control experiment, b) BLm25, and c) FAm25. The short Control experiment is the same as the original Control except it starts at 2 hours before dawn to match BLm25 and FAm25. 156
- 5.19 The resolved part of the horizontal-mean, time-mean vertical turbulent zonal momentum flux, $(\overline{w'u'})$, over the first two model hours and the whole domain for BLm25 in red, FAm25 in blue, and Control (short) in black. 157

5.20	New initial conditions following the spin-up for the wind experiments. The lines become darker red for increasing wind speed in each plot. Over the lowest 4 km, a) the u- and v- wind components, b) the relative humidity profiles, c) the potential temperature profiles, and d) the specific humidity profiles are presented.	160
5.21	As in figure 5.13 but for the suite of wind speed experiments. Here, U10 is the Control experiment.	162
5.22	As in figure 5.14 but for the suite of wind speed experiments. Here, U10 is the Control experiment.	163
5.23	Total daytime precipitation for the suite of wind experiments. The island coastline is highlighted as the black contour in each panel. The maximum total daytime precipitation, P_{max} , is reported in the title of each panel and its location is marked by a purple ‘x’ in the panel. The Control experiment is U10.	164
5.24	The daytime (model hour 6 to 18) cloud frequency on the left, and an instantaneous snapshot of the potential temperature anomaly, θ' , at hour 12. For a) and b) U05, c) and d) U05 with rain switched off, and e) and f) U05 with half the peak island sensible heat flux, $H_0 = 125 \text{ W m}^{-2}$	167
5.25	Comparison of a) the maximum of the mean low-level CT circulation strength derived from our experiments, v'_{CT} , and b) the length of the surface warm plume derived from our experiments, L_{WP} , with the scalings for the two parameters (see text above for more on how we derive these parameters from experiments and for the origins of the scalings). A line of best fit is included as a dashed black line in each panel, and the equation that describes the line is given near the bottom of each respective panel. Dots are coloured by background wind speed.	170
6.1	Boxplots of the horizontally-averaged a) LCL, b) boundary layer depth, c) cloud top height, and d) surface latent heat flux over the last four days of the spin-up simulations for the experiments with different grid lengths.	178

- 6.2 Horizontally-averaged vertical profiles of a) the u- and v- components of the in m s^{-1} where u is given as solid lines and v as dashed lines, and b) the potential temperature, θ , in K in solid lines and the specific humidity q_v in g kg^{-1} in dashed lines. These are a comparison across the equilibrium conditions for spin-up simulations with different horizontal grid lengths. 179
- 6.3 Horizontally-averaged vertical profiles of a) cloud liquid water content in g kg^{-1} , and b) cloud cover as a fraction of the whole domain. These are a comparison across the mean conditions during the four days used to define equilibrium state from the spin-up simulations with different horizontal grid lengths. 180
- 6.4 The daytime (0600 to 1800 hours) cloud frequency for the experiments which vary horizontal grid spacing. a) The original 100 m CT simulation from Chapter 5, b) the 200 m experiment, c) the 400 m experiment, d) the 800 m experiment, and e) the 1.6 km experiment. Recall that Chapter 5's 100 m CT simulation has a different inversion atop the cloud layer and so is just for reference. 183
- 6.5 An instantaneous snapshot of the liquid water path at 12PM in a) the 200 m experiment, b) the 400 m experiment, c) the 800 m experiment, and d) the 1.6 km experiment. Here, the cloud edge contour (i.e. where the liquid water path increases above 0 gm^{-2}) is marked with a blue contour, and the island coastline is marked with a red contour. 185
- 6.6 A snapshot of the across-flow component of the wind, n , at 10 m height and 12PM in a) DX0200, b) DX0400, c) DX0800, and d) DX1600. The island coastline is shown as a black contour in each plot. 186

6.7	The circulation mass flux, M_c , as defined in Chapter 5, Section 5.3.1. This is computed over the period from 9AM to 12PM in each of the experiments with increasing grid spacing. The 100 m experiment from Chapter 5 is included here as a square to give a sense of what the expected magnitude of M_c should be for these experiments. Recall that this experiment should not be considered as part of the main suite of experiments because it has different initial conditions.	188
6.8	A snapshot of the potential temperature anomaly, θ'_{sf_c} , at 12PM at the lowest model level (2 m) for a) the 200 m experiment, b) the 400 m experiment, c) the 800 m experiment, and d) the 1.6 km experiment. Here, θ' is defined as a deviation from the horizontal-mean at this time and height. In each panel, the island coastline is indicated with a black contour.	189
6.9	The total daytime (6AM to 6PM) precipitation for a) DX0200, b) DX0400, c) DX0800, and d) DX1600. The daytime rain total is noted in each panel's title as P_{max} and the location of this maximum is marked on each panel as a purple 'x'. The island coastline is marked with a black contour in each panel.	190
6.10	The hourly-mean vertical profile of the scaled cloud count, \hat{N}_{cld} for our suite of experiments at different grid lengths. These profiles are shown for a) an organised period during the peak island heating, and b) a disorganised period during the last hour of simulation.	197
6.11	The hourly-mean vertical profile of the cloud radius assuming the cloud area represents that of a circle, r_{cld} , for our suite of experiments at different grid lengths. These profiles are shown for a) an organised period during the peak island heating (11AM to 12PM), and b) a disorganised period during the last hour of simulation (11PM to 12AM). The solid line represents the median cloud radius, and the shading is the interquartile range.	199
6.12	As in 6.11, but for the whole daytime period from model hour 6 to 18 rather than for one-hour periods.	201

6.13	The daytime mean updraught (dashed line) and core (solid line) profiles of a) buoyancy excess measured in terms of θ_v , and b) total water excess. These profiles are computed for each experiment and coloured accordingly. See text for definition of “updraught” and “core”.	202
6.14	Time series of the cloud top height distributions for a) the 200 m experiment, b) the 400 m experiment, c) the 800 m experiment, and d) the 1.6 km experiment. The dashed black lines represent the 10th and 90th percentile, the grey shading is the interquartile range, the solid black line is the median, and the red ‘x’s mark the maximum at that time.	204
6.15	An overview of the CT system in the parametrised 1600 m simulation. a) The daytime cloud frequency, b) the equivalent potential temperature anomaly on the lowest model level (2 m), c) the 10 m n component of the wind, and d) the total accumulated precipitation.	205
6.16	Time series of the cloud top height distributions for a) the 1600 m experiment without parametrisation, and b) the 1600 m experiment with a shallow convection parametrisation. Symbology is as in Figure 6.14.	207

List of Tables

2.1	Hourly and sub-hourly ICAO/WMO standard observations taken on the airfield at TXKF. Heights of measurements are reported as either being made ‘above ground level’ (agl), or ‘above sea level’ (asl).	23
2.2	Coefficients used in the Simmons et al. (1999) approximation to the saturation vapour pressure curve. This approximation allows the smooth transition between vapour pressure with respect to water and vapour pressure with respect to ice. The coefficients and intercepts are chosen such that they match observed saturation vapour pressures (e.g. Buck, 1981 ; Alduchov and Eskridge, 1996).	26
3.1	Contingency tables for (a) CT classifications, (b) NT classifications, and (c) OB classifications. Using the CT classification as an example: top left are ‘hits’ where the algorithm and the manual classifications are both CT; top right are ‘false alarms’ where the algorithm classification is CT but the manual classification is not CT; bottom right are ‘Correct Negatives’ where the algorithm and the manual classifications are both not CT; and bottom left are ‘misses’ where the algorithm classification is not CT but the manual classification is CT.	46
3.2	a) Mean 10 m wind speed (m s^{-1}) at Bermuda’s L. F. Wade International Airport by algorithm classification. The number of observations in each classification is included in brackets. b) The percentage of scenes in each classification and the total number of classified scenes by wind direction.	59
4.1	Values of the weighting function, W_{1D} for the different grid spacing, Δx used in Chapters 4 , 5 , and 6 . For each grid spacing, W_{1D} is computed for the case where $z_{\text{turb}} = z_h = 700$ m.	75
4.2	Profiles of the initial conditions used in BOMEX a), and profiles of the large-scale forcing b). Where a dash is inserted at a height in the table, the values can be acquired through linear interpolation in height between adjacent given data points.	89

4.3	Single level or constant in height settings for the standard BOMEX configuration. This includes initial random noise, surface pressure, temperature, fluxes, and roughness lengths.	91
4.4	List of the experiments conducted to test sensitivity of shallow convection to the horizontal grid spacing using variations on the BOMEX case. All experiments have the same number of horizontal grid points (64×64). The Control experiment from Section 4.3 is marked in boldface font.	98
4.5	Profiles of a) the initial conditions for the Control experiment, and b) profiles of the large-scale forcing. Linear interpolation can be used to fill in data where there are dashes in the table.	113
4.6	Single level or constant in height settings for the spin-up simulation used to get the balanced reference state. This includes initial random noise, surface pressure, temperature, fluxes, and roughness lengths. Note that T_{sfc} is the temperature of the air on the lowest model level, and T_* is the sea surface temperature.	114
5.1	Experiment labels and the corresponding peak island sensible heat flux, H_0 , latent heat flux, E_0 , and Bowen ratio, $\beta = H_0/E_0$. The row in boldface text represents the settings of the Control experiment.	147
5.2	The boundary layer height, z_h , the height of the LCL, z_{LCL} , and the CT circulation mass flux, M_c , for the RH experiments. The original Control experiment is also included for reference.	158
5.3	Experiment labels and the corresponding geostrophic forcing, sea surface sensible heat flux, H_{sea} , and sea surface latent heat flux, E_{sea}	159
6.1	Experiment labels and the corresponding number of grid points in the x - and y - directions, N_x and N_y , and the total number of island grid points, N_{island} , for the Grey Zone resolution experiments. The corresponding dimensions for the control experiment are included, but this experiment is not repeated with the new initial conditions.	176

Chapter 1:

Introduction and Motivation

Islands are notorious for plaguing weather forecasters with difficult to predict local weather. Even in cases where the large-scale environment suggests settled and fair weather, a small, flat island can ruin a forecast after being heated in the daytime sun. Throughout the tropics and subtropics, islands regularly trigger the formation of bands of cumulus clouds which stream downwind. These are referred to as ‘cloud trails’ in this thesis.

In Bermuda (one such island) this phenomenon has become part of the local folklore dating back to local involvement in the American Revolutionary War. British forces in the Western Atlantic were enforcing an embargo on exports from the ‘rebel’ Continental Congress, and Bermuda (which relied heavily on food imports from the continent) was suffering as a result. Local sympathisers to the Congress organised to “trade” British gunpowder for food and supplies.

On an August night in 1775, Captain Morgan (a merchant captain associated with the Continental Congress, not the pirate of the previous century) safely made the exchange taking 100 barrels of gunpowder from the British stockpile at the east end of Bermuda - potentially influencing the outcome of the war.

No formal punishment for this treason was ever meted out as there were no forthcoming witnesses. Older Bermudians say that Captain Morgan’s ghost has returned to the scene of the crime whenever a dark-looking cloud looms over the east end of the island, and meteorologists at the Bermuda Weather Service colloquially refer to this cloud as “Morgan’s Cloud” - now knowing that this dark cloud is in fact a cloud trail

viewed from the ground and not a ghost!

This story is referenced in the Bermuda Weather Service’s glossary page on “Morgan’s Cloud” and is mentioned in “A Bermuda Journey” by [Zuill \(1946\)](#).

1.1 Background and Motivation

Cloud trails make for more than just an entertaining historical tale. In extreme cases, the typically shallow, benign convection associated with a cloud trail can deepen becoming thunderstorms associated with hazardous weather for aviation, mariners, and the general public.

What exactly is a cloud trail? From space, these appear as trailing cloud bands which anchor to their parent islands (e.g. red rectangle in [Fig. 1.1](#)). These cloud bands can easily extend more than 100 km downwind of an island ([Nordeen et al., 2001](#)). As implied by the tale of Morgan’s Ghost above, to an observer on the ground these cloud bands appear to have low, dark, and looming cloud bases (e.g. [Fig. 1.1](#) [inset]).

1.1.1 Circulation Generation

During the daytime, the sun heats the land surface more strongly than the sea surface. This is primarily due to the difference in specific heat capacity of the surface material. The sea has a much higher heat capacity than the land. This results in a smaller temperature increase for the same amount of absorbed solar radiation in water than on land. Additionally, heating can be dissipated more effectively in sea water which can move around whereas the land surface is generally motionless. As a result, a land-sea contrast in surface temperatures develops during the day. Without any background wind, this thermal contrast drives an atmospheric circulation with low-level inflow toward the island, rising motion over the island, and a return outflow within some layer aloft. This is known as the sea-breeze circulation and is part of a system that is well described and studied (for a comprehensive review of literature on the sea-breeze system, see [Miller et al., 2003](#)).

We can relate the horizontal land-sea contrast in surface temperatures to the gen-



Figure 1.1 True colour imagery from the Aqua Moderate resolution Imaging Spectroradiometer (MODIS) on the 23rd May 2017 at approximately 1230 Local Time. The islands of Bermuda and the shallow waters of the lagoon to the north of the island are seen in cyan and green colours, the deeper ocean surrounding the island appears as dark blue, and the clouds appear white. A red box highlights the location of the cloud trail at this time. The red star marks the eastern end of Bermuda, near where the Bermuda Weather Service is located. [Inset] Webcam picture of the cloud trail at about the same time as the satellite imagery. Camera is looking toward the north-northeast. Courtesy of the Bermuda Weather Service. The clouds with darker grey cloud bases associated with the cloud trail are in the left half of this picture.

eration of circulation. Starting from the vertical momentum equation, we can show that the more strongly heated land drives stronger upward vertical accelerations than over sea.

The vertical momentum equation can be approximated from the primitive equations:

$$\frac{Dw}{Dt} = -\frac{1}{\rho} \frac{\partial p}{\partial z} - g, \quad (1.1)$$

where, neglecting the Earth's rotation and frictional terms, w is the vertical velocity, ρ is the air density, p is the air pressure, and g is the acceleration due to gravity (Lackman, 2011). We can then linearise about a horizontally homogeneous, time-

invariant background state: $p = \bar{p}(z) + p'$; $\rho = \bar{\rho}(z) + \rho'$; and $w = \bar{w} + w'$, where variables ϕ have background state $\bar{\phi}(z)$, and deviations from the background ϕ' . Assuming that the background state is in hydrostatic balance, $\partial\bar{p}(z)/\partial z = -\bar{\rho}(z)g$, and that $\bar{w} = 0$, equation 1.1 can be re-written as:

$$\frac{Dw}{Dt} = \underbrace{-\frac{1}{\bar{\rho}} \frac{\partial p'}{\partial z}}_{\text{Forced}} \underbrace{-g \frac{\rho'}{\bar{\rho}}}_{\text{Free}}, \quad (1.2)$$

Thus vertical acceleration is the sum of acceleration due to vertical pressure perturbation gradients (i.e. forced ascent), and acceleration due to buoyancy (i.e. free ascent). Above, the anelastic approximation is applied to neglect products of perturbations (i.e. acoustic waves can be neglected by assuming the wind speed is much smaller than the speed of sound, $u \ll c$). Applying the same linearisation to the ideal gas law for moist air ($p = \rho R_d T_v$), the acceleration due to buoyancy becomes:

$$b = -g \frac{\rho'}{\bar{\rho}} \approx g \left(\frac{T'_v}{T_v} - \frac{p'}{\bar{p}} \right). \quad (1.3)$$

where $T_v = T(1 + c_{virt} m_v)$ is the virtual temperature - the temperature required for air containing water vapour to have the same density as dry air, $c_{virt} = (\frac{M_d}{M_v} - 1) \approx 0.61$, M_v is the molar mass of water vapour, M_d is the molar mass of dry air, and m_v is the water vapour mixing ratio. It can be shown (e.g. in Emanuel, 1994) that the pressure perturbation term is negligible so we can write:

$$b \approx g \frac{T'_v}{T_v}, \quad (1.4)$$

where significant condensate (liquid or ice water from cloud and precipitation) can be taken into account as a “hydrometeor loading” which reduces the buoyancy where present.

Now we can show that the more strongly heated land, with a larger T'_v , therefore results in stronger vertical accelerations than over the sea. To show the impact of this land-sea contrast on the circulation, we introduce relative vorticity (e.g. Holton and Hakim, 2013):

$$\boldsymbol{\omega} = \nabla \times \mathbf{u} = \left(\frac{\partial w}{\partial y} - \frac{\partial v}{\partial z} \right) \hat{\mathbf{i}}, \left(\frac{\partial u}{\partial z} - \frac{\partial w}{\partial x} \right) \hat{\mathbf{j}}, \left(\frac{\partial v}{\partial x} - \frac{\partial u}{\partial y} \right) \hat{\mathbf{k}}, \quad (1.5)$$

where the relative vorticity vector, $\boldsymbol{\omega}$, is mathematically the curl of the velocity vector, \mathbf{u} and represents the amount of “spin” in the flow. In cartesian coordinates, v is the velocity component in the y -direction, u is the velocity component in the x -direction, and w is the velocity component in the z -direction. The first two terms on the right are the horizontal relative vorticity, ω_x and ω_y , while the third term is the vertical relative vorticity, ζ .

From equation 1.5, if there is a horizontal gradient in w , there is horizontal vorticity. Now we have shown that a) horizontal temperature contrasts (i.e. gradients) translate to horizontal buoyancy gradients; b) those horizontal buoyancy gradients lead to horizontal vertical velocity gradients; and c) those horizontal velocity gradients are associated with horizontal components of the relative vorticity vector. To more directly link horizontal buoyancy gradients to horizontal relative vorticity, we introduce the relative vorticity tendency equation for an incompressible, inviscid fluid (modified from [Holton and Hakim, 2013](#)):

$$\frac{D\boldsymbol{\omega}}{Dt} = (\boldsymbol{\omega} \cdot \nabla)\mathbf{u} - \boldsymbol{\omega}(\nabla \cdot \mathbf{u}) + \frac{1}{\rho^2} \nabla \rho \times \nabla p, \quad (1.6)$$

where the first two terms on the right hand side represent vortex tilting due to across-flow gradients and vortex stretching due to along-flow gradients. The last term represents a source or sink of vorticity due to differential density and pressure gradients, the “baroclinic” or “solenoidal” term. Expanding in the $\hat{\mathbf{i}}$ direction, the vorticity equation in 1.6 becomes:

$$\frac{D\omega_x}{Dt} = \underbrace{-\omega_x \left(\frac{\partial v}{\partial y} + \frac{\partial w}{\partial z} \right)}_{\text{Vortex Stretching}} - \underbrace{\left(\frac{\partial u}{\partial y} \frac{\partial w}{\partial x} - \frac{\partial u}{\partial z} \frac{\partial v}{\partial x} \right)}_{\text{Vortex Tilting}} + \underbrace{\frac{1}{\rho^2} \left(\frac{\partial p}{\partial z} \frac{\partial \rho}{\partial y} - \frac{\partial p}{\partial y} \frac{\partial \rho}{\partial z} \right)}_{\text{Baroclinic}}, \quad (1.7)$$

where all terms are as defined above, and there is a similar formulation for the $\hat{\mathbf{j}}$ and $\hat{\mathbf{k}}$ directions. Introducing the vertical component of the Earth’s rotation would lead to

an additional source/sink term $f \frac{\partial u}{\partial z}$, where f is the coriolis parameter.

The baroclinic term can be re-written using the ideal gas law for moist air to replace density with virtual temperature as follows:

$$\frac{1}{\rho^2} \left(\frac{\partial p}{\partial z} \frac{\partial \rho}{\partial y} - \frac{\partial p}{\partial y} \frac{\partial \rho}{\partial z} \right) = -\frac{R_d}{p} \left(\frac{\partial p}{\partial z} \frac{\partial T_v}{\partial y} - \frac{\partial p}{\partial y} \frac{\partial T_v}{\partial z} \right), \quad (1.8)$$

where all variables are as previously defined. We assume that horizontal pressure gradients on the scale of small islands are small and therefore the second term in the brackets can be neglected. Then, given that pressure decreases with height, we see that if T_v increases northward (i.e. south of a heated island), the island acts to provide a source of positive ω_x , where this component of vorticity is positive pointing to the east.

From here, horizontal buoyancy gradients are linked to vorticity generation (and therefore circulation generation) at coastlines. This circulation is part of the sea breeze system. This can be expressed mathematically as $\frac{\partial \omega_x}{\partial t} \propto \frac{\partial b}{\partial y}$ and this scaling is explored in more detail in Chapter 5.

1.1.2 Atmospheric Convection

We've now introduced concepts of buoyancy and circulation, both of which are important for understanding atmospheric convection. It is useful to introduce some further important concepts that will recur during this thesis. We do so by means of "parcel theory" in which a "parcel" represents a thermodynamic system for which energy and moisture can be budgeted. A common definition of a parcel is to assume that it takes the properties of moist air at the surface (i.e. has the pressure, temperature, mixing ratio, and therefore density of surface air), and that it does not exchange mass or energy with the surroundings as it rises.

Such a parcel's buoyancy is then defined in terms of the difference in virtual temperature between the parcel and its environment:

$$b = g \frac{T_{v,p} - T_{v,e}}{T_{v,e}}, \quad (1.9)$$

where the parcel is positively buoyant if the parcel virtual temperature, $T_{v,p}$, is greater than the environmental virtual temperature, $T_{v,e}$. Equation 1.9 is equivalent to 1.4, where $T'_v = T_{v,p} - T_{v,e}$. If $T_{v,p} < T_{v,e}$, then the parcel is negatively buoyant, and if $T_{v,p} = T_{v,e}$, then the parcel is neutrally buoyant. As a parcel rises, the temperature of the parcel decreases, but the mixing ratio is roughly conserved. The rate at which a lifted parcel's temperature decreases is known as the “dry adiabatic lapse rate”:

$$\Gamma_d = -\frac{dT}{dz} = \frac{g}{c_{pd}}, \quad (1.10)$$

where c_{pd} is the specific heat capacity of dry air at constant pressure and Γ_d is the dry adiabatic lapse rate (about 10 K km^{-1}). We can use Γ_d to characterise the static stability of dry layers to lifted parcels by defining an environmental lapse rate, $\Gamma_e = -\frac{dT_e}{dz}$. If $\Gamma_e > \Gamma_d$ in a dry layer, then that layer is “absolutely unstable” (i.e. any lifted parcel will remain positively buoyant and continue to ascend freely through such a layer). If $\Gamma_e = \Gamma_d$, then the layer is “neutrally stratified”. Finally, if $\Gamma_e < \Gamma_d$, then that layer is “stably stratified”. Meanwhile, the cooling associated with parcel ascent results in a decrease in the saturation water vapour mixing ratio, $m_{v,s}$, through the Clausius-Clapeyron relation (Holton and Hakim, 2013). With m_v approximately constant, the relative humidity, $\text{RH} = m_v/m_{v,s}$, of our parcel therefore increases as the parcel rises.

At some height, the parcel rises and cools enough to become saturated ($\text{RH} \geq 100\%$). The height at which a parcel becomes saturated is known as its “lifting condensation level” or LCL. If the parcel continues to rise and cool, condensation occurs and the parcel becomes cloudy.

As condensation occurs, it releases latent heat. This partially offsets the cooling due to adiabatic expansion, and the parcel now cools at a reduced rate. If we assume the total condensate is retained by the parcel, this new rate can be given by the reversible moist adiabatic lapse rate:

$$\Gamma_m = g \frac{1 + \frac{L_v m_v}{R_d T}}{c_{cpd} + \frac{L_v^2 m_v \epsilon}{R_d T^2}}, \quad (1.11)$$

where Γ_m is the moist adiabatic lapse rate, L_v is the latent heat of vaporisation, and $\varepsilon = M_v/M_d \approx 0.622$ is the ratio of the molar masses of water vapour and dry air. Now we can discuss moist static stability of a layer to a rising parcel. A layer is “conditionally unstable” if $\Gamma_d > \Gamma_e > \Gamma_m$ (i.e. it is unstable to moist ascent but stable to dry ascent). If $\Gamma_e = \Gamma_m$, then that layer has “moist-neutral” stability. Finally, at the extremes, if $\Gamma_e < \Gamma_m$ then the layer is “absolutely stable”, and if the layer is saturated and $\Gamma_e > \Gamma_m$, then it is “moist absolutely unstable” layer.

The parcel might then ascend to a level at which it is positively buoyant and then will continue to rise due to its own buoyancy. This level is known as the “level of free convection”, or LFC. Eventually, the rising parcel will reach a height at which it becomes neutrally buoyant, this is known as the “level of neutral buoyancy” (LNB). Often this is in the vicinity of a temperature inversion where the environmental temperature increases with height through a layer.

Atmospheric convection is often broadly categorised by two modes: “shallow” and “deep”. This thesis focuses primarily on shallow convection which, in the tropics, refers to buoyant motions with origins near the surface that result in cumulus clouds which are typically confined below the freezing level. As a result, an individual shallow convective cell is incapable of producing extreme rain rates (Byers and Hall, 1955). Conversely, deep convection refers to any other surface-based convection - and is often used to refer more specifically to convection which has a LNB near the tropopause. Deep convection and shallow convection therefore have different vertical profiles of heating and moistening due to their vertical extent and the type and amount of condensate.

From the perspective of shallow cumulus, we can introduce “conserved variables” which do not change following a parcel. For energy, we introduce the “liquid potential temperature,” θ_l , and for moisture we introduce the total water mixing ratio, $m_T = m_v + m_{cl}$, where m_{cl} is the cloud liquid mixing ratio.

In a dry atmosphere, the effects of adiabatic cooling can be described using the potential temperature. This is the temperature an air parcel would have if it were moved vertically and dry adiabatically from its actual temperature and pressure to a reference pressure. This is expressed as:

$$\theta = T \left(\frac{p_0}{p} \right)^{R_d/c_{pd}}, \quad (1.12)$$

where θ is the potential temperature, $p_0 = 1000$ hPa is the reference pressure, and all other variables are as previously defined. The liquid potential temperature is a variant on this to account for the latent heat released through condensation and cloud formation, and is commonly given as:

$$\theta_l = \theta - \frac{\theta}{T} \frac{L_v}{c_{pd}} m_{cl}, \quad (1.13)$$

following [Betts \(1973\)](#) where θ_l is the liquid potential temperature, and all other variables are as previously defined. Cloudy air therefore has lower θ_l than the unsaturated environmental air which has $\theta_l = \theta$.

Through the lens of parcel theory, a shallow cumulus cloud might start life near the surface as a parcel with a small positive buoyancy perturbation. The parcel then ascends, cooling dry adiabatically until it reaches its LCL and becomes saturated. Latent heat release causes the parcel to cool less quickly with height at the moist adiabatic lapse rate as cloud forms. This parcel might then encounter a temperature inversion (such as the trade inversion in the Tropical Atlantic), cross its LNB, and become negatively buoyant and eventually cease rising. It should be noted that real convection might be better described as an entraining plume of air rather than a parcel, where the entrainment of cooler, drier, environmental air offsets the positive buoyancy of such a plume ([de Rooy et al., 2013](#)).

1.1.3 Island Convection in Literature

If there is weak or no background flow, the classical sea-breeze develops over islands. The sea-breeze provides a mechanical lifting mechanism which can “trigger” the initiation of convection. Over a heated circular island, the sea-breeze circulation is characterised by inward flow toward the island from all directions at low-levels, ascent over the island, and outward flow in all directions aloft. When background flow is introduced, the sea-breeze circulation is advected and tilted downwind becoming what

Estoque and Bhumralkar (1969) and later Savijarvi and Matthews (2004) refer to as the “heat island circulation”. In the along-wind direction, the leeward sea-breeze cell is advected downwind of the island where it weakens or dissipates, while the windward sea-breeze cell is advected over or just downwind of the island where it becomes phase-locked to the heated island. Furthermore, both leeward and windward sea-breeze cells are tilted downwind with height in the along flow direction. Conversely, in the across flow direction, the in-up-out circulation of the traditional sea-breeze is preserved. This across-flow circulation is continuously generated over the island and advected downwind (Fig. 1.2 for a visualisation of this circulation component). The upward motions in the across-flow circulation can trigger convection. The strength of background flow required for this transition from sea-breeze to heat-island circulation depends primarily on the island size, but may also depend on the strength of the island surface heating (Savijarvi and Matthews, 2004).

The band of ascent associated with the heat island circulation triggers convection in an organised band or cloud trail. This feature looks similar to an individual “cloud street” which are commonly found associated with horizontal convective rolls. However, horizontal convective roll circulations are the result of sheared flow instability, over a uniformly heated surface and their morphology may depend on the boundary layer stratification (Mason and Sykes, 1980, 1982). Conversely, cloud trails and associated circulations require a localised inhomogeneity in surface heating (i.e. a land-sea contrast on an island), and their morphology is therefore more dependent on island geometry.

Savijarvi and Matthews (2004) also noted that as the size of a heated island increased, the strength of the resulting circulation also increased. For islands with diameter greater than about 40 km, the circulation strength increased to be equivalent to that at a continental coastline. This is another piece of evidence demonstrating that small islands behave differently to larger islands. This difference in behaviour is repeated across several measures of convective intensity including lightning rates (Williams et al., 2004) and rainfall rates (Robinson et al., 2011; Sobel et al., 2011).

Due to their size, smaller islands generate a smaller volume of “island affected air”. Williams et al. (2004) provided a scale analysis for a typical deep convective plume and

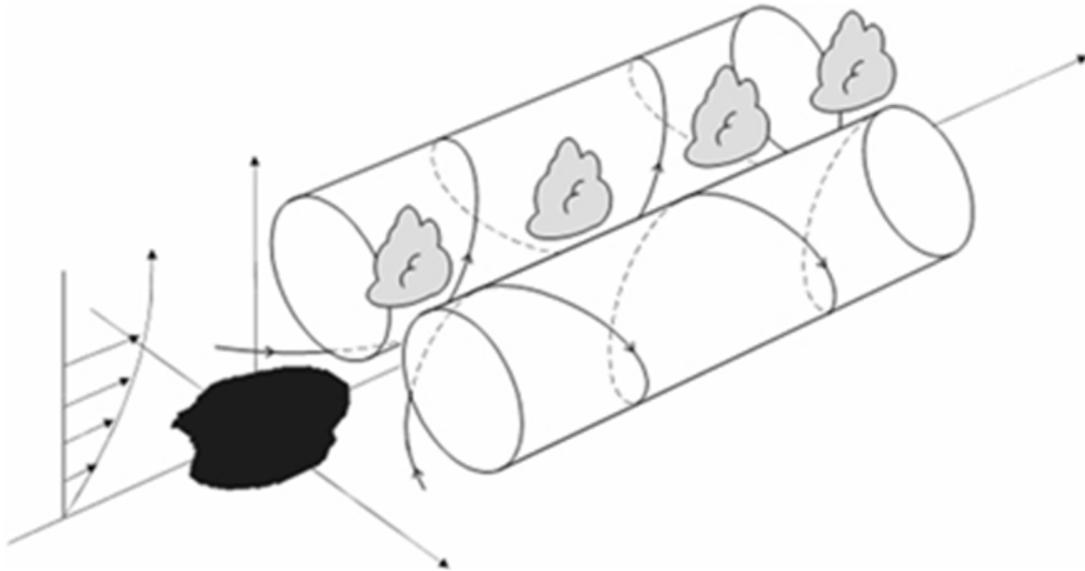


Figure 1.2 A schematic of the cloud trail circulation based on observations from the island of Nauru in [Matthews et al. \(2007\)](#). Schematic shows the island in black, streamlines of the full circulation. Areas where clouds preferentially form are indicated by small cloud shapes. The source of this material is the [COMET Website at <http://meted.ucar.edu/>](http://meted.ucar.edu/) of the University Corporation for Atmospheric Research (UCAR), sponsored in part through cooperative agreement(s) with the National Oceanic and Atmospheric Administration (NOAA), U.S. Department of Commerce (DOC). ©1997-2017 University Corporation for Atmospheric Research. All Rights Reserved.

compared the volume of air ingested by a plume over a typical cloud lifetime to the volume of air contained within the boundary layer over an island. For islands with area smaller than about 100 km^2 , the volume required by the convective plume exceeds the volume of island affected air and therefore the plume must ingest some oceanic air over its lifetime. This has implications for the makeup of aerosols within the cloud, and therefore the microphysics of the cloud, and its subsequent development.

Complex island geometries and mountainous islands introduce superimposed circulations on top of the heat island or sea-breeze circulation. [Crook and Tucker \(2005\)](#) showed that flow along the long axis of an ellipse-shaped island maximised the heating gradient and that the effect of heated terrain in general acted in the same sense as heating a flat island alone. Further, [Kirshbaum and Fairman \(2015\)](#) showed that a simple scaling of the heat flux for elevated heating can be used to estimate the strength of resulting island circulations for simulations over the Lesser Antilles. The scaling presented in [Kirshbaum and Fairman \(2015\)](#) also assumed ellipse-shaped islands and the

scaling increased the strength of island circulations for flow along the long axis of the island. Mountains also have a significant mechanical influence on the flow. Flow is either forced to ascend the slope of a mountain, or the flow splits and passes around the mountain. The exact behaviour of flow impinging on an isolated mountain can be split into regimes based on the non-dimensional mountain height which is related to the stability of the layer impinging on the mountain represented by the Brunt-Vaisala Buoyancy Frequency, N , the height of the mountain, H , and the speed of the flow, U . Combined, the non-dimensional mountain height, $H_m = NH/U$, can be used to summarise this behaviour. For $H_m > 1$ the layer is very stable, the mountain is very high, and/or the flow is weak. This allows more flow splitting around the mountain. For $H_m < 1$, the layer is less stable, the mountain not as high, and/or the flow is stronger. This results in more flow over the mountain. In this thesis, we focus on “flat” islands for which H_m is much less than one or zero in ideal cases.

As previously alluded to, cloud trails occur downwind of small heated islands around the world. [Nordeen et al. \(2001\)](#) quantified their occurrence at Nauru by manually classifying visible satellite imagery every hour at Nauru for a year. On that island, they found that cloud trails have a diurnal cycle in their occurrence and their life cycle, and grow in length at a rate consistent with advection. Meanwhile, a recent study based in the Lesser Antilles Islands also used a manual satellite-methodology to link thermodynamic profiles to cloud trail occurrence ([Kirshbaum and Fairman, 2015](#)).

Meanwhile, many case studies of cloud trails have been conducted for various purposes. For example early work by [Malkus and Stern \(1953\)](#) on Nantucket used aircraft observations to describe the stably stratified flow response as akin to a mountain wave response. Aircraft and targeted radiosonde observations revealed that the aerosol and microphysical properties of clouds within cloud trail convection matches those properties in atmospheric convection over the ocean despite origins over the island and the presence of a driving mesoscale circulation ([Malkus, 1963](#)). That associated circulation comes with an accompanying thermal perturbation ([Bhumralkar, 1973a](#); [Garstang et al., 1975](#); [Matthews et al., 2007](#)). Heated oceanic islands have been compared to urban “heat islands” ([Garstang et al., 1975](#)), and have some similarities to lake-effect snow bands ([Hsu, 1987a,b](#)) for their associated mesoscale circulations.

On Nauru, the impact of these island clouds on surface radiation has been quantified (McFarlane et al., 2005; Long and McFarlane, 2012), and multi-year studies in which cloud-trail-affected times are identified show some variability in their occurrence tied to changes in the large-scale environment due to the El Niño-Southern Oscillation’s impact on large-scale cloudiness in the tropical Pacific.

There are few long-term climatologies of cloud trail occurrence, and all limited to within the deep tropics. A further limitation of past satellite-based climatologies of cloud trails is that they all employ manual classification of imagery. A simple algorithm could be constructed to automate the same task and carry it out quickly and objectively. Furthermore, study of the downwind structure of the cloud trail system has been limited to case studies which are unable to adequately quantify the sensitivity of the cloud trail system to variations in the environment or island forcing.

1.1.4 Island Convection in Numerical Weather Prediction

Cloud trails provide a local forecasting challenge. Modern numerical weather prediction (NWP) tools can now provide guidance on the large-scale atmospheric state multiple days in advance with reasonably high forecaster confidence (Krishnamurti et al., 2003; Stern and Davidson, 2015). This allows conceptual models of how cloud trails work to be applied to the model guidance to give several days’ notice for locally impactful weather associated with cloud trails such as heavy rain or thunderstorms in some cases. A clearer link between supporting environments and cloud trail occurrence would help refine these conceptual models.

Furthermore, improvements in NWP have recently been focussed on increasing horizontal resolution and this has allowed a representation of small islands to emerge alongside a more direct numerical representation of convection. Historically, idealised simulations such as the 2-D experiments of Estoque and Bhumralkar (1969) and Bhumralkar (1973b) have provided insight into the circulation formation related to small islands. Meanwhile, more recent and sophisticated numerical approaches as in Kirshbaum and Wang (2014) and Kirshbaum and Fairman (2015) have been performed on mid-sized islands (Dominica). A full characterisation of the cloud trail system for a

small, flat island has yet to appear in literature.

Many of the high-resolution and Large Eddy Simulations (LES) relating to shallow convection have focussed on horizontally homogeneous cases (e.g [Siebesma et al., 2003](#); [Abel and Shipway, 2007](#)). These LES experiments have offered invaluable insight for the development of parametrisations for shallow convection. However, only recently have cases with heterogeneous surface boundary conditions been tested (e.g. for deep convection in [Harvey et al., 2019](#)). Introducing a small island in an LES-like framework has potential to yield new insight into the behaviour of organised shallow convection through the investigation of the case in which there are organising mesoscale circulations present. Small islands offer a predictable source of one form of mesoscale circulation.

Shallow cumuli have been shown to precondition the atmosphere for deep convection by moistening the lower free-troposphere where they occur. The increased humidity in the free troposphere buffers the evaporative effect of entraining environmental air into a nascent deep convective updraught ([Rio et al., 2009](#); [Wu et al., 2009](#); [Gentine et al., 2013](#)). The representation of shallow convection in NWP is important for setting the vertical temperature and moisture profiles, particularly in the trade-wind region over large parts of the tropical oceans ([Betts and Miller, 1986](#); [Zhang and McFarlane, 1995](#); [Johnson and Lin, 1997](#)). Furthermore the impact on the temperature structure can have serious knock-on effects for the prediction of impactful weather, for example tropical cyclone track and intensity ([Torn and Davis, 2012](#)). Furthermore, continuously forming shallow clouds in the trade wind region throughout the global tropics influences the radiation budget ([Medeiros and Nuijens, 2016](#)). Uncertainty in the representation of shallow cloud, and therefore the resulting cloud-radiation interaction, in coarser resolution NWP contributes to the overall uncertainty in estimates of climate sensitivity from climate models - tools which are becoming increasingly relevant as the climate changes.

In many weather models and most climate models, model horizontal grid spacing is much bigger than the horizontal scales of both shallow and deep convection. As a result of the increase in horizontal resolution (i.e. smaller grid spacing) in current and future NWP implementations, many of the assumptions underpinning traditional

parametrisation schemes are violated and the schemes may no longer behave as designed (Arakawa and Lamb, 1977; Arakawa and Wu, 2013). The ParaCon¹ project was born out of this need for new and fit-for-purpose convection schemes for current and future NWP. Some of the goals of this work closely parallel the goals of the ParaCon project. Climate models currently use grid lengths of $\mathcal{O}(100 \text{ km})$ (e.g. CMIP6 Sellar et al., 2019); global weather models use $\mathcal{O}(10 \text{ km})$ grid lengths; and some regional models such as the Met Office’s “United Kingdom Variable Resolution” model (UKV) and the experimental Met Office “London Model” use kilometre and sub-kilometre grid lengths (Bush et al., 2019). As a result, parametrisation schemes are used to predict the heating and moistening tendencies which would arise on the resolved scales if convection were present in climate and weather models, while some regional models attempt to explicitly represent convection. In addition to improving the representation of the vertical thermodynamic structure (Torn and Davis, 2012), including a convection scheme may also prevent significant numerical issues known as “grid-point storms”.

At climate model resolution (and even at regional weather model resolution), both shallow and deep convection are not well resolved at all. In fact, at climate model grid spacing, many small islands and their associated shallow cloud trail convection can fit inside a single grid cell (e.g. Fig. 1.3). Here, parametrisations for convection are still the norm. At the other extreme, at regional model resolution, convection can be partially resolved and some modelling centres choose not to employ a parametrisation. When a model’s horizontal grid spacing is close to the horizontal scales of convection, some aspects of convective motions can be supported by the grid and numerical implementation of the physics (i.e. the convection is partially resolved), and these models are termed “convection permitting”. An individual deep convective updraught might span several grid cells for a model with horizontal grid spacing of $\mathcal{O}(1 \text{ km})$, while shallow convection occurs on smaller horizontal scales of a few hundred metres (Ludlam, 1980; Kirshbaum and Grant, 2012) and would be poorly represented in the same model. The “grey zone” problem may therefore need to be treated differently for shallow and deep convection. For instance, tests performed on parametrisation schemes in the UM con-

¹ParaCon is a five-year programme of work jointly funded between the UK Natural Environment Research Council (NERC) and the Met Office with the aim of significantly improving the representation of convection across model scales ranging from $\mathcal{O}(1 \text{ km})$ to $\mathcal{O}(100 \text{ km})$.

vection scheme documentation (Stratton et al., 2017) suggest that a parametrisation scheme for shallow cumulus at “deep convection grey zone” resolutions might improve the representation of convective intensity (and therefore the hydrological cycle). In $\mathcal{O}(1$ km) models, small islands can be resolved, but we can expect their downstream impact to be poorly represented. Meanwhile, at intermediate $\mathcal{O}(10$ km) resolutions, the island might not be resolved at all, the expected circulation would span several grid cells, and the expected convection associated with that circulation would be poorly represented..

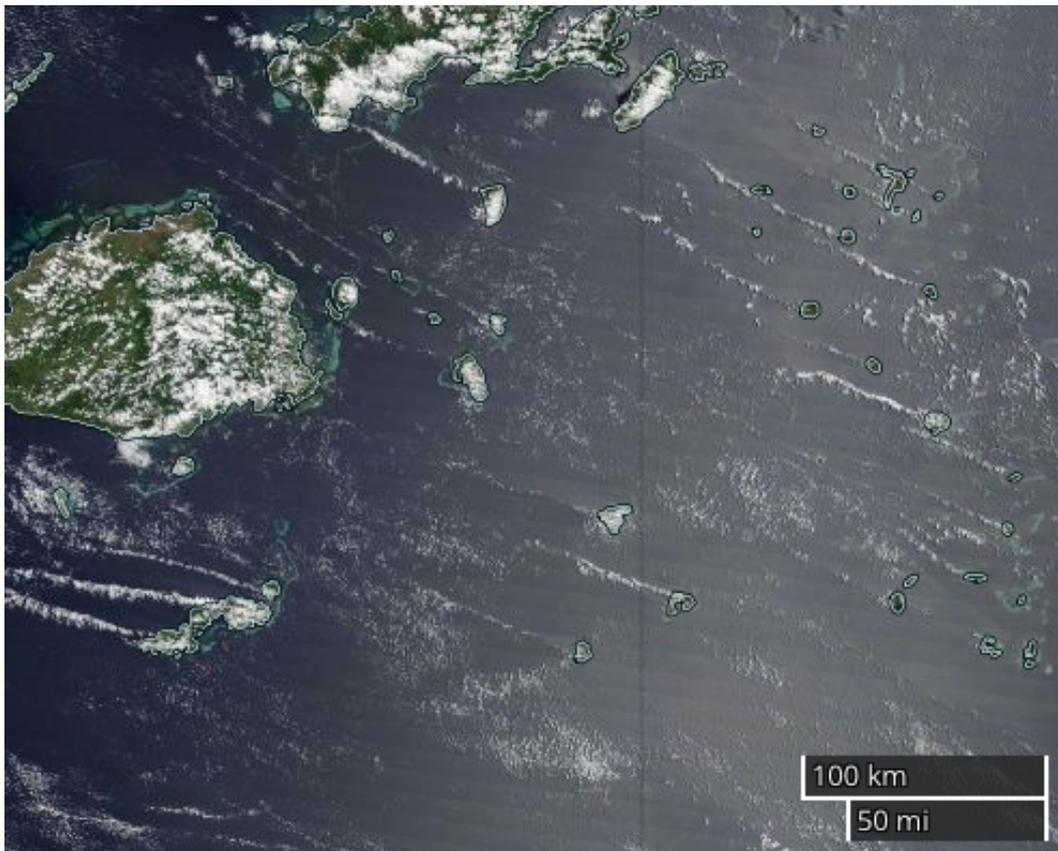


Figure 1.3 An example of many nearby and interacting cloud trails (along with other forms of island and non-island convection) as seen from Terra MODIS True Color imagery around Fiji on the 21st October 2019.

The numerical work that we set out to complete in this thesis expands upon existing work using high-resolution “truth” simulations to further our understanding of shallow convection from a new perspective. Where previously such simulations were conducted for sea-only conditions, we now look at what happens when an island is introduced.

1.2 Thesis Aims and Structure

In the above section, we highlight that past work on cloud trails has focussed on tropical islands and consists of single event field studies and climatologies which span relatively short periods of record. No work has employed automated cloud trail identification or has constructed a climatology longer than two years. Furthermore, no work has been published related to cloud trail occurrence or behaviour at Bermuda where the annual cycle is stronger and the island has a more complex geometry - past work primarily focussed on Nauru and Dominica with scattered case studies which include some more mountainous islands. Moreover, Bermuda is not traditionally part of the trade wind region on which the more focussed studies have centred. Bermuda's geographic isolation makes it an ideal case island for studying cloud trails as any surrounding surface heterogeneities are very far away. As a result, we can hypothesise what the annual and diurnal cycle of cloud trail occurrence at Bermuda could be, but little is known regarding the links between occurrence and the local or large-scale environmental conditions.

In addition, we note that small islands such as Bermuda are starting to be resolved in today's high-resolution weather forecast models. The vast body of literature indicates that the model resolution used is not sufficient to adequately represent all of the important convective processes explicitly, and that parametrisation schemes are required to do so. Research into the simulated small heated island problem exists either in 2-D frameworks, or highly idealised and coarse resolution 3-D models. Furthermore, parametrisations for shallow convection are based on assumptions which do not account for sub-grid scale organisation (e.g. a sub-grid cloud trail). No attempt at a thorough investigation of cloud trails in a more complex numerical model has yet been made, and research into the role of organisation in parametrisation is needed. Furthermore, cloud trails can be reliably introduced into simulations because they form in a given location through a known mechanism, whereas it is more difficult to know where other forms of forcing arise because they appear to arise from otherwise horizontally uniform environments (e.g. horizontal convective rolls).

We identify the following overarching research questions:

1. How does CT occurrence vary at Bermuda (a small isolated island known to have CTs), including diurnally, seasonally, and annually; and how do larger-scale environmental conditions modulate this occurrence?
2. How does the strength and structure of CTs respond to changes in island surface fluxes, environmental RH, and background flow for simulations in which shallow convection is well-resolved?
3. Is the numerical representation of CTs degraded for explicit convection at “grey zone” grid lengths; which properties are restored via parametrisation?

Research aimed at addressing each of these research questions is presented in this thesis. However, a complete analysis of these topics is beyond of the scope of this project.

First we introduce the data used throughout the thesis to answer the above overarching research questions in Chapter 2. This includes remotely sensed and in-situ observations of cloud trails, and re-analysis data which is used to investigate the larger-scale environment.

In Chapter 3 an automated and objective approach to constructing our climatology is outlined, tested, and used in this chapter. We then describe and discuss the resulting annual and diurnal cycle of cloud trail occurrence at Bermuda. The occurrence is then put into context of the large-scale and local environments as described through in-situ observations and re-analysis data. This work addresses our first research aim regarding observed cloud trail occurrence in Bermuda.

Next, we introduce the Met Office Unified Model in its idealised configuration (the UM) as a numerical tool for the evaluation of simulated shallow cumulus for different resolutions. Past numerical experiments on small island convection have been overly simple or at too coarse resolution to properly study shallow convection. Similarly, the more complex high resolution models employed in the literature have mainly been applied to uniform surface boundary conditions, or to larger islands and peninsulas. In this thesis we use a shallow non-precipitating cumulus case to investigate the horizontal grid spacing requirements and the behaviour of the UM when shallow convection is explicitly simulated. We use this case to introduce modifications for use in our small

island case. This modified framework is used in Chapters 5 and 6 to answer our last two main research questions.

In Chapter 5, we perform a control experiment and variants on this experiment to more completely describe the cloud trail system. These variant experiments help to address the relationship between the cloud trail system and the strength of the surface heating, environmental relative humidity, and background wind speed - questions for which the observations used in Chapter 3 cannot provide satisfactory answers.

Finally, cloud trails are viewed through the lens of the parametrisation community in Chapter 6, where model grid lengths are increased. Simulated cloud statistics are explored to characterise how a model without a parametrisation scheme at grey zone scales behaves in a cloud trail scenario. Chapter 6 also details how an existing shallow convection scheme behaves in that scenario, and discusses possible future parametrisation developments.

Finally, the novel contributions from the above chapters are summarised in Chapter 7 where potential avenues of future work are also discussed.

Chapter 2:

Data Sources

The present chapter describes the in-situ observations and other data used in the present thesis. The majority of this data is used in Chapter 3 where a climatology of CT occurrence and a morphological analysis of CTs is conducted. The results of Chapter 3, combined with the data described herein, are then used in modelling experiments in Chapters 4, 5, and 6.

2.1 Observations

2.1.1 Bermuda

This study focuses on Bermuda as a model island for examining the role of small, real-world surface inhomogeneities on the generation and organisation of convection. Bermuda is a small archipelago covering about 54 km² (CIA, 2017). Located near 32.3 °N, 64.8 °W, Bermuda is more than 1000 km away from the nearest land.

Theory and observations both indicate that significant orography (both upstream, and locally) can play a role in triggering, organising, or disrupting the organisation of convection. The isolation and size of Bermuda is assumed to minimise the role of effects from both upstream and local orography, therefore we expect the primary local influence on weather to be from Bermuda itself - and the land-sea thermal contrast (in the absence of large-scale weather systems).

Bermuda is comprised of low-rising hills. These hills make up the fish hook-shaped

island which is oriented from southwest to northeast. Bermuda's highest elevation above sea level is about 76 m (CIA, 2017). The geometry and terrain of Bermuda can be seen in Fig. 2.1

According to Williams et al. (2004) Bermuda would constitute a 'small' island. Where convection over small islands would typically need to ingest oceanic air over the convective lifetime. This makes small island convection more 'oceanic', and large island convection more 'continental'.

As Bermuda is the chosen field site, a combination of in-situ observations (Section 2.1.2) from Bermuda and remotely sensed observations (Section 2.1.3) centred on Bermuda are used in the present study. These are then supplemented by re-analysis data (Section 2.1.5) to give a more complete overview of the large-scale environment.

2.1.2 In-Situ Observations

As the national meteorological service for Bermuda, the Bermuda Weather Service (BWS) provides a full suite of conventional observations (Table 2.1). These are taken at the L. F. Wade International Airport (TXKF) to International Civil Aviation Organisation (ICAO) and World Meteorological Organisation (WMO) standards.

This study uses data from 2012 to 2016 inclusive. Sensors are located around the airfield at TXKF and include observations detailed in Table 2.1. These observations are taken regularly five minutes before the hour, every hour. Additional observations are made during changeable weather to comply with local ICAO regulations to provide aviation support.

Temperature, Relative Humidity, and Sea Level Pressure data are accessed via NOAA's National Climatic Data Centre - National Centres for Environmental Information (NCDC-NCEI) as decoded METeorological Aviation Report (METAR). The wind speed and direction are accessed directly from the BWS (personal communication, 2019).

The L. F. Wade International airport has one runway on which aircraft can land in either direction. The BWS therefore maintains anemometers on both ends of the runway. Standard operating procedure is to switch which runway and therefore which

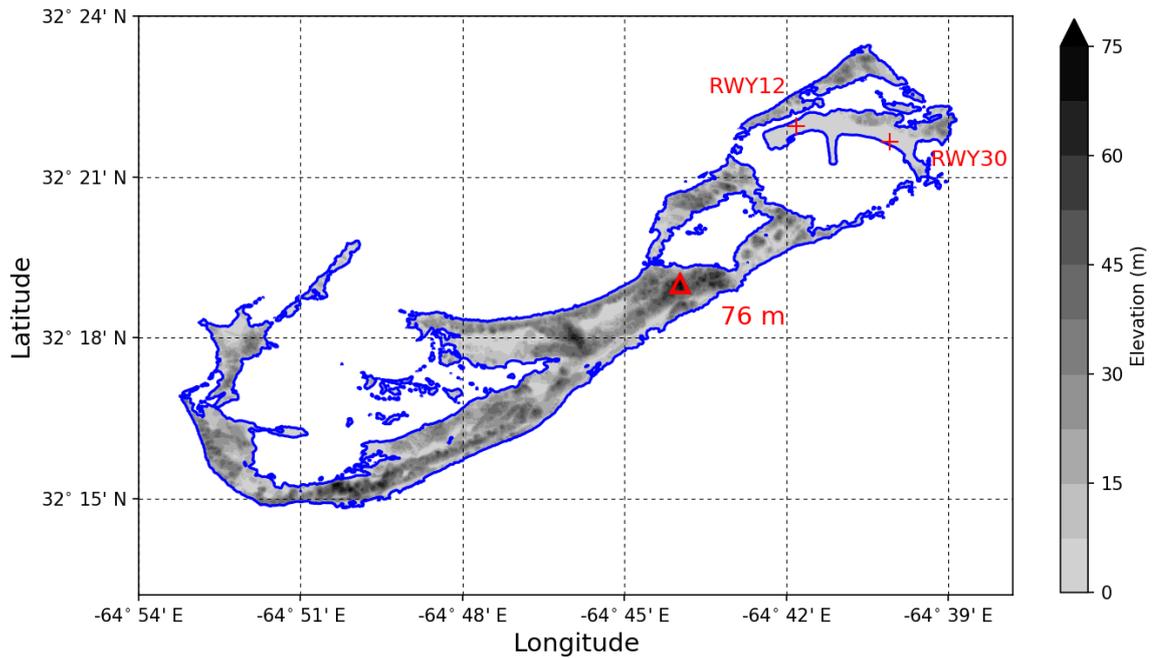


Figure 2.1 A Digital Elevation Model of Bermuda based on data from [Sutherland et al. \(2013\)](#). Data are at a horizontal grid spacing of 1 arc-second (about 30 m). The coast, as defined by the 0 m above mean sea level contour, is outlined in blue. The maximum elevation from this data set is also included as a red triangle with the elevation marked. The location of the anemometers at runway 12 (RWY12) and runway 30 (RWY30) are marked and labelled as red crosses. The anemometers mark each end of the runway at the L. F. Wade International Airport.

Table 2.1 Hourly and sub-hourly ICAO/WMO standard observations taken on the airfield at TXKF. Heights of measurements are reported as either being made ‘above ground level’ (agl), or ‘above sea level’ (asl).

Observation	Measurement Details	Source
Temperature	1.5 m (agl)	NCDC
Relative Humidity	1.5 m (agl)	NCDC
Wind Speed	10 m (agl), 10 min	BWS
Wind Direction	10 m (agl), 10 min	BWS
Sea Level Pressure	Adjusted from 3 m asl	NCDC

anemometer is in use based on the wind direction. This introduced biases in the reported wind speed and direction as found in the METAR. Wind observations on runway 30 are used because there are fewer periods of missing data.

The BWS also takes upper air observations using radiosondes. These include temperature, relative humidity, pressure, wind speed, and wind direction. The radiosonde data are accessed via NCDC’s Integrated Global Radiosonde Archive (IGRA) version 2 (e.g. [Durre et al., 2006, 2008](#)). The data are reported on WMO designated standard pressure levels. Additional intermediate pressure levels are reported so that linear interpolation can be used between levels. This property of the reported data is exploited to linearly interpolate data to a common 5-hPa resolution. This allows comparison and composites of radiosondes between days with different reported levels.

Radiosondes are regularly launched daily at 1200 UTC (0900 Atlantic Daylight Time (ADT)). Additional radiosondes are launched at 0000 UTC, but these are not regular during the period of interest. Therefore, in the present study, only the 1200 UTC radiosondes are used. Temperature and pressure are used to calculate the potential temperature following equation [1.12](#).

Similarly, the temperature, pressure, and relative humidity are used to calculate the specific humidity. A method which computes the saturation vapour pressure with respect to water for temperatures greater than 0 °C, and transitions to with respect to ice for temperatures below -23 °C, following [Simmons et al. \(1999\)](#) is used. This saturation vapour pressure is then converted to the actual vapour pressure using the relative humidity, and then converted to a specific humidity using the pressure:

$$e_s(T) = e_{si}(T) + [e_{sw}(T) - e_{si}(T)] \left(\frac{T - T_i}{T_0 - T_i} \right)^2 \quad (2.1a)$$

$$e_{sx}(T) = e_{s0} \exp \left(a_x \left(\frac{T - T_0}{T - b_x} \right) \right) \quad (2.1b)$$

$$e(T) = \frac{RH}{100} e_s(T) \quad (2.1c)$$

$$q(T, p) = \frac{\varepsilon e(T)}{p + e(T)(\varepsilon - 1)} \quad (2.1d)$$

where $e_s(T)$ is the saturation vapour pressure in hPa as a function of temperature, T ; $e_{si}(T)$ and $e_{sw}(T)$ are the saturation vapour pressures with respect to ice and water

respectively. $T_i = 250.16$ K is the temperature threshold below which $e_s(T) = e_{si}(T)$, and $T_0 = 273.16$ K is the reference temperature, above which $e_s(T) = e_{sw}(T)$. $e_{s0} = e_s(T_0)$ is the reference saturation vapour pressure at the reference temperature.

The saturation vapour pressures with respect to a given water phase x are determined by a pair of coefficients a_x and b_x which are given in Table 2.2 where subscript w refers to with respect to water, and subscript i with respect to ice. The final vapour pressure is found by multiplying the saturation vapour pressure by the relative humidity, $RH/100$, where RH is a percentage. Finally, the specific humidity is found in equation 2.1d.

Using the method outlined in Bolton (1980), the pressure of the lifting condensation level (LCL) can be computed from the surface temperature, dewpoint temperature, and pressure. In Chapter 3 we compute the LCL using the observations at the lowest height reported in radiosondes as follows:

$$T_{LCL} = \frac{1}{\frac{1}{T_d - 56} + \frac{\ln\left(\frac{T_{sfc}}{T_d}\right)}{800}} + 56 \quad (2.2a)$$

$$p_{LCL} = p_{sfc} \left(\frac{T_{LCL}}{T_{sfc}} \right)^{\frac{c_{pd}}{R_d}} \quad (2.2b)$$

$$T_d = \left[\left(\frac{1}{T_0} \right) - \left(\frac{R_v}{L_v} \right) \ln \left(\frac{e}{e_{s0}} \right) \right]^{-1} \quad (2.2c)$$

where p_{LCL} is the pressure of the LCL, T_{LCL} is the temperature of a parcel lifted to its LCL in Kelvin, T_d is the dewpoint temperature in Kelvin (inverted from the Clausius-Clapeyron relationship with respect to water, closely matching the Simmons et al. (1999) equation for temperatures typical at Bermuda), R_v is the gas constant for water vapour, $e_{s0} = 6.112$ hPa is the saturation vapour pressure at reference temperature $T_0 = 273.15$ K, and all other variables are as defined above.

Table 2.2 Coefficients used in the [Simmons et al. \(1999\)](#) approximation to the saturation vapour pressure curve. This approximation allows the smooth transition between vapour pressure with respect to water and vapour pressure with respect to ice. The coefficients and intercepts are chosen such that they match observed saturation vapour pressures (e.g. [Buck, 1981](#); [Alduchov and Eskridge, 1996](#)).

Coefficient	Value
a_w	17.502
b_w	32.19 K
a_i	22.587
b_i	-0.70 K

2.1.3 Satellite Data

Visible-channel (0.64 μm) passive satellite data from sensors aboard the 13th Geostationary Operational Environmental Satellite (GOES-13) are used. A geostationary orbit is designed such that the satellite orbits with the same angular velocity as the Earth. To achieve this, the sub-satellite point is directly above the equator and the satellite orbits at approximately 36,000 km altitude.

Orbiting at such great distance limits the horizontal resolution of the imaging sensors aboard the satellite. However, the fixed sub-satellite point allows frequent sampling of the same geographic location. This sensor offers routine observations over the Western Atlantic, including Bermuda, every 30 minutes with a nominal spatial resolution of 1 km.

The satellite data are accessed through the Comprehensive Large-Array Data Stewardship Service (CLASS) operated by NOAA and the US National Climate Data Centre (NCDC). The Northern Hemisphere Extended and Full Disk routine imagery is subset to the $4 \times 4^\circ$ box centred on Bermuda's coordinates is used.

For the purposes of this study, only daytime data from May through October, 2012-2016 are stored to reduce storage requirements. Only daytime imagery is meaningful for the visible-channel. This is a passive sensor which can only detect scattered/reflected or emitted electromagnetic radiation in the spectral band centred on 0.64 μm . Therefore, much of this signal comes from sunlight reflected to the satellite sensors by targets (e.g. clouds).

Alternatively, polar orbiting satellites orbit at much lower altitudes and can therefore achieve higher spatial resolution. However, the trade-off is that they sample the same spot infrequently. The time it takes for a satellite to revisit the same sub-satellite point in the Landsat series of polar orbiters is 16 days, but can achieve horizontal resolution of around 30 m.

2.1.4 Satellite Calibration

The GOES-13 visible-channel data provided by CLASS are in a 16-bit digital number format. The steps used to convert this to a nominal albedo are outlined through the following sequence of equations:

$$bit_{10} = bit_{16}/32 \quad (2.3a)$$

$$R_{nominal} = k \times bit_{10} - X_{space} \quad (2.3b)$$

$$R_{calibrated} = \kappa \times R_{nominal} \quad (2.3c)$$

$$\kappa = 0.00012j - 3.72315 \quad (2.3d)$$

$$\alpha_n = \frac{d^2}{\cos(\theta_z)} R_{calibrated} \quad (2.3e)$$

$$d = 1 + 0.017\cos(2\pi(N - 2)/365) \quad (2.3f)$$

$$\cos(\theta_z) = \sin(2\pi\phi/360) + \cos(2\pi\phi/360)\cos(\delta)\cos(h) \quad (2.3g)$$

$$\delta = -2\pi 23.44\cos(\alpha)/360 \quad (2.3h)$$

where the raw 16-bit data, bit_{16} , is converted to 10-bit data, bit_{10} , in equation 2.3a.

The 10-bit data is then converted to a raw reflectance using the linear pre-launch calibration fit in equation 2.3b. Here, $k = 0.001160$ is the slope of the fit, and $X_{space} = 29.0$ is the intercept.

To account for instrument drift during the lifetime of the sensor aboard the GOES-13 satellite platform, a post-launch calibration correction, κ , is applied (equation 2.3c). In the present study, the ‘Operational’ post-launch calibration correction is used. However, this correction is updated in monthly steps. This means that the final albedo

product might exhibit discontinuities. These are prevented by linear regression on the calibration coefficients so that there is a κ for each satellite image, j , (equation 2.3d). Here, j is defined as days since 01 Jan 1900.

Next, the calibrated and corrected reflectance is converted into a nominal albedo, α_n (equation 2.3e). This is done by multiplying by the Earth-Sun distance squared d^2 , and dividing by the cosine of the solar zenith angle, $\cos(\theta_z)$. The Earth-Sun distance in Astronomical Units (AU), d , is estimated from equation 2.3f. Here, N is the decimal day of year. This assumes that the Earth is closest to the Sun on 02 Jan each year, and furthest from the Sun on 02 Jul each year, which is accounted for through the ‘-2’ offset. The cosine of the solar zenith angle is estimated using equation 2.3g. $\cos(\theta_z)$ depends on the location’s latitude, ϕ , the solar declination angle δ , and the local hour angle, $h = 2\pi H_{solar}/24$, where $H_{solar} = H - 24((720 - 4\lambda)/1440)$, is the hour with respect to the solar cycle converted from the hour in UTC, H , based on the longitude, λ .

After applying the set of equations 2.3a-h to the raw data from CLASS, an internally consistent dataset of visible-channel imagery is created. A final check is applied to limit the retrieved albedo to within physical bounds (i.e. $[0, 1]$).

Using this method, the imagery appears particularly bright at low solar zenith angles. Therefore, satellite data that corresponds to times with solar zenith angles less than 75° are not included in analysis, and this choice is more thoroughly discussed in Chapter 3.

2.1.5 ERA-interim Reanalysis

European Re-Analysis Interim (ERA-interim) data (Dee et al., 2011) are additionally used to examine the large-scale environment with respect to cloud trail occurrence in Chapter 3. These data are also used in chapters 4, 5, and 6 to provide the large-scale subsidence forcing in idealised simulations of cloud trails. A re-analysis uses data assimilation methods to best combine observations of the Earth system with a modelled background state Earth system in a way that minimises errors from both observations and the model. A re-analysis will include observations and numerical methods that

might not have been available in real-time and can therefore provide a more accurate analysis than what is operationally available.

Re-analysis data for the warm season of the 2012-2016 period is used in this study. This data is provided with a horizontal grid spacing of 0.75° (approximately 80 km at the equator, or T255 spectral truncation). With this grid spacing, small islands like Bermuda do not appear as explicit land points in the re-analysis land-sea mask as the island is much smaller than the $0.75 \times 0.75^\circ$ grid box. Therefore, the influence of the island itself is missing in the re-analysis. In the vertical, ERA-interim data are provided on 60 pressure-levels from the surface to 0.1 hPa.

In the present study, the following ERA-interim variables are used: Temperature, specific humidity, u- and v-wind components, and pressure vertical velocity, ω , are used on pressure levels, while mean sea level pressure is also used.

Chapter 3:

Climatology of Cloud Trails at Bermuda

3.1 Introduction

Bermuda's location is fortuitous for the study of cloud trails (CTs) because of its isolation from upstream impediments to the flow. In this chapter, a review of existing literature on CTs is presented. Following this review, a simple novel classification algorithm is introduced. This algorithm objectively sorts visible-channel satellite imagery into one of three classifications: cloud trail (CT) if there is a cloud band downwind of the island; non-trail (NT) if there is no cloud band downwind of the island; or obscured (OB) if large-scale cloud obscures the island from view, or makes it impossible to determine the presence of a CT.

These classifications are then used to construct a climatology of CT occurrence for Bermuda. The combination of this climatology, radiosonde observations, and ERA-interim re-analysis data are used to investigate the local and larger-scale environments typical of each classification. This chapter is largely based on our published paper in *Monthly Weather Review* [Johnston et al. \(2018\)](#) (©Copyright 2020 AMS).

3.1.1 Motivation

This type of climatological study is designed to better understand the environments which are more supportive of CTs, and to determine differences between CT and NT days. At Bermuda, forecasters have intuition about the general environmental characteristics which are supportive of CTs. This study seeks to expand upon and quantify the environmental differences which decide the development, or not, of CTs.

CT climatologies have previously been constructed for short time periods (a year or less) and used manual classifications (e.g. [Nordeen et al., 2001](#); [Kirshbaum and Fairman, 2015](#)). In addition, no climatology has yet been constructed for an island with a significant annual cycle. This chapter seeks to produce a longer-term climatology based on an automated and objective satellite classifications.

3.1.2 Background

As introduced in Chapter 1, a CT is an organised convective phenomenon associated with low-level horizontal buoyancy gradients which force a low-level circulation. Surface heterogeneities such as small islands can result in large differences in surface heating over small horizontal distances.

Land has a lower heat capacity than the surrounding sea. Given the same daytime heating from the Sun, the temperature of the land surface would therefore increase more than the temperature of the sea surface. Mixing and diffusion within the sea water acts to buffer the already reduced temperature change. The daytime land surface temperature can be much higher than the surrounding sea and therefore air is heated more strongly over land, is less dense, and is more buoyant than over sea.

As undisturbed marine air flows over the strongly heated island, a thermal internal boundary layer (TIBL) forms. The TIBL is also referred to as a “convective internal boundary layer” (CIBL) and is more fully defined by [Miller et al. \(2003\)](#). This TIBL grows in depth following the flow across the island, and terminates in a turbulent region in the immediate wake of the island ([Matthews et al., 2007](#)).

The island-heated air continues to be advected downwind of the island where a

warm plume forms (e.g. [Garstang et al., 1975](#); [Matthews et al., 2007](#)). This supports a horizontal buoyancy gradient all along the edge of the island and the warm plume. The across-flow component of the CT circulation forms in this buoyancy gradient region and consists of: low-level inflow and convergence; ascent; and outflow and divergence aloft ([Matthews et al., 2007](#)).

The across-flow component of the CT circulation supports a band of ascent which is oriented parallel to the low-level wind direction ([Nordeen et al., 2001](#); [McFarlane et al., 2005](#)). This band of ascent acts as a structure on which convection can be triggered. This is generally what is observed in sequences of visible-channel satellite imagery, as CTs appear as bands of cloud anchored to and extending downwind of a small island. In some cases, the CT may be the only significant cloud present.

CTs are found to occur in the absence of significant orography on ‘flat’ islands such as Nantucket (e.g. [Malkus and Bunker, 1952](#); [Malkus and Stern, 1953](#)); Anegada ([Malkus, 1963](#)); the Bahamas Islands ([Bhumralkar, 1973a](#)); and Nauru (e.g. [Nordeen et al., 2001](#); [McFarlane et al., 2005](#); [Matthews et al., 2007](#)). A phenomenon similar to the flat-island-CT is observed downwind of heated islands with significant orography: for example, in the Lesser Antilles Islands (peaks near 1500m) (e.g. [Garstang et al., 1975](#); [Smith et al., 2007](#); [Kirshbaum and Fairman, 2015](#)); the Eastern Pacific Island of Guadalupe (1300 m) ([Dorman, 1994](#)); and Hawaii (over 4000 m) (e.g. [Smolarkiewicz et al., 1988](#); [Yang and Chen, 2008](#); [Yang et al., 2008b,a](#)). Orographically-generated flow perturbations are often larger than their thermally-generated counterparts. Flow blocking and mountain wave activity complicate the isolation of just the thermally-driven component as orographic forcing can act in the same sense as thermal forcing, and elevated heating can enhance the thermal circulation ([Crook and Tucker, 2005](#); [Kirshbaum and Wang, 2014](#)). Furthermore, transitions between sea-breeze driven convective triggering, and orographic convective triggering occur in relation to mountain height, low-level stability, and wind speed ([Wang and Sobel, 2017](#)).

In addition to orography, island size is important in determining the nature of any convection that develops in a heated flow. [Williams et al. \(2004\)](#) showed that convective vigour is indistinguishable from the background oceanic regime for islands with an area less than about 100 km² (i.e. small islands). Larger islands can result

in convection of similar intensity to that found in more continental regimes. Similar results were found using different metrics for convective intensity in [Robinson et al. \(2011\)](#) and [Sobel et al. \(2011\)](#), although these studies focussed on deep rather than shallow convection.

One mechanism for the increase in convective intensity over larger islands is the convergence of sea breeze fronts (e.g. [Crook, 2001](#)). However, under no background flow and for a given heating rate, two-dimensional simulations show that the strength of a sea breeze circulation decreases with island/peninsula width ([Savijarvi and Matthews, 2004](#)). Those authors also found that with some background flow, the sea breeze circulation is displaced and tilted downwind of the island. The sea-breeze circulation, which is already weaker for smaller islands, is transformed into a “steady heat island” circulation in which the leeward cell of the sea-breeze continues downwind and dissipates and the windward cell is shifted over or downwind of the island ([Estoque and Bhumralkar, 1969](#); [Savijarvi and Matthews, 2004](#)).

Analysis of intensive field campaign observations on Nauru (as part of the Nauru Island effect Study - NIES ([McFarlane et al., 2005](#))) revealed potential mechanisms for the initiation and maintenance of cloud trails observed there ([Savijarvi and Matthews, 2004](#); [Matthews et al., 2007](#)). These authors propose that a thermal internal boundary layer forms and grows as oceanic air advects across the heated island. This turbulent thermal layer then evolves into a warm, cloud-topped plume downwind of the island. This idea of a warm plume is also consistent with the observations at Barbados discussed by [Garstang et al. \(1975\)](#). [Matthews et al. \(2007\)](#) also suggested that the warm plume drives a thermal circulation that may be responsible for the maintenance of the cloud trails, which were found to extend on average 125 km downwind of Nauru ([Nordeen et al., 2001](#)).

Cloud trail climatologies at Nauru and in the Lesser Antilles Islands were made using visible-channel satellite imagery ([Nordeen et al., 2001](#); [Kirshbaum and Fairman, 2015](#)). Each hourly image was manually classified by [Nordeen et al. \(2001\)](#) as either ‘Cloud Plume’ - a line of cloud is seen extending downwind of and anchored to the island; ‘Non-plume’ - there is no evident band of anchored cloud; or ‘Obscured’ - the island is obscured from view by larger scale cloud phenomena. The current study

will follow this definition, but referring to ‘Cloud Trails’ (CT), ‘Non-trail’ (NT), and ‘Obscured’ (OB) scenes respectively.

Both [Nordeen et al. \(2001\)](#) and [Kirshbaum and Fairman \(2015\)](#) showed that strong surface heating during the day was important for CT development. CT occurrence was seen to peak in mid-afternoon both at Nauru and the Lesser Antilles Islands. At Nauru, this diurnal cycle in CT occurrence combined with the low elevation of Nauru (only rising to 30 m above sea level) reinforces the idea that cloud trails are primarily thermally driven by the difference in low-level heating between the island and the surrounding ocean ([Nordeen et al., 2001](#)).

In Section [3.2.1](#), the method for manually sorting imagery is outlined. The six-month warm season of 2012 is manually classified. This is used both to derive parameters for the algorithm and to compare to the algorithm’s classifications for the same period.

A simple automated method for sorting visible-channel satellite imagery at Bermuda into ‘Cloud Trail’, ‘Non-trail’, or ‘Obscured’ categories is outlined in section [3.2.2](#). An automated classification scheme is more readily reproducible, objective, and expedient when compared to manual classifications. Further, it has the potential to be quickly adapted for other locations or similar problems.

The algorithm is applied to 16,400 images over the warm seasons of five years from 2012 to 2016 to construct a climatology. We then use these classifications in conjunction with radiosonde observations and European Re-Analysis Interim (ERA-interim) data to describe the environments that coincide with each classification. Finally, we discuss some cloud trail behaviour at Bermuda and why some environments appear more favourable for CT formation.

3.2 Methods

3.2.1 Manual Classification Method

Initially, a manual classification of the first warm season (May through October, 2012) is performed to both aid in the design of an automated classification method, and

to validate the automated classifications. Scenes are classified using the three categories (CT, NT, and OB) outlined above. Here a scene refers to a visible-channel satellite image cropped to a $4 \times 4^\circ$ domain centred on Bermuda (as discussed in Chapter 2). The 10 m wind direction is used to determine where the downwind side of the island is in each scene. Scenes are classified as CT if the area around the island is not covered in cloud and a band or bands of cloud are seen downwind of, and apparently anchored to Bermuda (e.g. Fig. 3.1a). If the area around the island is not covered in cloud, but no band of cloud is seen anchored downwind of the island, the scene is classified as NT (e.g. Fig. 3.1b). Finally, if the scene is mostly cloud, particularly in such a way that covers much of the island from view, it is classified as OB (e.g. Fig. 3.1c).

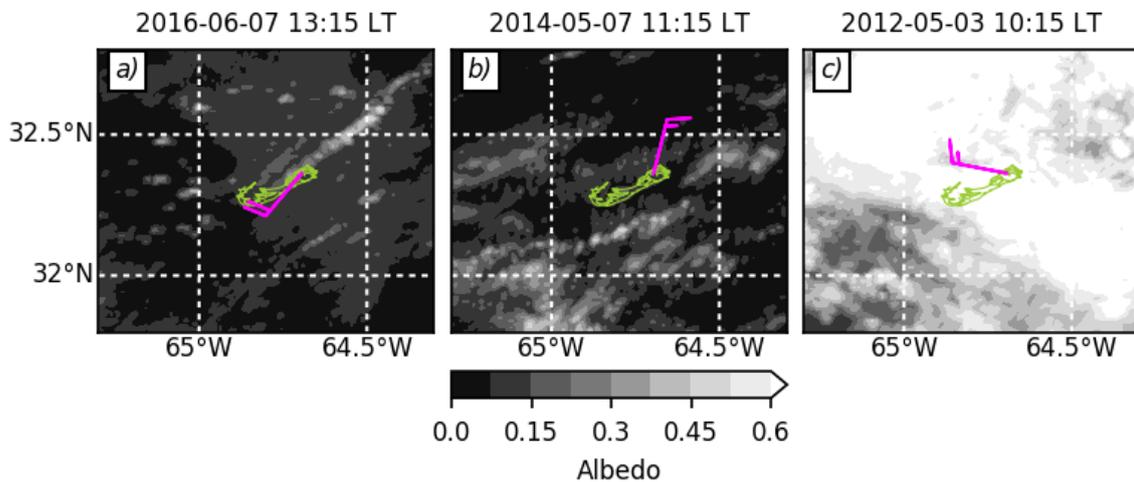


Figure 3.1 GOES-13 visible satellite imagery showing example scenes. (a) Cloud Trail scene where clouds organise into a band downwind of Bermuda as indicated by a southwest-northeast oriented band of higher albedo. (b) Non-Trail scene in which there are few clouds and some of the higher albedo near and over the island might be shallow water and land showing up rather than cloud. (c) Obscured scene where the island (and much of the surroundings) is obscured from view by widespread cloud as indicated by high albedo throughout the scene. In each example, a wind barb is plotted showing 10 m wind direction and speed in kts at TXKF provided by the BWS.

While this exercise is somewhat subjective, the majority of scenes were straightforward to classify. The biggest challenge was in distinguishing between downwind cloud bands that are CT vs. those that are associated with other phenomena such as low-level convergence not linked to the island, cold pools, etc. It is suspected that these features may also be incorrectly classified by an intentionally simplistic automated approach as outlined in Section 3.2.2 below.

3.2.2 The Classification Algorithm

From our manual classification and the previous work discussed above, we know that (i) CT are either absent or hidden from view by other large-scale cloud in OB scenes, and (ii) CT are characterised by more cloud downwind of the island than upwind. These two ideas are used to design an intentionally simplistic algorithm to automate the objective classification of thousands of satellite scenes into our three categories.

3.2.2.1 Cloud Masking

First, cloudy and cloud-free pixels must be identified. Each pixel is nominally 1×1 km. The albedo from the visible-channel satellite data is used to identify cloudy pixels (see a detailed description of how the measured satellite radiances are converted to albedo in Chapter 2, Section 2.1.4). For simplicity we create a binary cloud mask based on an albedo threshold. Pixels are called ‘cloudy’ and given a value of ‘1’ if the albedo is greater than this threshold. Remaining pixels with albedo less than this threshold are given a value of ‘0’.

Testing was carried out to assess the sensitivity of cloud identification to the choice of albedo threshold. Threshold values between 0.03 and 0.21 were considered for a mostly cloud-free day with a cloud trail in the afternoon hours according to observations from Bermuda’s L. F. Wade International Airport (ICAO: TXKF, WMO: 78016). These tests suggest that if our threshold is lower than 0.1, the land area of Bermuda and the shallow water surrounding Bermuda is falsely masked as cloud (e.g. Fig 3.2). Conversely, if the threshold is greater than 0.2, pixels that likely contain smaller cloud elements or that are part of optically thin cloud regions might be falsely masked as not cloudy. The albedo threshold is taken to be 0.15 as a compromise between these two ‘limits’, and this is the same value used in [Kirshbaum and Fairman \(2015\)](#).

A known issue with this simple masking method is that cloud cover remains ambiguous over land and coastal regions. Land provides a higher background albedo than the ocean. Land pixels may therefore still erroneously be masked as cloudy. Pixels over land, as well as pixels that are one pixel away from the coastline, are therefore

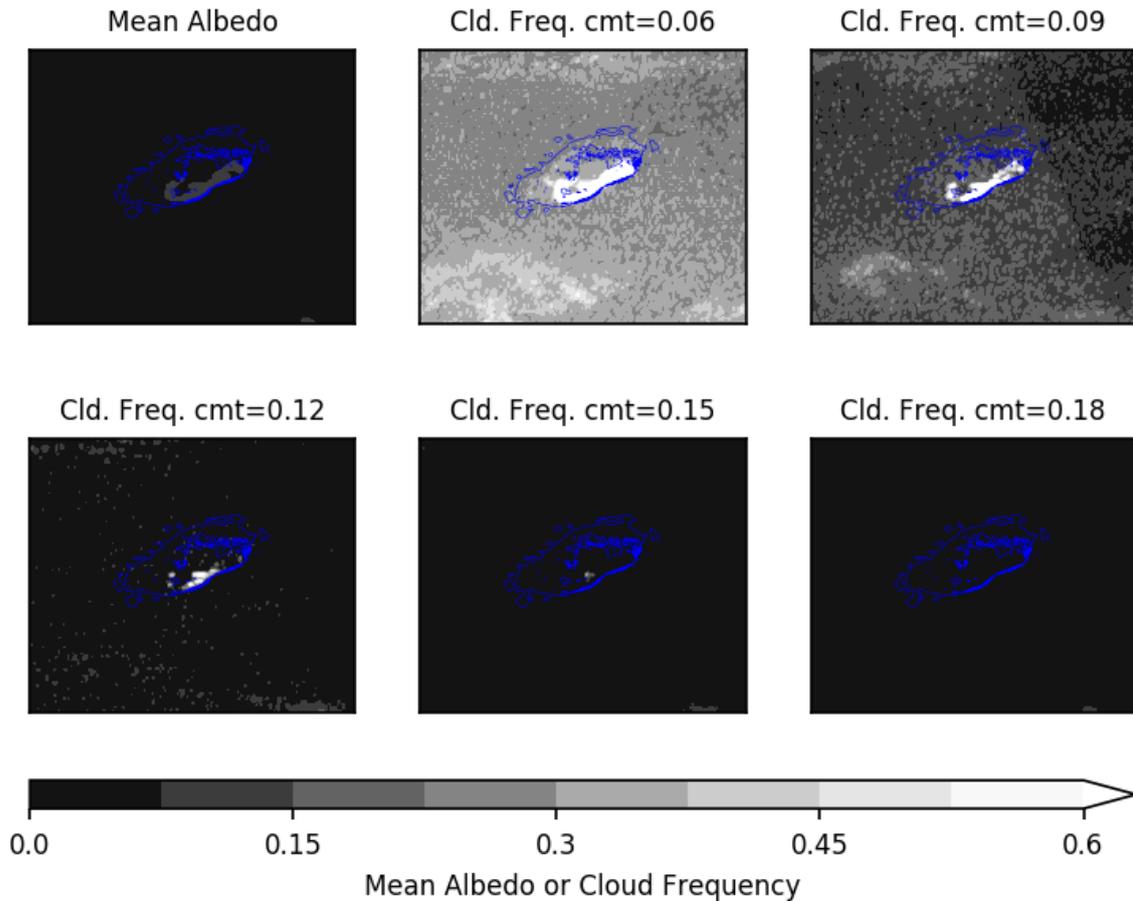


Figure 3.2 Part of the sensitivity analysis for cloud mask thresholds for 0700 - 1115 Local Time on 18 Aug 2016 (a mostly cloudless morning). The top left is the mean morning albedo. Mean albedo over the island is higher than the surrounding water. The remaining panels are the morning cloud frequency for increasing cloud mask thresholds (cmt). The title of each panel states the cmt used. The 15 m water depth is given by a thin blue contour to highlight the extent of the shallow waters of the lagoon to the north of Bermuda which can be seen in the cloud frequency for the $cmt = 0.06$ and $cmt = 0.09$ panels.

excluded from calculations to account for this issue.

Furthermore, for high solar zenith angles, spurious regions of high or low albedo can appear depending on the cloud cover. For instance, individual towering clouds might cast shadows on other cloudy regions - these shadowed areas are then falsely masked as cloud-free. Similarly, at sunrise and sunset (when the solar zenith angle is 90°) there can be a very bright line through the scene. Scenes must then be rejected if the maximum solar zenith angle is too high.

To determine what “too high” is, five sample scenes that were obviously impacted

by the above problems were chosen (e.g. Fig 3.3). All scenes with solar zenith angles less than some threshold are then discarded. This threshold is determined by taking a first guess of 90° and decreasing the threshold by 5° increments until the five sample scenes are discounted. We find that scenes with solar zenith angles $< 75^\circ$ are sufficiently illuminated to avoid these issues.

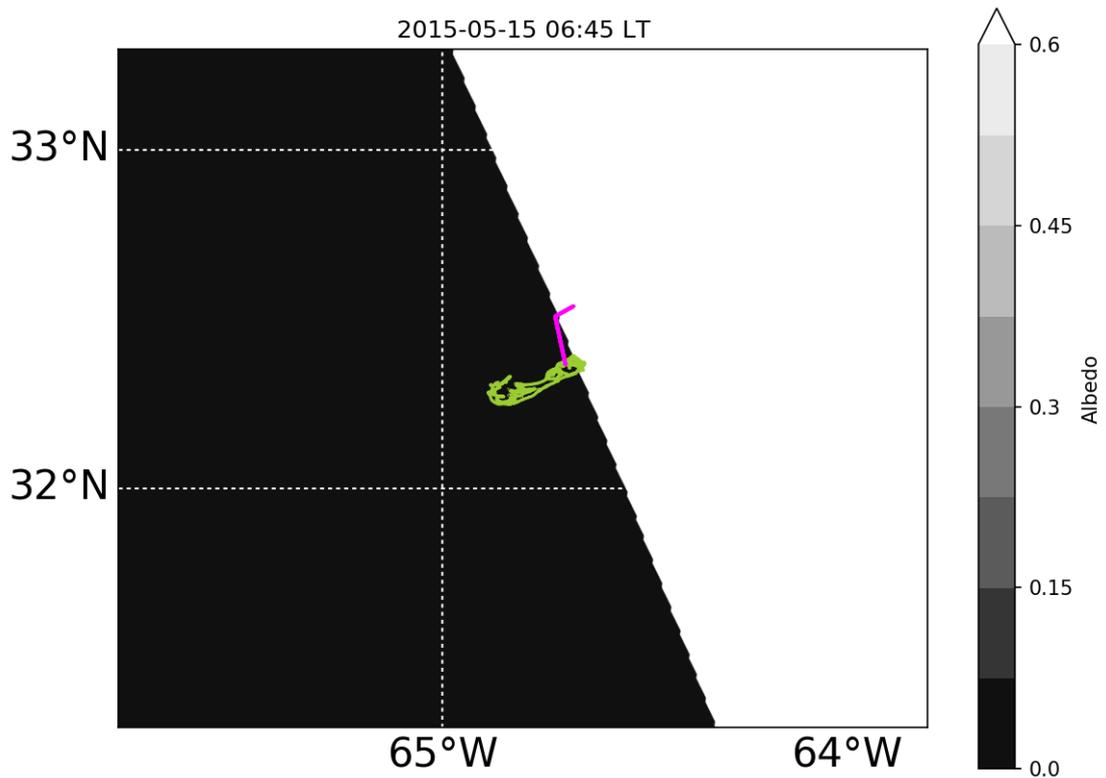


Figure 3.3 Visible-channel satellite imagery as in Fig. 3.1. An example of a case of high solar zenith angle at sunrise. The left half of this image has not yet had sunrise and appears black. The right half of this image has just had sunrise and so the incoming solar radiation is weak. The albedo (the ratio of the reflected solar radiation to incident solar radiation) is artificially high.

3.2.2.2 Classification Algorithm

A visible satellite scene must first pass the solar zenith angle test above (e.g. Fig. 3.4a). It is then converted to a cloud mask using the simple thresholding method discussed above (Fig. 3.4b). Once this is done, the scene can then be assessed for the presence of a CT. The image simplification starts by first classifying a scene as OB or non-OB (i.e. including both the CT and NT classifications). Then, the non-OB scenes

are further subdivided into CT or NT categories.

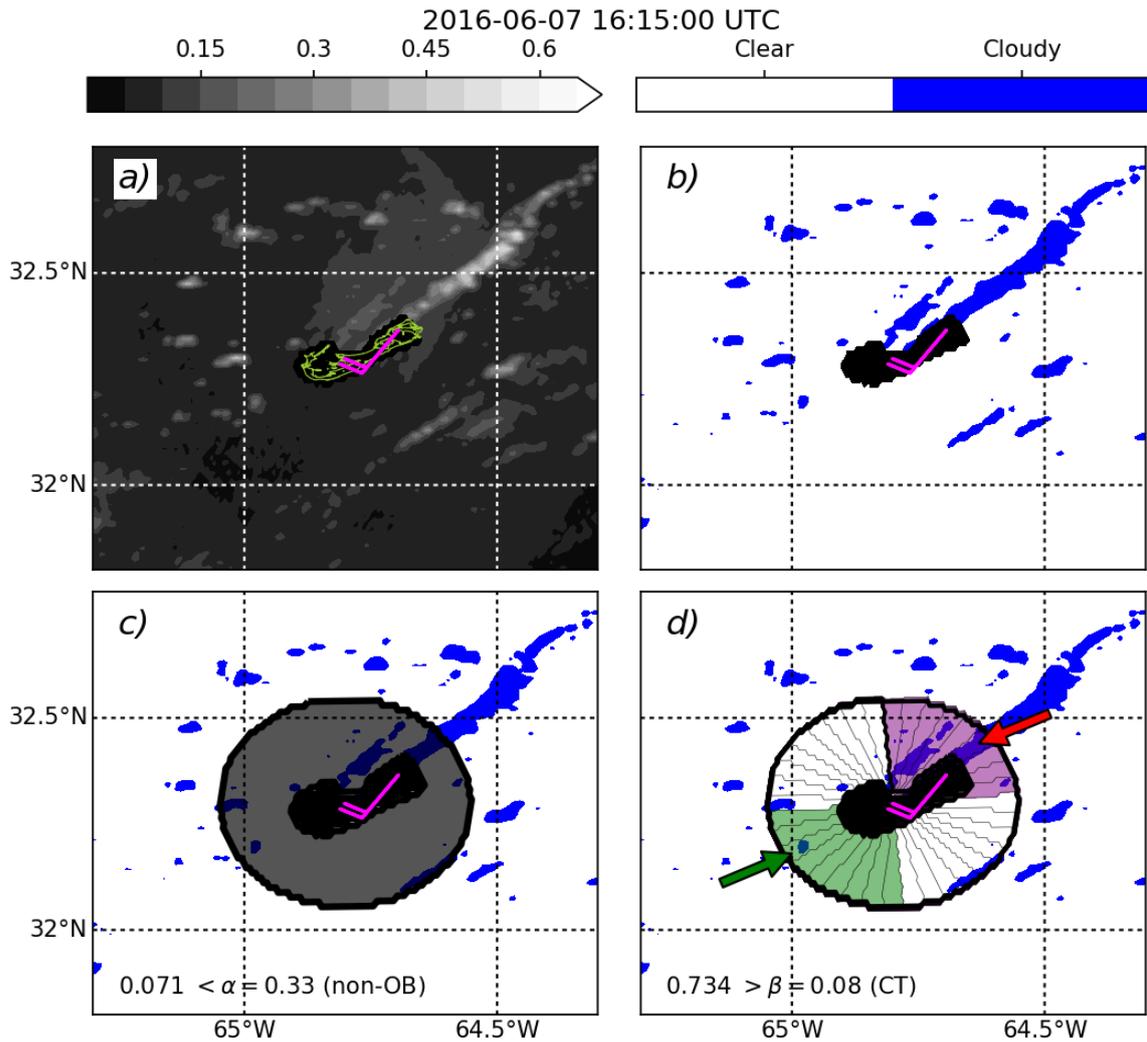


Figure 3.4 Walk-through of the satellite image classification algorithm. a) The visible-channel satellite scene is screened for solar zenith angle less than 75° to ensure it is well-illuminated; b) a cloud mask threshold is applied to the imagery; c) the cloud fraction in the grey circle, excluding the island is computed to determine whether the scene is obscured or not; d) if the scene is not obscured, the maximum cloud fraction in the upwind and downwind quadrants are compared. If the difference between the two is greater than some threshold then the scene is cloud trail, otherwise non-trail.

A circular region of interest with radius 0.25° (about 25 km) centred on Bermuda is considered. This circle contains the entire island and the edges of this circle are at least 0.10° (about 10 km) away from any land points (Fig. 3.4c). To determine how cloudy the scene is and therefore whether or not to classify it as OB, the cloud fraction in that circular area is calculated. Here, the cloud fraction is defined as the spatial mean of the cloud mask over a given area. Scenes with a cloud fraction in the circular

region greater than a threshold, α , are classified as OB. If the cloud fraction is less than α , then the scene is non-OB and tested further.

The non-OB scenes are further sorted into NT and CT classifications. As seen in loops of visible-channel satellite imagery and reported in the literature above, CT initiate at the island and extend downwind, forming a band of cloud anchored to the island. [Nordeen et al. \(2001\)](#) used cloud level (the mean over the 950 - 850-hPa layer) wind directions from once daily radiosondes to manually identify CT at Nauru. However, while it may be reasonable to assume that there is no significant change in wind direction during the day at Nauru because it is well embedded in the tropical Pacific trade wind region, Bermuda is near the axis of the Bermuda-Azores High and small changes in the position of this ridge axis could mean a reversal in the wind direction. Bermuda's more complex geometry make it more sensitive to wind direction as well.

[McFarlane et al. \(2005\)](#) showed that at Nauru, the surface wind direction compares well with the heading of identified CT. We have found that the 0900 LT pressure-weighted cloud level (950 - 850-hPa mean) wind direction from radiosonde ascents at Bermuda also compare well with the surface winds measured at the same time at Bermuda (not shown). For the purposes of this study, it therefore appears reasonable to use the half-hourly surface wind direction from TXKF to determine where upwind and downwind are with respect to Bermuda.

The same circular 0.25° area used to test for OB scenes is now divided into 10° sectors, thirty-six in total (Fig. 3.4d). The first sector is centred on north. We expect the CT signal to be strongest nearer to the island because of the anchoring described above. From sensitivity tests, if the radius of the circle is too large, the sectors start to become broader than the CT and so the signal in sector cloud fraction becomes damped (not shown). At distances of 0.25° , a sector is roughly 4.4 km wide.

As a result of Bermuda's geometry, some sectors contain more non-land pixels than others. Sensitivity to the number of pixels used for cloud fraction computations was tested by using sectors of different lengths and shapes to make them contain more similar numbers of pixels. The results are found to be generally insensitive to having a more equal pixel count in each sector.

Half-hourly 10 m wind direction observations are used to locate the upwind and downwind sectors. Nine sectors centred on the wind direction are considered. This is to account for fluctuations in the wind direction, differences between wind direction and CT heading, and CTs that span two sectors. Next, the cloud fraction is calculated for each of the nine sectors and the maximum cloud fraction of those nine sectors is taken to represent the upwind or downwind cloud fraction (marked as ‘U’ and ‘D’ respectively in Fig. 3.4d).

The difference between the downwind and upwind sector maximum cloud fraction (δF) is taken. Since CT are characterised by organised cloudiness downwind of islands, it follows that the cloud fraction should be higher downwind than upwind. δF must be > 0 to satisfy this condition. However, we also wish to exclude small differences which may be due to random sampling of the undisturbed upwind cloud field. Hence, δF must be compared to a threshold (β) based on cloud fraction statistics for the chosen definition of upwind and downwind sectors.

Values chosen for our algorithm parameters α and β are discussed in the next section.

3.2.3 Algorithm Parameters

Assuming the manual classification is the ‘truth’, it is used to estimate optimal values for α (the near-island cloud fraction above which a scene is considered OB) and β (the δF above which a non-OB scene is classified as CT).

For α , the cloud fraction is computed for each manually classified scene as described in the previous section for discriminating between non-OB and OB scenes. We then consider the cumulative distribution of this cloud fraction for the OB scenes, and the inverse cumulative distribution for non-OB scenes (Fig. 3.5). The cloud fraction value where these two distributions intersect is taken to be the optimal value for α . With our manual classifications, we find α to be 0.33.

To determine the value of β , we again refer to the manual classifications. We apply our method for discriminating between NT and CT scenes from the above section to the scenes which are manually classified as non-OB. We then consider the cumulative

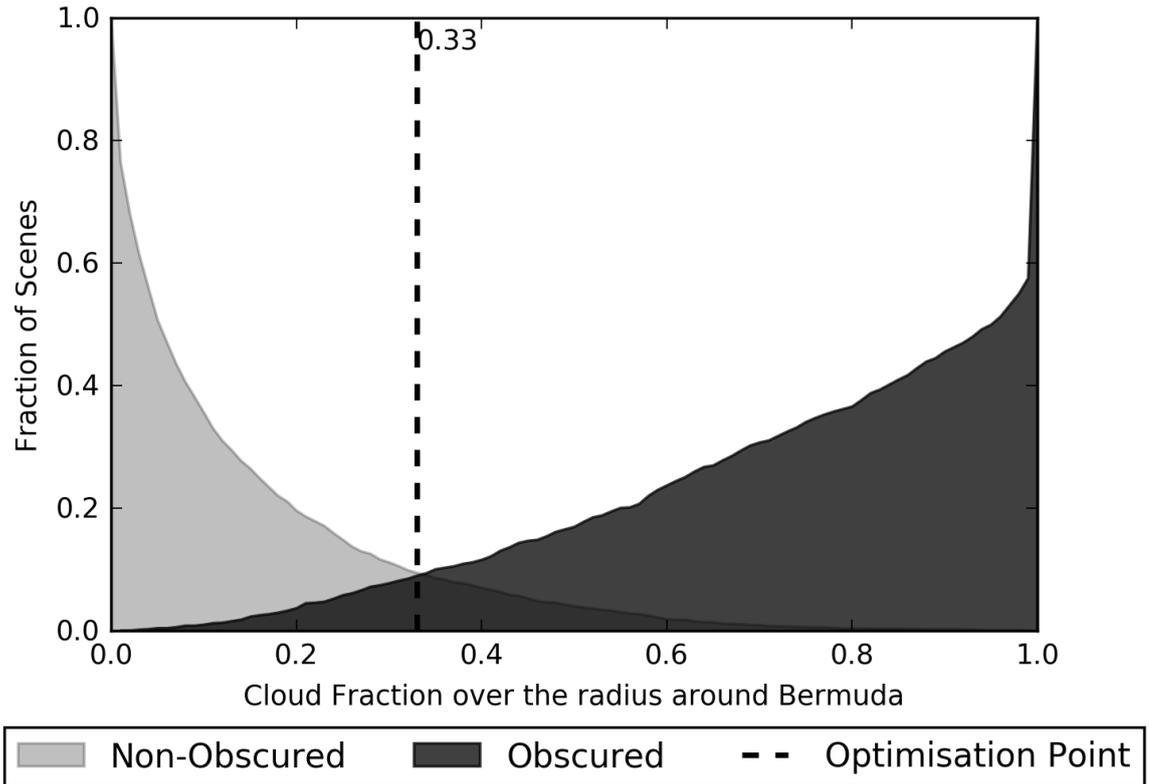


Figure 3.5 Cumulative distribution function of cloud fractions for manually classified non-OB scenes compared to the inverse cumulative distribution function of cloud fractions for manually classified OB scenes. The cloud fraction is calculated over the grey region illustrated in Figure 3.4c. By taking the cloud fraction where these two distributions intersect as the threshold to distinguish between OB and non-OB, we maximise the number of obscured scenes which are correctly classified.

distribution of δF for CT scenes, and the inverse cumulative distribution of δF for NT scenes (Fig. 3.6). Again, the δF corresponding to where these two distributions intersect, 0.08, is taken as the value for β .

Sensitivity tests (not shown) suggest that the overall patterns of the annual and diurnal cycle are not strongly sensitive to the choice of α and β . We also remark that $\alpha = 0.33$ is similar to the value for the mean cloud fraction for all days across our period of interest (0.342). Finally, a higher more conservative value for β might yield a more confident CT classification at the expense of rejecting the CT cases with more complex background cloud distributions.

In Section 3.3.1, the manual classifications are compared to the algorithm’s classifications for the warm season of 2012. This is the same period used to derive the

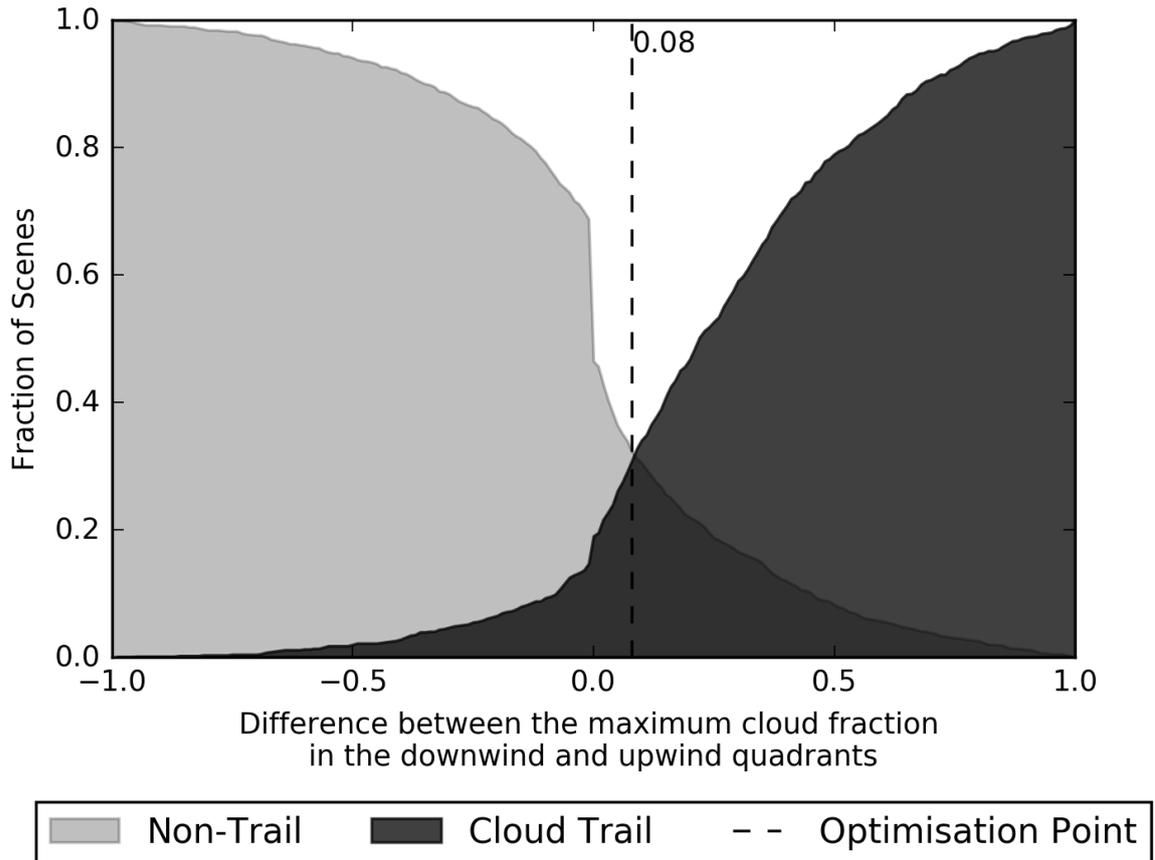


Figure 3.6 Cumulative distribution function of the maximum difference in cloud fraction between the upwind and downwind regions, δF , for manually classified NT scenes compared to the inverse cumulative distribution function of δF for manually classified CT scenes. The regions used to calculate δF are illustrated in Figure 3.4d. Only the scenes which are manually classified as non-OB are included here. By taking the δF where these two distributions intersect as the threshold for to separate non-OB scenes into NT and CT, we maximise the number of correctly classified CT scenes.

parameters which control the algorithm. The degree to which the manual classifications match the algorithm’s classifications is measured using skill scores.

Then environmental conditions are presented for each classification in Section 3.3.2. These are then discussed in terms of the larger-scale environment. Finally, CT morphology at Bermuda is compared with that reported in literature.

3.3 Results and Discussion

3.3.1 Algorithm Validation

Here a two-pronged approach is used to assess how well the algorithm classifies satellite imagery into our three defined categories. First, skill scores are applied to compare the manually classified 2012 warm season with the algorithm's classifications for that same warm season. Then side-by-side comparisons of select climatological figures are used to qualitatively assess the algorithm's performance.

3.3.1.1 Skill Scores

By our definitions for α and β , we are effectively maximising the Peirce Skill Score (PSS). This score ranges from -1 to 1, where 1 indicates a perfect classification, and 0 indicates no classification skill. [Peirce \(1884\)](#) defines it as follows:

$$PSS = \left(\frac{H}{H + M} \right) - \left(\frac{F}{F + C} \right), \quad (3.1)$$

where PSS is the Peirce Skill Score, H is the number of hits, M is the number of misses, F is the number of false alarms, and C is the number of correct negatives. This score describes the match when there are two possible outcomes, 'a' (event occurs) or 'b' (event does not occur). The corresponding algorithm-manual pairs are 'aa' for hit, 'ba' for miss, 'ab' for false alarm, and 'bb' for correct negative. For α , 'a' refers to an OB classification, and 'b' refers to a non-OB classification. At this stage, the algorithm has only made OB and non-OB classifications, the manual NT and CT classifications both count toward non-OB. For β , we ask what classification would the algorithm assign to the scenes that are manually classified as non-OB. In this case, 'a' refers to a CT classification and 'b' refers to an NT classification.

After running the algorithm with our new-found values for α and β , we can then apply the PSS to individual classifications. Following the 'a' and 'b' framework above for CT classifications as an example, 'a' now refers to a CT classification, and 'b' refers to either an NT or OB classification. In these cases a correct negative ('bb') can be any

combination of the two non-event scenes (e.g. NT-NT, NT-OB, OB-NT, or OB-OB).

For CT, the PSS is 0.51. It is 0.46 for NT classifications, and 0.82 for OB classifications. Overall, the algorithm has a score of 0.60. The reader must be reminded that the manual classification process is subjective, so while we take it for the truth in determining our algorithm parameters in the previous section, and for validation purposes in the present section, it should be understood that it is subject to human error and interpretation differences. Despite this, the manual classification is still instructive for building intuition for what to expect from algorithm classifications.

However, the PSS does not describe every aspect of the algorithm’s performance. Contingency tables have been produced to further aid in the quantification of the algorithm’s performance. Table 3.1(a) are the contingency table results for the CT classification. Included in the table are the number of ‘hits’, ‘misses’, ‘false alarms’, and ‘correct negatives’ as described above. The same four categories are shown for NT and OB classifications in Table 3.1b and c.

Table 3.1 Contingency tables for (a) CT classifications, (b) NT classifications, and (c) OB classifications. Using the CT classification as an example: top left are ‘hits’ where the algorithm and the manual classifications are both CT; top right are ‘false alarms’ where the algorithm classification is CT but the manual classification is not CT; bottom right are ‘Correct Negatives’ where the algorithm and the manual classifications are both not CT; and bottom left are ‘misses’ where the algorithm classification is not CT but the manual classification is CT.

Class		(a) CT	(b) NT	(c) OB			
hits	false alarms	649	371	743	275	1117	193
misses	correct negatives	320	2008	506	1824	111	1927

In all cases, there are more hits than either false alarms or misses. Commonly used metrics derived from such contingency tables are the ‘Hit Rate’ and ‘False Alarm Rate’ given by:

$$HR = 100 \times \frac{H}{H + M} \tag{3.2a}$$

$$FAR = 100 \times \frac{F}{F + C} \quad (3.2b)$$

respectively, and H , M , F , and C are as defined above.

Compared to the manual classification of the 2012 warm season, we find that the Hit Rate (Eq. 3.2a) is much higher than the False Alarm Rate (Eq. 3.2b) for each classification. These are 67.0% vs. 15.6% for CT, 59.5% vs. 13.1% for NT, and 91.0% vs. 9.1% for OB.

Finally, we consider the bias score for each of our algorithm classifications. This bias score is defined as:

$$BIAS = \frac{H + F}{H + M} \quad (3.3)$$

this quantity can range from 0 to ∞ , where a score of 1 represents a perfect classification, values less than 1 indicates the algorithm is biased against making a classification, and values greater than 1 indicate that the algorithm is biased toward making a given classification. Using the data provided in Table 3.1 we find the bias to be 1.05 for cloud trail classifications, 0.82 for non-trail classifications, 1.07 for obscured classifications, and 0.97 overall.

Taken all together, these scores suggest that the algorithm is skillful in making classifications as compared to our subjective manual classification. Furthermore, it is not strongly biased toward or against making any particular classification.

We therefore move to demonstrating that the algorithm preserves the qualitative features of the CT climatology. For instance, the algorithm could be systematically missing CT for specific scenarios, thus giving good skill scores but skewing the climatologies.

3.3.1.2 Algorithm-Manual Comparison for 2012

If the algorithm can also reproduce the annual and diurnal cycle reasonably well, it follows that conditionally averaging based on the algorithm classifications will match the manual classifications. Therefore, we choose to inspect the algorithm's temporal fidelity compared to the manual classification.

The algorithm is used to classify visible-channel satellite imagery for May through October of 2012. These classifications are then sorted by time of day and by month. The same sorting is done for the manual classifications for 2012. We show the percentage of scenes for a given time of day and in each month and compare the algorithm to the manual in Figure 3.7. This helps us illustrate both the annual and diurnal cycle of CT occurrence and how the diurnal cycle changes through the year.

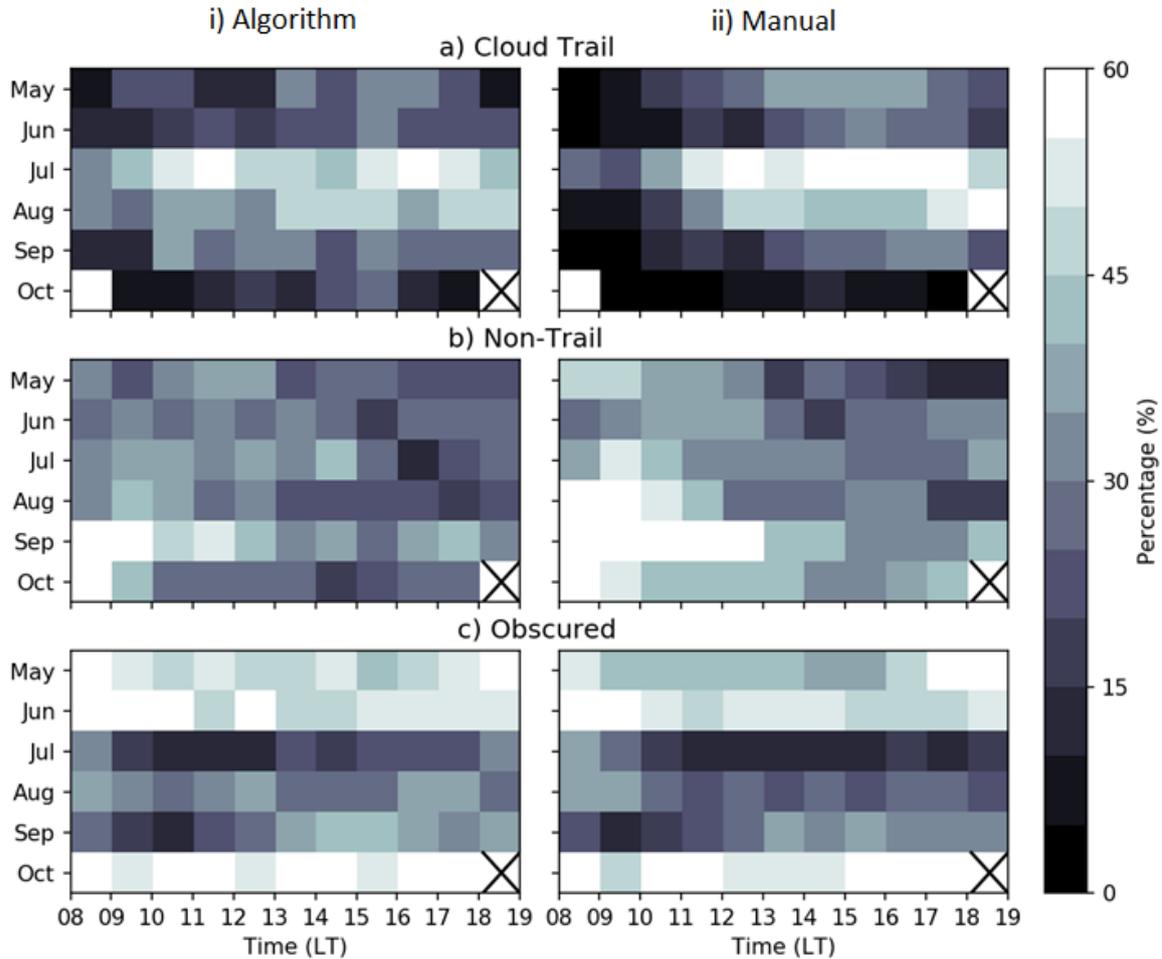


Figure 3.7 A comparison between i) the algorithm classifications for 2012 and ii) the manual classifications for 2012 of the percentage of scenes classified as a) CT, b) NT, and c) OB sorted by local time from left to right and month from top to bottom. A black ‘x’ is in the cell which has no classified scenes because of high solar zenith angles.

Comparing the algorithm to the manual classification for 2012, we see that the temporal structure of CT occurrence is broadly reproduced by the algorithm (Fig. 3.7i and ii). Overall the differences between the algorithm and manually classified data are small.

There are two time-frames with bigger differences. First, the algorithm has too few CT in early afternoon in July, and there are also too many OB scenes at this time. This suggests that the algorithm is having difficulty distinguishing between non-OB and OB for cases that are borderline OB/CT. Also of note, the algorithm has too many morning CT scenes compared to manual for July and August. This coincides with too few NT scenes and too few OB scenes suggesting that there are some borderline CT/NT and CT/OB scenes mis-classified as CT.

These bigger differences could be corrected by adjusting α and β , or adding more complexity to the algorithm. But these remedies might only serve to introduce differences of similar magnitude elsewhere, so we do not modify the algorithm based on this comparison.

We have shown that the algorithm performs well in both quantitative skill scores and it preserves the general structure of the statistics for 2012. We therefore apply the algorithm to the remaining period of record. Our final analysis period includes May through October of 2012 through 2016. Results from the algorithm classification for this period are presented in the next section.

3.3.2 CT Climatology

The five-year period in which imagery is classified makes this the longest satellite-based study of CT in the literature. Previous studies have only considered one- or two-year periods. This means that we can expand upon the discussion on the environmental factors which control the occurrence of CT at Bermuda in the context of both intra- and inter-annual variability.

3.3.2.1 Temporal Distribution of CT Occurrence

First, we return to the annual and diurnal cycle of CT occurrence in the algorithm for the five-year period (Fig. 3.8). Immediately, it is clear that expanding to five years has resulted in a much smoother temporal distribution.

Considering the annual cycle, the percentage of OB scenes decreases to a local minimum in July - the same month in which the percentage of CT scenes increases to

a maximum. For all times of day, the percentage of CT scenes nearly doubles from 24% in May to 44% in July, while OB scenes decrease from 45% to 22%. Meanwhile, the percentage of NT scenes is steadier between 29 and 35% (on the higher end in July and August).

Similarly, a clear diurnal cycle is recovered when considering rows in Figure 3.8. The morning is characterised by a higher percentage of NT than CT scenes, and this reverses in the afternoon and evening. As a seasonal average, the percentage of NT scenes declines from near 38% in the morning to near 28% in the afternoon. Concurrently, the percentage of CT scenes increases from near 24% in the morning to near 36% in the afternoon - offsetting much of the change in the percentage of NT scenes. The percentage of OB scenes varies less through the day; between 40% in the early morning and late evening, to near 34% around 1100 LT. This local minimum in OB scenes occurs around the same time as the decrease in the percentage of NT and increase in the percentage of CT scenes.

For days with CT scenes, the median first CT scene is identified at 0945 LT and median final scene at 1745 LT. The first CT scenes are identified earliest in the day in June, and the last scene is latest in the day in August. The first CT scene occurs later and the last occurs earlier in May and October. This implies a strong connection to the annual and diurnal cycle of solar heating.

Note that in this part of the analysis, only days with two or more consecutive CT scenes are considered (i.e. days with a CT lasting at least an hour). Cases in which CT occurred outside of the range of times with sufficiently low solar zenith angles are expected, and such scenes cannot be classified by the algorithm. This would result in a real start time that is earlier than detected in the algorithm. However, the algorithm's tendency toward more morning CT than manually classified might counteract this.

3.3.2.2 CT Intermittency and Duration

More than one period of consecutive CT scenes (or "CT periods") may be classified on a given day (Fig. 3.9). The algorithm classification for 2012 has an affinity for making CTs more intermittent than what a manual observer might classify. Figure

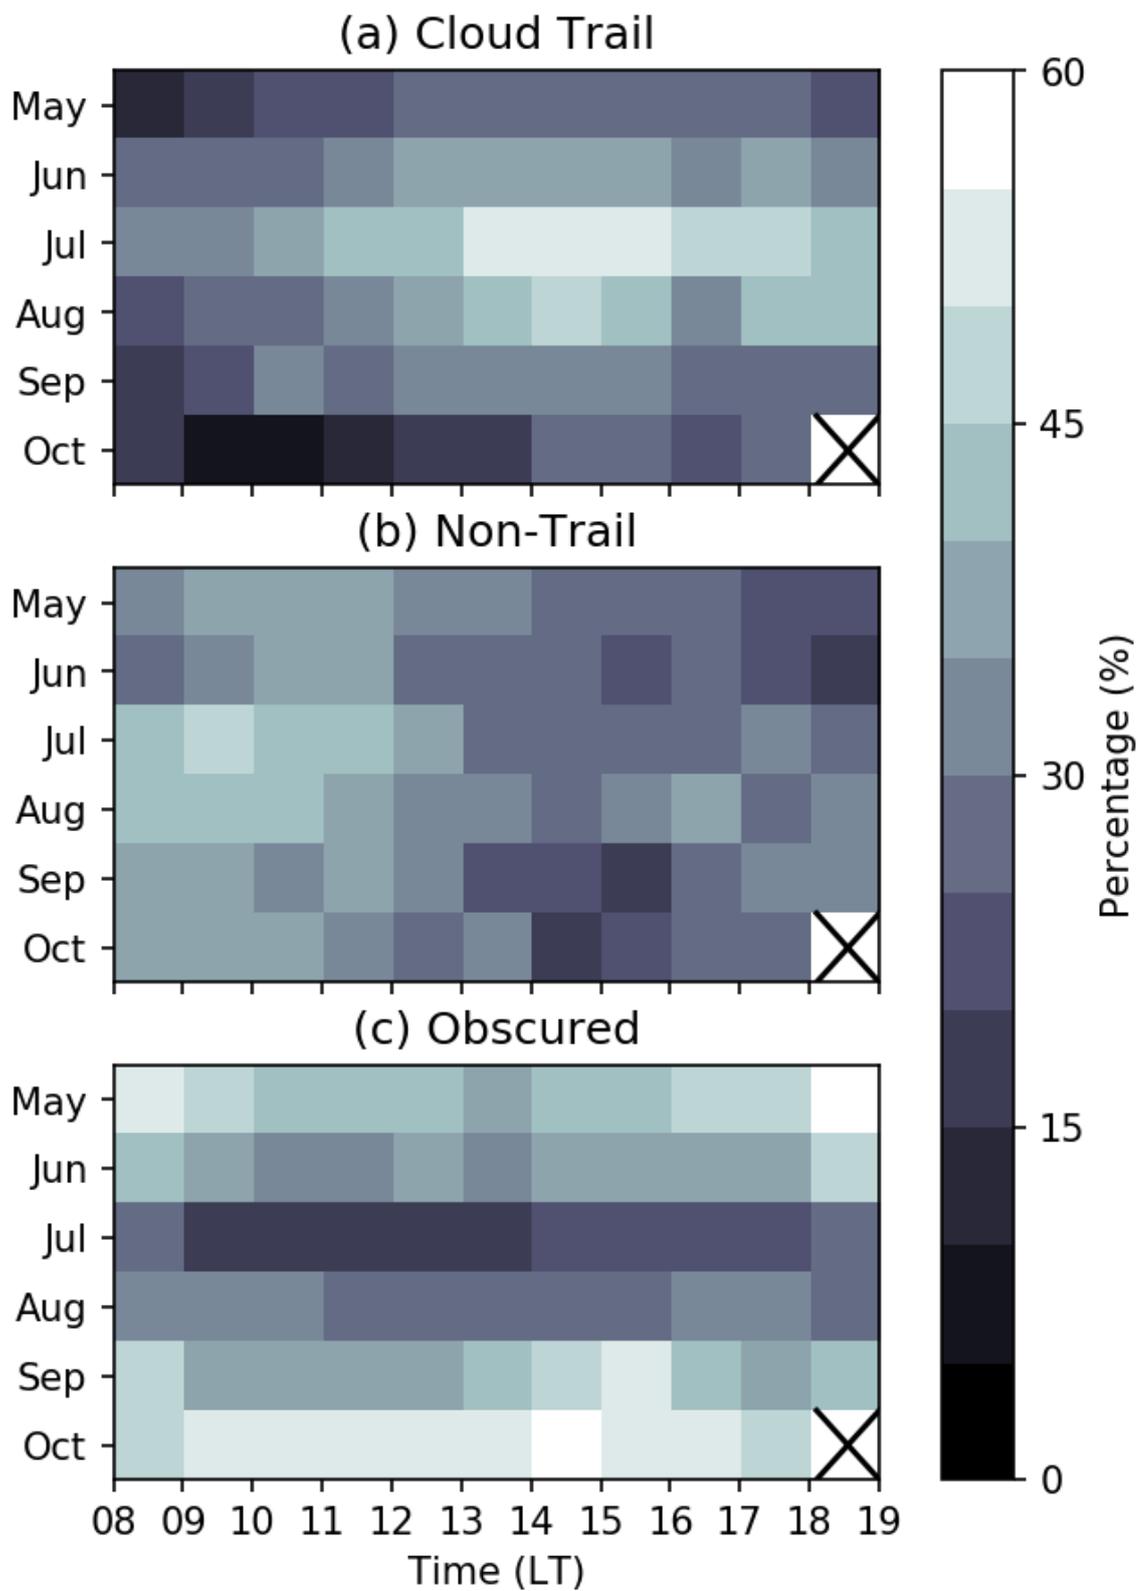


Figure 3.8 As in 3.7 but just for the algorithm applied to the whole period of record, 2012 - 2016.

3.9a shows that the algorithm most often has two CT periods (31 % of CT days) with a substantial portion of days having one (23 %) or three (25 %). Whereas our manual classification generally has one CT period (59 %), sometimes two (35 %), and rarely three (6 %, Fig. 3.9b). When all five years of the algorithm data is included, the general shape of this distribution unchanged (not shown).

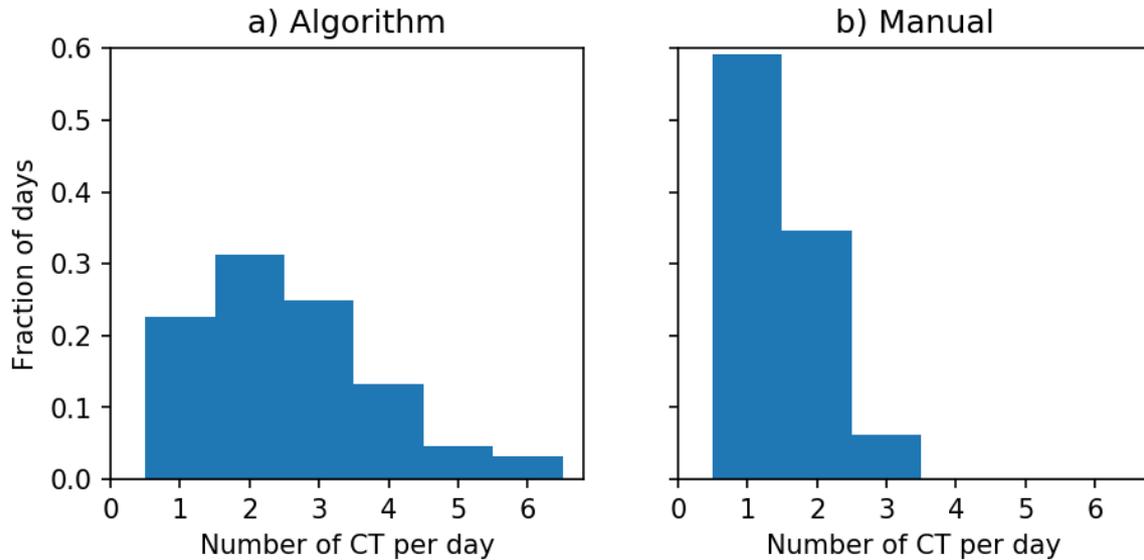


Figure 3.9 Frequency of different contiguous CT periods per day for the warm season of 2012. Results from a) the algorithm classification, and b) the manual classification.

Physically, we can explain more than one CT period as occurring when the environment is marginal for CT formation. A small change in the wind direction could mean flow is no longer parallel to the island’s long axis, the thermally driven circulation and therefore source of lift is then insufficient to reach the LCL. Similarly, a decrease in background low-level humidity might make cloud formation unobtainable by lift of the same strength. Either change, or some combination of both might cause the CT to temporarily dissipate before conditions return to allow it to redevelop. Alternatively, passing showers or upper-level cloud could temporarily obscure the island and/or disrupt a CT.

However, this intermittency is also partially an artefact of the algorithm. If the large scale cloudiness is near the threshold for distinguishing between OB and non-OB scenes (α), a small increase in cloud cover might fool the algorithm into classifying real CT as OB scenes. A manual classifier might continue using CT classification based

on knowledge from the previous classification. Similarly, if there is an increase in the ambient upwind cloud (e.g. due to an advancing front, or a cluster of isolated clouds), the algorithm might be fooled into making an NT classification as the downwind-upwind cloud fraction difference is no longer greater than our threshold β .

This 5-year algorithm-based climatology of the warm season CT occurrence reveals that during the period of peak CT occurrence in July, there is also a peak in their duration. The mean algorithm-classified duration increased from 01:58 hours in May to a peak of 03:17 hours in July before decreasing to 01:34 hours in October when our period of study ends. CTs that persist for just one scene are taken to have a duration of half an hour here, and each additional consecutive CT scene is counted as an additional half hour of duration. On days in which more than one period occurs, the longest duration is taken as the value for that day. Comparing the manual and algorithm classifications for just 2012, we find that the manually classified CT last roughly twice as long (peaking at 07:13 hours in July). The difference between manual and algorithm is because of the increased intermittency in the algorithm classification. The longer CT duration during the warm season peak is due to longer days and more settled weather which we suspect is enhanced by stronger solar forcing.

3.3.2.3 CT Length

For scenes that were manually classified as having CT (May through October 2012), we manually estimated the length of the CT in each scene by finding its end-point in the visible-channel satellite imagery and calculating the distance between that point and the centre-point of Bermuda (assuming that this is where the CT originated). This follows the methodology outlined by [Nordeen et al. \(2001\)](#).

We next consider the half-hourly mean CT length estimated for the manually classified imagery. First, there is a local maximum in CT length (78 km) at 0815 LT, and this then decreases to a local minimum (41 km) by 0915 LT. This early morning CT is consistent with the discussion above on the one or more short-lived CTs.

Second, the length tends to increase through the remainder of the day. CT on average grow from the local minimum length at 0915 LT until 1515 LT when they are

about 90 km long. The mean length then remains near 90-95 km through sunset (not shown).

Using the 10 m wind speed and the estimated CT length, we predict the length half an hour later assuming advection is the only factor controlling changes in length. These predictions are found to generally be in good agreement with the manually estimated lengths with a correlation coefficient of 0.75 ($R^2=0.55$) on 109 length predictions. Additional factors, such as precipitation or evaporation of cloud liquid due to entrainment of drier surrounding air might act against the increase in length due to advection. More investigation in Chapter 5 is aimed at understanding what is actually being advected given typical cloud lifetimes of around 30 minutes (Ludlam, 1980).

3.3.3 Environmental Characteristics

The next goal of this study is to describe and highlight differences between CT-, NT-, and OB-dominated environments at Bermuda. To do this, we composite radiosonde data and ERA-interim data by classification. As there is an annual cycle at Bermuda in both the classifications and the environmental conditions (as mentioned in Chapter 1 and in Section 3.1.2 of the current Chapter), we only consider the peak of the CT occurrence, June through August (JJA), to avoid reproducing the signal of the annual cycle in our composites.

On most days, there are satellite images in each of the three classifications. However, the radiosonde data is available once per day at 0900 LT. This presents the challenge of assigning a single classification to a day made up of several classifications.

Classifications are aggregated from half-hourly to daily using a threshold method. The threshold is chosen to be the 75th percentile of the daily percentage in a given classification. First, the percentage of each day in each classification is calculated for the entire period of record. Next, we find each classification percentage's 75th percentile. If the percentage of day in a classification is greater than the 75th percentile for that classification, then that day is assigned to that classification and included in its composite.

For example, the 75th percentile for CT is found to be 59 %. A CT day is taken

to be any day in which more than 59 % of the scenes are classified as CT. The 75th percentile thresholds for NT and OB are 50 % and 59 % respectively. The same day cannot be assigned to more than one classification by this method since any two of the three 75th percentile thresholds will always sum to greater than 100 % of a day.

An alternative method is to classify a day as CT if any scene on that day is classified as CT. In composites, this more strongly obscures the CT signal in the environmental profiles with the NT and OB signal.

3.3.3.1 The Local Environment

The interpolated IGRAv2 data (described in Chapter 2 Section 2.1.2) are used to consider the anomalies from the 0900 LT JJA 2012-2016 mean potential temperature and relative humidity profiles. Radiosonde ascents that are incomplete below 700-hPa are not included in the composites. The resulting radiosonde composites for each classification are shown in Figure 3.10. The mean profile of potential temperature and composite anomalies (top) and the mean profile of relative humidity and composite anomalies (bottom) are shown for the layer between the surface and 700-hPa. With this method, about 25 % of the JJA days are assigned to each classification's composite - leaving about 25 % of the days to be discarded.

For each composite, we assume that the mean anomaly is the centre of a normal distribution of anomaly profiles. We then estimate the uncertainty about the mean composite anomaly profiles as the standard error: $\frac{\sigma}{\sqrt{n}}$, where n is the number of observations in a given composite. CT-dominated days (black dash-dot profiles in Fig. 3.10b) have the highest surface potential temperature of the three classifications. Potential temperature anomalies (Fig. 3.10b) decrease in magnitude to near-zero above 950-hPa, indicating a less stable than normal boundary layer. The anomalies then increase to become positive again by 700-hPa, indicating that the layer aloft is slightly more stable than the climatology for this period. These conditions have long been regarded as favourable for shallow convection (e.g. [Malkus and Bunker, 1952](#)).

The composite profile for potential temperature anomalies for NT-dominated days (purple dashed profiles in Fig. 3.10) has a similar shape to that for the CT-dominated

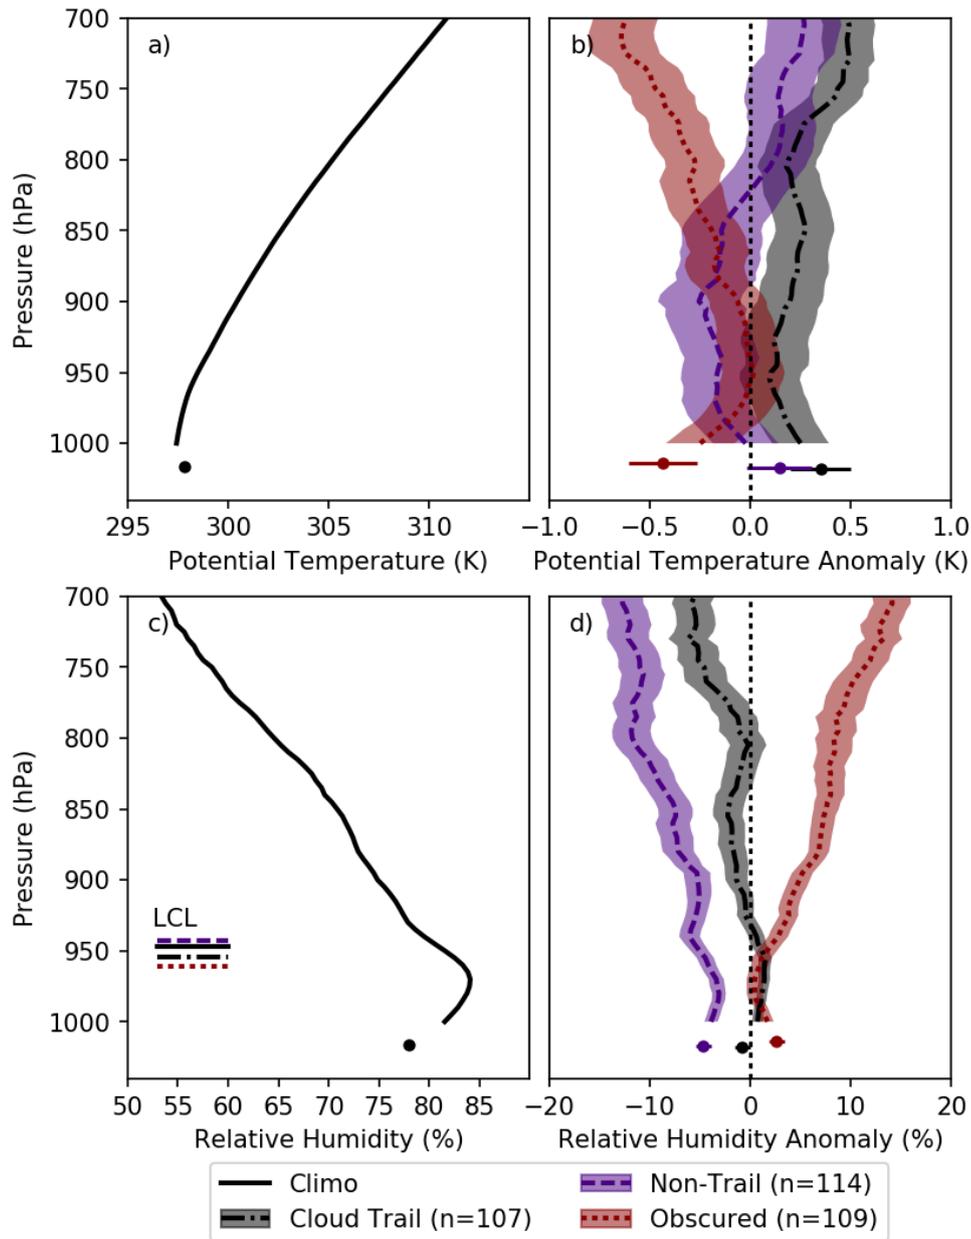


Figure 3.10 a) The climatological potential temperature (θ) profile from radiosondes at 0900 LT JJA, 2012-2016. b) The composite θ anomalies for CT (black dashdot line and gray shading), NT (purple dashed line and shading), and OB (dark red dotted line and shading). The shaded region represents the uncertainty about the mean anomaly: $\pm \frac{\sigma}{\sqrt{n}}$. All composites are of 0900 LT radiosondes for JJA in 2012-16. c) as in a) but for relative humidity (RH). The mean LCL for each classification and for the climatology are shown as horizontal line segments with corresponding line colours and styles. d) as in b) but for RH composites. In each panel the surface value is given by a dot with a range given by $\pm \frac{\sigma}{\sqrt{n}}$. LCL pressures are calculated using the radiosonde temperature and dewpoint at the lowest reported altitude and Eqs. 2.2a, 2.2b, and 2.2c from Chapter 2.

days, indicating a similar pattern of stability anomalies. However, the NT anomalies are smaller than the CT anomalies throughout the entire lower troposphere shown, and feature potential temperatures below normal between 1000 and 850-hPa (Fig. 3.10b).

A very different potential temperature regime is present for the OB-dominated days (the dark red dotted profiles in Fig. 3.10). The lowest 100-hPa is more stable than normal, while the profile is less stable than normal aloft. This pattern is consistent with the idea that the OB days are associated with large-scale cloudiness. Lower than normal near-surface potential temperatures are an expected consequence of cloud-shading and possible evaporative cooling from precipitation.

Similarly, a clear separation between the classifications is evident in the relative humidity anomaly composites (Fig. 3.10d). The OB-dominated days have the highest surface relative humidity, but the OB- and CT-dominated days have similar, near normal relative humidity between 1000 and 950-hPa. Above the 950-hPa level, the CT profile remains near normal, becoming slightly drier than normal while the OB profile becomes much more humid than normal. The NT profile again has a similar shape to the CT profile, but is drier than normal throughout this layer.

Overall, CT-dominated days are warmer than normal below 700-hPa with near normal relative humidity, while NT-dominated days are cooler and drier than normal below 700-hPa. The higher low-level relative humidity on CT-dominated days implies a lower LCL than for NT-dominated days. However, the warmer than normal low-levels on CT-dominated days implies a higher LCL than for the NT case. We therefore consider the height of the LCL (computed from radiosondes) to further examine this relationship.

The composite LCL pressure for each classification and the climatology are marked on Figure 3.10c. We find that the mean composite LCL pressure for CT-dominated days is 955-hPa. This corresponds to a lower height than both climatology (949-hPa) and NT-dominated days (941-hPa). Lower LCL heights suggest cloud formation is more readily achievable and these results provide a plausible explanation for why cloud trail formation is sensitive to low-level humidity.

In addition to thermodynamic profiles, wind speed is expected to influence the types of circulation that result from island surface heating by contributing to the

organisation of any cloud that forms. In dry idealised two- and three-dimensional simulations, [Savijarvi and Matthews \(2004\)](#), [Kirshbaum \(2013\)](#), and [Kirshbaum and Wang \(2014\)](#) show that for light or calm background wind regimes, thermally induced circulations form over a surface heat source (reminiscent of pure sea-breezes with cloud over the island). The role of low-level wind speed is discussed in the next section. For stronger background winds, thermally induced circulations form a band of ascent downwind of the heat source (reminiscent of CT with a band of cloud downwind of the island).

3.3.3.2 Low-level winds

The mean wind speed for each classification is shown in [Table 3.2a](#). NT scenes are associated with lighter winds and OB scenes are associated with stronger winds than both CT and all scenes. On calm NT days, or NT days with very light winds a sea-breeze circulation with cloud formation directly over the island could occur. On other NT days with more typical-strength background flow, the previous section suggests that a dry boundary layer is preventing the formation of the cloud bands characteristic of CTs. Its the combination of the calm days with cloud over the island and the dry days with no cloud band that give NT days an overall lower wind speed. The stronger winds associated with OB scenes are likely related to large-scale disturbances.

Furthermore, Bermuda is oriented such that its long axis runs approximately southwest-northeast. From a Lagrangian perspective, surface heating is maximised for low-level flow parallel to the long axis of an island as an air parcel remains over the island heat source longer. For southwesterly and westerly winds, more non-OB scenes are classified as CT (36-39 %) than NT (23-28 %) ([Table 3.2b](#)). For all other wind directions, including northeasterly flow which is also parallel to the long axis of the island, a greater proportion of non-OB scenes are classified as NT than CT. We suspect that this has to do with moisture availability. Northeasterly flow is climatologically associated with advective drying, while southwesterly flow is associated with advective moistening.

Table 3.2 a) Mean 10 m wind speed (m s^{-1}) at Bermuda’s L. F. Wade International Airport by algorithm classification. The number of observations in each classification is included in brackets. b) The percentage of scenes in each classification and the total number of classified scenes by wind direction.

a) Wind Speed	CT	NT	OB	All
	5.27	4.83	6.18	5.46
	(5,191)	(5,196)	(6,013)	(16,400)

b) Wind DIR	CT	NT	OB	All
N	27 %	44 %	29 %	1197
NE	26 %	39 %	35 %	1318
E	23 %	31 %	46 %	1578
SE	27 %	32 %	41 %	2000
S	33 %	33 %	34 %	2747
SW	36 %	23 %	41 %	2861
W	39 %	28 %	33 %	3399
NW	29 %	40 %	31 %	1300

3.3.3.3 Large-scale Environment

An examination of the larger-scale fields is expected to aid in explaining our above results linking wind direction to CT occurrence. We have already established that low-level moisture is a dominant control on CT formation. ERA-interim composites of the 0900 LT 1000-hPa specific humidity for JJA in Figure 3.11a, c, and e show that Bermuda lies in a moisture gradient pointing from northeast to southwest for all classifications. Given this background pattern, northeasterly flow would result in advection of drier low-level air on average. Such a flow therefore tends to make the environment less favourable for CT on average despite maximising low-level heating by maintaining a direction parallel to the long axis of the island.

Indeed, these composites are consistent with the radiosonde and surface composites in Fig. 3.10. They indicate that CT and OB days have similarly high low-level specific humidity while NT-days are less humid. Additionally, the western part of the Bermuda-Azores High (indicated by the 1020-hPa contour) is shown to extend its control on this region for much of JJA, and it retreats somewhat to the east on OB days. This pressure pattern implies a wind field dominated by westerlies and southwesterlies at Bermuda with lightest winds on NT days.

Finally, the composites show 500-hPa vertical motion between 0.00 and 0.01 Pa s^{-1}

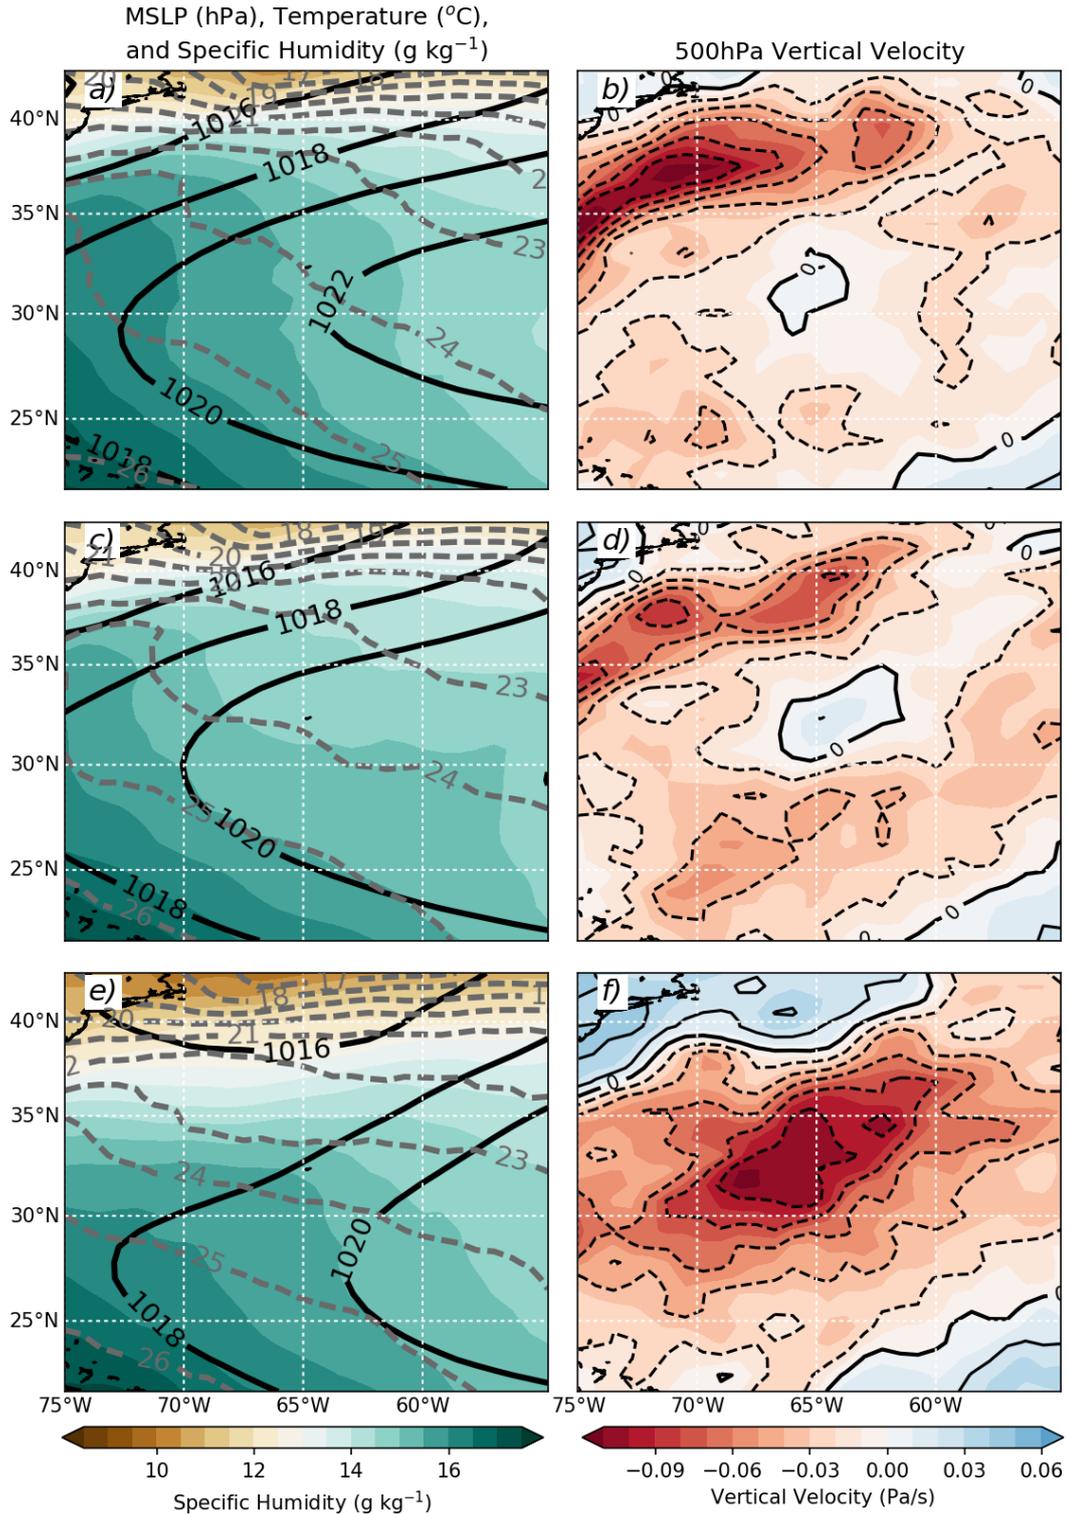


Figure 3.11 JJA 2012-2016 ERA-Interim reanalysis composites centered on Bermuda. Composites for a), b) cloud trails; c), d) non-trails; and e), f) obscured. (left) The mean sea level pressure in hPa (solid black contours), 1000-hPa temperature in $^{\circ}\text{C}$ (dashed grey contours), and 1000-hPa specific humidity in g kg^{-1} (brown to green-blue shading). (right) The 500-hPa vertical velocity in Pa s^{-1} , ascent (red) and subsidence (blue) are shown in dashed and solid contours, respectively.

near Bermuda on CT-dominated days (Fig. 3.11b), about 0.01 Pa s^{-1} on NT-dominated days (Fig. 3.11d), and about -0.10 Pa s^{-1} on OB-dominated days (Fig. 3.11f). Weak 500-hPa vertical motions or subsidence seen on CT- and NT-dominated days is consistent with the expected lack of large-scale cloudiness and favours shallow convection. This subsidence may also help to explain the drier and more stable than normal layer seen in the sounding composites for NT-dominated days.

3.4 Algorithm For Barbados

At the start of the current Chapter in Section 3.2.2, the algorithm is described as being intentionally simplistic. This is to aid the ease of application to other locations. This section seeks to briefly demonstrate that this is indeed the case.

GOES-13 visible-channel satellite imagery for Barbados over the same 5-year period from 2012-2016 is acquired via the same CLASS interface described in Chapter 2 Section 2.1.3. Wind speed and direction at 10 m, and radiosondes all from Grantley Adams International Airport (ICAO: TBPB, WMO: 78954), are used in the algorithm to construct the CT climatology for Barbados. These surface and radiosonde data come from the same source as the Bermuda data (described in Chapter 2 Section 2.1.2).

For this most basic demonstration, the algorithm parameters for Bermuda are reused ($\alpha = 0.33$ and $\beta = 0.08$). Performing a Barbados-specific manual classification and deriving new algorithm parameters for Barbados would likely produce more robust results.

To confirm that the algorithm generally picks out cloud features which look like cloud trails, we show the cloud frequency in each category (Fig. 3.12). Here, afternoon (12 - 15 LT) satellite images that have a classification are rotated so that the wind direction appears to be from the north.

We see that the features highlighted at both Bermuda and Barbados look somewhat like the example scenes from each category in Fig. 3.1. This suggests the algorithm is still able to detect CT reasonably well on islands of different sizes and shapes without changing any of the parameters which we ‘tuned’ for Bermuda. Barbados, with more significant topography, has a signature of higher cloud frequency over the

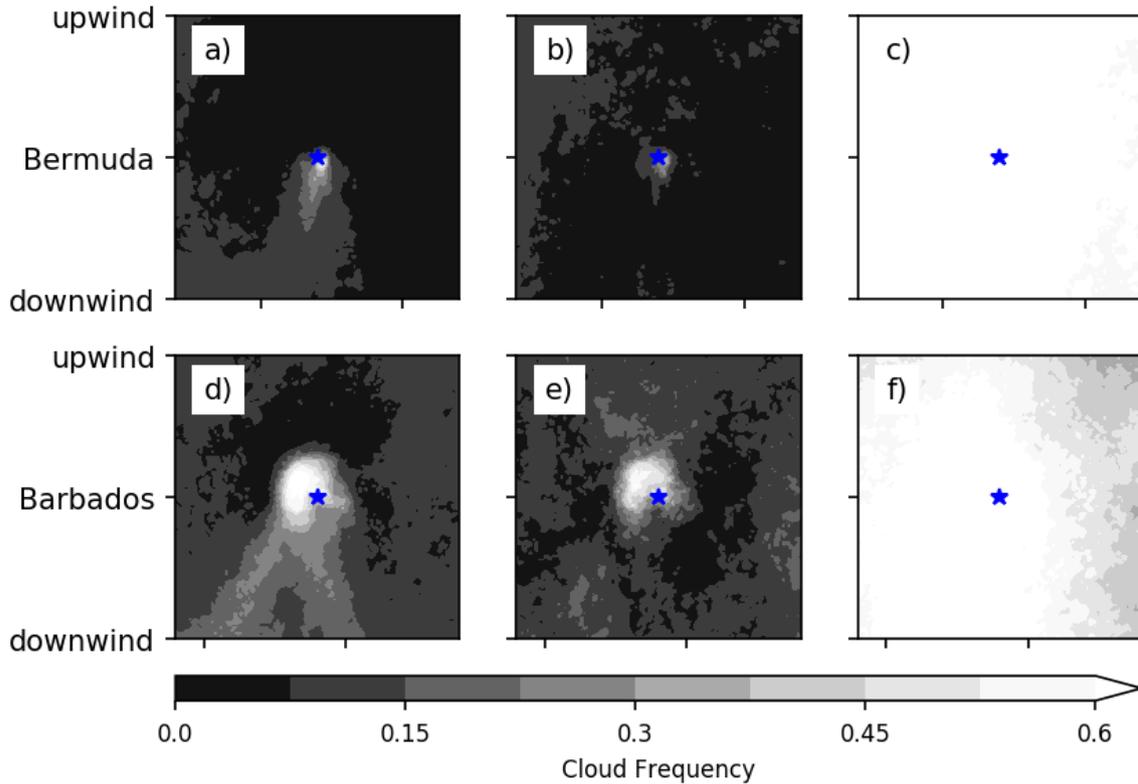


Figure 3.12 Cloud frequency by classification. a), b), and c) are the cloud frequency for the CT, NT, and OB categories respectively for Bermuda. Similarly, d), e), and f) are the same respective categories but for Barbados. A blue star marks the centre-point of the island in question. Each scene is cropped to be $\pm 0.5^\circ$ away from the island in either direction. See text for more information.

island itself for each classification due to mechanical lift by the island. The larger island size and lower solar zenith angles also allow more land pixels to be falsely masked as cloud. Furthermore, tuning the Barbados parameters to fit a manual classification of that data might make a stronger distinction between Fig. 3.12d and e.

We see that there is potential to easily apply the algorithm to other individual islands. Future observational work might use the algorithm to make a closer comparison of CT behaviour and response to thermodynamic profiles for different islands and in different geographic regions. Forecasting guidance could even be developed from a combination of these classifications, radiosonde soundings, and modelled soundings.

3.5 Summary and Conclusions

An algorithm is presented which automates the objective classification of visible-channel geostationary satellite imagery into Cloud Trail (CT), Non-Trail (NT), and Obscured (OB) categories at Bermuda. The algorithm first filters out morning and evening images with high solar zenith angles, and masks cloudy pixels. Next, the algorithm determines whether or not the scene is OB by considering the cloud fraction within a 0.25° radius of Bermuda. For non-OB scenes, the observed 10 m wind direction at the time of the satellite image is then used to define an upwind and downwind region with respect to the island. A scene is then determined to include a CT by considering the difference between downwind and upwind cloud fractions. Classification parameters are defined using a manual classification of images for May through October of 2012. α , the threshold for determining the cloud fraction around Bermuda above which a scene should be considered OB, is found to be 0.33. Then β , the downwind-upwind cloud fraction difference above which a CT is present, is found to be 0.08. An algorithm with these parameters is applied to May through October of 2012-2016 inclusive, totalling 16,400 classified images at Bermuda. We find that the resulting algorithm classifications are generally consistent with the manually classified subset of this period. However, the algorithm is too intermittent in CT detection, and it tends to misclassify imagery more in the mornings.

The resulting five-year climatology shows the expected annual and diurnal cycle in CT occurrence at Bermuda. CT occurrence is found to peak in July and between 1300 and 1500 LT. CT occurrence increases toward July at the same time as a decrease in the occurrence of OB scenes. This corresponds to a period in which the Bermuda-Azores High enforces increasingly settled weather across Bermuda in July. From July to October, the fraction of OB scenes increases and the fraction of CT scenes decreases. The large scale circulation modulates the annual cycle of CT occurrence.

Similarly, CT occurrence increases during the day while the fraction of NT classifications decreases. This is likely in response to the stronger solar heating present in the afternoon, and therefore stronger thermal forcing for the development of CT on otherwise non-OB days. The diurnal cycle of CT occurrence is controlled by the strength

of the solar heating and so follows changes in solar zenith angle with the annual cycle (i.e. fewer CT in May and October compared to July, with CT forming later in the day in those months compared to July).

These classifications are then used to explore the characteristics of the environments present on JJA days predominantly in each classification. Morning (0900 LT) radiosonde composites show that CT-dominated days are characterised by conditions that are near the mean potential temperature and mean relative humidity for the JJA period in 2012-2016. Surface observations and ERA-interim data show that lower than normal low-level potential temperature and relative humidity, lighter than normal 10 m winds, and weak 500-hPa subsidence are associated with NT-dominated days. Lower than normal low-level potential temperature, higher than normal low-level relative humidity, stronger than normal 10 m winds, and strong 500-hPa ascent are associated with OB-dominated days.

Differences in low-level humidity appear to be the most important factor in determining whether or not a non-OB day will have a CT. Days with higher low-level relative humidity have a lower LCL, and therefore the level to which turbulent mixing or lift must reach in order for cloud formation is lower. Given the day does not fall into the OB category, a more humid day is more likely to have a CT. Also of importance is the role of the low-level wind which controls the low-level heating following air across the island. Wind speed controls whether or not the buoyant production of turbulence and the induced thermally-driven circulation is confined to the island in light wind regimes, or whether these circulations can extend downwind of the island in regimes with background flow. Both wind speed and direction control the residence time of air as it crosses the heated island. This is particularly the case for non-circular islands as winds aligned with their long axis result in longer residence times. Longer residence time over the island results in more accumulated heating and therefore a stronger thermal circulation.

Some additional insight into the behaviour of CT through the day is gained. We find that CT can be transient, occurring in several periods through the day - each lasting on average about 3 hours. Other times, a single CT period is observed lasting from its morning formation potentially until after sunset. In an environment that

is only marginally conducive for CT, this transient behaviour might be explained by subtle changes in the environmental relative humidity or island circulation strength that temporarily stop CT formation. This observation is common to both the manual and algorithm classifications.

Based on these results, and ideas from past literature, non-OB days are likely to feature CT given the following:

1. Sufficient low-level humidity and therefore relatively low LCL to support cloud formation
2. Maximised low-level heating from a Lagrangian perspective. This is achieved via long-axis parallel low-level flow and low solar zenith angles.
3. Sufficiently strong low-level flow such that a pure sea breeze circulation isn't favoured over a steady heat island circulation.

These results take a step closer to fully characterising the environments that are conducive for CT formation, however, they also present further questions about their behaviour.

As numerical weather prediction systems currently in use around the world are approaching the ability to resolve phenomena on these scales, it is increasingly important to present a complete characterisation of CT behaviour. For instance, thermally driven circulations may still be present downwind of islands in NT cases, and a correct representation of the strength of this circulation will determine the model skill in cloud and precipitation over or downwind of the island and many island widths downwind.

An investigation of other behavioural features such as the initiation, persistence/-transience, and any transition from shallow to deep convective states are beyond the scope of conventional observations. Furthermore, we know that the strength of the island heating, and the strength of the wind speed are important controls on CT strength. CT length may appear a simple way to observe CT strength (stronger circulation, longer CT), but it also contains significant signal of the environment's ability to disrupt the trail. The observations presented here are therefore deemed insufficient to quantify CT strength, let alone discuss its sensitivity to variations in heating and wind speed.

In the following chapters we perform high resolution simulations to expand our understanding of the CT system. In Chapter 4 an idealised high resolution model is introduced and described (Section 4.2). Then the model is evaluated for its representation of shallow cumulus against a test case (Section 4.3). Finally, our experimental method for cloud trail simulations is described (Section 4.6.2).

This experimental method is then applied to several suites of experiments aimed at better understanding the CT system in Chapter 5. Finally, the representation of CT for different model grid lengths is discussed in the context of shallow convection parametrisation in Chapter 6.

Chapter 4:

Numerics of Shallow Convection: Resolution, Domain, and Idealisa- tions

4.1 Introduction

In Chapter 3, a combination of satellite-based and in-situ observations are used to describe the occurrence and morphology of Cloud Trails at Bermuda. However, these observations are not sufficient to fully explore the relationship between the Cloud Trail's (CT's) organised shallow convection and changes to environmental conditions. The satellite-based observations can only identify cloud locations, and the spatial coverage of conventional in-situ observations does not extend to the cloud trail region downwind of the island.

The land-sea contrast in surface buoyancy fluxes is the first-order driver of mesoscale circulations downwind of heated islands (e.g. [Miller et al., 2003](#); [Kirshbaum and Fairman, 2015](#)). Similarly, it is shown that the low-level relative humidity is important for determining whether or not cloud forms in association with said circulations (e.g. [Johnston et al., 2018](#), and Chapter 3).

Consider the updraught at cloud base in a cloud trail (w_{cb}) as a measure of the strength of the CT circulation. It contains information about the thermally-driven low-level convergence through mass continuity, and is relevant for triggering moist

convection. Quantities such as this have not been adequately measured with available observations for CT. Furthermore, the sensitivity of w_{cb} to changes in environmental relative humidity, strength of the land-sea contrast in buoyancy fluxes, and the ambient wind speed are left unanswered.

In Chapters 5 and 6, numerical experiments with the Met Office Unified Model in its idealised configuration (the idealised UM) are employed to consider quantities such as w_{cb} , and to expand upon their relationship to the physical mechanisms from which they originate. This idealised modelling framework is desirable as it conveniently lends to an evaluation of parameterisation schemes for convection because it reduces the work required to translate results into a more complex model used in operational NWP.

This chapter introduces the idealised UM, defends the idealisations used, and demonstrates that the idealised UM is fit-for-purpose through a steady-state well-resolved shallow convection case with explicit convection (i.e. without convection parameterisation). Finally, a method is outlined for devising appropriate combinations of initial conditions and large-scale forcing terms for the study of CT.

4.2 The Met Office Unified Model

The idealised configuration of the Met Office Unified Model at version 10.9 (the idealised UM) is used for numerical simulations. The idealised UM is chosen over a Large-Eddy Simulation model (LES) in order to take advantage of the direct comparison to the operational version of the model - this is particularly desirable for the evaluation of parameterisation schemes in shallow convection regimes in Chapter 6.

In LES, the energy-producing mechanisms and fluxes of heat and momentum occur on the resolved-scale, while the resolved turbulence energy is dissipated by sub-grid motions which must be parametrised (Mason, 1994). In the idealised UM, the relatively fine grid spacing of LES can be achieved but the dynamical equations are posed in a somewhat different way and are designed to represent a higher level of complexity than traditional LES for operational use.

These experiments are conducted on a limited area domain with bi-periodic bound-

ary conditions. This means that some important dynamical processes would not be represented because they arise from large-scale gradients which cannot be represented on this kind of domain. Idealisations are introduced in order to include these processes. Additionally, certain processes (e.g. radiation) are idealised to reduce their computational cost. Finally, idealisations are also introduced to aid in the isolation of factors of scientific interest. This allows a systematic analytical approach that is consequentially more straightforward. Chapter 5 will make use of such idealisations to explore the sensitivity of CT to environmental relative humidity, ambient wind, and surface forcing.

This section describes the dynamical core of the idealised UM in subsection 4.2.1. The choices for model physics (i.e. the treatment of sub-grid turbulence and microphysics) are discussed in subsection 4.2.2. Finally, the idealisations imposed are discussed in subsection 4.2.3.

4.2.1 ENDGame Dynamical Core

The UM’s ‘Even Newer Dynamics for General atmospheric modelling of the environment’ (ENDGame) dynamical core solves the following set of governing equations for a fully compressible, non-hydrostatic, deep-atmosphere system in cartesian coordinates (x, y, r) where x is the east-west direction, y is the north-south direction and r is vertical distance from the Earth’s centre (Davies et al., 2005).

The prognostic variables include the three wind components, $\mathbf{u} = (u, v, w)$, dry virtual potential temperature, $\theta_{vd} = \theta \left(1 + \frac{m_v}{\epsilon} \right)$, where θ is the potential temperature, ϵ is the ratio of the molar masses of dry air and water vapour, the dry air density ρ_d , and mass mixing ratios for water vapour, m_v , liquid water, m_l , and ice, m_i . These mixing ratios may be further split into cloud liquid, rain, cloud ice, snow, and graupel to reflect selected microphysics options.

$$\frac{D\mathbf{u}}{Dt} + 2\boldsymbol{\Omega} \times \mathbf{u} = - \frac{c_{pd}\theta_{vd}}{\left(1 + \sum_{x=v,cl,cf,\dots} m_x \right)} \nabla\pi + \mathbf{g} + \mathbf{S}^{\mathbf{u}} \quad (4.1a)$$

$$\frac{D\theta_{vd}}{Dt} = S^{\theta_{vd}} \quad (4.1b)$$

$$\frac{D\rho_d}{Dt} + \rho_d \nabla \cdot \mathbf{u} = 0 \quad (4.1c)$$

$$\frac{Dm_x}{Dt} = S^{m_x}, x = v, cl, cf, \dots \quad (4.1d)$$

$$\rho_d = \left(\frac{p_0}{R_d} \right) \frac{\pi^{\frac{1-R_d/c_{pd}}{R_d/c_{pd}}}}{\theta_{vd}} \quad (4.1e)$$

where Ω is the Earth's angular velocity. The first term on the right-hand side of equation 4.1a is the pressure gradient force written in terms of the Exner pressure $\pi = (p/p_0)^{\frac{R}{c_{pd}}}$, p is air pressure, p_0 is a reference air pressure, g is the apparent gravitational vector which is the sum of the actual gravity and the centrifugal force, $R_d \equiv c_{pd} - c_{vd}$ is the specific gas constant for dry air, and R is the specific gas constant including the effect of water vapour. The \mathbf{S} terms on the right hand side of the above equations represent source and sink tendencies from various parameterisation schemes, referred to as 'model physics'.

On the left hand side, the Lagrangian derivative (i.e. following the flow of fluid) is equal to the sum of the local rate of change and advection, $\frac{D}{Dt} = \frac{\partial}{\partial t} + \mathbf{u} \cdot \nabla$. In Cartesian coordinates, this translates as;

$$\frac{D}{Dt} = \frac{\partial}{\partial t} + u \frac{\partial}{\partial x} + v \frac{\partial}{\partial y} + w \frac{\partial}{\partial r} \quad (4.2)$$

where advection is treated with a semi-implicit, semi-Lagrangian numerical scheme.

The governing equations are discretised onto a Cartesian grid in the idealised model. This consists of an Arakawa C-type grid staggering in the horizontal, where u-wind component is on the east-west edges of each grid cell, and v-wind component is on the north-south edges of each grid cell, other variables are in the centre of grid cells (Arakawa and Lamb, 1977).

A Charney-Phillips vertical grid stagger is used (based on Charney and Phillips (1953)). This means that θ , w , and m_x exist on full model levels (theta levels), while ρ , π , u , and v are on half-levels (rho levels). The surface (i.e. the atmospheric level which is in contact with the rigid surface) is on a theta model level.

The formulation of the ENDGame dynamical form is given in UM Documentation Paper 016 in full (Wood et al., 2017). A full description of the UM and all available

model physics can be found at <https://code.metoffice.gov.uk/doc/um/>.

4.2.2 Model Physics

Chapter 5 aims to analyse ‘true’ CTs and the physical relationships between them and the larger-scale environment or forcing. Therefore the horizontal grid spacing is chosen to be fine enough that shallow convection is reasonably well-resolved, and a convection parameterisation scheme is not used. Whereas, in Chapter 6, the representation of CTs is tested at grid spacings for which shallow convection is not well resolved. Experiments with and without convection parameterisation are considered.

In both cases, however, small-scale turbulence and microphysics (i.e. regarding cloud and precipitation) are still not well-resolved. The boundary layer and sub-grid turbulence scheme chosen to treat the poorly resolved turbulence is detailed in subsection 4.2.2.1, the cloud scheme and is outlined in subsection 4.2.2.2, and idealisations which we employ for experiments are described in 4.2.3.

4.2.2.1 Boundary Layer and Sub-grid Turbulence Scheme

To represent the unresolved boundary layer motions and turbulence spectra, a 1-D ‘local’ flux gradient scheme with a ‘non-local’ component based on Lock et al. (2000) is blended with a ‘local’ 3-D Smagorinsky-Lilly type scheme (Smagorinsky, 1963; Lilly, 1967). The blending which is implemented in the present configuration of the idealised UM follows the approach of Boutle et al. (2014b).

The 1-D Lock et al. (2000) scheme treats local turbulent diffusion (i.e. diffusion at a given height) with a ‘first-order’ turbulence closure scheme (this is a common formulation used in global NWP applications), and a representation of non-local turbulent diffusion (i.e. diffusion which could occur across the depth of the boundary layer).

For scalars, χ , this is:

$$\overline{w'\chi'} = -K_h \frac{\partial \chi}{\partial z} + K_h^{NL} \gamma_\chi \quad (4.3)$$

and for vectors (i.e. momentum, \mathbf{u}):

$$\tau = K_m \frac{\partial \mathbf{u}}{\partial z} + \tau^{NL} \quad (4.4)$$

where χ is a conserved variable (either θ_l or q_T), K_h is the eddy diffusivity for scalars, and K_h^{NL} is the eddy diffusivity for non-local scalar fluxes. In the present implementation of the boundary layer scheme, the non-local term only applies to unstable boundary layers and is non-zero for surface-driven transport of θ_l . τ is the wind stress (i.e. $\overline{w'\mathbf{u}'}$), K_m is the eddy viscosity for momentum, and τ^{NL} is the non-local momentum flux. The exact specification of K_h , K_m , and τ depend on the diagnosed boundary layer type - further discussion on this can be found in [Lock et al. \(2000\)](#).

It is more common to represent the sub-grid turbulence in LES with an implementation of a Smagorinsky-style scheme. The UM implementation of this scheme diagnoses eddy viscosity and diffusivity as a function of the model grid spacing and the magnitude of the wind shear as follows:

$$\nu = (c_s \Delta x)^2 S, \quad (4.5)$$

where ν is the sub-grid eddy viscosity for momentum or sub-grid eddy diffusivity for scalars, $c_s = 0.2$ is the Smagorinsky constant, Δx is the horizontal grid spacing (assuming an approximation to isotropic grid is appropriate), and S is the rate of strain tensor which is related to the shear:

$$S = \frac{\|S_{ij}\|}{\sqrt{2}} = \left(\frac{1}{2} \sum_{i,j=1,3} S_{ij}^2 \right)^{\frac{1}{2}}, S_{ij} = \frac{\partial u_j}{\partial x_i} + \frac{\partial u_i}{\partial x_j} \quad (4.6)$$

where subscripts i and j represent Cartesian directions ($1 = x$, $2 = y$, and $3 = z$).

As mentioned earlier in this section, the operational UM blends the [Lock et al. \(2000\)](#) scheme with the Smagorinsky scheme. This blending is used for experiments later in this chapter and in Chapters 5 and 6. The 1-D [Lock et al. \(2000\)](#) scheme is blended with the Smagorinsky scheme as a linear combination of the two schemes with a weighting, W_{1D} , which allows for a smooth transition between the two schemes. As grid spacing decreases, boundary layer turbulence becomes increasingly well-resolved, the Smagorinsky scheme becomes more dominant, and the solution therefore becomes

increasingly like traditional LES. The present implementation of the UM defines the weighting function as follows:

$$W_{1D} = 1 - \tanh\left(\beta \frac{z_{turb}}{\Delta x}\right) \max\left[0, \min\left[1, r_f \left(l_0 - \frac{\Delta x}{z_{turb}}\right)\right]\right] \quad (4.7)$$

where W_{1D} is the weighting for the 1-D boundary layer scheme and takes a value of 1 for grid spacing where turbulence is poorly-resolved, and 0 for grid spacing where turbulence is well-resolved. β is a scaling parameter which controls the transition between fully resolved and fully-unresolved turbulence, z_{turb} is the appropriate length scale of the turbulence (discussed below), and Δx is the model horizontal grid spacing. r_f is a parameter which controls the rate at which the scheme transitions from a purely Lock et al. (2000) scheme to a purely Smagorinsky scheme and is defined as $r_f = \frac{1}{l_0 - l_1}$, $l_0 = 4$, $l_1 = 0.25$ - which is modified slightly from the relation used by Boutle et al. (2014b) by including the max function which forces the scheme to take the pure 1-D form for coarser resolutions. It can be shown that the scheme becomes purely 1-D for $\Delta x = l_0 z_{turb}$. β is defined as follows:

$$\beta = \begin{cases} \beta_{bl} & z < z_h \\ \beta_{bl} \frac{z_{fa} - z}{z_{fa} - z_h} + \beta_{fa} \frac{z - z_h}{z_{fa} - z_h} & z_h \leq z \leq z_{fa} \\ \beta_{fa} & z_{fa} < z \end{cases} \quad (4.8)$$

where $\beta_{bl} = 0.15$ is the value for heights less than the boundary layer height, z_h , and $\beta_{fa} = 1.0$ is the value above the height determined to be the free atmosphere, z_{fa} .

The Lock et al. (2000) scheme contains a representation of turbulence originating from decoupled stratocumulus cloud layers. However, this boundary layer type does not occur for our range of experiments. The turbulence length scale, z_{turb} is taken to be the height of the boundary layer, z_h .

In the simulations of Chapters 5, cumulus convection is expected to be well resolved. The option to blend everywhere is chosen - this is recommended for higher resolution configurations of the UM (Bush et al., 2019). For the simulations of Chapter 6, shallow cumulus is not well resolved, and convection parameterisation schemes are used. For these configurations, it is recommended to blend where not in cumulus

because of undesirable interaction between the boundary layer and convection schemes.

The boundary layer height is taken to be $z_h = z_{par}$ for grid columns which do not include cumulus convection, and $z_h = z_{cl}$ for grid columns which do include cumulus convection. Finally, z_{par} refers to a perturbed adiabatic surface parcel ascent.

The blended eddy diffusivity then has the following formulation:

$$K_{h,m}(Ri) = l^2 S f_{h,m}(Ri), \quad (4.9)$$

where subscript h refers to scalars, subscript m refers to vectors, l is the blended mixing length, S is the shear strain as defined above, and $f_{h,m}(Ri)$ is a stability-based correction function for the diffusivity where stability is measured using the Richardson Number, Ri .

The blended mixing length takes the form:

$$l = W_{1D} l_{bl} + (1 - W_{1D}) l_{smag} \quad (4.10a)$$

$$l_{bl}^{-1} = (kz)^{-1} + \lambda_0^{-1} \quad (4.10b)$$

$$l_{smag}^{-2} = (kz)^{-2} + (c_s \Delta x)^{-2} \quad (4.10c)$$

where l_{bl} is the mixing length for the 1-D scheme, $\lambda_0 = \max(40 \text{ m}, 0.15z_h)$ is the asymptotic mixing length (defined in this way to allow small mixing length for shallow boundary layers), and l_{smag} is the mixing length for the Smagorinsky scheme.

The final full eddy diffusivity is given by:

$$K_\chi = \max[W_{1D} K_{h,m}^{NL}, K_{h,m}(Ri)] \quad (4.11)$$

such that eddy diffusivity takes the down-weighted non-local diffusivity, $K_{h,m}^{NL}$, if the blended diffusivity, $K_{h,m}(Ri)$, is small. The full final turbulent flux is then calculated as:

$$F_h = -K_h \frac{\partial \chi}{\partial z} + W_{1D} F_h^{NL} \quad (4.12a)$$

$$F_m = -K_m \frac{\partial \mathbf{u}}{\partial z} + W_{1D} \tau^{NL} \quad (4.12b)$$

where F_h is the total flux of a conserved scalar, F_h^{NL} is the non-local flux of scalar, and F_m is the total flux of momentum. As the grid spacing decreases, W_{1D} goes to zero, the total fluxes tend to $-K_{h,m} \frac{\partial \chi}{\partial z}$, where $K_{h,m}$ is purely the Smagorinsky value. A more complete description of the boundary layer scheme is given in UM Documentation Paper 024 (Lock et al., 2017), and the subgrid turbulence scheme is given in UM Documentation Paper 028 (Halliwell et al., 2017).

For $z_{turb} = z_h = 700$ m, and grid spacing $\Delta x = 100$ m, W_{1D} is 0.29 for $z < z_h$ and decreases to 0.09 for $z > z_{fa}$. These weights are representative of experiments to be presented in Chapter 5. For Chapter 6, where grid spacing is increased to 800 m and 1.6 km, the weighting increasingly favours the 1-D scheme, this is summarised in Table 4.1.

Table 4.1 Values of the weighting function, W_{1D} for the different grid spacing, Δx used in Chapters 4, 5, and 6. For each grid spacing, W_{1D} is computed for the case where $z_{turb} = z_h = 700$ m.

Δx	$W_{1D}(z < z_h)$	$W_{1D}(z > z_{fa})$
50 m	0.103	0.076
100 m	0.290	0.092
200 m	0.579	0.128
400 m	0.793	0.241
800 m	0.912	0.527
1.6 km	0.974	0.834

4.2.2.2 Microphysics and Cloud Scheme

In the chosen configuration, an instantaneous condensation cloud scheme is used with a modified Smith (1990) parameterisation of the sub-grid variability of saturation specific humidity (q_{sat}). This is chosen over the prognostic cloud scheme as it more closely matches the microphysics methods used in traditional LES (e.g. Mellor, 1977; Sommeria and Deardorff, 1977). A prognostic cloud scheme is more appropriate for simulations using larger grid spacing.

This scheme assumes that conversion of vapour to cloud liquid occurs such that

any supersaturation is instantaneously removed:

$$q_{cl} = q_T - q_{sat}(T, p) \quad (4.13)$$

where q_{cl} is the local condensate, q_T is the local total water content (equal to the sum of condensate and vapour), and $q_{sat}(T, p)$ is the saturation specific humidity at a given temperature T and pressure p respectively. Saturation here is computed with respect to liquid water.

Equation 4.13 is then written in terms of conserved variables (q_T and T_l , the ‘liquid temperature’) through a Taylor Expansion. Assuming that $\frac{\partial q_{sat}}{\partial T}$ at constant pressure, is independent of temperature and that $\frac{\partial q_{sat}}{\partial p}$ is independent of pressure, the expression for condensate can be expanded into terms relating to the grid box mean and perturbations from the grid box mean of our conserved variables.

This reduces the problem to:

$$q_{cl} = Q_c + s \quad (4.14)$$

where Q_c represents mean grid box supersaturation ($\overline{q_T} - q_{sat}(\overline{T_l}, \overline{p})$), and s represents perturbations from the grid box mean with the condition that $q_{cl} \geq 0$ when $s \geq -Q_c$, and $q_{cl} = 0$ when $s < -Q_c$.

Within any particular grid box, a distribution of s occurs (with mean, by definition 0). Considering cloud to be where the water content is greater than zero (i.e. where $s > -Q_c$), the expression for mean condensate, $\overline{q_{cl}}$, using equation 4.14 to expand q_{cl} , is

$$\overline{q_{cl}} = \int_{s=-Q_c}^{\infty} (Q_c + s)G(s)ds \quad (4.15)$$

where $G(s)$ is the assumed distribution of perturbations within a grid box. Note that G is in terms of s therefore there is no need to know the variables q_T , T_L , and p here.

The $G(s)$ distribution is parameterised using a diagnostic symmetric triangular Probability Distribution Function (PDF) of s . The width of the triangle is determined so that condensation occurs when a critical relative humidity, RH_{crit} , is exceeded. For

all grid spacing used in this thesis, the resulting PDF widths are sufficiently narrow for this scheme to essentially be ‘all-or-nothing’.

For this implementation, RH_{crit} is diagnosed from boundary layer fluxes (computed from Section 4.2.2.1). Evidence from flight observations during the Rain in Shallow Cumulus Over the Ocean (RICO) Campaign (Rauber et al., 2007) and the Variability of the American Monsoon Systems (VAMOS) Ocean-Cloud-Atmosphere-Land Study Regional Experiment (VOCALS-REx) (Wood et al., 2011) field campaigns suggest that RH_{crit} decreases with increasing grid spacing, while the PDF width increases with increasing grid spacing (Figure. 4.1).

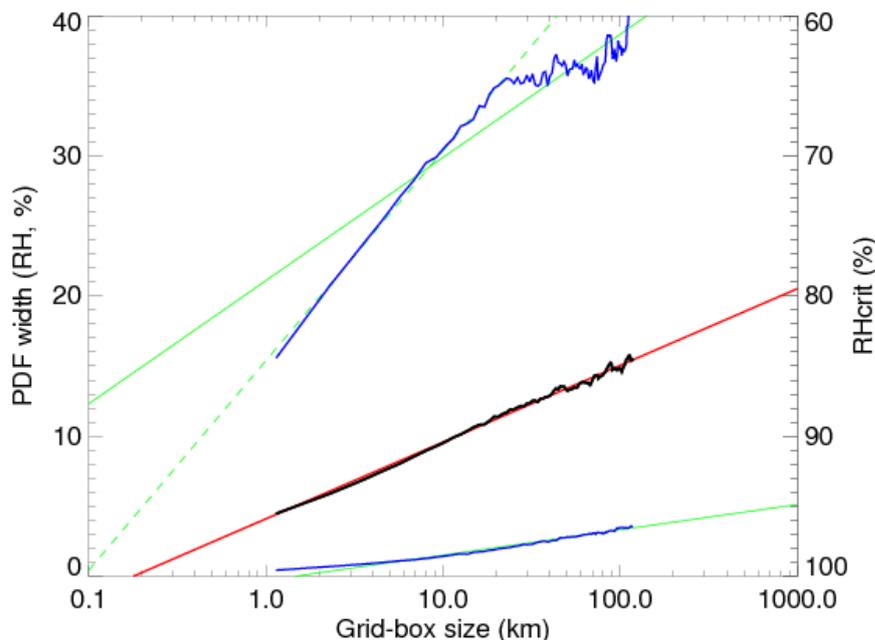


Figure 4.1 Critical relative humidity (RH_{crit}) plotted against horizontal grid spacing (i.e. flight-leg distance from field campaigns). The mean is in black, and the 5th and 95th percentiles are in blue. The 5th percentile is described by two lines of best fit (equation 4.16a), one for grid spacing less than 20 km as a dashed green line, and one for grid spacing greater than 20 km as a solid green line, however, because there is so much variability for grid spacing above 20 km, a constant value is used. The best fit to the mean is in a red solid line, and the best fit for the 95th percentile is a green solid line.

A TKE-based parameterisation is used to diagnose RH_{crit} . However, this can produce a wide range of values and so RH_{crit} is constrained using the above observations. The best fit for the observed 5th percentile marks the minimum and the best fit for the observed 95th percentile marks the maximum allowed RH_{crit} . The best fits are as

follows:

$$100 - RH_{crit_{min}} = \begin{cases} 6.50\ln\Delta x + 15.4 & \Delta x < 20km \\ 40 & \Delta x > 20km \end{cases} \quad (4.16a)$$

$$100 - RH_{crit_{mean}} = 2.38\ln\Delta x + 4.09 \quad (4.16b)$$

$$100 - RH_{crit_{max}} = 0.78\ln\Delta x - 0.30 \quad (4.16c)$$

where Δx is in km. Equation 4.16a corresponds to the dashed green line in Figure 4.1, Equation 4.16b corresponds to the solid red line in Figure 4.1, and Equation 4.16c corresponds to the solid green line in Figure 4.1. A full description of the large-scale cloud scheme is given in UM Documentatino Paper 029 (Wilson et al., 2017).

For grid lengths of 100 m (as used in the current chapter and in Chapter 5), an expected RH_{crit} is 101.4 % and RH_{crit} is limited to values between 99.6 % and 102.1 %. This is a narrow range near saturation which essentially reduces to an all-or-nothing scheme, with results comparable to the traditional methods from LES. For grid spacing of 1.6 km (the maximum used in this thesis in Chapter 6) RH_{crit} would range from 81.5 % to 99.9 %, with an expected value of 94.8 % and is the widest range.

4.2.2.3 Large Scale Rain

The parameterisation of precipitation is based on the mixed-phase, single-moment scheme of Wilson and Ballard (1999) with a prognostic formulation which allows precipitation variables to be advected rather than be removed at the end of each time step. As a result, hydrometeors can cross grid boxes as they fall which is potentially important in high-resolution experiments like those presented in this thesis. This scheme has three hydrometeor species: rain, “ice”, and graupel. However, in all of the experiments presented in this thesis, cloud formation occurs below the freezing level and thus all cloud and precipitation are in the liquid water phase and so only the rain formulation is described here.

A particle size distribution is used for the rain particles, and parametrised as a Gamma Function:

$$n(D) = n_0 D^\alpha e^{-\lambda D} \quad (4.17)$$

where n_0 is the intercept parameter, λ is the slope parameter, and $\alpha = 0.0$ is the constant shape parameter. The single-moment scheme approximates the intercept parameter as $n_0 = n_a \lambda^{n_b}$ where $n_a = 0.22$ and $n_b = 2.2$ are constants based on [Abel and Boutle \(2012\)](#).

A representation of warm-rain processes based on the formulation by [Boutle et al. \(2014a\)](#) is used. This assumes a log-normal distribution of the fractional standard deviation of q_v , f . Satellite and aircraft observations are used to parametrise f . Integrating the process rate over this distribution yields an analytic correction to the process rate as a function of f , and this allows correlations with other sub-grid moisture variables to ensure internal model consistency.

The UM large scale rain routines are capable of coupling to prognostic aerosol schemes, but this coupling (and the evolution of aerosols in general) are neglected in this thesis. See UM Documentation Paper 026 for more information on the large-scale rain scheme ([Forbes et al., 2017](#)). Furthermore, the majority of convection is expected to be non-precipitating for almost all of our experiments.

4.2.2.4 Representation of the Surface

The atmospheric interaction with the sea surface is represented in terms of the surface heat, moisture, and momentum fluxes. There are allowed to be interactive for some initial simulations described in section 4.6.2 of the current chapter. A more idealised representation of the surface interactions is described in 4.2.3.5 and is used for experiments in Chapters 5 and 6.

A Charnock relation is used to compute the roughness length used in the surface momentum flux calculations ([Charnock, 1955](#); [Smith, 1988](#)).

$$z_{0m} = \frac{1.54 \times 10^{-6}}{v_*} + \frac{\alpha}{g} v_*^2 \quad (4.18)$$

where $\alpha = 0.018$ is Charnock's coefficient, and $v_* \equiv |\tau_0/\rho_0|^{1/2}$ is the friction velocity,

interactively computed from the horizontal winds and the stability across the lowest vertical model level.

Similarly, a stability-corrected surface exchange scheme is applied for the surface heat and moisture fluxes. This method assumes that the surface is at the saturated specific humidity for the surface temperature. In all of the simulations with an interactive surface scheme, this surface temperature is prescribed, constant in time and space (more details are found in UM Documentation Paper 024 (Lock et al., 2017)).

The surface fluxes of heat, H_0 , and moisture, E_0 , are:

$$\frac{H_0}{c_{pd}\rho_0} = -\frac{c_H}{c_D^{1/2}}v_*\left(\Delta T + \frac{g}{c_{pd}}(z_1 + z_{0m} - z_{0h})\right) \quad (4.19a)$$

$$\frac{E_0}{\rho_0} = -\frac{c_H}{c_D^{1/2}}v_*\Delta q \quad (4.19b)$$

where c_H is the stability-corrected surface conductance (drag coefficient) for heat and moisture, c_D is the drag coefficient for momentum, ΔT is the difference between the sea surface temperature and the temperature of the lowest vertical model level (z_1), z_{0h} is the roughness length for scalars, Δq is the difference between the saturation specific humidity at the sea surface temperature and the specific humidity of the lowest vertical model level.

4.2.3 Idealisations

The numerical experiments presented herein are performed on a small domain and use boundary conditions which do not support large-scale gradients. Furthermore, high-resolution experiments can become computationally expensive using the full suite of model physics. It is therefore desirable to replace the effect of large-scale gradients which are deemed important to understanding the phenomena of interest, and to reduce the computational cost by replacing expensive physics with simple approximations. In LES experiments (for example in section 4.3) it is common to make similar idealisations.

This subsection discusses the primary idealisations implemented in the idealised UM to replace the effects of large-scale gradients in pressure, temperature, moisture,

and wind.

4.2.3.1 Domain and Numerical Idealisations

As mentioned above, the idealised experiments are conducted on a small-domain over a limited area. This is achieved using bi-periodic boundary conditions which means that flow out of the west boundary enters through the east boundary, similarly flow out of the north boundary enters through the south boundary of the domain.

The idealised UM uses a regular cartesian horizontal grid. In all experiments presented equal grid spacing is used in both the x- and y- directions. The domain is designed to be a channel, elongated in the east-west direction to align with the flow.

A final simplification is applied for Chapters 5 and 6 when an island is introduced to the domain. The island is defined as a circular ‘patch’ of a predefined area. This patch is located on the upwind (east) side of the domain. Finally, the island itself is highly simplified; it is completely flat (i.e. no orography), has a higher roughness length than the surrounding sea, and has a diurnal cycle of surface fluxes (see Section 4.2.3.5). All of which are applied uniformly over the island.

4.2.3.2 Geostrophic Forcing

In order to represent the role of large-scale pressure gradient forces, the large-scale is assumed to be in geostrophic balance. With this assumption, a ‘large-scale’ pressure gradient force can be defined in terms of the geostrophic wind and coriolis parameter (i.e. for the zonal momentum equation $\frac{1}{\rho} \frac{\partial p}{\partial x} = f v_g$).

This takes the form of a modification to the horizontal momentum tendency equations following Siebesma et al. (2003):

$$Q_u = f(v - v_g), \tag{4.20a}$$

$$Q_v = -f(u - u_g) \tag{4.20b}$$

where the Q terms are the geostrophic forcing, f is the Coriolis parameter of the f-plane, u and v are the zonal and meridional components of the wind, and u_g and v_g are

the zonal and meridional components of the prescribed geostrophic wind. A constant in time, vertical profile of geostrophic winds is prescribed and applied across the whole domain for the duration of the simulations.

4.2.3.3 Subsidence

Large-scale gradients in the horizontal winds result in vertical motions (i.e. ascent or subsidence) via mass continuity. A prescribed vertical velocity profile is applied to the horizontally averaged profiles of θ_l and q_T to represent the effects of large-scale vertical velocity as follows:

$$\frac{\partial \phi}{\partial t}_{subs} = -w_{subs} \frac{\partial \phi}{\partial z} \quad (4.21)$$

where ϕ represents either θ_l or q_T , $\frac{\partial \phi}{\partial t}_{subs}$ is the tendency due to the prescribed subsidence, and w_{subs} is the prescribed subsidence profile.

For the case of shallow subtropical convection, the vertical velocity typically provides a subsidence tendency which results in warming and drying of the column - forming part of the large-scale balance against the convective heating and moistening, surface heating and moistening, and radiative cooling.

The idealised implementation uses a single vertical profile of w_{subs} applied to the vertical gradient of the horizontal mean of the conserved variables to arrive at the tendencies. These tendencies are then applied across the whole domain.

4.2.3.4 Radiation

The previous idealisations are used to replace processes that cannot be supported on the domain configuration used. However, for the case of radiation parameterisation schemes, the calculations can be very computationally costly to integrate depending on their level of sophistication.

To avoid the added computational expense of including a full radiation scheme, the role of long-wave radiation is idealised as a prescribed increment to the potential temperature tendency as:

$$\frac{\partial \theta}{\partial t}_{rad} = Q_{rad} \quad (4.22)$$

where subscript $_{rad}$ indicates that this tendency is due to radiation. This is supplied as a single vertical profile across the whole domain, and constant in time.

In the case of (sub)tropical convection, radiation generally represents a net cooling tendency in the lower atmosphere. This forms part of the overall thermodynamic balance with subsidence, convective heating, and surface fluxes.

4.2.3.5 Idealisations of the surface

In the case of CT circulations, the land-sea gradient in buoyancy fluxes has already been demonstrated to drive circulation (e.g. [Kirshbaum and Wang, 2014](#); [Kirshbaum and Fairman, 2015](#)). This was also discussed in Chapter 3. The experiments are designed to test the dependence on this gradient by prescribing different surface fluxes across the domain.

Over sea, the surface sensible and latent heat fluxes are prescribed values constant in space and time. Over land, the sensible and latent heat fluxes are given a simplified diurnal cycle as follows:

$$H = \max \left[0, H_0 \cos \left(\frac{\pi t_0 - t}{2 \Delta t / 2} \right)^{1.5} \right] \quad (4.23a)$$

$$E = \max \left[0, E_0 \cos \left(\frac{\pi t_0 - t}{2 \Delta t / 2} \right)^{1.3} \right] \quad (4.23b)$$

where H_0 and E_0 are the peak sensible and latent heat flux respectively, t_0 is the time of the peak heat flux, t is the current model time, and Δt is the length of day. The max function forces the fluxes to zero at ‘night’ (i.e. when $t \leq t_0 - \Delta t$, or when $t \geq t_0 + \Delta t$), and allows the diurnal cycle during the ‘day’. The cosine functions are to the power 1.5 and 1.3 which follows a fit to observations of the diurnal cycle of surface fluxes over land from [Grabowski et al. \(2006\)](#). In addition, the peak sensible and latent heat fluxes are simplified to occur at the same time and at ‘midday’, which does not necessarily occur in reality. Full radiation and land surface schemes would be needed to account

for the variability in turbulent surface fluxes as they interact with net radiative forcing, winds, sub-surface storage, and cloud-shading.

In this framework, the idealised UM expects a sensible heat flux in the form $H = \rho c_p (\overline{w'T'_L})_{sfc}$ and a latent heat flux of $E = \rho L_v (\overline{w'q'_v})_{sfc}$.

In terms of the surface momentum flux, this is prescribed through a constant in time roughness length for sea $z_{0m,sea} = 0.0002$ m, and land $z_{0m,land} = 0.5$ m in all experiments. The sea surface roughness length is consistent with values typically obtained from test simulations using the an interactive implementation of Charnock's relationship (not shown), and from [WMO \(2018\)](#). For land, the values for roughness length are typical of areas with many, poorly spaced, low-rising obstacles ([WMO, 2018](#)). This is deemed appropriate for representing an island like Bermuda, given that the island is significantly built-up.

In section [4.3](#), this configuration of the idealised UM is tested using a shallow subtropical cumulus case. Results are then compared against an intercomparison of LES results from 10 different modelling centres. The purpose of these tests is to demonstrate that the modelling choices made in this section can explicitly represent shallow cumulus reasonably well.

4.2.3.6 Damping

In all of the numerical experiments presented in this thesis, the idealised UM is used with a rigid model lid at the top. As a result, realistic atmospheric waves (e.g. gravity waves) generated by the dynamics are reflected by the model lid whereas they would continue into and beyond the stratosphere in reality. These reflected waves can then interact with the phenomena of interest in undesirable ways.

To prevent and remove any of these spurious reflections, a damping layer is imposed near the model top. The damping is a relaxation of the vertical velocity back to zero over a prescribed time scale in the form:

$$\frac{\partial w}{\partial t}_{damp} = -\frac{w(x, y, z, t)}{\tau} \quad (4.24)$$

where the time tendency of vertical velocity due to damping, $\frac{\partial w}{\partial t}_{damp}$, is given as the

negative of the actual vertical velocity at a given location, $w(x, y, z)$, divided by the damping time scale, $\tau = 20$ s.

A vertical damping profile with a sine wave shape is applied. The peak of the sine is centred on the rigid model lid, and the damping decreases to zero at a prescribed height. For a model lid of 3 km, the damping decreases to zero at 2250 m, and for a model height of 40 km, the damping vanishes below 20 km.

4.3 Introduction to the BOMEX case

The Barbados Oceanographic and Meteorological EXperiment (BOMEX) refers to a field campaign northeast of Barbados in the subtropical Atlantic trade wind region (e.g. Fig.4.2). Specifically, the mean large-scale environmental conditions from part of Phase 3 of the campaign when largely settled weather was observed, characterised by non-precipitating cumulus (22 - 26 June 1969, [Siebesma and Cuijpers \(1995\)](#)).

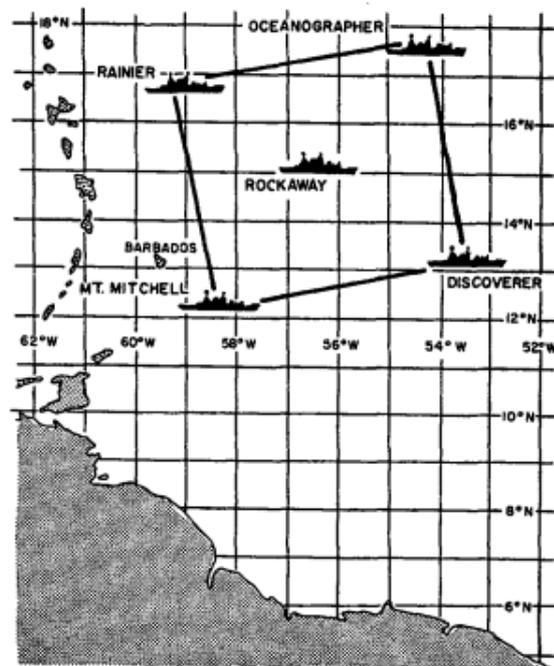


Figure 4.2 Map of the BOMEX ship array during the first three observation periods 3-15 May, 24 May to 10 June, and 19 June to 2 July, 1969; as in [Holland \(1972\)](#) ([©Copyright 2020 AMS](#)).

From this period, the large-scale environmental profiles of temperature, humidity, wind, and surface forcing are derived. [Cuijpers \(1994\)](#) demonstrated that balance

between the large-scale forcing terms can arise in settled conditions such as this. Furthermore, [Siebesma and Cuijpers \(1995\)](#) indicated that a short period was chosen because averaging over many days led to a smoothing of important vertical structure in thermodynamic profiles.

For experiments with island surface forcing, it is desirable for the large-scale conditions to not evolve through the duration of the experiment as this makes it challenging to distinguish between phenomena that arise due to the island, that which arise due to the transient large-scale regime, and those that are due to the presence of both. However, it is important to note that these transient regimes are the norm in reality.

The exact choices made to arrive at a representative BOMEX case are described in the following subsection. The UM is then assessed for its ability to reproduce the horizontal mean macroscopic cloud properties that are reported in [Siebesma et al. \(2003\)](#).

The sensitivity of these results to horizontal grid spacing is then evaluated by repeating the experiment with smaller or larger grid spacing. Finally, the domain lid is increased from 3 km to 40 km, and a stretched vertical level-set is implemented. At each step, the macroscopic cloud properties are assessed as compared to the “true” BOMEX case.

This understanding of the sensitivity of this model to initial conditions and large-scale forcing balance is used to outline a method for arriving at an internally consistent initial state and forcing to limit the amount of change in the background state.

4.3.1 The BOMEX setup

The LES intercomparison experiment outlined in [Siebesma et al. \(2003\)](#) prescribes a domain size and configuration for the participating models. This includes the following: a horizontal grid spacing of $100\text{ m} \times 100\text{ m}$ in the horizontal, and 40 m in the vertical. The model domain contains $64 \times 64 \times 75$ grid points in the x, y, and z directions respectively. This equates to a 6.4 by 6.4 km grid in the horizontal, and a model lid at 3 km in the vertical. Bi-periodic lateral boundary conditions are imposed.

Furthermore, several idealisations are applied. A geostrophic forcing profile is

applied to the momentum budget, constant surface fluxes are imposed, an advective drying profile is applied, and a radiative cooling profile is applied. The idealisations are designed to match the large-scale conditions observed and estimated from the BOMEX field campaign on the selected day. The formulation of these individual idealisations can be found in Section 4.2.3 of this chapter.

These idealisation and the initial conditions imposed are reproduced in more detail below. Note that there are some minor deviations from the outlined procedure in Siebesma et al. (2003) with the intent of achieving the closest vertical energy balances. Any differences are discussed below.

4.3.2 The Equilibrium state

The case outlined in the LES intercomparison study by Siebesma et al. (2003) is used as the basis for the BOMEX simulations presented within. The initial conditions (Fig. 4.3a) consist of a well-mixed layer from the surface to 500 m; a conditionally unstable ‘cloud layer’ between 500 and 1500 m; a statically stable layer with an inversion between 1500 and 2100 m; and a statically stable ‘free atmosphere’ above 2100 m. The initial moisture profile is provided as specific humidity in Siebesma et al. (2003), but the UM-accepted moisture variable is mixing ratio so the conversion is made to arrive at the values presented in Table 4.2a.

Winds are initialised as easterly and unidirectional throughout the column (Fig. 4.3b). Wind speed is constant with height in the 0 to 700 m layer. The wind speed then decreases linearly with height above 700 m, following the geostrophic forcing outlined next.

All of the initial conditions, including surface variables which are assumed to be constant in time, are summarised in Table 4.2a.

The geostrophic forcing profile (Fig. 4.3c) is imposed to maintain the free-atmospheric winds. Here, the geostrophic wind profile is easterly and unidirectional with height, and the geostrophic wind speed decreases with height from -10 m s^{-1} at the surface to 4.61 m s^{-1} at the model lid (3 km), expressed as $\langle u_g, v_g \rangle = \langle -10 + 1.8 \times 10^{-3}z, 0 \rangle$ (Siebesma et al., 2003).

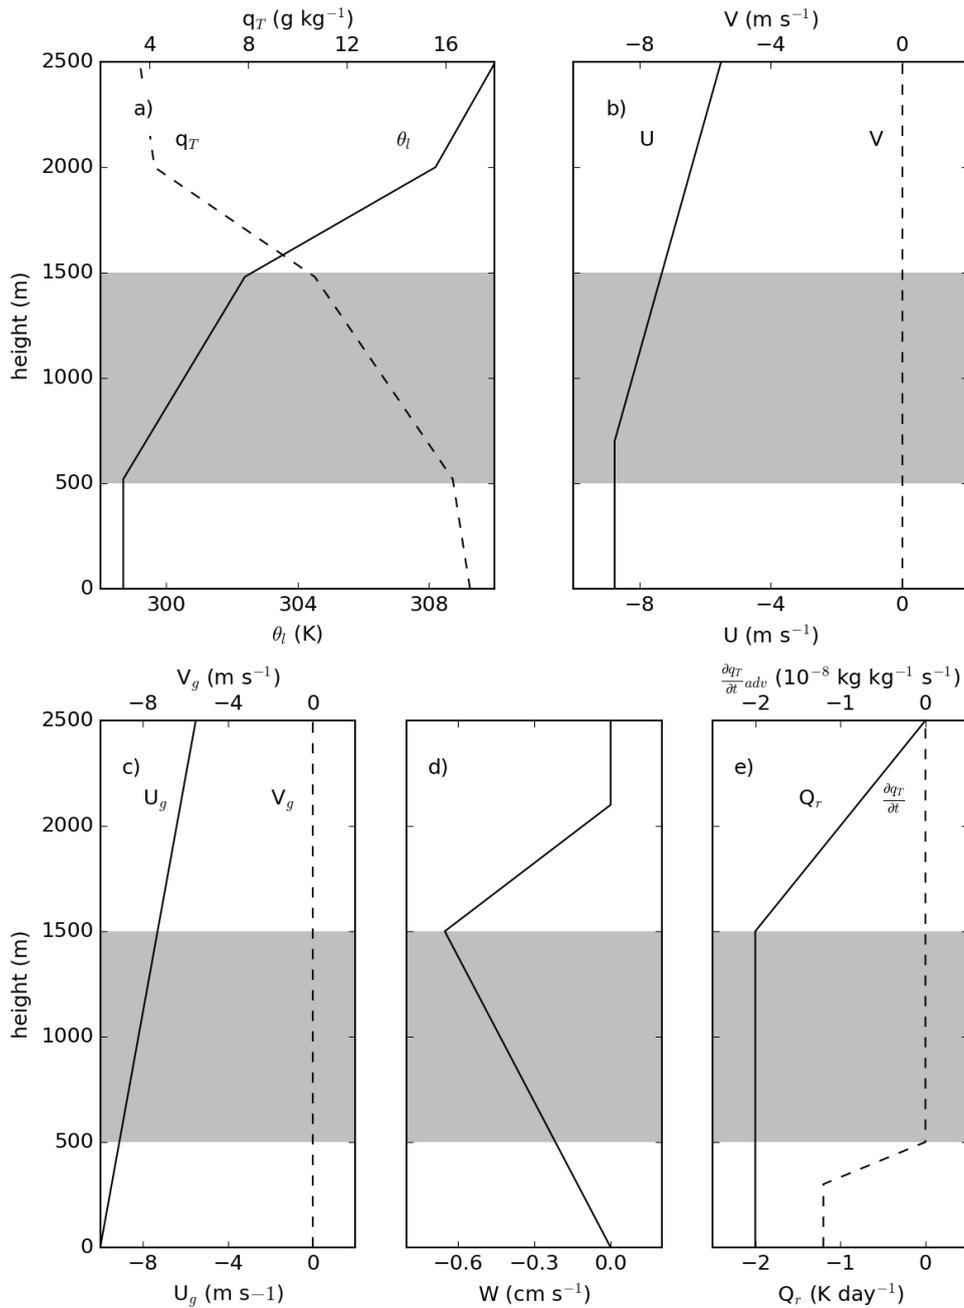


Figure 4.3 The initial conditions and large-scale forcing profiles applied in the BOMEX simulations following the set-up described in Siebesma and Cuijpers (1995) and Siebesma et al. (2003). a) The initial θ (solid line) and q_T (dashed line) profiles. b) the initial u (solid line) and v (dashed line) profiles. c) the geostrophic forcing profiles for u_g (solid line) and v_g (dashed line). d) the subsidence profile. e) the radiative cooling profile (solid line) and the advective drying profile (dashed line). The grey shading between 500 and 1500 m represents the expected cloud layer.

A subsidence forcing (Fig. 4.3d) is also applied with increasing subsidence from 0 cm s^{-1} at the surface to a peak of 0.65 cm s^{-1} at 1500 m. The subsidence decreases

Table 4.2 Profiles of the initial conditions used in BOMEX a), and profiles of the large-scale forcing b). Where a dash is inserted at a height in the table, the values can be acquired through linear interpolation in height between adjacent given data points.

a) Height (m)	θ (K)	\mathbf{m}_v (g kg ⁻¹)	\mathbf{u} (m s ⁻¹)	\mathbf{v} (m s ⁻¹)	
0	298.7	17.3	-8.75	0	
520	298.7	16.6	-	-	
700	-	-	-8.75	-	
1480	302.4	10.8	-	-	
2000	308.2	4.2	-	-	
3000	311.85	3.0	4.61	0	
b) Height (m)	$\frac{\partial\theta}{\partial t rad}$ (K day ⁻¹)	$\frac{\partial m_v}{\partial t adv}$ (g kg ⁻¹ day ⁻¹)	w (cm s ⁻¹)	u_g (m s ⁻¹)	v_g (m s ⁻¹)
0	-2	-1	0	-10	0
300	-	-1	-	-	-
500	-	0	-	-	-
1500	-2	0	-0.65	-	-
2100	-	-	0	-	-
2500	0	-	-	-	-
3000	0	0	0	-4.61	0

again to 0 cm s⁻¹ by 2100 m and remains 0 cm s⁻¹ through to the model lid at 3000 m. The subsidence is applied to the mean profiles of liquid potential temperature θ_l and total water q_T as described in subsection 4.2.3.3.

Finally, a vertical profile of ‘radiative cooling’ and a vertical profile of ‘advective drying’ are applied to the domain (Fig. 4.3e). The radiative cooling profile is used in place of a radiation scheme, which makes the simulations computationally cheaper and forces changes in static stability in a consistent way across different models. The radiative cooling profile used here is different to Siebesma et al. (2003). The profile we use decreases linearly to 0 K day⁻¹ at 2500 m rather than 3000 m following Siebesma and Cuijpers (1995). Idealised vertically integrated energy and moisture budgets balance more closely for the shallower radiative cooling profile. The advective drying profile represents the role of the ambient wind across large-scale moisture gradients which cannot be supported on the limited area, bi-periodic domain. This profile follows Siebesma

et al. (2003) with a small drying tendency in the lowest 300 m, and decreasing to 0 by 520 m.

For a closed system in equilibrium (which is the goal for this scenario) the vertically integrated radiative cooling and subsidence warming should be balanced by surface sensible heat fluxes and latent heating from convection. Similarly, the vertical integral of advective drying and subsidence drying should be balanced by surface latent heat fluxes.

The appropriate surface latent heat flux is computed as a residual of a moisture budget averaged across the observations from the BOMEX square. This yields a value of $170 \pm 25 \text{ W m}^{-2}$ (Holland, 1972). Holland (1972)'s estimate is consistent with later studies by Nitta and Esbensen (1974) and Holland and Rasmusson (1973). However, the surface sensible heat flux is a small residual of large terms in the energy budget and is very uncertain with that method. Measurements of the Bowen ratio during the campaign give β ranging from 0.07 to 0.14. Applying this to estimates of the latent heat flux suggests a sensible heat flux of $17 \pm 6 \text{ W m}^{-2}$.

Model simulations in Siebesma and Cuijpers (1995) give a latent heat flux of 150 W m^{-2} and a sensible heat flux of 4 W m^{-2} . However, those authors indicate that these mean fluxes are biased too low due to their model's over simplistic representation of air-sea fluxes, thus indicating that the surface heat fluxes are the least well constrained part of these budgets.

Siebesma et al. (2003) provides the participants of the model intercomparison with the following values: sensible heat flux, $(\overline{w'\theta'})_{sfc} = 8.0 \times 10^{-3} \text{ K m s}^{-1}$, and latent heat flux $(\overline{w'q'_T})_{sfc} = 5.2 \times 10^{-5} \text{ kg kg}^{-1} \text{ m s}^{-1}$. The surface pressure, $p_* = 1015 \text{ hPa}$, and the surface potential temperature, $\theta_* = 298.7 \text{ K}$, are given so that the surface dry air density can be found, $\rho_* = 1.1771 \text{ kg m}^{-3}$, from the ideal gas law:

$$\rho_* = \frac{p_*}{R_d T_*}, \quad (4.25)$$

where variables with subscript '*' denote surface values. The prescribed surface flux formulation for the UM requires fluxes in energy rather than kinematic form. Furthermore, for sensible heat flux this must be provided in terms of liquid temperature flux,

4.3. Introduction to the BOMEX case

$(\overline{w'T'_L})_{sfc}$, where $T_L = T - \frac{L_v}{c_{pd}}q_{cl}$ is the conserved energy variable. The energy flux can be recovered assuming no surface liquid water content using the following:

$$H_{sfc} = \rho_* \pi_* c_p (\overline{w'\theta'})_{sfc} \quad (4.26)$$

$$E_{sfc} = \rho_* L_v (\overline{w'q'_T})_{sfc} \quad (4.27)$$

where π_* is the surface Exner pressure which is used to convert potential temperature to temperature. This translates to $H_{sfc} = 9.5041 \text{ W m}^{-2}$, and $E_{sfc} = 153.0825 \text{ W m}^{-2}$. These are the values quoted in Table 4.3.

Table 4.3 Single level or constant in height settings for the standard BOMEX configuration. This includes initial random noise, surface pressure, temperature, fluxes, and roughness lengths.

Forcing	Height	Magnitude	Units
θ'_{rand}	0-1580 m	0.1	K
$m'_{v,rand}$	0-1580 m	0.025	g kg ⁻¹
p_*	0 m (Surface)	1015	hPa
T_{sfc}	0 m (Surface)	299.973	K
T_*	0 m (Surface)	300.4	K
H_{sfc}	0 m (Surface)	9.5041	W m ⁻²
E_{sfc}	0 m (Surface)	153.0825	W m ⁻²
z_{0m}	0 m (Surface)	0.0002	m
z_{0h}	0 m (Surface)	0.00002	m

The surface momentum fluxes are standardised by supplying a constant friction velocity (u_*) in Siebesma et al. (2003). However, that doesn't translate directly to the UM formulation for momentum transfer. Instead, this is represented by a constant in time and space roughness length (z_{0m}). In this case, a value of $z_{0m} = 0.0002 \text{ m}$ is prescribed as a sensible choice for open-ocean sea surface conditions (WMO, 2018).

Finally, to break the symmetry of the initial conditions, a field of random noise is imposed on the initial θ and q_v states from 0 to 1580 m. 0.1 K θ perturbations, and 0.00025 kg kg⁻¹ q_v perturbations are imposed following Siebesma et al. (2003).

All of the large-scale forcing profiles, and detailed values for the surface represen-

tation, are summarised in Tables 4.2b and 4.3 respectively.

4.4 Evaluation of the Control BOMEX Case

The BOMEX simulation is run for 24 hours in the idealised UM. The first 6 hours are compared with the results of the Siebesma et al. (2003) intercomparison. This comparison is performed in order to evaluate the general representation of explicit shallow convection in the idealised UM. The total cloud cover, horizontally averaged liquid water path, and horizontally averaged, vertically integrated turbulence kinetic energy (i.e. resolved macro-scale cloud properties) and the equilibrium state are used to quantify this shallow convection representation.

First, fields of liquid water path are inspected to consider more subjective qualities of the simulated clouds (Figure 4.4). Here, liquid water path is taken to be the vertical integral of cloud liquid water. Figure 4.4a shows many clouds with high liquid water content during the initial wave of clouds, or “spin-up”. The spacing between individual clouds increases and a greater mixture of clouds with higher and lower liquid water paths emerge by 180 minutes into the simulation. A range of different horizontal cloud scales start to appear here (Fig 4.4b). There is then some variability in the cloud amount and horizontal-mean liquid water path. Figure 4.4c shows a time with high horizontal-mean liquid water path, and we see that much of that signal comes from a single cloud and most of the clouds have lower liquid water paths. This suggests that bigger individual clouds could disrupt the “steady-state”. Or there is some quasi-steady behaviour marked by oscillations between periods of higher and lower mean cloud liquid.

4.4.1 Cloud Cover, Cloud Liquid, and TKE

Bulk cloud properties are now considered in time-series to quantify how the cloud field is varying through the first 6 hours of the simulation (Fig. 4.5). In each panel, the first 60 minutes of simulation are characterised by a distinct peak and subsequent trough in the three shown measures of cloud properties. This is consistent with numerical spin-up behaviour across the 10 models of the BOMEX intercomparison (black line

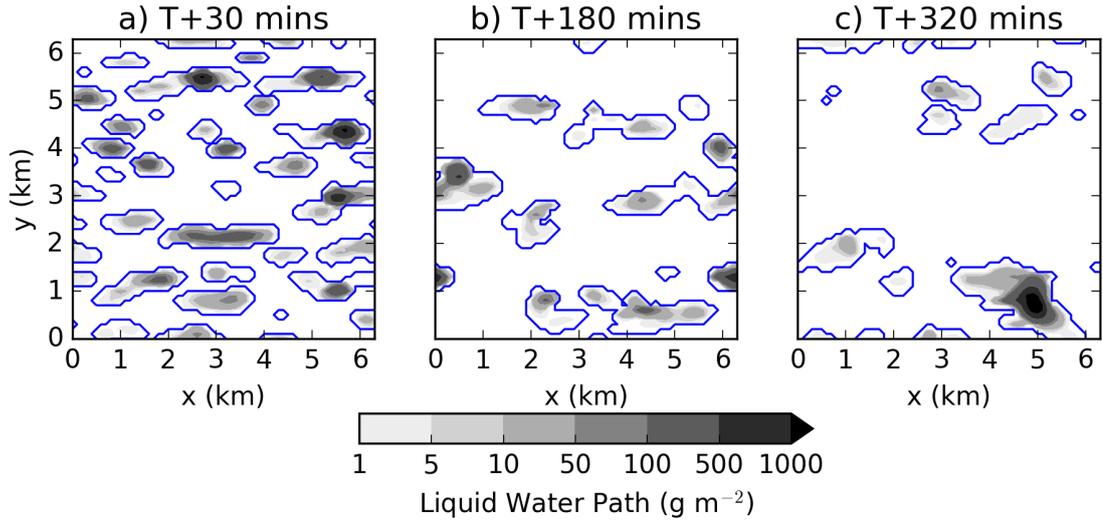


Figure 4.4 Liquid water path shaded in greys at a) during the spin-up, 30 minutes into the simulation; b) at the start of the quasi-steady-state, 180 minutes into the simulation; and c) during a maximum in domain mean liquid water path, 320 minutes into the simulation (see Fig. 4.5b). A blue contour is drawn to mark 0 g m^{-2} liquid water path.

and grey shading in Fig. 4.5). The remaining analysis will focus on activity occurring after this period.

We define the total cloud cover as the fraction of the domain in which liquid water path is non-zero. In general, the total cloud cover (Fig. 4.5a) is slightly high when compared to the Siebesma et al. (2003) intercomparison, but follows the same general pattern. Typical total cloud cover values in the intercomparison are 0.13 ± 0.06 , while a typical total cloud cover is the idealised UM is 0.18. Siebesma et al. (2003) notes that observed cloud cover values are near 0.25. However, the reported range is robust across the intercomparison’s models, and across different variations of the experiment in these models aimed at higher cloud covers.

The horizontal-mean liquid water path (Fig. 4.5b) is also comparable to the intercomparison. This is true in terms of the magnitude of the time-averaged liquid water path over the last three hours, near 5 g m^{-2} , and in terms of the general evolution of the time series. However, there is significant variability on these mean values, with local maxima in horizontal mean liquid water path up to $15 - 20 \text{ g m}^{-2}$ occurring in both the UM and the intercomparison.

Finally, we consider the vertically-integrated, horizontal-mean turbulence kinetic

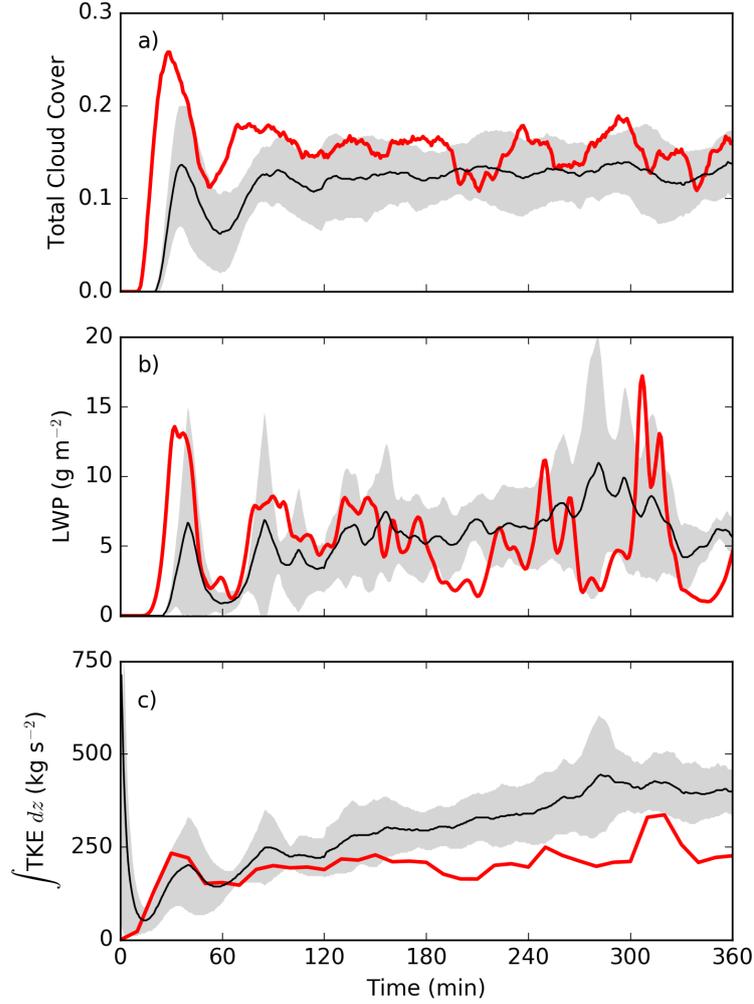


Figure 4.5 A time-series of macro-scale cloud properties for the first 6 hours of our BOMEX simulation. a) The total cloud cover as a fraction of the horizontal space with non-zero liquid water path; b) the horizontal mean liquid water path; and c) the horizontal mean, vertically integrated turbulence kinetic energy. The results from the UM simulation are in red. Estimated from [Siebesma et al. \(2003\)](#) Figure 2 are the mean (black line) and $\pm 2\sigma$ (grey shading) of the 10-model intercomparison.

energy, TKE :

$$TKE = \frac{1}{2} \int \rho(u'^2 + v'^2 + w'^2)^{0.5} dz \quad (4.28)$$

where u' , v' , and w' are the perturbations of the three wind velocity components from the horizontal mean. In the idealised UM, TKE is fairly steady after the spin-up period. Furthermore, it is of similar magnitude to the ([Siebesma et al., 2003](#)) results (4.5c). In the intercomparison, the vertically integrated TKE steadily and slowly increases to

near $350 \text{ kg m}^{-1} \text{ s}^{-2}$ through the 6 hours, while in the UM, it remains somewhat lower, near $250 \text{ kg m}^{-1} \text{ s}^{-2}$. It should be noted that our values for the UM do not include any sub-grid scale contribution to the TKE and so are expected to be lower.

4.4.2 Maintaining Equilibrium State

The BOMEX simulations are designed such that the large-scale forcing terms are balanced in the vertical. This should allow for only small changes to the horizontal mean thermodynamic profiles. The mean profiles over model hours three through six are shown in Figure 4.6), which also includes curves to indicate the initial conditions (dashed black lines). The UM maintains the equilibrium reasonably well, and deviations from the initial conditions are of the same order of magnitude as in the intercomparison.

The mean θ and q_v profiles are largely unchanged between the initial and 3-hour mean conditions (Fig. 4.6a and b). Meanwhile, the boundary layer winds have adjusted to the surface drag. Through balance with the geostrophic forcing profiles, a northerly component has developed within the boundary layer (Fig. 4.6c). Both the unchanged thermodynamic profiles and the development of a characteristic log-wind surface layer in the wind profiles are comparable to the changes noted in the intercomparison.

The cloud liquid mixing ratio, q_{cl} , behaves slightly differently. While the mean q_{cl} profile from the idealised UM maintains cloud base near the intercomparison, there are instances in which cloud tops exceed the range reported in Siebesma et al. (2003). There is also somewhat less q_{cl} in the idealised UM's cloud layer. All of this behaviour could be simply explained by the inclusion of a local maximum in LWP in the UM analysis presented here (see near 320 minutes in Fig. 4.5b), and a similar local maximum not being included in the averaging period for the BOMEX intercomparison as it occurs earlier in the mean of the intercomparison simulations. Averages over other times which do not include this local maximum in LWP have a cloud layer depth more similar to the intercomparison.

It should be noted that in the intercomparison, a one-hour averaging period is used, whereas a three-hour averaging period is presented here to reduce the noise in

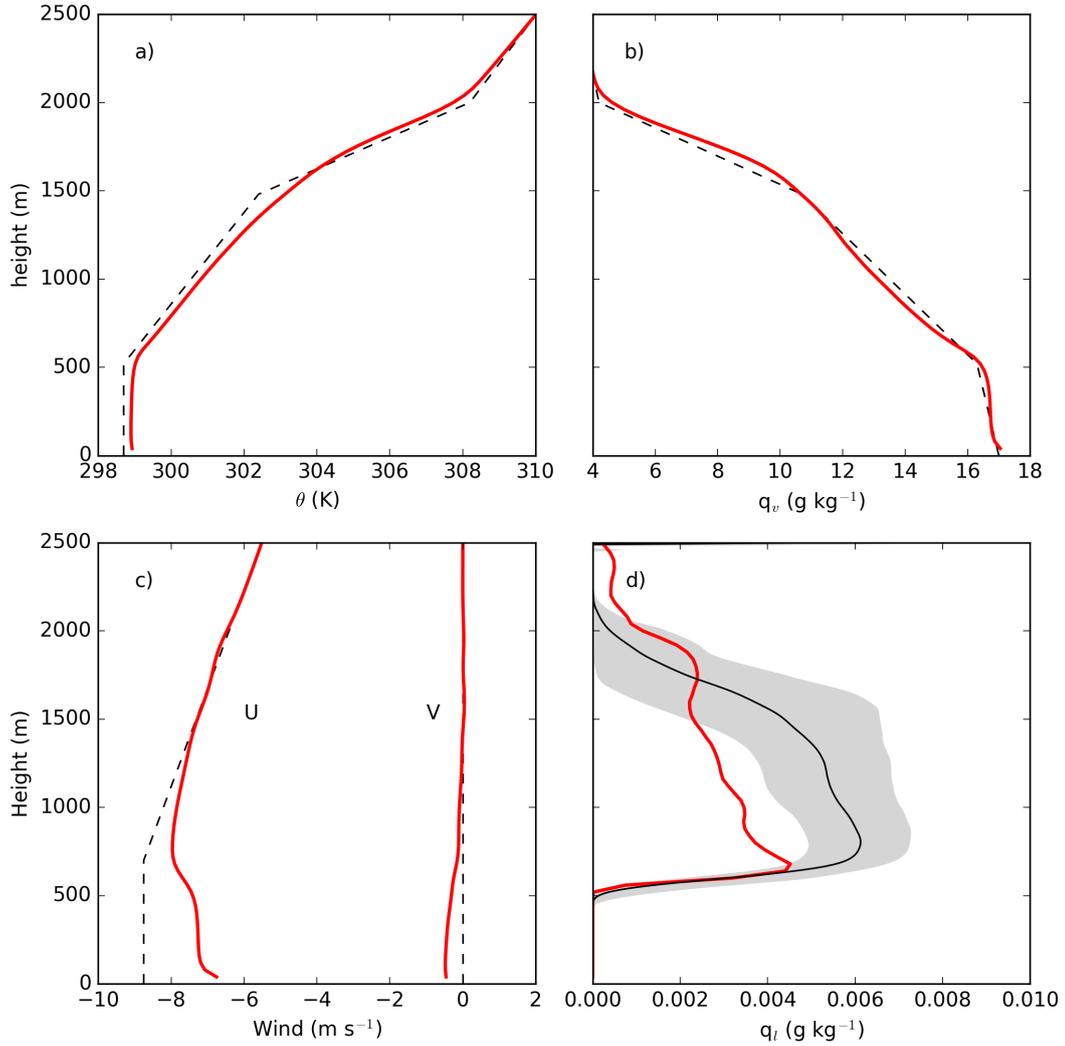


Figure 4.6 Horizontal-mean profiles over the last 3 hours of a) potential temperature, θ , b) water vapour specific humidity, q_v , c) the horizontal wind components, u and v , d) the cloud liquid water, q_l . The red solid lines represent the mean profile during this time-frame. In a), b), and c) the dashed black line represents the initial conditions. Estimated from Siebesma et al. (2003) Figure 3d are the mean (black line) and $\pm 2\sigma$ (grey shading) of the 10-model intercomparison for the cloud liquid water.

the cloud water content profile. The only profile presented that is significantly affected by this change is the q_{cl} profile as the others maintain smaller fluctuations about a steady-state.

4.5 Sensitivity to grid and domain changes

The idealised UM has been shown to reproduce the bulk cloud properties of the BOMEX experiment reasonably well when compared to the [Siebesma et al. \(2003\)](#) intercomparison. However, there are some numerical considerations which need to be addressed before experimenting with CT scenarios. For instance, the UKV (an operational convection-permitting configuration of the UM for regional modelling) does not employ the regular vertical grid prescribed by BOMEX, nor does it operate with such small grid lengths. These numerical properties of the UKV are currently common to other regional models in operation at meteorological centres globally.

We start by considering how increasing the horizontal grid length impacts the representation of shallow convection in the BOMEX case. The 100 m grid length used in the original BOMEX intercomparison brings inherent computational cost constraints with it. However, it is unclear whether grid lengths this small are necessary for shallow convection to be adequately represented in our implementation of the UM. This section therefore has dual aims: to determine whether a reasonable ‘truth’ CT simulation can be achieved with larger grid spacing (this becomes important in [Chapter 5](#)); and to provide a cursory look at the representation of shallow convection when it is poorly-resolved (this becomes important in [Chapter 6](#)).

Secondly, the BOMEX simulations prescribe a rigid model lid at 3 km height. This is problematic as gravity waves, for example, will reflect off of this lid where their reflections can subsequently have detrimental impacts on the representation of the phenomena of interest. To prevent these reflections in BOMEX, a damping layer is imposed near the domain top (see [Section 4.2.3.6](#)). BOMEX convection is shallow enough to not interact with this damping layer in general. However, a model lid at this height is inappropriate for attempting to represent deeper convection, while the evolution of the upper troposphere/lower stratosphere is important for the development of larger-scale weather systems which may trigger and organise convection. In the context of CTs, they have been observed to form in environments with a wide range of cloud layer depths, and transitions between primarily shallow to deep convection can occur.

In this section, the experiments to test sensitivity to horizontal grid spacing and the height of the rigid model lid are outlined and presented. At the end of this section, the results of these experiments are discussed.

4.5.1 BOMEX at Coarser Resolution

A set of experiments repeating the above BOMEX experiment using different horizontal grid lengths is performed. Only the horizontal grid spacing and the time step are changed. The chosen grid spacings are 50 m, 200 m, 400 m, and 800 m (these are summarised in Table 4.4). The 50 m experiment is repeated on a 128×128 grid to match the area of the Control (100 m) experiment and is indistinguishable from the original 50 m experiment (not shown).

Table 4.4 List of the experiments conducted to test sensitivity of shallow convection to the horizontal grid spacing using variations on the BOMEX case. All experiments have the same number of horizontal grid points (64×64). The Control experiment from Section 4.3 is marked in boldface font.

EXP	Grid Spacing (m)	Time Step (s)
DX050	50	1
DX100	100	3
DX200	200	6
DX400	400	12
DX800	800	24

The boundary layer and sub-grid turbulence scheme employs a blending between a 1-D boundary layer scheme (based on Lock et al. (2000)) and a 3-D Smagorinsky scheme. It smoothly transitions from the 1-D scheme to the 3-D scheme as the grid length decreases (more details in section 4.2.2.1). This blending compromise is meant to result in a sensible representation of sub-grid turbulence through the range of grid lengths for which neither the 1-D or 3-D scheme are intended. Therefore, there is no manual change of scheme for different Δx , but change in behaviour is expected.

Any resulting changes in the representation of sub-grid turbulence have consequences for the representation of clouds as sub-grid variability is used to diagnose an RH_{crit} distribution. For all grid lengths considered, this value is constrained by observations to a narrow range and remains essentially all-or-nothing. Furthermore, this has the advantageous effect of RH_{crit} varying in an internally-consistent manner across

our experiments based on grid spacing.

We return to the macro-scale cloud properties (introduced in Section 4.4.1) to see how they are impacted by changing grid spacing. We present box plots over model hours three through six off our BOMEX grid length experiments. We find that the total cloud cover decreases with increasing grid spacing (Fig. 4.7a). A comment from Siebesma et al. (2003) was that cloud cover was fairly low in the intercomparison compared to observed values. Participants in that study experimented with aspects of their sub-grid schemes and the initial conditions in unsuccessful attempts to reconcile this discrepancy - perhaps decreasing the grid spacing provides the answer.

A more complex relationship with liquid water path emerges. The median liquid water path increases with grid spacing between 50 and 400 m, but then falls off in DX800 (Fig. 4.7b). Meanwhile, the variability in liquid water path is large only for DX400. We suspect that this behaviour is related to the strength of resolved convective circulations which grow stronger with increasing grid length. This relationship is reflected in Figure 4.7c as the resolved vertically-integrated TKE increases substantially when grid lengths are increased to 400 and 800 m. Occasional very strong individual convective cells form and suppress much of the surrounding convection both in space and for some time following that burst of convection (not shown). Some such convective plumes extend to near the model lid in DX800 (Fig. 4.8). The maximum cloud top height reaches less than 1.9 km in DX0500, 2.5 km in DX100 and DX200, and 3 km in DX400 and DX800.

As introduced in Section 4.2.2.1, the model development expectation is that for the small grid spacing of LES, the energy-producing motions and fluxes of heat and momentum occur on the resolved scales, while the sub-grid scheme dissipates the energy (Mason, 1994). With increasing grid spacing, instability must accumulate in a nonphysical way so that motions can arise on the resolved-scale. This results in resolved-scale motions and therefore turbulent motions which are too strong.

Notably, DX800 and DX050 have similar liquid water paths, but they have vastly different total cloud cover and vertically-integrated TKE ranges. This implies fewer, wetter clouds supported by stronger convective updraughts in DX800. The lack of clouds is offset by their water content in order to arrive at similar horizontal-mean

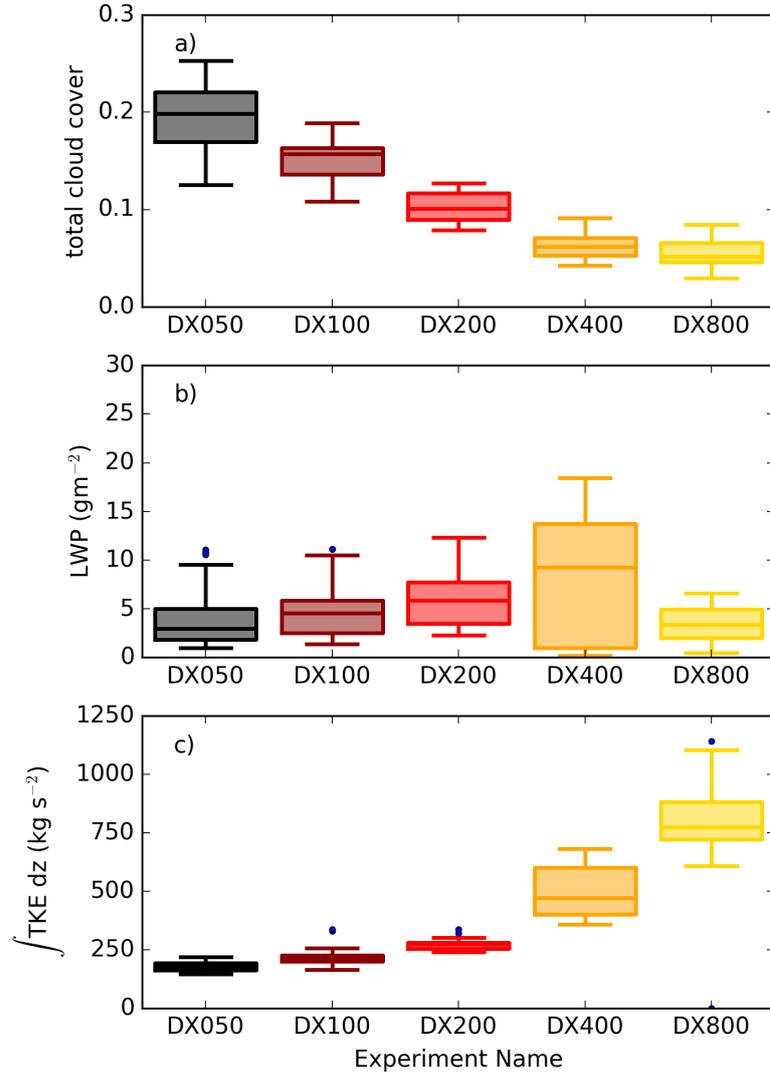


Figure 4.7 Box plots of the horizontal-mean conditions over model hours three through six of the BOMEX grid length experiments summarised in Table 4.4. Presented are the total cloud cover in a), the liquid water path in b), and the vertically-integrated TKE in c). For each box, the middle line represents the median value, the top and bottom of the box represent the interquartile range, and the top and bottom of the whiskers represent the interdecile range. Small blue dots either side of the whiskers mark outliers.

liquid water paths.

Increasing the grid spacing also increases the delay to the onset of convective initiation in these experiments (Fig. 4.9a). We define convection initiation to be the earliest time that total cloud cover is greater than or equal to 0.05. A delayed onset is consistent with the concept that greater resolved convective instability is needed to initiate explicit convection with larger grid lengths in the shallow cumulus grey

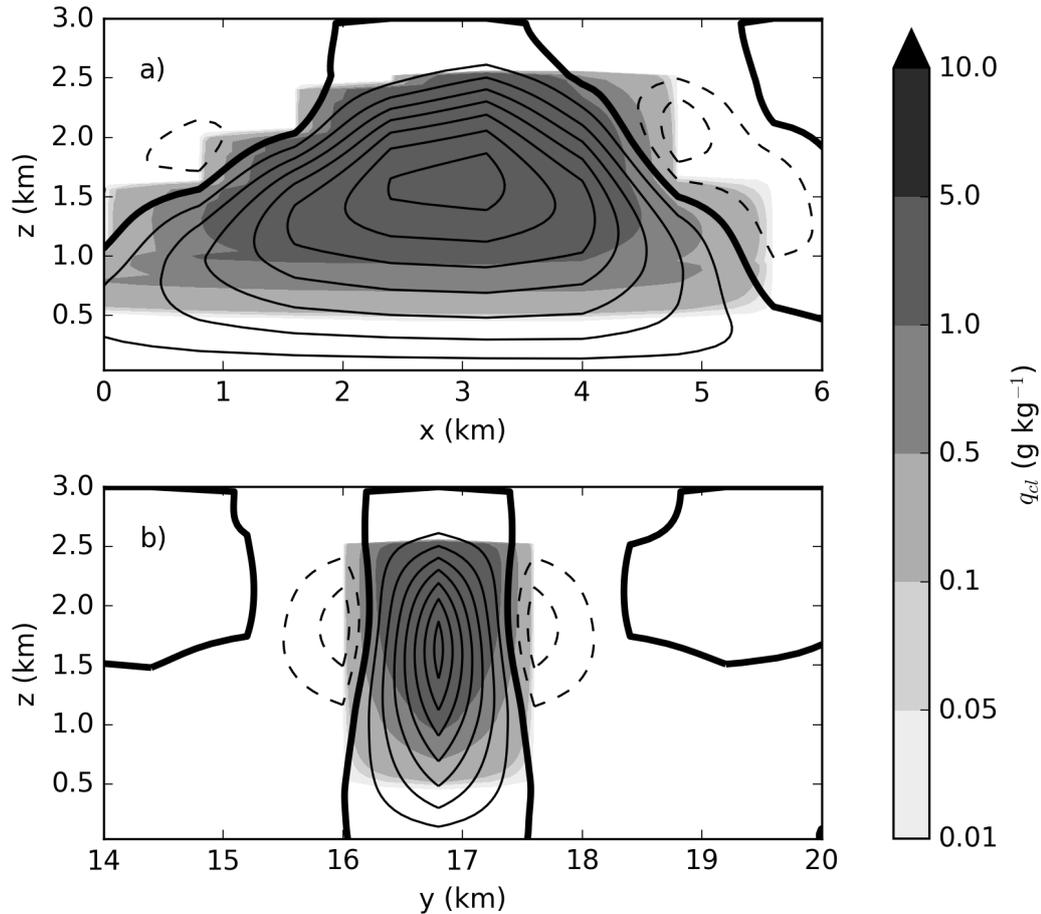


Figure 4.8 Vertical cross sections of through a particularly strong convective plume in DX800. a) an x-z cross section, and b) a y-z cross section. Cloud liquid water, q_{cl} , is presented in grey shading with black contours of vertical velocity every 1 m s^{-1} , w , overlaid (solid where positive, dashed where negative, a heavy black line marks the zero vertical velocity).

zone. We find that this delay in initiation increases by roughly 1 second for every 1 m increase in the grid spacing. The convection making up this “first wave” covers a smaller fraction of the domain (Fig. 4.9b). However, the fewer convective cells are much wetter leading to a higher peak in liquid water path during this time, and they are associated with stronger circulations associated with a much higher peak in vertically-integrated TKE at this time (Fig. 4.9c and d). These results are consistent with expectations for delayed convective onset.

Finally, the experiments with the largest grid spacing tend to drift furthest away from the initial conditions. In hour 5 of DX400 and DX800, the horizontal-mean, time-mean θ and q_v profiles are furthest away from their initial states. We introduce a mean

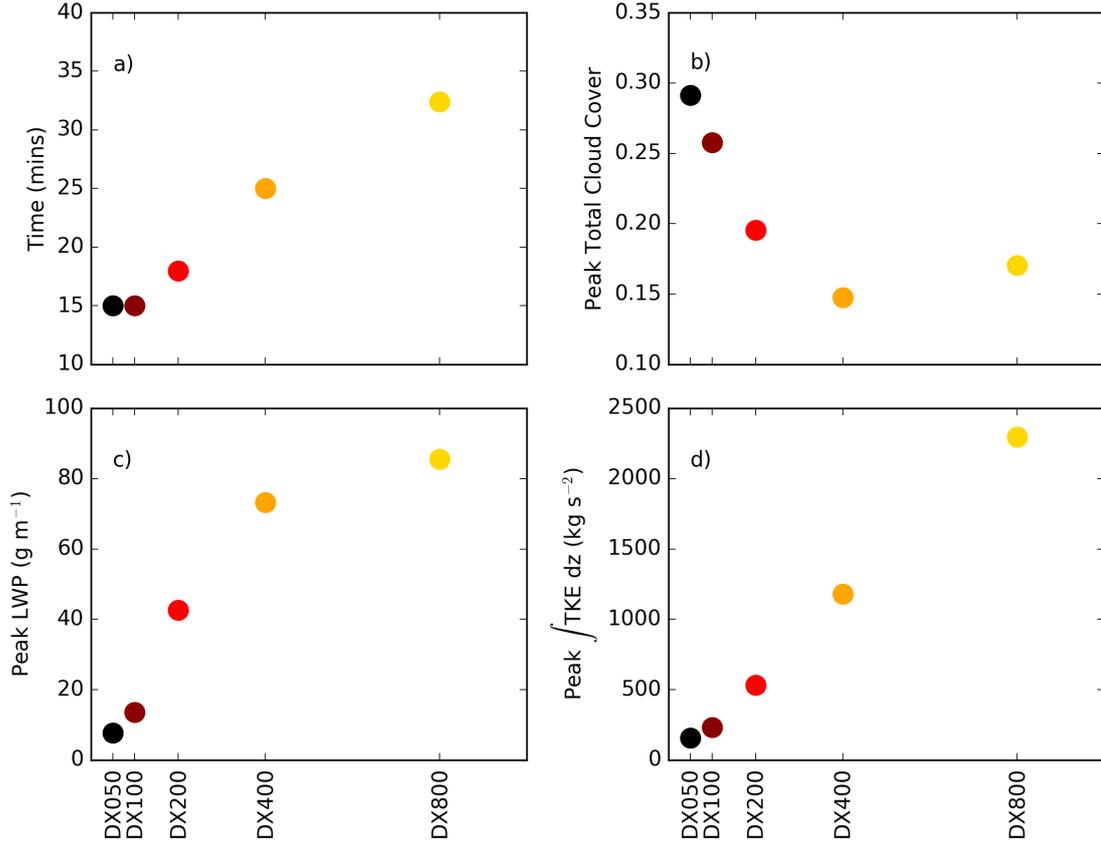


Figure 4.9 A summary of the spin-up behaviour across the BOMEX grid length experiments. a) the time at which total cloud cover first increases to 0.05, b) the peak total cloud cover during the spin-up, c) the peak liquid water path during the spin-up, and d) the peak vertically-integrated TKE during the spin-up. In each panel dots are coloured by grid length.

percentage change to quantify this drift:

$$\delta_0 = 100 \times \frac{1}{z_{max}} \int_0^{z_{max}} \frac{\langle \phi(z) \rangle - \langle \phi_0(z) \rangle}{\langle \phi_0(z) \rangle} \partial z \quad (4.29)$$

where δ_0 is the vertically averaged percentage change between a quantity, $\phi(z)$, horizontally averaged over the whole domain and temporally averaged over the last simulated hour, $\langle \phi(z) \rangle$, and the initial horizontally averaged quantity, $\langle \phi_0(z) \rangle$.

For θ , δ_0 is small for all experiments and is not shown. Larger drifts emerge for the q_v profiles. In DX050 and DX100, δ_0 is a small 1.2 and 2.6 % respectively. δ_0 then increases dramatically with grid spacing to 12.8 % for DX800 where the stronger convective plumes are moistening near the top of the model - in a layer that has no prescribed forcing to balance these tendencies. The change in q_v profiles and associated

δ_0 are presented in Figure 4.10. These changes imply stronger cloud top detrainment for increasing grid length, and less mixing/detrainment within the cloud layer.

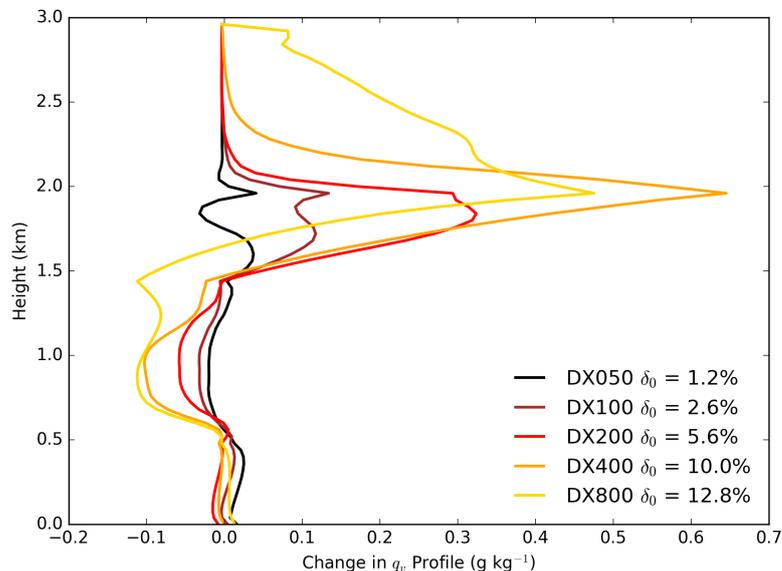


Figure 4.10 Difference between the initial q_v profile and the horizontal-mean, 3-hour time-mean q_v profiles for model hour three through six for the BOMEX grid length experiments. The percentage change from the initial conditions is reported next to the corresponding legend entry.

We find that shallow convection is not entirely numerically converged at 100 m. Total cloud cover increases quite a bit when grid length is halved from 100 to 50 m. However, convective circulations appear to be mostly converged based on horizontal-mean liquid water paths, vertically-integrated TKE, and respect for the large-scale equilibrium. Similarly, a crude representation of shallow cumulus is likely possible for grid lengths of 200 m, but equilibrium is not as well maintained which we suspect is related to stronger resolved circulations. We conclude that 100 m grid spacing is sufficiently fine for the purposes of this study. The next consideration is how the vertical configuration impacts the representation of shallow convection.

4.5.2 BOMEX with rigid model lid at 40 km

We have now shown that grid lengths are sufficient for shallow convection in our implementation of the idealised UM. We now want to determine how our choice of vertical domain structure impacts the representation of shallow convection. We repeat

our 100 m grid length Control BOMEX simulation with the model lid placed at 40 km. In place of the regularly spaced vertical levels used in the Control experiment, we introduce a stretched set of 140 levels as is used in operational configurations of the UM with sub-kilometre-scale grid lengths. The stretched level-set consists of small separation between levels near the surface and the separation between vertical levels increases with height (Fig. 4.11). This potentially supports a better representation of the surface layer and perhaps the overall boundary layer compared to the regular spacing employed in the Control BOMEX experiment. The spacing of the new vertical levels is such that the number of levels between the surface 3 km height is similar to the control experiment.

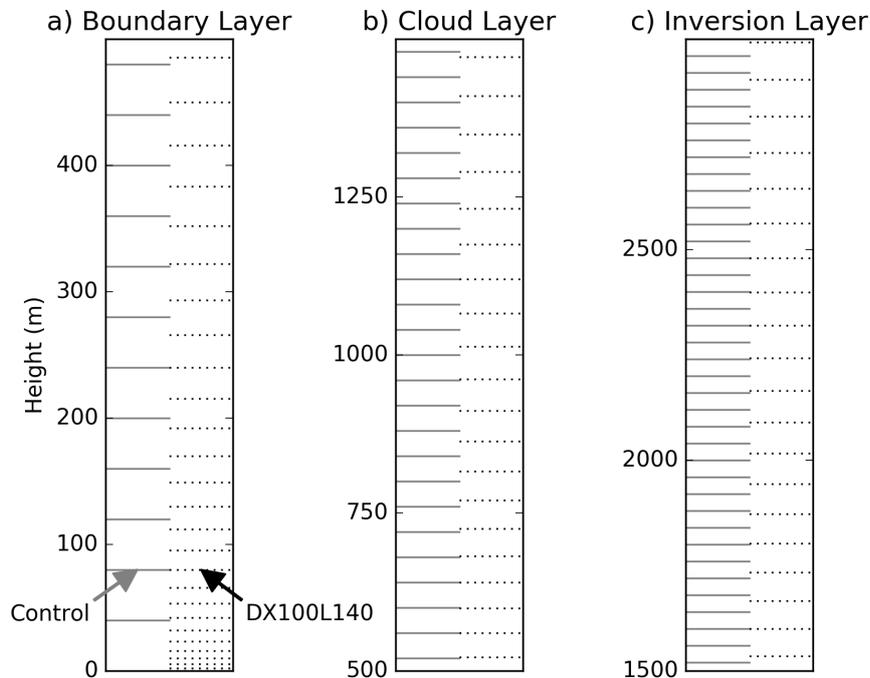


Figure 4.11 An illustration of the spacing between vertical levels for the control experiment (solid grey lines) and DX100L140 (dotted black lines). Three layers in the lowest 3 km are shown.

To complete the initial conditions between 3 km and 40 km, the temperature profile follows a moist pseudoadiabat to 15 km, then linearly increases with height to 1200 K at the model lid. The mixing ratio decreases linearly with height to the moisture resetting value ($q_{v,min} = 10^{-8} \text{ kg kg}^{-1}$) near 5500 m. The magnitude of both the zonal geostrophic forcing is allowed to continue the original linear decrease from -10 m s^{-1} at the surface to zero (at 5500 m), it is then kept at zero above that level. The zonal

wind initial profile is unchanged below 3 km, then follows the geostrophic wind profile above 3 km. All other large-scale forcing profiles are kept at zero above 3 km, and all other initial condition and forcing profiles are unchanged.

The improved representation of the surface layer in our new experiment, “DX100L140”, could have implications for how the prescribed surface fluxes are vertically redistributed. Large changes here may impact the representation of shallow convection, and the large-scale tendencies that result.

Indeed, notable changes to the macro-scale cloud properties occur in DX100L140. The initial wave of convection has higher total cloud cover, liquid water path, and vertically-integrated TKE than the control (Fig. 4.12). The magnitude of the differences in TKE is comparable to the differences between the control experiment and DX200 or DX400 during their respective spin-up period. Note that for a fair comparison between the DX100L140 and the previous experiments, vertical integrals are computed between the surface and 3 km only.

In addition, the boundary layer has warmed and dried in DX100L140 when comparing the horizontal-mean, time-mean profiles of θ and q_v over model hours three through six (not shown). This results in higher cloud bases, and less q_{cl} . These also manifest as $\delta_{0,\theta}$ of 0.242 % and δ_{0,q_v} of 6.9 % which are comparable to DX200. Furthermore, the wind profile is further away from the geostrophic wind profile within the cloud layer than in the control experiment or the intercomparison.

These changes to the boundary layer are reflected subtly in the turbulent flux profiles. In the boundary layer, $\overline{w'q_T'}$ is nearly constant with height, and the fluxes are even lower in the cloud layer - related to decreased cloud amount and liquid content. However, because the cloud layer is deeper, these profiles also extend through a deeper layer. Some of the differences in cloud layer depth could be explained by moving the damping layer away from the top of the cloud layer. The remaining differences in the turbulent fluxes are assumed to be due to the change in vertical levels, in particular, better resolution of the surface layer.

We suspect that the problem lies with the increased resolution within the surface layer. In the control, the 40 m separation between vertical levels only allows a weak thermodynamic surface layer to form whereas a shallow superadiabatic layer (i.e.

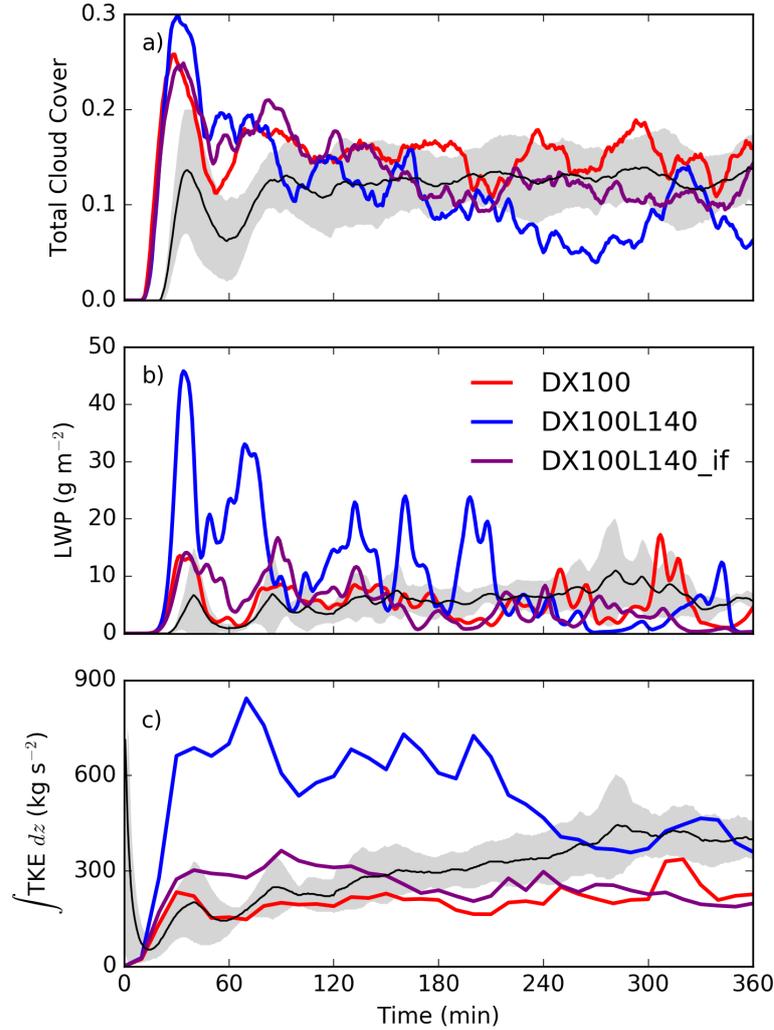


Figure 4.12 As in Fig. 4.5 with the DX100L140 values added to the figure in blue. An experiment with interactive surface fluxes, DX100L140if (to be introduced below), is also included in purple.

$\partial\theta/\partial z < 0$) with a sharp hydrolapse (i.e. large $\partial q_v/\partial z$) forms in DX100L140 (Fig. 4.13). This affects the communication of the prescribed surface fluxes to the boundary layer turbulence scheme through changes in near-surface vertical gradients. We suspect the prescribed surface fluxes are too high to maintain equilibrium given these changes to the surface layer. To test this hypothesis, we re-run DX100L140 with interactive surface fluxes, “DX100L140_if”.

In DX100L140_if, the macro-scale time series collapses back to match the control experiment without the notable decrease in total cloud cover with time in DX100L140, and with horizontal-mean liquid water paths and vertically-integrated TKE of similar

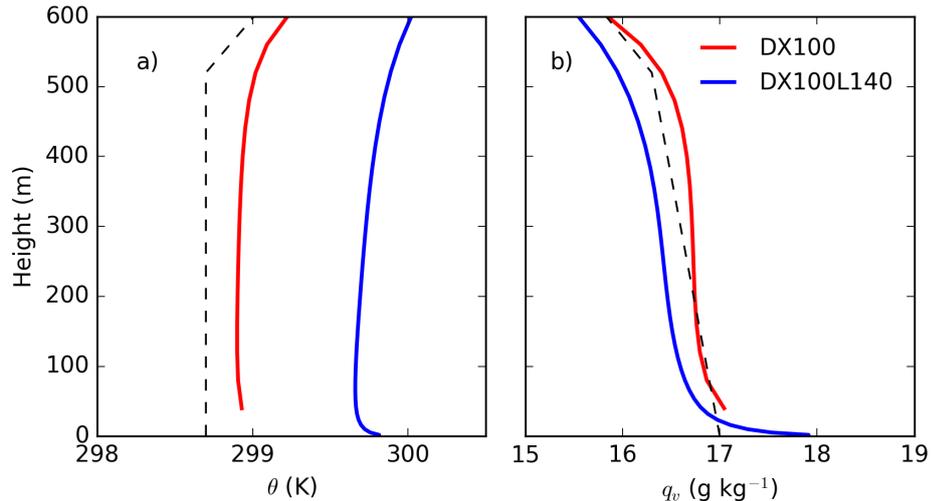


Figure 4.13 As in Fig. 4.6 panels a and b, but only the lowest 600 m. The DX100L140 values are added to the figure in blue.

magnitude to the control experiments and the intercomparison. This is associated with a much smaller change in the thermodynamic profiles away from their initial conditions. The surface layer still develops as in DX100L140, but the remainder of the boundary layer profiles in DX100L140-if match DX100 more closely. Finally, the horizontal-mean, time-mean turbulent fluxes from model hours three through six settle on lower surface moisture fluxes, and similar surface heat fluxes. The pattern in and above the cloud layer exhibits the same behaviour associated with lower cloud liquid in this layer as seen in all our experiments.

An alternative hypothesis is that our initial random noise is driving too-strong initial eddy turnover which alters the thermodynamic profiles enough so that vertical heat and moisture budgets do not balance and a model drift occurs. To test this, we repeat DX100L140 with smaller initial random noise, “DX100L140_sp”. This change slightly delays the cloud initiation and reduces the magnitude of the initial numerical spin-up but has little impact on the evolution of the simulation after six hours, and does not bring DX100L140 closer to the control (not shown). These changes in cloud initiation for DX100L140_sp are consistent with expectations. The weaker initial random noise results in weaker initial eddies which require longer time for their ascending branches to reach the LCL.

4.6 Method for initialising Cloud Trail simulations

Given results from the previous sections tests using BOMEX, CT simulations are initialised using an equilibrium state achieved through a spin-up simulation performed on a small domain for multiple days. The 140-level model configuration appears to be particularly sensitive to the choice of surface fluxes for maintaining large-scale equilibrium. This spin-up method is inspired by [Siebesma and Cuijpers \(1995\)](#) who used LES with the BOMEX test case with interactive surface fluxes and compared surface fluxes which were modelled to surface fluxes estimated from observations. Our model formulation for surface fluxes and the multiple observational and re-analysis sources for our environmental profiles are not going to be consistent with each other and so we cannot expect them to initially be in balance.

Equilibrium conditions are desirable for the CT experiments presented in this thesis. Maintaining an equilibrium allows a clear analysis of any changes due to the presence of an island or the change in numerical grid length. For equilibrium conditions, the wind profiles must be in balance with the surface drag, while temperature and moisture profiles must be in a state of balance with the forcing profiles; as a result, surface fluxes are steady.

[Siebesma and Cuijpers \(1995\)](#) highlight that averaging over many cases smooths and removed vertical structure to the thermodynamic profile. This is because features such as inversions appear at different heights on different days. Therefore initial conditions must be chosen carefully to ensure these features are preserved. Retaining vertical structure is crucial for supporting an equilibrium state for its role in controlling boundary layer height when combined with a subsidence forcing. Without the vertical structure, the boundary layer grows continually deeper, entraining dry higher momentum air from above and triggering inertial oscillations and strong intermittency in convection. Such behaviour is exaggerated for surface sensible heat fluxes that are too high as these fluxes control boundary layer growth to first order.

Furthermore, we expect there to be some spin-up at the start of all simulations. This is because our idealised simulations are initialised with horizontally homogeneous profiles with random noise to break the symmetry. However, we aim to reduce the

length of this spin-up to avoid excessive computational costs. This is especially true for the large domain required for the 100 m simulations in Chapter 5. To do this, a multi-day simulation dedicated to the adjustment of the horizontally averaged profiles towards an equilibrium state is performed. This “spin-up” simulation is performed on a small domain with the same grid length and time step as each corresponding large domain with an island. These spin-up simulations are described in this section, and the profiles are used to initialise simulations in Chapters 5 and 6.

4.6.1 Choosing Initial Conditions

The initial conditions for the reference simulation are chosen based on the thermodynamic conditions on a chosen CT day. A mean of many CT day profiles cannot be used because important meteorologically-relevant vertical structure is not preserved during averaging. The classification algorithm from Chapter 3 is used to identify days which are predominately CT at Bermuda. From these identified days, the 16th July 2015 is chosen because of the steady ambient conditions (Radiosonde in Fig. 4.14). Note that this day has stronger background flow than is typical of CT days.

The 1200 UTC (0900 LT, morning) radiosonde profiles from Bermuda on this day are then simplified into the following meteorologically relevant layers: a well mixed near-surface layer, a shallow inversion to cap the mixed layer, a conditionally unstable cloud layer terminating below the freezing level, and an inversion between the top of the cloud layer and the freezing level (Fig. 4.15). Above the inversion a strong hydrolapse is imposed where relative humidity is decreased to 20 % and remains at this value through the remainder of the column. The low-level relative humidity increases from 77.5 % at the surface to 89.0 % near the top of the mixed layer. A unidirectional initial wind profile consisting of easterly flow at 10 m s^{-1} from the surface to 9 km is imposed. The wind speed then decreases with height, vanishing at and above 15 km. The observed wind profile contains a layer of vertical wind shear near the surface that is ignored for simplicity and wind shear is not expected to be important in controlling whether or not a cloud trail forms. This information is summarised in Table 4.5a.

This same simplified wind profile serves as the geostrophic forcing profile. The

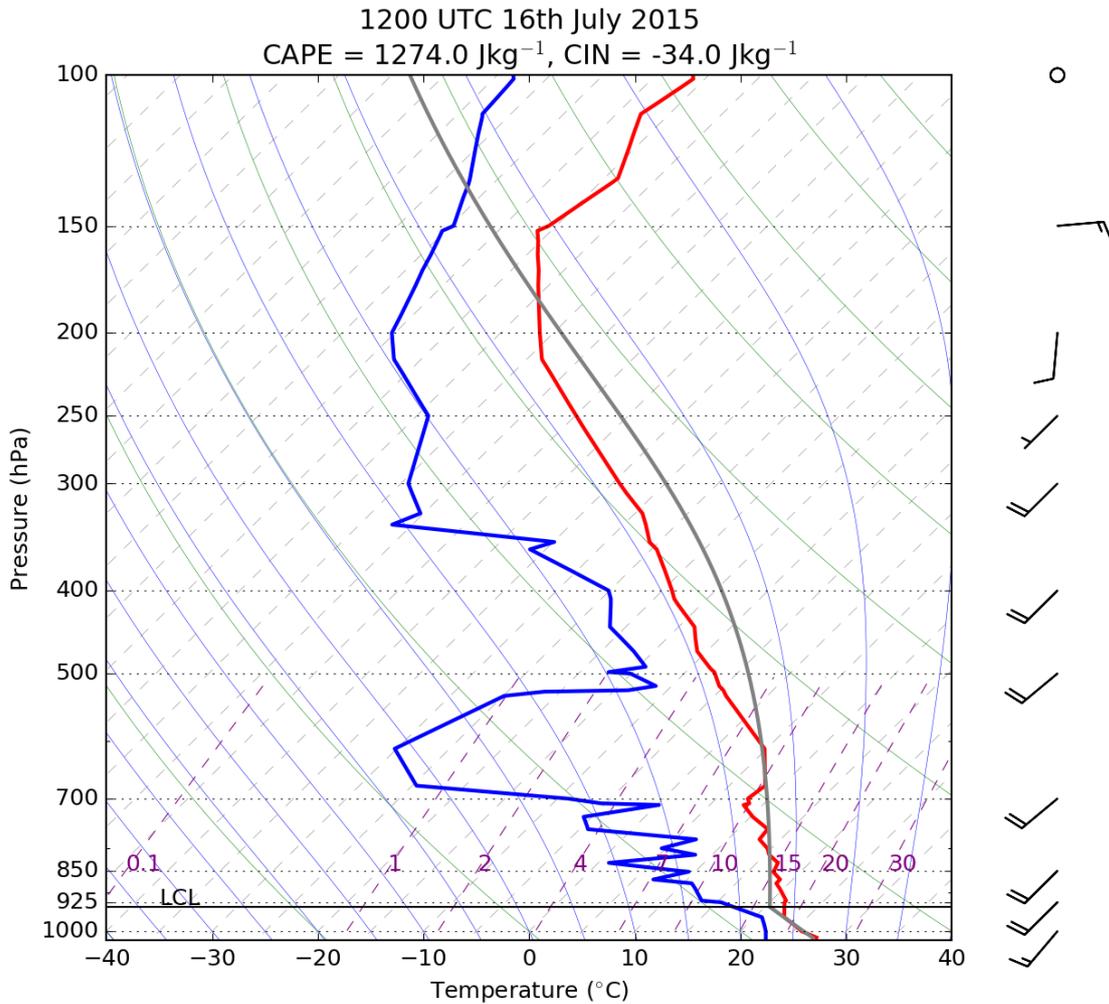


Figure 4.14 Skew-T log-P diagram of the observed temperature in a solid red curve and observed dew point in a solid blue curve. A parcel ascent (based on the mean properties of the lowest 500 m) is shown in a solid grey curve. The pressure of the LCL is indicated as a solid black line. Pressure, in hPa, is marked by dotted grey lines; temperature, in °C, is marked by dashed grey lines; mixing ratio, in g kg^{-1} , is marked by dashed purple lines; dry adiabats, in °C, are marked by thin green lines; and moist adiabats, in °C, are marked by thin blue lines. Winds barbs are used to note the wind speed and direction with height on the right, where a full barb indicates 10 kts, a half barb indicates 5 kts, and a circle indicates calm winds.

horizontal gradients in moisture and temperature are assumed to be negligible and therefore idealised tendencies due to horizontal advection are also assumed to be negligible.

The ERA-interim re-analysis subsidence profile from the same time as the radiosonde and at the grid point nearest Bermuda is used. This profile is simplified so that only the subsident layer between the surface and about 4 km height is retained.

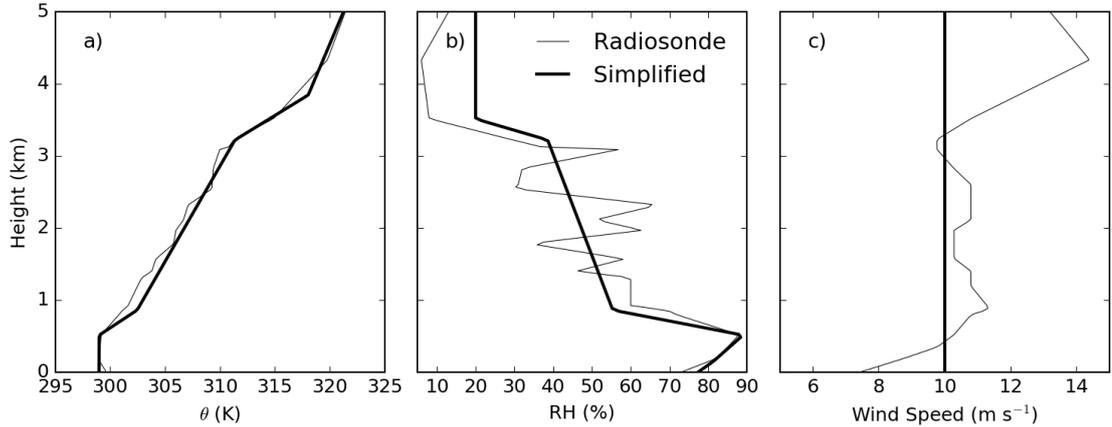


Figure 4.15 A comparison between the radiosonde and simplified initial condition profiles of a) the potential temperature, b) the relative humidity, and c) the wind speed.

This results in a profile with subsidence in the upper boundary layer and throughout the cloud layer. Subsidence then decreases to zero through the inversion, reaching zero near 4 km. Our subsidence profile is then kept at zero above that level. We only intend to force the lower part of the domain where shallow convection is expected, so we choose to retain only the part of the profile acting in this layer. This gives us forcing profiles similar to BOMEX. This subsidence profile drives idealised heating and moistening tendencies as described in section 4.2.3.3.

Finally, a radiative cooling tendency profile is imposed to span the same range of heights as the subsidence. This profile has a 2 K day^{-1} cooling tendency from the surface to the base of the inversion near 3.2 km. The idealised cooling profile goes to zero between 3.2 km and the freezing level (near 4.3 km), then vanishes above the freezing level. This is chosen to be similar to the BOMEX radiative cooling profile for simplicity. We note that this does not attempt to balance the idealised subsidence warming; any imbalance that results is corrected by small changes to the temperature profile during the spin-up simulation described in section 4.6.2. All forcing profiles are summarised in Table 4.5b.

As mentioned above, the sum of the idealised radiative cooling and idealised subsidence warming result in a net cooling in the cloud layer. This makes the cloud layer more unstable initially. The lower part of the temperature profile will be altered during the initial numerical spin-up (as the wind profile adjusts to surface drag and bound-

ary layer eddies emerge). This will change the implied tendency due to subsidence and introduce different regions of idealised net heating and cooling that are closer to balanced at each height. The approach chosen here is to perform dedicated spin-up simulations with the primary goal of achieving balanced wind profiles. The spin-up is allowed to run long enough for both the wind and thermodynamic profiles to adjust such that the large-scale forcing tendencies are more in balance.

4.6.2 Spin-up Simulations

Each spin-up simulation is performed on a 32×32 horizontal grid at the intended resolution of the experiment. The simplified initial condition profiles and idealised forcing profiles are used in the simulations (see Table 4.5). Random noise with mean of 0 and standard deviation of 0.01 K for potential temperature, and 0.025 g kg^{-1} for mixing ratio, is applied to the lowest 1 km to break the initial symmetry of the problem and allow boundary layer overturning motions. Finally, the surface is prescribed as sea-only with a fixed sea surface temperature of 302.3 K based on observations from the Bermuda Weather Service on this day. For the spin-up simulations, interactive surface fluxes and roughness are used. These settings are summarised in Table 4.6.

Each spin-up simulation is then run for ten days. This duration is chosen to be long enough for the wind profiles to adjust to the surface drag, the thermodynamic profiles in the boundary layer to then come into balance with the interactive surface fluxes, and then to allow some time in the new near-equilibrium state to achieve robust statistics. The wind profiles are expected to require at least two days to come into balance based on simple damped inertial oscillation arguments. The thermodynamic profiles are expected to take longer to reach equilibrium because of their dependence on the wind profiles and vertical fluxes of heat and moisture out of the boundary layer.

The Control spin-up simulation provides initial conditions both for a reference simulation, and experiments varying the island surface energy fluxes in Chapter 5. This spin-up simulation is described below. Further spin-up simulations are described in Chapter 5 Section 5.3.3 when different geostrophic forcing profiles are imposed. In Chapter 6, a further variant of these spin-up simulations is described. There, the

Table 4.5 Profiles of a) the initial conditions for the Control experiment, and b) profiles of the large-scale forcing. Linear interpolation can be used to fill in data where there are dashes in the table.

a) Height (m)	θ (K)	RH (%)	u (m s ⁻¹)	v (m s ⁻¹)
0	299.00	77.5	-10	0
148	299.00	81.4	-	-
509	299.01	89.0	-	-
857	302.47	55.4	-	-
3217	311.38	38.6	-	-
3509	-	20.0	-	-
3840	318.05	-	-	-
5931	323.86	-	-	-
7635	328.95	-	-	-
9000	-	-	-10	-
9714	335.98	-	-	-
10963	339.10	-	-	-
12408	342.42	-	-	-
14189	354.57	-	-	-
15000	-	-	0	-
16649	402.01	-	-	-
40000	1200.00	20.0	0	0
b) Height (m)	$\frac{\partial\theta}{\partial t}$ rad (K day ⁻¹)	w (cm s ⁻¹)	u_g (m s ⁻¹)	v_g (m s ⁻¹)
0	-2	0.000	-10	0
148	-	-0.156	-	-
392	-	-0.379	-	-
628	-	-0.588	-	-
857	-	-0.603	-	-
1119	-	-0.498	-	-
1381	-	-0.384	-	-
1644	-	-0.337	-	-
1906	-	-0.315	-	-
2168	-	-0.311	-	-
2430	-	-0.291	-	-
2693	-	-0.300	-	-
3217	-2	-0.418	-	-
3867	-	-0.208	-	-
4326	0	0.000	-	-
9000	-	-	-10	-
15000	-	-	0	-
40000	0	0.000	0	0

Table 4.6 Single level or constant in height settings for the spin-up simulation used to get the balanced reference state. This includes initial random noise, surface pressure, temperature, fluxes, and roughness lengths. Note that T_{sfc} is the temperature of the air on the lowest model level, and T_* is the sea surface temperature.

Forcing	Height	Magnitude	Units
θ'_{rand}	0-1000 m	0.01	K
$m'_{v,rand}$	0-1000 m	0.0025	g kg ⁻¹
p_*	0 m (Surface)	1017	hPa
T_{sfc}	0 m (Surface)	301.0	K
T_*	0 m (Surface)	302.3	K
H_{sfc}	0 m (Surface)	<i>interactive</i>	W m ⁻²
E_{sfc}	0 m (Surface)	<i>interactive</i>	W m ⁻²
z_{0m}	0 m (Surface)	<i>interactive</i>	m
z_{0h}	0 m (Surface)	<i>interactive</i>	m

spin-up simulation is used to initialise simulations with different grid lengths.

4.6.2.1 The Equilibrium State

For the Control CT spin-up simulation, near-equilibrium is reached around day six (marked in Figure 4.16 as a vertical dotted grey line). Recall that one of the original purposes of spin-up simulations is to arrive at balanced wind profiles. This occurs by day 3 in these simulations (Fig. 4.16a). The initial inertial oscillation which coincides with the model numerical spin-up peaks during the first day. The amplitude of this oscillation is subsequently damped, becoming negligible by the end of the second day.

The boundary layer height is a useful indicator of the overall equilibrium because it is controlled by multiple components of the system in which we are interested. These components include the wind profiles, thermodynamic profiles, and surface fluxes. Throughout this thesis, the boundary layer height, z_h , is defined the Richardson Number in the following form:

$$Ri(z) = \frac{g}{\theta_s} \frac{(\theta(z) - \theta_s)}{(u(z) - u_s)^2 + (v(z) - v_s)^2} (z - z_s), \quad (4.30)$$

where Ri is the Richardson Number, g is the acceleration due to gravity, θ_s is the

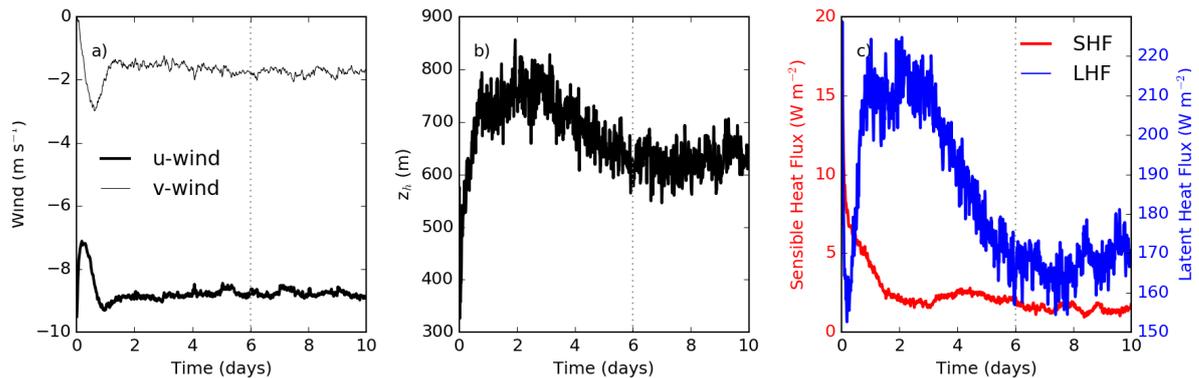


Figure 4.16 Time series of a) The mass-weighted boundary layer mean u- and v-wind components, b) the boundary layer height, z_h , which is diagnosed as the height at which the local Richardson Number, Ri , increases to above a threshold critical Richardson Number, $Ri_c = 0.25$, and c) the surface energy fluxes for the ten-day spin-up simulation. The start of day six is marked on each plot with a vertical dotted grey line, the region to the right of this line (i.e. the last four simulated days) is considered to be in ‘equilibrium’.

potential temperature on the lowest model level, u_s and v_s are the wind components on the lowest model level, and z_s is the height of the lowest model level. All variables are interpolated to be on θ model levels. Within convective boundary layers, Ri as defined above is negative. The height where Ri increases to a critical Richardson Number, $Ri_c = 0.25$, is taken to be the height of the top of the boundary layer and is generally near cloud base.

The boundary layer height grows over the first day (Fig 4.16b). This growth coincides with the strong initial surface fluxes associated with the surface layer and developing wind and thermodynamic profiles (not shown). The strong positive heat and moisture fluxes at the surface over the first few hours are seen in Fig. 4.16c. The sensible heat flux levels off in the $1-3 \text{ W m}^{-2}$ range by the end of day 2 - coincident with the stabilisation of the winds.

Once a well-developed surface layer forms, the boundary layer grows by about 100 m (Fig. 4.16b) and ingests drier air from above. The drying of the boundary layer results in a greater air-surface difference in humidity, Δq , and is responsible for a surge in surface latent heat fluxes over days two through four (Fig. 4.16c). Given near constant wind speed and prescribed constant sea surface temperature, the surface latent heat fluxes are controlled by the humidity of the boundary layer during this

time.

On days 4 and 5 the cloud layer settles into the more humid state of the spin-up (Fig. 4.17b). Now, if the boundary layer height is perturbed upwards, it does not ingest significantly drier air from aloft. As a result, there are no longer significant changes to the air-surface Δq , and therefore the latent heat flux. The latent heat flux is our last indicator to reach equilibrium state.

The surface energy fluxes corresponding to this equilibrium state are a required output from the spin-up. The boundary layer height shows an initial increase in depth during the first day as the capping inversion (near 500 m, solid line Fig. 4.15a) is eroded. This is followed by a slow decrease in the boundary layer height over the next two days before an equilibrium is reached. In terms of the surface energy fluxes, there is an initial peak in fluxes during which time a surface layer is established. The surface fluxes then begin to level off as the flow comes into balance (Fig. 4.16c).

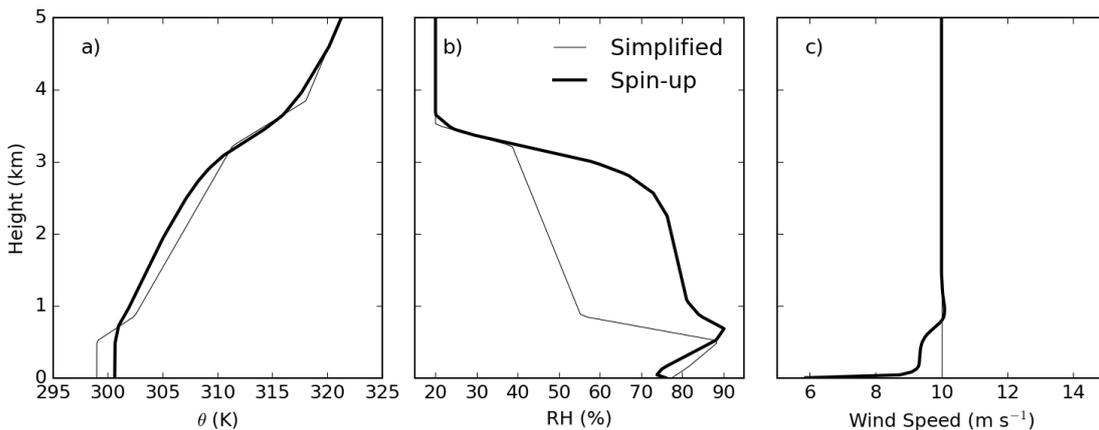


Figure 4.17 Comparison of the simplified profiles used to initialise the spin-up simulations with the equilibrium profiles from the Control spin-up simulation. As in Fig. 4.15

The horizontally-averaged wind and thermodynamic profiles over the last four days are therefore assumed to represent the quasi-steady “equilibrium state”. A multi-day period is used for time-averaging to obtain statistical convergence on these mean values. The horizontal averaging performed for the wind profile is more nuanced in order to smoothly adjust it toward the geostrophic forcing values in the cloud layer. The final wind profiles, \mathbf{u}_f are a linear combination of the wind profiles every ten minutes from the last four days of the spin-up simulation, \mathbf{u}_e , and the geostrophic forcing, \mathbf{u}_g , where

\mathbf{u} indicates the horizontal vector wind. This goes as $\mathbf{u}_f = a\mathbf{u}_g + (1 - a)\mathbf{u}_e$, and a is defined as:

$$a = \begin{cases} 0 & z < z_h \\ (z - z_h)/2z_h & z_h \leq z \leq 2z_h \\ 1 & z > z_h \end{cases} \quad (4.31)$$

so that the final wind profile is completely the result from the spin-up within the boundary layer, below z_h . It then transitions to be completely the geostrophic wind profile at and above $2z_h$. The \mathbf{u}_f profiles are then averaged in time to obtain the initial conditions used in the Control CT simulation.

While this procedure only has a small effect, it does remove small deviations from the geostrophic wind within the cloud layer due to cloud-scale motions which would be inconsistent with the cloudless initial state. This is also important for our final step in which the number of points that describe these equilibrium profiles is decreased, or ‘point-reduced’.

Some notable changes have occurred between the simplified initial conditions and the final Control spin-up equilibrium. Firstly, there have been notable boundary layer and cloud layer changes to the θ profile. The boundary layer has warmed from θ of about 299 to 301 K, the low-level capping inversion (near 500 m) has been eroded and smoothed, and the cloud layer (600 - 3200 m) has cooled slightly (Fig. 4.17a).

Secondly, the cloud layer has become substantially more humid (Fig. 4.17b). The relative humidity in the layer from 600 to 3200 m has increased from near 50 % to near 75 %. Part of that increase in relative humidity is due to the cooling in this layer, but the primary contribution comes from an increase in specific humidity in this layer from net vertical moisture fluxes out of the boundary layer (not shown).

Finally, the wind speed is reduced within the boundary layer with the formation of a characteristic profile near the surface (Fig. 4.17c). This change is expected as the simplified wind profile has no ‘knowledge’ of the surface drag which drives this structure. Our resulting initial condition profiles look more like boundary layer wind profiles from observations, particularly near the surface (Fig. 4.15c, light grey line).

While there have been changes to the meteorologically relevant layers (significant changes in the case of cloud-layer RH), the end environment is still expected to allow cloud trail formation. Looking back to Chapter 3 Fig. 3.10, these changes following the spin-up simulation actually bring the profile closer to the mean of the CT cases in the boundary layer through warming there. The more humid free-troposphere might be more strongly correlated with OB days, but in the absence of a large-scale disturbance in these idealised simulations, the CT should not be obscured or disrupted by organised precipitation.

Finally, the equilibrium profiles are ‘point-reduced’ using an iterative end-point fit algorithm (Ramer, 1972; Douglas and Peucker, 1973). This algorithm (referred to as the ‘RDP algorithm’) describes a curve using a reduced number of points and chooses which points to retain based on some error tolerance. First, the algorithm considers the start and end point of the profile and draws a straight line between them; then, if any of the points on the profile are more than some error tolerance ε_y away from the straight line in the y-direction, the line is redrawn from the start point to the point with a y-distance furthest from the line. This is repeated iteratively until a line is drawn for which none of the original profile points deviate by more than ε_y . Then the start and end of that line are retained, and the process is restarted using the end of the previous line segment and the end of the original profile. The final curve is made by connecting the retained points. By doing this, the number of levels in the profile is reduced from 140 to approximately 40 levels given $\varepsilon_y = [0.01 \text{ K}, 0.1 \%, \text{ and } 0.01 \text{ m s}^{-1}]$. The process of this algorithm is illustrated in Fig. 4.18.

This point-reduction is necessary as the idealised UM can only be initialised with profiles made up of fewer than 100 points. The RDP algorithm described above results in profiles with the same shape but described by fewer points. These profiles can then be used with the idealised UM.

The surface energy fluxes taken from this equilibrium state are then $H = 1.5139 \text{ W m}^{-2}$ and $E = 167.2327 \text{ W m}^{-2}$. These are the horizontally-averaged values over the last four days of the spin-up simulation (the same period over which the profiles are averaged).

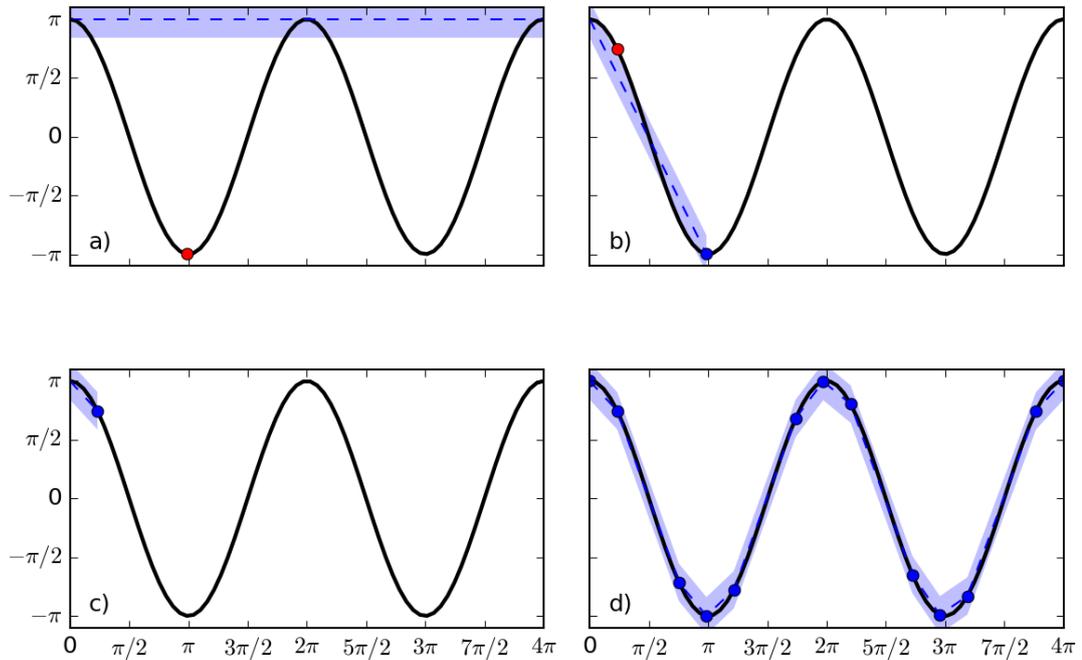


Figure 4.18 An illustration of how the iterative end-point fit algorithm reduces a highly resolved example curve ($\pi\cos(x)$) with 70 points to a curve with 13 points and a maximum error of $\varepsilon_y = 0.5$. In each panel the well-resolved curve is marked in black. a) The first guess curve (blue dashed line), the tolerance envelope (blue shaded region), and point with greatest distance outside that envelope (red dot). If more than one point has the same distance, the first occurrence is considered. b) The red dot from a) becomes the end point of the line for this iteration (blue dot), but there is still a point outside the tolerance (red dot). c) the red dot from b) becomes the end of the line for this iteration (blue dot), but now no points lie outside the tolerance, and the next segment is initiated from the red point in c) to the end of the well-resolved curve. This is repeated until the whole curve is represented, as in d).

4.7 Summary and Conclusions

The idealised UM has been introduced. A set of parameterisation schemes have been chosen to reconcile the benefits of LES and operational considerations, and these account for sub-grid turbulence, microphysics, and surface processes. Finally, idealisations which represent large-scale processes, reduce computational cost, and facilitate scientific experimentation are outlined and rationalised. Many of the imposed idealisations follow methods traditionally used in LES.

The well-studied BOMEX shallow, non-precipitating cumulus case is used to demonstrate that this configuration of the idealised UM can sensibly reproduce shallow

convection statistics as reported from an LES intercomparison. The BOMEX case is typical of conditions in the subtropical Atlantic, including near Bermuda, during the warm season. Given the success of the idealised UM, it is deemed fit for the purpose of expanding upon observations from Chapter 3 with credible simulated CTs consisting of shallow convection in Chapter 5.

The BOMEX case is repeated across a range of grid spacings from 50 m to 800 m to explore the shallow convective ‘grey zone’. Further, these grid spacings approach the 1.5 km grid length currently in use in the UKV. These experiments demonstrate that the representation of shallow convection does indeed break down at larger grid lengths with delayed convective onset and too-deep convection. In fact, results in DX050 suggests the 100 m representation is only just treating shallow cumulus sensibly. In Chapter 6, the horizontal grid spacing is again increased towards these same values for CT cases. These BOMEX results test sensitivity to the grid spacing and build expectations for simulated CT in the grey zone, and therefore identify the behaviour which a convection scheme might correct.

In Chapter 5, the required grid spacing for successful shallow convection (100 m) makes the simulations particularly computationally costly. To reduce some of this cost, the spin-up is handled on a smaller domain from which adjusted thermodynamic profiles, balanced wind profiles, and equilibrium surface fluxes are derived. This method will be repeated in Chapter 6 for coarser grid spacing experiments and produces initial thermodynamic and wind profiles that are subtly different.

Chapter 5:

Idealised Simulations of Cloud Trails

5.1 Introduction

Chapter 3 demonstrates that Cloud Trails (CTs) can be found downwind of non-mountainous islands and uses Bermuda as an example. Unlike horizontal convective rolls which have a preferred heat flux/wind regime (Weckworth et al., 1997), CTs may occur under a wide range of environmental conditions given that the large-scale environment is weakly disturbed, the boundary layer is sufficiently humid, and there is an island to provide a land-sea surface heating contrast and therefore source of circulation generation. Chapter 3 also shows that CTs occur most often with strong solar heating and when the wind direction is oriented along the long-axis of Bermuda so that the land-sea heating contrast is maximised.

In Chapter 4, the BOMEX field campaign (Holland, 1972) and subsequent LES intercomparison (Siebesma et al., 2003) are used to demonstrate that the idealised UM is fit-for-purpose for explicitly simulating shallow convection. We then introduce spin-up simulations to combine simplified observations, re-analysis data, and imposed forcing. The spin-up simulation yields wind profiles in balance with surface drag and the imposed geostrophic forcing; thermodynamic profiles in balance with imposed radiative cooling and subsidence forcings; and quasi-steady turbulent surface fluxes which we intend to impose over the “sea” part of the experiments in the current chapter.

This balanced state from the spin-up at the end of Chapter 4 is used to construct a Control CT simulation and subsequent suites of experiments. Each CT simulation

features a circular idealised representation of an island. The land surface properties of the island are introduced in Section [4.2.3.5](#).

Results from the Control experiment are compared to literature on island-driven circulations. The Control experiment is meant to further explore the cloud trail system in ways that are inaccessible to observations. The simulated CT system is then perturbed to address the questions posed in Chapter [3](#): How does the CT circulation strength vary with the strength of the island surface flux? Is the boundary layer humidity a stronger control on cloud formation in CT than the free-atmospheric humidity? How does the strength of the background wind speed influence the CT system?

The current chapter is therefore split into four sections dealing with: the Control simulation, the relationship between surface fluxes and CT strength, the relative role of boundary layer humidity versus free-atmospheric humidity, and the role of wind speed on CT characteristics.

5.2 The Cloud Trail System

Based on observations from Chapter [3](#), the simulated CT is expected to take the form of a cloud band oriented parallel to the low-level wind direction (e.g. [Nordeen et al., 2001](#); [McFarlane et al., 2005](#); [Johnston et al., 2018](#)). Furthermore, we expect to find a warm and dry plume in the immediate wake of the island, where the strong island surface heat fluxes heat air which is advected downwind into the plume, and locally stronger mixing entrains drier air from above the boundary layer ([Garstang et al., 1975](#); [Matthews et al., 2007](#)).

The heated island and warm plume introduce a horizontal buoyancy gradient. This gradient is a source for horizontal vorticity generation either side of the warm plume and island, and therefore the formation of an in-up-out across-flow circulation. As it is advected downwind, this circulation supports the characteristic organised cloud band associated with CT.

The across-flow CT circulation is carried downwind by the background flow and primarily maintained by its own inertia. The forcing generating the across-flow CT circulation ends with the warm plume. Beyond this point, the circulation can carry

tens or hundreds of kilometres downwind.

5.2.1 Details of the Control Simulation

Chapter 4 demonstrated that the idealised UM is able to simulate plausible shallow convection when used with grid lengths of 100 m or less, and so 100 m grid spacing is used for the Control experiment in the present chapter. For our reference CT experiment, the horizontal domain size is 1184×320 grid points ($118.4 \text{ km} \times 32.0 \text{ km}$) over 140 vertical levels from the surface to 40 km. Simulations using 100 m grid lengths are computationally costly for this domain size. Idealisations from Section 4.2.3 are used to reduce computational cost and facilitate experimentation, while a less costly spin-up simulation on a smaller domain is performed to produce initial conditions. Specifics of the initialisation process and this spin-up simulation are outlined in the next section below.

5.2.1.1 Initial Conditions and Idealised Forcing

The initial conditions for the reference simulation are based on an observed temperature, relative humidity, and wind profile from a 0900 LT morning radiosonde ascent (Chapter 4, Fig. 4.14). These profiles are first simplified to retain meteorologically-relevant layers. A subsidence profile from ERA-interim is used for idealised subsidence tendencies, and a radiative cooling profile based on the BOMEX radiative cooling profile is stretched in the vertical to cover our Control Simulation’s deeper cloud layer. The wind profile is also simplified to be uniform with height below 9 km, and taper to calm above 15 km with an identical geostrophic wind profile for geostrophic forcing. The simplified initial condition profiles and the idealised forcing profiles are detailed in Chapter 4, Table 4.5.

The simplified radiosonde profiles and the idealised forcing profiles are supplied to a 10-day spin-up simulation (described in Chapter 4, Section 4.6.2). This spin-up simulation allows the simplified thermodynamic and wind profiles to achieve equilibrium with the idealised forcing on a smaller (32×32 grid point) domain without an island rather than the larger domain required for CT simulations with an island. The

primary goals of this spin-up simulation are to adjust the boundary layer wind profile to the surface drag, and to compute internally consistent surface fluxes. Therefore an interactive surface flux scheme is used to allow time evolution of the surface sensible heat flux, H , and latent heat flux, E , to achieve an equilibrium state. The equilibrium values are found to be $H = 1.5 \text{ W m}^{-2}$ and $E = 167.2 \text{ W m}^{-2}$. We note that the cloud layer becomes significantly more humid during the spin-up simulation, but this change is not expected to prevent CT formation in the reference simulation. The simplified radiosonde profiles are shown as thin lines, and the balanced profiles are shown as thick lines in Chapter 4, Fig. 4.17.

The balanced conditions which result from the spin-up simulation are used to initialise and force the reference CT simulation on the larger domain with a circular island. This larger simulation will now have a much reduced numerical spin-up. Eddy-scale circulations need time to form within the boundary layer from random perturbations applied to otherwise horizontally-homogeneous initial conditions. Recall that these random perturbations are applied to the θ and q_v fields in the lowest 1 km layer. These perturbations break down into boundary layer eddies over the first hour or so of simulation.

We would like to highlight that the wind speed used in the Control experiment, 10 m s^{-1} , is on the high end of what is observed to support CTs at Bermuda. The day on which the initial conditions is based does have a CT with these winds. We therefore still do expect a CT.

5.2.1.2 Domain and Island Dimensions

As mentioned in Chapter 2, Bermuda occupies a land area of about 54 km^2 , with a maximum elevation near 76 m. This is simplified as a circular 50 km^2 island without orography. The idealised island is centred at $x = 108 \text{ km}$, $y = 16 \text{ km}$ within the domain (see Fig. 5.1). As the imposed geostrophic forcing is easterly, the domain is chosen to be a channel elongated in the general direction of the flow and the island is placed on the upwind (east) side of the domain. The length of this channel is chosen to be 118.4 km. This just accommodates the island and a CT as a typical CT is about 100

km long (Nordeen et al., 2001; Johnston et al., 2018). Bi-periodic lateral boundary conditions are imposed on our domain, so we attempt to choose the smallest domain that allows the whole CT to exist without re-entering the upwind boundary to keep computational costs down.

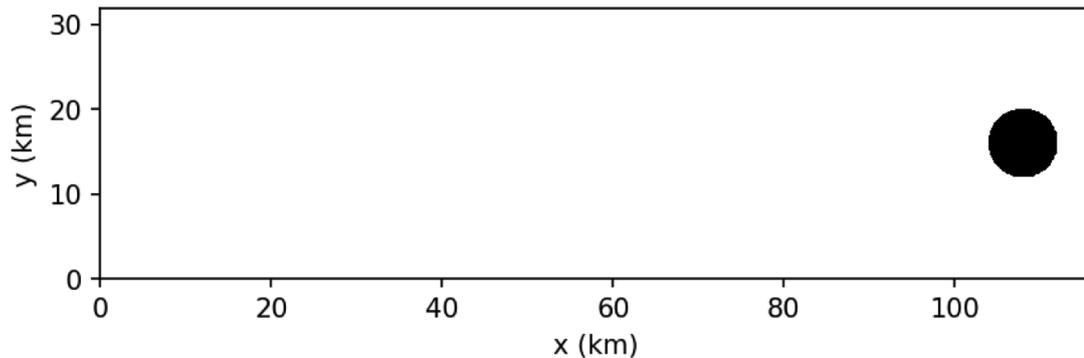


Figure 5.1 Schematic where “sea” grid points are shown in white and “land” grid points are shown in black for the experiments with 100 m horizontal grid spacing. The land points correspond to the 50 km² island which is centred at $x = 108$ km, $y = 16$ km.

In the reference simulation, a diurnal cycle of surface energy fluxes (H_{sfc} and E_{sfc}) is prescribed for the “land” grid points and constant surface energy fluxes are prescribed for the “sea” grid points. The diurnal cycle follows the cosine form given in equations 4.23a and 4.23b from Chapter 4, subsection 4.2.3.5. Both the peak surface sensible heat flux, H_0 , and the peak surface latent heat flux, E_0 , peak at 250 W m⁻² at midday. At the midday peak, this yields a Bowen ratio ($\beta = H/E$) of 1.0 over the island. This Bowen ratio is chosen for simplicity due to lack of surface flux observations at Bermuda, note that $\beta = 1.0$ is on the high side of typical land surface values (Kalanda et al., 1980; Ching et al., 1983). Other values of the Bowen ratio are explored later in Section 5.3.1.

Finally, this Control setup is run for 24 hours. In hours 0 - 6, the initial random noise breaks the initial symmetry and allows clouds to develop (as the island has zero surface energy flux). In hours 6-18, the diurnal cycle of surface fluxes is applied to the island - this is the focus period for studying the CT. Finally, in hours 18-24 the island is no longer heated, and the remains of the CT can be analysed.

5.2.2 Results from the Control Simulation

Previously discussed literature suggests that there are three overarching components to the CT system: the obvious and easily observed changes within the downwind cloud field; a thermodynamic response including a turbulent warm plume; and a dynamical response featuring the CT circulation. These are natural categories into which results from the Control simulation will be discussed in the following sections.

5.2.2.1 The Cloud Field Evolution

An initial domain-wide burst of scattered clouds develops during the initial numerical spin-up in the first simulated hour. At this time, the island is a zero energy flux patch with higher roughness lengths, and over the first six hours of the simulation. This patch leaves a more statically stable, decelerated wake and boundary layer circulations that form here are weaker than those that form in the upwind boundary layer. As a result, the island leaves a cloudless wake whereas over the sea clouds develop unabated during this spin-up period (the remnants of this can be seen in Figure 5.2a). This more statically stable, cloud-free wake behaviour can be seen downwind of relatively cool lakes during horizontal convective roll events over continents in summertime.

The scattered individual cloud elements settle out onto a horizontal separation scale that is larger than that produced by the initial random noise. This process occurs 1 hour to 3 hours into the simulation. We estimate the separation scale as $\approx \sqrt{AN_{cld}}$, where A is the area of the domain in km^2 , and N_{cld} is the number of clouds in the domain. The first wave of clouds have a separation scale of about 1.4 km. Once the initial random noise has been dispersed by boundary layer overturning, the clouds end up with a separation scale near 2.1 km. The scattered cloud field, seemingly undisturbed by the island's roughness patch can be seen in Figure 5.2a, b, and c.

There is no obvious response from the cloud field during the first hour of island heating (Fig. 5.2c). A patch of cloud then forms in the immediate wake of the island between hours 7 and 9 (1 to 3 hours after the island is first heated). This patch is distinct from the ambient cloud field as cloud elements here have higher liquid water content, higher cloud tops, and are more cluster just downwind of the island (Fig. 5.2d,

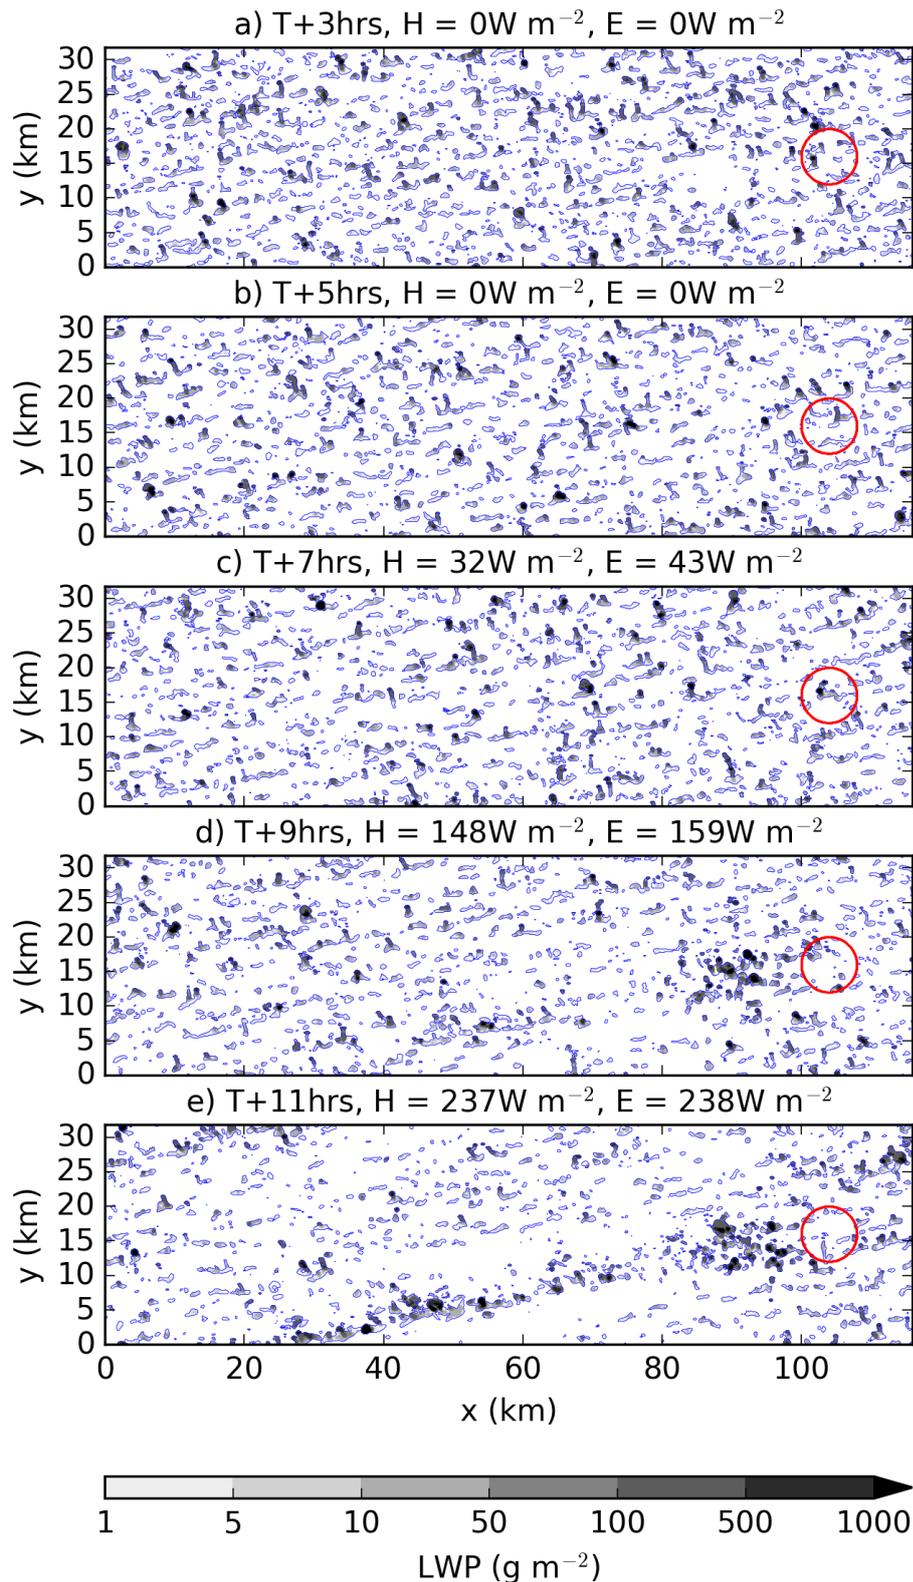


Figure 5.2 Snapshots of the liquid water path in $g m^{-2}$ in greyscale shading. A contour of $LWP = 0 g kg^{-1}$ is outlined in blue, and the island coastline outlined in red. Snapshots are taken every two hours from 3 - 11 inclusive. Surface fluxes, H and E , over the island at the time of the snapshots are noted in each panel's title.

near $x=90$ km, $y=15$ km). This cloud patch is associated with the turbulent warm plume described in [Matthews et al. \(2007\)](#), and discussed in the next sub-subsection.

Finally, the cloud band emerges from the cloud field between hours 9 and 11. By hour 11, the cloud band is well-formed and clearly visible extending west-southwestward from the island in [Figure 5.2e](#). The cloud band continues for another few hours as the island heating peaks and then starts to wane.

Another way to view the CT's influence on the cloud field is to consider the daytime cloud frequency. Here, "daytime" is defined as the 12-hour period from model hour 6 to 18, inclusive, in which the prescribed surface heat fluxes, H and E , over the island are non-zero (i.e. the imposed diurnal cycle has started). The liquid water path, LWP , (i.e. the vertically integrated cloud liquid water content) is used to define a cloud mask on 1-minute time output. Grid cells in which $LWP > 0 \text{ g m}^{-2}$ are masked as cloudy. The daytime cloud frequency is then the time-mean cloud mask between hour 6 and 18.

We now introduce the "CT region", a rectangular area extending downwind of the island in which the majority of the CT system is contained. The CT region is defined by rotating a rectangle so that its long axis is aligned with the mean boundary layer wind direction. The boundary layer height, z_h , is defined by where the Richardson Number increases to a critical value, $Ri_c = 0.25$ (see [Chapter 4](#), [equation 4.30](#) for more information on how Ri is computed). The rectangle is positioned so that it extends 100 km downwind of the leeward-most point of the island and 4 km upwind of the windward-most point, and is 6 km wide. This yields a 112×6 km rectangle running from east-northeast to west-southwest.

From the CT region we define coordinates x' in the along-rectangle and therefore along-wind direction, and y' in the across-rectangle and therefore across-wind direction. Here, x' measured as distance downwind of the leeward side of the island (negative values of x' refer to distance upwind of the leeward side of the island), and y' is measured as distance to the right of the island assuming an observer is looking downwind (negative values of y' refer to distance to the left of the island).

Now we consider the daytime cloud frequency in the rectangular CT region defined above. Cloud frequency is averaged in the y' direction to consider variation along

the CT. Immediately downwind of the island is the turbulent plume, identified by a maximum in cloud frequency. This is followed by a local minimum cloud frequency marked “void” in Fig. 5.3a. The turbulent cloud patch ends about 15 - 20 km downwind of the island, and the persistent cloud band which is characteristic of CTs appears about 40 km downwind of the island with an increase in cloud fraction between 20 and 40 km. This can also be seen in Fig. 5.2e and even more so in panel f.

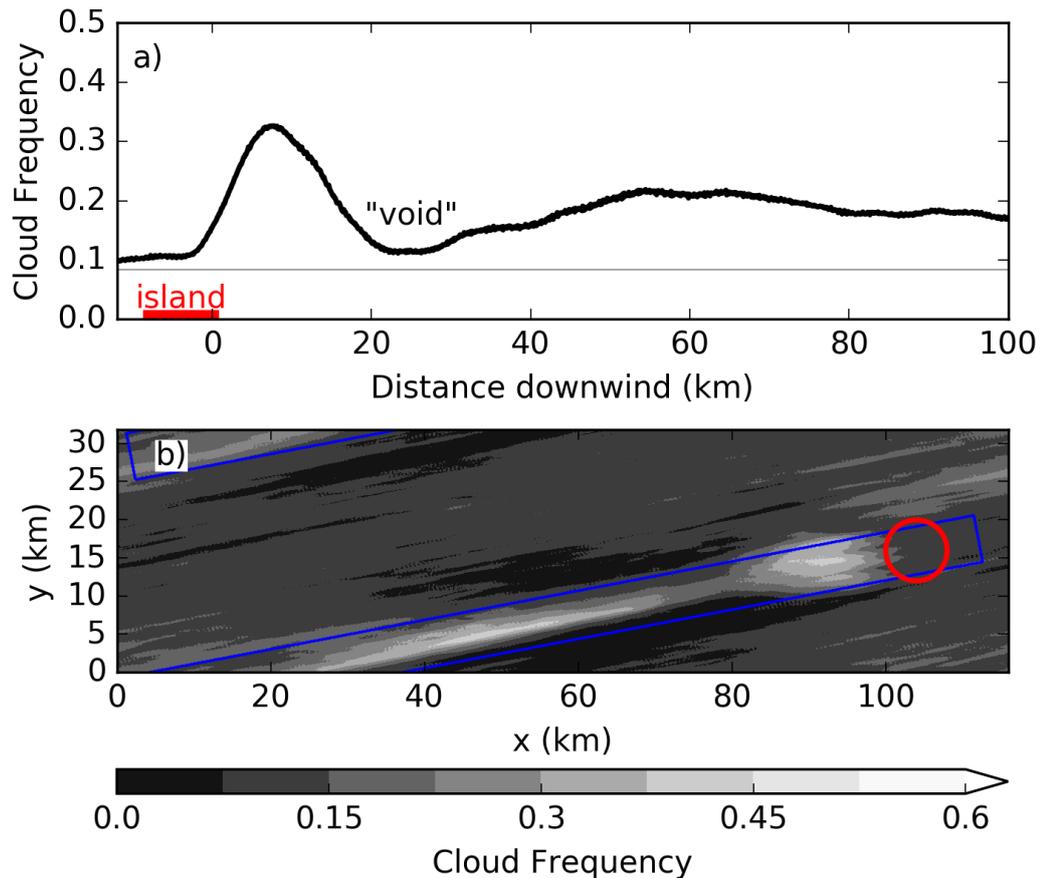


Figure 5.3 The daytime (hour 06 - 18 model time) cloud frequency for the reference simulation. a) as a function of distance away from the leeward edge of the island and b) the spatial distribution of cloud frequency. The cloud frequency in a) is computed by taking the average in the y' direction across the CT region marked as a blue rectangle in b). In a) a grey horizontal line denotes the spatial mean daytime cloud frequency outside the CT region and represents the ambient and suppressed cloud frequency. The island is highlighted in red in both panels.

After 11 hours, the CT begins to re-enter the upwind lateral boundary - a consequence of the bi-periodic lateral boundary conditions. The first half of the “day” is characterised by the cloud band formation and therefore there is minimal impact from

re-circulation across the periodic boundary conditions. The second half of the day has significant re-circulation of the CT system. Comparing the first twelve hours to the last, it appears that there is no significant augmentation of the CT system from the re-circulation of the cloud band. A separate test experiment in which the domain is half as long is done to aggravate the re-circulation issue (not shown). This test suggests that even with increased re-circulation, the description of the CT system is preserved.

Despite no obvious detriment, it is still desirable to avoid any complication from the re-circulation of the cloud band. Therefore the majority of this Chapter's analysis will focus on the first half of the day. It should be noted that there are many real-world situations in which upwind islands produce CT systems which interact with downwind islands without significantly disrupting the satellite appearance of the CT downwind (e.g. recall downwind of the islands of Fiji in Chapter 1, Figure 1.3). Since our focus is on the CT system on one island in isolation, we point out where we think the choice of boundary conditions is having an important impact on the presented results.

Ignoring potential self-interactions, after “sunset”, when island surface fluxes return to zero, the cloud band drifts downwind becoming un-anchored to the island. It retains the banded structure for at least an hour before decaying into the background scattered cumulus field.

5.2.2.2 The thermodynamic response

While there is a suppression of cloud in the wake of the island in the first few simulated hours, there is no potential temperature anomaly co-located with this region near the surface (Fig. 5.4a). The island heating begins after model hour 6. In response, there is an immediate increase in the temperature over and immediately downwind of the island at the lowest model level (e.g. comparing Fig. 5.4b and c). This locally decreases the boundary layer static stability as the daytime surface layer develops over the island. This surface layer, less than 100 m in depth, becomes super-adiabatic as the heating nears the midday peak (at 12 PM, not shown). Non-heated upwind marine air is advected over this island heat source continuously, and a balance between the surface heating and advective cooling limits the peak surface temperatures, with 2 m

potential temperature anomalies peaking at 4.1 K at model hour 12.5.

During the first three hours of island heating, between model hours 6 and 9, a turbulent plume forms (e.g. Fig. 5.4c and d). This is advected and extends downwind from the island to form the turbulent warm plume. The strong turbulent motions in this region: eject convective plumes forming the cloud patch (mentioned in the previous section above); locally dry the boundary layer through entrainment of dry air from aloft and detrainment of moisture from below; and mix higher momentum air from the cloud layer into the boundary layer. The warm turbulent plume continues downwind. By model hour 9, the warm turbulent plume has reached its full length and stays near this length, but grows warmer as the island heating increases on approach to midday.

Along the coast of the island, where the horizontal buoyancy gradients are strongest, is where the CT circulation is primarily formed. The low-level portion of this circulation acts to converge the low-level warm plume into a characteristic triangular shape (Fig. 5.5). This triangular surface warm plume coincides with the triangular-shaped turbulent cloud patch immediately downwind of the island, but the main cloud band extends well beyond the end of the surface warm plume.

Aloft, the warm plume spreads out and downwind in response to the divergent part of the CT circulation and stronger winds away from the surface. This leads the warm plume to overrun the more dense, cooler marine air which also undercuts the warm plume due to the inward component of the low-level across-flow CT circulation. This locally increases the stability of the boundary layer by introducing a “cap” of warmer, drier, island air over cool, humid marine air (Fig. 5.6). This capped region coincides with the “void” between the turbulent cloud patch and the cloud band mentioned earlier.

We visualise this capped region in Figure 5.6 by looking at the difference in the potential temperature between a surface-based parcel and the environment (“parcel perturbation”, $\theta - \theta_{sf,c,p}$) for a parcel that rises dry adiabatically through the lowest 2.5 km in panel a. This is averaged across the same blue rectangle from Figure 5.3b. In Figure 5.6b, the parcel perturbation anomaly is considered by subtracting the horizontal mean over the whole domain at each height from a.

The ascending branch of the across-flow CT circulation eventually erodes the cap

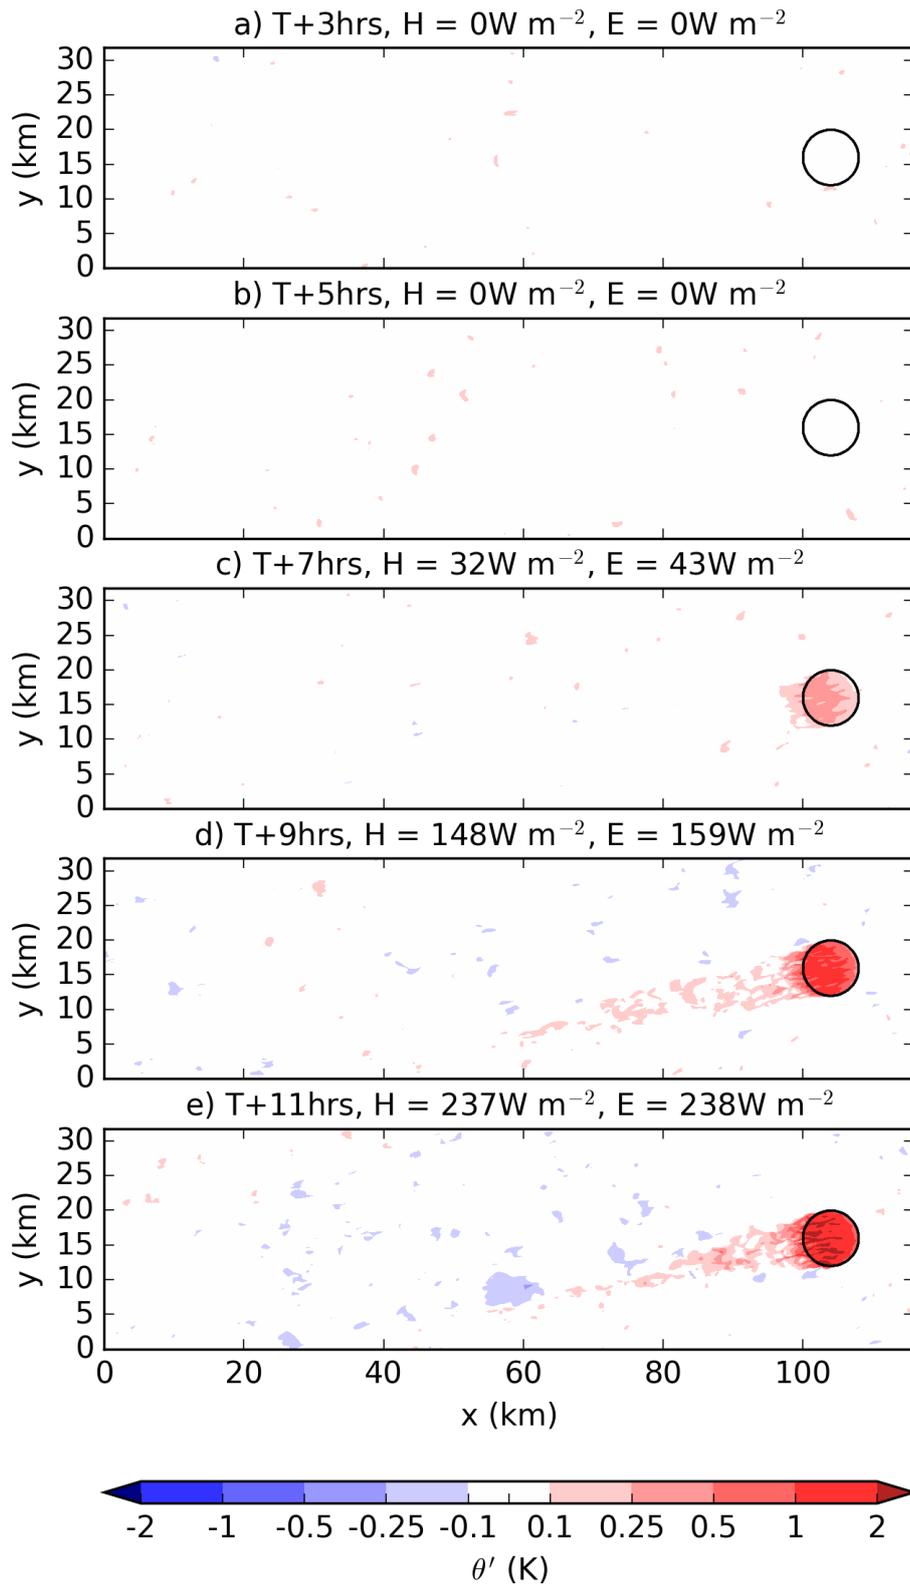


Figure 5.4 Snapshots of the potential temperature anomaly from the horizontal mean at each time, θ' , in K. θ' is taken at the lowest model level, 2 m. Snapshots are taken every two hours from hour 3 to 11 inclusive. The island sensible heat flux, H , and the island latent heat flux, E , at the time of the snapshot are noted in the title.

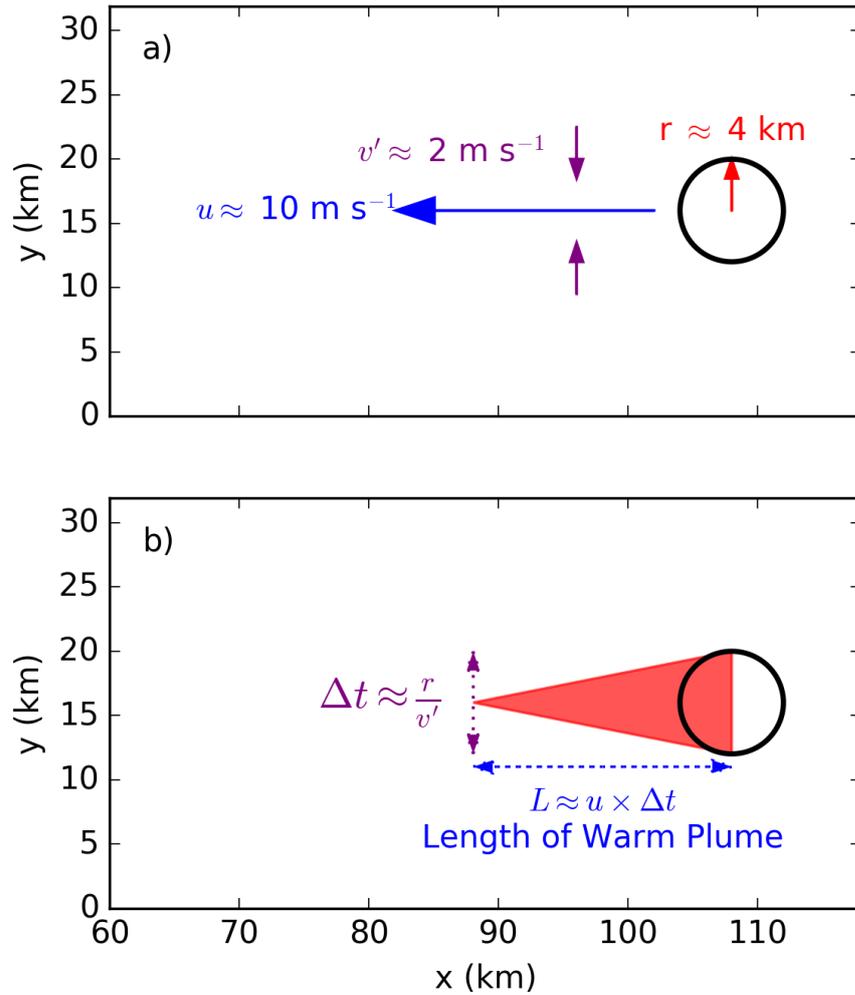


Figure 5.5 Schematic illustrating the relationship between the ambient flow, u , the thermally-forced flow perturbations v' , and the island geometry, given by island radius, r , combine for the triangular shape of the low-level warm plume. a) Example values for u , v' , and r . b) The triangular warm plume, red shaded region is determined by how long it takes v' to advect the edges of the warm plume inward where they meet in the middle after some time Δt . During Δt , the whole system has also moved downwind with u leading to the characteristic triangular warm plume shape in red shading. This process also controls the length of the warm plume.

(Fig. 5.6). Downwind of where the surface warm plume terminates, the part of the warm plume aloft is eroded away in the middle and the warm plume splits into a “V” shape (see evolution of this in Fig. 5.7a through c). In fact, in the middle of the “V” a slight cool anomaly develops near the top of the boundary layer which coincides with the persistent ascent, and the start of cloud formation where persistent ascent and re-evaporation of cloud liquid cools the layer. Some subsidence within the two arms of the “V” warm plume aloft also plays a role in its warming.

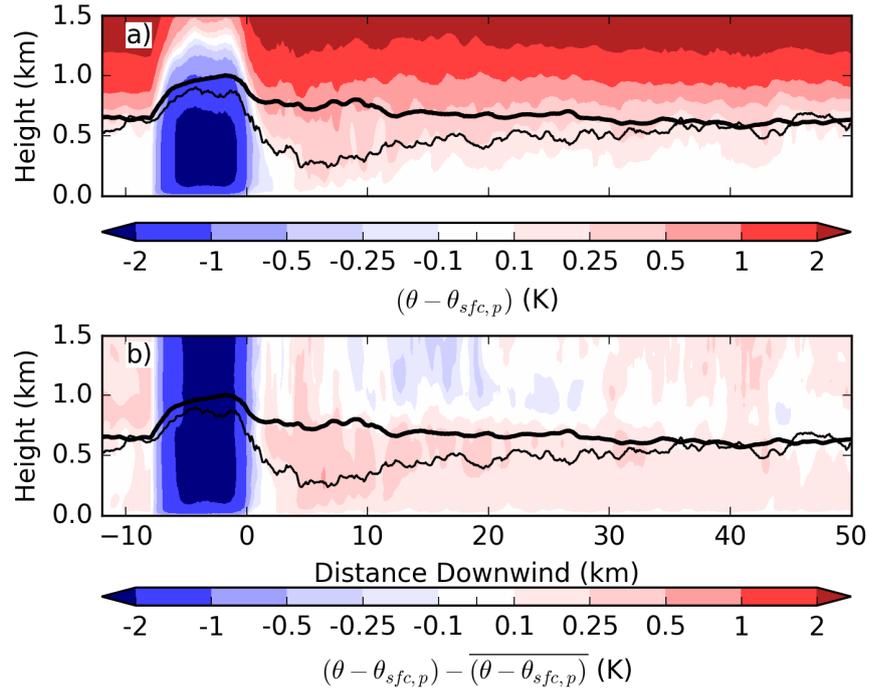


Figure 5.6 Snapshot at hour 12 of a) the difference between θ of a surface parcel and θ of the environment along the CT region, and b) the anomaly of that difference along the CT region from the horizontal mean. In both panels, the bold line is the mean height of the LCL and the thin line is the mean boundary layer height, z_h , both are along the CT region. The capping influence of the warm plume is highlighted as positive anomalies in b) and z_h tracks the bottom of this capping warm layer where the difference between the LCL and z_h is largest.

Either side of the main cloud band, there are lower cloud frequencies (Fig. 5.3b). In addition to the gently subsident branch of the across-flow circulation, the spreading out the warm plume aloft helps to suppress the shallow convection in the either side of the main cloud band. This spread-out warm layer aloft locally increases boundary layer static stability reducing cloud formation. The circulations that arise from the heated island are discussed more in the next subsection.

5.2.2.3 The dynamical response

There is an immediate dynamical response to the island heating. Prior to island heating, the higher surface drag and increased stability over the island reduces the wind speed. This forms a wake of lighter winds that extends about 15 km downwind of the island (Fig. 5.8a).

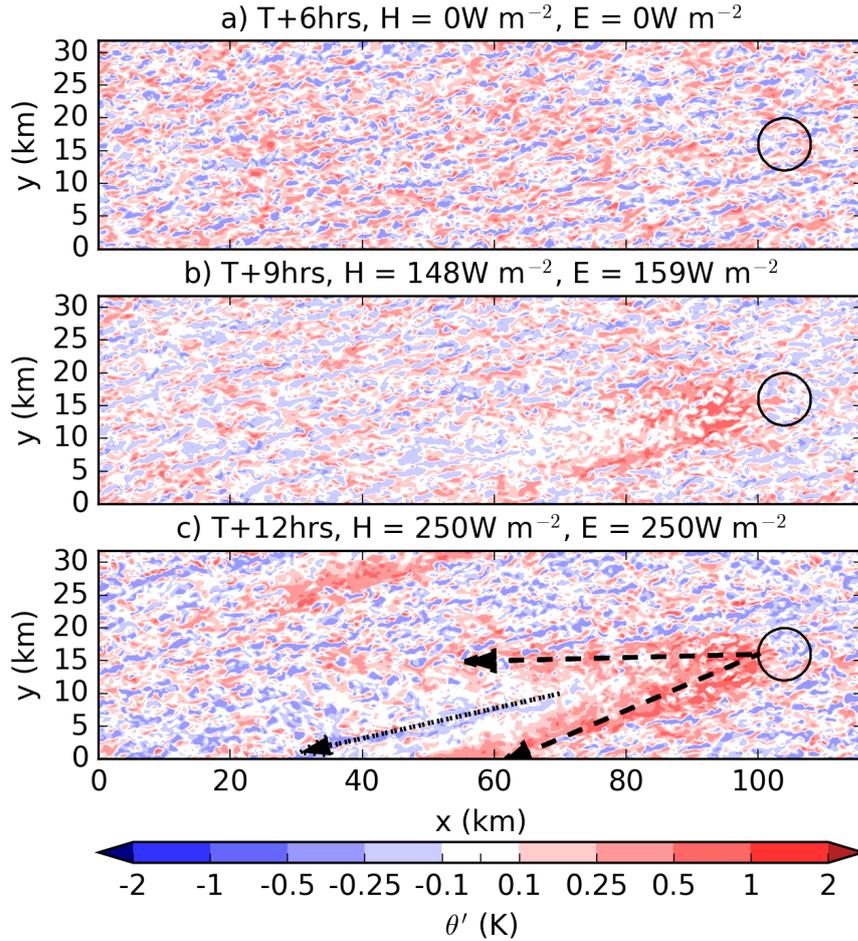


Figure 5.7 Snapshots of the potential temperature anomaly, θ' at the model level nearest to the horizontal mean boundary layer height at each time a) $z_h = 622$ m at model hour 6, b) $z_h = 609$ m at model hour 9, and c) $z_h = 586$ m at model hour 12. The ‘V’-shaped part of the warm plume is marked with two arrows with dashed lines that parallel the arms of the ‘V’, and the cool anomaly aligned with the cloud band is marked by an arrow with a dotted line.

Once the island is heated, this wake of lighter winds is almost immediately erased (Fig. 5.8b). The heated island quickly builds a neutral or convectively unstable boundary layer which, through strong turbulent mixing, entrains higher momentum air from aloft. This mixing, in combination with the along-flow component of the thermally-forced circulation to be discussed next, leaves a wake with *stronger* winds than the undisturbed environment occupying the same 15 km downwind area as the original wake (Fig. 5.8c).

It is now useful to decompose the wind into its components in the along-flow direction, s , and the across-flow direction, n . Using this decomposition, the horizontal

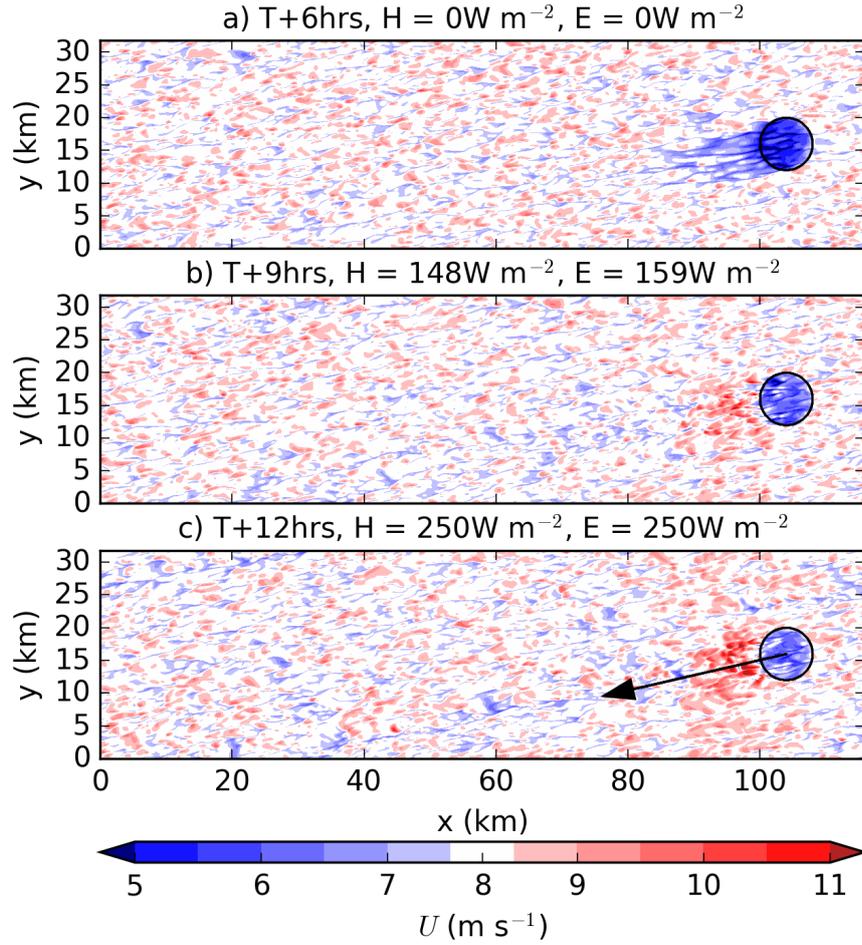


Figure 5.8 Snapshots of the 10 m wind speed, U , for a) model hour 6 (at dawn), b) model hour 9, and c) model hour 12. The horizontal-mean wind direction vector at this level points from east-northeast to west-southwest and is given as a black arrow in c).

mean of s , \bar{s} , is exactly equal to the horizontal mean wind speed, U . The horizontal mean of n , \bar{n} , is exactly zero, and therefore anomalies from the horizontal mean in n are equal to local values, $n = n'$. We also consider anomalies in s from the horizontal mean, s' . s' and n' highlight the thermally-driven circulation in the along-flow and across-flow directions respectively. The properties of this decomposition are summarised as:

$$s = \bar{s} + s', \text{ where } \bar{s} = U \quad (5.1a)$$

$$n = \bar{n} + n', \text{ where } \bar{n} = 0 \quad (5.1b)$$

In the along-flow direction the traditional sea-breeze circulation attempts to form.

This traditional circulation consists of the positive low-level along-flow wind anomalies on the windward side of the island, and negative along-flow wind anomalies on the leeward side of the island. In the case of light background flow, the low-level flow might reverse on the leeward side of the island as a result of this circulation. However, here, there is substantial background flow. For these strong flow cases, the along-flow components of the circulation are advected downwind as shown in 2-D simulations performed by Savijarvi and Matthews (2004). The windward cell becomes phase-locked over or downwind of the island. What would have been the leeward cell is advected even further downwind, and broadens into the flow becoming hard to identify. This along-wind circulation pattern was found to form quickly in Savijarvi and Matthews (2004) and becomes part of a steady “heat island circulation”.

To inspect the circulation anomalies, we return to the CT region rectangle and consider these along- and across-flow anomalies. To highlight the circulation we focus only on the region extending 50 km downwind of the island. We can examine the along-flow “windward-leeward” circulation by averaging s' across the CT region. Furthermore, we can consider the across-flow “inward-upward-outward” part of the circulation by averaging n' along the flow. These are presented in Figure 5.9a and b respectively.

The windward cell of the along-flow circulation downwind of the island appears as positive low-level s' between 0 and 10 km downwind and negative s' beyond 5 km downwind near the height of the boundary layer (Fig. 5.9a). This is a tilted and elongated windward cell of the sea-breeze circulation, augmented by strong momentum fluxes within the turbulent warm plume. The across-flow circulation more closely resembles the traditional sea-breeze over an island but extends tens or hundreds of kilometres downwind of the island (Fig. 5.9b). The across-flow circulation is somewhat asymmetric and this is potentially due to the beginnings of self-interaction due to the bi-periodic lateral boundary conditions.

The across-flow component of the CT circulation is of primary interest because it provides the organised lift to produce the cloud band. This circulation forms rapidly in response to island heating (Fig. 5.10). By hour 9, strong horizontal buoyancy gradients associated with the heated air over the island and in the warm plume have driven

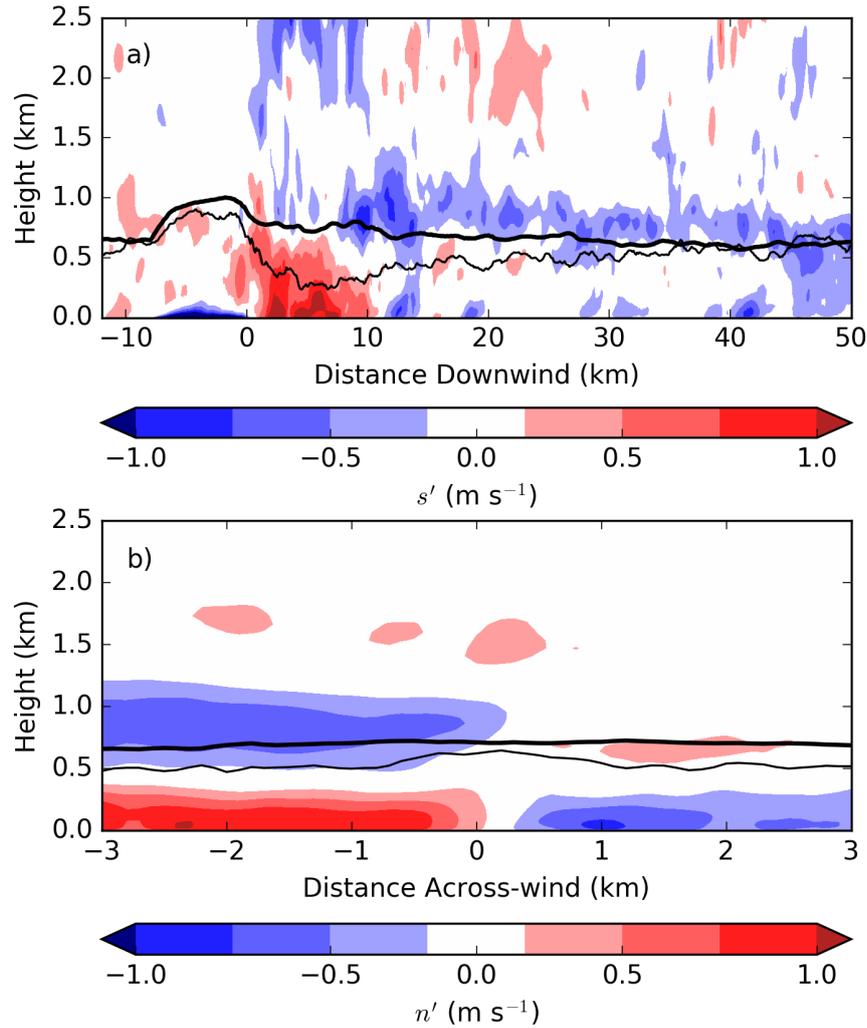


Figure 5.9 Midday (model hour 12) snapshots of a) the along-flow wind anomaly, s' , averaged across the CT region and b) the across-flow wind anomaly, n' , averaged along the CT region. In a), downwind is to the right, and in b), the perspective is centred on the flow with the right side of the plot to the right of the flow, and the left side of the plot to the left of the flow. In both panels, positive anomalies point to the right.

enough accumulated circulation production for the recognisable convergent pattern of low-level n' to appear. This across-flow component of the CT circulation primarily forms over the island where these buoyancy gradients are strongest, but continues to grow along the entire length of the warm plume. The background flow carries the circulation downwind from the warm plume where it eventually decays about 100 km downwind.

Near the surface, the across-flow circulation narrows the warm plume. It draws in cooler, more dense marine air from either side of the warm plume. This cooler air

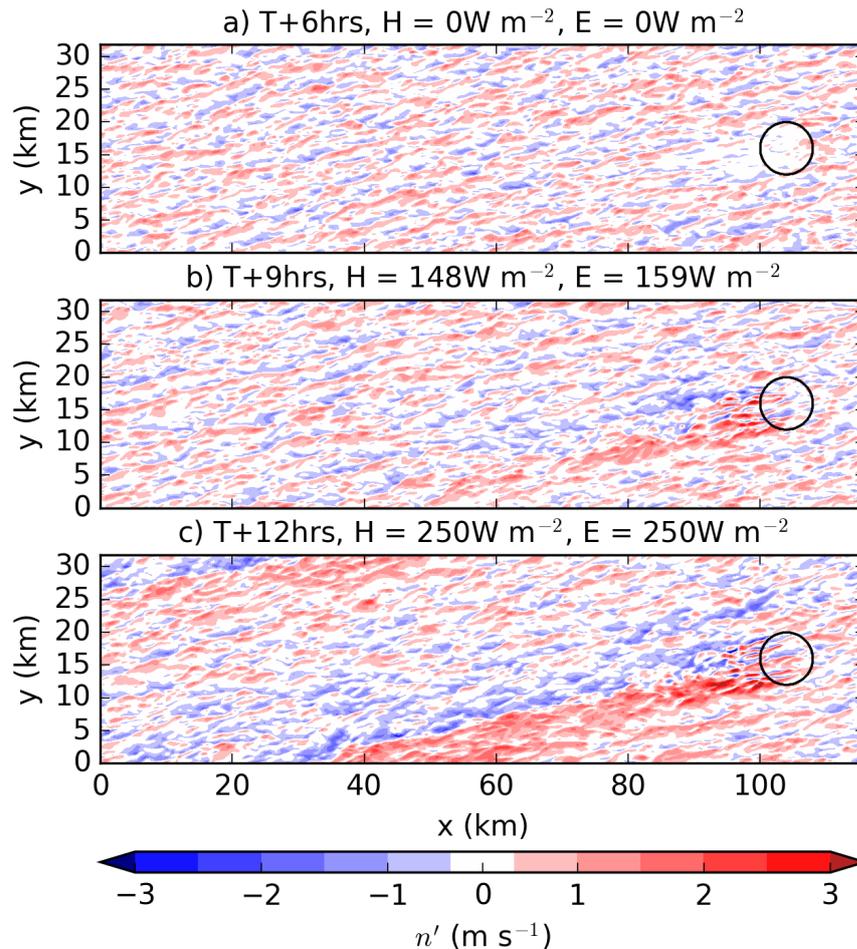


Figure 5.10 As in Fig 5.8 but for the 10 m across-flow component of the wind, n' . Here, positive denotes flow from the south-southeast, and negative from the north-northwest.

undercuts the warm plume. Here, the circulation helps to increase the stability of the boundary layer downwind of the island by imposing a cool anomaly below a warm anomaly. Near the top of the boundary layer, the divergent upper branch of the across-flow circulation forces the warm plume to spread out. The overall effect is to shoal and broaden the warm plume, helping to explain the simulated CT structure. Within the turbulent part of the warm plume, low-level buoyant motions are sufficiently energetic to be expelled from the boundary layer and form clouds. At some point downwind of the island, this is no longer the case. Further downwind, the across-flow circulation is strong enough, and the static stability imposed by the warm plume weak enough for the circulation to lift air for cloud formation. Between these two regions, there exists a transition region wherein the local boundary layer stability to surface-based ascent is too strong for the CT circulation to lift low-level air to saturation, and the warm

plume air is not buoyant or humid enough for cloud formation and a relative “void” of clouds forms.

This cloud “void” region is not void of strong boundary layer updraughts. The simulated vertical velocity here has a boundary layer mean around 0.9 m s^{-1} . Given the horizontal mean boundary layer height of near 610 m, it would take about 665 seconds (roughly 11 minutes) for a parcel to ascend the depth of the boundary layer. Given a mean boundary layer wind speed of about 8.5 m s^{-1} for the Control simulation, a parcel would travel about 5.6 km downwind as it ascends the depth of the boundary layer. This is a much shorter distance than that which is implied by the size of the “void” between the end of the turbulent cloud patch and the start of the cloud band. While it does simply take time for near-surface parcels to ascend to saturation, this is generally more than enough time to do so within the cloud void. This lends further support to there being some hindrance to a parcel’s ability to rise to the LCL within this region.

The across-flow circulation acts as an organising structure, forcing ascent in the cloud band and favouring deeper cloud there, while subsidence in the surroundings suppresses cloud growth. In effect, the island acts to spatially redistribute the clouds. Prior to island heating there is a fairly spatially uniform mixture of cloud heights throughout the domain, but during and after island heating the clouds with higher tops are concentrated in the cloud band, while clouds with lower tops are concentrated in the environment. There is little change in the statistical distribution of cloud top heights (Fig. 5.11). Here, we take the liquid water path field to create a cloud mask where non-zero liquid water path is considered cloudy, then find the maximum cloud top height of each cloud. These statistics form the cumulative distribution in Figure 5.11a and b.

While the role of bi-periodic boundary conditions does not appear to be significant for the interpretation of the presented results, it should be addressed. In these experiments, the weakened tail-end of the CT system re-enters the domain to the right of the island. This introduces an asymmetry in the later evolution of the CT system which starts to become more obvious in the experiments in Section 5.3. For the most part, the analysis of these experiments focuses on the initial development of the CT

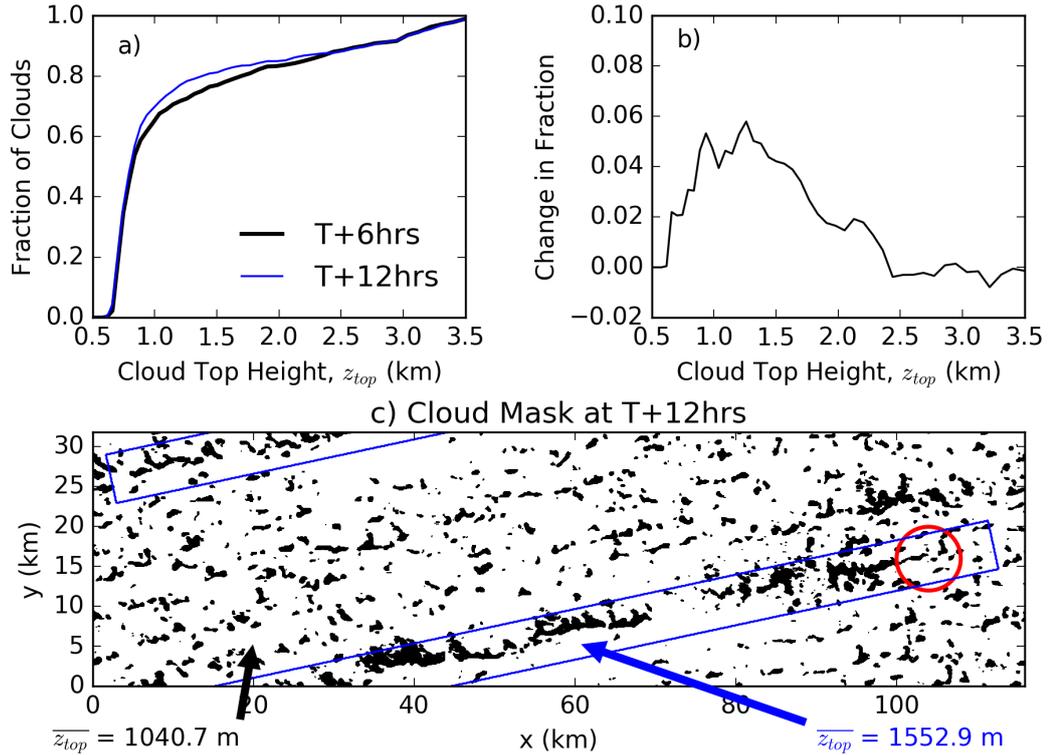


Figure 5.11 A snapshot of a) cumulative distribution function (CDF) of cloud top heights, z_{top} , at model hour 6 (in black) and 12 (in blue), b) the difference between the midday CDF and dawn CDF, and c) the midday cloud mask with a blue rectangle containing the CT region. Annotated are the mean midday cloud top heights, $\overline{z_{top}}$ for the CT region in blue, and elsewhere in black.

system which avoids much of this issue altogether. However, in the daytime cloud fields and snapshots of instantaneous fields in the 9-12 hour range, the wrap-around can be identified.

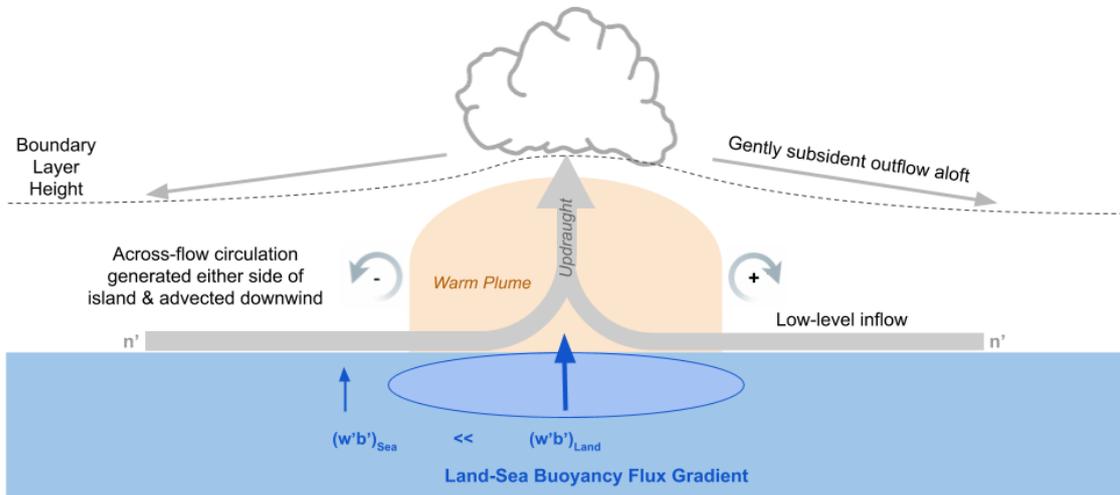
5.2.2.4 Cloud Trail System

The changes to the cloud field, the thermodynamic response to the heated island, and the dynamic response to the heated island make up the “Cloud Trail System”. These three components are summarised in Figure 5.12 and are considered from the across-flow response in a and the along-flow response in b.

In the across-flow perspective, Figure 5.12a, the island (blue circle with dashed outline in the blue sea region) heats the air above leading to the formation of a warm plume, horizontal buoyancy gradients, and subsequent circulation generation. In this

a) Cloud trail system: looking downwind toward an island

i.e. Background flow into the page



b) Cloud trail system: looking across-wind at the island

i.e. Background flow to the left

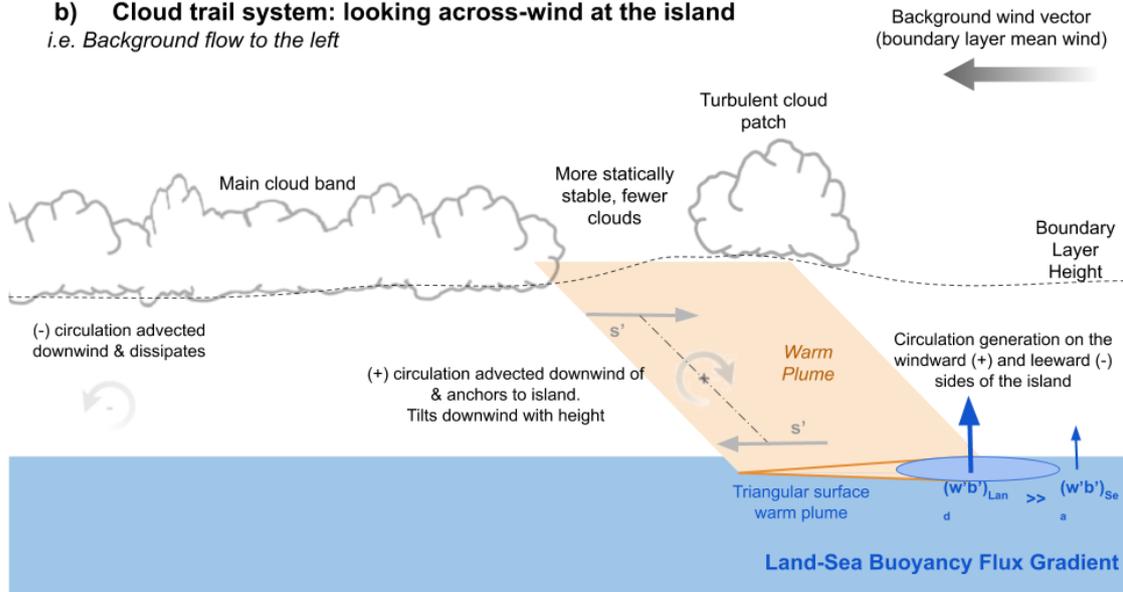


Figure 5.12 A schematic diagram of the main features of the Cloud Trail System as identified in this section. a) The cross-flow system where the land-sea contrast is surface buoyancy fluxes generates a convergent low-level circulation at and downwind of the island similar to the traditional sea-breeze. Cloud forms in the turbulent warm plume that forms over and downwind of the island, and atop the ascending branch of the across-flow circulation. b) The along-flow system where the traditional sea-breeze circulation is advected downwind and transformed into the “heat island circulation”. The main cloud band is supported by the across-flow circulation which is advected downwind from the island, while, at least in the Control simulation, the cloud band and turbulent warm plume cloud feature are separated due to along-flow variations in static stability and dominant mechanisms for ascent.

perspective, the warm plume's buoyant convection and forced circulation associated with the ascending branch of the across-flow circulation support cloud formation. Here, the presented anomalies can be interpreted as extending some distance downwind. The wind anomalies associated with the circulation extend along the full length of the CT downwind, while the warm plume narrows and dissipates some distance downwind determined by the strength of the across-flow circulation and the background wind speed. The boundary layer height is depicted to be higher over the island reflecting our diagnostic using near-surface parameters and local stability in a Richardson Number which, by definition, has to be higher over the warmer surface of the island.

In Figure 5.12b, the along-flow system should be interpreted keeping in mind that the anomalies associated with the across-flow system are superimposed. Similarly, the heated island leads to the formation of a warm plume, and horizontal buoyancy gradients. However, now, there is along-flow circulation generation depicted. It is also advected and extends downwind but the leeward circulation becomes indistinct far downwind while the windward circulation becomes locked downwind of the island and tilted with height. The warm plume narrows to a triangular region near the surface due to the across-flow circulation from a). Aloft, the warm plume overruns some cooler marine air for a more statically-stable region and a relatively cloud-free "void" separating a turbulent cloud patch from the main cloud band.

In the next sections of this chapter, this CT system is explored through idealised simulations.

5.3 Exploring the CT response

In this section, three sets of experiments are outlined, performed and analysed. These experiments are inspired by Chapter 3 of this thesis which found that CTs are more likely to be observed during summertime and in the afternoon (i.e. when the land surface fluxes are strongest, and therefore the land-sea buoyancy gradient is strongest), when the low levels are more humid (so it is more likely for a lifted low-level parcel to reach saturation), and when the wind is strong enough for circulation to be advected downwind rather than remain over the island as in a traditional sea breeze.

Assuming the island surface heating is immediately well-mixed into the boundary layer, the buoyancy perturbation by the island, b' , can be approximated as:

$$b' \approx g \frac{\Delta T}{T_0}, \quad (5.2)$$

where g is the acceleration due to gravity, $T_0 = 300K$ is a reference temperature, and ΔT is the difference in temperature between the island-heated boundary layer and the ambient oceanic boundary layer. ΔT depends on how strongly the island is heated, H , and the residence time of air over the island, Δt :

$$\Delta T \approx \frac{H}{\rho c_{pd} z_h} \Delta t, \quad (5.3)$$

where ρ is the air density and c_{pd} is the specific heat capacity of dry air at constant pressure. If we assume that the residence time over the island is given by the strength of the flow across the island, U , and a characteristic width of the island in the along-flow direction, l_x , such that $\Delta t \approx l_x/U$. Equation 5.2 then becomes:

$$b' \approx \frac{g H l_x}{\rho c_{pd} z_h T_0 U}, \quad (5.4)$$

The CT circulation forms in the across-flow direction, y' , and the horizontal buoyancy gradient in that same direction drives the CT circulation generation that is of interest (as discussed in Section 1.1.1). This buoyancy gradient, $\partial b'/\partial y'$, is due to the land-sea contrast which we can approximate using the across-flow island geometry, $\partial y' \approx l_y$:

$$\frac{\partial \omega_x}{\partial t} \propto \frac{\partial b'}{\partial y'} \approx \frac{g H l_x}{\rho c_{pd} z_h l_y T_0 U}, \quad (5.5)$$

where ω_x is the horizontal component of the vorticity pointing in the along-flow direction, l_y is the characteristic width of the island in the across-flow direction. All other variables are as defined above. It should be noted that this resembles [Kirshbaum and Fairman \(2015\)](#), equation 10 with a factor of $1/z_h^2$.

Given that our experimental island is circular, the ratio of across- to along-flow

characteristic island width is $l_x/l_y = 1$ in equation 5.5. In future instances, these characteristic island widths are taken to $l_x = l_y = r$, where r is the island radius. The terms, g , ρ , c_{pd} , z_h , and T_0 are assumed to be constant across our experiments. The across-flow buoyancy gradient then reduces to $\frac{\partial b'}{\partial y'} \propto C \frac{H}{U}$, where C is some constant containing the characteristics of the background thermodynamic profile and the island geometry.

In subsection 5.3.1 the sensitivity of the CT circulation to the strength of surface fluxes over the island is tested. From the Control experiment in the previous section, we have shown that circulation generation occurs in step with the development of a land-sea buoyancy gradient. The surface sensible heat flux is the primary contributor to local changes in buoyancy. Therefore, a range of surface sensible heat fluxes and total energy fluxes is tested in 5.3.1 to quantitatively link the buoyancy gradient to the circulation strength.

Chapter 3 also suggests that CTs only occur for scenarios in which the low levels are sufficiently humid to result in a low LCL. Conversely, Kirshbaum and Fairman (2015) showed that for the Lesser Antilles Islands, the free tropospheric relative humidity was an important distinction. While we seek clarity on this point, we also seek to confirm the downwind response to the heated island under ‘non-trail’ conditions (i.e. in the absence of a CT, but without being obscured). In subsection 5.3.2, the original initial conditions used in the reference simulation are modified to reduce the low-level relative humidity in one experiment, and reduce the free tropospheric relative humidity in the other.

Chapter 3’s observational analysis proved insufficient to satisfactorily answer some questions about changes to CT behaviour for different background winds because of concurrent significant changes in the background thermodynamic profiles. In subsection 5.3.3, the background wind speed is decreased from the Control experiment’s 10 m s⁻¹ to 5 m s⁻¹ at the lower end of our tested range.

From the theoretical scaling above, and our Control experiment, we can show that the role of wind speed is two-fold. First, wind speed controls the residence time of low-level air over the heated island. Light winds result in a longer residence time which strengthens the low-level buoyancy gradient (this relationship is implied by equation

5.5). A stronger buoyancy gradient results in a stronger CT circulation.

Secondly, increasing the strength of the CT circulation increases the intensity of the convection in the cloud band. We expect that for the lower wind speed experiments, the resulting CT circulation will be stronger. This may result in deeper convection marked by higher cloud tops and more intense precipitation.

5.3.1 Island Surface Fluxes

Chapter 3 showed that CT occurrence at Bermuda is strongly tied to the strength of the solar heating and therefore to the strength of the surface buoyancy fluxes.

The numerical experiments of this section all include a diurnal cycle over the island as described in Chapter 4 Section 4.2.3.5. Our experimental design allows us to prescribe the peak surface energy fluxes of the diurnal cycle, H_0 and E_0 . We can relate the surface sensible and latent heat fluxes to the surface buoyancy flux which is of primary interest. Following the UM formulation (repeated from the UM Document Paper 024, Lock et al. (2017)):

$$\overline{w'\theta'_v} = c_{virt}\bar{\theta}(\overline{w'q'_v}) + (\overline{w'\theta'})[1 + c_{virt}\bar{q}_v], \quad (5.6)$$

where the buoyancy flux is $(g/\bar{\theta}_v)(\overline{w'\theta'_v})$, and all other terms are as previously defined and are taken at the surface. The first term in equation 5.6 contains the contribution from the latent heat flux; the second term contains the contribution from the sensible heat flux. While the UM is written in terms of the total water and the liquid temperature kinematic fluxes, it can be shown that these fluxes either reduce to or can be approximated as the water vapour flux and the potential temperature heat flux at the surface as there is no liquid and the surface pressure is very near the reference pressure, p_0 .

For the reference simulation, $\bar{\theta} = 302.3$ K, $\bar{q}_v = 0.017$ kg kg⁻¹, $\rho = 1.17$ kg m⁻³, $\rho c_{pd}(\overline{w'\theta'}) = 250$ W m⁻², and $\rho L_v(\overline{w'q'_v}) = 250$ W m⁻². The kinematic fluxes become $(\overline{w'\theta'}) = 0.213$ K m s⁻¹, and $(\overline{w'q'_v}) = 8.54 \times 10^{-5}$ kg kg⁻¹ m s⁻¹.

The contribution from latent heat flux in equation 5.6 is 1.57×10^{-2} K m s⁻¹. Meanwhile, the contribution from the sensible heat flux is an order of magnitude higher

at $21.52 \times 10^{-2} \text{ K m s}^{-1}$.

The experiments in this subsection focus on varying the peak island surface sensible heat flux, H_0 , from 50 to 400 W m^{-2} in 50 W m^{-2} increments, without changing the latent heat flux, E_0 . The initial thermodynamic and wind profiles, and sea surface energy fluxes are identical to the Control experiment's for all buoyancy flux experiments.

To confirm that E_{sfc} is a negligible contribution to the overall cloud trail system (not just that its contribution to the surface buoyancy flux is small), two additional experiments are conducted in which the Bowen ratio, $\beta = H_0/E_0$, is varied. To do this, the total energy flux of the Control simulation, 500 W m^{-2} , is re-partitioned for Bowen ratio of 1/3 and 3 (experiments H125E375 and H375E125 respectively). This suite of experiments, including labels used to refer to the experiments later, is summarised in Table 5.1.

Table 5.1 Experiment labels and the corresponding peak island sensible heat flux, H_0 , latent heat flux, E_0 , and Bowen ratio, $\beta = H_0/E_0$. The row in boldface text represents the settings of the Control experiment.

Experiment ID	H_0 (W m^{-2})	E_0 (W m^{-2})	Bowen Ratio
H050E250	50	250	0.2
H100E250	100	250	0.4
H125E375	125	375	1/3
H150E250	150	250	0.6
H200E250	200	250	0.8
H250E250	250	250	1.0
H300E250	300	250	1.2
H350E250	350	250	1.4
H375E125	375	125	3.0
H400E250	400	250	1.6

In fields of daytime cloud frequency, the turbulent plume and cloud band of the CT are evident across this suite of experiments (Fig. 5.13). This appears as higher cloud frequency in the immediate wake of the island associated with the turbulent plume. Then a line of higher cloud frequency parallel to the low-level wind direction extends to the west-southwest of the island forming the cloud band. The signal in the cloud frequency field grows stronger with increasing H_0 (Fig. 5.13a through h).

The relationship between changes in the cloud fields and the strength of the island

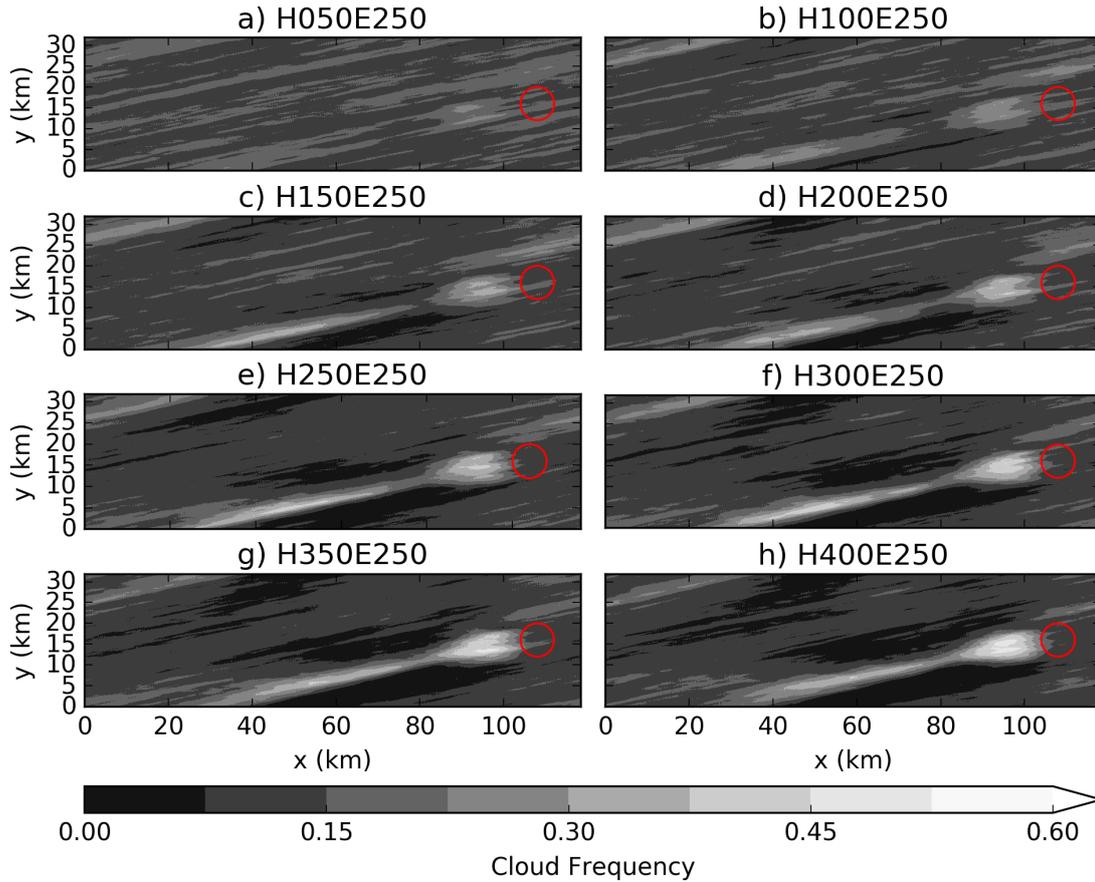


Figure 5.13 Daytime cloud frequency for the eight heat flux experiments in which only the peak island surface sensible heat flux is varied (i.e. excluding H125E375 and H375E125). Here the Control experiment, H250E250, is shown in panel e).

heating is a monotonic increase. The weakest island heating primarily results in a turbulent plume cloud patch only (Fig. 5.13a), while the strongest island heating not only has a pronounced turbulent plume cloud patch and cloud band, but also strongly suppresses the cloud field in the surroundings either side of the cloud band (Fig. 5.13h).

Earlier in this chapter we discussed the shape and extent of the warm plume, and linked it to the strength of the island circulation and the background wind speed (e.g. Fig. 5.5). This suite of heat flux experiments reinforces that thinking. A diffuse and weak warm plume is present for H050E250 (Fig. 5.14a), and this transitions to a pronounced and triangular-shaped surface warm plume for the intermediate experiments H150E250 to H300E250 (Fig. 5.14c to f). By H400E250, with the strongest heat fluxes, the surface warm plume is also notably shorter than the control experiment (H250E250) in addition to taking on this triangular shape (Fig. 5.14h).

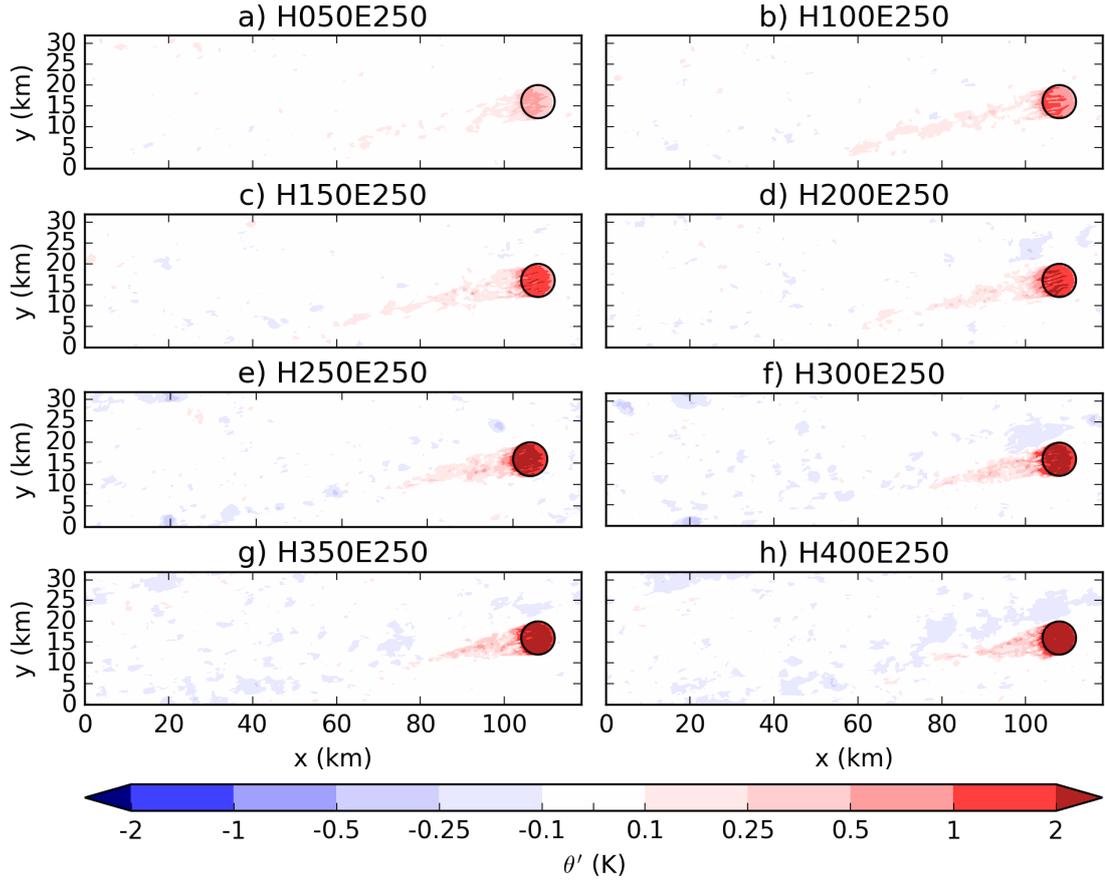


Figure 5.14 Instantaneous potential temperature anomaly from the horizontal mean, θ' , at hour 12 (midday), on the lowest model level (2 m). In each panel the island coastline is marked in black. Here, the Control experiment, H250E250, is shown in panel e).

Furthermore, the experiments with the strongest surface sensible heat fluxes also have a smaller transition region between the turbulent plume’s cloud patch and the cloud band. This transition region is co-located with the break in the CT cloud frequency signature noted in Figure 5.3. For instance this “void” is smallest for the H400E250 experiment, and the void grows wider and the cloud band signature wanes as H decreases.

In order to measure the strength of the CT circulation, we introduce the CT mass flux, M_c . M_c is defined as the updraught mass flux over a modified CT region. The CT region, which we initially defined in section 5.2.2.1, is modified so that it starts one island diameter downwind from the leeward-most point on the island, and the rectangle’s leeward side is found by multiplying the mean boundary layer wind speed by

the averaging period (three hours). This excludes the majority of the turbulent warm plume and only includes the portion of the cloud band that would have experienced heating during the length of the averaging period (three hours, resulting in a CT region which extends roughly 95 km downwind of the island). The CT region is the same width (6 km across) as previously defined. We choose to evaluate M_c at the height of the horizontal-mean LCL averaged over the whole domain, and the time-mean over the period between hours 9 and 12 inclusive when the cloud band is maturing, but before there is any meaningful wrap-around. M_c is defined as:

$$M_c = \int_{y'} \frac{1}{X'} \int_{x'} \langle \rho w \rangle dx' dy', \quad (5.7)$$

where X' is the length of the CT region, $\langle \rho w \rangle$ is the time-mean updraught mass flux, x' and y' are as previously defined. Note that M_c has units of $\text{kg m}^{-1} \text{s}^{-1}$. This measure reduces to an approximation for the mean updraught mass flux at cloud base along the length of the CT multiplied by the width of the cloud band (an example of the modified CT region and the mass flux can be seen in Figure 5.15d).

Given that the initial wind speed is not changed across this suite of experiments, it can be shown from equation 5.5 that the circulation strength should be proportional to H through its control over the horizontal buoyancy gradient. Here, the circulation strength is measured using M_c . We can therefore expect $M_c \propto H$.

Indeed, stronger surface heating results in stronger M_c (Fig. 5.15a and b). The CT circulation strength increases linearly with surface heating over the range explored through this suite of experiments. In Figure 5.15a, the mean M_c over the three hours from 9 AM to 12 PM is shown against the maximum total net surface energy flux imposed over the island, all of the experiments lie close to the line of best fit, except for two; H125E375, and H375E125 - the experiments where the energy flux is re-partitioned by Bowen ratio.

With these Bowen ratio experiments, we see that for the same total net surface energy flux, a range of different M_c are achieved. Earlier we linked the CT circulation strength back to the horizontal buoyancy gradients associated with the land-sea contrast. Those buoyancy fluxes were shown to be primarily driven by the sensible

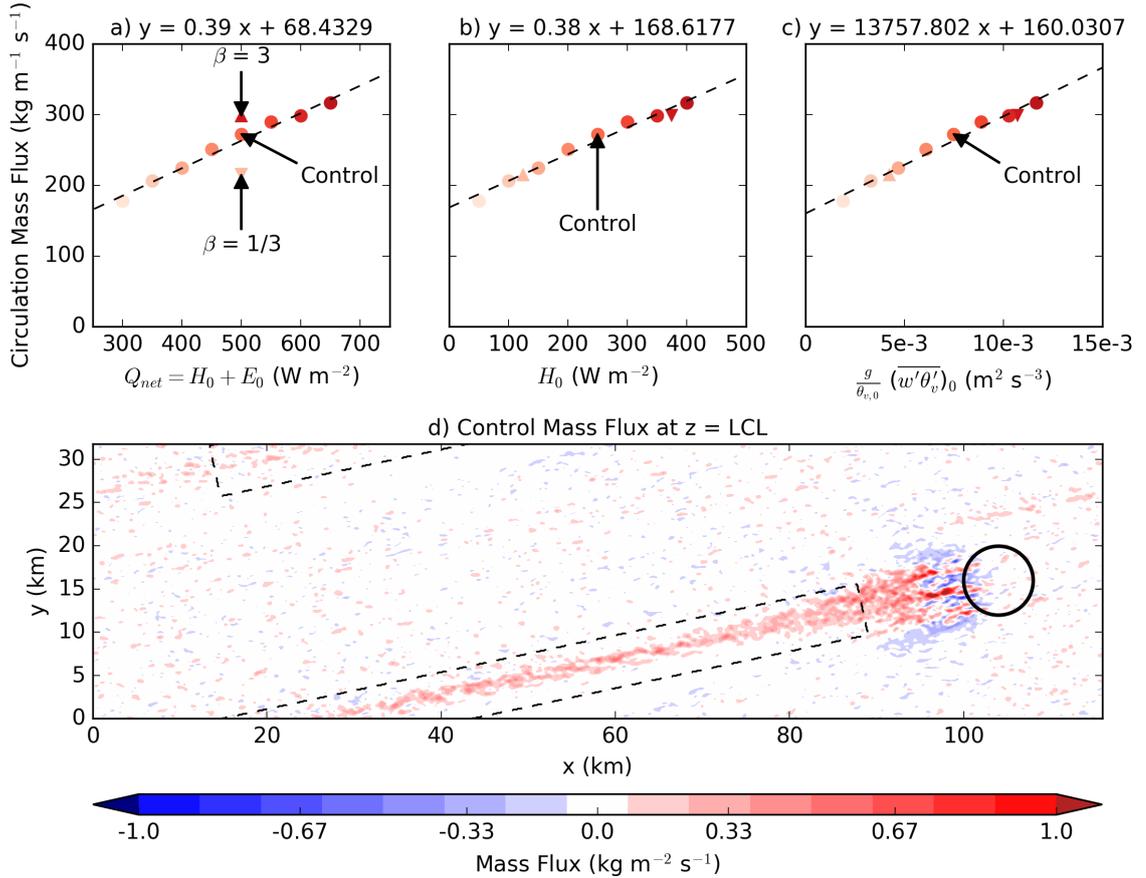


Figure 5.15 The circulation mass flux, M_c , against a) peak total island turbulent surface energy flux, Q_{net} , b) peak island surface sensible heat flux, H_0 , and c) peak island surface buoyancy flux. A line of best fit is drawn as a dashed black line and the equation for the line of best fit is reported in each panel’s title. The symbols are coloured with a darker red for higher H_0 in each experiment, and are circular for experiments where only H_0 is changes, and triangular for experiments where the Bowen ratio, β , is changed but Q_{net} is kept constant. An example of the time-mean mass flux between 9 AM and 12 PM at the height LCL is shown in the shading in d). In d), a bold black circle represents the island coastline, and a dashed rectangle represents the area over which M_c is computed.

heat flux. Re-plotting M_c against the surface sensible heat flux (Fig. 5.15b), those two Bowen ratio experiments fall more in line with the line of best fit. This demonstrates that the dominant control on the M_c does indeed come from the sensible heat flux, and that the latent heat flux plays a much smaller role. For completeness, we shown M_c against the surface buoyancy flux in Figure 5.15c.

We note that our line of best fit does not cross zero in any of our relationships from Figure 5.15a-c. We can show that the intercept represents updraught mass flux

related to the updraughts in the ambient cloud. We find that the ambient cloud field for model hours 3 to 6 has a mean updraught mass flux of $0.48 \text{ kg m}^{-2} \text{ s}^{-1}$ and this translates to M_c of $166 \text{ kg m}^{-1} \text{ s}^{-1}$.

5.3.2 The importance of Relative Humidity

Two experiments are conducted to investigate the role of relative humidity (RH) and cloud in the CT system. Kirshbaum and Fairman (2015) indicated that for robust CTs past the Lesser Antilles Islands, the free atmospheric RH was much lower than for NT events. Conversely, Chapter 3 of this thesis shows that CT days are primarily distinguished from NT days by having higher *low-level* RH.

To explore this further, we conduct an experiment in which the free atmospheric RH is decreased by 25 %. We expect that this drier free atmosphere will lead to a drier turbulent plume, and shallower cloud elements both in the ambient cloud field and in the cloud band. For our purposes, the free atmosphere is defined as being above the boundary layer ($z_h \approx 600 \text{ m}$). Wherever the RH would be less than 20 % after imposing our 25 % decrease, we reset it as:

$$RH(z) = \text{MAX}(20, RH(z) - 25), \quad (5.8)$$

where $RH(z)$ is the spun-up initial RH profile for the Control simulation. This experiment is labelled FAm25.

In the second experiment, the RH within the boundary layer is reduced by 25 % as above. This experiment is labelled BLm25. In BLm25, the drying has the effect of increasing the height of the LCL from 566 m in the Control simulation to 1232 m. All else equal, the CT circulation would need to be stronger in order to lift near-surface air to saturation in BLm25 than in the Control experiment. This concept comes from observations-based results in Chapter 3. In BLm25, we expect CT cloud band formation to be delayed compared to the Control experiment - if a cloud band forms at all. Should this experiment be cloud-free, it might provide an opportunity to investigate the role the clouds themselves play in the CT system. The humidity profiles for these experiments are shown in Figure 5.16.

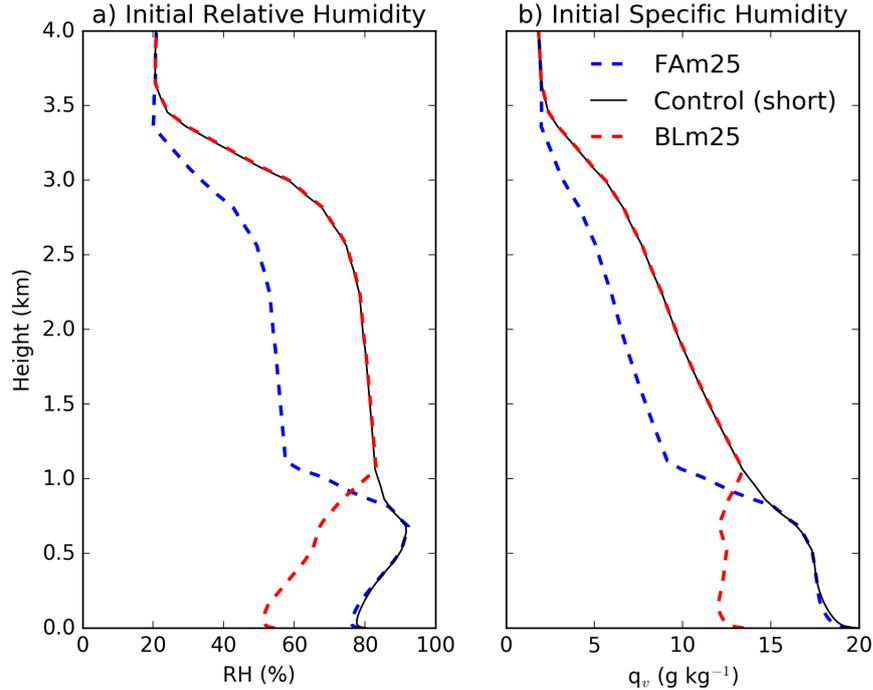


Figure 5.16 Initial profiles of a) RH, b) specific humidity, q_v , for BLM25 (red dashed line), FAM25 (blue dashed line), Control (short) (black solid line). The remainder of the profile above 4 km remains unchanged between the initial conditions.

These reductions in RH are deliberately large to exemplify the role of RH in each targeted layer. Radiosonde observations in Chapter 3, Figure 3.10 show a typical standard error (σ/\sqrt{N}) of about 2 % for a sample size, $N \approx 100$. This yields a standard deviation, σ , of about 20 %. Our reductions in RH are therefore within the tail of the statistical range of observations. However, the model state is not in equilibrium following these reductions in RH. We expect significant model drift away from these initial conditions and toward the more humid reference profiles and this may include changes to the boundary layer height and wind profiles. From a vertically-integrated budget analysis, we estimate that it would take the surface fluxes 0.9 model days for the imposed surface fluxes to completely restore the column integrated water vapour deficits imposed in BLM25, and 1.5 days for FAM25.

We want to explore the CT system’s response to these imposed RH profiles rather than some intermediate profile as it transitions back toward the reference state. To help ensure that we are doing this, we reduce the amount of time in the pre-dawn stage from six to two hours. This gives enough time for the ambient boundary layer

eddies to have several overturning periods and stabilise (not shown) so that the initial numerical spin-up should not be occurring during initial CT formation. An additional sensitivity experiment reducing the pre-dawn model period to two hours in the Control simulation, “Control (short)”, is used for comparison, and we find that the shorter pre-dawn period has a negligible effect on the CT system compared to the Control experiment (not shown).

Furthermore, targeting specific layers introduces unrealistic discontinuities in the relative humidity profiles (Fig. 5.16). These discontinuities are quickly smoothed by the model prior to the start of the imposed island diurnal cycle. Despite this smoothing, the imposed structures are not severely altered (see Fig. 5.17 for a comparison of the RH profiles at the start of island heating). There is an appreciable amount of the original imposed structure remaining. For instance, the cloud layer RH profile remains near the control RH profile for BLM25, but the cloud layer RH remains much below the control profile in FAM25 (Fig. 5.17a). The original structure is more strongly retained by BLM25 as less mixing occurs between the boundary layer and cloud layer. Surface parcels in BLM25 are less able to penetrate into the cloud layer as they do not reach saturation and therefore latent heating cannot contribute to their buoyancy. In contrast, this is not the case in FAM25 and mixing between the boundary layer and what would have been the cloud layer dries the boundary layer and moistens the cloud layer.

In BLM25, the prescribed idealised subsidence tendency is the primary mechanism through which the free atmospheric RH changes. This subsidence imparts a drying tendency on the cloud layer. Without any cloud formation in BLM25 (implied by the meagre cloud fields in Fig. 5.18b), the upward flux of moisture in cloudy plumes which originate in the boundary layer is absent. It is this upward flux of moisture from the boundary layer which acts as a source to replenish what is lost via subsidence and this is how we achieve equilibrium in our Control experiment.

The results of these two experiments are summarised and compared to the reference simulation with the short pre-dawn period in Figure 5.18a-c. FAM25 has very little ambient cloud and a narrow cloud band forms further downwind from the island than in the Control experiment. While still present, the turbulent cloud patch is smaller than

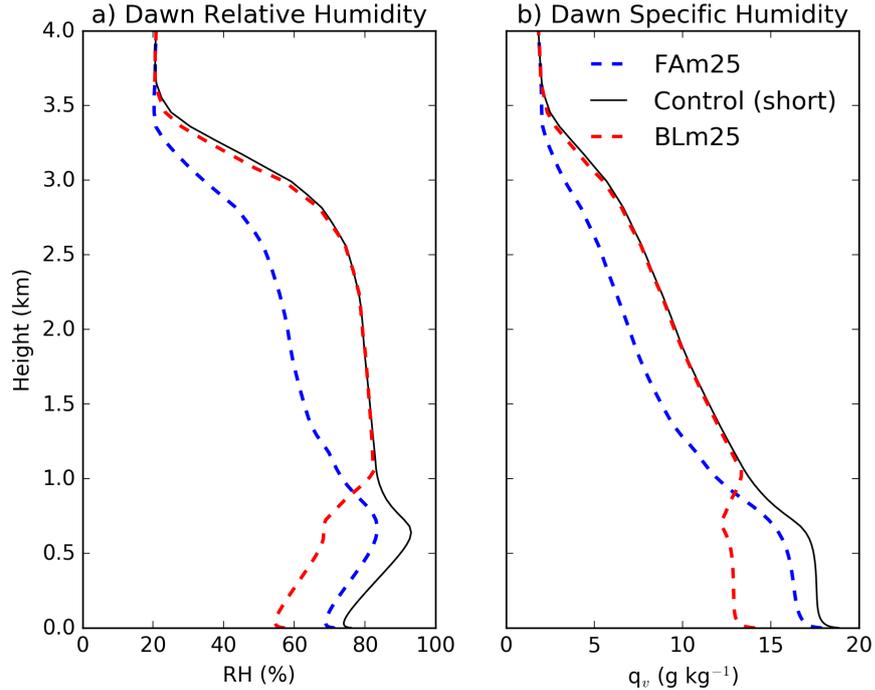


Figure 5.17 As in Fig. 5.16 but for the conditions at dawn, after two model hours in these experiments.

in the Control experiment. Interestingly, the cloud band in FAm25 is oriented more strongly east-west than either the reference or BLm25 experiments. This implies that the mean boundary layer wind direction is more easterly. We show the horizontal-mean, time-mean zonal momentum flux profiles over the first two model hours and the whole domain in Figure 5.19. The implied tendencies from these fluxes (i.e. $-\partial(\overline{w'u'})/\partial z$) indicate less weakening of the boundary layer zonal wind in FAm25 than in BLm25, and that the momentum fluxes do not extend into the free atmosphere as much in BLm25. However, the Control experiment has the smallest implied reduction in zonal winds, yet the orientation of the cloud band is similar to BLm25. Another process must be occurring here to explain the differences in cloud band orientation, but a more complete investigation of these differences is beyond the scope of this thesis.

This relatively cloud-free behaviour in the environment is consistent with our expectations for a dry free atmosphere. Boundary layer eddies are strong enough to lift air to saturation, but they quickly entrain the dry air aloft and re-evaporate. In FAm25, the whole column is somewhat drier than our Control experiment because of the exchange of moisture between the boundary layer and the free atmosphere which

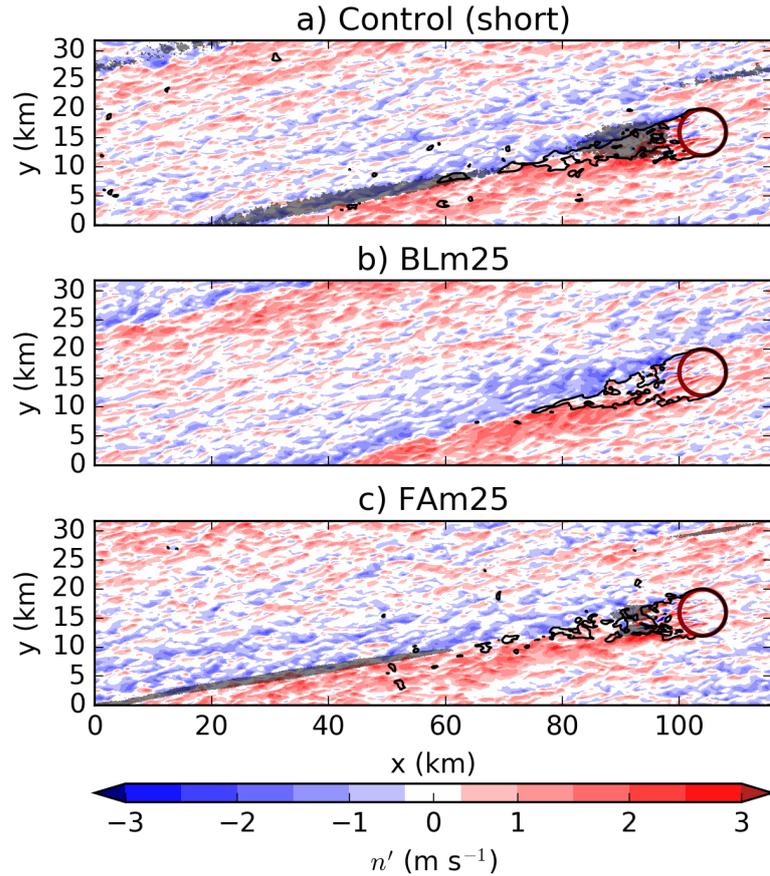


Figure 5.18 The instantaneous across-flow wind component, n' , at 10 m (blue-white-red shading), area covered by daytime cloud frequency greater than 25 % (grey shading), and the instantaneous warm plume extent as defined by the 0.1 K θ' contour at 2 m. The n' and the warm plume extent are taken at model hour 12. Fields are presented for a) Short Control experiment, b) BLm25, and c) FAm25. The short Control experiment is the same as the original Control except it starts at 2 hours before dawn to match BLm25 and FAm25.

has been artificially dried.

Even more dramatically, the BLm25 experiment has very little cloud in either the ambient area or the CT region. Only a small region of transient cloud associated with the turbulent plume forms (not shown). The LCL in this experiment is so high (> 1000 m) that a much stronger surface perturbation is required to trigger moist convection. Even when moist convection is triggered, there is limited support from boundary layer circulations. The CT circulation which would typically support the CT cloud band is not strong enough to support continuous cloud formation in that region and further downwind given the high LCL.

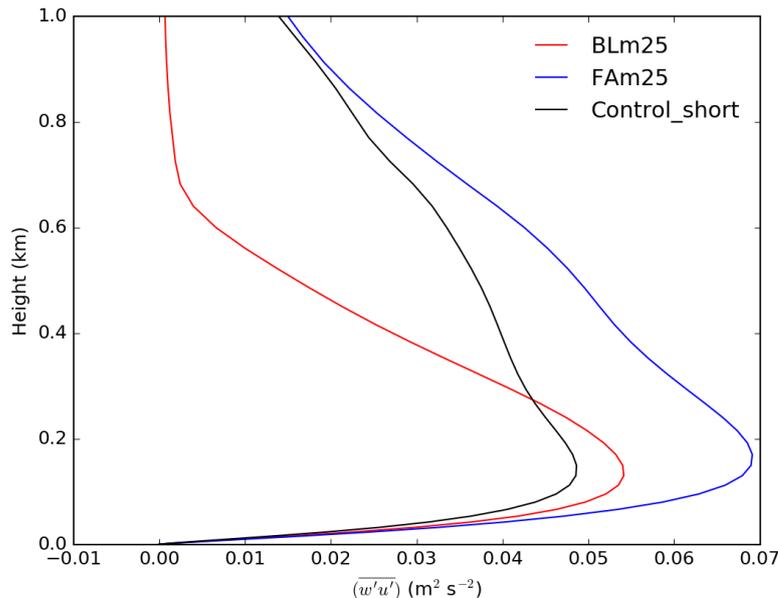


Figure 5.19 The resolved part of the horizontal-mean, time-mean vertical turbulent zonal momentum flux, $\overline{(w'u')}$, over the first two model hours and the whole domain for BLM25 in red, FAM25 in blue, and Control (short) in black.

The extent and intensity of the warm plume is mostly unchanged across experiments. In FAM25, the warm plume extends slightly less far downwind (the black contour in Fig. 5.18). However, the last few kilometres of the warm plume at this height and distance downwind consist of θ' values that are only just above the 0.1 K threshold chosen, and a stricter choice (i.e. threshold greater than 0.1 K) would exclude such fringe regions of the warm plume but would match more closely across experiments (not shown).

The low-level inward part of the CT's in-up-out circulation (shown in Figure 5.18) is again fairly similar across these experiments. However, our measure of CT strength, M_c , indicates that Control (short) and FAM25 have similar strength circulations (432 and 415 $\text{kg m}^{-1} \text{s}^{-1}$ respectively) while BLM25 has the weakest by far with M_c of only 111 $\text{kg m}^{-1} \text{s}^{-1}$. In all cases, the majority of the CT circulation is confined to the boundary layer and the main ascending branch continues some distance above the top of the boundary layer through its inertia. The height of the LCL is within 150 m of (near) the height of the top of the boundary layer in FAM25 and Control (short), whereas the LCL is much higher than the top of the boundary layer in BLM25 (these values are summarised in Table 5.2). The ascending branch of the CT circulation is

therefore not truly captured by this measure in BLm25. Evaluating M_c at z_h instead of the height of the LCL, z_{LCL} , in BLm25 yields M_c of $294 \text{ kg m}^{-1} \text{ s}^{-1}$ which still indicates a weaker circulation - even when M_c is recalculated for z_h in the other RH experiments (not shown).

Table 5.2 The boundary layer height, z_h , the height of the LCL, z_{LCL} , and the CT circulation mass flux, M_c , for the RH experiments. The original Control experiment is also included for reference.

Experiment ID	z_h (m)	z_{LCL} (m)	M_c ($\text{kg m}^{-1} \text{ s}^{-1}$)
Control	598	633	297
Control (short)	601	629	432
BLm25	587	1031	111
FAm25	812	957	415

We note that the shorter pre-dawn period coincides with a stronger M_c (comparing Control with Control (short) experiments in Table 5.2). There could still be some adjustment between the prescribed forcing and initial conditions occurring during the six pre-dawn hours of the Control experiment and that somehow results in a weaker CT circulation than for the two pre-dawn hours in Control (short).

Results from these experiments demonstrate that small heated islands will produce meaningful circulations in their wake - even in the absence of cloud. Satellite images classified as NT in Chapter 3 would likely still have been affected by other components of the CT circulation. These simulations suggest that NT circulations are, however, weaker.

In this subsection we have shown that the characteristic CT cloud band is still achievable for drier free-atmospheres. We have further shown that the boundary layer relative humidity is the key control on cloud formation, given a CT circulation. This is somewhat at odds with Kirshbaum and Fairman (2015), who report that for the Lesser Antilles, the free atmospheric relative humidity is a key distinguishing factor for CT and NT, where it is much drier for CT. This may be a local factor associated with the regional climatology. For instance, it is possible that the largely settled conditions required for significant persistent CT past the Lesser Antilles predominately occur for regimes in which the free atmosphere is much drier than normal, whereas this does not seem to be the case for Bermuda.

Furthermore, the real-world environment is rarely in equilibrium. The implications of the above results indicate that a moistening trend through the day could result in sudden transition from NT to CT as the dry circulation is already present to provide the low-level triggering structure.

5.3.3 Wind Speed

In this subsection, the role of the ambient wind speed on the CT system is inspected. Here, particular attention is paid to the wind speed's impact on the strength of the CT circulation and cloud characteristics. In a final suite of five additional experiments, the wind speed is decreased from the 10 m s^{-1} used in the Control simulation to 5 m s^{-1} in 1 m s^{-1} increments (summarised in Table 5.3).

Table 5.3 Experiment labels and the corresponding geostrophic forcing, sea surface sensible heat flux, H_{sea} , and sea surface latent heat flux, E_{sea} .

Experiment ID	U_g	H_{sea} (W m^{-2})	E_{sea} (W m^{-2})
U05	5	4.3802	135.7834
U06	6	4.5102	143.7222
U07	7	3.7983	151.9601
U08	8	3.1569	157.1509
U09	9	2.2969	161.6953
U10	10	1.5139	167.2327

Whenever the background wind speed is changed, a new spin-up simulation is required. To reduce the wind speed, the simplified wind profile from Table 4.5 in Chapter 4, Section 4.6.2, and the geostrophic forcing profile are re-scaled so the amplitude matches the new wind speed. The new spin-up delivers initial conditions and surface fluxes consistent with the new wind profiles. These new initial conditions can be seen in Figure 5.20 while the new surface fluxes are reported in Table 5.3. The reference simulation is then repeated with the new background winds and geostrophic forcing.

As the wind speed changes there are changes to the background profiles (Fig. 5.20). But in all cases the profiles keep the same basic shape. As wind speed increases, there is an increase in E_{sea} , an increase in boundary layer relative humidity, and consequently a lower LCL. There is also a general trend of decreasing H_{sea} for increasing

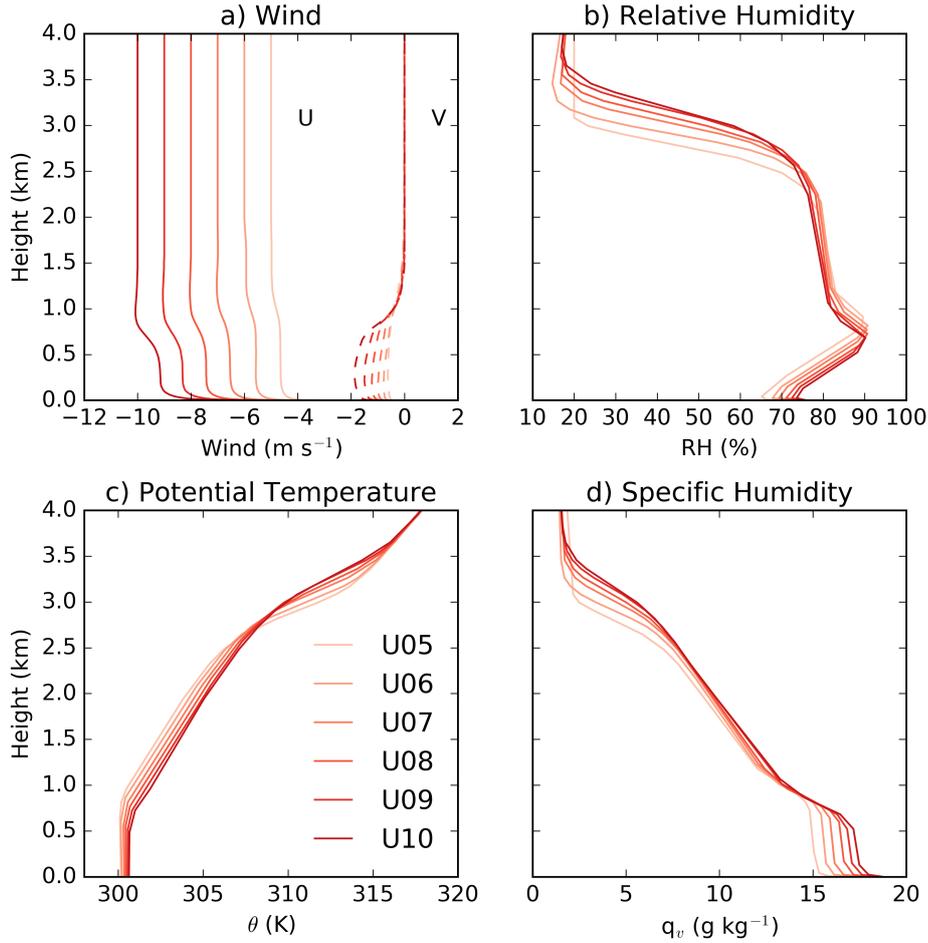


Figure 5.20 New initial conditions following the spin-up for the wind experiments. The lines become darker red for increasing wind speed in each plot. Over the lowest 4 km, a) the u- and v- wind components, b) the relative humidity profiles, c) the potential temperature profiles, and d) the specific humidity profiles are presented.

wind speed. For stronger winds, the temperature profiles more strongly tend toward a warmer boundary layer with near surface temperatures that nearly match the sea surface temperature. A bulk surface flux representation (e.g. the below adaptation from [Stull, 1988](#); [Fairall et al., 2003](#)) can be used to explore this further:

$$(\overline{w'\theta'})_s = C_H U (\theta_s - \theta_m), \quad (5.9a)$$

$$(\overline{w'q'})_s = C_E U (q_{v,s} - q_{v,m}), \quad (5.9b)$$

where $(\overline{w'\theta'})_s$ is the surface sensible heat flux, $(\overline{w'q'})_s$ is the surface moisture flux,

C_H is the bulk transfer coefficient for heat, and C_E is the bulk transfer coefficient for moisture, U is the mixed layer mean wind speed, subscript s denotes surface values, and subscript m denotes mixed layer values. Bulk transfer coefficients are dimensionless and have typical values which range from 1×10^{-3} to 5×10^{-3} .

Taking U10 (the control experiment), this experiment initially has the strongest surface fluxes because it has the strongest wind speed, U . The increased surface buoyancy flux (due to the combined effects of increased heat and moisture fluxes) is an increased source of energy for boundary layer eddies. These eddies are then stronger and grow the boundary layer. During this process, they entrain drier air with higher potential temperature from aloft into the boundary layer - resulting in warming and drying.

The warmed boundary layer reduces $(\theta_s - \theta_m)$ and, from equation 5.9a, this means the surface sensible heat flux decreases with time as the boundary layer grows. This decrease in heat flux allows a new balance to be reached between the surface forcing (growing the boundary layer) and the imposed idealised subsidence (opposing that growth). Similarly, the boundary layer has dried during its initial growth toward a new equilibrium. This results in a larger $(q_{v,s} - q_{v,m})$ and, from equation 5.9b, this means larger surface moisture fluxes. The boundary layer, now no longer entraining as much dry air, undergoes net-moistening which decreases surface moisture fluxes until an equilibrium is reached between the now weaker drying due to entrainment at the top of the boundary layer, and the surface moisture fluxes. These changes are consistent with results varying wind speed in the RICO case (Nuijens et al., 2009; Nuijens and Stevens, 2012).

Furthermore, as wind speed increases the cloud top height decreases, and cloud cover increases (not shown). The combination of lower LCL and lower cloud top height results in a cloud depth and LWP which are roughly the same across the wind speed regimes. Because the spin-up simulations are performed on small domains, only a handful of clouds are supported at any given time (not shown). There are therefore periods in which there are no (or few) substantial clouds in the domain. Such periods are shorter-lived for increasing wind speed, and therefore domain-mean, time-mean cloud cover also increases with the wind speed.

As wind speed decreases, the simulated cloud trail transitions from a turbulent cloud patch with a cloud band downwind to a wider cloud band connected more closely to the turbulent cloud patch, then finally to a wider cloud band, indistinguishable from the turbulent cloud patch, with a ‘V’-shaped cloud structure extending from the end of the cloud band. By U05, the cloud band includes a stratiform layer formed by the spreading of cloud tops at the inversion near 3.5 km height which makes the cloud band appear wider. These changes are all shown in Figure 5.21.

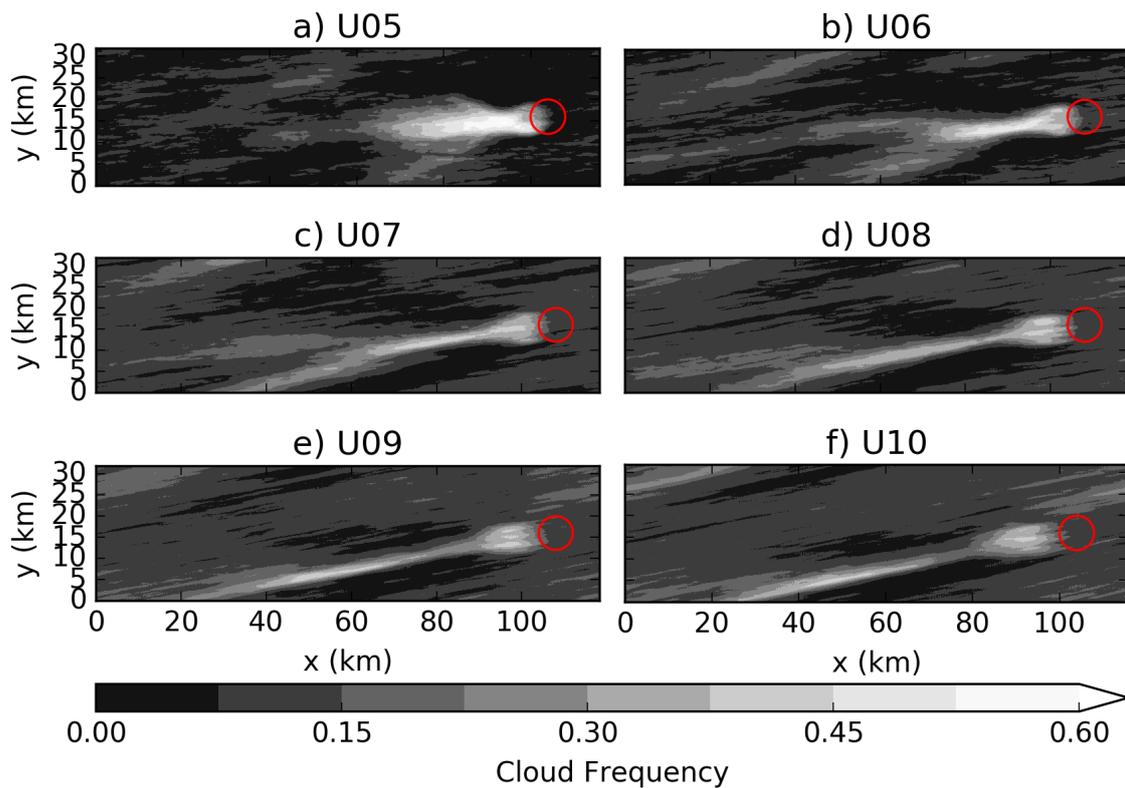


Figure 5.21 As in figure 5.13 but for the suite of wind speed experiments. Here, U10 is the Control experiment.

From equation 5.5, we can see that the horizontal buoyancy gradient is stronger for lower wind speeds. Therefore we expect the CT circulation strength, the increase of which is controlled by $\partial b' / \partial y$, to increase as we decrease the wind speed. Decreasing the wind speed increases the residence time of low-level parcels over the island heat source. All else being equal, a stronger CT circulation results in stronger low-level convergence.

A stronger CT circulation has two consequences: firstly, it results in stronger

convective triggering in the cloud band; secondly, it results in a cloud band which forms closer to the island. The ascending branch of the CT circulation is stronger, so it does not take as long to erode through the warm plume cap (despite this also being stronger), and can lift parcels to saturation much closer to the island.

As wind speed decreases, the CT system forms more closely to the island, particularly in U05 and U06. This has real-world consequences for inhabitants on small islands - particularly should precipitation, or hazardous weather, occur as a result of the CT system.

The downwind extent of the triangular surface warm plume decreases for decreasing wind speed (Fig. 5.22). This change can be explained through compounding effects of a stronger CT circulation. All else being equal, a stronger circulation means that the surface warm plume is pinched off into a triangle faster. But also, lighter winds compound that by carrying the plume less far downwind in the first place. Both of these effects act in unison to decrease the downwind extent of the surface warm plume.

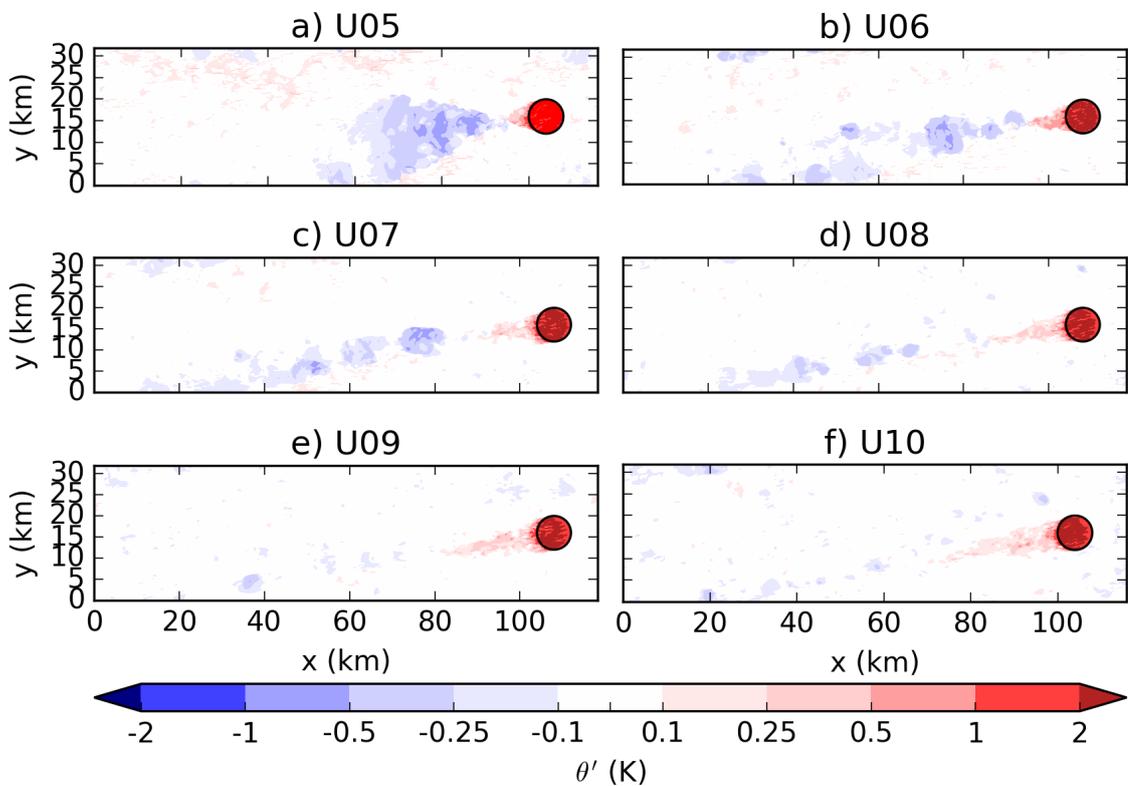


Figure 5.22 As in figure 5.14 but for the suite of wind speed experiments. Here, U10 is the Control experiment.

As the wind speed decreases and the CT circulation strength increases, the intensity of the resulting moist convection also increases. Here, we take a more quantitative look at this. We consider the total accumulated rainfall during daytime (between model hour 6 and 18). Peak daytime precipitation totals in the CT region increase for decreasing wind speed. This is the result of a combination of stronger convection with higher rainfall rates and raining convective cells persisting or reforming over the same spatial location for longer (see Fig. 5.23). It should be noted that while the convection is stronger, it never penetrates through our 3.5 km inversion to become deep convection.

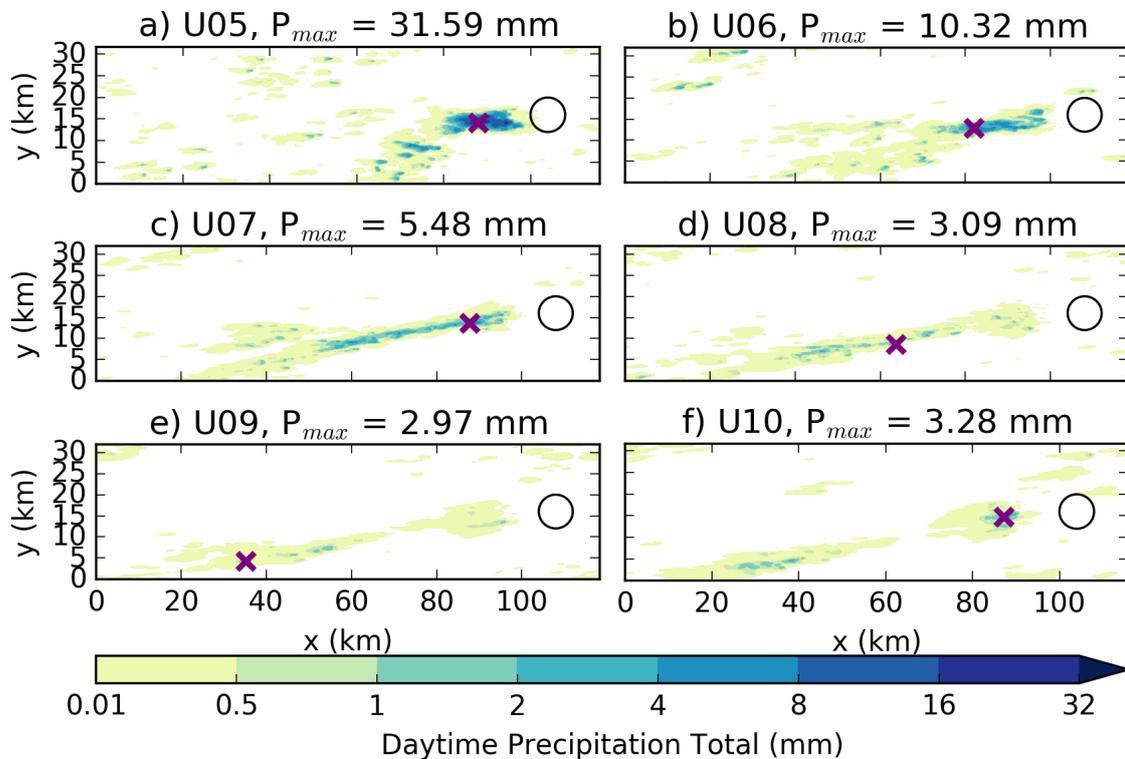


Figure 5.23 Total daytime precipitation for the suite of wind experiments. The island coastline is highlighted as the black contour in each panel. The maximum total daytime precipitation, P_{max} , is reported in the title of each panel and its location is marked by a purple ‘x’ in the panel. The Control experiment is U10.

The precipitation partially evaporates as it falls into the unsaturated sub-cloud layer. This evaporation contributes to the formation of cold pools which can be seen in Figure 5.22. In U10 there are a few weak and transient ‘cold pools’ which seem to move downwind with the background flow. These features become more intense, with stronger more well-defined negative anomalies in near-surface fields of θ , for de-

creasing wind speed. By U05, the cold pool becomes a near stationary feature with little progression downwind. This quasi-stationary behaviour is explained by repeated triggering of new convection in a region of targeted low-level convergence (not shown). Where this low-level convergence is strongly tied to the CT circulation. New convection in this region is preferentially triggered at the tip of the surface warm plume where the low-level branch of the CT circulation from either side of the warm plume meets.

There is clearly evaporative cold pool formation in many of these experiments (e.g. Fig. 5.22), and the intensity of the cold pools increases for decreasing wind speed. The horizontal vorticity within the ambient surface layer and the horizontal vorticity within the upwind edge of the cold pool are acting in the same direction. There is therefore no cancellation of flow relative to the cold pool gust front through vorticity arguments, and therefore convection which is triggered here must tilt up and over the cold pool. Rotunno et al. (1988) suggest that this upright triggering along the upwind edge of the cold pool is an important aspect of persistent convective systems as tilted low-level updraughts have to overcome the negative effect of increased entrainment to persist, precipitate, and initiate new convection.

That being said, the slow downwind pace of the cold pools can be discussed in terms of the cold pool phase speed, c . This can be estimated as $c = \sqrt{g'h}$, where $g' = g \frac{\Delta\theta_v}{\theta_v}$ and h is taken to be of the same order of magnitude as the boundary layer depth ($h = 500$ m). The highest phase speed might be near 4 m s^{-1} , but the phase speed averaged over each cold pool (which is a more representative measure of how it might interact with the environment) is generally 2 to 3 m s^{-1} for U05, and decreases for increasing wind speed as $\Delta\theta_v$ gets smaller in magnitude because convection and precipitation are weaker for stronger ambient winds.

In U05, the cold pool phase speed is closest to the boundary layer mean wind speed which is about 4.5 m s^{-1} . The leading edges of the cold pool travel with $u \pm c$ ($u + c$ when the wind and cold pool gust front's motion are in the same direction, $u - c$ when they are in opposite directions), then U05's cold pool fronts are the slowest of all the experiments presented here. This combined with repeated triggering of new convection at the tip of the warm plume gives the impression of a single stationary precipitating cell, where in actuality, there are many cells triggering and precipitating into the same

region, resupplying the cold pool with fresh evaporatively cooled air (this is seen in animations of cloud and vertical velocity fields, not shown).

The leading edge of the cold pool gust front forms a cold pool triangle as it expands with distance downwind from the warm plume. Along the leading edges of this triangle, low-level convergence supports triggering of new convection. This partially explains the ‘V’ shaped cloud region extending away from the cloud band in Figure 5.21. The rest of this ‘V’ shape, particularly in U05, is possibly due to asymmetries introduced to the system introduced as a result of the bi-periodic lateral boundary conditions. This favours suppressed conditions to the right of the CT band with respect to the flow (seen north of the cloud band in Figure 5.21a).

As U05 produces the most extreme changes to the CT system, two additional experiments are performed to help understand the vastly different behaviour. Firstly, U05 is repeated but for $H_{sfc} = 125 \text{ W m}^{-2}$, U05H125. Next U05 is repeated with large-scale rain processes disabled, U05_noRa. The cloud field and the surface warm plume for these experiments are shown in Figure 5.24.

In U05_NoRain, the cloud frequency is not as high in the cloud band and the two appendages that we have associated with the upwind edge of the cold pool from U05 are not as pronounced (Fig. 5.24c). Switching off the rain processes does not seem to have impacted how far downwind the warm plume extends before diminishing to a point due to the convergent low-level branch of the CT circulation, but it has removed any indication of the cold pool that was dominant in U05 (Fig. 5.24).

The buoyancy gradient scaling, the effect of which is captured by H/U from equation 5.5, is the same in U05H125 as in the Control experiment. This implies a CT circulation of similar strength. In U05_H125, the CT system does not return to the Control simulation’s regime. The CT circulation is still stronger than in the reference simulation, the CT band itself is more strongly precipitating, and it is oriented more strongly east-west. U05_H125 is closer to U05, rather the Control (U10), but with a shorter cloud band.

From the results presented in Figure 5.24, our buoyancy gradient scaling, H/U , only describes the CT circulation strength. To start to explain additional observed changes in the CT system, we need to include information about the warm plume

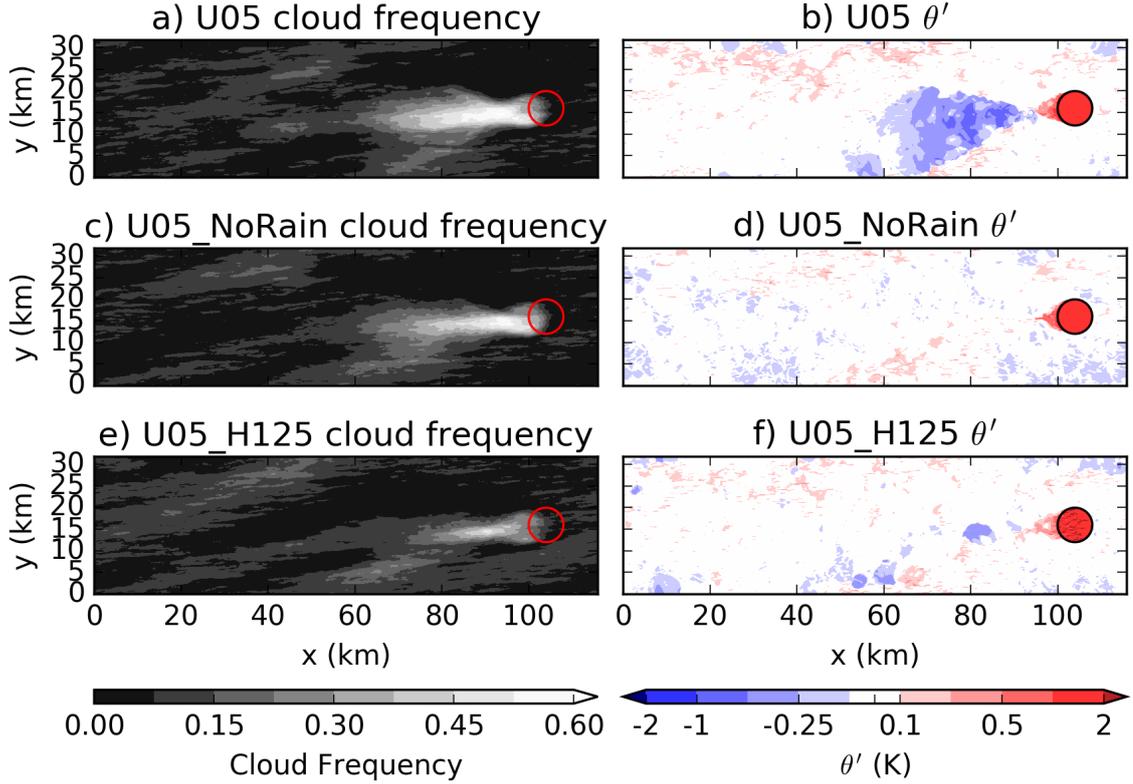


Figure 5.24 The daytime (model hour 6 to 18) cloud frequency on the left, and an instantaneous snapshot of the potential temperature anomaly, θ' , at hour 12. For a) and b) U05, c) and d) U05 with rain switched off, and e) and f) U05 with half the peak island sensible heat flux, $H_0 = 125 \text{ W m}^{-2}$.

geometry (i.e. the convergence of the surface warm plume into a triangle). This is controlled by the island-forced low-level across-flow component of the wind, v'_{CT} .

First, we return to our approximation for buoyancy perturbation (equation 5.4). This buoyant acceleration is acting as the flow crosses the island, and the low-level convergent across-flow wind component starts narrowing the low-level warm plume beyond the leeward edge of the island. We therefore are interested in the wind generated by this buoyancy once the flow has crossed the island. To approximate this we integrate in time by multiplying by another factor of $\Delta t \approx l_x/U$ to get the vertical velocity:

$$w_{CT} \approx \frac{gHl_x^2}{\rho c_{pd} z_h T_0 U^2}, \quad (5.10)$$

where all terms are as previously defined. From continuity, we can relate w_{CT} to the thermally-generated across-flow component of the wind by using:

$$\frac{\partial w_{CT}}{\partial z} = -\left(\frac{\partial v'}{\partial y'} + \frac{\partial u'}{\partial x'}\right), \quad (5.11)$$

where u' is the anomalous along-flow component of the wind, and x' is the coordinate in the direction of the flow. We assume that $\frac{\partial v'}{\partial y'} = \frac{\partial u'}{\partial x'}$, $\frac{\partial v'}{\partial y'} \approx \frac{2v'}{l_y}$, and integrate in the y' direction which we approximate by multiplying by l_y , then rearranging for v' :

$$v' = -\frac{1}{4} \frac{\partial w_{CT}}{\partial z} l_y, \quad (5.12)$$

Next we assume that the vertical gradient of the vertical velocity can be approximated as w_{CT}/z_h to get our CT across-flow wind component:

$$v'_{CT} = -\frac{1}{4} \frac{gHl_x^2 l_y}{\rho c_{pd} z_h^2 T_0 U^2}, \quad (5.13)$$

From here, the length of the surface warm plume can be estimated as the time it takes for the across-flow wind component to travel the characteristic island width in the across-flow direction, l_y/v'_{CT} , multiplied by the along-flow wind speed, U :

$$L_{WP} \approx \frac{l_y U}{|v'_{CT}|} = \frac{4\rho c_{pd} z_h^2 T_0 U^3}{gHl_x^2}, \quad (5.14)$$

where L_{WP} is the downwind extent of the warm plume and all other variables are as previously defined. We can now gather the factors which are kept constant across our suite of experiments into a single constant (e.g. thermodynamic constants and island geometry) to get:

$$v'_{CT} \propto C_v \frac{H}{U^2} \quad (5.15a)$$

$$L_{WP} \propto C_{WP} \frac{U^3}{H}, \quad (5.15b)$$

where C_v and C_{WP} are constants containing parameters which can be assumed constant across our suite of experiments as they describe the island geometry and mostly unchanged environmental conditions.

By this scaling, a two-fold increase (i.e. from 5 to 10 m s⁻¹) in wind speed would result in a 4-fold decrease in island-forced across-flow wind and an 8-fold increase in warm plume length. Therefore simply halving the surface fluxes is not sufficient to match the change in warm plume length without changing the wind speed.

In our suite of wind speed experiments, the scaling can be simplified further as H is constant across these experiments. We therefore predict that $v'_{CT} \propto 1/U^2$ and $L_{WP} \propto U^3$. We take the time-mean between model hours 9 through 12 of the n' field at 10 m height, and find the maximum magnitude of this field in the rectangular CT region to use as the estimate of v'_{CT} from our experiments (where the CT region is as originally defined in section 5.2.2.1). We then take the time-mean between model hours 9 through 12 of the θ' field at 2 m height, and find a continuous region of $\theta' > 0.1$ K which overlaps the island and take this to be the surface warm plume. The maximum distance downwind of this warm plume region is taken to be our estimate for L_{WP} from our experiments. We then test our scalings by plotting $1/U^2$ against our estimate of v'_{CT} from the experiments, and U^3 against our estimate of L_{WP} from the experiments in Figure 5.25a and b where we find that the scaling is consistent with the experimentally derived quantities in both cases.

We have thus shown that different overall CT responses exist for the same horizontal buoyancy gradient. U05.H125 and U10 have the same H/U but the CT response downwind is still different because of how the resulting CT circulation organises the downwind cloud field with distance away from the island. This organisation is described by the ratio U^3/H which approximates the cloud trail length. Using both ratios provides a more complete description of the resulting CT system.

Secondary to that, the magnitude of the wind shear between the boundary layer and the cloud layer decreases with decreasing wind speed. This results in a less dramatic tilt of convective updraughts with height for experiments with lower wind speeds. The more upright convection that results is more persistent. We speculate that the decreased tilt reduces the surface area of a given cloud and therefore decreases the entrainment of sub-saturated air from the environment. This implies that each combination of U and H_{sfc} results in a unique CT response described by H/U in the across-flow direction and U^3/H in the along-flow direction. Also, cold pool-relative

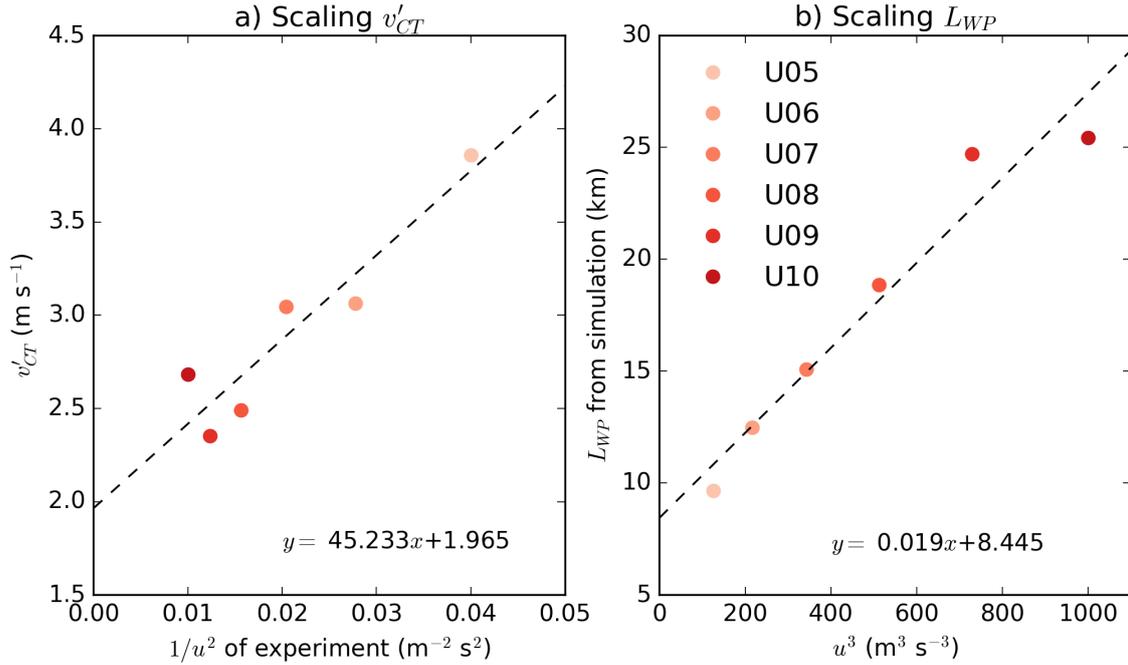


Figure 5.25 Comparison of a) the maximum of the mean low-level CT circulation strength derived from our experiments, v'_{CT} , and b) the length of the surface warm plume derived from our experiments, L_{WP} , with the scalings for the two parameters (see text above for more on how we derive these parameters from experiments and for the origins of the scalings). A line of best fit is included as a dashed black line in each panel, and the equation that describes the line is given near the bottom of each respective panel. Dots are coloured by background wind speed.

speed depends on U and precipitation. However, similar behaviour can evolve in the absence of cold pools as shown in U05_noRain.

5.4 Summary and Conclusions

This chapter demonstrates that the idealised UM can produce convincing simulated cloud trails given our idealised experimental set-up. We present non-hydrostatic representations of CT on an idealised small island with a grid length of 100 m. Our simulations support the basic understanding of the CT system with new insights on the interaction between the turbulent warm plume and circulation components. We have defined a transition region between the turbulent cloud patch and characteristic CT cloud band in which buoyant processes decline and mechanical processes take over. This is a combination of the time taken to lift near-surface parcels to their saturation

by the ascending branch of the CT circulation and the extent of capping due to the warm plume which presents an obstacle to updraughts reaching the LCL in this region. Furthermore, we show that combinations of wind speed and surface fluxes can yield the same net island heating, yet also result in CT with meaningfully different morphologies.

A series of open questions from Chapter 3 are addressed through three suites of experiments. In the first suite, the island forcing is perturbed. First, our experiments show that the island surface sensible heat flux is the important control on the CT response. We then find that there is a straightforward linear relationship between the strength of the island forcing and the CT strength - supporting our original hypothesis from Chapter 3.

Next, we investigated the role of relative humidity in the CT system. Is it just a matter of turning the clouds on or off? We show that to first order, yes. The warm plume and CT circulation are still present in BLm25 (no cloud band) but the circulation is not as robust for BLm25, and further work is required to fully separate the role of clouds and the changes to the boundary layer in this difference.

Furthermore, we demonstrate that reduced free-tropospheric relative humidity does not prevent the formation of a cloud band, nor make it stronger. [Kirshbaum and Fairman \(2015\)](#) showed that free-tropospheric relative humidity was significantly lower for well-formed CT past the Lesser Antilles Islands. We speculate that this result might be due to features of the large-scale environment specific to this region which are correlated to conditions supportive of CT described by the remainder of the profile. Alternatively, we offer that the larger, more mountainous islands of the Lesser Antilles could force stronger circulations which require a dry free-troposphere in order to prevent strong precipitation over a given island which would disrupt the shallow CT circulation. In a real-world case, cloud cover and precipitation over land would also modulate the strength of the surface fluxes.

Finally, we ask what is the role of the wind speed in the CT system? Literature and results from Chapter 3 indicate that CT circulation strength is inversely proportional to the wind speed. Our results support that understanding. However, we demonstrate that non-linear responses related to cold pools driven by evaporation of precipitation,

and the shape/size of a warm plume becomes increasingly important as the wind speed decreases in our simulations. At the lower wind speeds tested, where forcing is strong enough for precipitation, the low-level RH might then become important for controlling the strength of the cold pools in addition to controlling the height of the LCL as already noted.

This change in behaviour with wind speed has implications from a forecasting perspective. As wind speed decreases and CT circulation strength increases, there is an increasing likelihood for precipitation within the CT system. We show that should the wind speed become low enough for the precipitating part of the system to become quasi-stationary, significant (>25 mm) daily rainfall accumulations in the lee of the island are possible. We suggest that there may be a wind speed-island geometry configuration for which the precipitating part of the CT system would travel upwind, bringing precipitation over the island itself. Personal communication with the Bermuda Weather Service suggests that while rare, precipitating CT that occasionally grow into thunderstorms, typically occur in Bermuda for southwesterly winds of 8-12 kts (i.e. $4-6 \text{ m s}^{-1}$ winds, parallel to the long axis of the island). Cloud cover forming over the island itself likely reduces H during the day compared to our simulations, while island geometry (Bermuda is not circular) would likely counteract this effect for long-axis parallel flow.

It should be noted that the CT system response is well understood and is fundamentally a simple problem to reproduce under a wide range of environmental conditions - both transient and in equilibrium as suggested by the observations from Chapter 3 and early numerical studies of the thermal response to heated islands (e.g. [Estoque and Bhumralkar, 1969](#); [Mahrer and Pielke, 1976](#); [Savijarvi and Matthews, 2004](#)).

Chapter 6 takes the Control experiment from the current chapter and attempts to perform it for larger grid spacing. We then take a convective cloud census for the Control and larger grid length experiments. To put them into the context of the UM shallow cumulus parametrisation, we analyse and compare the resulting cumulus demographics between experiments.

Chapter 6:

Cloud Trails in the Shallow Convective ‘Grey Zone’

6.1 Introduction

In Chapter 5, we presented a control experiment (Section 5.2.1) and several permutations of that experiment which aid in further understanding the response to the shallow cumulus field to the introduction of a small heated island. In the previous Chapter’s control experiment, the island’s presence rearranges clouds spatially with only small changes to the cloud statistics themselves.

In Chapter 4 Section 4.5.1 we used the BOMEX (Siebesma et al., 2003) inter-comparisons to indicate that increasing the grid length used increases the strength of overturning eddies and convective updraughts in the UM. Doubling grid spacing to 200 m, results in updraughts that are too strong and too penetrative. These updraughts start to interact with the rigid model lid at 3 km height when grid spacing is increased to 400 m or more. During the initial spin-up, stronger overturning for larger grid lengths grows, warms, and dries the boundary layer. After the numerical spin up, we are left with a deeper boundary layer, a higher LCL, and altered vertical thermodynamic profiles such that the tendencies resulting from our prescribed forcing profiles might not still be in equilibrium.

We take this to be descriptive of the “shallow cumulus grey zone”, and we revisit this grey zone in the current chapter. Horizontal grid lengths are increased from 100

m to 1.6 km. This spans the range of numerical regimes from reasonably well-resolved shallow convection, to shallow convection which is not truly resolved at all. Furthermore, the 1.6 km grid length case is of the same order magnitude as that used in the current implementation of the Met Office’s UKV regional model. We note that the Met Office currently do not employ any form of convection scheme for this model (personal communication).

Traditionally, 1.6 km grid length is considered to be ‘convection permitting’. With this grid spacing, deep convection can be crudely represented by the resolved motions and this is allowed in the UKV. However, shallow convection is not represented by resolved motions and isn’t treated through a convection scheme. The shallow convective motions must either be missing, projected onto the resolved scales (resulting in too-strong motions), or unintentionally treated by another scheme e.g. the boundary layer turbulence scheme.

The current chapter discusses naively moving to higher resolution from the relatively coarse UKV, and the benefits of explicitly representing shallow convection. This is done through a cloud census and the subsequent evaluation of cloud population statistics for the control cloud trail simulation repeated for different horizontal grid lengths. We take the 100 m control experiment from Chapter 5 as the “truth” and consider the other experiments through the perspective of a shallow convection scheme.

6.2 Methods for the Grey Zone

We saw in Chapter 4 that poorly-resolved shallow convection grows too strong and too deep. The control experiment of the previous chapter uses initial conditions which we derive from a small-domain, sea-only, spin-up simulation (see Chapter 4, Section 4.6.2 for more detail). Equilibrium initial conditions are used for these new CT experiments to simplify analysis by minimising the time evolution of the large scale state, and are achieved using the methods outlined and argued in Chapter 4.

6.2.1 Choice of Initial Profiles

In Chapter 5, the initial conditions used for experiments include an inversion near 3 km height which “caps” convection to a cloud layer. These initial conditions were found using a spin-up method which we outlined in Chapter 4 Section 4.6.2. However, we have shown in Chapter 4 Section 4.5.1 that resolved convection is deeper for coarser resolution. The aforementioned inversion near cloud top contains the convection to a layer between 600 and 3200 m in our well-resolved experiments. However, an undilute surface parcel ascent would not reach its level of neutral buoyancy until the tropopause (e.g. Chapter 4, grey profile in Figure 4.14). Given that poorly-resolved cases already tend to have too-deep convection, we choose to implement a stronger inversion at cloud top to prevent this behaviour. We increase the temperature at 3.2 km by 5 K and then follow a pseudo adiabat until it intersects the original profile near the tropopause.

Because we have made changes to the initial conditions, new spin-up simulations are required to achieve balanced states for our CT simulations. A new spin-up is performed for initial conditions specific to each experiment’s grid length. These spin-up simulations continue to have 32×32 horizontal grid points and so the domain size grows with increasing grid length. The initial RH, wind, and large-scale forcing profiles are unchanged from the control experiment. Interactive surface fluxes are allowed, and the simulation is run for ten days (as in the control spin-up).

6.2.2 Design of the CT Experiments in the Grey Zone

Our grey zone island experiments with different grid lengths are conducted by changing as little as possible about the physical set up. The domain size is unchanged from the control experiment (118.4×32 km), and the island remains defined as a 50 km^2 patch centred at $x = 108$ km, $y = 16$ km. The experiments are labelled and the resulting number of grid points in each direction, and number of island points are summarised in Table 6.1.

We remind the reader that while no changes to the model physics are imposed between experiments with different grid lengths, some subgrid parametrisation schemes adjust to grid length automatically. The blended boundary layer and sub-grid turbu-

Table 6.1 Experiment labels and the corresponding number of grid points in the x - and y - directions, N_x and N_y , and the total number of island grid points, N_{island} , for the Grey Zone resolution experiments. The corresponding dimensions for the control experiment are included, but this experiment is not repeated with the new initial conditions.

Experiment ID	N_x	N_y	N_{island}
Control	1184	320	5000
DX0200	592	160	1250
DX0400	296	80	312
DX0800	148	40	78
DX1600	74	20	19

lence scheme is primarily weighted toward 3-D Smagorinsky at 100 m and is heavily weighted toward 1-D Lock et al. (2000) at 800 m (see Chapter 4, Section 4.2.2.1 for a more complete description). This has knock-on effects on the cloud scheme which uses sub-grid turbulence to diagnose the distribution of sub-grid moisture perturbations. In addition, the constraints on the RH_{crit} also change with grid length (see Chapter 4, Section 4.2.2.2).

It is useful to point out to the reader that the results in this chapter will be somewhat specific to these initial conditions. If we were to use the original 100 m CT initial conditions from the previous chapter, too-deep convection erodes the cloud top inversion and the domain is dominated by deep convection which penetrates to the tropopause and is strongly precipitating (particularly in DX0800 and DX1600). In addition, there is no prescribed idealised forcing above our cloud top inversion and so deep convection reaching above those heights becomes problematic for maintaining balanced moisture and energy budgets. It is unclear whether a CT would still emerge for these cases.

6.3 Sensitivity of Spin-up Simulations to Grid Length

We perform the spin-up simulations as outlined above for grid lengths of 100 m, 200 m, 400 m, 800 m, and 1.6 km. The 100 m spin-up is for comparison with the control experiment of the previous chapter. In the following analysis, we refer to the

100 m spin-up of the previous chapter as “DX0100C5”, and the 100 m spin-up of the present chapter as “DX0100C6”.

First, we consider changes to the boundary layer. As mentioned in Chapter 4, the boundary layer height is a good bulk diagnostic for the overall equilibrium as it depends on multiple components of interest. There are some monotonic (but small) changes to the system with increasing grid length (Fig. 6.1). The LCL and the boundary layer height both increase with increasing grid length (Fig. 6.1a and b). These patterns are consistent with our previously established conceptual model of stronger boundary layer overturning with increasing grid spacing. During the spin up, the stronger overturning and associated entrainment of air from above yields a warmer, drier, deeper boundary layer with a higher LCL for larger grid spacing. The median cloud top height increases slightly with increasing grid length (Fig. 6.1c). However, this combined with cloud bases increasing (implied by increasing LCL) means that the total cloud depth is mostly unchanged. Finally, the surface latent heat flux increases with grid length (Fig. 6.1d). Given that the near-surface wind speeds (to be discussed next) are almost identical across experiments, this must be controlled by the near-surface moisture deficit. Here, we argue that this is because the boundary layer grows deeper, drier, and warmer with increasing grid length.

Changes to the initial thermodynamic profiles that result from the spin-up are mostly small and are shown in Fig. 6.2. There are some small differences in the mean wind profiles, these are mostly around the surface layer (lowest 100 m) and within 100 m of the top of the boundary layer (Fig. 6.2a). These differences can presumably be explained through differences in the boundary layer height and the strength of the surface sensible heat fluxes (Fig. 6.2b). These drive mixing within the boundary layer of different intensity and also control the momentum exchange between the boundary layer and the free atmosphere.

Meaningful changes to the stability in the lower cloud layer occur, particularly for the DX0800 and DX1600 spin-up experiments. In these coarser experiments, the 750 m to 1 km layer becomes weakly capped during the spin up, while this layer has near-constant stability for the DX0200 and DX0400 spin-up experiments. There are then related changes to the moisture profile. The clouds in the DX0800 and DX1600

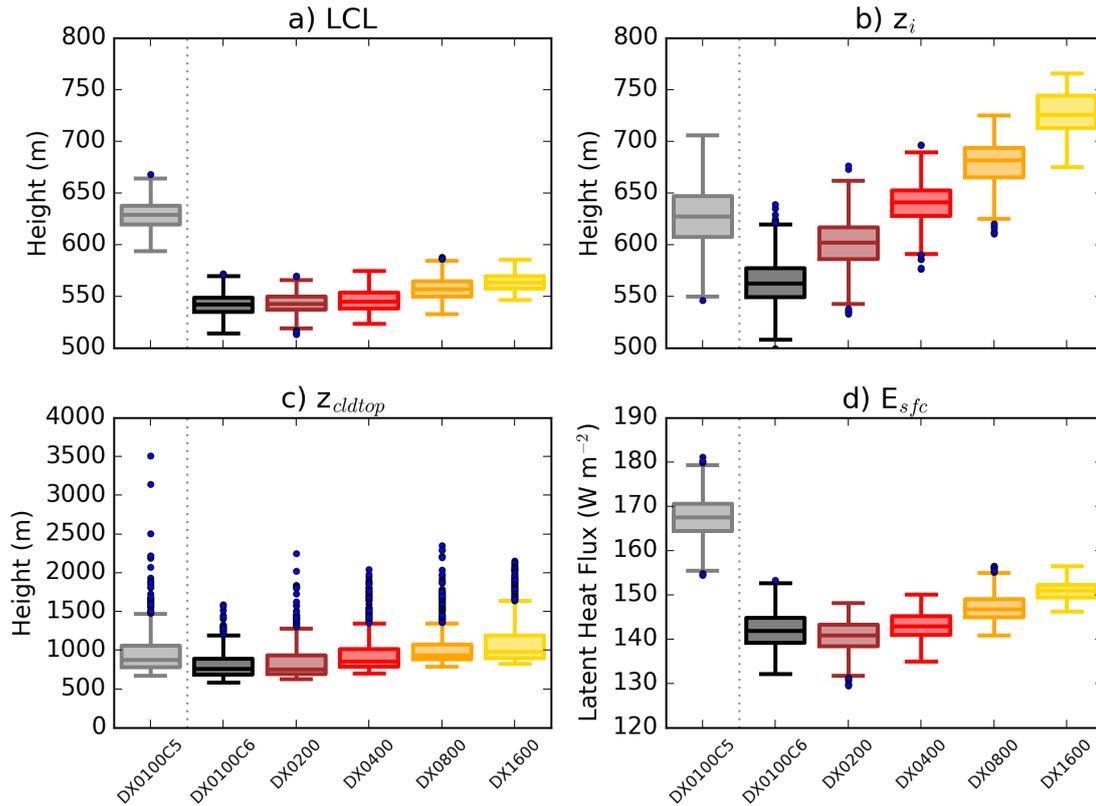


Figure 6.1 Boxplots of the horizontally-averaged a) LCL, b) boundary layer depth, c) cloud top height, and d) surface latent heat flux over the last four days of the spin-up simulations for the experiments with different grid lengths.

spin-up experiments do not reach into the upper cloud layer (1 - 3 km) as frequently as a result of this capping. Less moisture is then able to reach these heights. Meanwhile, the prescribed idealised subsidence continues to dry this layer resulting in a drier upper cloud layer than in DX0100C6, DX0200 or DX0400 spin-up experiments (Fig. 6.2b).

These changes to the stability in Fig. 6.2b can also help to explain changes to the cloud itself. In terms of horizontally-averaged cloud liquid water, DX0100C6, DX0200 and DX0400 spin-up experiments look very similar throughout with a fairly evenly distributed vertical profile. Conversely, the DX0800 and DX1600 spin-up experiments look similar between 1 km and 3 km height, but are very different below 1 km. We find that a doubling in grid spacing corresponds to a doubling in cloud liquid water at cloud base (Fig. 6.3a). This increase does not change the total effective cloud cover in a meaningful way except for in DX1600 in which it is significantly higher (not shown). The explicit representation of convection with 800 m and 1.6 km grid lengths results

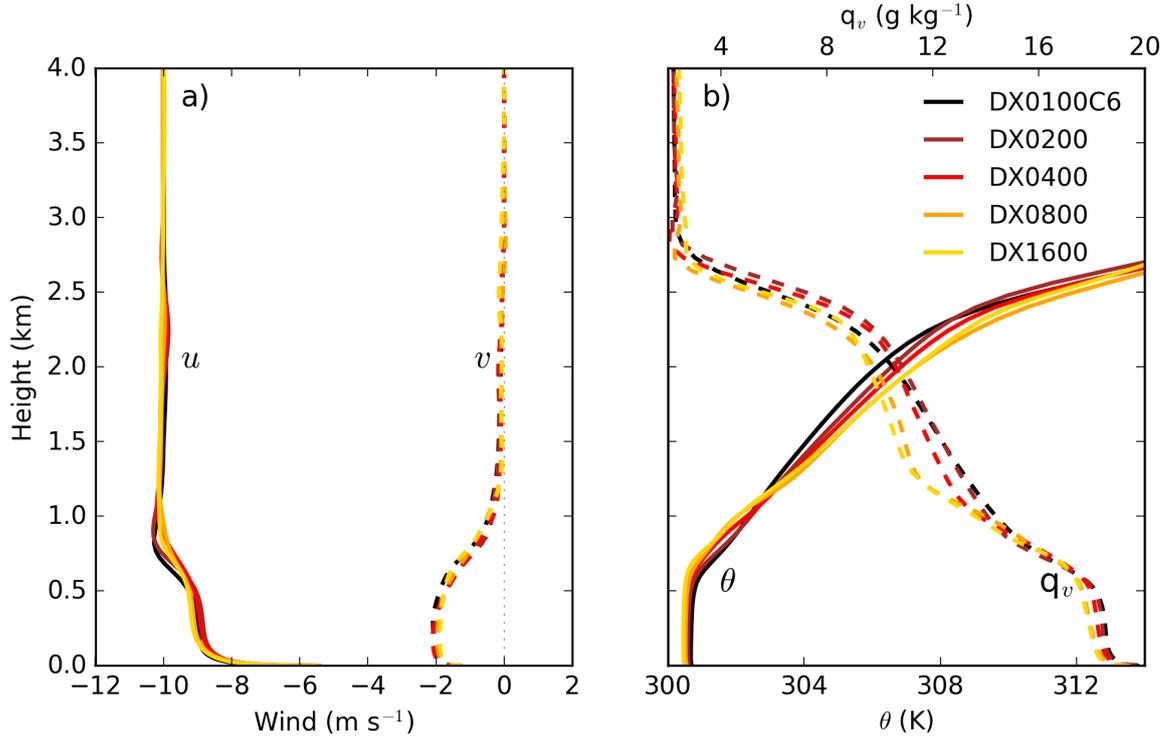


Figure 6.2 Horizontally-averaged vertical profiles of a) the u- and v- components of the wind in m s^{-1} where u is given as solid lines and v as dashed lines, and b) the potential temperature, θ , in K in solid lines and the specific humidity q_v in g kg^{-1} in dashed lines. These are a comparison across the equilibrium conditions for spin-up simulations with different horizontal grid lengths.

in a different response to the capping of the lower cloud layer.

Assuming all of the convection is triggered at the surface, a fraction of the positively buoyant convective plumes remain trapped within the boundary layer. A proportion of these plumes rise above their LCL (and often exit the boundary layer in the process) to become cloudy plumes. Of these cloudy plumes, a percentage of them are negatively buoyant above the boundary layer or entrain too much environmental air and dissipate. The remaining positively-buoyant cloudy plumes continue to rise through the bulk of the cloud layer. The most weakly buoyant cloudy plumes terminate lower in the cloud layer. Some plumes are so strongly buoyant that they rise unimpeded through the entirety of the cloud layer and they encounter our imposed inversion at about 3 km height. Here, all of the plumes become negatively buoyant and eventually terminate. At this 3 km inversion, some of the cloudy plumes spread out horizontally, becoming anvil-like features. The above conceptual model is useful

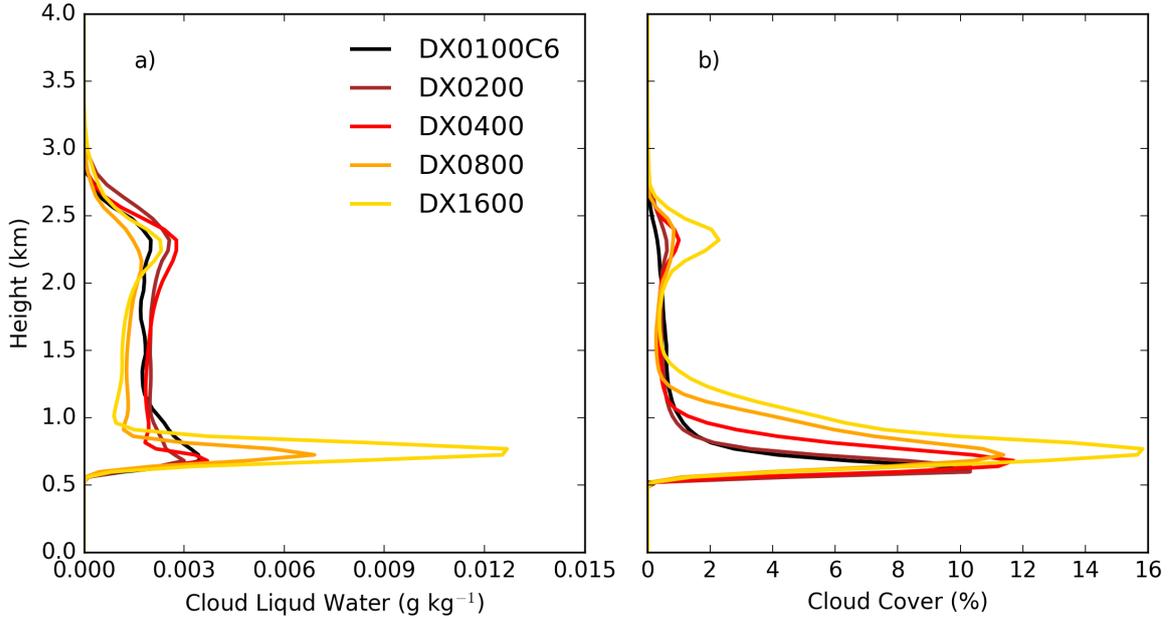


Figure 6.3 Horizontally-averaged vertical profiles of a) cloud liquid water content in g kg⁻¹, and b) cloud cover as a fraction of the whole domain. These are a comparison across the mean conditions during the four days used to define equilibrium state from the spin-up simulations with different horizontal grid lengths.

for understanding the vertical structure in both panels of Fig. 6.3.

Above we have shown that there are mostly minor differences between the final equilibrium states of these spin-up simulations for different grid lengths. It is likely that if we were to conduct our experiments with the same initial conditions, the differences they would evolve toward the solution in the spin-up simulation. But this would be on top of the CT and any effect it has. Moreover, we might expect to see similar differences emerge if using a data assimilation system in an operational setting. We therefore choose to cautiously proceed with using these equilibrium states as initial conditions in the reference CT simulation. Conceptually, the differences between CT simulations for different grid lengths can be decomposed as follows:

$$\Delta\phi_{total} = \Delta\phi_{ic} + \Delta\phi_{\Delta x} \quad (6.1)$$

where $\Delta\phi_{total}$ is the total difference between two CT experiments, $\Delta\phi_{ic}$ is the difference due to changes in the environment which result from changing grid length and the resulting spin-up, and $\Delta\phi_{\Delta x}$ represents the difference due to changes in the representation

of the CT due to changes in grid length.

In the following sub-section, we use the initial conditions from the above spin-up experiments and apply them to CT simulations to arrive at reference simulations for our range of grid lengths, 200 m, 400 m, 800 m, and 1.6 km. Note that we do not repeat the 100 m experiment with the new initial conditions as we expect the CT evolution to be largely unchanged by these changes in initial conditions for this Δx . In cases where a comparison is useful, we include the control experiment from Chapter 5 with alternate symbology.

6.3.1 Results for the Grey Zone

Here, some expectations from both the BOMEX experiments performed with larger grid lengths, and the spin-up simulations from above are informing the expectations for the changes in behaviour in the CT system for the experiments below.

In the BOMEX experiments, we saw stronger convective boundary layer eddies for larger grid spacing. In contrast to the ‘random’ eddies of the BOMEX simulations, we now introduce a forcing (the island) which triggers an organised boundary layer structure. Given the behaviour with increasing horizontal grid length in the BOMEX experiments, we expect that the CT circulation strength will increase with increasing grid spacing.

This increase in CT circulation strength is then expected to have downstream implications. First, the cloud band is expected to consist of deeper convection for larger grid spacings. As a result, the cloud band would contain convective cells which are more strongly precipitating. The stronger precipitation may trigger stronger cold pools which can start to disrupt the CT circulation and make the cloud band more intermittent.

6.3.1.1 The Cloud Field

As in Chapter 5, we start with an overview of the CT system for these experiments with increased grid spacing. We first consider the daytime cloud frequency for each of the grid length experiments of the current chapter and include the 100 m CT simulation

from Chapter 5, Section 5.2.1 for comparison. Recall that a direct comparison between the 100 m and the rest of the grid length experiments cannot be made because the increase in inversion strength does impact the ambient cloud behaviour through a reduction in the number of convective plumes that penetrate into the inversion and how far they do penetrate.

The daytime cloud frequency is shown in Figure 6.4. We expect three main regions based on the behaviour of the CT in the 100 m experiment (Fig. 6.4a): a turbulent cloud patch in the immediate wake of the island; a less cloudy “void” downwind of that; and finally a cloud band oriented with the low-level wind direction. In DX0200, these features are mostly preserved (Fig. 6.4b). However, in DX0400 the turbulent cloud patch has mostly been lost, and in DX0800 and certainly in DX1600, this turbulent cloud patch is absent (Fig. 6.4c, d, and e).

A general pattern of changes in the mean behaviour of the cloud trail system emerges. With increasing grid spacing, the turbulent cloud patch vanishes as the cloud ‘void’ region migrates upwind toward the island, and the cloud band region is populated with increasing cloud frequency.

Furthermore, the turbulent cloud patch already shows signs of disappearing in DX0200, a band of higher cloud frequency appears in the turbulent cloud patch on the right hand side with respect to the flow in Figure 6.4b. In animated sequences of the cloud mask field and of the 10 m across-flow, n -component of the wind (not shown), the structures that emerge in the turbulent cloud patch arise from existing upwind boundary layer eddies that are invigorated as they pass the island. In DX0200, this occurs most strongly to the north side of the island, but we do not investigate this any further.

In addition, in DX0800 and more obviously in DX1600, the domain size comes into question. The cloud band is nearly 1.5 times the length of the original 100m CT simulation and the DX0200 cloud bands. This results in strong interactions with the tail of the cloud band and the forcing over the island starting just before midday. In Figure 6.4d and e, the cloud band is seen extending west-southwestward from the island and then re-entering the domain in the top right corner passing just north of the island. It should be noted that the low-level circulation is likely centred on the southern edge

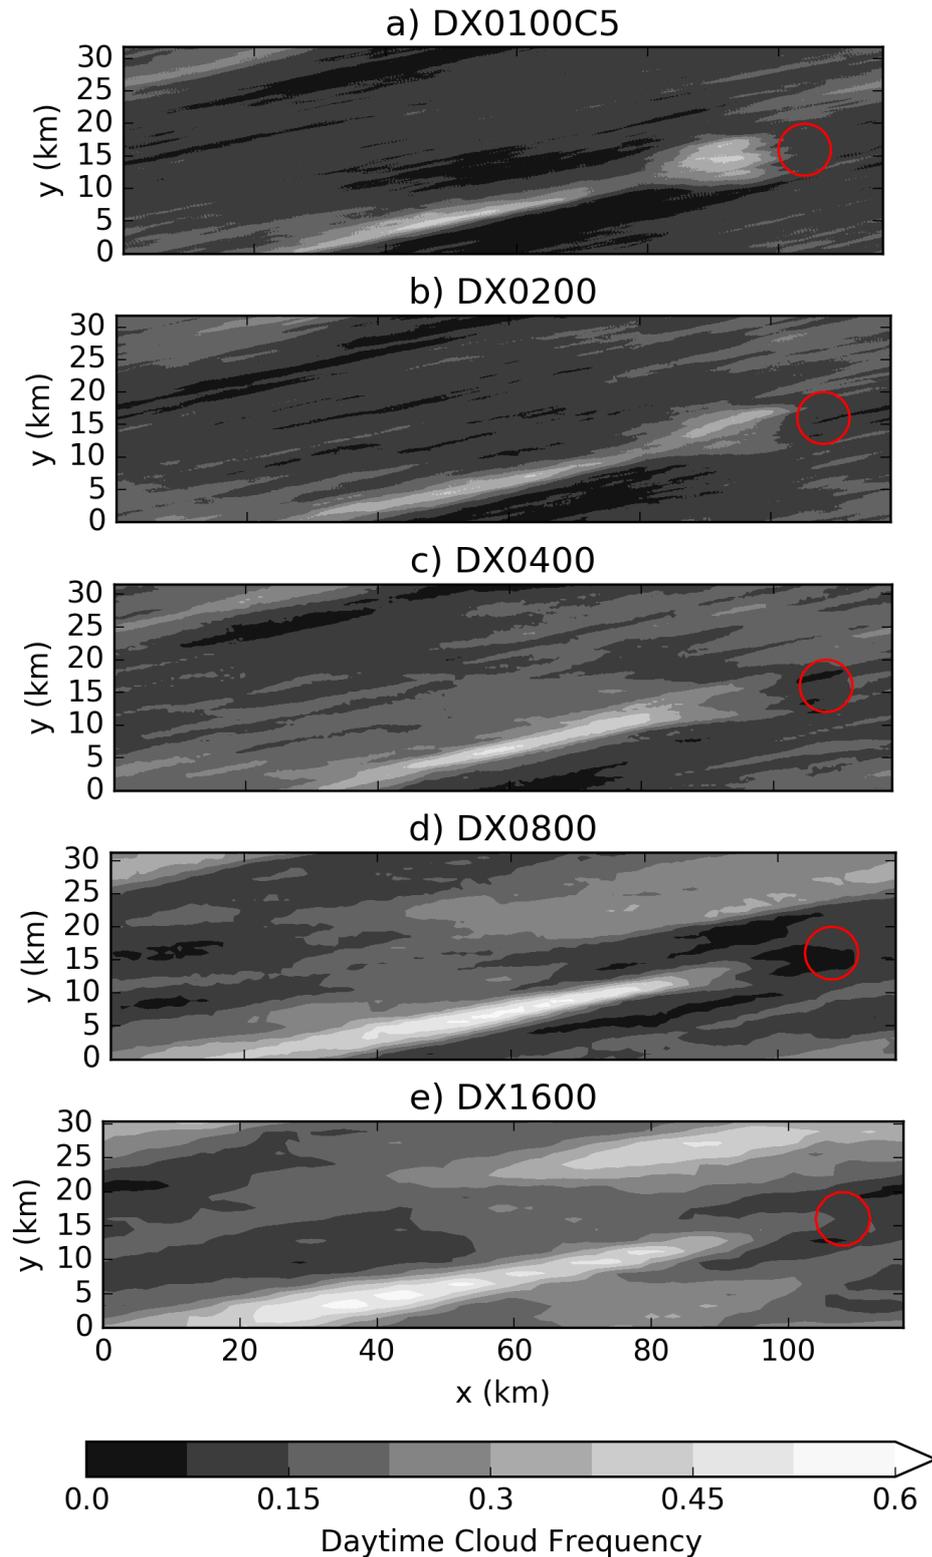


Figure 6.4 The daytime (0600 to 1800 hours) cloud frequency for the experiments which vary horizontal grid spacing. a) The original 100 m CT simulation from Chapter 5, b) the 200 m experiment, c) the 400 m experiment, d) the 800 m experiment, and e) the 1.6 km experiment. Recall that Chapter 5’s 100 m CT simulation has a different inversion atop the cloud layer and so is just for reference.

of this cloud band as the vertical wind shear across the LCL is to the right (north) of the orientation of the cloud band. Furthermore in DX0800 and DX1600, as mentioned before cloud are more strongly split into a group extending just above the boundary layer, and a group which reaches to the cloud top inversion near 3 km height. The cloud band looks like it is oriented more east-west in DX1600 than DX0200, but this is because the cloud band consists more of the deeper clouds which are advected westward with the cloud layer wind. At cloud base, the CT therefore has the same orientation in all experiments.

In the spin-up simulations, the equilibrium conditions and the bulk cloud profiles suggest that two groups of behaviour emerge in this range of grid lengths: the more well-resolved cases (200 and 400 m) and the poorly resolved cases (800 m and 1.6 km). However, there is a smooth transition in the realised CT structure in the cloud frequency fields as dx is increased from 200 m to 1.6 km, rather than the regime shift between 400 m and 800 m that one might suspect from the spin-up simulations.

It should be noted that these cloud frequency fields are good for highlighting the cloud band by design. However, an instantaneous daytime snapshot of the cloud field might not see such a well-defined cloud band structure. For example, Figure 6.5 shows the midday snapshot of the liquid water path for each experiment. Here, we get a sense that there is some structure and organisation to the cloud field, but a distinct cloud band is not necessarily present. Rather, it is the preferential formation of longer-lived cloudy elements within the cloud band region which allows the structure to emerge when averaged into the cloud frequency field.

Given that the cloud band re-enters the domain and strongly interacts with the island, and this behaviour becomes stronger with increasing grid length, we continue to choose to focus analysis on the first half of the simulation period for the bulk of the remaining analysis. The reader should recall that real-world groups of islands do have CT cloud bands which interact with each other, and so this simulated behaviour has elements of realism for those cases. However, this form of interaction is not within the the scope of this study.

Next we consider the peak circulation and circulation strength for the suite of experiments. We discuss how this varies with increasing grid spacing. The changes

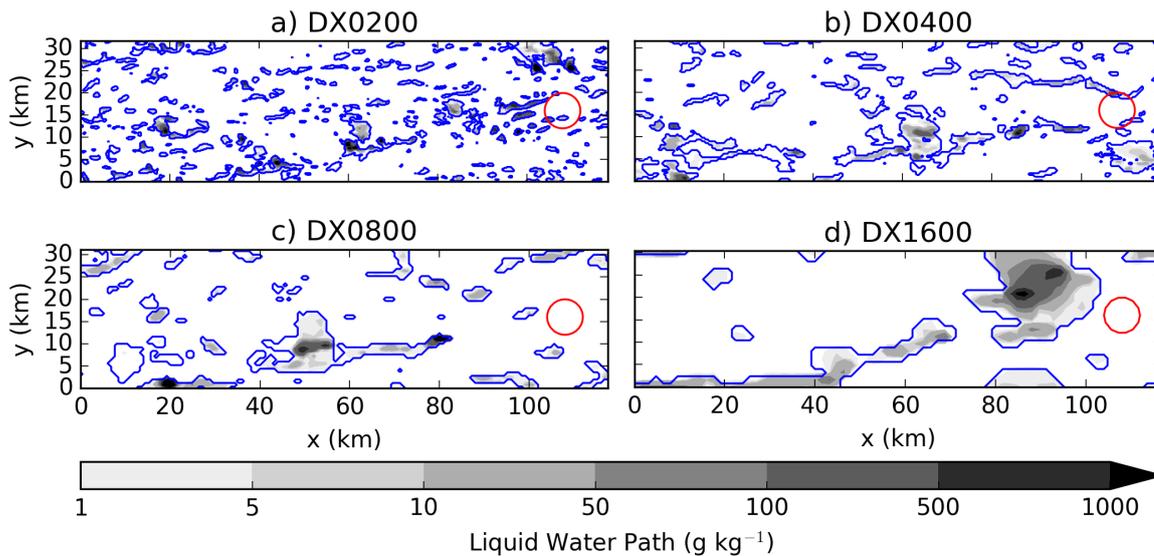


Figure 6.5 An instantaneous snapshot of the liquid water path at 12PM in a) the 200 m experiment, b) the 400 m experiment, c) the 800 m experiment, and d) the 1.6 km experiment. Here, the cloud edge contour (i.e. where the liquid water path increases above 0 gm^{-2}) is marked with a blue contour, and the island coastline is marked with a red contour.

to the circulation are then discussed in the context of the warm plume which we have shown to be a diagnostic for the underlying drivers of the CT system in Chapter 5.

6.3.1.2 The CT circulation strength

From our analysis of the BOMEX experiments using larger Δx , we expect the boundary layer eddies to be over-done (i.e. too strong) in the current experiments with larger grid spacing. While our augmented initial conditions prevent the too-strong convection penetrating past our inversion, we still expect the convective motions (e.g. eddies and the CT circulation) to be too strong relative to the experiments with smaller grid spacing.

We start by inspecting the low-level circulations. As in Chapter 5, Section 5.2.2.3 we decompose the wind into a component perpendicular to the flow, n , and a component parallel to the flow, s , where the ‘flow’ is defined as the domain mean wind vector at that time and height. Recall that n horizontally averages to zero, and s horizontally averages to the horizontal-mean wind speed by definition. Considering these wind components is useful for assessing the strength and extent of the low-level circulation

part of the CT system.

Looking at n at 10 m height for a snapshot at model hour 12 (midday) in each experiment, we see that the main low-level convergent feature associated with the CT circulation is present for all experiments (Fig. 6.6). However, there are some major differences in the smaller scale resolved structures within that low-level circulation and in the space either side. In DX0200, there are still many small-scale circulations with low-level convergence present in the background, and as they move across the island (as seen in animated sequences of this field, not shown), they result in small perturbations to the dominant CT circulation (e.g. Fig. 6.6a). These background perturbations become fewer and larger for increasing grid length. These stronger background eddies are amplified as they cross the island and occasionally trigger deeper convection near the island in DX1600 resulting in a few large undulations in the convergence line (e.g. Fig. 6.6d).

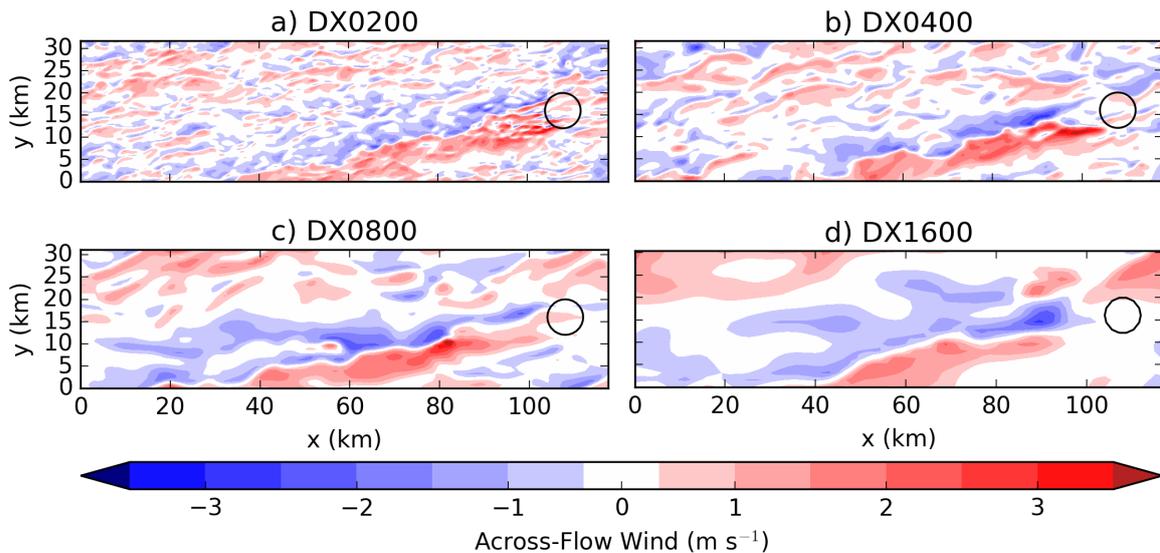


Figure 6.6 A snapshot of the across-flow component of the wind, n , at 10 m height and 12PM in a) DX0200, b) DX0400, c) DX0800, and d) DX1600. The island coastline is shown as a black contour in each plot.

Also, the CT circulation itself becomes increasingly dominant in the domain. In Figure 6.6d, the CT circulation is essentially the only low-level convergent structure present, with maybe two or three other smaller circulations in the middle of the domain. The CT circulation is seen wrapping back around extending a region of positive n over the island again (the red shading over the top right portion of the island in Fig. 6.6d).

We speculate that the stronger circulation in DX1600 decays more slowly with distance downwind from the island than in the more well-resolved experiments. This slower decay allows the CT circulation to persist long enough to be carried back across the bi-periodic lateral boundary conditions.

The interaction of the ambient eddies in DX0800 (and more so at DX1600) help to explain some of the intermittent behaviour of the cloud band mentioned in the previous section. When the CT system is more mature, the already too-strong eddies are enhanced as they pass over the island (in DX0800 and DX1600, not shown). The cloud band breaks up into cloud line segments with areas of enhanced and suppressed low-level convergence along the main CT band associated with these eddies (particularly in DX1600, somewhat apparent from Fig. 6.5c and d). At the tail of the cloud band, where the low-level circulation starts interacting with the island due to the bi-periodic lateral boundary conditions, precipitation starts to produce low level motions which are likely interacting with the circulation formation region over the island (Fig. 6.5d and Fig. 6.6d). The increasing importance of precipitation and changes to the representation of the convection itself as we increase the grid spacing are both discussed in more detail in Section 6.3.1.3.

We have now shown that there are changes to the CT circulation, but how does M_c , our measure of overall CT strength defined in Section 5.3.1 of Chapter 5, respond to those changes? M_c is an approximation for the mean vertical mass flux at the horizontal-mean LCL in the CT region, multiplied by the cloud band width. As expected, the CT circulation strength increases with increasing Δx (Fig. 6.7). However, the pattern of increase is not straightforward.

As in the well-resolved wind experiments of Chapter 5 Section 5.3.3, we see the complicating role of precipitation once again (this is discussed more below in Section 6.3.1.3). While M_c does continue to increase, some of this measure of CT circulation strength is offset by increasingly strong downward motions induced by evaporation of precipitation and overlap between the end of the CT and the island in DX1600.

The general pattern of the warm plume is preserved across the range of grid lengths tested. At the surface, it consists of a region of positive temperature anomalies which narrows with distance downwind from the island (Fig. 6.8). As discussed in Chapter

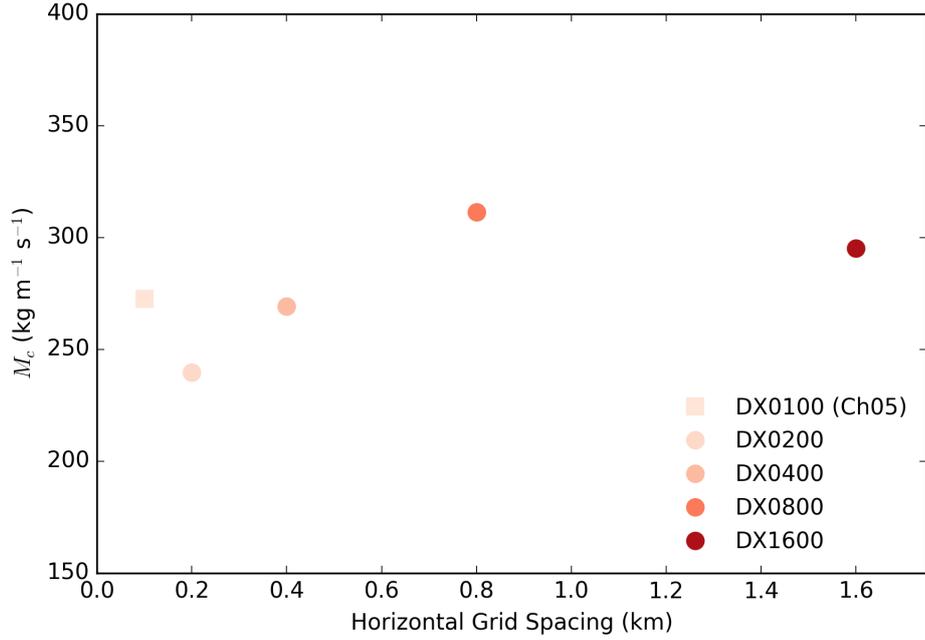


Figure 6.7 The circulation mass flux, M_c , as defined in Chapter 5, Section 5.3.1. This is computed over the period from 9AM to 12PM in each of the experiments with increasing grid spacing. The 100 m experiment from Chapter 5 is included here as a square to give a sense of what the expected magnitude of M_c should be for these experiments. Recall that this experiment should not be considered as part of the main suite of experiments because it has different initial conditions.

5 Section 5.2.2.2, this narrowing behaviour is due to the convergent branch of the low-level CT circulation.

Of note, the instantaneous warm plume extent is more or less unchanged between DX0200 and DX1600 (Fig. 6.8a-d). The low-level n wind component is comparable across each experiment. Therefore the downwind extent of the surface warm plume is not significantly altered by changing the grid length. What does change is the introduction of patches of stronger negative θ'_{sf_c} to the right (north) side of the warm plume. Despite having similar low-level convergence strength by eye, the experiments with larger grid lengths have stronger M_c . We speculate that the vertical structure to the convergence within the boundary layer is different across these experiments. Through mass continuity, the vertically integrated convergence below the height of the LCL gives the vertical velocity at the LCL. Hypothetically, stronger low-level convergence at a height above the level analysed, or stronger divergence near the top of the surface to LCL layer, could support the changes seen in M_c . Stronger updraughts

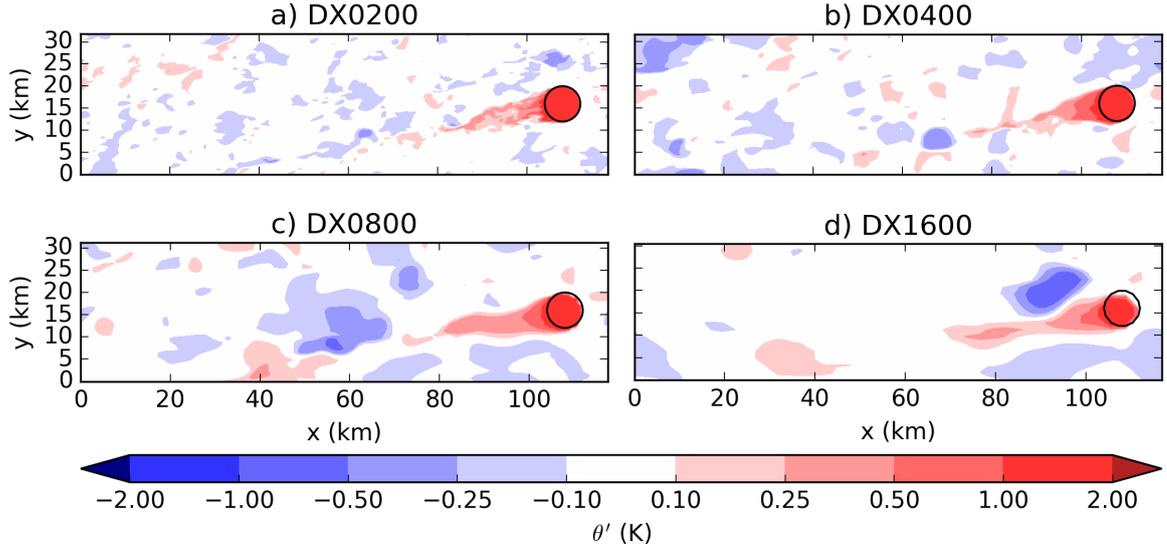


Figure 6.8 A snapshot of the potential temperature anomaly, θ'_{sfc} , at 12PM at the lowest model level (2 m) for a) the 200 m experiment, b) the 400 m experiment, c) the 800 m experiment, and d) the 1.6 km experiment. Here, θ' is defined as a deviation from the horizontal-mean at this time and height. In each panel, the island coastline is indicated with a black contour.

then produce more strongly precipitating convective plumes in the coarser resolution experiments. We see the effects of stronger precipitation appear as cold anomalies in Figure 6.8d. The following subsection takes a closer look at how precipitation changes with increasing grid length.

6.3.1.3 Changes to convective vigour: Precipitation Intensity

Interestingly, if we choose to use the daytime surface precipitation accumulation as our measure of convective vigour, this increases with grid length (Fig. 6.9). But it does not do so monotonically. There is a transition from precipitation primarily in the turbulent cloud patch and the cloud band at fine resolutions (100, and 200 m) to being in the convective elements in the environment, and on the edges of the cloud band at 400 and 800 m. Finally, by 1.6 km, there is no precipitation in the environment, but there is significant precipitation at the tail of the cloud band.

Intermittent pulses of convection which reach the cloud top inversion at the tail of the cloud band produce intense rain rates. In DX1600, the majority of the 10 mm daytime total falls in less than 10 minutes. A strong pulse of convection along the tail

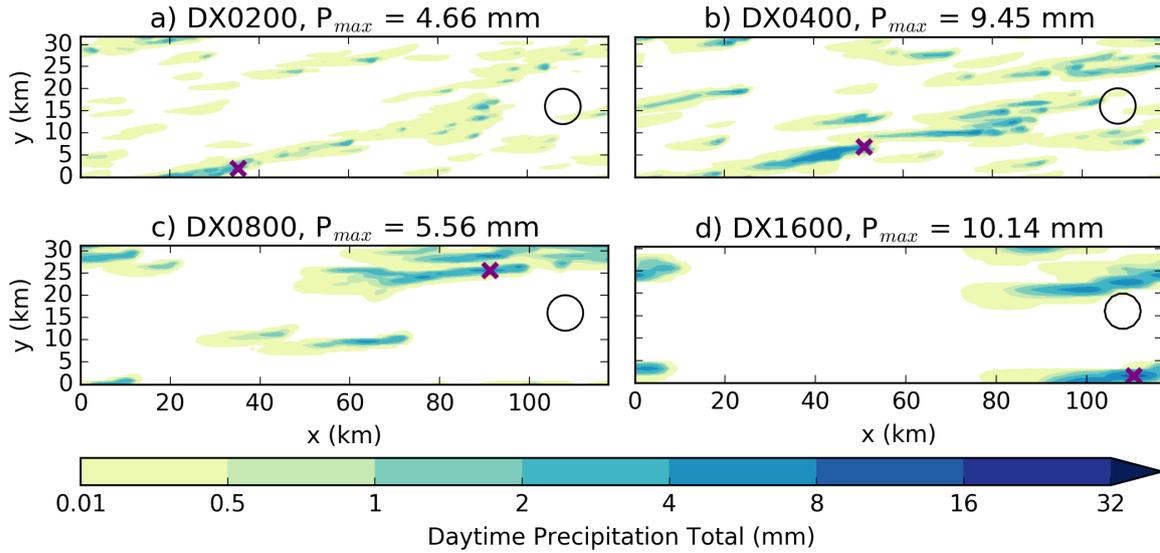


Figure 6.9 The total daytime (6AM to 6PM) precipitation for a) DX0200, b) DX0400, c) DX0800, and d) DX1600. The daytime rain total is noted in each panel’s title as P_{max} and the location of this maximum is marked on each panel as a purple ‘x’. The island coastline is marked with a black contour in each panel.

of the cloud band in DX1600 is responsible. This pulse consists of the deeper cloud and moves westward with the cloud layer flow and so the precipitation falls to the right (north) of the line of low-level convergence. Convective cells which track with the cloud layer flow and any convergence around the edges of cold pools helps to broaden the cloud band we see in Figure 6.4d and e.

Overall, even at grid lengths of 1.6 km the response to including a diurnal cycle on a small flat island is able to reproduce many aspects of the CT system. If a forecaster is concerned about the details of the CT system (e.g. will there be significant rain?), the DX1600 experiment will not suffice.

6.4 The Shallow Convection Parametrisation Perspective

6.4.1 The Convection Scheme

Here, a shallow convection scheme is introduced which we use on CT experiment DX1600. The UM uses a mass flux convection scheme based on [Gregory and Rowntree](#)

(1990) in pressure coordinates and with many modifications which include a CAPE closure for deep convection, changes to the coupling to the boundary layer scheme, and adaptive detrainment. A mass flux approach for parametrising convection is commonly used (Swann, 2001). A fundamental principle of a mass flux scheme is conservation of mass. However, energy is not necessarily conserved. This thesis focuses on the shallow scheme, but the basic mass flux principles are the same between shallow and deep schemes.

For the purposes of the scheme, ‘shallow’ refers to convection starting from the boundary layer and stopping below the freezing level, below 2.5 km, or below an inversion with descent or weak ascent above. In addition, there cannot be strong resolved-scale ascent above the cloud top. The shallow mass flux scheme was developed based on observations from the BOMEX campaign (which we describe in Section 4.3). All other convection originating from the boundary layer is considered ‘Deep’.

First, a convective diagnosis is performed to determine whether convection originating from the boundary layer is possible and whether it is shallow or deep. If the surface layer is unstable (i.e. there is a positive surface buoyancy flux), then the grid box is ‘convective’. An undilute parcel ascent that retains condensate is calculated from the LCL for convective grid boxes and is assumed to have properties from the top of the surface layer. The parcel ascent terminates where the parcel is no longer buoyant. This ascent is used to determine whether convection is shallow or deep. In the shallow mass flux scheme, it is assumed that all parametrised convection is ‘shallow’. In addition, wherever the resolved vertical velocity at the LCL is above 0.05 m s⁻¹, the scheme is suppressed.

The parametrisation links to the large-scale equations by contributing to the source/sink terms in the model formulation. These terms include any sub-grid transport terms due to convection. e.g. for θ :

$$S_{conv}^{\theta} = \left(\frac{\partial \theta}{\partial t} \right)_{conv} = g \frac{\partial \overline{\rho w' \theta'}}{\partial p} + \frac{L_c}{c_{pd} \pi} Q_{conv} \quad (6.2)$$

where S_{conv}^{θ} is the source or sink of θ due to the sub-grid convection scheme, $\overline{w' \theta'}$ is the sub-grid energy transport due to convection, and Q_{conv} represents the conversion

of water vapour into liquid or ice condensate due to convective processes, and L_c is the latent heat of the condensate (i.e. equal to L_v , the latent heat of vaporisation, when it is in liquid phase, and $L_s = L_v + L_f$, latent heat of sublimation, when in ice phase). Recall that π is the Exner pressure, given in Chapter 4 Section 4.2. A corresponding equation exists for each of the model prognostic fields, θ , q , q_{cl} , q_{cf} , u , or v . The sub-grid transports for a field, ϕ , can be written as:

$$\overline{w'\phi'} = \sigma_u \overline{w}_u (\overline{\phi}_u - \overline{\phi}_e) + \sigma_d \overline{w}_d (\overline{\phi}_d - \overline{\phi}_e) \quad (6.3)$$

where σ is the grid box area fraction of each draught type, ϕ is a model prognostic field, subscript u indicates updraughts, d downdraughts, and e the environment. This approach assumes $\sigma_u \ll \sigma_e$, $\sigma_d \ll \sigma_e$ and $|\overline{w}_e| \ll |\overline{w}_u|$ and $|\overline{w}_e| \ll |\overline{w}_d|$. In other words, that the area of the grid occupied by the convective up and downdraughts is small compared to the environment, and that the magnitude of the up and downdraughts are large relative to the large-scale vertical velocity. Equation 6.2 becomes:

$$\begin{aligned} \left(\frac{\partial \theta}{\partial t}\right)_{conv} &= g \frac{\partial \{ \rho \sigma_u \overline{w}_u (\overline{\theta}_u - \overline{\theta}_e) + \rho \sigma_d \overline{w}_d (\overline{\theta}_d - \overline{\theta}_e) \}}{\partial p} + \frac{L_c}{c_{pd}\pi} Q_{conv} \\ &= \frac{\partial}{\partial p} \left(\underbrace{M^{UD}\theta^{UD}}_{\text{updraught}} + \underbrace{M^{DD}\theta^{DD}}_{\text{downdraught}} \right) \\ &\quad - \underbrace{\frac{\partial}{\partial p} (M^{UD}\theta^E)}_{\text{compensating subsidence}} - \underbrace{\frac{\partial}{\partial p} (M^{DD}\theta^E)}_{\text{compensating ascent}} \\ &\quad + \underbrace{\frac{L_c}{c_{pd}\pi} Q_{conv}}_{\text{adiabatic processes}} \end{aligned} \quad (6.4)$$

where the mass fluxes are in Pa s^{-1} , and terms take the form $M^{UD} = g\rho_u\sigma_u\overline{w}_u$. Similar equations are applied for the other terms.

The layer is split into two; the parcel and the environment. The parcel (i.e. the updraught) is characterised by ascent and is usually saturated. The environment is characterised by descent and it is usually unsaturated. The parcel and the environment interact via entrainment (environmental air is transferred into the parcel), and

detrainment (parcel air is transferred into the environment). It is assumed that the updraught is exactly balanced by subsidence in the environment and the net vertical motion is zero.

A bulk parcel model is used in the scheme and it is derived by summation over an ensemble of convective cloud types. Exchanges between each parcel ‘ I ’ and the environment are treated through an entrainment rate, $E_I = \epsilon_I M_I^P$, a mixing detrainment rate, $N_I = \mu_I M_I^P$, and a forced detrainment rate, $D_I = \delta_I M_I^P$ following [Turner \(1963\)](#), where ϵ , μ , and δ are fractional mass entrainment and detrainment coefficients. The parcel equations are similar to those presented by [Yanai et al. \(1973\)](#), except here there are two detrainment processes: mixing detrainment (detrainment of parcel air through turbulent mixing at parcel edges); and forced detrainment.

It is assumed that there are no phase changes of water within the environment. There is no change in storage in the cloud and therefore changes to the total storage must be due to the environment:

$$M^E = -M^P \quad (6.5a)$$

$$\frac{\partial \phi^E}{\partial t} + \frac{\partial M^P \phi^E}{\partial p} = -M^P [\epsilon \phi^E - \mu \phi^N - \delta \phi^R] \quad (6.5b)$$

The fractional rates ϵ and μ vary with height, while δ is a constant profile based on assumed cloud radii for shallow convection. The ϵ and μ profiles depend on the parcel buoyancy excess, sub-cloud fluxes, and environmental relative humidity ([Grant and Brown, 1999](#); [Grant and Lock, 2004](#); [Derbyshire et al., 2004, 2011](#)).

For shallow convection, a cloud base mass flux closure based on [Grant \(2001\)](#) is used where the cloud base mass flux, m_b^{sh} is given by:

$$m_b = m_b^{sh} = 0.03w_* \quad (6.6)$$

where w_* is the sub-cloud convective velocity scale, the mass flux is strictly ρm_b . The sub-cloud velocity scale is calculated by the boundary layer scheme as:

$$w_* = (z_h \overline{w'b'})_{sfc}^{\frac{1}{3}} \quad (6.7)$$

where z_h is the boundary layer depth and $\overline{w'b'}_{sfc}$ is the surface buoyancy flux (as defined in Chapter 4, Section 4.2.2.1). The initial parcel properties are given by the environmental potential temperature and moisture plus a perturbation based on the environmental buoyancy gradient. The cloud base parcel is assumed to be at saturation.

A ‘grey zone’ shallow cumulus parametrisation option is available which computes a deep convective mass flux with a *CAPE* closure. The final ‘grey zone mass flux’ is a linear combination of the shallow cumulus scheme and the deep convective mass flux scheme where the weighting used to combine the two is grid length dependent with a similar formulation to the blended boundary layer scheme.

A complete description of all of the Convection schemes implemented in the UM can be found in UM Documentation Paper 027 ([Stratton et al., 2017](#)).

6.4.2 Cloud Census and Population Statistics

For this mass flux scheme, an ensemble of convective cloud types is represented as a single ‘bulk’ plume with the mean vertical profile of the ensemble. However, this operates under the assumption that the convective cloud elements are uniformly spread through a grid cell. Does organising the cloud field change these profiles?

Furthermore, the cloud spacing and organisation of the cloud field may have a role in the evolution of future or neighbouring clouds. The ‘humidity halos’, cloud particles, and hydrometeors left behind by predecessor clouds have been shown to play a role in the development of future clouds (e.g. [Perry and Hobbs, 1996](#); [Moser and Lasher-Trapp, 2018](#)). Could there be a change to the number, size, or internal properties of clouds between disorganised and organised convective times? Is this behaviour robust across our experiments varying grid spacing?

We consider how profiles of convection parametrisation-relevant cloud population statistics are changed by the introduction of a flat diurnally heated island and how they vary for the same problem with increasing grid lengths into the shallow convection grey zone regime.

6.4.2.1 Cloud Statistics

To begin an inspection into how the representation of convection changes with increasing grid spacing, we consider two summary statistics about the simulated clouds in each experiment - the scaled number of clouds, \hat{N}_{cld} and the effective cloud radius, r_{cld} . These are defined as follows:

$$\hat{N}_{cld} = \frac{N_{cld}}{N_{max}} \quad (6.8a)$$

$$r_{cld} = \sqrt{\frac{A_{cld}}{\pi}} \quad (6.8b)$$

where a grid cell is considered ‘cloudy’ if it has cloud water content, $q_{cl} > 0$. The number of clouds, N_{cld} , is divided by the number of grid points in the domain, $N_{max} = (L_x/dx)(L_y/dy)$ to get a scaled number of clouds, \hat{N}_{cld} . Where L_x is the length of the domain in the x direction, and L_y is the length of the domain in the y direction. This scaling follows that used in [Tobin et al. \(2012\)](#). The number of clouds is scaled so that a more direct comparison can be made between the experiments which vary grid spacing presented here. The cloudy area, A_{cld} , is defined as the number of cloudy grid cells in each identified cloud object multiplied by the area of a grid cell (Δx^2). The effective cloud radius, r_{cld} , is defined by assuming the cloud has a circular horizontal cross-section and using the cloud area, A_{cld} . Vertical profiles of the above quantities can be constructed by considering $q_{cl} > 0$ to describe regions at a given height that are cloudy, while the cloud footprint can be considered using lwp . There is no cloud ice, q_{ci} , to include because all of our clouds terminate below the freezing level (q_{ci} would need to be included in the total cloud water calculations for deeper convection).

Because the convection scheme represents an ensemble of clouds it is useful to have a sense of how many clouds there are and the depths over which they extend. Furthermore, the scheme is representing the mean of this ensemble, and entrainment and detrainment processes are modelled based on assumed relationships with the cloud radius, parcel buoyancy, and environmental RH. These three variables form the core focus of this section.

In addition to the above variables, we also consider the statistical and spatial distribution of cloud top heights through these experiments. As noted in the description of the convection scheme, the cloud top height is diagnosed by where an undilute condensate-loaded parcel ascent reaches its level of neutral buoyancy (LNB). The LNB is near the same height (about 3 km) in all of our experiments because of our imposed cloud-top inversion. Our analysis instead diagnoses cloud top height as the height of the highest model level at which there is $q_{cl} > 0$. Note that model levels become further apart with height on this stretched vertical grid; they are near 50 m apart near cloud base and 100 m apart near cloud top.

6.4.2.2 Scaled Cloud Count

Here we choose to consider two one-hour periods: an ‘organised’ convective period (11AM to 12PM) when the CT system is fully mature in all simulations and the island heating is at its peak; and a ‘disorganised’ convective period (11PM to 12AM) when the island hasn’t been heated for 5 hours, and the majority of organisation due to the island has dissipated in all simulations. We choose not to consider a ‘pre-dawn’ period as an example for the ‘disorganised’ convection case as there are still signs that DX1600 is spinning-up convective scale motions even at that late stage. DX1600 exhibits somewhat different behaviour compared to our other experiments. We will highlight these differences at the end of this subsection.

We find that \hat{N}_{cl} increases with increasing grid length (Figure 6.10). This result is repeated for both times considered. Furthermore, there is no major change comparing the organised to disorganised period of a given experiment (comparing Fig. 6.10a to b). We therefore focus on how these profiles change with grid length.

As grid length increases the whole profile increases and each increase in grid length results in a profile that is almost entirely higher than the previous experiment (Fig. 6.10). In Figure 6.10a, the DX1600 experiment is close to the DX0800 experiment at this height. We speculate that the grouping of DX0200 and DX0400 and the grouping of DX0800 and DX1600 arises due to the blended boundary layer scheme which is weighted toward the 3-D Smagorinsky turbulence in DX0200, is more evenly spilt in

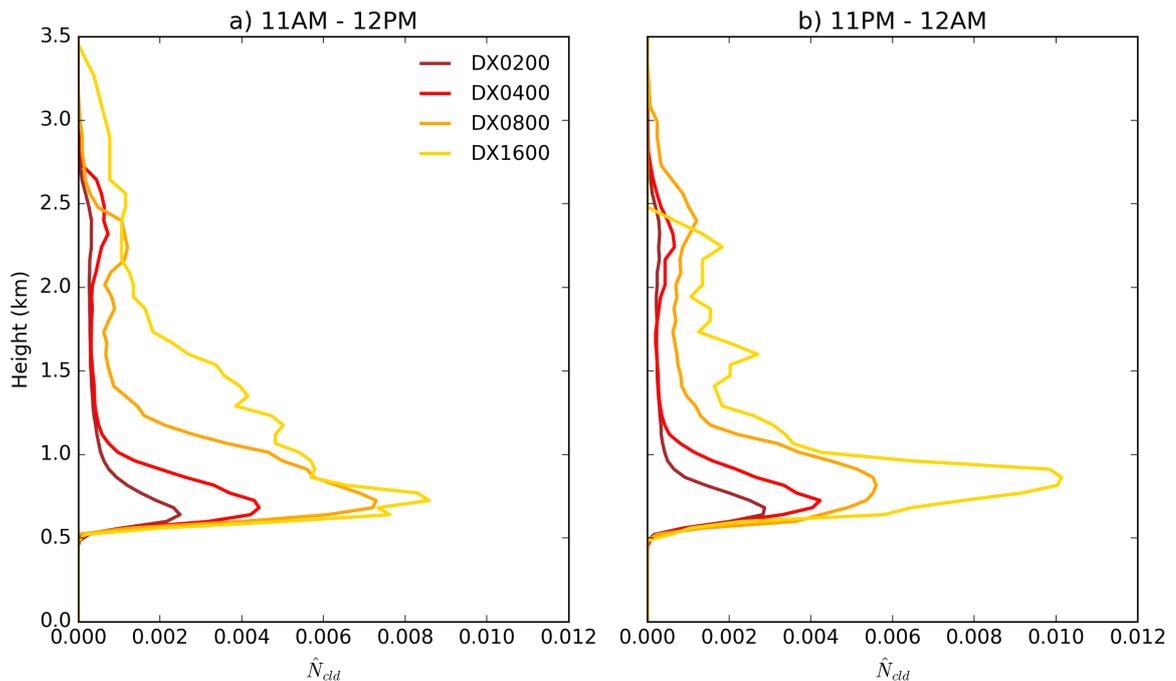


Figure 6.10 The hourly-mean vertical profile of the scaled cloud count, \hat{N}_{cld} for our suite of experiments at different grid lengths. These profiles are shown for a) an organised period during the peak island heating, and b) a disorganised period during the last hour of simulation.

DX0400, and is weighted more toward the 1-D Lock in DX0800/DX1600 (the 1-D and 3-D boundary layer and sub-grid turbulence schemes are described in Chapter 4, Section 4.2.2.1).

In the mid-cloud layer, DX0200 and DX0400 are nearly identical between 1.25 and 2 km (Fig. 6.10). We interpret this as more individual plumes are overshooting the LCL in DX0400, but the same amount are able to make it through the entirety of the cloud layer. In DX0800, more plumes are both overshooting the LCL and making it through the cloud layer, and even more do for DX1600.

Almost all of the clouds terminate between 2.5 and 3 km in DX0200 and DX0400 (Fig. 6.10). We remind the reader that a surface based parcel becomes negatively buoyant at 3 km in these experiments. This suggests that the ascending parcels in DX0200 and DX0400 are only weakly positively buoyant through the conditionally unstable cloud layer and do not reach the cloud top inversion with enough momentum to continue rising much farther. Whereas in DX0800, these ascending plumes sometimes terminate in that same height range (2.5 to 3 km in Fig. 6.10a), but sometimes go

deeper (Fig. 6.10b). This implies that updraughts which make it through the cloud layer are more strongly positively buoyant. This reasoning helps explain why more plumes apparently make it higher than 1 km and why \hat{N}_{cld} is higher in the mid-cloud layer. Similar behaviour is observed for DX1600 with some periods of more penetrating updraughts, and others with much more suppressed cloud top heights.

In general, there is also an increase in \hat{N}_{cld} with increasing grid spacing in the 2 to 2.75 km layer at both times (Fig. 6.10). This layer consists of actively rising updraughts, and the decaying cloud tops left over to spread out at the base of the inversion. We speculate that mixing detrainment is less effective with larger grid lengths, so these DX0800 updraughts can carry more cloud liquid to the cloud top inversion where it spreads out and lingers before evaporating.

In DX1600, the local maximum in \hat{N}_{cld} near cloud top is missing in both Figure 6.10a and b. In Figure 6.10a, too many plumes rise too high into the inversion where they all terminate at 3.5 km without spreading out or lingering. This gives the DX1600 profiles the same shape as the higher resolution experiments, but with higher \hat{N}_{cld} values. However, in b, a similar amount of clouds reach through the cloud layer, but they are not able to reach as high, terminating below 2.5 km.

We know that there are definite behavioural differences in the shallow convection at 1.6 km grid lengths. There appear to be two main “modes” - a very shallow boundary layer topping cumulus, and a ‘deep’ cumulus; the deep cumulus are clouds that would have continued to ascend to the tropopause had we not imposed a strong inversion which restricts our cloud top. This too-strong convective mode triggers more intermittently than the other experiments presented. We speculate that the intermittency in DX1600 is related to the too-strong convective motions occurring on too-large horizontal scales. There are not enough resolved eddies to dissipate the low-level instability steadily. It instead builds up to be released intermittently. More investigation is needed to fully understand this behaviour, perhaps with variants on the large-scale forcing or size of the domain.

6.4.2.3 Cloud Radius, Parcel Buoyancy, and Environmental RH

We have shown that there is no strong difference between the number of clouds during organised and disorganised periods during a CT day for all of our tested grid lengths. We now consider whether there are changes to the typical cloud radii and the range of cloud radii between our organised and disorganised periods, and with increasing grid length. These results are shown in Figure 6.11.

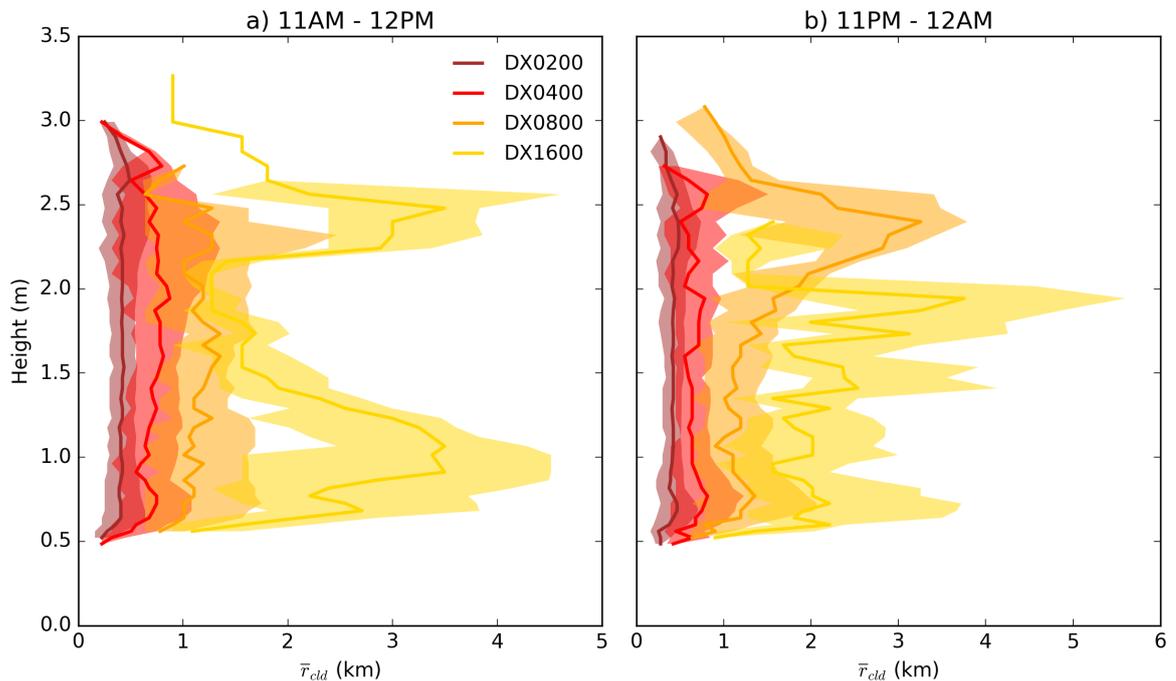


Figure 6.11 The hourly-mean vertical profile of the cloud radius assuming the cloud area represents that of a circle, r_{cld} , for our suite of experiments at different grid lengths. These profiles are shown for a) an organised period during the peak island heating (11AM to 12PM), and b) a disorganised period during the last hour of simulation (11PM to 12AM). The solid line represents the median cloud radius, and the shading is the interquartile range.

The overall story is much the same for cloud radius, r_{cld} , as for \hat{N}_{cld} . The median r_{cld} increases for increasing grid length. This makes physical sense and is an expected result as the grid length limits how narrow a cloud can get. For example, the smallest possible cloud radius in the DX1600 experiment is about 900 m (i.e. a cloud occupying one grid cell). There are however some interesting changes to the profile of cloud radius with increasing grid length, and the differences between the organised period in Figure

6.11a and the disorganised period in b.

In DX0200 and DX0400, the cloud radius is nearly constant with height from cloud base through cloud top during both time periods (Fig. 6.11a and b). This implies that the very shallow cumulus which only just overshoot the LCL have about the same size as the deeper cumuli which extend through the entire cloud layer. The cloud radius in DX0800 is also nearly constant with height in the organised period (Fig. 6.11a). The vertically-averaged cloud radius for the organised period is near 0.4 km in DX0200, 0.8 km in DX0400, and 1.3 km in DX0800. These numbers are interesting as it appears that in each experiment, the median cloud radius settles on approximately $2\Delta x$ (slightly less for DX0800, and DX1600 varies significantly with height). Because there is more variability and fewer clouds in DX0800 and DX1600 we also consider the mean cloud radius over the daytime period (model hour 6 through 18), and find that DX1600 does not settle as closely to typical radii of $2\Delta x$. Instead, the vertical-mean r_{cld} in DX1600 is near 2.6 km (ranging from a low-level peak of 3.5 km to a minimum near 2 km height of 1.3 km in Fig. 6.12). For $\overline{r_{cld}} = 2\Delta x$, this implies a cloud occupying 12-13 grid cells in the horizontal.

Profiles of r_{cld} look the same between the organised period, the disorganised period, and the daytime period in DX0200 and DX0400, while DX0800 and DX1600 are more statistically converged onto smoother profiles in Figure 6.12. The cloud radii in DX0800 and DX1600 have pronounced maxima near cloud base and cloud top where updraughts are presumably encountering (and penetrating through) the more stable layer near 1 km height and terminating at the cloud top inversion. While the clouds are bound to be too big in DX1600 due to grid size restrictions on the lower limit, the clouds in DX0200 are still too big compared to other LES studies (e.g. Kirshbaum and Grant, 2012, report a mean radii of 0.15 to 0.5 km in BOMEX- and RICO-like experiments which we believe to be similar to the current experiments). Perhaps there is more horizontal mixing in our experiments which is enhancing the moistening around the clouds and supporting their horizontal growth. More analysis is required to fully explain this result.

For parcel buoyancy and environmental RH, we continue to focus on the daytime period to ensure robust statistics (particularly for the DX1600 experiment). We define “updraughts” as cloudy rising air parcels, i.e. anywhere $q_{cl} > 0$ and $w > 0$, and also

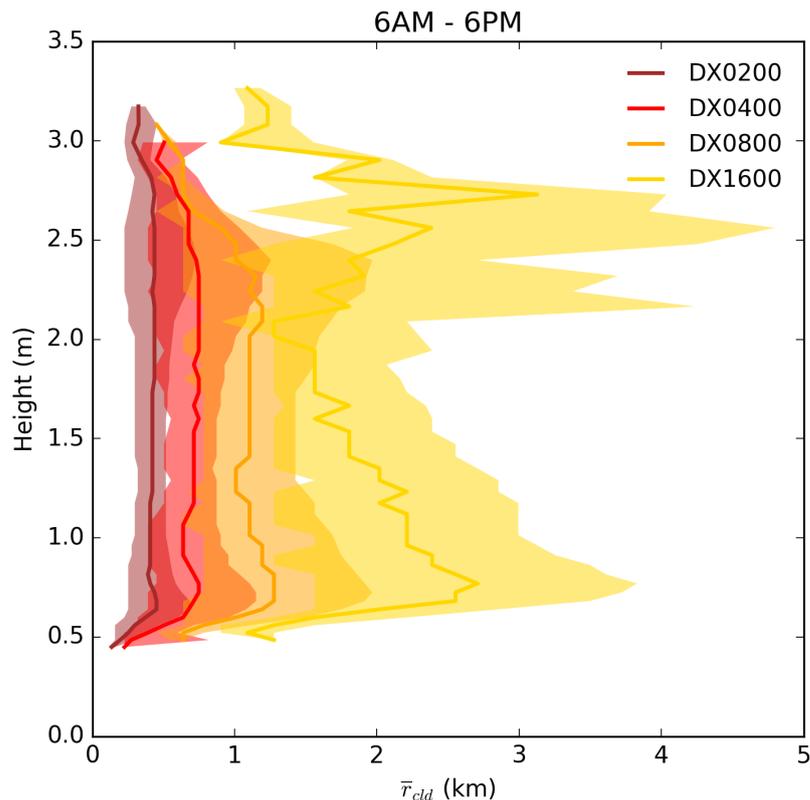


Figure 6.12 As in 6.11, but for the whole daytime period from model hour 6 to 18 rather than for one-hour periods.

define “cores” as buoyant cloudy updraughts, which have the additional criteria that $\theta'_v > 0$ following Siebesma and Cuijpers (1995) and Grant and Brown (1999). Then the mean buoyancy excess of all updraughts is given by, $\langle \theta_v^{u,c} \rangle$ where the angle brackets $\langle \rangle$ denote the mean over all updraughts, superscript u is for updraughts, and superscript c is for cores. The environmental RH is considered through the total water excess of updraughts and cores, $\langle q_T^{u,c} \rangle$. Note that there is no requirement that the updraughts are spatially coherent.

We find that the profiles of buoyancy excess all broadly have the same magnitude across all experiments (Fig. 6.13a). The peak core buoyancy excesses reported from LES for variants of the BOMEX shallow cumulus experiment are near 0.4 K for cloud tops reaching 2 km - about half of the value reported here (0.8 K). Updraughts are weakly buoyant near cloud base, and generally grow more buoyant through the cloud layer to about 2 km height. Above this level, the updraughts become less buoyant with height until they become negatively buoyant above about 2.5 km. Interesting

differences emerge between the experiments near cloud base. The height at which the updraughts become more strongly buoyant increases with increasing grid length (dashed lines in Fig. 6.13a). Meanwhile, while the cores all become strongly buoyant just above cloud base (i.e. above $z > 0.75$ km). In a layer near 1.0 km height, the cores of DX0800 and DX1600 are less buoyant than above or below, and are less buoyant than DX0200 and DX0400 (solid lines in Fig. 6.13a). We attribute this slight decreased buoyancy in these two experiments to the development of an inversion at this height during their respective spin-up simulations.

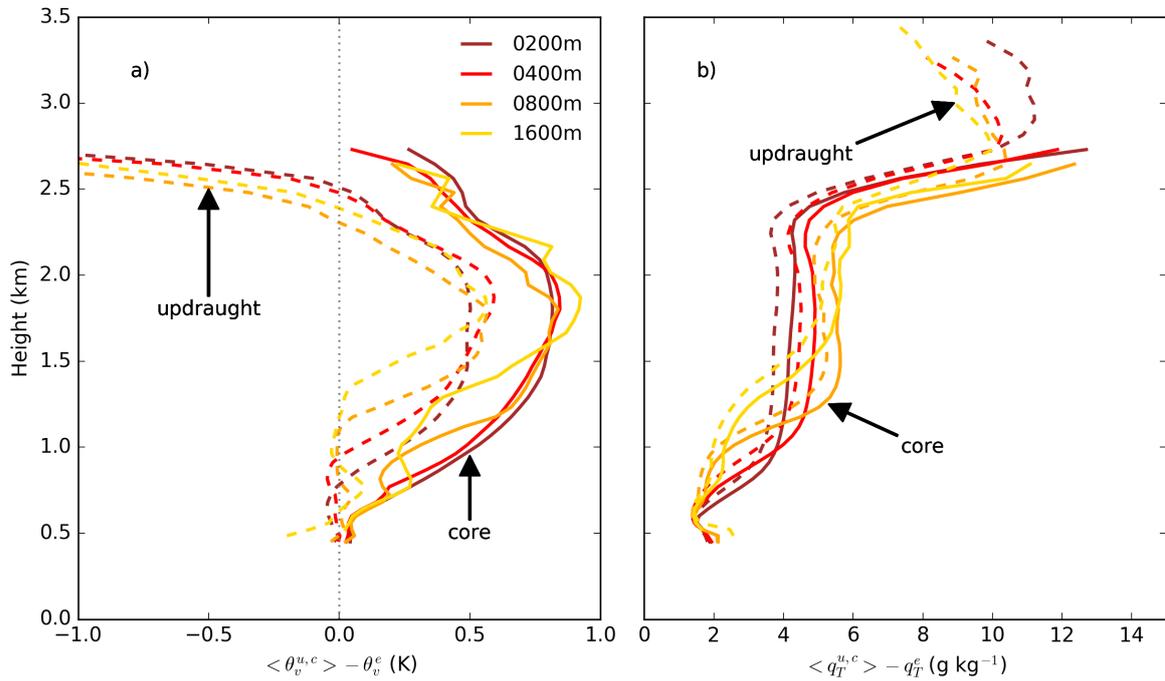


Figure 6.13 The daytime mean updraught (dashed line) and core (solid line) profiles of a) buoyancy excess measured in terms of θ_v , and b) total water excess. These profiles are computed for each experiment and coloured accordingly. See text for definition of “updraught” and “core”.

In all experiments the updraughts are roughly 0.5 g kg^{-1} drier than the cores (difference between solid and dashed lines in Fig. 6.13b). In addition, they all start the ascent with a q_T excess near 1 g kg^{-1} . The profiles of the q_T excess exhibit clearer differences among the experiments than the buoyancy excess profiles. These differences in q_T excess highlight the effect of differences in the vertical structure that come with the different grid spacing - the profiles all grow in q_T excess, but this occurs over a deeper and lower layer in DX0200 and DX0400 than in DX0800 and DX1600. The q_T

excess is then maintained, and it is larger for larger grid spacing, until updraughts reach about 2.5 km. Above this height, the q_T excess roughly doubles and the cores terminate near 2.75 km. The remaining updraughts have slowly decaying q_T excess with height above 2.75 km but then the updraughts terminate before reaching 3.5 km (dashed lines in Fig. 6.13b). This sharp increase at cloud top is due to updraughts and cores punching into the dry layer aloft (see initial conditions in Fig. 6.3). Again, because of the more stable layer that develops near 1 km height in DX0800 and DX1600, fewer clouds penetrate to the cloud top inversion. The cloud layer is therefore drier between 1.5 and 2.5 km height because of the prescribed subsidence forcing. Any updraughts or cores that do occur in this layer therefore have higher q_T excesses. This implies q_T in the clouds is roughly unchanged by model grid spacing, but rather q_T excess is controlled by the environmental relative humidity.

6.4.2.4 Cloud Top Height

We next consider the time evolution of cloud top heights in our domain. This analysis defines cloud top height as the highest model level on which q_{cl} is greater than 0. Cloud objects are first identified as any continuous area with liquid water path greater than 0 g m^{-2} . The maximum cloud top height of each cloud object is found and retained. This is done on ten-minute model output. The time series of the statistical distribution of cloud top heights in each experiment is reported in Figure 6.14.

Firstly, we see that our measures to maintain a cloud layer within our inversion were successful (Fig. 5.11). We see that the statistical distribution of cloud top heights is mostly unchanged through the day in DX0200 (Fig. 5.11a). However, a small signal for higher cloud tops in the daytime appears for DX0400 and DX0800 (interquartile range in Fig. 5.11b and c). The convective onset in DX1600 has a very different pattern and comes in two stages: first, the very shallow cumulus develops and slowly gets deeper; second, there is a regime shift around hour 6 in which deeper convection starts to form (Fig. 6.14d). This shift occurs around the time the island is first heated, and we suspect that this is by coincidence because it occurs near the same time in this experiment's spin-up simulation (not shown). There are bursts of deeper convection roughly every six hours which have cloud tops higher than any of the other

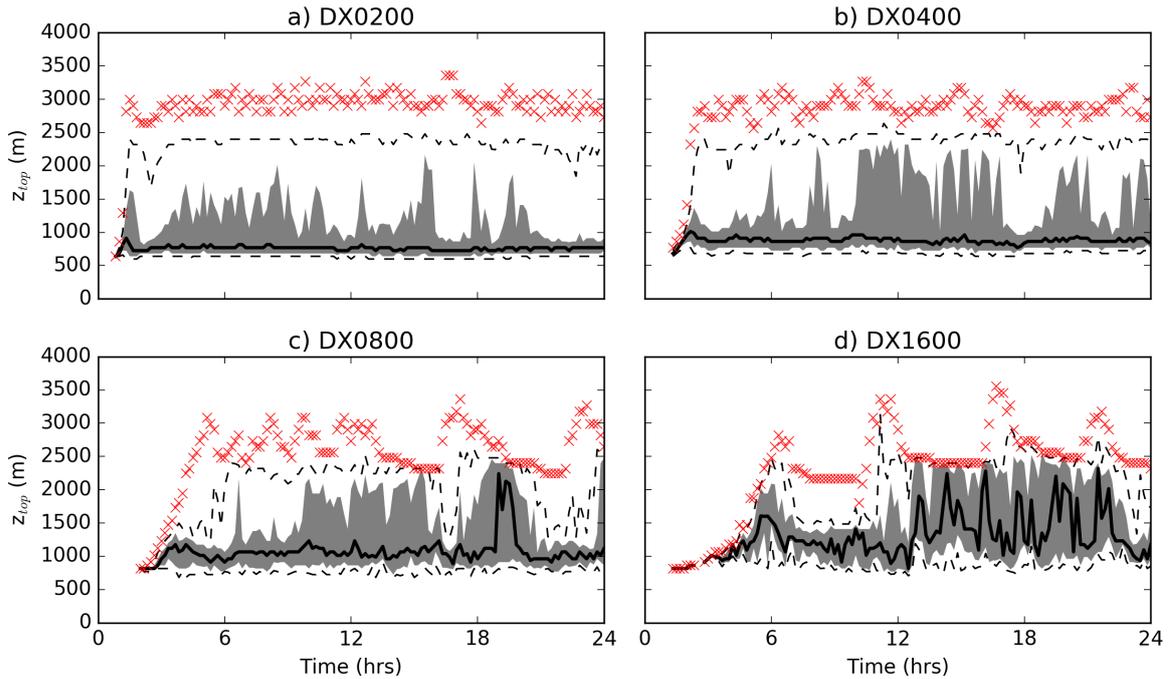


Figure 6.14 Time series of the cloud top height distributions for a) the 200 m experiment, b) the 400 m experiment, c) the 800 m experiment, and d) the 1.6 km experiment. The dashed black lines represent the 10th and 90th percentile, the grey shading is the interquartile range, the solid black line is the median, and the red ‘x’s mark the maximum at that time.

experiments. At any given time, there are typically about 350 clouds in DX0200, 100 clouds in DX0400, 30 clouds in DX0800, and about 10 clouds in DX1600. The statistics are therefore more robust for the finer resolution experiments.

6.4.3 Parametrised Results

The body of work presented in the current chapter all supports the idea that the explicit representation of shallow convection for grid lengths of 1.6 km is problematic. Furthermore, this chapter repeats the conclusion that shallow convection is not adequately explicitly represented until grid lengths decrease to 200 m where a crude representation of shallow convection can be made. This result is repeated from Chapter 4 and numerous times in literature (e.g. Bryan et al., 2003, and others). Yet, the current implementation of the UKV does not employ a shallow convection parametrisation scheme.

Here we begin a discourse on the use of a shallow parametrisation scheme in the

idealised UM with 1.6 km grid lengths. This is comparable to the grid lengths used in the UKV (1.5 km) and uses similar sub-grid turbulence and cloud formulations (Bush et al., 2019). We use the standard settings of an existing shallow convection scheme built for the UM as described in Section 6.4.1 above. Without performing a spin-up simulation we repeat our 1.6 km CT experiment starting from the same initial conditions and label the new experiment “DX1600S”. Note that this section is not intended to be a thorough examination of a shallow convection scheme, but is solely meant to determine whether introducing a shallow convection scheme modifies the CT system in a meaningful way. The following analysis of the parametrised experiment is intentionally superficial and serves to signpost the challenges faced in current and future work in parametrisation development for shallow convection in the grey zone.

Keeping all else the same, we then impose the equilibrium conditions and the shallow convection scheme into our 1600 m grid spacing experiment for the CT domain with an island (a new spin-up simulation is not performed). Presented in Figure 6.15 is an overview of the CT system.

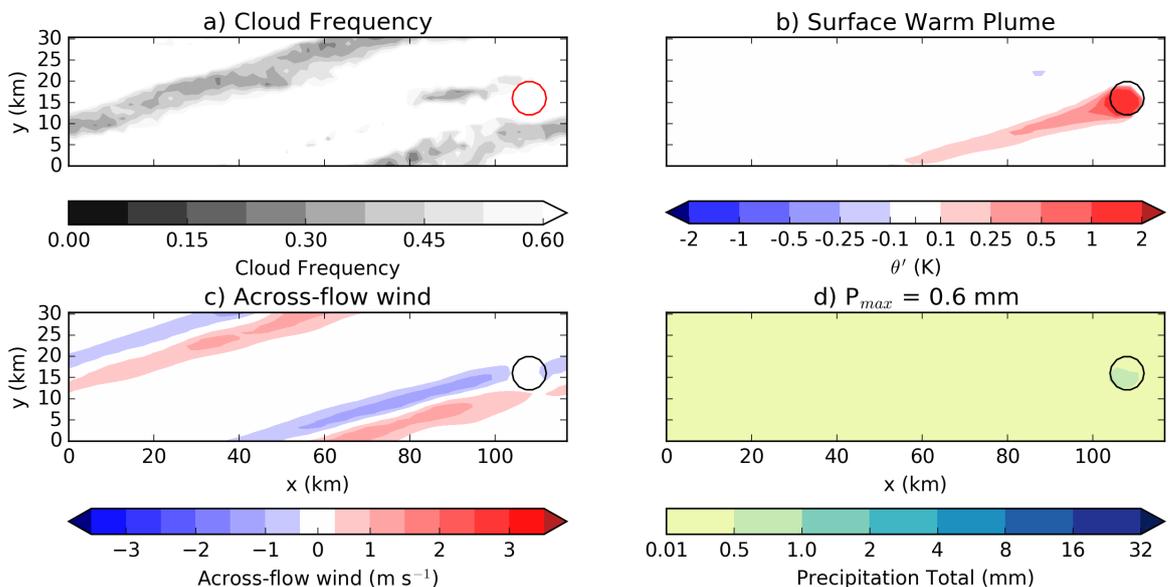


Figure 6.15 An overview of the CT system in the parametrised 1600 m simulation. a) The daytime cloud frequency, b) the equivalent potential temperature anomaly on the lowest model level (2 m), c) the 10 m n component of the wind, and d) the total accumulated precipitation.

The cloud frequency field is higher everywhere and a cloud band is harder to pick out from the ambient cloud frequency (white areas in Figure 6.15a). We suspect that

this is because we prescribe the surface sensible and latent heat fluxes (proportional to the surface buoyancy fluxes) in our experiment to be ≥ 0 everywhere and the convection scheme’s triggering condition is that the surface buoyancy flux must be positive. This also highlights a subtlety in using cloud frequency as a measure here. This diagnostic of cloud frequency assumes that if there is any cloud in a grid box that the grid box is cloudy. In future, including the shallow convective cloud area fraction in our calculations would provide a better estimate of the cloud frequency - potentially highlighting the cloud band better.

We find that the warm plume is still present in the parametrised experiment (Fig. 6.15b). This plume extends farther downwind in DX1600S than in DX1600 (see Fig. 6.8). The warm plume fails to narrow with distance downwind from the island, and has lower potential temperature anomalies over the island itself in DX1600S than in any of the other experiments in this chapter.

We suspect that the lower potential temperature anomalies are driving weaker circulation generation over the island. We can see this weaker circulation in Figure 6.15c where the n -wind component in DX1600S is near 1.5 m s^{-1} and near 2.5 m s^{-1} without the parametrisation in DX1600.

Every grid cell contains weakly precipitating parametrised convective cloud during this daytime period (Fig. 6.15d). While there is light precipitation everywhere in DX1600S, the heavy precipitation in the cloud band that we saw in DX1600 does not form with this parametrisation. We expect some light precipitation in the cloud band based on the 100 m experiment from Chapter 5 which we take as the truth, but the precipitation in DX1600 was too strong, too far downwind, and therefore having a negative impact of the low-level CT circulation structure there.

The cloud tops in this experiment all go to the base of the cloud-top inversion (Fig. 6.16b). This is because the parametrisation sets the cloud depth at the level of neutral buoyancy (LNB) for an undilute surface-based parcel. Therefore any convecting grid cell will be given a cloud top near the cloud-top inversion by the scheme. The reader should be reminded that the scheme treats a *bulk convective plume* and therefore if all of the cloud is convective, all of the spread of the statistical distribution of cloud top heights is averaged into this one plume. The scheme is representing the role of

those very shallow clouds, but they are occurring sub-grid and are cannot separately be represented in this diagnostic. That being said, the scheme is still triggering convection nearly everywhere and with the same cloud depth. This is clearly not realistic and might be remedied with a different diagnostic for cloud top (i.e. a dilute parcel ascent) or with a representation of stochastic behaviour.

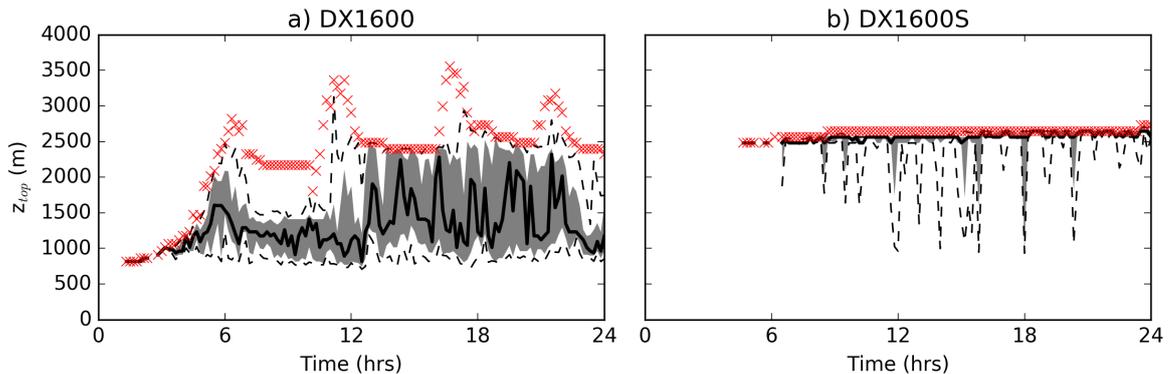


Figure 6.16 Time series of the cloud top height distributions for a) the 1600 m experiment without parametrisation, and b) the 1600 m experiment with a shallow convection parametrisation. Symbology is as in Figure 6.14.

Also of note is that the resolved-scale variability is almost completely erased with this parametrisation 6.15. Where DX1600 attempts to allow some of the larger boundary layer eddies to emerge, only the forced circulation is present for DX1600S, and even that is weaker than in DX1600.

We conclude that using this parametrisation does demonstrate some benefits over having no scheme at all - even without a careful choice of parameters, and neglecting potentially important features which can be included in the scheme such as convective momentum transport. A more thorough investigation of this existing shallow parametrisation scheme might reveal which components of the scheme are performing well compared to our well-resolved experiment - and which are not.

The loss of resolved structure presented here is common to both recent stochastic shallow cumulus mass flux schemes (e.g. Nie and Kuang, 2012), and “scale-aware” schemes (e.g. Brast et al., 2012). However, those studies were performed with homogeneous surface conditions. Perhaps introducing stochastic or scale-aware variants to the mass flux scheme here might return some of the strength of the CT circulation in DX1600S. This thinking is in line with the consensus of the break-out group on

Parametrising Grey Zone Convection in the Future of Convection Workshop: Challenges and Progress, July 2019.

Furthermore, of particular interest to the ParaCon group (currently developing new parametrisation schemes in the UM) and the convection community in general, is how this scheme behaves in the grey zone. A grey zone implementation of the shallow convection scheme which scales the behaviour between the shallow and deep scheme depending on grid length exists, but it is unclear what the nature of the grey zone formulation should be (i.e. should a scheme simply be “turned down” or behave somewhat differently?).

6.5 Summary and Conclusions

When simulating shallow convection at grey zone resolutions, the vertical convective motions are projected onto the resolved scale. This projection results in convective plumes that are too broad and too intense, thereby making the vertical cloud structure very sensitive to how well the vertical structure in the thermodynamic profile is resolved. The sensitivity to thermodynamic vertical structure is a result repeated from experiments using increasing grid length with the BOMEX test case in Chapter 4, Section 4.5.1.

For grid lengths of 800 m or greater we demonstrate that the cloudy rising plumes can broadly be separated into two categories: those which only just exit the boundary layer forming a layer of very shallow cumulus, and those which can traverse the entire depth of the troposphere with cloud tops near the tropopause (or a significant inversion). However, as grid length increases an inversion that would cause well-resolved updraughts to terminate might not be sufficiently strong to do so for poorly-resolved cases and we begin to move toward the territory of ‘grid point storms’ and their associated numerical issues.

To study CTs at coarser resolution, we increase the strength of the cloud top inversion used in Chapter 5’s experiments. We then conduct spin-up simulations for grid lengths of 200 m (DX0200), 400 m (DX0400), 800 m (DX0800), and 1.6 km (DX1600) as described in Chapter 4 Section 4.6.2, with modifications and experiment numbers

summarised in section 6.2. These spin-up simulations are also performed at 100 m for comparison to the environment. We show that in doing this, we largely preserve the overall vertical thermodynamic structure with some minor modification remaining unavoidable for DX0800 and DX1600. The resulting equilibrium profiles are used in CT simulations and show that there are systematic changes to the representation of CT as the grid length is increased.

Overall, the general features of the CT system are largely preserved as grid lengths are increased from 200 m to 1.6 km. The turbulent cloud patch which we described in our well-resolved experiments in Chapter 5 Section 5.2.2.1 disappears when grid lengths are increased beyond 400 m. In addition, the CT circulation grows stronger as grid length increases. The cloud band convection is then too intense and associated with strong precipitation along the end of the cloud band.

Next we consider this chapter's experiments from the perspective of a shallow convection scheme. We note that the shallow mass flux scheme implemented in the UM is based on an ensemble of convective clouds, and has parcel exchanges which are related to the cloud radius, parcel buoyancy, and environmental relative humidity. These same properties of the clouds in our experiments are then examined. We find that the vertical profile of the scaled number of clouds increases with increasing grid length as more clouds form overall and more penetrative clouds (reaching the cloud top inversion) form. We find that the cloud radius increases with grid spacing across all of our experiments, but remains roughly $2\Delta x$.

The distribution of cloud top heights highlights increased intermittency in the amount of the deeper cloud as grid length increases. In DX0800, distinct periods where predominantly very shallow cloud start to occur, and in DX1600 those periods become full breaks. Furthermore we find that after the breaks in DX1600, the deeper cloud briefly has tops higher than in the other experiments. We speculate that the intermittency and deeper clouds in DX1600 are related to the stronger convective motions which lead to a charge-discharge of low-level instability and therefore an oscillation in the cloud field.

Finally, the mean updraught buoyancy and humidity are considered. The results for DX1600 exhibit the most buoyant cloud cores and has clouds that reach highest

into the inversion and so become most negatively buoyant. DX1600 also has the largest updraught excess in q_T through the whole cloud layer. We suggest that in these poorly resolved cases, updraughts are less diluted as the mixing motions along their ascent cannot be resolved explicitly. We therefore think these experiments will benefit the most from the introduction of a shallow convection scheme.

Our “first look” at a shallow convection scheme is mixed. We find that the scheme successfully reduces the intensity of the convection and prevents strong precipitation from developing. However, the scheme also produces lightly precipitating parametrised convection in all grid cells and has a weaker CT system overall than in any of the explicit cases. While precipitation is not key to CT formation, heavy precipitation is one potential impact from CTs. Furthermore, evaporatively-driven cold-pool dynamics related to heavy precipitation interact with the CT circulation to result in different behaviour than in non-precipitating CTs. The ability to capture whether a CT will be strongly precipitating is therefore of interest. These results indicate that using a shallow convection scheme has the potential to improve NWP at $> \mathcal{O}(1 \text{ km})$ grid spacings. We speculate that further improvements are possible using additional modifications toward physically-based, stochastic, and scale-aware convection parametrisation.

Chapter 7:

Conclusions and Future Work

7.1 Summary of Work

This thesis has investigated the occurrence of cloud trails (CTs) at Bermuda, the modulation of that occurrence by the larger-scale environment, and the representation of CTs in Numerical Weather Prediction models (NWP). CTs occur downwind of small islands (with areas less than 100 km²) around the world in response to a strong land-sea contrast in surface heating and background flow. The land-sea contrast drives an across-flow circulation which organises cumulus into a cloud band oriented parallel to the low-level flow and anchored to the island. The CT system can be viewed as a modified sea-breeze system. The background flow advects the along-flow component of the circulation over and downwind of the island. The windward cell becomes phase-locked to the island, and the leeward cell carries on downwind and weakens. Current implementations of NWP, and implementations in development, employ sufficiently high horizontal spatial resolution to start representing these small islands, and therefore the CT system is being treated in some manner. Furthermore, at these resolutions, many modelling centres are not using convection parametrisation schemes (which have traditionally been required to represent sub-grid scale convection and maintain numerical stability). It is therefore increasingly important to have a more complete understanding of the CT system and its representation in poorly-resolved situations in which a parametrisation scheme is not used. From this background, this thesis asks the following questions:

1. How does CT occurrence vary at Bermuda (a small isolated island known to have CTs), including diurnally, seasonally, and annually; and how do larger-scale environmental conditions modulate this occurrence?
2. How does the strength and structure of CTs respond to changes in island surface fluxes, environmental RH, and background flow for simulations in which shallow convection is well-resolved?
3. Is the numerical representation of CTs degraded for explicit convection at ‘grey zone’ grid lengths; which properties are restored via parametrisation?

The following section will address the above questions based on the work detailed in Chapters 3 through 6.

7.1.1 Cloud Trail Climatology at Bermuda

In order to explore CT occurrence, a novel algorithm for identifying CTs based on visible-channel satellite imagery and 10 m wind direction is introduced. This algorithm returns a classification of either cloud trail (CT) if there is a cloud band downwind of the island; non-trail (NT) if there is no cloud band downwind of the island; or obscured (OB) if large-scale cloud obscures the island from view, or makes it impossible to determine the presence of a CT. Half-hourly satellite imagery from the warm seasons (May through October, inclusive) of five years from 2012 to 2016 inclusive are classified, for a total of 16,400 classifications. Compared to a manual classification of imagery for the warm season of 2012 (taken to be the “truth”), the algorithm successfully captures the annual and diurnal cycle of CT occurrence. The algorithm has a tendency to result in several CT periods per day where a manual classifier might identify one continuous CT period. Despite this, the comparison between algorithm and manual classifications yields a hit rate of 67 %, false alarm rate of 15.6 % and a bias score of 1.05, where a bias of 1.00 is a perfectly unbiased classification.

CT occurrence is found to peak in summertime and in the afternoon. This is consistent with our understanding that CT formation is linked to the strength of island surface heating. The mean wind speed for CT scenes is near the mean for all warm

season times in our five-year period. The wind speed is lower for NT scenes and higher for OB scenes. In addition, CTs occur most frequently for westerly and southwesterly winds. Winds from these directions are both long-axis-parallel for Bermuda and are associated with more humid environments. The combination of wind speed and direction controls the residence time of flow over the island heat source, the strength of the land-sea buoyancy gradient, and therefore the strength of the CT system. A more quantitative evaluation linking wind speed, surface heating, and CT strength does not follow from this analysis because of insufficient downstream observations to assess the strength of the CT, and absent surface flux observations to link the CT strength back to the forcing.

Using the resulting classifications, we construct June-July-August (JJA) composite morning (0900 LT) radiosonde anomalies of potential temperature, θ , and relative humidity, RH, for Bermuda. These composites show that the typical CT day has a profile that is warmer than normal, more humid than normal in the boundary layer with near-normal, or lower than normal RH in the free-troposphere. In contrast, on NT days the profile is slightly cooler than normal with a drier than normal boundary layer, and similar free-tropospheric RH. OB days are much more humid than normal in the boundary layer and free-troposphere, with a more stable than normal boundary layer, and less stable than normal free-troposphere.

These differences between the classifications are repeated in composites of ERA-interim fields over the Northwestern Atlantic centred on Bermuda for JJA at 0900 LT. The re-analysis fields reveal subtle changes in the prevailing Bermuda-Azores high pressure system that are associated with changes to the low-level flow and therefore changes in the low-level RH. The re-analysis helps to identify the large-scale patterns in which CT occurrence is higher. Furthermore, the re-analysis reveals that both CT and NT days are nearly identical in terms of the mean vertical velocity profiles with weak 500-hPa subsidence over the island. OB days are associated with strong ascent at 500-hPa.

7.1.2 Idealised Simulations: BOMEX and Shallow Convection

Next, we take the idealised version of the Unified Model (the idealised UM) and assess its representation of shallow convection. The idealised UM has the functionality to employ the same parametrisation schemes for boundary layer turbulence, cloud, large scale rain, and convection as implementations of the UM which are in use at meteorological centres. In addition, simplifying idealisations can be employed to test aspects of simulated phenomena for research and model development purposes. We choose the Barbados Oceanographic and Meteorological Experiment (BOMEX), a non-precipitating trade cumulus test case, to evaluate the representation of shallow convection in the idealised UM. The BOMEX case has been well-evaluated with Large Eddy Simulations (LES) for horizontal grid lengths of 100 m, and this grid length is also used here. This case includes a rigid lid at 3 km height, bi-periodic lateral boundary conditions, prescribed surface fluxes, large-scale drying and warming due to subsidence, drying tendency due to idealised horizontal advection, imposed radiative cooling, and geostrophic wind forcing.

Of primary interest are the macro-scale cloud properties and the maintenance of large-scale equilibrium. We find that both of these aspects hold for the BOMEX test case. The evolution of total cloud cover, horizontally averaged liquid water path (LWP), and horizontally averaged turbulence kinetic energy (TKE) over the experiment closely follows that of the LES intercomparison. Similarly, there is only small drift away from the initial conditions in the horizontally averaged θ , specific humidity (q_v), and wind profiles. While the idealised UM cloud base matches the LES intercomparison, the cloud layer has less horizontally averaged cloud liquid water than the intercomparison. These differences in the cloud layer occur over a layer of higher inter-model disagreement, and we attribute them to differences in the turbulence and cloud schemes between the idealised UM and traditional LES.

BOMEX is then repeated using larger grid lengths. These experiments start with 50 m grid length, and this is doubled until the grid length reaches 800 m. We demonstrate that the idealised UM has not yet converged for 100 m grid lengths as significant changes in macro-scale cloud continue to occur when grid length is decreased to 50

m. For 200 m grid length updraughts start to approach the rigid model lid, and for 400 and 800 m grid lengths, they strongly interact with the model lid. Eddies generally become stronger for increasing grid length and this is highlighted in time series of horizontal-mean vertically integrated TKE. Relatedly, the model also more strongly drifts away from the initial conditions with increasing grid length. We conclude that 100 m grid spacing is able to represent shallow convection reasonably well given that the main change for 50 m grid spacing is an increase in total cloud cover.

To translate the BOMEX case into something less idealised, we impose a rigid model lid at 40 km over 140 stretched vertical levels and show that this change does introduce more significant drift away from the initial conditions. This is attributed to the increase in vertical resolution in the surface layer which impacts how heat and moisture are redistributed vertically. Given that we do not have a field study from which to derive well-constrained initial condition and idealised forcing for our intended CT simulations we outline a method to derive these. We use a simplified radiosonde ascent for initial conditions and modified BOMEX forcing profiles on a small, sea-only domain with interactive surface fluxes for ten days and refer to this as a “spin-up” simulation. The spin-up simulations are performed on this small domain to limit computational costs. We show that after six days, the vertically integrated heat and moisture budgets are in an equilibrium comparable to the BOMEX experiment. We therefore use the horizontal-mean environmental profiles and surface fluxes from this method to prescribe as initial conditions and constant sea surface forcing in our later CT simulations.

7.1.3 Idealised Simulations: The Well-Resolved CT System

For the well-resolved case, we use the initial conditions and surface fluxes from the spin-up simulation in the idealised UM with 100 m grid spacing. Then we run a simulation using bi-periodic boundary conditions over a 118.4×32 km channel domain with a 50 km^2 island centred near the right-hand boundary (this puts the island on the upwind side of the domain given imposed easterly flow). This forms our Control experiment. First we show that the UM can produce CT like behaviour in this configuration and with these prescribed conditions. We further show that the cloud

band, warm plume, and circulation of the CT system as described from observations and simulations in literature are reproduced in our control experiment.

Our control experiment reveals further structure to the CT region. We identify two main regions: a turbulent plume where buoyant free convection forms short-lived clouds that are advected downwind, and a cloud band where the convergent low-level across-flow component of the CT circulation lifts low-level air to saturation. A transition region between these two parts of the CT has less cloud activity and we propose that this is related to increased boundary layer stability where the warm plume over-runs the cooler surface oceanic air. Over the island and in the turbulent plume, several convective rolls form, and more investigation is required to understand their role in the CT system and their relation to upwind boundary layer eddies, and the large-scale environmental profiles.

Chapter 3, left unanswered questions related to the relationship between CT strength and the strength of the island thermal forcing and the background wind speed. We also wanted to consider what happens in the absence of the cloud band and the relative importance of free-tropospheric and boundary layer RH. Three suites of experiments are presented to address these topics.

In the first suite of experiments we introduce a scaling for CT circulation strength which we use to show is proportional to the surface sensible heat flux, H . We then introduce an updraught mass flux measure of CT strength, M_c . Based on the relationship between surface heat fluxes and CT circulation modified from past literature, we hypothesise that CTs grow stronger for stronger island surface sensible heat fluxes. To test this hypothesis, we increase the maximum surface heat flux over the island's diurnal cycle keeping all else equal. We find that M_c increases with the surface buoyancy flux which is predominantly determined by H . Additional experiments in which the island sensible and latent heat fluxes are re-partitioned for a different Bowen ratio (H/E) are performed. These experiments show that the role of island surface latent heat flux is negligible. For our experiment with the lowest H , only traces of a CT system appear, while for even our intermediate H , a robust CT forms and grows stronger with increasing H , not just in terms of M_c , but also in cloud frequency and potential temperature anomaly (θ').

Next, we perform two experiments to explore the role of RH in different vertical layers. In Chapter 3, we pointed out that reducing the boundary layer RH increases the height of the LCL and would require a stronger convergent low-level across-flow component of the CT circulation in order to provide sufficient lift for the cloud band to form. This is at odds with studies on islands in the Lesser Antilles which find that there is some link between the free-tropospheric RH and CT formation. In the first of these two experiments, we reduce the boundary layer RH by 25 %, while we reduce the free-tropospheric RH by 25 % in the second experiment. All else equal, we find that reducing the boundary layer RH not only prevents the cloud band from forming, it also prevents any cloud formation in the whole domain. Similarly, reducing the free-tropospheric RH prevents domain-wide cloud formation, but a cloud band is still able to form as the low-level RH is high enough to support a low enough LCL for the resulting CT circulation to lift air to saturation. Interestingly, the absence of a cloud band does mean the absence of a CT circulation. This result potentially has implications for the formulation of convection triggering terms for convection parametrisation schemes. In the case of this CT circulation, the thermally-driven circulation which might trigger convection occurs non-locally from the surface buoyancy perturbation.

Finally, we investigate the role of wind speed in the CT system. We show that the wind speed controls the residence time over the island heat source and therefore plays a crucial role in the evolution of the CT system. In this suite of experiments, we test ambient wind speeds between 5 and 10 m s⁻¹. We expect that the lighter wind speeds will have a stronger response as our scaling shows that CT strength varies as $1/U$. We find that there is a gradual shift toward the deeper convective regime for decreasing wind speed, rather than a regime shift at a threshold speed. However, this suite of experiments becomes more complicated than our simple scaling used to describe the first suite of experiments. As wind speed decreases, the cloud band forms closer to the island (closing the gap between the turbulent plume and cloud band) and becomes more strongly precipitating. At the lowest wind speed tested, this manifests as a quasi-stationary system just downwind of the island. We propose that the cold pool dynamics start to become important at this wind speed as the upwind propagating gust front becomes slow moving allowing newly triggered precipitating CT convection to replenish

the cold pool. A scenario in which the wind speed decreases through the day might manifest as a shallow CT transitioning to a deeper, more strongly precipitating one with impacts tracking upwind toward land. Furthermore, in environments with a more weakly capped cloud layer, a decrease in wind speed might result in a transition from shallow to deep convection with its own circulations that might disrupt or interact with the characteristic CT system described in Chapter 5. Identifying the potential for transitions, let alone the timing of the transition or intensity of the resulting convection presents a challenge both for forecasters and model developers.

7.1.4 Idealised Simulations: The Poorly-Resolved CT System

In this chapter, we look at how the CT system is changed by increasing model grid spacing up to 1.6 km. Here, we find that every tested grid length allows a representation of the CT system. At 1.6 km, the 50 km² island consists of less than 20 grid cells, yet a warm plume, low-level circulation, and cloud band still appear (at least in means of cloud fields). However, we no longer expect a reasonable representation of shallow convection for grid lengths greater than about 200 m, and this is indeed the case here. The turbulent plume part of the CT system is lost when grid length is increased to 400 m, and strongly precipitating cells start to appear. By 1.6 km, the “cloud band” consists of the preferred triggering region for convection. M_c increases with increasing grid length following our picture of stronger convective plumes established with our BOMEX work.

More generally, for grid lengths of 800 m and 1.6 km, convection tends to either be very shallow (reaching only a few hundred metres above the LCL), or penetrative (rising until they encounter their level of neutral buoyancy (LNB) from a surface parcel perspective. This means convective plumes are generally too strong and too deep in these experiments. The ramifications of this on the CT system are that strong precipitation starts to develop where precipitation is weak in the well-resolved cases. As the shallow convection becomes increasingly poorly resolved, the proportion of very shallow clouds and the horizontally averaged cloud liquid in these clouds increases. Furthermore, the penetrative convective plumes reaching their LNB in the 800 m and 1.6 km experiments spread out and linger increasing cloud cover and cloud counts at the top of the cloud

layer (near 3 km). While this result might be specific to these experiments, this general behaviour at 800 m and 1.6 km would have widespread implications for the heat and moisture budgets (through the vertical and temporal distribution of latent heating and the horizontal distribution of precipitation), the radiation budgets (through cloud cover and microphysical properties related to cloud liquid water content), and the momentum budget (through convective momentum transports, particularly in downdraughts).

We take this as strong motivation for requiring some better treatment of shallow convection in these poorly-resolved regimes. We therefore take the perspective of convection parametrisation developers and consider the number and size of clouds in our island domain during a time when convection is not organised and during the cloud band period near midday. While some aspects of the CT system are degraded with increasing grid length, the property that the CT only redistributes the clouds spatially is preserved. We show this in vertical profiles of cloud counts and radii which are unchanged by the presence of a CT. However, significant changes do occur with increasing grid length. In terms of cloud radii, this increases with increasing grid length, but does not seem to have converged by 200 m. The cloud radius controls the amount of surface area of a given cloud exposed to the unsaturated environment. In the UM's convection parametrisation, the mixing entrainment rate is based on an assumed profile of cloud radii, therefore a better estimate of this profile might yield a better mixing entrainment rate profile.

A shallow convection scheme is introduced in our most poorly-resolved case with grid length of 1.6 km. A cursory look at the resulting CT system shows that this parametrisation has an overall homogenising effect. Lightly precipitating shallow convection is triggered in every grid cell. The low-level CT circulation is much weakened and the characteristic cloud band is hard to identify in our previously used cloud frequency metric. The weaker circulation also results in a warm plume which extends farther downwind. While aspects of the CT system appear suppressed following the introduction of this scheme, a more careful selection of parameters might yield a more satisfactory compromise between resolved and sub-grid scale processes and therefore a "better" representation of the CT system. Furthermore, parametrisations designed for use in the shallow convection grey zone might be more appropriate here.

7.2 Future Work

There are several natural paths on which we might extend the work in each chapter of this thesis. We hint at this potential by applying our algorithm for a cursory look at CTs at Barbados, mechanism denial experiments for our well-resolved CT simulations, and introducing convection parametrisations for our poorly-resolved CT simulations.

Our algorithm in Chapter 3 was applied to the warm season which we knew a priori to be the most likely time to observe CT in Bermuda. Our algorithm coverage could easily be extended to include the remaining months to identify off-season cases for individual study and potentially identify unexpected behaviour. Furthermore, new higher spatial, temporal, and spectral resolution geostationary satellite imagery (available since 2016) can now be used to consider the evening decay of CTs.

Our analysis of the environments associated with CT occurrence naturally leads to the development of forecaster guidance. While NWP is approaching the ability to resolve small islands (and indeed some are now partially resolved), Bermuda is still too small for even partial representation by most global NWP. This means that the island itself, let alone its impact on the atmosphere, is not represented. Assuming NWP is able to give an adequate representation of the large scale, an analysis of these CT environments from a forecasting perspective could directly produce additional guidance for predicting different impacts from CTs. This could take the form of a decision tree as in [Niziol \(1987\)](#) for Lake Effect Snow, or perhaps post-processing of model output to arrive at probabilistic guidance.

The well-resolved CT experiments that we present are performed in the context of a small, flat, isolated island. There are many permutations of the experiment design which could be tested to explore other aspects of the CT system. Of most interest might be to extend the wind experiments to lower wind speeds. We note that the wind speeds tested in this thesis cover the higher-end of the range of observed wind speeds for CTs on Bermuda. We found that a decrease in wind speed allowed slow-moving cold pools (we showed that our diagnosed cold pool phase speeds were partially cancelling the background flow at 5 m s^{-1}). More work is needed to fully understand how precipitation interacts with the CT system in terms of both the cold pools and

the role of convective latent heat release in generating/maintaining circulation.

Furthermore, this thesis was unable to fully explore the behaviour of CTs after sunset. Our observational analysis in Chapter 3 is limited to daytime CTs because of the use of visible-channel satellite imagery. Similarly, our high-resolution simulations use bi-periodic lateral boundary conditions which potentially led to issues interpreting the CT development later in the day. New satellite spectral channels (e.g. the 3.9 μm channel aboard GOES-16, Schmit et al. (2017)) could be used to detect cloud at night. Alternate idealised experiment design, perhaps with radiative lateral boundary conditions, or more complex model grid nesting, or variable horizontal resolution could be employed to better address CT behaviour through the evening and overnight. Such observations and numerical experiments could lead to new insight into the controls of CT length and lifetime.

We take a cursory look at several aspects of the CT system and how they vary as model grid length is increased in Chapter 6. That chapter is aimed at leading to many paths of future work regarding NWP regimes in which shallow convection is poorly resolved. We advocate for the use of shallow convection schemes even at “convection permitting” resolutions. Furthermore, suggested research avenues targeting the improvement of shallow convection parametrisation schemes including more sophisticated entrainment profiles, a physically-constrained representation of stochasticity, and “scale-aware” behaviour where the scheme behaves differently as convection is increasingly well resolved. We showed in both Chapters 4 and 6 that the representation of the vertical thermodynamic structure in models is crucial for an accurate simulated cloud depth. In addition, other projects under the remit of ParaCon include a large island which focuses on the propagation of deep convection from land to sea over the diurnal cycle. Experiments like that could be compared or combined with our island experiments which focus on shallow convection to gain insight into more mixed, shallow-deep environments.

There are many steps that could be taken towards a more realistic CT simulation, for instance, our island could have more realistic geometry and topography, or we could introduce interactive radiation and surface fluxes. We could go so far as to use a real case in a global model nested down to higher resolution over an island to capture time

evolving large-scale environments. It would be of interest to the climate community to have a better understanding of the impact of these small islands on larger-scale circulations. NWP used for climate prediction largely employs grid lengths far too large to represent the small islands responsible for CTs. However, an individual CT, extending more than 100 km downwind of these islands in observations, might still extend across several model grid cells. In addition, convection is usually parametrised for these grid lengths. We show that one small, flat, island in isolation under equilibrium conditions does little more than change the spatial distribution of the clouds. Does the interaction between several islands start to become important for the larger scale environment? Do we need to have some representation of sub-grid islands in convection parametrisation schemes? There are parts of the world with many small islands, for instance the Caribbean Sea and the Tropical Pacific. Combination field-numerical experiments such as the EUREC4A (Elucidating the role of Clouds-Circulation Coupling in Climate) project ([Bony et al., 2017](#)) which is planned for upwind of and over Barbados are expected to result in new insight into the role of mesoscale organisation of shallow cumulus in the trade wind region, and several types of organisation have already been identified unrelated to island heating ([Stevens et al., 2019](#)). Such field campaigns have more resources to conduct targeted in-depth observational investigations of interesting convective systems and have the potential to produce model-constraining observations.

Finally, we note that the work presented here not only initiates further commentary on the future development of models, but also on what might be required for future model development in an era marked by increased public demand for “hyper-local”, “hyper-accurate” forecasts with frequent updates. The introduction of higher model resolution is allowing new phenomena tied to surface inhomogeneities to be represented explicitly. However, this increase in resolution is not a panacea. Even at kilometre-scale grid lengths, shallow convection is not well represented and therefore we cannot expect its influence locally and on larger-scales to be adequately represented. In order to meet these growing public demands and reduce model uncertainty, further work is needed on convection parametrisation (as is ongoing within the ParaCon community), and in understanding the behaviour of higher resolution models more generally.

Bibliography

- Abel, S. J. and I. A. Boutle, 2012: An Improved Representation of the Rain Drop Size Distribution for Single-moment Microphysics Schemes. *Quart. J. Roy. Meteor. Soc.*, **138**, 2151–2162.
- Abel, S. J. and B. J. Shipway, 2007: A Comparison of Cloud-Resolving Model Simulations of Trade Wind Cumulus with Aircraft Observations taken during RICO. *Quart. J. Roy. Meteor. Soc.*, **133**, 781–794.
- Alduchov, O. A. and R. E. Eskridge, 1996: Improved Magnus form Approximation of Saturation Vapour Pressure. *J. Appl. Meteor.*, **35**, 601–609.
- Arakawa, A. and V. R. Lamb, 1977: Computational Design of the Basic Dynamical Processes of the UCLA General Circulation Model. *Methods in Comp. Phys.*, **17**, 174–265.
- Arakawa, A. and C. Wu, 2013: A Unified Representation of Deep Moist Convection In Numerical Modelling of the Atmosphere: Part I. *J. Atmos. Sci.*, **70**, 1977–1992.
- Betts, A. K., 1973: Non-Precipitating Cumulus Convection and its Parameterization. *Quart. J. Roy. Meteor. Soc.*, **99**, 178–196.
- Betts, A. K. and M. J. Miller, 1986: A New Convective Adjustment Scheme, II: Single Column Tests Using GATE wave, BOMEX, and Arctic Air-mass Datasets. *Quart. J. Roy. Meteor. Soc.*, **112**, 693–709.
- Bhumralkar, C. M., 1973a: An Observational and Theoretical Study of Atmospheric Flow over a Heated Island: Part I. *Mon. Wea. Rev.*, **101**, 719–730.
- 1973b: An Observational and Theoretical Study of Atmospheric Flow over a Heated Island: Part II. *Mon. Wea. Rev.*, **101**, 731–745.

- Bolton, D., 1980: The Computation of Equivalent Potential Temperature. *Mon. Wea. Rev.*, **108**, 1046 – 1053.
- Bony, S., B. Stevens, F. Ament, S. Bigorre, P. Chazette, S. Crewell, J. Delano, K. Emanuel, D. Farrell, C. Flamant, S. Gross, L. Hirsch, J. Karstensen, B. Mayer, L. Nuijens, J. H. J. Ruppert, I. Sandu, P. Siebesma, S. Speich, F. Szczap, J. Totems, R. Vogel, M. Wendisch, and M. Wirth, 2017: A Field Campaign to Elucidate the Couplings Between Clouds, Convection and Circulation. *Surveys in Geophys.*, **38**, 1529–1568.
- Boutle, I. A., S. J. Abel, P. G. Hill, and C. J. Morcrette, 2014a: Spatial Variability of Liquid Cloud and Rain: Observations and Microphysical Effects. *Quart. J. Roy. Meteor. Soc.*, **140**, 583–594.
- Boutle, I. A., J. E. J. Eyre, and A. P. Lock, 2014b: Seamless Stratocumulus Simulation across the Turbulent Grey Zone. *Mon. Wea. Rev.*, **142**, 1655–1668.
- Brast, M., V. Schemann, and R. A. J. Neggers, 2012: Investigating the Scale Adaptivity of a Size-Filtered Mass Flux Parametrisation in the Gray Zone of Shallow Cumulus Convection. *J. Atmos. Sci.*, **75**, 1195–1214.
- Bryan, G., J. C. Wyngaard, and J. M. Fritsch, 2003: Resolution Requirements for the Simulation of Deep Moist Convection. *Mon. Wea. Rev.*, **131**, 2394–2416.
- Buck, A., 1981: New Equation for Computing Vapour Pressure and Enhancement Factor. *J. Appl. Meteor.*, **20**, 1527–1532.
- Bush, M., T. Allen, C. Bain, I. Boutle, J. Edwards, A. Finnenkoetter, C. Franklin, K. Hanley, H. Lean, A. Lock, J. Manners, M. Mittermaier, C. Morcrette, R. North, J. Petch, C. Short, S. Vosper, D. Walters, S. Webster, M. Weeks, J. Wilkinson, N. Wood, and M. Zerroukat, 2019: The First Met Office Unified Model/JULES Regional Atmosphere and Land configuration, RAL1. *Geosci. Model Dev. Discuss.*, **in review**, 1–47.
- Byers, H. R. and R. K. Hall, 1955: A Census of Cumulus-Cloud Height Versus Precip-

BIBLIOGRAPHY

itation in the Vicinity of Puerto Rico During the Winter and Spring of 1953-54. *J. Meteor.*, **12**, 176–178.

Charney, J. G. and N. A. Phillips, 1953: Numerical Integration of the Quasi-Geostrophic Equations for Barotropic and Simple Baroclinic Flows. *J. Meteor.*, **10**, 71–99.

Charnock, H., 1955: Wind Stress on a Water Surface. *Quart. J. Roy. Meteor. Soc.*, **81**, 639–640.

Ching, J. K. S., J. F. Clarke, and J. M. Godowitch, 1983: Modulation of Heat Flux by Different Scales of Advection in an Urban Environment. *Bound. -Layer. Meteor.*, **25**, 171–191.

CIA, 2017: The CIA World Factbook: Bermuda.

URL <https://www.cia.gov/library/publications/the-world-factbook/geos/bd.html>

Crook, N. A., 2001: Understanding Hector: The Dynamics of Island Thunderstorms. *Mon. Wea. Rev.*, **129**, 1550–1563.

Crook, N. A. and D. F. Tucker, 2005: Flow over heated terrain. Part I: Linear theory and idealized numerical simulations. *Mon. Wea. Rev.*, **133**, 2552–2564.

Cuijpers, J. W. M., 1994: *Large-Eddy Simulation of Cumulus Convection*. Ph.D. thesis, Delft University of Technology, The Netherlands, 158 pp.

Davies, T., M. J. P. Cullen, A. J. Malcolm, M. H. Mawson, A. Staniforth, A. A. White, and N. Wood, 2005: A new dynamical core for the Met Office's global and regional modelling of the atmosphere. *Quart. J. Roy. Meteor. Soc.*, **131**, 1759–1782.

de Rooy, W. C., P. Bechtold, K. Fröhlich, C. Hohenegger, H. Jonker, D. Mironov, A. P. Siebesma, J. Teixeira, and J. I. Yano, 2013: Entrainment and Detrainment in Cumulus Convection: An Overview. *Quart. J. Roy. Meteor. Soc.*, **139**, 1–19.

Dee, D. P., S. M. Uppala, A. J. Simmons, P. Berrisford, P. Poli, S. Kobayashi, U. Andrae, M. A. Balmaseda, G. Balsamo, P. Bauer, P. Bechtold, A. C. Beljaars, L. van de

- Berg, J. Bidlot, N. Bormann, C. Delsol, R. Dragani, M. Fuentes, A. J. Geer, L. Haimberger, S. B. Healy, H. Hersbach, E. V. Hólm, L. Isaksen, P. Kållberg, M. Köhler, M. Matricardi, A. P. McNally, B. M. Monge-Sanz, J. J. Morcrette, B. K. Park, C. Peubey, P. de Rosnay, C. Tavalato, J. N. Thépaut, and F. Vitart, 2011: The ERA-Interim reanalysis: Configuration and performance of the data assimilation system. *Quart. J. Roy. Meteor. Soc.*, **137**, 553–597.
- Derbyshire, S. H., I. Beau, P. Bechtold, J. Y. Grandpeix, J. M. Piriou, J. L. Redelsperger, and P. M. M. Soares, 2004: Sensitivity of moist convection to environmental humidity. *Quart. J. Roy. Meteor. Soc.*, **130**, 3055–3079.
- Derbyshire, S. H., A. V. Maidens, S. F. Milton, R. A. Stratton, and M. Willet, 2011: Adaptive Detrainment in a Convection Parametrisation. *Quart. J. Roy. Meteor. Soc.*, **137**, 1856–1871.
- Dorman, C. E., 1994: Guadalupe Island Cloud Trail. *Mon. Wea. Rev.*, **122**, 235–242.
- Douglas, D. and T. Peucker, 1973: Algorithms for the Reduction of the Number of Points Required to Represent a Digitized Line or its Caricature. *The Canadian Cartographer*, **10**, 112–122.
- Durre, I., R. S. Vose, and D. B. Wuertz, 2006: Overview of the Integrated Global Radiosonde Archive. *J. Climate*, **19**, 53–68.
- 2008: Robust automated quality assurance of radiosonde temperatures. *J. Appl. Meteor. Climatol.*, **47**, 2081–2095.
- Emanuel, K. A., 1994: *Atmospheric Convection*. Oxford University Press: New York, NY, USA, 581pp pp.
- Estoque, M. A. and C. M. Bhumralkar, 1969: Flow over a localized heat source. *Mon. Wea. Rev.*, **97**, 850–859.
- Fairall, C. W., E. F. Bradley, J. E. Hare, A. A. Grachev, and J. B. Edson, 2003: Bulk Parametrisation of Air-Sea Fluxes: Updates and Verification for the COARE Algorithm. *J. Clim.*, **16**, 571–591.

- Forbes, R., J. Wilkinson, D. Wilson, I. Boutle, and S. A. Smith, 2017: The Large-Scale Precipitation Parametrisation Scheme. *Unified Model Documentation Paper*, **026**, pp30.
- Garstang, M., P. D. Tyson, and G. D. Emmitt, 1975: The Structure of Heat Islands. *Rev. Geophys. Space*, **13**, 139–165.
- Gentine, P., A. A. M. Holtslag, F. D’Andrea, and M. Ek, 2013: Surface and Atmospheric Controls on the Onset of Moist Convection over Land. *J. Hydrometeor.*, **14**, 1443–1462.
- Grabowski, W. W., P. Bechtold, A. Cheng, R. Forbes, C. Halliwell, M. Khairoutdinov, S. Lang, T. Nasuno, J. Petch, W.-K. Tao, R. Wong, X. Wu, and K.-M. Xu, 2006: Daytime Convective Development Over Land: A Model Intercomparison Based on LBA Observations. *Quart. J. Roy. Meteor. Soc.*, **132**, 317–344.
- Grant, A. L. M., 2001: Cloud-Base Fluxes in the Cumulus-Capped Boundary Layer. *Quart. J. Roy. Meteor. Soc.*, **127**, 407–421.
- Grant, A. L. M. and A. R. Brown, 1999: A Similarity Hypothesis for Shallow-Cumulus Transports. *Quart. J. Roy. Meteor. Soc.*, **125**, 1913–1936.
- Grant, A. L. M. and A. P. Lock, 2004: The Turbulent Kinetic Energy Budget for Shallow Convection. *Quart. J. Roy. Meteor. Soc.*, **130**, 402–422.
- Gregory, D. and P. R. Rowntree, 1990: A Mass Flux Convection Scheme with Representation of Cloud Ensemble Characteristics and Stability-Dependent Closure. *Mon. Wea. Rev.*, **118**, 1483–1506.
- Halliwell, C., I. Boutle, and K. Hanley, 2017: Subgrid Turbulence Scheme. *Unified Model Documentation Paper*, **028**, pp30.
- Harvey, N., C. Daleu, S. Woolnough, and R. S. Plant, 2019: Understanding the Impact of Heterogeneity on the Diurnal Cycle of Deep Convection. *Conference Poster. Convection Parametrisation: Progress and Challenges Workshop*.
- Holland, J. Z., 1972: Comparative Evaluation of some BOMEX Measurements of Sea Surface Evaporation, Energy Flux, and Stress. *J. Phys. Oceanogr.*, **2**, 476–486.

BIBLIOGRAPHY

- Holland, J. Z. and E. M. Rasmusson, 1973: Measurements of the Atmospheric Mass, Energy, and Momentum Budgets Over a 500-Kilometer Square of Tropical Ocean. *Mon. Wea. Rev.*, **101**, 44–55.
- Holton, J. R. and G. J. Hakim, 2013: *An Introduction to Dynamic Meteorology. 5th Edition*, Academic Press.. Academic Press, Elsevier Science.
- Hsu, H., 1987a: Mesoscale Lake-effect Snowstorms in the Vicinity of Lake Michigan: Linear Theory and Numerical Simulations. *J. Atmos. Sci.*, **44**, 1019–1040.
- Hsu, H. M., 1987b: Study of Linear Steady Atmospheric Flow above a Finite Surface Heating. *J. Atmos. Sci.*, **44**, 186–199.
- Johnson, R. H. and X. Lin, 1997: Episodic Trade Wind Regimes over the Western Pacific Warm Pool. *J. Atmos. Sci.*, **54**, 2020–2034.
- Johnston, M. C., C. E. Holloway, and R. S. Plant, 2018: Cloud Trails past Bermuda: A Five-Year Climatology from 2012 to 2016. *Mon. Wea. Rev.*, **146**, 4039–4055.
- Kalanda, B. D., T. R. Oke, and D. L. Spittlehouse, 1980: Suburban Energy Balance Estimates for Vancouver, B. C., Using the Bowen Ratio-Energy Balance Approach. *J. Appl. Meteor.*, **19**, 791–802.
- Kirshbaum, D. J., 2013: On Thermally Forced Circulations over Heated Terrain. *J. Atmos. Sci.*, **70**, 1690–1709.
- Kirshbaum, D. J. and J. G. Fairman, 2015: Cloud Trails Past the Lesser Antilles. *Mon. Wea. Rev.*, **143**, 995–1017.
- Kirshbaum, D. J. and A. L. M. Grant, 2012: Invigouration of Cumulus Cloud Fields by Mesoscale Ascent. *Quart. J. Roy. Meteor. Soc.*, **138**, 2136–2150.
- Kirshbaum, D. J. and C.-C. Wang, 2014: Boundary Layer Updrafts Driven by Airflow over Heated Terrain. *J. Atmos. Sci.*, **71**, 1425–1442.
- Krishnamurti, T. N., K. Rajendron, T. S. Vijaya Kumar, S. Lord, Z. Toth, X. Zou, S. Cocke, J. E. Ahlquist, and I. M. Navon, 2003: Improved Skill for the Anomaly Correlation of Geopotential Heights at 500 hPa. *Mon. Wea. Rev.*, **131**, 1082–1102.

BIBLIOGRAPHY

- Lackman, G., 2011: *Midlatitude Synoptic Meteorology: Dynamics, Analysis & Forecasting*. American Meteorological Society: Boston, MA, USA, 345pp pp.
- Lilly, D. K., 1967: On the Application of the Eddy Viscosity Concept in the Inertial Subrange of Turbulence. *NCAR Ms. No 123, Boulder, Co.*
- Lock, A., J. Edwards, and I. Boutle, 2017: The Parametrisation of Boundary Layer Processes. *Unified Model Documentation Paper*, **024**, pp30.
- Lock, A. P., A. R. Brown, M. R. Bush, G. M. Martin, and R. N. B. Smith, 2000: A New Boundary Layer Mixing Scheme. Part I: Scheme Description and Single-Column Model Tests. *Mon. Wea. Rev.*, **128**, 3187–3199.
- Long, C. N. and S. A. McFarlane, 2012: Quantification of the Impact of Nauru Island on ARM Measurements. *J. Appl. Meteor. Climatol.*, **51**, 628–636.
- Ludlam, F., 1980: *Clouds and Storms: The Behavior and Effect of Water in the Atmosphere*. Pennsylvania State University Press: University Park, PA, USA, 405pp pp.
- Mahrer, Y. and R. A. Pielke, 1976: Numerical Simulation of the Airflow Over Barbados. *Mon. Wea. Rev.*, **104**, 1392–1402.
- Malkus, J. S., 1963: Tropical Rain Induced by a Small Natural Heat Source. *J. Appl. Meteor.*, **2**, 547–556.
- Malkus, J. S. and A. F. Bunker, 1952: Observational Studies of the air flow over Nantucket Island During the Summer of 1950. *Pap. Phys. Oceanogr. Meteor.*, **12**, 1–50.
- Malkus, J. S. and M. E. Stern, 1953: The Flow of a Stable Atmosphere Over a Heated Island, Part I. *J. Meteor.*, **10**, 30–41.
- Mason, P. J., 1994: Large-Eddy Simulation: A Critical Review of the Technique. *Quart. J. Roy. Meteor. Soc.*, **120**, 1–26.

- Mason, P. J. and R. I. Sykes, 1980: A Two-Dimensional Numerical Study of Horizontal Roll Vortices in the Neutral Atmospheric Boundary Layer. *Quart. J. Roy. Meteor. Soc.*, **106**, 351–366.
- 1982: A Two-Dimensional Numerical Study of Horizontal Roll Vortices in an Inversion Capped Planetary Boundary Layer. *Quart. J. Roy. Meteor. Soc.*, **108**, 801–823.
- Matthews, S., J. M. Hacker, J. Cole, J. Hare, C. N. Long, and R. M. Reynolds, 2007: Modification of the Atmospheric Boundary Layer by a Small Island: Observations from Nauru. *Mon. Wea. Rev.*, **135**, 891–905.
- McFarlane, S. A., C. N. Long, and D. M. Flynn, 2005: Impact of Island-Induced Clouds on Surface Measurements: Analysis of the ARM Nauru Island Effect Study Data. *J. Appl. Meteor.*, **44**, 1045–1065.
- Medeiros, B. and L. Nuijens, 2016: Clouds at Barbados are Representative of Clouds Across the Trade Wind Regions in Observations and Climate Models. *Proceedings of the National Academy of Sciences*, **113**, E3062–E3070.
- Mellor, G., 1977: The Gaussian Cloud Model. *J. Atmos. Sci.*, **34**, 356–358.
- Miller, S. T. K., B. D. Keim, R. W. Talbot, and H. Mao, 2003: Sea breeze: Structure, Forecasting, and Impacts. *Rev. Geophys.*, **41**, 1011.
- Moser, D. H. and S. Lasher-Trapp, 2018: Cloud-Spacing Effects upon Entrainment and Rainfall along a Convective Line. *J. Appl. Meteor. Climatol.*, **57**, 1865–1882.
- Nie, J. and Z. Kuang, 2012: Responses of Shallow Cumulus Convection to Large-Scale Temperature and Moisture Perturbations: A Comparison of Large-Eddy Simulations and a Convective Parametrisation Based on Stochastically Entraining Parcels. *J. Atmos. Sci.*, **69**, 1936–1956.
- Nitta, T. and S. Esbensen, 1974: Heat and Moisture Budget Analyses Using BOMEX Data. *Mon. Wea. Rev.*, **102**, 17–28.
- Niziol, T., 1987: Operational Forecasting of Lake Effect Snowfall in Central and Western New York. *Wea. Forecasting*, **2**, 310–321.

- Nordeen, M. K., P. Minnis, D. R. Doelling, D. Pethick, and L. Nguyen, 2001: Satellite Observations of Cloud Plumes Generated by Nauru. *Geophys. Res. Lett.*, **28**, 631–634.
- Nuijens, L. and B. Stevens, 2012: The Influence of Wind Speed on Shallow Marine Cumulus Convection. *J. Atmos. Sci.*, **69**, 168–184.
- Nuijens, L., B. Stevens, and A. P. Siebesma, 2009: The Environment of Precipitating Shallow Cumulus Convection. *J. Atmos. Sci.*, **66**, 1962–1979.
- Peirce, C. S., 1884: The numerical measure of the success of predictions. *Science*, **4**, 453–454.
- Perry, K. D. and P. V. Hobbs, 1996: Influences of Isolated Cumulus Clouds on the Humidity of their Surroundings. *J. Atmos. Sci.*, **53**, 159–174.
- Ramer, U., 1972: An Iterative Procedure for the Polygonal Approximation of Plane Curves. *Computer Graphics and Image Processing*, **1**, 244–256.
- Rauber, R. M., B. Stevens, H. T. Ochs, C. Knight, B. A. Albrecht, A. M. Blyth, C. W. Fairall, J. B. Jensen, S. G. Lasher-Trapp, O. L. Mayol-Bracero, G. Vali, J. R. Anderson, B. A. Baker, A. R. Bandy, E. Burnet, J.-L. Brenguier, W. A. Brewer, P. R. A. Brown, R. Chuang, W. R. Cotton, L. Di Girolamo, B. Geerts, H. Gerber, S. Gke, L. Gomes, B. G. Heikes, J. G. Hudson, P. Kollias, R. R. Lawson, S. K. Krueger, D. H. Lenschow, L. Nuijens, D. W. O’Sullivan, R. A. Rilling, D. C. Rogers, A. P. Siebesma, E. Snodgrass, J. L. Stith, D. C. Thornton, S. Tucker, C. H. Twohy, and P. Zuidema, 2007: Rain in Shallow Cumulus Over the Ocean: The RICO Campaign. *Bull. Amer. Meteor. Soc.*, **88**, 1912–1928.
- Rio, C., F. Hourdin, J. Y. Grandpeix, and J. P. Lafore, 2009: Shifting the diurnal cycle of parameterized deep convection over land. *Geophys. Res. Lett.*, **36**, L07809.
- Robinson, F. J., S. C. Sherwood, D. Gerstle, C. Liu, and D. J. Kirshbaum, 2011: Exploring the Land-Ocean Contrast in Convective Vigour Using Islands. *J. Atmos. Sci.*, **68**, 602–618.

- Rotunno, R., J. B. Klemp, and M. L. Weisman, 1988: A Theory for Strong, Long-Lived Squall Lines. *J. Atmos. Sci.*, **45**, 463–485.
- Savijarvi, H. and S. Matthews, 2004: Flow over Small Heat Islands: A Numerical Sensitivity Study. *J. Atmos. Sci.*, **61**, 859–868.
- Schmit, T. J., P. Griffith, M. M. Gunshor, D. J. M., S. J. Goodman, and W. J. Lebar, 2017: A Closer Look at the ABI on the GOES-R Series. *Bull. Amer. Meteor. Soc.*, **98**, 681–698.
- Sellar, A. A., C. G. Jones, J. Mulcahy, Y. Tang, A. Yool, A. Wiltshire, F. M. O’Connor, M. Stringer, R. Hill, J. Palmieri, S. Woodward, L. de Mora, T. Kuhlbrodt, S. Rumbold, D. I. Kelley, R. Ellis, C. E. Johnson, J. Walton, N. L. Abraham, M. B. Andrews, T. Andrews, A. T. Archibald, S. Berthou, R. Burke, E. Blockley, K. Carslaw, M. Dalvi, J. Edwards, G. A. Folberth, N. Gedney, P. T. Griffiths, A. B. Harper, M. A. Hendry, A. J. Hewitt, B. Johnson, A. Jones, C. D. Jones, J. Keeble, S. Liddicoat, O. Morgenstern, R. J. Parker, V. Predoi, E. Robertson, A. Siahayan, R. S. Smith, R. Swaminathan, M. T. Woodhouse, G. Zeng, and M. Zerroukat, 2019: UKEMS1: Description and Evaluation of the UK Earth System Model. *J. Adv. Modelling. Earth Sys.*, **11**, manuscript accepted.
- Siebesma, A. P., C. S. Bretherton, A. Brown, A. Chlond, J. Cuxart, P. G. Duynkerke, H. Jiang, M. Khairoutdinov, D. Lewellen, C. H. Moeng, E. Sanchez, B. Stevens, and D. E. Stevens, 2003: A Large Eddy Simulation Intercomparison Study of Shallow Cumulus Convection. *J. Atmos. Sci.*, **60**, 1201–1219.
- Siebesma, A. P. and J. W. M. Cuijpers, 1995: Evaluation of Parametric Assumptions for Shallow Cumulus Convection. *J. Atmos. Sci.*, **52**, 650–666.
- Simmons, A., A. J. and Untch, C. Jakob, P. Källberg, and P. Undén, 1999: Stratospheric Water Vapour and Tropical Tropopause Temperatures in ECMWF Analyses and Multi-Year Simulations. *Quart. J. Roy. Meteor. Soc.*, **125**, 353–386.
- Smagorinsky, J., 1963: General Circulation Experiments with the Primitive Equations. *Mon. Wea. Rev.*, **91**, 99–164.

- Smith, R. B., A. C. Gleason, P. A. Gluhosky, and V. Grubisic, 2007: The Wake of St. Vincent. *J. Atmos. Sci.*, **54**, 606–623.
- Smith, R. N. B., 1990: A Scheme for Predicting Layer Clouds and their Water Content in a General Circulation Model. *Quart. J. Roy. Meteor. Soc.*, **116**, 435–460.
- Smith, S. D., 1988: Coefficients for Sea Surface Wind Stress, Heat Flux, and Wind Profiles as a Function of Wind Speed and Temperature. *J. Geophys. Res.*, **93**, 15467–15472.
- Smolarkiewicz, P. K., R. M. Rasmussen, and T. L. Clark, 1988: On the dynamics of Hawaiian Cloud bands: Island forcing. *J. Atmos. Sci.*, **45**, 1872–1905.
- Sobel, A. H., C. D. Burleyson, and S. E. Yuter, 2011: Rain on Small Tropical Islands. *J. Geophys. Res.*, **116**, D08102.
- Sommeria, G. and J. W. Deardorff, 1977: Subgrid-Scale Condensation in Models of Non-Precipitating Clouds. *J. Atmos. Sci.*, **34**, 344–355.
- Stern, H. and N. E. Davidson, 2015: Trends in the skill of weather prediction at lead times of 1-14 days. *Quart. J. Roy. Meteor. Soc.*, **141**, 2726–2736.
- Stevens, B., S. Bony, H. Brogniez, L. Hentgen, C. Hohenegger, C. Kiemle, T. S. L'Ecuyer, A. K. Naumann, H. Schulz, P. A. Siebesma, J. Vial, D. M. Winker, and P. Zuidema, 2019: Sugar, Gravel, Fish, and Flowers: Mesoscale Cloud Patterns in the Trade Winds. *Quart. J. Roy. Meteor. Soc.*, **Early View**, 1–12.
- Stratton, R., M. Willet, S. Derbyshire, R. Wong, M. Whitall, and G. G. Rooney, 2017: Convection Schemes. *Unified Model Documentation Paper*, **027**, pp30.
- Stull, R. B., 1988: *An Introduction to Boundary Layer Meteorology*. Kluwer Academic Publishers, Dordrecht, 666pp pp.
- Sutherland, M. G., S. J. McLean, M. R. Love, K. S. Carignan, and B. W. Eakins, 2013: Digital Elevation Models of Bermuda: Data Sources, Processing and Analysis. *Prepared for the NOAA Pacific Marine Environmental Laboratory by NOAA National Geophysical Data Center*, 1–7.

- Swann, H., 2001: Evaluation of the Mass-Flux Approach to Parametrising Deep Convection. *Quart. J. Roy. Meteor. Soc.*, **127**, 1239–1260.
- Tobin, I., S. Bony, and R. Roca, 2012: Observational Evidence for Relationships Between the Degree of Aggregation of Deep Convection, Water Vapour, Surface Fluxes, and Radiation. *J. Clim.*, **25**, 6885–6904.
- Torn, R. D. and C. A. Davis, 2012: The Influence of Shallow Convection on Tropical Cyclone Track Forecasts. *Mon. Wea. Rev.*, **140**, 2188–2197.
- Turner, J. S., 1963: The Motion of Buoyant Elements in Turbulent Surroundings. *J. Fluid Mech.*, **16**, 1–16.
- Wang, S. and A. H. Sobel, 2017: Factors Controlling Rain on Small Tropical Islands: Diurnal Cycle, Large-Scale Wind Speed, and Topography. *J. Atmos. Sci.*, **74**, 3515–3532.
- Weckworth, T. M., J. W. Wilson, R. M. Wakimoto, and N. A. Crook, 1997: Horizontal Convective Rolls: Determining the Environmental Conditions Supporting their Existence and Characteristics. *Mon. Wea. Rev.*, **125**, 505–526.
- Williams, E., T. Chan, and D. Boccippio, 2004: Islands as Miniature Continents, Another Look at the Land-Ocean Lightning Contrast. *J. Geophys. Res.*, **109**, D16206.
- Wilson, D., C. Morcrette, and I. Boutle, 2017: The Large-Scale Cloud Scheme and Saturated Specific Humidity. *Unified Model Documentation Paper*, **029**, pp30.
- Wilson, D. R. and S. P. Ballard, 1999: A Microphysically-Based Precipitation Scheme for the UK. *Quart. J. Roy. Meteor. Soc.*, **125**, 1607–1636.
- WMO, 2018: *Guide to Instruments and Methods of Observation. Part I: Measurement of Meteorological Variables. Chapter 5: Measurement of Surface Wind..* WMO, 167–183 pp.
- Wood, N., A. Staniforth, A. White, J. Thuburn, T. Allen, T. Davies, M. Diamantakis, M. Dubal, M. Gross, T. Melvin, C. Smith, and M. Zerroukat, 2017: ENDGame Formulation. *Unified Model Documentation Paper*, **016**, pp30.

- Wood, R., C. R. Mechoso, C. S. Bretherton, R. A. Weller, B. Huebert, F. Straneo, B. A. Albrecht, H. Coe, G. Allen, G. Vaughan, P. Daum, C. Fairall, D. Chand, L. Gallardo Klenner, R. Garreaud, C. Grados, D. S. Covert, T. S. Bates, R. Krejci, L. M. Russell, S. de Szoeke, A. Brewer, S. E. Yuter, S. R. Springston, A. Chaigneau, T. Toniazzo, P. Minnis, R. Palikonda, S. J. Abel, W. O. J. Brown, S. Williams, J. Fochesatto, J. Brioude, and K. N. Bower, 2011: The VAMOS Ocean-Cloud-Atmosphere-Land Study Regional Experiment (VOCALS-REx): goals, platforms, and field operations. *Atmospheric Chemistry and Physics*, **11**, 627–654.
- Wu, C. M., B. Stevens, and A. Arakawa, 2009: What Controls the Transition from Shallow to Deep Convection? *J. Atmos. Sci.*, **66**, 1793–1806.
- Yanai, M., S. Esbensen, and J. H. Chi, 1973: Determination of Bulk Properties of Tropical Cloud Clusters from Large Scale Heat and Moisture Budgets. *J. Atmos. Sci.*, **30**, 611–627.
- Yang, Y. and Y.-L. Chen, 2008: Effects of Terrain Heights and Sizes on Island-Scale Circulations and Rainfall for the Island of Hawaii during HaRP. *Mon. Wea. Rev.*, **136**, 120–146.
- Yang, Y., S.-P. Xie, and J. Hafner, 2008a: The Thermal Wake of Kauai Island: Satellite Observations and Numerical Simulations. *J. Clim.*, **21**, 4568–4586.
- Yang, Y., S. P. Xie, and J. Hafner, 2008b: Cloud Patterns Lee of Hawaii Island: A Synthesis of Satellite Observations and Numerical Simulation. *J. Geophys. Res. Atmos.*, **113**, D15126.
- Zhang, G. J. and N. A. McFarlane, 1995: Sensitivity of Climate Simulations to the Parametrisation of Cumulus Convection in the Canadian Centre General Circulation Model. *Atmos. Ocean*, **33**, 407–446.
- Zuill, W., 1946: *Bermuda Journey*. Coward-McCann, inc (1946), 426 pp.