Convection in the Unified Model

Martin S. Singh

Honours Thesis submitted as part of the B.Sc. (Honours) degree
in the School of Mathematical Sciences, Monash University.

Supervisor: Prof. Christian Jakob
Co-supervisor: Prof. Michael Reeder

Submitted 20th October, 2008.

Acknowledgments

This research was conducted under the supervision of Prof. Christian Jakob and Prof. Michael Reeder in the School of Mathematical Sciences, Monash University. Charmaine Franklin and John Edwards provided valuable technical support for the UM single column model. The assistance of Laura Davies on many practical aspects of modelling with the SCM was greatly appreciated by the author. The Monash University atmospheric science honours class of 2008 were a constant source of moral support during the writing of this thesis.

Contents

Chapter 1

Introduction

The Earth's atmosphere is profoundly influenced by moist convection. Convective heating is a major contributor to the energy budget of global circulations such as the Hadley cell (Emanuel *et al.*, 1994), and, on smaller scales, is responsible for whether it will be a sunny or stormy day. Despite its obvious importance for both weather forecasts and climate simulation, many aspects of moist convection are still not well understood, or well represented in atmospheric models. In atmospheric general circulation models (GCMs), convection is described as an ensemble of many different convective clouds at different points in their life-cycle. In general, this description contains three components; a trigger, which determines whether the atmosphere can support convection, a cloud model, which describes the distribution of convective activity in the vertical, and a closure, that deduces the overall strength of convection. In this thesis, the representation of convection in a particular GCM – the Unified Model is investigated within an idealised framework of large-scale conditions.

The Unified Model (UM), is a coupled ocean atmosphere GCM developed by the U.K. Meteorological Office. In the simulations used in this study, it is run in single column mode. This means that instead of simulating the entire atmosphere, one vertical column of grid-boxes is selected, and these become the model domain. This approach has two advantages, firstly, since only a fraction of the full domain is being integrated, the computational costs of running the model are dramatically reduced. This allows a large number of simulations to be performed, and thus a number of different scenarios can be tested. Secondly, since there is no horizontal interaction between grid-boxes, all large scale dynamics must be prescribed. This allows for simulations under simplified conditions, in order to examine the behaviour of key parameterisations without the complications of dynamical interactions. Similarly, as the model is run in atmosphere only mode, conditions at the lower boundary – the sea-surface temperatures (SST) – must also be given to the model. In this study a set of simple large-scale conditions are imposed in order to investigate convective activity within the UM single-column model (SCM). A vertical wind profile with a maximum in the mid-troposphere, and a value for the SST are chosen. The model is then integrated with constant SST and vertical wind for a period of 25 days. A number of simulations are run with different values for the vertical wind maximum, and the SST, corresponding to the range typically encountered over the tropical oceans. This allows the behaviour of the model over this 2-D parameter space of vertical wind strength and SST to be investigated.

Surprisingly, with such a simple set of boundary conditions the model produces a rich

set of phenomena. It is found that the convective response of the UM is of a highly variable convective precipitation rate. A burst of convective activity occurs for a single timestep, which is followed by a number of timesteps in which convection is suppressed. This behaviour is not characteristic of an ensemble of convective cells that make up a GCM grid-box. Rather, a build-up and release of potential energy describes the response that might be expected from an individual convective cloud. More importantly, however, the highly variable convection observed in the model is not even consistent with the paradigm the model itself is based on – the quasi-equilibrium hypothesis.

Before we delve more deeply into the model results however, we must understand the ideas and paradigms on which these models are based. Thus the rest of this chapter is devoted to a brief review of convection and convection modelling. Chapter 2 gives an overview of the Unified Model and its particular convection parameterisation, as well as a description of the specific simulations used in this study. In Chapter 3 we present results of the simulations, with a particular emphasis on time averaged behaviour. The strong variability of convection is discussed in chapter 4, and the convective trigger and closure are investigated as possible causes. In the final chapter, a discussion of the results is given, and it is argued that the convective trigger should be modified if the variability in convective rain rates is to be reduced.

1.1 Moist atmospheric convection

In the most general context, convection occurs when density gradients result in non-zero buoyancy forces in some regions of a fluid. This force imbalance accelerates the flow, and facilitates the transport of energy and mass we generally associate with convection. In the atmosphere these density gradients are due to differences in temperature, while in the oceans concentrations of various solutes are also important. In the field of meteorology, convection tends to be framed in the language of parcel theory (see e.g. Wallace & Hobbs, 2006). We consider a hypothetical 'parcel' of air existing in some background environment. The parcel is displaced infinitesimally in such a way that it exchanges no energy with the environment. The parcel buoyancy at the new position is calculated, and if the resultant force serves to amplify the initial disturbance the atmosphere is called unstable. This instability causes a rearrangement of fluid in the vertical which re-instates atmospheric stability. We call this rearrangement atmospheric convection. So far this is no different to our original description – density gradients can cause regions of the fluid to be buoyant, and spontaneously rise. However, an important characteristic of atmospheric convection that sets it apart from convection in other fluids is the presence of water vapour. When water changes phase, a large amount of latent heat is either absorbed or released. This fundamentally changes the thermodynamics of the fluid. Incorporating this effect into the parcel description of convection can result in a situation termed *conditional instability*, in which the atmosphere is stable to small perturbations, but unstable to displacements of finite amplitude (Emanuel, 1997).

Conditional instability is most clearly demonstrated on a thermodynamic diagram known as a skew-T log-p diagram, an example of which is given in figure 1.1. The vertical co-ordinate is the log of pressure, while the other axis is temperature, plotted at and angle of 45° to the horizontal (hence the name skew-T log-p). The bold and dashed trace correspond to the vertical profile of temperature and dew-point temperature (as determined by an atmospheric sounding for instance). The stability of the hypothetical parcel can now be assessed by giving

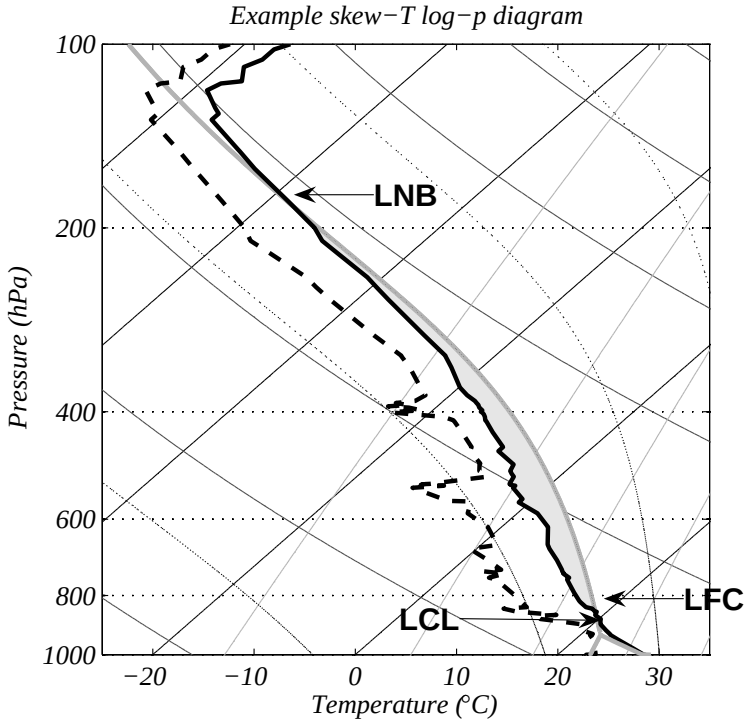

Figure 1.1: A skew-T log-p diagram for the atmosphere over Florida in September. The solid and dashed black lines show the profile of temperature and dew-point temperature as collected by a radiosonde attached to a weather balloon. The vertical co-ordinate is the log of pressure, while isotherms are the straight solid lines, plotted at an angle to the horizontal. Also plotted are lines of constant potential temperature (curved solid), equivelant potential temperature (curved dotted), and saturation mixing ratio (grey). The thick grey lines are the profile of a hypothetical parcel lifted from the surface to assess the stability of the atmosphere. The energy that would be released by this process, the convective available potential energy (CAPE), is shaded in yellow.

it the properties of the sounding at the surface and lifting it adiabatically. At low levels, the parcel is unsaturated, and thus maintains a constant potential temperature, following the curved solid lines on the figure. Assuming it does not mix with the environment, the parcel will also maintain a constant water vapour mixing ratio, and thus the parcel dew point temperature follows the grey lines as the parcel lifts. Eventually, these two lines will meet, and saturation will occur. This is known as the lifting condensation level (LCL). If the parcel rises further, it will begin to condense water, and thus release latent heat. This condensation warming changes the lapse rate so that the parcel now follows a line of constant equivelant potential temperature (the dotted lines on figure 1.1). An example of such a parcel ascent is shown in on figure 1.1. It is important to note that if the parcel is lifted high enough, the change in lapse rate associated with condensation allows it to become warmer then the environment. The level at which this occurs is called the level of free convection (LFC), and above this, the parcel is buoyant and will begin to rise of its own accord. This spontaneous release of energy continues to accelerate the parcel upwards until it again becomes negatively buoyant at the level of neutral buoyancy (LNB). Thus, with the inclusion of the thermodynamic effects of water vapour, it can be seen that the atmosphere can be stable to small displacements of fluid, but still contain large amounts of potential energy that is released with a sufficiently large

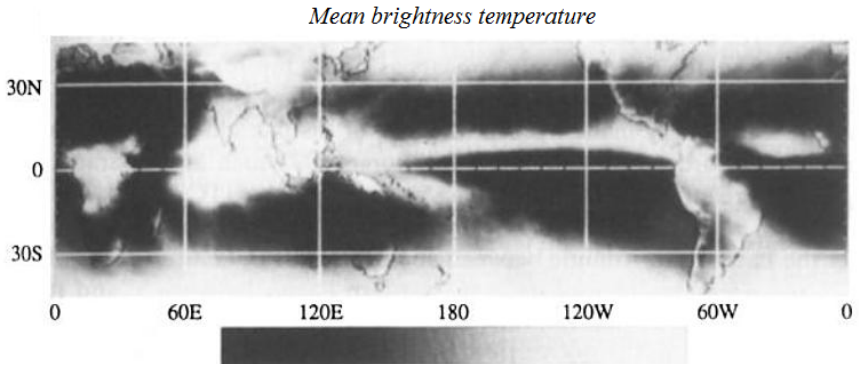

Figure 1.2: Average brightness temperature for August 14 to December 17, 1983. Brightness temperature is a measure of the temperature at cloud top – lower temperatures imply a colder cloud top. Low values of brightness temperature are often used as indicators for the existence of deep convection. Source: Holton (2004).

disturbance. The amount of potential energy that can be released in this way is called the convective available potential energy (CAPE). This quantity is the maximum kinetic energy gained as a parcel rises from its LFC to its LNB, and is proportional to the shaded area on figure 1.1. CAPE is often used as an indicator of how likely, and how intense convective development might be (Wallace & Hobbs, 2006), and, as will be seen, is also a key variable used in the modelling of convection in general circulation models.

The simple analysis we have performed here is obviously highly idealised. The assumption that the saturated parcel follows a moist adiabat implicitly assumes that any condensate is immediately removed as precipitation. In reality, much of this liquid water, and indeed solid ice, is carried upwards along with the parcel, and reduces its buoyancy, altering the amount of energy available. Thus, CAPE as we have defined it is sometimes called irreversible CAPE, since mass is irreversibly removed from the system via precipitation (Emanuel, 1994). The other extreme is to assume all condensate remains with the parcel, slowing its ascent. The energy the parcel can acquire in this case is called reversible CAPE[1]. The parcel method also requires that there is no mixing with the environment. This is also a poor assumption. More comprehensive models of convection must take mixing – entrainment and detrainment – into account. Nevertheless, following a parcel on its idealised ascent through the atmosphere gives a picture of how moist processes allow large amounts of potential energy to build up in the atmosphere, something that cannot occur in a purely dry circulation.

1.2 Convection on the Earth

The distribution of moist convection is far from uniform on the Earth's surface. Figure 1.2 shows the average brightness temperature measured by satellite over a period of four months in 1983. Brightness temperature represents the effective emission temperature of the cloud top, and is thus related to cloud top height. Values less than 240 K are often used as an indicator of

[1]The amount of CAPE diagnosed in this case also depends on assumptions about the freezing of liquid water within the cloud (Emanuel, 1994).

persistent deep convection[2]. The inter-tropical convergence zone, the Asian monsoon, and the position of land masses in the tropics, can be clearly seen to have signatures in the brightness temperature. However, in modelling convection, we are interested in what forces convective activity on a slightly smaller scale – the scale of GCM grid-boxes. Depending on the particular application, these grid-boxes can be on the order of tens or hundreds of kilometers across. On these scales, the amount and strength of convection is seen as being mainly influenced by two parameters; the sea-surface temperature[3] (SST) and the so called large-scale convergence (Tompkins & Craig, 1999). The SST largely determine the local thermodynamic properties of the atmosphere, while wind convergence can be related to mean vertical motion through conservation of mass. The term 'large-scale' in this context refers to scales that can be resolved by the GCM. This terminology implicitly assumes a separation of scales at the GCM grid-box level. That is, we assume that while the motion of individual clouds may not be predictable, the statistical properties of an ensemble of clouds can be predicted (Emanuel, 1997). This thesis will investigate the response of the Unified Model to changes in SST and grid-mean vertical wind. To contextualise these model results we first review some of the observational studies on the relationships between mean vertical wind, SST and tropical convection.

Graham & Barnett (1987) examined the effect of SST and large-scale vertical motion on convection over the tropical oceans. Using satellite observations of outgoing long-wave radiation as a proxy for convective activity they found that SSTs greater than 27.5°C were required for significant deep convection to occur. Below this critical threshold little convection was observed, and above it, the amount of convection was relatively insensitive to further SST increases. Instead, the sensitivity was in the large-scale convergence. Implied upward motion was highly correlated with strong and consistent convection. Under conditions of subsidence, on the other hand, convection was suppressed, even if the SST was above 28°C (Graham & Barnett, 1987). A number of further studies using more comprehensive data sets have also reported a threshold type behaviour in the dependence of amount of convection on SST (e.g. Arking & Ziskin, 1994; Fu *et al.*, 1994), with the change in regime consistently occurring at SSTs of around 27 - 28°C. More recently, however, a view has developed that the effect of SST on convection is in fact weak, and the strongest effect in due to vertical motion (Tompkins, 2001). Lau *et al.* (1997) showed that the sharp increase in convective activity at SSTs of 27°C corresponds to an increase in the number of observations with positive vertical motion that also occurs at this SST. Thus when stratified by vertical motion, the SST effect becomes much weaker.

1.3 Modelling atmospheric convection

One of the biggest challenges in general circulation modelling is that of accurately representing moist convection. The difficulty arises because the building blocks of convection – individual cells or updrafts – are at most a few kilometers across (Emanuel, 1994). Typical GCM resolutions are of the order of tens of kilometers or more, and thus, even with a considerable increase in computational power, convection will remain an unresolvable process[4].

[2]The precise value used as the threshold depends on the time period over which the observations are averaged (Salby *et al.*, 1991).

[3]Over land the relevant parameters are the surface temperature, and moisture availability.

[4]There is, however, some work in embedding cloud ensemble models in GCMs. See Randall *et al.* (2003) for a summary.

This means that the effects of convection are not calculated explicitly by the model, but must be parameterised in terms of the large-scale resolved variables. The task of cumulus parameterisation is thus to relate the distribution of an ensemble of convective clouds that may exist in a grid-box, to the large-scale variables that are explicitly calculated by the model (Arakawa, 1993). In reality, this is a task that involves the entire spectrum of atmospheric physics – from cloud microphysics, to boundary layer meteorology, to meso-scale dynamics. Atmospheric models, however, generally have discrete subroutines for the different time and space scales – the convection scheme, the microphysical scheme or the clouds scheme. Thus, a model's *representation* of convection does not simply depend on its convection *parameterisation*, but all the model physics that interact with it. In this section we will describe the conceptual models that modern convective parameterisations are based on, while noting that the full representation of convection by the model involves all of its physics.

The earliest attempts to consider the effect of convective motions on the atmosphere in large scale modelling were simple adjustment schemes. In the radiative convective equilibrium model of Manabe & Strickler (1964), for instance, convection was assumed to instantaneously adjust the lapse rate of the atmosphere back to a predefined value whenever that value was exceeded by radiative processes. While this is sufficient for calculations of the equilibrium profile under the action of radiation and convection, prognostic models require a more general approach. One such approach, which has had a large impact on subsequent work in the field, is the mass-flux representation of convection. These types of models were first developed by a number of authors in the 1970s (see e.g. Yanai *et al.*, 1973), however, a complete mass-flux parameterisation was presented by Arakawa & Schubert in 1974. Modern implementations of this parameterisation have three components, a trigger model, that determines whether the thermodynamic state of the atmosphere is able to support convection, a cloud model, that calculates the vertical distribution of convective heating, and a closure, that governs the overall strength of the convective response. We will begin by describing the cloud model.

Clouds are described as entraining plumes, with a height dependent mass flux M_i. The total mass-flux within the grid-box is then the sum over all the clouds. Figure 1.3 shows how this is implemented in the model. The change in upward mass-flux between two layers is determined by the amount of entrainment and detrainment within the layer. That is, assuming the plumes are always in a quasi-steady state, we have,

$$\frac{\partial M_i}{\partial z} = \epsilon_i M_i - \delta_i M_i \tag{1.1}$$

where ϵ, and δ, are the fractional entrainment, and detrainment rates respectively. Applying a steady state plume model amounts to assuming that the timescale for individual clouds to form or dissipate is much smaller than the timescale over which the ensemble of convection within the grid-box responds to changing large-scale conditions. Thus within any grid-box there would be a number of clouds, each at a different stage in its life cycle. For modelling purposes, we need only consider cloud properties averaged over the life of a cloud. With the added assumption that the fractional area covered by convection is much smaller than the grid-box, equations similar to (1.1) can be derived for the profiles of temperature, humidity and wind within the cloud. These equations depend only on (1) the entrainment and detrainment rates, and (2) the cloud base mass-flux. The cloud base mass flux is calculated via the convective closure, while entrainment and detrainment are parameterised directly. Arakawa

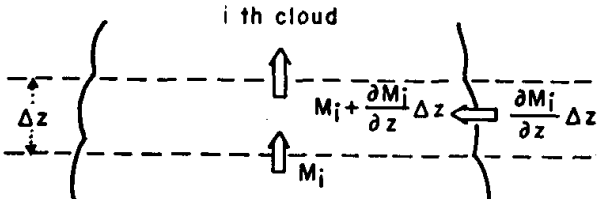

Figure 1.3: Basis of the cloud model in the Arakawa-Schubert parameterisation of convection. The cloud mass-flux, M_i is modified in each model layer by entrainment and detrainment. Source: Arakawa & Schubert (1974).

& Schubert originally proposed a constant fractional entrainment rate, with detrainment only occurring at cloud top (Randall *et al.*, 1997*a*). However their model was based on a cloud spectrum, with many cloud types (with different cloud top heights) each with a potentially different entrainment rate. Due to computational expense, most applications of the mass flux scheme to convective parameterisation problems, use one entraining plume with a variable entrainment rate (Randall *et al.*, 1997*a*).

Using the cloud model, the vertical distribution of convective heating and drying can be deduced in terms of the cloud base mass-flux. It is this mass-flux that must be determined in order to close the model. The convective closure is generally expressed in terms of the energy released by clouds in the process of convection. However, at its root is the concept of convective quasi-equilibrium (QE). The QE hypothesis asserts that the consumption of conditional instability by convective processes close to balances its production by non-convective processes (Emanuel *et al.*, 1994). Specifically, Arakawa & Schubert assumed that the change in the 'cloud work function' with respect to time was much smaller than the changes due to the large-scale forcing alone. The cloud work function, $A(\lambda)$, in their formulation is an integral measure of the buoyancy in an entraining plume with entrainment rate λ. The authors note that this approximation amounts to the assumption that the timescale over which convection adjusts to the large scale forcing, is much smaller than the timescale over which the forcing changes. Thus the convection is always in a state near to equilibrium with the 'slow' processes (Arakawa & Schubert, 1974). The validity of this assumption is obviously dependent on what the temporal and spatial scales which we are considering, as well as exactly we mean by 'non-convective processes' Randall *et al.* (1997*b*).

In most operational models, a single entrainment rate is used rather than a spectrum of cloud types. Thus the variable to which the QE hypothesis is applied is the convective available potential energy (CAPE) (see previous section). Additionally, strict QE, in which the amount of CAPE is constant in not enforced. Instead convection is assumed to relax the atmosphere back to an equilibrium over a specified timescale. Thus the statement of quasi-equilibrium becomes,

$$\left.\frac{\partial \mathrm{CAPE}}{\partial t}\right|_{conv} = -\frac{\mathrm{CAPE}}{\tau_{CAPE}}. \tag{1.2}$$

Here, CAPE is the grid-box mean convective available potential energy, the tendency is that due to convective processes, and τ_{CAPE} is the CAPE timescale – the timescale over which convection brings the environment back to equilibrium. While the QE hypothesis has been verified in observational studies with some success (see Arakawa & Schubert, 1974; Lord & Arakawa, 1980), the CAPE timescale is difficult to constrain. Gregory & Park (1997) note that this timescale must be resolution dependent, as the mean grid-box vertical velocity is also resolution dependent. This is physically realistic – the timescale over which convection

stabilises the grid-box environment should depend on the size of that grid-box.

The third component to modern convective parameterisations based on the mass-flux concept is the trigger model. This consists of an initial determination of the atmospheric stability, to assess whether the convection should occur in a grid-box. Usually an ascent is performed in which a parcel is lifted adiabatically from the surface, and its potential to become positively buoyant and undergo free convection is determined. If this potential is great enough the convection scheme is called to determine the amount and depth of convection. The convective trigger thus prevents spurious convection occurring above a very stable boundary layer, and allows a saving of computational time when the convection scheme is not needed. The details of how the parcel ascent is constructed are highly dependent on the particular model in use. Thus we leave the specifics of the trigger model for the next section where the Unified Model is described in full.

While the mass-flux framework described above is the basis of many modern convective parameterisations, there have been other types put forward. Two important parameterisations are the Kuo scheme (Kuo, 1974) and the Betts-Miller scheme (Betts, 1986). The Kuo scheme bases the strength of convection on the large-scale convergence of moisture within the model. That is, convection is in equilibrium with the supply of moisture by the large-scale flow, rather than the supply of instability as in the Arakawa-Schubert parameterisation. The Betts-Miller parameterisation takes a simpler approach. The temperature and moisture profiles of the atmosphere are relaxed back to observed equilibrium profiles. In this way, the scheme is similar to the very earliest convective parameterisations such as that of Manabe. However in this case, the relaxation profiles are not necessarily neutrally stable, and potentially could depend on the type of convective regime (Betts, 1986).

The appearance of so many different convective parameterisations gives some idea of the difficulty of the task. Other than assuming a statistical equilibrium, or at least quasi-equilibrium, between convection and the large-scale flow on which it exists, the approach of all cumulus parameterisations are rather different. Arakawa (2004) notes that despite the different reasoning used to construct them, the climatologies of convection produced by these parameterisations are comparable – at least if the SST are fixed. He attributes this to the negative feedbacks between large-scale destabilisation, surface temperature and the convection scheme that exist in all modern parameterisations. However, it is his view, and that of this author, that simply producing plausible climatologies is not enough to claim success in the cumulus parameterisation problem. For climate models especially, what is most important is that all the relevant physical processes are accurately represented. The more the representation of convection in numerical models can be justified on physical grounds, rather than simply observational evidence, the more confidence can be given to their applicability in a changed climate. To this end, the cumulus parameterisation problem is one of the most difficult and important issues in atmospheric science today.

Chapter 2

The Unified Model

The purpose of this study is to investigate the convective response of the Unified Model under a set of idealised large-scale conditions. This chapter is devoted to describing the Unified Model and the particular implementation used in this study. In section 2.1 we describe the Unified Model and its convective parameterisation in detail. Section 2.2 gives an overview of the application of single-column modelling to GCM evaluation, as well as describing the particular forcing and initial conditions for the simulations in this study.

2.1 The Unified Model physics

The Unified Model (UM) is a the name given to a suite of Earth-system models used and developed by the U.K. Meteorological Office. In this study we run the UM atmosphere component (version 6.2) in single-column mode. In a model such as this there is no ocean atmosphere interaction – the sea surface temperatures are prescribed, and all large scale dynamics are given as boundary conditions to the column. The active part of the model when run in SCM mode is its physical parameterisations. As convective activity is the main focus of this study we describe the UM convection scheme in detail in the next section. For details of other physical parameterisations, the interested reader is referred to documentation papers compiled by the U.K. Meteorological Office (1999), and the references therein (these papers apply to version 4.4 of the model).

2.1.1 Convective parameterisation

Convection in the UM is parameterised via a mass flux scheme based on that outlined in Gregory & Rowntree (1990). This study uses version 4a of the scheme in which there are three regimes of convection; deep, shallow and mid-level. The mid-level and deep regimes are parameterised the same way, however, the deep convection is triggered by instability to the lifting of surface parcels, while mid-level convection is triggered by instability above the boundary layer. The shallow convective scheme is triggered when an initial parcel ascent indicates convection will terminate below the freezing level, and the large scale flow implies subsidence. The main difference between the shallow scheme and the deep scheme are that the shallow scheme includes more mixing with the environment, and uses the boundary-layer turbulence kinetic energy closure of Grant (2001). In the model simulations used in this

study, subsiding vertical velocities are never encountered, and thus the shallow scheme is never called. The focus instead is on deep and mid-level convection.

As outlined in section 1.3, the mid-level and deep regimes have three components – a trigger model, a cloud model and a closure. The closure relates the cloud-base mass flux to the CAPE within the column by applying equation (1.2) (note that this is not the case in Gregory & Rowntree (1990)). However, in the 4a scheme the timescale τ_{CAPE} is not constant, but depends on the distribution of water vapor in the column, such that,

$$\tau_{CAPE} = \tau_{MAX} \frac{(1 - RH_{av})}{0.4}, \tag{2.1}$$

provided,

$$t_{sub} \;\leq\; \tau_{CAPE} \;\leq\; \tau_{MAX}.$$

Here, RH_{av} is the mass weighted average relative humidity of the cloud environment, and τ_{MAX} is specified. The lower limit, t_{sub} is length the sub-timestep. The convection scheme (along with the boundary layer scheme) in the UM is run multiple times for every model timestep, and we refer to this shorter time interval as the sub-timestep. If the function defined by (2.1) does not lie in between t_{sub} and τ_{MAX}, the CAPE timescale is simply set to either the lower or upper limit. The value of τ_{MAX} used in the majority of this study is 1 hour, and we refer to this as the 'control' value. In section 4.1 the effect of changing this parameter is investigated.

The cloud model of the convection parameterisation in the UM is a based on the entraining plume model outlined in the previous section – but with only one cloud type. Thus, the ensemble of convective clouds is made up of entraining plumes, with imposed entrainment and detrainment characteristics. The entrainment rate depends on height, and, based on the estimate of Simpson (1971) is given by,

$$\epsilon = 3A_E p/p_*^2, \tag{2.2}$$

where ϵ is the entrainment rate per unit mass flux, p is the pressure, $p_* = 10^5$ Pa and A_E is set to 1.5 Pa for all model levels except the lowest, in which $A_E = 1$. Detrainment is produced by two mechanisms – mixing through the edges of clouds, and forced detrainment that occurs when a given cloud becomes negatively buoyant. Mixing detrainment occurs at all levels, with the fractional detrainment rate, η given by,

$$\eta = (1 - A_E^{-1})\epsilon. \tag{2.3}$$

Forced detrainment only occurs in clouds within the ensemble that have reached their level of neutral buoyancy. It is calculated as follows; if, on being lifted from level k to $k + 1$, the entraining parcel has a buoyancy less than a threshold, b_{min}, forced detrainment begins to occur. The threshold b_{min} is dependent on the layer thickness, but is always at least 0.2 K. The strength of the forced detrainment is calculated to give the plume a buoyancy of b_{min}. Thus, forced detrainment represents the termination of weaker members of the ensemble, while allowing the stronger updrafts to continue upwards.

The mass-flux cloud model described above can only be called, however, if either the convection trigger finds the boundary layer suitable for convection, or the mid-level convection scheme finds instability in the free troposphere. The triggering of surface forced convection

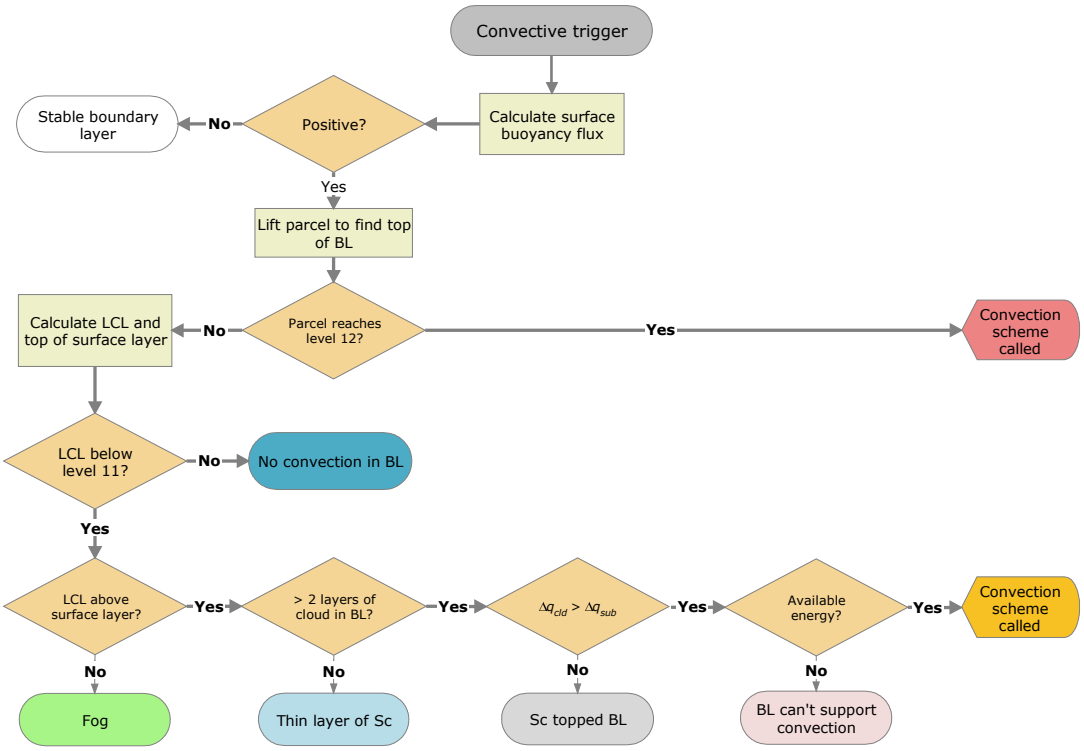

Figure 2.1: Flowchart of the decision tree used to trigger the mass-flux convection scheme.

is based on the boundary layer parameterisation of Lock *et al.* (2000). A test parcel ascent is used to determine the structure of the boundary layer, which is then diagnosed as being in one of seven distinct boundary layer regimes. The different regimes determine the values of the diffusivity profiles for heat and momentum used in the parameterisation. Also diagnosed by the test parcel ascent, is whether or not the boundary layer is cumulus topped. If it is, the convection scheme is called, in either its deep or shallow mode. Note, that the convection trigger is only called once per model timestep, and its diagnosis applies to all sub-timesteps within.

The parcel is initialised with the properties of the environment above the surface layer. This is taken to be the first model level at which the environmental profile is stable to moist adiabatic displacements, or one tenth of the boundary layer height at the previous timestep – whichever is lowest. Before the parcel is lifted, however, the pressure and temperature at the lifting condensation level (LCL) are calculated using an approximate formula derived by Bolton (see Lock *et al.*, 2000), and bulk aerodynamic formulae are used to calculate the surface buoyancy flux, F_B. If F_B is negative, the boundary layer is diagnosed as stable, and convection does not occur. The flux of buoyancy is parameterised by the difference between the surface and lowest model layer virtual potential temperature. As such, it is similar to the surface sensible heat flux (see eq. 3.2a), but with a correction for the effect of water vapour on the density of air. Generally, the sign of F_B is the same as the sensible heat flux, unless the latent heat flux is very large in magnitude in comparison. If the surface buoyancy flux

is positive, the parcel is lifted reversibly in order to find the top of the boundary layer. This is diagnosed via one of two conditions. In clear air the top of the boundary layer is assumed to be the level at which the parcel becomes negatively buoyant by an amount greater than a threshold, θ_{pert}. Like the threshold for forced detrainment b_{min}, this is variable, but has a minimum value of 0.2 K. If the parcel reaches above the LCL, the boundary layer top can also be taken as the level at which the excess buoyancy between of the parcel over the environment is a maximum. Thus instead of using the absolute value of the buoyancy excess, the cloud and environment buoyancy gradients are compared.

Using the information from the parcel ascent, the convective state of the boundary layer can be determined. Boundary layer cloud is diagnosed if the parcel reaches above the LCL. If this cloud is determined to be cumulus, the convection scheme is called, otherwise it is assumed to be stratocumulus, and mixing is performed by the diffusion profiles of the boundary layer parameterisation. For cumulus cloud to exist, there must be at least 2 model layers of cloud, and the cloud base (LCL) must be below model level ten (≈ 1.8 km at the current resolution). Additionally, the lifting condensation level must be above the surface layer, otherwise a fog layer is diagnosed, and convection does not occur. If all these conditions are satisfied, stratocumulus decks are differentiated from cumulus regions by comparing gradients of water content in the cloud and sub-cloud layer. Based on the assumption that stratocumulus cloud decks are relatively well-mixed in comparison to regions of convection, it is assumed that in-cloud gradients of moisture are larger in cumulus clouds than in stratocumulus (Lock *et al.*, 2000). Thus, cumulus convection is only diagnosed if,

$$\left.\frac{\partial q_w}{\partial z}\right|_{cld} > 1.1 \left.\frac{\partial q_w}{\partial z}\right|_{sub}, \tag{2.4}$$

where q_w is the total water content, and the gradients are averages over the cloud (cld) and sub-cloud layer (sub). The exception to this rule is if the parcel reaches the highest model level passed to the boundary layer scheme (currently level 12). In this case, surface processes must affect the atmosphere higher than can be parameterised by the boundary layer scheme, and thus the convection scheme is called in the hope that it will fulfill the required vertical transports indicated by the parcel ascent.

Finally, one more condition must be satisfied if the convection scheme is to be called in a given timestep. The integral of the parcel buoyancy within the cloud layer is required to be positive. That is, a reversible, undilute parcel ascent must result in the net release of available potential energy for the boundary layer to be seen as able to support cumulus convection. A logical flowchart of the conditions required for a cumulus capped boundary layer to be diagnosed is shown in figure 2.1. If the parcel ascent passes all tests, the convection scheme is called, either in its deep or shallow mode. If not, only mid-level convection can exist. The mid-level scheme can be triggered at any sub-timestep by a simple stability criterion. A parcel from model level k is lifted to level $k+1$ pseudo-adiabatically (i.e. irreversibly.). Upon reaching level $k+1$ the mid-level convection scheme is allowed to run if its equivelant potential temperature satisfies,

$$(\theta_e)_{par} - (\theta_e)_{env} > -\Delta\theta, \tag{2.5}$$

where θ_e is the equivelant temperature of the parcel (par) and environment (env) at level $k+1$.

The maximum stability, $\Delta\theta$ is set to 0.5 K. This stability test is applied to all levels above the boundary layer and above any surface driven convection to determine the occurrence of mid-level convection.

The convection diagnosis scheme outlined above will be shown to have a strong influence on the model response at the timestep to timestep scale. The reasons for this sensitivity are examined in detail in chapter 4. In the next section the single-column model is described, along with the particular forcing used to drive the model in this study.

2.2 Description of model simulations

Climate model evaluation has two main purposes; to determine how well a model simulates the real Earth-system – verification, and to analyse how the model could be improved – model development. The most obvious way to perform model verification is to run a model for a certain time period, and compare it to observations over the same time period. Thus the model is directly compared to the 'reality'[1] it is attempting to simulate. This approach, while no doubt useful, has some disadvantages. Firstly, modern general circulation models consume huge amounts of computational resources and time. Thus, only a small number of simulations can be done, and the number of scenarios under which the model can be tested is limited. Secondly, it can be very difficult to determine *why* a model behaves the way it does from such a comparison (Randall *et al.*, 1996). For this purpose, methods that are more focused on particular aspects of the model are needed. One such technique is the use of a single-column model. As the name suggests, a single model simulates only a vertical slice of the atmosphere. A full general circulation model integrates the fluid equations of motion forward in time on a three dimensional grid. Each grid-box has associated with it a set of all the model variables (such as the winds, temperature, pressure etc.) at each timestep. In a single-column model, rather than simulate the whole atmosphere, one column of grid-boxes is simulated. This has the advantages of greatly reducing the computational resources required, and allowing focus to be put on a small number of parameterisations that make up the entire GCM. The main disadvantage is that, since only one column is simulated, the large-scale dynamics must be entirely prescribed. Hence, model response that depends on interactions between the physical parameterisations and the dynamics cannot be simulated. Additionally, a choice must be made on how to impose the conditions on the boundary of the column. In this section we describe the methodology and rationale behind the particular sets of boundary conditions applied in order to study convective activity in the UM single-column model.

The purpose of this study is to examine the convective response of the Unified Model under idealised large-scale conditions. The single-column model is thus an ideal setting, as it allows complete control of these conditions. Along with the winds and the conditions at the sea-surface, the single-column model requires the specification of various advective tendencies of heat and moisture in order to be run. To examine these requirements we follow Randall & Cripe (1999) and consider the transport of an arbitrary scalar variable ϕ governed by,

$$\frac{\partial \phi}{\partial t} = -\mathbf{V} \cdot \nabla_H \phi + w \frac{\partial \phi}{\partial z} + S. \tag{2.6}$$

Here, $\mathbf{V}$ is the horizontal velocity field (u, v), w is the vertical velocity, and S represents

[1] Or at least the best approximation to reality we have.

sources and sinks of ϕ due to unresolved processes. The purpose of the model is to calculate the tendency on the left hand side. The single-column model is capable of calculating the unresolved processes, and, given a vertical velocity, the vertical transports. However, it cannot calculate the horizontal advection terms – these must be prescribed. There are a number of ways in which this can be done. If a sufficiently high quality dataset is available, observations can be used to derive the advective terms directly (called 'revealed forcing' by Randall & Cripe). However, this study aims to investigate convection under idealised forcing scenarios, and thus the horizontal advection terms for moisture and heat are simply set to zero. For horizontal momentum, relaxation forcing, as outlined in Randall & Cripe (1999) is used. Here, rather than explicitly impose the value of the advective terms, their effect is encapsulated in a relaxation back to observations. Thus equation 2.6 becomes,

$$\frac{\partial \mathbf{V}}{\partial t} = -\frac{\mathbf{V} - \mathbf{V}_{obs}}{\tau_{adv}} + S, \tag{2.7}$$

where τ_{adv} is an advective time-scale and $\mathbf{V}_{obs}$ and $\mathbf{V}$ are the observed and modelled values of the horizontal velocity respectively. The use of relaxation forcing has the advantage that, while it allows the model to be run prognostically (i.e. calculating its own tendencies), it also constrains the model not to stray too far from reality. Again, since this study is idealised in nature, we take a slightly simpler approach, and let $\mathbf{u}_{obs}$ be given by the initial velocity profile, which is described below. The advective timescale τ_{adv} is set to one hour.

The two conditions left to specify are the vertical wind profile and sea surface temperature (SST) at each step of the model integration. This study examines the sensitivity of the model's convection scheme to changes in these boundary conditions, and as such a range of different values are used. However, the wind profile and SST is always kept constant within each model run, in an attempt to determine the equilibrium response of the model. The SST is varied between 296 and 304 K, representing the full range of SSTs that can be observed in the tropical oceans. For convenience, the sea-surface temperature will be often referred to by whole numbers of degrees Celsius, despite all model simulations being run at whole numbers of Kelvin. Thus, when referring to a simulation with a SST of 27°C, for instance, the actual value is 297 K, or 26.85°C.

Following Betts & Jakob (2002), the imposed vertical wind profiles are sinusoidal, with a maximum in the mid troposphere, such that,

$$w(z) = w_R \left\{ 1 - \cos\left(\frac{2\pi(p(z) - p_T)}{p_s - p_T} \right) \right\}. \tag{2.8}$$

Here p, p_s, and p_T are the pressure at height z, the surface and model level 25 at the beginning of the run. This model level is at approximately 200 hPa, giving a maximum vertical wind at 600 hPa. w_R is the parameter varied in the sensitivity studies. Values between zero and 0.2 ms^{-1} are examined. These were chosen based on an estimate of the strongest grid-box mean vertical winds observed in the active period of the monsoon during the Tropical Warm Pool International Cloud Experiment (TWP-ICE).

As well as boundary conditions, the SCM must be be given an initial state. The initial profiles of temperature, moisture and horizontal wind were set based on profiles of the tropical atmosphere during the monsoon from TWP-ICE. Figure 2.2 summarises the initial state of the atmosphere. The temperature and moisture profiles are plotted as a skew-T log-p, while

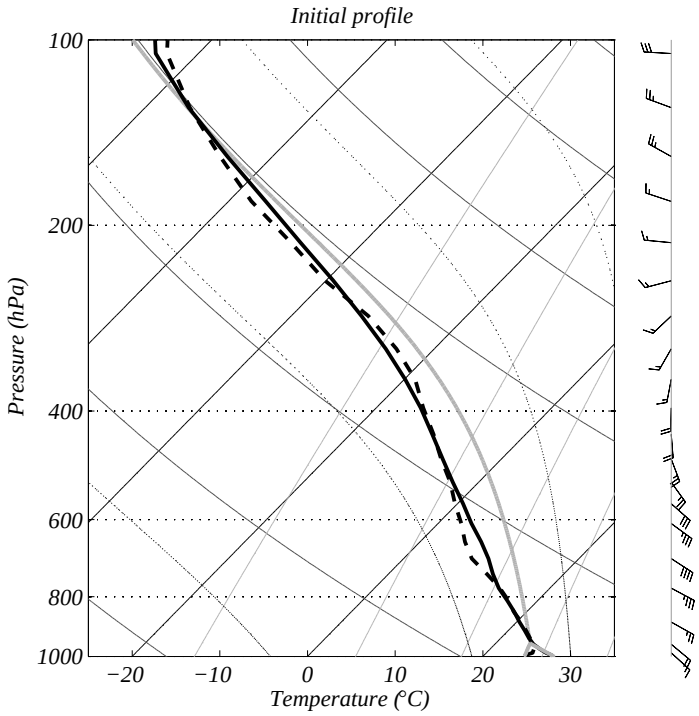

Figure 2.2: Thermodynamic and wind profile of the atmosphere used to initialise all simulations.

the horizontal winds are marked with wind-barbs on the right hand side. Notice that the profile is relatively warm and moist. The temperature profile has an associated sea surface temperature of approximately 28.5°C, and the profile is close to saturation throughout the free troposphere. Thus, especially for the cooler SSTs imposed, the model must equilibrate by cooling the atmosphere, and condensing out some water vapour. This is accounted for by allowing some adjustment time when calculating statistics about the simulations as a whole, however, in some cases remnants of the initial conditions persist throughout the simulation. These effects are discussed in detail in section 3.2.

It should be noted that whilst the horizontal advective tendencies of temperature and humidity have been set to zero, this does not mean there is no implied horizontal transport by the model. This can be seen by again considering the conservation of a scalar ϕ. Without the horizontal advection terms equation (2.6) becomes,

$$\frac{\partial \phi}{\partial t} + w\frac{\partial \phi}{\partial z} = S, \tag{2.9}$$

Integrating over the atmospheric column,

$$\frac{\partial}{\partial t}\int_0^\infty \phi\, dz + \int_0^\infty w\frac{\partial \phi}{\partial z}\, dz = \int_0^\infty S\, dz.$$

By integrating by parts, and noting that the vertical velocity at the surface and top of the

atmosphere is zero, this can be written,

$$\frac{\partial}{\partial t}\int_0^\infty \phi\, dz = -\int_0^\infty \phi\frac{\partial w}{\partial z}\, dz = \int_0^\infty S\, dz.$$

With the simplified mass continuity condition of non-divergence[2], this becomes,

$$\frac{\partial}{\partial t}\int_0^\infty \phi\, dz = \int_0^\infty \phi\nabla_H \cdot \mathbf{V}\, dz = \int_0^\infty S\, dz. \tag{2.10}$$

That is, the total amount of the scalar ϕ within the column can be changed by local sources S, or horizontal wind convergence induced by the imposed vertical velocity profile. Wherever the vertical wind is divergent, and thus removing mass from the grid-box, there must be a compensating convergence in the horizontal wind field. This mass convergence will carry with it tracers such as water vapour. Since, for the vertical wind profiles outlined above, the implied horizontal flow is convergent in the lower atmosphere, and divergent higher up, water vapour will be drawn in at lower levels and ejected aloft. Water vapour mixing ratios are highest in the lower atmosphere, and thus this process results in a net transport of water vapour into the atmospheric column. Thus, for strong vertical motion, this process can maintain rain rates far in excess of the surface evaporation rate.

All model runs were 29 days and 18 hours long (594 hours), beginning on the 19th of January 2006 at 03:00 UTC. The model is run over the Western tropical pacific at 2° N 156° E. Since the sea surface temperatures are fixed, these settings only effect the model through cloud short-wave radiation interactions. The model timestep used in all simulations is 30 minutes, while the boundary layer and convection schemes are run on a sub-timestep of one third of this. A summary of the physical and numerical conditions used in the single-column model simulations is shown in table 2.1.

[2]This form of the conservation of mass equation is only strictly true for incompressible fluids. However, we use it here simply to demonstrate the coupling between the horizontal and vertical motion fields.

Condition	Formula	Value
Physical conditions:		
Beginning of run	t_0	26th January 2006 03:00 UTC
Length of run	t_{len}	24 days 18 hours (594 hours)
Timestep	Δt	30 minutes
Sub-timestep	t_{sub}	10 minutes
Latitude	lat	-2°N
Longitude	lon	156°E
Initial Conditions:		
Temperature	T	profile derived from TWP-ICE data
Humidity	q	profile derived from TWP-ICE data
Horizontal wind	(u, v)	profile derived from TWP-ICE data
Vertical wind	w	same as boundary condition
Sea surface temperature	SST	same as boundary condition
Large-Scale Forcing:		
Horizontal wind profile	(u, v)	relaxation to initial profile over timescale of one hour
Horizontal heat advection	$\mathbf{V} \cdot \nabla_H T$	zero
Horizontal moisture advection	$\mathbf{V} \cdot \nabla_H q$	zero
Vertical wind	w	sinusoidal profile with maximum in midtroposphere of magnitude w_R ms^{-1}
Sea surface temperature	SST	constant value, between 296 - 304 K

Table 2.1: Summary of forcing data used to run the single-column model.

Chapter 3

Results

In the previous chapter, model simulations under large-scale conditions representing an idealised tropical atmosphere were described. In this chapter, the results of these model integrations are shown, with particular reference to features relating to convective activity. First, an example simulation using moderate vertical motion and sea-surface temperature is presented to describe some of the features common to all runs. In section 3.2 the model response from a time averaged point of view is considered. The effect of changing the SST and vertical forcing on average rainfall, evaporation and surface fluxes is analysed. Time dependent model behaviour is examined in the next chapter, in which the convection trigger is implicated in driving much of the timestep to timestep model variability.

3.1 Typical forcing

Figure 3.1 shows the precipitation as a function of time for the case in which the sea-surface temperature is fixed at SST = 27°C, and the vertical wind maximum is $w_R = 0.06$ ms^{-1}. The blue line and green shading correspond to deep and mid-level convective precipitation respectively– i.e rain produced by the model's mass flux convection scheme. The red line is *large-scale* rain – that produced by grid-box mean saturation. The top panel shows the entire 25 day simulation, with output plotted at each half-hour timestep. The most obvious feature to note is the strong variability in rain rates between timesteps. Convective events have a magnitude that is at times in excess of 200 mm day^{-1}, while the mean precipitation rate is only 44 mm day^{-1}. In this case, and throughout this thesis we take the mean as an average over the last fifteen days of the simulation using the notation μ_{15}. This neglect of the beginning of the run is required to limit the effect of the initial conditions on the results. The temperature profile used to initialise the simulations is an arbitrary tropical profile, which has an associated SST of 28.5°C. When different sea-surface temperatures are used, it follows that the model must begin with a thermodynamic state far from equilibrium. A period of approximately ten days is taken to be the timescale over which the initial profiles adjust to the imposed boundary conditions. This can be seen in figure 3.2 for the example simulation. Apart from some regions of the atmosphere near the tropopause, the temperature structure is relatively stable after 240 model hours. There are of course still fluctuations due to the short timescale variability in convection, but the temporally averaged structure is relatively steady. Note that this does not mean the model has reached an equilibrium. As will be

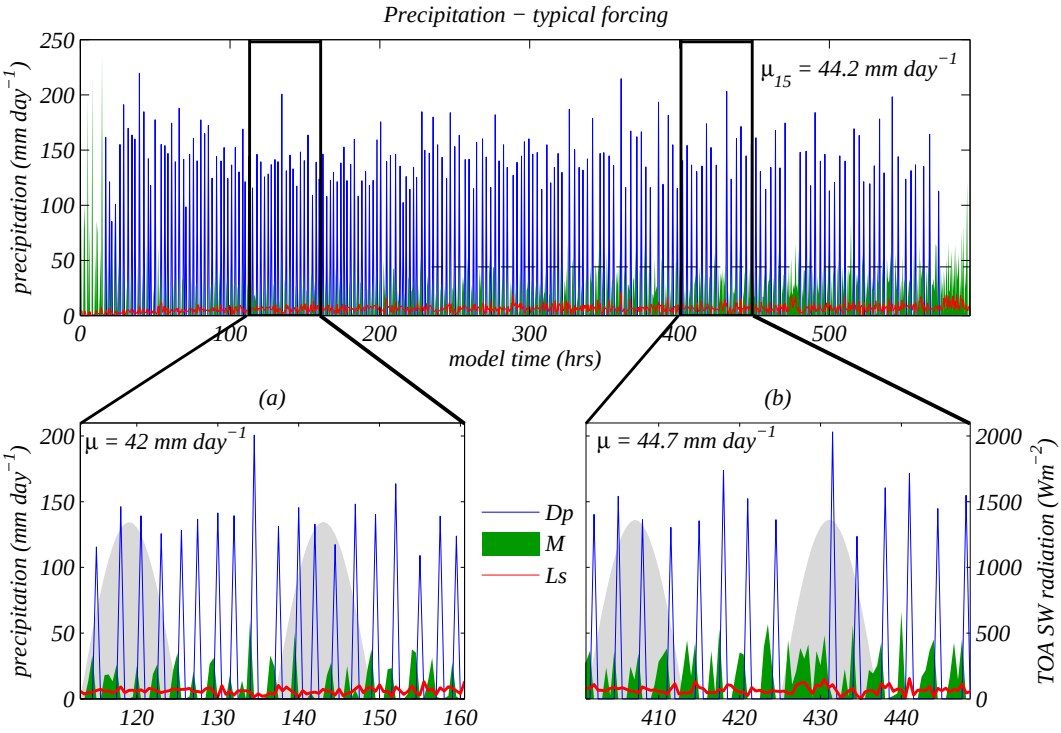

Figure 3.1: Time series of precipitation in the model for the case; $w_R = 0.06$ ms^{-1} and SST = 27°C. The main panel shows precipitation at each time-step of the model integration. This is broken into three components – rain due to deep convection (blue), mid-level convection (green) and large-scale precipitation (red). Insets (a) and (b) show details of two 48 hour periods. Also shown in the insets is the top of the atmosphere solar flux, shaded in grey.

discussed later, the model does not always reach a strict equilibrium, in the sense that all model variables achieve statistical stationarity. Similarly, especially for cases of weak forcing, equilibria can be transient, and regime shifts are possible. Notwithstanding these caveats, this method will be used throughout this thesis when referring to averages over simulations.

The insets in figure 3.1 each show rainfall over two days of the simulation in detail. Both periods show quasi-periodic behaviour in the deep convection, in which single timestep bursts of convective activity are followed by several timesteps in which there is no deep convective rain. In between the deep convective bursts mid-level convection occurs, bringing some precipitation to the ground. The mid-level convection is less noisy, and smaller in magnitude than the deep bursts, but also has considerable variability. These model results imply an ensemble of convection being almost non-existent for some timesteps, before switching on for thirty minutes, while the large-scale conditions remain constant. For grid-boxes on the order of 100 km across, this is not realistic. Even more importantly, however, this behaviour is not consistent with the paradigm of quasi-equilibrium. Under this framework, a situation with constant large-scale forcing would be expected to yield convective activity that is also relatively constant between timesteps. The model instead responds more akin to what one could expect from an individual convective cloud – the atmosphere becomes conditionally unstable in the suppressed periods, and then releases the built up potential energy relatively rapidly.

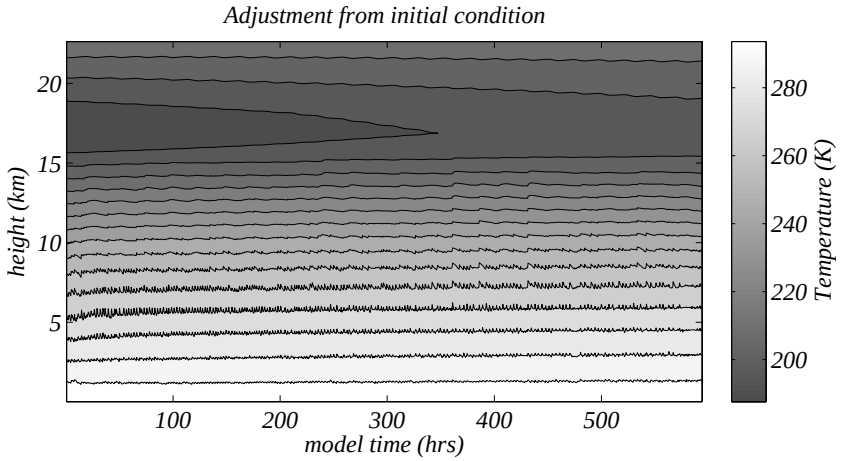

Figure 3.2: Time height cross-section of temperature for the case, $w_R = 0.06$ ms^{-1}, SST = 27°C. Convective motions allow the lower atmosphere adjusts to the sea-surface temperature faster than the upper troposphere, which is governed more by radiative processes.

This perhaps suggests that the inherent spatial scale associated with the convection scheme is smaller than the grid-box we are attempting to simulate. Physically, this scale is set by the magnitude of the vertical velocity given to the model. The SCM itself, however, also has a scale associated with the size of the ensemble – the CAPE closure time-scale. Thus, the fact that the model responds as if its inherent spatial scale is too small could be due to an unrealistically small value of τ_{CAPE}. This hypothesis is investigated in section 4.1.

Comparing the two insets of figure 3.1, it can be seen that the frequency of the cycle in convective activity changes. In (a) the mean period is approximately 2.5 hours, while in (b) it is closer to 3.5 hours. On close inspection of the main figure, this change in regime can be seen to occur at around model hour 230. Thus, even with constant forcing, the model can respond with surprising qualitative changes of behaviour. These regime changes are more common at weaker forcing; a dramatic example is the case of zero vertical wind, and a sea-surface of 30°C shown in figure 3.3. The model clearly has two periods with which the convection can oscillate, and switches between them with little perturbation. Again, interpreting this through the framework of an ensemble of convective clouds implies an extreme sensitivity of the statistical properties of the ensemble on an almost constant large-scale forcing. If we are to accept that the ensemble properties of convection are inherently predictable, then clearly such sensitivities should not exist in reality. Note also, that the average rain rate (including both convective and large-scale precipitation) is almost equal for the periods in which the model is in each regime. Similarly, in insets 3.1a and 3.1b the rain rate is relatively constant. This is to a large degree a consequence of the imposed boundary conditions. This can be seen by considering the conservation of total column water content, written in the form of equation (2.10),

$$\frac{\partial}{\partial t}\int_0^\infty q_w\, dz = \int_0^\infty q_w \nabla\cdot\mathbf{V}\, dz + \int_0^\infty (E-P)\, dz, \tag{3.1}$$

where q_w is the total water content, P is the precipitation, E is the surface evaporation[1],

[1]Evaporation of liquid water within the atmosphere does not change the *total* water content, only the fraction that is in liquid phase.

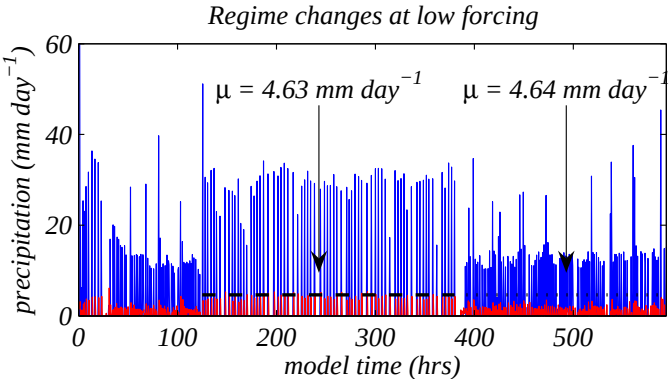

Figure 3.3: As in main panel to figure 3.1 but for the case $w_R = 0.0$ ms^{-1}, SST = 30°C. Two different convective regimes can clearly be seen. However, the total rainfall rate in each is almost equal.

and the integrals are over the entire atmospheric column. Now, the evaporation largely depends on the surface wind-speed and the sea-surface temperature, which are both held relatively constant by the boundary conditions. Similarly, conservation of mass requires that the horizontal wind divergence is strongly coupled to the the vertical wind profile, which is prescribed. Thus, there can only be a significant change in precipitation as a result of changes to the distribution of water vapour in the column. Evidently, this distribution is tightly constrained (for a given set of boundary conditions at least), and the water vapour profile cannot vary too greatly.

The insets in figure 3.1 also show the top of the atmosphere solar flux, shaded in grey. In neither of the plots does the diurnal cycle appear to have a significant effect on convective activity – the majority of the variability is associated with much shorter periods. This has been confirmed by running selected cases under perpetual night conditions. In these simulations some variability is reduced – periods of quasi-periodicity become more regular, for instance, but in general the qualitative conclusions remain the same. The largest effect is on points of regime change. These do occur under perpetual night conditions, but with reduced frequency. We thus retain the diurnal cycle in all simulations, whilst noting that it results in some extra variability.

The features outlined above are common to many of the different combinations of sea-surface temperature and vertical wind tested in the model. Almost all simulations displayed highly variable convection rates at the timestep to timestep level. Many showed quasi-periodic behaviour, while at low wind forcing regime changes within model runs were observed. The causes of some of this behaviour are examined in detail in chapter 4. In the next section, these temporal characteristics are set aside, and the effect of changing forcing on the time averaged properties of the simulations is considered.

3.2 Time averaged response

3.2.1 Precipitation

Figure 3.4 shows the dependence of the mean convective and large-scale rain rates over the parameter space of sea-surface temperatures and vertical forcing. The model was run for

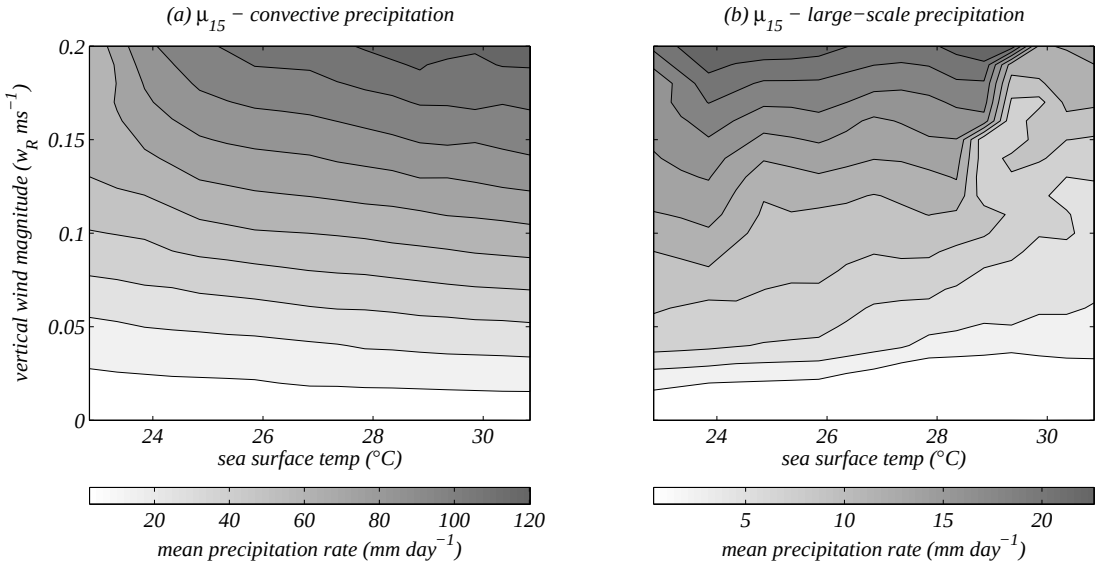

Figure 3.4: Contour plots of the mean convective rain rate (a), and large-scale rain rate (b), over the parameter space of simulations. The means are taken from the last fifteen days of each simulation to allow for some spin-up time of the model.

values of w_R between 0 and 0.2 ms^{-1}, at intervals of 0.02 ms^{-1}, and sea-surface temperatures of 296 - 304 K, at intervals of 1 K. At SSTs above 300 K, additional simulations were performed at each 0.5 K and 0.01 ms^{-1} to resolve some of the features in this parameter region in more detail. The average rain rates for each simulation were calculated using the final 15 days of output. These μ_{15} values were used to create the contour plots in figure 3.4. The main trend in convective rain rates is that due to vertical wind; stronger ascent rates are associated with higher convective precipitation rates. The effect of SST is smaller, with high SSTs implying stronger convection. These two trends are remarkably linear for most of the parameter range tested. Only at the low sea-surface temperatures – 23 - 24°C and strong ascent is there a significant change. Under these conditions, the effect of vertical forcing becomes weak, and the amount of convection is driven by the SST. It is very rare, however that sea-surface temperatures of 23°C are correlated with strong vertical motion (Bony *et al.*, 1997), and thus this region of the parameter space would be rarely encountered in observations, or presumably, full general circulation model integrations. The general model behaviour is consistent with observational studies which have shown that tropical convection is forced primarily by vertical motion, with a smaller effect due to SSTs (see section 1.2). However, the model does not capture the threshold behaviour seen in many of these studies, in which convective activity increases sharply when the SST changes from 26 - 28°C. Given the difficulty in characterising this behaviour in observations, and the idealised nature of the simulations in this study, this should not necessarily be interpreted as a model deficiency. From the discussion of the column water budget in the previous section, it could be argued that the dependence of the model on vertical winds is entirely determined by the boundary conditions. This is to some degree true – the only way the model could have a rain rate not proportional to the mean convergence is if the column moisture distribution is fundamentally changed. At the very least then, figure 3.4a confirms that the moisture structure within the model remains relatively realistic. Additionally, the column water budget described by equation (3.1) relates to total

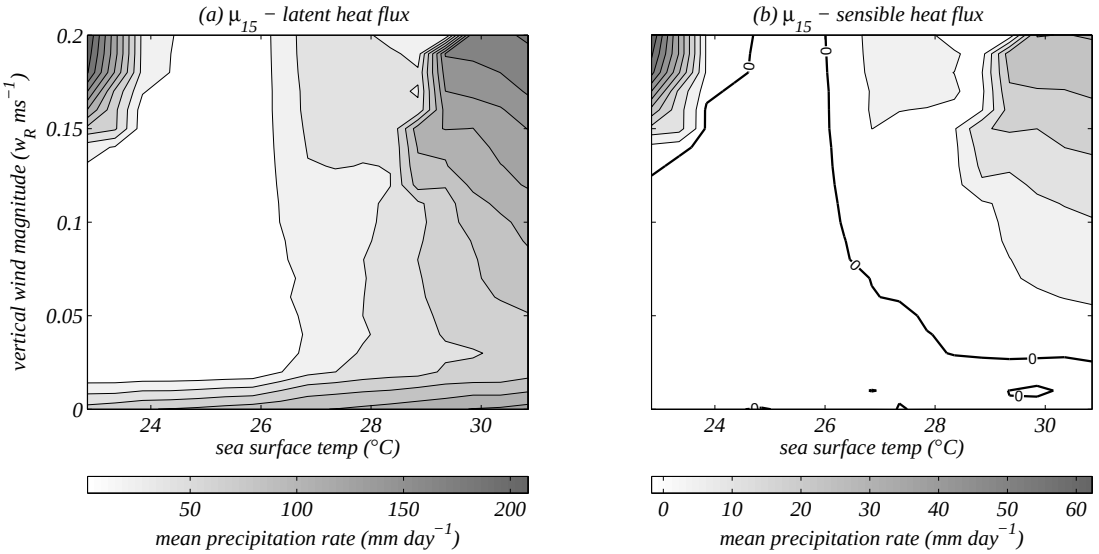

Figure 3.5: Contour plots of the mean (a), latent and (b), sensible surface flux over the parameter space of simulations. Positive fluxes are upwards.

precipitation. If the model's convective regime was not producing enough precipitation, this could be offset by the large-scale rain rate. This kind of behaviour is seen in the model when the CAPE timescale is increased in section 4.1.

The average large-scale precipitation (panel (b)) shows a much less coherent dependence on the SST and vertical wind. Like convection, there is an increase of large-scale precipitation with ascent for SSTs less than 29°C. This is unsurprising – upward vertical motion transports air with low potential temperature and high specific humidity upwards, thus cooling and moistening the free troposphere. Saturation of the grid-box mean humidity is thus strongly correlated to vertical motion. The dependence on sea-surface temperatures is more complicated. At lower SSTs there is little dependence, however, at around 29°C, there is a strong decrease in large-scale precipitation. This feature is also observed in the surface fluxes shown in the next section, but remains barely detectable in the convective rain rates. The cause of such a regime change is not clear at this time. Further experimentation is needed to understand this model behaviour.

3.2.2 Surface fluxes

The Earth's surface can be both a significant source or sink in the atmospheric energy budget. Energy transport at the surface can occur through sensible heat flux – direct transfer of heat – or latent heat flux – transfer of energy embodied in the phase of water. Positive values of the sensible and latent heat fluxes correspond to heating and moistening of the lowest model layer by the surface. In atmospheric models these energy fluxes generally calculated via simple bulk aerodynamic formulae (see Garratt, 1992, ch. 3). In the Unified Model, the sensible heat flux, F_H is taken to be,

$$F_H = -\rho K_H C_p \left(s_1 - s_{surf}\right), \tag{3.2a}$$

where, ρ is the air density at the surface, C_p is the specific heat of dry air at constant pressure, and s_1 and s_{surf} are the dry static energies[2] of the first model layer and the surface respectively. Dry static energy is conserved for dry, adiabatic, hydrostatic adjustments. If the surface layer is statically stable, s will increase with height, while, more commonly over warm oceans, the surface layer is super-adiabatic, and the resultant sensible flux is positive. The exchange co-efficient K_H depends primarily on the wind speed, and also on the thermodynamic structure of the surface layer. Since the surface winds are constrained tightly by the relaxation forcing, and the surface temperature is fixed, F_H is effectively determined by the temperature of the lowest model level, and more generally the boundary layer thermodynamic structure. Similarly, under this bulk transfer formulation, the latent heat flux, F_L is determined via,

$$F_L = -\rho K_H L_{vap} \left((q_w)_1 - (q_w)_{surf} \right). \tag{3.2b}$$

Here, the relevant gradient is that of total water content q_w. For an ocean surface, the surface humidity is taken to be the saturation value of the surface temperature. Again, since this is fixed, the main dependence left in the equation is on the specific humidity of the first model layer.

Figure 3.5 shows the trends with SST and vertical forcing of the mean surface fluxes, calculated similarly to figure 3.4. As is to be expected for an ocean surface, the magnitude of the latent heat flux is substantially greater than the sensible heat flux under most conditions. The relative changes in each panel, however, are rather similar. The main difference is that, for low forcing values, there is a decrease of average latent heat flux with increasing vertical wind, while this trend does not exist in the sensible heat flux. Both surface fluxes show a sharp increase associated with the increase in large-scale rain described in the last section, although this is most pronounced in the latent heat flux. Both fluxes also have a sharp increase at very low SSTs and high vertical forcing. This is possibly related to the change in regime seen in the convective rain in this region of the parameter space.

Given that an increase in surface temperature gives both an increase in dry static energy and saturation specific humidity, one might expect that there should be a strong correlation between the surface fluxes and the SST. However, the mean flux shown in figure 3.5 is measured after the model has had a ten day adjustment time. Thus, it is the way in which the near surface layer adjusts to this surface temperature that determines the magnitude of the surface fluxes. It is surprising then, that there exists a large region of the parameter space in which the sensible heat flux is negative. The zero contour is marked in bold in figure 3.5b, and it can be seen that beneath an SST of 26°C, the average sensible heat flux is negative, regardless of the vertical forcing. This implies two things. Firstly, by equation (3.2a), the potential temperature of the first model layer must be greater than that of the sea-surface. To maintain this situation, the near surface layer must be fed by a source of heat from above, as it is constantly losing energy to the surface. Secondly, unless the latent heat flux is very large, the buoyancy flux must also be negative, thus forcing the convective trigger to shut off deep convection. In fact, many of the simulations with negative average sensible heat flux have no deep convection at all – all convective adjustment is performed by the mid-level scheme.

This stable surface layer profile is somewhat artificial, and likely to be due to the initial conditions. The initial thermodynamic profile is of a warm moist atmosphere overlying an ocean surface of 28.5°C. Thus at the beginning of the simulation, the lower atmosphere would

[2]Dry static energy is given by $s = C_p T + gz$, where gz is the geopotential.

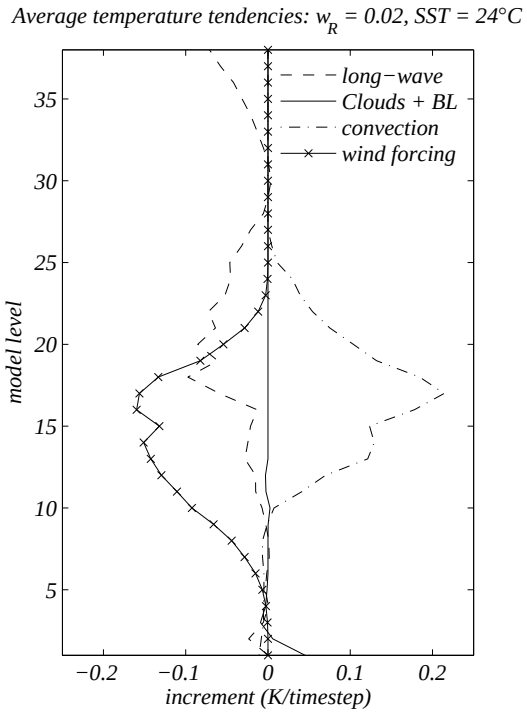

Figure 3.6: Mean (μ_{15}) temperature increments due to various processes in the atmosphere as a function of model level. The simulation shown is one that maintains a stable surface layer throughout the run. Particular processes shown are long-wave cooling (dashed), the combination of clouds and boundary layer processes (solid), and convection (dashed dot). the solid line with crosses corresponds to cooling by vertical advection of low potential temperature air by the imposed wind field.

be much warmer than the surface, and the surface layer highly stable. A more plausible steady state might be reached if the initial atmospheric state was more consistent with the surface temperature. Nevertheless, from a model evaluation point of view, these stable surface layer simulations are interesting, as any convection must occur through the mid-level scheme. Comparing these simulations with simulations in which the deep scheme is called allows an analysis of the differences in behaviour of the mid-level and deep convective parameterisations. The cloud model of these two schemes are identical, and hence differences in behaviour occur due to the triggering mechanism. This is discussed in the next chapter.

The other interesting question that arises from the stable surface layer cases is of the processes that maintain such a high temperature in the lowest model layer. Experiments with perpetual night simulations show that some of the warming is caused by solar radiation. When the sun is turned off, the region in parameter space over which the average sensible flux is negative shrinks, but does not disappear. Thus, while some simulations appear to be maintained in a stable state by warming of the atmosphere via short-wave radiation, others remain even in the absence of this heat source. Figure 3.6 shows the mean (μ_{15}) temperature increments due to various processes in the model for a simulation with a consistently stable surface layer (time series for this simulation are shown in figure 4.1). Temperature tendencies due to convection, the large-scale uplift, and long-wave cooling are shown. The solid line is the sum of the heating due to boundary layer processes, and clouds (both convective and

large-scale). It is clear from the figure that that it is this term in the budget that is responsible for heating the lowest layers of the atmosphere. This low-level heating appears at first to be due to boundary layer mixing. However, it should be noted that diffusion by the boundary layer scheme can only redistribute thermal energy, it cannot warm the atmosphere as a whole. Thus the temperature increment due to boundary layer processes must sum to zero over the whole profile. It then, that most of the low-level warming is due to the condensation warming of clouds. As will be shown later (figure 4.1) there is indeed substantial amounts of cloud at the lowest levels of the atmosphere for these stable surface layer cases. If this heat source is to act throughout the simulation, it requires a source of moisture. This can come either from the surface directly, by latent heat flux, or through the convergence of moisture implied by the vertical wind field (see section 2.2). It appears that both of these sources contribute to maintaining a stable surface layer. In simulations with no vertical wind (and therefore no moisture convergence) a stable surface layer cannot be maintained throughout the entire run. However, when the vertical wind profile was altered so that the lowest five model layers did not converge moisture, but the rest of the profile did, a stable surface layer was maintained. The conclusion is thus, that positive latent heat fluxes, and vertical motion within the boundary layer are ultimately what maintain a stable surface layer! This odd situation is no doubt partly due to the rather pathological initial condition used – a very warm, moist atmosphere above a cool SST. Indeed this configuration seems to be a rather unstable one. In simulations close to the zero contour on figure 3.5b the simulation often begins with a stable surface-layer. If the buoyancy flux is allowed to become slightly positive, however, deep convection is able to occur, destroying the stable layer for the rest of the simulation. It is thus unlikely that this behaviour would occur significantly often in the full GCM, or indeed the SCM with observationally based forcing. It would nevertheless be worthwhile to confirm this hypothesis with further modelling studies.

Chapter 4

Triggering of convection

In the previous chapter it was shown that convective activity in the single-column model has an 'on-off' nature. Oscillations can be set up in which bursts of convection occur that last one timestep, followed by a number of timesteps in which convection is suppressed. Here, we examine this behaviour in detail. Two hypotheses for the reasons for this variability are tested. In section 4.1 we examine the closure to the convection scheme, and its role in producing variability in the convective rain rates. Specifically, the CAPE timescale is lengthened, and the effect on convection is examined. Section 4.2 analyses the behaviour of the convective trigger, and the conditions under which deep convection is allowed to occur. These considerations lead to the examination of a specific case in section 4.3 that appears to be strongly sensitive to the CAPE timescale and the trigger, in order to understand the interaction between the two. Before this, however, the noisy nature of convection in the model is documented for a variety of forcings.

Figure 4.1 shows time series of four different model simulations with combinations of weak and strong forcing (w_R = 0.02 & 0.14 ms^{-1}), and low and high SST (SST = 24 & 30°C). The upper panel shows precipitation for the entire model simulation, broken up into deep-convective (blue), mid-level convective (green shading) and large-scale rain. Beneath this panel are three insets that focus on a particular 24 hour period in the simulation. Panel (a) repeats the precipitation information of the main panel, with the addition of the solar flux shaded grey. Panel (b) shows the surface latent and sensible fluxes, and panel (c) shows a height time cross-section of the *area cloud fraction*. This is a prognostic variable, produced by the models cloud scheme. It corresponds to the area of each grid-box covered by cloud at any given time.

None of the simulations presented show a strong dependence on the diurnal cycle. This confirms the claim made in chapter 2 that solar heating is a smaller source of variability than the internal model variability itself. The four simulations show the diverse behaviour produced by this internal variability. Both simulations at the lower SST have consistently negative surface fluxes, and as a result are dominated by mid-level convection. At low vertical forcing this convection is intermittent, with some timesteps convecting and others not. With a higher ascent rate, however, convective activity, while variable in magnitude, occurs every timestep. Thus the mid-level scheme responds to stronger forcing by both an increase in the frequency and intensity of convection. This is in contrast to the deep convection parameterisation. In the high SST cases, convection is slightly less frequent at low forcing. However the bulk of the

increase in mean rain rate is due to the increase in magnitude of individual convective events. Even at very high vertical ascent rates, deep convection displays on-off behaviour, while small amounts of mid-level convection fill the gaps in between. This quasi-periodic behaviour is characteristic of the deep scheme in most of the simulations of this study. Thus we see that the mid-level regime has less variability than the deep – especially at high forcing. Since the difference between the deep and mid-level convection parameterisations in the UM is simply in the triggering conditions, this suggests that the convective trigger is certainly part of the reason for such variability.

The surface fluxes have quite different characters for low and high sea-surface temperatures. At high SST, the latent heat flux is strongly positive, with a much smaller sensible heat flux. In general, the sensible heat flux is greater than zero, however, in the particular low forcing case shown here F_H is in fact negative for a significant portion of the simulation. There appears to be a strong correlation between the variability in the latent heat flux and convective events. When convection occurs, the latent heat flux increases between 5 and 10%, with a similar fractional increase in sensible heat. This may be a response to the convective activity – since both convective and boundary layer schemes are run on sub-timesteps, they are able to respond to each other 'instantly'. Another possibility is that the change in fluxes is associated with the change in boundary layer regime that occurs during convective events. When convection occurs, the boundary layer must be in a cumulus capped regime (see section 2.1), and this changes the way heat transport is parameterised through the boundary layer, with a possible effect on the fluxes.

At low SST, the fluxes are more comparable in size. The surface sensible heat flux is consistently negative, corresponding to the 'stable surface layer' cases discussed in the previous section. Timestep to timestep variability in the fluxes is not as strongly correlated to convection, and is not related to changing of boundary layer regime (the boundary layer is in its stable regime throughout). A large portion of this variability is related to the partitioning of fluxes between sensible and latent. The dashed line on the surface flux plots in figure 4.1 correspond to the net flux – sensible plus latent. Especially for the weak forcing case, this is much less variable than the fluxes themselves. Thus the total surface flux is relatively constant, while the ratio of latent to sensible heat flux – the Bowen ratio, varies significantly. It is not clear what the cause of these fluctuations in the Bowen ratio is, and further study is needed to understand it.

The third panel of the insets in figure 4.1 show time-height cross-sections of the *area cloud fraction*. This is the fractional area of each grid-box covered by cloud. All simulations display cirrus decks from 10 km to the tropopause ($\approx$ 17 km) with almost 100% of the sky covered. Since this feature exists for all simulations, it is most likely a remnant of the moist upper troposphere of the initial condition (see figure 2.2). With no large scale vertical motion in this region of the atmosphere, these clouds, once produced, can remain present for the entire simulation. In the the high SST cases the upper level cloud is broken by plumes of low cloudiness caused by the deep convection scheme. Convection causes precipitation, and as a result significantly dries the atmosphere. Hence, convective activity *decreases* the amount of cloud in the grid-box – at least where the convective clouds are entraining. In the detraining part of the clouds near their top, moist air is being ejected, and convection contributes to moistening.

Beneath the upper level cirrus, the cloud fraction varies considerably both between and within each run. For the high SST cases, variability associated with the drying by convective

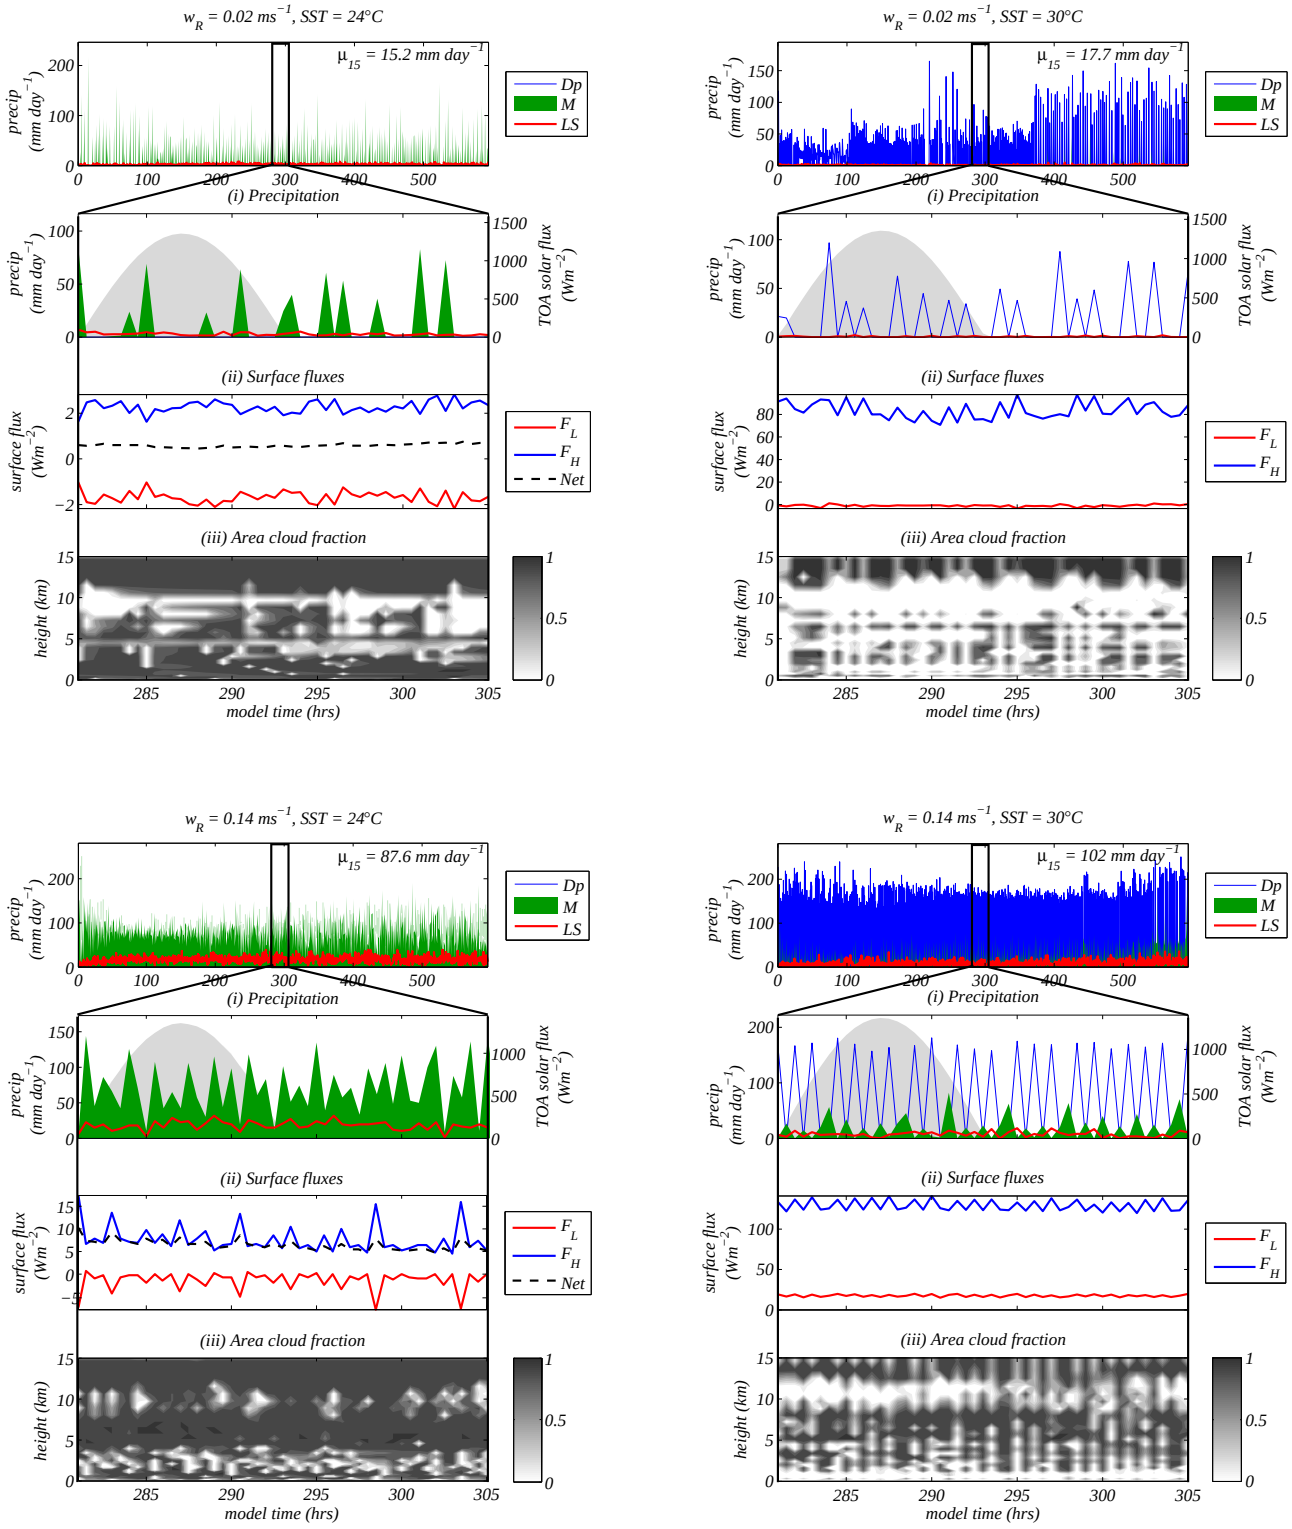

Figure 4.1: Time series of model simulations under conditions of weak and strong ascent, as well as low and high SST. See text for details.

events dominates. The low SST cases undergo only mid-level convection, which has less effect on the cloud field. As is to be expected, the vertical forcing also has a strong effect on cloudiness. Strong upward motion induces adiabatic cooling, and thus increases the grid-box mean relative humidity. This increases the amount of *large-scale cloud*, which, like large-scale rain is caused by the gird-box mean humidity reaching a threshold value, rather than the effects of parameterised convection.

The case of $w_R = 0.02$ ms^{-1} and SST = 30°C has a layer like vertical structure of clouds in the vertical. Between 0 – 10 km there are alternating regions of high and low cloud fraction. This layering does appear to some degree in the thermodynamic structure of the atmosphere during this period. However, the effect is slight, and it is more likely that the cause of the cloud layering is the mass-flux convection scheme, and its parameterisations of entrainment and detrainment. This is purely speculative, however, and more study into this behaviour is needed to understand it fully.

4.1 Effects of CAPE closure timescale

The mass-flux convection scheme in the Unified Model uses a closure based on convective available potential energy (CAPE) to determine the strength of convection that occurs. As outlined in section 2.2, the amount by which convection heats and dries the atmosphere per timestep is effectively determined by the CAPE timescale, τ_{CAPE}. The recharge-discharge type cycles of convection described in previous sections suggest the effect of convection per timestep may be too high. That is, the CAPE timescale may be too short. In this section we investigate this hypothesis, and analyse the effect of increasing the CAPE timescale.

Simulations were performed over the parameter space of w_R and SST for three different CAPE timescales. Figure 4.2 shows results from a representative group. The panels each show the precipitation for one full simulation. As in previous figures the blue line is deep convection, the green is mid-level convection, and the red is large scale rain. Note, however, that in order to show the changes in variability the mid-level convection is plotted as a line rather than shaded as was done previously. Each row corresponds to one forcing scenario. The simulations shown are for the same boundary conditions as figure 4.1, representing all the combinations of weak and strong ascent, with high and low SST. Each column corresponds to the use of a different CAPE timescale. The first column, is using the original CAPE closure that has been used thus far. In the second column, the value of the parameter τ_{MAX} in equation 2.1 is set to three hours, rather than one – the CAPE timescale is tripled. In both of the first two columns the CAPE timescale is governed by equation 2.1, and has an actual value much less than τ_{MAX}. As an example, for the case of an SST of 30°C, and a weak vertical wind ($w_R = 0.02$ ms^{-1}) the average value of the CAPE timescale is 24 minutes – not even a full timestep. For stronger vertical forcing, the vertical transport of water vapour results in a more saturated profile, and thus the CAPE timescale is even shorter. For $w_R = 0.14$ ms^{-1} τ_{CAPE} is only 11 minutes! For this reason, a third group (shown in the third column) of simulations were performed, this time without the models relative humidity based CAPE timescale. Instead, a constant value of $\tau_{CAPE} = 3$ hours was used. The results for this group are comparable to using the relative humidity based timescale with τ_{MAX} of the order of 6 - 8 hours.

The results summarised in figure 4.2 describe a varied response to changing the CAPE

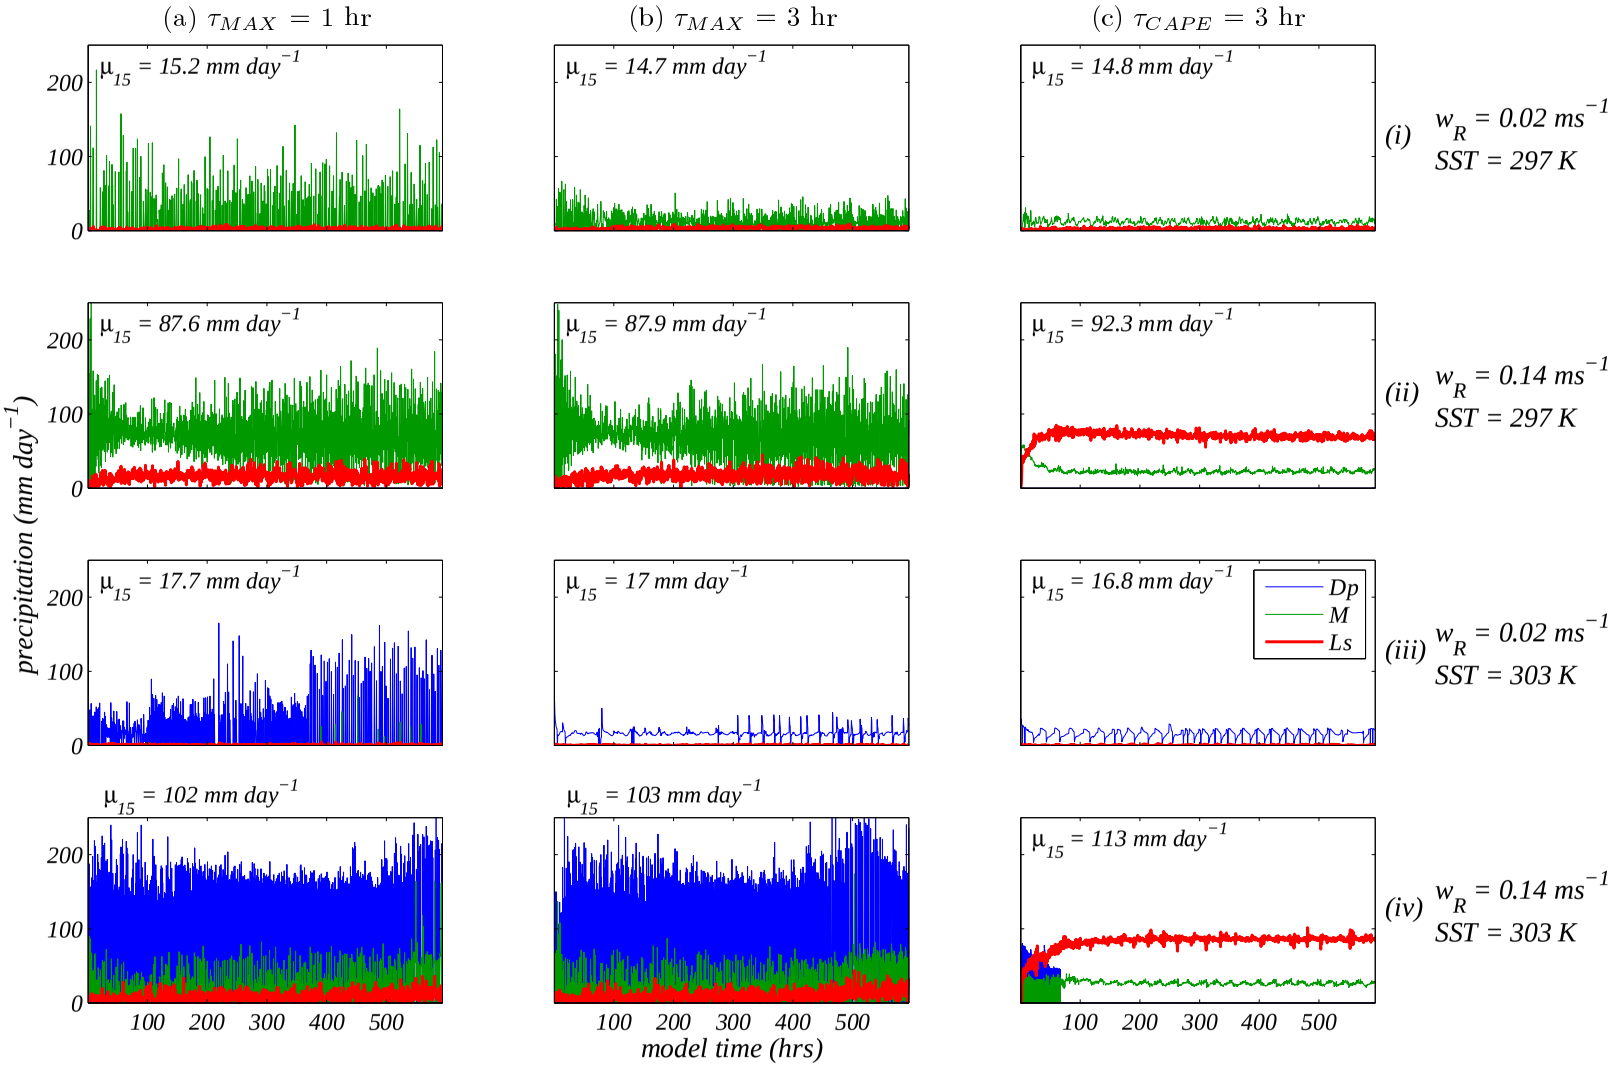

Figure 4.2: Convective and large scale rain rates for four simulations with different CAPE timescales. Rows correspond to individual simulations – strong and weak forcing, with low and high SST – while each column is a different value of τ_{CAPE}. The first two columns use the models relative humidity based CAPE timescale diagnosis, while the third sets a fixed value of τ_{CAPE}. Precipitation is broken down into deep convection (blue), mid-level convection (green) and large-scale rain (red).

timescale. The two solutions at low SST (i & ii) both remain in the stable surface layer regime (see previous chapter) for all values of τ_{CAPE}. In the low forcing case, the mid-level convection shows a consistent reduction in variability as the CAPE timescale is increased, and a corresponding decrease in the intensity of individual convective events. The average rainfall, as has been discussed previously, must remain relatively constant, as it is constrained by the supply of moisture from the large scale forcing. For the higher forcing case (ii), there appears to be almost no difference between the first two panels, while the third shows a dramatic change. The variability in the mid-level convection, as well as its average value decrease, while the magnitude of the large scale rain increases significantly. The lack of difference between panel (aii) and (bii) can be explained with reference to the equation governing the CAPE timescale (2.1 on page 10). At high forcing, the convergence of moisture into the column is strong, and the atmosphere is close to saturation. The value of τ_{CAPE} calculated by the relative humidity based scheme is often smaller than the sub-timestep of $t_{sub} = 10$ min. The CAPE timescale is thus set to t_{sub} for a significant portion of the model run. When τ_{MAX} is increased to 3 hours, the value of the relative humidity function triples, but due to the lower limit of t_{sub}, the CAPE timescale itself increases only marginally. When the CAPE timescale is fixed to three hours, the jump is much larger, and hence the difference in character of the solution is also greater. The significant increase of large-scale rain for the high CAPE timescale case can be understood in terms of the boundary conditions. At high forcing, there is a significant convergence of moisture within the column which must be removed by precipitation. A long CAPE timescale limits the amount that can be precipitated by the convection scheme per timestep. Hence another mechanism must remove the moisture in the model – namely the production of large-scale rain. This response shows that the CAPE timescale selected is probably too long, at least for this particular model setup.

Apart from a few notable exceptions, the high SST cases (iii) and (iv) behave similarly to the cooler sea-surface runs. The rainfall rates are smoothed out as the CAPE timescale is increased, with the average rainfall remaining relatively constant. Again, at high forcing, the difference between panels (a) and (b) is small, while the large scale rain increases dramatically in (c). Interestingly, at strong ascent, and high SST – conditions most conducive to convection – we see an example where the atmosphere transits *to* a stable surface layer (civ). In all the cases seen in the previous chapter of this behaviour, the stable surface layer was the result of a rather unrealistic initial condition. Here, the atmosphere is initialised with a sea-surface temperature that is thermodynamically consistent, and yet it still manages to revert to a rather unphysical state. The most likely explanation for this is that the long CAPE timescale does not allow sufficient pumping of moisture vertically via convection, and thus the surface layers become saturated. The condensation of this low level cloud is enough to stabilise the surface layer, which then shuts off deep convection, feeding back on the entire process. This again relates to the issue of convection not being strong enough when such a long CAPE timescale is imposed. It seems that, at least for the implementation of the model used, a CAPE timescale of three hours is too long.

Panel (biii) of figure 4.2 also presents interesting results. In this simulation, the modelled convection loses its 'on-off' character almost entirely. The rain rate becomes relatively constant, with some variability occurring every 20 hours or so. This slow period occurs also in the longest CAPE timescale case (ciii). Perpetual night simulations show that this periodicity is an interaction between the diurnal cycle and a natural model period of 18 hours. The relatively smooth convective rain rate shown in case (biii) does not occur in other sim-

ulations with this CAPE timescale. This case, along with its corresponding simulation with $\tau_{MAX} = 1$ hour, allow for an examination of the causes of the high frequency oscillations in the convective activity of the model. In section 4.2 we compare these two simulations to try to understand why one oscillates wildly and the other does not.

It has been shown that increasing the CAPE timescale does indeed reduce the timestep to timestep variability associated with convection in the UM single-column model. However, the CAPE timescales required for this to occur are many times the value used in the control simulations. This would no doubt change substantially the behaviour of convection in a full dynamical simulation. Even in the cases above, where the boundary conditions limit the freedom of the model problems with these very long CAPE timescales can be seen. At high forcing, the long CAPE timescale limits the amount of convection, and therefore convective precipitation that can occur. The column water budget, however, requires that the implied convergence of moisture by the ascent profile be either balanced by precipitation, or cause large changes to the distribution of moisture within the column. The vertical profile of convective moistening is determined by the bulk cloud model – this remains the same in all simulations. Thus, large changes to the moisture profile in the model do not occur. Instead, the strong convergence of moisture leads to precipitation via the large-scale cloud scheme. This can be seen as a sharp increase in the large-scale rain rates for the high forcing cases. Interestingly, in both the high forcing, long CAPE timescale cases there is no deep convection after the initial adjustment period. In both cases the surface layer becomes stable, thus preventing deep convection, and remains that way for almost the entire simulation. This effect has been discussed at length for the cool SST cases, but it is very surprising to see it at such a high sea-surface temperature. Clearly, increasing the CAPE timescale so much has large effects on the behaviour of the model, in addition to removing the short period variability. These effects do not, in themselves, suggest that a CAPE timescale of three hours is too long. If one parameter in a model is changed so dramatically, it will no doubt require a change in other parameters whose value was set based on the model response with a short CAPE timescale. However, care must be taken not to create other deficiencies in order to solve this one problem. Notwithstanding these difficulties, the original value of τ_{CAPE} used in this study is likely to be too short. Under strong moisture convergence, the CAPE timescale is often no bigger than the sub-timestep of ten minutes. this is clearly not desirable on numerical grounds let alone physically. If the timescale under which convection adjusts to the large-scale flow is truly so short, then either the timestep should be decreased, or the CAPE closure implemented differently. One such possibility would be to use a strict equilibrium formulation. This involves assuming a complete invariance of the convective available potential energy, rather than a relaxation over some timescale. Thus instead of using equation (1.2) the closure condition becomes that convection and large-scale processes exactly balance in the production and destruction of CAPE. This kind of closure is that originally proposed by Arakawa & Schubert (1974). Whether a strict equilibrium is assumed, or not, better observational constraints on the true timescales on which ensembles of convection adjust would no doubt improve our ability to model atmospheric convection.

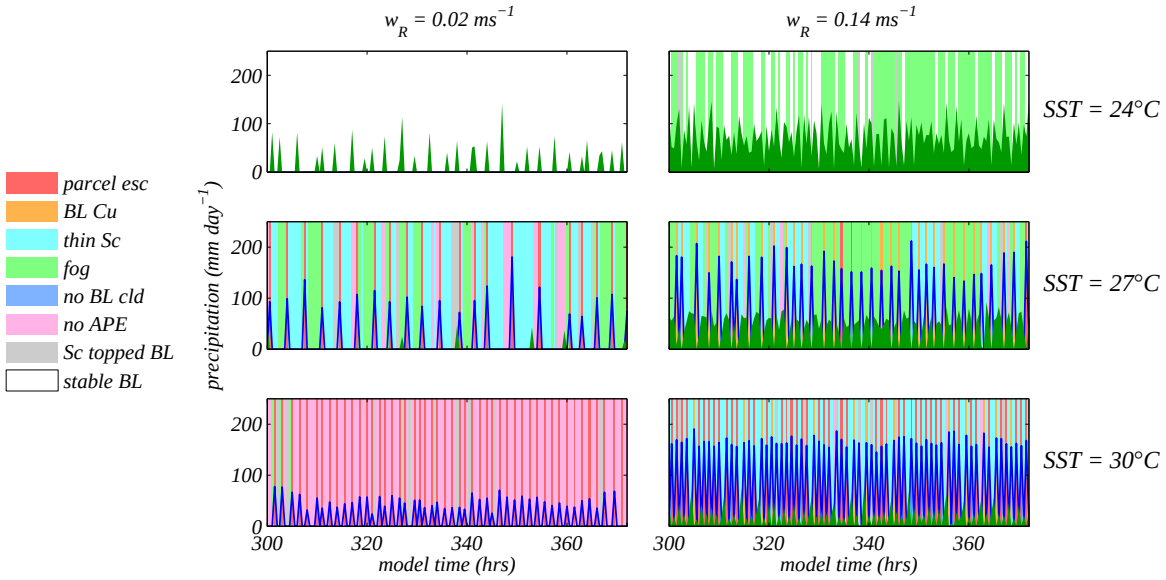

Figure 4.3: Outcome of the convection diagnosis routine for arbitrary three day periods in nine simulations. The columns correspond to weak and strong vertical forcing at three different SSTs. In the foreground the rain rates are plotted in the same fashion as figure 4.1. The background colour denotes the determination of the boundary layer conditions by the convective trigger. Colours correspond to the colours in the flowchart on page 11.

4.2 Diagnosis of a convective boundary layer

In the last section, it was shown that one possible reason for the highly variable convective rain rates seen in the model is that the convection scheme is overactive. This over-activity can be corrected by increasing the timescale over which convection is assumed to adjust to the large-scale forcing. However, another possibility is that the mass flux determination is not flawed, and rather it is that the convection scheme is not called when it is appropriate. That is, convection is not triggered when the atmospheric conditions are able to support it. Investigating the behaviour of the convection trigger model is the subject of this section.

The trigger mechanism for deep convection is outlined in detail in section 2.1. The flowchart on page 11 shows the logical sequence of tests that are undertaken to determine whether the boundary layer is cumulus topped, and thus whether deep convection will be triggered. There are two ways in which this can occur. The first is via a series of conditions on a test parcel ascent through the atmosphere to the inversion. The parcel must reach its LCL, and have a layer of cloud at least two model levels thick. Additionally, the cloud layer must be sufficiently heterogeneous so as not to be diagnosed as stratocumulus, and there must be a net energy release by the parcel as it condenses. The second pathway to trigger surface driven convection is if the parcel reaches above level 12 in the model. This condition exists simply because the boundary layer scheme is only active on the bottom 12 model levels, and thus represents a limitation of combining the boundary layer and convection diagnosis schemes. Nevertheless, it will be shown that this condition is satisfied as often as the parcel being diagnosed as convective 'legitimately'.

Figure 4.3 summarises the behaviour of the convection trigger for different ascent rates

and SSTs. Each of the six panels corresponds to an arbitrary three day period in a standard 25 day run. Both weak and strong ascent is represented ($w_R = 0.02$ & 0.14 ms^{-1}), as well as sea-surface temperatures of 24, 27 and 30°C. In the foreground precipitation rates of the two types of convective rain – deep (blue) and mid-level (green) – are plotted, while the background colour denotes the outcome of the convection diagnosis routine. Each colour corresponds to a different end state of the flowchart in figure 2.1. Red and orange represent convection occurring by the parcel reaching level 12, or by the model diagnosing boundary layer cumulus respectively. Cyan represents a thin layer of stratocumulus (Sc) – if only one model level is cloudy beneath the top of the parcel ascent then it is automatically diagnosed as Sc. Green corresponds to an LCL below the top of the surface layer, and thus a diagnosis of fog. A blue colour indicates that the LCL is higher than the top of the parcel ascent. Pink represents convection being shut off because the buoyancy integral I_B defined by equation (4.1) is negative. A grey background indicates the cloud layer was diagnosed as well mixed, and hence assumed to be Sc, while a white background indicated the surface buoyancy flux is negative, and hence the surface layer is stable. Figure 4.3 thus, allows us to examine not only when convection occurs, but *why* the model did or did not diagnose convection at any given time.

At low ascent and an SST of 24°C the stable surface layer regime that has been discussed in detail previously can be seen. When the ascent rate is increased, deep convection is still suppressed, but generally not because the surface layer is stable. Instead convection is not allowed to occur because the cloud base (LCL) is below the top of the surface layer. Thus, rather than surface driven convection, a 'fog' layer is diagnosed. While a layer of cloud just above the surface is consistent with the diagnosis by the model's cloud scheme (see panels (iii) in figure 4.1, it is not likely to be realistic over such a warm SST. The prevalence of this saturated layer could be in part due to the initial conditions. Placing a cool sea surface temperature below a warm, moist atmosphere would initially cause saturation of the lowest levels of the atmosphere. The trigger will then diagnose fog, preventing ventilation of the system via deep convection, thus limiting the ability of the atmosphere to disperse moisture. However, a low LCL is also a significant preventer of convection in the moderate SST cases (27°C). Here, surface forced convection does occur at some timesteps, and thus ventilation is not an issue. In this case, single timestep bursts of convection are separated by a number of timesteps of 'fog'. This implies that either convection moistens the boundary layer through evaporation of rainfall so much that it takes five or six timesteps to recover. Or, more plausibly, the diagnosis of fog by the convective trigger is very sensitive to small changes in the near surface environment. We do not seek to answer this specific question in this thesis, and instead consider in more detail the high SST simulations.

At high SST (30°C) the number of timesteps between convective events shortens considerably. Both at low and high ascent rates there is a quasi-periodic oscillation with a period between two and three timesteps. Toward the end of the period shown, the low-forcing case in particular exhibits a very regular cycle of convective activity that repeats every three timesteps. The cycle begins with a deep convective event, most often triggered by the parcel reaching above model level 12. The next two timesteps are then non-convective, but for different reasons. The first is a 'thin stratocumulus' case. Here, there is only one layer of cloud beneath the top of the parcel ascent. In general, stratocumulus decks are differentiated from cumulus clouds by comparing in-cloud and sub-cloud gradients in humidity. When there is only one model level that is cloudy, this is impossible, and the diagnosis defaults to stratocu-

mulus, shutting off convection. On the second suppressed timestep, the convection scheme is not called due to a lack of potential instability. The convection trigger requires that for cumulus cloud to exist, the net energy release by the parcel ascent must be positive. This condition is expressed as a requirement that the buoyancy integral,

$$I_B = \int_{LCL}^{z_h} (\theta_\rho)_{par} - (\theta_\rho)_{env}\, dz \tag{4.1}$$

is positive. Here θ_ρ is the density potential temperature[1], and the integral is from the lifting condensation level to the top of the boundary layer. Since the buoyancy is allowed to be slightly negative before the parcel ascent is terminated, in a very neutral boundary layer, a situation could arise in which a parcel is negatively buoyant for almost its entire ascent, but still rises above the LCL. The condition in (4.1) is designed prevent this from resulting in the diagnosis of cumulus cloud. In the next section we compare the parcel and environment at the three different timesteps in this particular convective cycle, in order to understand the changes that cause this cyclic behaviour to occur. While this analysis is rather specific, applying to only part of one single simulation, it will provide insight to the triggering mechanism and how it is so readily able to produce this type of behaviour in almost all the simulations.

4.3 A three timestep convection cycle

Figure 4.4 shows details of the three timestep convection cycle referred to in the last section. The upper panel displays the same information as figure 4.3. Panel (b) shows the model level corresponding to the of top of the parcel ascent (green), the LCL (red), and the top of the surface layer (blue). The dotted line is at level 12, the highest model level that the boundary layer scheme has knowledge of. The third plot shows the value of the integral in (4.1), which is related to the parcel's available potential energy[2]. It can be seen that throughout the period shown, convection is only triggered by the 'parcel escape' mechanism. In non-convective timesteps the parcel reaches level four or five (250 - 410 m), while on a convecting timestep the top of the ascent is at the tropopause near model level 30. This suggests that while the model level 12 cutoff is somewhat arbitrary, the convection trigger is not sensitive to precisely how many levels the boundary layer scheme acts on, and that there is a real change in the parcel or its environment when convection occurs. Whether this insensitivity applies in general is not so clear. Over land, one can observe boundary layers of the order of 3 - 4 km – well in excess of the 2.6 km that corresponds to model level 12. Under these conditions it may well be that the precise height of the cutoff becomes important.

The top of the surface layer, and the lifting condensation level, on the other hand are constant, at level two and three respectively, throughout the period shown. Note, however, that if the surface layer grows, or the LCL drops by a single model layer, convection will be shut off due to fog. The LCL is calculated via an approximate formula based on surface layer characteristics, rather than included in the parcel ascent. Thus, the trigger is strongly sensitive to a variable that is calculated in a rather simple way. It is not unlikely, therefore,

[1]This is the potential temperature of dry air that would have the same density as air containing some amount of water (both liquid, solid, and vapour). See Emanuel (1994) for details.

[2]The integral itself is not an energy – it has units Kelvin meters. However for a relatively well mixed boundary layer I_B is proportional to the available potential energy of the parcel.

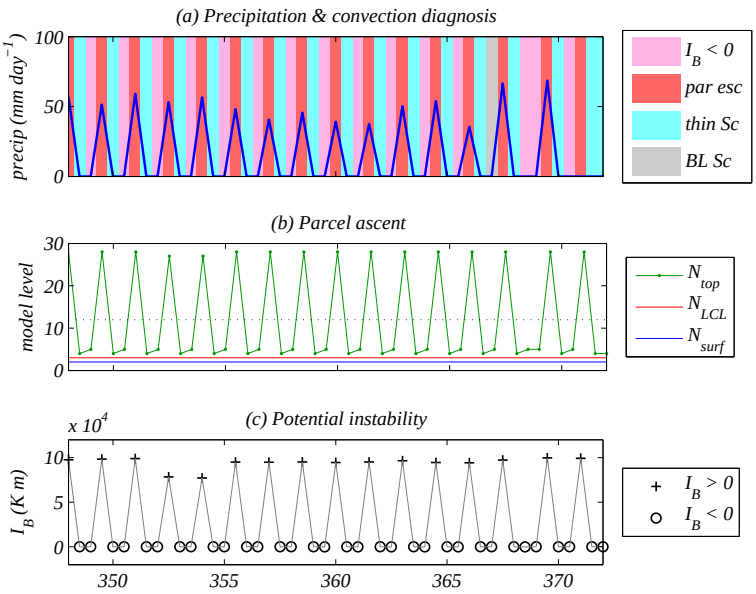

Figure 4.4: Details of the regular three timestep cycle that occurs in the case $w_R = 0.02$ ms^{-1}, SST = 30°C. Panel (a) shows the convection diagnosis in the same way as figure 4.3, (b) displays the model level corresponding to the top of the parcel ascent, N_{top}, the lifting condensation level, N_{LCL}, and the top of the surface layer, N_{surf}. The third panel shows the value of the buoyancy integral I_B.

that when the model diagnoses a fog layer, the intrinsic atmospheric profile may be quite suitable for convection, at least in the warm SST simulations of this study.

Another possible sensitivity of the trigger mechanism is the 'thin Sc' condition. When only one model layer of the parcel ascent is cloudy, the gradient of total water content within the cloud cannot be calculated, and the model is forced to diagnose stratocumulus (and thus no convection). In figure 4.4 this condition appears every three timesteps. However, each time, at the next time step the cloud layer has thickened to two model levels, and subsequently passes the gradient test. Thus, cloud that would otherwise be diagnosed as cumulus, is not due to the resolution of the model. In the period shown in figure 4.4 this is inconsequential. I_B is negative both when the cloud is one or two model layers thick. This may not always be the case however, and such a sensitivity could be responsible for some of the variability seen in the model.

Figure 4.5a shows the parcel temperature excess as a function of height for the three timesteps in the cycle. Plotted is the difference in parcel and environment density temperature, T_ρ, averaged over the timestep before, after and during convection for twelve cycles. Density temperature is the temperature required for a dry parcel to have the same density as a moist, cloudy parcel of air (Emanuel, 1994). Thus each line in the figure corresponds to the parcel buoyancy as it is lifted through the atmosphere. It can be seen that at the timestep in which convection occurs, the atmosphere is close to neutral toward the parcel. Immediately after a convective event, the atmosphere is stabilised, especially around model level 6. The model then takes two timesteps to recover, before being able to convect again. The differences between the three curves are small – on the order of 0.2 K. Nevertheless,

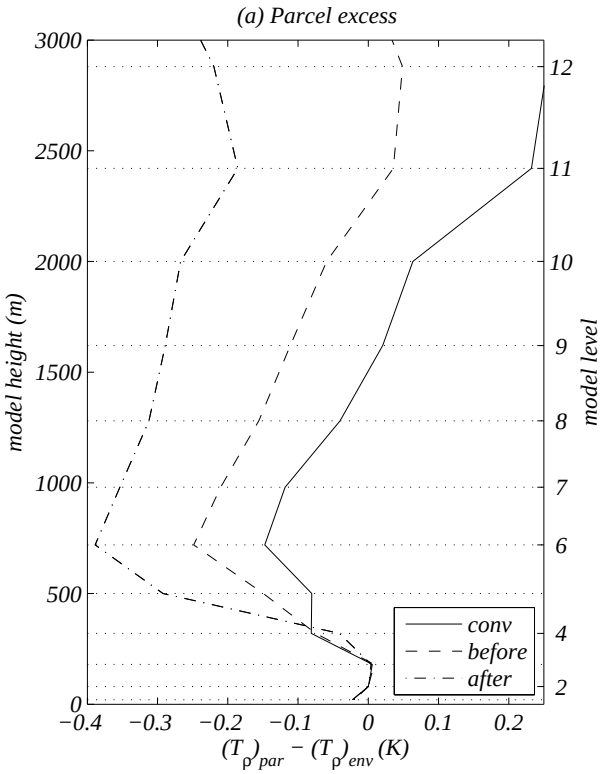

Figure 4.5: Excess buoyancy of the parcel over the environment. Twelve of the cycles shown in figure 4.4 were used to find the average density temperature excess for the convective timestep (conv) and the timestep immediately before and after in the cycle.

these small differences cause the convective parcel to rise well into the upper troposphere, while in the timestep before and after, the ascent terminates below 500 m. This result has the potential to be quite sensitive to the details of the parcel ascent, and in particular, the amount of negative buoyancy the parcel is allowed to have before the ascent terminates. This value, denoted θ_{pert} is generally variable, but for the time period in question, remains fixed at 0.2 K. From figure 4.5a it can be seen that this would only need to be increased slightly, and the parcel before the convective timestep would rise above level 12, allowing convection to occur. Doubling θ_{pert} would allow all timesteps to be convective. The value of θ_{pert}, which is essentially the perturbation to the surface potential temperature within boundary layer thermals, is not well constrained observationally or theoretically. In fact, even in the UM, the value used in a similar diagnosis in the mid-level scheme is 0.5 K. Thus, the sensitivity the model exhibits to this parameter, is not ideal, and may be responsible for much of the timestep to timestep variability.

While the changes in the parcel ascent are small between the convective and non-convective timesteps, it is still interesting to analyse the forcing that creates them. Figure 4.6 shows separately how the parcel and environment change through the three timestep convection cycle. Plotted on the left is the parcel density temperature at the non-convective timesteps, as a perturbation to the convective parcel profile. Thus a negative value indicates the parcel

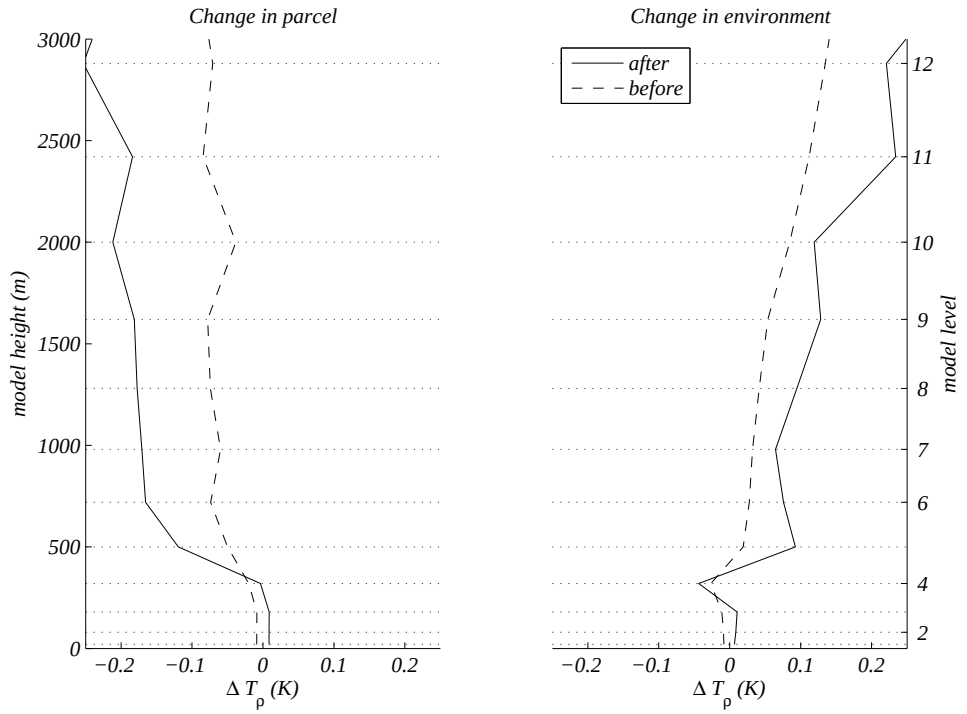

Figure 4.6: Difference in (a) parcel temperature, and (b) environment temperature between the convective timestep and the non-convective timestep. The solid line is the average difference between the density temperature in the timestep after convection compared with the convective timestep. The dashed line is similar, but for the timestep before convection.

is cooler (denser) in the non-convective timestep than the convective timestep. The figure on the right is similar, but for the environmental temperature. The reason for the change in parcel ascent is clear. In the profile immediately after convection, there is sudden increase in environment buoyancy, and a corresponding decrease in parcel buoyancy around model level four and five. This is exactly where the ascent terminates in the non-convective timesteps. Thus it appears as if the parcel and environment contribute roughly equally to the extra stability encountered in the timestep immediately after convection. The pattern for the second non-convective timestep in the cycle is similar, but less pronounced, with a slight bias toward the effect of the parcel.

It can seen in figure 4.6 that the temperature near the surface is almost constant throughout the period in question. The parcel density temperature is initialised with the environmental value at the top of the surface layer (level 2), and hence this value is also very similar for each ascent. As the parcel lifts, its buoyancy is changed by two processes – adiabatic cooling, and condensation. The adiabatic cooling term is essentially the same for any parcel. Hence the changes with height seen in figure 4.6a must be due to condensation. The reason for the change in parcel buoyancy thus must be a change in humidity of the near surface layer. This can be seen in figure 4.7a. The profile of humidity increments, averaged over the convective timesteps, is shown for the relevant physical processes. Convection causes a strong drying of the lowest four layers of the atmosphere. The result is that, in the timestep after convection,

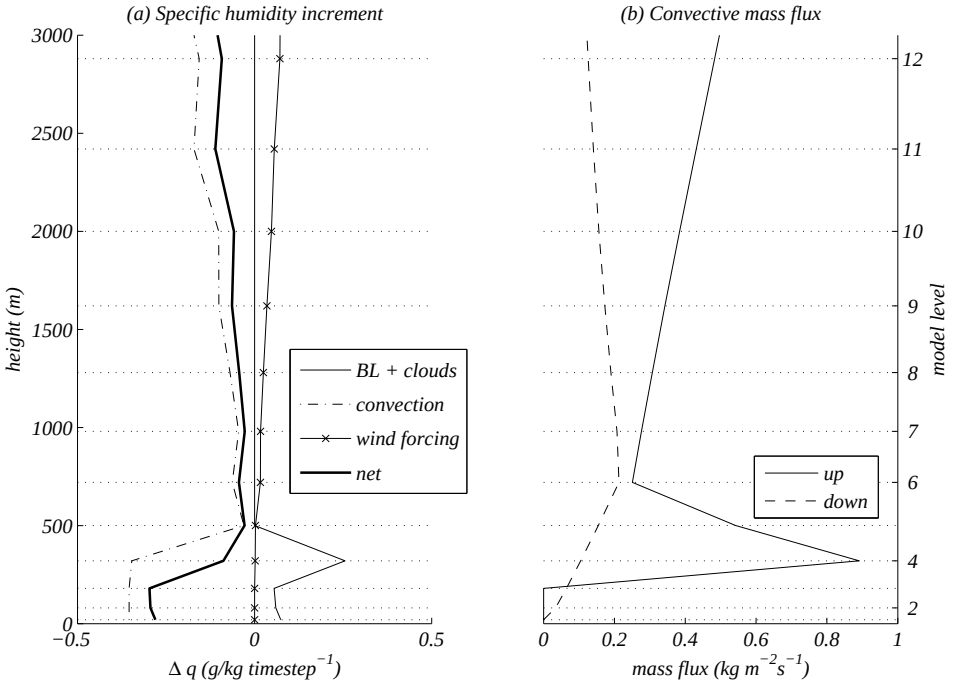

Figure 4.7: (a) The average temperature increment due to various processes in the timesteps in which convection acts. Drying by convection dominates in the lower levels. (b) Average profile of convective mass flux over twelve periods of the the three-timestep oscillation. The solid line is the upward mass flux, and the dashed line the mass flux associated with down-drafts.

the parcel is initialised with a lower humidity. This is in spite of the moistening of the surface layer via evaporation of rain. This term is contained in the 'boundary layer and clouds' term on the figure. It can be seen that there is some positive forcing due to evaporative effects, but the drying by convection dominates.

The reason for convection inducing such drying is partly explained by panel (b) of figure 4.7. This plot shows the average convective mass-flux for both up-drafts and down-drafts over the convective timesteps in question. There is a very strong upward mass flux at cloud base, which decreases substantially for the next two model layers. It is only above model level six that the mass-flux takes a shape more reminiscent of an entraining plume – the basis for this parameterisation. In section 1.3 the physics governing such a plume was described. Equation (1.1), on page 6 shows that the upward mass flux increases in response to entrainment, and decreases in response to detrainment. Detrainment is broken into two types – mixing detrainment and forced detrainment. Mixing detrainment occurs at all heights within a cloud, and its strength is related to the entrainment rate. Forced detrainment, on the other hand, only occurs when clouds reach their level of neutral buoyancy. It describes the ejection of mass from the top of these clouds into the environment. The strong decrease in mass-flux in levels five and six are unlikely to be caused by mixing detrainment, the primary contributor is most likely forced detrainment. This mass flux profile physically describes an ensemble of clouds with a large number of small cumulus clouds, and relatively few towers powerful enough to undergo deep convection. The result for the model is a large amount of convective warming

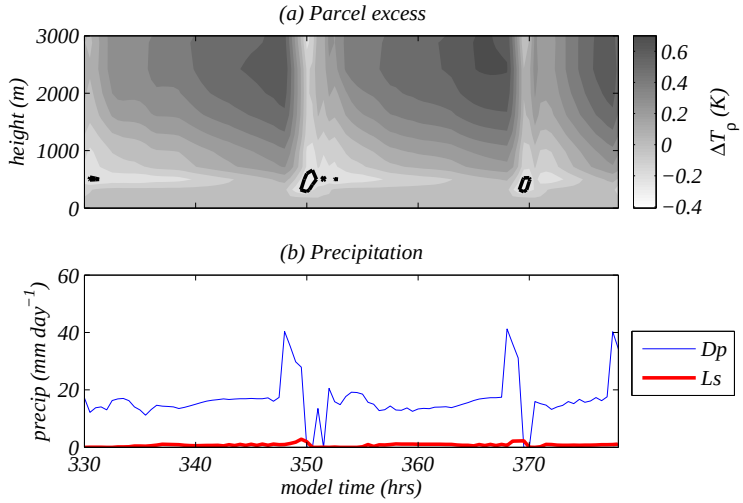

Figure 4.8: The uncharacteristic case of a smooth temporal distribution of convective rain. The case shown is for $w_R = 0.02$ ms^{-1}, SST = 30°C, and a CAPE timescale governed by $\tau_{MAX} = 3$ hours. The upper panel shows a contour of the parcel buoyancy excess over an arbitrary two day period of the simulation. The bold line is the position of the -0.2 K contour. Panel (b) shows the corresponding convective and large-scale rain rates.

(due to large scale compensating subsidence) and drying (due to precipitation) of the environment in the levels just above the LCL. This process explains the increase in environmental temperature at these levels. However, the parcel is only affected by the atmosphere below this, in the surface layer. Here, the updraft mass-flux is zero, and yet there is considerable drying of the surface layer air. The dominant process, therefore, must be the down-drafts. These carry relatively dry air from mid-levels downwards, where they are detrained near the surface. Returning to the description of an entraining plume, equation (1.1) can be applied to a downdraft. It can be seen that detrainment is proportional to the vertical derivative of the mass flux. The downdraft mass-flux in figure 4.7 can be seen to be consistently decreasing in the lowest six model layers. The associated detrainment of dry air reduces the humidity of the surface layer, thus preventing surface driven convection at the next timestep.

This behaviour can be contrasted with the case where the forcing is kept the same, but the CAPE timescale is increased. In figure 4.2 it can be seen that when the SST is kept at 30°C, the wind forcing at 0.02 ms^{-1}, but the CAPE timescale is set so that $\tau_{MAX} = 3$ hr, the variability is dramatically reduced. This behaviour is not characteristic of of this CAPE timescale, or these forcings – it appears to be a rather isolated case. Nevertheless, we see that rather than rapid oscillatory convection, the rainfall rate is relatively constant. Every twenty hours or so, some variability breaks out, but this only lasts a few timesteps before the model returns to a constant rainfall rate. This indicates that the constant rainfall regime is stable, and the model physics pushes toward this equilibrium when perturbed. This is shown in more detail in figure 4.8. A 48 hour period is plotted, with the upper panel showing the parcel buoyancy excess, and the lower showing the precipitation. It can be seen that while the rain rate is constant, the parcel slowly gains buoyancy with respect to the environment. Eventually, and rather suddenly, this slight instability induces a large increase in rain rate. The increased amount of convection stabilises the surface layer again, preventing convection

for a short time. This is shown on figure 4.8a by the line contour representing a value of buoyancy of -0.2 K. Once the convection resumes, it is at its original intensity, and the cycle continues. Note, that the period for the cycle is around 20 hours. While it is influenced by the diurnal cycle, a similar periodicity exists under perpetual night conditions.

It is clear that the reason the model oscillates wildly at the shorter CAPE timescale, but is relatively smooth at the longer, is that in the latter case, convection does not stabilise the environment so much as to prevent it from occurring at the next timestep. However, the two cases studied indicate that there can be considerable sensitivity to the precise profile of convective heating and drying. Had the drying in the short time-scale case occurred at higher levels, the parcel would have been less affected, and convection may not have been shut off. This sensitivity to the precise nature of the convective heating in most cases prevents convection from reaching an equilibrium. Only in the long-timescale case, when the convective equilibrium is neutral or slightly unstable with respect to the parcel ascent is any sort of equilibrium reached. It can be argued therefore that the convective trigger in the UM appears to be too prescriptive to the profile of the atmosphere required for convection for an equilibrium to be reached. In most cases the equilibrium structure that convection moves toward is not consistent with the parcel ascent. The trigger and the closure, therefore 'push' the model profile in different directions, causing the oscillatory behaviour in the convective response.

Chapter 5

Discussion & conclusions

In this study, the representation of convection in the U.K. Meteorological Office Unified Model (UM) was investigated. The model was run as a single-column model (SCM), in which only a column of grid-boxes are used, and the large-scale conditions prescribed. Simulations were performed under a series of idealised scenarios with fixed sea-surface temperatures and large-scale ascent rates. The range of sea-surface temperatures (SSTs) used to force the model are those that may be typically found over the tropical ocean. Similarly, a profile of vertical motion with a maximum, w_R in the mid-troposphere was prescribed, and the model run at a range of different w_R values. Conditions such as these, in which the large scale environment stays constant for long periods of time, are almost never observed. Modes of variability of the atmosphere are always present in one form or another, and in this sense, the experimental set up is somewhat unrealistic. Nevertheless, the results from these types of idealised simulations can be very instructive into the behaviour of a model, and can lead to new understanding of model deficiencies in more realistic simulations. The response of the Unified Model single-column model to the boundary conditions outlined is a case in point.

Convective activity exhibits an 'on-off' nature in almost all of the simulations. That is, while the large-scale conditions are kept constant, the amount of convection as measured by the convective precipitation, or the convective mass-flux, varies greatly from timestep to timestep. In particular, convection occurs for a single timestep, and then is shut-off for several timesteps after this. Evidently, the heating and drying of the atmosphere due to a convective event stabilises the environment too much to allow convection to be diagnosed in the next few timesteps. In a global model, this kind of behaviour would represent an ensemble of convection within a grid-box switching on for thirty minutes, and then being suppressed for some hours. This is not very realistic behaviour for an ensemble of convective clouds that may have a horizontal scale of 100 km or so. It is rather more characteristic of a single convective cell, which have lifetimes of the order of one timestep (Emanuel *et al.*, 1994). If the purpose of this model was to simulate a convective cloud, we may well expect the short period cycles that are observed. In this case, instability builds up over a number of hours, before the atmosphere is stabilised via convective heating. However, the physics of a GCM grid-box are rather different. At these scales, a single column would contain a large number of convective clouds, some dissipating, some intensifying (Arakawa & Schubert, 1974). The convective activity of a cumulus ensemble does not behave as a stabilising response to the build up of conditional instability. Emanuel *et al.* (1994) argue that instead, convection should

be seen as being in statistical equilibrium with the large scale flow. This is the paradigm on which the UM convection scheme is based. The CAPE closure also, requires a balance between dissipation of available potential energy by convection, with its production by large-scale processes. Thus, while the highly variable convective rain rates seen in the model are of dubious realism, the more important point is that this behaviour violates the very assumption the model is based on. Clearly, in almost all of the simulations, especially those with high SST and strong vertical ascent, the large scale conditions present a production of available potential energy. By the assumption of quasi-equilibrium, this should result in a balancing stabilisation of the atmosphere by convection. After an initial adjustment, we should expect the model's convection ensemble to reach statistical equilibrium and the convection rates to remain relatively constant. That this is not observed shows that *convection* in the model – a result of an interaction between all model processes – is not consistent with the assumptions that were used to create the model's *convection parameterisation* – the actual 'convection' subroutines in the model code.

A natural question to ask is how much these deficiencies in the representation of convection affect the overall model performance. It was shown in section 3.2 that the time averaged convective rain-rates are relatively smooth, and behave as expected. Vertical ascent has a strong effect on the total amount of convection observed, while there is a weaker correlation to the SST. This result is, to some degree, a corollary of the constraints imposed by the boundary conditions (see section 2.2), however, it at least confirms that the model produces plausible moisture structures, and partitions precipitation into large-scale and convective rain consistently. It is possible, therefore, that, when averaged over some timesteps, the on-off convection behaviour has a relatively small effect on the output. Thus, one must ask the question of whether effort should be spent on correcting this particular problem, at the expense of other issues. Convection parameterisation is a difficult problem – there are many more deficiencies in modern convective parameterisations than have been outlined in this thesis. It is the opinion of this author, however, that regardless of the effect on the overall output, the problem of large variability at the timestep to timestep level is an important one. This issue goes to the heart of one of the main assumptions behind modern convection schemes – the concept of an equilibrium between convection and the large-scale. It is of paramount importance for climate simulation, and even weather prediction, that we not only reproduce accurate climatologies of the various atmospheric parameters, but that the basic model physics is sound. A model that produces convection with such noise – despite the best efforts of model developers to ensure an equilibrium exists – does not inspire confidence in its robustness under changing conditions.

The time averaged response of other variables to the changing boundary conditions do present some unexpected results. At strong vertical ascent, both the sensible and latent surface heat fluxes, exhibit a threshold type of increase as the SST is increased above 29°C. Similarly, the large-scale rain rate decreases substantially in the same region of the parameter space. This behaviour needs to be investigated further before its causes can be understood. At lower sea-surface temperatures, there is a strong tendency for the model to exhibit a stable surface layer in which the sensible heat flux is negative, or, to diagnose a large amount of low-level cloud – a fog layer. Sea-surface temperatures of 24°C are much higher than would be expected to produce these types of conditions in reality. The reason this occurs is most likely due to the initialisation of the model, rather than a deficiency of the model itself. The profiles of temperature and moisture used as initial conditions to all simulations are taken from a

typical thermodynamic profile of the Tropical Warm Pool International Cloud Experiment (TWP-ICE) (see May *et al.*, 2008). The corresponding sea surface temperature is 28.5°C. Thus the initial profile of the simulations is very warm and moist, in comparison to what would be observed over a 23 - 24°C ocean. The atmosphere therefore adjusts, via radiative cooling, to a state more in equilibrium with the lower SST. This causes low level saturation of the atmosphere, and resulting low level cloud. Continuous convergence of moisture via the vertical wind profile feeds this cloud. In the stable situations, the condensation associated with this cloud is enough to maintain the surface layer at a higher temperature than the sea-surface, and hence the surface heat flux remains negative. A situation such as this is unlikely to occur in the real atmosphere except possibly in regions with strong advection of warm moist air over a much cooler sea surface. Further investigation, with varying initial conditions is required to confirm that this bias toward a stable, moist surface layer, is indeed due to the initial conditions and not a property of the model.

In chapter 4 the attention was turned back to the temporal structure of the simulations, and some explanation for the variability in convection was sought. Two possible reasons for the on-off convective behaviour were investigated. The first is to do with the convective closure. As mentioned above, the convective behaviour of the model is akin to what one would expect for an individual convective cloud, which has a much faster timescale than an ensemble of convective elements. However, the vertical velocities with which the model is forced are grid-box mean values. Thus the physical scale of the problem is much greater than that of an individual convective cloud. Aside from the magnitude of the vertical velocity, the only parameter relating modelled convection to any time or space scale is the CAPE timescale. A shorter timescale implies a stronger convective response per timestep. If the response is too strong, it might overstabilise the atmosphere, thus producing oscillations like those observed. The CAPE timescale used in the '4a' convective scheme of the UM is allowed to vary, and is governed by equation (1.2). Using the control value of $\tau_{MAX} = 1$ hour results in a CAPE timescale that averages half that for low forcing, and a mere 11 minutes under conditions of strong moisture convergence. It has been suggested that this is too short on physical grounds (Jakob, 2008, personal comm.). Even on numerical grounds however, to use a CAPE timescale that is so close to the timestep on which the convection scheme is integrated forward is not ideal. If this is truly the timescale on which convection adjusts to the large-scale flow, either the sub-timestep must be reduced, or a 'strict' quasi-equilibrium (QE) may be more appropriate. Strict QE closure was that originally proposed in the Arakawa-Schubert model, and involves assuming the amount of convective available potential energy remains constant with time. Thus there is always a strict balance between convective motions and the large-scale production of instability.

Two different CAPE timescales were tested, in addition to the control value. Using the relative humidity formulation, the parameter τ_{MAX} in equation (1.2) was increased to 3 hours. Additionally, a fixed value of $\tau_{CAPE} = 3$ hours was also investigated. Tripling the CAPE timescale, in general had a surprisingly small effect on the convection in the model – especially at high forcing. This can be attributed to the way in which the timescale is allowed to vary. In the case of $\tau_{MAX} = 1$ hour, the value calculated using the relative humidity function is often much less than 10 minutes. However, because it is not allowed to drop below the length of the sub-timestep, the CAPE timescale is set to ten minutes. When the value of τ_{MAX} is tripled then, this has only has a small effect on the final value of the CAPE timescale. A much stronger effect is observed when the CAPE timescale is

fixed at three hours. The convection is certainly less variable, and generally does not have an on-off character. However, this represents a very large change from the control value. In simulations with strong ascent, and thus strong convergence of moisture, the large-scale rain begins to dominate, as the convection scheme can no longer provide enough precipitation to remove moisture fast enough. If the full model was run with this CAPE timescale there would undoubtedly be even larger changes to the behaviour of convection. Clearly, more work is required, both in modelling and in observational studies, in order to constrain the timescales which are relevant to convection on different spatial scales.

The variability in convection is also strongly dependent on the convection trigger model. The Unified Model uses a parcel ascent to determine the state of the boundary layer in order to diagnose whether surface driven convection is possible. This diagnosis scheme predicts the height of the boundary layer, and the existence and type of cloud within it. If the boundary layer is cumulus topped, convection is allowed to occur, and the model's deep convection scheme is called. If not, convection can only occur above the boundary layer via mid-level instability. The boundary layer diagnosis scheme, as well as diagnosing convection, must determine the type of boundary layer to be input to the boundary layer scheme. As such, the parcel ascent used is quite detailed, and the conditions for a cumulus capped boundary layer are very specific. It is found that this complexity makes the diagnosis of convection very sensitive to small changes in the thermodynamics of the lowest model layers, and correspondingly drives much of the timestep to timestep variability. In the specific case studied, a three timestep cycle existed in which on only one of the timesteps was convective. In the other two, strong drying of the lowest model levels by down-drafts stabilised the parcel ascent, and prevented deep convection from occurring. The actual change in the parcel is modest – the reduction in buoyancy is in the order of 0.1 - 0.2 K. However, due to the strict conditions required for a cumulus topped boundary layer to be diagnosed, this is enough to prevent convection.

It is very likely that a relaxation in the conditions necessary for the boundary layer to be diagnosed as convective would remove much of the noise associated with convection in the model. However, the question that must be asked is whether these changes to the gird-box mean environment *should* be enough to shut of convection. This, to some degree, depends on the precise purpose of the convective trigger. Does the trigger mechanism in a model aim to predict, in detail, the ascent of a buoyant plume from the surface, and thus determine whether convection *will* occur? Or, is the trigger model a simple test – a back of the envelope calculation – as to whether convection *can* occur? The Unified Model is closer to the former. The problem that this creates is that the 'equilibrium' that the mass-flux scheme relaxes to, in this quasi-equilibrium formulation, may not be neutral with respect to the parcel ascent of the trigger. This is what appears to cause the oscillations seen in section 4.3. When the atmosphere reaches a state that is close to neutral with respect to the trigger model, convection occurs and re-stabilises it – preventing more convection occurring at the next timestep. In this author's opinion, only the second type of trigger is consistent with the assumption of quasi-equilibrium. If we accept that convection does not act as a reactive response to conditional instability, but as a force that is in equilibrium with the large-scale flow, we must allow this equilibrium a chance to be reached. Additionally, in a grid-box of the order of 100 km across, there will no doubt be some updrafts that are able to penetrate through the boundary layer, even if it is weakly stable. Thus, the convective trigger should only prevent the convection scheme from being called if the atmosphere is highly unsuitable for convection. In the weakly

stable case, one would hope that the cloud base mass-flux diagnosed by the closure scheme would be small, thus reflecting that there are only a few convective clouds in the grid-box.

The results from this study show that even under the simple case of constant large-scale forcing, the Unified Model single-column model responds in unexpected and complex ways. It is not entirely clear, however, the degree to which these responses occur under more realistic conditions. The two main features of the model response reported on in this thesis are the strong temporal variability of convection, and the dominance of stable boundary layer, and low level fog regimes at low SST. It has been argued that the former, due to its appearance in almost all simulations, is a genuine response of the model, and would likely exist both in SCM simulations forced by observations, and the GCM itself. The latter effect, that of a stable surface layer being maintained, is more likely due to the very warm thermodynamic profile used in the initialise simulations. In order to confirm that these hypotheses are true, further modelling, both with the SCM and the full GCM are required. These tests could also be used to better characterise these problems, and evaluate the possible solutions – an increase in CAPE timescale, or relaxation of convection triggering conditions – in more detail.

Bibliography

Arakawa, A. 1993 Closure assumptions in the cumulus parameterisation problem. In *The representation of cumulus convection in numerical models* (ed. K. Emanuel & D. Raymond), *Meteorological Monographs*, vol. 24, chap. 1, pp. 1 – 16. American Meteorological Society.

Arakawa, A. 2004 The cumulus parameterization problem: past, present, and future. *Journal of Climate* **17** (13), 2493–2525.

Arakawa, A. & Schubert, W. 1974 Interaction of a cumulus cloud ensemble with the large-scale environment, part I. *Journal of the Atmospheric Sciences* **31** (3), 674–701.

Arking, A. & Ziskin, D. 1994 Relationship between clouds and sea-surface temperatures in the western tropical Pacific. *Journal of Climate* **7** (6), 988–1000.

Betts, A. 1986 A new convective adjustment scheme. Part 1: Observational and theoretical basis. *Quarterly Journal of the Royal Meteorological Society* **112** (473), 677–691.

Betts, A. & Jakob, C. 2002 Evaluation of the diurnal cycle of precipitation, surface thermodynamics, and surface fluxes in the ECMWF model using LBA data. *J. Geophys. Res* **107**, 8045.

Bony, S., Lau, K. & Sud, Y. 1997 Sea-surface temperature and large-scale circulation influences on tropical greenhouse effect and cloud radiative forcing. *Journal of Climate* **10** (8), 2055–2077.

Emanuel, K. 1994 *Atmospheric convection*. Oxford University Press, USA.

Emanuel, K. 1997 Overview of atmospheric convection. In *The physics and parameterization of moist atmospheric convection* (ed. R. Smith), *NATO ASI series: Mathematical and physical sciences*, vol. 505, chap. 1, pp. 1 – 28. Kluwer Academic Publishers.

Emanuel, K., Neelin, J. & Bretherton, C. 1994 On large-scale circulations in convecting atmospheres. *Quarterly Journal of the Royal Meteorological Society* **120** (519), 1111–1143.

Fu, R., Del Genio, A. & Rossow, W. 1994 Influence of ocean surface conditions on atmospheric vertical thermodynamic structure and deep convection. *Journal of Climate* **7** (7), 1092–1108.

Garratt, J. 1992 *The atmospheric boundary layer*. Cambridge University Press Cambridge.

Graham, N. & Barnett, T. 1987 Sea-surface temperature, surface wind divergence, and convection over tropical oceans. *Science* **238** (4827), 657–659.

Grant, A. 2001 Cloud-base fluxes in the cumulus-capped boundary layer. *Quarterly Journal of the Royal Meteorological Society* **127** (572), 407–421.

Gregory, D. & Park, S. 1997 The mass flux approach to the parameterization of deep convection. In *The physics and parameterization of moist atmospheric convection* (ed. R. Smith), *NATO ASI series*, vol. 505, chap. 12, pp. 297–320. Kluwer Academic Publishers.

Gregory, D. & Rowntree, P. 1990 A mass flux convection scheme with representation of cloud ensemble characteristics and stability-dependent closure. *Monthly Weather Review* **118** (7), 1483–1506.

Holton, J. 2004 *An introduction to dynamic meteorology*, 4th edn. Academic Press.

Kuo, H. 1974 Further studies of the parameterization of the influence of cumulus convection on large-scale flow. *Journal of the Atmospheric Sciences* **31** (5), 1232–1240.

Lau, K., Wu, H. & Bony, S. 1997 The role of large-scale atmospheric circulation in the relationship between tropical convection and sea-surface temperature. *Journal of Climate* **10** (3), 381–392.

Lock, A., Brown, A., Bush, M., Martin, G. & Smith, R. 2000 A new boundary layer mixing scheme. Part I: scheme description and single-column model tests. *Monthly Weather Review* **128** (9), 3187–3199.

Lord, S. & Arakawa, A. 1980 Interaction of a cumulus cloud ensemble with the large-scale environment. Part II. *Journal of the Atmospheric Sciences* **37** (12), 2677–2692.

Manabe, S. & Strickler, R. 1964 Thermal equilibrium of the atmosphere with a convective adjustment. *Journal of the Atmospheric Sciences* **21** (4), 361–385.

May, P., Mather, J., Vaughan, G., Jakob, C., McFarquhar, G., Bower, K. & Mace, G. 2008 The Tropical Warm Pool International Cloud Experiment. *Bulletin of the American Meteorological Society* **89** (5), 629–645.

Randall, D. & Cripe, D. 1999 Methods for the specification of observed forcing in single column models and cloud system models. *Journal of Geophysical Research* **104**, 527 – 545.

Randall, D., Ding, P. & Pan, D. 1997*a* The Arakawa-Schubert parameterization. In *The physics and parameterization of moist atmospheric convection* (ed. R. Smith), *NATO ASI series*, vol. 505, chap. 11, pp. 281–296. Kluwer Academic Publishers.

Randall, D., Khairoutdinov, M., Arakawa, A. & Grabowski, W. 2003 Breaking the cloud parameterization deadlock. *Bulletin of the American Meteorological Society* **84** (11), 1547–1564.

Randall, D., Pan, D., Ding, P. & Cripe, D. 1997*b* Quasi-Equilibrium. In *The Physics and Parameterization of Moist Atmospheric Convection* (ed. R. Smith), *NATO ASI series*, vol. 505, chap. 14, pp. 359–385. Kluwer Academic Publishers.

Randall, D., Xu, K., Sommerville, R. & Iacobellis, S. 1996 Single-column models and cloud ensemble models as links between observations and climate models. *Journal of Climate* **9**, 1683 – 1697.

Salby, M., Hendon, H., Woodberry, K. & Tanaka, K. 1991 Analysis of global cloud imagery from multiple satellites. *Bulletin of the American Meteorological Society* **72** (4), 467–480.

Simpson, J. 1971 On cumulus entrainment and one-dimensional models. *Journal of the Atmospheric Sciences* **28** (3), 449–455.

Tompkins, A. 2001 On the relationship between tropical convection and sea-surface temperature. *Journal of Climate* **14** (5), 633–637.

Tompkins, A. & Craig, G. 1999 Sensitivity of tropical convection to sea-surface temperature in the absence of large-scale flow. *Journal of Climate* **12** (2), 462–476.

U.K. Meteorological Office 1999 Unified model documentation papers. [Available online at http://ugamp.nerc.ac.uk/um/], viewed 19/11/2008.

Wallace, J. & Hobbs, P. 2006 *Atmospheric science: an introductory survey*, 2nd edn. Academic Press.

Yanai, M., Esbensen, S. & Chu, J. 1973 Determination of bulk properties of tropical cloud clusters from large-scale heat and moisture budgets. *Journal of the Atmospheric Sciences* **30** (4), 611–627.