Atmos. Chem. Phys. Discuss., 8, 20537–20564, 2008
www.atmos-chem-phys-discuss.net/8/20537/2008/

Atmospheric
Chemistry
and Physics
Discussions

Statistical properties of cloud lifecycles in cloud-resolving models

R. S. Plant

Department of Meteorology, University of Reading, Earley Gate, Reading, P.O. Box 243,
RG6 6BB, UK

Received: 15 August 2008 – Accepted: 13 October 2008 – Published: 9 December 2008

Correspondence to: R. S. Plant (r.s.plant@reading.ac.uk)

Published by Copernicus Publications on behalf of the European Geosciences Union.

Abstract

A new technique is described for the analysis of cloud-resolving model simulations, which allows one to investigate the statistics of the lifecycles of cumulus clouds. Clouds are tracked from timestep-to-timestep within the model run. This allows for a very simple method of tracking, but one which is both comprehensive and robust. An approach for handling cloud splits and mergers is described which allows clouds with simple and complicated time histories to be compared within a single framework. This is found to be important for the analysis of an idealized simulation of radiative-convective equilibrium, in which the moist, buoyant, updrafts (i.e., the convective cores) were tracked. Around half of all such cores were subject to splits and mergers during their lifecycles. For cores without any such events, the average lifetime is 30 min, but events can lengthen the typical lifetime considerably.

1 Introduction

In recent years Cloud Resolving Models (CRMs) have become an increasingly important tool for the study of convective phenomena. CRMs should not be regarded as providing surrogates for observations; rather, they allow idealized but realistic simulations to be produced which provide a laboratory for the careful diagnostic analysis of generic convective systems. Such analysis is a distinctive methodology that is necessary to improve our understanding of the basic phenomena and to develop improved parameterization methods for larger-scale models.

This paper describes and illustrates the use of a novel analysis technique for CRM data, which allows one to investigate statistical properties of the lifecycles of clouds produced during CRM simulations.

Current analyses of CRM data often focus on determining and understanding the spatial and temporal average properties of the full ensemble of convective clouds that are produced in the model in response to some specified external forcing (e.g. Petch et al., 2007). Rather less attention has been devoted to the lifecycle behaviour of individual clouds. There are currently many simulations (whether labelled as CRM or otherwise) which are being performed with convection represented explicitly but at rather coarse resolution (~ 1 to 5 km) (e.g Petch et al., 2002; Done et al., 2004; Khairoutdinov et al., 2005). In such simulations a deep convective cloud may occupy only a small number of model gridpoints. Thus, although the results may provide genuine value relative to their lower-resolution counterparts with parameterized convection (e.g. Roberts and Lean, 2008), it is far from clear that the simulations will provide a good representation of individual clouds. A statistical investigation into cloud lifecycles could therefore be valuable in order to reveal which aspects of the lifecycles are well or poorly captured at these model resolutions. Even assuming a high-resolution simulation, however, statistical information on the cloud lifecycle would be useful to test the realism of the model clouds, and to allow one to examine the detailed effects of model parameterizations, such as the microphysics. A good recent study of cloud lifecycles in a CRM simulation is that of Zhao and Austin (2005a,b). However, practical constraints limited that study to an investigation of six clouds, making it difficult to assess whether the results are generic.

Statistical investigations into the lifecycles of cumulus clouds could also allow improvements to be made to existing convective parameterizations. Cho (1977) considered the effects of incorporating a cloud lifecycle into a cumulus-ensemble mass-flux framework, and showed that the effects on the apparent heating Q_1 were negligible. However, an additional contribution arises in Q_2 compared to a steady-state cloud model, due to mixing of air from the decayed cloud with its environment. Another example comes from the popular Kain and Fritsch parameterization (Kain, 2004) for mesoscale models. A rudimentary lifecycle is included by assigning to the convective plumes a (somewhat arbitrary) lifetime which extends over multiple timesteps. The parameterization is a mass-flux scheme which considers a single plume to be representative of all convection occuring within a model grid box. Based on the pioneering study of Arakawa and Schubert (1974), some other parameterizations consider a spectrum of convective plumes (Plant and Craig, 2008, is a recent example). Future parameterizations might seek to combine these two features: multiple cloud types and a simple cloud lifecycle. However, this is not possible at present, essentially because there is a lack of available information about how the cloud lifecycle varies with cloud type (and forcing regime).

An automated method is presented which first identifies and then tracks the development of individual clouds in a CRM simulation. Its most important characteristic is that it is run online, at every timestep, alongside the model simulation. By exploiting the high temporal resolution available in a model, it is possible to devise a tracking method that is at once simple, comprehensive and robust. The method is fully described in Sect. 2 and results obtained from tracking moist, buoyant updrafts in a CRM simulation are discussed in Sect. 3. Conclusions are drawn in Sect. 4.

2 Methodology

The purpose of the tracking algorithm is to capture the complete time evolution of each cloud produced in a numerical model simulation. The algorithm can be divided into three main parts, which are described in turn below.

2.1 Identify clouds at a given timestep

Cloud identification requires, first, a determination of the cloudy points, and second, joining such points together appropriately into distinct clouds. A wide variety of criteria have been used in the literature for the identification of cloudy points. Analyses of satellite observations often employ a brightness-temperature threshold (e.g. Kuo et al., 1993; Carvalho and Jones, 2001; Machado and Laurent, 2004). Other methods are based on radar echoes (e.g. Foote and Mohr, 1979; Dixon and Wiener, 1993; Theusner and Hauf, 2004) and even the visual inspection of photographs (e.g. Plank, 1969; Hozumi et al., 1982).

In model simulations, the identification of cloudy grid points is less constrained by the character of the data and it is possible to threshold model variables that are arguably more directly related to the presence of clouds. One popular choice (e.g. Xu and Randall, 2001; Cohen and Craig, 2006) is to use a vertical velocity criterion ($w>1\,\mathrm{ms}^{-1}$ anywhere in the column) in order to pick out strong updrafts. This approach has its origin in analyses of aircraft observations (LeMone and Zipser, 1980; Zipser and LeMone, 1980). Other methods are simply to use model variables for cloud water and/or ice content (e.g. Cohen and Craig, 2006), or to consider the convective transport of boundary layer air by means of a passive tracer (Zhao and Austin, 2005a). Siebesma and Cuijpers (1995) compared three identification methods for simulated shallow cumulus, which they referred to as the cloud decomposition (positive cloud water), the updraft decomposition (positive cloud water and vertical velocity) and the cloud-core decomposition (positive cloud water, vertical velocity and buoyancy). The cloud-core decomposition produced the best agreement between the mass flux representation of turbulent fluxes (assumed by many parameterizations) and the actual model fluxes. For related discussions, see also Swann (2001); Siebesma et al. (2003); Yano et al. (2004).

It would be wrong to view any particular cloud definition as intrinsically correct. Rather the different definitions allow one to focus attention on different aspects of the cloud field. In Sect. 3 we will use the same "cloud-core" definition as in Siebesma and Cuijpers (1995), but it would be straightforward to implement other choices.

Once the "cloudy" gridpoints have been determined, it remains to connect adjacent points together into distinct clouds. As discussed by Kuo et al. (1993) for example, either a four-segmented or eight-segmented method can be used, the former considering only those adjacent gridpoints which share a gridbox edge, whereas the latter also allows connections to neighbouring gridpoints along a diagonal. In Sect. 3 an eight-connected method will be used. For example then, the cloud structure G_3 shown in Fig. 1 would be considered to contain six model gridpoints. Although Kuo et al. (1993) obtained similar results from the two methods, some differences occur in the numbers of one-point and two-point clouds in coarse-resolution simulations (Lennard, 2004).

For a cloud to be included in the statistics presented below, it is required to contain at least two cloudy gridpoints. Thus the very smallest clouds, such as the group G_1 in Fig. 1, are ignored. One would not expect these to be well represented by the model. A further requirement on the cloud is that it should persist for at least 5 min. The combination of these two conditions helps to ensure that the final statistics should not be overly sensitive to the precise definition of cloudy gridpoints, since any isolated, short-lived fluctuations above the thresholds are excluded.

2.2 Relationships with clouds at previous timestep

The purpose of the second part of the algorithm is to establish the relationships between clouds present at the current timestep and those present at the previous one. Many tracking methods have been developed for determining the evolution of features in data of relatively low temporal resolution (e.g. Dixon and Wiener, 1993; Carvalho and Jones, 2001; Machado and Laurent, 2004, and references therein). Given two time slices, each of which contains one or more features of interest, the aim is to establish the features that are in common between the two slices, essentially satisfying oneself that a feature in the first slice is highly likely to have evolved into some feature(s) in the later slice. Often the method will involve forming some estimate of the propagation speed of the feature. On occasion, the relationships between features in the two time slices may not be entirely clear. (Data errors, such as radar clutter, can also produce some tracking errors, as noted by Weusthoff and Hauf (2008) for instance.)

Here the tracking algorithm is applied online, as a diagnostic component of the numerical model simulation. Because data is available to the tracking algorithm with very high temporal resolution, it is possible to establish the relationships between clouds at adjacent timesteps using a method that is both simple and comprehensive. It is assumed that in a single timestep there is no motion further than a single horizontal gridlength. This is a numerical stability requirement for many of the advection schemes used by CRMs, including the LEM used in Sect. 3. In order for there to be a meaningful relationship between two clouds at adjacent timesteps, we assert that at least some of the air constituting the cloud at the previous timestep must be present in the cloud at current timestep. It follows that all of the relationships required can be found by looking for clouds present (at least in some part) at the previous timestep within a halo region for each current cloud. As illustrated in Fig. 1, the halo is comprised of those gridpoints that either overlap or are adjacent to the cloudy gridpoints of the current cloud.

From the set of all relationships, the character of the relationships between previous and current clouds can be determined. Construct the maximum possible number of subsets of relationships, subject to the constraint that each cloud at the current or previous timestep appears in one and only one subset. A useful property of each such subset can be denoted by $p \rightarrow c$, where p and c are the total number of clouds in the subset from the previous and current timesteps respectively. This property allows the following characteristic relationships to be distinguished.

- 0 → 1 (i.e., there is no relationship to a cloud at the current timestep from any of the clouds present at the previous timestep) signifies the birth of a new cloud.
- 1 → 0 signifies the death of a cloud.
- 1 → 1 (by far the most common occurrence in practice) signifies a straightforward continuation of a pre-existing cloud.
- 1 → 2+ signifies the splitting up of a pre-existing cloud.
- 2+ → 1 signifies a merger of pre-existing clouds to form a single cloud structure.
- 2+ → 2+ signify more complicated relationships, which might occur, for example, if a pre-existing cloud simultaneously both breaks-up and absorbs another pre-existing cloud. Such happenings are extremely rare, but nonetheless must be accounted for.

Note that 2+ has been used to denote two or more clouds. It is convenient to be able to distinguish between the births, deaths and straightforward continuations on the one hand, and the splits, mergers and complicated relationships on the other. In order to do so, we will henceforth refer to the latter types of relationship as “events”.

2.3 Compile timeseries data for each cloud

For each cloud a timeseries is stored of relevant data, including such properties as the cloud size (number of gridpoints), the precipitation rate and mass fluxes. The procedure for updating and organizing the timeseries depends upon the character of relationships to the clouds present at the previous timestep. Births, deaths and straightforward continuations are easy to deal with, signalling respectively the start of a new timeseries, the output of a completed series, and the addition of a new entry to a pre-existing series.

For any event, the timeseries of all pre-existing clouds contributing to the event are terminated and archived into a library. New timeseries are begun for all of the current clouds involved in the event. The full time history for current clouds can thus be reconstructed by means of references to the library. If a current cloud has undergone a single event during its time history we describe it as being a second-generation cloud. Higher orders of generation are also possible: for example, if a cloud splits into two parts, one of which later splits again. Higher orders can be incorporated by extending the above procedure to allow inter-library references. In this way, each current cloud can be followed back through all of its contributing elements.

As well as the character of each event, it is useful also to save parameters which estimate the relative contributions of the various clouds involved. Specifically, we calculate the quantities f_i^c which represent the fraction of a cloud i from the previous timestep that can be linked to the current cloud c. (The fraction should be interpreted as zero if there is no relationship between i and c.) For multi-generational clouds, this is generalized to a fractional association a_n^c with some cloud element n held in the library. The association is given by a product of the form $f_n^m f_m^l \ldots f_k^j f_j^i f_i^c$, summed over all possible sequences of relationships that lead from n to c.

The determination of the fraction f_i^c is based on the areas occupied by the clouds concerned. In a 2 → 1 merger of clouds i and j to produce cloud c, the fractions are trivially

$$f_i^c = f_j^c = 1 \tag{1}$$

while for a 1 → 2 split of cloud i into clouds c and d we have

$$f_i^c = \frac{A^c}{A^c + A^d}; \qquad f_i^d = \frac{A^d}{A^c + A^d} \tag{2}$$

where A denotes the cloud area.

A generalization of the approach to encompass other events is given by the equation below, allowing fractions to be determined for potentially complicated events involving multiple clouds from both the previous and current timesteps. Specifically:

$$f_i^c = N_i \left(r^c + \frac{A_i}{I_i} \right) \tag{3}$$

where I is the total number of relationships from a particular (subscripted) cloud at the previous timestep to all clouds at the current timestep. The reduced area r^c may be positive or negative and is intended to provide an indication of any portion of the current cloud c that is not linked to clouds from the previous timestep. (In a merger, for example, a positive value would indicate a current cloud that is larger than its consitutents from the previous timestep.) It is defined by

$$r^c = A^c - \sum_i{}' \frac{A_i}{I_i}, \tag{4}$$

where the primed summation extends over those clouds i from the previous timestep that have a relationship to the current cloud c. To complete the specification of the fraction, it remains to define the quantity N_i. This is a normalization factor, chosen to ensure that all clouds from the previous timestep that are involved in events are fully linked to current clouds. Hence, N_i is such that

$$\sum_c f_i^c = 1 . \tag{5}$$

As a check on the formula in Eq. (3), it is easy to confirm that for a 2 → 1 merger and a 1 → 2 split, the fractions reduce to those given in Eqs. (1) and (2) respectively.

In order to compare clouds with events during their time history alongside clouds without any events, it is useful to be able to construct single timeseries for the clouds that are subject to events. We define the lifetime of such clouds as extending backwards from the death of a cloud to the birth of its first contributing element. For an extensive cloud property E, the timeseries is obtained from

$$E^c = \sum_n a_n^c E_n \tag{6}$$

the sum extending over all contributing cloud elements n, with the understanding that this includes c itself and that a_c^c=1. For an intensive cloud property I, the product of the association with the size of a cloud element is considered to provide a weighting factor, so that

$$I^c = \frac{\sum_n a_n^c A_n I_n}{\sum_n a_n^c A_n} . \tag{7}$$

Finally, we note two restrictions on the multi-generational cloud library that are imposed for purely practical reasons. If the library becomes very large, or if clouds extend back through many generations, then searching through the library can become a time-consuming operation, considerably slowing the model simulation. Therefore, we remove from the library any archived cloud elements n with associations a_n^c that are less than 0.05 for all of the current clouds c. Moreover, we do not allow a cloud to extend backwards for more than 10 generations. Some diagnostics characterizing the removed cloud elements are output in order to allow various checks that the removals do not have significant adverse effects on the final cloud statistics. For example, in the simulation results to be presented in Sect. 3 the average duration of removed elements is short: under 2 min, which compares with an average duration of around 14 min for retained library elements. Various test runs with different values for the removal criteria have also been performed to check explicitly for any effects of the removals.

3 Results from a CRM simulation

The tracking algorithm described in Sect. 2 has been tested in both a cloud-resolving model and in artificial dynamical systems of cellular automata (based on variations of the game-of-life rules). The advantage of the artificial system is that its rules can be altered to test various aspects of the algorithm: for example, allowing events to be extremely rare or else frequent and complex. Explicit timestep-to-timestep validations have been performed to check that the algorithm is robust and functions as designed.

We present results for a simulation of radiative-convective equilibrium performed with the Met. Office Large Eddy Model (LEM) (Gray et al., 2001). The setup is not dissimilar to simulations that have previously been studied by Cohen and Craig (2004, 2006). The convection is forced by cooling the troposphere at 4 K day^{-1}, over a sea-surface which has its temperature held fixed at 300 K. The simulation domain is a doubly-periodic grid of size 64×64×20 km^3 with a horizontal gridlength of 2 km and 76 staggered vertical levels. The Coriolis parameter is set to zero and no mean wind is imposed, so that the convection is not expected to be organized by the large-scale state. In fact, limited self-organization does occur in such conditions, as discussed by Cohen and Craig (2006); Davies (2008).

The LEM has a variable timestep, which can change during the simulation. This is to ensure good behaviour of the subgrid model (Derbyshire et al., 1994), and also that the CFL stability condition for advection remains satisfied throughout. The mean timestep is 0.51 s.

The simulation is run for 36 model days, of which the first 19.5 days are used to spin-up from a rather arbitrary initial condition to the equilibrium state. The domain-mean model state does not vary in time once equilibrium is reached, apart from fluctuations attributable to the finite size of the domain (Cohen and Craig, 2006). Statistics are presented for the lifecycles of 4617 completed convective cores tracked during the remainder of the simulation. Table 1 summarizes some basic statistics of interest. Although there are some isolated single-gridpoint cores present, it is clear that the part of the domain containing moist, buoyant, updrafts is well captured. Complicated events are seen to be rare, as anticipated, but splits and mergers are not unusual.

The statistics for the proportions of various events are potentially very sensitive to the removal criteria applied to the cloud library (Sect. 2.3). However, in test runs allowing up to 40 generations and reducing the required associations to 0.01, the same proportions were produced to within 5%.

3.1 Convective core lifetimes

Figure 2 shows the distribution of convective core lifetimes, both for all cores tracked (panel a), and for those whose time histories do not contain any events (panel b). 54.2% of core lifecycles do not contain any events, and these have a mean lifetime of 29.7 min. Their lifetime distribution is found to be approximately exponential up to ~45 min, with a small peak for around 60 to 75 min. Including the core lifecycles which do contain events enhances the proportion of longer-lived cores, raising the mean lifetime to 54.6 min. In Sect. 3.2 we discuss further the effects of events on the convective core lifetime.

Examining timeseries for particular cores shows that (as expected) the pattern of evolution of each core is rather specific to the core in question. Nonetheless, it is possible to draw out some general properties of the simulated, convective-core lifecycle by normalizing the timeseries for each core and then compositing the cores to produce an averaged lifecycle. Time is normalized using the core lifetime, and each core property is normalized for each core by its time-mean value across the lifecycle. Figure 3 shows such composites for various core properties, and for cores with different ranges of lifetime. Panel (d) shows the centre of mass, h, which is defined as the first moment of the condensate, C.

$$h = \frac{\int C z \rho dz}{\int C \rho dz} \tag{8}$$

In order to demonstrate the variability between cores and across lifecycles, Fig. 4 shows probability distribution functions of normalized core properties for selected lifecycle stages.

Clearly the longer-lived cores have more pronounced lifecycles. The short-lived convective cores (with lifetimes less than 30 min) receive little support for time development, the vertically-integrated mass flux decreasing monotonically through their composite lifecycle (Fig. 3c). As a result, the core area and centre of mass remain almost constant (Fig. 3a,d). Measures of updraft strength (the area and mass flux) for the longer-lived cores (lifetimes larger than 30 min) exhibit clear peaks towards the later part of the lifecycle. These cores increase their centre of mass throughout their composite lifecycles (Fig. 3d). This may occur due to vertical transport of the normalized condensate, or else because the production of condensate occurs at progressively higher levels through the course of the lifecycles. The positions of the peaks in the mass flux probability distributions (Fig. 4b) are consistent with the composite lifecycles. The distributions have a broad spread, comparable to or perhaps somewhat larger than the variations seen in the composites across the lifecycle. This suggests that the convective-core composites do indeed have value in describing the lifecycle, but that caution should be used in interpreting them as generic lifecycles. Note also that the probability distributions tend to become more spread as the lifecycle progresses.

The normalized precipitation rate shows a considerable increase across the lifecycle and the rates remain relatively large when the convective core dies. This is consistent with the notion that the lifetime of a convective cell is similar to the time required for precipitation to develop (Rogers and Yau, 1989). It also highlights the importance of differences between convective cloud definitions. For example, Fig. 4a shows that the use of a precipitation threshold would not capture many of the convective cores during the first 10% of their lifecycles. Certainly then, it would not be appropriate to compare directly the statistics of cores obtained here to statistics obtained by, say, tracking radar echoes. Observationally-based "cloud" identification and tracking methods (such as those cited in Sects. 2.1, 2.2) are strongly constrained by the nature of the available data. Systematic studies of the sensitivity of lifecycle statistics to cloud definition are needed, both to inform comparisons between clouds observed with different systems and to allow comparisons between modelled and observed clouds.

3.2 The role of lifecycle events

It was shown in Sect. 3.1 that the presence of "events" in the lifecycle of a convective core considerably lengthens the lifetime of the cores. Here we consider the role of events in more detail.

Figure 5 shows the distribution of times that separate consecutive events identified by the tracking algorithm. For separations larger than ∼5 min, the distribution is roughly exponential. However, many of the events picked-out by the algorithm are quickly followed by other events, often within tens of seconds. Indeed 49.1% of all event separations are less than 1 min. The interpretation is that the joining together or breaking up of convective cores is rarely a clean process that happens once only at a single timestep. Rather, the joining (for example) of two cores is more typically a somewhat messy affair, perhaps with some portion of the combined core becoming temporarily detached as the constituent parts coalesce to produce what may ultimately become a unified entity.

In order to examine the effects of cores joining up or splitting, it would therefore not be appropriate to rely on the total number of events found by the tracking algorithm as a useful measure. The total reflects not only the number of incidents occurring in a lifecycle, but also how clean or messy those events are. Instead, we prefer to define "separated events" as those events satisfying the following criteria.

- The event must not take place within 5 min of the start or end of the lifecycle.
- The event must take place at least 5 min after the previous “separated event”.

The choice of 5 min is a reasonable but somewhat arbitrary one. It is equivalent to a relatively high time resolution that might be available in the data from current operational radar networks (Weusthoff and Hauf, 2008). We have checked that our conclusions are not qualitatively affected by reasonable changes of this choice.

Some statistics based on separated events are shown in Figs. 6 and 7. There are 2113 lifecycles that contain separated events, and the effects on the convective-core lifetimes are considerable. Each event increases the mean lifetime by about 15 min, or around half the mean lifetime of a lifecycle that does not contain any separated events.

A simple-minded explanation for this increase can be provided if it is supposed that each convective core initiated has a mean lifetime of ≈30 min (irrespective of whether there is another core close by), and that if cores are to be initiated in the vicinity of a pre-existing core then the characteristics and lifecycle-stage of the pre-exisiting core have no effect on the initiation process. Whether there is any truth in this explanation we leave as a topic for future work. There are at least hints in Figs. 6a and 7 that the idea may not be entirely unreasonable. Statistical independence of the cores would imply that the number of events satisfies a zero-truncated Poisson distribution: such a distribution is overlaid on Fig. 6a. The timing of separated events within the lifecycles is shown in Fig. 7. Although such events are less likely to occur towards the beginning or end of a lifecycle, the likelihood of events for most of the lifecycle is fairly uniform.

4 Conclusions

This paper describes the design and implementation of a novel method for analysing the results of CRM simulations. The algorithm developed operates as an almost self-contained diagnostic suite that is plugged into the model simulation. The objective is to identify clouds from the CRM results and track their evolution. By examining clouds on a timestep-to-timestep basis it is possible to exploit the high temporal resolution data available to an online diagnostic system. In conjunction with a numerical stability condition that is satisfied by the advection schemes typically used in CRMs, a simple methodology is sufficient to provide comprehensive and robust tracking. The algorithm has been designed to be as generic as possible. Alternative identification criteria for cloudy gridpoints would be trivial to implement, and most of the decisions about how to examine the lifecycle data are deferred to the postprocessing. Thus, the methodology would be straightforward to adapt in order to track other features online in other models.

The use of the methodology was demonstrated for an idealized simulation of radiative-convective equilibrium. The statistics obtained allow one to quantify the behaviour of individual convective cores in ways that are not accessible to other analysis methods. Around half of all cores tracked were subject to merging and splitting during their lifecycles. While this fact complicates the analysis, we were able to demonstrate how cores with simple and complicated time histories can be considered within a single framework and to demonstrate the considerable impact of "events" on the core lifetimes.

Acknowledgement. This work was supported by a Royal Society grant.

References

Arakawa, A. and Schubert, W. H.: Interaction of a cumulus cloud ensemble with the large-scale environment. Part I, J. Atmos. Sci., 31, 674–701, 1974. 20539

Carvalho, L. M. V. and Jones, C.: A satellite method to identify structural properties of mesoscale convective systems based on the maximum spatial correlation tracking technique (MASCOTTE), J. Appl. Meteorol., 40, 1683–1701, 2001. 20540, 20542

Cho, H.-R.: Contributions of cumulus cloud life-cycle effects to the large-scale heat and moisture budget equations, J. Atmos. Sci., 34, 87–97, 1977. 20539

Cohen, B. G. and Craig, G. C.: The response time of a convective cloud ensemble to a change in forcing, Q. J. Roy. Meteor. Soc., 130, 933–944, 2004. 20547

Cohen, B. G. and Craig, G. C.: Fluctuations in an equilibrium convective ensemble. Part II: Numerical experiments, J. Atmos. Sci., 63, 2005–2015, 2006. 20541, 20547, 20548

Davies, L.: Self organisation of convection as a mechanism for memory, Ph.D. thesis, University of Reading, UK, 165 pp., 2008. 20547

Derbyshire, S. H., Brown, A. R., and Lock, A. P.: The Meteorological Office Large-Eddy Simulation Model, Met O (APR) Turbulence and Diffusion Note No. 213, Met Office, UK, 1994. 20547

Dixon, M. and Wiener, G.: TITAN: thunderstorm identification, tracking, analysis and nowcasting – a radar-based methodology, J. Atmos. Ocean Tech., 10, 785–797, 1993. 20540, 20542

Done, J., Davis, C. A., and Weisman, M.: The next generation of NWP: Explicit forecasts of convection using the weather research and forecasting (WRF) model, Atmos. Sci. Lett., 5, 110–117, 2004. 20539

Foote, G. B. and Mohr, C. G.: Results of a randomized hail suppression experiment in northeast Colorado. Part VI: Post hoc stratification by storm intensity and type, J. Appl. Meteorol., 18, 1589–1600, 1979. 20540

Gray, M. E. B., Petch, J., Derbyshire, S. H., Brown, A. R., Lock, A. P., Swann, H. A., and Brown, P. R. A.: Version 2.3 of the Met Office Large Eddy Model. Part II: Scientific documentation, Met O (APR) Turbulence and Diffusion Note No. 276, Met Office, UK, 2001. 20547

Hozumi, K., Harimaya, T., and Magono, C.: The size distribution of cumulus clouds as a function of cloud amount, J. Meteorol. Soc. Jpn., 60, 691–699, 1982. 20540

Kain, J. S.: The Kain–Fritsch convective parameterization: An update, J. Appl. Meteorol., 43, 170–181, 2004. 20539

Khairoutdinov, M., Randall, D., and DeMott, C.: Simulations of the atmospheric general circulation using a cloud-resolving model as a superparameterization of physical processes, J. Atmos. Sci., 62, 2136–2154, 2005. 20539

Kuo, K.-S., Welch, R. M., and Weger, R. C.: The three-dimensional structure of cumulus clouds over the ocean. 1: Structural analysis, J. Geophys. Res., 98, 20 865–20 711, 1993. 20540, 20541

LeMone, M. A. and Zipser, E. J.: Cumulonimbus vertical velocity events in GATE. Part I: Diameter, intensity and mass flux, J. Atmos. Sci., 37, 2444–2457, 1980. 20541

Lennard, R.: The distribution of cumulus cloud sizes, MSc thesis, University of Reading, UK, 67 pp., 2004. 20541

Machado, L. A. T. and Laurent, H.: The convective system area expansion over Amazonia and its relationships with convective system life duration and high-level wind divergence, Mon. Weather Rev., 132, 714–725, 2004. 20540, 20542

Petch, J. C., Brown, A. R., and Gray, M. E. B.: The impact of horizontal resolution on the simulations of convective development over land, Q. J. Roy. Meteor. Soc., 128, 2031–2044, 2002. 20539

Petch, J. C., Willett, M., Wong, R. Y., and Woolnough, S. J.: Modelling suppressed and active convection. Comparing a numerical weather prediction, cloud-resolving and single-column model, Q. J. Roy. Meteor. Soc., 133, 1087–1100, 2007. 20538

Plank, V. G.: The size distribution of cumulus clouds in representative Florida populations, J. Appl. Meteorol., 8, 46–67, 1969. 20540

Plant, R. S. and Craig, G. C.: A stochastic parameterization for deep convection based on equilibrium statistics, J. Atmos. Sci., 65, 87–105, 2008. 20540

Roberts, N. M. and Lean, H. W.: Scale-selective verification of rainfall accumulations from high-resolution forecasts of convective events, Mon. Weather Rev., 136, 78–97, 2008. 20539

Rogers, R. R. and Yau, M. K.: A Short Course in Cloud Physics, 3rd Ed., Butterworth-Heinemann, Woburn, MA, USA, 304 pp., 1989. 20549

Siebesma, A. P. and Cuijpers, J.: Evaluation of parametric assumptions for shallow cumulus convection, J. Atmos. Sci., 52, 650–666, 1995. 20541

Siebesma, A. P., Bretherton, C. S., Brown, A., Chlond, A., Cuxart, J., Duynkerke, P. G., Jiang, H., Khairoutdinov, M., Lewellen, D., Moeng, C. H., Sanchez, E., Stevens, B., and Stevens, D. E.: A large eddy simulation intercomparison study of shallow cumulus convection, J. Atmos. Sci., 60, 1201–1219, 2003. 20541

Swann, H.: Evaluation of the mass-flux approach to parametrizing deep convection, Q. J. Roy. Meteor. Soc., 127, 1239–1260, 2001. 20541

Theusner, M. and Hauf, T.: A study on the small scale precipitation structure over Germany using the radar network of the German weather service, Meteorol. Z., 13, 311–322, 2004. 20540

Weusthoff, T. and Hauf, T.: The life cycle of convective-shower clouds under post-frontal conditions, Q. J. Roy. Meteor. Soc., 134, 841–857, 2008. 20542, 20551

Xu, K.-M. and Randall, D. A.: Updraft and downdraft statistics of simulated tropical and midlatitude cumulus convection, J. Atmos. Sci., 58, 1630–1649, 2001. 20541

Yano, J.-I., Guichard, F., Lafore, J.-P., Redelsperger, J.-L., and Bechtold, P.: Estimations of mass fluxes for cumulus parameterizations from high-resolution spatial data, J. Atmos. Sci., 61, 829–842, 2004. 20541

Zipser, E. J. and LeMone, M. A: Cumulonimbus vertical velocity events in GATE. Part II: Synthesis and model core structure, J. Atmos. Sci., 37, 2458–2469, 1980. 20541

Zhao M. and Austin P. H.: Life cycle of numerically simulated shallow cumulus clouds. Part I: Transport, J. Atmos. Sci., 62, 1269–1290, 2005a. 20539, 20541

Zhao M. and Austin P. H.: Life cycle of numerically simulated shallow cumulus clouds. Part II: Mixing dynamics, J. Atmos. Sci., 62, 1291–1310, 2005b. 20539

Table 1. Statistics of the convective-core lifecycles tracked during a CRM simulation of radiative-convective equilibrium. Gridpoints containing moist, buoyant, updrafts are referred to as core points, while two-or-more connected core points constitute a convective core (Sect. 2.1). The proportions of births, deaths, splits, mergers and complicated events are expressed as fractions relative to the number of straightforward continuations. All quantities are computed from data at each model timestep between 19.5 and 36 days.

Quantity	Mean	Standard deviation
Core points	52.4	6.9
Core points not part of convective cores	7.0	2.8
Number of convective cores	10.0	2.0
Proportion of births and deaths	3.0×10^{-4}	–
Proportion of splits	2.4×10^{-4}	–
Proportion of mergers	1.7×10^{-4}	–
Proportion of complicated events	4.2×10^{-6}	–

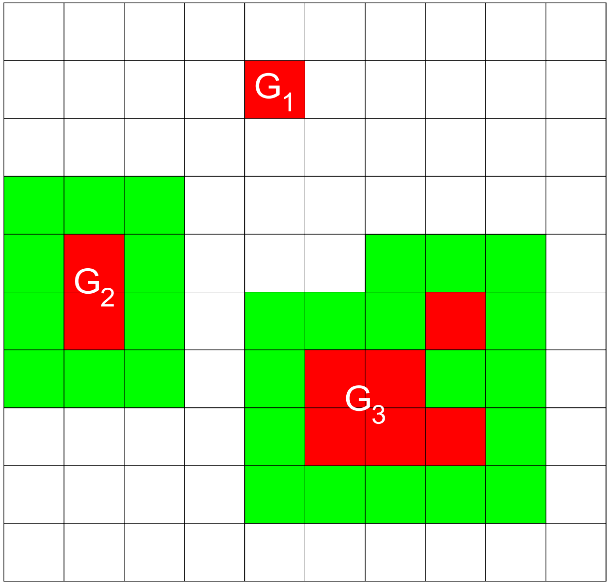

Fig. 1. Schematic diagram showing a portion of the horizontal domain used by a numerical model. Gridpoints identified as cloudy are shown in red, whilst green indicates the halo of gridpoints to be considered when determining relationships with the clouds present at the previous timestep. G_1, G_2 and G_3 label groups of cloudy gridpoints.

Title Page
Abstract Introduction
Conclusions References
Tables Figures
◀◀ ▶▶
◀ ▶
Back Close
Full Screen / Esc
Printer-friendly Version
Interactive Discussion

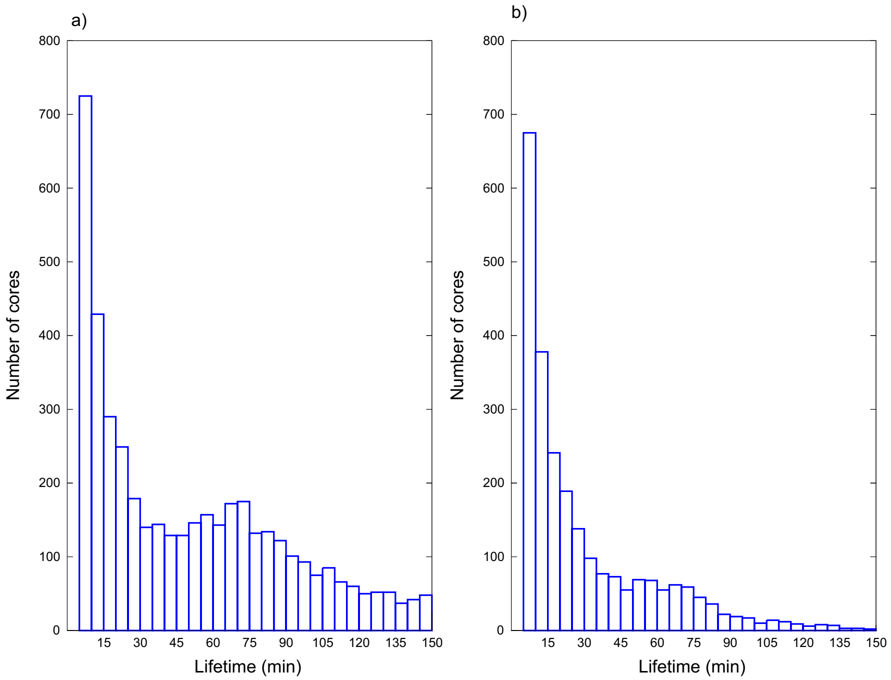

Fig. 2. Convective core lifetime distribution for **(a)** all cores, and **(b)** cores which do not contain events in their time histories. The lifetimes are binned into intervals of 5 min.

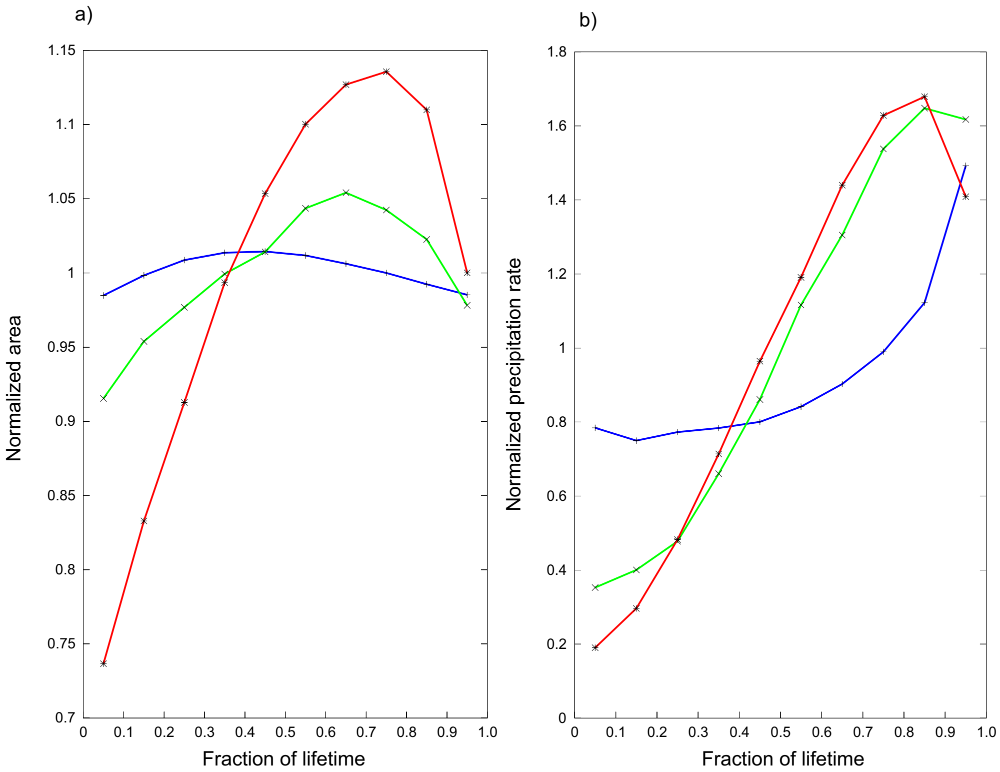

Fig. 3. Timeseries of composited convective cores, showing the mean evolution across the core lifecycle of **(a)** core area, **(b)** precipitation rate, **(c)** vertically-integrated mass flux, and **(d)** the centre of mass. The normalized lifecycle has been divided into bins of 0.1. Each panel shows three lines, each of which corresponds to a composite constructed from cores having a particular range of lifetimes: 5 to 30 min (blue), 30 to 60 min (green) and longer than 60 min (red).

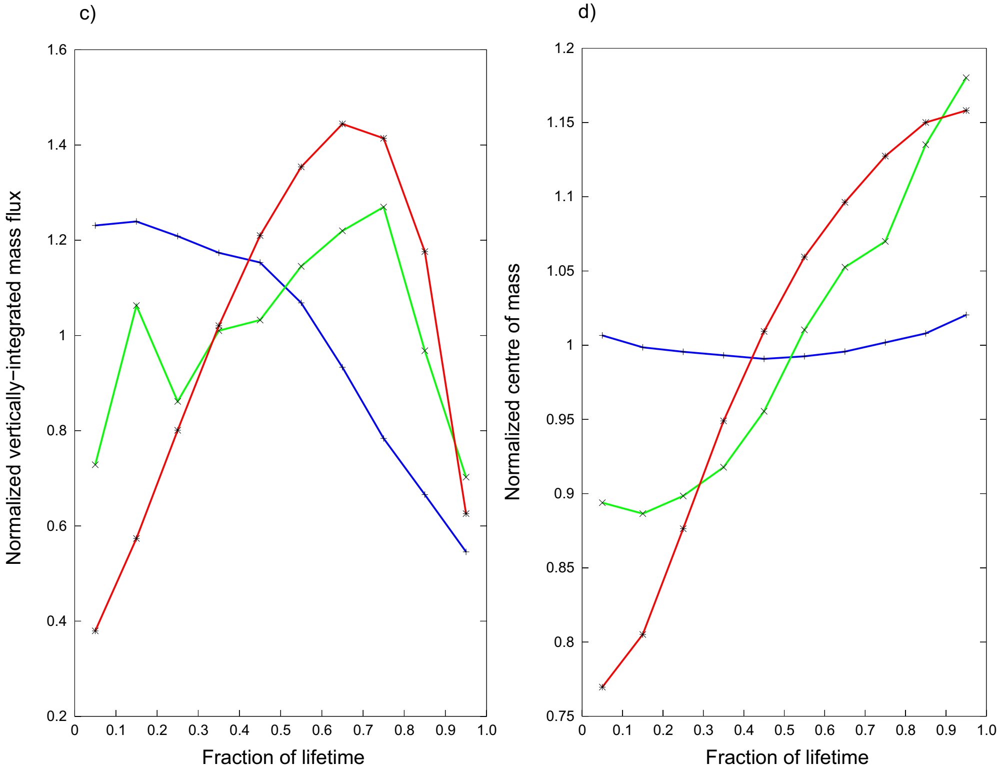

Fig. 3. Continued.

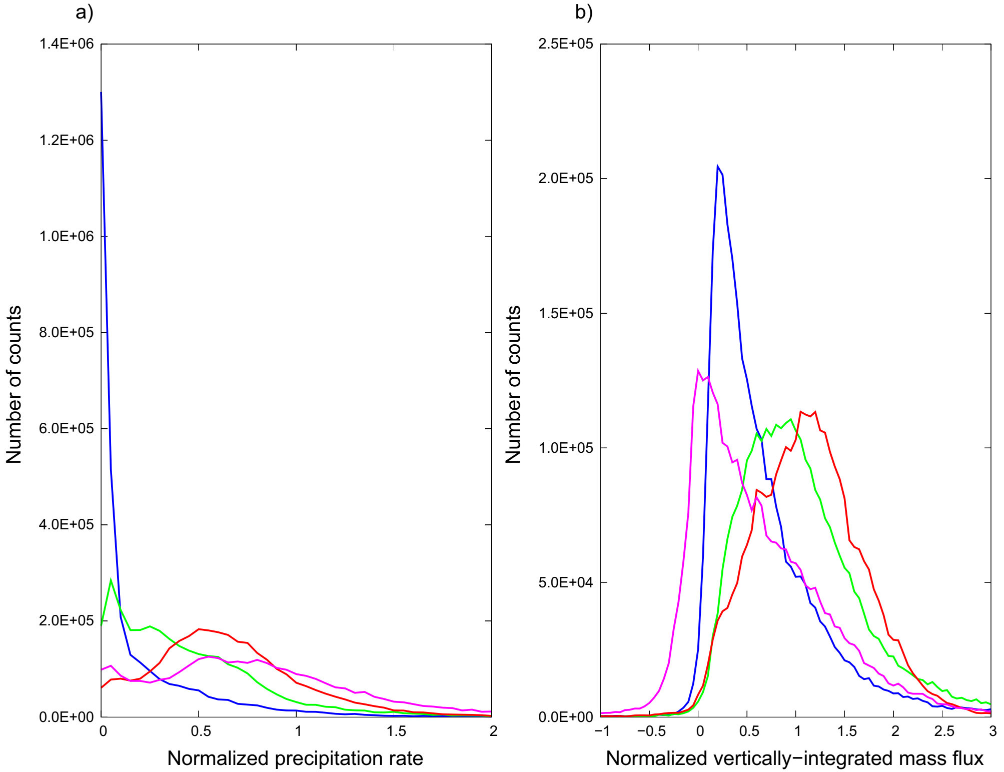

Fig. 4. Probability distribution functions at various stages of the convective-core lifecycle of the normalized **(a)** precipitation rate, and **(b)** vertically-integrated mass flux. Each panel shows four lines, each of which corrresponds to a range of the normalized lifetime: 0 to 0.1 (blue), 0.3 to 0.4 (green), 0.6 to 0.7 (red) and 0.9 to 1.0 (magenta).

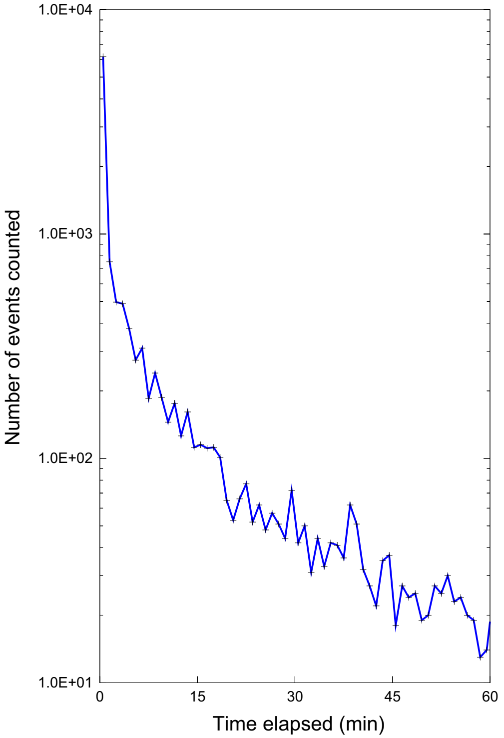

Fig. 5. Distribution of the times separating consecutive events in the lifecycles of tracked convective cores. The vertical scale is logarithmic, and the separation bin size is 1 min.

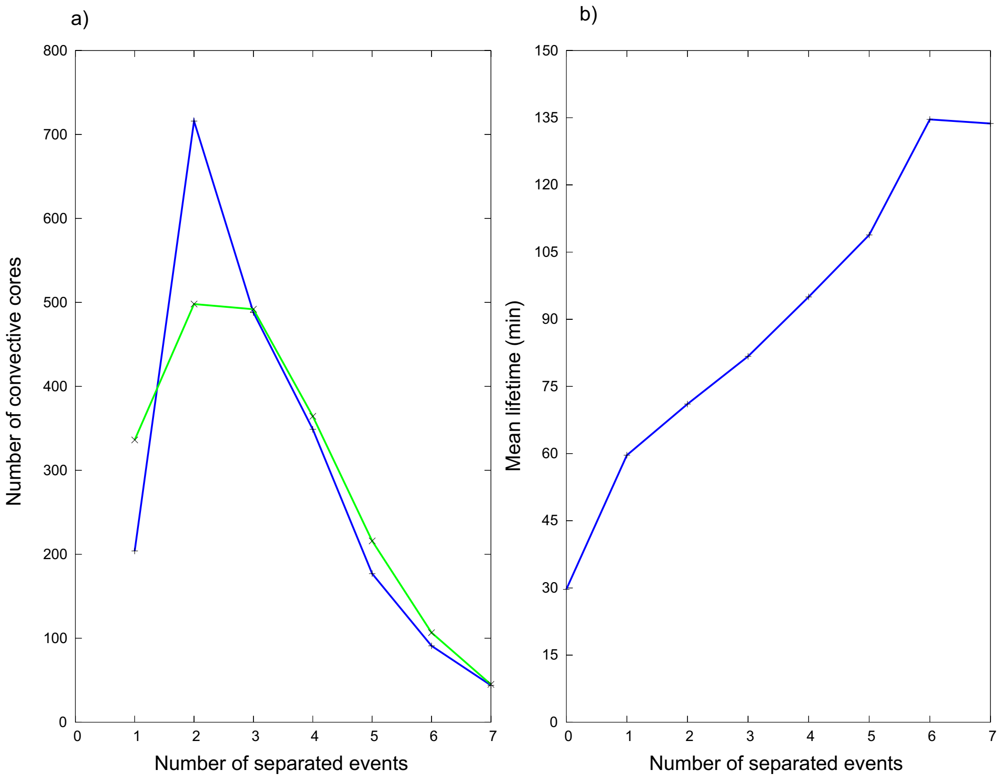

Fig. 6. Panel **(a)** shows the number of convective core lifecycles that contain a given number of separated events (blue). Also shown (green) is a zero-truncated Poisson distribution for the same mean number of separated events. Panel **(b)** shows the mean lifetime of the convective cores for given numbers of separated events.

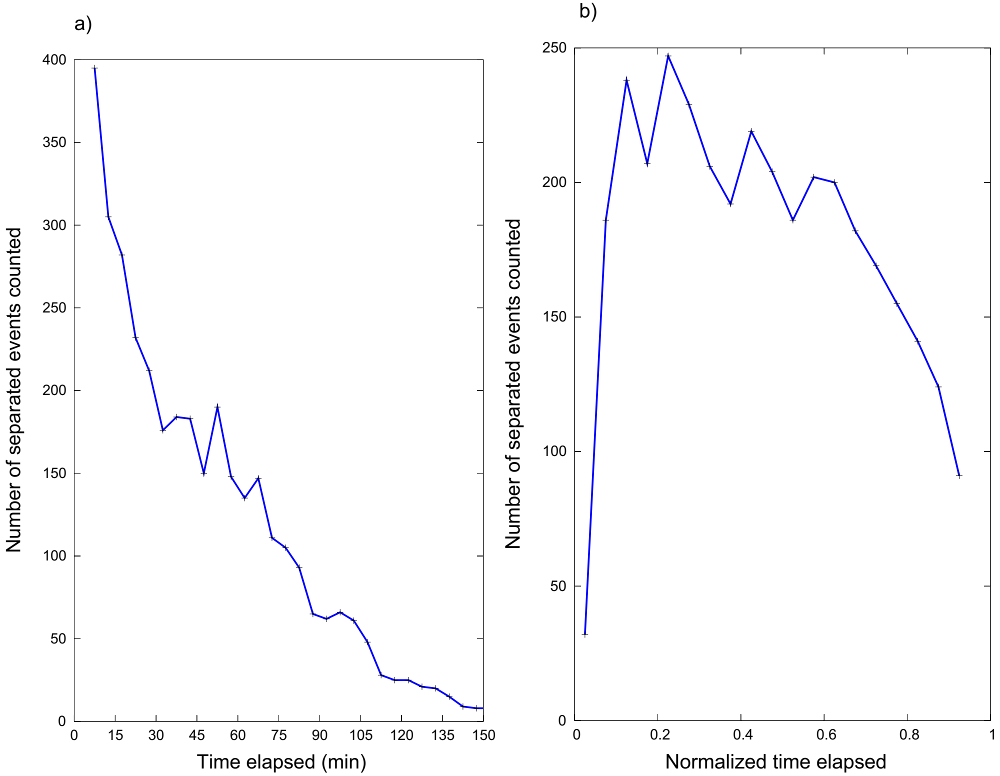

Fig. 7. Distribution of the timings of separated events, both as **(a)** absolute times after the start of a lifecycle, and as **(b)** relative times, with the timings being normalized by the convective core lifetime. The bin size is 5 min in **(a)** and 0.05 in **(b)**.