

Tripleclouds: An Efficient Method for Representing Horizontal Cloud Inhomogeneity in 1D Radiation Schemes by Using Three Regions at Each Height

JONATHAN K. P. SHONK AND ROBIN J. HOGAN

Department of Meteorology, University of Reading, Reading, United Kingdom

(Manuscript received 7 March 2007, in final form 20 July 2007)

ABSTRACT

Radiation schemes in general circulation models currently make a number of simplifications when accounting for clouds, one of the most important being the removal of horizontal inhomogeneity. A new scheme is presented that attempts to account for the neglected inhomogeneity by using two regions of cloud in each vertical level of the model as opposed to one. One of these regions is used to represent the optically thinner cloud in the level, and the other represents the optically thicker cloud. So, along with the clear-sky region, the scheme has three regions in each model level and is referred to as "Tripleclouds." In addition, the scheme has the capability to represent arbitrary vertical overlap between the three regions in pairs of adjacent levels. This scheme is implemented in the Edwards–Slingo radiation code and tested on 250 h of data from 12 different days. The data are derived from cloud retrievals using radar, lidar, and a microwave radiometer at Chilbolton, southern United Kingdom. When the data are grouped into periods equivalent in size to general circulation model grid boxes, the shortwave plane-parallel albedo bias is found to be 8%, while the corresponding bias is found to be less than 1% using Tripleclouds. Similar results are found for the longwave biases. Tripleclouds is then compared to a more conventional method of accounting for inhomogeneity that multiplies optical depths by a constant scaling factor, and Tripleclouds is seen to improve on this method both in terms of top-of-atmosphere radiative flux biases and internal heating rates.

1. Introduction

Clouds have an important role in the radiation budget of the earth (Liou 1986). In a study of marine stratiform cloud by Randall et al. (1984), it was stated that a mere 4% increase in global cloud cover could counteract the warming caused by a doubling of carbon dioxide. This implies that the radiation balance of the earth in general circulation models (GCMs) is very sensitive to the representation of clouds. However, the size of a GCM grid box is never small enough to accurately capture the finescale structure present in most cloud types. Instead, clouds are typically represented as horizontally homogeneous across the fraction of each level of the grid box that they occupy [as in, e.g., all seven forecast models evaluated by Illingworth et al. (2007)]. Using this so-called plane-parallel representation introduces significant biases in the top-of-atmosphere radiative

fluxes (e.g., Cahalan et al. 1994; Pomroy and Illingworth 2000). These biases are caused by the nonlinearity of the relationship between optical depth and either shortwave albedo or longwave emissivity and, in the shortwave, their magnitude is seen to vary with solar zenith angle and surface albedo (Carlin et al. 2002). The challenge is to find computationally efficient methods of accounting for these biases.

From data used by Albrecht et al. (1988), a 15% overestimate was reported in top-of-atmosphere shortwave albedo when the inhomogeneity of stratocumulus was neglected, and a value of between 10% and 15% was reported by Barker and Davies (1992) in an investigation using stochastically modeled, broken cumulus cloud. Carlin et al. (2002) investigated the effect of inhomogeneity on the radiative properties of high cloud. They used cloud radar data to infer a bias in shortwave albedo of 11% when averaged over all solar zenith angles. A further positive bias was identified in longwave emissivity of high cloud in the study of Pomroy and Illingworth (2000), when inhomogeneity was neglected.

In theory, to produce reasonable representations of

Corresponding author address: Jonathan K. P. Shonk, Department of Meteorology, University of Reading, Earley Gate, P.O. Box 243, Reading RG6 6BB, United Kingdom.
E-mail: j.k.p.shonk@reading.ac.uk

the top-of-atmosphere radiative fluxes, low and perhaps unphysical ice or liquid water content values have to be used in the GCM (Harshvardhan and Randall 1985). Despite this, other factors, such as the overlap assumption, may act to offset the bias (Hogan and Kew 2005). These results have implications for the formation of precipitation (Jakob and Klein 1999) and could affect the reliability of its simulation of future climate scenarios.

Methods have been proposed that overcome this plane-parallel bias without significant additional computational cost in the radiation scheme of GCMs. The simplest method is to scale the optical depth of the clouds in the grid box by a certain factor and use this "effective optical depth" in the radiation calculations in place of the true optical depth. A number of values for this factor have been suggested. Cahalan et al. (1994) studied marine stratocumulus and suggested using a reduction factor of 0.7. They recognized that the optimum value would vary for different times of the year and locations on the earth. This scaling factor method has been implemented in the model of the European Centre for Medium-Range Weather Forecasts (Tiedtke 1996), with the factor applied to both ice and liquid clouds over the whole globe, in both the shortwave and longwave. Davis et al. (1990) suggested raising the optical depths to the power of 0.8. Methods involving scaling with a single factor, however, tend to underperform in certain areas of the world and under certain weather regimes, and it is now well recognized that a single value is inappropriate for global use. The diverse values of scaling factor were investigated in the shortwave by Oreopoulos and Cahalan (2005), who used global satellite data from different seasons. They found values from 0.65 to 0.8, depending on time of year, cloud particle phase, and global location, although with more extreme values at certain locations. Scaling factor is also a function of grid box size, as shown by Pomroy and Illingworth (2000). They considered longwave emissivity for ice clouds and evaluated the scaling factor at different grid box sizes. They found scaling factors varying from 0.93 for a 10-km grid box down to 0.47 for a 300-km grid box. Another problem with scaling factor approaches is that scaling factors derived from satellite data or vertically integrated radar measurements will not necessarily be appropriate for application to each model level.

More recently, advanced cloud schemes have been developed that statistically simulate variation in cloud across model grid boxes. Both Tompkins (2002) and Bushell et al. (2003) propose methods to determine cloud fraction by prognosing the variance of water content. The challenge is to use this information in the

radiation scheme. A gamma-weighted radiative transfer scheme was proposed by Barker (1996), who weighted the optical depth across a grid box using a gamma distribution. He found this method to be no more than four times more computationally expensive than the method using homogeneous, plane-parallel clouds with a two-stream approximation. Carlin et al. (2002) considered the effect that the gamma-weighted transfer scheme had on the shortwave albedo. They found that the significant biases introduced by using the plane-parallel method were markedly reduced by applying the scheme. Rossow et al. (2002) also found a gamma distribution to be suitable when modeling optical depth in their observational study, and the suitability of the gamma distribution was compared with that of a lognormal distribution by Hogan and Illingworth (2003), using radar data. However, both the gamma and lognormal distributions have the conceptual disadvantage that they are unbounded at the upper limit, implying a small but nonzero probability of optical depths tending toward infinity. Tompkins (2002) used a beta distribution outside the radiative transfer scheme to describe relative humidity and cloud fraction. This choice has two advantages over the gamma distribution. First, the distribution is bounded, with finite upper and lower limits. Second, it can be both positively or negatively skewed [distributions of relative humidity were shown to skew both positively and negatively by Larson et al. (2001)].

An alternative method is the Monte Carlo independent column approximation (McICA) method (Pincus et al. 2003). A multicolored independent column approximation (ICA) can be used to simulate radiative transfer through inhomogeneous clouds in a GCM grid box, when coupled to a "cloud simulator" that generates columns drawn from the probability distribution function (PDF) predicted by the cloud scheme (Räisänen et al. 2004). The challenge, however, is to achieve a similar accuracy without greatly increasing the computational cost of the scheme. One of the problems with the ICA is the effective double integral; that is, the fact that we are integrating in both space and wavelength when calculating the radiative fluxes, which is very time consuming. The McICA seeks to solve this issue by generating a number of columns containing possible cloud arrangements and performing the radiative calculation for each spectral band using a different column. This means that the integral across both space and wavelength is evaluated at the same time, thereby drastically increasing efficiency. The advantage of the McICA is that, as it effectively represents the independent column case, the output fluxes should have zero bias with respect to the ICA, but with no notable in-

crease in computer run time. However, as the full radiation scheme is not run on each column, random (but unbiased) errors are introduced in the instantaneous fluxes. Despite implementation in a number of models, opinions are still divided as to whether these errors have a significant effect on forecast skill.

In this investigation, we present Tripleclouds, a new method of representing cloud inhomogeneity in radiation schemes. The method is introduced in section 2. It works by using two cloudy regions in each vertical level and a single clear-sky region. One of the cloudy regions is used to represent the optically thinner cloud in the level, while the other represents the optically thicker cloud. For this to work in a standard radiation code, changes must be implemented in both the two-stream solver and the cloud overlap component of the code, as described in section 3. We then perform sensitivity tests on the new scheme. The data and method are discussed in section 4, and the results are presented in section 5. There then follows a comparison between Tripleclouds, the plane-parallel method, and the scaling factor method proposed by Cahalan et al. (1994) in section 6. The results are summarized in the final section.

2. The Tripleclouds method

The goal of the various methods discussed in the previous section has been to account for inhomogeneity in the GCM without drastically increasing computer run time. Using the ICA would give radiative fluxes that are close to reality, within the limitations discussed by O'Hirok and Gautier (2005). However, the computer time required to perform such a run is high, and inappropriate for use in a GCM. In the new method, the continuous PDF of water content in a vertical level of a grid box is modeled using a number of discrete values. Two cloudy regions of different optical depths were used diagrammatically in order to demonstrate the effect of inhomogeneity in longwave emissivity by Pomroy and Illingworth (2000) and in the shortwave albedo by Carlin et al. (2002). If these two regions were actually used to represent the cloud in the radiation scheme, it may therefore be sufficient to remove the plane-parallel biases in both shortwave and the longwave. We now test the hypothesis of whether it is possible to represent the PDF by just two cloudy regions and one clear-sky region.

Figure 1 shows a schematic of how the Tripleclouds method works. Consider first the column of atmosphere within a single GCM grid box from which a single vertical level is extracted. For now, it is assumed that we know exactly how the cloud is distributed across the level in the form of the PDF of water content. In practice, within a GCM, the distribution will

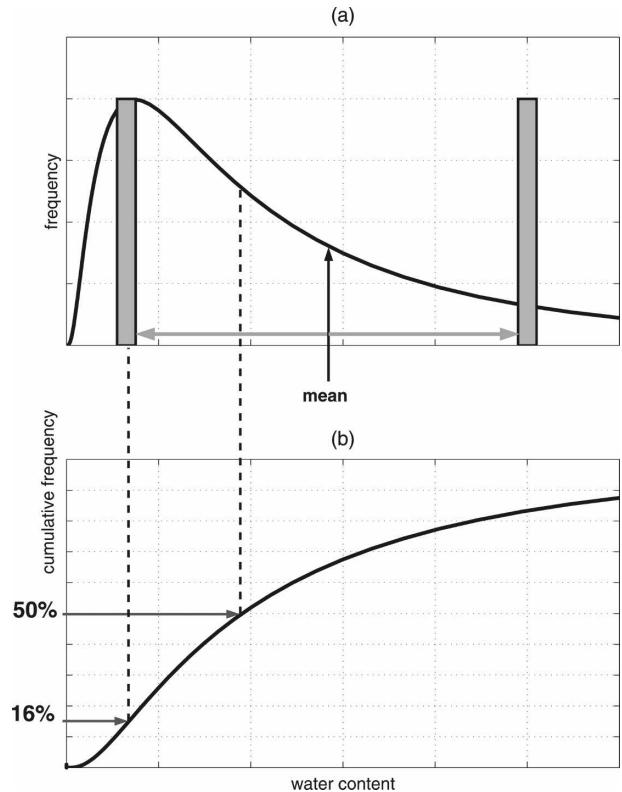


FIG. 1. Schematic showing the method behind Tripleclouds: (a) the continuous PDF of water content and (b) the cumulative frequency. The lines demonstrate how the Tripleclouds calculation is implemented; the two shaded boxes in (a) represent the resulting two-point PDF. See text for a full explanation.

have to be determined either using cloud water variance predicted by an advanced cloud scheme (Tompkins 2002; Bushell et al. 2003) or diagnosed using empirical relationships (e.g., that of Hogan and Illingworth 2003). The Tripleclouds scheme is then applied to the nonzero water content values in the PDF, as shown in Fig. 1. The distribution is divided into two regions through a given percentile of the distribution. This “split percentile” is chosen to be the median, or 50th percentile. This equally splits the cloud into optically thinner cloud (lower water content) and optically thicker cloud (higher water content). A value of “lower percentile” is then chosen. In Fig. 1, this is chosen to be the 16th percentile. The water content value corresponding to this percentile, W_1 , is found, and all the water content values representing the thinner cloud are set to this value. The water content value for the thicker cloud region, W_2 , is chosen such that the mean of W_1 and W_2 is equal to the mean water content of the cloudy part of the level, \bar{W} .

The lower percentile and the split percentile are the two tunable parameters in Tripleclouds, and varying

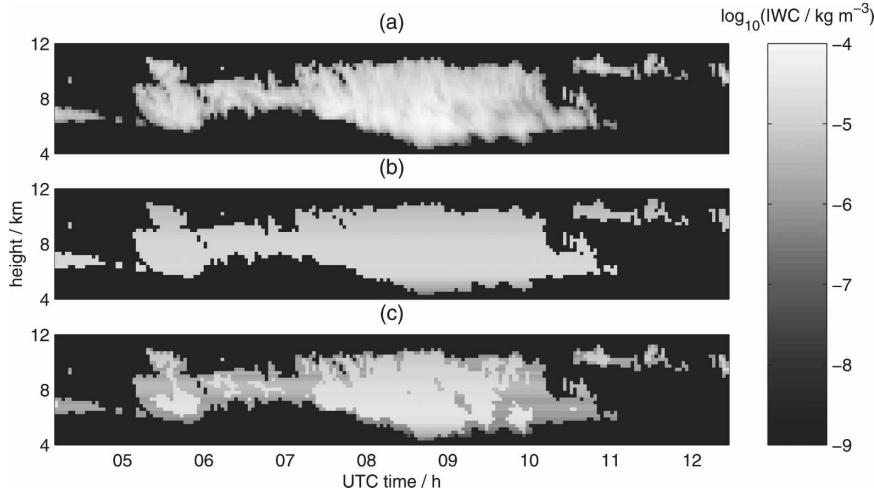


FIG. 2. Comparison of the plane-parallel and Tripleclouds schemes when applied to observational data: (a) just over 8 h of radar retrievals of ice water content taken from Chilbolton, southern United Kingdom, on 28 Jun 2003 and (b), (c) the same data when the plane-parallel and Tripleclouds approximations are applied directly without any parameterizations or assumptions that affect overlap.

these will have the effect of changing the amount of represented cloud inhomogeneity. The values chosen here are the control values and will be justified *a posteriori* by their performance on real data. The choice of 50 for the split percentile is made so that the two cloudy regions both have the same cloud fractions. The value of 16 for the lower percentile has the appealing property that, for a Gaussian PDF, the standard deviation of the resulting two values is the same as that of the full PDF. In practice, this Gaussian fit is only appropriate for cases with low fractional cloud variance, on account of the small but finite probability of negative water content. It should be noted that, for low fractional variance, the Gaussian distribution is a good approximation to both the lognormal and gamma distributions. Assuming we have two water content values W_1 and W_2 , whose mean is \bar{W} , then

$$W_1 + W_2 = 2\bar{W}. \quad (1)$$

Also, for the two water content values to represent the standard deviation σ_w , they must be separated by two standard deviations:

$$W_2 - W_1 = 2\sigma_w. \quad (2)$$

Eliminating W_2 leaves an equation in the lower water content value W_1 :

$$W_1 = \bar{W} - \sigma_w. \quad (3)$$

A Gaussian distribution of W is defined as

$$p(W) = \frac{1}{\sqrt{2\pi}\sigma_w} \exp\left[-\frac{(W - \bar{W})^2}{2\sigma_w^2}\right]. \quad (4)$$

The fraction of the distribution with water content lower than W_1 is therefore given by

$$f_1 = \int_{-\infty}^{W_1} p(W) dW, \quad (5)$$

and, when this is written in terms of standard integrals, it is found that

$$f_1 = \frac{1}{2} \left[1 + \operatorname{erf}\left(-\frac{1}{\sqrt{2}}\right) \right] = 0.159, \quad (6)$$

corresponding to a lower percentile of just under 16.

In Fig. 2, both the Tripleclouds and plane-parallel approximations are applied to radar data from Chilbolton in Hampshire, southern United Kingdom. The data were measured by the 94-GHz vertically pointing radar on 28 June 2003 and have an equivalent horizontal extent of about 450 km (the mean wind speed in the cloud level is of order 15 m s^{-1}). The distribution of ice water content across the domain is as shown in Fig. 2a and is derived from radar reflectivity factor and temperature using the empirical method of Hogan et al. (2006). Figures 2b and 2c show the same data when the plane-parallel and Tripleclouds approximations are applied to them. The raw data show a region of thicker cloud in the center of the domain (between 0800 and 0900 UTC) that is not represented in the plane-parallel case. Using Tripleclouds, however, a region of thicker cloud is clearly represented in this area. The performance of

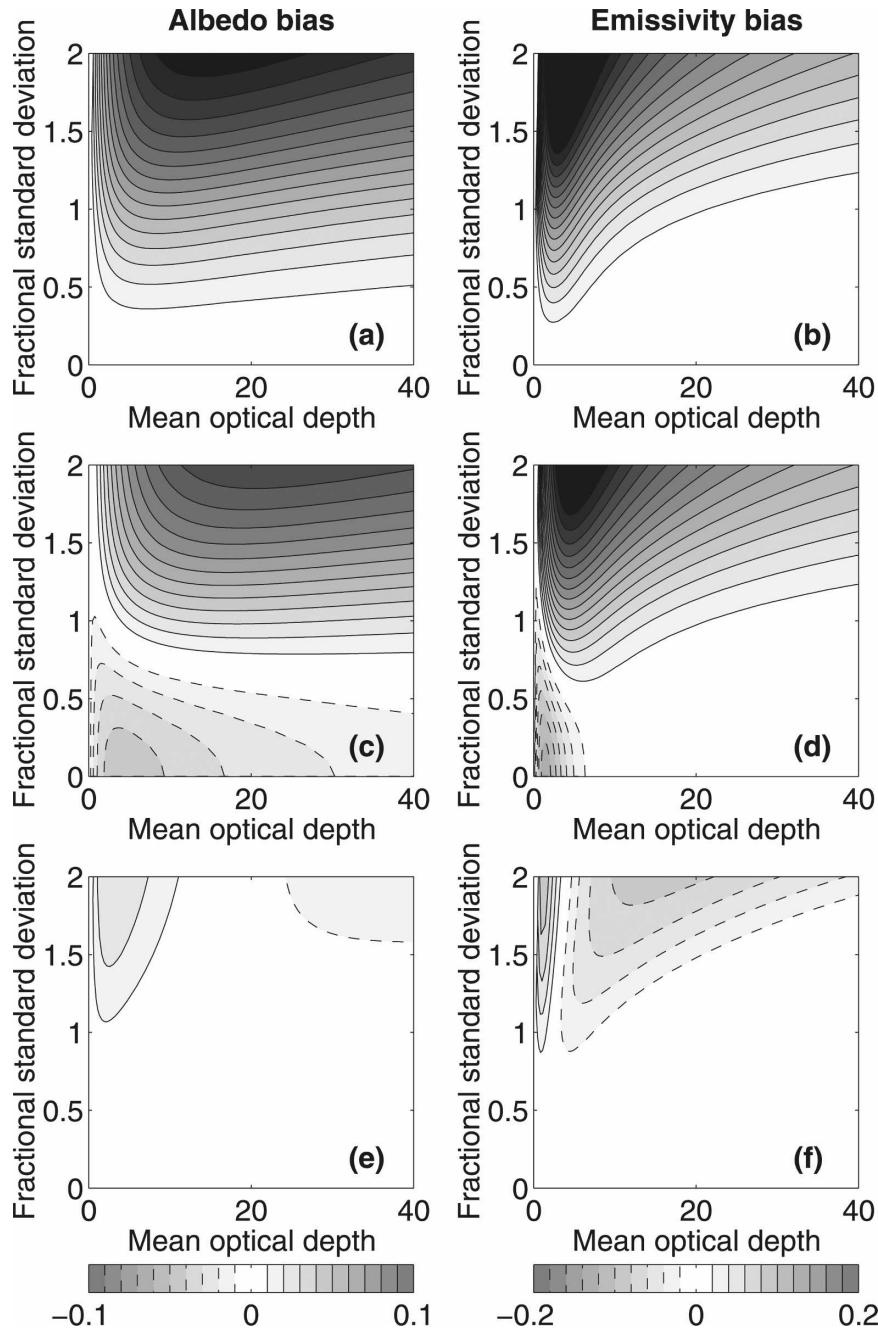


FIG. 3. (left) Shortwave cloud albedo bias and (right) longwave emissivity bias calculated for idealized single-layer ice clouds with a cloud fraction of unity and a lognormal distribution of optical depth, calculated as a function of the mean optical depth of the cloud and the fractional standard deviation of optical depth across the grid box (equal to the standard deviation of the natural logarithm of optical depth). (a), (b) The biases for the plane-parallel approximation; (c), (d) the biases for the scaling factor method with a value of 0.7; and (e), (f) the biases for Tripleclouds with a lower percentile of 16 and a split percentile of 50. Contour interval is 0.01 for albedo bias and 0.02 for emissivity bias. For the purposes of this figure, albedo α has been calculated from optical depth δ using $\alpha = 0.2 + 0.525\delta/(\delta + 3.5)$, which reproduces the albedo predicted by the Edwards–Slingo code to within 0.02 for the control conditions considered in section 5 (the same approximation was used in Fig. 11 of Hogan and Kew 2005). Likewise, we follow Pomroy and Illingworth (2000) and calculate emissivity ε from the absorption optical depth δ_a (assumed to be equal to $\delta/2$) using $\varepsilon = 1 - \exp(-1.66\delta_a)$, which also closely reproduces the full calculations of the Edwards–Slingo code.

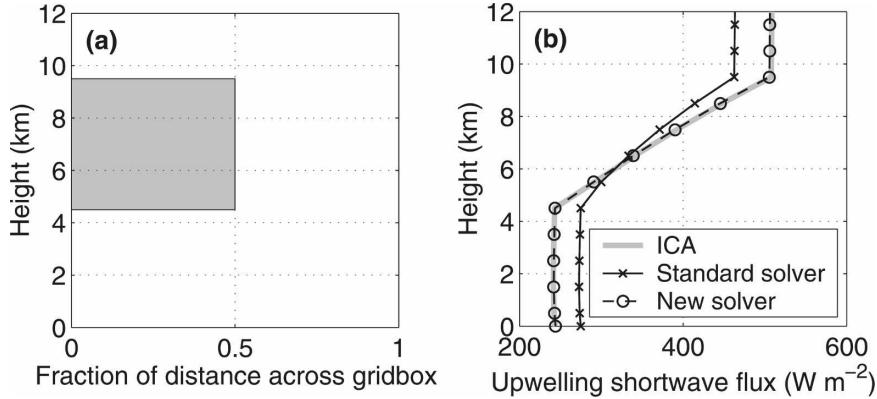


FIG. 5. (a) Demonstration of the problem of anomalous horizontal photon transport, and its solution, for a single level of ice cloud with a visible optical depth of 20, an effective radius of 50 μm , and a cloud fraction of 0.5, as illustrated by the gray region. (b) The corresponding upwelling shortwave fluxes calculated by the Edwards–Slingo code for a solar zenith angle of 45°, a surface albedo of 0.5, and the *U.S. Standard Atmosphere, 1976* (McClatchey et al. 1972). Results are shown for the ICA (in which the code was run independently on the clear and cloudy halves of the column), the “standard” solver utilizing (14) and (15) to describe the relationship between the fluxes, and the “new” solver described at the end of section 3b.

$i + 1$ (and similarly for $V_{i+1/2}^{aa}, V_{i+1/2}^{ab}$, etc.). Likewise, $U_{i+1/2}^{ab}$ is the fraction of upwelling radiation leaving region a at level $i + 1$ that enters region b at level i . Thus the equation defining the downwelling flux in region a just above a half-level is

$$F_{i+1/2}^{a-} = S_i^{a-} + T_i^a(V_{i-1/2}^{aa}F_{i-1/2}^{a-} + V_{i-1/2}^{ba}F_{i-1/2}^{b-} + \dots) + R_i^a(V_{i+1/2}^{aa}F_{i+1/2}^{a+} + V_{i+1/2}^{ba}F_{i+1/2}^{b+} + \dots), \quad (14)$$

and for the upwelling flux:

$$F_{i-1/2}^{a+} = S_i^{a+}U_{i-1/2}^{aa} + S_i^{b+}U_{i-1/2}^{ba} + \dots + T_i^aU_{i-1/2}^{aa}F_{i+1/2}^{a+} + T_i^bU_{i-1/2}^{ba}F_{i+1/2}^{b+} + \dots + R_i^aU_{i-1/2}^{aa}(V_{i-1/2}^{aa}F_{i-1/2}^{a-} + V_{i-1/2}^{ba}F_{i-1/2}^{b-} + \dots) + R_i^bU_{i-1/2}^{ba}(V_{i-1/2}^{ab}F_{i-1/2}^{a-} + V_{i-1/2}^{bb}F_{i-1/2}^{b-} + \dots) + \dots, \quad (15)$$

where the ellipses indicate similar terms corresponding to third and further regions. The reason that (14) and (15) have different forms is that $F_{i+1/2}^{a+}$ and $F_{i+1/2}^{a-}$ are the fluxes just above a half-level, that is, at the bottom of a model level. Therefore, the downwelling flux at the bottom of region a depends primarily on the properties of region a (R_i^a , etc.), whereas the upwelling flux depends on the properties of all of the regions in the level below, each weighted by the appropriate transfer coefficient. If we had defined the fluxes as being just below a half-level, then the general form of (14) and (15) would have been reversed.

It can be seen that for N regions, each component of the flux depends on $2N$ other flux components, so each row and column of the matrix to be solved will have $2N + 1$ nonzero elements. Naturally, this is much more computationally expensive to solve than a tridiagonal system, but it will now be shown that a modification can be made that results in both much closer agreement with the independent column approximation (ICA) and more efficient code.

The principal point to note is that, in (15), terms involving $F_{i-1/2}^{b-}$ correspond to downwelling radiation at the base of region b that is immediately reflected back up into region a . This is “anomalous horizontal photon transport”; for the horizontal resolution of most global forecast and climate models, only a very small fraction of such transport would occur, and in the ICA this transport is zero. The corresponding error in shortwave domain-mean upwelling flux for a very simple scenario is demonstrated in Fig. 5, which shows a comparison of ICA and the standard two-region solver in the Edwards–Slingo code. By setting the offending terms in (15) to zero, the matrix problem would be more efficient to solve, but this would not significantly reduce the error in the shortwave because the downwelling radiation beneath the clear-sky region of a grid box is still reflected back into both cloudy and clear-sky halves of the grid box, when in the ICA it would remain within the clear half and lead to a larger upwelling flux at the top of the atmosphere. Note that the longwave errors are much smaller because of the much reduced role for scattering.

This problem is solved by appealing to the physical

interpretation of Gaussian elimination presented in section 3a. We recursively define $\alpha_{i-1/2}^a$ as the albedo of the atmosphere below region a of level i by adapting (9) to multiple regions:

$$\begin{aligned} \alpha_{i-1/2}^a = & V_{i+1/2}^{aa} [R_i^a + (T_i^a)^2 \beta_i^a \alpha_{i+1/2}^a] \\ & + V_{i+1/2}^{ab} [R_i^b + (T_i^b)^2 \beta_i^b \alpha_{i+1/2}^b] \\ & + \dots, \end{aligned} \quad (16)$$

and similarly for $G_{i-1/2}^a$ (and likewise for regions b, c , etc.).

Thus, the complex expression for the upwelling flux in each region given by (15) is replaced by the much simpler expression (12) applied separately to each region. Physically, this means that, at each height, the downwelling radiation in a particular region is either reflected back up into the same region or absorbed; none is reflected up into another region. Therefore, anomalous horizontal photon transport is almost completely eliminated. This new scheme is also more efficient since, rather than solving a dense matrix problem, only two passes are required: the first up through the atmosphere to calculate the α and G terms, followed by a downward pass to calculate the fluxes.

The new scheme has been implemented in the Edwards–Slingo code for two and three regions at each height. Figure 5 demonstrates the much-improved performance of the new scheme compared to ICA. The small residual error is due to the fact that anomalous photon transport is only eliminated when it is associated with downward radiation being reflected back upward. Upward radiation may still be reflected downward into a different region, although generally this results in a much smaller error.

c. Overlap considerations

As well as the amendments in the previous subsection, the Edwards–Slingo code needed further modifications before it could be used to compare the performance of Tripleclouds with that of other methods. The cloud component of the code currently applies a form of maximum-random overlap (Geleyn and Hollingsworth 1979). Regions of both cloud and clear sky in adjacent levels are assumed to overlap maximally, and any cloud or clear sky unaccounted for by maximum overlap is assumed to overlap randomly. It would be a trivial exercise to extend this overlap method to a cloud scene with three regions at each height. As this investigation is exclusively considering representations of inhomogeneity, biases introduced by any other potential sources should be minimized. Hence, we require the ability to recreate any arrangement of overlap for any

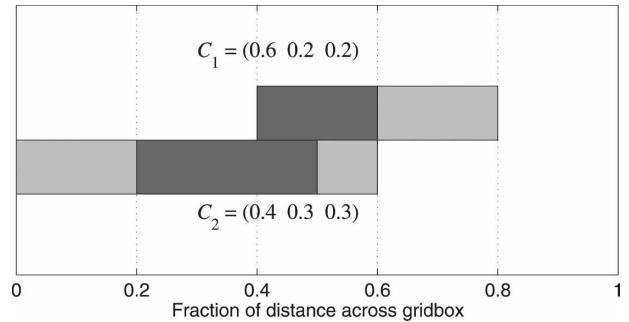


FIG. 6. An example of overlap between two levels of cloud. The pale gray region represents the thinner cloud; the dark gray region represents the thicker cloud. The vectors C_1 and C_2 are the cloud fractions in each level, where the first number in the vector represents the clear-sky fraction, the second represents the thinner cloud fraction, and the third represents the thicker cloud fraction. The resulting overlap matrix for these two levels is given in Eq. (18).

cloud scene, and the maximum-random overlap method currently applied in the Edwards–Slingo code is incapable of this.

For a pair of adjacent levels, each with three regions, the overlap can be described by a three-by-three “overlap” matrix:

$$\mathbf{O} = \begin{pmatrix} X_{c,c} & X_{c,l} & X_{c,h} \\ X_{l,c} & X_{l,l} & X_{l,h} \\ X_{h,c} & X_{h,l} & X_{h,h} \end{pmatrix}, \quad (17)$$

where $X_{a,b}$ represents the areal fraction of the column covered by both region a in the upper level and region b in the lower level. The labels c, l , and h refer to the clear sky, thinner cloud (lower water content), and thicker cloud (higher water content) regions, respectively. An example of cloud overlap between two levels is shown in Fig. 6 along with the cloud fraction vectors C_1 and C_2 . In this case, the overlap matrix is

$$\mathbf{O} = \begin{pmatrix} 0.2 & 0.2 & 0.2 \\ 0.2 & 0.0 & 0.0 \\ 0.0 & 0.1 & 0.1 \end{pmatrix}. \quad (18)$$

The Tripleclouds scheme currently outputs three quantities: the fractions of each level that are assigned to each cloud region; the water content values of each of these regions in each level, and a three-by-three overlap matrix for each interface, as in (18). These can then be used as inputs for the radiation code that determine the exact arrangement of the clouds without any need for overlap approximations, and are used to calculate the transfer coefficients U and V in (14) and (15). The overlap matrices themselves are entered in the code in the form of a four-column vector, contain-

ing four of the elements of the overlap matrix at each interface. From these four values, in addition to the vector C in each level, it is possible to calculate the remaining terms in the matrix, using the condition that, for a given interface, the sums of the columns of the overlap matrix correspond to the region fractions in the level above and the sums of the rows give the fractions in the level below. There is also the condition that, by definition, the elements of the matrix must all add up to one.

Amendments have been made to the Edwards–Slingo code that allow the overlap to be read in and used in this way. The modified version of the code will be used throughout this investigation.

4. Observational data and experimental method

For the purposes of this investigation, we use fields of ice and liquid water content (IWC and LWC) derived from a combination of radar, lidar, and microwave radiometer measurements as part of the Cloudnet project. The retrieval processes used in Cloudnet are summarized by Illingworth et al. (2007): IWC is derived from radar reflectivity and temperature (Hogan et al. 2006), while LWC is derived from measurements of cloud base, cloud top, and liquid water path. These datasets are all derived from the 94-GHz vertically pointing radar, the infrared lidar ceilometer, and microwave radiometers at Chilbolton and have a vertical resolution of 60 m and temporal resolution of 30 s. The raw water content data are then averaged both horizontally and vertically. Horizontal averaging is used to reduce the time taken to perform the ICA runs by a factor of 5; vertical averaging is performed so that the resolution is comparable to that of a GCM. The resulting resolutions are 240 m in the vertical and 2.5 min in the horizontal. A total of 250 h of data are used in this study, taken from 12 different days in the period from January 2003 to March 2004. Each daylong period is split into 10 scenes of length 125 min, starting from midnight. The resulting 120 scenes are all single phase: half of them contain only ice cloud, half of them contain only liquid water cloud. Assuming a horizontal wind of 5 m s^{-1} in the boundary layer and 15 m s^{-1} in the free atmosphere, this results in scene sizes of between 40 and 120 km. The ice scenes feature predominantly frontal cases, selected to represent a very wide range of cloud geometries, while the liquid scenes feature mainly stratus and stratocumulus clouds with physical depths of between 500 m and 2 km, a range of optical depths, and varying degrees of inhomogeneity. Figure 7a shows histograms of mean IWC (pale gray) and LWC (dark gray) for the cloudy regions in each of the

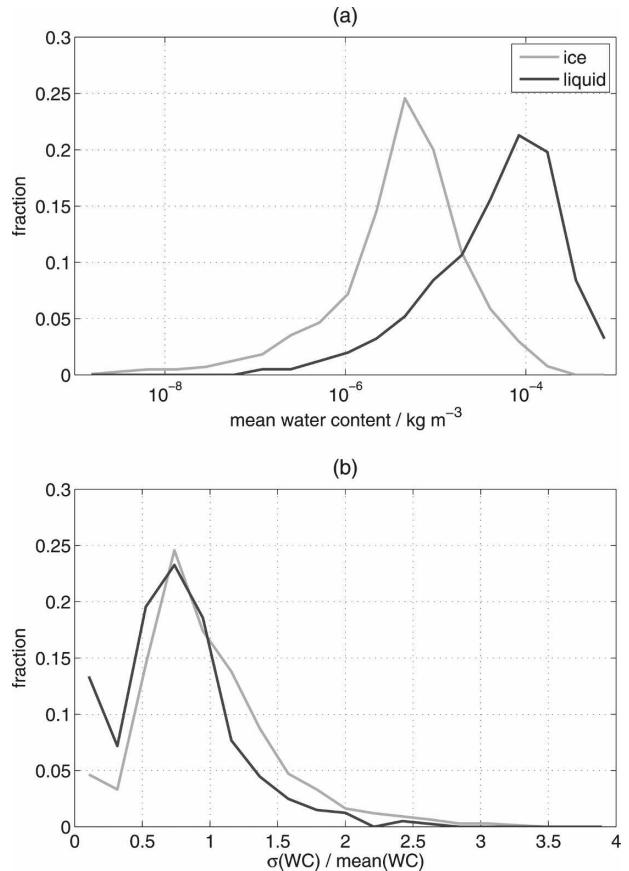


FIG. 7. Histograms showing the variation in the statistical properties of water content across the scenes: The distributions of (a) mean water content and (b) fractional standard deviation of water content are shown for the regions containing cloud in each vertical level in each of the 120 scenes.

levels, while Fig. 7b shows the spread of the same quantities in the form of fractional standard deviation.

A number of calculations were performed on each scene, using both the ICA and a number of “single column” runs. A full list of all runs performed on each scene is given in Table 1. The ICA runs integrate each high-resolution column of the observational data individually, using one homogeneous region in each vertical level (i.e., no partial cloudiness), then average the fluxes across the scene. The single-column runs are made to mimic the operation of a GCM. The Tripleclouds outputs (cloud fractions and water content for each region in each level and overlap at each interface; see section 3c) are used as inputs to the Edwards–Slingo code in such a way that the code processes the entire scene in a single integration.

5. Initial experiments and sensitivity tests

In this section, we investigate the sensitivity of the performance of Tripleclouds to lower percentile, split

TABLE 1. List of all the runs performed on each of the scenes for the rest of this experiment.

Independent column	ICA control run ICA Tripleclouds run: "TCICA"
Two region, single column	Plane-parallel run: "PP" Single-column scaling factor, with a scaling factor of n : "SF(n)"
Three region, single column	Tripleclouds runs, with varying percentiles, vertical resolution, and horizontal gridbox size: "TC(n)," where n is lower percentile

percentile, horizontal grid box size, vertical resolution, and solar zenith angle. A single-column Tripleclouds run is performed on each of the scenes, then the biases are evaluated with respect to the ICA runs (using the ICA runs as "truth"). Each of the variables is then changed to test the sensitivity. For this experiment, certain values are held constant: surface albedo is set to 0.2, and effective radius values for ice particles and liquid droplets are set to 50 and 5 μm , respectively. For the thermodynamic profile, the *Standard U.S. Atmosphere* is used (McClatchey et al. 1972). The shortwave calculation uses five spectral bands; the longwave calculation uses nine spectral bands. The "control" values of lower and split percentile are 16 and 50, respectively. For vertical resolution, the control value is 240 m; for solar zenith angle, it is 60°.

Before testing sensitivity, we calculate the value of mean plane-parallel shortwave albedo bias across the scenes. All biases throughout this experiment are calculated in terms of the change to the "cloud radiative forcing." First, the upward clear-sky longwave and shortwave fluxes are calculated at the top of the atmosphere. These upward top-of-atmosphere fluxes are then subtracted from the clear-sky fluxes across each of the scenes, leaving only the change in net top-of-atmosphere flux that is due to the presence of clouds. This is called the cloud radiative forcing (CRF). By definition, this quantity will be negative in the shortwave and positive in the longwave. Throughout this paper, however, when referring to CRF in the shortwave, we shall take the modulus such that it is always a positive quantity. This enables the longwave and shortwave CRFs to be more easily compared. CRF is calculated both for the control ICA runs and the experimental runs, and the percentage error for each scene is calculated as

$$\text{percentage error} = 100 \times \frac{\text{CRF}(X) - \text{CRF}(ICA)}{\text{CRF}(ICA)}, \quad (19)$$

where $\text{CRF}(X)$ is the cloud radiative forcing calculated in experiment "X." The percentage biases are then calculated across the scenes as

$$\text{percentage bias} = 100 \times \left(\frac{\overline{\text{CRF}(X)} - \overline{\text{CRF}(ICA)}}{\overline{\text{CRF}(ICA)}} \right). \quad (20)$$

Percentages will be used so that our results can be directly compared to those derived from Albrecht et al. (1988). The scenes with little or no cloud are excluded from the analysis, as they could give a high percentage bias in CRF for a very small absolute bias. A domain-mean optical depth threshold of 0.01 is used to discriminate between the scenes with sufficient cloud to use, which leaves 98 "cloudy" scenes.

a. Plane-parallel albedo bias

Figure 8 shows a scatterplot comparing the errors in shortwave and longwave CRF for the 98 cloudy scenes against the domain-mean optical depth at 550 nm. The scenes are separated into the ice cases (filled circles) and liquid cases (open circles). A positive plane-parallel CRF error is seen for almost all of the scenes, with values up to nearly 30% being calculated. The mean bias is found to be of order 8%, which is lower than the value of 15% from Albrecht et al. (1988) in their field study, although the order of magnitude is comparable. However, we do expect a lower value for our study: we are considering approximately GCM-gridbox-sized scenes, while Albrecht et al. (1988) used a much longer time series of cloud data. In contrast, the mean bias using Tripleclouds is less than 1%. The other points in Fig. 8 are described in section 6.

b. Sensitivity to lower percentile

In this experiment, we vary the lower percentile from 16 and run Tripleclouds on each of the 98 cloudy scenes at each value of lower percentile. The results from this are shown in Fig. 9. The black lines show the mean biases at each lower percentile value when averaging over all the cloudy scenes. The solid gray lines show separately the mean biases over the cloudy ice and liquid scenes, and the dashed gray lines show the edges of the range within the 10th and 90th percentiles of the errors for the individual scenes and each lower percentile value. This range is hereafter referred to as the

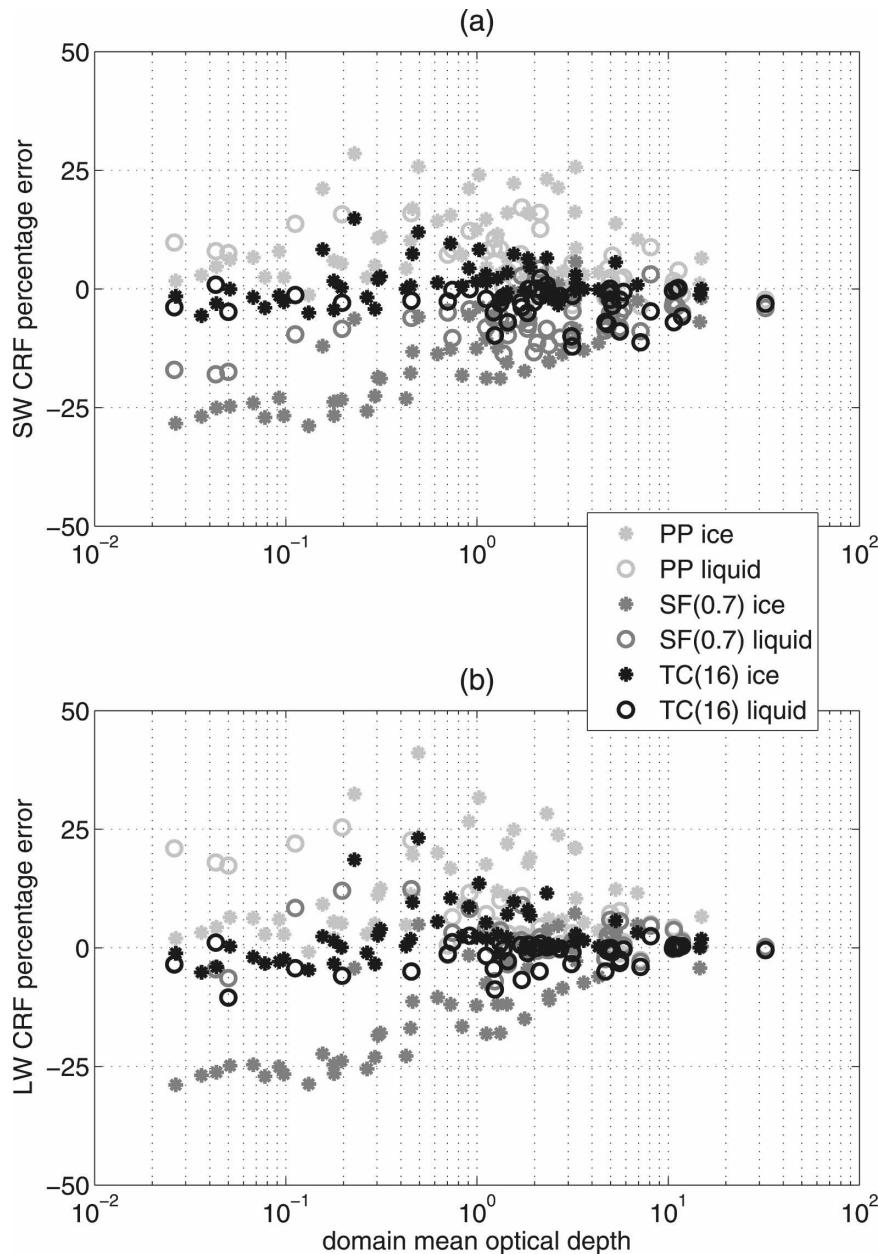


FIG. 8. Plot of the cloud radiative forcing errors of the 98 cloudy scenes as a function of domain-mean optical depth using different schemes: abbreviations for the runs defined in Table 1.

“spread.” The results show that there is a difference in mean biases across the ice scenes and the liquid scenes. This is caused by the difference in mean optical depth across the cases: the liquid cases have a higher mean optical depth, as shown in Fig. 7. A similar result is found when applying the Tripleclouds scheme to the idealized case used in Fig. 3. At high values of fractional variance, we see that there is a switch in sign of the biases when optical depth increases.

We also find that there is no one unique value of lower percentile that creates a zero bias in both shortwave and longwave CRF, although it should be noted that the size of the relative biases generated by varying percentile is small with respect to the plane-parallel biases. For all values of percentile shown, the mean bias never exceeds the 8% plane-parallel bias found in the previous subsection. In fact, in the range of lower percentile from 10 to 20, the mean bias never exceeds a

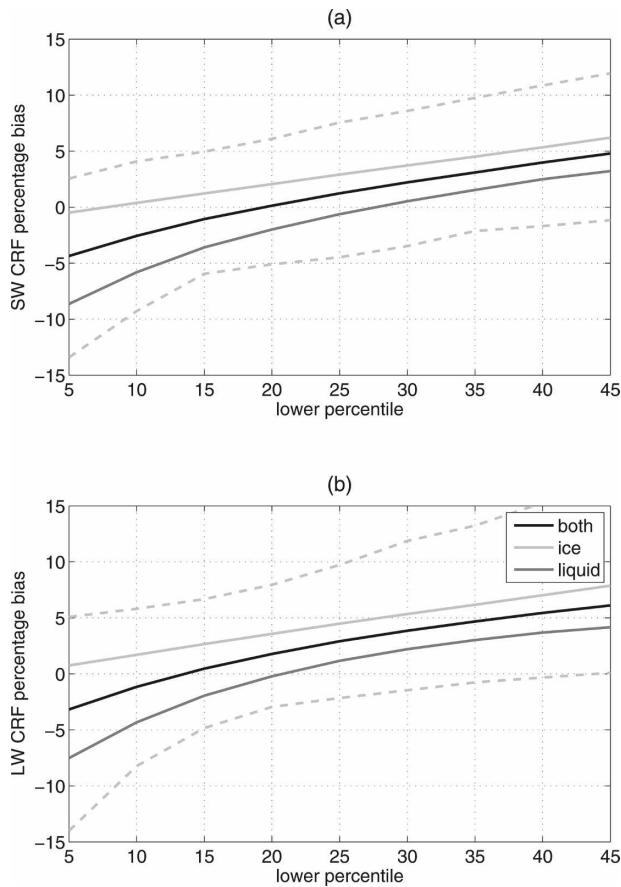


FIG. 9. Variation of mean cloud radiative forcing bias with varying lower percentile when the 98 cloudy scenes are approximated using Tripleclouds: (a) the variation of percentage bias in shortwave CRF with respect to the ICA run and (b) the variation of percentage bias in longwave CRF. The black lines show the mean biases over the scenes at each lower percentile value, and the gray dashed lines show the edges of the spread in error bounded by the 10th and 90th percentiles at each lower percentile value. The solid gray lines show the mean biases over just the ice scenes (pale gray) and the liquid scenes (dark gray).

magnitude of 2%, and the spread never gives an error above 10%.

The optimum value of lower percentile that minimizes the mean biases seems slightly different for the shortwave and longwave. In the shortwave, the optimum is about 19, while the longwave optimum is about 13. Using a value of 16 for both spectral regions, therefore, should not cause significant biases to be introduced into the radiative calculations; in fact, much less than 1% in terms of CRF. There is also an increase in bias with rising lower percentile. This is as expected: increasing the lower percentile has the effect of making the representation more homogeneous, and positive biases imply a more homogeneous case. The value of 16, proposed in section 2 on theoretical grounds, will be

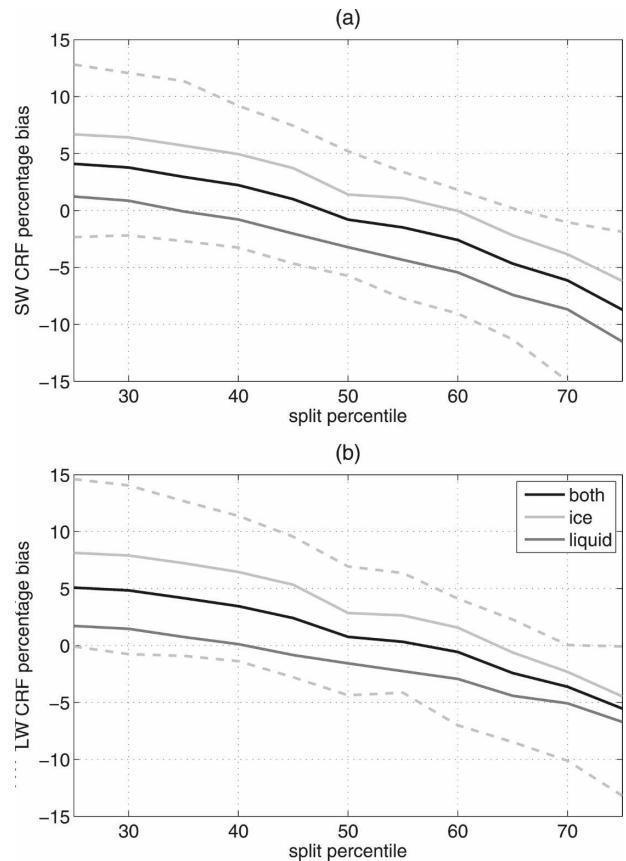


FIG. 10. As in Fig. 9, but for variation of cloud radiative forcing bias with varying split percentile using Tripleclouds.

retained as the lower percentile on account of the fact that, under a lower percentile of 16, the Tripleclouds scheme performs well in both longwave and shortwave parts of the spectrum and for scenes containing ice and liquid water.

c. Sensitivity to split percentile

Next, we investigate the sensitivity of the CRF biases to the split percentile, while the lower percentile is held constant. Variation of the split percentile has more than one effect on the cloud representation. Raising the split percentile allocates more of the cloud to the thinner cloud region, so increasing the weighting of the lower water content value. The water content for the thicker region will then not only be weighted less, but will also be forced to a higher value to ensure conservation of the mean. The results of the split percentile sensitivity experiment are shown in Fig. 10.

Increasing the split percentile is seen to have the effect of making the biases more negative. So, a higher value of split percentile implies a less plane-parallel behavior and, hence, more inhomogeneous representa-

tion of the cloud. Considering the mean biases over all ice and liquid cases, varying split percentile between 25 and 75 gives mean relative biases not exceeding 8%, and for the chosen value of 50 the mean bias is not more than 1%. This indicates that the choice of 50 for the split percentile is suitable. Again, the biases for the ice scenes suggest a slightly higher value, and the liquid scenes suggest a slightly lower value. The overall average biases point toward an optimum value of 49 in the shortwave and of 56 in the longwave. However, the scenes that were used were chosen for the fact that they represent a wide variety of different cloud types. On account of this wide representivity, we can conclude that, for most cases, the optimum value of split percentile is not significantly different from 50.

Lower percentile and split percentile are not independent of each other in terms of CRF bias. In these two sensitivity tests, it was found that 16 and 50 were a suitable combination, but the dependence implies that, for any value of split percentile that we choose, a value of lower percentile can be found that minimizes the biases. In reality, the tuning of the lower percentile to 16 is only valid for the chosen split percentile value of 50, and this second sensitivity experiment merely confirms that a split percentile of 50 is appropriate for use with a lower percentile of 16.

d. Sensitivity to horizontal gridbox size

We next consider the effect of varying horizontal gridbox size on the CRF biases from both Tripleclouds and the scaling factor method. The vertical resolution of the input data is held at 240 m, and the horizontal resolution of the input radar data is kept at 2.5 min. The period from each of the 12 days containing the 98 scenes is then divided into different numbers of scenes with varying horizontal size. As the horizontal scale of the observational data is temporal, the length in space of the scenes will be different in the free atmosphere from that in the boundary layer as the distance–time relationship for ice and liquid cases will be different. For this reason, the horizontal axis is expressed in terms of time. Using values of 5 and 15 m s^{-1} as typical wind speeds for the liquid cloud and ice cloud levels, respectively, the time axis can be considered to run from about 10 to 180 km for the ice cases and 4 to 60 km for the liquid cases. For comparison purposes, runs are also performed with the scaling factor method using a factor of 0.9.

Figure 11 shows the results. In the longwave, the bias for the scaling factor method is seen to decrease with increasing gridbox size, in agreement with the findings of Pomroy and Illingworth (2000). In the shortwave, it is found that the same is also true: there is a depen-

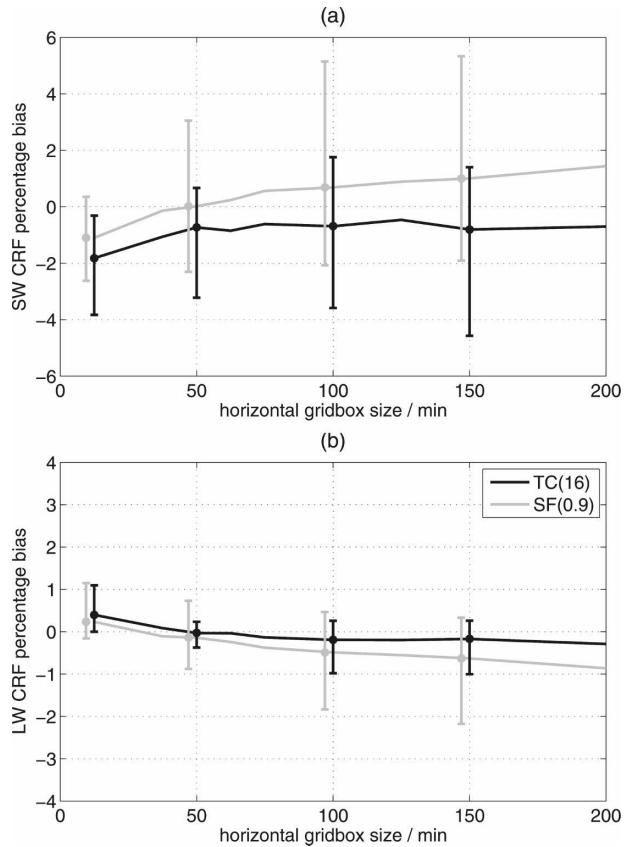


FIG. 11. Variation of cloud radiative forcing bias when the 12 days of data are evaluated using Tripleclouds with lower percentile of 16 and the scaling factor method with a scaling factor of 0.9 with varying horizontal gridbox size. The lines represent the mean percentage bias across all the scenes at each gridbox size for Tripleclouds (black) and scaling factor (gray). The error bars show the spread of individual scenes bounded by the 10th and 90th percentiles. The scaling factor error bars are offset slightly for clarity.

dence of the bias on gridbox size, although the bias is seen to increase with increasing gridbox size. In contrast, we see that Tripleclouds has a weak dependence on horizontal gridbox size except at lower gridbox sizes, where there is a dependence, with decreasing gridbox size increasing bias in the longwave and decreasing bias in the shortwave. The mean bias for Tripleclouds never exceeds 2% through the range shown, and the spread lines never exceed 5% in the shortwave or 2% in the longwave.

e. Sensitivity to vertical resolution

Next, the scenes are evaluated with varying values of vertical resolution, and the mean biases evaluated as before. Variation of vertical resolution is seen to have little effect on the biases in CRF, both in the shortwave and longwave for the Tripleclouds scheme; hence the

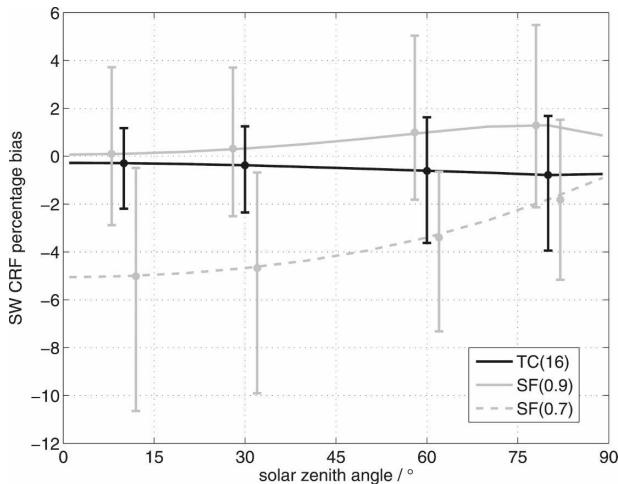


FIG. 12. Variation of shortwave cloud radiative forcing bias, when the 98 cloudy scenes are evaluated using Tripleclouds and the scaling factor method with varying solar zenith angle; lines as described in Fig. 11. The solid gray line uses a scaling factor of 0.9; the dashed gray line uses a scaling factor of 0.7.

figure for this is not shown. The magnitude of the biases in the longwave is small with Tripleclouds, with the mean bias never exceeding 1% and the spread never exceeding 3%. In the shortwave, the result is similar, with negligible biases introduced by considering varying resolution. Even the outliers in the range rarely introduce biases over 1%. A similar result is obtained from the scaling factor method, with the mean bias over all scenes at about 1% in the shortwave and -0.5% in the longwave when a scaling factor of 0.9 is used. There is no significant dependence of the biases on vertical resolution.

f. Sensitivity to solar zenith angle

Finally, we consider the effect of changing solar zenith angle on the relative biases. The scenes are all evaluated as before at differing values of solar zenith angle. For comparison, two scaling factor runs are also performed using 0.9 and 0.7 as the scaling factor. The results are shown in Fig. 12 (for shortwave CRF only, as the longwave fluxes have no dependence on the solar zenith angle). Across the range of solar zenith angles, for constant lower and split percentiles, the percentage biases are insignificant for Tripleclouds: the mean biases do not exceed 1% at any value of solar zenith angle and the spread never exceeds 3%. In fact, there is no real shift in biases with changing solar zenith angle in the Tripleclouds scheme. Comparison with the scaling factor method, however, shows that there is a dependence of bias on solar zenith angle. For the case where 0.7 is used as the factor, the dependence is strong, with

a bias of a much larger magnitude at smaller solar zenith angles than that at larger angles, perhaps suggesting that a different value of scaling factor is required for each solar zenith angle. Despite this, the dependence is weaker for a scaling factor of 0.9, which performs better at lower solar zenith angles, with the mean bias never exceeding 2% at higher angles. A lower value of scaling factor seems appropriate at higher solar zenith angles, while a higher value is better for lower solar zenith angles. In Tripleclouds, there is a slight variation in bias with solar zenith angle, to a comparable magnitude to that of the scaling factor method using 0.9. However, the error bars indicate that the spread of the errors using Tripleclouds is reduced by a factor of about 1.5, with a much larger reduction at low angles. It is seen that this combination of split percentile and lower percentile can produce a reliable, nonbiased output irrespective of solar zenith angle at the time.

6. Comparison of Tripleclouds and the plane-parallel methods

The different cloud representation methods are now compared. We start by performing a simple calculation to estimate the difference in computer run time between the plane-parallel (two-region) and Tripleclouds (three-region) representations. Consider a model that has n vertical levels, m of which contain clouds. The time taken for the code to evaluate one region in one level will be a constant value, irrespective of whether the region contains clear sky or cloud. Hence, the time taken for a single two-region integration is proportional to $n + m$ (n clear-sky regions, m cloudy regions). As Tripleclouds introduces an extra region into each of the m cloudy levels, the time for this run is proportional to $n + 2m$. So, the increase in run time can be written as

$$\text{run time increase} = \frac{n + 2m}{n + m}. \quad (21)$$

The version of the ECMWF model used in the study of Illingworth et al. (2007) has 60 model levels. Nineteen of these lie in the height range in which most clouds form, here taken to be the range from 1 to 10 km. Substituting $m = n/3$ into Eq. (21) predicts a run-time increase of 25%, which is approximately what is found empirically.

The 98 cloudy scenes are then evaluated using the Edwards–Slingo code at the default values described in the section 5. We perform eight runs on each scene, all of which are described below. The resultant CRF values are compared with the ICA runs when none of the methods are applied, which we again use as the truth. The results are shown in Figs. 8 and 13.

Figure 8 compares the individual CRF errors for each

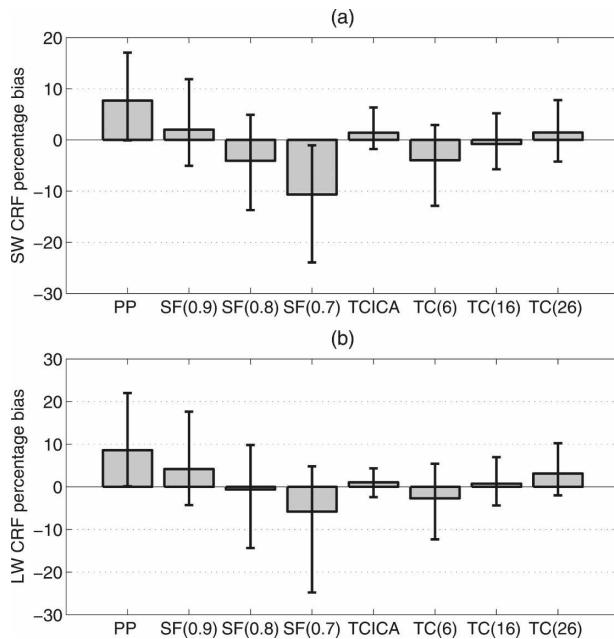


FIG. 13. Variation of the cloud radiative forcing of the scenes using different cloud approximation methods: the (a) shortwave and (b) longwave CRF percentage bias. Abbreviations for the runs are as defined in Table 1. The bars show the mean biases over the scenes for each method, and the error bars show the edges of the spread in error bounded by the 10th and 90th percentiles for each method.

cloudy scene when using the plane-parallel, the scaling factor, and the Tripleclouds schemes. As mentioned before, the plane-parallel scheme results in positive errors in nearly all the cases, with the majority of the points giving errors of under 10% in the shortwave and longwave, but with outliers giving errors of up to 30% in the shortwave and 40% in the longwave. The scaling factor with a value of 0.7 seems to have the opposite problem, with errors being mostly negative. At higher optical depths, the errors are of similar order (but of opposite sign) to those of plane-parallel, but at decreasing values of optical depth these errors become increasingly negative, reaching magnitudes as high as 30% in both the shortwave and the longwave. The reason for the high percentage biases at low optical depth is because of the shape of the relationship between either albedo or emissivity and optical depth. At low optical depths, this is approximately linear, so application of the same scaling factor is inappropriate. This is not a problem suffered by Tripleclouds, whose points have errors with much smaller magnitudes and a lower spread.

The first bar in Fig. 13 shows the mean plane-parallel bias introduced by using a conventional plane-parallel approximation. It results in a bias in the CRF of about 8% in both shortwave and longwave. The next three

bars show the effects of applying the scaling factor to the optical depth using a scaling factor of 0.7 [as proposed by Cahalan et al. (1994) and used by the ECMWF; Tiedtke (1996)] and, for comparison, values of 0.8 and 0.9. Despite its operational use, the factor of 0.7 is clearly inappropriate for the scenes that we are using, as it overcorrects the CRF biases, leaving a mean bias that is negative in both the shortwave and longwave. Similar results were found regarding the suitability of 0.7 as the scaling factor (e.g., Barker et al. 1996; Pincus et al. 1999; Rossow et al. 2002). A value of scaling factor between 0.8 and 0.9 seems more realistic.

The last four bars show the results from various Tripleclouds runs. The three bars on the right show the runs from Tripleclouds, using lower percentiles of 6, 16, and 26. An independent-column run using Tripleclouds is also performed for comparison (TCICA). This is evaluated in the same way as the control ICA run (as described in section 4), but with the Tripleclouds approximation applied to the observational data as shown in Fig. 2c. This run is performed to check that the overlap code for the single-column runs (described in section 3c) behaves correctly. There are small differences in the biases for the TCICA run with the TC(16) run in both shortwave and longwave. This therefore validates the code changes in section 3 for much more complicated cloud scenes than in Fig. 5. In fact, the mean CRF biases for Tripleclouds are seen to be small with respect to the biases for both the operational PP and SF(0.7) runs, with values of below 1% in both the shortwave and the longwave.

Overall, the Tripleclouds TC(16) method performs well. First, the amount of random error is reduced using the Tripleclouds method. Even in comparison to the best of the scaling factor methods, the reduction in spread is apparent: the spread for the SF(0.9) run is about one and a half times larger than that of the TC(16) run. Second, the optimum lower percentile in Tripleclouds is not largely dependent on solar zenith angle or resolution, while the optimum scaling factor has a relatively strong dependence on solar zenith angle. As solar zenith angle is constant in this experiment, its variation would cause an increase to the size of the scaling factor error bars. Third, the biases have a much lower sensitivity to the lower percentile in Tripleclouds than to the scaling factor over a reasonable range of values.

The heating rates are compared in Fig. 14. A profile of domain-mean heating rate is obtained for each of the cloudy scenes, and the root-mean-squared heating rate error is calculated with respect to the ICA run, using the 240-m-thick levels from the surface up to 12 km. The rms ensures that the variation in sign of the error

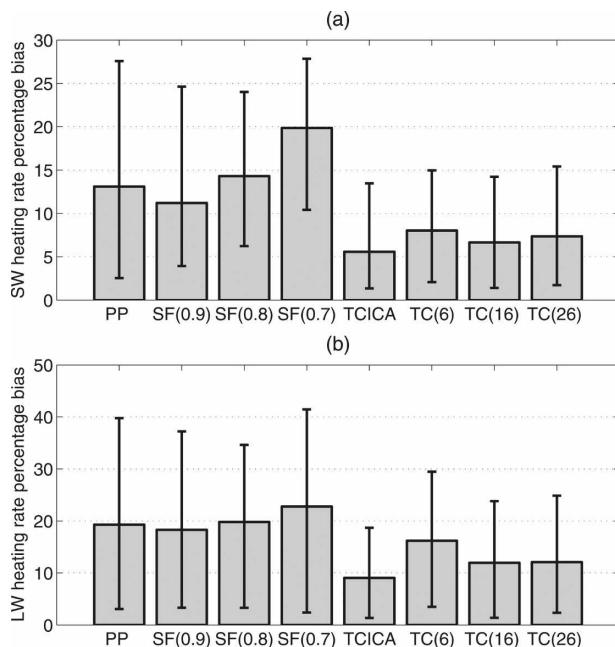


FIG. 14. Variation of mean rms heating rate biases in both (a) shortwave and (b) longwave for the 98 cloudy scenes: bars and error bars defined in Fig. 13.

does not affect the results. The rms errors are calculated as a percentage of the rms heating rate, and the overall mean rms biases are then found by averaging these percentage errors across all the scenes.

We see in Fig. 14 that the plane-parallel method gives an rms error of 13% in the shortwave and 20% in the longwave. Using the scaling factor, the representation of the heating rates deteriorates as the factor reduces, a result also found by Oreopoulos and Barker (1999). The Tripleclouds cases, however, perform much better. The independent-column Tripleclouds run performs best, while the single-column runs do not introduce much extra error in heating rate. Comparing TC(16) with SF(0.8) and SF(0.9), the two runs which produced the best results when minimizing the CRF biases, it is apparent that Tripleclouds significantly outperforms the scaling factor method. Tripleclouds, therefore, not only reduces biases in top-of-atmosphere radiative properties, it reduces biases in the atmospheric distribution of heating.

7. Discussion and conclusions

This paper presents Tripleclouds, a method of accounting for cloud inhomogeneity in the radiation scheme of GCMs. It makes use of the ability of the Edwards–Slingo code to include an extra cloudy region in the radiative transfer calculation, leading to a repre-

sentation that has a single clear-sky region and two cloudy regions. These two cloudy regions are assigned to optically thinner cloud and optically thicker cloud.

The values of water content that are applied to these two cloudy regions are determined using statistical considerations. If the distribution of water content in a single level were to be modeled as Gaussian, then using the 16th percentile of the distribution gives an accurate measure of its standard deviation. In reality, the distribution of water content is more lognormal (Hogan and Illingworth 2003), although it turned out that using the 16th percentile was an adequate method of defining the water content in the thinner cloud region. The water content in the thicker region can then be chosen to conserve mean water content in the level. The two cloudy regions are weighted equally, so the regions represent the thinner half and thicker half of the clouds in the domain level.

Tripleclouds is seen to be an improvement on the widely used plane-parallel cloud representation. Over the 12 days of radar data analyzed here, containing a range of different cloud types, the plane-parallel albedo bias was 8%. Using Tripleclouds reduced this bias to less than 1%. The scaling factor of 0.7, proposed by Cahalan et al. (1994) and applied worldwide in the ECMWF model by Tiedtke (1996), is shown here to be inappropriate for the midlatitude ice and liquid clouds analyzed. The shortwave albedo bias using this scaling factor was larger in magnitude than that of the unscaled plane-parallel representation, with a value of order 11%.

Running the radiation code with different values of scaling factor showed that the value of 0.7 is a significant overcompensation, with a value between 0.8 and 0.9 seeming more appropriate for the cloud scenes considered. However, it should be borne in mind that, throughout the scheme comparison carried out in section 6, all of the variables were held constant, and it was seen that, while not a function of vertical resolution, the scaling factor had a dependence on solar zenith angle and a weak dependence on horizontal gridbox size. So, that some value of scaling factor between 0.8 and 0.9 may seem to be most appropriate, in fact only implies that it is appropriate for the particular solar zenith angle in question. Another complication in the choice of scaling factor occurs for lower values of optical depth, where its relationship with both shortwave albedo and longwave emissivity becomes close to linear; hence the scaling factor needs to become much nearer to one.

In contrast, the Tripleclouds scheme is seen to have no such dependence on solar zenith angle or optical depth. It is also applicable to different values of vertical

resolution and horizontal gridbox size without any need for major recalibration of the percentiles. Variation of the lower percentile (and also split percentile) of Tripleclouds results in a small change in bias in the radiative properties, although the sensitivity of the bias to the Tripleclouds parameters is much smaller than that of the scaling factor.

It was also found that, for both the shortwave and longwave heating rates, applying the scaling factor to the plane-parallel approximation results in a degradation in quality of the representation; that is, the random errors increase in magnitude, and again there is a strong dependence of the errors on scaling factor. In Tripleclouds, not only are the heating rate errors insensitive to the lower percentile, the errors overall are reduced with respect to the plane-parallel heating rate biases.

Further steps are still required in the verification of Tripleclouds as a viable cloud representation scheme suitable for use in weather and climate models. So far, we have been running the radiation code on radar data, in which we know exactly where the clouds are in our scene. In reality, the weather or climate model only has knowledge of cloud fractions and gridbox-mean ice or liquid water contents, and certainly no knowledge of cloud arrangement. The overlap parameterization still needs to be considered and implemented into the Edwards–Slingo code in place of the current overlap scheme that reads in a prescribed overlap from the source data. Current parameterizations of overlap that exist include random overlap and maximum-random overlap (Geleyn and Hollingsworth 1979). However, we intend to make use of a newer form of overlap that uses a decorrelation pressure to calculate overlap, which creates a range of overlaps varying from maximum random to total random, depending on separation of adjacent layers of clouds. Investigations into ways of determining this decorrelation pressure have already been made, for example, that of Hogan and Illingworth (2003). This study also considers ways of modeling cloud water variance, which is another variable that is required for Tripleclouds to be applied to model data. Using variances typical of different regions of the world, the Tripleclouds scheme can then be used to model data output and test the applicability of the scheme at all latitudes and world weather regimes.

Comparisons of Tripleclouds with more sophisticated methods than the scaling factor method are also possible future steps in the Tripleclouds investigation. Methods such as the statistical approach of Oreopoulos and Barker (1999), the single scattering property renormalization process of Cairns et al. (2000) and the Monte Carlo independent column approximation method of Pincus et al. (2003) would provide an insightful com-

parison of the performance of Tripleclouds with respect to other state-of-the-art cloud schemes. It is also possible to extend the three-region Tripleclouds scheme to atmospheric quantities beyond cloud variables, such as aerosol optical depth. Surface properties, such as surface albedo and emissivity, can also be inhomogeneous on scales smaller than that of a GCM grid box. This is particularly true at coastlines, where the properties of the land and the sea surfaces can be very different. There is often a tendency for boundary layer clouds to form preferentially over the land or the sea. Implementation of a multiregion representation of such surface properties could provide more realistic behavior of fluxes into and out of the surface.

Acknowledgments. The first author is funded by the University of Reading under an RETF studentship.

REFERENCES

- Albrecht, B. A., D. A. Randall, and S. Nicholls, 1988: Observations of marine stratocumulus clouds during FIRE. *Bull. Amer. Meteor. Soc.*, **69**, 618–626.
- Barker, H. W., 1996: A parameterization for computing grid-averaged solar fluxes for inhomogeneous marine boundary layer clouds. Part I: Methodology and homogeneous biases. *J. Atmos. Sci.*, **53**, 2289–2303.
- , and J. A. Davies, 1992: Solar radiative fluxes for stochastic, scale-invariant broken cloud fields. *J. Atmos. Sci.*, **49**, 1115–1126.
- , B. A. Wielicki, and L. Parker, 1996: A parameterization for computing grid-averaged solar fluxes for inhomogeneous marine boundary layer clouds. Part II: Validation using satellite data. *J. Atmos. Sci.*, **53**, 2304–2316.
- Bushell, A. C., D. R. Wilson, and D. Gregory, 2003: A description of cloud production by non-uniformly distributed processes. *Quart. J. Roy. Meteor. Soc.*, **129**, 1435–1455.
- Cahalan, R. F., W. Ridgway, W. J. Wiscombe, T. L. Bell, and J. B. Snider, 1994: The albedo of fractal stratocumulus clouds. *J. Atmos. Sci.*, **51**, 2434–2455.
- Cairns, B., A. A. Lacis, and B. E. Carlson, 2000: Absorption within inhomogeneous clouds and its parameterization in general circulation models. *J. Atmos. Sci.*, **57**, 700–714.
- Carlin, B., Q. Fu, U. Lohmann, G. G. Mace, K. Sassen, and J. M. Comstock, 2002: High-cloud horizontal inhomogeneity and solar albedo bias. *J. Climate*, **15**, 2321–2339.
- Davis, A., P. Gabriel, S. Lovejoy, D. Schertzer, and G. Austin, 1990: Discrete angle radiative transfer. III: Numerical results and meteorological applications. *J. Geophys. Res.*, **95**, 11 729–11 742.
- Edwards, J. M., and A. Slingo, 1996: Studies with a flexible new radiation code. I: Choosing a configuration for a large-scale model. *Quart. J. Roy. Meteor. Soc.*, **122**, 689–719.
- Geleyn, J.-F., and A. Hollingsworth, 1979: An economical analytical method for the computation of the interaction between scattering and line absorption of radiation. *Contrib. Atmos. Phys.*, **52**, 1–16.
- Harshvardhan, and D. A. Randall, 1985: Comments on “The parameterization of radiation for numerical weather prediction and climate models.” *Mon. Wea. Rev.*, **113**, 1832–1833.

- Hogan, R. J., and A. J. Illingworth, 2003: Parameterizing ice cloud inhomogeneity and the overlap of inhomogeneities using cloud radar data. *J. Atmos. Sci.*, **60**, 756–767.
- , and S. F. Kew, 2005: A 3D stochastic cloud model for investigating the radiative properties of inhomogeneous cirrus clouds. *Quart. J. Roy. Meteor. Soc.*, **131**, 2585–2608.
- , M. P. Mittermaier, and A. J. Illingworth, 2006: The retrieval of ice water content from radar reflectivity factor and temperature and its use in evaluating a mesoscale model. *J. Appl. Meteor. Climatol.*, **45**, 301–317.
- Illingworth, A. J., and Coauthors, 2007: Cloudnet—Continuous evaluation of cloud profiles in seven operational models using ground-based observations. *Bull. Amer. Meteor. Soc.*, **88**, 883–898.
- Jakob, C., and S. A. Klein, 1999: The role of vertically varying cloud fraction in the parameterization of microphysical processes in the ECMWF model. *Quart. J. Roy. Meteor. Soc.*, **125**, 941–965.
- Larson, V. E., R. Wood, P. R. Field, J.-C. Golaz, T. H. Vonder Haar, and W. R. Cotton, 2001: Small-scale and mesoscale variability of scalars in cloudy boundary layers: One-dimensional probability density functions. *J. Atmos. Sci.*, **58**, 1978–1994.
- Liou, K.-N., 1986: Influence of cirrus clouds on weather and climate processes: A global perspective. *Mon. Wea. Rev.*, **114**, 1167–1199.
- McClatchey, R. A., R. W. Fenn, J. E. A. Selby, F. E. Volz, and J. S. Garing, 1972: Optical properties of the atmosphere. 3rd ed. Air Force Cambridge Research Laboratories Rep. AFCRL72-0497, 108 pp.
- Meador, W. E., and W. R. Weaver, 1980: Two-stream approximations to radiative transfer in planetary atmospheres: A unified description of existing methods and a new improvement. *J. Atmos. Sci.*, **37**, 630–643.
- O'Hirok, W., and C. Gautier, 2005: The impact of model resolution on differences between independent column approximation and Monte Carlo estimates of shortwave surface irradiance and atmospheric heating rate. *J. Atmos. Sci.*, **62**, 2939–2951.
- Oreopoulos, L., and H. W. Barker, 1999: Accounting for subgrid-scale cloud variability in a multi-layer 1D solar radiative transfer algorithm. *Quart. J. Roy. Meteor. Soc.*, **125**, 301–330.
- , and R. F. Cahalan, 2005: Cloud inhomogeneity from MODIS. *J. Climate*, **18**, 5110–5124.
- Pincus, R., S. A. McFarlane, and S. A. Klein, 1999: Albedo bias and the horizontal variability of clouds in subtropical marine boundary layers: Observations from ships and satellites. *J. Geophys. Res.*, **104**, 6183–6191.
- , H. W. Barker, and J.-J. Morcrette, 2003: A fast, flexible, approximate technique for computing radiative transfer in inhomogeneous cloud fields. *J. Geophys. Res.*, **108**, 4376, doi:10.1029/2002JD003322.
- Pomroy, H. R., and A. J. Illingworth, 2000: Ice cloud inhomogeneity: Quantifying bias in emissivity from radar observations. *Geophys. Res. Lett.*, **27**, 2101–2104.
- Räsänen, P., H. W. Barker, M. F. Khairoutdinov, J. Li, and D. A. Randall, 2004: Stochastic generation of subgrid-scale cloudy columns for large-scale models. *Quart. J. Roy. Meteor. Soc.*, **130**, 2047–2067.
- Randall, D. A., J. A. Coakley Jr., D. H. Lenschow, C. W. Fairall, and R. A. Kropfli, 1984: Outlook for research on subtropical marine stratiform clouds. *Bull. Amer. Meteor. Soc.*, **65**, 1290–1301.
- Ritter, B., and J.-F. Geleyn, 1992: A comprehensive radiation scheme for numerical weather prediction models with potential applications in climate simulations. *Mon. Wea. Rev.*, **120**, 303–325.
- Rossow, W. B., C. Delo, and B. Cairns, 2002: Implications of the observed mesoscale variations of clouds for the earth's radiation budget. *J. Climate*, **15**, 557–585.
- Stephens, G. L., P. M. Gabriel, and P. T. Partain, 2001: Parameterization of atmospheric radiative transfer. Part I: Validity of simple models. *J. Atmos. Sci.*, **58**, 3391–3409.
- Tiedtke, M., 1996: An extension of cloud-radiation parameterization in the ECMWF model: The representation of subgrid-scale variations of optical depth. *Mon. Wea. Rev.*, **124**, 745–750.
- Tompkins, A. M., 2002: A prognostic parameterization for the subgrid-scale variability of water vapor and clouds in large-scale models and its use to diagnose cloud cover. *J. Atmos. Sci.*, **59**, 1917–1942.
- Zdunkowski, W. G., W.-G. Panhans, R. M. Welch, and G. J. Korb, 1982: A radiation scheme for circulation and climate models. *Beitr. Phys. Atmos.*, **55**, 215–238.