



Monte Carlo Spectral Integration: a Consistent Approximation for Radiative Transfer in Large Eddy Simulations

Robert Pincus¹ and Bjorn Stevens²

¹ University of Colorado and NOAA Earth System Research Laboratory, Boulder, Colorado, USA

² Department of Atmospheric and Oceanic Sciences, University of California Los Angeles, California, USA and Max Planck Institute for Meteorology, Hamburg, Germany

Manuscript submitted 28 June 2008; in final form 3 November 2008

Large-eddy simulation (LES) refers to a class of calculations in which the large energy-rich eddies are simulated directly and are insensitive to errors in the modeling of sub-grid scale processes. Flows represented by LES are often driven by radiative heating and therefore require the calculation of radiative transfer along with the fluid-dynamical simulation. Current methods for detailed radiation calculations, even those using simple one-dimensional radiative transfer, are far too expensive for routine use, while popular shortcuts are either of limited applicability or run the risk of introducing errors on time and space scales that might affect the overall simulation.

A new approximate method is described that relies on Monte Carlo sampling of the spectral integration in the heating rate calculation and is applicable to any problem. The error introduced when using this method is substantial for individual samples (single columns at single times) but is uncorrelated in time and space and so does not bias the statistics of scales that are well resolved by the LES. The method is evaluated through simulation of two test problems; these behave as expected. A scaling analysis shows that the errors introduced by the method diminish as flow features become well resolved. Errors introduced by the approximation increase with decreasing spatial scale but the spurious energy introduced by the approximation is less than the energy expected in the unperturbed flow, i.e. the energy associated with the spectral cascade from the large scale, even on the grid scale.

DOI:10.3894/JAMES.2009.1.1

1. How much approximation is acceptable in radiative transfer treatments for large eddy simulation?

The term “large-eddy simulation” (LES) is sometimes used to describe a class of numerical models with high spatial resolution, but it more properly refers to a class of calculations: those in which the large energy-rich eddies are directly simulated, as opposed to being modeled (meaning represented in the abstract), and which are not biased by errors introduced through the modeling of the small scales. This modeling is required for flows whose Reynolds numbers are too large to allow for an explicit, or direct, representation of all the scales of motion. One expects LES to converge to such direct representations (*i.e.*, direct numerical simulation) in the limit of sufficiently small grid-spacing. In practice, imperfect parameterizations of physical processes such as radiative transfer, cloud microphysical processes, boundary processes, *etc.* make this limit difficult to realize, even in principle. Current computational capabilities limit LES to domain sizes of order 10 km and grid spacings of order 10 m.

Radiative transfer is the best understood of the subgrid scale processes in that a well-established theory allows radiative fluxes to be computed to arbitrary accuracy given a well-characterized medium and unbounded computational resources. Radiative transfer can be quite important; many flows of interest (*e.g.*, stratocumulus) are driven directly by local heating resulting from the divergence of radiative fluxes. The computational costs of the most detailed calculations are far too high for use in LES, however. Given the computational burden of exact treatments of radiative transfer, and the interest in how such fluxes couple and co-evolve with the flow, there is a long tradition of defining approximate methods for estimating radiative heating rates.

To whom correspondence should be addressed.

Robert Pincus, Cooperative Institute for Research in Environmental Sciences, University of Colorado/NOAA Earth Systems Research Laboratory Physical Sciences Division, 325 Broadway, R/PSD1, Boulder, CO 80305, USA

Robert.Pincus@colorado.edu

Nearly all existing methods start with a gross simplification, namely the assumption of one-dimensional radiative transfer, wherein radiation is assumed to travel within, but not among, model columns. This assumption is known to be poor if the grid spacing is less than the radiative smoothing scale (Marshak et al., 1995), as is usually the case in LES. Neglecting three-dimensional radiative transfer certainly leads to errors in local flux and intensity (radiance) calculations. But errors in these quantities are irrelevant for LES, which responds most directly to heating rates on resolved space and time scales. Heating rate errors caused by the 1D assumption are small in stratocumulus clouds in the domain mean (see, *e.g.*, Zuidema and Evans, 1998; DiGiuseppe and Tompkins, 2003) though they may be larger in cumulus. In a general sense, one-dimensional radiative transfer can be thought of as an approximation producing heating rate errors that are small in stratiform clouds, somewhat larger in cumuliform clouds, and short-lived in both cases.

But it turns out that even detailed one-dimensional radiative transfer calculations are so prohibitively expensive that they are rarely employed in LES. As one example, when we coupled one reasonably efficient one-dimensional radiative transfer code to a hydrodynamic code and used it to compute heating rates in every column at every time step, our solution time increased by a factor of about fifty. One strategy for avoiding this high computational cost is to simply specify the radiative heating rates, perhaps in combination with other large-scale forcing (*e.g.*, Siebesma et al., 2003). This approach is most useful when radiation drives the flow indirectly, for instance by helping sustain surface fluxes, but makes it impossible to study the coupling of the flow with the radiation. Some amount of interactivity can be gained for individual case studies by developing a parametric fit to detailed radiative calculations (*e.g.*, Stevens et al., 2005). Fits can be quite fast but necessarily introduce some amount of bias. More importantly, they are most successful if their range of applicability can be limited *a priori*. This means they must be revisited when new situations are to be simulated, which may occur on the timescale of the evolution of the flow itself. As LES is applied to a wider range of problems and to increasingly large-scale and heterogeneous environments (see, *e.g.*, Khairoutdinov and Randall, 2006) such parameterizations of radiative transfer become increasingly unattractive.

An alternative is to do a full (one-dimensional) radiation calculation at reduced temporal and/or spatial resolution relative to the rest of the calculation (*e.g.*, Ackerman et al., 2004). Although spatial and temporal sampling may be implemented independently, they are fundamentally equivalent: in a flow with grid spacing δx and time step δt related by a characteristic velocity $u_* \propto \delta x / \delta t$, computing radiation every N time steps amounts to averaging over a spatial scale proportional to $N\delta x$, with the size of the averaging errors correlated with the local velocity. If N is large enough,

heating rate errors can be introduced at scales approaching those being explicitly simulated. Because the heating rates are sampled in a manner that depends on the flow itself such approaches can introduce unwelcome biases (*e.g.*, Xu and Randall, 1995; Pauluis and Emanuel, 2004).

Here we introduce an approximate method that is fast enough to be practical, unbiased, and has numerical convergence properties consistent with the philosophy of large-eddy simulation. The method is demonstrated using a benchmark of one-dimensional radiative transfer but should apply more generally. An example implementation is shown to introduce noise at levels and on scales that does not affect the simulation of large eddies. We then consider the general question of how errors introduced by approximate radiative treatments can be expected to influence LES, concentrating on the class of methods that are unbiased but introduce uncorrelated random noise.

2. Monte Carlo spectral integration

Full radiative transfer calculations are computationally expensive because they integrate over a large portion of the electromagnetic spectrum. The main challenge in performing spectral integration is that the absorption cross-section of gases can change by orders of magnitude over very small spectral intervals. The current state-of-the-art solution is to use a “correlated k -distribution” (Lacis and Oinas, 1991; Fu, 1992) to calculate radiative fluxes, whose spatial divergence then provides the local heating rate. To build a single-layer k -distribution one chooses a set of B relatively broad spectral bands within which Rayleigh scattering by molecules and the optical properties of aerosols and clouds can be considered uniform. Within each band similar values of absorption coefficient k are grouped into G “ g -points” within which $k \approx k(g)$. The broadband flux $F(x, y, t)$ within the column centered at the point (x, y) and at time t can then be computed as

$$F(x, y, t) = \sum_b^B w_b \sum_g^{G(b)} w_{g(b)} F_{b,g}(x, y, t) \quad (2.1)$$

where the g -point weights w_g are the fraction of each band accounted for by each g -point, so that $\sum w_g = 1$ for all b . In practice, (2.1) is usually applied separately to the visible and infrared portions of the spectrum. For infrared bands $w_b = 1$, because all radiation is emitted within the domain, while for solar bands w_b denotes the amount of solar energy within band b . In most implementations B and G are both $O(10)$ so a broadband calculation requires as many as several hundred pseudo-monochromatic calculations. A “correlated k -distribution” extends this single-layer idea by further assuming that the relative strengths of absorption lines within a band are correlated with height in the atmosphere so that the mapping between k and g derived for one level of the atmosphere can be applied to all levels.

This mapping the technique to be used for multiple levels in a vertically inhomogeneous atmosphere.

We propose using the flux computed for a single randomly-chosen band and g -point as a proxy for the full calculation, choosing a different band and g -point for each column and at each dynamical time step. That is, we approximate (2.1) using randomly-chosen values of b' and g' as

$$F(x, y, t) \approx F_{MC}(x, y, t) = B w_{b'} F_{b', g'}(x, y, t) \quad (2.2)$$

where the probability of choosing a given value of b' and g' is given by the weight with which the spectral interval contributes to the overall sum, *i.e.*,

$$p(b') = 1/B \quad \text{and} \quad p(g') = w_{g'(b')}. \quad (2.3)$$

Formally, F_{MC} is a single-sample Monte Carlo estimate of the complete spectral integration F ; repeated application of (2.2) might be called ‘‘Monte Carlo spectral integration.’’ Individual estimates using (2.2) will contain substantial random error relative to (2.1) but this noise decreases as $1/\sqrt{n}$ where n is the number of samples. The two estimates converge as the number of samples increases whether those estimates are accumulated in space (*i.e.*, across an LES domain), in time (*i.e.*, over the course of a simulation), or both. In addition, given the frequency with which radiation is typically calculated (every 30–50 time steps) and the total number of g -points in correlated k -distribution schemes (of order 100–300), (2.2) represents even less computational work than infrequent radiation calculations using (2.1).

3. Example large eddy simulations using Monte Carlo spectral integration

3.1. An example implementation

We have implemented Monte Carlo spectral integration in a radiative transfer code and coupled this radiation solver to a large eddy simulation code. Our treatment of radiative transfer follows Fu (1992) and Fu et al. (1997). The spectrum is divided into 6 bands in the solar (50000 to 2500 cm^{-1}) and 12 in the infrared (2500 to 0 cm^{-1}). The code includes parameterizations for the optical properties of cloud water and ice and a model for absorption by the water vapor continuum within each band. A k -distribution accounts for the effects of water vapor, ozone, methane, and nitrous oxide using as many as 12 g -points per each band. Radiative transfer is computed with a δ -scaled four-stream solver. In our implementation (2.2) is applied separately to the solar and thermal infrared portions of the spectrum. This treatment increases the total computational cost by about 50% relative to simulations that use a parametric fit designed to capture the sensitivity of radiative heating rates to cloud water.

The code used to perform the large-eddy simulation is the most recent version of the UCLA LES (Stevens et al., 2005). For these simulations the model is configured with 192 grid-points in each horizontal direction and 131 points in the vertical. The horizontal grid spacing δx is 35 m; the vertical grid varies from 10 m near the surface to 5 m in a zone that spans the inversion. The model time step varies so as to maintain the peak CFL (Courant-Friedrichs-Lewy) number within a fixed range and is typically 1–2 s. Sub-grid scalar fluxes are carried by the numerics using a method (*i.e.*, UCLA-1) shown to yield realistic results in Stevens et al. (2005).

3.2. Simulations of stratocumulus

We use the LES and radiation package to simulate nocturnal stratocumulus based on the first research flight (RF01) of DYCOMS-II (Stevens et al., 2003). This is a challenging case because the large-scale flow is sensitive to biases on the small scales (Stevens et al., 2005).

Initial profiles, surface specifications, and large-scale forcing follow Stevens et al. (2005), with two exceptions. First, because the profiles of subsidence in Stevens et al. (2005) were tailored to the parametric representation of radiative transfer, subsidence is simply left out of the current simulations. This allows cloud top to deepen by about an additional 10 m hr^{-1} relative to simulations with subsidence but does not affect the simulation’s sensitivity to errors in radiative heating rates. Secondly, during the radiative transfer calculation a summer subtropical sounding is grafted onto the state of the lower atmosphere as represented by the LES.

We simulate four hours of evolution starting from quiescent initial conditions with a horizontally uniform cloud whose water content increases adiabatically from cloud base at about 600 m to cloud top near 840 m. We show results from nocturnal simulations but we have performed simulations centered on local noon and find no important differences.

The noise introduced by Monte Carlo spectral integration has one unintended but desirable result: model fields spin up from a cold start somewhat more quickly and gently in the presence of noise in the radiative heating rates. It is standard practice to initialize LES with uncorrelated random noise in the temperature and/or humidity fields in order to break the symmetry of the initial conditions and allow the flow to develop. Using (2.2) to compute radiative heating rates provides a continuous source of small scale variability. The flow spins up somewhat more rapidly as a result, leading to the accumulation of less instability and a less violent initial overturning of the flow. Similar but more pronounced effects are associated with the addition of a stochastic perturbation to the sub-grid flux (*e.g.*, Mason and Thomson, 1992; Weinbrecht and Mason, 2008).

Once the flow is spun up, however, noise in the radiation field does not affect the simulation noticeably. The dashed




Figure 1. Evolution of a nocturnal cloud-topped boundary layer computed using fully-resolved (black lines) and Monte Carlo-sampled (dashed blue lines) one-dimensional radiation calculations. Bulk properties of the flow, including domain-averaged turbulence kinetic energy (top panel), liquid water path (middle), and cloud top height (bottom) are essentially indistinguishable in the two simulations, despite the large amount of uncorrelated noise introduced in radiation calculations by the Monte Carlo sampling.

blue lines in Fig. 1 show results from the last 110 min of a 4 h simulation made using (2.2), while the black lines show a simulation that uses (2.1) to compute the heating rates (at each gridpoint and timestep) starting at $t = 2.5$ h. The evolution of the boundary layer in the two simulations is, statistically speaking, indistinguishable.

Profiles accumulated over the last half-hour of the simulation (Fig. 2) show a similar insensitivity. Spatial and temporal variability in cloud structure causes some variability in radiative fluxes, but this is dwarfed by the noise introduced by Monte Carlo sampling (see the bottom two panels of Fig. 2). Nonetheless, profiles of properties dynamic and thermodynamic properties (shown in the top three rows of Fig. 2) are essentially identical. The single exception is the third moment of vertical velocity, which increases in the presence of Monte Carlo noise. This increase, however, is of the order of sampling errors (as measured, for example, by the difference between independent realizations of identically forced flows), so we conclude that no quantity is affected significantly by the noise introduced by (2.2).

3.3. Simulations with small net forcing and large noise

Results in section 4, particularly (4.10), suggest that random noise might affect the flow simulated by LES when the mean

forcing is small relative to the amount of noise. To assess this possibility we make two simulations of a weakly-driven flow. The problem specification follows the DYCOMS-II RF01 case, as above, but without the mean radiative driving. Figures 3 and 4 mirror Figs. 1 and 2 in showing time series and profiles from the last half-hour of the simulations, respectively. Here, however, the black lines show simulations in which no radiation is calculated at all, so that the flow is driven entirely by the weak surface fluxes and evaporative cooling at cloud top, while the dashed blue lines show a calculation in which radiative driving is first computed in each column using (2.2), then modified by removing the domain-mean radiative flux, leaving only the column-by-column noise. (We omit profiles of radiative flux and its variability since these are identically zero with one exception: the standard deviation of radiative flux for the Monte Carlo spectral integration, which follows the bottom right panel of Fig. 2.) In this case the noise has a barely noticeable effect on the spin-up, most likely because the evolution is driven by the (unperturbed) surface forcing.

In the absence of strong driving the circulations are weak (compare, for example, the variance of vertical velocity in Fig. 4 with the values in Fig. 2) but, despite the fact that perturbations in radiative heating rates are very large relative to the small mean forcing, the noisy simulation does not substantially depart from its noise-less counterpart. Hence, even in a situation constructed to be as sensitive as possible, the deleterious effects of the grid-scale noise that arises from sampling of the radiative fluxes are minor.

4. Error, scale, and evolution in large eddy simulations

The results in Section 3 demonstrate that the radiative transfer approximation described in Section 2 allows for realistic large eddy simulations in two very different regimes. In this section we seek to understand more deeply *why* the approximation works, in part so that we may understand any limits of applicability.

What degree of error is acceptable in an approximate treatment of radiative transfer? The answer depends on the application; for an LES we might insist that the approximation does not change the simulation of the large eddies relative to some expensive benchmark calculation. In this section we first estimate the amount of error that might be tolerated by an LES by comparing the amount of energy in eddies expected at a given spatial scale with the perturbation introduced by some error arising from an approximate treatment of radiative transfer. We then apply this analysis to the particular case of heating rate errors which, like those introduced by Monte Carlo spectral integration, are uncorrelated in space and time. As in Section 3 we use the simulation of nocturnal stratocumulus as a prototype problem. The simple nature of the flow (*i.e.*, its stationarity, homogeneity, and the dominance of a single underlying

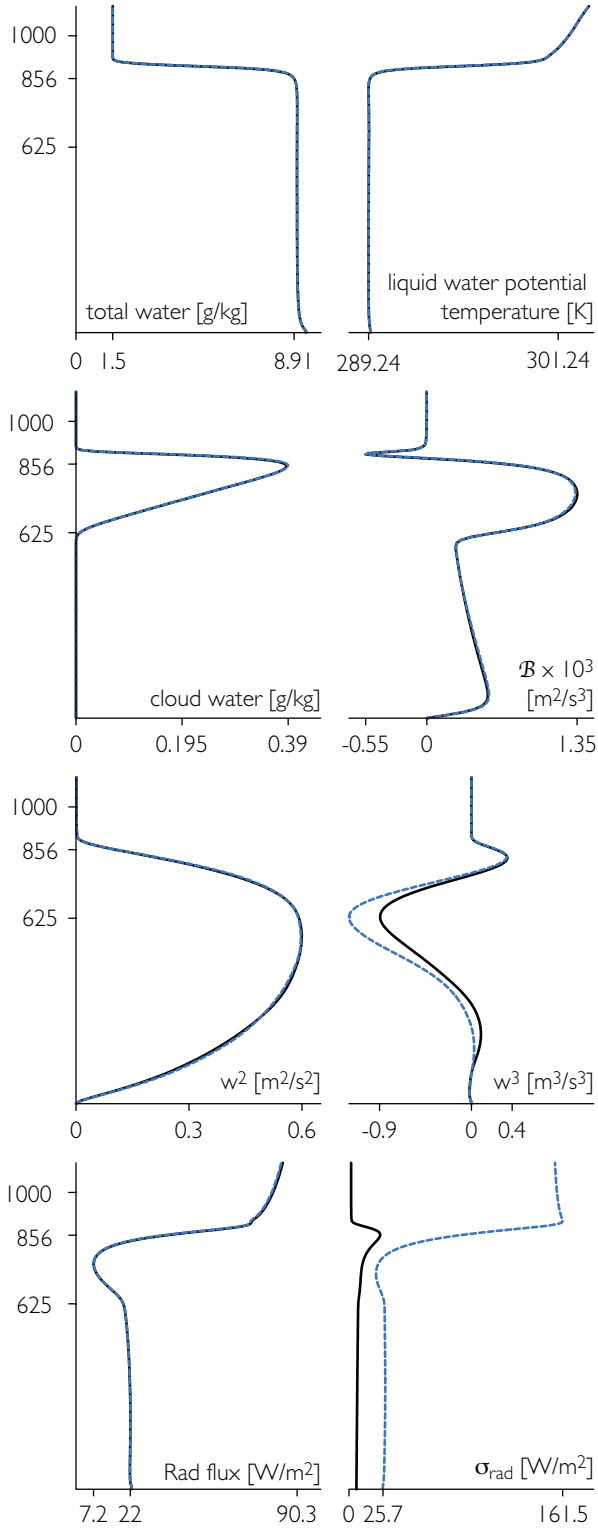


Figure 2. Profiles of dynamic, thermodynamic, and radiative properties averaged over the last half hour of a four-hour simulation of a nocturnal cloud-topped boundary layer using fully-resolved (black lines) and Monte Carlo-sampled (dashed blue lines) radiation calculations. All the dynamic and thermodynamic profiles are essentially the same, indicating that

using sparse radiation calculations does not affect the simulation. The one exception is the third moment of vertical velocity which is subject to large sampling errors. The mean radiative flux profiles are likewise identical, as they must be subject to sampling error and systematic differences in cloud structure, and the variability caused by Monte Carlo sampling is much larger than the variability introduced by spatial variability in cloud structure.

ing length-scale) facilitate scaling estimates of the effects of noise.

4.1. How sensitive is the flow to approximation error?

Consider a boundary layer of depth h driven by radiative cooling of net magnitude ΔF at cloud top. Cooling occurs over a length scale $\lambda \ll h$. The cooling rate may be expressed as a buoyancy flux \mathcal{B}_h

$$\mathcal{B}_h = \frac{g}{\rho c_p \Theta_0} \Delta F \quad (4.1)$$

where ρ is the air density, Θ_0^{-1} the coefficient of thermal expansion, g gravity and c_p the isobaric specific heat capacity.

If there are no other sources of turbulence kinetic energy one expects this cooling to drive convective circulations with large eddies of spatial scale h (hence the subscript on the left hand side of (4.1)). By balancing the production of turbulence kinetic energy with its dissipation on the grid scale one can estimate that the vigor, or specific turbulence energy, of boundary layer scale eddies following Deardorff (1970), *i.e.*,

$$\bar{e}_h \propto (\mathcal{B}_h h)^{2/3}. \quad (4.2)$$

Similarly, energy in eddies of arbitrary, but bounded, size $\ell \ll h$ (yet still much larger than the Kolmogorov scale) should conform to an inertial cascade of turbulence, such that

$$\bar{e}_\ell \propto \bar{e}_h \left(\frac{\ell}{h} \right)^{2/3}. \quad (4.3)$$

The eddy velocity scale w_* is related to the turbulence kinetic energy density as $w_{*,\ell} \propto \sqrt{\bar{e}_\ell}$ so that eddies at scale ℓ have a turn-over time of

$$\tau_\ell \propto \ell w_{*,\ell}^{-1} \propto \ell \bar{e}_\ell^{-1/2}. \quad (4.4)$$

Next consider an approximate estimate for ΔF which, compared to some accurate but computationally expensive benchmark estimate, introduces a gradient of magnitude $\Delta F'_\ell$ at scale ℓ . This is equivalent to introducing a spurious perturbation forcing \mathcal{B}'_ℓ so that the net forcing becomes

$$\mathcal{B}_\ell = \bar{\mathcal{B}}_\ell + \mathcal{B}'_\ell. \quad (4.5)$$

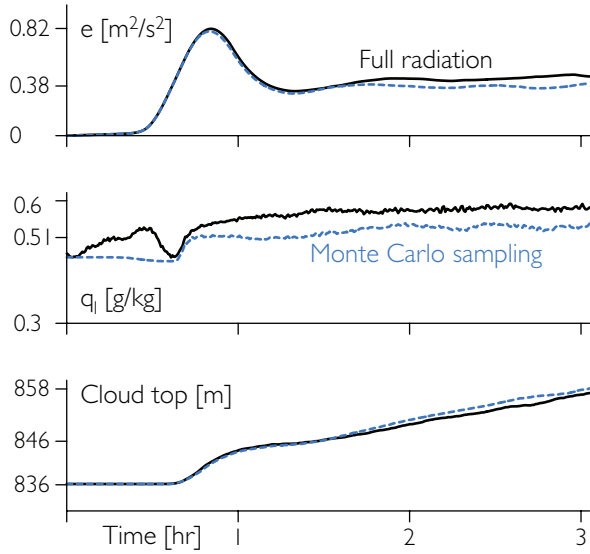


Figure 3. Evolution of a cloud-topped boundary layer computed in the absence of any radiative driving (black lines) and driven only by uncorrelated random noise in radiative heating rates with zero mean (dashed blue lines). Surface fluxes and evaporative cooling drive very gentle circulations that should, in principle, be more susceptible to noise in radiative heating rates. The noise, shown in blue dashed lines, is much larger than the mean forcing but nonetheless does not degrade the simulation.

This spurious forcing will introduce an energy perturbation at scale ℓ . If the perturbations are to organize eddies on this scale and systematically change the flow being simulated, two conditions must be met:

1. the perturbations must persist for time τ_ℓ , and
2. they must be large enough so that perturbations to the energy density at scale ℓ are commensurate to the nominal value, *i.e.*, e'_ℓ/\bar{e}_ℓ is $O(1)$ or larger.

4.2. The impact of random noise

The magnitude and scaling behavior of B'_ℓ depend, in general, on the approximation used to compute radiative transfer. In the case of Monte Carlo estimates of ΔF (or any estimate which is unbiased but contains uncorrelated random noise of magnitude σ_F), the scaling behavior is known: B'_ℓ decreases as $1/\sqrt{n}$, where n is the number of independent samples. This number is determined by the scale ℓ , the grid spacing δx (which we assume to be isotropic in the horizontal) and time step δt , as well as the amount of noise in individual estimates of ΔF . The numbers of samples on scale ℓ can be determined as

$$n_\ell = n_{\ell,xy} n_{\ell,t} \approx \left(\frac{\ell}{\delta x}\right)^2 \frac{\tau_\ell}{\delta t} \approx \left(\frac{\ell}{\delta x}\right)^3 \left(\frac{h}{\ell}\right)^{1/3} \quad (4.6)$$

where we have employed (4.4) and assumed that the time step is determined by the CFL condition $\delta t \approx \delta x / \sqrt{e_h}$. Note that δt scales with the vigor of the large eddies.

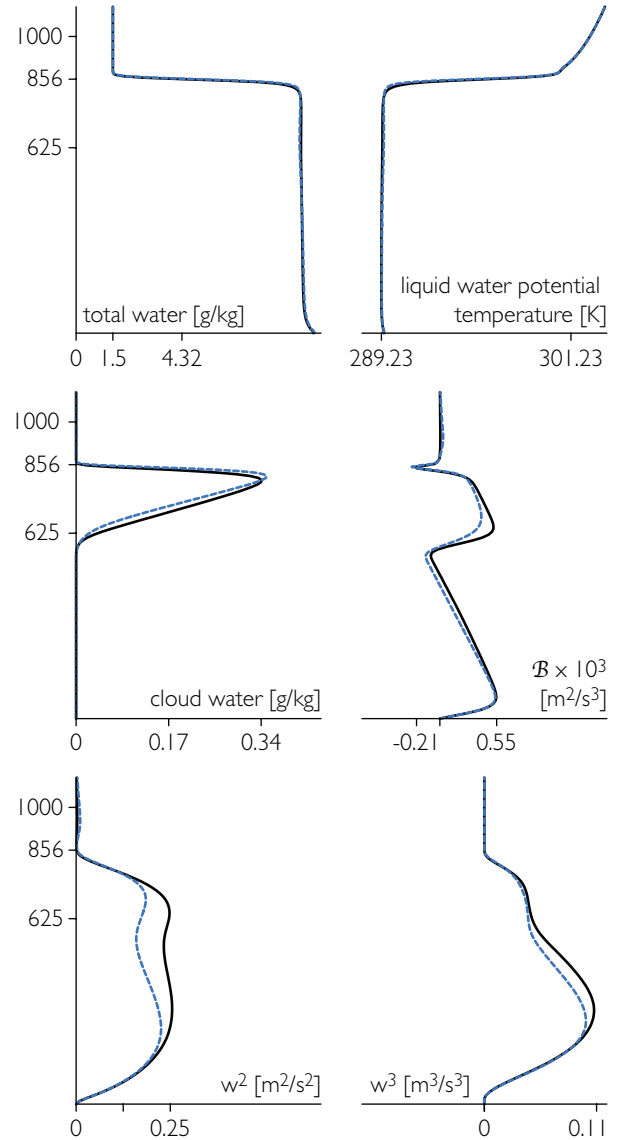


Figure 4. Profiles of dynamic and thermodynamic properties averaged over the second hour of simulations of a cloud-topped boundary layer with no net radiative forcing. Black lines show a simulation forced only by weak surface fluxes, while dashed blue lines show a simulation in which only the noise from Monte Carlo spectral integration is applied in each grid column at each time step. Turbulence in the noisy simulation is reduced because the layer deepens slightly more quickly; this stabilizes the flow more than the radiative perturbations destabilizes it. Even in this weakly-forced flow, however, large amounts of noise applied at small time and space scales do not fundamentally disrupt the simulation.

If we further assume that the magnitude of the random error in each estimate of ΔF and hence B is proportional to the benchmark (expected) value then the perturbation to the buoyancy flux can be expressed as

$$B'_\ell = \overline{\sigma_B} / \sqrt{n_\ell} \quad (4.7)$$

where σ_B^2 measures the sample variance of \mathcal{B}_h . Although our analysis is not restricted to this limit, we envision scenarios where the ratio $\alpha \equiv \sigma_B / \overline{\mathcal{B}_h}$ is of order unity.

In the limit of $\sqrt{n_\ell} \gg 1$ we say that the radiation field is “well-sampled”. But for an inertial range within which $\mathcal{B}_\ell = \mathcal{B}_h$, this limit also implies that

$$\mathcal{B}'_\ell \ll \overline{\mathcal{B}_\ell} \quad (4.8)$$

which is to say that systematic perturbations to the buoyancy flux are also small. In this limit the specific energy attributable to the spurious forcing is

$$e'_\ell = (\mathcal{B}'_\ell \ell)^{2/3} = \left(\frac{\sigma_B}{\sqrt{n_\ell}} \ell \right)^{2/3} = (\sigma_B \ell)^{2/3} \frac{\delta x}{\ell} \left(\frac{\ell}{h} \right)^{1/9} \quad (4.9)$$

so the standard error in the specific energy of an eddy of scale ℓ is

$$\frac{e'_\ell}{\bar{e}_\ell} \propto \alpha^{2/3} \frac{\delta x}{\ell} \left(\frac{\ell}{h} \right)^{1/9}. \quad (4.10)$$

Thus, even when unbiased noise in individual radiative flux estimates is as large as the net cooling rate, the turbulence kinetic energy introduced at resolved spatial scales (*i.e.*, those for which $\ell > N\delta x$, where N is roughly in the range 4–10) is a small fraction of the energy expected due to the downscale cascade of energy from the largest eddies.

Let us examine this result at two limiting values of ℓ : the scale of the largest eddies h and the grid scale δx . For the former $\ell = h$ and (4.10) reduces to

$$\frac{e'_h}{\bar{e}_h} \propto \alpha^{2/3} \frac{\delta x}{h}. \quad (4.11)$$

In LES $h \gg \delta x$ by definition, so that the velocity error at scale h is much less than one even if α is of order unity. This implies that sampling noise will have a negligible effect on the simulation of large eddies.

For the latter, $\ell = \delta x$, we expect

$$\frac{e'_{\delta x}}{\bar{e}_{\delta x}} \propto \left[1 + \alpha \left(\frac{\delta x}{h} \right)^{1/6} \right]^{2/3} - 1, \quad (4.12)$$

where, because we expect a larger bias, we have not made the small perturbation assumption in our scaling. Indeed, the amount of noise introduced into the smallest scales is not necessarily negligible. On the other hand, spurious energy introduced on these scales is efficiently dissipated. Based on the behavior at these two limits, we suggest that even large amounts of random noise in radiative heating rate calculations are unlikely to systematically affect large eddy simulations.

Approximations introducing uncorrelated random noise have errors that diminish with increasing grid resolution: regardless of scale, the anomalous energy introduced into the

flow is proportional to δx to some power, and approaches 0 as $\delta x \rightarrow 0$ (albeit slowly at the smallest scales). In the language of numerical analysis this property means that the approximation embodied in (2.2) is a consistent one.

In two-dimensional simulations the sampling introduces somewhat more noise, both because the flow is less efficient at transporting large-scale energy to small scales, and because an eddy of a given physical size encompasses fewer grid elements, and hence fewer samples of the radiative flux. Even so, experiments using (2.2) in two dimensions (not shown) evince little effect of the sampling noise on the large-scale statistics.

Finally we note that these estimates may also be applied to flows whose turbulence is also forced by non-radiative processes such as surface heat fluxes by defining α as the ratio of the random radiative driving relative to the mean driving by the sum of well sampled processes.

5. Matching tools to tasks

Monte Carlo spectral integration (McSI), as embodied by the approximation in (2.2), is particularly attractive for use in hydrodynamic calculations because the sampling of radiative heating rates is done in a space orthogonal to the phase space of the flow. This ensures that errors it introduces will not systematically correlate with the flow, and will therefore remain unbiased. While errors may arise due to the spurious behaviors of the small scale, our analysis suggests that such errors are unlikely to be any larger than those already associated with biases introduced by the sub-grid scale model, local truncation error, or the uncertain representation of physical processes.

When implemented in LES, (2.2) has the further advantage of being a consistent approximation, in that errors scale with the size of the grid and so can be made arbitrarily small as the mesh is arbitrarily refined. This means that using Monte Carlo spectral integration as the radiative transfer method for LES is both practical, because it is computationally efficient, and correct, because it is consistent with the principles underlying large-eddy simulation. We expect the radiative transfer approximation to be useful for coarser resolution cloud-resolving models even if the formal justification for such calculations is less refined.

McSI is similar in spirit to the Monte Carlo Independent Column Approximation (McICA; Pincus et al., 2003) used in global modeling. There are two key differences. First, McICA is used to account for subgrid-scale spatial structure within each model column, including horizontal variability and cloud overlap, rather than operating on homogeneous columns as we do here. Secondly, McICA introduces substantially less noise than does (2.2). Experiments with a range of global models (*e.g.*, Räisänen et al., 2005; Pincus et al., 2006; Barker et al., 2008) have shown that model evolution is unaffected by uncorrelated high-frequency, small-scale perturbations. To date the justification for McICA has simply been that the models don't respond to

sampling noise; we suspect that an extension of the reasoning in section 4 helps explain why this is so. But the success of both McICA and the present method illustrate that hydrodynamic simulations are far more sensitive to systematic errors in radiative calculations than to random noise.

McSI can also be thought of as a special case of a full three-dimensional broadband Monte Carlo radiative transfer calculation with two very large simplifications: 1) deterministic spatio-temporal sampling (i.e. one calculation per grid column), as opposed to randomly-selected samples, and 2) the use of the one-dimensional, as opposed to three-dimensional, radiative transfer equation. One can imagine various levels of approximation between one-dimensional and fully three-dimensional radiative transfer, and Monte Carlo spectral integration is independent of how the pseudo-monochromatic radiative transfer equations are solved. In particular, (2.2) could be applied to solution methods that adjust fluxes computed in one dimension to account for three-dimensional radiative transfer effects including first-order shadowing (Várnai and Davies, 1999), higher-order radiative smoothing (Marshak et al., 1998), or both (Wapler and Mayer, 2008).

Based on the reasoning in Section 4, however, we argue that three-dimensional radiative transfer is warranted only if local heating rate anomalies relative to one-dimensional radiative transfer are large enough and persist for long enough to affect the flow. It would be straightforward to assess the magnitude and time scale of these differences using a high-frequency sequence of snapshots from an LES. This seems an important check to perform before coupling a three-dimensional solver to an LES, especially given results in two dimensions showing an insensitivity of flow characteristics to multidimensional radiative transfer (Mechem et al., 2008). One-dimensional radiative transfer is known to be a poor approximation at LES scales, but it may turn out to be a perfectly useful approximation for large-eddy simulations.

Acknowledgments: This work has been supported by the National Science Foundation Science and Technology Center for Multi-Scale Modeling of Atmospheric Processes, managed by Colorado State University under cooperative agreement No. ATM-0425247. BS also acknowledges support from the DFG Metsröm Program. We thank Drs. Yu Gu and Kuo-Nan Liou for making their radiation code available to us; our work is based on a new solver inspired by and using elements of their original code. The comments of an anonymous reviewer helped us clarify our arguments. RP thanks Jonathan Gold, Laurie Ochoa, and BS and his family for hospitality during a visit to UCLA to work on this project.

References

- Ackerman, A. S., M. P. Kirkpatrick, D. E. Stevens, and O. B. Toon, 2004: The impact of humidity above stratiform clouds on indirect aerosol climate forcing. *Nature*, **432**, 1014–1017, doi:10.1038/nature03174.
- Barker, H. W., J. N. S. Cole, J.-J. Morcrette, R. Pincus, P. Räisänen, K. von Salzen, and P. Vaillancourt, 2008: The Monte Carlo Independent Column Approximation: An assessment using several global atmospheric models. *Quart. J. Roy. Meteor. Soc.*, **134**, 1463 – 1478, doi:10.1002/qj.303.
- Deardorff, J. W., 1970: Convective velocity and temperature scales for the unstable planetary boundary layer and for Rayleigh convection. *J. Atmos. Sci.*, **27**, 1211–1213, doi:10.1175/1520-0469(1970)027<1211:CVATSF>2.0.CO;2.
- DiGiuseppe, F. and A. M. Tompkins, 2003: Effect of spatial organisation on solar radiative transfer in three dimensional idealised stratocumulus cloud fields. *J. Atmos. Sci.*, **60**, 1774–1794, doi:10.1175/1520-0469(2003)060<1774:EOSOOS>2.0.CO;2.
- Fu, Q., 1992: On the correlated k -distribution method for radiative transfer in nonhomogeneous atmospheres. *J. Atmos. Sci.*, **49**, 2139–2156 doi:10.1175/1520-0469(1992)049<2139:OTCDMF>2.0.CO;2.
- Fu, Q., K.-N. Liou, M. Cribb, T. P. Charlock, and A. Grossman, 1997: Multiple scattering parameterization in thermal infrared radiative transfer. *J. Atmos. Sci.*, **54**, 2799–2812, doi:10.1175/1520-0469(1997)054<2799:MS PITI>2.0.CO;2.
- Khairoutdinov, M. and D. A. Randall, 2006: High-resolution simulation of shallow-to-deep convection transition over land. *J. Atmos. Sci.*, **63**, 3421–3436, doi:10.1175/JAS3810.1.
- Lacis, A. A. and V. Oinas, 1991: A description of the correlated k -distribution method for modeling non-grey gaseous absorption, thermal emission, and multiple scattering in vertically inhomogeneous atmospheres. *J. Geophys. Res.*, **96**, 9027–9063, doi:10.1029/90JD01945.
- Marshak, A., A. B. Davis, R. F. Cahalan, and W. J. Wiscombe, 1998: Nonlocal independent pixel approximation, direct and inverse problems. *IEEE Trans. Geosci. Remote Sens.*, **36**, 192–205, doi:10.1109/36.655329.
- Marshak, A., A. B. Davis, W. J. Wiscombe, and R. F. Cahalan, 1995: Radiative smoothing in fractal clouds. *J. Geophys. Res.*, **100**, 26247–26261, doi:10.1029/95JD02895.
- Mason, P. and D. J. Thomson, 1992: Stochastic backscatter in large-eddy simulations of boundary layers. *J. Fluid Mech.*, **242**, 51–78, doi:10.1017/S0022112092002271.
- Mechem, D. B., Y. L. Kogan, M. Ovtchinnikov, A. B. Davis, K. F. Evans, R. G. Ellington, and E. E. Takara, 2008: Multi-dimensional longwave forcing of boundary layer cloud systems. *J. Atmos. Sci.*, **65**, 3963–3977, doi:10.1175/2008JAS2733.1.
- Pauluis, O. and K. Emanuel, 2004: Numerical instability resulting from infrequent calculation of radiative heating. *Mon. Wea. Rev.*, **132**, 673–686, doi:10.1175/1520-0493(2004)132<0673:NIRFIC>2.0.CO;2.

- Pincus, R., 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:10.1209/2002JD003322.
- Pincus, R., R. Hemler, and S. A. Klein, 2006: Using stochastically-generated subcolumns to represent cloud structure in a large-scale mode. *Mon. Wea. Rev.*, **34**, 3644–3656, doi:10.1175/MWR3257.1.
- Räisänen, P., H. W. Barker, and J. N. S. Cole, 2005: The Monte Carlo Independent Column Approximation's conditional random noise: Impact on simulated climate. *J. Climate*, **22**, 4715–4730, doi:10.1175/JCLI3556.1.
- Siebesma, A. P., C. S. Bretherton, A. Brown, A. Chlond, J. Cuxart, P. Duynkerke, H. Jiang, M. Khairoutdinov, D. Lewellen, C.-H. Moeng, E. Sanchez, B. Stevens, and D. E. Stevens, 2003: A large-eddy simulation intercomparison study of shallow cumulus convection. *J. Atmos. Sci.*, **60**, 1201–1219, doi:10.1175/1520-0469(2003)60<1201:ALESIS>2.0.CO;2.
- Stevens, B., D. H. Lenschow, C.-H. M. Ian Faloona, D. K. Lilly, B. Blomquist, G. Vali, A. Bandy, T. Campos, H. Gerber, S. Haimov, B. Morley, and D. Thorton, 2003: On entrainment in nocturnal marine stratocumulus. *Quart. J. Roy. Meteor. Soc.*, **129**, 3469–3492, doi:10.1256/qj.02.202.
- Stevens, B., C.-H. Moeng, A. S. Ackerman, C. S. Bretherton, A. Chlond, S. de Roode, J. Edwards, J.-C. Golaz, H. Jiang, M. Khairoutdinov, M. P. Kirkpatrick, D. C. Lewellen, A. Lock, F. Müller, D. E. Stevens, E. Whelan, and P. Zhu, 2005: Evaluation of large-eddy simulations via observations of nocturnal marine stratocumulus. *Mon. Wea. Rev.*, **133**, 1443–1462, doi:10.1175/MWR2930.1.
- Várnai, T. and R. Davies, 1999: Effects of cloud heterogeneities on shortwave radiation: Comparison of cloud-top variability and internal heterogeneity. *J. Atmos. Sci.*, **56**, 4206–4224, doi:10.1175/1520-0469(1999)056<4206:EOCHOS>2.0.CO;2.
- Wapler, K. and B. Mayer, 2008: A fast three-dimensional approximation for the calculation of surface irradiance in large-eddy simulation models. *J. Appl. Meteor. Climatol.*, **47**, 3061–3071, doi:10.1175/2008JAMC1842.1.
- Weinbrecht, S. and P. J. Mason, 2008: Stochastic backscatter for cloud-resolving models. Part I: Implementation and testing in a dry convective boundary layer. *J. Atmos. Sci.*, **65**, 123–139, doi:10.1175/2007JAS2166.1.
- Xu, K.-M. and D. A. Randall, 1995: Impact of interactive radiative transfer on the macroscopic behavior of cumulus ensembles. Part I: Radiation parameterization and sensitivity tests. *J. Atmos. Sci.*, **52**, 785–799, doi:10.1175/1520-0469(1995)052<0785:IOIRTO>2.0.CO;2.
- Zuidema, P. and K. F. Evans, 1998: On the validity of the independent pixel approximation for boundary layer clouds observed during ASTEX. *J. Geophys. Res.*, **103**, 6059–6074, doi:10.1029/98JD00080.