



# The Subgrid Importance Latin Hypercube Sampler (SILHS): a multivariate subcolumn generator

V. E. Larson and D. P. Schanen

Department of Mathematical Sciences, University of Wisconsin – Milwaukee, Milwaukee, WI, USA

Correspondence to: V. E. Larson (vlarson@uwm.edu)

Received: 18 January 2013 – Published in Geosci. Model Dev. Discuss.: 19 March 2013

Revised: 30 August 2013 – Accepted: 12 September 2013 – Published: 29 October 2013

**Abstract.** Coarse-resolution climate and weather forecast models cannot accurately parameterize small-scale, nonlinear processes without accounting for subgrid-scale variability. To do so, some models integrate over the subgrid variability analytically. Although analytic integration methods are attractive, they can be used only with physical parameterizations that have a sufficiently simple functional form. Instead, this paper introduces a method to integrate subgrid variability using a type of Monte Carlo integration. The method generates subcolumns with suitable vertical correlations and feeds them into a microphysics parameterization. The subcolumn methodology requires little change to the parameterization source code and can be used with a wide variety of physical parameterizations.

Our subcolumn generator is multivariate, which is important for physical processes that involve two or more hydrometeor species, such as accretion of cloud droplets by rain drops. In order to reduce sampling noise in the integrations, our subcolumn generator employs two variance-reduction methods, namely importance and stratified (Latin hypercube) sampling. For this reason, we name the subcolumn generator the Subgrid Importance Latin Hypercube Sampler (SILHS).

This paper tests SILHS in interactive, single-column simulations of a marine stratocumulus case and a shallow cumulus case. The paper then compares simulations that use SILHS with those that use analytic integration. Although the SILHS solutions exhibit considerable noise from time step to time step, the noise is greatly damped in most of the time-averaged profiles.

## 1 Introduction

Large-scale models of the atmosphere typically parameterize subgrid-scale physical processes, such as microphysics or radiative transfer. A typical physical parameterization estimates the rate of a process at a point in space, whereas a coarse-resolution model instead needs the process rate averaged over a grid-box scale. Therefore, there is a benefit to upscaling point process rates to grid-box averages by accounting for subgrid-scale variability. Accurate upscaling is of potential importance for any process that involves significant small-scale spatial variability and is highly nonlinear. For instance, upscaling is expected to benefit nonlinear microphysical processes such as the autoconversion of cloud droplets to raindrops and the accretion of liquid cloud water by falling raindrops.

In many cases, the problem of upscaling reduces to the problem of performing integrals such as the following:

$$\langle f(x, y) \rangle = \int P(x, y) f(x, y) dx dy. \quad (1)$$

Here,  $x$  and  $y$  are fields that vary over a grid box, such as liquid cloud water and rain mixing ratios;  $P(x, y)$  is the subgrid probability density function (PDF) that provides the probability density that particular values of  $x$  and  $y$  occur within a particular grid box and time step; and  $f(x, y)$  is a function representing the process to be integrated, such as autoconversion or accretion. For illustration, we have written Eq. (1) as a bivariate function of  $x$  and  $y$ , but in principle, any number of variates could be used.

The function  $f(x, y)$  represents either an analytic function or a numerical subroutine. Such integrals may be performed by several alternative quadrature methods. As examples, we

cite two methods, namely analytic integration and Monte Carlo integration:

1. *Analytic integration.* If both the PDF,  $P(x, y)$ , and the process rate function,  $f(x, y)$ , are sufficiently simple analytic functions, then the integral  $\langle f(x, y) \rangle$  may be performed analytically. Numerous papers have performed such integrals in order to compute the liquid cloud fraction using an assumption of instantaneous saturation adjustment (e.g., Sommeria and Deardorff, 1977; Mellor, 1977; Smith, 1990; Lewellen and Yoh, 1993; Bony and Emanuel, 2001; Tompkins, 2002). Other papers have used analytic integration in order to upscale microphysical process rates (e.g., Zhang et al., 2002; Morrison and Gettelman, 2008; Cheng and Xu, 2009; Larson and Griffin, 2013; Griffin and Larson, 2013; Weber and Quaas, 2012).

Advantages of analytic integration include the facts that the integrals so obtained are exact and entail relatively little computational expense. Disadvantages include the facts that deriving such integrals is time-consuming and the facts that complicated process rate functions, including numerical algorithms, cannot be integrated analytically.

2. *Monte Carlo integration.* A second way to perform the integral in Eq. (1) is to use Monte Carlo integration. That is, one may draw a sample of  $x$  and  $y$  that is distributed according to the PDF  $P(x, y)$ , evaluate the process rate  $f(x, y)$  at those points, and average the resulting sample of process rates. Monte Carlo integration has been applied to radiative transfer parameterizations by the Monte Carlo independent column approximation (McICA) (e.g., Barker et al., 2002; Pincus et al., 2003; Räisänen et al., 2004; Räisänen and Barker, 2004; Räisänen et al., 2005; Pincus et al., 2006). A type of Monte Carlo integration has also been recommended for use with microphysics (Larson et al., 2005; Larson, 2007). A similar but deterministic sampling strategy is used in the Tripleclouds parameterization (Shonk and Hogan, 2008; Shonk et al., 2012).

A disadvantage of Monte Carlo integration is the sampling noise that is inherent in the use of a necessarily limited number of sample points. However, these sampling errors have been found to be of relatively little consequence in tests of McICA (Pincus et al., 2003; Räisänen et al., 2005; Pincus et al., 2006; Räisänen et al., 2007; Barker et al., 2008; Räisänen et al., 2008). Nevertheless, in order to reduce the sampling errors, variance-reduction techniques have been introduced for the McICA (Räisänen and Barker, 2004; Hill et al., 2011). Another method to reduce noise is the use of stratified sampling, in which sample points are chosen such that they are spread out and do not clump (Larson et al., 2005).

An advantage of Monte Carlo integration is that it can be applied quite generally to complicated process rates  $f(x, y)$ , including those that are numerical algorithms. Furthermore, Monte Carlo integration is non-intrusive; in other words, it allows the integral to be performed without requiring modification to the computer code that implements the point process rate  $f(x, y)$ . This is possible because Monte Carlo integration separates and distinguishes the representation of subgrid variability from the calculation of local physical processes. Therefore, Monte Carlo integration might be particularly useful for models like the Weather Research and Forecasting (WRF) model (Skamarock et al., 2005) that contain many options for physics parameterizations, and also useful for models whose physics packages are updated frequently.

Another advantage of some Monte Carlo integrators is that they allow a variety of assumptions about the correlation of sample points in the vertical, including both random and maximal overlap of vertical points. Flexibility of the vertical correlations is useful for driving radiative transfer parameterizations and diagnostic parameterizations of precipitation.

This paper describes and presents tests of a new Monte Carlo integration method called the Subgrid Importance Latin Hypercube Sampler (SILHS). SILHS generates a distribution of sample points that permits efficient Monte Carlo integration of integrals such as Eq. (1). The process rate function  $f(x, y)$  must be supplied separately by, for instance, a microphysics scheme. In addition, the PDF  $P(x, y)$  in Eq. (1) must be provided separately by, for instance, a cloud and turbulence parameterization such as the Cloud Layers Unified By Binormals (CLUBB) parameterization (Golaz et al., 2002; Larson and Golaz, 2005). Unlike previous Monte Carlo generators, which focused mostly on calculating cloud overlap and generating profiles of liquid cloud water, SILHS is used here to generate profiles of rain water and vertical velocity, along with profiles of liquid cloud water.

A key capability of SILHS is that it generates multivariate samples, unlike, for instance, the method of Räisänen et al. (2004), which does allow two variates to depend diagnostically on one another, but not to vary independently. The use of multivariate samples is important for processes such as accretion, which depends on the correlation of liquid cloud water and rain water. If the raindrops fall through cloud water, they accrete cloud water, whereas if they fall through clear air, they evaporate. Multivariate samples can also be generated by the use of copulas (Larson, 2007; Norris et al., 2008). Copulas are quite flexible and promising, but they do have the drawback of allowing only the specification of rank correlations among variables within a grid level and not the more familiar linear correlations.

SILHS generates vertical profiles of sample points, allowing the integral in Eq. (1) to be computed at each grid level. SILHS therefore may be called a subcolumn generator. SILHS uses both importance and stratified sampling in order to perform variance reduction, that is, to reduce sampling noise. SILHS provides an option to choose sample points preferentially within an important region, namely clouds containing liquid. SILHS also spreads out (i.e., stratifies) the sample points within the cloudy and clear regions separately using Latin hypercube sampling (McKay et al., 1979; Press et al., 1992; Owen, 2003; Gentle, 2003; Larson et al., 2005).

The remainder of this paper is organized as follows. In Sect. 2, we discuss the methodology behind SILHS, including the treatment of vertical overlap of clouds, reduction of noise, and multivariate correlations. In Sect. 2 we also briefly outline two other numerical models we use in conjunction with SILHS, namely CLUBB and a microphysics scheme (Khairoutdinov and Kogan, 2000). In Sect. 3, we compare the results of simulations that use Monte Carlo integration of the microphysics with those that use analytic integration for both a cumulus and a stratocumulus case. In Sect. 4, we present conclusions and provide a future outlook.

## 2 Methodology

This section describes the algorithm behind SILHS, first providing an overview of the role that SILHS plays within a single-column model, and then providing more detail on how subcolumns are generated.

The following overview describes the four major steps that our single-column model undertakes in a single time step:

1. *Construct the multivariate PDF of subgrid-scale variability (performed by CLUBB).* In our approach, sample points are drawn from an underlying PDF. This PDF is computed separately from SILHS and provided as input to SILHS. In this paper, the PDF is calculated by CLUBB, which is described in Golaz et al. (2002) and Larson and Golaz (2005). In this paper, CLUBB's PDF contains the following variates: vertical velocity  $w$ ; total water mixing ratio (vapor + liquid)  $r_t$ ; liquid water potential temperature  $\theta_l$ ; rain mixing ratio  $r_r$ ; and rain drop number mixing ratio  $N_r$ . The two rain variates are distributed according to a single lognormal, whereas the other variates are distributed according to a two-component normal mixture. (A two-component normal mixture is the sum of two normal distributions, i.e., a double Gaussian PDF. The normal-mixture terminology is widely used in statistics.)
2. *Draw a sample of subcolumns from the subgrid PDF (performed by SILHS).* Once the PDF has been calculated by CLUBB, a sample of subcolumns for each

grid column is drawn from it. The generation of subcolumns is done by SILHS. SILHS is designed so that in the limit of an infinite number of sample columns per grid column and time step, the sample statistics converge to the desired vertical correlations, to the desired marginal PDF at each grid level, and to the desired horizontal (i.e., within grid box) correlations between variates. (The marginal PDF of, for example, variate  $x$  is the PDF that remains when the full multivariate PDF is integrated over all variates other than  $x$ .) Because CLUBB computes a separate multivariate PDF at each vertical grid level (rather than a single grid-column PDF, as in Larson, 2007), CLUBB does not provide any information about vertical correlations. Instead, the vertical overlap of the variates is handled by SILHS.

3. *Feed subcolumns into microphysics and compute microphysical tendencies (performed by a microphysics parameterization).* Each subcolumn is fed, one by one, into a microphysics parameterization, which again is separate from CLUBB and SILHS. The microphysics parameterization remains unaltered and does not “know” that the profiles being fed to it are from subcolumns rather than grid-box means. The microphysics simply computes local quantities (such as autoconversion, accretion, and evaporation) on the assumption that there is no heterogeneity within the subcolumn (i.e., that at each grid level in a subcolumn, all fields are horizontally uniform). The microphysics calculates a separate set of microphysical tendencies for each subcolumn.
4. *Average microphysical tendencies from each subcolumn in order to form a grid-box average.* Finally, the microphysical tendencies from all the subcolumns within a grid column are averaged together, using a weighted mean if necessary (described further below). The resulting grid-box-averaged profile of microphysical tendencies is fed back into the host model (in this case CLUBB) and used to update the grid-box means. In this way, the subcolumns provide a means for the microphysics to interact with the subgrid variability provided by CLUBB.

Background to the present paper can be found in Larson et al. (2005), which describes a predecessor to SILHS. In that paper, Latin hypercube samples were drawn from CLUBB and fed into a simple autoconversion formula, but the simulations described there did not use a complete microphysics parameterization and had no feedback from the autoconversion back to CLUBB.

We now discuss further each of the four steps in the aforementioned overview. However, because the focus of this paper is the sampling that is carried out by SILHS, namely step 2, that is the only step that we will describe in detail.

## 2.1 Constructing the multivariate PDF of subgrid-scale variability

In order to compute the subgrid PDF from which SILHS draws sample points (step 1 in Sect. 2 above), this paper uses CLUBB (Golaz et al., 2002; Larson and Golaz, 2005; Larson et al., 2012). CLUBB computes the subgrid PDF by prognosing various higher-order central moments and making an assumption about the shape of the PDF. The prognostic equations contain much of the physics of the combined CLUBB–SILHS model. The prognosed higher-order moments include variances of turbulent and thermodynamic quantities ( $w$ ,  $r_t$ , and  $\theta_1$ ), covariances of those same quantities, and the third-order moment of vertical velocity. The framework of the prognostic equations is derived directly from the governing equations of fluid flow by forming higher-moment equations and Reynolds-averaging them. However, many of the individual terms within those equations are parameterized. Because these equations are prognostic, they contain some degree of memory of the state of the flow at a prior time step. The variances determine the width of the PDF; the covariances are related to the correlations between variates, and the third-order moment determines whether the PDF is skewed to low or high values. The moments, along with the assumption that the PDF has a normal-mixture/lognormal functional form, determine the PDF for each grid box and time step. We omit further details about CLUBB’s formulation because it has been described in previous papers (Golaz et al., 2002; Larson and Golaz, 2005; Larson et al., 2012).

The CLUBB (and SILHS) source code, along with separate documentation, is freely downloadable for non-commercial use at <http://clubb.larson-group.com/>. CLUBB and SILHS are written entirely in Fortran 95 and can be compiled using a number of Fortran compilers on the Linux operating system. Both CLUBB and SILHS are contained within a single subversion (svn) code repository. The simulations described and shown in this paper were created with revision 4895 of the source code. Further documentation and instructions for running CLUBB and SILHS are contained in several README files that are included along with the source code. These files serve collectively as a user manual. Readers who are interested in details of the algorithm are encouraged to peruse the source code. We have strived to write well-structured code that is annotated generously with code comments.

## 2.2 Drawing a sample of subcolumns from the subgrid PDF

In this subsection, we detail step two in Sect. 2, namely, the method by which sample points are drawn from the PDF of subgrid variability. In overview, the sampling strategy is to generate a uniformly distributed sample at a particular altitude of interest within the domain; then, starting from that altitude of interest, choose a vertically correlated profile

of sample points; and finally, transform the uniformly distributed sample points to the PDF at each level calculated by CLUBB.

In our case, the PDF of interest is a single multivariate PDF with normal-mixture marginals for  $w$ ,  $r_t$ , and  $\theta_1$ , and single lognormal marginals for  $r_r$  and  $N_r$  (Larson and Griffin, 2013). A normal-mixture marginal  $P(x)$  of a generic variate  $x$  has the following functional form:

$$P(x) = aP_1(x) + (1 - a)P_2(x), \quad (2)$$

where  $a$  is the mixture fraction and a normal component is given by

$$P_i(x) = \frac{1}{(2\pi)^{1/2}\sigma} \exp\left[-\frac{(x - \mu)^2}{2\sigma^2}\right]. \quad (3)$$

Here  $\mu$  and  $\sigma^2$  are the mean and the variance, respectively, of  $x$ . The mean and variance may differ between components.

A lognormal marginal  $P(x)$  of  $x$  has the form

$$P(x) = \frac{1}{(2\pi)^{1/2}\sigma_x} \exp\left[-\frac{(\ln x - \mu)^2}{2\sigma^2}\right]. \quad (4)$$

Here  $\mu$  and  $\sigma^2$  are the mean and the variance, respectively, of  $\ln x$ , and  $\ln x$  is normally distributed. A lognormal shape was chosen because it contains only positive values, which is appropriate for hydrometeors. It has plausible tails, and the mathematics to handle an arbitrary number of correlated lognormal variates is tractable. A multivariate lognormal is tractable because a multivariate normal is straightforward mathematically, and the normal and lognormal distributions are related by simple exponentiation (Garvey, 2000).

As seen in Eq. (4), a 1-D lognormal functional form is specified by two parameters, which can be related by a simple transformation to the lognormal’s mean and variance (Garvey, 2000, Appendix B). The means of  $r_r$  and  $N_r$  are calculated by the microphysics parameterization in use. SILHS diagnoses the variances of  $r_r$  and  $N_r$  as being proportional to the mean squared. For example, we set  $\overline{r_r^2}/\overline{r_r}^2$  equal to a constant everywhere throughout the simulation. The values of the constants of proportionality that we choose, for the marine stratocumulus case discussed below, are listed in Table 1 of Larson and Griffin (2013). These values are based on a large-eddy simulation. In order to improve the results for the shallow cumulus case discussed below, we increase  $\overline{N_r^2}/\overline{N_r}^2$  within cloud to 2.2; below cloud, we decrease  $\overline{N_r^2}/\overline{N_r}^2$  and  $\overline{r_r^2}/\overline{r_r}^2$  to 1. A simple proportionality seems unlikely to be accurate in all cases, particularly when the mean is large, but using a proportionality is computationally inexpensive and has some support from satellite observations (M. Lebsock, personal communication, 2012).

A multivariate subgrid PDF is associated with subgrid (within-grid-box or “horizontal”) correlations between the

variates. In nature, these correlations vary with time and space. However, in order to maintain simplicity and reduce computational cost, we have prescribed the horizontal correlations in the present paper (see further details below). The values of the horizontal correlations that this paper uses are based loosely on large-eddy simulations. For the marine stratocumulus case discussed below, we use the values of correlations listed in Table 1 of Larson and Griffin (2013). We use the same correlation values for the shallow cumulus case discussed below, except that, to improve the results for cumulus, the correlation between the extended liquid water,  $s$  (Larson et al., 2005), and the droplet number is set to zero. Prescribing constant horizontal correlations in the simulations is not ideal because correlations are not constant in nature (Larson et al., 2011). In future work, one could attempt to diagnose the correlations at each time step, using a methodology such as that described in Larson et al. (2011).

The procedure by which SILHS generates subcolumns can be divided into the following five tasks:

1. *Choose a starting grid level for sampling*

As discussed below, we impose vertical correlations by choosing a sample value at one grid level, moving to an adjacent vertical level, choosing a new but similar sample value, and so forth. The starting grid level can be any grid level (e.g., the lowermost level). However, if the starting grid level is far from a liquid cloud layer, then, because sample points gradually de-correlate with vertical distance, we cannot expect to choose sample points preferentially within liquid cloud, as we desire.

In order to mitigate this problem partly, we choose the starting grid level to be the level with the greatest liquid water mixing ratio. This does not help when a grid column contains, for instance, a large-liquid, stratocumulus layer and a small-liquid cumulus layer. However, it does help for grid columns that contain, for instance, a single cumulus layer. If no liquid is present at any level, then we choose the grid level that is half the maximum grid level. Although our tests seem to indicate that our criterion for choosing the starting grid level is acceptable, this is an area that is in need of further experimentation, particularly for grid columns with multiple, distinct cloud layers. In such cases, one could choose different starting grid levels for different subcolumns, thereby sampling preferentially to some degree within multiple layers.

2. *Generate an uncorrelated multivariate sample at the starting grid level*

Once the starting grid level is chosen, SILHS generates for that level an uncorrelated, multivariate sample. The distribution lies between (0,1) and is uniform except for weightings discussed below. Stated differently,

SILHS generates a vector of independent, uniformly distributed sample points, each element of which corresponds to a separate variate, such as  $w$ . Once the starting grid level is chosen, SILHS generates for that level an uncorrelated, multivariate sample. (The uncorrelated, uniformly distributed samples are transformed to correlated, normally distributed ones in task 4 below.)

Care must be taken in sampling the variates that have two mixture components, namely  $w$ ,  $r_t$ , and  $\theta_1$ . In order to ensure that the two mixture components are sampled in an unbiased way, SILHS assigns each sample point to one mixture component or the other according to the relative weights of the two mixture components. For concreteness, suppose that the weight of the first component is `mixt_frac` and the weight of the second is `1-mixt_frac`. To assign the sample point to a component, SILHS generates an extra random variate,  $X_{u\_dp1}$ , that is uniformly distributed between 0 and 1 (see p. 56 of Johnson, 1987, and Larson et al., 2005). Then SILHS chooses the first component if  $X_{u\_dp1} < \text{mixt\_frac}$  and chooses the second component otherwise. All further discussion in this subsection pertains to a single mixture component.

Although the points are drawn from a uniform distribution, the sampling procedure of SILHS is complicated slightly by the fact that SILHS uses both importance sampling and stratified sampling in order to reduce noise.

First, if liquid cloud fraction is small at the starting grid level, then SILHS contains the option to sample preferentially within an important region, namely cloud containing some liquid (Räisänen and Barker, 2004); therefore SILHS does importance sampling (Gentle, 2003). SILHS concentrates sample points within liquid/mixed-phase cloud because regions with liquid often contain considerable variability, and much interesting microphysics occurs there. In particular, if liquid cloud fraction exceeds 0.5, then SILHS does not sample preferentially within cloud, but if liquid cloud fraction lies between 0.001 and 0.5, then SILHS chooses an equal number of sample points within liquid cloud (i.e., where  $s > 0$ ) and outside of liquid cloud (i.e., where  $s < 0$ ) for each mixture component. Because the samples are drawn preferentially from liquid cloud in grid boxes with low cloud fraction, SILHS must weight the sample points appropriately by liquid cloud fraction in order to ensure that grid-box averages are unbiased. Namely, the weights outside of liquid cloud are given by

$$\frac{1 - (\text{cloud fraction at starting grid level})}{(\text{number of samples})/2}, \quad (5)$$

and the weights within liquid cloud are given by

$$\frac{\text{cloud fraction at starting grid level}}{(\text{number of samples})/2}. \quad (6)$$

At this initial step in the algorithm, the sample points are generated according to a uniform distribution within liquid cloud and a uniform distribution outside of cloud. To generate uniformly distributed points, SILHS uses the Mersenne twister algorithm (Matsumoto and Nishimura, 1998).

When liquid cloud fraction at the starting grid level lies between 0.001 and 0.5, SILHS chooses half the sample points within liquid cloud. If one were interested only in within-cloud processes, such as auto-conversion and accretion, one could choose all points within cloud. However, some interesting processes do occur outside of cloud, such as evaporation of raindrops falling alongside sheared cumulus clouds and clear-sky radiative transfer. Therefore, SILHS chooses some sample points outside of cloud.

In a second measure to reduce noise, SILHS employs stratified sampling. Stratified sampling spreads sample points in order to avoid the clumping of points that naturally arises due to statistical chance. In particular, SILHS uses Latin hypercube sampling separately in both the cloudy and cloud-free regions. The methodology behind Latin hypercube sampling is thoroughly described in many sources (e.g., McKay et al., 1979; Press et al., 1992; Owen, 2003; Gentle, 2003; Larson et al., 2005). In most applications, Latin hypercube sampling guarantees that sample points are spread out (i.e., stratified), but CLUBB–SILHS does not necessarily prohibit clumping of sample points. The reason is the following. CLUBB uses a two-component mixture for some variates. For those variates, SILHS guarantees that the points assigned to particular component are stratified, but SILHS does not guarantee that the collection of points from both components, taken together, are stratified. In other words, SILHS does not attempt to ensure that, say, a high-percentile sample value chosen from one component with a low mean exceeds a low-percentile sample value chosen from the high-mean component. Nevertheless, SILHS's Latin hypercube sampling does usually increase stratification of sample points.

### 3. Generate vertically correlated profiles of sample points

Cloud parcels at different altitudes within a grid-column-sized volume are said to be “vertically overlapped” if one cloud parcel resides vertically above the other. The vertical overlap of clouds strongly influences cloud albedo (e.g., Morcrette and Fouquart,

1986). Two possible overlap assumptions are maximum overlap, in which cloud fields are stacked vertically to the greatest extent possible, and random overlap. These assumptions have the drawback of introducing spurious dependence on the vertical grid spacing. In recent years, several authors have assumed that clouds are overlapped according to a weighted average of maximum and random overlap (Hogan and Illingworth, 2000; Bergman and Rasch, 2002; Räisänen et al., 2004; Pincus et al., 2005; Barker, 2008; Shonk et al., 2010; Oreopoulos et al., 2012). The weight of the maximum overlap component is assumed to decrease exponentially with distance between the layers, with an e-folding length that needs to be determined. The e-folding length may vary with hydrometeor type (Pincus et al., 2005), meteorological regime (Pincus et al., 2005), and/or latitude (Shonk et al., 2010; Oreopoulos et al., 2012).

SILHS adopts a related but different approach. SILHS does not assume that the points are a weighted average of maximum and random overlap but does choose sample points that decorrelate exponentially with increasing vertical distance between sample points:

$$\text{vert\_corr} = \exp(-\text{vert\_decorr\_coef} \Delta z / \text{Lscale}). \quad (7)$$

Here `vert_corr` is not the correlation itself between points in the vertical, but `vert_corr` does increase with increasing correlation, as will become clear momentarily. The quantity  $\Delta z$  is the distance between grid levels, `Lscale` is CLUBB's turbulent mixing length, and `vert_decorr_coef` is a constant number that allows SILHS to adjust the degree of vertical overlap and is provisionally set to 0.03 in this paper. Given a uniformly distributed random number,  $X_u(k)$ , between 0 and 1 at grid level  $k$ , SILHS chooses the value at an adjacent grid level (e.g.,  $X_u(k+1)$ ) according to a uniform distribution that is centered on the value  $X_u(k)$  and has a half-width of  $1 - \text{vert\_corr}$ . If  $X_u(k+1)$  lies outside  $(0, 1)$ , then the value is folded back so that it does lie within  $(0, 1)$ . That is, if  $X_u(k+1) > 1$ , then we set  $X_u(k+1) = 2 - X_u(k+1)$ ; if  $X_u(k+1) < 0$ , then we set  $X_u(k+1) = -X_u(k+1)$ . Such correlated vertical profiles are built for each variate, including the extra  $X_{u\_dp1}$  variate that determines the mixture component. For simplicity, each variate uses the same value of `vert_corr`, even though in nature the vertical correlations of all quantities may not necessarily be equal.

Given that the half-width equals  $1 - \text{vert\_corr}$ , `vert_corr` = 1 corresponds to maximum overlap, `vert_corr` = 0 to random overlap, and  $0 < \text{vert\_corr} < 1$  to intermediate overlap. The degree of vertical correlation depends on CLUBB's buoyant mixing length,

Lscale (see Eq. 7). When Lscale is large, then parcels are allowed to travel further in the vertical direction (Golaz et al., 2002), which we assume is indicative of higher vertical coherence. This link between CLUBB's mixing length and the decorrelation length introduces an approximate but interactive, dynamical aspect into the diagnosis of vertical overlap.

Ensuring that the samples are drawn from the desired means, variances, and covariances at each grid level is accomplished later by the transformation from a uniform distribution to a normal-mixture/lognormal distribution described below. Note that SILHS directly influences only the vertical correlation of uniformly distributed points and not the normal-mixture/lognormal points. Although the vertical correlations of the normal-mixture/lognormal points are related to those of uniformly distributed points, they are usually not equal.

#### 4. Transform uncorrelated, uniformly distributed points to correlated, normally distributed points

At this point in the algorithm, SILHS has generated multiple subcolumns of uncorrelated, uniformly distributed (but weighted) sample points. Furthermore, SILHS has assigned each sample point to a mixture component by the aforementioned  $X_{u\_dp1}$  variate. SILHS now transforms these samples to a normal-mixture distribution with the desired horizontal correlations. Even the hydrometeor variates, which are ultimately transformed to single lognormal distributions, are transformed here to a normal-mixture distribution (with identical means and variances in each mixture component so that the distribution collapses to a single normal).

First, SILHS transforms the sample points for each mixture component from an uncorrelated uniform distribution to an uncorrelated standard normal distribution  $\mathbf{x}_{\text{std}}$  (i.e., a normal distribution with zero mean and unit variance). This transformation is accomplished by use of the inverse cumulative distribution function (Johnson, 1987; Larson et al., 2005). Second, we transform  $\mathbf{x}_{\text{std}}$  to a normally distributed sample,  $\mathbf{x}_{\text{nonstd}}$ , with the desired subgrid covariance matrix. The purpose is to take into account the (horizontal) correlations among variates. The covariance matrix depends on the mixture component to which the sample point is assigned, but this mixture component is identified by the  $X_{u\_dp1}$  variate. To transform the sample points from standard normal to non-standard normal, we use the commonly adopted formula (Johnson, 1987, pp. 52–54)

$$\mathbf{x}_{\text{nonstd}} = \mathbf{L}_{\text{cov}}\mathbf{x}_{\text{std}} + \boldsymbol{\mu}, \quad (8)$$

where  $\boldsymbol{\mu}$  is the vector of means of the normal-mixture component and the matrix  $\mathbf{L}_{\text{cov}}$  is the Cholesky decomposition of the matrix of subgrid (horizontal) covariances among all variates, including  $w$ ,  $r_t$ ,  $\theta_l$ , and all hydrometeors. For the lognormal variates (i.e., the hydrometeors),  $\boldsymbol{\mu}$  is not the vector of means of the lognormal variates but rather the means of  $\ln \mathbf{x}_{\text{nonstd}}$ , such that when  $\mathbf{x}_{\text{nonstd}}$  is exponentiated, the desired means of the lognormal variates are recovered (see Eq. 4). To do so, we use formulas from Garvey (2000) and Larson and Griffin (2013). In a similar fashion, for lognormal variates, the covariances associated with  $\mathbf{L}_{\text{cov}}$  have been transformed from the lognormal space to a normal space (Garvey, 2000; Larson and Griffin, 2013).

Computing the Cholesky decomposition of the covariance matrix,  $\mathbf{L}_{\text{cov}}$ , at each time step and for each grid box would be excessively expensive. To avoid this cost and simplify the formulation, we prescribe the horizontal correlations between hydrometeors as fixed constants for all grid boxes and time steps. To do so, given the correlations, we compute the Cholesky decomposition of the correlation matrix,  $\mathbf{L}_{\text{corr}}$ . Because the correlations are constant,  $\mathbf{L}_{\text{corr}}$  needs to be computed only once at the beginning of a simulation. We transform  $\mathbf{L}_{\text{corr}}$  to the Cholesky decomposition of the covariance matrix,  $\mathbf{L}_{\text{cov}}$ , by row multiplication of  $\mathbf{L}_{\text{corr}}$  by the standard deviation,  $\sigma(i)$  of each variate  $i$ :

$$\mathbf{L}_{\text{cov}}(i, j) = \sigma(i)\mathbf{L}_{\text{corr}}(i, j). \quad (9)$$

This multiplication does need to occur at each time step, but the multiplication is much cheaper than performing a Cholesky decomposition.

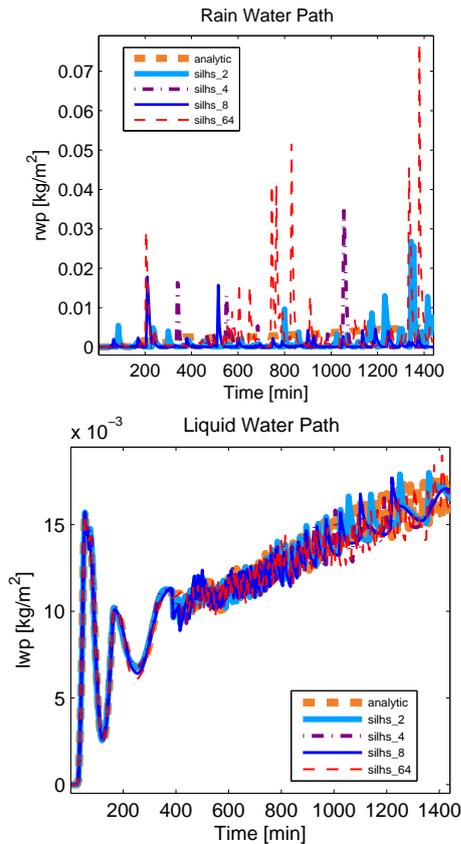
#### 5. Exponentiate certain hydrometeor variates in order to transform them to lognormal distributions

CLUBB and SILHS assume that hydrometeor variates such as  $r_r$  and  $N_r$  obey single lognormal distributions. These variates are transformed from normal to lognormal distributions by exponentiation of the sample points. A single, rather than double, lognormal results because previously we assumed the same mean for each of the two components of a given hydrometeor species, and also the same variance.

### 2.3 Feeding subcolumns into microphysics and computing microphysical tendencies

The subcolumns can, in principle, be fed into a variety of physics parameterizations if so desired, thereby driving the parameterized physical processes with subgrid variability. However, in this paper, we feed the subcolumns only into a microphysics parameterization.

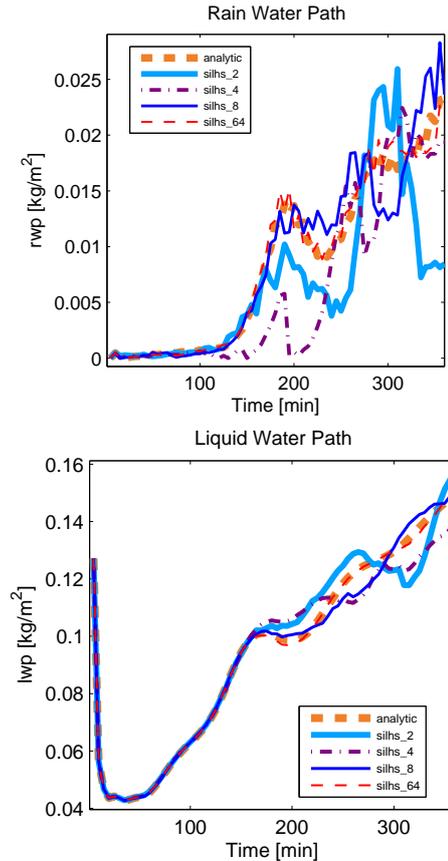
In order to compute microphysical process rates (step 3 in Sect. 2), this paper uses the drizzle parameterization



**Fig. 1.** Time series of rain water path (upper) and liquid water path (lower) from the RICO shallow cumulus case. Five different integrations of the microphysics are overplotted: analytic (orange dashed line), and SILHS-based sampling with 2 points (light blue solid line), 4 points (purple dot-dashed line), 8 points (dark blue solid line), and 64 points (red dashed line) per grid box and time step.

of Khairoutdinov and Kogan (2000). The Khairoutdinov–Kogan (KK) parameterization was formulated based on a detailed simulation of a shallow marine stratocumulus cloud and does not contain ice. In our simulations, we prescribe cloud droplet number mixing ratio. The KK parameterization prognoses both rain mixing ratio and rain number mixing ratio. One could compute the sedimentation *term* in the rain budget equation separately for each subcolumn, but for simplicity, this paper computes the sedimentation term only once using grid-box mean profiles. However, the grid-mean sedimentation *velocity*, which is used in sedimentation term, is obtained by averaging sedimentation velocities computed separately for each subcolumn.

The advantage for this paper of using the KK microphysics parameterization is that the KK process rates are formulated in terms of power laws, which can be integrated analytically over CLUBB’s PDF (Larson and Griffin, 2013; Griffin and Larson, 2013). Simulations that use analytic integration will provide a reference against which we will compare



**Fig. 2.** As in Fig. 1, but for the DYCOMS-II RF02 marine stratocumulus case. Using more sample points reduces the noise.

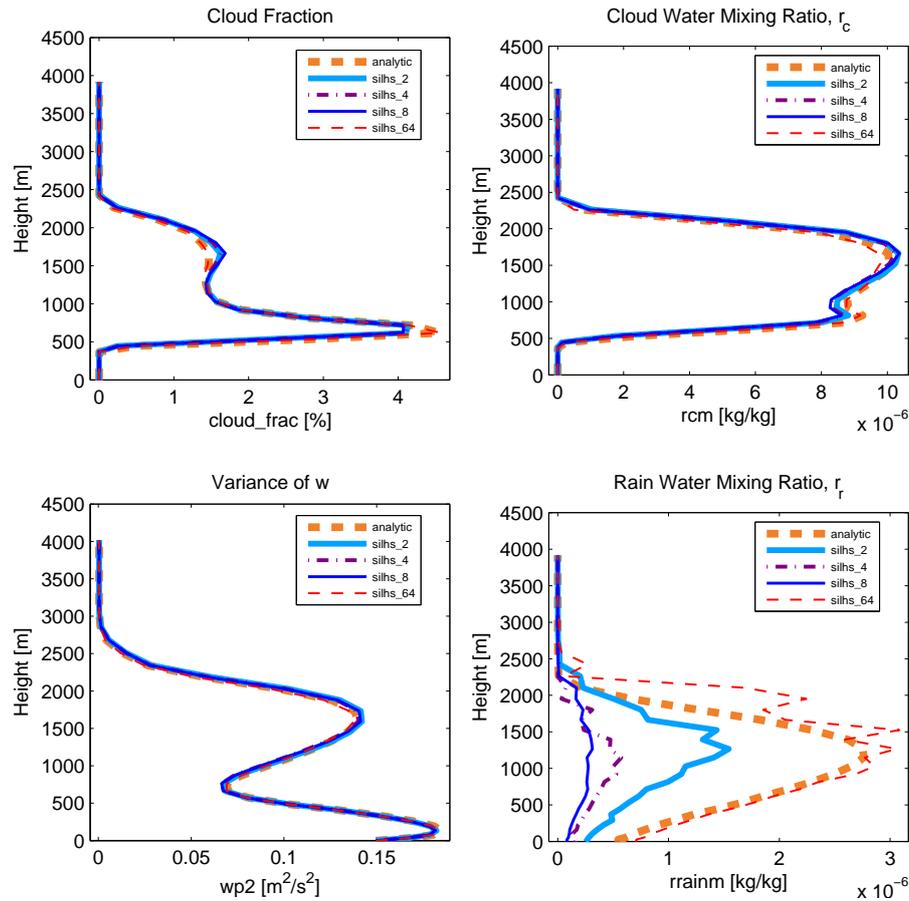
simulations that use the SILHS Monte Carlo integrations. In the limit of an infinite number of sample points, the SILHS integration should converge to the analytic integration.

#### 2.4 Averaging microphysical tendencies from each subcolumn in order to form a grid-box average

To allow SILHS to be interactive, the subcolumn microphysical tendencies must be fed back into the grid-box mean equations of the host model, which, in this case, is CLUBB (see step 4 in Sect. 2). The averaging of the microphysical tendencies from each sample point is a straightforward calculation that needs to be inserted into the host model. Because subcolumns are chosen preferentially within cloud, the averages must be appropriately weighted according to cloud fraction.

### 3 Results

In order to illustrate features of the solutions obtained by SILHS, we simulate two cloud cases using a single-column model that combines CLUBB and SILHS in an interactive fashion (“CLUBB–SILHS”). One case involves drizzling shallow cumulus observed during the Rain In Cumulus over

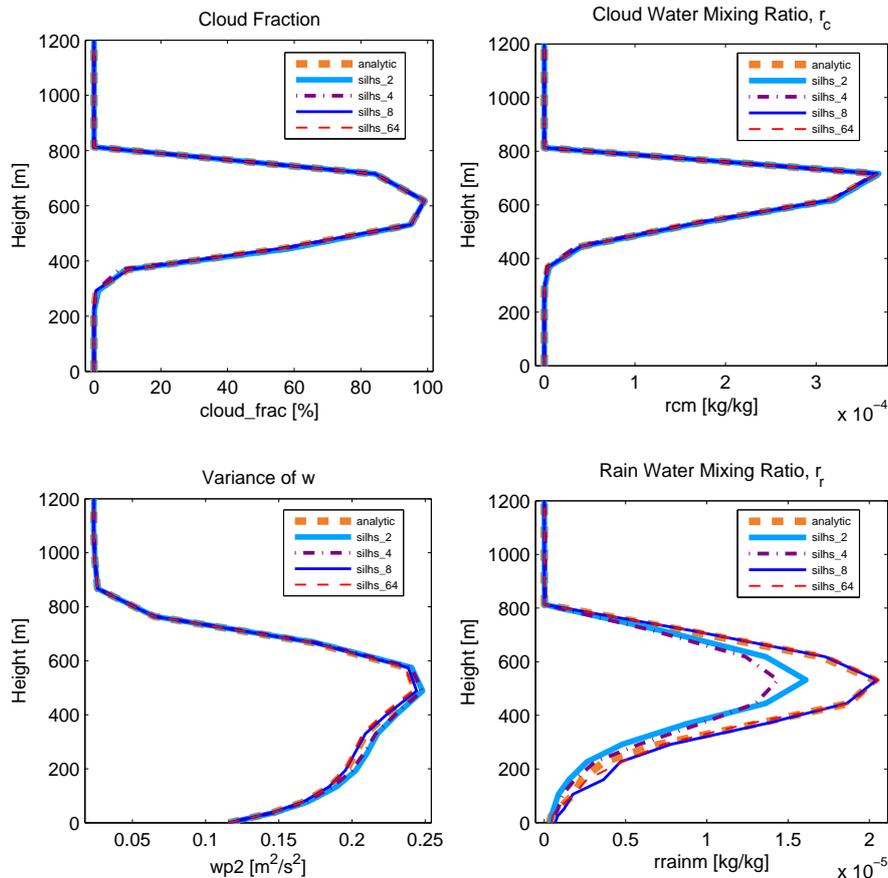


**Fig. 3.** Profiles of cloud fraction (upper left), cloud liquid mixing ratio (upper right), variance of vertical velocity (lower left), and rain mixing ratio (lower right) from the RICO shallow cumulus simulation. Five different integrations of the microphysics are overplotted: analytic (orange dashed line), and SILHS-based sampling with 2 points (light blue solid line), 4 points (purple dot-dashed line), 8 points (dark blue solid line), and 64 points (red dashed line) per grid box and time step. Profiles are averaged over the last 12 h of the simulation. Averaged over 12 h, the noise manifested in the time series (Fig. 1) is significantly smoothed, except in rain mixing ratio itself.

the Ocean (RICO) field experiment (Rauber et al., 2007). The initial profiles, forcings, and surface fluxes of our configuration follow the GEWEX Cloud System Study (GCSS) specifications for the RICO case (vanZanten et al., 2011). The other case involves a drizzling marine stratocumulus cloud observed during Research Flight 2 (RF02) of the Second Dynamics and Chemistry of Marine Stratocumulus (DYCOMS-II) field experiment (Stevens et al., 2003). Our configuration again follows GCSS specifications (Wyant et al., 2007). As per GCSS specifications, for the RICO simulation, surface latent and sensible heat fluxes are computed interactively in terms of surface wind speed, temperature, and water vapor mixing ratio, whereas for the DYCOMS-II RF02 simulation, the surface fluxes are prescribed. Both the RICO and DYCOMS-II RF02 cases were simulated using a 5 min time step, and both used a stretched vertical grid with 128 levels and a vertical grid spacing of about 100 m at 1000 m altitude.

Figures 1 to 4 display results from 5 single-column simulations that are configured identically, except that one uses 2 sample points per grid box and time step (light blue solid line), another 4 points (purple dot-dashed line), another 8 points (dark blue solid line), another 64 points (red dashed line), and the final forgoes SILHS and instead integrates the microphysics analytically over the PDF (orange dashed line). The CLUBB simulations with analytic integrations serve as reference simulations to which the corresponding CLUBB–SILHS solutions should converge. Comparing the CLUBB–SILHS simulations to the CLUBB simulations with analytic integration allows us to isolate the effect of noise in the simulations from unrelated model errors in CLUBB. The analytic integration methodology is described in Larson and Griffin (2013) and Griffin and Larson (2013).

The time series of rain simulated by CLUBB–SILHS shows considerable noise, but the noise is reduced in liquid water path (LWP) (see Figs. 1 and 2). Rain exhibits



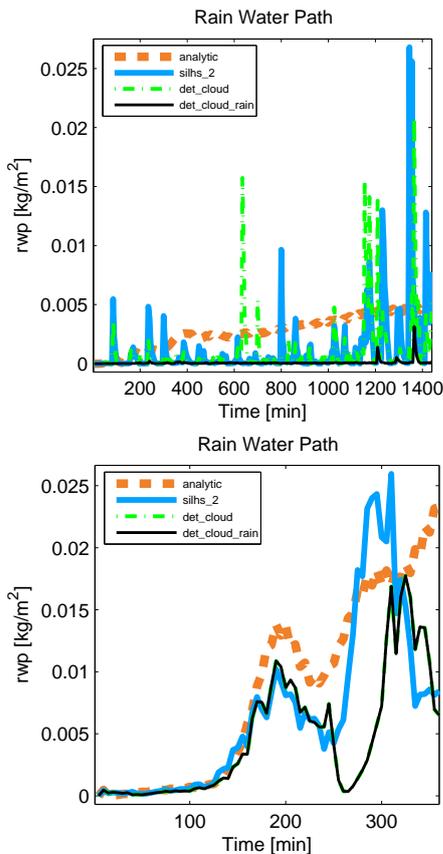
**Fig. 4.** As in Fig. 3, except for the DYCOMS-II RF02 marine stratocumulus simulation. Profiles are averaged over the full duration of the simulation (6 h). Averaged over 6 h, the noise manifested in the time series (Fig. 2) is considerably smoothed in all fields.

substantial noise in part because it is updated directly by SILHS's microphysical tendencies, and this injection of noise at each time step is not overcome by negative feedbacks and diffusion of rain. Liquid water, on the other hand, is influenced only indirectly by noise from subcolumns. As a result, subcolumn noise is felt more strongly by rain than liquid water. Rain exhibits diminished but still considerable noise even when 64 sample points per grid box and time step are used (red dashed line). However, large-eddy simulations of shallow cumulus clouds with  $\sim 5 \text{ km} \times 5 \text{ km}$  domains also exhibit noise in horizontally averaged liquid water path and rain water path.

Although the time series exhibit considerable noise from time step to time step, this noise is not evident in most profiles when they are averaged over long time periods. Figure 3 shows profiles from RICO averaged over 12 h, and Fig. 4 shows profiles from DYCOMS-II RF02 averaged over 6 h. We did not use longer averaging periods because longer periods would span too much of the diurnal cycle. In both cases, the noise is greatly reduced in cloud fraction, liquid cloud water mixing ratio ( $\overline{r_c}$ ), and variance of vertical

velocity ( $\overline{w^2}$ ). One exception is profiles of rain water mixing ratio for RICO, which do not match the analytic solution even when 64 points are used. RICO fields are noisier than those in DYCOMS-II RF02 because RICO, being a cumulus case, has greater horizontal variability. Nevertheless, the noise in rain, which is directly updated by the subcolumn microphysics tendencies, does not infect the long time averages of cloud and turbulence fields, which are not directly updated by the subcolumns. Moreover, the plots (and other unshown tests) show that the SILHS solution gradually approaches the analytic solution as more sample points are used, which is an important property. We note that if a simulation uses a time step longer than 5 min we used, then a longer averaging time will be necessary to reduce sampling noise.

Does sampling the within-cloud value of liquid water mixing ratio with only 2 sample points improve the results as compared to using a deterministic average of the within-cloud liquid water or grid-averaged rain? We have performed calculations in which liquid is set to the within-cloud average, rather than being sampled, in a call to the microphysics, and in which both liquid and rain are averaged deterministically rather than sampled randomly. The result is shown



**Fig. 5.** Time series of rain water path from the RICO shallow cumulus simulation (upper panel) and from the DYCOMS-II RF02 marine stratocumulus simulation (lower panel). Shown are an analytic integration of the microphysics (orange dashed line), SILHS-based sampling with 2 points (light blue solid line), deterministic use of the within-cloud liquid cloud water mixing ratio (green dashed line), and deterministic use of within-cloud liquid and rain mixing ratio (black solid line). In the RICO case, treating liquid deterministically but not rain still permits sampling noise in the rain; treating both liquid and rain deterministically leads to an underprediction of rain.

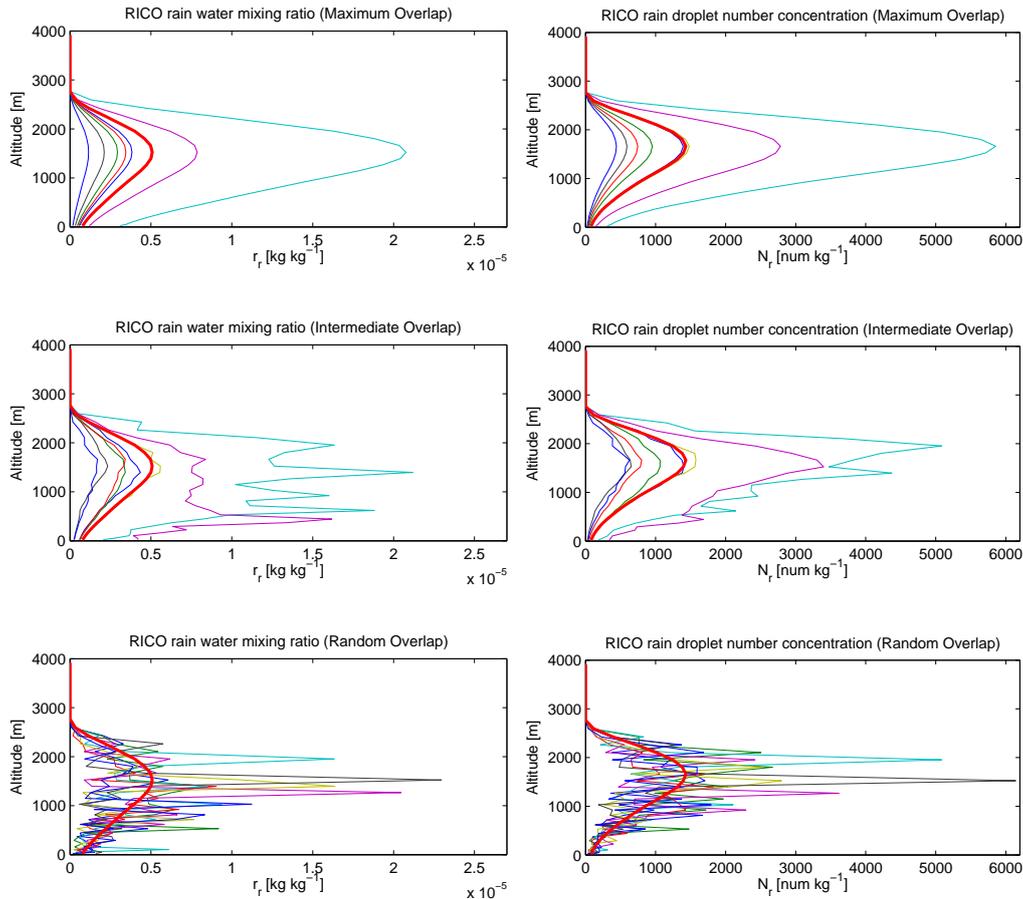
for RICO and DYCOMS-II RF02 in Fig. 5. For DYCOMS-II RF02, all results are similar, because this cloud is fairly homogeneous. For RICO, however, calculating liquid deterministically but sampling rain leads to as much noise as sampling both liquid and rain. However, calculating both liquid and rain deterministically leads to an underprediction of rain. This indicates that accounting for the non-zero variance of rain is important in the RICO case because of nonlinearity in the rain processes.

Examples of the vertical overlapping of profiles produced by SILHS are shown in Fig. 6. This figure shows profiles of rain mixing ratio,  $r_r$ , and rain number mixing ratio,  $N_r$ , for RICO. The maximally overlapped profiles (upper panel), which use  $\text{vert\_corr} = 1$  in Eq. (7), are smooth because SILHS attempts to maximize the vertical correlations

while preserving the grid-box means and variances at each grid level. The randomly overlapped profiles (lower panel) show no correlation between adjacent vertical grid levels, except the correlation that is inherent in following the vertical trend of the mean profiles. The intermediately overlapped profiles (middle panel), which use  $\text{vert\_decorr\_coef} = 0.03$  in Eq. (7), show an intermediate degree of vertical correlation. A separate, desirable feature of these profiles is that, collectively, they preserve the strong positive correlation between  $r_r$  (left figure column) and  $N_r$  (right figure column) that has been imposed by Eq. (8). The correlation between  $r_r$  and  $N_r$  is depicted by the color coding of the subcolumns, so that, for instance, the blue lines correspond to the  $r_r$  and  $N_r$  variates of the same (multivariate) subcolumn. We see that for each subcolumn, when  $r_r$  is large,  $N_r$  tends to be large as well, as desired for positively correlated variates. Similar vertical overlap properties are exhibited for DYCOMS-II RF02 (see Fig. 7). For ease of interpretation, in these plots, SILHS does not preferentially sample within cloud but instead chooses each subcolumn with equal probability, so that each thin line on the plot has equal weight.

In comparison to the profiles of  $r_r$  and  $N_r$ , the vertical overlapping of profiles of vertical velocity,  $w$ , has different properties owing to the fact that  $w$  is distributed according to a two-component normal mixture (see Fig. 8 for RICO and Fig. 9 for DYCOMS-II RF02). Namely, the profiles with maximal overlap are smooth except for a small number of discrete jumps. These jumps in the profiles occur when the mixture component from which the sample point is drawn switches from one component to the other. The switch in components occurs because, in a maximally overlapped profile, the extra variate that determines the choice of mixture component,  $X_{u\_dp1}$ , is constant with altitude, but the weight of the first normal component,  $\text{mixt\_frac}$ , varies with altitude. If the profile of  $\text{mixt\_frac}$  crosses the value of  $X_{u\_dp1}$ , then a jump in the profile will occur. The jumps are unrealistic but are difficult to avoid if an unbiased distribution and a two-component mixture are to be preserved.

In SILHS's cheapest available configuration, which generates two sample points per grid box and time step, the computational cost of SILHS is almost as large as the cost of CLUBB and nine times larger than the cost of the (inexpensive) Khairoutdinov–Kogan microphysics parameterization. For comparison, CLUBB occupies roughly 20 % of the total runtime of atmospheric climate models (Bogenschütz et al., 2013). The cost of SILHS is large in part because SILHS assumes that the marginal PDFs are normal-mixture or log-normal functions. These PDF shapes, while plausible for atmospheric applications, require the computation of special functions such as the exponential function and the inverse error function. Computing these special functions is expensive. Additional costs are incurred by the computation of the exponential decay of vertical correlation, the matrix multiplication needed to generate correlated samples, and the generation of sample points that are stratified. More details on the



**Fig. 6.** A sample of eight equally weighted SILHS-generated profiles (thin lines) and the analytic grid-box mean profile (thick red line) drawn from a single time step during the RICO shallow cumulus simulation. (Importance sampling is foregone for the purpose of illustration in Figs. 6 through 9.) The vertical correlation within each profile is specified to be maximal (upper panels), intermediate (middle panels), or random (lower panels). Plotted are the rain mixing ratio (left figure column) and rain number mixing ratio (right figure column) fields from the same eight profiles. The color coding of the sample profiles is the same in both figure columns. By comparing lines of the same color in both the right-hand and left-hand halves of the figure, one may see that the profiles of rain mixing ratio and number mixing ratio are highly correlated to each other. This result is expected because high correlation is specified in these simulations. The degree of vertical overlap in SILHS can be adjusted as desired between maximal and random.

relative costs of SILHS, CLUBB, and the microphysics are presented in Table 1.

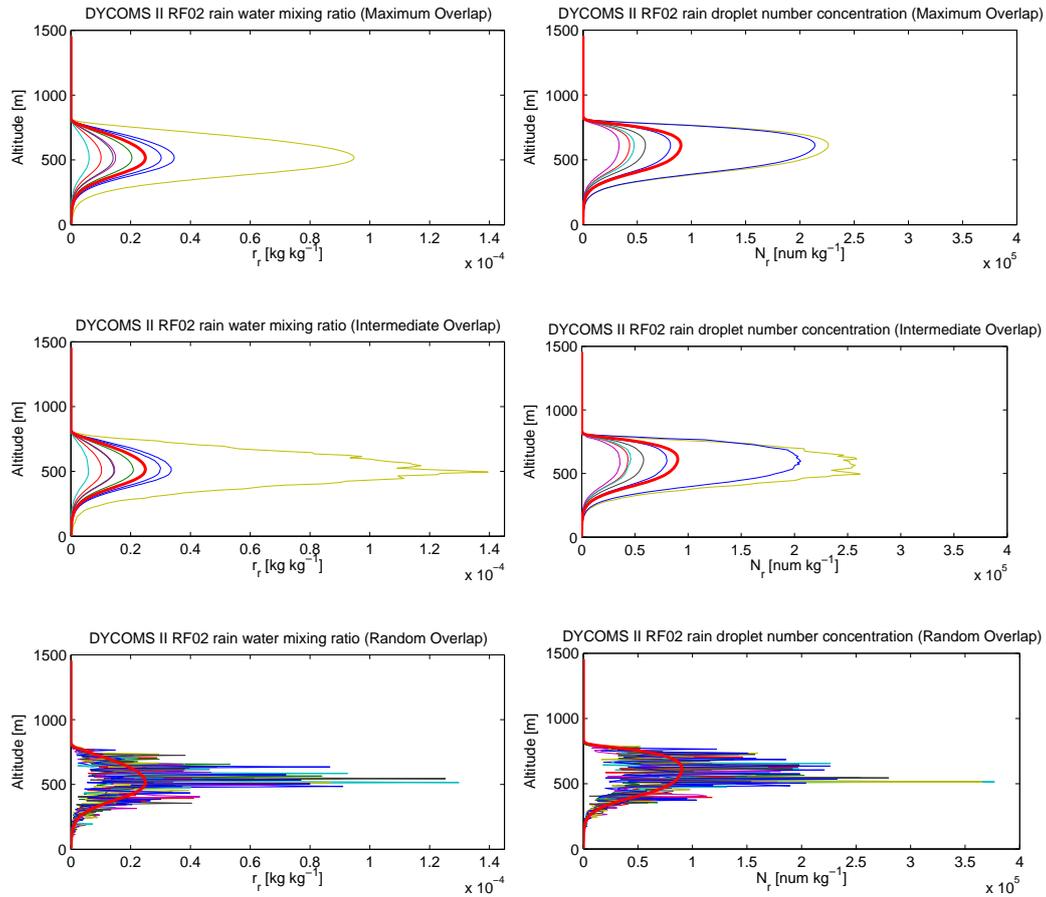
#### 4 Discussion and conclusions

In this paper, we have presented a new method, implemented in a software package called “SILHS”, that generates subcolumns for atmospheric models. The subcolumns, once generated, are fed into a microphysics parameterization, and the microphysical tendencies so produced are averaged and used to update the grid-box mean fields. We have tested SILHS in 1-D simulations of a shallow cumulus layer and a stratocumulus layer. In these two simulations, although SILHS does introduce noise into rain water mixing ratio, this noise does

not significantly degrade the time averages of liquid cloud water, cloud fraction, or other fields.

SILHS allows users to choose the number of sample points per grid box and time step. Increasing the number of sample points reduces statistical noise but increases the computational cost. We suspect that most users of SILHS with computational constraints will choose to use two sample points per grid box and time step. Even with only two sample points, SILHS allows the microphysics to sample both the clear-sky and within-cloud variability, thereby avoiding errors such as the systematic biases that result when convex or concave functions are fed grid-box means (e.g., Cahalan et al., 1994; Larson et al., 2001).

SILHS has disadvantages and advantages over other possible approaches to generating subcolumns, several of which are listed here:



**Fig. 7.** As in Fig. 6, but for the DYCOMS-II RF02 marine stratocumulus simulation. The profiles of rain mixing ratio and number mixing ratio have only moderate correlation, as specified by the user.

**Table 1.** Relative computational expense of SILHS, CLUBB, and the KK microphysics. Two sample points per grid box and time step are drawn. The cost of SILHS is broken into various sub-costs. This simulation was run on a single processor. The code was compiled using the SunStudio Fortran compiler under the Linux operating system, and the timing was performed by Sun Analyzer 7.9.

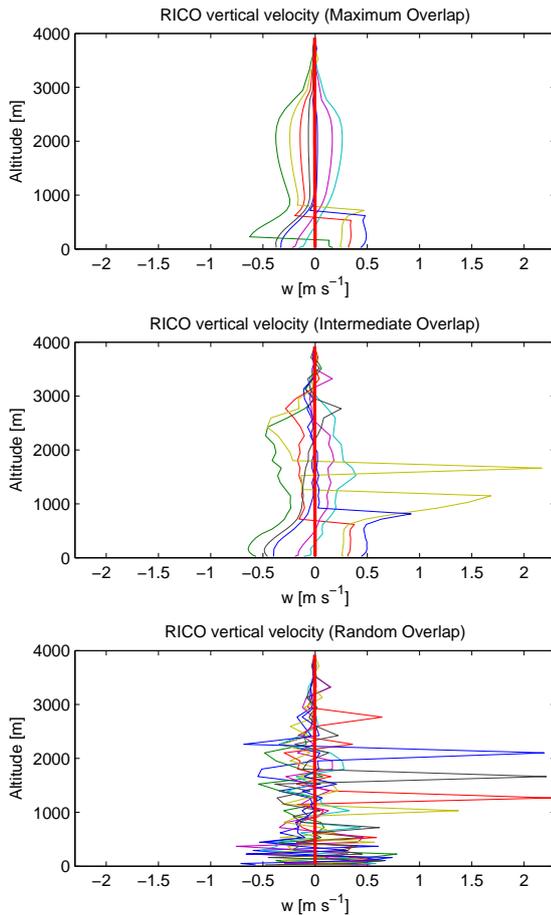
Code portion	% computational expense
CLUBB	50 %
Microphysics	5 %
SILHS	45 %
Stratify sample points	3 %
Compute vertical correlations	3 %
Compute cumulative distribution function	5 %
Multiply Cholesky matrix	5 %
Convert to lognormal means	2 %
Exponentiate lognormal variates	2 %
Miscellaneous SILHS	25 %

*Disadvantages:*

1. As noted earlier, the computational cost of SILHS is large.
2. SILHS preferentially samples within liquid clouds but does not preferentially sample within ice clouds.

*Advantages:*

1. SILHS is multivariate. That is, SILHS generates samples that have the desired prescribed correlations between variates, such as liquid cloud water and rain. Allowing non-zero correlations is useful for modeling processes (such as accretion of cloud droplets by rain drops) that depend on the correlation between two hydrometeor species.
2. SILHS contains methods to reduce sampling noise. In particular, SILHS includes importance sampling because it samples preferentially within clouds when cloud fraction is small. SILHS also includes stratified

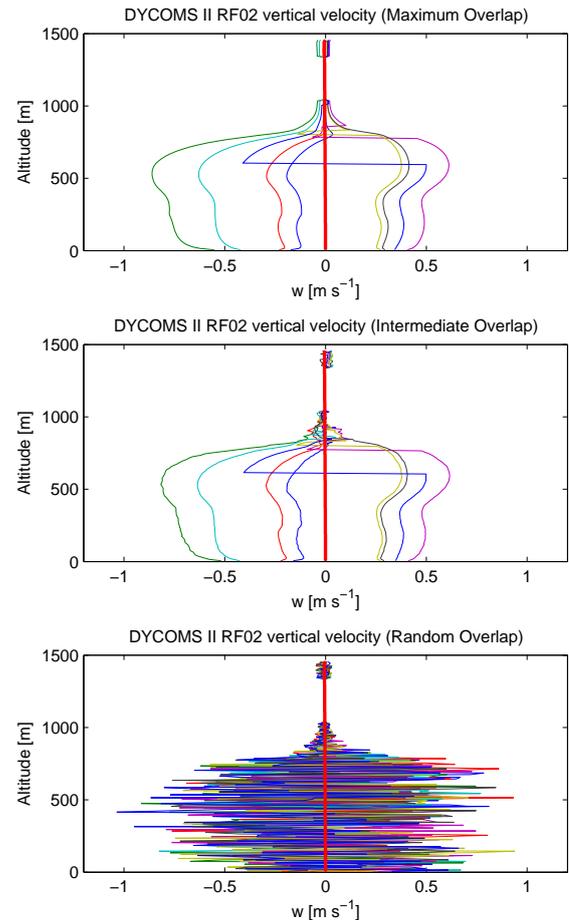


**Fig. 8.** A sample of eight equally weighted SILHS-generated profiles of vertical velocity (thin lines) and the analytic solution (thick red line) drawn from a single time step during the RICO shallow cumulus simulation. The vertical correlation within each profile is specified to be maximal (upper panel), intermediate (middle panel), or random (lower panel). In the maximal and intermediate overlap cases, a profile may sometimes jump suddenly in the vertical if the sampler switches from one normal component to the other. (For random overlap, even if the component remains the same, jumps can occur due to random noise.)

sampling because it uses Latin hypercube sampling to help spread out the sample points both within cloud and within clear air.

Inherent in Monte Carlo methods, such as SILHS, is statistical noise due to small sample sizes. Although most aspects of the two simulations we present are not significantly degraded by sampling noise, this conclusion is based solely on the cases that we tested and has the following caveats:

- This paper presents tests from a cumulus and stratocumulus case that contain weak to moderate drizzle. It is unclear whether the sampling noise will be tolerable under a wider variety of meteorological conditions,



**Fig. 9.** As in Fig. 8, but for the DYCOMS-II RF02 marine stratocumulus simulation.

particularly if we use only 2 sample points per grid box and time step. For instance, sampling noise may be greater in deep convective cases, which precipitate more heavily.

- Our tests use a time step of 5 min. If sample points are generated less frequently, then the time-averaged sampling noise is expected to be greater.
- The simulations that we present are one-dimensional. In three-dimensional simulations, however, the sampling noise may conceivably propagate in undesirable ways. Three-dimensional tests have been performed for the McICA radiative sampler (Pincus et al., 2003; Räisänen et al., 2005; Pincus et al., 2006; Räisänen et al., 2007; Barker et al., 2008; Räisänen et al., 2008), and the results are encouraging, but radiation is usually a slower process than microphysics, and McICA uses dozens or hundreds of subcolumns, instead of the two that we use.

In conclusion, we now list several research topics that may be worth exploring in future work:

1. In our tests, the degree of correlation in the vertical was chosen somewhat arbitrarily. In the future, the vertical correlation methodology should be tested and refined with the aid of observations or fine-resolution 3-D reference simulations.
2. Although the present paper assumes that the correlations are prescribed and constant, in nature the correlations vary in space and time. In future work, the correlations could be diagnosed using a method such as that of Larson et al. (2011).
3. Microphysical processes influence not only the mean of the PDF but also the higher-order moments. However, in this first implementation, the microphysical tendencies from SILHS directly influence only the grid-box means. In future work, the tendencies from SILHS could be made to interact with the higher-order moments.
4. Although SILHS is expensive, its computational burden may be blunted by the use of massively parallel computers. It is true that there is computational overhead associated with choosing multivariate sample points, a computation that is not embarrassingly parallel. However, computing microphysics and other physical processes in multiple subcolumns may be a suitable problem for parallelization because, in our formulation, the subcolumns do not communicate information with each other. Thus, for calculations whose dynamics cannot efficiently use all available processors, the unused processors could be used to compute physics in subcolumns.
5. In order to produce the most accurate integration feasible, SILHS attempts to *reduce* sampling noise. On the other hand, some stochastic sampling techniques deliberately *add* noise to ensemble weather forecasts in order to increase the spread of the ensemble (e.g., Buizza et al., 1999; Teixeira and Reynolds, 2008; Berner et al., 2008; Palmer and Williams, 2008). Adding noise has been shown to increase the ensemble spread, especially if the noise is correlated between grid columns. Other studies have added noise for the purpose of improving the mean climatology of climate models (e.g., Lin and Neelin, 2000; Berner et al., 2008, 2012). For such applications, accurate forecasting may not require further reduction of the noise left by SILHS. That being said, there is no reason to suppose that the noise produced by SILHS represents model error.
6. Although this paper has focused on the problem of treating subgrid variability in microphysical calculations,

the subcolumn methodology could, in principle, be applied more widely. For example, in future work, one could envision feeding subcolumns into parameterizations of microphysics, radiative transfer, aerosol physics, atmospheric chemistry, and possibly other physical processes. The use of subcolumns or similar quadrature techniques could help the climate modeling community transition from the development of cloud parameterizations, which are limited in scope, to the development of a more general unified parameterization of subgrid variability.

*Acknowledgements.* The authors are grateful for financial support under grant 0968640 from the National Science Foundation and grant DE-SC0006927 from the SciDAC program of the US Department of Energy. The manuscript benefited from helpful comments by Rachel Storer, Minghuai Wang, Peter Caldwell, Petri Räisänen, and an anonymous reviewer. We thank Eric Raut for preparing some of the figures.

Edited by: J. Annan

## References

- Barker, H.: Representing cloud overlap with an effective decorrelation length: An assessment using CloudSat and CALIPSO data, *J. Geophys. Res.*, 113, D24205, doi:10.1029/2008JD010391, 2008.
- Barker, H. W., Pincus, R., and Morcrette, J.-J.: The Monte Carlo Independent Column Approximation: Application within large-scale models, in: Proceedings of the GCSS workshop, Kananaskis, Alberta, Canada, Amer. Meteor. Soc., 2002.
- Barker, H. W., Cole, J. N. S., Morcrette, J.-J., Pincus, R., Räisänen, P., von Salzen, K., and Vaillancourt, P. A.: The Monte Carlo Independent Column Approximation: An Assessment using Several Global Atmospheric Models, *Q. J. Roy. Meteor. Soc.*, 134, 1463–1478, 2008.
- Bergman, J. W. and Rasch, P. J.: Parameterizing vertically coherent cloud distributions, *J. Atmos. Sci.*, 59, 2165–2182, 2002.
- Berner, J., Doblas-Reyes, F. J., Palmer, T. N., Shutts, G., and Weisheimer, A.: Impact of a quasi-stochastic cellular automaton backscatter scheme on the systematic error and seasonal prediction skill of a global climate model, *Phil. Trans. R. Soc. A*, 366, 2561–2579, doi:10.1098/rsta.2008.0033, 2008.
- Berner, J., Jung, T., and Palmer, T. N.: Systematic model error: the impact of increased horizontal resolution versus improved stochastic and deterministic parameterizations., *J. Climate*, 25, 4946–4962, doi:10.1175/JCLI-D-11-00297.1, 2012.
- Bogenschutz, P. A., Gettelman, A., Morrison, H., Larson, V. E., Craig, C., and Schanen, D. P.: Higher-Order Turbulence Closure and Its Impact on Climate Simulations in the Community Atmosphere Model, *J. Climate*, doi:10.1175/JCLI-D-13-00075.1, 2013.
- Bony, S. and Emanuel, K. A.: A parameterization of the cloudiness associated with cumulus convection; Evaluation using TOGA COARE data, *J. Atmos. Sci.*, 58, 3158–3183, 2001.

- Buizza, R., Miller, M., and Palmer, T. N.: Stochastic representation of model uncertainties in the ECMWF Ensemble Prediction System, *Q. J. Roy. Meteor. Soc.*, 125, 2887–2908, 1999.
- Cahalan, R. F., Ridgway, W., Wiscombe, W. J., Bell, T. L., and Snider, J. K.: The albedo of fractal stratocumulus clouds, *J. Atmos. Sci.*, 51, 2434–2455, 1994.
- Cheng, A. and Xu, K.-M.: A PDF-Based Microphysics Parameterization for Simulation of Drizzling Boundary Layer Clouds, *J. Atmos. Sci.*, 66, 2317–2334, 2009.
- Garvey, P. R.: Probability methods for cost uncertainty analysis, Marcel Dekker, 2000.
- Gentle, J. E.: Random Number Generation and Monte Carlo Methods, Springer, 2nd edn., 2003.
- Golaz, J.-C., Larson, V. E., and Cotton, W. R.: A PDF-based model for boundary layer clouds. Part I: Method and model description, *J. Atmos. Sci.*, 59, 3540–3551, 2002.
- Griffin, B. M. and Larson, V. E.: Analytic upscaling of local microphysics parameterizations, Part II: Simulations, *Q. J. Roy. Meteorol. Soc.*, 139, 58–69, 2013.
- Hill, P. G., Manners, J., and Petch, J. C.: Reducing noise associated with the Monte Carlo Independent Column Approximation for weather forecasting models, *Q. J. Roy. Meteorol. Soc.*, 137, 219–228, 2011.
- Hogan, R. J. and Illingworth, A. J.: Deriving cloud overlap statistics from radar, *Q. J. Roy. Meteorol. Soc.*, 126, 2903–2909, 2000.
- Johnson, M. E.: Multivariate Statistical Simulation, John Wiley and Sons, 1987.
- Khairoutdinov, M. and Kogan, Y.: A new cloud physics parameterization in a large-eddy simulation model of marine stratocumulus, *Mon. Weather Rev.*, 128, 229–243, 2000.
- Larson, V. E.: From cloud overlap to PDF overlap, *Q. J. Roy. Meteorol. Soc.*, 133, 1877–1891, doi:10.1002/qj.165, 2007.
- Larson, V. E. and Golaz, J.-C.: Using probability density functions to derive consistent closure relationships among higher-order moments, *Mon. Weather Rev.*, 133, 1023–1042, 2005.
- Larson, V. E. and Griffin, B. M.: Analytic upscaling of local microphysics parameterizations, Part I: Derivation, *Q. J. Roy. Meteorol. Soc.*, 139, 46–57, 2013.
- Larson, V. E., Wood, R., Field, P. R., Golaz, J.-C., Vonder Haar, T. H., and Cotton, W. R.: Systematic biases in the microphysics and thermodynamics of numerical models that ignore subgrid-scale variability, *J. Atmos. Sci.*, 58, 1117–1128, 2001.
- Larson, V. E., Golaz, J.-C., Jiang, H., and Cotton, W. R.: Supplying Local Microphysics Parameterizations with Information about Subgrid Variability: Latin Hypercube Sampling, *J. Atmos. Sci.*, 62, 4010–4026, 2005.
- Larson, V. E., Nielsen, B. J., Fan, J., and Ovchinnikov, M.: Parameterizing correlations between hydrometeors in mixed-phase Arctic clouds, *J. Geophys. Res.*, 116, D00T02, doi:10.1029/2010JD015570, 2011.
- Larson, V. E., Schanen, D. P., Wang, M., Ovchinnikov, M., and Ghan, S.: PDF parameterization of boundary layer clouds in models with horizontal grid spacings from 2 to 16 km, *Mon. Weather Rev.*, 140, 285–306, 2012.
- Lewellen, W. S. and Yoh, S.: Binormal model of ensemble partial cloudiness, *J. Atmos. Sci.*, 50, 1228–1237, 1993.
- Lin, J. W. B. and Neelin, J. D.: Influence of a stochastic moist convective parameterization on tropical climate variability, *Geophys. Res. Lett.*, 27, 3691–3694, 2000.
- Matsumoto, M. and Nishimura, T.: Mersenne Twister: A 623-dimensionally equidistributed uniform pseudo-random number generator, *ACM Trans. Model. Comput. Simul.*, 8, 3–30, 1998.
- McKay, M. D., Beckman, R. J., and Conover, W. J.: A comparison of three methods for selecting values of input variables in the analysis of output from a computer code, *Technometrics*, 21, 239–245, 1979.
- Mellor, G. L.: The Gaussian cloud model relations, *J. Atmos. Sci.*, 34, 356–358, 1977.
- Morcrette, J.-J. and Fouquart, Y.: The overlapping of cloud layers in shortwave radiation parameterization, *J. Atmos. Sci.*, 43, 321–328, 1986.
- Morrison, H. and Gettelman, A.: A New Two-Moment Bulk Stratiform Cloud Microphysics Scheme in the Community Atmosphere Model, Version 3 (CAM3). Part I: Description and Numerical Tests, *J. Climate*, 21, 3642–3659, 2008.
- Norris, P. M., Oreopoulos, L., Hou, A. Y., Tao, W.-K., and Zeng, X.: Representation of 3D heterogeneous cloud fields using copulas: Theory for water clouds, *Q. J. Roy. Meteorol. Soc.*, 134, 1843–1864, 2008.
- Oreopoulos, L., Lee, D., Sud, Y. C., and Suarez, M. J.: Radiative impacts of cloud heterogeneity and overlap in an atmospheric General Circulation Model, *Atmos. Chem. Phys.*, 12, 9097–9111, doi:10.5194/acp-12-9097-2012, 2012.
- Owen, A. B.: Quasi-Monte Carlo Techniques, in: *Siggraph 2003*, Course 44, San Diego, CA, Association for Computing Machinery, 2003.
- Palmer, T. N. and Williams, P. D.: Introduction. Stochastic physics and climate modelling, *Phil. Trans. R. Soc. A*, 366, 2421–2427, doi:10.1098/rsta.2008.0059, 2008.
- Pincus, R., Barker, H. W., and Morcrette, J.-J.: A fast, flexible, approximate technique for computing radiative transfer in inhomogeneous cloud fields, *J. Geophys. Res.*, 108, 4376, doi:10.1029/2002JD003322, 2003.
- Pincus, R., Hannay, C., Klein, S. A., Xu, K.-M., and Hemler, R.: Overlap assumptions for assumed probability distribution function cloud schemes in large-scale models, *J. Geophys. Res.*, 110, D15S09, doi:10.1029/2004JD005100, 2005.
- Pincus, R., Hemler, R., and Klein, S. A.: Using stochastically-generated subcolumns to represent cloud structure in a large-scale model, *Mon. Weather Rev.*, 134, 3644–3656, 2006.
- Press, W. H., Teukolsky, S. A., Vetterling, W. T., and Flannery, B. P.: *Numerical Recipes in C: The art of scientific computing*, Cambridge University Press, 2nd Edn., 1992.
- Räisänen, P. and Barker, H. W.: Evaluation and optimization of sampling errors for the Monte Carlo Independent Column Approximation, *Q. J. Roy. Meteorol. Soc.*, 130, 2069–2085, 2004.
- Räisänen, P., Barker, H. W., Khairoutdinov, M. F., Li, J., and Randall, D. A.: Stochastic generation of subgrid-scale cloudy columns for large-scale models, *Q. J. Roy. Meteorol. Soc.*, 130, 2047–2067, 2004.
- Räisänen, P., Barker, H. W., and Cole, J. N. S.: The Monte Carlo Independent Column Approximation's conditional random noise: Impact on simulated climate, *J. Climate*, 18, 4715–4730, doi:10.1175/JCLI3556.1, 2005.
- Räisänen, P., Järvenoja, S., Järvinen, H., Giorgetta, M., Roeckner, E., Jylhä, K., and Ruosteenoja, K.: Tests of Monte Carlo Independent Column Approximation in the ECHAM5 Atmospheric

- GCM, *J. Climate*, 20, 4995–5011, doi:10.1175/JCLI4290.1, 2007.
- Räisänen, P., Järvenoja, S., and Järvinen, H.: Noise due to the Monte Carlo independent-column approximation: short-term and long-term impacts in ECHAM5, *Q. J. Roy. Meteorol. Soc.*, 134, 481–495, 2008.
- Rauber, R. M., Rauber, R. M., Stevens, B., Ochs, H. T., Knight, C., Albrecht, B. A., Blyth, A. M., Fairall, C. W., Jensen, J. B., Lasher-Trapp, S. G., Mayol-Bracero, O. L., Vali, G., Anderson, J. R., Baker, B. A., Bandy, A. R., Burnet, E., Brenguier, J.-L., Brewer, W. A., Brown, P. R. A., Chuang, P., Cotton, W. R., Girolamo, L. D., Geerts, B., Gerber, H., Göke, S., Gomes, L., Heikes, B. G., Hudson, J. G., Kollias, P., Lawson, R. P., Krueger, S. K., Lenschow, D. H., Nuijens, L., O’Sullivan, D. W., Rilling, R. A., Rogers, D. C., Siebesma, A. P., Snodgrass, E., Stith, J. L., Thornton, D. C., Tucker, S., Twohy, C. H., Zuidema, P., Sperber, K. R., and Waliser, D. E.: Rain in Shallow Cumulus Over the Ocean: The RICO Campaign, *Bull. Amer. Meteor. Soc.*, 88, 1912–1928, 2007.
- Shonk, J. K. P. and Hogan, R. J.: Tripleclouds: An efficient method for representing horizontal cloud inhomogeneity in 1D radiation schemes by using three regions at each height, *J. Climate*, 21, 2352–2370, 2008.
- Shonk, J. K. P., Hogan, R. J., Mace, G. G., and Edwards, J. M.: Effect of improving representation of horizontal and vertical cloud structure on the Earth’s global radiation budget, Part I: review and parametrisation, *Q. J. Roy. Meteorol. Soc.*, 136, 1191–1204, 2010.
- Shonk, J. K. P., Hogan, R. J., and Manners, J.: Impact of improved representation of horizontal and vertical cloud structure in a climate model, *Clim. Dynam.*, 38, 2365–2376, 2012.
- Skamarock, W. C., Klemp, J. B., Dudhia, J., Gill, D. O., Barker, D. M., Wang, W., and Powers, J. G.: A description of the Advanced Research WRF Version 2., Tech. Report NCAR TN-468 STR, Natl. Cent. for Atmos. Res., Boulder, CO, 2005.
- Smith, R. N. B.: A scheme for predicting layer clouds and their water content in a general circulation model, *Q. J. Roy. Meteor. Soc.*, 116, 435–460, 1990.
- Sommeria, G. and Deardorff, J. W.: Subgrid-scale condensation in models of nonprecipitating clouds, *J. Atmos. Sci.*, 34, 344–355, 1977.
- Stevens, B., Lenschow, D. H., Vali, G., Gerber, H., Bandy, A., Blomquist, B., Brenguier, J.-L., Bretherton, C. S., Burnet, F., Campos, T., Chai, S., Faloon, I., Friesen, D., Haimov, S., Laursen, K., Lilly, D. K., Loehrer, S. M., Malinowski, S. P., Morley, B., Petters, M. D., Rogers, D. C., Russell, L., Savic-Jovicic, V., Snider, J. R., Straub, D., Szumowski, M. J., Takagi, H., Thornton, D. C., Tschudi, M., Twohy, C., Wetzel, M., and Van Zanten, M. C.: Dynamics and chemistry of marine stratocumulus – DYCOMS-II, *Bull. Amer. Meteor. Soc.*, 84, 579–593, 2003.
- Teixeira, J. and Reynolds, C. A.: Stochastic Nature of Physical Parameterizations in Ensemble Prediction: A Stochastic Convection Approach, *Mon. Weather Rev.*, 136, 483–496, 2008.
- Tompkins, A. M.: 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, 2002.
- vanZanten, M., Stevens, B., Nuijens, L., Siebesma, A., Ackerman, A., Burnet, F., Cheng, A., Couvreur, F., Jiang, H., Khairoutdinov, M., Kogan, Y., Lewellen, D., Mechem, D., Nakamura, K., Noda, A., Shipway, B., Slawinska, J., Wang, S., and Wyszogrodzki, A.: Controls on precipitation and cloudiness in simulations of trade-wind cumulus as observed during RICO, *J. Adv. Model. Earth Syst.*, 3, M06001, doi:10.1029/2011MS000056, 2011.
- Weber, T. and Quaas, J.: Incorporating the subgrid-scale variability of clouds in the autoconversion parameterization using a PDF-scheme, *J. Adv. Model. Earth Syst.*, 4, M11003, doi:10.1029/2012MS000156, 2012.
- Wyant, M. C., Bretherton, C. S., Chlond, A., Griffin, B. M., Kitagawa, H., Lappen, C.-L., Larson, V. E., Lock, A., Park, S., de Roode, S. R., Uchida, J., Zhao, M., and Ackerman, A. S.: A Single-Column-Model Intercomparison of a Heavily Drizzling Stratocumulus Topped Boundary Layer, *J. Geophys. Res.*, 112, D24204, doi:10.1029/2007JD008536, 2007.
- Zhang, J., Lohmann, U., and Lin, B.: A new statistically based autoconversion rate parameterization for use in large-scale models, *J. Geophys. Res.*, 107, 4750, doi:10.1029/2001JD001484, 2002.