A NEW APPROACH TO SIMULATING A REALISTIC SPATIAL STRUCTURE OF CLOUD DROPLETS

Michael L. Larsen, Michigan Technological University, Houghton, MI
(*mllarsen@mtu.edu*)
Mentor: Dr. Alexander Marshak, NASA GSFC, Code 913

Abstract. *Physically motivated theories of cloud evolution and microphysics require an accurate characterization of cloud texture. In particular, the statistical properties of the spatial distribution of cloud particles must be known for a suitably accurate treatment of radiative transfer and droplet size distribution evolution. Most current theories dealing with these phenomena implicitly assume a perfectly random spatial distribution of droplets, despite observations that contradict this assumption. Many different mechanisms for quantifying cloud droplet spatial variability exist in the literature, but nearly all have undesired properties. One of the largely unexamined statistical tools that could be used for quantifying cloud non-uniformity makes use of a fairly recent mathematical subfield often labelled as "fractal mathematics".*

In this report, we discuss a cascading stochastic construction as a possible tool for simulating cloud non-uniformity and its implications. To our knowledge, this is the first simulation method that does not ignore deviations from perfect spatial randomness as a function of droplet size (e.g. particles of larger size are more clustered than particles of small size, which is consistent with the empirical data. Visual inspection reveals that the simulated cloud exhibits striking qualitative similarity with in situ data. The easily constructed nature and self-similar properties of the simulated stochastic distributions provide us with a convenient framework for studying realistic 3-D cloud radiative transfer with Monte Carlo techniques.

INTRODUCTION

The spatial distribution of cloud particles has been a fairly active research area for the last few decades. It has now been widely accepted that departures from perfect spatial randomness are ubiquitous among cloud particles. Nevertheless, the magnitude, scale, and physical nature of the non-uniformity in droplet spatial positions are very actively debated. That these departures from randomness can significantly influence microphysical processes has been firmly established; for example (Kasper, 1984; Larsen et al., 2003) discussed implications on the theory of coagulation, (Borovoi, 1984, 2002; Davis and Marshak, 1997; Knyazikhin et al., 2002; Kostinski, 2001, 2002; Marshak et al., 1998; Shaw et al., 2002) and many others have dealt with radiative transfer through a non-uniform random medium, and (Kostinski et al., 2003) has discussed the stability of a so-called droplet size distribution equilibrium in light of these departures. In order to predict the magnitude of these influences, all of these treatments require some quantitative measure of how far the true distribution resides from perfect randomness.

There is no shortage of methods used to identify this departure from pure randomness; common methods include using the curvature of the Power Spectral Density of concentration measurements (e.g. (Pinsky and Khain, 2001)), the so-called coefficient of variation (which is essentially the same as the clustering index or Fishing statistic) (e.g. (Baker, 1992; Borrmann et al., 1993; Chaumat and Brenguier, 2001) and many others), the auto-correlation function of concentration measurements (too many authors to mention), the pair-correlation function and its deviates (e.g. (Kostinski and Jameson, 2000; Kostinski and Shaw, 2001)), and fractal dimension (e.g. (Davis et al., 1999; Wiscombe et al., 2003)).

A careful examination, however, reveals that since the power spectral density and autocorrelation function can only be applied to continuous measurements, their use for detecting spatial structure of cloud particles (discrete entities) is dubious at best. Additionally, the Fishing statistic, Power Spectral Density, and some of the deviates of the pair-correlation function have scale-memory and hence do not give a scale-localized description of departures from pure randomness - a necessary requirement in calculating the effects on most of the microphysical processes discussed above (see (Shaw et al., 2002) for an explanation and a more extensive review and commentary on the cloud particle clustering literature). Some insight would doubtlessly be obtained if we had certain knowledge of the cause of the departures from pure randomness. Unfortunately, there is no consensus about this, either. Some investigators (e.g. (Shaw et al., 1998)) have suggested that cloud particles cluster preferentially in regions of low vorticity -avoiding the high-strain regions created by high vorticity. Other studies have suggested that the clustering is likely an inertial response to the superimposed flow-field. Finally, other scientists seem to advance the hypothesis that there are regions of the sky that - due to the micro-meteorological state - are favorable for producing cloud drops and that advection and turbulent diffusion are unable to eliminate the clustering effect caused by the localized \birth" process. Little formal work has been done along these lines to date, so a definitive answer to the question of the origins of the departures from pure randomness may take quite some time. Despite this, however, we conclude that since deviations from pure randomness exist for the aerosols that form cloud particles (Larsen et al., 2003; Priening, 1983) and the raindrops that some of them become (Jameson and Kostinski, 1999a and b, 2000; Kostinski and Jameson, 1997, 1999; Kostinski et al., 2003), the clustering process is not likely due to an effect that only influences cloud drops of a specific size, mass, or composition. Feasible theories regarding the origin of this spatial clustering should be forthcoming as the studies in these cognate fields further refine their measurements of deviations from pure randomness.

The existence of deviations from pure randomness over a broad region of masses, size scales, compositions, and other physical properties entices one to hypothesize there may be some degree of scale-invariance in the governing dynamics. Some scientists that use fractal and multifractal methods to analyze these systems argue that if there is some semblance of a self-similar mechanism causing the clustering, it may be prudent to expect the statistics of the particle positions themselves to have fractal properties. Though the author has previously used the pair-correlation function almost exclusively to detect and classify “correlations” (or clustering), the potentially self-similar character discussed above has motivated this report. The next 2 sections of this report is devoted to giving a very brief overview of the terminology and concepts behind fractal analysis, which will be utilized later in the report, but can safely be skipped by those with a solid working knowledge of fractals and related mathematical constructs. These sections are followed by an analysis of some *in situ* data collected during an ARM intensive field campaign. After demonstrating that there is some evidence for a fractal structure in this data, there is a tutorial section on how to construct a multiplicative multifractal cascaded system in a 1, 2, or 3-dimensional embedding dimension and qualitative comparison of the construction to the ARM intensive data. Finally, in the final section we will discuss work that we intend to carry out after the summer program at NASA has terminated - a Monte-Carlo analysis and discussion of the radiation transmission through the simulated medium.

Elementary introduction to fractal analysis

Since the majority of people that will read this report will either (i) know more about fractal analysis than the author, or (ii) would be bored with a long formalistic tract, this introduction will be exceedingly informal and makes no claims to completeness or any level of rigor. Nearly all of the following explanation is a digested and rephrased version of information that was distilled from Knyazikhin (2003; personal communication); Marhsak (2003; personal communication); Peitgen et al., 1992 and, to a lesser extent, from Harte, 2001; Stoyan and Stoyan, 1994; Vicsek, 1989). Despite the extensive aid received, the responsibility for any errors of omission or misinterpretation remains that of the author.

A quick glance through any introductory text on fractal mathematics will quickly assert that the fundamental property of a fractal object is some level of self-similarity. Loosely, this means that certain properties of the entire system are retained if one examines successively smaller subsets of the object (the properties are independent of spatial scale, or scale-invariant). This is in contrast to more familiar experience which suggests that examining smaller and smaller scales would smooth out sharp gradients and eventually render everything uniform or homogeneous. In the case of deterministic fractals, the self-similar property often indicates that the system is made up of smaller copies of itself (e.g. see Figure 1). Because of the existence of these gradient structures that do not smooth with scale, fractals often appear “rough” or “jagged”, despite deep symmetries and similarities in their constructive structure. There are many other interesting and useful properties of deterministic fractals, but they are beyond the scope of this report.

It probably comes as no surprise that the self-similar property of a fractal system carries with it a great deal of information. In fact, if we have a self-similar figure, we can infer all essential information about the system given any arbitrarily small finite convex subset of the system. This is a truly novel idea, for if we invoke Shannon's information theoretic formalism, this implies that all of the information in a system is in each part, no matter how small. Perhaps the most accessible way to interpret the notion of fractal dimension1 is that it represents the amount of independent coordinates needed to span the space of particle positions. The notion that this quantity could be a fractional number is reminiscent of information theory where a signal can have a non-integer number of bits of information; there are several ways to visualize such a notion but describing them would take us on a lengthy tangent.

For our purposes, we will be working with random or stochastic fractals. Here, following (Vicsek, 1989), we will call stochastic fractals scale invariant instead of self similar to emphasize that only the statistical properties of the system are scale-invariant. (One could not expand part of the system and end up with a copy of the system as a whole, except in a statistical sense). It is the notion of statistical invariance that we will use in the construction of our model. In essence, our model uses a stochastic algorithm to generate a system with specified fractal and topological dimensions, number of particles in the system, and geometrical extent (size) of the system.

The Random Fractal Structure of the Cloud Model

At this point, it may not yet be apparent how the notions of fractal geometry and deviations from pure randomness of cloud particle positions could be related.

Empirically, we observe that systems with a smaller fractal dimension look more clustered (e.g. see Figure 2). It certainly seems like we should be able to use fractal dimension as a proxy for the deviations from pure randomness - but why is this the case?

Although I do not currently have a wholly satisfactory answer to that question in a formal sense, we can examine the connection by looking at the notion of a "concentration" (or density) (Davis (2003; personal communication); Knyazikhin (personal communication); Marhsak (personal communication); Wiscombe et al., 2003). At the risk of insulting the reader, the basic notion of a concentration or density goes back to the existence of N particles in a volume V . One assumes that the volume chosen was representative of the sample as a whole and the ratio N/V can be used as an estimator for the concentration of the distribution. In principle, if one could carry out many concentration measurements, the limit of large numbers would permit the ensemble of inferred sample concentrations to converge to a Gaussian distribution in which the average would give the "true" value. Taking this ensemble-estimated average concentration and multiplying by any volume you desire, then, should yield an unbiased estimator for the number of particles in the volume.

In a fractal system, however, there exists a range of scales where the expected number of droplets in a volume does not scale as $N \propto V$. For this range of scales, $N \propto V^{D/d}$ where D and d are the fractal and topological dimensions of the system, respectively. We call this system scale-invariant because, though N and V are not proportional, a power law scaling is expected to be constant through the pertinent range. It is in this manner that we can begin to interpret a lower dimension as corresponding to a more clustered distribution.6 If $N = cV^{D/d}$ we classically compute $dN/dV = cD/d\ (V^{D/d-1})$, which has a concave up structure, blowing up to infinity at V = 0 and decaying to zero at V = ∞. Since this is an everywhere decaying curve, if one chooses an arbitrary spot (origin) in the region of interest, the probability of particles being a specific distance from that position decreases as the distance from that position increases. From the definition of "probability", particles tend to be placed where they are probable to be positioned - hence close to the origin. Since they tend to be close to the chosen point, they also tend to be close to each other (i.e. clustered). If D/d is lower, the curve has a steeper slope and the effect described here is larger, hence an even larger tendency for particles to be close to the chosen point and clustered. The final inference, then, is that as D/d decreases, clustering increases. Since d is fixed for any system, lowering the fractal dimension D is the way to obtain greater clustering in a fractal system.

Aside on Droplet Size Distribution Structure

Although the introduction of the notion of fractal dimension into the atmospheric sciences is relatively new, it is not without several notable precursors. Of special interest is a series of papers written by Liu and Hallett (Liu et al., 2002; Liu and Hallett, 1997, 1998), which examine the size distribution structure using Bayesian estimation methods.

An oft-used concept in various areas of atmospheric science is the notion of a drop size distribution (DSD) (this is called a droplet size distribution for cloud particles, a size distribution for aerosol particles, or a raindrop size distribution for rain). The DSD can be written --depending on the application one has in mind - in terms of (i) the number of drops of a specific size relative to the number of drops of all sizes

$$\left(\text{e.g. } \frac{[n(r, r+\delta r)]\delta r}{\int_0^\infty n(r)dr}\right),$$

or (ii) the number of particles of a specific size per unit volume. It is the latter of these that relies on the classical notion of concentration as defined earlier and has fallen under previous scrutiny.

In (Liu et al., 2002; Liu and Hallett, 1997, 1998), a theoretical statistical model was used to examine collected DSD data from a cloud scope operating on an airplane in Louisiana. The theory suggests (and the data seems to support) the claim that the DSD shape is strongly dependent on the size of the volume used to calculate it. More specifically, the DSD is shown to continually broaden with increasing sampling size until a characteristic "saturation volume" is reached beyond which little fluctuation occurs.

The notion that the "number of particles of a specific size per unit volume" depends on the volume used to measure the quantity *is not commensurate* with the classical notion of concentration as discussed above. No longer could one take an ensemble of equivalent volumes and determine, through use of the limit of large numbers, the "true" parameters of the DSD since the properties of the agglomerated distribution are not the same as those of the collection of samples. This notion is counter-intuitive, and the fallacy relies on the ultimate structure (or texture) of the collection of particles.

To put a similar observation in a different way, there are three fundamental spatial scales that must be widely separated for concentration as traditionally defined to be a meaningful quantity - the mean inter-particle distance, the scale on which concentration fluctuations, and the scale on which the concentration is measured (Kostinski et al., 2003). When these scales are not separated, one must take care to define what one means by "concentration".

There are at least two ways out of this conflict, however. In earlier work, some investigators (including the author) have decided to sidestep this issue by using the other definition of a DSD and using this alternative notion to obtain information about the spatial distribution and collective properties of the droplets (Kostinski and Jameson, 1999; Kostinski et al., 2003; Shaw et al., 2002). Another way out of the conundrum, that used here, requires one to apply some of the tools of fractal analysis to get a "scaling theory of concentration" and, ultimately, a volume dependent DSD from the data at hand.

On Assigning a Fractal Dimension to a System

Before launching head-long into the simulation, however, an explanation as to how fractal dimension is estimated is in order. For deterministic fractals, there often is an analytic result that gives the fractal dimensions of the system. For stochastic fractals, the method commonly referred to as "box-counting" gives a fairly simple method of estimating the fractal dimension. (Other explanations of this method can be found in nearly any text or article that relies on numerical computation of a fractal dimension. Of particular interest to the author are Davis (2003; personal communication); Knyazikhin (2003; personal communication); Lovejoy, 1982; Marhsak (2003; personal communication); Peitgen et al., 1992). As noted by Davis (2003; personal communication), when the fractal dimension is less than the topological dimension of a system, the points making up the system are characterized by a hierarchy of holes or gaps. If the system of interest is broken into boxes of characteristic size ² (so that a small volume is size ²d), it turns out that the number of boxes needed to \cover" the set of droplet positions (how many boxes that are not "holes") goes as:

$$N(\in) \propto \in^{-D} = V^{-D/d}. \quad (1)$$

What is meant by "cover" is best illustrated with a simple example. Figure 3 displays a distribution of 80 points distributed perfectly randomly inside a square (d = 2). At each grid resolution scale ($\in$), the number of squares that contain at least one particle are counted and recorded as "N" for that ϵ. When plotted on a log-log graph, one can infer D for this system by looking at the slope of the resulting line since $\log N(\epsilon) = -D\log^{\epsilon}$ (see Figure 4). If the grid boxes that are counted are shaded in, then every point is "covered" by the shading.

The Fractal Properties of ARM Data

In an earlier section, it was argued that there may be some degree of scale-invariance in clouds. This argument, however, was due to a fairly tenuous inference from the clustered behavior of aerosols and precipitation. There have been many studies (e.g. (Chaumat and Brenguier, 2001; Kostinski and Shaw, 2001; Pinsky and Khain, 2001; Shaw et al., 2002, 1998)) that conclude from simulations or empirical measurements that cloud particles are clustered. However, if the attempt will be made to capture this clustering by specifying the fractal dimension it stands to reason that the fractal dimension of cloud particles should be measured and used to try and generate distributions that are statistically similar to the cloud.

Data from the ARM Spring 2000 SGP Cloud IOP was downloaded8 to try and determine the fractal dimension for environmental clouds. The instrument used was a FSSP (Forward Scattering Spectrometer Probe) operating at 1 Hz and detecting particles in 15 size bins, between about 5 and 60 microns in diameter. For the purposes of this paper, the flights on March 3rd and 17th will be analyzed. (I was pleased to find out that several other investigators have independently chosen those dates as looking particularly steady.) An example of the data in raw form is given in Figure 5. A quick glance at the data is likely sufficient to convince most readers that there is significantly less clustering in the upper panel than the lower panel.

When (Knyazikhin et al., 2002) carried out a similar analysis on a subset of the data, a plot similar to that shown in the left side of Figure 6 was utilized. This way of displaying the data enables one to see the intermittency in the different size bins and get an intuitive feel for the clustering that occurs.

Now that we've identified that clustering exists in this system, however, we have to determine whether we obtain a reliable fractal dimension of the system. Figure 7 shows an example of how this information was extracted from the data, and Figure 8 and table 1 summarize the calculated dimensions for both flight dates.

In an earlier section, it was mentioned that by using a fractal analysis method, a "scaling theory of concentration" could be derived to come up with a modified DSD. For the volume defined by a sub-sample of the FSSP in the ARM-IOP, one can actually carry out this analysis and the results are shown in Figure 9; note the significant change in the large droplet sizes.

Simulating a distribution with a nontrivial fractal dimension

To really understand the theoretical basis for the algorithm developed, a fairly lengthy foray into some mathematical topics that may be foreign to many of the readers would be required. Rather than handle this formally, a phenomenological explanation follows.

There are basically two classes of stochastic particle distributions that are fractal in character -- either you can build the scale-invariant property into the particle distribution itself, or you can make a fractal *support* for the system. Roughly, a point x is a member of the support set X for random process y if and only if $p_y(x)$ - the probability that the member of a realization of y exists at x - is greater than zero (e.g. the support is the region in which there is nonzero probability of finding something). In the case where this support itself has fractal structure, a realization of a random process carried out on this support has fractal characteristics itself and is typically called a "multi-fractal" process or, more precisely, a process on a multi-fractal measure.

The specific model that we implemented was based on work that my mentor and collaborators did some time ago when trying to examine bounded cascade models (Marshak et al., 1994). Rather than talk about the mathematical properties in detail, it is easier to visualize the construction process itself (see Figures 10 and 11 and their captions for the step-by-step construction technique). For simplicity, just a one-dimensional cascade is shown. In 2d, the similar graphs have to be three dimensional (two directional dimensions and a dimension to show the magnitude at that (*x,y*) coordinate; see Figure 12.) In 3D, no convenient easy visualization exists since it would require a 4 dimensional Figure on a 2 dimensional piece of paper.

As mentioned in the figure caption, we can control the resulting fractal dimension of the system by increasing or decreasing the assigned \threshold" level. In reality, data analysis reveals that small droplets are distributed in a manner very close to Poisson. Large droplets are very rare in space and we see, as expected, that their respective fractal dimension is nearly zero. (This is simulated by just putting a very few particles into the distribution at randomly. Such particles end up behaving like delta-functions in any subsequent theoretical analysis). For intermediate sized droplets (i.e. 30μm to 100μm), we see a dimension that is close to neither 0 nor 1. The hope is that, in using the multiplicative multi-fractal cascade to simulate their respective positions, we end up generating a realistic distribution. (A comparison between drops distributed in 3d perfectly randomly and with a fractal dimension of 1/2 is shown in Figure 14.) In a paper my Mentor and collaborators are currently working on, it is shown that the existence of particles in this intermediate clustering regime has a very interesting influence on the radiative properties of the system.

An example of a pseudo-realistic composite system like that described above is shown in Figure 15. Therein, a large number of (small) particles (colored black) are distributed perfectly randomly while a smaller number of intermediate sized (red) particles are distributed with fractal dimension near 1.5 throughout the volume (e.g. D/d = 1/2). Note that this system (i) has the same dimension for the two simulated size intervals as the raw data, and (ii) there is qualitative similarity in the intermittency of the measured FSSP data and that simulated for both size ranges. Because of this, we suggest that careful study of this simulated system (where we know the exact positions of each drops and hence can do exact theoretical analysis on radiative transfer, collision/coalescence, drop size distribution evolution, etc.) could potentially yield rough quantitative insight into the effects of the three dimensional correlated structure of cloud particles.

DISCUSSION

It seems justifiable to ask why we introduced the fractal formalism for the reconstruction method discussed above; wouldn't one be better to make use of a method that uses more familiar tools? One finds that most other tools used for detecting deviations from pure randomness are ill-suited for a reconstruction of this type due to their scale-memory. Fractal structures, of course, do not fall to this criticism because they are inherently scale-invariant. The only other statistic that may yield similar sorts of insight is the pair-correlation function. However, using the pair-correlation function in this context may be troubling because taking 1-dimensional data and inverting it to three dimensions has proved a daunting task. Analytical results now exist along these lines that could aid with the inversion (Hotzer and

Collins, 2002), but they are pretty complicated to work with. The fractal approach, with the power-law pair correlation function that they are associated with, just happen to reduce to a form in which the "reduced" fractal dimension does not change no matter which dimension the system is sampled in (see Figure 16.)

This report has shown that we now have a way to simulate cloud particle positions that retain at least some of the statistical properties that we observe in an in-situ environment. As mentioned in the introduction, this construction was motivated to try and develop a more accurate treatment of the microphysics of cloud droplet interactions - in particular, radiative transfer.11 There have been many studies that have investigated the effects of correlations or clustering on three-dimensional radiative transfer. The conclusion that has been drawn by the community is that the Beer-Lambert law is only strictly obeyed in a perfectly random distribution; attenuation with depth is always faster than exponential in a positively correlated medium. A number of numerical simulations with prescribed pair-correlation functions have verified this, but the magnitude of the deviations in nature have been difficult to classify due to the uncertain magnitude of correlations and the lack of an analytical connection between correlation magnitude and deviations from the exponential decay law. The latter of these obstacles will be removed with a manuscript that my mentor is currently assembling, and the former is an extension of the work completed here that I hope to pursue in the coming months. In particular, we hope to develop a simple Monte-Carlo "shadowing" code, very similar to that in (Shaw et al., 2002), to determine what corrections need to be accounted for in a realistic three-dimensional cloud.

It also bears mentioning that, even though the correlations may be mild in nature, a small correction to the radiative properties of clouds could have a very significant effect on the modelling and climate communities. Over recent years, an extraordinary collection of papers have been written on the indirect effect and other climate forcing mechanisms that yield changes on the order of a watt per square meter. Even if this work in clouds yields a very small change in expected extinction (say, a few percent), the effect of clouds on climate is so large that the in°uence of non-exponential attenuation could be as large or larger than these other effects. Finally, even though there may not be any obvious applications for this knowledge, a lot of people in the atmospheric physics community would like to find out what the process is that causes these correlations in the first place. Perhaps, if we determine that a scale-invariant treatment yields the most accurate and reliable predictions about real clouds, we can say something about the mechanism that clustered them in the first place.

CONCLUSION

During this summer program, my mentor and I were able to generate a realistic simulation of cloud particle spatial structure that retains some of the measured characteristics of cloud texture. We are hopeful that by examining our simulated system we can wean some information about the characteristics of real clouds. At this point, we have made a solid step in the direction of handling radiative transfer in clouds more accurately, but much more work needs to be done.

ACKNOWLEDGEMENTS

I would like to thank A. Marshak for helping me become familiar with a totally foreign body of literature and patiently answering all of my annoying questions. Further, I would like to thank my advisor,A. Kostinski, for giving me the opportunity to come down and do this work this summer, J. Kwiatkowski, V. Martins, B. Menighini, and A. Tokay for interesting discussions on unrelated research projects I worked on while here, and A. Brooks, M. Delaney, and the rest of the UMBC/GEST staff that made the summer much more enjoyable than it had to be. My work is supported by the GEST/UMBC Graduate Summer Student Program, NSF Grant Number ATM-0106271, and the Michigan Space Grant Consortium.

REFERENCES

Baker, B. (1992). Turbulent entrainment and mixing in clouds: A new observational approach. Journal of the Atmospheric Sciences 49, 387- 404.Borovoi, A. (1984). Radiative transfer in inhomogeneous media. Dok. Akad. Nauk SSSR 276, 1374-1378.

Borovoi, A. (2002). On the extinction of radiation by a homogeneous but spatially correlated random medium: Comment. Journal of the Optical Society of America, 19, 2518-2520.

Borrmann, S., R. Jaenicke, and P. Neumann (1993). On spatial distributions and inter-droplet distances measured in stratus clouds with in-line holography. Atmospheric Research 29, 229-245.

Chaumat, L. and J. Brenguier (2001). Droplet spectra broadening in cumulous clouds. part ii: Microscale droplet concentration heterogeneities. Journal of the Atmospheric Sciences 58, 642-654.

Davis, A. and A. Marshak (1997). Levy kinetics in slab geometry: Scaling of transmission probability. In M. Nowak and T. Dewey (Eds.), Fractal Frontiers, Singapore, pp. 63-72. World Scientific.

Davis, A., A. Marshak, H. Gerber, and W. Wiscombe (1999). Horizontal structure of marine boundary layer clouds from centimeter to kilometer scales. Journal of Geophysics Research 104, 6123-6144.

Harte, D. (2001). Multifractals: Theory and Applications. Boca Raton, Florida: Chapman and Hall/CRC Press.

Hotzer, G. and L. Collins (2002). Relationship between the intrinsic radial distribution function for an isotropic field of particles and lower-dimensional measurements. Journal of Fluid Mechanics 459, 93.

Jameson, A. and A. Kostinski (1999a). Fluctuation properties of precipitation. part iv: Fine scale clustering of drops in variable rain. Journal of the Atmospheric Sciences 56, 82-91.

Jameson, A. and A. Kostinski (1999b). Fluctuation properties of precipitation. part v: On the distribution of rain rates - theory and observations in clustered rain. Journal of the Atmospheric Sciences 56, 3920-3932.

Jameson, A. and A. Kostinski (2000). Fluctuation properties of precipitation. part vi: Observations of hyperfine clustering and drop size distribution structures in three-dimensional rain. Journal of the Atmospheric Sciences 57, 373-388.

Kasper, G. (1984). On the coagulation rate of aerosols with spatially inhomogeneous particle concentrations. Journal of Colloid and Interface Science 102, 1929-1933.

Knyazikhin, Y., A. Marshak, W. Wiscombe, J. Martonchick, and R. Myneni (2002). A missing solution to the transport equation and its effect on estimation of cloud absorptive properites. Journal of the Atmospheric Sciences 59, 3572-3585.

Kostinski, A. (2001). On the extinction of radiation by a homogeneous but spatially correlated random medium. Journal of the Optical Society of America, A 18, 1929-1933.

Kostinski, A. (2002). On the extinction of radiation by a homogeneous but spatially correlated random medium: Reply to comment. Journal of the Optical Society of America, A 19, 2521-2525.

Kostinski, A. and A. Jameson (1997). Fluctuation properties of precipitation. part i: On deviations of single size drop counts from the poisson distribution. Journal of the Atmospheric Sciences 54 (17), 2174-2186.

Kostinski, A. and A. Jameson (1999). Fluctuation properties of precipitation. part iii: On the ubiquity and emergence of the exponential size spectra. Journal of the Atmospheric Sciences 56 (1), 111-121.

Kostinski, A. and A. Jameson (2000). On the spatial distribution of cloud particles. Journal of the Atmospheric Sciences 57 (7), 901-915.

Kostinski, A., M. Larsen, and A. Jameson (2003). Exploring the stochastic micro-structure of rain: Scale-dependence of spatial correlations. Journal of Hydrology, Conditionally Accepted.

Kostinski, A. and R. Shaw (2001). Scale-dependent droplet clustering in turbulent clouds. Journal of Fluid Mechanics 434, 389-398.

Larsen, M., W. Cantrell, J. Kannosto, and A. Kostinski (2003). Detection of correlations among aerosol particles. Aerosol Science and Technology 37, 476-485.

Liu, Y., P. Daum, and J. Hallett (2002). A generalized systems theory for the effect of varying fluctuations on cloud droplet size distribution. Journal of the Atmospheric Sciences 59, 2279-2290.

Liu, Y. and J. Hallett (1997). The \1/3" power-law between effective radius and liquid water content. Quarterly Journal of the Royal Meterology Society 123, 1789-1795.

Liu, Y. and J. Hallett (1998). On size distributions of cloud droplets growing by condensation: A new conceptual model. Journal of the Atmospheric Sciences 55, 527-536.

Lovejoy, S. (1982). The area-perimeter realation for rain and clouds. Science 216, 185-187.

Marshak, A., A. Davis, R. Cahalan, and W. Wiscombe (1994). Bounded cascade models as nonstationary multifractals. Physical Review E 49 (1), 55-69.

Marshak, A., A. Davis, W. Wiscombe, and R. Cahalan (1998). Radiative effects of sub-mean-free-path liquid water variability observed in stratiform clouds. Journal of Geophysical Research 103, 9557-19567.

Peitgen, H., H. Juergens, and D. Saupe (1992). Chaos and Fractals: New Frontiers of Science. New York: Springer.

Pinsky, M. and A. Khain (2001). Fine structure of cloud droplet concentation as seen from the fast-fssp measurements. part i: Method of analysis and preliminary results. Journal of Applied Meteorology 40, 1515-1537.

Priening, O. (1983). Optical single-particle counters to obtain the spatial inhomogeneity of particulate clouds. Aerosol Science and Technology 2, 79-90.

Royden, H. (1988). Real Analysis (3rd ed.). Prentice Hall.

Saffman, P. and J. Turner (1956). On the collision of drops in turbulent clouds. Journal of Fluid Mechanics 1, 16-30.

Shaw, R., A. Kostinski, and D. Lanterman (2002). Super-exponential extinction of radiation in a negatively correlated random medium. Journal of Quantitative Spectroscopy and Radiative Transfer 75, 13-20.

Shaw, R., A. Kostinski, and M. Larsen (2002). Towards quantifying droplet clustering in clouds. Quarterly Journal of the Royal Meteorological Society 128, 1043-1057.

Shaw, R., W. Reade, L. Collins, and J. Verlinde (1998). Preferential concentration of cloud droplets by turbulence: Effects on the early evolution of cumulous cloud droplet spectra. Journal of the Atmospheric Sciences 55, 1965-1976.

Stoyan, D. and H. Stoyan (1994). Fractals, Random Shapes, and Point Fields. Great Britain: Wiley.

Vicsek, T. (1989). Fractal Growth Phenomena. New Jersey: World Scientific.

Wiscombe, W., Y. Knyazikhin, and A. Marshak (2003). A new look into the treatment of small-scale drop variability in radiative transfer. ARM Annual Report.

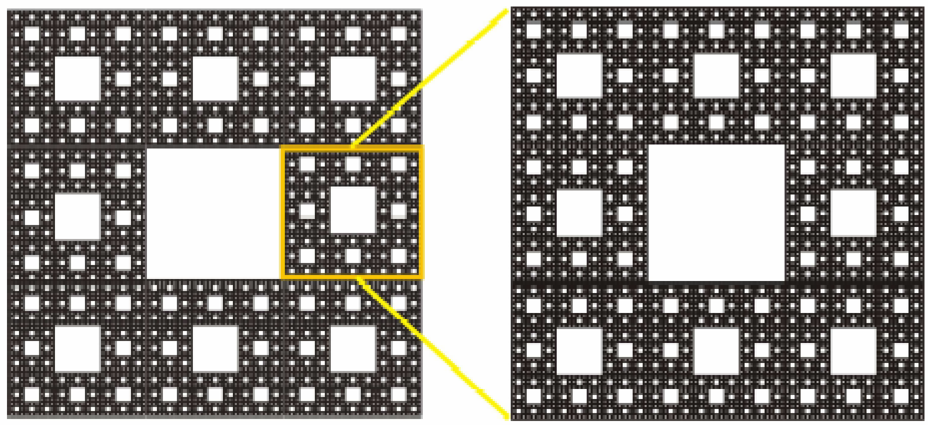

FIG. 1 A Sierpinski Carpet – a famous example of a figure with fractal properties. (Here, of course, the printer can only print with a limited degree of precision, so the figure is merely an approximation of the fractal object.) One generates the fractal in the following way. Start with a pre-shaded 3x3 grid. Erase the shading in the middle square. Now, subdivide each of the remaining squares into 3x3 grids and remove the centers of each. If you repeat the process of subdividing and erasing the middle square an infinite amount of times, the collection of points that remain shaded comprise the Sierpinski Carpet. Note that "zooming in" on a part of the figure, as is done in the right panel, results in an exact copy of the figure as a whole, hence it is self-similar.

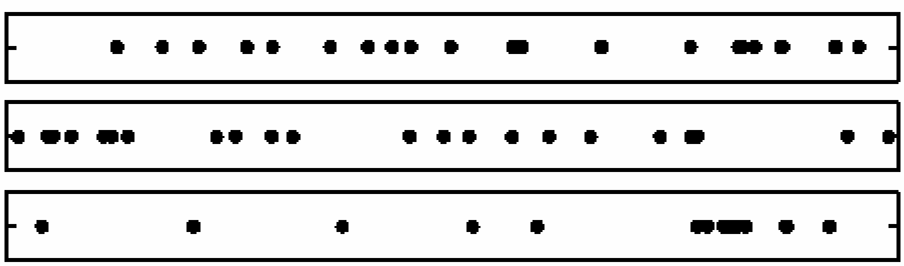

FIG. 2 Distributions in topological space of dimension 1. Each panel has the same amount of particles (22). In the top panel, they are distributed perfectly randomly. As one moves to successively lower panels, the fractal dimension of the system is decreased. We see that this results in more and more clustered distributions. In the limit of infinitely clustered, the entire system could be described by a single point (the position at which all the particles coincide). That would be a zero-dimensional geometric figure and, in fact, such a distribution would be described by a fractal dimension of zero.

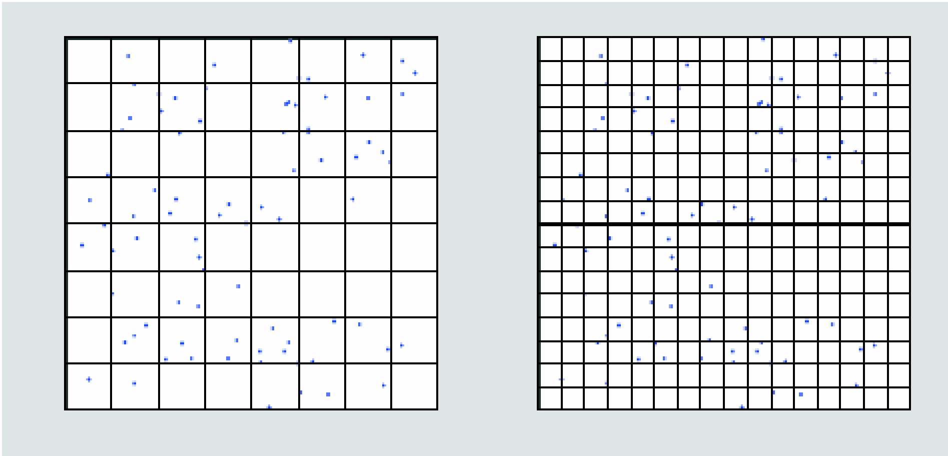

FIG. 3 80 particles distributed randomly in a square. Both panels have the same distribution, but the grid in the left hand box has a grid spacing twice as large as the right box. If one counts the amount of boxes that are occupied in the left hand panel, one finds 43 occupied boxes. In the right hand panel, 73 boxes are occupied. These measurements (and many more like them) are used to develop a figure like figure 4.

FSSP Bin	Diameter (μm)	03/03/00 Dimension	03/17/00 Dimension
1	5.3	0.93	0.94
2	8.6	0.94	0.95
3	11.7	0.93	0.94
4	14.2	0.93	0.93
5	17.0	0.93	0.93
6	20.1	0.93	0.92
7	23.6	0.92	0.91
8	27.2	0.89	0.91
9	31.5	0.78	0.88
10	36.2	0.60	0.84
11	40.6	0.56	0.81
12	45.2	0.56	0.75
13	49.9	0.50	0.71
14	54.7	0.41	0.68
15	59.8	0.39	0.66

TABLE I The data graphed in figure 8; calculated dimensions for the FSSP data-sets from the ARM IOP in March, 2000.

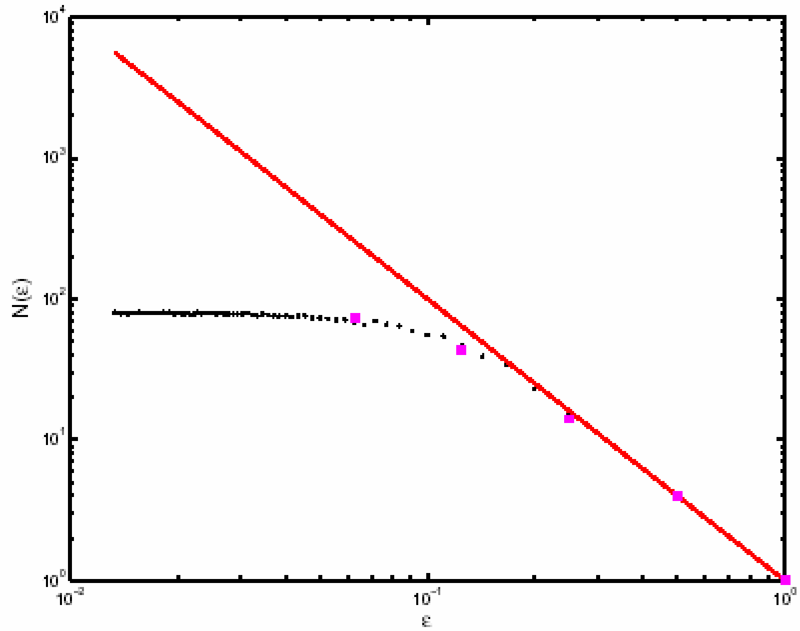

FIG. 4 A log-log graph displaying spatial scale of grid and number of occupied boxes. The red line indicates $N = \epsilon^{-2} = V^{-1} = V^{D/d}$, the upper bound for this curve (if every grid box is occupied). Since the distribution has finite size, at large enough spatial scales the curve must converge to the red line. For ϵ much less than the box size, a horizontal line should result since eventually all particles will be in their own box and the number of boxes needed to cover the distribution will equal the amount of particles in the distribution. In-between these two spatial scales, a linear relationship indicates fractal behavior. In the case of a perfectly random distribution, the intermediate regime and the large scale regimes both behave as ϵ^{-2} (e.g. number in this system *does* scale with volume to the first power, and concentration is a valid notion to use). Because of this, the curve above looks as if there are only two ranges here. The maroon squares are the points obtained from counting in figure 3; the black dots are from an automated counting method developed early on in the project.

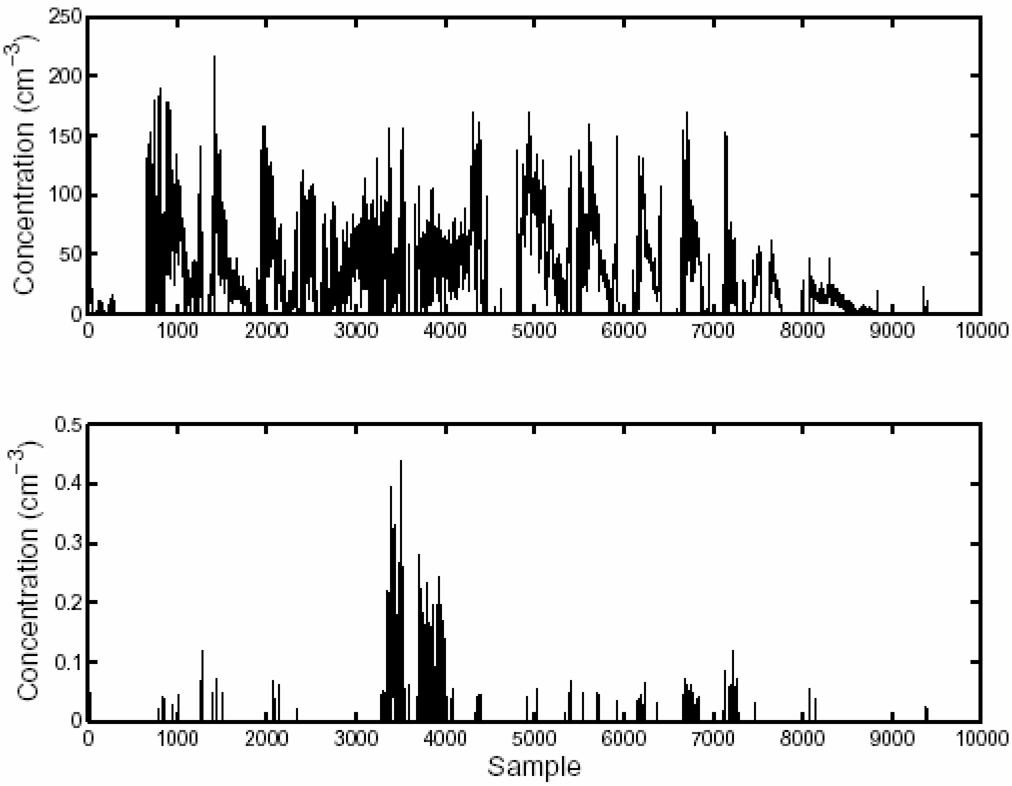

FIG. 5 Raw data from the ARM IOP FSSP. These data were taken on 3/3, from bins 4 (top) and 12 (bottom) of the FSSP. Bin 4 corresponds to particle diameters around 14 μm and bin 12 corresponds to about 45 μm. Analysis reveals that the fractal dimension of the 14 μm bin is near 1 (0.93 – e.g. fractal analysis indicates only a small degree of clustering, which may be due to nothing more than gaps in the cloud-cover). The 45 μm bin gives a fractal dimension of 0.56. The flight was approximately 9600 s (2.67 hours) long, and the FSSP sampled a total volume of about $0.3m^3$ in that time (essentially a tube 770 kilometers long with a cross sectional area of $4 \times 10^{-7}m^2$).

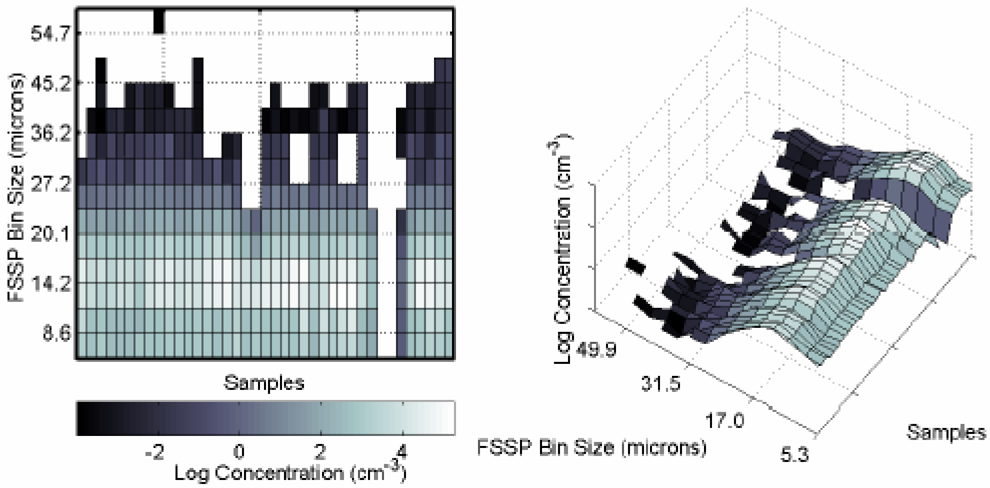

FIG. 6 An insightful way of examining small subsets of the data. The left hand panel displays the FSSP size bins and their respective occupation levels. In the previous analysis by (Knyazikhin et al., 2002) as well as here, all the information used for fractal dimension analysis was whether or not a particular bin has any particles in a particular sample. Another way to visualize the data is shown at right, where the three-dimensional character of the data is a little more apparent. The clustering of large particles can be noted by looking at the relatively large fluctuations between neighboring samples for large size bins.

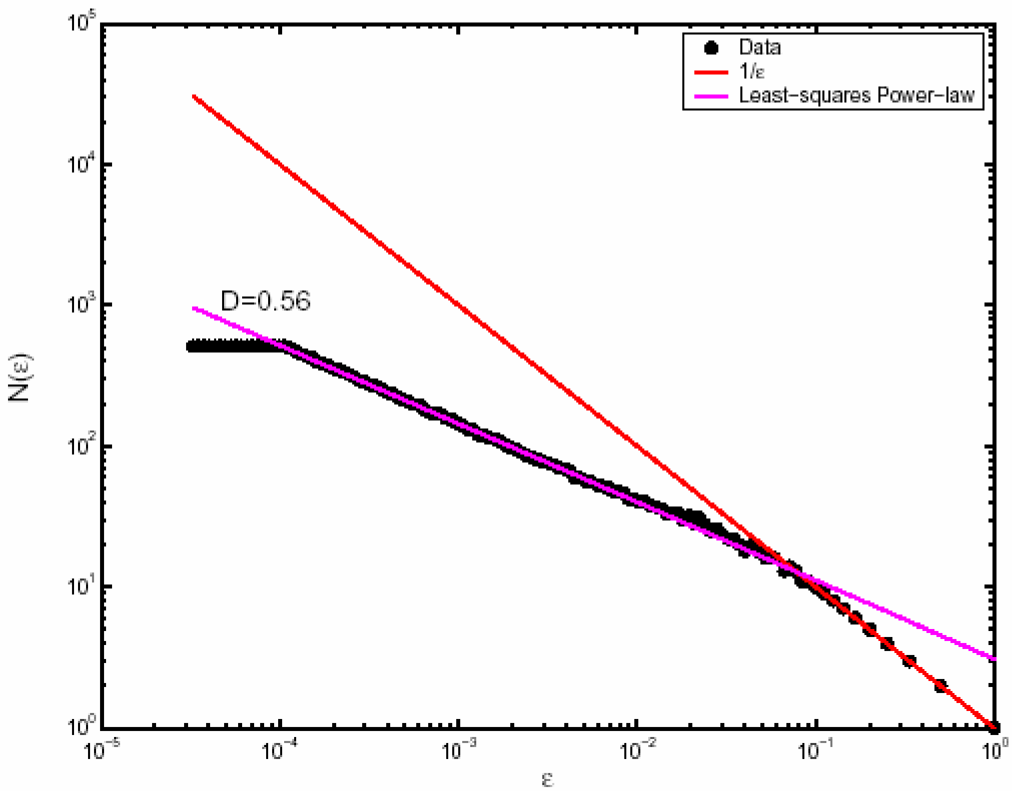

FIG. 7 A figure similar to 4, the three linear regions are clearly visible. The intermediate straight line was obtained from a non-linear least-squares power-law fit over a subrange of the data, giving a dimension of 0.56 for the 45 μm bin on March 3rd. Note that this entire analysis, however, is quite sensitive to the length of the data-set used to determine the points on this graph. Using only the first half of the distribution, for example, significantly changes the inferred value of D as well as the scales that the system appears scale invariant. For the whole of the data, however, it appears that the fractional scaling is appropriate for scales from about 30 cubic centimeters (10^{-4} on the graph) to about 30 Liters ($10^{(} - 1)$ on the graph).

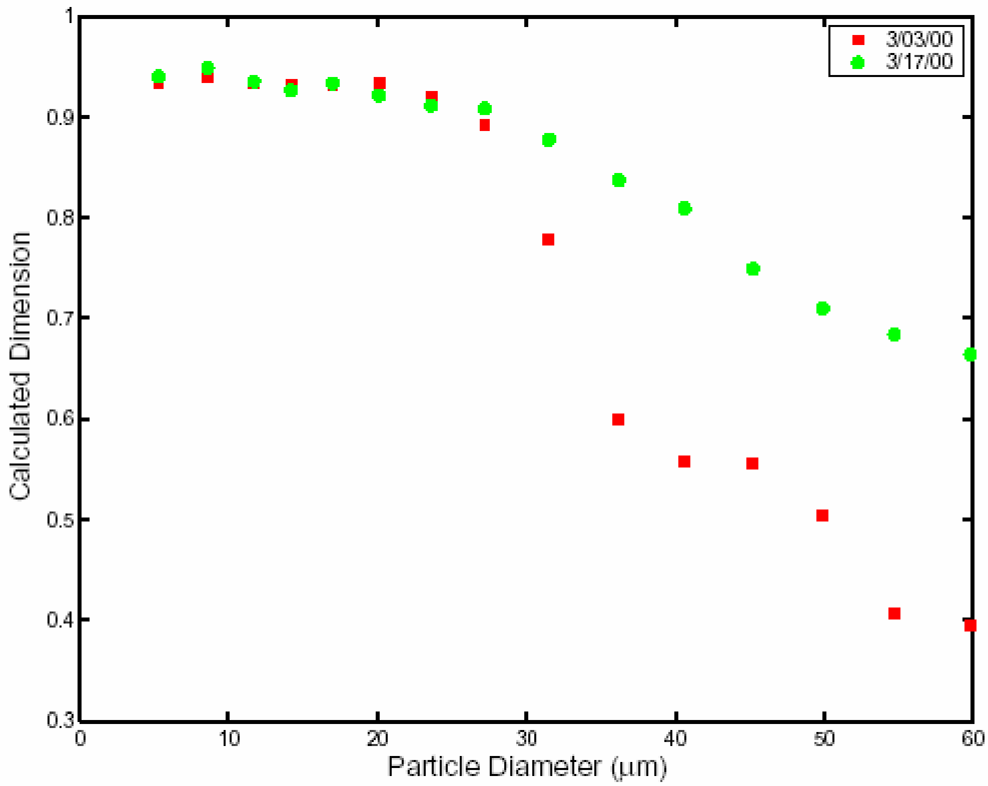

FIG. 8 Inferred fractal dimensions for all size bins for the flights on March 3rd and March 17th, using the entire flight as the sample set. Note the nearly monotonic decay in dimension with increasing particle diameter and the significant deviations from pure randomness for larger droplets.

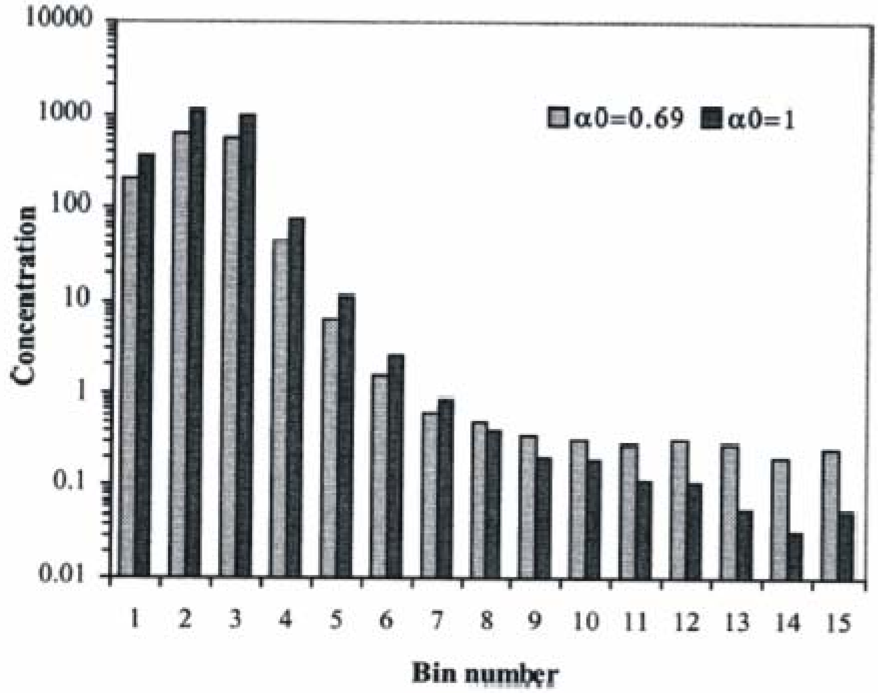

FIG. 9 A demonstration of the modified droplet size distribution given considerations of the fractal distribution of large droplets. Figure adapted from (Knyazikhin et al., 2002), which also gives more extensive description of how this was found. The parameter α_0 isn't particularly natural in this context, but it can be thought of as a surrogate for fractal dimension in an ensemble sense. (Lower values of α_0 means that there is more clustering/spottiness of the occupation of drops in the larger size bins.) The key observation is to note the increased contribution from large droplets in the system that displays "intermittency".

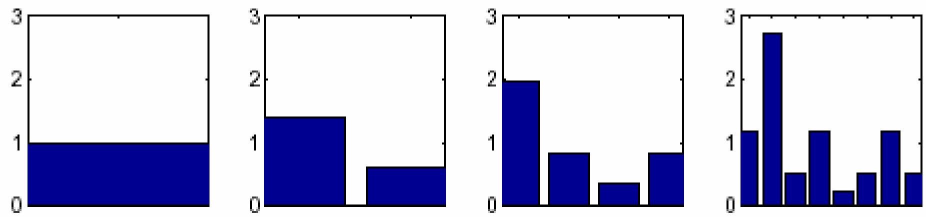

FIG. 10 The first steps of the cascade process that is used to generate the particle distributions. These examples are in 1-dimensional topological space for simplicity. In one-dimension, take a parameter $0 \leq p \leq 1$ and the unit interval (leftmost panel). Then, assign the value $1+p$ to half of the interval and $1-p$ to the other half at random (second panel). Now, repeat the process and in the half that had $1+p$ in the first step, half of that interval gets $(1+p)(1+p)$ while the other half gets $(1+p)(1-p)$ (this can be seen in the 2 leftmost bars in the third panel from the left). You then just repeat this subdivision routine until you have a system of the size you want. The system becomes more "intermittent" with more cascade steps and higher values of p. In 2 and 3-dimensions this gets more complicated, because the subdivisions in the system require 2^{d-1} parameters. The general principle works the same way, except that each area (2d) or volume (3d) gets broken into 2^d subunits.

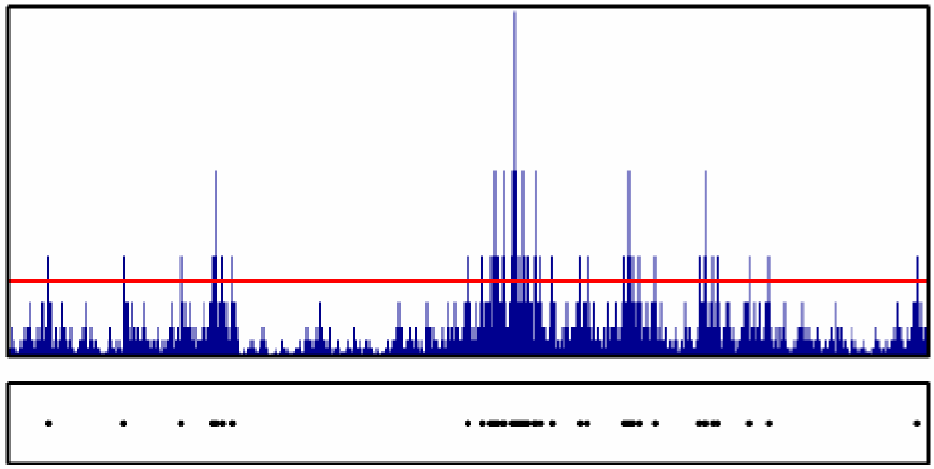

FIG. 11 The result of a 1 dimensional simulated cascade that was allowed to run for some time. The support for this system is found by taking all values of the system that exceeded a user-defined threshold value (specified by the red line in the top panel). If each point with nonzero measure is assigned a particle, the distribution in the lower panel results; obviously with a clustered structure. Increasing the threshold value in the upper panel creates larger correlations; for very small threshold values, one retains a perfectly random distribution. One can use a simulation like this to obtain data comparable to the FSSP if one envisions an airplane flying from left to right on the lower panel; encountering a "dot" is analogous to encountering a measurement with nonzero concentration in the bin of interest.

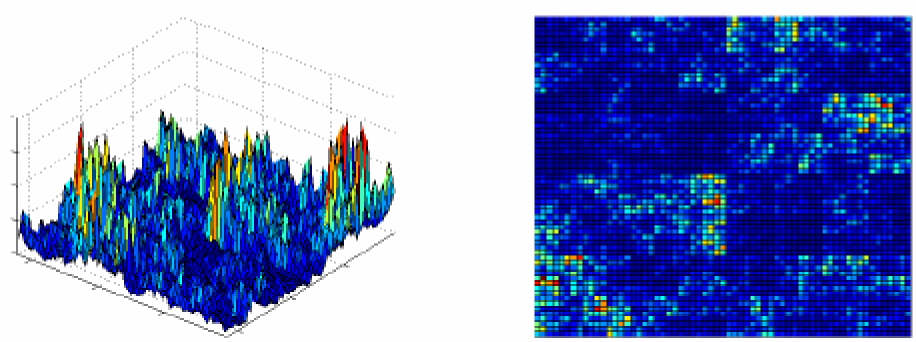

FIG. 12 A two-dimensional analogue to the top panel of figure 11. Here, one would determine where to place the particles in the plane by using a plane parallel to the (x, y) plane to "cut the bottom off" the distribution. The remaining areas (i.e. non-blue colors in the right hand panel) comprise the support for the system.

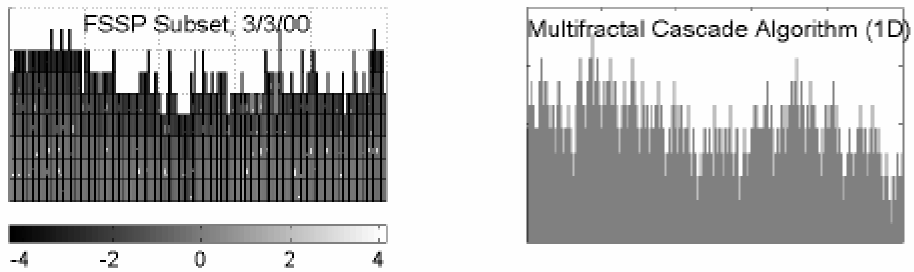

FIG. 13 A suggestive comparison between data and a simulation. In the caption for figure 11 it was noted that lower threshold cutoffs result in more uniform distributions. When the data from the FSSP is displayed in the manner akin to the left panel here, we see the same effect. The notion that larger particle clustering can be associated with nothing more than a larger threshold is suggestive based off the qualitative similarity you can see here. The data is a subset about 20 kilometers long, sampling a total volume of about 8 liters.

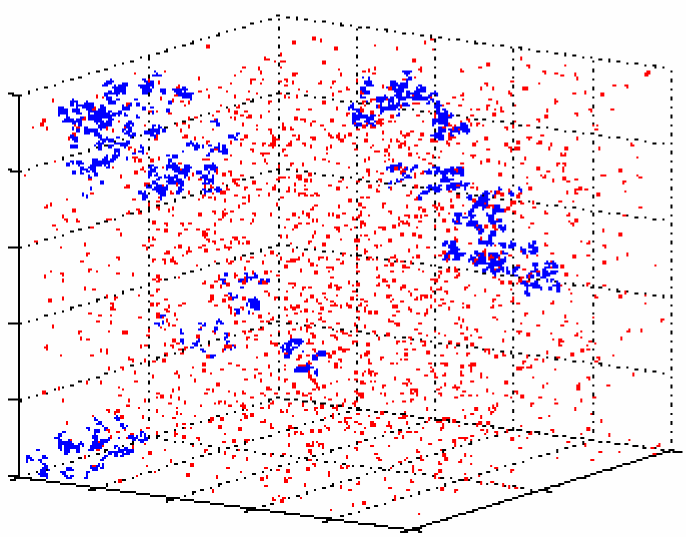

FIG. 14 A demonstration of the space-filling properties of distributions with different dimensions. In this figure there are an equal amount of particles that are colored red and blue. The red particles are distributed perfectly randomly throughout the space, whereas the blue particles are distributed in a way such that their fractal dimension is significantly less than 1. An intermittency not unlike that which we see with the data manifests itself in the blue-colored dots.

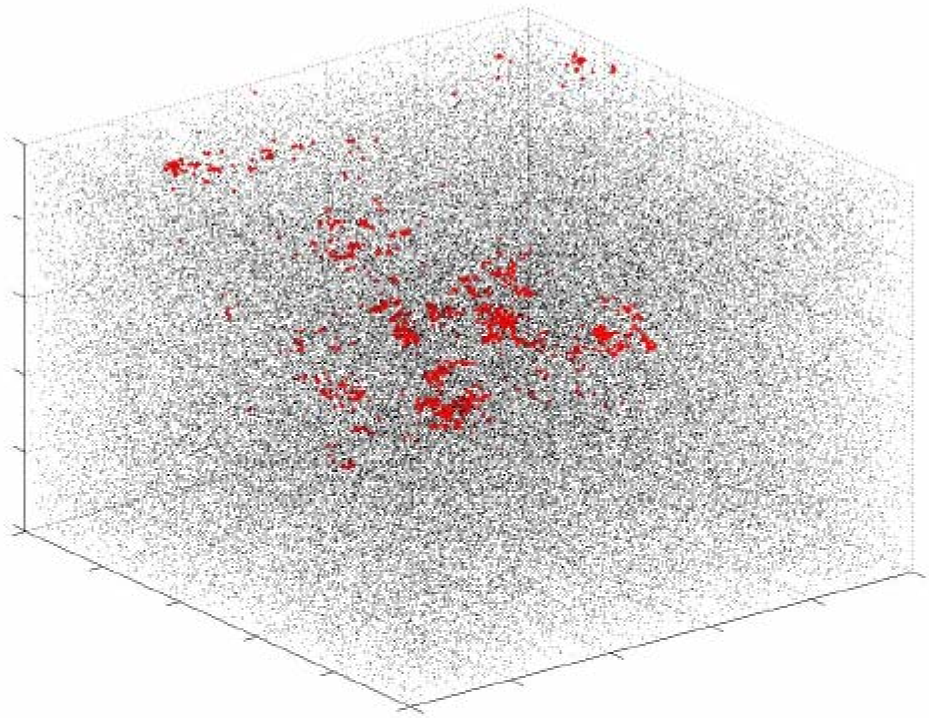

FIG. 15 An attempt to construct a pseudo-realistic "cloud" in a cubical volume. The numerous, smaller (black) particles are distributed totally at random. The slightly sparser, larger (red) particles were given a fractal dimension around 0.5 to simulate drops that are around $50\mu m$. In a system this size, with these relative numbers of particles, there may or may not even be any large particles. If there were, the radiative properties of this volume would be altered much more than the mean radius of the distribution, or even this volume, would be. (This was part of the point the authors of (Knyazikhin et al., 2002) tried to get across). Note that if you take a one-dimensional sample out of this cube, it seems believable that measurements like those in figure 6 and the left hand panel of figure 13 would be obtained.

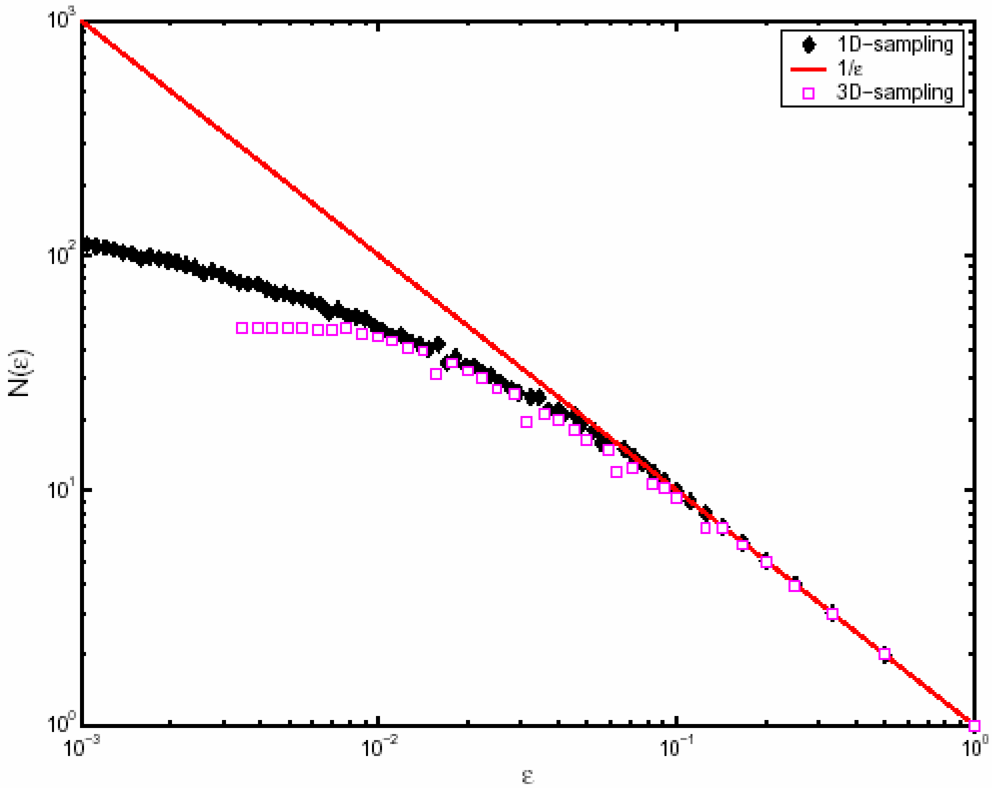

FIG. 16 A demonstration that one dimensional sampling of a three dimensional distribution does not change the inferred dimension of the system. Here, a fully 3-d simulation was carried out and the graph above was obtained. Then several 1-d passes through the cube were pieced together to obtain the diamond curve above. Note that they seem to have the same slope (and hence same dimension) in their scale-invariant region. (Note that to put both of these on one graph, I had to take the cube root of N and plot it against ϵ, since it is really D/d that remains invariant.)