ELSEVIER
Atmospheric Research 59–60 (2001) 163–197
ATMOSPHERIC RESEARCH
www.elsevier.com/locate/atmos

The HYDROP experiment: an empirical method for the determination of the continuum limit in rain

Nicolas Desaulniers-Soucy[a], Shaun Lovejoy[b,*], Daniel Schertzer[c]

[a]*EMS Technologies Canada Ltd., Canada*
[b]*Department of Physics, McGill University, 3600 University St., Montreal, Quebec, Canada H3A 2T8*
[c]*LMM, Paris, France*

Accepted 23 July 2001

Abstract

We describe the HYdrometeor Detection and Ranging using stereO-Photography (HYDROP) experiment, which is the first designed to measure the position and size of hydrometeors in a sufficiently large volume so that the continuum limit (large number of particles) can be examined directly. The experiment is instrumented by three medium-size cameras, two high-powered flash lamps and custom-made reflectors. The hydrometeors in approximately 500 scenes were photographed; six triplets of rain were digitized and analyzed in this study. Some details of the image processing and raindrop position and size reconstructions are outlined. We made the first determination of the sizes and positions of raindrops (diameters > 0.2 mm) in a significant volume (approximately 10 m^3). Preliminary estimates of the errors involved in the reconstructions yield position uncertainties of 3 mm in the directions perpendicular to the field of view, 3 cm along the field of view, as well as 0.5-mm-diameter uncertainty. We perform some preliminary analysis of six three-dimensional reconstructions—with 11 000–22 000 raindrops each and with nominal rain rates in the range of 8–15 mm/h. Since the mean interparticle distance was 7–8 cm and the largest sphere of observation has a radius of the order 1 m, our range of scales is roughly a factor of 10. © 2001 Elsevier Science B.V. All rights reserved.

Keywords: Hydrometeor; Scaling; Fractals; Multifractals; Cloud physics; Raindrops

1. Introduction

A basic problem in cloud and rain physics is the very large range of dynamically significant scales in both space and time. In order to account for this, two approaches have

---

[*] Corresponding author. Fax: +1-514-398-8434.
E-mail address: lovejoy@physics.mcgill.ca (S. Lovejoy).

PII: S0169-8095(01)00115-6

been developed. The traditional phenomenological approach is based on a series of strong homogeneity assumptions, each applied over limited ranges of scale. Due to the huge range of scales—centimeters to thousands of kilometers—of cloud and precipitation phenomena, this requires many distinct dynamical mechanisms; as an extreme example, Houze (1981) postulates one for each factor of two or so in scale.

More recently, a scaling approach has been developed based on a unique, but inhomogeneous dynamical mechanism able to accurately account for cloud and rainfall variability over a wide range of scales. At first (Lovejoy, 1981), this was a fractal mechanism, but following the development of multifractals in turbulence, the scaling mechanism was identified as a multiplicative cascade, (Schertzer and Lovejoy, 1985, 1987; Lovejoy and Schertzer, 1985, 1986) and the multifractal implications became clear. Indeed, it has become increasingly evident that if the dynamics respect a scale-invariant symmetry principle, then the generic result is a multifractal field. Theoretical developments, such as the scaling gyroscopes cascade (Chigirinskaya and Schertzer, 1996; Schertzer et al., 1997), has shown that these conclusions are also likely to apply to high Reynolds number Navier–Stokes turbulence in both two and three dimensions. From scales as large as 5000 km (Lovejoy et al., 2001) down to meters (Sachs et al., in press) of cloud radiances have indeed been shown to be multifractal (Lovejoy and Schertzer, 1990b; Cahalan, 1994; Davis et al., 1996), as has been shown by numerous analyses of rainfall data from gauges and radar over various scale ranges (e.g. Lovejoy et al., 1987; Gupta and Waymire, 1990; Tessier et al., 1993, 1994, see also the review of Lovejoy and Schertzer, 1995a,b).

Today, several studies of scales from 3 m (Lovejoy and Schertzer, 1991) and up have shown that rain statistics are indeed multiscaling as predicted. However, what has been most prominently missing is the study at the small-cale end: as we consider larger and larger volumes of the atmosphere containing more and more hydrometeors (denote this by N), how does the particle description converge to a field description? And what is the nature of the continuum limit obtained as $N \rightarrow \infty$? This is exactly analogous to the classical problem of the continuum limit in continuum mechanics. In the latter case, the main condition for the classical limit being the existence of a scale separation between the micro (statistical physics)-scales of the order of the mean free path, and the macroscopic (continuum mechanical) larger scale variability. When this classical "continuum hypothesis" holds, then the continuum limit is homogeneous. However, in the analogous rain case, the mean interparticle distance is typically many centimeters or larger, i.e. it is much larger than the turbulent dissipation scale (typically of the order of millimeters). Since the turbulent wind field itself is highly intermittent (it is apparently multifractal), we may indeed expect that the continuum limit is likewise nonclassical. Multifractility has already been verified for cloud droplets (e.g. Lovejoy and Schertzer, 1995a), we here study the analogous question for drops. Indeed, the main a priori criticism of this argument is that due to raindrop inertia, the rain does not immediately respond to the turbulence, i.e. the particles and turbulence may be "decoupled" below some small length. Although this effect will primarily affect the smallest scales and the biggest drops in the most turbulent storms, it may be enough to introduce a length scale of several meters below which some homogeneity is realized.

The central motivation of this study is therefore, to empirically study the approach to the continuum limit and to directly determine its scaling properties. In a scaling process, each statistical moment has a different scaling exponent. If this exponent function is nonlinear, then the field/time series is a multifractal. Mathematically, multifractal measures have densities which are singular with respect to the usual (Lebesgue) measures (length, area, volume, etc.), in physics jargon, the classical scaling of the latter being termed “trivial”. The homogeneous (classical) continuum limit is the special case of trivial scaling. The spatial distribution of liquid water—both raindrops and snowflakes—was chosen as a basic measurable field relevant to cloud physics. These hydrometeor fields are fundamental in estimating rain from radar or areal rainfall from rain gauges. In order to fulfill these goals, the HYdrometeor Detection and Ranging using stereO-Photography (HYDROP) experiment was designed.

The second section of this paper describes the series of hardware challenges encountered and the solutions employed to photograph a number of raindrops sufficient for statistical analysis. The third section describes the software strategies used to implement stereo-photography for the determination of raindrop positions using three photographs of the same scene. The fourth section concentrates on the raindrop size determination procedure and its calibration. The fifth section presents the raw data analyzed in this paper as well as some measurements of meteorological interest, such as the raindrop size distributions and nominal rainfall rates estimated from the recovered data sets. The analyses relevant to the determination of the continuum limit in rain will be presented in a future paper by the same authors.

1.1. A review of studies of rainfall

Many different types of rainfall analysis have been conducted in order to find evidence of scaling over a wide range of scales. With few exemptions, they were performed on scales too large to resolve raindrops, and their limitations were important in motivating the present study. Since the focus in this paper is the study of the scale-by-scale properties of rainfall at the smallest scales, the literature review listed in the remainder of this section will share the same focus. Previous studies are discussed according to the number of dimensions of their observing spaces: time series, one-dimensional cross-sections, two-dimensional intersections and, finally, three-dimensional reconstructions.

1.2. Time series

Out of a four-part series by Kostinski and Jameson, three papers are especially relevant to this study: Jameson and Kostinski (1998, 1999), Kostinski and Jameson (1997). In the first two cited papers (1998 and 1999), the Poisson temporal distribution was tested over two time series collected using a Joss–Waldvogel (JW) disdrometer (Joss and Waldvogel, 1967) with a temporal resolution of 1 min. The data consists of the number of raindrop hits for each 1-min interval, arranged in a series of size bins corresponding to raindrop diameter (the size resolution was not indicated).

The first data set of Jameson and Kostinski (1998, 1999) consisted of a 900-min JW time series, whereas the second is 33 min long and was gathered by List et al. (1988). In the first part of their study (Kostinski and Jameson, 1997), temporal fluctuations of drop counts in a single size bin were analyzed and in the second part (Jameson and Kostinski, 1998), at correlations between drop counts from different size bins were considered. The fourth part involved a video disdrometer used by the Iowa Institute of Hydraulic Research of the University of Iowa (Schönhuber et al., 1994).

In the first two parts, Jameson and Kostinski found that the "Poisson-mixture", in which the mean is computed over "patches" (the interested reader is referred to the original paper for details), fits the data the most accurately during the passage of the squall line and the pure Poisson accurately represents the data subsequently. A time-dependent Poisson distribution fitted the List data set better than the pure Poisson.

The size of raindrops and their corresponding time of arrival through an optical disdrometer with a 100 cm^2 area have been recorded by Lavergnat and Golé (1998). They quote a size resolution of 10^{-3} mm (for diameters larger than 0.3 mm) and a temporal resolution of 1 ms. The longest continuous experiment with this high-resolution data was 4 months and the combined data sets span 14 months. They also recorded 1-min diameter–velocity dual histograms for a continuous period of 25 months.

They noted that the histograms of the time intervals between drop counts do not follow the exponential survival law predicted by a Poisson distribution, but are closer to a power law (Pareto) behavior. Without explicit theoretical justification (power laws are associated with scaling), they picked the latter as the basis of their model and fitted the distribution parameters to their data.

In addition, they computed a (mono)fractal dimension associated with the high-resolution time series and obtained $D_f = 0.82$ over a scaling range of 1 s to 4 months, with a slight break in their scaling at the scale of 1 day. They generated a time series of raindrop arrival times whose box-counting analysis revealed a behavior extremely close to that of the experimental data, including the slight scaling break. They could even model two breaks (at 30 min and 1 week) in the scaling of the analysis of the Olsson et al. (1993) tipping bucket time series. Based on this evidence, they concluded that rain is a renewal process. However, they did not put forward a physical interpretation of the observed temporal behavior of rain.

Uijlenhoet et al. (1999) have tested the classical Poisson homogeneity hypothesis in rainfall on a 35 min time series of 10 s raindrop size spectra collected with a 50 cm^2 optical disdrometer. Even for stationary time series, the total raindrop arrival rate showed larger fluctuations than the Poisson distribution allows. By analyzing separately the raindrops belonging to different diameter intervals, they concluded that the deviation from the Poisson distribution was evident only for raindrops with diameters smaller than 1.14 mm. Note that this could be consistent with the "decoupling" of the largest drops from the turbulence at the very small scales.

1.3. One-dimensional cross-sections

So far, only one type of empirical measurement of rainfall has provided spatial 1-D cross-sections at small scales: lidar studies (Lovejoy and Schertzer, 1991). On the other hand, there are many instances where researchers have mounted a droplet detector on an aircraft and recorded data while passing through clouds. This yields a 1-D section through space-time. The high speed of the aircraft relative to that of the clouds as well as the small sample size of the detector makes the section nearly spatial. Of course, reality is more complicated since the aircraft trajectory, just like the object of study, is affected by turbulence (and may even be fractal!). Since the vertical and horizontal cloud structures are quite different due to stratification, the uncontrolled vertical fluctuations in the aircraft mix these different structures in a nontrivial way.

A frequently cited paper concerning the study of the variability of cloud droplet populations is that of Austin et al. (1985), in which the deviation from Poisson Statistics was computed. In order to quantify this deviation, the cloud droplet spectra were measured using 50 and 10 Hz Forward-Scattering Spectrometer Probes (FFSP) during several flights through continental cumulus clouds. By thresholding the ratio of the standard variation of the local droplet number concentration with respect to a “steady condition” standard deviation, they classified portions of the flights as “steady” and “mixing”.

Following the same philosophy, Paluch and Baumgardner (1989) measured cloud droplet concentrations at 20-ms intervals using a FSSP during flights through young nonprecipitating cumulus clouds. The corresponding spatial resolution was approximately 2 m as a consequence of an aircraft speed of 100 m/s. They checked the validity of the probability distribution function of the distance between two particles in a Poisson distribution and when it did not adequately fit the data, they conducted that some mixing occurred in the region during the time of measurement.

Baker (1992) used the same apparatus as Paluch and Baumgardner (1989) to explicitly test the Poisson distribution hypothesis. The main difference is that Baker recorded the time of arrival of each droplet with a temporal resolution of 4 μs. The airplane velocity was roughly constant at 80 m/s, which corresponds to a spatial resolution of 0.32 mm. Baker defined a “Fishing” test based on the fact that the mean and variance of the Poisson Distribution are equal and applied it to FSSP data. Regions where the Poisson statistics failed to represent the data were identified at different scales, from 5 mm to 100 m, as well as regions where the fluctuations were less important.

Brenguier (1993) modified the electronics of the data acquisition system of the FSSP to virtually eliminate dead times and lower the temporal resolution to 62.5 ns. This corresponds to a spatial resolution of approximately 5 μm for an aircraft flying at 80 m/s! Brenguier did not attempt any statistical analysis of his data. He observed that there were extremely sharp edges in terms of droplet concentrations at scales of the order of 5 mm.

A complete multifractal study has been performed by Brösamlen (1994) from the cloud Liquid Water Content (LWC) data of the FIRE project (5-m resolution), as reported by Albrecht et al. (1988), some of which is published in Lovejoy and Schertzer (1995a,b). Five different data sets showed scaling over a range of 5 m to 330 km. The universal multifractal parameters were estimated to be $\alpha = 1.99 \pm 0.01$; $C_1 = 0.08 \pm 0.02$; $H = 0.29$. The latter two parameters were also measured on nearly the same cloud data set by Davis et al. (1996).

FSSP data were used by Malinowski and Zawadzki (1993) in a fractal analysis. Droplet number concentration and liquid water content were recorded at 64 Hz at 68 m/s which yields a spatial resolution of roughly 1 m. They computed the box-counting dimension and the correlation dimension on thresholded fields of number concentration and LWC. They concluded in favor of a monofractal field of dimension between 0.48 and 0.63, which they extrapolated (assuming both monofractality and isotropy, ignoring stratification) to a fractal dimension of clouds of 2.55, in a range extending at least from 10 m to 1 km.

Jameson et al. (1998) explicitly tested the Poisson distribution hypothesis using airborne two-dimensional array probes (PMS 2D-C). Droplets from 0.025 to 0.8 mm were imaged by the detector and histograms were computed at each second. The sample volume was 10^{-3} m^3 (130 m long). They chose 50-μm droplets for their analyses. They showed that the variance in droplet number concentration could be as high as 100 times the variance of the corresponding Poisson Distribution with the same mean. They also measured a two-point excess correlation function significantly larger than zero for nonzero lag. They explained these deviations by a "Poisson-mixture", as introduced in Jameson and Kostinski (1998, 1999) and Kostinski and Jameson (1997).

A multifractal study has been performed using the lidar data (Lovejoy and Schertzer, 1991). A 10-Hz pulsed YAG laser was fired approximately 5000 times and the returning light intensity was recorded. After a correction for $1/r^2$ range fall-off, each 3-m segment along 540 m was thresholded to eliminate noise. A total of 500 × 180 segments were analyzed and trace moments (multiresolution statistical moments) were plotted in space, time, and "squares" of the space-time domain. Excellent multiscaling over three orders of magnitude in time and two in space were found.

1.4. Two-dimensional "intersections"

Lovejoy and Schertzer (1990a) used chemically treated blotting paper to record raindrop positions in a two-dimensional surface of 128 × 128 cm. Although the intent of the authors was to achieve planar intersection of space-time rain, the exposure was unfortunately long enough to produce a nontrivial temporal projection with potentially a different fractal dimension. A total of 452 raindrops were collected and digitized manually. The correlation dimension was found to be 1.83, the codimension being 0.17. The codimension C is the dimension of the embedding space minus the dimension of the set; $C=0$ implies trivial scaling. The small number of raindrops has lead to criticism of the results (Jameson and Kostinski, 1998) and prompted us to complete a series of rain measurement experiments culminating in the present study.

In a series of experiments using "a raindrop camera", raindrops were photographed in a volume of 0.143 m^3, with seven different exposures separated by 1.5 s, for a combined volume of 1 m^3 (Jones and Dean, 1953). The raindrop diameter in as many as 26 865 such volumes were manually recorded and the raindrop spectra analyzed in Jones (1992). Although the projected position of all these raindrop on a plane could potentially be analyzed, the position information is not reported in the three papers quoted above and was not made available to the research community, unlike the diameter information.

1.5. Three-dimensional reconstructions

Up until now, holography is the only method used to measure a three-dimensional distribution of hydrometeors. The major limitation of this method are the relatively small sampling volumes. It is extremely difficult to "spread" a laser beam enough to cover a large volume while simultaneously maintaining its quality. This is why studies performed so far have been limited to sampling volumes of the order of a cubic centimeter. In order to gather enough statistics, small and numerous cloud and fog droplets were studied.

The first study aimed at testing the Poisson distribution was performed using fog droplets by Kozikowska et al. (1984). The drops in a volume of 22.5 cm^3 were localized and their sizes recorded. The droplets were classified in four size bins and three cell sizes were studied. They concluded that:

> The measurements show systematically greater numbers of empty cells and greater numbers of cells with high droplet content than would follow from purely random distributions, as represented by Poisson curves. This suggests some clustering mechanism.

At the other extreme, Borrmann et al. (1993) found that the deviations from Poisson were negligible. They located 998 droplets of a stratus cloud which was low enough to be sampled from a small mountain in West Germany. The sampling volume was 8 cm^3. They computed the droplet number concentrations in subcells of the sampling volume and the mean interparticle distances. Note that the inner scale of the turbulence (the viscous scale) can readily be of the order of 1 mm so that some clustering may be expected for these cloud drops at these scales.

Note that an analysis done by the author of this paper of the same data yielded a correlation dimension of the order of 0.5. Due to the low number of particles, no other scaling analysis was performed on their data. This shows that a proper empirical analysis of rain should be made without first assuming trivial scaling. Even then, it is obviously impossible to empirically prove that the codimension of the liquid water density is exactly zero (i.e. the Poisson homogeneity assumption). The most that can be done by experiment is to put tight bounds on it.

In conclusion, none of the studies presented in this literature review convincingly supported the Poisson distribution as a good description of rain statistics. Since this result has profound consequences on cloud physics, many authors have tried to diminish its importance by:

- minimizing the significance of slight discrepancies from the Poisson distribution;
- introducing ad hoc "correction" laws such as the Poisson-mixture;
- considering two manifestations of rain: in turbulent (apparently too difficult to handle) and in homogeneous regimes.

Others introduced ad hoc laws to explain their results such as the bi-Pareto and simple fractal. In contrast, a multifractal LWC density would result if the nonlinear dynamics respect a scale invariance symmetry. The corresponding physical model is of coupled cascade processes involving various conserved fluxes (see Schertzer and Lovejoy, 1987 for more details). We, therefore, conclude that the cascade model at the basis of the multifractal formalism is still the simplest and most physically based answer to the quest for a characterization of rain over a wide range of scales, and the many scaling laws as found in numerous rain studies support this approach.

2. Experimental setup

2.1. Experimental setup

To illustrate the challenge we faced in determining the position and size of hydrometeors, consider a 1-mm-diameter raindrop scatterer. In the geometry indicated in Fig. 1, this would produce an intensity at the camera equivalent to only 10^{-10} of the intensity incident on the raindrop! One of the reasons for the low cross-section of raindrops is that typically the light is scattered in a direction significantly different from forward scattering and the cross-section falls off rapidly away from the forward direction. An additional

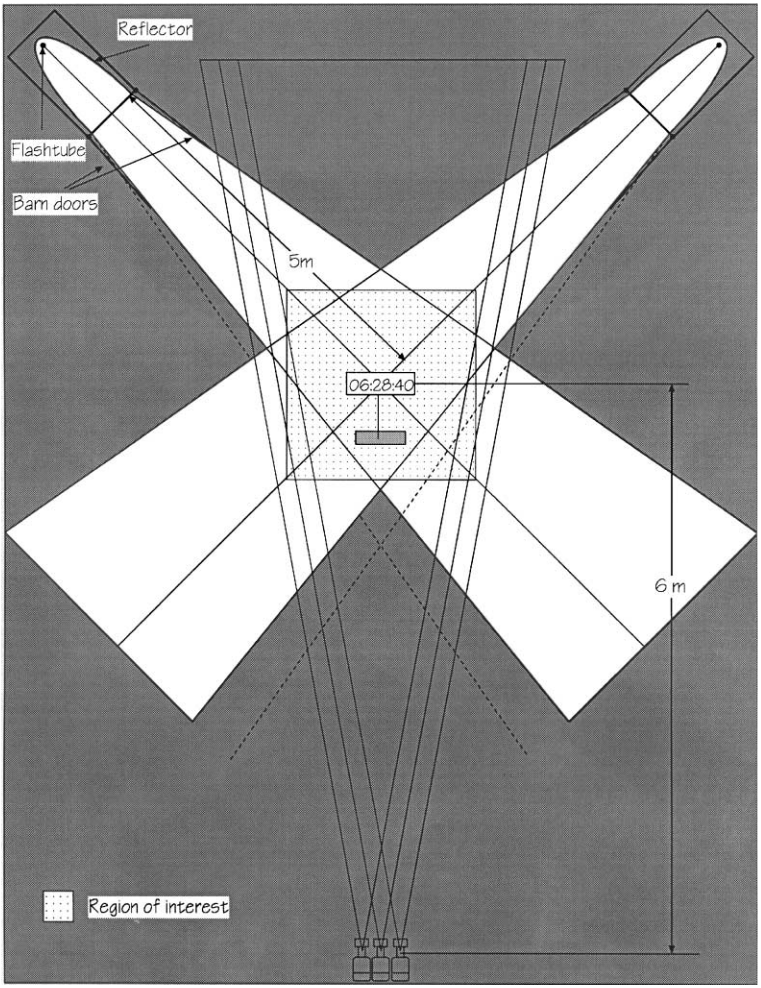

Fig. 1. Geometry of the different optical elements.

constraint is imposed by the fact that raindrops fall relatively rapidly (up to 10 m/s, e.g. Beard, 1976); hence, the light must be delivered in a very short time. In order to "freeze" the image of the raindrop on photographic recording media, the light source must be extremely powerful.

To this end, we selected a high-voltage discharge through two xenon-filled flashtubes connected in series. These flashtubes are qualified as UV-rich, fast extinguishing. In each flashtube, up to 1 kJ is dissipated in a 12-cm-long arc over a period of roughly 50 μs. Each flashtube can hold off 5 kV and the assembly requires a high-voltage trigger of at least 25 kV applied simultaneously around each flashtube to be triggered into conduction.

These flashtubes were powered by a discharge power supply consisting of a 110 V to 10 kV transformer, halfwave charging a set of two 10-kV, 20-μF capacitors. The output voltage of the transformer can be varied by controlling the input voltage using a variable autotransformer. A total of 2 kJ of electrical energy is dissipated during each discharge and according to the manufacturer of the flashtubes, approximately half of this energy is transformed into visible light. The instantaneous power can reach 40 MW with a peak current larger than 10 kA. Hence, special high-voltage connectors and RG-213 coaxial cable became necessary to connect the power supply to the flashtubes.

Even such a high power of light proved to be insufficient for the photographic requirements of this study. A custom-made reflector was designed by optimizing both the efficiency of collecting the light where it was needed as well as the uniformity of the light intensity across the region where the raindrops had to be detected (in a cube centered 5 m from the light source). The theoretical efficiency calculated from the reflector shape was of the order of 30%. By comparison, the efficiency, η_n, of the isotropic illumination when no reflector is present is:

$$\eta_n = \frac{\text{surface of backplane}}{\text{surface of sphere @ 6 m}} = \frac{4}{4\pi 6^2} \approx 0.008, \tag{1}$$

or 0.8%, i.e. an improvement of a factor of approximately 35.

The reflector was inserted into a high-voltage housing system capable of evacuating the heat generated by the discharges. The solution implemented resembles the schematics shown in Fig. 2. This system is secured by a 0.7 × 1 × 1.2 m wooden structure on wheels, which allows easy displacement and orientation of the light sources.

2.2. Optical geometry

The objective of the project was to record hydrometeor positions and sizes in a volume of approximately 10 m^3, which corresponds to a cube of side slightly larger than 2 m on a side. The distance from the camera to the plane of focus, or center of the Region Of Interest (ROI) was chosen to be 6 m. The optimized optical geometry of the setup (as seen in Fig. 1) is constrained by this distance as well as by the distance between the light source and the center of the region of interest. The angle between the camera and the light sources remained to be determined and it was done, in the most part, by trial and error analysis.

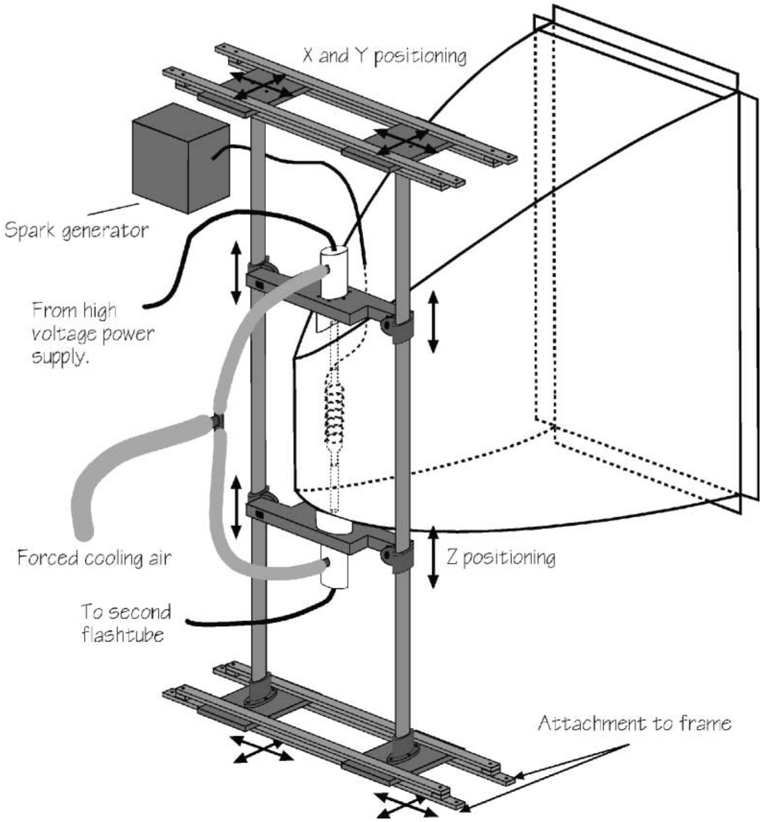

Fig. 2. Schematics of the flashtube mounting system.

At the center of the region of interest lies a clock. The reasons for its presence are manifold:

- it serves as a visual guide for focusing the cameras;
- it appears in the photograph being taken making sure that the correct photographs are being correlated;
- it guides the orientation of a photograph by indicating where the bottom of the image is;
- it tells the time at which the photographs were taken for future comparison with meteorological measurements.

Although it partly blocks the field of view, the advantages outweigh the disadvantages.

The camera separation was chosen to be 20 cm. Admittedly, no attempt was made to optimize it. In order to better resolve ambiguities, it was decided that three cameras would be used instead of two. Ambiguities arise when two or more raindrop images in one photograph are such that either one could be the image of a single raindrop in another photograph. Moreover, the center camera was elevated with respect to the others, allowing the ambiguities to be resolved even more reliably.

The vertical orientation of all three cameras was either horizontal or raised by 12°, just enough to get the lower edge of their field of view horizontal. An illustration of the geometry of these setups can be seen in Fig. 3. Please note also how the distance from the center of the region of interest is increased by the elevation from 1.2 to 2.1 m. Although the first experiments with rain were carried out with a colinear, horizontal camera orientation, it was changed to be convergent, elevated mode after a few months.

2.3. Imaging

Many factors, including the exceptional optical quality of the Zeiss lens, lead to the choice of the Hasselblad 500 EL/M camera with a 150-mm lens. The EL/M is the motorized version of classic 500 CM. The lens aperture was set to f/22, for which the radius of the circle of confusion due to diffraction and to focusing a point at the edge of the region of interest are equal. It is roughly equal to the resolving power of a film such as the

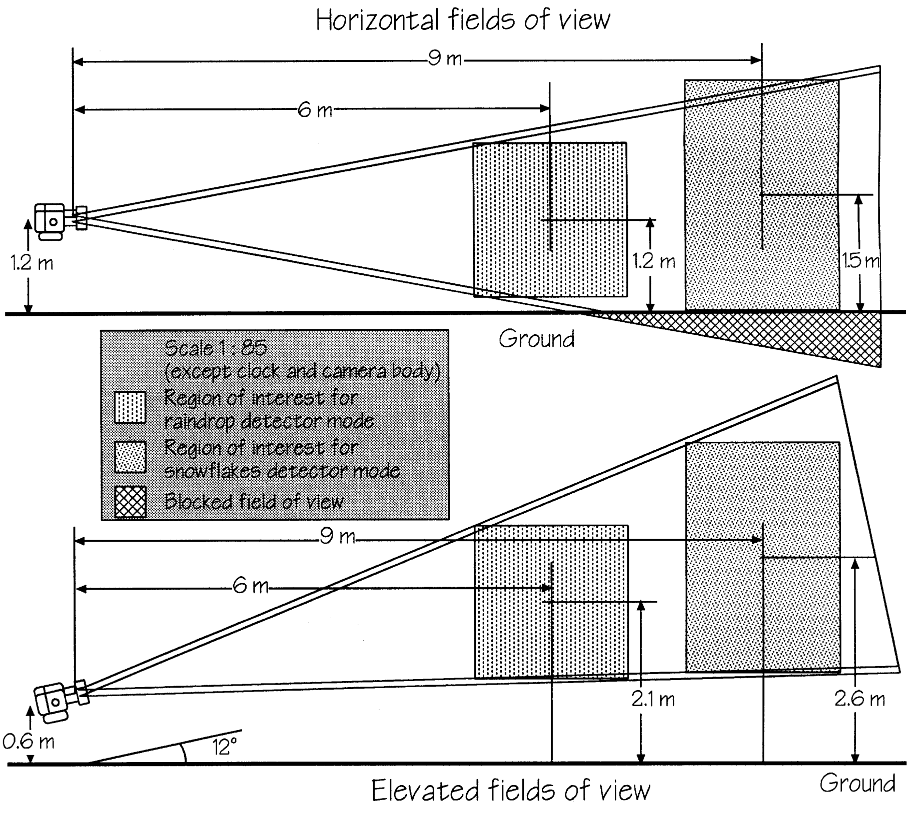

Fig. 3. Comparison between the vertical orientations of the two setups used.

popular T-MAX 100 ASA, which has a high exposure index (ASA) compared to other films with comparable resolving power. This film was chosen after extensive testing. Field experiments were made necessary by the subtle differences between photographic emulsions.

The essential aspect of the camera synchronization is actually the fast discharge time of the flashtubes (50 μs, which can be translated to 1/20 000 s in photographic terms). The photographs were taken at night in order to make the task of the image processing algorithm easier by reducing the background noise. Although this considerably reduces the need for precise synchronization, the shutters of the three cameras, nevertheless, had to be opened at the same instant to trigger the flashtubes. This was taken care by a custom-made circuitry connecting each camera and the spark generator igniting the flashtubes.

The light characteristics of the photographic negative were translated into visual information treatable numerically using the microdensitometer of the National Research Council of Canada Astronomy Centre, called the PDS. The PDS measures the transmitted light intensity of a laser beam directly through the negative and quantizes it using 12 bits. From there, it performs a logarithmic conversion to find the density of the negative at any given point. The PDS was set to a resolution of 6 μm. Since the Hasselblad camera exposes an area of (approximately) 5.5 × 5.5 cm of the photosensitive emulsion, the resulting range of scales in all directions is 9200.

The only problem with this method is the very low speed at which the PDS scans: it takes 30 h to scan a single negative, or 90 h to scan the three negatives necessary to reconstruct a scene. The number of scans made by the PDS will, therefore, always be limited. It could, nevertheless, serve as an excellent standard against which we may compare the quality of the possibly more numerous scanned photographs.

3. Raindrop position determination

3.1. Raindrop image processing

Considering the optical properties of a raindrop (see Appendix A of Desaulniers-Soucy, 1999) and the distance between the region of interest and the sources of light (see Fig. 2), these will be considered as point-like for the remainder of this paper. The image of a point-like light source on the photographic negative will differ depending on the position of the light source with respect to the plane where images are correctly focused. Goodman (1996) quantified the severity of the focusing error by the parameter, W_m. The use of the dimensionless parameter w was preferred, which is given by the following equation in the case of a square aperture of width $2W$ defined as (following Eq. (6-39) of Goodman, 1996):

$$w = \frac{W_\mathrm{m}}{\lambda} = \frac{1}{2}\left(\frac{1}{z_a} - \frac{1}{z_i}\right)W, \tag{2}$$

where λ is the wavelength of the light used for imaging, z_a is the distance between the object and the lens, z_i is the distance between the lens and the point of perfect focus.

The larger w, the more the image is out of focus (w is 0 when the image is perfectly in focus). According to numerical calculations detailed in Appendix C of Desaulniers-Soucy

(1999), the cross-sectional intensity profile of a raindrop illuminated by a point-like light source is given in Fig. 4 for different values of *w*.

The digitized image of a typical large raindrop (in focus) is shown in Fig. 5. Two very distinct bright areas are visible, each one of them being the image of a light source, as seen through the optics of the raindrop. The bright area at left corresponds to the right light source. It is brighter because at the position of the raindrop, the scattering angle of the light rays emerging from the right light source was smaller than the corresponding angle from the left light source.

Since two distinct areas horizontally separated are present in the image of a large raindrop, vertical cross-sections through one of them can be matched with the numerical cross-sections of Fig. 4. A typical vertical cross-section for a large raindrop in focus and out of focus is shown in Fig. 6. The cross-sections correspond well to the predicted profiles of Fig. 4, with $w=0$ for the image in focus and w is between 0.6 and 0.8 for the image out of focus.

It should be observed from Fig. 4 that the maximum intensity of a raindrop image out of focus will be smaller than the corresponding intensity for a raindrop in focus. In addition, the raindrop image intensity varies as the square of the raindrop radius and as the inverse square of the raindrop–camera distance.

In light of these important variations, it is important to find a density threshold which is higher than the noise floor of the digitized photographic negative and low enough to detect small raindrops that are out of focus, at relatively high scattering angles. In fact, two

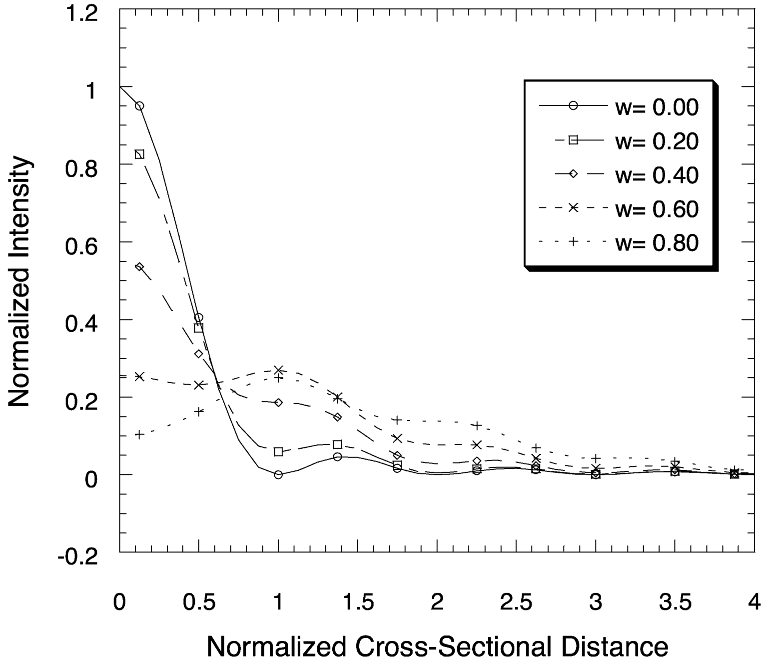

Fig. 4. Intensity profiles across half a water sphere for different values of the defocusing parameter w.

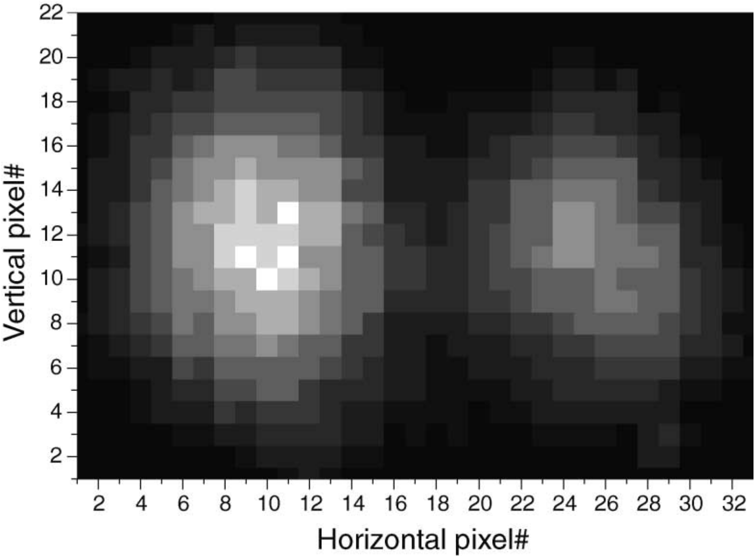

Fig. 5. Density plot of a typical image from a large raindrops in focus. The brighter the shade of gray, the more exposed the negative was at that position (pixel).

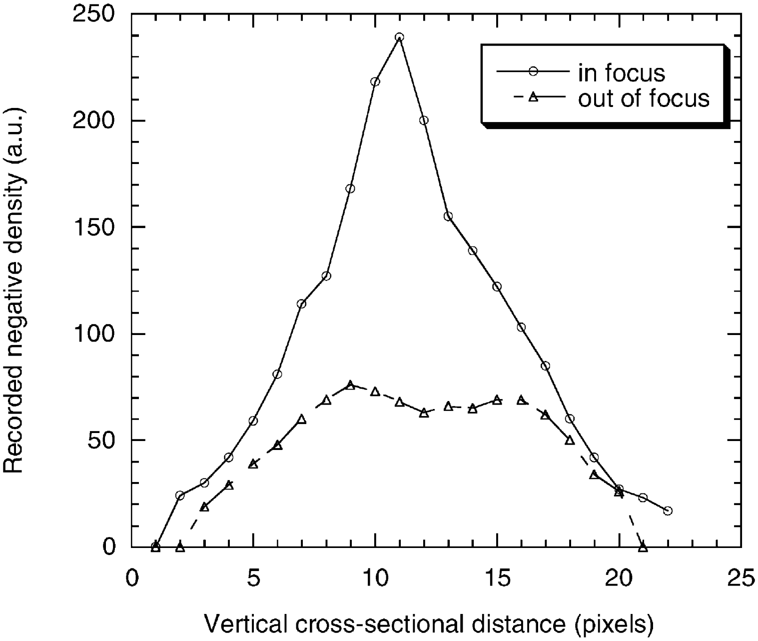

Fig. 6. Vertical cross-sectional density profiles through the image of a large raindrop, in and out of focus.

thresholds were used, in order to further decrease the number of false detections (using the high threshold) and increase the area of the raindrop image (using the low threshold). An example of a cross-section through an entire photograph is shown in Fig. 7 with the two thresholds superposed.

Once a bright region of a digitized photograph has been singled out, it is subject to a number of tests whose goal is to reject images other than raindrops. A raindrop image is then parameterized using a number of fitting functions whose shape in a three-dimensional plot would be a cone whose top is cut at different height. In the profiles of Fig. 6, the cone would be hardly cut for the image of a raindrop in focus and cut pretty low for the out of focus image. The number of cones fitted is variable, but preference was set for one (small drop for which the image of both light sources are indistinguishable or large raindrop out of focus) or two (for large raindrops in focus). If more fitting functions are necessary, it is primarily in the case of overlapping raindrop images. The details of the parameterization are given in Desaulniers-Soucy (1999).

3.2. Raindrop image correlation and position computation

The raindrop images were listed and parameterized from all three photographs of the same scene. Before computing the three-dimensional position of each raindrop using its image in all three photographs, the stereographic optical parameters must be calibrated. The details of this calibration are beyond the scope of this paper; however, the methodology followed closely the principles laid down by Faugeras (1993). Suffice it to say that nine

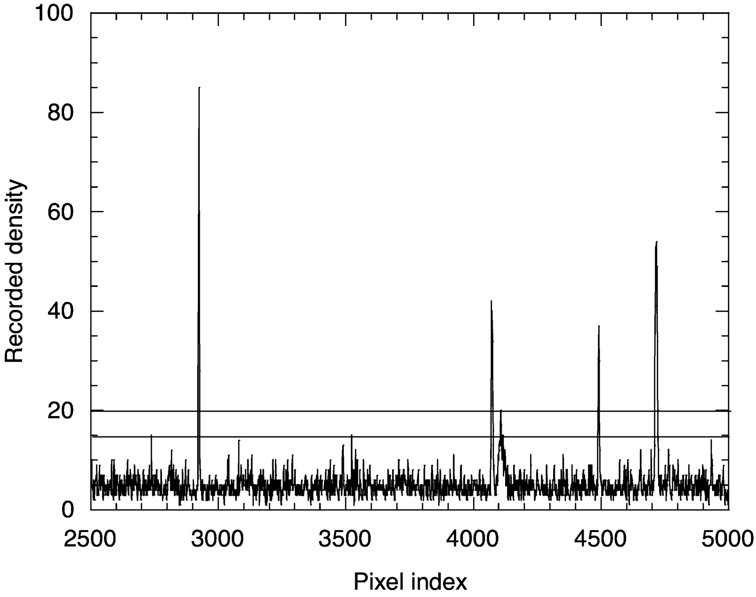

Fig. 7. A typical cross-sectional profile across the centre region of a photograph digitized by the PDS. A typical value of the upper (20) and lower (15) thresholds are shown across the figure. The large spikes are clearly due to raindrop images.

parameters per camera must be calibrated. With digital cameras at a fixed position through many experiments, it would make sense to calibrate these parameters beforehand, using a specially designed geometrical object. In our case (traditional photography), large raindrops were chosen from the photographs, in such a way that its image in all three photographs could unambiguously be identified. Using the image positions of 25 to 50 of these raindrops, a nonlinear regression was used to converge to a solution for each of the scenes studied.

Once the calibration is done, the list of parameterized raindrop images in all photographs is scanned several times in the search for images corresponding to the same unique raindrop. Using the position of two or three of these images, the three-dimensional position of the raindrops were reconstructed. Because of the large number of the images in each list and the different number of fitted functions in one image, the automated process has to consider many possibilities; hence, it cannot be expressed simply. The details are described in Desaulniers-Soucy (1999). At the end of this process, we were able to match 90% of the fitting functions. The extent to which this is a good or bad performance would require further analyses since there might be many reasons for having more fitting functions than necessary (e.g. dust, noise raindrop image profile, rain drops at the edge of one photograph, but invisible in others).

For those raindrops for which three images have been identified, there exist three reconstructed positions (one per photographic pair). In such a case, the final reconstructed position of the raindrop corresponds to the center of gravity of the three positions calculated from stereo-photography. The difference between the final position and the three intermediate ones led to a computation of a standard deviation, which was the basis for the estimate of the uncertainty of the position reconstruction: 3 cm for the direction along the field of view and 3 mm for the two other perpendiculars.

4. Size determination

As described in Appendix A of Desaulniers-Soucy (1999), ray optics has been used to relate the diameter of a spherical scatterer to the distance between the center of the two bright spots visible at the surface of its image, visible on Fig. 5, and the position of the scatterer with respect to the optical elements of the setup. This relationship has been used to compute the diameter of raindrops for which the two spots are distinguishable with an accuracy of a quarter of a millimeter. This method was empirically tested using calibrated glass beads as surrogates for raindrops (Desaulniers-Soucy, 1999).

However, between 90% and 95% of raindrops detected for this study had indistinguishable bright spots (because they were either too small or too far from the plane of focus), while larger, better focused raindrops had distinguishable spots. Hence, another method was needed to measure the size of the majority of the raindrops.

4.1. Deriving an expression for the raindrop diameter

As derived in Appendix A of Desaulniers-Soucy (1999), the intensity incident on the camera, I_{real}, coming from a spherical scatterer is proportional to the square of the scatterer diameter, a, the intensity, I_0, incident on the scatterer, the inverse square of the distance, r, between the scatterer and the camera and a function of the scattering angle $\gamma(\Theta)$ in the simple case where only the rays with zero and one internal refraction are considered:

$$I_{\mathrm{real}} = \frac{a^2}{r^2} I_0 \gamma(\Theta), \tag{3}$$

where the phase function $\gamma(\Theta)$ corresponds to the sum over the two polarizations and the number of internal reflections, p, of the following quantity (Desaulniers-Soucy, 1999):

$$\gamma(\Theta) = \sum_{p} (\varepsilon_{1,p}^2 + \varepsilon_{2,p}^2) D, \tag{4}$$

where $\varepsilon_{1,p}^2$ is the total energy squared of an incident pencil of light after p internal reflections, with polarization 1 (similarly for polarization 2) and D is sometimes called the divergence.

Eq. (4) is relatively straightforward to evaluate numerically. However, the detector used for this study—three cameras with photosensitive film—do not record the real intensity. In fact, due to diffraction, focusing and other lens defects, the optics of the camera spreads the intensity over several pixels on the negative. In addition, the film has a response curve which can be approximated, at best, by a logarithmic curve with a minimum detectable threshold. The threshold is the single effect that will most complicate the evaluation of I_{real} from the measured intensity, I_{meas}. As a very crude approximation, we calculate I_{meas} as the sum of the exponential of the pixel digitized value:

$$I_{\mathrm{meas}} = \sum_{\text{all pixels}} \mathrm{e}^{\frac{d}{64}}, \tag{5}$$

where d is the digitized pixel value (from 0 to 255) and the factor of 1/64 is introduced to make the conversion to the standard photographic densities (from 0 to 4). Note that the negative density is not a function of the intensity incident on the camera lens, but rather of the energy flux over a given area of the photosensitive emulsion, for a given exposure time. Photographers use the unit of lux for this quantity. If the camera lens, aperture, exposure, film and development parameters are the same throughout this experiment, we can keep an intensity-based derivation, the proportionality constant between the energy flux and the intensity being absorbed in the multiplicative factor, G, introduced below.

Since two flashes illuminate the scatterer, the real intensity at the camera is the sum of the intensities generated by the left and right flash, I_{l} and I_{r}, respectively. The other approximation that we make is that the measured and real intensities are related by a multiplicative factor, G, which is only a function of the position of the scatterer:

$$I_{\mathrm{meas}} = (I_{\mathrm{l}} + I_{\mathrm{r}}) G(\vec{x}), \tag{6}$$

where $\vec{x}$ is the position of the scatterer. If there is an obvious bias with respect to the diameter, it could be corrected after the multiplicative factor has been determined. This is a simple way to formulate the complex problem of identifying and calibrating many different factors influencing the measured intensity such as the shape of the reflectors, their reflective properties and the attenuation from the camera optics. These, if defined at every position in the ROI through G, are constant from one exposure to the next. Hence, the determination of G is an efficient empirical way to determine the coupled effects of all these factors.

From Eq. (3), we can write an expression for I_{meas}:

$$I_{\mathrm{meas}} = \frac{a^2}{r^2}\{I_{0,\,\mathrm{l}}(\vec{x})\gamma[\Theta_{\mathrm{l}}(\vec{x})] + I_{0,\,\mathrm{r}}(\vec{x})\gamma[\Theta_{\mathrm{r}}(\vec{x})]\}G(\vec{x}), \tag{7}$$

where $\Theta_{\mathrm{l}}(\vec{x})$, $\Theta_{\mathrm{r}}(\vec{x})$ are the scattering angles made by the camera, the scatterer and the left and right flashes, respectively, and I_{l}, I_{r} are the intensities incident on the scatterer produced by the left and right flashes, respectively.

Two factors in this equation are impossible to work out analytically: the I_0's and G. We can, nevertheless, say that the factor $(I_{0,\mathrm{l}}(\vec{x})\gamma[\Theta_{\mathrm{l}}(\vec{x})]+I_{0,\mathrm{r}}(\vec{x})\gamma[\Theta_{\mathrm{r}}(\vec{x})])$ is a function of position only (although the reflector was designed to provide an approximately constant illumination, it is in fact only approximately constant) and so is $G(\vec{x})$. We can, therefore, combine both functions and define an effective incident intensity, I_{eff}, as follows:

$$I_{\mathrm{eff}}(\vec{x}) \equiv \{I_{0,\,\mathrm{l}}(\vec{x})\gamma[\Theta_{\mathrm{l}}(\vec{x})] + I_{0,\,\mathrm{r}}(\vec{x})\gamma[\Theta_{\mathrm{r}}(\vec{x})]\}G(\vec{x}) = \frac{I_{\mathrm{meas}}r^2}{a^2}, \tag{8}$$

where $\vec{x}$ is the position of the scatterer. Using this equation, an expression for the scatterer diameter, a, can be readily found:

$$a = r\sqrt{\frac{I_{\mathrm{meas}}}{I_{\mathrm{eff}}}}. \tag{9}$$

4.2. Evaluating the effective intensity

We evaluated I_{eff} across the field of view using raindrops that show two peaks on the negative. When possible, we used the diameter measurement from the distance between the image of the light sources on the surface of large raindrops in Eq. (8) to estimate I_{eff}.

For each raindrops at a position $\vec{x}$, there is one measurement of the effective intensity; hence, one measurement of G. In addition, the drops showing two peaks are not uniformly distributed across the region of interest small drops will only show two peaks in the central region where they are in focus. Outside this region the measured intensity will fall below the threshold: in the front, the drop will be too blurry and on the sides, one intense peak will “hide” the other, we can nevertheless yield hope to find enough two peak drops even in the regions where they are scarce in order to have accurate estimates of I_{eff}. However, the discrete estimates of I_{eff} vary considerably across the ROI due to large fluctuations in the phase functions, γ. Building a field from such highly fluctuating discrete samples was not possible; thus, it was decided to remove some of the variability from the data by defining the field $G'(\vec{x})$ linked to I_{eff} in the following way:

$$I_{\mathrm{eff}}(\vec{x}) = \{\gamma[\Theta_{\mathrm{l}}(\vec{x})] + \gamma[\Theta_{\mathrm{r}}(\vec{x})]\}G'(\vec{x}), \tag{10}$$

where:

$$G'(\vec{x}) = G(\vec{x})\frac{I_{0,\,\mathrm{l}}(\vec{x})\gamma[\Theta_{\mathrm{l}}(\vec{x})] + I_{0,\,\mathrm{r}}(\vec{x})\gamma[\Theta_{\mathrm{r}}(\vec{x})]}{\gamma[\Theta_{\mathrm{l}}(\vec{x})] + \gamma[\Theta_{\mathrm{r}}(\vec{x})]}. \tag{11}$$

We observe that the ratio in Eq. (11) varies significantly from 1 only when the incident intensities, $I_{0,\mathrm{l}}$, $I_{0,\mathrm{r}}$, greatly differ. Since the reflectors were designed to provide approximately constant illumination over the ROI, we can hope to have reduced the variability in the field to be measured experimentally, namely G' .

From the set of discrete measurements of I_{eff}, the set of discrete measurements of G' has been calculated. Among the latter will be values of G' , which deviate greatly from the real value and once out of 20 times (statistically) where the diameter measurement will be greatly in error. In addition, there could be errors in the position of the drop due to an incorrect matching of images. These errors are outliers and do not obey a Gaussian distribution around the real value. They would greatly affect the mean if we were to include these estimates. In the process of constructing the field of effective intensity, we will have to pay a special attention to this fact.

The software built for analyzing the data collected during the study (Desaulniers-Soucy, 1999) was adapted to compute the field of effective intensity in order to take advantage of its viewing capacities. However, the field is defined in a pyramidal frustum, the unit cells where the field is computed share that shape. Each side of the ROI is divide by 8; therefore, there are $8^3 = 512$ unit cells. The set of values of G' in each cell is ranked and, if the number permits, the lowest and highest values of G' in the cell are discarded. The average over the remaining values is computed and the cell is given that value. This process is illustrated in the first part of Fig. 8.

This procedure is not complete since typically many cells remain empty and outliers might still exist in weakly populated cells. In order to correct for this, a running average was performed over the different cells. Each cell was given the average over its value and that of the 26 other nonempty neighbors. Again, when the number of nonempty cells permitted, the largest and smallest values were discarded from the average. For the sake of simplicity, each cell was given a uniform weight in the computation of the average. This is illustrated for one cell in the second part of Fig. 8.

The last step after the computation of G' before the computation of the raindrop diameter is to find an interpolated value of G' at the position of the raindrop. The field was created in such a way that the position of the center of the boundary cells lies at the boundary of the ROI. In other words, the field is defined inside a volume which is slightly larger than the ROI.

Since the resulting field of G' is smooth, we are not bound to choose an interpolation method that emphasizes smoothness over simplicity. Therefore, a combination of cubic spline interpolation and linear interpolation was used to compute G' for each raindrop.

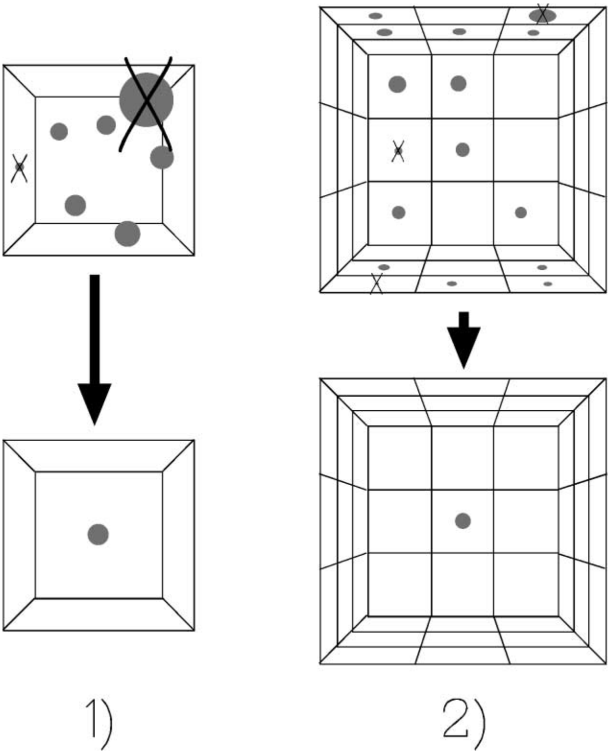

Fig. 8. Illustration of the process of (1) choosing discrete values of G' in the process of computing the average G' for a unit cell. (2) Choosing G' from neighboring cells before computing the running average.

It is interesting to compare how the two different diameter estimation methods perform on the same raindrops. Obviously, only raindrops showing two peaks can be used for this comparison, since one of the methods relies on the distance between two peaks. The diameter estimate from the distance between the two light source images is considered fairly accurate and can be used as the reference value against which to compare the effective intensity method. One point of caution remains, however. Although the interpeak distance can be measured quite accurately, it could happen that noise is interpreted as a peak or that two raindrops whose images are neighbors on the photograph are confused with the image of a single raindrop with two peaks. These situations will produce outliers and, therefore, we should not be surprised to find some of them.

The comparison takes the form of a scatter plot of all diameter estimations with the estimate from Eq. (9) as the *X*-axis coordinate and the estimate from the interpeak distance as the *Y*-axis coordinate. Ideally, the points on this plot should be very close to the $y=x$ line. If there is a systematic discrepancy from that line, it may be possible to correct for it by introducing a diameter-dependent factor in our formulation of the problem.

The scatter plot for one data set of the 0° elevation-angle geometry is shown in Fig. 9 with the ideal $y = x$ line. The most striking aspect is the large spread of the data points. This is slightly misleading since the vast majority of the points are quite close to the ideal line. Moreover, we have to consider the possibility of outliers to explain the amplitude of the spread. Another obvious feature is the scarcity of points for low diameters. This is to be expected, however, since it takes at least three horizontal pixels to resolve a two-peak raindrop (one high, one low and one high). At 6 m from the camera, this situation can be fulfilled with drops that have a minimum diameter of 0.75 mm.

The only systematic discrepancy from that line can be seen for raindrops with diameters as estimated from I_{eff} between $3/4$ and 1 mm. The interpeak distance yields a diameter estimate systematically lower. There are two good reasons to avoid correcting for this effect. Firstly, a correction would probably yield a negative intercept, meaning that raindrops barely showing on the photograph would have a negative diameter! If anything, it should be positive. Secondly, this bias might be created by an artifact of the detection threshold. If we consider raindrops with a small diameter, those that will show two peaks are those that receive the most illumination. Therefore, based on the measured intensity, they will appear larger and would, therefore, be shifted right on Fig. 9. This could explain this bias.

From the scatter plots of both geometries, the uncertainty on the raindrop diameter as estimated from I_{eff} can be approximated visually as 0.5 mm, 19 times out of 20. Due to the

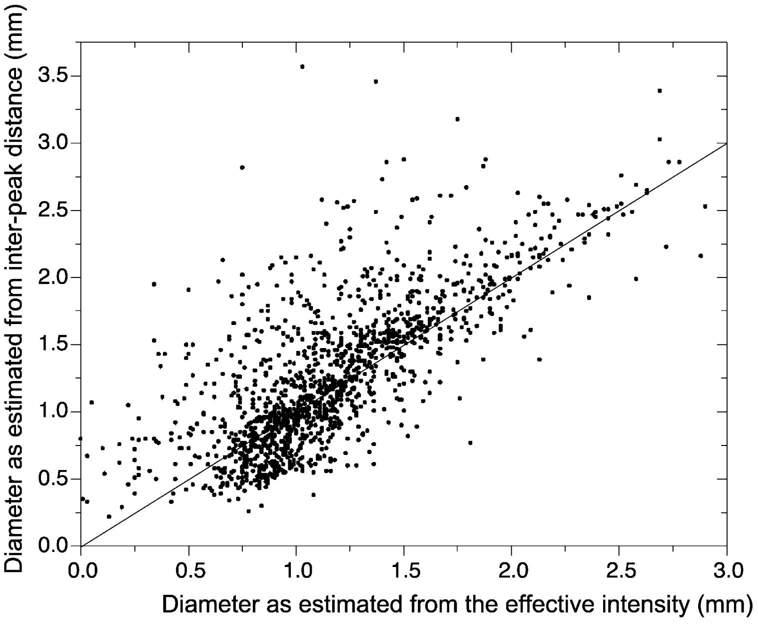

Fig. 9. Comparison of the two diameter-evaluation methods for a typical data set. The solid line represents the ideal situation when both methods estimate the same diameter.

uncertainty associated with the interpeak distance method and the possibility of mistakes, no further elaboration of the method was used to evaluate the uncertainty.

5. Results

5.1. The data sets

The first high enough quality photographs were taken on August 23, 1996, the last ones on September 27, 1998. The objective was to take as many different rain and snow events as possible. However, a minimum threshold on precipitation level had to be met in order to justify taking photographs of an event, since otherwise there would not be many raindrops above the detection level and the statistical description of this event would be poor. The experimental setup allowed pictures to be taken only at night, restricting further the number of events that could be photographed, especially in the middle of the summer—due to the shortness of the night.

Nevertheless, approximately 200 films were taken during the 2-year period of the experiment. Considering that there are three films taken simultaneously (one for each of the three cameras) and that synchronization and other problems produced a success rate conservatively estimated at 2/3, there is the possibility of analyzing 200 [films] × 2/3 × 12 [exp/film]/3 [exp/scene] = 533 scenes of rain and snow altogether. Among these scenes, some additional problems may have occurred, which would only be visible upon digitization and analysis.

A subset of a typical photograph is shown in Fig. 10. It has to be blown up; otherwise, the raindrop images would be indistinguishable from printing defects. The alarm clocks

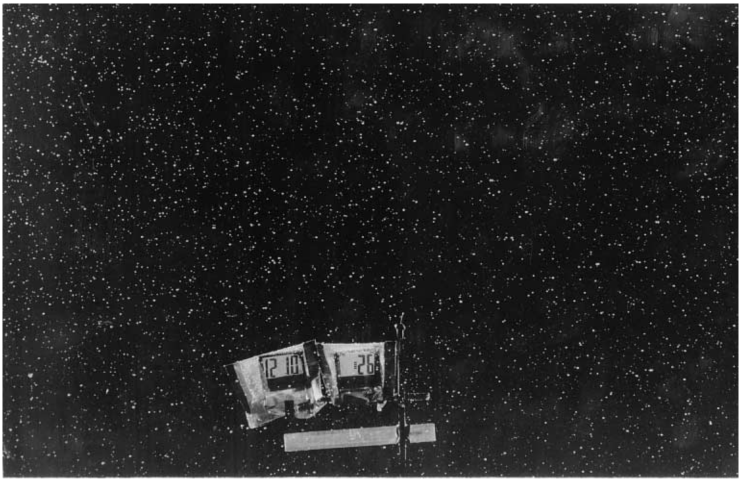

Fig. 10. Blown-up section of a photograph of a typical rain event, including the alarm clocks. The clocks were approximately 12 cm wide and were placed at the center of the region of interest.

are shown in order to provide a relative scale. The summary of the reconstructed data from the six original data sets is given in Table 1. While the top, bottom, left and right boundaries were set by the imaging equipment, no physical limit exists along the field of view, other than the plane in which the cameras lie. A trade off between the most uniform lighting conditions inside the ROI and the size of the hydrometer-full volume to be analyzed was made by truncating the ROI along the Z-axis (in a direction parallel to the field of view). The coordinates of the planes bounding the ROI along Z is indicated in Table 1.

5.2. Visualization

As an example, two of the data sets of Table 1 will be shown. Since the three-dimensional field of liquid water density and the set of raindrop positions and sizes are all available, the possibilities for visualization are numerous. The field of liquid water content is shown in color in (Desaulniers-Soucy, 1999). It was decided, however, mainly because of legibility, to restrict the illustrations to the position of the raindrops with symbols providing some idea of the size of each raindrop in the region of interest.

It was decided to show the data sets as projections because it enables an easy visualization of geometrical biases in the raindrop diameter-estimation method. In this respect, the projection along vertical axis is particularly revealing, because all the interesting optics lie in an horizontal plane. The first data set shown is representative of the f142 geometry (Fig. 11). Note that the limits of the field of view of the cameras are clearly visible.

In order to identify geometrical biases, one must be careful when looking at the images of raindrops. This is because the ROI is wider at larger distances from the camera. Hence, it is normal to find more raindrops at a larger distance from the camera.

The data set f207_12h36min25s will serve as an example of the f207 geometry. The projection along the vertical axis is illustrated in Fig. 12. The most striking difference in the f142 projections is that there are many more raindrops with “visible” size (diameter between 1 and 2 mm). This impression is increased by the choice of the limits of the ROI: from 4.5 to 7 m instead of 4 to 6.5 m for the f142 data sets.

Table 1
Final reconstruction data for all six data sets studied

Data set identification	Total no. of raindrops	ROI boundaries along Z	No. of raindrops in ROI	ROI volume (m^3)
f142_07h01min28s	44938	[4 m, 6.5 m]	14047	8.90
f142_07h05min49s	27730	[4 m, 6.5 m]	11064	8.98
f142_07h06min10s	30059	[4 m, 6.5 m]	11311	8.55
f207_12h35min29s	26850	[4.5 m, 7 m]	12752	10.19
f207_12h36min25s	29889	[4.5 m, 7 m]	14777	10.69
f207_12h38min50s	48269	[4.5 m, 7 m]	22146	10.66

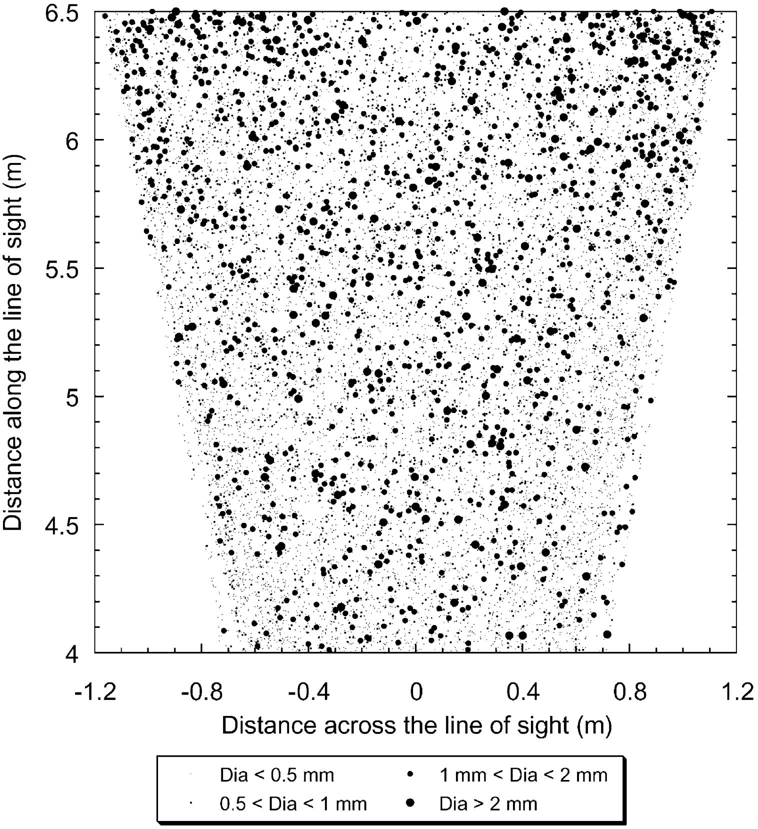

Fig. 11. Projection along the vertical axis of the raindrop positions inside the region of interest, from $z=4$ m to $z=6.5$ m, for the f142_07h01min28s data set. Circles of varying diameters are meant to provide an idea of the size of the raindrops.

5.3. Size distributions

One of the standard tests of raindrop detectors is to determine the histogram of the raindrop size distribution. Due to the historical importance of the Marshall–Palmer raindrop size distribution (Marshall and Palmer, 1948), $N(D)=N_0e^{-\Lambda D}$, where $N(D)\mathrm{d}D$ is the number of drops per unit volume with diameters between D and $D+\mathrm{d}D$ and Λ is a parameter dependent upon the nominal rainfall rate, the histograms are usually plotted with semilog axes. However, we prefer to present the data as exceedance probability histograms plotted with log–log axes. The reason behind this choice comes from the following considerations.

First, express the probability distribution in the following way:

$$\Pr(V > v) = \mathrm{e}^{-c(\log v)}, \qquad (12)$$

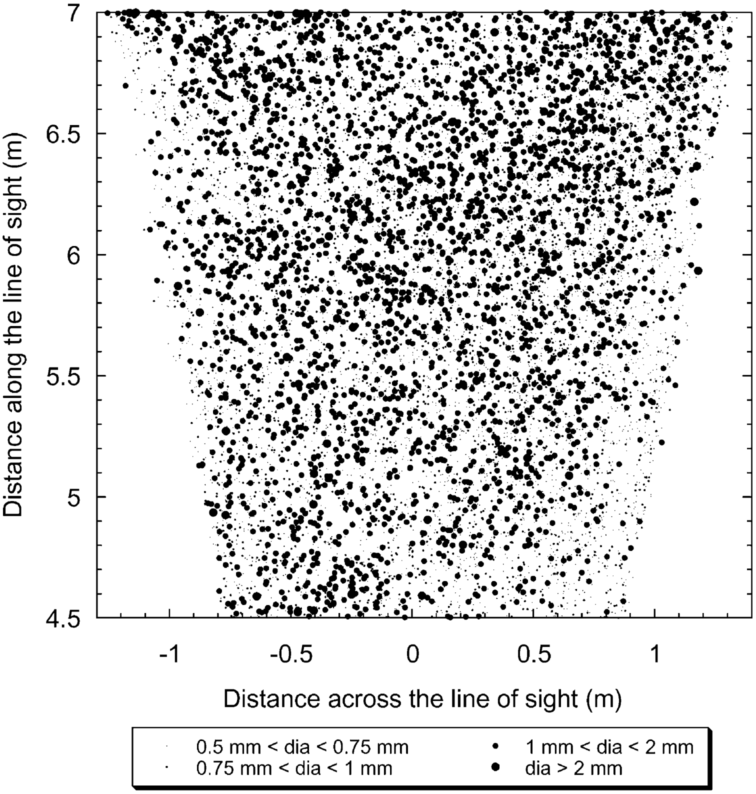

Fig. 12. Projection along the vertical (x) axis of the raindrop positions inside the region of interest, from $z = 4.5$ m to $z = 7$ m, for the f207_12h36min25s data set. Circles of varying diameters are meant to provide an idea of the size of the raindrops. Raindrops whose diameter was smaller than 0.5 mm were too numerous to be displayed.

where v is the drop volume, so that c' is the logarithmic derivative ("Pr" indicates "probability"). The qth moment is:

$$\langle v^q \rangle = \int v^q \mathrm{dPr} = \int c' \, \mathrm{e}^{q\log v - c(\log v)} \mathrm{d}\log v. \tag{13}$$

We now use the "moving maximum" method (Bender and Orszag, 1978), which requires that the maximum is sufficiently "peaked". We conclude that the maximum contribution to the qth moment comes from raindrops of volumes corresponding to $q = c'$. This is why we have superimposed on the histograms a line of slope equal to -1 in order to give an estimate of the raindrops contributing most to the mean. In addition, Eq. (13) is valid irrespective of the resolution for moment $q = 1$. A line of slope -2 was also added to the graph in order to estimate the number of raindrops contributing mostly to the variance.

The exceedance probability histogram for the combined f142 data sets can be seen in Fig. 13. Following the previous discussion, tangents with slope equal to -1 and -2 have been added to the graph. It yields two conclusions as follows.

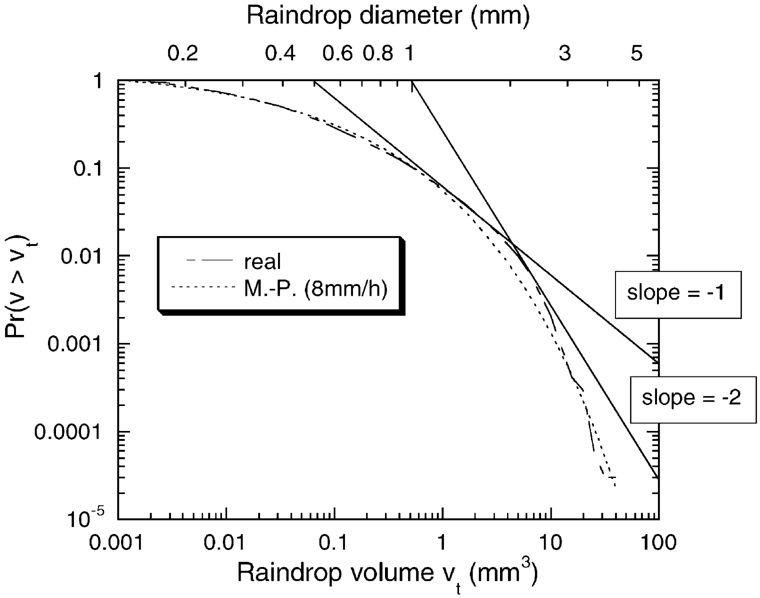

Fig. 13. Exceedance probability histogram representing the set of raindrops in the ROI for the three combined f142 data sets. The raindrop volume and diameter are plotted with a logarithmic scale on the bottom and top axes, respectively. Tangents with slopes equal to -1 and -2 have been added to the graph. For comparison, a Marshall–Palmer distribution with average rain rate equal to 8 mm/h has been plotted on the same graph.

- The raindrops that contribute the most to the mean volume have a diameter equal to approximately 1.5 mm. Raindrops with this (or larger) diameter constitute approximately 3% to 4% of the total.
- The raindrops that contribute the most to the variance have diameters close to 2 mm. The latter does not make up much more than 0.5% of the total, roughly 175 raindrops among the three data sets. The highest moment, q_s, that can be reliably estimated from these data sets, given by the most negative slope on the graph, is approximately 3.5.

These numbers have a practical relevance in the measurement and computation of quantities like Z_e or Z in radar meteorology (Desaulniers-Soucy, 1999), which involve the second power of the raindrop volumes. As a consequence of what is presented here, they roughly share the same convergence properties as the variance. These notions are explained in Desaulniers-Soucy (1999). In addition, it can be seen that the maximum observed raindrop diameter is close to 4 mm.

The famous Marshall–Palmer (MP) distribution has been computed with an average rainfall rate of 8 mm/h. and the resulting cumulative probability distribution is plotted on Fig. 13, together with the real distribution of the f142 data sets. The MP distribution is defined as follows (Rogers and Yau, 1989):

$$N(D) = N_0 e^{-\varsigma D}, \qquad (14)$$

where $N(D)\mathrm{d}D$ is the number of drops per unit volume with diameters between D and $D+\mathrm{d}D$ and N_0 is 0.08 cm^{-4}. Marshall and Palmer found that the function, ς, depends only on rainfall rate, R_r, and is given by:

$$\varsigma(R_r) = 41R_r^{-0.21}, \tag{15}$$

where ς has units of cm^{-1} and R_r is measured in mm/h. It can be noted that the size distribution of the f142 data sets corresponds quite closely to a MP distribution. Since deviations from the MP distribution have often been observed, this criterion is not a requirement. However, it was nice to observe that the size estimation performance of HYDROP produced a distribution that is plausible.

We also checked the consistency of the observed size distribution across the individual data sets of f142. To this end, all three individual distributions have been plotted on Fig. 14. The distributions of the data sets taken within 30-s interval (07h05min49s and 07h06min10s) show remarkable similarity; the results seems to be robust. The third data set, whose photographs were taken 4 min earlier, yields a distribution which is quite close to the other two. Hence, it can be said that the size distribution of the combined data set is representative of each of its constituents.

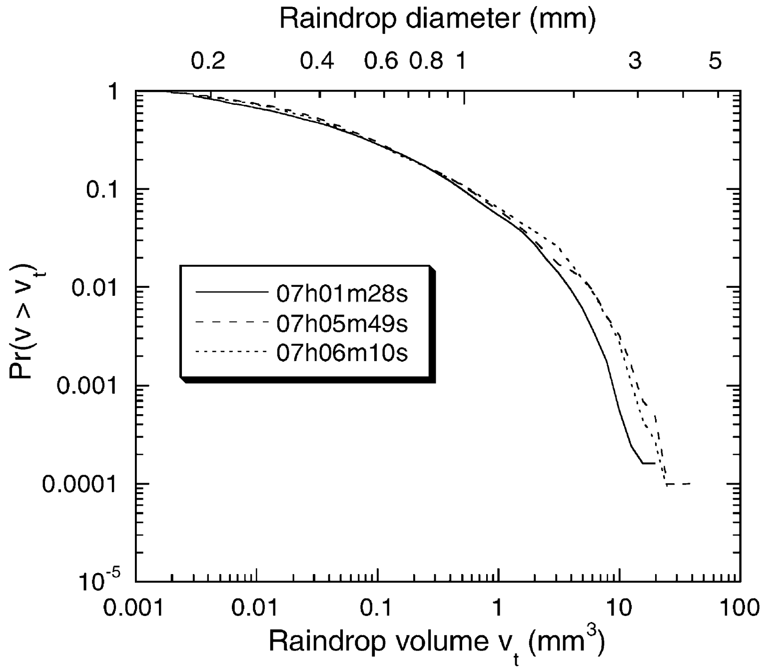

Fig. 14. Exceedance probability histogram representing the set of raindrops in the ROI for each individual f142 data set. The raindrop volume and diameter are plotted with a logarithmic scale on the bottom and top axes, respectively.

The exceedance probability distribution of the f207 data sets differs slightly from that of the f142 data set, as Fig. 15 shows. The mean is dominated by raindrops of diameter slightly larger than 1 mm and these (or larger) represent approximately 15% of all raindrops. As a consequence, we expect the fluctuations around the mean to be less wild than those of f142 and the convergence of the mean should be improved.

However, there is a bump in the distribution for very large raindrops, with diameters between 3 and 5 mm. There are only a handful of them, but they will make the determination of the variance more problematic. The highest moment that can be determined reliably from these data sets is approximately $q=2$. In order to verify if the relatively high occurrence of the very large raindrops is shared among the three f207 data sets, the three size distributions have been plotted together in Fig. 16. The answer is quite obvious: suspiciously, only one data set has this feature. Therefore, if some inconsistencies are to be found in the analyses, it could be worth comparing the results of the combined analyses of two (without the suspicious data set) and three data sets.

5.4. Nominal rainfall rate

A nominal rainfall rate was estimated using the assumption that all raindrops fall at their terminal velocity in still air. The semiempirical expressions for terminal velocities

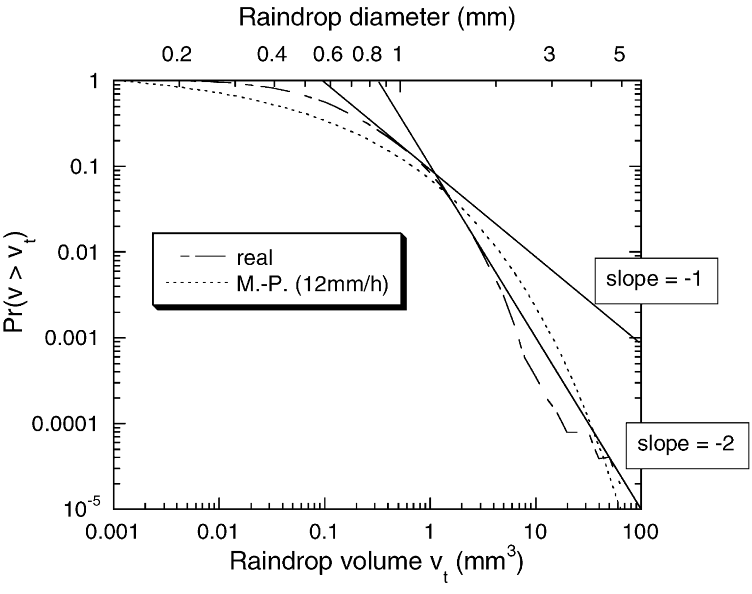

Fig. 15. Exceedance probability histogram representing the set of raindrops in the ROI for the three combined f207 data sets. The raindrop volume and diameters are plotted with a logarithmic scale on the bottom and top axes, respectively. Tangents with slope equal to − 1 and − 2 have been added to the graph. For comparison, a Marshall–Palmer distribution with average rain rate equal to 12 mm/h has been plotted on the same graph. According to this figure, the variance is not well defined.

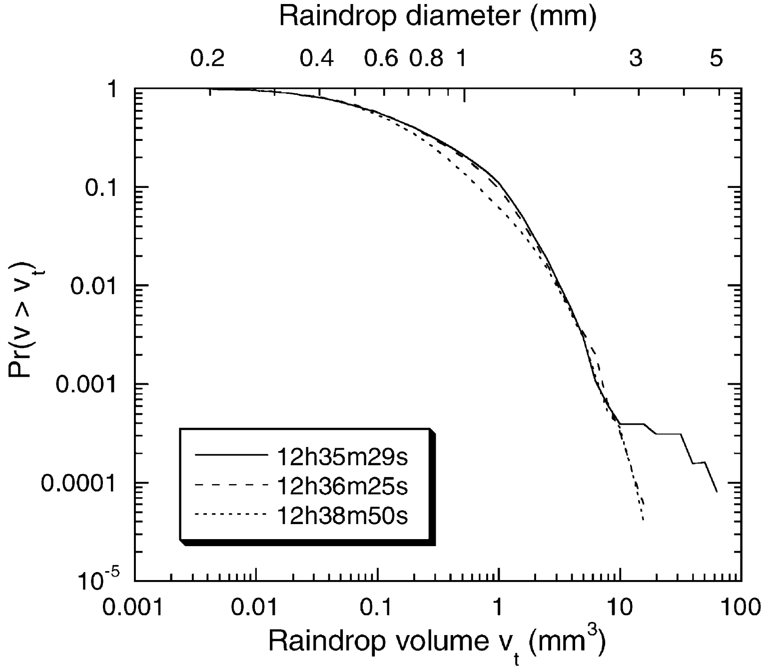

Fig. 16. Exceedance probability histogram representing the set of raindrops in the ROI for each individual f207 data set. The raindrop volume and diameter are plotted with a logarithmic scale on the bottom and top axes, respectively.

used in this study were taken from (Rogers and Yau, 1989) and can be summarized as follows:

$$u = k_3 r \qquad 40\ \mu\text{m} < r < 0.6\ \text{mm} \tag{16}$$

$$u = k_2 r^{1/2} \qquad 0.6\ \text{mm} < r < 2\ \text{mm},$$

where r is the raindrop radius and $k_2 = 2.01 \times 10^3$ cm$^{1/2}$ s^{-1}, $k^3 = 8 \times 10^3$ s^{-1} (Beard, 1976).

As was shown in the previous section, the maximum raindrop radius encountered in this study is 2 mm; therefore, we should not be worried about the upper limit of validity of Eq. (16). The lower limit is not a problem either since raindrops of radius smaller than 40 μm, although present in the data set, are not numerous enough to contribute significantly to the rainfall rate. Therefore, the first expression of Eq. (16) was applied to any raindrop with radius smaller than 0.6 mm.

The procedure used to estimate the rainfall rate from the data sets f142 and f207 is described as follows:

- A cube is defined inside the ROI for each data set. This cube is centered along Z with respect to the clock position ($z = 5.85$ m for f142 and $z = 5.95$ m for f207).
- The height of the bottom of this cube is computed.

- The time taken for each drop to fall to that height is computed.
- The time axis is divided into bins and a histogram is made from these fall times.
- The amount of water that fell through the bottom of the cube is computed for each bin, and the rainfall rate is computed from this sum and the area of the bottom face of cube. It is further converted to units of mm per hour.

Note that there is an upper time limit to the validity of the nominal rainfall rate since, eventually, the highest drops will have fallen through the bottom boundary. The validity limit is given by the time required by a drop of diameter equal to the maximum diameter found in the size histogram, located at the top of the cube, to fall through the cube bottom boundary.

The nominal rainfall rate calculated from a data set in the f142 geometry is shown in Fig. 17. The cube side was equal to 1.57 m and 6256 raindrops were located inside this cube. The average rainfall rate is approximately 6.5 mm/h. The other two data sets have similar nominal rainfall rates. An example of the rainfall rate in the f207 geometry is given in Fig. 18. The cube side is equal to 1.72 m and there were 7996 raindrops in that cube. An average rainfall rate for this data set is approximately 11 mm/h. One of the other two f207 data sets shows a similar average rainfall rate and the other,

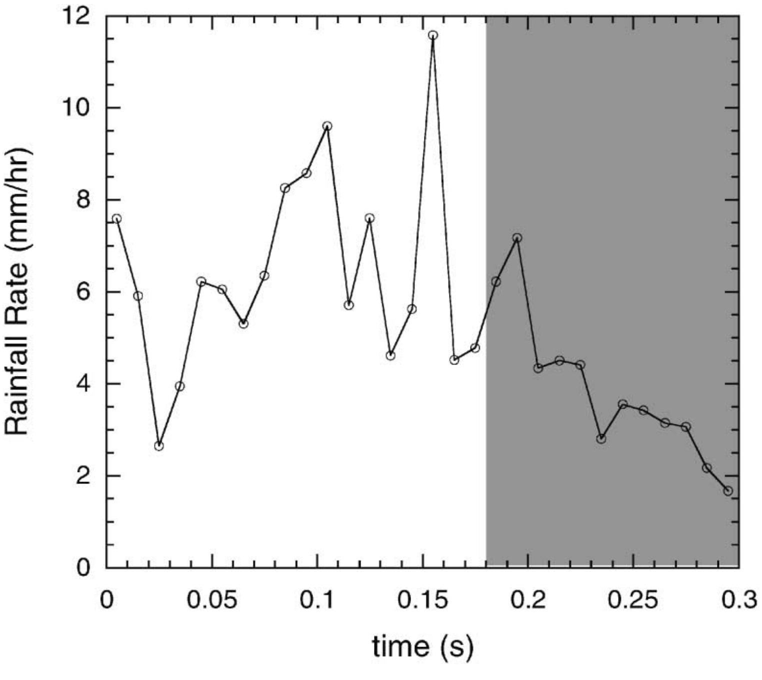

Fig. 17. Rainfall rate as calculated from the raindrops detected in a cube of side 1.57 m inside the ROI of f142_07h01min28s. Considering a maximum raindrop terminal velocity of 8.83 ms^{-1} (for 4-mm diameter), the rainfall rate is only valid for times smaller than approximately 0.18 s.

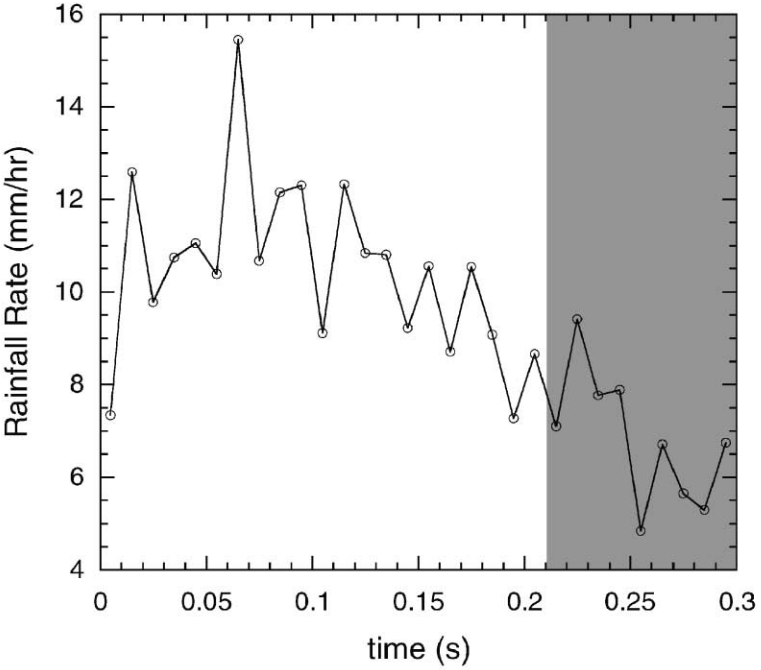

Fig. 18. Rainfall rate as calculated from the raindrops detected in a cube of side 1.72 m inside the ROI of f207_12h36min25s. Considering a maximum raindrop terminal velocity of 8.26 ms^{-1} (for 3-mm diameter), the rainfall rate is only valid for times smaller than approximately 0.21 s.

f207_12h38min50s, has a slightly larger rainfall rate, 14 mm/h, with many more raindrops in the cube: 12 408.

6. Conclusions

We designed a novel experiment to determine the position and size of hydrometeors in a region with a large enough number of raindrops such that the continuum limit of rain should reasonably have been reached (approximately 10 m^3, more than 10^4 particles). While previous empirical determinations of the properties of rainfall have been from the large scale down to at best several meters, the scales from the individual drops on up to meters have received comparatively little attention. In order to study rainfall at these scales, a detector was designed and built, the HYdrometeor Detection and Ranging by stereO-Photography (HYDROP) detector, whose operation enabled the determination of a large number of raindrop positions and volumes. The detector had to have a large enough resolution to lower the raindrop diameter-detection threshold as much as possible at the same time as recording this data in a large enough volume to study the convergence properties explained above.

As a result, HYDROP successfully detected the positions and sizes of tens of thousands of raindrops in a volume of ≈ 10 m^3. The uncertainty on raindrop position measurements was estimated to be 3 mm in the plane perpendicular to the line of sight and to be 3 cm along the line of sight. The uncertainty in raindrop diameter measurements is 0.25 mm for the raindrops whose size was estimated by the interpeak distance and 0.5 mm for the estimates done through the effective intensity method. Note that, in the case of the effective intensity method, the uncertainties quoted here are valid only for raindrops whose size is larger than 1 mm. No reliable estimates of the uncertainties for smaller raindrops was possible using the calibration method adopted in this paper.

The HYDROP detector is capable of taking 12 exposures at 30-s intervals before reloading the photographic film. The experiments must be conducted at night since darkness is important to resolve raindrops on the film (by minimizing background light). The position of the flashes can be adjusted to optimize the size detection threshold, the volume of the Region of Interest (ROI) or the different types of hydrometeors. Each element of the apparatus is mobile; therefore, it can be changed from one geographic location to another relatively easily (in spite of its size).

The key step in the position determination is the identification of the same drop on two or more of the negatives. This is achieved by a correlation algorithm. The performance of the correlation algorithm has been tested successfully with rainfall rates ranging from 5 to 15 mm/h. The critical parameters for this performance are the density of rain drop images on the photographs and the precision of the projection matrix calibration. The calibration method used was based on large raindrops present in all photographs and one absolute distance from the camera to a fixed point in the field of view. The algorithms for correlation images of the same raindrops and for computing their position was designed to be effective even with photographs of intense rain events.

More than 500 scenes were recorded (on triplets of negatives) during several rain and snow events, but only six triplets of rain were analyzed in the preliminary stage of this study. The six scenes came from two different rain events, one in November 1996 (f142) and one in July 1997 (f207). The number of raindrops detected in the ROIs varied from 11 064 to 22 146. The rainfall rate was estimated from each of these data sets. The exceedance size probability distribution of the f142 data sets were very close to the Marshall–Palmer (MP) distribution with a rainfall rate of 8 mm/h. The distribution of f207 raindrop volumes showed that small raindrops were less abundant than the prediction of MP with a rainfall rate of 12 mm/h (and inversely for larger raindrops).

The HYDROP experiment has made the first determinations of raindrop size information in a significant sample volume. Although the original motivation was to study the approach of the continuum limit from the small scales to the large—the subject of a paper in preparation—this new type of information has many other possible uses. For example, we expect that it should be possible to borrow methods originally developed in high-energy physics to perform searches for products of drop collisions and, hence, to do cloud physics experiments which until now had only been possible in the laboratory. It will also be possible to directly simulate the radiative properties of rain at various wavelengths including microwaves and visible; this is important for remote-sensing of rain.

Acknowledgements

We are extremely grateful to Mr. Peter B. Stetson of the Dominion Astrophysical Observatory in Victoria for digitizing the photographic negatives used in this project. We are equally grateful to Mr. Robert Nowak for relentless support at various levels during the duration of this study.

The financial assistance from the Faculty of Graduate Studies and Research, the Atmospheric Environment Services, the National Science and Engineering Research Council and the Fonds pour la formation des chercheurs et d'aide à la recherche was greatly appreciated.

References

Albrecht, B.A., Randall, D.A., Nicholls, S., 1988. Observations of marine stratocumulus clouds during FIRE. Bull. Am. Meteorol. Soc. 69, 618–626.

Austin, P.H., Baker, M.B., Blyth, A.M., Jensen, J.B., 1985. Small scale variability in warm continental cumulus clouds. J. Atmos. Sci. 42 (11), 1123–1138.

Baker, B.A., 1992. Turbulent entrainement and mixing in clouds: a new observational approach. J. Atmos. Sci. 49 (5), 387–404.

Beard, K.V., 1976. Terminal velocity and shape of cloud and precipitation drops aloft. J. Atmos. Sci. 33, 851–864.

Bender, C.M., Orszag, S.A., 1978. Advanced Mathematical Methods for Scientists and Engineers. McGraw-Hill, San Francisco.

Borrmann, S., Jaenicke, R., Neumann, P., 1993. On spatial distributions and inter-droplet distances measured in stratus clouds with in-line holography. Atmos. Res. 29, 229–245.

Brenguier, J.-L., 1993. Observations of cloud microstructure at the centimeter scale. J. Appl. Meteorol. 32, 783–793.

Brösamlen, G., 1994. Radiative Transfer in Lognormal Multifractal Clouds and Analysis of Cloud Liquid Water Data. M. Sc. Thesis, McGill University.

Cahalan, R.F., 1994. Bounded cascade clouds: albedo and effective thickness. Nonlinear Proc. Geophys. 1, 156–167.

Chigirinskaya, Y., Schertzer, D., 1996. Cascade of scaling gyroscopes: Lie structure, universal multifractals and self-organized criticality in turbulence. In: Molchanov, S.A., Woycynski, W.A. (Eds.), Stochastic Models in Geosystems. Springer, New York, pp. 57–81.

Davis, A., Marshak, A., Wiscombe, W., Cahalan, R., 1996. Scale invariance of liquid water in marine stratocumulus: Part I: Spectral properties and stationary issues. J. Atmos. Sci. 53, 1538–1560.

Desaulniers-Soucy, N., 1999. Empirical test of the multifractal continuum in rain, PhD Thesis, McGill University, Montréal.

Faugeras, O.D., 1993. Three-dimensional computer vision. A Geometric Viewpoint. MIT Press, Boston, MA.

Goodman, J.W., 1996. Introduction to Fourier Optics, 2nd edn. McGraw-Hill, New York.

Gutpa, V.J., Waymire, E., 1990. Multiscaling properties of spatial rainfall and river flow distributions. J. Geophys. Res. 95, 1999–2010.

Houze, R.A., 1981. Structures of atmospheric precipitation systems: a global survey. Radio Sci. 16, 671–689.

Jameson, A.R., Kostinski, A.B., 1998. Fluctuation of properties of precipitation: Part II. Reconsideration of the meaning and measurement of raindrop size distributions. J. Atmos. Sci. 55, 283–294.

Jameson, A.R., Kostinski, A.B., 1999. Fluctuation properties of precipitation: Part IV. Fine scale clustering of drops in variable rain. J. Atmos. Sci. 59, 82–91.

Jameson, A.R., Kostinski, A.B., Black, R.B., 1998. The texture of clouds. J. Geophys. Res. 103 (D6), 6211–6219.

Jones, D.M.A., 1992. Raindrop spectra at the ground. J. Appl. Meteorol. 31, 1219–1225.

Jones, D.M.A., Dean, L.A., 1953. A Raindrop Camera. Res. Rep. 3, US Army contract DA-36-039 SC-42446, Ill. State Water Surv., Urbana.

Joss, J., Waldvogel, A., 1967. Ein Spektrograph für Niederschlagstrofen mit Automatischer Auswertung. Pure Appl. Geophys. 68, 240–246.

Kostinski, A.B., Jameson, A.R., 1997. Fluctuation properties of precipitation: Part I. On deviations of single-size drop counts from the Poisson distribution. J. Atmos. Sci. 54, 2174–2186.

Kozikowska, A., Haman, K., Supronowicz, J., 1984. Preliminary results of an investigation of the spatial distribution of fog droplets by a holographic method. Q. J. R. Meteorol. Soc. 110, 63–73.

Lavergnat, J., Golé, P., 1998. A stochastic raindrop time distribution model. J. Appl. Meteorol. 37, 805–818.

List, R., Hudak, D., Nissen, R., Tung, N.P., Soo, S.K., Kang, T.S., 1988. Raindrop size distributions in warm rain at Penang. In: Theon, J.S., Fugono, N. (Eds.), Tropical Rainfall Measurement. Deepak Publishing, New York, pp. 271–278.

Lovejoy, S., 1981. Analysis of rain areas in terms of fractals. 20th Conf. on Radar Meteorology. AMS, Boston, pp. 476–484.

Lovejoy, S., Schertzer, D., 1985. Generalised scale invariance and fractal models of rain. Water Resour. Res. 21, 1233–1250.

Lovejoy, S., Schertzer, D., 1986. Scale invariance, symmetries, fractals and stochastic simulations of atmospheric phenomena. Bull. AMS 67, 21–32.

Lovejoy, S., Schertzer, D., 1990a. Fractals, raindrops and resolution dependence of rain measurements. J. Appl. Meteorol. 29 (11), 1167–1170.

Lovejoy, S., Schertzer, D., 1990b. Multifractals, universality classes, satellite and radar measurements of clouds and rain. J. Geophys. Res. 95, 2021–2034.

Lovejoy, S., Schertzer, D., 1991. Multifractal analysis techniques of rain and cloud fields from 10^3 to 10^6 m. In: Schertzer, D., Lovejoy, S. (Eds.), Scaling, Fractals and Non-Linear Variability in Geophysics. Kluwer Academic Publishing, Dordrecht, pp. 111–144.

Lovejoy, S., Schertzer, D., 1995a. How bright is the coast of Brittany? In: Wilkinson, G. (Ed.), Fractals in Geoscience and Remote Sensing. Office for Official Publications of the European Communities, Luxemburg, pp. 102–151.

Lovejoy, S., Schertzer, D., 1995b. Multifractals and rain. In: Kundzewicz, A.W. (Ed.), New Uncertainty Concepts in Hydrology and Hydrological Modelling. Cambridge Univ. Press, Cambridge, pp. 62–103.

Lovejoy, S., Schertzer, D., Tsonis, A.A., 1987. Functional box-counting and multiple elliptical dimensions in rain. Science 235, 1036–1038.

Lovejoy, S., Schertzer, D., Stanway, J.D., 2001. Direct evidence of atmospheric cascade dynamics from 1 to 20,000 km. Phys. Rev. Lett., submitted for publication.

Malinowski, S.P., Zawadzki, I., 1993. On the surface of clouds. J. Atmos. Sci. 50 (1), 5–13.

Marshall, J.S., Palmer, W.McK., 1948. The distribution of raindrops with size. J. Meteorol. 5, 165–166.

Olsson, J., Niemczynowicz, J., Berndtsson, R., 1993. Fractal analysis of high-resolution rainfall time-series. J. Geophys. Res. 98 (D12), 23265–23274.

Paluch, I.R., Baumgardner, D.G., 1989. Entrainment and fine-scale mixing in a continental convective cloud. J. Atmos. Sci. 46 (2), 271–278.

Rogers, R.R., Yau, M.K., 1989. A Short Course in Cloud Physics, 3rd edn. Pergamon, New York.

Sachs, D., Lovejoy, S., Schertzer, D., 2001. The multifractal scaling of cloud radiances from 1 m to 1 km. Fractals (in press).

Schertzer, D., Lovejoy, S., 1985. Generalised scale invariance in turbulent phenomena. Phys. Chem. Hydrodyn. 6, 623–635.

Schertzer, D., Lovejoy, S., 1987. Physically based rain and cloud modelling by anisotropic, multiplicative turbulent cascades. J. Geophys. Res. 92, 9692–9714.

Schertzer, D., Lovejoy, S., Schmitt, F., Chigirinskaya, Y., Marsan, D., 1997. Multifractal cascade dynamics and turbulent intermittency. Fractals 5, 427–471.

Schönhuber, M., Urban, H., Poiares Baptista, J.P.V., Randeu, W.L., Riedler, W., 1994. Measurement of precipitation characteristics by a new distrometer. Proceedings, Atmospheric Physics and Dynamics in the Analysis and Prognosis of Precipitation Fields, Rome, Italy.

Tessier, Y., Lovejoy, S., Schertzer, D., 1993. Universal multifractal in rain and clouds: theory and observations. J. Appl. Meteorol. 32, 223–250.

Tessier, Y., Lovejoy, S., Schertzer, D., 1994. Multifractal analysis and simulation of the global meteorological network. J. Appl. Meteorol. 33 (12), 1572–1585.

Uijlenhoet, R., Stricker, J.N.M., Torfs, P.J.J.F., Creutin, J.-D., 1999. Towards a stochastic model of rainfall for radar hydrology: testing the Poisson homogeneity hypothesis. Phys. Chem. Earth (B) 24, 747–755.