

Accuracy of rainfall measurement for scales of hydrological interest

S.J. Wood, D.A. Jones and R.J. Moore

Centre for Ecology and Hydrology, Wallingford, Oxon, OX10 8BB, UK
e-mail for corresponding author: rm@ceh.ac.uk

Abstract

The dense network of 49 raingauges over the 135 km² Brue catchment in Somerset, England is used to examine the accuracy of rainfall estimates obtained from raingauges and from weather radar. Methods for data quality control and classification of precipitation types are first described. A super-dense network comprising eight gauges within a 2 km grid square is employed to obtain a “true value” of rainfall against which the 2 km radar grid and a single “typical gauge” estimate can be compared. Accuracy is assessed as a function of rainfall intensity, for different periods of time-integration (15 minutes, 1 hour and 1 day) and for two 8-gauge networks in areas of low and high relief. In a similar way, the catchment gauge network is used to provide the “true catchment rainfall” and the accuracy of a radar estimate (an area-weighted average of radar pixel values) and a single “typical gauge” estimate of catchment rainfall evaluated as a function of rainfall intensity. A single gauge gives a standard error of estimate for rainfall in a 2 km square and over the catchment of 33% and 65% respectively, at rain rates of 4 mm in 15 minutes. Radar data at 2 km resolution give corresponding errors of 50% and 55%. This illustrates the benefit of using radar when estimating catchment scale rainfall. A companion paper (Wood *et al.*, 2000) considers the accuracy of rainfall estimates obtained using raingauge and radar in combination.

Keywords: rainfall, accuracy, raingauge, radar

Introduction

Hydrological modelling and forecasting require rainfall data as one of their most important inputs. Often this will be required as an areal average estimate, the two most relevant scales of interest being the catchment scale and the size of the square grid used by weather radar (here, 2 km). In general, these areal average rainfall values will be estimated from a point measurement, for example using a raingauge, or an equivalent areal observation such as can be obtained from weather radar. The accuracy of all forms of rainfall measurement is of importance, particularly to hydrological modelling applications (Larson and Peck, 1974; Peck, 1980; Moore, 1998), and this work aims to quantify the accuracy of areal rainfall measurement by raingauges and by radar. For small flashy catchments a small time-step is necessary to model the flood response and even for larger catchments benefits can be gained from frequent model updating using flow observations (Austin and Moore, 1996). A time-step of 15 minutes forms the focus of attention here since this is commonly used in the UK for flow forecasting purposes. However, it is of fundamental importance to know how accurately rainfall can be measured over different time-scales and this is also investigated.

The investigation of rainfall accuracy reported here is made possible through the use of the specially designed

HYREX raingauge network and associated weather radars. This dense network of 49 tipping-bucket raingauges installed across the Brue catchment, in Somerset, South-West England, is shown in Fig. 1. Continuously scanning C-band radars at Wardon Hill, 30 km south of the catchment, and at Cobbacombe Cross, 70 km to the west, are used here as sources of radar rainfall data. The configuration of the raingauge network is such that 19 of the gauges are located singly within 2 km radar grid squares. There are two SW-NE lines of 4 squares each containing at least two gauges, with the alignment chosen to be along the prevailing storm direction and orthogonal to the topography. In addition, there are two higher density grid squares, each containing 8 gauges, in areas of low and high relief. The two high density grid squares are designed to give best estimates of mean rainfall over a 2 km pixel. The size of the tipping-bucket is 0.2 mm and the time-of-tip is recorded to the nearest ten seconds. In addition, the Bridge Farm raingauge site (Gauge 8) also includes an automatic weather station and a soil moisture station on a permanent basis. The associated relief is shown in Fig. 2 which indicates that the higher ground is both to the east and to the north of the catchment. This higher ground lies in the path of weather systems which move in a north-easterly direction and can induce low level growth of precipitation, sometimes referred to as orographic enhancement.




Fig. 1. The HYREX dense raingauge network over the Brue catchment.

An important step in the use of the HYREX dataset for investigating the accuracy of rainfall measurements from raingauges and radar is the initial quality control of the data. Indeed, instrument malfunction can be viewed as an im-

portant facet of the overall rainfall measurement accuracy problem. Data quality control is discussed in detail in the next section before proceeding to the main part of the paper dealing with the quantification of rainfall measurement




Fig. 2. Relief over the Brue catchment.

accuracy at different scales in space and time. Quality control extends to the identification of precipitation type, to allow later analyses to differentiate accuracy in stratiform and convective rain, and to suppress the use of data at times when precipitation falls in solid form.

Data quality control

The dense network of raingauges over the Brue catchment, whilst providing an excellent experimental facility in support of rainfall research, presents a formidable challenge in terms of data quality control. The integrity of data obtained from HYREX will be of vital importance to subsequent analyses. Periods of invalid data, when a gauge is not performing properly, are difficult to identify. Diagnostic reports were made by field staff when the data were downloaded from each gauge approximately once a month. Typical problems reported were funnels being blocked by debris and mouse damage to electrical cables. The exact time that the fault occurred within the month is thus difficult to determine. If a gauge was found to be blocked, then the corrective procedure followed was first to remove the blockage and then to let the remaining water run through and be recorded. This reduces the error in the total amount of rain recorded between successive download dates although it can cause difficulties when the data need to be analysed on a finer time-scale.

A conservative estimate of the periods of invalid data could be obtained simply by neglecting all data prior to the time when a fault was found at a particular raingauge, back-dated to when the gauge was previously visited and found to be working satisfactorily. However, this assumes that all the faults are both found and recorded: this may not always be possible, especially if the gauge rectifies itself before the next download. The use of cumulative rainfall totals over a month to highlight malfunctioning gauges has provided the best method for data quality control. Data from clusters of 10 gauges, all relatively close to one another, have been plotted together as cumulative hyetographs over a month to attempt to pinpoint any anomalous gauge behaviour. This approach, combined with the diagnostic reports and knowledge of when each gauge was visited, can lead to accurate (nearest day) assessments of periods of inoperation. It is worth noting that fewer than 50% of the faults discovered in this way were recorded in the diagnostic reports. Since the quality control method is somewhat subjective, experience has indicated that it is important not to be over-zealous in data removal: local fluctuations are to be expected, especially in convective rainfall situations. The cumulative hyetograph of Fig. 3 reveals a host of unusual features for the month of September 1994. In particular, the gauges at Towers Field (Gauge 40) and Dropping Lane Farm (Gauge 36) are seen to record apparent dramatic rainfall on the 4th and 6th respectively. Both gauges were visited on the 6th, although diagnostic reports have no

mention of faults with either gauge. The plots expose the obvious presence of gauge blockages in both cases: these periods were therefore defined as having invalid data for these gauges. The gauge at Redlynch Farm Stump (Gauge 38) illustrates a period where the trace has the gradually rising behaviour associated with a blocked gauge with water "trickling" through. This is clearly evident on and around 15 September. On this occasion a blockage is reported in the diagnostic report, dated 4 October 1994, and the data period is again defined as having invalid data.

Potential timing errors were investigated by calculating correlation coefficients of 15 minute gauge totals with either coincident radar or neighbouring raingauge totals, for a range of time shifts. In nearly all cases, a zero time shift produced the highest correlation coefficient serving to confirm the absence of longstanding timing errors in the data. However, the procedure did identify a case where timing errors had arisen, seemingly because the internal clock of a gauge had begun to "drift" over a period of several months. The availability of tip-times in the raw data allowed a correction to be made for this drift.

Figure 4 provides a summary of the valid data periods (as determined by the above quality control procedures) for each raingauge, plotted as a bar chart over time with the vertical axis indicating the gauge number. The period of record analysed here is from 1 September 1993 to 30 June 1996. For each gauge the start date of operation is taken to be the date *after* the date that the first tip was recorded. It is thus seen that the first gauges came into operation on and around 19 September 1993. The two super-dense networks of 8 gauges within a 2 km square – gauges 5 to 12 (low relief) and 34 to 41 (high relief) – are highlighted on the chart. August 1994 through September 1994 is a period which was particularly "patchy" in terms of coverage by valid gauges. Residue from the mowing of the grass in the summer months within the gauge compound is a known cause of problems, which was rectified in later years by clearing the grass manually.

Snow and hail can be expected to activate the tipping mechanism of a raingauge in a different way to regular rainfall. In addition, the response of the radar will be sensitive to the type of hydrometeors present. As a consequence, part of the quality control process aimed to identify unusual weather events. Using the UK Met Office's Daily Weather Summaries, hourly weather identifiers were created to distinguish convective and stratiform type rainfall as well as solid precipitation such as snow and hail. Periods where solid forms of precipitation were detected are, in the main, omitted from the analysis carried out here since the concern is to quantify *rainfall* measurement accuracy. Figure 5 details an event on 22 February 1994 when snow fell throughout the night, stopping around mid-morning. It shows the 15 minute rainfall estimates from the radars at Cobbacombe and Wardon Hill scanning over the Brue catchment, as well as the raingauge at Bridge Farm (Gauge 8). The presence of snow storms is shown clearly by both




Fig. 3. Cumulative hyetographs used for quality control of the HYREX dataset: raingauges 32 to 44 over the period 1 to 30 September 1994.

radars although there is no response at the raingauge. The choice of the Bridge Farm gauge is due to its proximity to the Automatic Weather Station although its behaviour is typical of all other gauges in the catchment. Around 0900 the gauge begins to register a number of tips, probably signalling the melting of snow collected in the gauge funnel overnight. This is in agreement with the temperature at the automatic weather station rising above 0°C at approximately 0900. Rain which fell from midday onwards appears to be recorded satisfactorily by both radar and raingauge.

Another form of unusual precipitation is freezing rain. Only one event of this nature was identified in the HYREX dataset analysed here. This event, discussed by Pike (1996), occurred late in the evening of 30 December 1995. Figure 6 shows how both radars recorded the presence of the storm. The anomalous spike recorded by the Bridge Farm rain-gauge is representative of similar occurrences at many other gauges in the network but the cause of this behaviour is unknown.

Having addressed the important issue of data quality control, the resulting HYREX raingauge dataset can be used with confidence to address the central question of rainfall measurement accuracy at different space and time scales.

The substantial periods of missing data resulting from what is a carefully maintained set of raingauges would seem to have important consequences for other networks of operational raingauges. These will not enjoy typically the benefit of the spatial gauge density which has allowed the quality control procedures used here to be applied with confidence.

Besides the exclusion of snowfall events identified from Meteorological Daily Weather summaries, certain limited quality-control procedures were applied to the radar data. These were based on a comparison of monthly total rainfalls as measured by the radar and by raingauges at a pixel-scale. The analysis did not show any sudden change in the relative amounts recorded by the two sources, although it did reveal a consistent bias in the radar data treated as an estimate of the rainfall that would be recorded by a raingauge on the ground. It also revealed a small number of periods of anomalous propagation affecting recorded radar values over the Brue catchment. Values from these periods do not affect the results presented here since the estimated "true" rainfall is zero in these cases and such cases are excluded from the analysis. This does highlight the fact that the comparison between radar and raingauge-based estimates excludes cases




Fig. 4. Bar chart indicating the periods of valid data for each rain gauge in the HYREX network for the period 1 September 1993 to 30 June 1996; the two super-dense networks of 8 gauges within a 2 km square – gauges 5 to 12 (low relief) and 34 to 41 (high relief) – are highlighted.

where radar performs very badly. The analyses here do not make any correction for the long-term bias in the radar data, although this is considered in the accompanying paper (Wood *et al.*, 2000). This seems reasonable in that the radar data used here are essentially the same as would have been available for operational use in flood warning centres during the period being analysed. It should be noted that no rain-gauge correction procedure has been applied to the radar data analysed here.

Evaluating the accuracy of rainfall estimates

For hydrological applications, major uses of rain gauge data are to provide estimates of rainfall input to a catchment, and to provide a comparison with radar-derived estimates of

rainfall, typically relating to a 2 km grid square. Measurement of rainfall at both catchment and 2 km grid-square scales will be carried out typically by a single, or at most two, rain gauges in an operational network. For many studies, the position of the rain gauges which perform best in measuring the areal rainfall average are of interest, since this could influence future siting of rain gauges. These questions are usually addressed by a spatial correlation function (or semi-variogram) approach which can provide a formula for estimation accuracy in terms of the number and potential location of rain gauges (see, for example, Morrissey *et al.*, 1995). The correlation function relates to the treatment of rainfall as a random function in space. A review of areal rainfall estimation accuracy using random function theory is provided by Bras and Rodriguez-Iturbe (1993) while papers by Anagnostou *et al.* (1999) and Ciach and Krajewski (1999) apply similar methods to comparisons of rain gauges and




Fig. 5. Data quality control during unusual events using radar and raingauge hyetographs: snowfall event on 22 February 1994 as measured by the Cobacombe and Wardon Hill radars and by the raingauge at Bridge Farm, which records only as snow starts to melt in the funnel.




Fig. 6. Data quality control during unusual events using radar and raingauge hyetographs: freezing rain event late in the evening of 30 December 1995 as measured by the Cobacombe and Wardon Hill radars and by the raingauge at Bridge Farm, giving an anomalous spike.

radar. However, a correlation-based approach has the major disadvantage that it obscures a potentially important consideration: namely, that the accuracy with which rainfall is estimated can be expected to vary with the intensity of rainfall. Anagnostou *et al.* (1999) apply the correlation theory to the logarithm of rainfall which has the effect of treating estimation accuracy as varying proportionally to rainfall intensity. A further difficulty is that the estimation accuracy depends on the strength of spatial dependence which itself may change with rainfall intensity. A correlation function approach produces an answer which is averaged across rainfall intensities. This can be overcome partly by judicious pre-selection of the time frames included in the analysis (O'Connell *et al.*, 1979). Thus different numerical answers would be obtained from correlation analyses of:

- (i) all time frames;
- (ii) time frames with non-zero rainfall recorded at least one raingauge;
- (iii) time frames with "widespread rainfall" according to some specification; and
- (iv) time frames within or close to a "rainfall event" according to a given definition.

The correlation analysis could therefore be manipulated to provide values for the accuracy of estimation for a range of cases. For example, one could define the overall intensity of rain in a time frame as the average of the intensities recorded at all raingauges and arrive eventually at values for the accuracy of estimation of rainfall over a particular area when the overall intensity is in each of a number of ranges, say 0–1 mm h⁻¹, 1–2 mm h⁻¹, 2–3 mm h⁻¹, etc. There are clearly a number of difficulties here including:

- (a) specifying the ranges considered so that the data provide enough time frames for the estimation of the spatial covariance structure; and
- (b) the need to frame the selection criteria so that the assumption of spatial stationarity remains reasonable.

The dense network of raingauges installed for the HYREX project enables a more direct approach to be taken; one which might be termed "empirical". In this approach the dense network is used to provide a good approximation to the true rainfall against which the accuracy, say, of a single gauge estimate may be compared for different rainfall intensities; further details are given below. The approach does have the disadvantage that the results are limited to the particular cases studied, in contrast to the correlation function approach in which, once the correlation function has been estimated, a single formula supplies the answer to how well the rainfall over any-sized area is estimated by any configuration of any number of raingauges. In addition, the results are subject to errors of unknown size, but ones which can reasonably be assumed to be small: these errors are related to the use of a reasonably good, but not perfect, estimate of what the true rainfall would be.

The empirical approach has the advantage of providing a

direct way of examining how the size of estimation error varies with the amount of rain actually observed. This approach is most appropriate when there is a large number of gauges available to form an estimate of the "true" rainfall which is itself to be estimated by a simpler procedure. It is thus reasonable to apply this procedure to radar pixels having 8 raingauges in the pixel, and in cases where the rainfall over the Brue catchment is to be estimated.

In the case of a 2 km pixel, the "ground truth" or "best estimate", T , is defined as the mean of the available raingauges (up to the maximum of 8) within the pixel. This quantity can be defined at both the high and the low relief pixels and is deemed invalid if less than 6 of the raingauges are working satisfactorily for that particular time-frame. Similarly, in the case of catchment rainfall, the network provides up to 49 raingauges which are available for use in a weighted average to form a value, T , for the catchment average rainfall. The essence of the approach is the assumption that there are so many raingauges used in calculating T that it is essentially the same as the unknown true rainfall for the pixel or catchment. It is then of interest to see how accurately T can be reproduced using either the value from a single raingauge or a radar estimate of rainfall. For a given time-frame, the estimate R can be compared with T to form a raw estimate of the mean square error

$$S^2 = (R - T)^2. \quad (1)$$

Such raw estimates of mean square error are obviously poor and it would be standard practice to form an average across a large number of time-frames. For this study a slightly different approach is taken driven by the underlying interest in how the size of the error S varies with the amount of rainfall T (or of $\log S$ against $\log T$), where each time-frame supplies a pair of values, and a smooth curve is fitted to estimate the relationship. This averages across time-frames, but in such a way as to reveal the dependence of the size of S on T . Effectively this estimates $E(S^2)$ as a function of $\log T$ and plots are made of $\frac{1}{2} \log E(S^2)$ against $\log T$. It is convenient to think of this as $\log S$ against $\log T$.

In this approach there is an underlying assumption that the values for T are good estimates of the true rainfall, so that values of S are a good guide to how well the simpler estimates perform. The next section considers the question of how good the values of T are for the 8 gauge radar-pixels.

Accuracy of estimates of 2 km grid square rainfall

For the two super-dense network squares, "ground truth" is established by forming the simple arithmetic average of the 8 gauges arranged in a "diamond within a square" configuration. The HYREX gauge network was designed such that one of these squares occupies an area of low relief whilst the other lies in an area of high relief, providing the ability to




Fig. 7. Rainfall variability within the low relief 2 km pixel, obtained by multiquadric interpolation of the HYREX raingauge network dataset, for the 15 minute interval ending 2230 3 August 1994.

assess rainfall variability over these two area types. An example of the variability of rainfall experienced over a 2 km pixel is given in Fig. 7, for the 15 minute period ending 22:30 3 August 1994.

The accuracy with which rainfall over a 2 km square pixel can be measured is of importance for hydrological application since this spatial scale is often the basic unit of integration for weather radar data. In turn, distributed rainfall-runoff models at the catchment scale have been configured using grids of this size and using a 15 minute period of time integration (for example, Bell and Moore, 1998a,b). It is therefore important to know both how well radar estimates rainfall on a 2 km pixel basis and how well a single raingauge should agree with the radar estimate of rainfall for the square in which it lies.

The method used to quantify the accuracy of estimates of rainfall over a 2 km grid square is based on a model which assumes a spatial correlation function which is constant over the range of the pixel (that is 2.83 km), apart from a possible sharp drop over very small distances. The advantage of this model is that simple estimates for the mean square error of estimation are immediately available. The model is that the i th raingauge value for a single time-frame, R_i , can be decomposed into three additive components. The first is a random variable variable, X , representing a *field value* common to all points within the pixel; the second is an identically distributed random variable, Y_i , representing the *measurement error* at gauge i ; and the third is a random variable, Z_i , representing a *local fluctuation* in the rainfall

field at the i th gauge. Thus

$$R_i = X + Y_i + Z_i. \quad (2)$$

The *true rainfall* at any point, including non-observation points, is of the same form but excluding the measurement error term, Y . The random variables X , Y and Z are assumed independent among themselves and, apart from X , from site to site. In particular, the local fluctuations Z at different locations are assumed independent no matter how close the points considered. The inclusion of the common value X at all locations leads to the measured rainfalls at different sites having a positive correlation (and also the true rainfalls). Under these assumptions, the true average rainfall intensity for the pixel is exactly X for the time-frame in question. The data values $\{R_i, i=1,2,..n\}$ from the n raingauges available for the time-frame may be used to construct the sample mean and sample variance

$$\bar{R} = \frac{1}{n} \sum_{i=1}^n R_i \quad (3)$$

$$S^2 = \frac{1}{n-1} \sum_{i=1}^n (R_i - \bar{R})^2. \quad (4)$$

In practice, a rule is introduced which leaves the mean and variance undefined if fewer than six of the eight gauges are working for the time-frame.

The quantity S^2 is capable of three interpretations. Firstly, it can be viewed as being a measure of how different

the values are likely to be of rainfalls recorded at different sites within a pixel: thus

$$E(R_i - R_j)^2 = 2E(S^2), \quad i \neq j \quad (5)$$

which follows from the independence assumptions made in the model. Secondly, it provides a measure of how well any single gauge estimates the true pixel average rainfall (X here, or T in the more general notation used earlier)

$$E(R_i - X)^2 = E(S^2). \quad (6)$$

This follows since S^2 is essentially an estimate of the variance of $(Y_i + Z_i)$. Thirdly, S^2 can be used to provide an estimate of how well the estimate given by the average of a number, m say, of gauges within a 2 km pixel can be expected to measure the true pixel average rainfall. Thus if $\bar{R}^{(m)}$ is the average of m gauges

$$E(\bar{R}^{(m)} - X)^2 = \frac{1}{m} E(S^2) \quad (7)$$

and so

$$s^2 = \frac{1}{m} S^2 \quad (8)$$

provides an estimate of the mean square error for $\bar{R}^{(m)}$.

There is a need to combine a large number of framewise pairs (T, S^2) in order to examine how the size of estimation error relates to the amount of rainfall that has actually fallen. The same procedure that was used in the section headed "Evaluating the accuracy of rainfall estimates" is used except that, here, a single value of S^2 is derived from the several raingauge values for the grid square. The adjusted averaging in Eqn. (4) is the usual formula used for estimating a variance which assures an unbiased estimate under the assumptions outlined above. The adjustment reflects the contribution made by a single gauge to the estimated value of the "true" quantity T estimated by $\bar{R}$. It is not possible to make a standard adjustment for this effect in other cases, but, when enough gauges are available to form T , the effect is small as has been found in cases where analyses have been repeated with the candidate gauge removed from contributing to T .

Values of S^2 and T calculated for each 15 minute time-frame can be used to establish an empirical relation between estimation accuracy and rainfall intensity. The mean square error of estimation, given by S , is plotted against the true rainfall estimate, given by T , as a scatter plot using logarithmic axes. Figure 8 presents an example plot calculated for the grid square in the area of low relief. A random dither factor has been added to each plotted point to emphasise the number of points which are overplotted, which arise through discretisation effects. Different rainfall types are distinguished using the symbols $\bullet$ and $\circ$ to denote frontal and convective respectively. Data influenced by solid forms of precipitation are omitted from this analysis. Figure 8 includes a smoothed plot of S against T which is created




Fig. 8. Standard error, S , of rainfall estimate from a single gauge as a function of total rainfall in a 15 minute interval over the low relief square. Vertical lines indicate the 90% range in sampling uncertainty and the symbols $\bullet$ and $\circ$ denote frontal and convective rainfall types respectively.

by the method of locally weighted least squares outlined in Appendix A: the same method is used in presenting other results. The vertical lines in Fig. 8 indicate a 90% range in the sampling distribution for the smoothed line, obtained as outlined below. The relation corresponding to Fig. 8, but for the high relief square, is very similar in overall distribution and behaviour although the confidence limits are broader.

An assessment of the uncertainty inherent in the smoothed line is obtained by a Monte Carlo bootstrap method (Efron and Tibshirani, 1986) in which the original sample of data is repeatedly resampled on the basis of sampling with replacement of blocks of 3 days. This effectively allows for both spatial and temporal correlation in the rainfall fields. Each resample is used to calculate a smoothed line by the method of Appendix A and the distribution of the resampled y values for a given x gives an indication of the sampling uncertainty. Where a comparison between the accuracies of different estimators is to be made, it is possible to apply the same combination of a smoothed line and bootstrap resampling to the differences of the squared-errors to allow an assessment of whether one estimator is better than the other. This approach automatically takes into account the statistical dependence of the errors in the estimators for the same time-point. Space limitations preclude inclusion of an example of this type of analysis.

Figure 9 shows how the accuracy of an 8-gauge estimate (derived via Eqn. 8) compares with that of a typical single gauge for the low relief square. Also shown is the accuracy associated with the 2 km radar estimate using Wardon Hill

radar. One interpretation of the values here is that, with certain assumptions, for a 15 minute rainfall of 10mm, errors in the radar estimate have an apparent mean squared error of 16 mm^2 , of which 1 mm^2 may have arisen by use of the 8-gauge estimate in the assessment procedure rather than the unknown true rainfall. Thus the value of S for radar might really be 3.87 mm ($\sqrt{(16-1)}$) rather than 4 mm , which shows that the effect is small. A further interpretation relates to the situation with more typical networks of operational gauges where radar data can only be compared to a single raingauge in isolated radar-pixels. Figure 9 indicates that even if radar were to provide a perfect estimate of the true pixel-average rainfall, the radar data and raingauge data can be expected to have a root mean square difference of about 3 mm if the true rainfall total is 10 mm in 15 minutes.

Figure 10 shows the effect of temporal scale on the accuracy of a typical gauge as an estimator of 2 km grid square rainfall. Hourly accumulations of rainfall are seen to be more accurate than 15 minute ones for higher rainfalls, whilst for lower rainfalls the 15 minute totals are more accurate. Similar effects at high and low rainfalls are seen when moving to one day totals.

The results obtained suggest that a straight line can be used for the accuracy relationship on a log-log graph, although further investigation might be required at low rainfall intensities. A slope of near unity for 15 minute accumulations suggests use of a logarithmic transformation would lead to almost constant accuracy on this scale. Figure 10 indicates that a power transformation with a positive exponent near zero would be needed to gain the same effect for hourly and daily accumulations.




Fig. 9. Comparison of the standard errors of estimation for 15 minute rainfall in the low relief square for a single gauge, for 8 gauges in the square and for radar estimates.




Fig. 10. Standard error, S , of rainfall estimate obtained from a single gauge as a function of total rainfall over the Brue catchment for time periods of 15 minutes, 1 hour and 1 day.

Accuracy of estimates of catchment average rainfall

Estimation of rainfall over a typical catchment is of importance since this spatial scale is usually taken as the basic unit for rainfall-runoff modelling, with lumped catchment models often providing the most practical solution to flow forecasting problems. It is again of interest to see how errors of estimation vary with rainfall intensity on a 15 minute timescale. An example of the variability of rainfall over the Brue catchment is given in Fig. 11, which shows the distribution of rainfall for the 15 minute interval ending 22:30 3 August 1994. When forming a value, T , to represent the true average rainfall for the Brue catchment it is clear that a simple arithmetic mean of the 49 raingauges in the HYREX network will introduce a tendency to weight the result towards values for the squares containing 8 and 2 gauges. Instead, an estimate was formed by calculating the arithmetic mean rainfall for each of the 28 squares containing one or more gauges and then taking the average of these values. Specifically the catchment average rainfall for a 15 minute interval is calculated as

$$T = \bar{R} = \frac{1}{m} \sum_{j=1}^m \left(\frac{1}{n_j} \sum_{i=1}^{n_j} R_{ij} \right) \quad (9)$$

with the estimation variance for a typical gauge being estimated by

$$S^2 = \frac{1}{m} \sum_{j=1}^m \left(\frac{1}{n_j} \sum_{i=1}^{n_j} (R_{ij} - \bar{R})^2 \right) \quad (10)$$

where m is the number of squares containing at least one operating gauge, n_j is the number of gauges in square j with




Fig. 11. Rainfall variability over the Brue catchment, obtained by multiquadric interpolation of the HYREX raingauge network dataset, for the 15 minute interval ending 22:30 3 August 1994.

valid data for the time-interval, and R_{ij} is the rainfall at gauge i in square j . Equation (10) is based on the same principal of not giving undue prominence to squares containing two or more gauges. Only quality controlled rainfall values are used in forming the averages.

Rainfall data for each time-interval are used to construct the sample mean, T , and sample variance, S^2 , and a scatter plot of S versus T for all time-frames produced. Figure 12 shows how the accuracy of measurements from two particular raingauges compare with those from a "typical" single gauge as estimates of the catchment average rainfall. The gauge at Kilkenny Fields (Gauge 2) located on the edge of the catchment is seen to provide a poor estimate of the catchment average rainfall for all rainfall intensities. In contrast, the gauge at Crabtree Lane Field (Gauge 25) situated near the centre of the catchment provides an estimate of the catchment average rainfall that is comparable to that of a typical gauge. Also shown is the accuracy obtained using the 2 km radar data from Wardon Hill: this proves more accurate than the typical gauge estimate for all but the highest rainfalls.

Conclusions

Assimilation of a variety of data sources into the analysis of rainfall distributions, and their hydrological effects, can introduce its own problems. Intensive data quality analysis has been carried out on over 2 years of raingauge data.

Overall, this has shown that the HYREX dense raingauge network has performed well apart from some problems in the first summer (in late Summer 1994) when a large number of gauges suffered blockages, due especially to mown grass. The summers of 1995 and 1996 saw a marked improvement in this respect. Meteorological factors have




Fig. 12. Standard error, S , of the catchment rainfall estimate obtained from a single gauge or from radar as a function of total rainfall in a 15 minute interval over the Brue catchment.

been considered by studying Daily Weather Summaries. This allowed snow and other solid precipitation events to be removed from the majority of subsequent analyses, and studied separately for particular purposes. In addition, it has been possible to categorise stratiform and convective rainfall situations.

The main aim has been to quantify the accuracy of estimates of rainfall over 2 km grid square and catchment scales, focussing on the 15 minute timescale as the most relevant to rainfall-runoff modelling for flood forecasting. An empirical approach has been pursued which has involved treating each 15 minute time-frame separately. For each time-frame the "true" rainfall was set to the mean rainfall over a large number of gauges and the estimate and error were calculated. Scatter plots of these quantities allowed smoothed representative lines to be produced quantifying the relationship between measurement accuracy and rainfall intensity.

Two super-dense gauge networks each comprising of 8 gauges within a 2 km side square, one in an area of low relief and the other in high relief, were used to estimate the "true" rainfall as a simple arithmetic average. Estimation errors for individual gauge estimates, and radar estimates, of this true mean rainfall over the square have been calculated. Errors of around 33% at 4 mm were observed at the grid square situated in low relief, with larger errors of 45% observed at the high relief grid square. However, the orography in the catchment is not varied enough to allow an extensive investigation of its effect on rainfall estimation accuracy. The effect of the time-scale over which measurements are made was also investigated. As the time interval becomes larger, the measurement accuracy improves due partly to the reduction in the effects of tipping-bucket discretisation errors. Whilst convective conditions in the period analysed have led to higher intensity rainfall than during frontal events, it has not been possible to distinguish different rainfall estimation accuracies for these.

The HYREX raingauge network of 49 gauges over the Brue catchment has also been used to estimate the "true" catchment rainfall as a weighted average of the available gauge values as part of a similar assessment of the accuracy of estimates of catchment rainfall. A brief comparison of the various types of estimate can be stated as follows. For 4 mm of rain in 15 minutes the standard error of a single gauge estimate varies from 33% for the 2 km square in which it lies to 65% when treated as a catchment estimate, noting that the area ratio involved is 1:34. In these circumstances, 8 gauges in a 2 km square would provide an estimate with a standard error of 11.5%. Further, radar at a 2 km resolution has a standard error of 50% for estimating rainfall on a 2 km square and 55% for catchment rainfall. Radar at 5 km resolution has estimates of catchment rainfall with a 60% standard error. These results for radar relate to data available operationally in real-time. A companion paper (Wood *et al.*, 2000) considers the benefit of correcting for the long-term bias present in these data.

Appendix A

SMOOTHING USING WEIGHTED LOCAL LEAST SQUARES

Scatter plots produced using the empirical approach are smoothed using locally weighted least squares. Considering a target point along the x-axis, the corresponding y value is determined as a function of the local points having similar x-values. The importance to be attached to each of the local points is determined by a weighting function $K(t)$ defined as:

$$K(t) = \begin{cases} 1 - t^2 & |t| < 1 \\ 0 & \text{otherwise} \end{cases} \quad (\text{A.1})$$

which has the property that $K(t)$ is large only over the range $\pm 5^{-1/2}$. The value $5^{-1/2}$ arises as the standard deviation of a probability density proportional to $K(t)$. If h is the required effective smoothing width, centred on the target x value, then h^* is required such that $K(x/h^*)$ is large over the range of x from $-h$ to $+h$: hence, $h^* = \sqrt{5} h$. A typical value for h is $(\log 1 - \log 0.75)$.

For the present smoothing application the sample points x_i and y_i denote the logarithm of the true rainfall and the mean square error of the estimate (ie. $\log T$ and S^2), respectively, for one time-frame. Smoothing proceeds as follows. For a target point x^* , and data (x_i, y_i) , $i = 1 \dots N$, weights for each point of the sample set are calculated as

$$w_i = K\left(\frac{x_i - x^*}{h^*}\right). \quad (\text{A.2})$$

Now, the local weighted least squares regression equation is

$$y = y_w + \beta(x - x_w) \quad (\text{A.3})$$

giving the smoothed value of y at x^* as

$$y^* = y_w + \beta(x^* - x_w) \quad (\text{A.4})$$

where

$$x_w = \frac{\sum w_i x_i}{\sum w_i} \quad (\text{A.6})$$

$$y_w = \frac{\sum w_i y_i}{\sum w_i} \quad (\text{A.7})$$

$$\beta = \frac{\sum w_i (y_i - y_w)(x_i - x_w)}{\sum w_i (x_i - x_w)^2} \quad (\text{A.8})$$

and β is treated as undefined unless the denominator $\sum w_i (x_i - x_w)^2 > 1$. The final smoothed line is obtained from the plot of $\frac{1}{2} \log y^* (= \log \sqrt{y^*} = \log S)$ against $x^* (= \log T)$.

References

- Anagnostou, E.N., Krajewski, W.F. and Smith, J., 1999. Uncertainty quantification of mean-areal radar-rainfall estimates. *J. Atmos. Oceanic Technol.*, 16, 206–215.

- Austin, R.M. and Moore, R.J., 1996. Evaluation of radar rainfall forecasts in real-time flood forecasting models. *Quaderni Di Idronomia Montana*, 16, 19–28.
- Bell, V.A. and Moore, R.J., 1998a. A grid-based distributed flood forecasting model for use with weather radar data: Part 1. Formulation. *Hydrol. Earth System Sci.*, 2, 265–281.
- Bell, V.A. and Moore, R.J., 1998b. A grid-based distributed flood forecasting model for use with weather radar data: Part 2. Case studies. *Hydrol. Earth System Sci.*, 2, 283–298.
- Bras, R.L. and Rodriguez-Iturbe, I., 1993. *Random functions and hydrology*. Dover, New York, 559 pp.
- Ciach, G.J. and Krajewski, W.F., 1999. On the estimation of radar rainfall error variance. *Advances in Water Resources*, 22, 585–595.
- Efron, B. and Tibshirani, R., 1986. Bootstrap methods for standard errors, confidence intervals, and other measures of statistical accuracy. *Statistical Sci.*, 1, 54–77.
- Essery, C.I. and Wilcock, D.N., 1991. The variation in rainfall catch from standard UK Meteorological Office raingauges: a twelve year case study. *Hydrol. Sci. J.*, 36, 23–34.
- Larson, L.W. and Peck, E.L., 1974. Accuracy of precipitation measurements for Hydrologic Modeling. *Water Resour. Res.*, 10, 857–863.
- Moore, R.J. (Ed.), 1998. *Requirements and applications of weather radar data to hydrology and water resources*. World Meteorological Organisation Technical Reports in Hydrology and Water Resources No. 70, WMO/TD – No. 934, Geneva, Switzerland.
- Morrissey, M.L., Maliekal, J.A., Greene, J.S. and Wang, J., 1995. The uncertainty of simple spatial averages using rain gauge networks. *Wat. Res. Res.*, 31, 2011–2017.
- O'Connell, P.E., Gurney, R.J., Jones, D.A., Miller, J.B., Nicholass, C.A. and Senior, M.R., 1979. A case study of a raingauge network in southwest England. *Water Res. Res.*, 15, 1813–1822.
- Peck, E.L., 1980. Design of precipitation networks. *Bull. Am. Meteorol. Soc.*, 61, 894–902.
- Pike, W.S., 1996. The Ice Storms of late December 1995. *J. of Meteorology*, 21, 207–217.
- Wood, S.J., Jones, D.A. and Moore, R.J., 2000. Static and dynamic calibration of radar data for hydrological use. *Hydrol. Earth System Sci.*, 4, 545–554.