defra
Department for Environment Food and Rural Affairs

Environment Agency

delivering benefits through evidence

Reservoir Safety – Long Return Period Rainfall

Volume 1 Technical Report (Part 1)

Project: FD2613 WS 194/2/39

Flood and Coastal Erosion Risk Management Research and Development Programme

Statement of use
This document provides details of the analysis undertaken during project WS 194/2/39 entitled *Reservoir safety: long return period rainfall* funded by Defra. The report describes a new model of rainfall depth-duration-frequency applicable to the UK, which is proposed as a replacement to that published in the Flood Estimation Handbook (IH, 1999). The report is intended to inform Defra and Environment Agency staff, reservoir panel engineers, consultants, contractors and other agencies and organisations involved in hydrological frequency estimation about the new model. Further work to develop a software implementation of the new model is ongoing.

Dissemination status: Publicly available

Keywords: Reservoir safety, extreme rainfall, Flood Estimation Handbook, rainfall frequency estimation, depth-duration-frequency model, spatial dependence, standardisation, annual maxima, FORGEX, PMP

Research contractor:
Ms Lisa Stewart, Centre for Ecology & Hydrology, Maclean Building, Crowmarsh Gifford, Wallingford, Oxon. OX10 8BB (ejs@ceh.ac.uk)

Peter Dempsey, Met Office, FitzRoy Road, Exeter, Devon EX1 3PB (peter.dempsey@metoffice.gov.uk)

Professor Christopher Collier, National Centre for Atmospheric Science, School of Earth & Environment, University of Leeds, LS2 9JT (c.g.collier@leeds.ac.uk)

Professor Clive Anderson, Dept of Probability & Statistics, University of Sheffield, Sheffield S3 7RH (c.w.anderson@shef.ac.uk)

Defra project officer: John Goudie, Flood Management Division, Ergon House, Horseferry Road, London SW1P 2AL (john.goudie@defra.dsi.gov.uk)

Publishing organisation
Department for Environment, Food and Rural Affairs
Flood Management Division,
Ergon House,
Horseferry Road
London SW1P 2AL

Tel: 020 7238 3000 Fax: 020 7238 6187

www.defra.gov.uk/environ/fcd

Published by the Department for Environment, Food and Rural Affairs (2013).

Joint Defra/Environment Agency Flood and Coastal Erosion Risk Management Research & Development Programme

Reservoir Safety – Long Return Period Rainfall

R&D Technical Report WS 194/2/39/TR
Volume 1 (Part 1)

Submitted: March 2010

Authors:

E. J. Stewart[1], D. A. Jones[1], C. Svensson[1], D. G. Morris[1], P. Dempsey[2], J. E. Dent[2], C. G. Collier[3] & C. W. Anderson[4]

[1] Centre for Ecology & Hydrology, Wallingford
[2] Met Office, Exeter
[3] National Centre for Atmospheric Science, School of Earth & Environment, University of Leeds
[4] Department of Probability & Statistics, University of Sheffield

Foreword

This report presents the results of a major project to develop a new statistical model of point rainfall depth-duration-frequency (DDF) for the UK. This is intended to replace both the Flood Estimation Handbook (FEH) model (Faulkner, 1999) and the present guidance given to Defra panel engineers (Defra, 2004) that the FEH rainfall estimates should not entirely replace the old Flood Studies Report (FSR) estimates of 1975.

Acknowledgements

Rainfall data have been supplied by the Environment Agency for England and Wales, by the Scottish Environment Protection Agency for Scotland, by the Met Office for locations throughout the United Kingdom, and by Met Éireann for locations in the Republic of Ireland near the border with Northern Ireland. The use of all these data is gratefully acknowledged.

The members of the Project Steering Group are thanked for their helpful comments on all aspects of the work involved in the project:

John Gosden, Jacobs (Chairman)
Dr Geoff Baxter, Environment Agency
Sir David Cox, University of Oxford
Paul Ditchfield, Defra
Marcus Francis, Jacobs
Fiona Geddes, Environment Agency
Dr Peter Mason, MWH Global
Professor Paul Samuels, HR Wallingford
Peter Spencer, Environment Agency
Professor Howard Wheater, Imperial College

Besides those listed as authors, CEH staff who made substantial contributions to the project included Fai Fung and Sonja Folwell.

Executive Summary

This report presents the results of a major project to develop a new statistical model of point rainfall depth-duration-frequency (DDF) for the UK to replace the current Flood Estimation Handbook (FEH) DDF model. The new model was constructed for estimating rainfall depths falling over durations ranging from 1 hour to 192 hours (8 days) for return periods ranging from 2 years to 10,000 years. However, in many locations it is capable of producing indicative estimates for higher return periods, up to 100,000 years.

The project was commissioned in response to concerns, expressed by reservoir engineers, about the apparently high estimates produced by the FEH DDF model when it was applied to return periods in excess of its recommended upper limit of 1,000 years. In many locations, the FEH model was giving 10,000-year estimates considerably higher than the Flood Studies Report (FSR) probable maximum precipitation (PMP). This is used to calculate the probable maximum flood as a statutory part of the spillway design procedure for major reservoirs.

In this project, the framework of the FEH approach to rainfall modelling has been retained, but a number of key elements have been substantially revised. In particular, the Focused Rainfall Growth Curve Extension (FORGEX) methodology has been reformulated, and the dataset of annual maximum rainfalls, to which the final model is fitted, has been updated. The main improvements are:

- Data – the number of suitable hourly raingauges is more than double the total available to the FEH, giving much improved coverage in Scotland and south-west England. The data have been subjected to extensive quality control.

- Standardisation – by using SAAR and northing in addition to FEH's RMED, the new model removes more of the location-dependent variation in rainfall before combining maxima from networks of raingauges.

- Spatial dependence model – the new model allows for a reduction in spatial dependence (i.e. greater independence) between raingauges as return period increases.

- FORGEX – the FEH method of deriving rainfall growth curves has been improved to give a better fit to the network maxima and more gradual variation between locations.

- DDF model – the new model is more flexible than the FEH model and is better able to represent the output from FORGEX across the full range of durations and return periods. Unlike FEH, the new model does not increase exponentially if extrapolated beyond the range of return periods represented in the observed datasets.

The report presents comparisons of rainfall estimates from the new model with those from the FSR and the FEH models for 71 sites across the UK, 35 of which are close to impounding reservoirs. These show that, generally, estimates for the longer return periods are lower, especially in comparison with FEH. However, in Scotland estimates for the shortest durations have increased. These changes are due, respectively, to the improved spatial dependence model and improvements to the hourly rainfall dataset. In the majority of cases, the new 10,000-year estimates are lower than the FSR PMP.

Also presented in the report is a comparative assessment of the return periods of 26 major UK rainfall events from the period 1880 to 2006 derived from FSR, FEH and the new model. The estimated return periods from the new model are generally substantially higher than from the earlier models, but statistical arguments suggest that they are closer to what may be expected for the largest point-wise return period observed across a large number of years and a large number of raingauges.

Implementing the new model will be computationally intensive and will require new, detailed digital maps of median annual maximum rainfall (RMED) to be developed for durations ranging from 1 hour to 8 days.

The report sets out an implementation programme and makes wide-ranging suggestions for future research.

Contents Part 1

Contents Part 2

Appendices: Volume 2

Glossary

AREA	Notional area covered by a network of gauges computed from average inter-gauge distance.
Areal reduction factor	A factor applied to a point rainfall for a given return period to give an areal average rainfall with the same return period (*note that several other definitions exist*).
D	Duration (length of the period over which rainfall is accumulated).
DDF	Depth-duration-frequency.
Discretisation conversion factor	A factor applied to annual maxima extracted from data at 1-hour or 1-day resolution to adjust them to correspond to 'fully-sliding' values.
F	Non-exceedance probability.
FEH	Flood Estimation Handbook (five volumes) (*Institute of Hydrology, 1999*).
Fixed duration maxima	Largest of all rainfall totals, where the duration considered coincides with the time resolution of the underlying data.
FORGEX	Focused Rainfall Growth Curve Extension methodology described in FEH Volume 2.
FSR	Flood Studies Report (five volumes) (*NERC, 1975*).
Fully-sliding maxima	The value of annual maximum that would be found by considering all possible time periods of the given duration, not limited by time resolution of data.
GEV	Generalised Extreme Value distribution.
ICE	Institution of Civil Engineers.
Index-flood method	A simple approach to standardisation, assuming that division of values by a site-specific value will explain all differences between distributions at different sites.
L-kurtosis	A measure of kurtosis (tendency to have long tails) of a distribution, based on L-moments.
L-moment ratios	Statistics such as L-skewness, constructed as the ratio of two L-moments.

L-moments	Summary statistics of a distribution, computed as linear combinations of the ranked observations.
L-scale	A measure of the spread of a distribution derived from L-moments.
L-skewness	A measure of skewness (lack of symmetry) of a distribution, based on L-moments.
lcmed	An L-moment based version of the coefficient of variation (L-scale/median).
M5	5-year return period rainfall (mm).
MCS	Mesoscale convective system, defined as a continuous cloud system of thunderstorms associated with a wide area of precipitation (*Houze, 1993*).
N	Actual number of gauges in a network.
N_e	Effective number of independent gauges in a network.
Netmax	Network maximum point (the highest annual maximum standardised rainfall within a network of raingauges).
ngy	The northing component of the National Grid Reference of a location, expressed in 1000 km units.
Orography	For a given location, variations in altitude and slope in the near and distant neighbourhood of the location.
PMF	Probable maximum flood.
PMP	Probable maximum precipitation.
Raingauge network	Any collection of raingauges from the available set of gauges, not restricted to any one original measuring authority, usually within a specified distance of a particular location.
ReFH	Revitalised Flood Hydrograph model (*Kjeldsen, 2007*).
RMED, *RMED*	Median of annual maximum rainfalls at a given site, different values for different durations.
rsaar1	A constructed variable based on SAAR.
SAAR, *SAAR*	Standard average annual rainfall (1961-1990) (mm).

Semi-sliding maxima	Contrasts with ‘fully-sliding maxima’, used to recognise that, in practice, rainfall data are not available for an arbitrarily fine resolution.
Sliding duration maxima	The FEH term for fully-sliding maxima. The maximum is taken across all totals of the given duration.
Spatial dependence	A general statistical quality corresponding to the tendency for high or low values to occur simultaneously at a collection of sites.
Spatial dependence model	In this project, a model where the effect of spatial dependence on the distribution of the maximum value across a collection of sites is represented.
Standardisation	Methods for creating a processed set of values in which the effects of explainable differences between sites have been eliminated.
Standardised rainfall	In this project, standardised rainfalls are created by adjusting to have a common median and to reduce differences in spread based on values of SAAR.
T	Return period (years).
x	Standardised rainfall.
y	Gumbel reduced variate.

Entries in italics are symbols used in equations

1. Introduction

This report presents the final results of Research and Development project WS 194/2/39: *Reservoir Safety – Long return period rainfall*, funded by Defra and managed by the Environment Agency. The objective of the project was to devise an improved method of estimating the magnitude of rainfall of return periods from 100 to 10,000 years (corresponding to annual exceedance probability from 10^{-2} to 10^{-4}) relevant to reservoir safety in the UK.

During the course of the project, the Institution of Civil Engineers' Reservoir Safety Advisory Group (RSAG), which is made up of experts from government, industry and academia, requested that the upper limit of return period should be increased to 100,000 years to provide inputs for quantitative risk assessment.

The research has led to a number of revisions to the FORGEX methodology, which underpins the model of rainfall frequency estimation published in Volume 2 of the Flood Estimation Handbook (FEH) (*Faulkner, 1999*). A new model of rainfall depth-duration-frequency has also been developed, which can be applied at any location in the UK where there are sufficient rainfall data. Further work will be needed so that the model can be applied at ungauged sites, and to develop a software implementation to replace the one currently available on the FEH CD-ROM.

1.1 Background to floods and reservoir safety

There are about 2,600 reservoirs that fall under the Reservoirs Act 1975, the main law relating to reservoir safety in Great Britain. The legislation builds on the Reservoirs (Safety Provisions) Act 1930 in establishing a system of inspection for large raised reservoirs holding at least 25,000 m^3 of water impounded above natural ground level. The 1975 Act places various public safety obligations on reservoir undertakers (operators, users and owners) and applies in England, Wales and Scotland.

Current guidance on floods and reservoir safety (*ICE, 1996*) divides dams into four categories in terms of the potential hazard to life and property downstream, as shown in Table 1.1. Where there is a severe threat of loss of life and extensive damage, greater security against dam failure is needed than where the threat is lower. Therefore, the design standards for dams of category A to D range from the probable maximum flood (PMF) to the 150-year return period flood. Figure 1.1 shows the locations of category A and B dams in Great Britain and the locations of large raised reservoirs in Northern Ireland.

The assessment of reservoir safety requires a complete design inflow hydrograph, which is derived from a rainfall-runoff method, in which an appropriate design storm and associated antecedent conditions are applied to a hydrological model of the catchment (*NERC, 1975; Reed & Field, 1992; Institute*

Table 1.1 Flood standards by dam category (adapted from ICE, 1996)

Dam category	Potential effect of a dam breach	Reservoir design flood inflow: General	Reservoir design flood inflow: Minimum standard if overtopping tolerable
A	Where a breach could endanger lives in a community	Probable Maximum Flood (PMF)	10,000-year flood
B	Where a breach (i) could endanger lives not in a community (ii) could result in extensive damage	10,000-year flood	1,000-year flood
C	Where a breach would pose negligible risk to life and cause limited damage	1,000-year flood	150-year flood
D	Special cases where no loss of life can be foreseen as a result of a breach and very limited additional flood damage would be caused	150-year flood	Not applicable

of Hydrology, 1999). In the case of the FSR/FEH rainfall-runoff method (*Houghton-Carr, 1999*), the appropriate design storm ranges from the 193-year rainfall (used in the synthesis of the 150-year flood) to the probable maximum precipitation (PMP), which is used to generate the PMF. The storm duration critical to reservoir safety varies from a few hours to several days. This reflects the wide range of catchment areas of UK reservoirs, which vary from less than one square kilometre to hundreds of square kilometres.

Until the publication of the FEH in 1999, UK practice in flood estimation was based on the methods outlined in the Flood Studies Report (FSR) published by NERC (1975) and refined in a number of supplementary reports over the following decade. The methodologies for rainfall frequency estimation described in both the FSR and the FEH are based on the analysis of annual maximum rainfalls for a range of durations. Both methods are based on a two-stage procedure involving the estimation of an index variable to which growth factors are applied to yield an estimate of rainfall depth of the required duration and return period.

The FEH (*Institute of Hydrology, 1999*) introduced a new set of procedures for estimating rainfall and flood frequency in the UK. Two particularly innovative features of the FEH approach were the use of digital catchment information

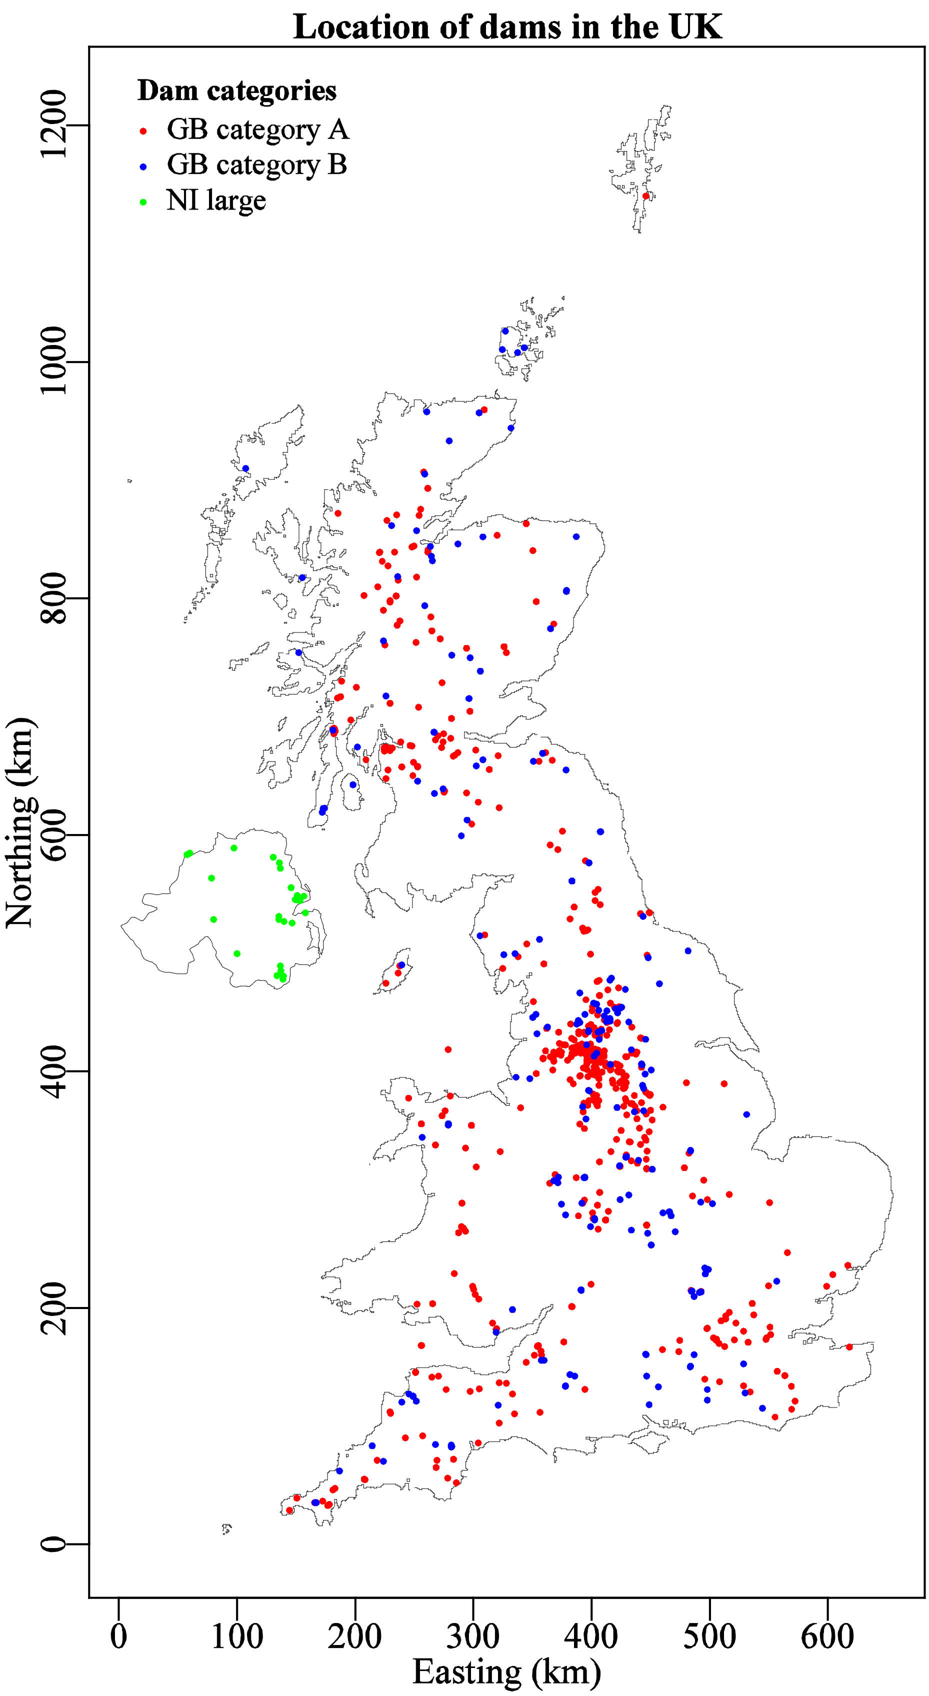

Figure 1.1 Locations of category A and B dams in Great Britain (Tedd *et al.*, 1992), and locations of large dams in Northern Ireland (Cooper, 1987)

to help estimation at ungauged sites, and the introduction of flexible regionalisation schemes, which allow the extent of data pooling to be tailored to the target return period. In the case of the FEH rainfall analysis, data from many raingauge sites were pooled together to define rainfall frequency curves using the Focused Rainfall Growth Curve Extension (FORGEX) methodology (*Reed et al., 1999*) which in turn formed the basis of a new depth-duration-frequency (DDF) model.

The FEH methods have been widely adopted for a variety of applications, including the design and appraisal of flood defence works and the mapping of flood risk. However, there has been concern expressed within the reservoir industry over the results obtained from the DDF model for high return periods used for assessing reservoir flood safety. In some cases, for example, (*MacDonald & Scott, 2001*) it has been noted that the FEH 10,000-year return period rainfall exceeds the estimate of PMP derived from the FSR.

As a result of this, and pending the outcome of the current project, Defra issued interim guidance to reservoir engineers, which recommended that the FEH should not be used for the assessment of 1 in 10,000-year return period rainfall.

1.2 Project specification and approach

On behalf of Defra, Cox (2003) made an independent assessment of the adequacy of the FEH methodology for estimating extreme rainfall for reservoir safety, and made a number of recommendations. The most important included:

- Further analysis of major rainfall events experienced in the UK should be made to extend the study of Collier *et al.* (2002).
- The nature of spatial dependence at high rainfall levels and its effect on the plotting positions used in FORGEX should be reviewed.
- The precise mode of extrapolation to high levels should be re-examined.

These recommendations formed the basis of the project specification. The project has investigated alternative ways of estimating extreme rainfall relevant to reservoir safety for the whole of the UK. The study has also compared the resulting frequency estimates with the currently recommended PMP values derived from the FSR. Rainfall durations from 1 hour to 8 days were considered to encompass the range of typical storm durations relevant to reservoir flood design in the UK.

The analytical approach within the current project has been to reassess the key elements of the FORGEX method and the fitting of the DDF model, and to consider possible improvements and refinements. Therefore, the basic framework of the methodology has been retained. Other possible approaches to estimating rainfall extremes are outlined in the appendices to an earlier project report (*Svensson et al., 2006*), but these remain outside the scope of this project.

The use within FORGEX of standardisation by RMED has been reconsidered after examining whether the assumption of a common underlying distribution for the whole of the UK is valid. Various climatological and topographical descriptors have been analysed and a revised standardisation has been proposed.

The current analysis has retained the concept of focusing, whereby rainfall growth curves are centred on a particular location and the amount of information used from wider raingauge networks is related to the return period of the required estimate. Network maximum points, that is the highest standardised annual maximum value over a given raingauge network, are still being used to define the rainfall growth curve at higher return periods, but the spatial dependence model used to assign plotting positions to these points has been revised. In the final stage of the analysis, a new DDF model has been developed that is specifically targeted to estimating long return period rainfall.

1.3 Structure of the report

This report presents details of the analysis leading to a revised FORGEX procedure for estimating rainfall in the UK for return periods from 100 to over 10,000 years. The report also describes the development of a new model of rainfall depth-duration-frequency that can be applied throughout the UK, and presents example applications of the model at a number of sites. Further information about data sources and the details of particular parts of the analysis are presented in the appendices.

Section 2 gives further information about estimating extreme rainfall and reservoir safety, and outlines the key features of the FEH approach to estimating rainfall frequency. Details of the datasets compiled within this project are given in Section 3, and Section 4 describes some of the initial analyses of annual and seasonal maximum data. The analyses that led to two key aspects of the FORGEX procedure being revised; the standardisation and the model of spatial dependence, are presented in Sections 5 and 6. Section 7 describes the changes proposed to the FORGEX method, together with results from a number of raingauge sites across the UK. The proposed new model of rainfall depth-duration-frequency is discussed in Section 8, and example results and comparisons with other models are presented in Section 9. Section 10 sets out the main conclusions of the project, together with recommendations for further research.

2. Extreme rainfall estimation and reservoir safety

2.1 Introduction

Reservoirs are constructed for several reasons, such as water supply, hydropower and to protect against floods. The safe operation of reservoirs needs to take into account the risk of overtopping of the dam structure from extreme water levels. Whilst design peak flows for other hydrological applications can be estimated directly from observed data using statistical methods, reservoir safety assessment requires the derivation of the entire flow hydrograph to estimate the total volume of water that could enter the reservoir during an extreme rainfall event. The inflow hydrograph is generally obtained through hydrological modelling, where the design rainfall is a key input.

The design rainfall used as input to the hydrological model needs to be appropriate to the catchment, taking into account both the spatial and temporal characteristics of the rainfall event. A simple design rainfall may consist of an estimate of the total event rainfall and an associated temporal distribution, assuming a spatially uniform distribution over the catchment.

There are some 2,600 reservoirs in Britain covered by the Reservoirs Act 1975 (BRE, 1994). These are classified into four categories depending on the potential risk to communities downstream in the event of a dam failure (ICE, 1996). The reservoir design flood inflow for the highest category, Category A, is the probable maximum flood (PMF), with the lower categories B-D having design flow return periods between 10,000 and 150 years.

Therefore, current reservoir design criteria require two different types of design rainfalls; those based on statistical probability and those based on an estimation of what can be considered to be the maximum possible rainfall. The former is addressed through rainfall frequency estimation (Section 2.2) and the latter through probable maximum precipitation (PMP) estimation (Section 2.3).

2.2 Rainfall frequency estimation

Rainfall frequency estimation is concerned with estimating the total event rainfall of a particular duration (regardless of the temporal rainfall distribution within the event) that corresponds to a particular return period. The return period is a common way of expressing the average frequency with which rainfall of a particular magnitude occurs. When the annual maximum rainfall series is analysed, it refers to the average number of years between occurrences of rainfall events exceeding that particular magnitude at a site. The inverse of the return period corresponds to the probability that a particular rainfall is exceeded in any one year.

Frequency estimation methods are underpinned by statistical theory. For hydrological applications, there are important extensions that can be made to traditional extreme value theory, which provide methods to incorporate more data into the analysis than a traditional annual maximum approach. Extended modelling at a single site may include using peak-over-threshold data (that is selecting all independent peaks above a certain threshold, which may include several or no peaks in any one year), or using monthly or seasonal maxima. The single-site estimate can also be improved by using data from other sites with similar characteristics. This is called ‘data pooling’ or ‘regionalisation’. Depending on which method is used, spatial dependence between the sites in the pooling-group (region) may need to be taken into account.

The United Kingdom is one of several countries in the world that uses a variation of the ‘index-flood’ method, where a local estimate of an index variable (typically, the mean or median annual maximum rainfall) is multiplied by a regionally derived dimensionless growth curve to obtain a design rainfall estimate (for example, *Stedinger et al., 1993*). Regionalisation generally involves outlining similar regions, selecting a regional probability distribution function, and estimating parameters. Similar regions may be connected, but they do not have to be. The boundaries of the regions can be fixed, or the regions can be centred on the point of interest so that the boundaries vary for each subject site. Different statistical distributions and fitting methods are used in different countries, but the Generalised Extreme Value (GEV) distribution is often used. The nationwide methods used in nine different countries are described in a separate project report (*Svensson et al., 2006*). Another literature review (*Svensson, 2007*), carried out as part of the project, deals with different methods for deriving areal reduction factors. These are used for converting estimates of extreme rainfall from point values to areal values.

The method of rainfall frequency estimation which is currently recommended for hydrological studies in the UK is outlined in Volume 2 of the FEH (*Faulkner, 1999*). Since it was first published, several papers have been written on the discrepancy in results between the FEH and the earlier design standard, the FSR (*NERC, 1975*). These papers have mainly criticised the extrapolation in the FEH DDF model (fitting a log-Gumbel distribution), and the fact that the dependence model assumed the dependence between rainfall at different sites to be constant rather than decreasing with increasing return period (for example *Babtie Group, 2000; MacDonald & Scott, 2001; Cox, 2003*). Because of these concerns, Defra gave the guidance to panel engineers that the FEH rainfall estimates should not entirely replace the old FSR estimates (*Defra, 2004*). Therefore, both methods have been part of the design standard for the UK since 2004 for different return periods; the FEH for the 193-year return period (to be used when estimating the 150-year flood event), the higher of the two methods for the 1000-year return period, and the FSR only for the 10,000-year return period.

The basis of the FEH method of design rainfall estimation is a six-parameter depth-duration-frequency model. This model was derived from the results of applying the FORGEX (Focused Rainfall Growth Curve Extension) methodology across the UK. The DDF model is implemented on the FEH CD-ROM (*Centre* *for Ecology & Hydrology, 2009*), and the software provides design rainfall estimates for both individual points and catchment areas.

The earlier UK method is outlined in Volume 2 of the FSR (*NERC, 1975*), and also involved the mapping of key rainfall variables. Both the FSR and FEH methods were developed specifically for UK conditions.

Although attempts have been made to transpose both flood and rainfall estimation methods to non-UK conditions, for example the FSR approach has been used in Nigeria and Sri Lanka (*Sutcliffe, 1980*), Zimbabwe (*Bullock, 1987*), Indonesia and other countries (*Meigh et al., 1997*), they have not become established as international practice. This is largely a result of the significant amount of detail and analysis which underlies the variables, and the parameters and relationships that constitute the steps of the methods. Few other countries have prepared this level of information. The semi-empirical relationships between the hydrological model parameters and the catchment and climate descriptors presented in the FSR and the FEH cannot be assumed to apply outside the UK.

The FEH and FSR rainfall frequency estimation methods are described below. A fuller account of different approaches to design rainfall estimation is given in *Svensson et al. (2006*). These include methods used in New Zealand, Australia, Canada, Sweden, France, Germany, the United States and South Africa.

2.2.1 Flood Estimation Handbook

The Flood Estimation Handbook (FEH) uses the FORGEX (Focused Rainfall Growth Curve Extension) method for rainfall frequency estimation (*Reed et al., 1999*) followed by the fitting of a depth-duration-frequency (DDF) model (*Faulkner, 1999*). FORGEX is a development of the FORGE method (*Reed & Stewart, 1989*). The method combines data from circular regions centred on the point of interest, or focal point. FORGEX is an index-flood method, and the frequency curve is obtained by multiplying an index variable specific to the site of interest with a regionally derived growth curve. Because the regions vary slightly for each site, the resulting growth curves vary relatively smoothly across the country. The method is suitable for estimating rare rainfalls with return periods up to 2,000 years. The FEH model was originally calibrated up to a return period of 1,000 years (FEH Volume 2), and subsequently the recommended upper limit was adjusted to 2,000 years (FEH Volume 1). The software on the FEH CD-ROM (Centre for Ecology & Hydrology, 2009) now gives a warning message when estimates for return periods over 1,000 years are supplied.

The index variable

The index variable used is the median annual maximum rainfall at the site, RMED. Values of RMED for eight rainfall durations between 1 hour and 8 days were estimated at stations with record lengths of at least nine years. The estimates will still be susceptible to climatic fluctuations, but the chosen minimum record length is inevitably a compromise to allow adequate spatial coverage of the UK. RMED was mapped over the UK (Figure 2.1) on a 1-km grid using georegression (*Faulkner & Prudhomme, 1998*). This is an extension of the interpolation method known as kriging, which is based on the theory of geostatistics (*Journel & Huijbregts, 1978*). The georegression of RMED uses variables such as the distance from the coast, continentality, elevation and other topographic variables.

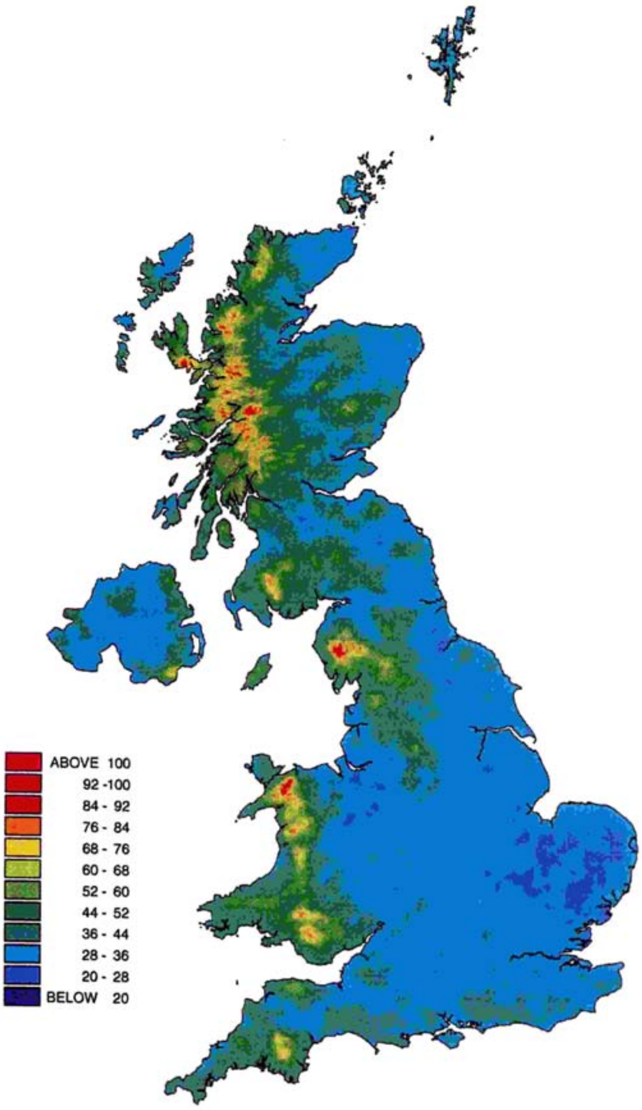

Figure 2.1 1-day RMED over the United Kingdom (mm) (from Faulkner, 1999)

The growth curve
FORGEX combines data from a hierarchy of expanding circular regions centred on the point of interest (Figure 2.2). Data from smaller networks are used to estimate the growth curve for short return periods, whilst data from the larger networks are used for the longer return periods. The size of a region is a compromise between keeping a region small, and therefore relevant, and avoiding excessive extrapolation beyond the information available in the region

to obtain high return period rainfall estimates. An upper limit of 200 km is used for the largest network radius.

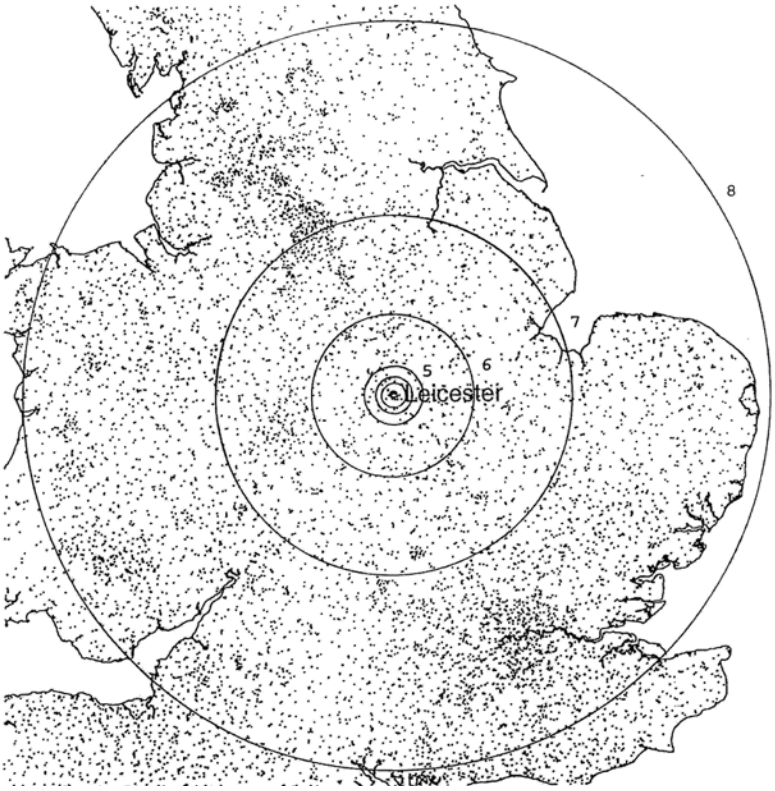

Figure 2.2 Networks (regions) of daily raingauges focused on Leicester (from Faulkner, 1999)

A complex empirical approach is used to produce the growth curve. This combines a regional frequency analysis with an analysis of network maximum points whose plotting positions are shifted to allow for spatial dependence of extremes.

Standardised annual maxima (standardised through division by RMED) from individual records with a minimum length of 10 years are ranked and plotted on a Gumbel reduced variate scale using the Gringorten plotting position. Note that data for different sites are superposed on the plot (Figure 2.3), rather than being pooled into a single long series.

The Gringorten plotting position is expressed as

$$F(i) = \frac{i - 0.44}{N + 0.12} \ ,$$

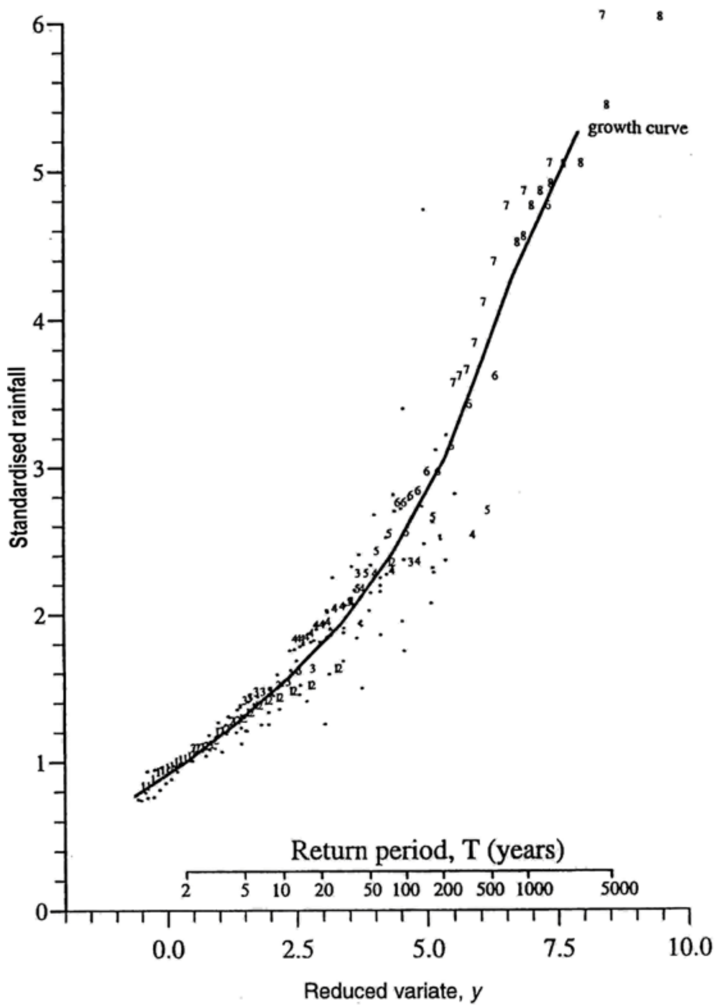

Figure 2.3 Individual points (dots), netmax points (numbers) and fitted empirical growth curve for 1-day rainfalls focused on Leicester (after Faulkner, 1999)

where $F(i)$ is the non-exceedance probability, i is the rank of the observation in increasing order, and N is the number of annual maxima.

The Gumbel reduced variate, y, is defined as

$$y = -\ln(-\ln F).$$

Not all points are plotted, and only points which are plotted are used to fit the growth curve. Data from points within each progressively larger region are used to estimate sections of the growth curve corresponding to progressively greater return periods. Thus, data points from within the jth network are only plotted if their plotting position falls within the jth section of the growth curve. Each section, or y-slice, has width 1.0 on the Gumbel reduced variate scale, starting at y = 0.3665, which corresponds to the position of the median (T = 2 years). These points are plotted as dots in Figure 2.3. Because few stations have more than about 100 years of records, there are few points plotted at a return period larger than 200 years.

The network maximum (netmax) series is defined as the annual maximum series of the largest standardised value recorded anywhere within the region. Therefore, there is one netmax value per year. These points are plotted as numbers in Figure 2.3, the numbers referring to the number of the region (Figure 2.2). Because of spatial dependence in the network of raingauges, the plotting positions for the netmax points have been modified using a spatial dependence model (see next section).

An empirical growth curve consisting of concatenated linear segments is fitted to the plotted points (both individual and netmax) through a least-squares routine (Figure 2.3). Finally, a depth-duration-frequency (DDF) model is fitted to avoid any contradictions between durations or return periods. This is an exponential curve on the Gumbel reduced variate scale, fitted to all durations (1, 2, 6, 12 hours, and 1, 2, 4, 8 days) and return periods (2, 5, 10, 20, 50, 100, 200, 500, 1000 years) at once, using a least-squares criterion.

The spatial dependence model

Dales & Reed (1989) show that the distribution of the network maximum from *N* independent and identically distributed Generalised Extreme Value (GEV) distributions lies exactly ln*N* to the left of the regional growth curve (the 'typical' growth curve for a station in the region) on a Gumbel reduced variate scale (Figure 2.4). Reed & Stewart (1994) note that this result is not restricted to the GEV. In practice, because of inter-site dependence in annual maxima, the netmax growth curve is found to lie a shorter distance to the left. Dales & Reed label this distance $\ln N_e$, terming N_e the effective number of independent gauges. Therefore, spatial dependence can be assessed from the relationship between the typical and network maximum growth curves. Conversely, the fitting of the regional growth curve can be helped by information on spatial dependence. If an estimate of N_e is available, the top part of the netmax series can provide valuable information to guide the extension of the regional growth curve to long return periods.

Dales & Reed (1989) developed a model of spatial dependence to estimate the offset distance, $\ln N_e$, for any raingauge network. The UK model relates dependence to the geometry of the raingauge network:

$$\frac{\ln N_e}{\ln N} = 0.081 + 0.085 \ \ln AREA - 0.051 \ \ln N \ - 0.027 \ \ln D$$

where *N* is the number of gauges and D the rainfall duration in days. AREA is a nominal area spanned by the network.

For each network in turn, the netmax series is constructed and the values ranked. If the same gauges were available for each year of record, the plotting position of each point would effectively include a shift to the right by $\ln N_e$. Because the number of operational raingauges, *N*, has varied greatly over the past two centuries, a more complicated determination of the plotting positions is needed, as described by Jones (1997).

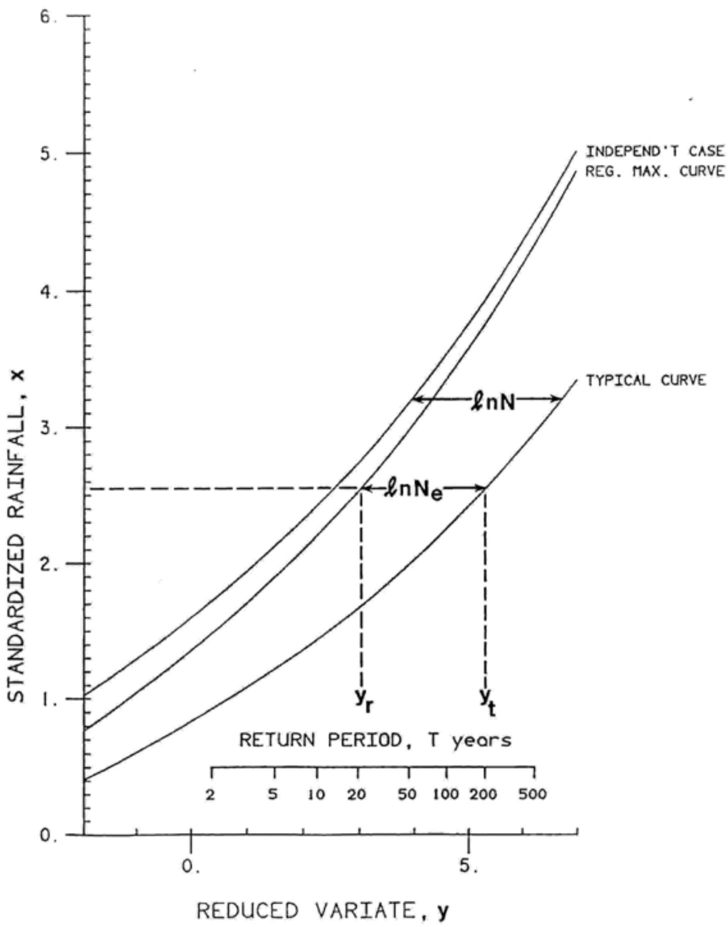

Figure 2.4 Growth curves for the independent case (the network maximum, provided that rainfalls at all the stations in the network are independent of each other), the regional maximum curve (the network maximum, for the realistic case where there is dependence between rainfalls at the different stations in the network), and the typical curve (the regional growth curve) (from Dales & Reed, 1989)

2.2.2 Flood Studies Report

Results obtained using methods in the Flood Studies Report (*NERC, 1975*) are often compared with results calculated using the newer FEH methodology. Below is a brief outline of the main ideas behind the rainfall frequency estimation method used in the Flood Studies Report (FSR).

The FSR method involves mapping an index variable, the 5-year return period rainfall, and producing regional growth curves. The design rainfall is obtained by multiplying the value of the index variable at the site of interest with the regional growth factor for the desired return period. The UK is divided into two regions: England and Wales, and Scotland and Northern Ireland. Dividing the UK into only two regions means that rainfall frequency estimates show little spatial variability, and this method has been criticised for not capturing local variations well.

The index variable

The 5-year return period rainfall, M5, for durations of 2 and 25 days was estimated for more than 600 stations with an average of 60 years of record. This was supplemented by estimates at another 6,000 stations with records from 1961 to 1970. The M5 was determined using a quartile analysis (outlined later) of the ordered annual maximum rainfall series. It was noted that the 5-year return period rainfall is very close to the value of the geometric mean of the upper two quartiles (the upper half) of the series of annual maximum rainfall. Therefore, the M5 was estimated as the geometric mean of the upper half of the annual maximum rainfall series.

Maps were prepared for the average annual rainfall (AAR), 2-day M5, the ratio (2-day M5)/AAR, and the ratio (25-day M5)/AAR. For intermediate durations, *D*, the best estimates of M5 rainfall are obtained by linear interpolation on a plot of the logarithm of M5 against the logarithm of *D*. For durations between 60 minutes and 48 hours, the M5 rainfall is obtained through a relationship with the ratio (60-minute M5)/(2-day M5). This ratio is derived using relationships with the M5 value of precipitable water, the number of days per year with thunder and the 2-day M5 rainfall.

The growth curve

The standardised growth curves were derived from non-standardised rainfall frequency curves for the two regions (England and Wales, and Scotland and Northern Ireland). First, frequency curves were fitted graphically using the results of quartile analyses of the annual maximum rainfall series, as described in the next section. It was noted that when the logarithm of the rainfall was plotted against the Gumbel reduced variate, the frequency curve is linear. Frequency curves for different durations and classes (range of M5 values within the same duration) were plotted on the same diagram and it was noted that the slope of the curve decreases with increasing M5 rainfall amounts in a systematic manner regardless of duration. Therefore, the shape of the curve does not depend on duration, but only on the rainfall amount (the M5 value). Finally, dimensionless growth curves were derived by dividing the frequency curves by the M5 rainfall.

Fitting a frequency curve using quartile analysis

If a series of *N* annual rainfall maxima are arranged in ascending order, and plotted using plotting positions, a graphically fitted smooth frequency curve can be drawn. The FSR uses the plotting position suggested by Chegodayev (1953), $(m-0.31)/(N+0.38)$, where m is the order of the maxima. However, it is also noted that summary statistics from a quartile analysis of (long) series can be used to fit a frequency curve, making it easier to combine many datasets into a comprehensive regional analysis.

The ordered series of annual maxima are divided into four quarters, or quartiles. The geometric means of each quartile can be shown to correspond to particular return periods. For example, the mean of the highest quartile corresponds to the 10-year rainfall. In addition to the quartile means, the means of two halves of the series are also calculated. The mean of the middle half (the middle two quartiles) corresponds to the 2-year rainfall (M2) and the mean of the upper half corresponds to the 5-year rainfall (M5, which is used as the index variable). The four highest observations in each series are also noted (denoted H1, H2, H3 and H4) together with the number of years in each record.

These statistics were calculated for all the stations with long records (average 60 years) in each region. The series were divided into classes, depending on the magnitude of the M5 rainfall. For example, there are 175 stations in the England and Wales region with a 2-day M5 between 40 and 50 mm. The quartile statistics and four highest observations were calculated for all these 175 series separately. Subsequently, the median value of each statistic/observation was obtained, reflecting regional 'average' statistics/observations. It is worth noting that the FSR definition of the 'median' differs from most current textbooks. The 'median' here refers to the mean of the middle half in a quartile analysis. These 'averaged' quartile statistics were used to plot points on the lowest part of the frequency curve. The average record length was used to calculate plotting points for the 'averages' of the highest four observations, which plot slightly higher on the frequency curve.

The upper part of the frequency curve is informed by quartile/highest observation statistics derived from the series obtained when selecting only the largest value from each annual maximum series. In the example above, that means a series of 175 H1 values. Plotting positions for the quartile means are estimated (different to those for the original series, above) and added to the diagram. If the network of stations is not too dense, the four largest values in the H1 array can be plotted using the station-year method. That is, if observations are independent, space can be substituted for time, and the largest four values can be assumed to be the largest in a very long, pooled, time series of 60 years × 175 stations = 10,500 station years. Therefore, points can be plotted corresponding to return periods of several thousands of years.

2.3 Probable maximum precipitation estimation

Values of probable maximum precipitation (PMP) are required for the design of reservoirs in category A, the highest safety class (ICE, 1996). In this and similar contexts, these values are somewhat loosely defined as the largest rainfall amount that is likely to fall. This might be interpreted as meaning that there may be a very small but positive probability that the value might be exceeded. However, many formal definitions of what is meant by a PMP value imply that it is a rainfall amount that can never be exceeded (WMO, 1986), at least if the true value of PMP could be known. In practice, the values used for PMP are only estimates and, for this reason, there may be a positive probability of an estimated PMP value being exceeded. This probability of exceedance arises from the estimation uncertainty associated with the PMP value.

This section relates to methods for estimating PMP treated as a strict upper bound. The next section (Section 2.4) relates to analyses which start from the idea that there is an upper bound and attempt to quantify the probability of that amount of rainfall actually occurring. This probability has a different character from that discussed above in relation to the uncertainty associated with the estimated upper bound. Theoretically, it would be possible to envisage extreme types of climate in which rainfalls close to PMP occur every year. To a certain extent, if one accepts the idea of a fixed upper bound to the amount of rainfall, the frequency of occurrence of PMP is unimportant for reservoir design because it is only necessary to design for a slightly larger rainfall which, by definition, can never occur. Treatment of the potential occurrence of a closely spaced sequence of storms which might lead to a reservoir failure is notionally covered by using PMP values for a range of durations. However, this raises the more difficult question of assessing the frequency of occurrence of such sequences of storms.

Section 2.5 relates to approaches to defining a value for PMP which do not necessarily consider that there is an upper bound to the amount of rainfall that might fall, but instead proceed by defining a statistical distribution for the amount of rain and then derive the rainfall amount which has a very small probability of being exceeded according to this distribution. Technically this is not a PMP value, as it is not an upper bound, but values derived in this way are often called PMPs. To distinguish these values from PMPs which specifically represent an upper bound, this report will refer to 'frequency-based values of PMP'.

The exceedance probability used to define the value of PMP in the frequency-based approach is similar to the probability discussed above for how often a PMP-sized event will occur when PMP is treated as an upper bound. However, the true probability of exceedance of the PMP value would differ from this because the threshold (in this case the PMP) has a value which is contaminated by estimation uncertainty. The uncertainty ranges around the PMP estimates (in contrast to the occurrence rates discussed above) are similar for both the frequency-based value of PMP and an upper-bound PMP, because, in the former case, the statistical distribution from which the value is derived must be estimated.

Since methods for finding frequency-based values of PMP can be used to determine the exceedance probability of any given rainfall amount, they can also be used to evaluate the exceedance probability of a value of PMP determined in any other way, and in particular that of a PMP value which has been evaluated on the basis that it is an upper bound to the rainfall amount.

In the rest of this subsection, a true PMP value is treated as an upper bound to the amount of rainfall that can occur. The frequency of occurrence of upper-bound PMPs is discussed in Section 2.4, while Section 2.5 deals with frequency-based values of PMP.

The basic reference for the estimation of PMP was produced by the World Meteorological Organization (WMO, 1986). This covers storm maximisation by reference to meteorological structures, storm transposition and statistical extrapolation. Storm maximisation uses the physical structure of storms in the atmosphere, particularly vertical temperature and humidity relationships, and extrapolates these to physical limits. The methods were thoroughly discussed and demonstrated in Wiesner (1970), which focuses on practice in the USA.

Much work has been carried out in the UK using weather radar to investigate the detailed structure of storms, for example by Browning & Hill (1984). The establishment and regular operation of the radar rainfall network in the UK as part of the observation system routinely produces information about the areal extent and intensity of rainfall-producing systems. Some detailed research has also been carried out to improve forecast capability, and also particular studies of PMP, e.g. Austin *et al*. (1995).

The general lack of studies into the detailed physical structure of major storms in the UK, from the point of view of flood risk, may be because such events are comparatively rare. However, there is potential for these events to occur most summers in the UK. The existence of mesoscale convective systems (MCSs) was recognised after the publication of the FSR and over 30 probable MCSs between 1981 and 1997 were identified by Gray & Marshall (1998). Although the present study does not re-evaluate PMP, the dataset of 50 extreme rainfall events compiled by Collier *et al*. (2002) has been updated to include 63 events occurring from 1880 to 2006. This work shows that there are clear links between synoptic rainfall type and rainfall depth and duration.

2.3.1 Methods for estimating probable maximum precipitation

Techniques used for estimating PMP have been listed and discussed by Wiesner (1970), and summarised by Austin *et al*. (1995):

(a) The maximisation and transposition of actual storms.
(b) The use of generalised data or maximised depth, duration, area data from storms. These are derived from thunderstorms or more general longer duration storms.
(c) The use of empirical formulae determined from maximum depth, duration, area data, or from theory.
(d) The use of empirical relationships between the variables in a particular valley (for example relating intensity to wind velocity and surface dew point using past extreme storms in the same valley (Wiesner, 1970, Section 18.7)).
(e) Statistical analyses of extreme rainfalls.
(f) The storm model approach.

Methods (a) and (b) involve the classification of storms by calculating the corresponding storm efficiency (E), which is defined as the ratio of maximum observed rainfall to the amount of precipitable water in the representative air column during the storm (NERC, 1975). Collinge *et al*. (1992) describe this approach.

If no vertical soundings are available, it is assumed that the air mass in the storm is saturated and the vertical humidity profile is represented by the screen level dew point, following the saturated pseudo-adiabatic lapse rate. The precipitable water can then either be calculated by using tephigrams to determine the mixing ratio, or using tables which directly give the precipitable water as a function of height and dew point (see Wiesner, 1970). PMP is estimated using the climatological maximum dew point at the surface located at particular places.

The return period associated with PMP is discussed in Section 2.4.

Generalised estimates of PMP

Probably the oldest and most widely used method of estimating PMP is storm maximisation and transposition. Alexander (1963) outlines a method for determining the probability of a transposed storm occurring, centred over the study catchment. Transposition works best in flat areas where orographic effects can be neglected.

Normally, base maps are produced for each study region which depict a number of isohyets of different size and duration for specific events. The base maps are standardised for maximum persisting dew point temperature so that it is simple to adjust the totals (adjust and maximise moisture availability) for application to the catchment. Choice of the suitable base map is a function of the size and response time of the catchment being studied.

This approach has been used in many countries, and has led to the following type of formula being produced:

$$\mathrm{PMP} = \mathrm{FAFP}\left\{\frac{T}{C} - m^2\left(\frac{T}{C} - 1\right)\right\},$$

where FAFP = Free Atmosphere Forced Precipitation

- m = storm intensification factor, which is the ratio of the storm core rainfall to the total rainfall in the storm. It is a measure of the concentration of the rainfall in the core phase of the storm and usually has values around 0.6
- T = the 24 hour rainfall (or other duration), which has a 100 year return period
- C = the convergence component of T

Factors which contribute to the FAFP are:

- convergence
- forcing from mesoscale weather systems
- fronts and baroclinic waves
- condensation efficiency

Recent developments of this approach have been described by the Australian Bureau of Meteorology (*Bureau of Meteorology, 2003*).

Empirical formulae

Empirical formulae have been derived representing local or world maximum values of precipitation – usually point values rather than areal averages. In the UK, the most extreme rainfalls for varying durations have been analysed by Hand *et al*. (2004). For durations up to 2 hours for convective rainfall, an almost linear relationship between depth and duration can be found on a plot with log-log scales. This may be represented by the following approximate formula for point rainfalls:

$$R = 100\sqrt{D},$$

where R is the rainfall in millimetres and D is the storm duration in hours.

The FSR procedure

The FSR presents maps of estimated maximum rainfall for durations of 2 and 24 hours for the UK and Ireland. Catchment specific values can be obtained by calculating the average over the catchment. For specific duration PMP between 2 and 24 hours, a straight line is plotted on a graph of PMP rainfall depth versus the logarithm of PMP duration. An areal reduction factor, based on catchment area and duration, is applied. The choice of PMP storm profile is based upon historic data, but is assumed to be symmetrically distributed according to the extreme profile such that the estimated maximum occurs in every duration centred on the peak of the storm profile. These PMP estimates were not updated for the FEH.

Areal reduction factors

If a value of PMP has not been evaluated to correspond directly to an areal average rainfall accumulation, but instead relates to a point rainfall, the value will need to be adjusted. Areal reduction factors (ARFs) are usually presented in graphical form (Area versus Duration), for example in the FSR (NERC, 1975). Care must be taken, however, as the empirical basis for such factors can mean that they do not necessarily apply to PMP. Some recent research (for example, *Skaugen, 1997*) has suggested that ARF should decrease with increasing return period: see also Svensson (2007).

Storm models

The determination of the PMP from a storm model requires a maximisation of all the variables which cause precipitation. The rainfall intensity in a storm model is normally a function of the geometry of the catchment, inflow velocity of the air, a moisture factor and a convergence factor. The theoretical relationship between convergence, vertical motion and condensation to form precipitation is well known. Storm models describe a simplification of these processes. Collier & Hardaker (1996) describe an example of this type of technique. They found that such a model generated PMP values very similar to the FSR for storm durations less than 11 hours, but for durations between 11 and 24 hours the FSR PMP values were exceeded. It was assumed that rainfall lasting between 11 and 24 hours was due to mesoscale convective systems (MCSs) as described by Browning & Hill (1984).

2.3.2 Uncertainty analysis of PMP estimates

Dent (2008) revisited the PMP estimates of the FSR (NERC, 1975), and assessed the uncertainties in this and some other methods of estimating PMP.

The FSR lists three methods for PMP estimation. Dent (2008) gives a rough estimate of the uncertainties in the PMP estimates by assuming possible error bands for the variables used as input for the calculations. For a south London location, he finds that 2-hour storms give a wide spread of results, from 168 mm to 230 mm. The ‘growth factor’ and ‘quick’ methods produce lower estimates than the method using maximisation of precipitable water. The latter is similar to the commonly used WMO method (*WMO, 1986*). Only two estimates were obtained for 24-hour storms, and these were both very similar; 265 mm and 267 mm. This is presumed to reflect that both estimates rely on using the same index variable (M5) and similar growth factor relationships. It can be noted that, for both durations, the lower ends of the ranges of estimates are exceeded by observed rainfall events.

The above uncertainty estimates relate to PMP values calculated for a single location. However, Dent (2008) also revisited the seven short-duration storms located throughout the UK that were used for the FSR analysis. To these he applied the standard WMO (1986) method of estimating PMP, which maximises the precipitable water available for the storm. The resulting PMP estimates are shown as red dots in Figure 2.5, together with the magnitudes of extreme events observed between 1900 and 1998 (from Collier *et al*., 2002). The new PMP estimates all exceed the magnitudes of storms observed so far for the corresponding durations.

2.4 Frequency of occurrence of PMP

This section summarises some analyses that have attempted to estimate how often an event with the size of the PMP will occur; a fuller treatment is given in Appendix M. Two assessments have been produced, both based on the same underlying approach. The common approach is to start from the storm model of Collier & Hardaker (1996), which was designed to estimate values of PMP (as an upper bound). Then, first identify the components of the model which will most important in extreme rainfall events, and, second, consider the combination of meteorological conditions which will lead to high values for this component. Finally, probabilities are assigned to these combinations of conditions.

The first assessment applies to thunderstorms, while the second applies to mesoscale convective systems (MCSs). In each case, the analysis centres on estimating how often events close to PMP in size will occur rather than on estimating a value for PMP. Thunderstorms are relatively frequent and affect limited areas, while mesoscale convective systems are rarer and affect rather larger areas.

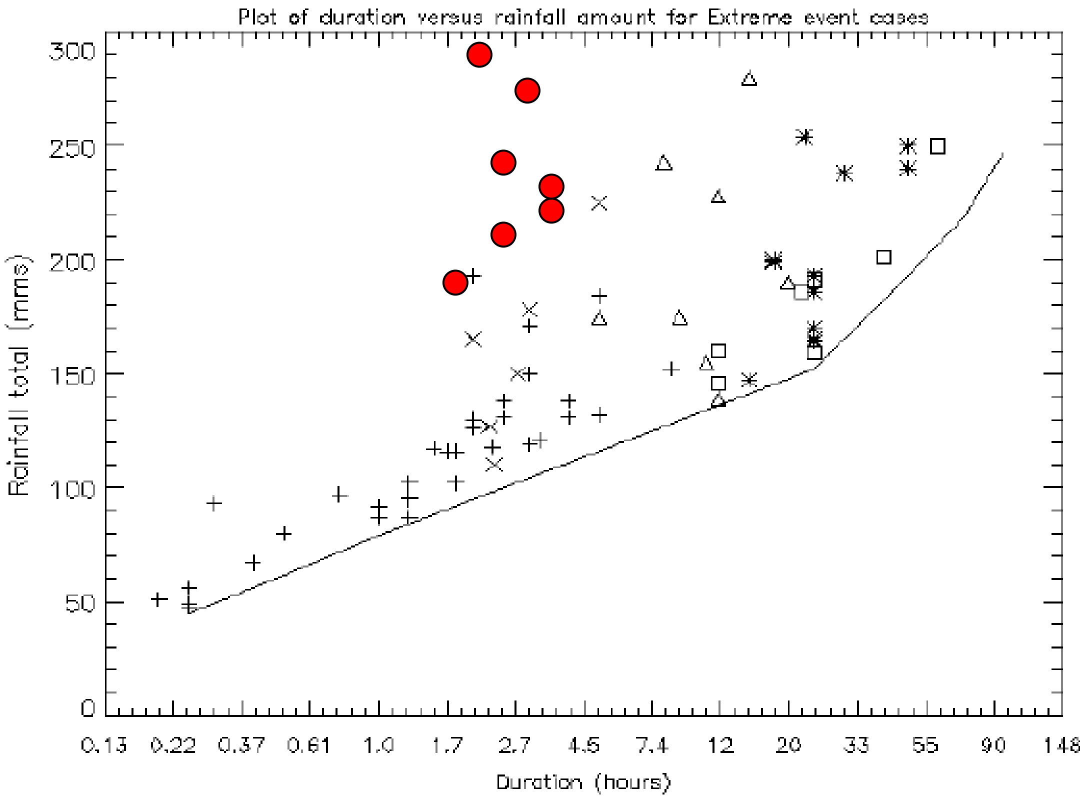

Figure 1b. Extreme rainfall threshold line, all 20th and 21st Century Extreme cases represented by a symbol according to rainfall duration and amount and synoptic type. '+' = convective, 'X' = convective with frontal forcing, ж = orographic, Δ= frontal with embedded convection and □ = frontal.

Figure 2.5 Storm maximisation estimates of FSR events in relation to Defra Extreme Events threshold and events observed 1900-1998 (from Collier *et al.*, 2002)

2.4.1 Thunderstorms

The conclusions found in Appendix M are as follows. They are based on the observed frequency of severe thunderstorms in England and Wales, and assumptions about the likelihood of (i) their occurrence over a particular catchment, (ii) the occurrence of wind speeds in the direction of the maximum orographic gradient in the catchment and (iii) the occurrence of maximum surface heating. For 10-hour rainfall approaching PMP over an individual 200 km^2 catchment, a return period of about 2×10^5 years is estimated. For an individual point, the estimated return period is about 1×10^4 years.

2.4.2 Mesoscale convective systems (MCSs)

A summary of the analysis in Appendix M is as follows. Considering the frequency of only MCS events, and making assumptions as above, the 15-hour rainfall approaching PMP over an individual large (1000 km^2) catchment has an estimated return period of about 5×10^5 years.

2.4.3 Other estimates of return period

These estimates can be compared with the return period of PMP of about 5×10^5 years which Lowing & Law (1995) estimated from the FSR (*NERC, 1975*) growth curves, and the 3.5×10^4 years return period of the envelope line of observed events at the time of the FSR study (FSR II 4.3.2).

2.5 Frequency-based values of PMP

As discussed earlier (Section 2.3), if values of PMP are only considered to be rainfall amounts that are very unlikely to be exceeded, they can be estimated from a statistical distribution representing the full range of possible outcomes. Perhaps the simplest approach would be to fit an empirical distribution to an observed record of annual maximum rainfall. For example, Chow (1951) (*see also Hershfield, 1961, 1965*) suggested the estimation of the magnitude of extreme rainfalls using a relationship between the value of rainfall (X_T) for a given duration which has a given return period (T,) as a function of the average annual maximum (m), the standard deviation of the annual maximum series (s) and a frequency factor (K_T) which is a function of the climate of the region. The relationship used is:

$$X_T = m + K_T s .$$

Possible outliers in the observed data and postulation of the correct distribution of annual maximum rainfall (which implies K_T) both cause problems for this method. Early work used a Gumbel distribution to define the frequency factor K_T, but distributions such as the GEV have also been used. Many possible methods of fitting distributions to observed rainfall amounts are available, including POT methods. While initially attractive, this general approach might be considered likely to lead to erroneous results because of the use of a distribution fitted to typical values of rainfall (albeit of annual maximum) to extrapolate out to very rare return periods. Because of this doubt, other approaches, which concentrate on modelling the most extreme events, have been developed and one of these, the storm transposition method which has been applied in the USA, is described in Appendix M. However, the main topic of the present study is developing a method for estimating how often a rainfall amount of any given size will be exceeded, anywhere in the UK, and therefore this can be applied to existing values of PMP. The results of the model when applied to PMP are reported in Section 9.3. The approach attempts to make the maximum possible use of rainfall information across all gauges in a region, and across several durations, to reduce the extrapolation effect.

If a rainfall frequency model is to be used to determine a typical extreme value of rainfall for use in design studies as a surrogate or alternative to PMP, there is no definitive value for the annual exceedance probability that has found universal acceptance as giving a rainfall amount which is broadly equivalent to the PMP, although it is sometimes taken to be about 1 in 10^5. When used in the context of reservoir design (*ICE,1996*), which is summarised by Table 1.1, it is clear that the return period should be greater than 10^4 years.

2.6 Summary

This section has described the rainfall frequency estimation methods used in the UK (FSR and FEH) and which represent a starting point for both existing knowledge about extreme rainfalls in the UK and the methods used for estimating these extreme rainfalls. The work in this report builds on this foundation.

The section has also summarised several methods of estimating the value of probable maximum precipitation (PMP) which, while it is not itself a subject for the present project, is important in reservoir design, as it is incorporated into established design recommendations. This section has also discussed existing work on assessing how often PMP-sized events will occur.

The outcome of the work in the present study is a model which can be used to assess how often a given rainfall amount will be exceeded, and so it can also be used to assess the return period of pre-defined values of PMP. The outcomes of such calculations are presented later in Section 9.3. In particular, Figures 9.14-9.30 (even numbers) and Tables 9.2-9.10 show comparisons of estimated rainfalls resulting from both the new model and existing models with the PMP values evaluated for the Flood Studies Report (the main source of PMP values for the UK). Tables 9.12 (a) and (b) allow a comparison of values of PMP against the sizes of some of the most extreme rainfall events experienced in the UK.

3. Data available for the study

Data from a number of different sources were made available for this study. The main statistical analysis approach has been to adopt a similar strategy to that used to develop the FEH rainfall frequency model (*Faulkner, 1999*). Therefore, datasets of annual rainfall maxima for a series of durations ranging from 1 hour to 8 days have been gathered for as many sites as possible in the UK. In addition, wherever possible, summer and winter maxima have been abstracted for the same rainfall durations, which has allowed an evaluation of seasonal rainfall frequency across the UK. In this report, the raw data, from which these sets of annual and seasonal maxima have been derived, are referred to as the 'systematic' data record.

The project has also used meteorological information on 63 notable storm events experienced in the UK since the late nineteenth century. This database of extreme events has been used as a 'reality check' against which the final estimates of rainfall frequency can be compared.

A summary of the issues surrounding data collation and the abstraction of maxima is given here, while further details are available in a separate project report and appendices (*Svensson et al., 2009*)

3.1 Sources of systematic data and annual maxima

The analysis carried out to develop the FORGEX method of rainfall frequency estimation is described in Volume 2 of the FEH (*Faulkner, 1999*). Annual maxima for durations of 1, 2, 4, 6, 12, 18 and 24 hours and 1, 2, 4 and 8 days were analysed. These data originated from a variety of sources, some taken from continuous hourly and daily records, and some from archive material used in the FSR analysis carried out in the early 1970s. The annual maxima used in the FEH analysis are still available in computerised form at CEH Wallingford and were readily available to the current study. However, not all of the continuous hourly raingauge records, from which annual maxima were abstracted, still exist following several changes of computer system and recent developments in storage media and data handling. Therefore, at the start of the project, it was decided to use as much of the original information from the FEH as possible, and to update the continuous hourly and daily raingauge records where feasible.

3.1.1 Daily rainfall data

Daily rainfall depths for a dense network of over 6,000 gauges in the UK are supplied to CEH by the Met Office on an annual basis and these data were available for use within the current project. The daily accumulations are recorded at 9 am GMT each day and the longest records in this dataset start in 1961 and go up to the present day. The daily gauge network is shown in Figure 3.1. The project was also able to use daily data from a smaller

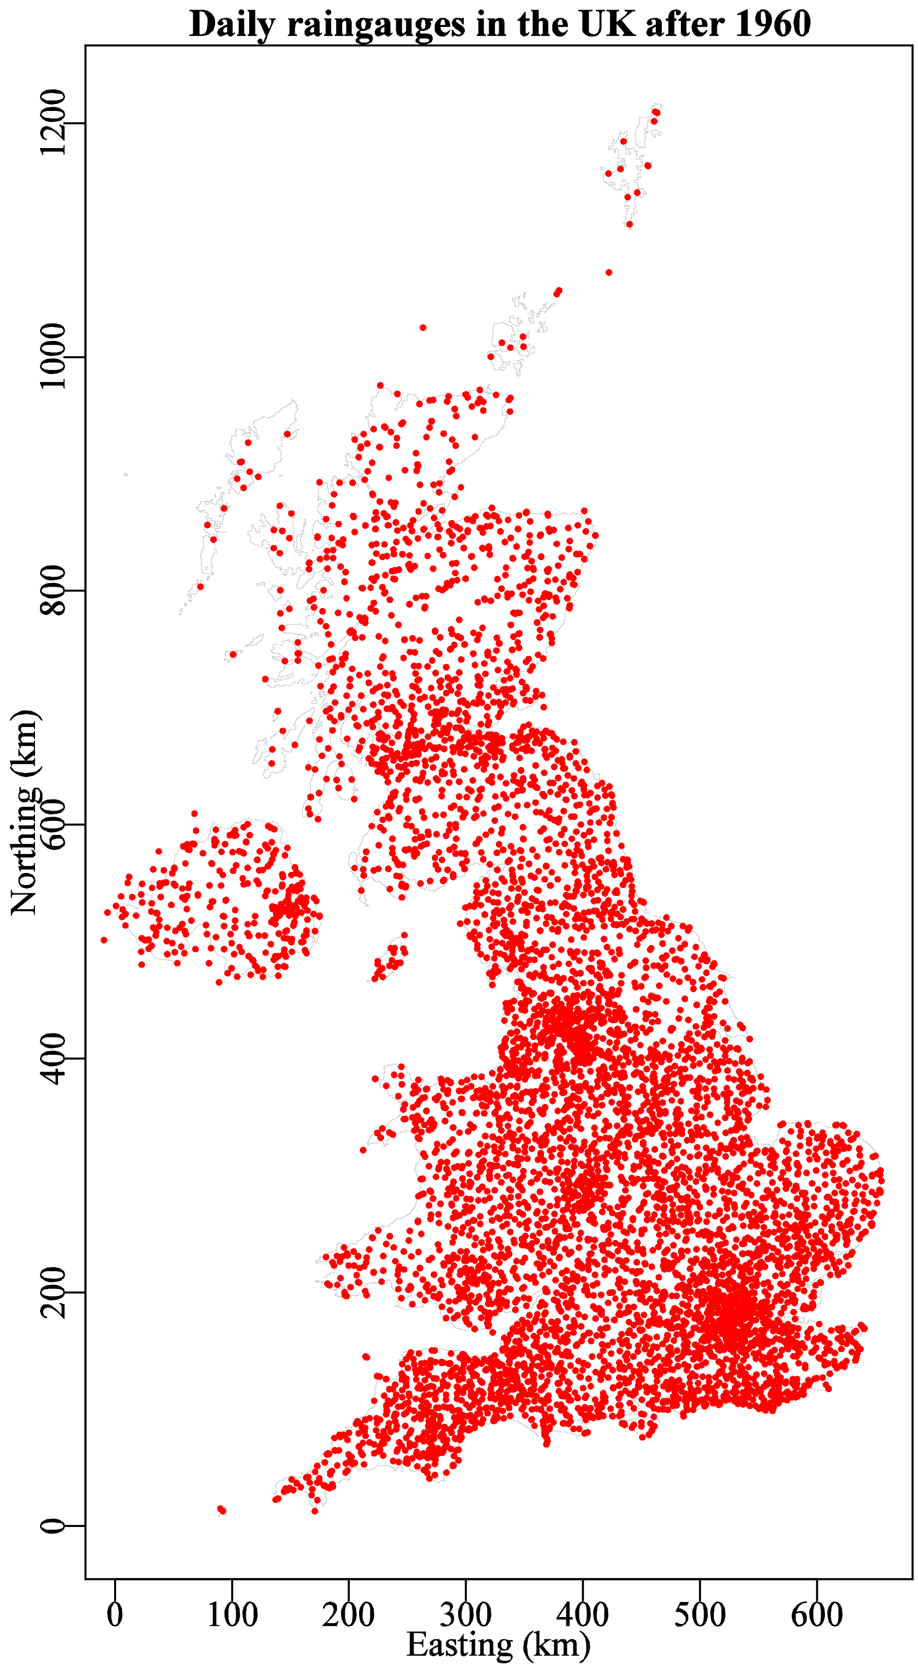

Figure 3.1 Locations of daily raingauges in the UK and the Republic of Ireland with digitised records in the period 1961 - 2005

network of raingauges with digitised records in the period 1853 to 1960 (*see Figure 3.2*). Each gauge within this set of over 500 long-term records has at least 40 years of data. Using these datasets, it was possible to update the FEH annual maxima to 2004 for durations of 1, 2, 4 and 8 days.

Met Éireann also made continuous daily rainfall records available for 30 gauges in the Republic of Ireland, located near the border with Northern Ireland. The earliest of these records starts in 1941 and the latest ends in 2005. A total of 13 gauges have records that start before 1961.

3.1.2 Hourly rainfall data

The Environment Agency provided continuous hourly rainfall data for 1,005 gauges in England and Wales for the project. Continuous hourly rainfall data for another 136 gauges in Scotland were supplied by the Scottish Environment Protection Agency (SEPA). The team was advised that no similar records for gauges in Northern Ireland existed other than for gauges forming part of the Met Office network. The Met Office supplied continuous hourly records for 472 raingauges throughout the UK, providing a total of over 19,000 station-years (Figure 3.3). The hourly data comprise rainfall depths recorded on the clock hour. The records received span 1961 to 2006 for the Environment Agency, 1983 to 2006 for SEPA and 1949 to 2007 for the Met Office data.

3.2 Instrumentation

The raingauge data used in this study came from a variety of instruments depending on the supplier of the data and the relevant time period. Various types of storage gauge were used to obtain daily accumulations of precipitation, whereas recording gauges were used for sub-daily records. The latter include tilting siphon (Dines) and tipping bucket (TBR) gauges (for example, *Met Office, 1982*).

Information on the type of gauge used at each site and time period is not available for all the records used in this project. Instruments may also have been replaced on the same or a similar site. This means that there are likely to be inconsistencies in the capture and recording of the rainfall data, particularly for the recording gauges.

All collection-type raingauges (as opposed to for example optical) experience losses in high winds depending on how sheltered their position is and the shape of the gauge (for example, *Hughes et al., 1993*). Other problems include splashing in and out of the gauge and evaporation. Tilting siphon gauges experience losses when they tilt and siphon, particularly during heavy rainfall, although the loss is reduced by the use of a rain trap (for example *Shaw, 1988*). Various studies have found that compared with storage gauges, tipping bucket gauges under-record the amount of rain falling during heavy events, because water flows into the buckets while they are tipping (*Niemczynowicz, 1986; Nystuen, 1999; Hodginson et al., 2004*). Better precision can be obtained

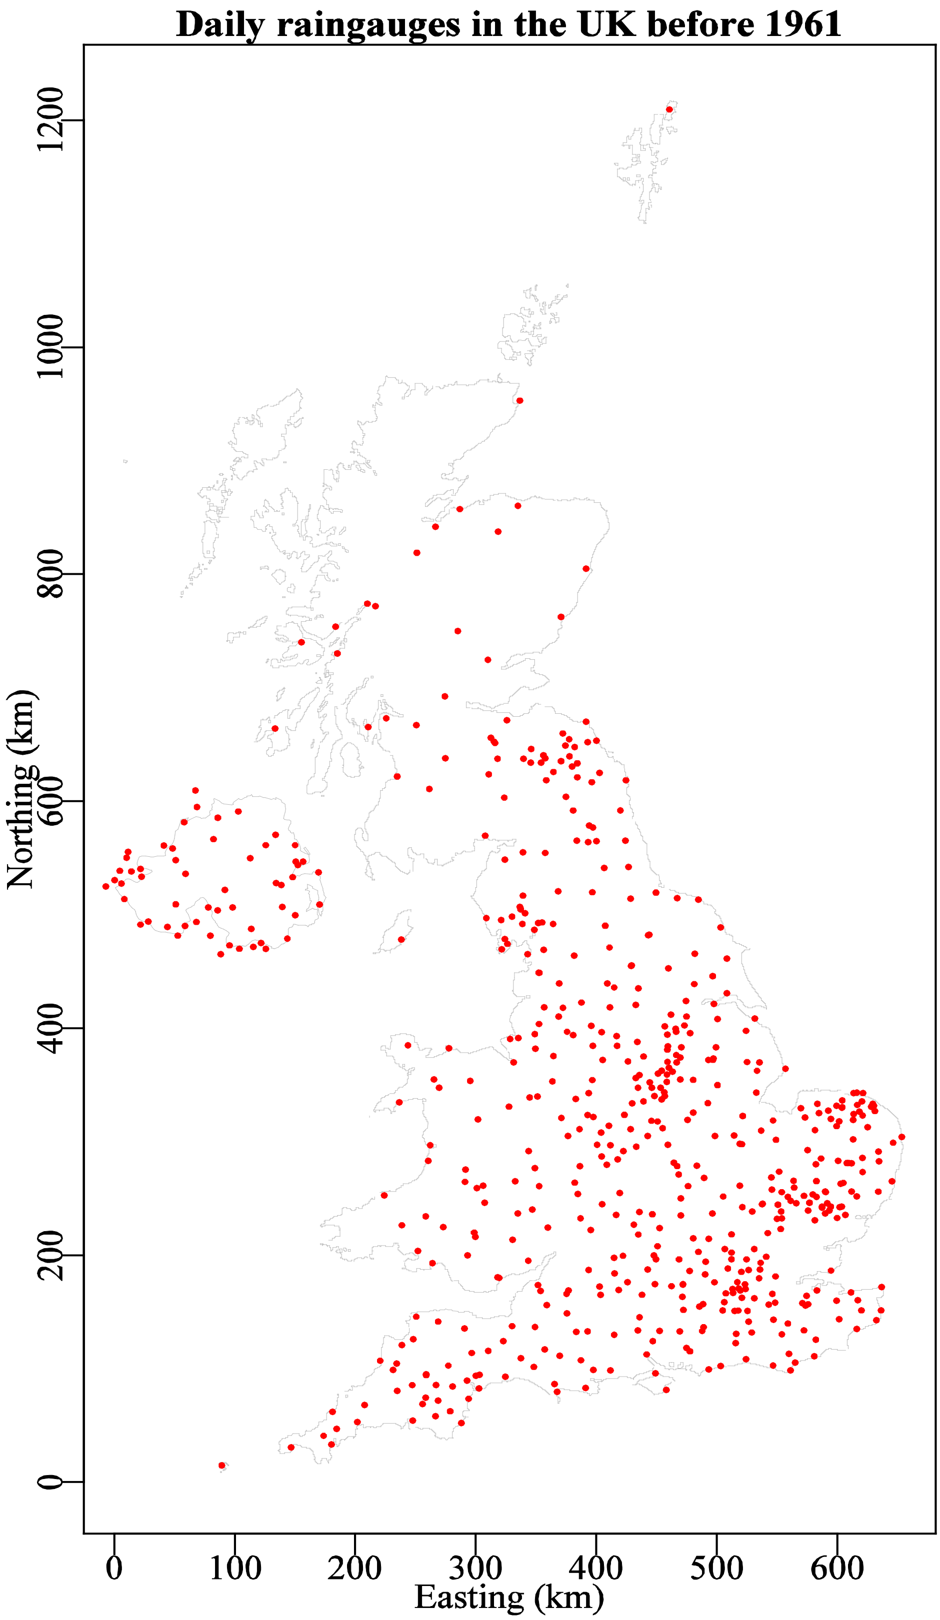

Figure 3.2 Locations of long-term daily raingauges in the UK and the Republic of Ireland with digitised records within the period 1853 - 1960

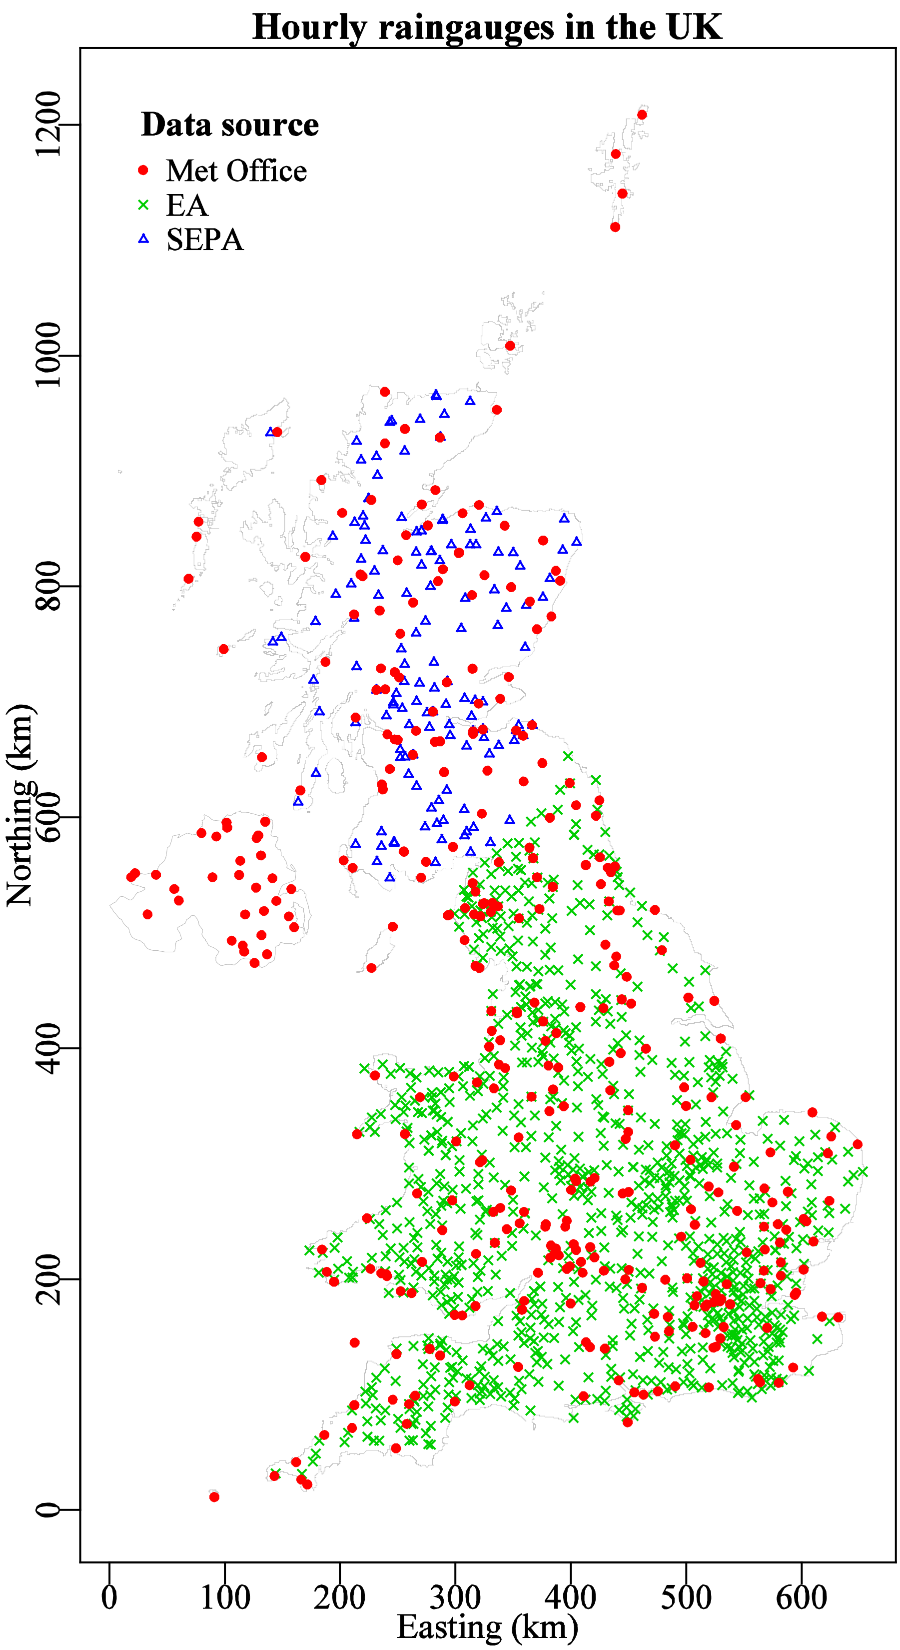

Figure 3.3 Locations of hourly raingauges for which continuous records were provided for use in this project

by using a full dynamic calibration. Rather than using a single conversion factor to convert the number of tips to rainfall amount, the gauge is calibrated for a number of different rainfall rates.

However, Nystuen (1999) found that the discrepancies mainly occur for rainfall rates exceeding 100 mm h^{-1} (for a Belfort model 382 gauge), and Fankhouser (1998) suggests that a precise static calibration is more important than applying a dynamic calibration because the latter affects only high intensities. In a review of eight tipping bucket gauges used in the UK, Hodgkinson *et al*. (2004) found that dynamic calibration did not greatly improve the rainfall catch, and suggested that the tipping bucket totals should be adjusted using monthly totals from a co-located standard Met Office storage gauge. They found at least a five per cent smaller total for the recording gauges for a one-year accumulation, and proposed the differences occurred as a consequence of the different sizes and shapes of the various gauges. This affects the wind field and therefore the rainfall catch. In contrast, Zaidman & Lamb (2005) found that, for accumulations less than 25 mm, TBRs seemed to be slightly biased to record more rain than the check gauge, whilst for accumulations of more than 100 mm, TBRs appeared to record less rain than the check gauge. For intermediate conditions, their findings show TBRs to be almost unbiased.

In the present project, hourly rainfall totals from recording gauges are used for accumulations up to 24 hours. There has been no attempt to correct any potential under-recording of hourly rainfall amounts for the following reasons. The type of gauge that has been used is not always known, and size and shape is not known even if it is clear that a gauge is, for example, a type of tipping bucket. It was considered too time consuming to investigate each recording gauge individually, particularly as there is not necessarily a daily storage gauge on the same site on which to base a comparison.

3.3 Quality control of the systematic rainfall record

Full details of the quality control procedures applied to each of the available sets of systematic hourly data are given in a previous project report and associated appendices (*Svensson et al., 2009*), whilst only a summary is presented here. The data from different sources were archived separately. The daily data held at CEH had already been through rigorous quality control procedures and therefore only particularly high values were checked in case they represented accumulations over a number of days. The hourly rainfall records from the Environment Agency and SEPA tended to be relatively short, with each one coming from a single gauge. In contrast, it was possible to obtain relatively long series from the Met Office data by joining records from different gauges which operated on the same site during different time periods. Quality controlling the data included rudimentary initial data checks, followed by more detailed checks of individual sequences of values. For example, values exceeding 999.9 mm were removed, and the data examined for gaps, negative rainfall values and invalid date-time fields. An audit trail of the changes made to the original data has been maintained.

3.4 Abstraction and quality control of maxima

Once the systematic data had been checked and consolidated, annual maxima and associated date and time information for durations of 1, 2, 4, 6, 12, 18 and 24 hours and 1, 2, 4 and 8 days were abstracted and archived, mirroring the analysis undertaken for the FEH. The annual maxima were abstracted using calendar years, again to be consistent with the FEH dataset. Once this exercise was complete, the newly available annual maxima were combined with the pre-existing FEH dataset, and, regardless of source, were joined into longer series where possible. For the daily data, only records for gauges with the same National Grid reference (to the nearest 100 m) were joined, whereas for the hourly records, a separation distance of up to 300 m was allowed.

The current study has also considered aspects of seasonal rainfall frequency, and therefore summer and winter maxima were abstracted from the continuous records available. Following the FSR analysis, summer was defined as the six-month period from May to October, and winter from November to April. After the abstraction of the seasonal maxima, similar rules to those outlined above were applied to allow some records to be joined.

3.5 Summary of the final dataset of annual and seasonal maxima

The current project has been able to use more spatially extensive and up-to-date annual and seasonal maxima across the UK than were available at the time of the FEH analysis. However, many of the new hourly rainfall records are relatively short. It was therefore decided that the FEH criterion for including a gauge record in the frequency analysis, which stipulated that at least 10 annual maxima should be available, should be relaxed slightly to include nine-year records.

Table 3.1 is reproduced from FEH Volume 2 and shows the number of station-years of data that were available to the FEH study compared to the FSR analysis. It can be seen that substantially more data were available for the FEH study, with more than three times the number of station-years used in the FSR for the 1-hour duration.

For comparison, a summary of the annual and seasonal maxima available to the current study is presented in Table 3.2. It can be seen that the number of daily raingauges with at least nine valid annual maxima has increased slightly, but the corresponding number of hourly gauges available has more than doubled since the early 1990s, when the FEH study was carried out. Table 3.2 also indicates that fewer station-years of seasonal maxima are available compared to annual maxima. The differences are only slight in the case of the daily data and reflect the fact that valid winter maxima require data from two consecutive calendar years. In the case of the hourly data, some of the FEH annual maxima came from tables held at the Met Office and it was not possible to abstract corresponding seasonal maxima.

Table 3.1 Summary statistics of data used in FEH and FSR analyses (from Faulkner, 1999)

Duration	No. of gauges	No. of station-years	Approx. no. of station-years used in FSR analysis
1 day	6,106	150,245	96,000
1 hour	375	7,389	2,300

Table 3.2 Summary of data available to the current study

Duration	No. of gauges	No. of station-years		
		Annual	Summer	Winter
1 day	6,504	171,904	171,588	164,278
1 hour	969	17,010	13,105	13,300

The spatial distributions of the final daily and hourly annual maximum datasets are shown in Figures 3.4 and 3.5. All but the most remote areas of the UK are well covered by the annual maximum dataset for the daily durations. However, Figure 3.5 shows that adequate hourly maximum data are not available in some regions, notably south-west England, parts of the south coast and Kent. This is generally because, although recording raingauges are operated in these areas, often for use in flood forecasting, their record lengths were not sufficient to provide the nine annual maximum values required for inclusion in the analysis. Therefore, it should be possible to update the dataset of annual and seasonal maxima in the future, when the raingauges have been operating continuously for nine years or more. Also, it is known that there are older non-digitised records of sub-daily rainfalls available for south-west England. Digitising these would provide a valuable extension to the database.

The evolution of the networks of daily and sub-daily raingauges in the UK is illustrated by the maps in Figures 3.6 and 3.7. Each map shows the gauges which meet the record length criterion (that is at least nine annual maximum values) operating in selected calendar years. For the daily data, Figure 3.6 shows that much of Scotland and some parts of Wales are particularly poorly represented in the digitised records for the pre-1961 period (although the national network has comprised over 3,000 daily gauges since 1900 (*Svensson*

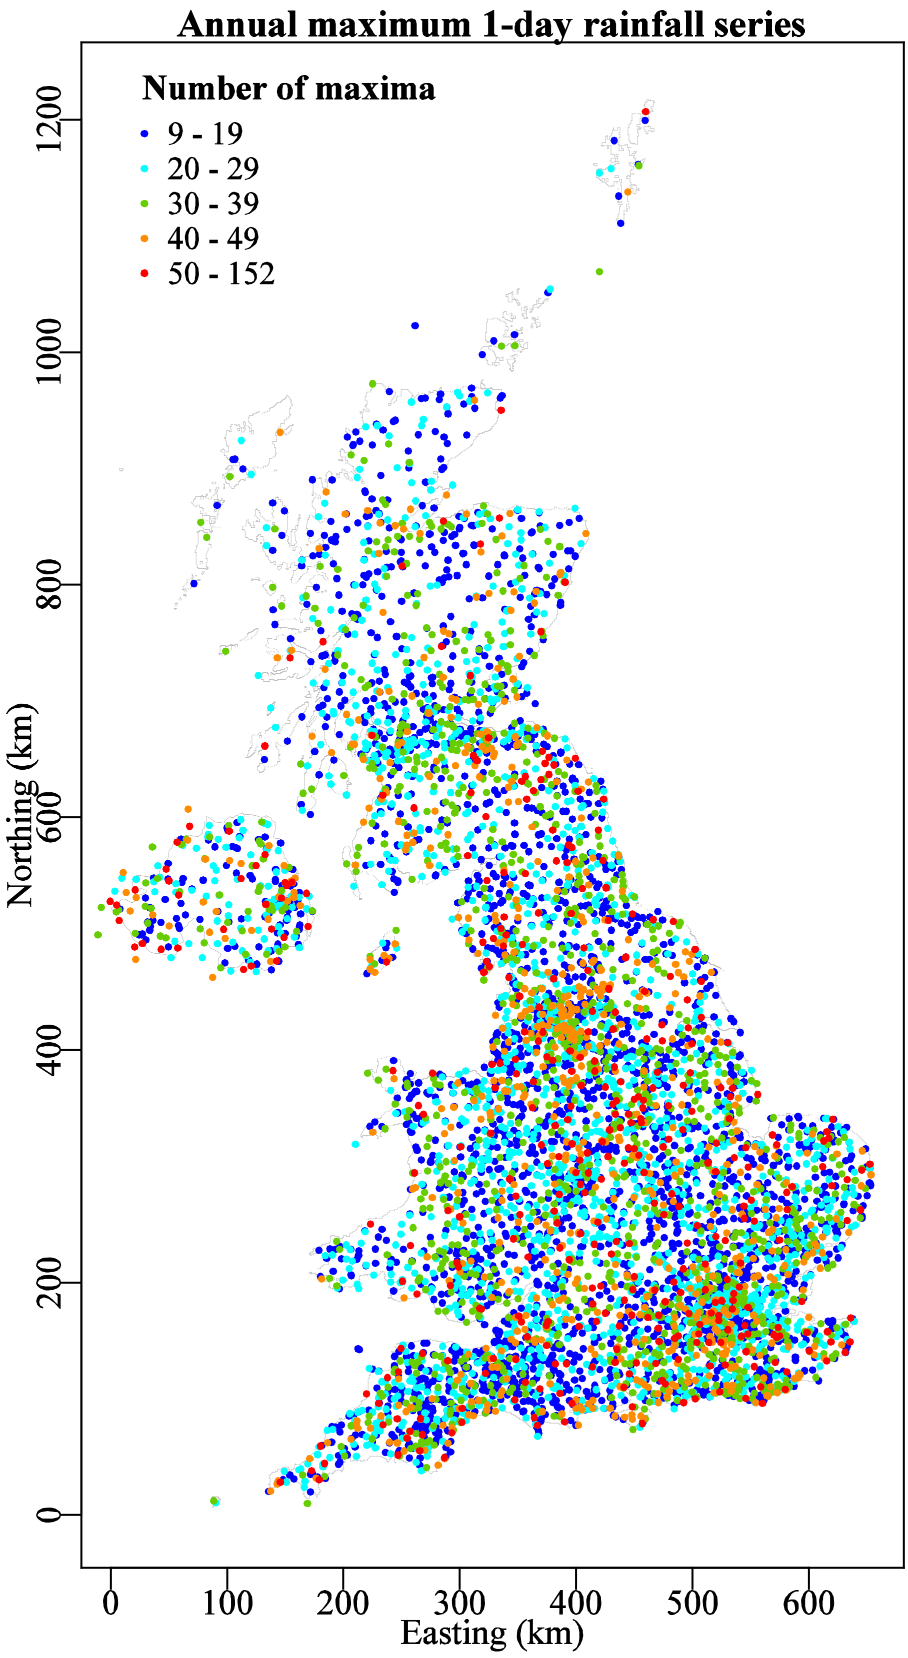

Figure 3.4 Locations of daily raingauges with at least nine annual maxima

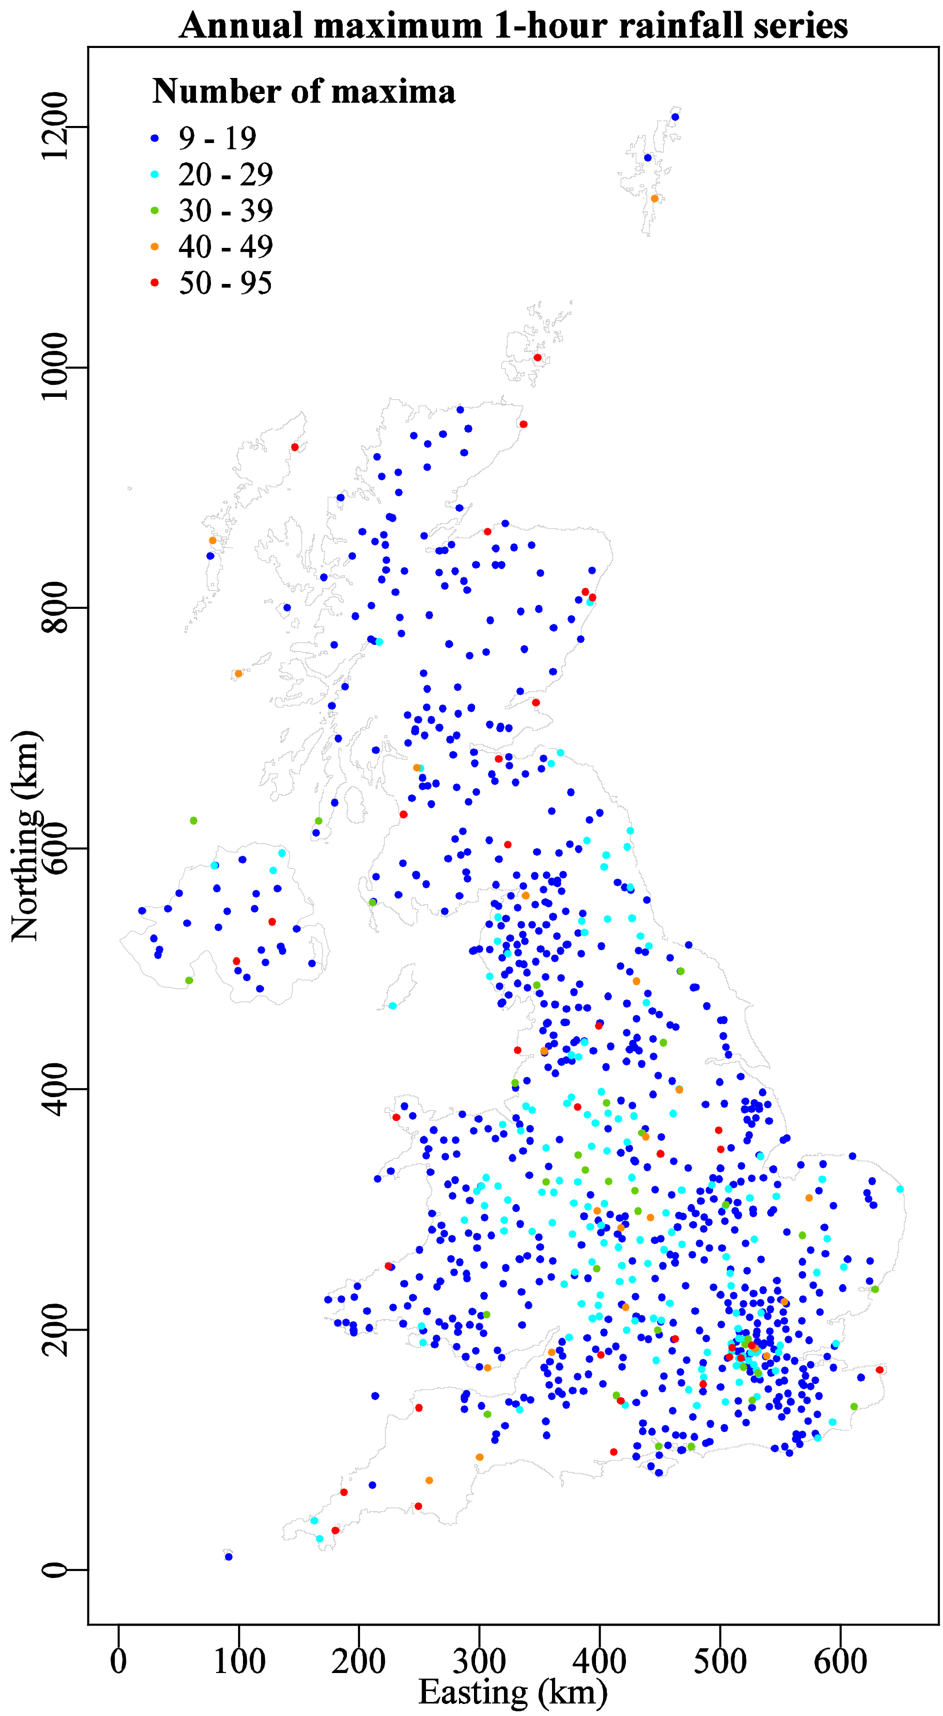

Figure 3.5 Locations of hourly raingauges with at least nine 1-hour annual maxima

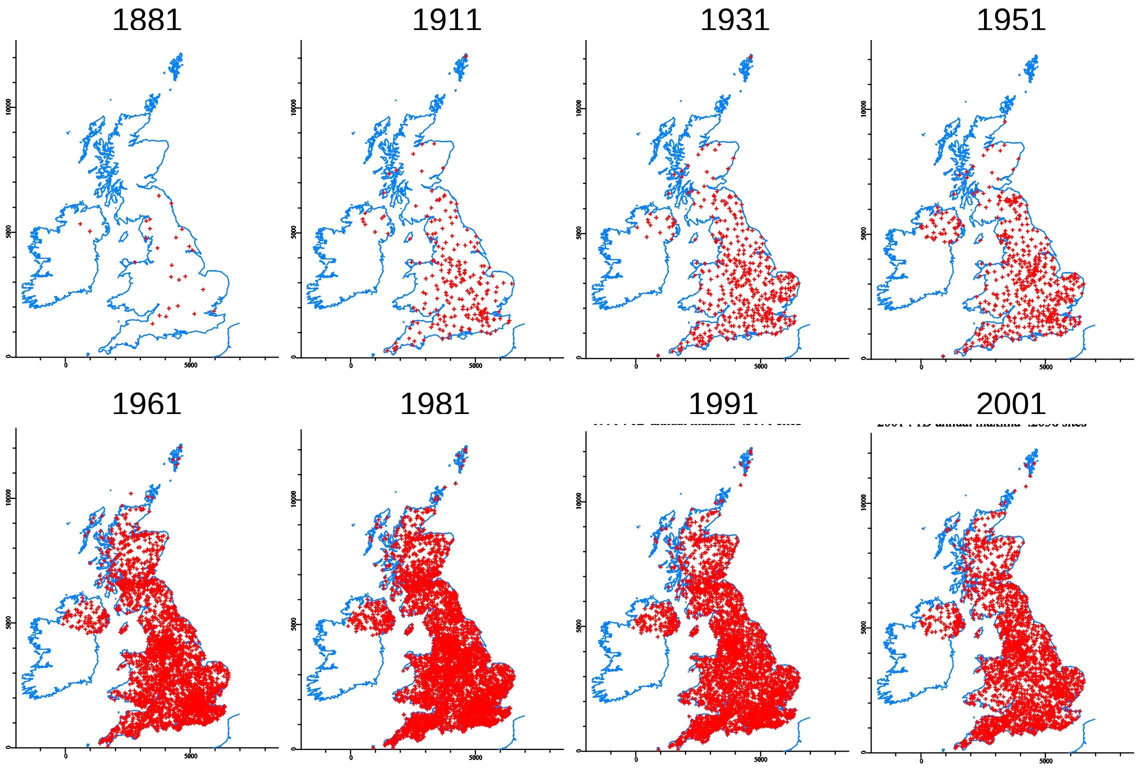

Figure 3.6 Locations of gauges with daily annual maxima for selected calendar years

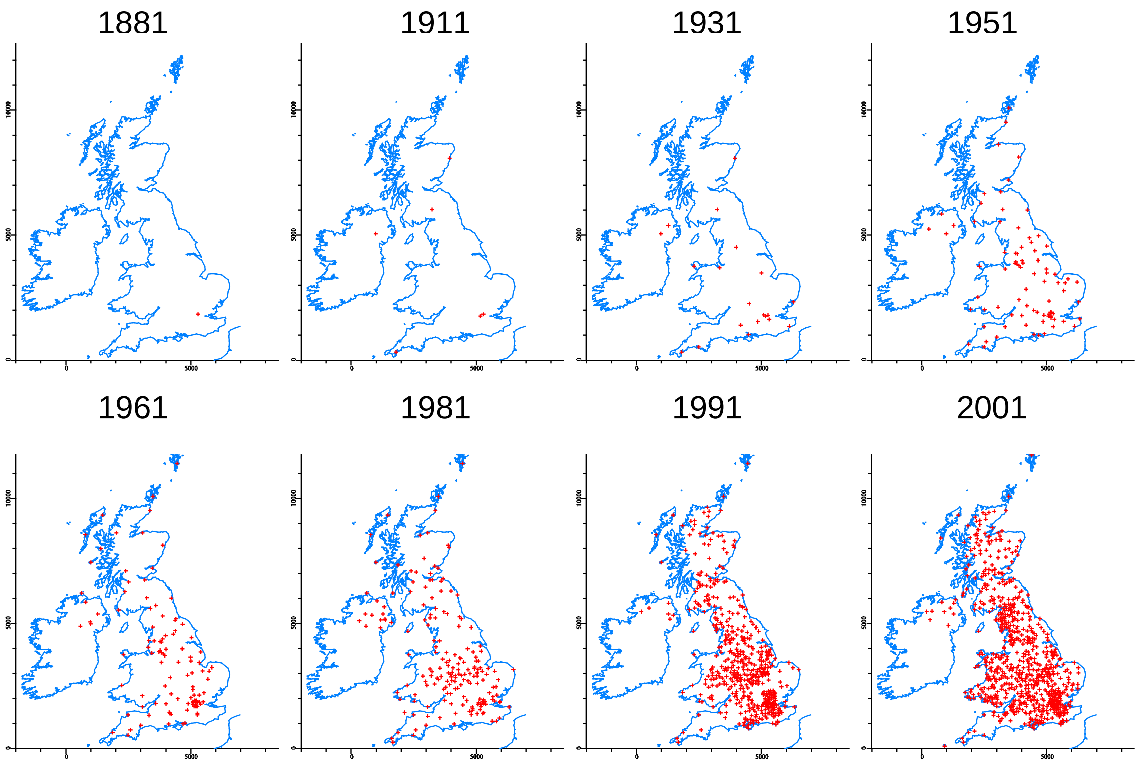

Figure 3.7 Locations of gauges with hourly annual maxima for selected calendar years

et al., 2009; Eden, 2009), only a small proportion of the pre-1961 observations have been digitised). In the case of the hourly dataset (Figure 3.7), while the coverage of England is reasonably good by 1951, it is not until 1981 that the network covers most of Northern Ireland and Wales, and coverage of Scotland does not improve until 1991. Some rationalisation of the raingauge network took place after 1999, with instrument failure in the year 2000 being a particular problem.

The two longest hourly records both run for 95 years: the record for Kew Observatory (287049) starts in 1886 and runs until 1980, while that for Eskdalemuir (610122) starts in 1910 and ends in 2006. There are 92 gauges with at least 30 annual maxima and 33 gauges with at least 50 annual maxima for the 1-hour duration.

Details of the sub-daily raingauges used in the statistical analysis are given in Appendix A.

3.6 Adjusting for the effects of discretisation

The basic data from which annual and seasonal maxima were abstracted represent clock hour accumulations in the case of durations from 1 to 24 hours, and rain day depths (measured at 0900 GMT) in the case of durations from 1 to 8 days. These rainfall durations are often termed '*fixed'* durations. However, for most flood design applications there is a requirement for rainfalls to be aggregated over a duration that can start at any time. This is referred to in the FEH as a '*sliding'* duration. For this reason, both the FSR and the FEH made allowances for the effects of discretisation by applying adjustment factors (*discretisation conversion factors*) to convert frequency estimates from fixed to sliding durations.

Here, 'discretisation' relates to the fact that the underlying rainfall records are held at one of two time resolutions; either 1-hour or 1-day. As in the FEH, the convention is adopted that rainfall totals of different durations derived from the 1-hour resolution dataset are labelled by '1 hour', '2 hours', etc., while rainfall totals derived from the 1-day resolution dataset are labelled by '1 day', '2 days' *etc.* The concept of 'sliding durations' arises in connection with finding the largest rainfall total (in a year) over a duration which is longer than the resolution at which the underlying data are held. Deriving such values from the data involves finding the totals for all time intervals of the given duration, where such an interval must start at the beginning of a time period at the underlying resolution. This restriction has led to annual maxima determined in this way being described in this report as either 'semi-sliding durations' or 'fixed durations' (where the latter refers to the case where the required duration is the same as the underlying resolution). This is to distinguish the values obtainable in practice from the idealised case where a rainfall total considered as potentially being the annual maximum might start at any time, not restricted to a grid of starting-times. These idealised values are referred to as relating to 'fully-sliding durations'.

For almost all practical uses of estimates of extreme rainfalls, it is the fully-sliding duration rainfall that is of interest, as there is no role for a grid of possible starting times. There are exceptions in the case of certain insurance applications, where a payment would be made if the rainfall with a fixed time period (that is between two given times of day) exceeded a stated amount. However, the ‘fixed duration’ case as considered here would only correspond to a few such applications. For example, the maximum rainfall in a calendar week (seven days, Sunday to Saturday) would align with the ‘fixed duration’ case, while a working week (five days, Monday to Friday) would not.

In the current study, the annual and seasonal maxima for fixed and semi-sliding durations have been analysed extensively and a set of discretisation conversion factors has then been applied in the final stages when the DDF model was fitted to the outputs of the revised FORGEX methodology (see Table 3.3). These factors are used to convert fixed and semi-sliding duration annual maximum rainfalls to equivalent values for fully-sliding durations. It should be noted that all the results of the revised FORGEX procedure presented in Section 7 of this report refer to fixed and semi-sliding durations. The conversion to fully-sliding durations is made in the analysis in Section 8. This leads to the frequency estimates derived from the revised DDF model (discussed in Section 9) also being aligned to the fully-sliding durations. Table 3.3 records the conversion factors used for this study, labelled as ‘revised’, in comparison with the values used in the FSR and the FEH. Further information is given in Section 8 and Appendix J.

Table 3.3 Discretisation conversion factors used in the FSR, FEH and the current study

Duration	Resolution	Adjustment factors FSR	 FEH	 Revised
1 hour	1 hour	1.15	1.16	1.16
2 hour	1 hour	1.06	1.08	1.08
6 hour	1 hour	1.015	1.019*	1.019
1 day	1 day	1.11	1.16	1.146
2 days	1 day	1.06	1.11	1.072
4 days	1 day	1.03	1.05	1.043
8 days	1 day	1.015	1.01	1.025

* By interpolation between 2-hour and 8-hour values

3.7 Extreme rainfall event database

As part of an earlier Defra-funded project (FD2201 *Extreme rainfall and flood event recognition*), a set of 50 extreme storms experienced in the UK since 1900 was selected and relevant meteorological information was compiled (*Collier et al., 2002*). The aim of the study was to investigate the nature of very extreme rainfall events and the meteorological situations that caused them. The study found that the estimates of probable maximum precipitation (PMP) published in the FSR had been exceeded in several cases. This suggested that further research on the nature of extreme rainfall events was needed. One of the key requirements of the current project was to extend and supplement this dataset of extreme storms to establish an authoritative database of all major rainfall events relevant to reservoir safety, and to include information about storm profile, spatial extent, storm type, annual probability and data quality. While this database has not been used in the main statistical analysis, it provides an opportunity to compare statistical frequency estimates with the characteristics of observed extreme storm events in the UK.

Information has been compiled for a total of 63 events with durations of at least 1 hour that were recorded in the UK between 1880 and 2006. A list of the events included is provided in Appendix B, together with summary information and estimates of frequency derived from the FEH DDF model. Further details of these events, including the selection criteria, are given by Dempsey & Dent (2009).

There are four events in Appendix B that do not exceed the threshold. Of these, one event occurred at Inverness in 1915 (No. 13) and is considered to be a key historic event for Scotland. Another is the North Yorkshire Castleton event in 1930 (No. 20a), which is represented also with a longer duration exceeding the threshold. The two remaining events are short-duration convective events occurring in 1890 and 1895 (events No. 2 and 4). These were included because the coverage of raingauges prior to 1900 was sparse.

3.8 Other non-systematic data

Other readily available information about extreme rainfall events was collated at the outset of this project. The annual publication *British Rainfall* contains tables of the largest 1-day rainfalls recorded during each calendar year. The tables in the yearbooks from 1900 to 1960 were digitised with a view to using the data to increase the historical non-systematic data on extreme events. A total of 791 observations were archived in a spreadsheet and the intention was to use them as additional network maxima within the revised FORGEX procedure. However, it was subsequently decided that the data would add relatively little additional information to the systematic daily records, and this, together with time constraints, meant that the idea was not implemented.

Including the additional network maximum points within FORGEX would have required an indication of the extent of the daily raingauge networks across the UK in each year. For this reason, estimates were made of daily gauge density every five years from 1900 to 2000. Again, these data were archived in a spreadsheet but were not used in the final statistical analysis. Further details of the data have been given by Svensson *et al.* (2009).

4. Initial analyses of annual and seasonal maxima

This section gives an overview of the initial statistical analysis of the final dataset of annual and seasonal maxima. The main aim of this work was to assess possible families of distributions for describing the maxima for the full range of durations (1 hour to 8 days). This is particularly relevant in developing the revised model of rainfall depth-duration-frequency (DDF) discussed in Section 8. The results of an analysis of seasonal rainfall frequency are also presented, together with the results of a limited study of trend within the daily annual maximum dataset.

4.1 Choice of model for the distribution of annual maxima

Initial analyses of the dataset considered alternative models for the distribution of the annual maxima, with the aim of assessing the most appropriate way of extrapolating frequency curves to very high return periods. Another reason for investigating distributional form was to consider the validity of the Gumbel distribution for modelling segments of the rainfall growth curve within the FEH FORGEX procedure.

A substantial part of the statistical analyses carried out within this project has used quantities known as L-moments and L-moment ratios. Appendix I gives some background on these. These quantities are analogues of ordinary moments. They are numerical measures of distributional shape which are derived from linear combinations of ordered sample values.

Figure 4.1 shows an example of an L-moment ratio plot, which is one of the graphical tools that are commonly used when considering the choice of the family of distributions for extreme value problems. See for example Hosking & Wallis (1997, *Section A.13*). In this figure, L-kurtosis is plotted against L-skewness for each gauge within the dataset for the 1-day duration. A standard feature of such plots is that they include both curves and individual points which represent several different possible families of probability distributions. The curves represent families which have three parameters, while families of distributions having only location and scale parameters each plot as single points. In addition, there is a curve labelled ‘OLB’ (overall lower bound) which, for each value of L-skewness, shows the lowest possible value of L-kurtosis for any possible theoretical distribution, allowing the corresponding family to be chosen as the most appropriate distribution. However, note that sample estimates of L-skewness and L-kurtosis do not necessarily obey this bound. A point is plotted for each record being considered, showing the sample values of L-skewness and L-kurtosis. The L-moment ratio plot is typically suggested as a basis for choosing a family of distributions for either a single record or for records from many different sites in a network of flow or rainfall stations. In typical analyses, the cloud of points representing the statistics for the individual

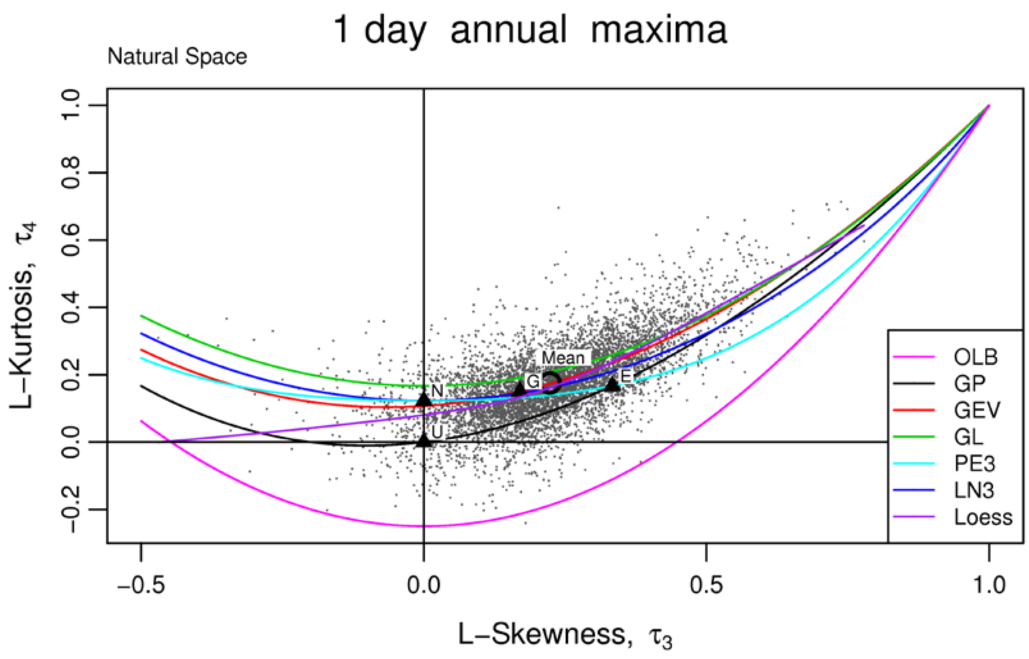

Figure 4.1 L-moment ratio plot for 1-day annual maxima showing sample points and theoretical curves and points for selected distributions (see text). 'Mean' indicates the average of the sample points, while the 'Loess' line is a smooth curve through the sample.

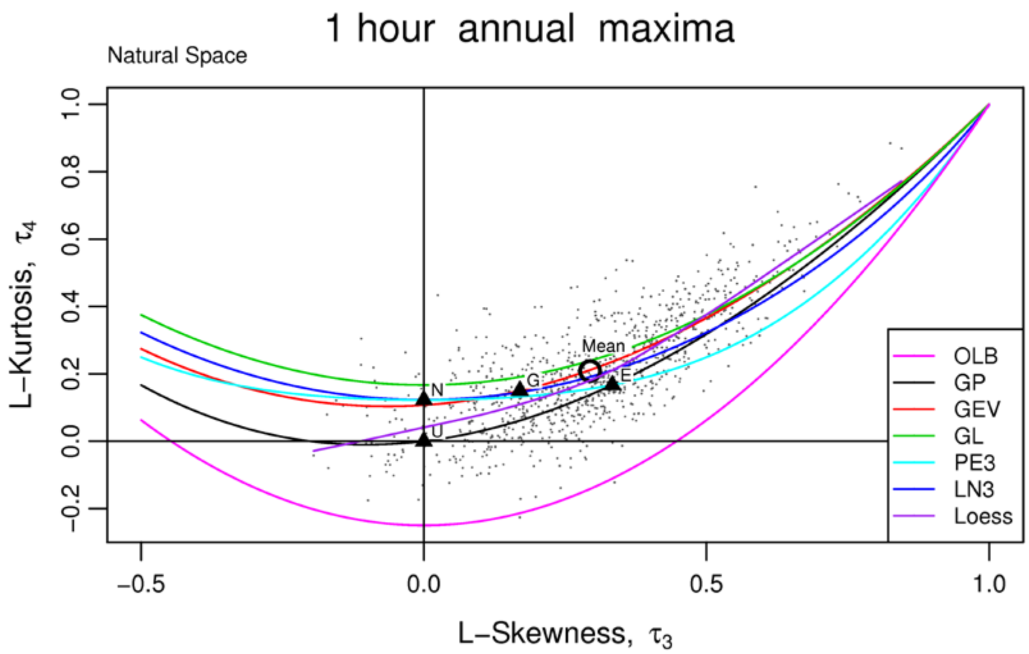

Figure 4.2 L-moment ratio plot for 1-hour annual maxima (see Figure 4.1 for description)

sites will centre about one of the curves, allowing the corresponding family to be chosen as the most appropriate distribution. The existence of an overall lower bound for theoretical distributions means that the cloud of points is likely to have a curved configuration related to this bound, although the numbers of sites available may not always reveal this. In the present application, where distributions are typically skewed to the right, the cloud of points generally shows an upward trend matching the lower bound curve.

In this standard type of L-moment ratio plot, a number of extreme value distributions typically used in hydrological frequency analysis are considered as possible candidates and are included as lines or points. These distributions include the Uniform (U), Normal (N), Gumbel (G) and Exponential (E) distributions, all of which have two parameters and which are shown as special points on the plot. Here, the Exponential distribution has two parameters when the lower bound of possible values is considered to be a parameter. The three-parameter distributions considered are the Generalised Extreme Value (GEV), Generalised Pareto (GP), Generalised Logistic (GL), Pearson Type 3 (PE3) and Log-Normal (LN3). For the Log-Normal distribution, the lower bound of possible values is treated as a parameter. This form of distribution is also sometimes known as the Generalised Normal.

In the case of Figure 4.1, which shows the L-moment ratio plot for 1-day annual maxima, the cloud of points lies closest to the curve for the GEV distribution. Points with the largest values of L-skewness have a tendency to plot higher than any of the curves for specific distributions. While a possible explanation might be bias in the sample estimates of the L-moment ratios, this is excluded by the results given by Hosking & Wallis (1997, *Figure 2.7*). Figure 4.2 shows a similar L-moment ratio plot for the 1-hour duration. In both Figures 4.1 and 4.2, the point for the mean of the L-moment ratios across all raingauges plots to the right of the point representing the Gumbel distribution. This indicates that rainfall frequency curves would have a tendency to curve upwards when plotted on a Gumbel reduced variate scale.

It was concluded that, of the three-parameter distributions, the GEV distribution is the most appropriate for the annual maximum dataset, although there are indications of a lack of fit in cases where the skewness is high.

4.2 Possible transformations of rainfall

Simple transformations were applied to the annual maximum dataset to check whether this improved the conformance of the data to any of the families of statistical distributions. L-moment ratio plots were derived after applying the transformations and examples for durations of 1 hour, 6 hours, 24 hours and 1 day are shown in Figures 4.3 to 4.6. These figures suggest that, compared with the ‘natural’ (untransformed) data, the transformed data points do not lie any closer to the GEV curve, or indeed to the curves for any of the other distributions, indicating that the data can be used without transformation.

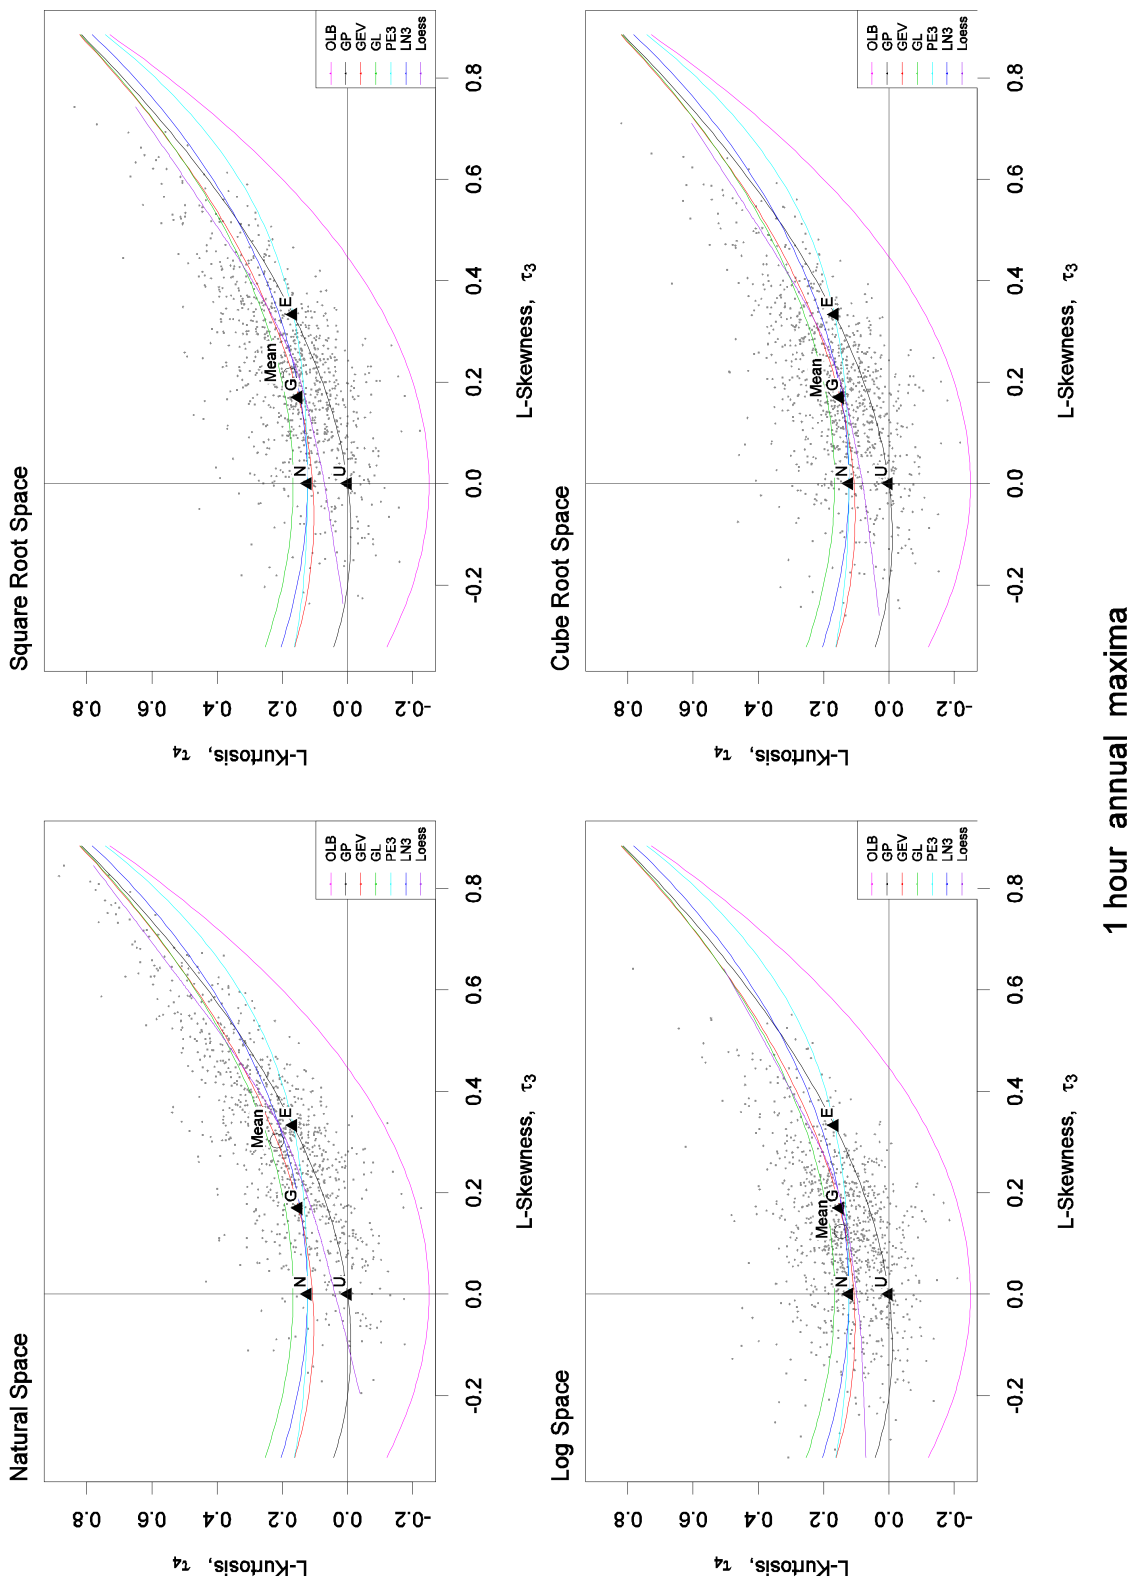

Figure 4.3 L-moment ratio diagrams for 1-hour maxima in natural space and following simple transformations

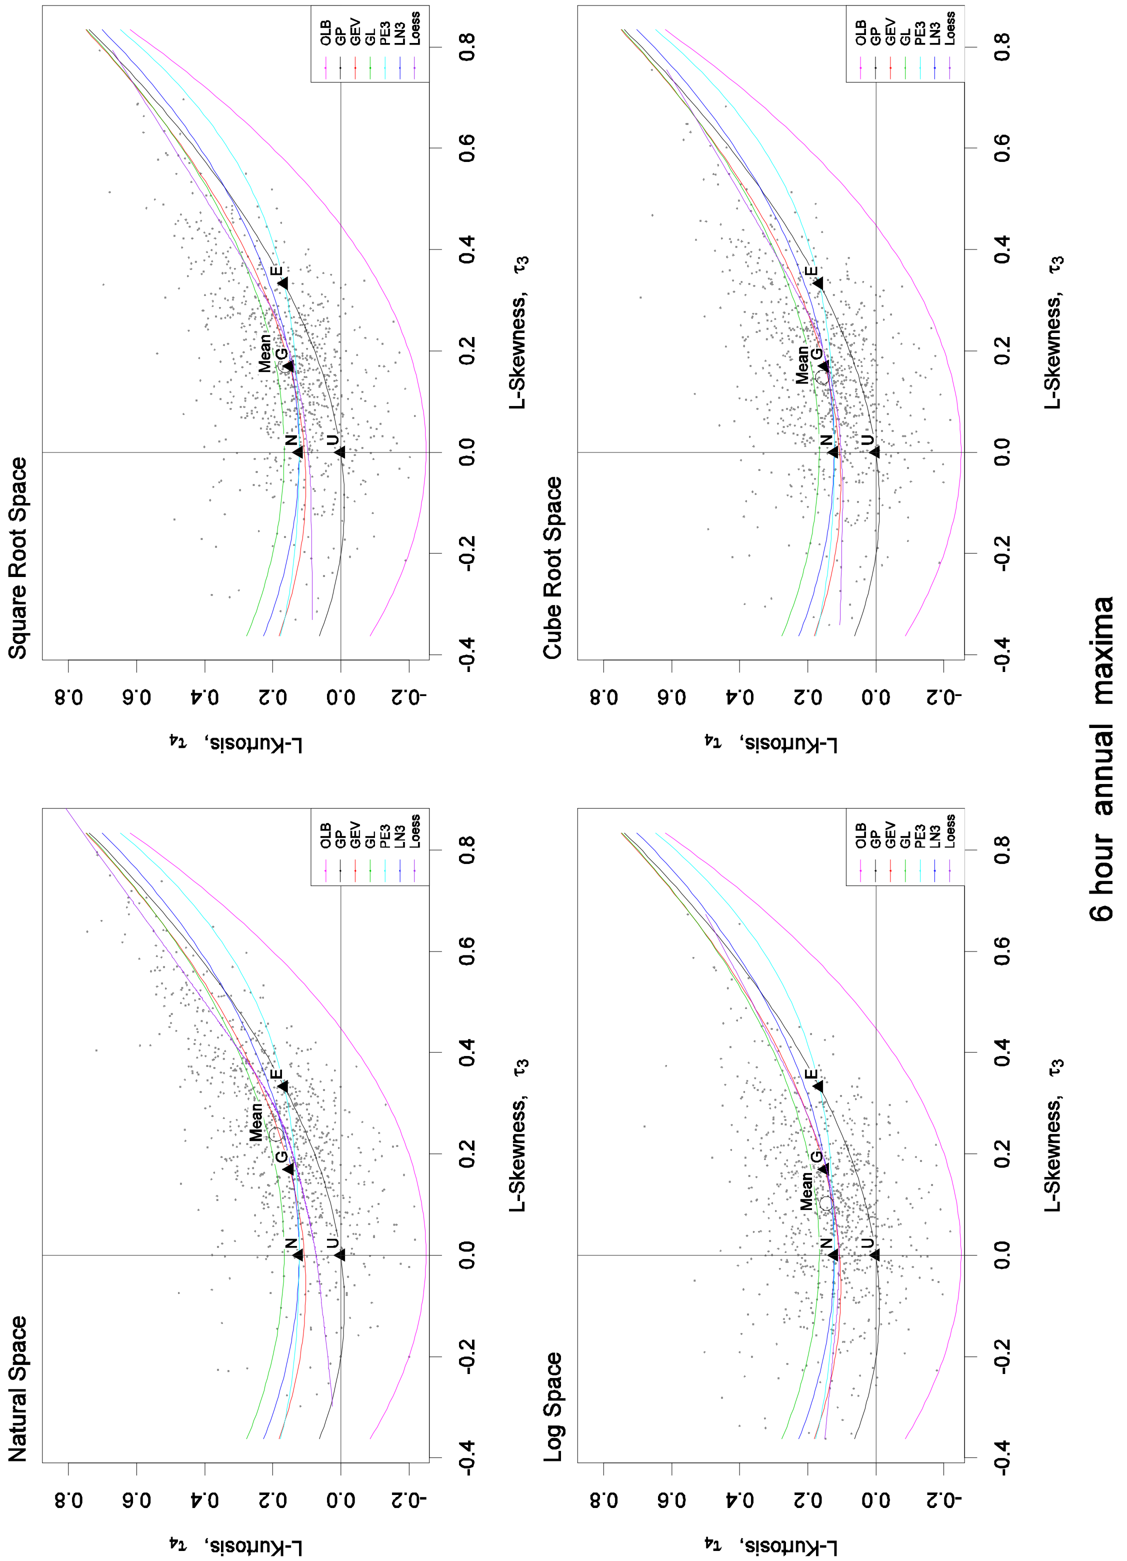

Figure 4.4 L-moment ratio diagrams for 6-hour maxima in natural space and following simple transformations

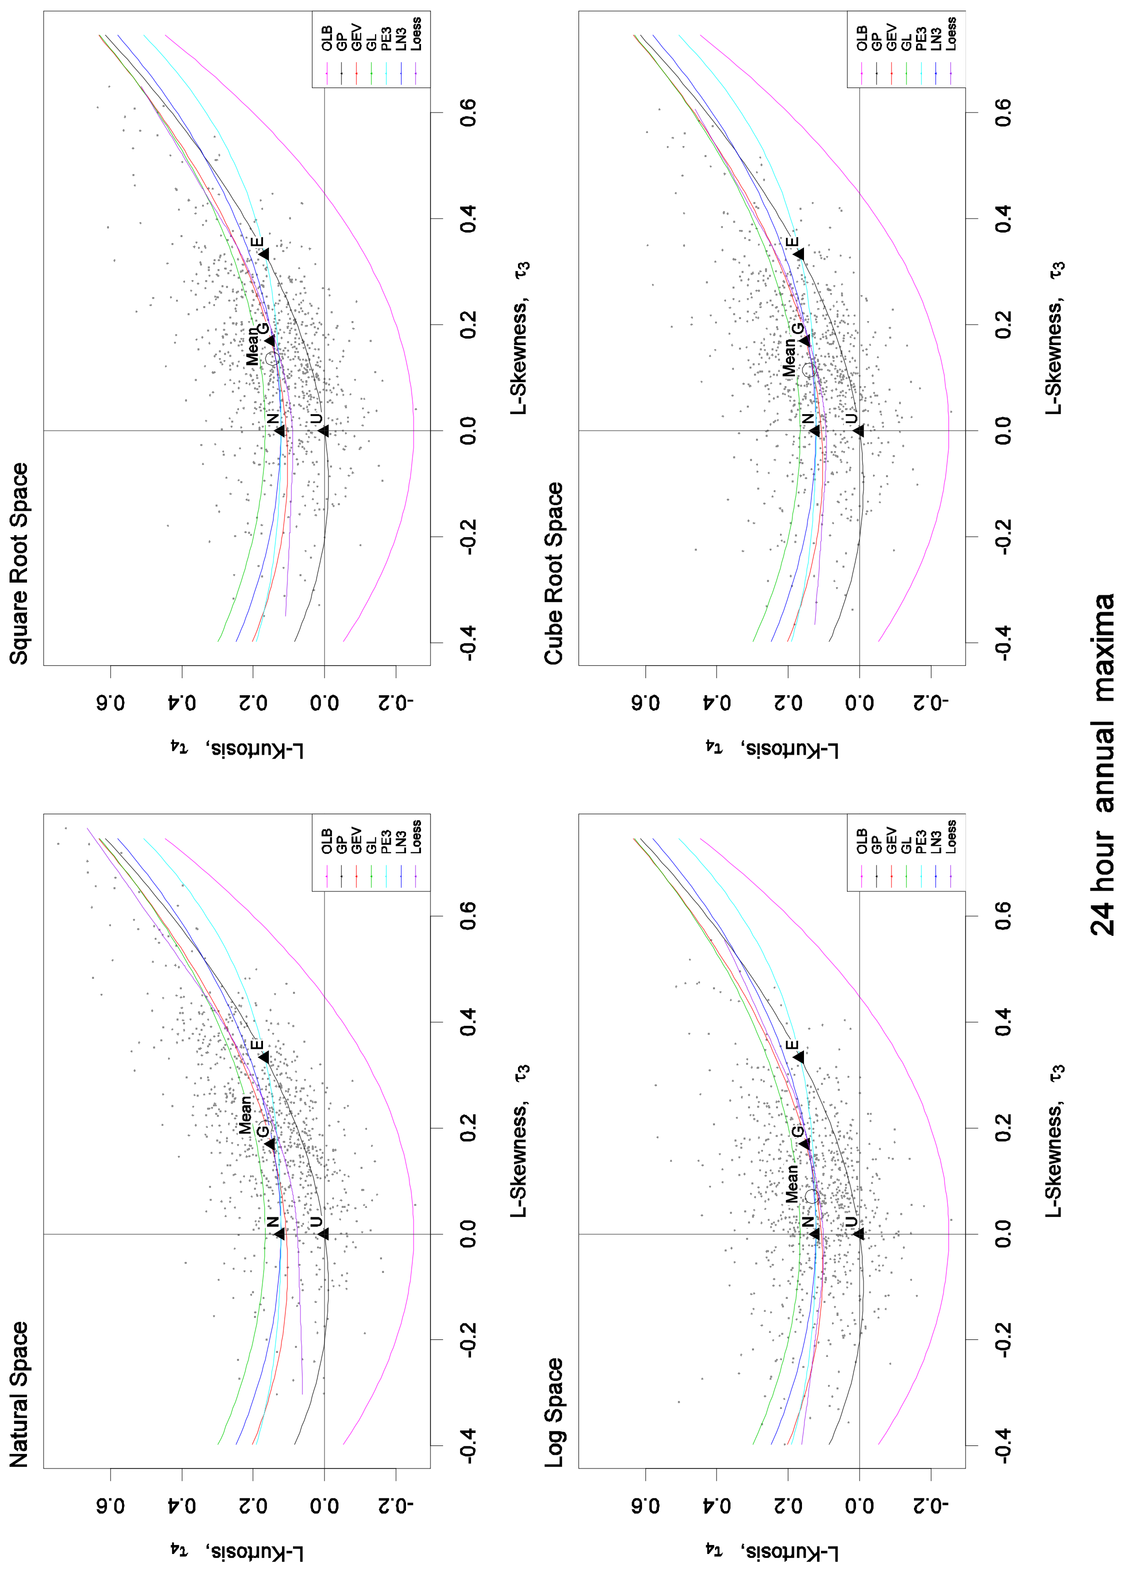

Figure 4.5 L-moment ratio diagrams for 24-hour maxima in natural space and following simple transformations

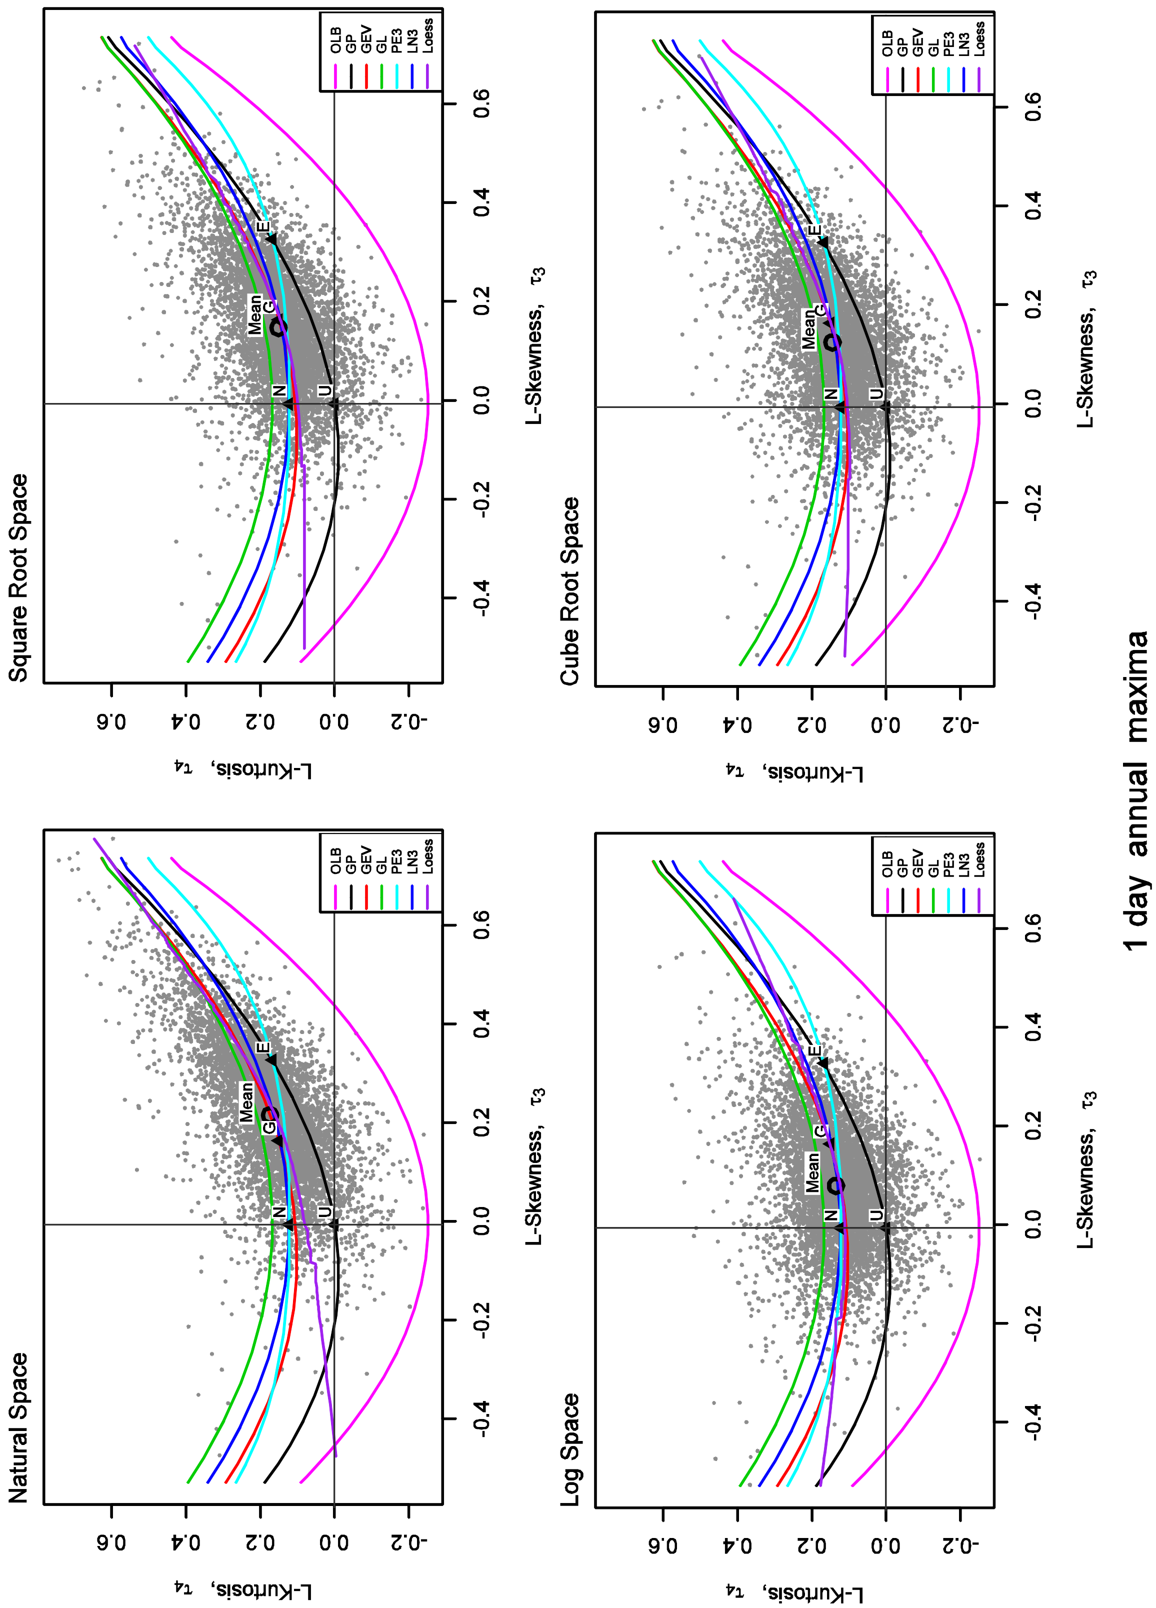

Figure 4.6 L-moment ratio diagrams for 1-day maxima in natural space and following simple transformations

One point that arises from the comparison of Figures 4.3 to 4.6 across the different durations again relates to the position of the ‘mean’ point relative to the ‘Gumbel’ point. The change in this relative position indicates that the rainfall frequency curve plotted on a Gumbel scale for the 1-hour duration will be curved upwards (as noted earlier), but less curved for higher durations. Similar results were found for durations longer than 1 day, and for the 8-day duration the analysis indicated that a fitted GEV curve would have slight downward curvature, and therefore a finite upper bound. This difference in curvature between the frequency curves of different durations was identified as a potential problem for the final depth-duration-frequency modelling stage of the project, since it might cause curves to cross and therefore introduce inconsistency into the final estimates. This is discussed in more detail in Section 8.

The conclusion at this stage is that applying a transformation to the data does not solve the problem of finding families of distributions which will ensure consistency of rainfall amounts between durations when the distributions are used as a DDF (depth-duration-frequency) model. This problem was therefore carried forward to the depth-duration-frequency modelling stage of the project, the results of which are reported in Section 8. The estimates of rainfall derived as the final result of this project are not based on assuming a GEV distribution. However, the studies of spatial dependence, reported in Section 6, do make this assumption as part of estimating a model for spatial dependence. It is not thought that departures from a GEV distribution in those analyses will have an important carry-over effect into the final results.

4.3 Seasonal rainfall frequency

Systematically abstracted seasonal maxima were not readily available to the FEH team and therefore the most recent comprehensive study of seasonal rainfall frequency in the UK is that presented in the Flood Studies Report (FSR) in 1975. The continuous hourly rainfall data made available to the project offered the opportunity to abstract seasonal maxima and to carry out a short analysis of seasonal rainfall frequency across the full range of durations.

During the development of the Revitalised FSR/FEH rainfall-runoff (ReFH) method of design hydrograph analysis (Kjeldsen *et al*., 2005), a limited study of seasonal rainfall maxima was conducted and seasonality was introduced into the design rainfall inputs via the application of a set of simple correction factors. The appropriate factor is applied to the all-year design rainfall depth of the required duration to yield the winter or summer design estimate. The data used within the ReFH study were not sufficient to investigate any possible variation in the correction factors with return period. Details of the seasonal rainfall analysis are given by Kjeldsen *et al.* (2006).

4.3.1 Definition of seasons

The ReFH analysis based its definition of seasons on the FSR and this has been adopted in the current study. Therefore, two seasons are distinguished;

winter (November to April) and summer (May to October). It should be noted that the annual maxima under analysis have been abstracted using calendar years, following the approach adopted by both the FSR and FEH. This does lead to some inconsistencies in the data because the winter maximum associated with a particular year may be experienced in the previous calendar year (between October and December). However, a full re-analysis of seasonal and annual maxima for the 12-month period from October to September was not considered to be feasible within the current project.

4.3.2 Variation with site descriptors and return period

An analysis was carried out using the whole dataset of seasonal maxima to look for patterns in the variability of seasonal ratios with return period. A series of GEV distributions were fitted to the full set of seasonal maxima using the method of L-moments (Hosking & Wallis, 1997). A separate distribution was fitted to the record for each gauge with at least nine seasonal maxima for each of the key durations. Similarly, GEV distributions were also fitted to the annual maxima at each of these gauges. Estimates of winter, summer and annual rainfall of each duration were derived for return periods of 2, 5, 10, 20 and 100 years.

Values of winter/annual and summer/annual rainfall were calculated for each duration and return period, an exploratory analysis was carried out of the variation of these ratios with the available site descriptors (National Grid co-ordinates, altitude, and standard average annual rainfall over the period 1961 to 1990 (SAAR)) was examined. After considering various combinations of these descriptors, it was concluded that simple predictive models for the ratios could be obtained using SAAR as a single explanatory variable. As an initial step, a transformed variable *rsaar1* was constructed where

$$rsaar1 = 1.5 - 1000/\text{SAAR}. \qquad (4.1)$$

However, the final model developed uses a modified quadratic function of this transformed variable as the final predictor of the ratios. Further details of the analyses are given in Appendix C.

Figures 4.7 to 4.8 show some of these fitted relationships for the ratios as functions of SAAR, duration and return period. It can be seen that the ratios vary most markedly with SAAR, although there are also more minor variations with duration and return period. Some moderately extensive tables of these results are included in Appendix C to provide a basis for practical applications of these ratios.

In conclusion, the results of this limited study confirm those of the ReFH analysis which found that seasonal correction factors vary with SAAR across the UK. The analysis has not found evidence of a strong return period effect, although an analysis of longer records, if available, would be valuable in the future.

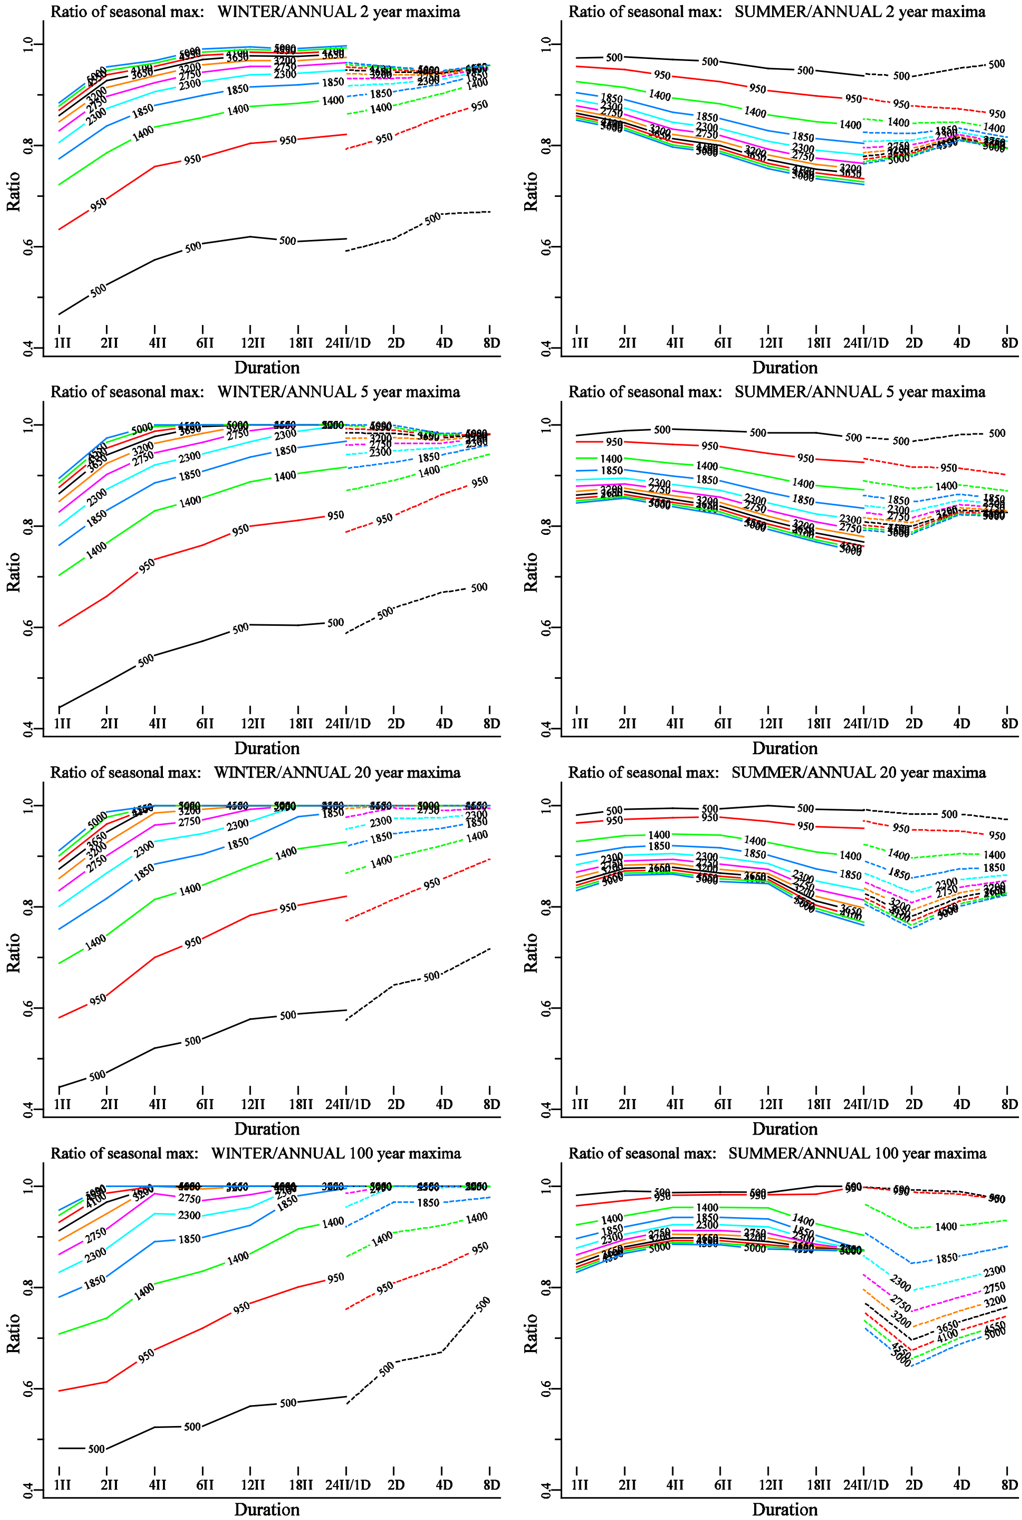

Figure 4.7 Ratios of seasonal to annual rainfall for 5 return periods as a function of duration. Each line is labelled with a value of SAAR (mm).

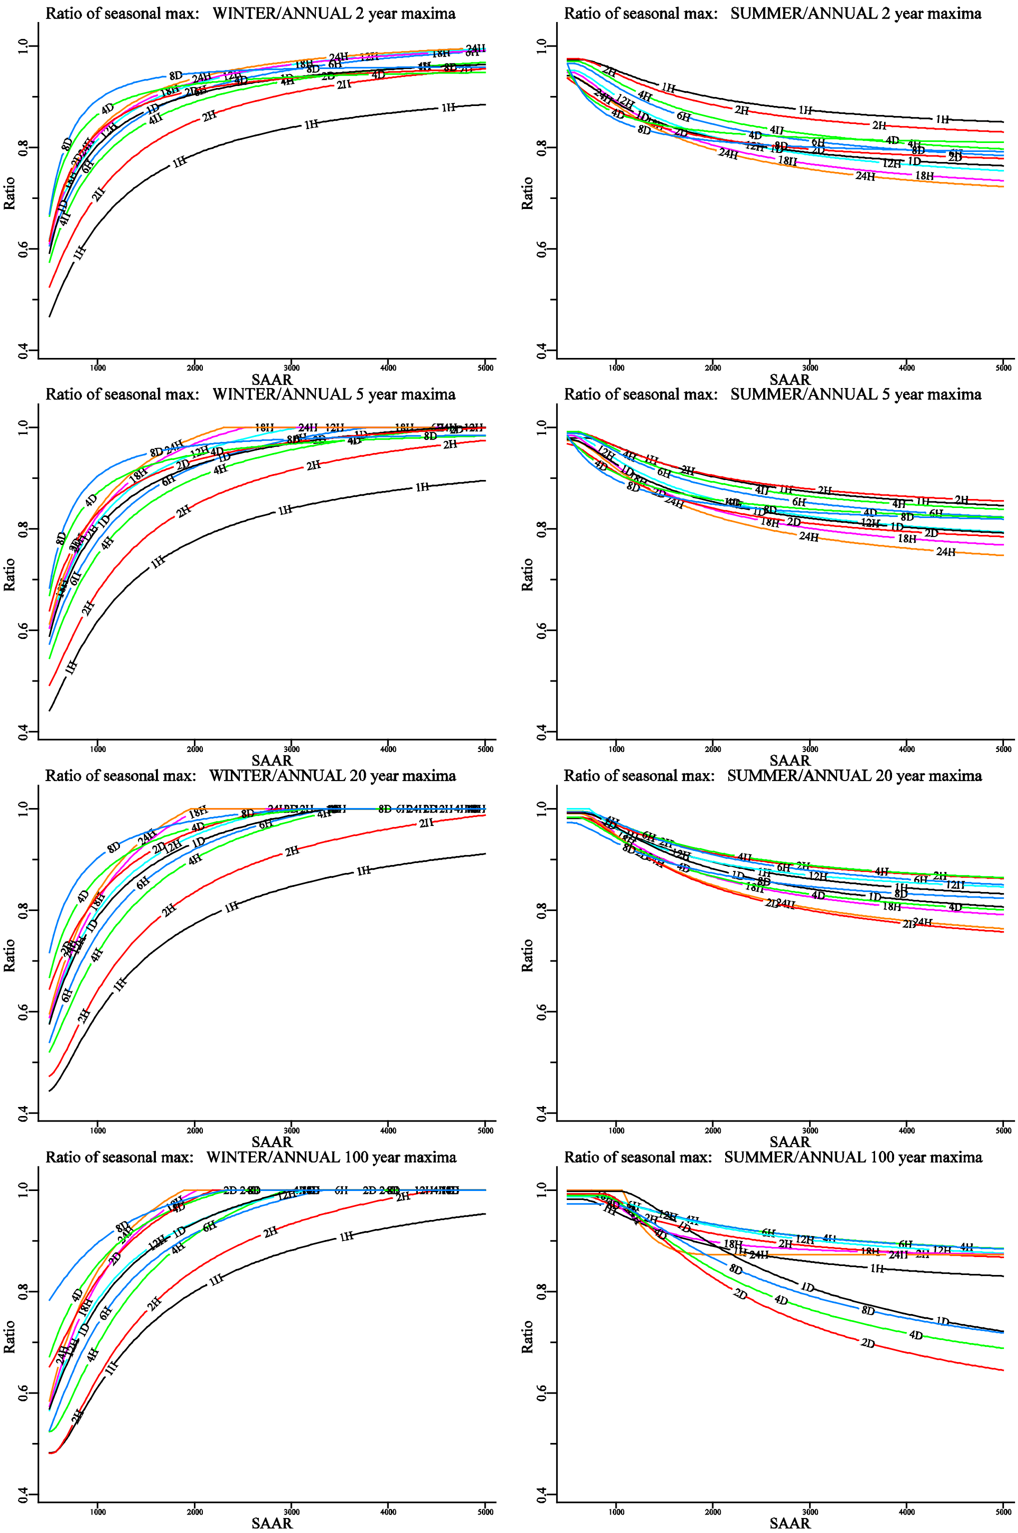

Figure 4.8 Ratios of seasonal to annual rainfall for five return periods as a function of SAAR (mm). Each line is labelled using the duration.

Developing a fully seasonal rainfall DDF model was beyond the scope of the current study. However, a set of seasonal correction factors has been derived to use alongside the all-year model. This will allow seasonal rainfall estimates to be used within rainfall-runoff models, if appropriate. Further details can be found in Appendix C.

4.4 Analysis of possible trend in daily annual maxima

Simple analyses of possible trend with time were carried out relatively early in the project using the daily annual maximum dataset. Figure 4.9 shows the temporal behaviour of the annual maxima for three gauges selected on the basis of their long records. These plots show the data values, together with a trend line fitted by regression and a locally fitted trend line using ‘loess’ with a high degree of smoothing.

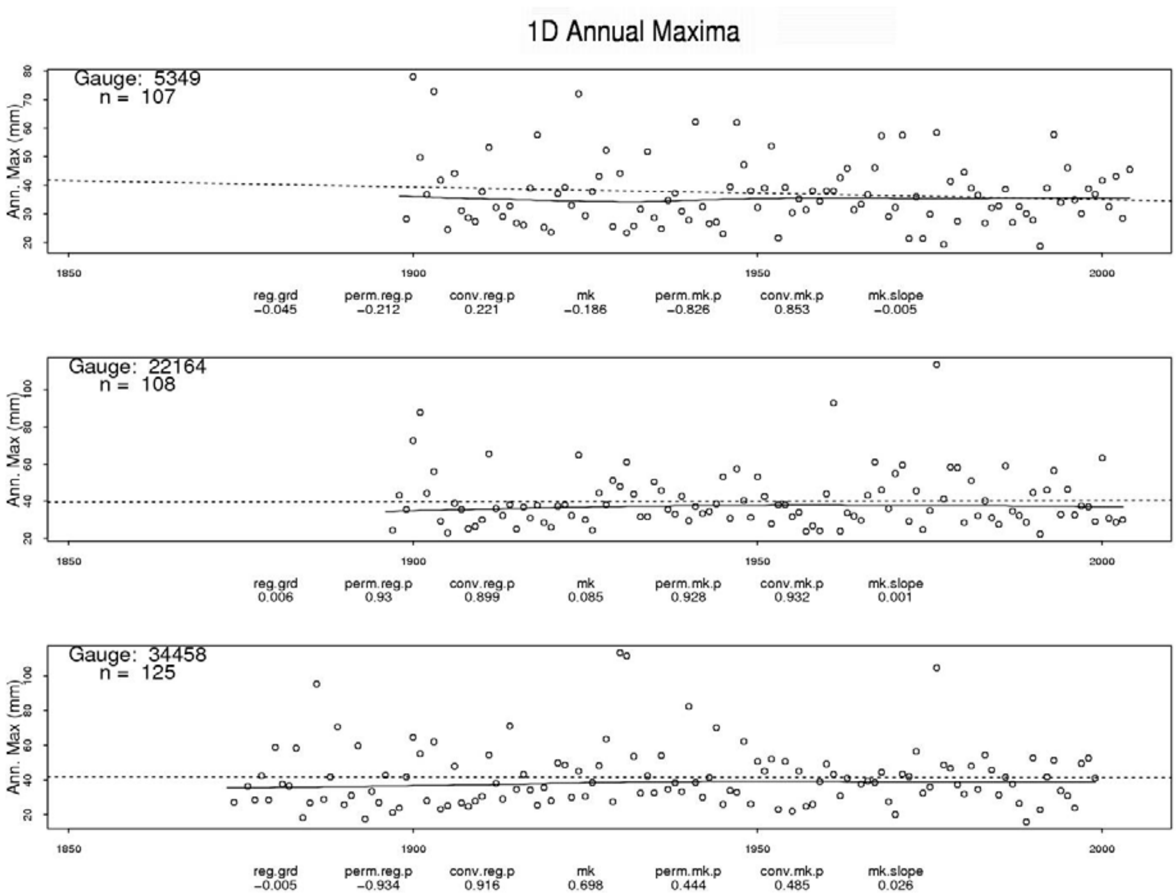

Figure 4.9 Temporal variation of annual maxima for some example gauges selected because they have long records. The dashed lines are fitted linear trend lines, while the solid lines are locally weighted trend lines.

Figure 4.10 presents some results of a simple trend analysis for all the available daily-read raingauges which have 100 or more years from the beginning to the end of the data record. Any years where the annual maximum is missing would mean that the number of data values might be less than this. This criterion yielded 56 raingauges. Of these, 39 showed a positive (increasing) trend, while 17 had a negative (decreasing) trend. Figure 4.10 shows the geographical distribution of these fitted trends. The results where obtained using a linear trend line fitted by regression.

When applied separately to each raingauge, a formal test of ‘no trend’ (against the alternative of a linear trend with time) gave 10 raingauges having significant trend at the 10 per cent level. These results are for a two-sided test using a permutation test of the coefficient of linear trend in order to avoid making assumptions about the statistical distributions. The construction of the test relies on the assumption that there is no statistical dependence between the annual maxima for successive years. In each of the 10 cases where the trend is formally significant, the trend is positive (increasing). The number 10 here may be compared with the number of cases expected to be judged to have a significant trend if there really is no trend; that would be 5.6 out of 56. Any more complicated assessment of this number would need to take into account the fact that the tests for the individual raingauges are not independent of one another. This dependence arises from the large statistical dependence of rainfalls at neighbouring raingauges in the same day, a dependence which carries through to the annual maxima.

Nonetheless, it is interesting to note that Figure 4.10 does not show substantial groupings of positive and negative trends. This argues against the existence of a broad-scale trend. In addition, in some instances where a large trend is fitted to the available data record for a gauge, it has been possible to find a shorter but partially overlapping record at a neighbouring raingauge which indicates that the longer record started when the annual maxima were unusually high or low and that, had even a few of the earlier years’ conditions been included, the fitted trend would have been rather different.

There are of course other considerations involved in judging whether an apparent trend is real. In the case of raingauge data, there are questions regarding changes of exposure over time. In the present instance, the longest data records start when common standards for instrumentation and recording practices were only just beginning to be developed.

The simple analyses of trend carried out within this project have not revealed any evidence that would invalidate analyses of annual and seasonal maxima. This is based on the assumption that there are no trends in extreme rainfall amounts over the period for which data are available. This assumption was made for the FEH analysis of rainfall and is retained here. There remains scope for further exploration of trend-related possibilities, such as that trends might exist only for sites of relatively high or low rainfall.

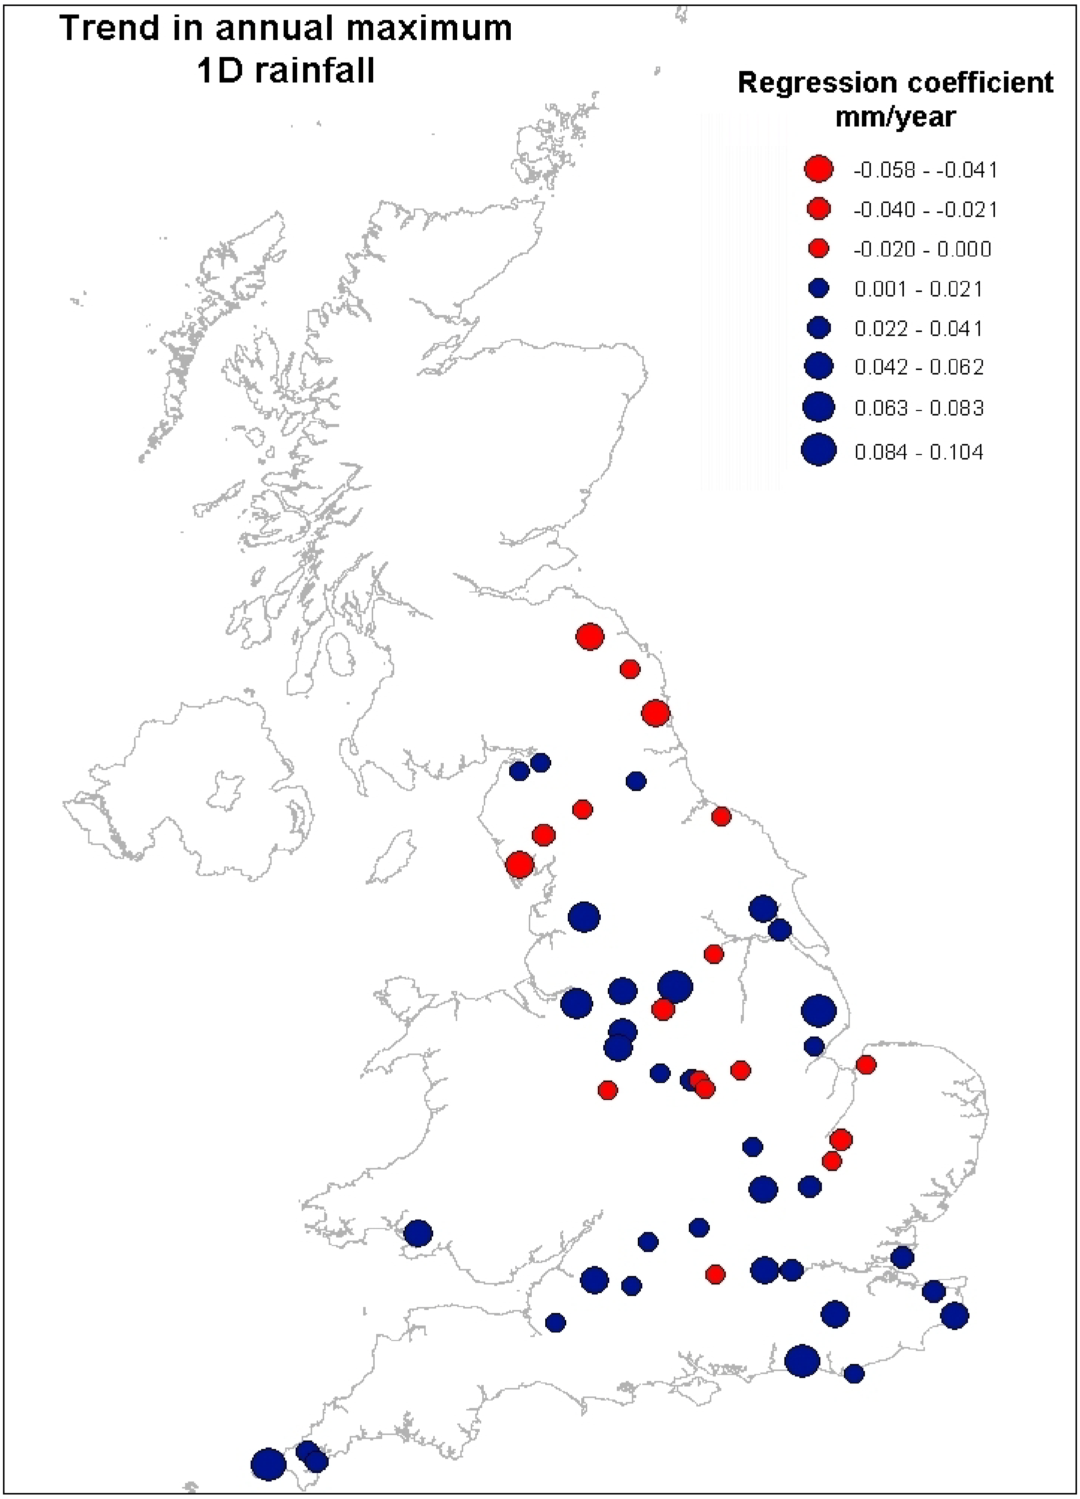

Figure 4.10 The geographical pattern of estimated trends fitted to the annual maxima of daily rainfall for 56 raingauges selected to have long records

Osborn *et al.* (2000) investigated trends in daily rainfalls binned into 10 different categories depending on rainfall amount. Their main dataset comprised 110 UK stations with data from 1961 to 1995, but they also used a smaller set of 11 gauges with records going back to 1908. The analysis was carried out for four seasons. Averaged over the UK, heavy rainfall was found to increase during the winter and decrease during the summer in the period 1961 to 1995, with trends being less clear in the spring and autumn. When put in the context of decadally smoothed rainfalls from the longer dataset (1908-1995), the decrease in summer rainfall during 1961-1995 appears to be a return to a longer-term average from an anomalous period in the 1960s. The winter rainfall does not show any obvious long-term trend (*Figure 12b* of *Osborn et al. (2000)*), but the level of significance of the linear trend estimate for the period 1908-1995 is not presented in the paper. The Osborn study does not include trend estimates for the year as a whole, so comparison with trends in the annual maxima of the present study is not straightforward. However, it can be noted that for the 1-day duration, 64 per cent of the annual maxima occur in the summer half-year between 1 May and 31 October (83 per cent for the 1-hour duration).

Maraun *et al.* (2007) updated the trend study by Osborn *et al.* (2000) to 2006. They also derived new historical spatial average precipitation series, using varying numbers of available gauges in each time period. For the whole of the UK, the number of gauges varies from 37 in 1900-1920 to around 540 in 1960-2006. The varying number of gauges in the different time periods means that some caution is needed when interpreting the results, as there may be issues of inhomogeneity. Using the new series of spatial average rainfalls, they found a long-term increase in winter precipitation intensity, whereas the summer rainfall intensities are more consistent with interdecadal variability than with an overall trend.

Fowler *et al.* (2005) suggest that the HadRM3H regional climate model may be used with some confidence to estimate future extreme rainfall distributions, since the model has good predictive skill in estimating statistical properties of extreme rainfall during the baseline period 1961-1990. Using this model, and the IPCC SRES scenario A2, Ekström *et al.* (2005) indicate that 1-day rainfalls of a 50-year return period may be about 10 per cent greater in 2070-2100 than in the baseline period across the UK.

5. Revision of the FEH standardisation

5.1 Introduction

As already discussed, this project has adopted the basic approach taken by the FEH analysis, using the two-stage procedure of standardisation and growth curve construction to develop sets of rainfall frequency curves for the whole of the UK. This has been followed by the development of a depth-duration-frequency (DDF) model to provide rainfall estimates for any location, duration and return period. Each of the steps in the process has been reviewed, and the aspects of the FEH procedures identified as offering the best prospects for improvement have been revised. This section considers the basis of the standardisation used within the FORGEX methodology and presented within FEH Volume 2.

5.2 Existing FEH standardisation

The FORGEX method used in developing the FEH DDF model adopted the index-flood approach, deriving rainfall-frequency estimates from the product of the index variable (RMED, the median annual rainfall depth of the required duration) and a set of growth factors relating RMED to *T*-year values. For each duration analysed, at-site values of RMED were used to standardise individual annual maxima, with the aim of removing differences in the underlying statistical distribution of maxima between sites. Therefore, following standardisation, annual maxima at individual sites were assumed to have the same statistical distribution. Using RMED as the standardising variable had the advantage of simplicity, since its value has an exceedance probability of 0.5, corresponding to a return period of 2 years. It also unified the statistical approaches to frequency estimation across the FEH methods for rainfall and floods.

It is clear that the standardisation method selected plays an extremely important part in the fundamental reasoning behind FORGEX, in that there is an intrinsic assumption that, after standardisation, the statistical distributions of rainfalls at different sites are the same. The locally-targeted nature of the FORGEX procedure means that this need only hold within a radius of 200 km, which is the maximum distance from the focal point at which annual maxima are permitted to contribute to the growth curve. Even this can be weakened somewhat to allow larger departures from the assumption at greater distances. However, the fundamental idea in FORGEX is to merge together information from a number of raingauge locations as if the standardised rainfalls all have the same distribution. Ideally, the standardisation method selected should remove any predictable differences between the distributions. Here ‘predictable’ relates to any differences between the distributions for ungauged sites which can be foreseen on the basis of known descriptors such as location or SAAR (average annual rainfall); that is any mappable variable. Preliminary analyses for this project found that there were differences between the distributions of rainfall (standardised by RMED) at different sites, some of which can be dealt with by the revised standardisation approach outlined in Sections 5.3 and 5.4.

The standardisation applied within FORGEX to each annual maximum at a given site can be expressed as follows:

$$R_{standardised} = \frac{R}{RMED} = 1 + \frac{R - RMED}{RMED}. \qquad (5.1)$$

Here, R denotes the annual maximum value for a given year, $RMED$ denotes the median of the annual maxima for the particular site and $R_{standardised}$ denotes the standardised value. The FEH FORGEX procedure made the assumption that most of the variation in the distributional shape of the rainfall values between sites is removed by this transformation. The pooling of data from raingauges within successive radii centred on a target site is a limited recognition that the standardisation does not fully work in terms of achieving a common distributional shape. However, the procedure could only be expected to cope with slowly varying changes in distributional shape over the country.

5.3 Revised standardisation

This project has introduced a revised type of standardisation of the following form:

$$R_{revised} = 1 + \frac{R - RMED}{f \times RMED} = 1 + \frac{1}{f}\left(R_{standardised} - 1\right). \qquad (5.2)$$

Here, $R_{revised}$ is the revised standardised rainfall, other quantities are as before, and f is a new scaling factor which varies from site to site. The revised standardisation is based on the premise that, although the simple FEH standardisation is effective in bringing the distributions of rainfall together so that they have the same location parameter as defined by the median, other differences still exist for which adjustments can also be made by applying the site-specific factor, f. This factor reflects the differences between sites of the spread of the distributions of the standardised values. Since the final method for estimating rainfall needs to be applied at ungauged locations, the scaling factor needs to be specified in terms of readily available geographical variables, rather than being derived directly from rainfall data.

The formula for the revised standardisation can readily be reversed to allow the determination of a design rainfall from the corresponding value of the standardised value:

$$R = RMED \times \left\{1 + f\left(R_{revised} - 1\right)\right\}. \qquad (5.3)$$

5.4 Relationship with site descriptors

The specification of the standardisation factor, f, was explored by defining a median-based version of the coefficient of variation, *lcmed*, as:

$$lcmed = \frac{\lambda_2}{RMED} \qquad (5.4)$$

and by analysing possible relationships between *lcmed* computed at raingauge sites and a number of site descriptors. Here, λ_2 is a scale parameter commonly used in the context of L-moments, the 'L-scale' (see Appendix I for some discussion). For both λ_2 and $RMED$, direct sample-based estimates are used here. Note that the designation '*lcmed*' is used in order to distinguish this quantity from two mean-based coefficient of variations which, when using ordinary moments and L-moments, are commonly denoted by 'CV' and 'L-CV' respectively.

The site descriptors available were National Grid coordinates (easting and northing), altitude and SAAR (Standard Annual Average Rainfall in mm). The analysis was conducted for the 11 key durations ranging from 1 hour to 8 days for which annual maxima were available, with each duration being treated separately. After considering various combinations of the site descriptors noted above, a decision was made to construct a transformed variable, *rsaar*1, where

$$rsaar1 = 1.5 - \frac{1000}{SAAR}\,. \qquad (5.5)$$

This transformation was selected so that *rsaar*1 increases for increasing SAAR (for ease of understanding), and so that zero is within the range of transformed values. Another reason for selecting this transformation of SAAR was that the relationship between *lcmed* and SAAR could be represented reasonably well by a linear relationship between *lcmed* and *rsaar*1, albeit with substantial error. Figure 5.1 shows examples of the fitted relationships for annual maxima.

Analyses of the residuals from this single-variable model suggested that, for the higher durations, an improvement could be made by using a predictive model of the form

$$lcmed = \alpha + \beta \times rsaar1 + \gamma \times ngy\,. \qquad (5.6)$$

Here, *ngy* is the National Grid northing expressed in 1000km units, giving values that lie between 0 and about 1.2. The parameter γ is fixed at zero for the sub-daily durations but is fitted for durations of 1 day and above. Appendix D presents some plots which show a comparison of the sample estimates of *lcmed* with the modelled values. Plots are also shown of the residuals from the model.

The final stage of the analysis was to replace the set of equations for *lcmed* with an equivalent set of equations for a standardisation factor based on *lcmed*. An

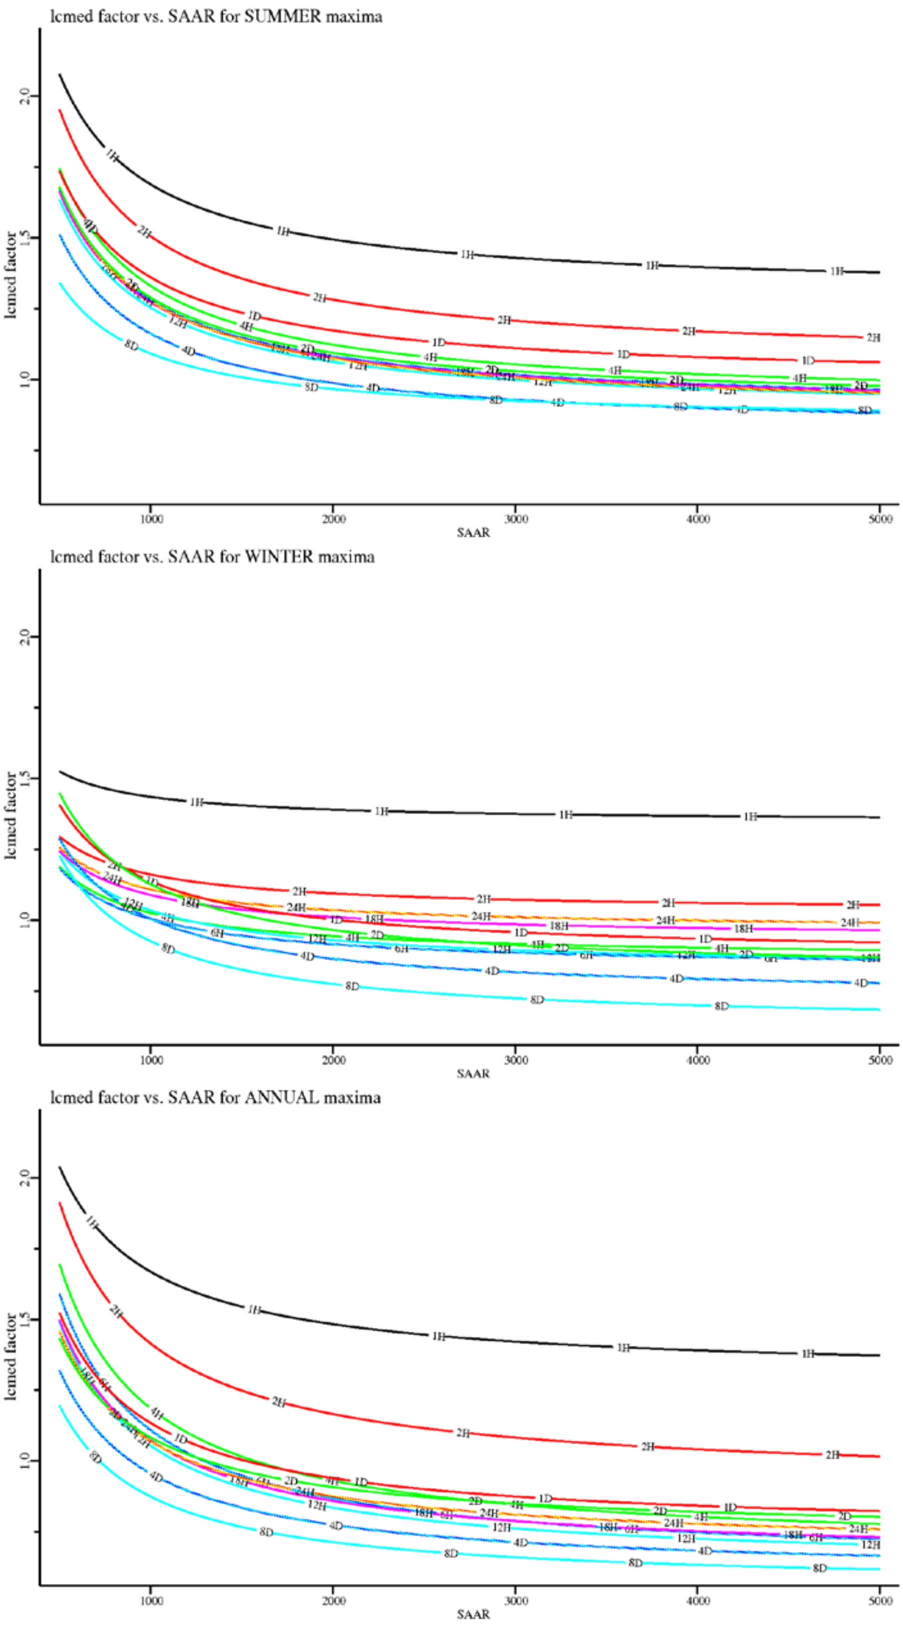

Figure 5.1 Fitted relations between lcmed and SAAR. (Each line is labelled with the duration of the total being analysed. In each case a value of *ngy*=0.5 is used.)

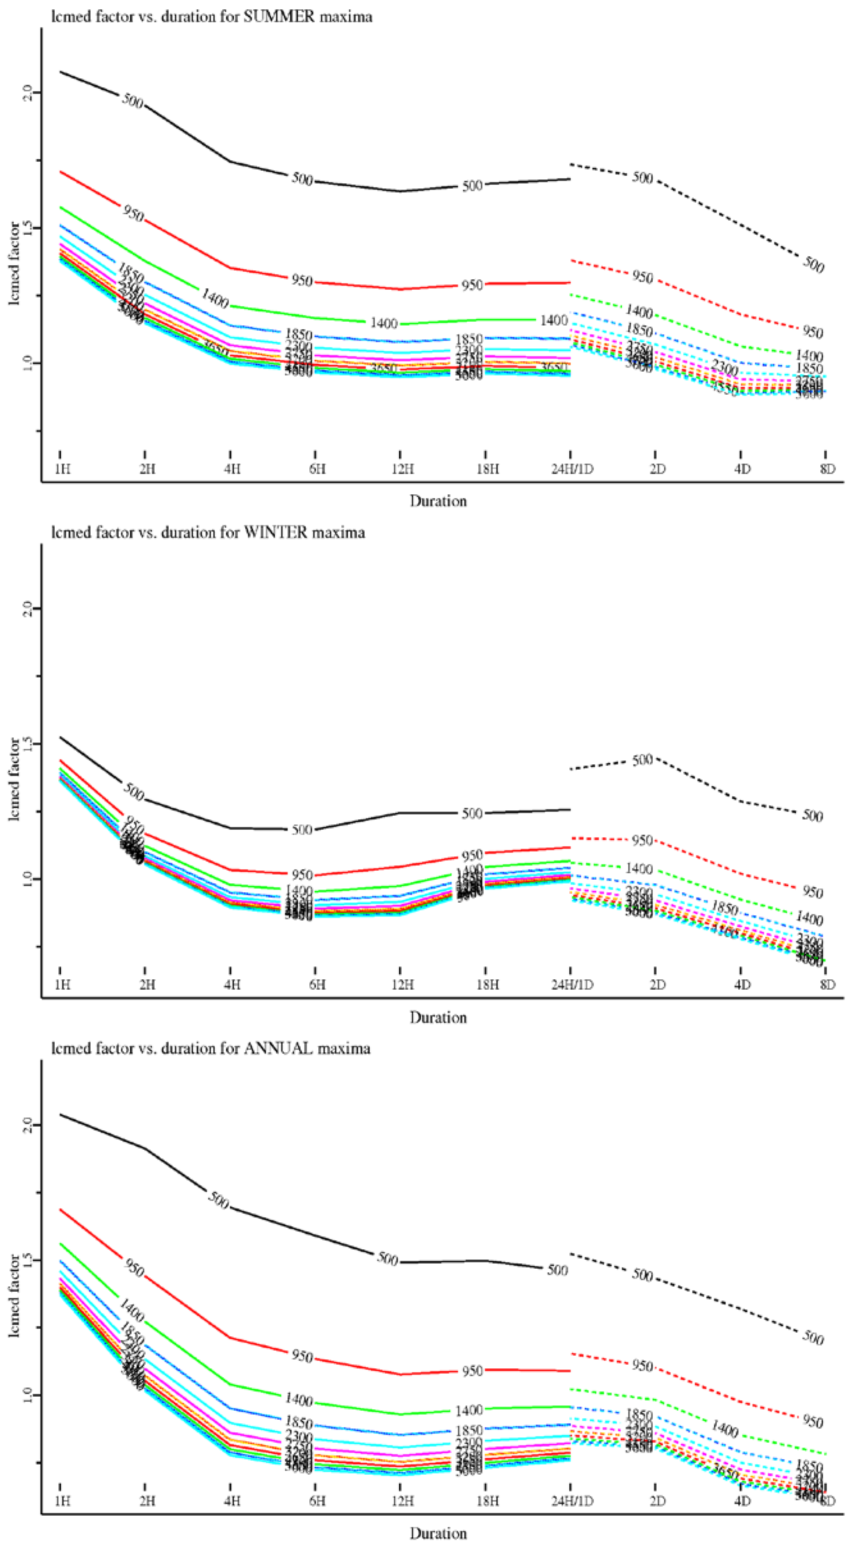

Figure 5.2 Fitted relations between lcmed and the duration of the total being analysed. (The duration scale is ordinal and 24 hours and 1 day are plotted at the same location. Each line is drawn for a selected value of SAAR which is shown in the label for the line. In each case a value of *ngy*=0.5 is used.)

overall scaling constant was used so that the L-scale of the standardised data would be approximately constant across the range of durations. The scaling factor was selected to give values of the standardisation factor in the region of 1.

The additional standardisation factor f which appears in Equation 5.2 is defined in terms of coefficients *a, b* and *c* as

$$f = a + b\frac{1000}{SAAR} + c \times ngy \,. \qquad (5.7)$$

The values of the coefficients are derived from the regression coefficients for *lcmed* by dividing these by 0.15. Values of the coefficients that are used for standardising annual maximum rainfall are shown in a simplified form in Table 5.1. Appendix D includes tables of the coefficients to the higher precision actually implemented for all three cases of summer, winter and annual maxima.

Table 5.1 Coefficients for the additional standardisation factor f for annual maximum rainfalls for different durations

Duration	a	b	c
1 hours	1.30	0.37	0
2 hours	0.92	0.50	0
4 hours	0.68	0.51	0
6 hours	0.63	0.48	0
12 hours	0.63	0.44	0
18 hours	0.65	0.43	0
24 hours	0.68	0.38	0
1 day	0.71	0.39	0.07
2 days	0.63	0.35	0.20
4 days	0.44	0.36	0.30
8 days	0.43	0.32	0.26

Figures 5.1 and 5.2 show how the relationship between the standardisation factor, f, and SAAR changes with duration. These figures have been drawn for the case of a location with *ngy*=0.5.

At some later time it may be appropriate to re-evaluate this analysis of standardisation with the possible intention of providing adjustment coefficients which vary smoothly across durations. Later analysis may also allow other geographical variables to be brought into account. The different bases of the maximum duration totals of, for example, the 24-hour and 1-day datasets, would also need to be considered when contemplating the variation of the coefficients across durations. For the purposes of the present project, it is worth noting that the effects of the revised standardisation are largely contained within sub-analyses which treat a single duration at a time. The analysis progresses by, for each duration separately:

(i) standardising the data;
(ii) formulating a model for spatial maxima (Section 6.4);
(iii) applying the model for spatial maxima (Section 6.5);
(iv) combining information from spatial maxima across networks of different sizes (Section 7);
(v) fitting size-frequency relationships;
(vi) removing the effects of the standardisation, to form rainfall-frequency relationships.

It is only after the effects of standardisation have been removed that information from the different durations is brought together in the DDF (Depth-Duration-Frequency) model (Section 8).

5.5 Effects of the revised standardisation

The effect of the revised standardisation is shown in Figures 5.3 to 5.5. These figures show rainfall frequency plots of annual maximum rainfall, of median-standardised rainfall and of rainfall with the revised standardisation. These plots are shown on the usual Gumbel scale and each line represents the 1-day annual maxima for a single raingauge. Comparison of Figures 5.4 and 5.5 shows that there has been a noticeable reduction in the spread of the rainfall frequency curves as a result of the revised standardisation.

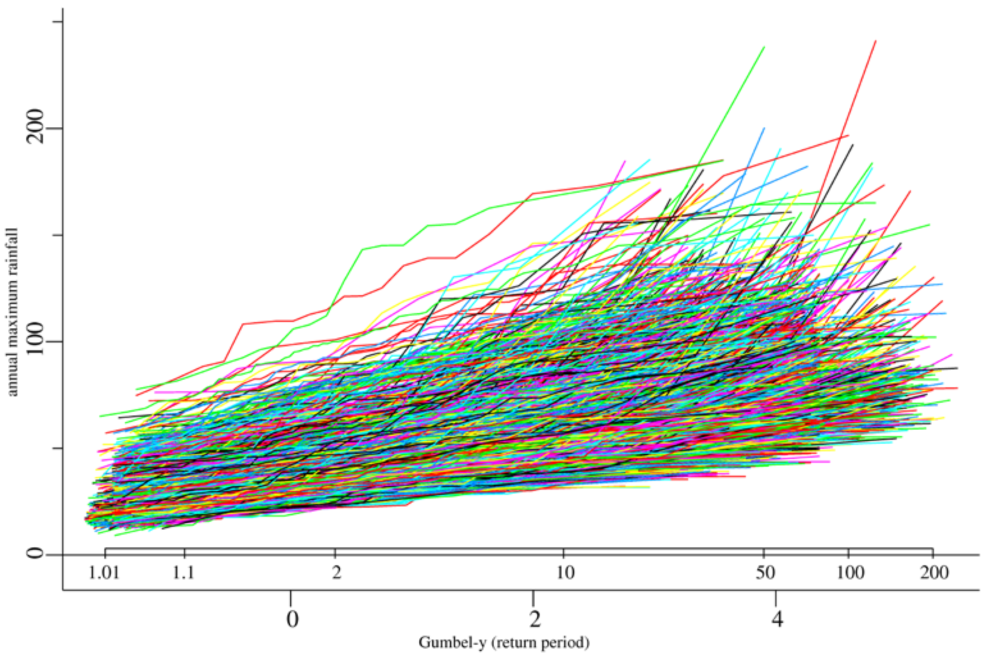

Figure 5.3 Rainfall frequency curves for 1-day rainfall for the UK (Gumbel scale with return period in years)

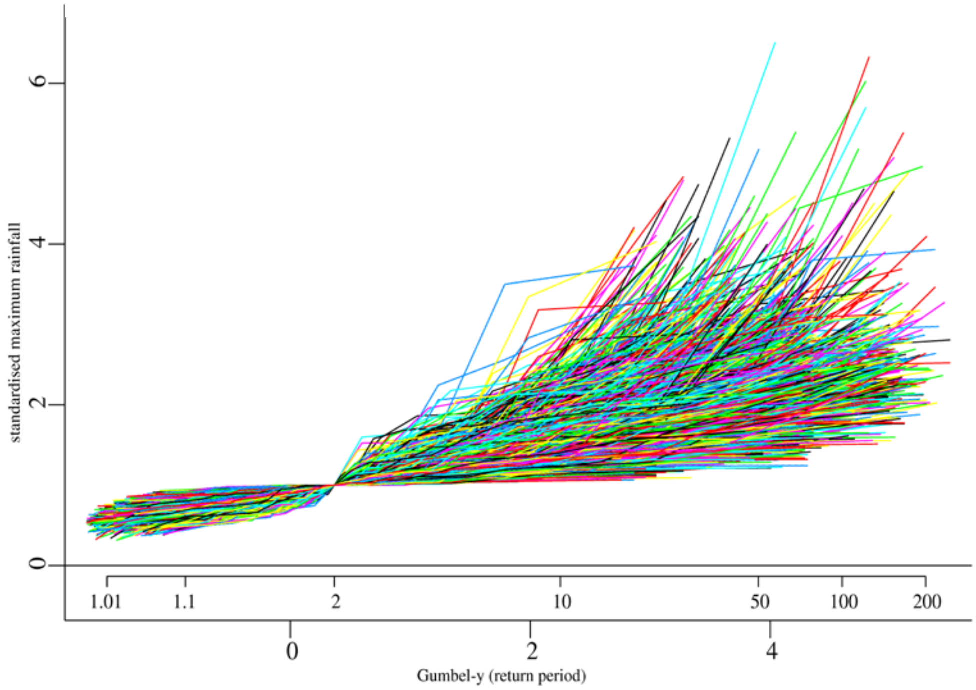

Figure 5.4 Rainfall frequency curves for median-standardised 1-day rainfall for the UK (shown on Gumbel scale with return period in years)

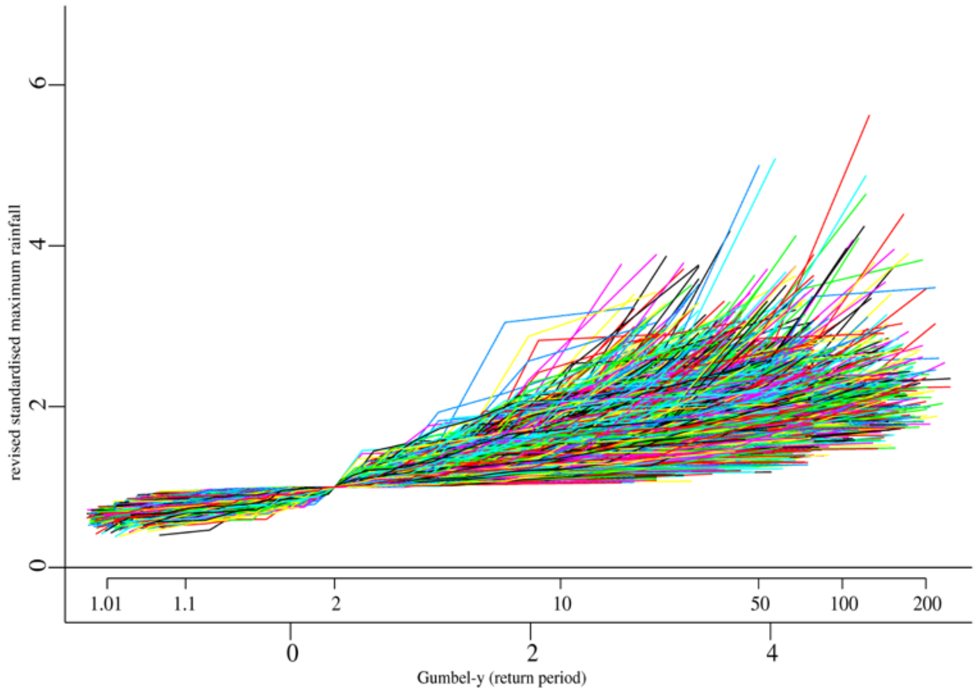

Figure 5.5 Rainfall frequency curves for 1-day rainfall with the revised standardisation for the UK (shown on Gumbel scale with return period in years)

Figures 5.6 and 5.7 show the results of applying the existing FEH standardisation and the revised standardisation on typical rainfall growth curves across the UK for durations of 1 hour, 6 hours, 1 day and 8 days. The growth curves shown are average GEV distributions fitted to all the annual maxima at gauges within 52.5 km of the centre of each box.

One effect of the revised standardisation is to ensure that the slopes of the growth curves for different durations are closer together than for the simple standardisation by the median. However, this is just an artefact of the way in which the standardisation coefficients in Table 5.1 have been defined and is unimportant to the FORGEX procedure. Of more importance is the reduction in the extent of the variation across the country of the slopes of the frequency curves. This can be seen in a comparison of Figures 5.6 and 5.7.

One important feature that can be seen in Figure 5.7 is the departure of the curve for the 1-hour duration from the curves for other durations that occurs in Scotland and to some extent in Northern Ireland. For these cases, as elsewhere, the slopes of the curves are fairly similar in the vicinity of the median (the vertical dashed line in each box) and it is the curvature of the line which explains the visual difference between the curves when judged across the country. The different curvatures correspond to different skewnesses and the revised standardisation has no effect on the skewness of the distributions.

The primary reason for introducing the revised standardisation has been to reduce the variation between the frequency curves of standardised rainfall. Comparisons of Figures 5.4 and 5.5 and of Figures 5.6 and 5.7 show that there has been a substantial improvement.

5.6 Summary

The standardisation step within the FORGEX procedure, whereby annual maxima at each raingauge site are divided by the at-site median value of the appropriate duration, RMED, has been modified within the current project. The revised standardisation, as expressed by Equation 5.2, involves a standardisation factor which varies from site to site and can be calculated from Equation 5.7, using the coefficients presented in Table 5.1. The effects of the revised standardisation on the rainfall growth and frequency curves produced by FORGEX are discussed in Section 7.

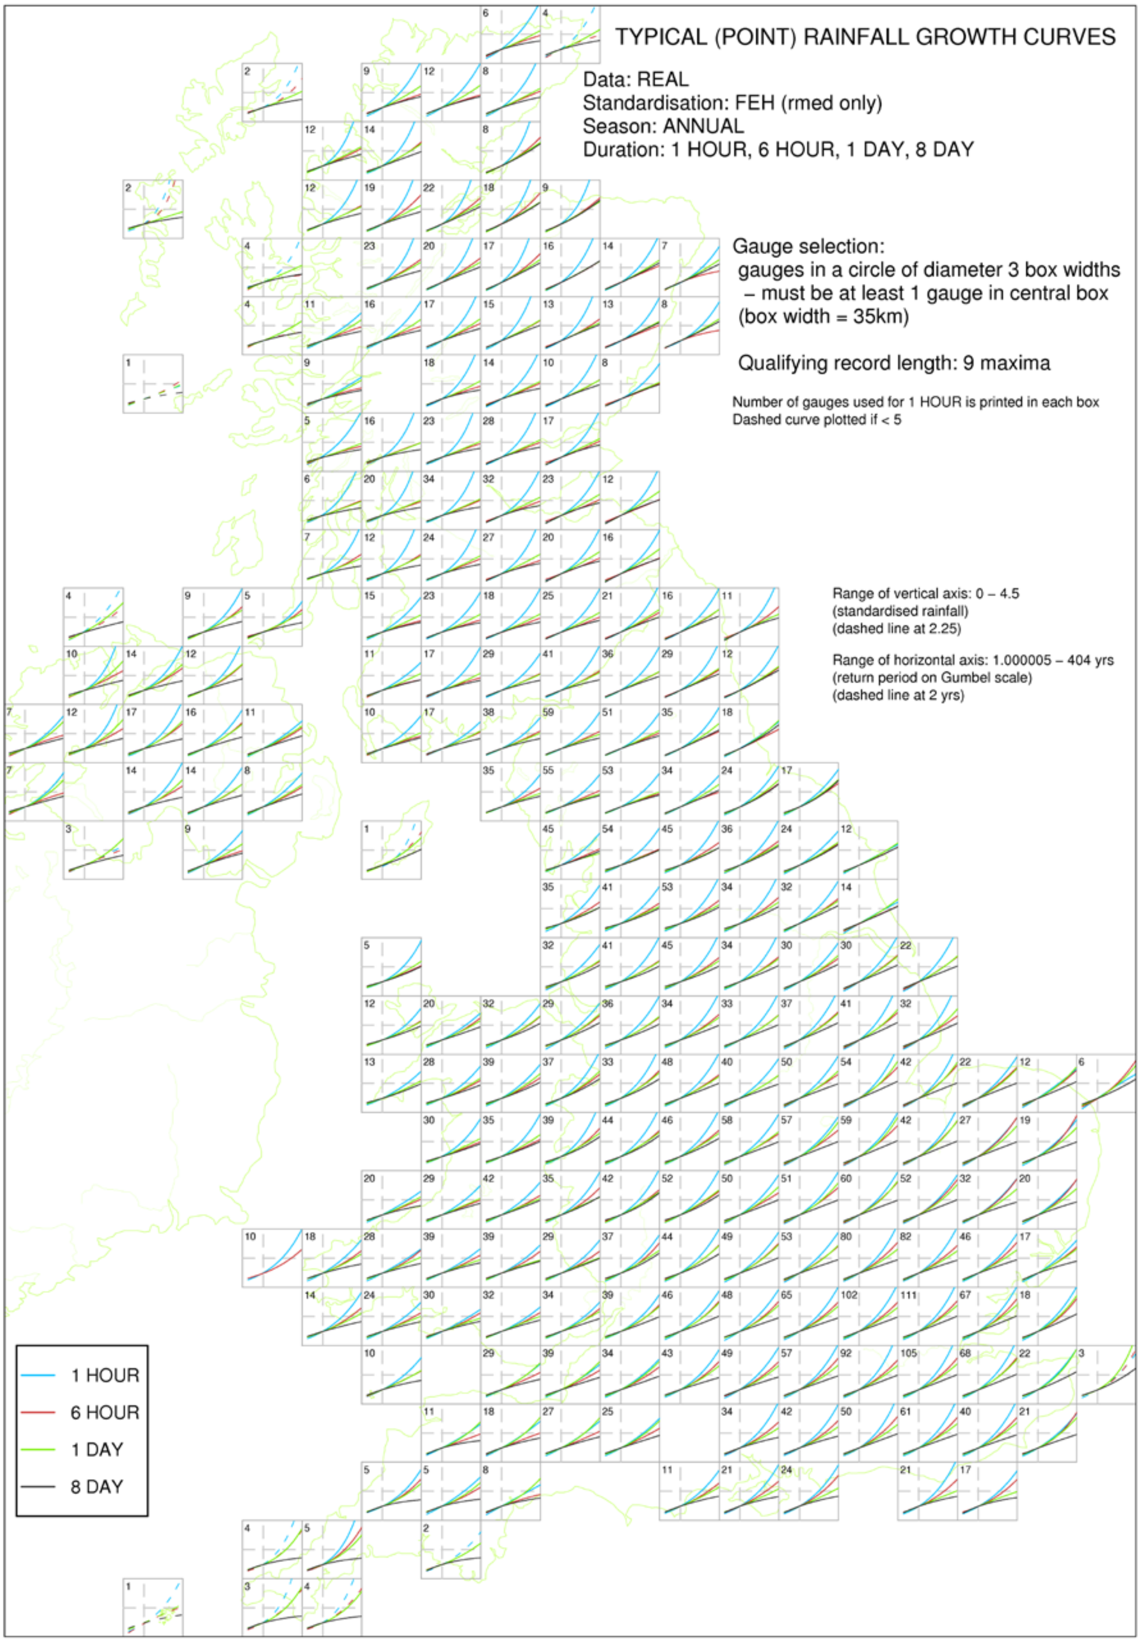

Figure 5.6 Typical growth curves for various durations (fitted GEV distributions and FEH standardisation)

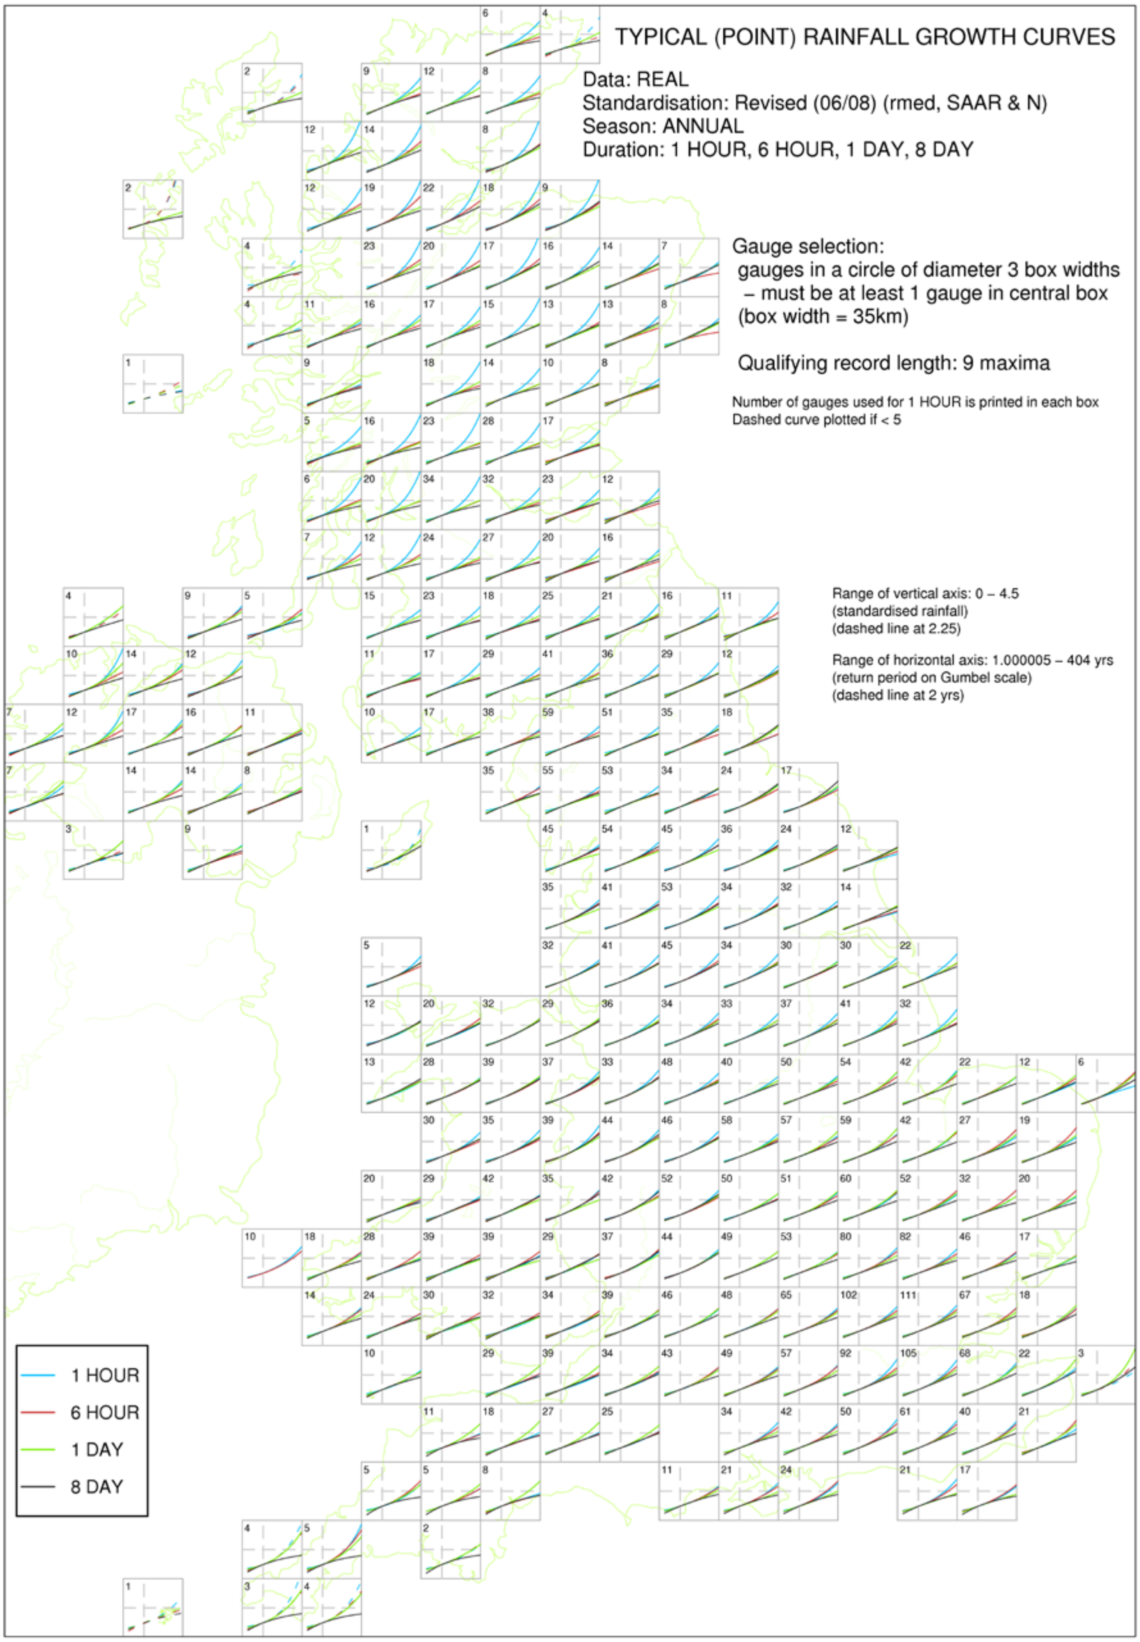

Figure 5.7 Typical growth curves for various durations (fitted GEV distributions and revised standardisation)

6. A revised model of spatial dependence in rainfall extremes

6.1 Introduction

The second key component of the FEH FORGEX methodology that has been revised during the course of the current project is the form of the spatial dependence model which underpins the use of network maximum points in the construction of rainfall growth curves. This section outlines the analysis leading to the development of a revised model of spatial dependence in rainfall extremes which includes the effect of return period. Further details can be found in Appendix F.

6.2 Dales & Reed model

The FORGEX method incorporates a model of spatial dependence in rainfall extremes developed by Dales & Reed (1989) to define the plotting positions of the netmax points. Spatial dependence is exhibited when a rainstorm of large areal extent is experienced over a network of gauges. In some cases, this may be the annual maximum event over the gauge network. Dales & Reed show that if the distribution of a network maximum calculated from N independent and identically distributed values each having a common Generalised Extreme Value (GEV) distribution is considered, this distribution lies exactly $\ln N$ to the left of the regional growth curve on a Gumbel reduced variate scale. In practice, because of inter-site dependence in annual maxima, the netmax growth curve is found to lie a shorter distance to the left. Dales & Reed label this distance $\ln N_e$, terming N_e the effective number of independent gauges.

Dales & Reed (1989) developed a model of spatial dependence to estimate N_e for any raingauge network. Model parameter values were estimated for different regions of the UK and for national average conditions. It is the latter form of the model that is incorporated into the version of FORGEX used in the FEH:

$$\frac{\ln N_e}{\ln N} = 0.081 + 0.085 \ln AREA - 0.051 \ln N - 0.027 \ln D , \qquad (6.1)$$

where N is the number of gauges, $AREA$ is a nominal area spanned by the network and D is the rainfall duration in days. Dales & Reed (1989, Section 4.4) provide a formula for calculating $AREA$:

$$AREA = 2.5 \times \bar{d}^{\,2} , \qquad (6.2)$$

where $\bar{d}$ is the mean of the pairwise inter-gauge distances in the network being considered.

6.3 Spatial dependence within FORGEX

One of the important features of the FORGEX procedure is the idea of using the behaviour of the distribution of the maximum rainfall within a network of gauges (at moderate return periods) to estimate the distribution of rainfalls at individual raingauges (at high return periods). This idea is discussed in the present section, together with some possible difficulties with the approach as previously applied.

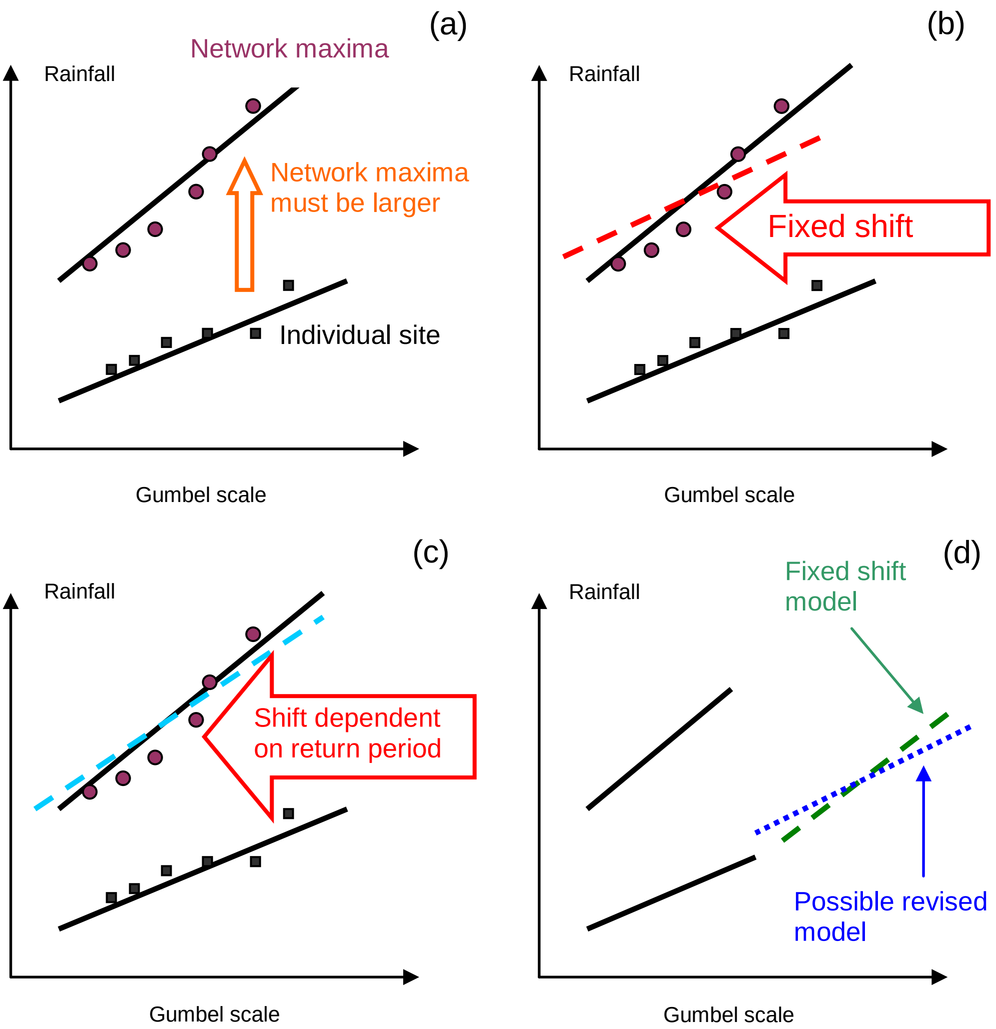

Figure 6.1 Basic ideas of the original (FEH) FORGEX procedure and of the procedure developed here (see text)

Figure 6.1 shows a pictorial representation of the relationship between the distributions involved in the FORGEX procedure. The plots here are fictitious and are shown in the form of rainfall frequency plots using a Gumbel reduced variate scale. This latter scale is an important feature of the model underlying the FORGEX procedure. The two frequency curves involved are those for the annual maxima of rainfalls at a single raingauge and the annual maxima of the largest of the rainfalls observed across a fixed network of raingauges. The immediate point of Figure 6.1(a) is that the frequency curve for the network maxima will plot higher than the curve for a single raingauge.

The basic model used for spatial dependence within the FORGEX procedure used by FEH is illustrated in Figure 6.1(b). Note that this is not a complete model for spatial dependence, but rather involves only a direct representation of the effects of spatial dependence on the relationship between the rainfall frequency curve for network maxima and that for a single raingauge. The fixed shift model is that the frequency curve for a single raingauge is shifted to the left, with the same shift applying everywhere, giving the dashed line in Figure 6(b). This model is not entirely arbitrary. If there really were full independence between the rainfalls recorded at different raingauges, then this model holds with a shift whose size depends only on the number of gauges. It is clear also that, in the special case of full dependence between raingauges, the distribution of network maximum and single raingauge rainfalls would be identical, and therefore the model applies with a zero shift. A shift of an intermediate size therefore has a direct interpretation as being intermediate between full independence and full dependence.

There have been some concerns that the FORGEX procedure, as implemented for the FEH study, does not accommodate the sort of behaviour that some have thought it reasonable to expect, namely a dependence that varies either with size of rainfall or with return period. Specifically, dependence is thought to be lower for rarer events. Figure 6.1(c) shows a representation of this sort of behaviour.

The dependence model within the FORGEX procedure is used in a way that effectively leads to the network maximum curve being shifted to the right to counteract or reverse the left-shift of the dependence model itself. This shift to the right is illustrated in Figure 6.1(d). If the fixed-shift model is incorrect, and a variable shift to the left would be better, as indicated in Figure 6.1(c), then using the simpler model is likely to lead to estimates of rainfall for the single raingauge curve being too high at high return periods.

The FORGEX procedure of FEH does not strictly involve a shift of an overall estimate of the frequency curve for the network maximum. Instead, it works in an integrated way to calculate revised plotting positions for rainfall values from the network maximum dataset, so as to be appropriate for the single raingauge curve. This procedure automatically accommodates differences in the raingauge density available for different years for which the network maxima are available. The extra complication of the shift reversal stage is represented in Figure 6.2.

Figure 6.3 shows an example of the relationship between the distributions of annual maximum rainfall for a single gauge and for the network maximum. This example relates to the raingauge at Oxford (256225). The lower line shows the frequency plot for the standardised annual maximum 1-day rainfall (that is, divided by the median of the annual maxima) for this particular raingauge, together with

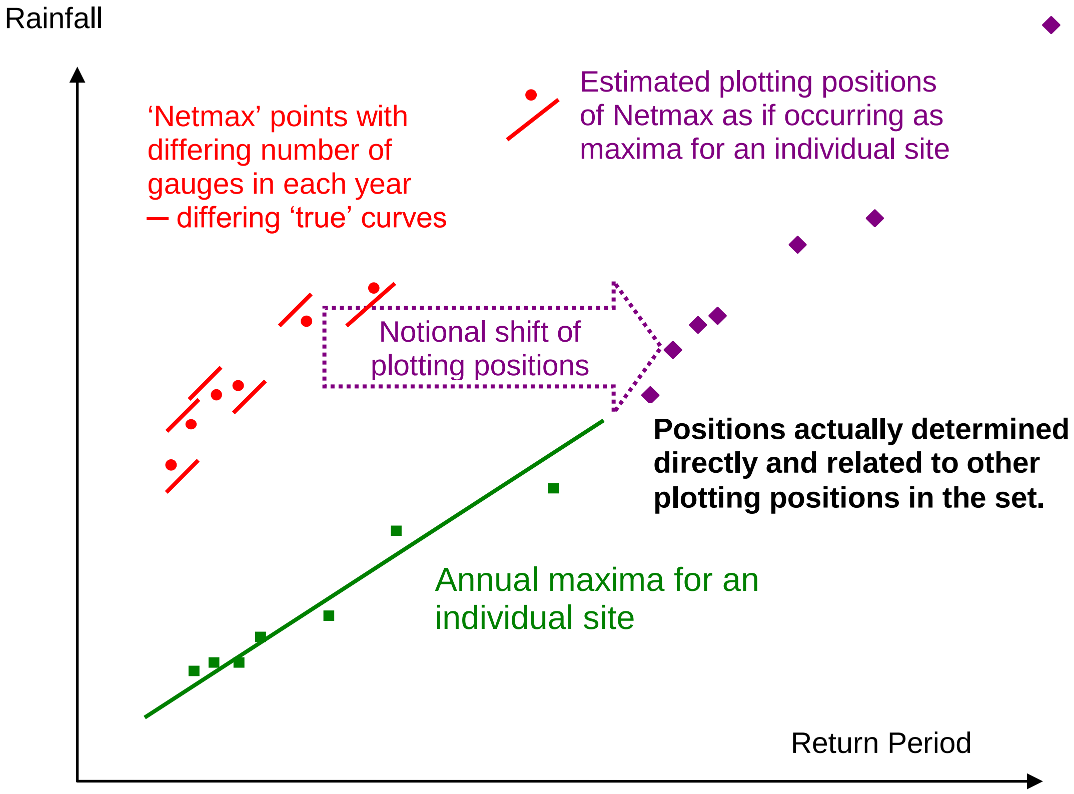

Figure 6.2 Effects of the right shift step in more complicated situations where the numbers of gauges in the network varies for different years

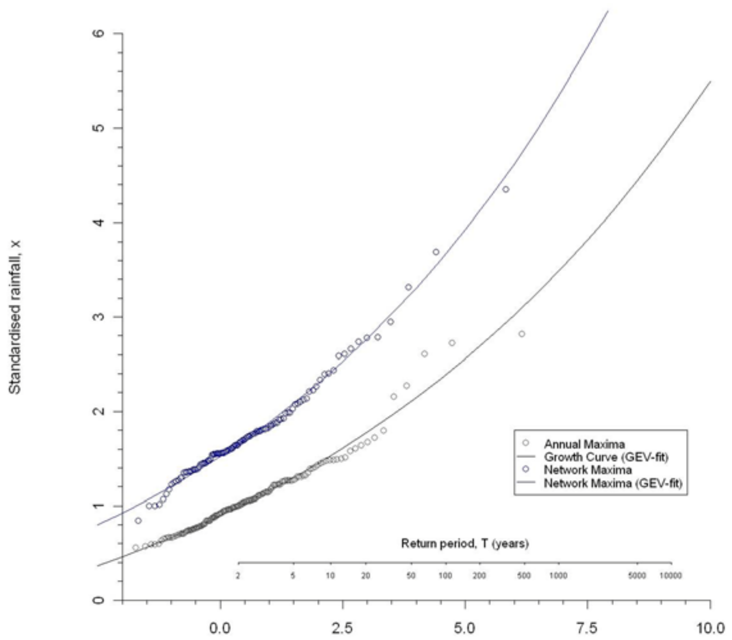

Figure 6.3 Comparison of the frequency curves for standardised rainfall at a single gauge and for the network maximum for a network containing 16 raingauges in a 100-km radius of the Oxford raingauge (256225)

the annual maxima for individual years plotted using the usual Gringorten plotting position formula. The upper line relates to network maxima derived so as to achieve a long record of network maxima corresponding to a network having 16 raingauges within a 100 km radius of the Oxford raingauge. The network maximum record is constructed on a year-by-year basis, where (for each calendar year) a random selection is made of 16 of the gauges with valid annual maxima for that year and, for each year, the largest of the standardised versions of these values is put into the constructed dataset.

Figure 6.4 is an extended version of Figure 6.3. In this case, the single randomly selected set of network maxima has been extended to include 50 randomly selected networks (using the same criteria as before), giving a set of 50 points at each plotting position. In addition, the medians of each of these sets have been joined to give a good estimate of the frequency curve for the network maximum of 16 gauges in a 100 km radius. The curve labelled as the *typical curve* differs from the lower curve in Figure 6.3 since, in this instance, it is constructed as an average curve for all the raingauges included in the randomly selected networks. The curve corresponds to a GEV distribution fitted by L-moments to the average of the L-moments of the standardised rainfalls at the selected raingauges. The final curve in Figure 6.4, labelled as the *independent curve*, is the result obtained from the assumption that the values at the 16 raingauges are statistically independent, taking the *typical curve* as the assumed common distribution. Therefore, the upper smooth curve is the result of the constant shift model, with the amount of shift calculated for the assumption of 16 gauges in the network having independent annual maxima. It can be seen that the median line from the resampled network maxima is rather lower than the *independent curve* for low return periods (up to about 3 years), but matches it rather well for higher return periods. This behaviour is essentially what has been described as a tendency to independence at high return periods. An alternative way of assessing the behaviour of the median line is to consider the slope of the curve in the region of the centre of the network maximum values (near the median). The slope here is larger than the slope of the *typical curve* and, more particularly, larger than the slope of the *independent curve*. This corresponds to the curve rising from ‘positive dependence’ (a position corresponding to an effective number of gauges smaller than the actual number of gauges) to ‘independence’. The work reported in Appendix E uses the slope of the frequency curve for the network maximum as a way of assessing the adequacy of the constant shift model.

Figure 6.5 shows an additional example. In this case, the networks which are included are restricted to have coverage areas (as measured by the Dales & Reed criterion) which are within a specific range of allowable areas. Here, the range is 8,000 to 16,000 km^2 and networks have 128 gauges. A further difference between this figure and Figure 6.4 is that the values derived from observed raingauge data are the standardised values using the revised standardisation discussed in Section 5 (that is including both SAAR and NGR northing in the standardisation). In general, the effect of using the revised standardisation is (for data for a 1-day duration) to reduce the slopes of the ‘typical’ and ‘independence’ curves and to reduce the variations of these as the central location changes. While there are some detailed changes in how the

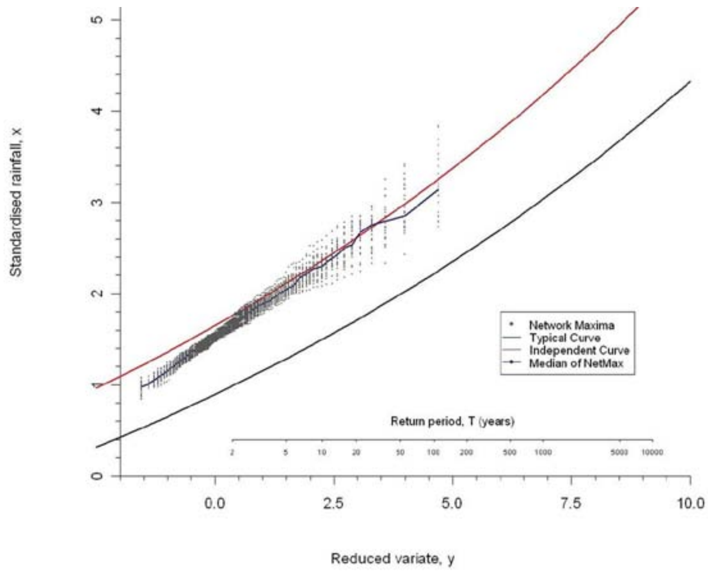

Figure 6.4 Comparison of the frequency curves for median-standardised rainfall at a typical gauge and for the network maximum for a network containing 16 raingauges in a 100-km radius centred on the Oxford raingauge (256225). (*The points relate to 50 resamples of the constructed network maximum series. The upper (red) smooth curve is the theoretical result obtained for the network maximum of 16 independent raingauges where the individual values have the lower (black) curve for their frequency curve.*)

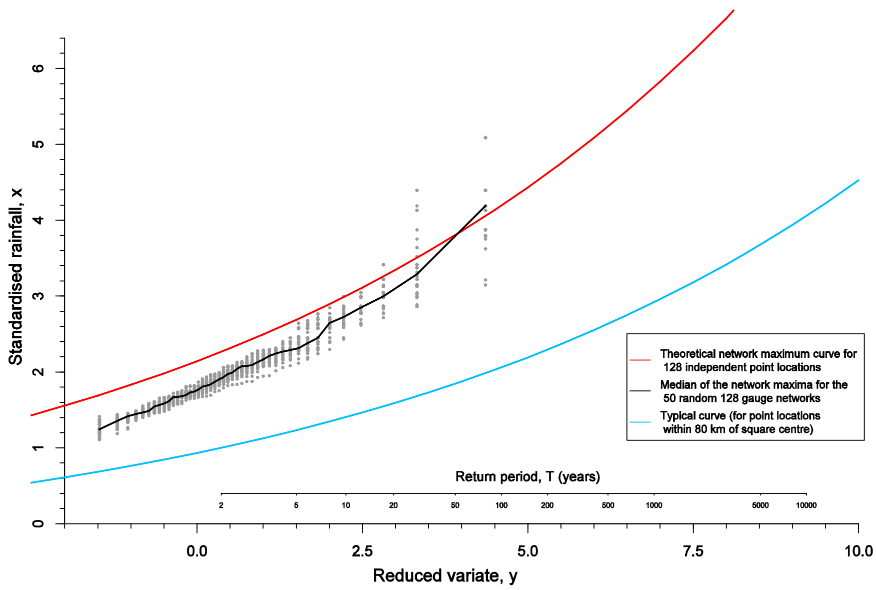

Figure 6.5 Comparison of the frequency curves for standardised rainfall (revised standardisation) at a typical gauge and for the network maximum for a network containing 128 raingauges in a 80-km radius of a selected location, and where the networks are chosen to cover areas between 8,000 and 16,000 km^2. (*The points relate to 50 resamples of the constructed network maximum series. The upper (red) smooth curve is the curve for the network maximum of 128 independent raingauges where the individual values have the lower (blue) curve for their frequency curve.*)

'network maxima' points (and the derived median line) plot between the two types of standardisation, there are no strong systematic effects apart from those which arise in relation to the changes in the 'typical' and 'independence' curves.

6.4 Model for spatial dependence

6.4.1 Evaluation of the constant shift model

As indicated earlier, the FEH's use of a constant shift model as part of the FORGEX procedure has been subject to doubt. As part of the present project, the question has been re-assessed in two ways. Firstly, analyses similar to that presented in Figure 6.5 have been carried out on a systematic basis. This has been done for various areas and numbers of gauges in networks, and in ways which allow possible variations in behaviour across the country to be considered. Secondly, a more formal assessment has been made, which uses a quantitative approach to assessing the validity of the model and which provides a convenient way of synthesising the results for the whole country in a single plot. These studies are described in more detail in Appendix E. The essential conclusion is that the constant shift model has been shown to be invalid.

6.4.2 Model definition

This project has developed a new model of spatial dependence in rainfall extremes. Note that this is not a full model for the joint distribution of annual maximum rainfalls at collections of raingauges, but instead specifies $F_{\max}(x)$ as a function of $F(x)$ (for the same value of x), where

- $F(x)$ is the distribution function of (standardised) annual maximum rainfall at a typical site,
- $F_{\max}(x)$ is the distribution function of maximum (standardised) rainfall across a network of raingauges.

The fact that annual maximum rainfalls at nearby raingauges are statistically dependent is reflected in the relationship of $F_{\max}$ to F. The specification of the spatial dependence model relates $F_{\max}$ at a given 'x' to F at the same 'x'. Therefore, it is sometimes convenient that the specification specifies $F_{\max}$ in terms of F without specifically using 'x' in the notation. The number of gauges, denoted by N, plays an explicit part in the specification of the model, because there is a need to ensure that the model behaves reasonably compared with the simplest case where the rainfalls at all gauges are independent. In the independent case

$$F_{\max}(x) = \{F(x)\}^N . \qquad (6.3)$$

This section presents a general overview of the model that has been developed and the thinking behind it, while Appendix F (Section F.1) gives a more detailed treatment of some technicalities.

It is understood that the model for the relation between the distributions of network maxima and individual site annual maxima is to be applied within a framework where only limited descriptive information about the network is available. There is an assumption that, apart from a few such descriptors, the networks concerned have similar characteristics. The descriptors used in the FEH study were the area of coverage and the number of gauges in the network. In this project a model has been developed which uses SAAR (standard average annual rainfall) as an additional characteristic.

The discussion in Section 6.3 has been in terms of data plotted in a standard way on a Gumbel scale, and of horizontal shifts with respect to this scale. This scale is of relevance to extreme value theory because it provides a simple representation for the effect of taking the maximum of a number of independent random variables. In the present situation, it provides a basis for constructing a model for spatial dependence based on transformed versions of F, $F_{\max}$ and F^N (the N'th power of F).

For a given value, x, of rainfall, the value of the distribution function, $F(x)$, for a typical gauge would be represented by the point y on the horizontal axis, where y is given by

$$y = -\ln(-\ln F(x)). \qquad (6.4)$$

In standard plots using a Gumbel scale, the value of rainfall (x) is plotted on the vertical axis.

The model of complete independence is represented by saying that, for the distribution of the maximum of N independent values for this distribution, this point moves to a new location, y_N, where y_N is given by

$$y_N = -\ln\left(-\ln\left\{F(x)^N\right\}\right) = -\ln N - \ln(-\ln F(x)) = -\ln N + y. \qquad (6.5)$$

The model for the effects of spatial dependence that is used here is framed so as to define the location $y_{\max}$ for the distribution function for the network maximum, $F_{\max}(x)$, where $y_{\max}$ is given by

$$y_{\max} = -\ln(-\ln F_{\max}(x)). \qquad (6.6)$$

Note that $y_{\max}$ and y_N are the transformed versions of the distribution functions of:

(i) ($y_{\max}$) the maximum of N dependent outcomes from distribution function F, where the dependence is unknown but reflects the actual spatial dependence of rainfall;

(ii) (y_N) the maximum of N independent outcomes from distribution function F, which has distribution function F^N.

The spatial dependence model is formulated in parallel to the relation for independence, which explicitly shows y_N as a function of y: specifically,

$$y_N(y) = -\ln N + y. \qquad (6.7)$$

Therefore, the model is expressed initially by treating y_{max} as a function of y, denoted by $y_{max}(y)$, and this can then be re-expressed as a rule for calculating $F_{max}(x)$ from $F(x)$.

Given the reasonable notion that the effects of spatial dependence in the rainfall data will be somewhere between complete dependence and complete independence (encompassing the notion of 'positive dependence'), a requirement of the model is that $y_{max}(y)$ lies between y and $y_N(y)$. The notion of 'less dependence at high return periods' means that $y_{max}(y)$ should be closer to $y_N(y)$ for large values of y than for low values. The model used here is given by

$$y_{max}(y) = (-\ln N)H(y;a,b) + y, \qquad (6.8)$$

where $H(y;a,b)$ is the logistic distribution function with parameters a and b, which is given by

$$H(y;a,b) = \left\{1 + \exp\left(-\frac{y-a}{b}\right)\right\}^{-1}. \qquad (6.9)$$

This model leads to a function $y_{max}(y)$ which starts from y for large negative y (*i.e.* is asymptotic to y for return periods near 1), and approaches $y_N(y)$ for large positive y (*i.e.* is asymptotic to $y_N(y)$ for large return periods). In addition, there is symmetry in the behaviour for large positive and negative values of y. These features may be considered to be unrealistic or to be making unproven assumptions. However, it is not presently feasible to work with more parameters or to explore other approaches which might remove these constraints. The model does have more flexibility than the 'constant shift' model used by the FEH, which would be represented by

$$y_{max}(y) = (-\ln N)\rho + y, \qquad (6.10)$$

where ρ is a single parameter between 0 and 1. The aim would be to use the two parameters a and b to obtain a good match over the range of data available. If the data contained no strong evidence of an approach to independence for high return periods the values of a and b would reflect this and the fitted model would be such that $y_{max}(y)$ is close to $y_N(y)$ for values of y relevant to the model's use within FORGEX (that is for values of y only moderately larger than those used for model fitting).

For implementation, the parameters a and b are replaced by somewhat more meaningful parameters that relate more directly to the location of the function $y_{max}(y)$ between y and $y_N(y)$. In either case, the parameters relate to the strength of dependence within a given network of gauges, in terms of how this affects the distribution of the spatial maximum. Part of the model fitting process entails relating these parameters to the characteristics of the network such as the areal extent, number of gauges, etc.

For completeness, the above model (in terms of the Gumbel reduced variate, y) is equivalent to the following rule for calculating $F_{max}(x)$ from $F(x)$:

$$F_{\max}(x) = \{F(x)\}^{N^{K(x;a,b)}}, \quad (6.11)$$

where

$$K(x;a,b) = \left\{1 + \exp\left(\frac{a}{b}\right)\{-\ln F(x)\}^{1/b}\right\}^{-1}. \quad (6.12)$$

However, as noted in Appendix F, the basic form of the model has been modified slightly in order to ensure that the model defines a proper distribution function.

6.4.3 Modified parameterisation

The following alternative parameterisation has been used in place of the parameters a and b used in the logistic distribution. The revised parameterisation may be easier to interpret but, more importantly, is considered to be more appropriate for modelling when trying to construct a relationship between the dependence parameters and descriptors of the networks. In particular, the parameters γ_1 and γ_2 relate directly to the relative location of function $y_{\max}(y)$ between y and $y_N(y)$ .For this reason, they have been called the relative-location parameters. These parameters are discussed more fully in Appendix F. The overall model-fitting procedure consists of obtaining raw estimates of the relative-location parameters γ_1 and γ_2 for many sizes of networks, and then finding relationships to network size that will allow these relative-location parameters to be predicted. The simplest possible situation would be to have no network descriptors, in which case, the model-fitting procedure would produce a model for dependence equivalent to taking a weighted average of the values of γ_1 and γ_2 and then using the relationship which constructs the distribution of network maxima based on these average values. The actual procedure implemented uses equations which determine γ_1 and γ_2 as functions of network area, number of gauges and average SAAR value across the network.

The two relative-location parameters γ_1 and γ_2 are defined so that they each measure dependence on a scale from 0 to 1, where a value of '0' means independence and '1' means full dependence. The parameters relate to the distance at which the curve $Y = y_{\max}(y)$ lies between the two curves $Y = y_N(y)$ (independence) and $Y = y$ (full dependence). The relevant distance between the curves is then $1 - H(y;a,b)$. To obtain two relevant values, this distance value is evaluated at two values of y which are determined as follows. Two initial values y_1, y_2 ($y_1 < y_2$) are chosen and these remain fixed across all data sets. However, for a network of a given size, these values are converted to values y_1^* and y_2^* given by

$$y_1^* = y_1 + \ln N,$$
$$y_2^* = y_2 + \ln N, \quad (6.13)$$

where N is the number of gauges in the network being considered. The effect of this is to ensure that the places at which the relative-location parameters are determined remain aligned with the values of the reduced variate y for which information can be gained from the available data. This is presented in more detail in Appendix F.

The relative location parameters γ_1 and γ_2 are chosen to be the value of $1-H(y;a,b)$ at $y=y_1^*$, and at $y=y_2^*$. Therefore,

$$\gamma_1 = 1-H(y_1^*;a,b)\,,$$
$$\gamma_2 = 1-H(y_2^*;a,b)\,. \qquad (6.14)$$

Once a model relating γ_1 and γ_2 to the network configuration has been found, it can then be used to determine γ_1 and γ_2 for the networks required within the FORGEX procedure. These can then be converted to values of a, b for use within the new FORGEX procedure, where this conversion is achieved using Equations (F.10-11) of Appendix F.

When implementing the above:

- the values of y_1^* and y_2^* are determined for the number of gauges, N, in the network being considered;
- the values of γ_1 and γ_2 are determined for the specific network being considered which, for this project, means as functions of the number of gauges, area of coverage and the average SAAR for the network.

For the present project, the specific choices of $y_1=0$ and $y_2=4$ have been made.

6.4.4 Model fitting

The structure and parameters of a model relating the relative-location parameters γ_1 and γ_2 to descriptors of a network of raingauges has been investigated in a two-stage procedure. The two stages are as follows:

- Evaluation of raw estimates of the relative-location parameters derived from fitting the spatial dependence model to a wide range of cases based on subsets of the project database. These subsets consist of networks within a given radius of selected target points, each containing a given number of raingauges.
- Exploration of the relationship of the raw estimates of the relative-location parameters to descriptors of networks for each case.

Both of these stages are described in detail in Appendix F. For the first stage, the raw estimates of the relative-location parameters are determined by minimising an Anderson-Darling measure of fit between the modelled distribution and the empirical distribution of the network maxima extracted for each specific subset of the project database. This estimation procedure was chosen because the fitting criterion is directly related to a distance between the modelled and empirical distribution functions. The Anderson-Darling measure of fit is used as a test statistic in the context of testing the goodness of fit of an assumed distribution. In comparisons with other candidate test-statistics, such as the Kolmogorov-Smirnov and Cramer-von Mises statistics, it has been judged best at detecting departure from the assumed mode in the tail of the distribution (Pearson & Hartley, 1972, Section 30.2). This seems a relevant finding for applications to extreme values. Estimation of parameters by minimising goodness of fit measures, known as minimum distance estimation, has been studied by, for example, Parr & Schucany (1980) and Boos (1982).

Raw estimates of the spatial dependence parameters were obtained for a large set of network configuration cases. In brief, these consisted of all combinations of the following factors, restricted by the ability to find networks with those characteristics (for example there are no networks with large numbers of gauges within small radii for any target location, but this availability varies radically between the hourly and daily durations):

- Durations: from 1 hour to 8 days (11 different durations)
- Radius: radii in kilometres 2,4,8,16,32,64,128,200,300,600,1300
- Number of gauges: 2,4,8,16,32,64,128,256,512,1024,2048
- Central locations: on regular grids covering the country, grid spacing 10km for radii up to 32km, spacing 20km for radii 64km and over.

For each of these target configurations, the following procedure was applied. The network of gauges available within the given radius of the central location was examined to establish firstly whether or not the required configuration could be met at all, but also to establish a range of acceptable sub-network areas that could reasonably be extracted from the records being used. The network area is always defined here in the way established by Dales & Reed and used by the FEH; that is, using Equation (6.2). The range of acceptable areas always has a ratio of 2 to 1 for the upper and lower bounds. The procedure mirrors the Dales & Reed approach of selecting random sub-networks of all the gauges available in a given year, subject to the area being within the specified range. As part of the information extracted from the estimation procedure, the average of all the areas of the selected sub-networks was used to define the 'area' associated with the estimated spatial dependence parameters. In addition to this descriptor, a number of others were carried forward as being potentially useful in the later stage of modelling, which attempts to relate the spatial dependence parameters to descriptors of the network from which the maximum is extracted. These were:

- the central location, and the average of the locations of gauges;
- the altitude of the central location, and the average of the altitudes of the gauges;
- the value of SAAR at the central location, and the average of the SAAR values for the gauges.

The exploration of the relation of the spatial dependence parameters to network descriptors is described in Appendix F (Section F.3). The model which was fitted and used within the FORGEX procedure (Section 6.5) is that each of the relative-location parameters γ_1 and γ_2 is determined by a relationship of the basic form

$$\gamma = a + b\ln AREA + c\frac{\ln N}{1+\frac{1}{2}\ln N} + d\frac{SAAR}{1000}. \tag{6.15}$$

In the above equation, $SAAR$ refers to the average SAAR across the gauges in the network of N gauges and $AREA$ is the area of the network based on the Dales & Reed definition (Equation (6.2)). Equation (6.15) includes a function of N (the number of gauges) which is modified from its usual simpler form of just having $\ln N$. This was included because values of the coefficient c are sometimes negative and using the simpler form would then mean that values of γ (dependence) would decrease indefinitely with an increasing number of gauges, which might be thought counter-intuitive. The modified term moderates the influence of the number of gauges in very extreme cases. In practice, including the modified term meant the model performed slightly better at the fitting stage.

This basic form is supplemented by restricting the values of γ_1 and γ_2 to lie between zero and one, and to have $\gamma_2 \le \gamma_1$. If the test that $\gamma_2 \le \gamma_1$ fails the value of γ_2 is reset to that of γ_1.

Values of the coefficients in the above relation are given in Tables 6.1 and 6.2. The same values are included with an increased precision in Appendix F.

Table 6.1 Model for the spatial dependence parameter γ_1

	Coefficients for γ_1			
Duration	a	b	c	d
1 hour	0.173	-0.013	-0.044	0.075
2 hours	0.228	-0.015	-0.042	0.035
4 hours	0.357	-0.026	-0.023	0.000
6 hours	0.412	-0.027	-0.023	-0.029
12 hours	0.540	-0.033	-0.048	-0.051
18 hours	0.615	-0.039	-0.039	-0.054
24 hours	0.599	-0.032	-0.070	-0.077
1 day	0.764	-0.052	0.027	-0.101
2 days	0.789	-0.055	0.035	-0.079
4 days	0.791	-0.061	0.060	-0.051
8 days	0.839	-0.058	0.027	-0.058

Table 6.2 Model for the spatial dependence parameter γ_2

	Coefficients for γ_2			
Duration	a	b	c	d
1 hour	0.055	-0.006	-0.076	0.109
2 hours	0.104	-0.007	-0.084	0.078
4 hours	0.223	-0.016	-0.082	0.040
6 hours	0.252	-0.016	-0.078	0.011
12 hours	0.308	-0.016	-0.100	-0.011
18 hours	0.390	-0.023	-0.115	-0.011
24 hours	0.333	-0.009	-0.139	-0.050
1 day	0.481	-0.030	-0.029	-0.063
2 days	0.522	-0.032	-0.038	-0.055
4 days	0.523	-0.040	-0.007	-0.029
8 days	0.544	-0.036	-0.041	-0.031

If the coefficient a is taken as a representative measure of strength of dependence, the results in Tables 6.1 and 6.2 show:

- increasing dependence with duration, and
- less dependence as measured by γ_2 at its typical return period than measured by γ_1 at its lower typical return period.

The coefficient b indicates that dependence decreases with increasing area, as expected. Surprisingly, the coefficient c (for the effect of the number of gauges) is, in most cases, negative. Intuitively, one might expect 'dependence' to increase as the number of gauges increases. However, the relevance of this point is not straightforward because of two considerations:

(i) the return period on the typical curve at which this dependence is measured also varies with the number of gauges;

(ii) the overall effect of the dependence model is not contained in the dependence measure γ itself, but rather in $(1-\gamma)\times \ln N$, which is the left-shift on the Gumbel scale from the 'typical' curve to the network maximum curve. –'Intuition' requires that this left-shift should increase with N.

Unfortunately, these considerations interact. The situation is summarised in Figure F.5. The variation of the coefficient d with duration indicates that, for shorter durations (up to around 4 or 6 hours) the dependence increases with SAAR, but for longer durations the dependence decreases with SAAR.

The model proposed here may be compared with the model of Dales & Reed (1989) by noting that the latter model is a special case of this model, with the parameters specified by

$$(1-\gamma_1)=\frac{\ln N_e}{\ln N}=a+b\ln AREA+c\ln N+d\ln D\,, \qquad (6.16)$$

and $\gamma_2=\gamma_1$. See Equation (6.1). Here N_e is the 'equivalent number of independent gauges' treated by Dales & Reed, and a, b, c and d are the regression coefficients in their model for the ratio of logarithms using the variables $AREA$, number of gauges (N) and duration (D).

6.5 Using the spatial dependence model within FORGEX

The part of the FORGEX procedure that makes use of the spatial dependence model is that which assigns plotting positions to network maxima. For this, all the gauges within a given distance of the central location, which have data for a given year, are found and the maximum of the standardised annual maxima is used as the network maximum for that year. For each year, the network maximum has associated with it values of the following quantities:

- the number of gauges available for the year;
- the area (as defined by the Dales and Reed procedure) associated with the actual set of gauges available for the year;
- the average SAAR for the region covered by the network.

These associated values are used to specify the spatial dependence parameters for the network using the model outlined in Section 6.4. This allows the dataset to be reduced to a set of pairs $\{x_{(i)}, s_i; i = 1, \cdots, n\}$, where n is the number of years, with at least one gauge in the network having a valid annual annual maximum, the values $\{x_{(i)}; i = 1, \cdots, n\}$ are the ordered set of network maxima and $\{s_i; i = 1, \cdots, n\}$ are the corresponding (re-ordered with the network maxima) values of the auxiliary information describing the spatial dependence for each year. Specifically, for the year j which supplies the *j*'th highest value of maximum standardised rainfall,

$$s_j = (\gamma_{1,j}, \gamma_{2,j}, N_j),$$

where $\gamma_{1,j}$ and $\gamma_{2,j}$ are the values of the relative location parameters determined for the year and N_j is the number of gauges in the network with valid annual maxima for that year.

The method for determining plotting positions used in the FEH procedures was found to have problems when used with the non-constant shift model for spatial dependence that is used here. This has, therefore, been revised. The revised procedure determines the plotting positions by using a modified joint likelihood function for all of the plotting positions jointly, whereas the FEH procedure used a separate likelihood function for each plotting position. Details of this are given in Appendix F (Section F.4). The joint likelihood function depends on both:

- the structure of the spatial dependence model as outlined in Section 6.3 and described in more detail in Appendix F;
- the relation of the relative-location parameters γ_1 and γ_2 to the network descriptors as outlined in Section 6.4 and described in more detail in Appendix F.

The role of the procedure for determining plotting positions is essentially similar to its role within the version of FORGEX used in the FEH. That is it provides a formal way of determining the results of the notional 'rightward shift' of the points relating to network maxima as described informally in Section 6.3 (Figure 6.2). One improvement over the FEH procedure is the accommodation of the revised model for the effects of varying spatial dependence. A consequence of this revised model is that the plotting positions for the highest rainfalls should be further to the right than would result from the FEH model, as described informally in Section 6.3 (Figure 6.1(d)).

6.6 Summary

This section has described the basis of the spatial dependence model used within this project, as part of an updated FORGEX procedure, to draw conclusions about rainfall at high return periods. Analyses have been shown which justify moving away from the simpler model used in the Flood Estimation Handbook. An improved model has been fitted which describes the extent of spatial dependence of rainfall for any given duration in terms of two parameters. This allows the relationship between the distributions of annual maximum rainfall at a single site, and for the maximum within a network of gauges, to be described in terms of the network area, the number of gauges and the typical rainfall for the network, as measured by an average SAAR value.

7. Development of a revised FORGEX methodology

7.1 Introduction

As outlined earlier, the Focused Rainfall Growth Curve Extension (FORGEX) method was introduced in Volume 2 of the FEH (*Faulkner, 1999*) to estimate a rainfall frequency curve from raingauge annual maxima for a single duration for any individual location ('focal point') in the UK. At each focal point, FORGEX was applied independently for a range of durations between 1 hour and 8 days and the output was used in the production of a depth-duration-frequency (DDF) model for that point. In the FEH, two types of data were used to fit the curves: standardised annual maxima from individual raingauges, and network maxima (or 'netmax' points) consisting of the series of observations obtained by taking the highest value (after standardisation) in each year from all gauges within a given radius of the focal point. Plotting positions, which represent the non-exceedance probability, of both annual maxima and netmax points were estimated using the Gringorten formula (see Section 2), in the latter case taking into account the degree of spatial dependence between the gauges.

The changes made within this project to two key elements of the FEH FORGEX procedure, i.e. the standardisation and the spatial dependence model, have already been described in some detail. This section demonstrates the influence that these changes have on rainfall frequency curves produced by the original FORGEX method, as well as examining the effect of the new annual maximum dataset compiled for this project. In addition, the FORGEX methodology itself has been revised to improve both the selection of data and the method of curve fitting. Of these revisions, the biggest departures from FEH practice are that the maxima from individual raingauges (known as 'pooled points' in FEH) are no longer used, and spatial smoothness is promoted by the introduction of additional networks, with reduced weights, up to a maximum radius of 300 km.

A set of raingauge sites with long hourly and daily records of annual maxima was selected for testing the revised FORGEX methodology. The relatively long data records allowed at-site RMED values to be reliably estimated. Originally this set comprised 34 sites, but it was later extended to include a further two sites to allow the effects of altitude to be studied. Subsequently, the set of test sites was extended to 71 to include as many sites close to large raised reservoirs as possible and to obtain a reasonable spatial cover. Full details of the test sites are given in Appendix G.

7.2 Summary of the FEH FORGEX procedure

In order to summarise the effects of the changes that this project has made to FORGEX, it is helpful to review the detail of the method as presented in the FEH. For each duration, the following steps were taken:

1. Identify all raingauges within 200 km of the focal point that have at least 10 annual maxima. Note that in the current project, this threshold has been reduced to nine for consistency with other parts of the analysis and to increase the number of eligible gauges.

2. Standardise the observations (in the FEH, this was achieved by dividing by the gauge median for the relevant duration).

3. For each gauge, rank the annual maxima and estimate their probabilities, and hence return periods, using the Gringorten plotting position formula,

$$F_i = (i\text{-}0.44) / (n\text{+}0.12) \,, \qquad (7.1)$$

   where F_i is the non-exceedance probability attributed to the i^{th} member of n annual maxima ranked in ascending order. The return period, T, in years is given by

$$T = 1/(1\text{-}F) \qquad (7.2)$$

   Table 7.1 shows return periods obtained using Equations 7.1 and 7.2 for the highest four values for record lengths of nine years (the lower limit adopted by the current project), 152 years (the longest record used) and two typical intermediate lengths.

4. Express the return period in terms of the Gumbel reduced variate, y, using

$$y = \text{-ln}(\text{-ln}(1\text{-}1/T)) \qquad (7.3)$$

   (see Table 7.2 for the relationship between return period and y).

5. Select hierarchical raingauge networks, each comprising all gauges within a radius of the focal point as follows:
   a) Select gauges in order of increasing distance from the focal point until there are at least 20 annual maxima with return period between 1.179 and 2 years (y between -0.6335 and 0.3665). This set of gauges is the first network.
   b) Bring in more gauges until there are at least 20 annual maxima with return period between 2 years and 4.44 years (y between 0.3665 and 1.3665). This extended set of gauges forms the second network.
   c) Continue to form larger networks for 'y-slices' of width 1.0 until all the gauges are included (i.e. the 200 km limit has been reached). Note that the pooled points can at most reach y-values of about 5.5 because of their limited record length.

6. For each of these networks, estimate the return period of the highest standardised annual maximum in each calendar year as follows:
   a) For each year, find the highest standardised rainfall from all the gauges in the network that have data for that year.

b) Rank this set of annual 'network maxima', then estimate their return periods after first estimating the number of equivalent independent gauges in each year. In the FEH the conversion from actual number of gauges to the number of equivalent independent gauges is based on the mean intergauge separation for the gauges with data in each year.

7. If necessary, define intermediate networks with radii between that of the last network containing at least 20 pooled points (from step 5) and 200 km, so that no networks have a *y*-span greater than 1.

8. Construct the rainfall growth curve as a connected sequence of linear segments. Each segment is fitted to the pooled points from the corresponding network and the network maxima with *y*-values within the span of the segment from the corresponding network and any lower networks. The fitting procedure minimises a combination of the sum of the squares of the vertical distances from the points to the segments (pooled points and network maxima being given equal weight) and the difference between gradients of adjacent segments. The final results of the FEH FORGEX procedure consist of the fitted segments up to the y-value of the third highest network maximum.

A fitted FORGEX curve is illustrated in Figure 7.1, using an example from central England for the one-day duration. Note that the pooled points are indicated by dots and the network maxima by their network number.

Table 7.1 Return period (*T*) in years estimated using the Gringorten formula for the highest four values in an *n*-year record

n	T_n	T_{n-1}	T_{n-2}	T_{n-3}
9	16.3	5.8	3.6	2.6
25	44.9	16.1	9.8	7.1
50	89.5	32.1	19.6	14.1
152	271.6	97.5	59.4	42.7

Table 7.2 Relationship between return period (*T*) and Gumbel reduced variate, *y*

Return period, *T* (years)	Gumbel reduced variate, *y*
2	0.367
10	2.250
50	3.902
100	4.600
500	6.214
1,000	6.907
5,000	8.517
10,000	9.210
50,000	10.820
100,000	11.513

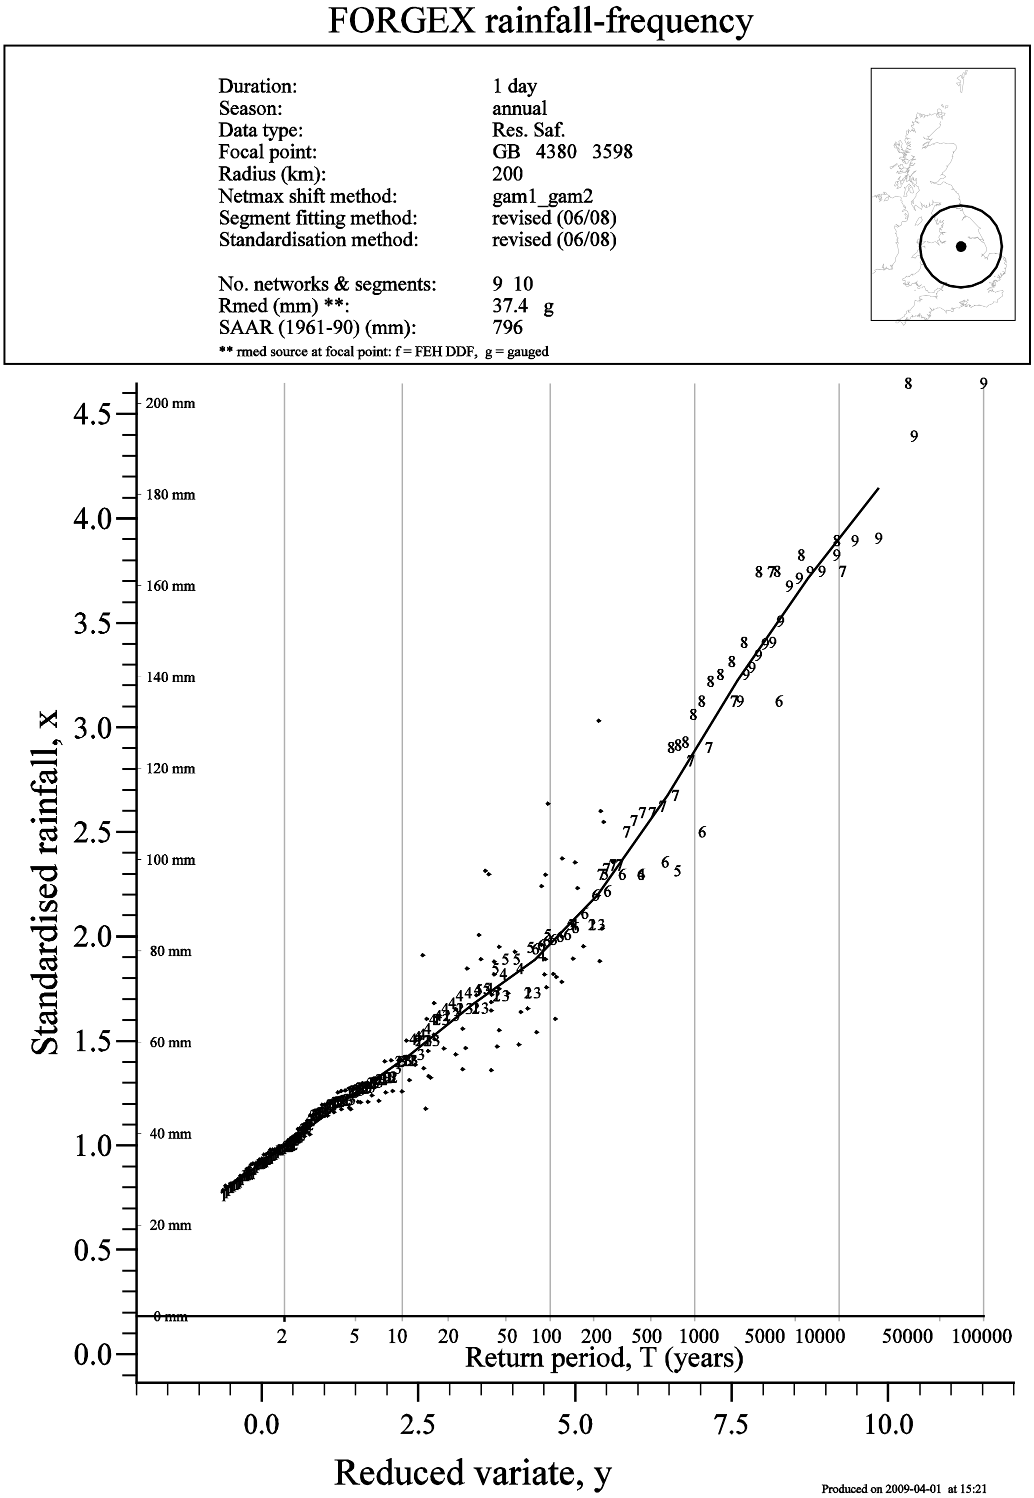

Figure 7.1 Example of FEH FORGEX rainfall growth curve for a focal point in central England (Ogston Reservoir)

7.3 Corrections to the FEH method and their effect

During this project, three errors were discovered in the programming of the FEH FORGEX procedures. The first two concern the implementation of the FEH spatial dependence model, and the third concerns the segment fitting routine.

- The geometric mean intergauge distance was used when computing the effective number of gauges in a network, but according to published descriptions (for example, *Dales & Reed, 1989*) the arithmetic mean should have been used.
- In FEH Vol. 2, Equation (8.3) states a rule for computing the shift in y for a network maximum. It includes the variable D, which is defined in day-equivalents, to be computed for hourly durations as the number of hours divided by 18 for durations shorter than 15 hours. The program that was used had divided by 24 for all durations.
- There was an error in the way that network maxima from the final network were allocated to segments as part of the fitting procedure.

To some extent the first two errors cancelled each other, although their net effect was to increase the estimated rainfall for a given return period. The effect of the third was small, and in some cases undetectable by eye.

An example of the typical effects of correcting these errors is shown in Figure 7.2.

7.4 Effect of the new dataset on rainfall frequency curves

As described in Section 3, a new dataset of annual maxima for the UK for 11 key durations from 1 hour to 8 days was assembled for this project. For the shorter durations, this dataset provides considerably longer data series and better spatial coverage than were available to the FEH team. Specifically, the number of hourly annual maximum records available to the current study is more than double that used in the FEH analysis. The most obvious effect of applying the corrected FEH FORGEX procedure to the extended dataset is that the upper return period limit of rainfall frequency curves derived from hourly annual maxima is increased, typically to over 1,000 years in the most densely gauged parts of the UK. In addition, in most cases the effect of the new dataset is to steepen the rainfall frequency curves for the 1-hour duration, particularly in Scotland. For the daily durations, the relative increase in record length is more modest, and hence the differences between the frequency curves derived from the FEH and updated datasets are relatively small. Examples for 1-hour and 1-day durations for three locations, Dingwall (northern Scotland), Ogston Reservoir (central England) and St Mawgan (south-west England) are shown in Figures 7.3(a) to 7.3(f).

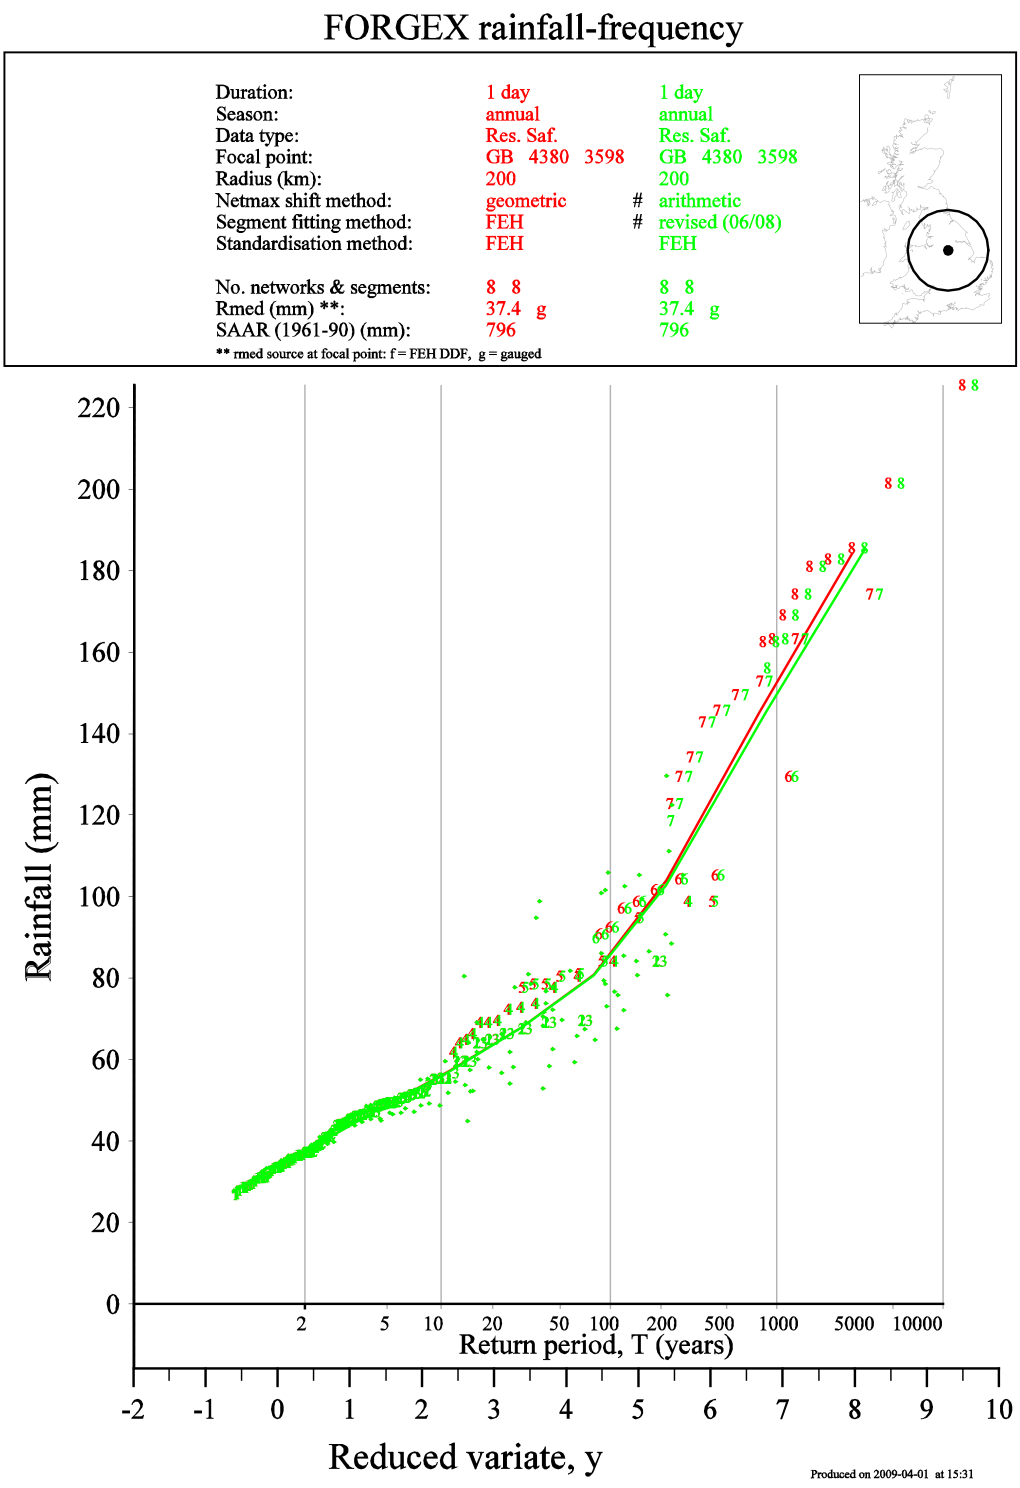

Figure 7.2 Example of the typical effects of correcting the errors in the FEH FORGEX method (the corrected curve is shown in green)

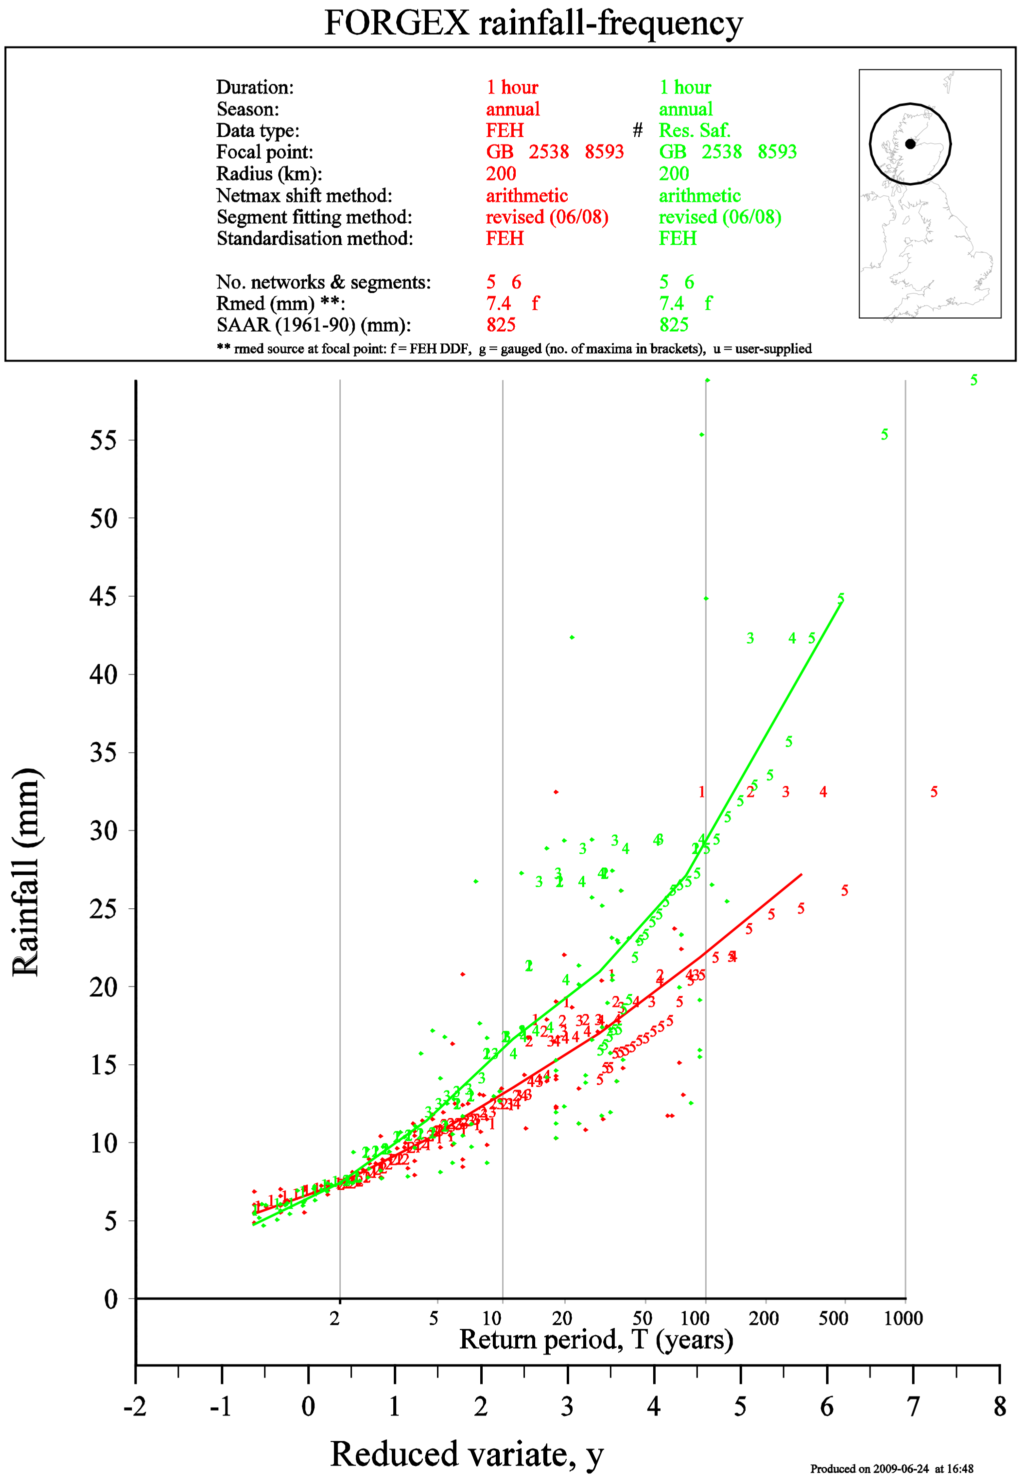

Figure 7.3(a) Example of the effect of the new dataset on rainfall frequency curves derived from the corrected FEH FORGEX method for northern Scotland – 1 hour duration (the curve for the new dataset is shown in green)

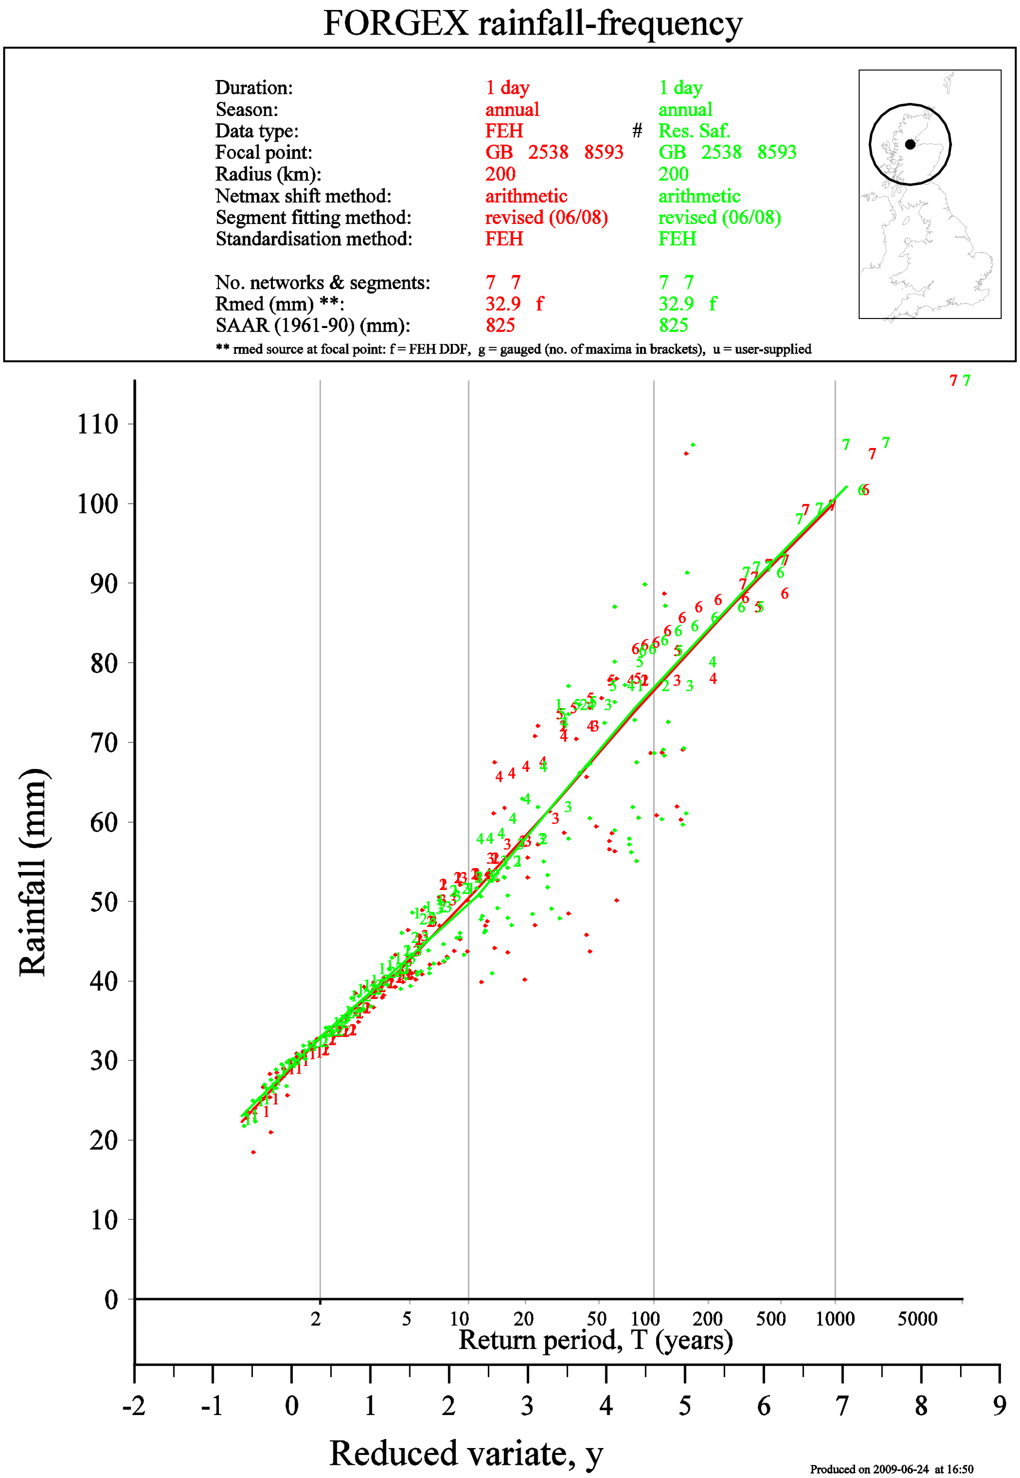

Figure 7.3(b) Example of the effect of the new dataset on rainfall frequency curves derived from the corrected FEH FORGEX method for northern Scotland – 1 day duration (the curve for the new dataset is shown in green)

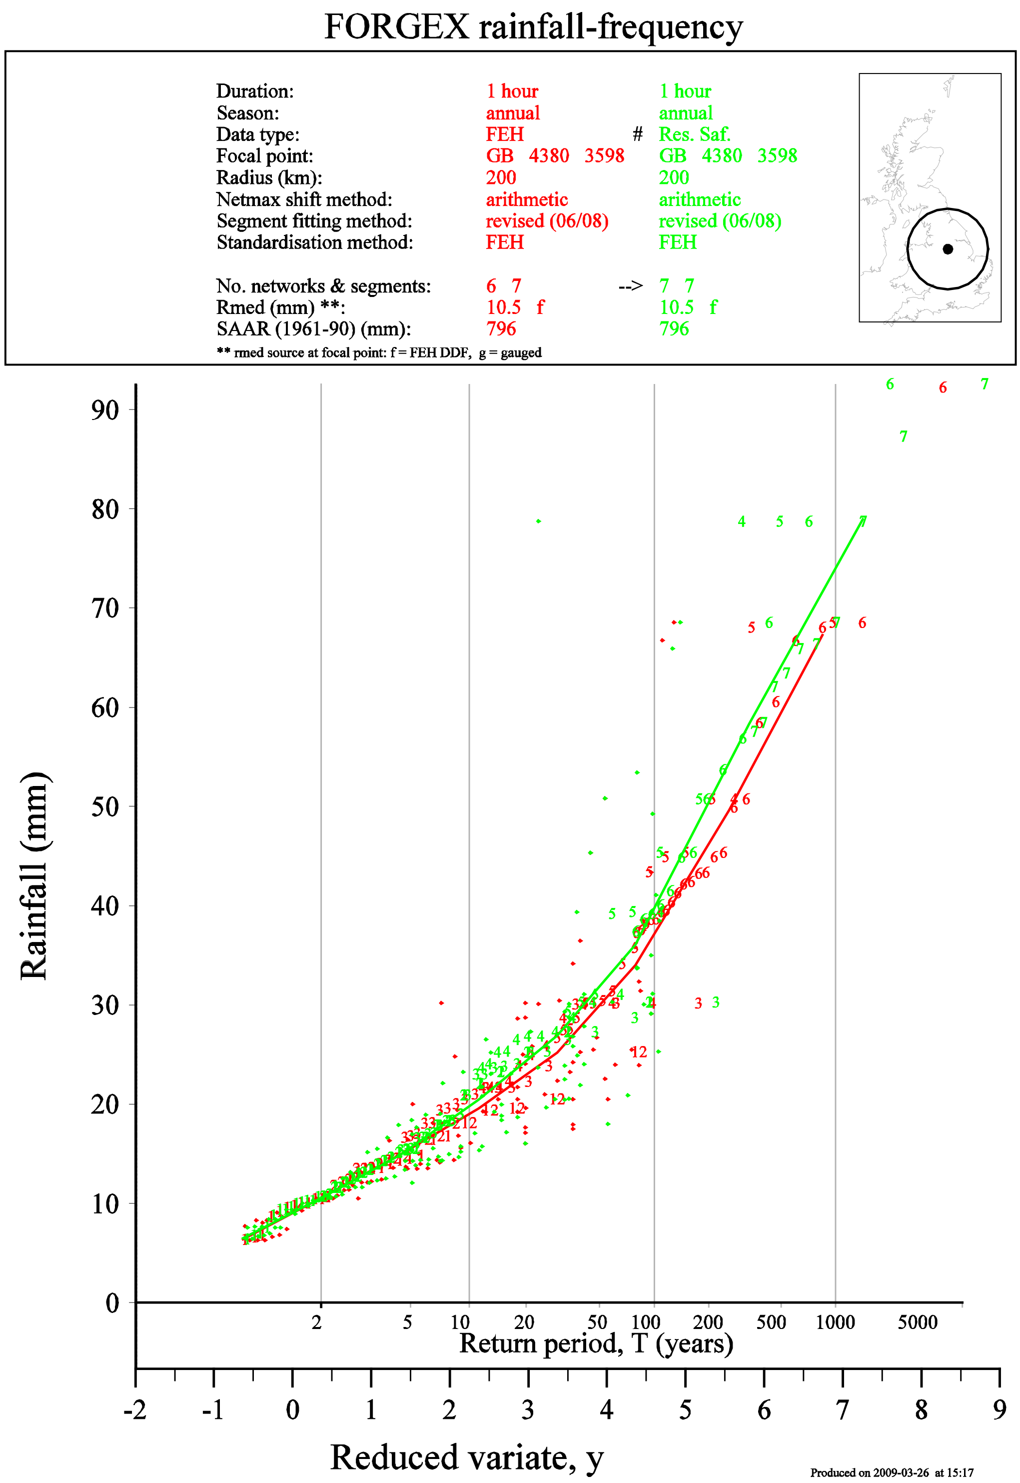

Figure 7.3(c) Example of the effect of the new dataset on rainfall frequency curves derived from the corrected FEH FORGEX method for central England – 1 hour duration (the curve for the new dataset is shown in green)

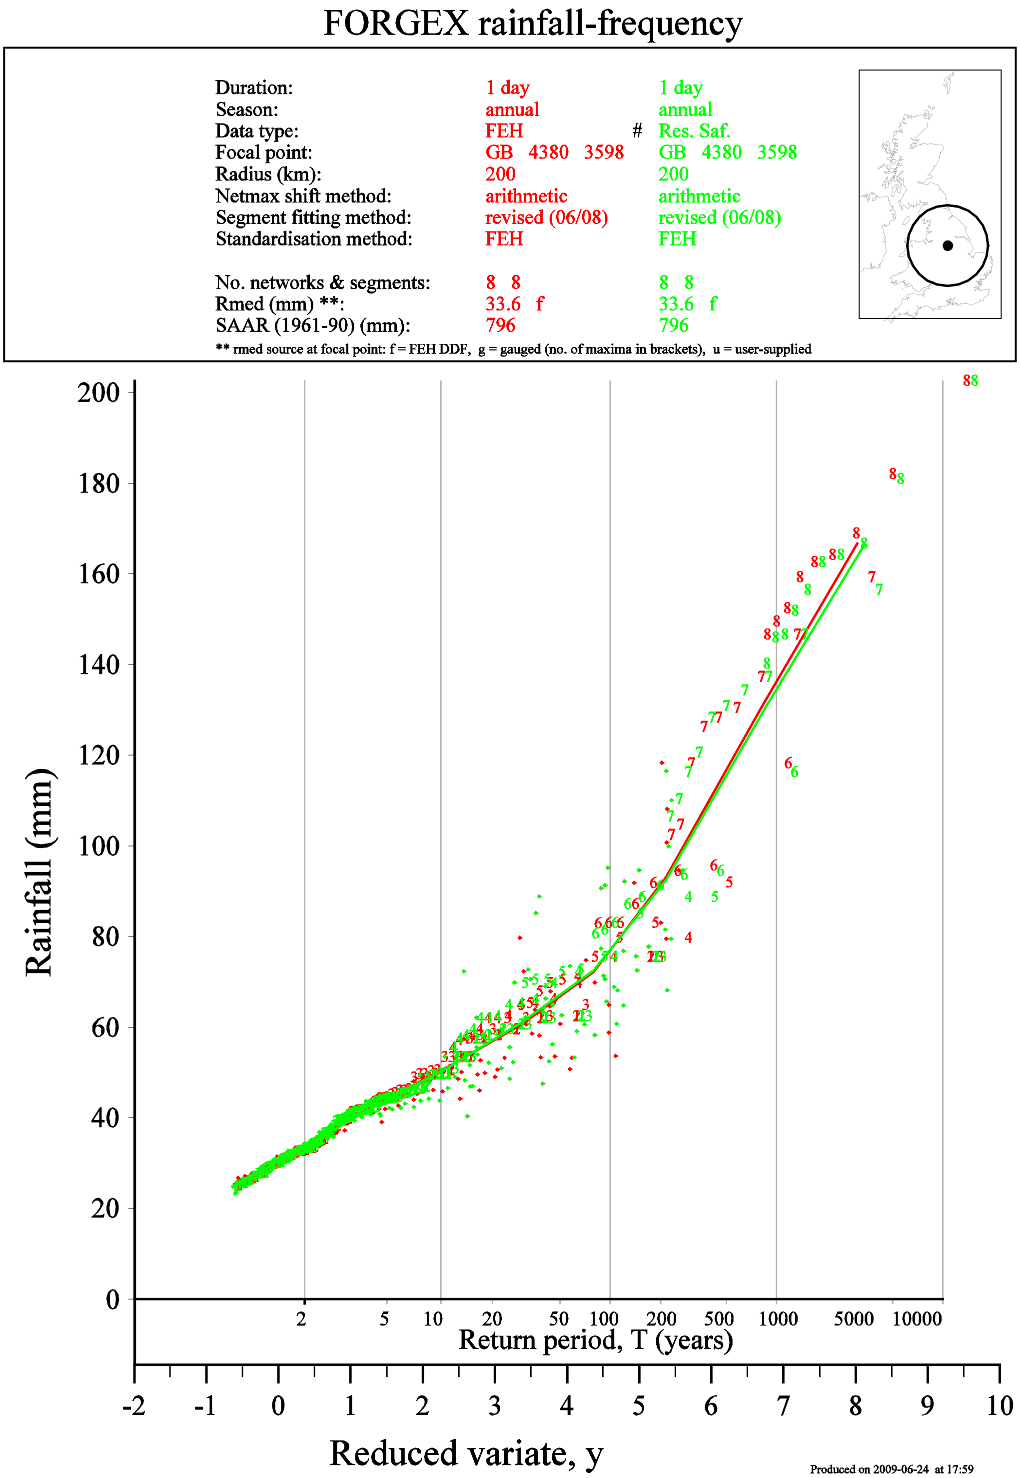

Figure 7.3(d) Example of the effect of the new dataset on rainfall frequency curves derived from the corrected FEH FORGEX method for central England – 1 day duration (the curve for the new dataset is shown in green)

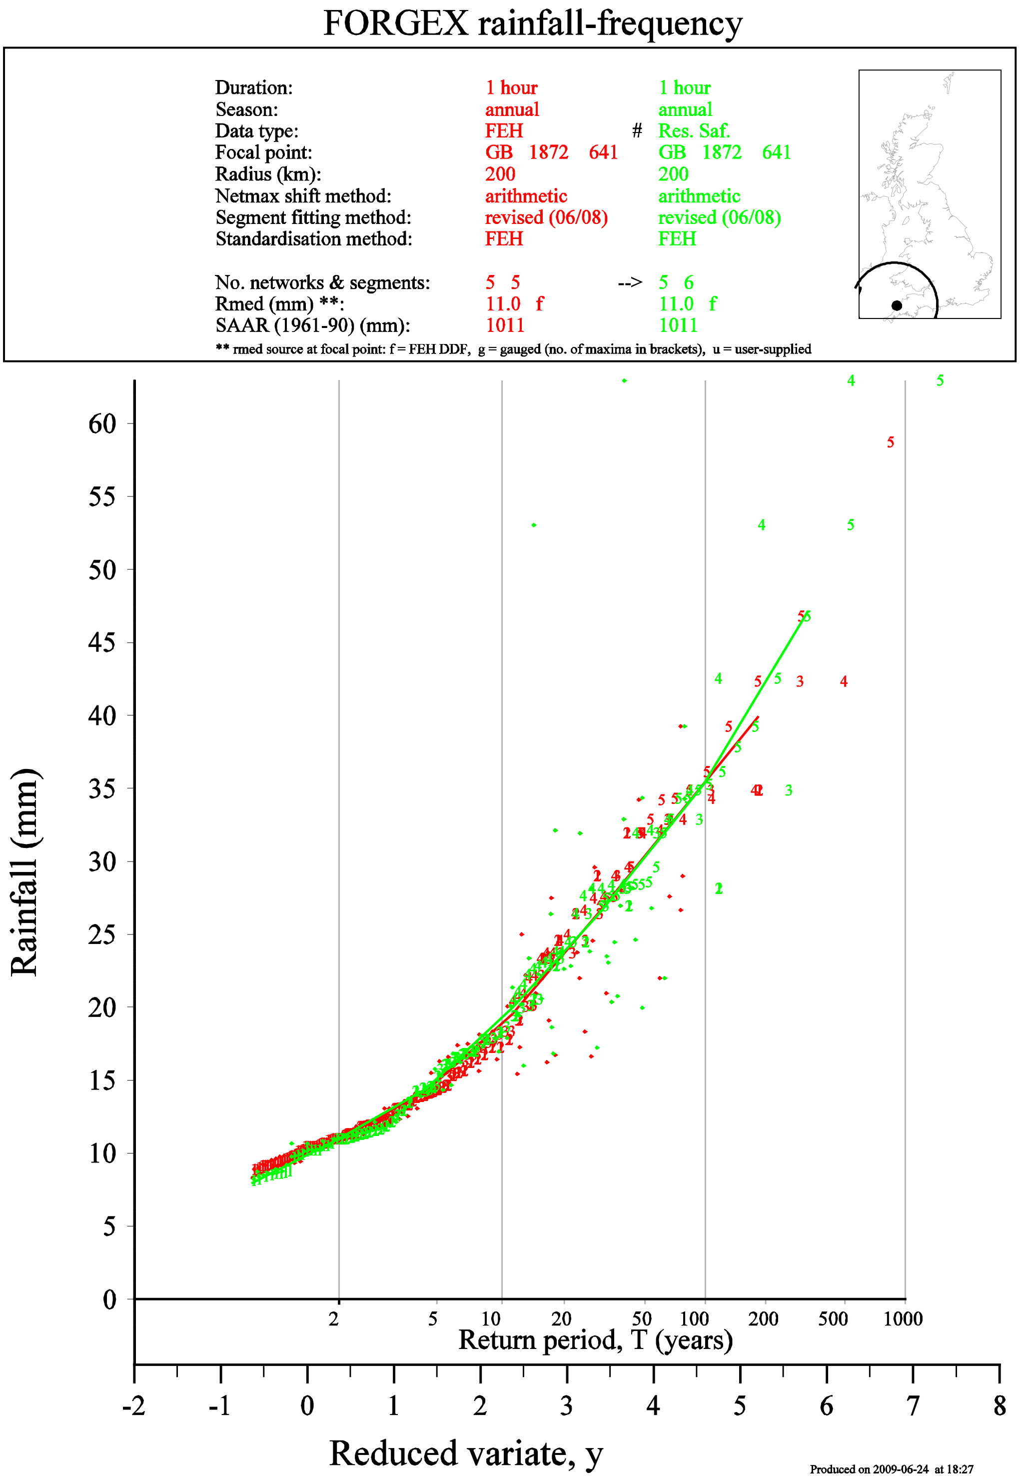

Figure 7.3(e) Example of the effect of the new dataset on rainfall frequency curves derived from the corrected FEH FORGEX method for south-west England – 1 hour duration (the curve for the new dataset is shown in green)

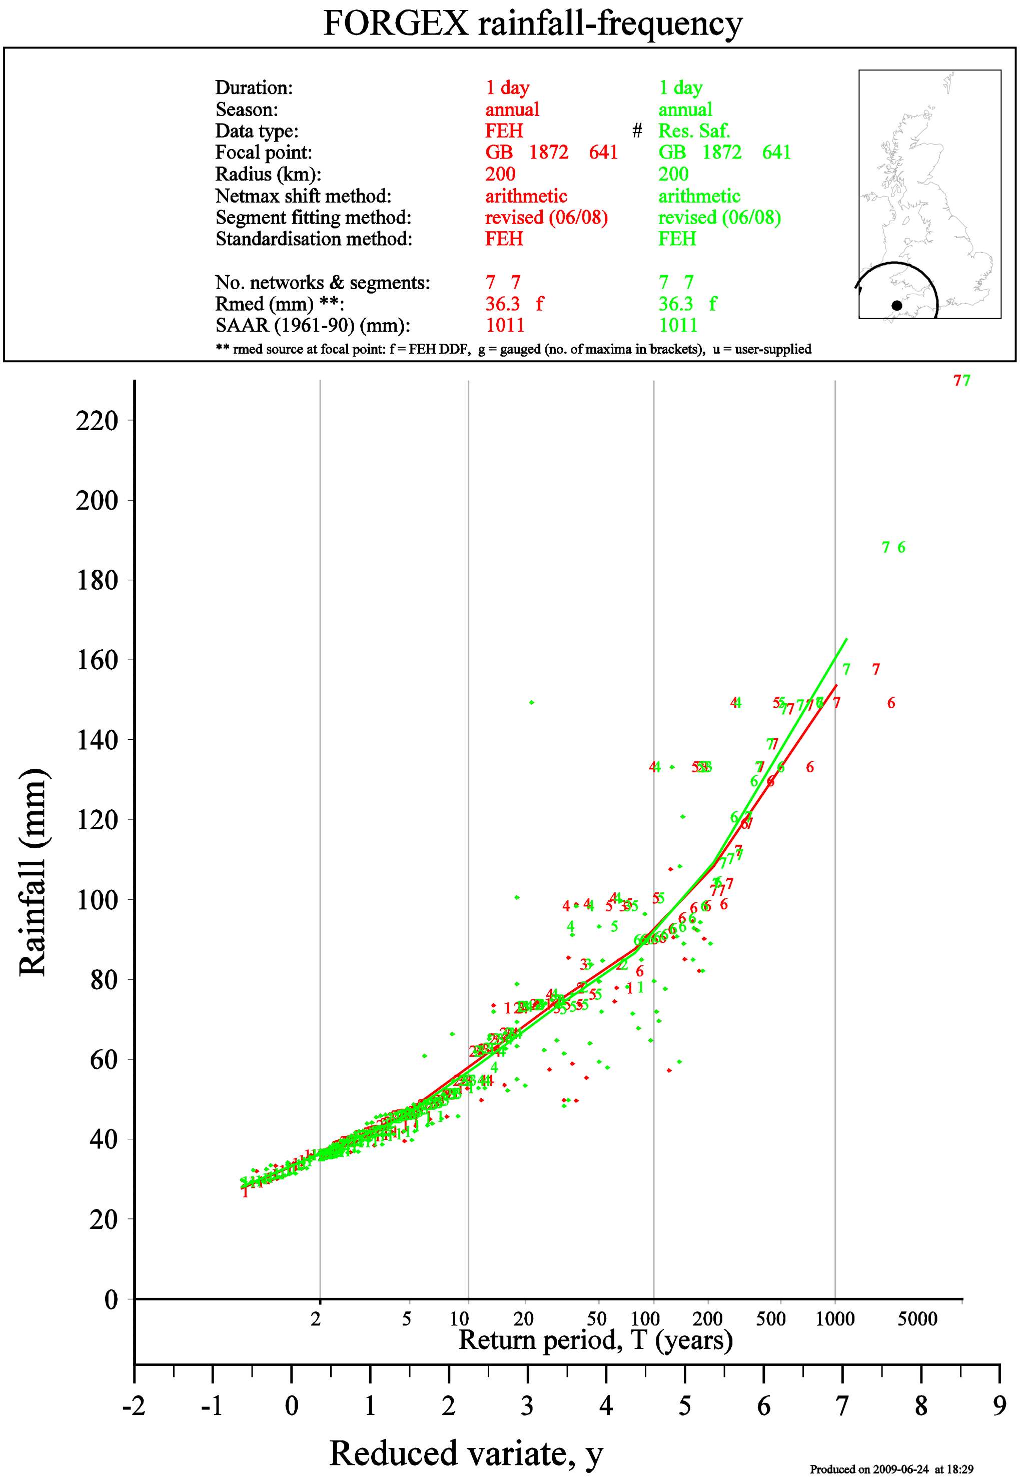

Figure 7.3(f) Example of the effect of the new dataset on rainfall frequency curves derived from the corrected FEH FORGEX method for south-west England – 1 day duration (the curve for the new dataset is shown in green)

7.5 Effect of the revised standardisation on rainfall frequency curves

The effect of the new standardisation is illustrated in Figures 7.4 and 7.5. Given the form of the revised standardisation, the largest changes compared with the FEH's simpler form would be expected to arise in cases where the SAAR for the target location is different from the values of SAAR for other gauges in the region contributing to the FORGEX procedure. Therefore, for high SAAR locations, there is a tendency to reduce rainfall estimates for a given return period. This is because the contributions to the standardised growth curves from sites with lower SAAR values are made comparatively less steep than with the original standardisation. Similarly, for low SAAR locations there is a tendency for estimates to be increased where annual maxima from gauges in high SAAR locations play an important part in the FORGEX procedure. This is illustrated by the 1-day FORGEX curves for two contrasting locations, Honister Pass in the Lake District (high SAAR) and Jesmond Dene near Newcastle (low SAAR) (Figures 7.4 and 7.5). Note that despite their separation of 115 km, the land area, and hence the raingauge network, covered by the 200 km radius circles for these two sites, is very similar.

7.6 Effect of the revised model of spatial dependence on rainfall frequency curves

For the 1-hour duration, the effect of applying the revised model of spatial dependence described in Section 6 together with the revised standardisation is almost universally to move rainfall frequency curves to the right, as illustrated by Figure 7.6 for the same focal point in central England as before (Ogston Reservoir). The shift is less marked for focal points in western Scotland, and in Shetland there is a very slight leftward shift in the 1-hour rainfall frequency curve. The rightward shift is also less marked for focal points in Northern Ireland, as illustrated by Figure 7.7. In the case of the 1-day duration, for all the focal points considered, the revised frequency curves plotted to the right of their FEH counterparts.

7.7 Revisions to the FORGEX methodology

This sub-section presents the revisions that have been made to the FORGEX methodology in this project, with regard to curve fitting and data selection. These comprise:

- Frequency/growth curves are fitted to network maxima only (no use is made of pooled points).
- New rules for the definition of network radii have been introduced.
- Additional networks are used, up to a radius of 300 km.
- New rules for the selection of network maxima have been introduced in the segment fitting procedure.
- Network maxima are now weighted.

7.7.1 Fitting to network maxima only

This change was introduced as a result of unrealistic rainfall frequency curves being produced by the FEH method at a location near the Kent coast (NGR 6324, 1661)). Figure 7.8 shows the curve produced by applying the FEH FORGEX procedure (note that this has used the new data, standardisation procedure and spatial dependence model). The curve has been lifted to such an extent by the cloud of pooled points at around 20-120 years return period, that it subsequently falls in the transition to the upper part of the curve, which is determined solely by network maxima. Investigation showed that the majority of these points were from a single event (the September 1973 West Stourmouth storm – see Appendix B), which was recorded at numerous raingauges; these have been highlighted on the plot. The plot indicates that the return period of this event when assessed using network maxima, which have taken into account spatial dependence between the gauges, varies from about 90 to 50,000 years depending on the radius of the network. Fitting to the network maxima alone produces a smoother curve (Figure 7.9). Note that this has taken into account that the West Stourmouth event could truly have a return period of about 90 years in this locality (the network 1 point), but it equally allows for the possibility of the event being rarer when assessed in terms of progressively larger networks. The revised approach is one whose result is not biased by the number of gauges that happen to be operating in the vicinity of a major storm.

To test the effects of fitting to network maxima only, comparative plots were produced at each of the 35 test locations listed in Appendix G for each of the durations 1 hour, 1 day and 8 days. In the majority of cases, it was found that excluding pooled points made little difference (Figure 7.10 is typical). Where differences did occur, they were in the mid-curve region, but none were as pronounced as the Kent example. Figure 7.11 shows the location (in south-west England) where the change made the second biggest difference, and Figure 7.12 shows the case (in the London area) where the change caused the biggest increase in the curve. Note that in these diagrams the red line has been fitted to all the data and the green line has been fitted only to the network maxima.

It is therefore considered that this change will be beneficial in situations similar to the Kent example (locations that are close to the largest events on the database, where these events have been recorded at numerous gauges), and in other circumstances it will make only a small difference.

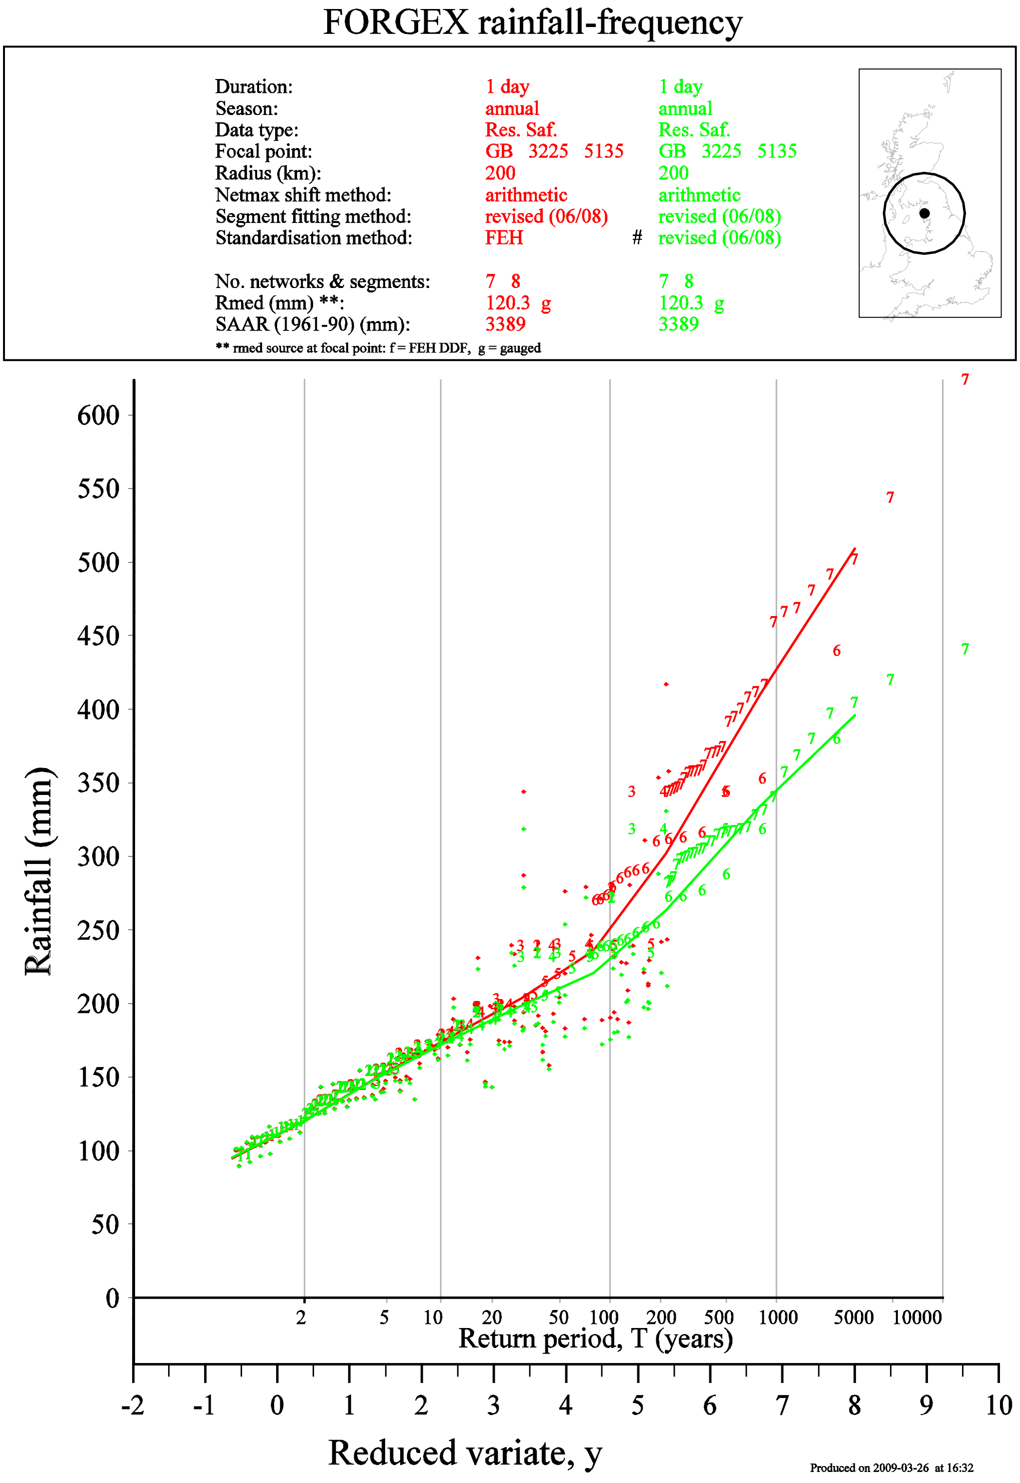

Figure 7.4 Example of the effect of the revised standardisation on the 1-day rainfall frequency curve a high SAAR site, Honister Pass in north-west England (the revised curve is shown in green)

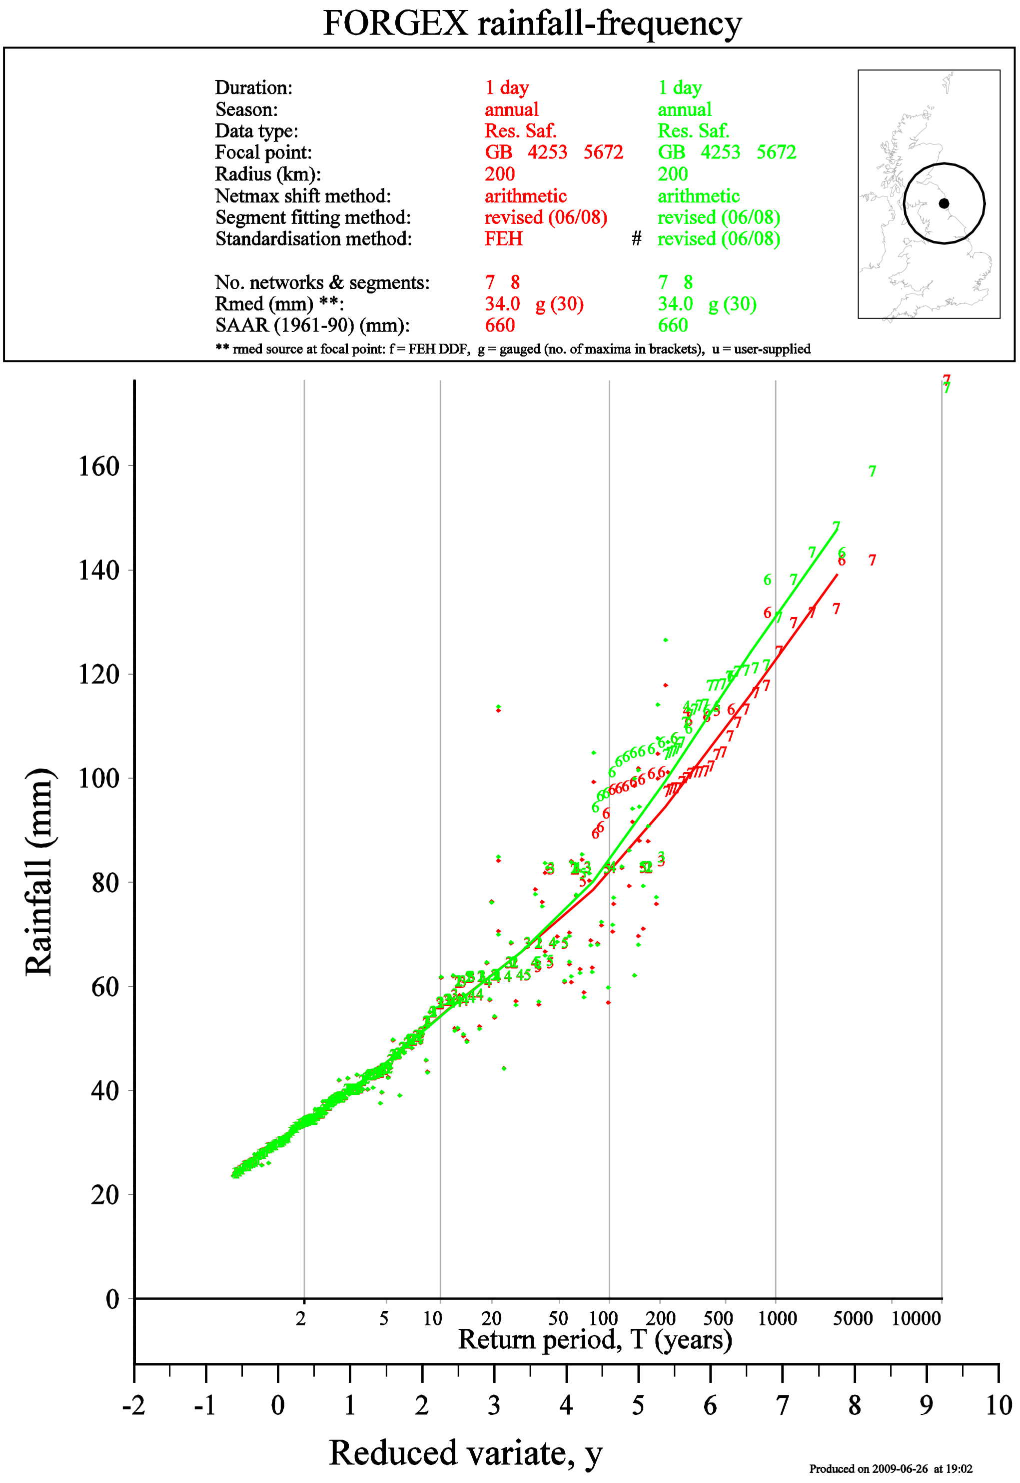

Figure 7.5 Example of the effect of the revised standardisation on the 1-day rainfall frequency curve a low SAAR site, Jesmond Dene in north-east England (the revised curve is shown in green)

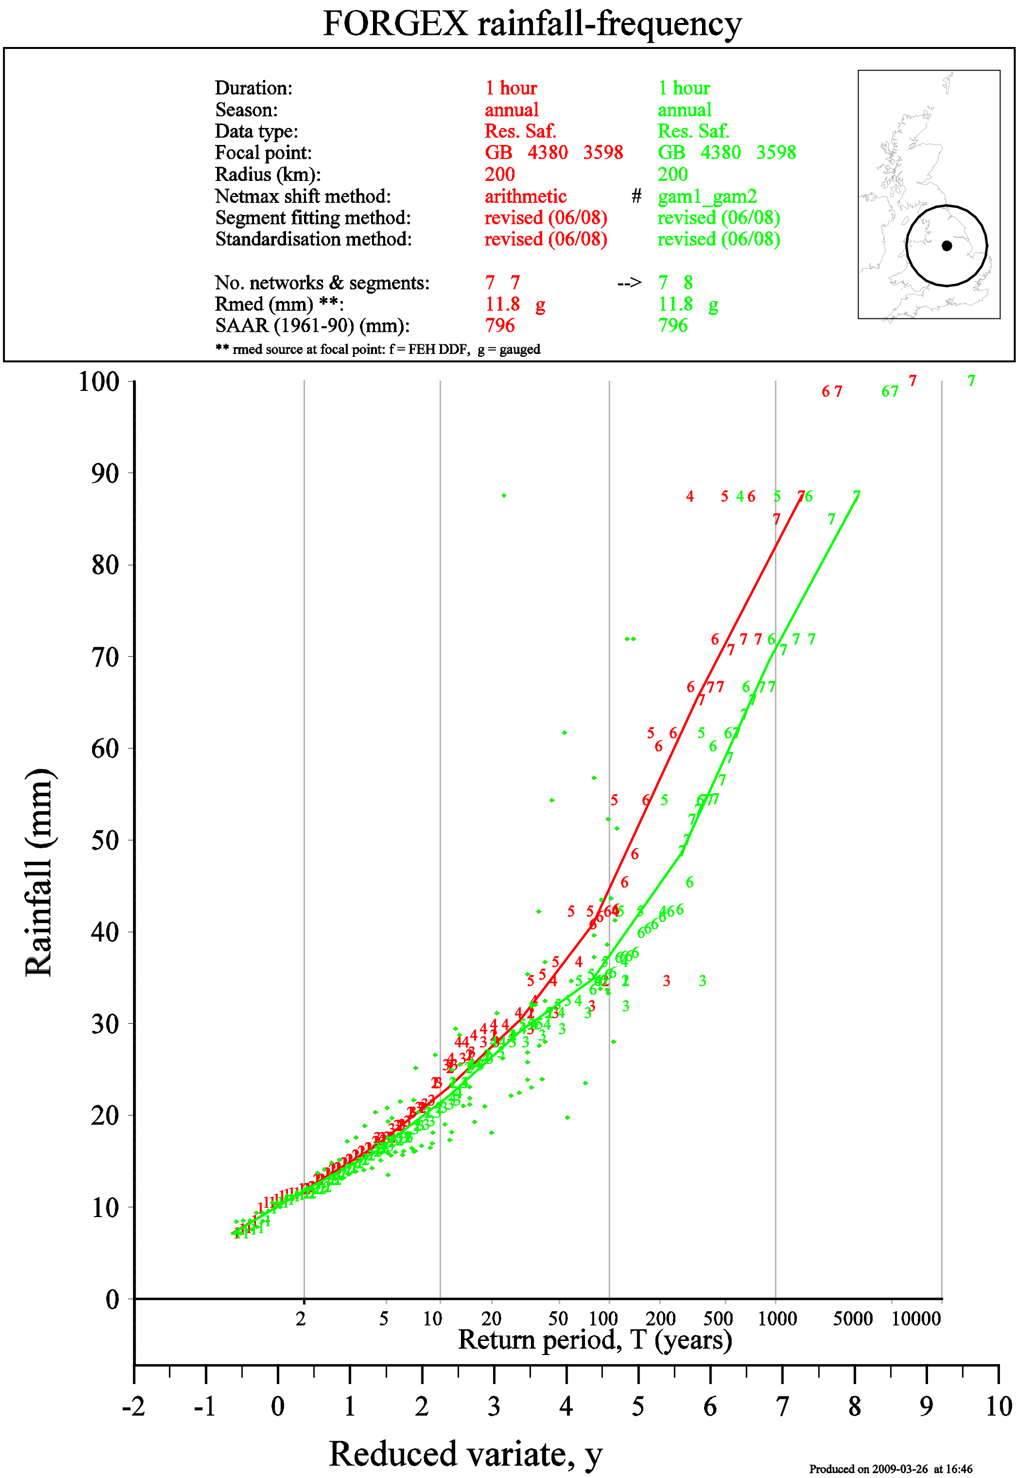

Figure 7.6 Example of the effect of the revised spatial dependence model and revised standardisation on a 1-hour rainfall frequency curve for a site in central England (the revised curve is shown in green)

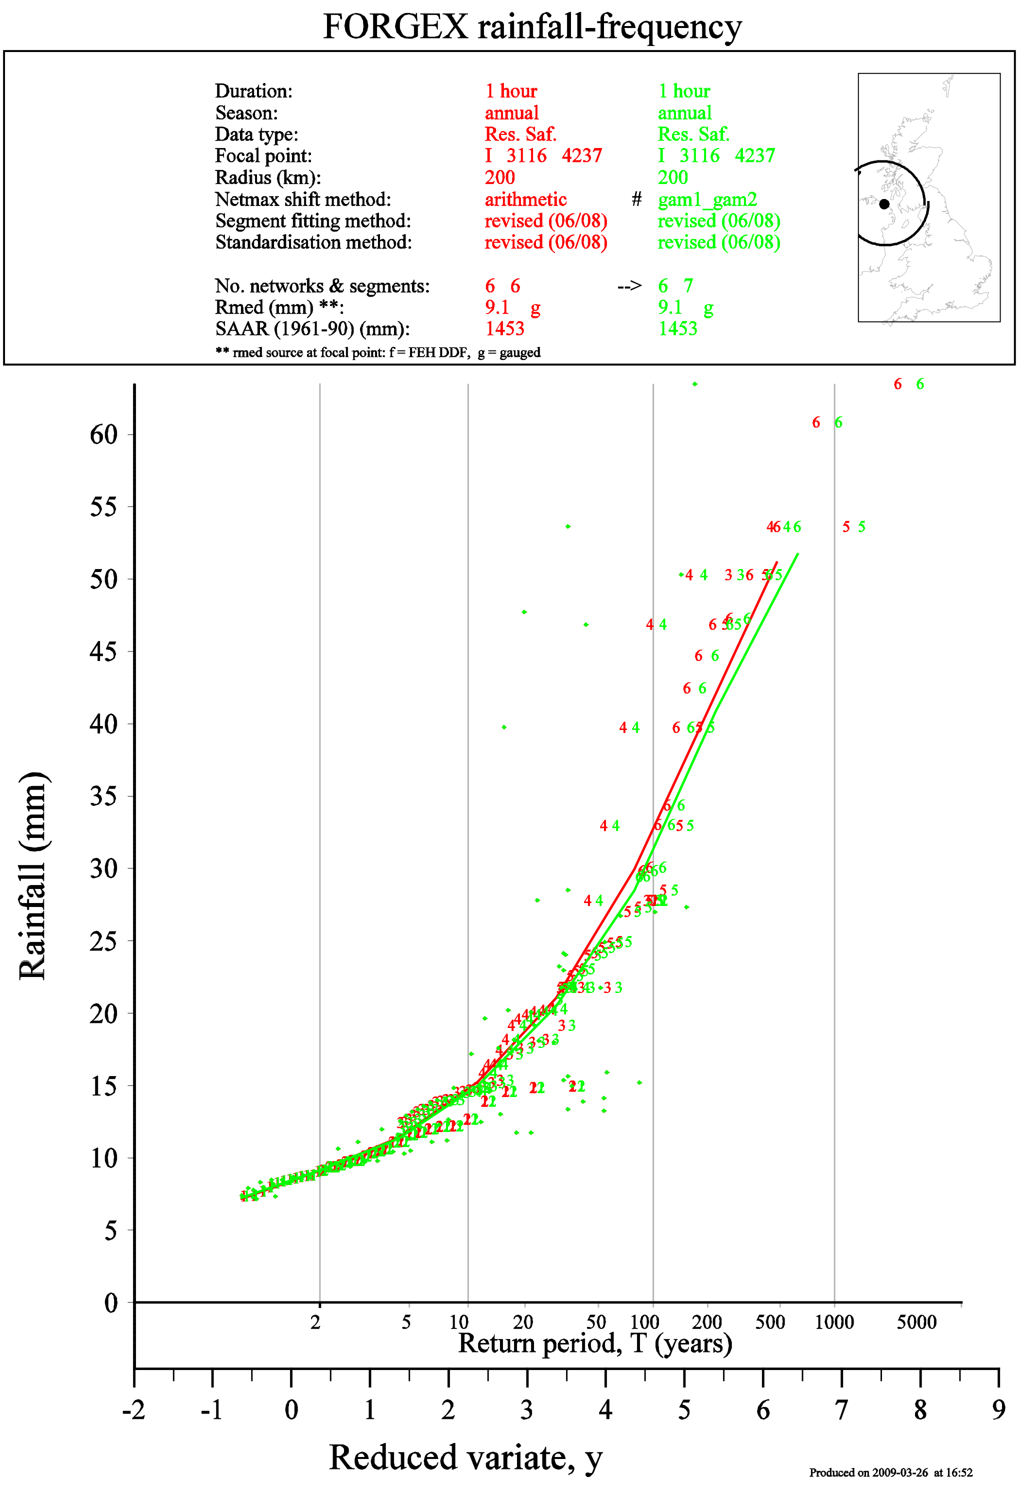

Figure 7.7 Example of the effect of the revised spatial dependence model and revised standardisation on a 1-hour rainfall frequency curve for a site in Northern Ireland (the revised curve is shown in green)

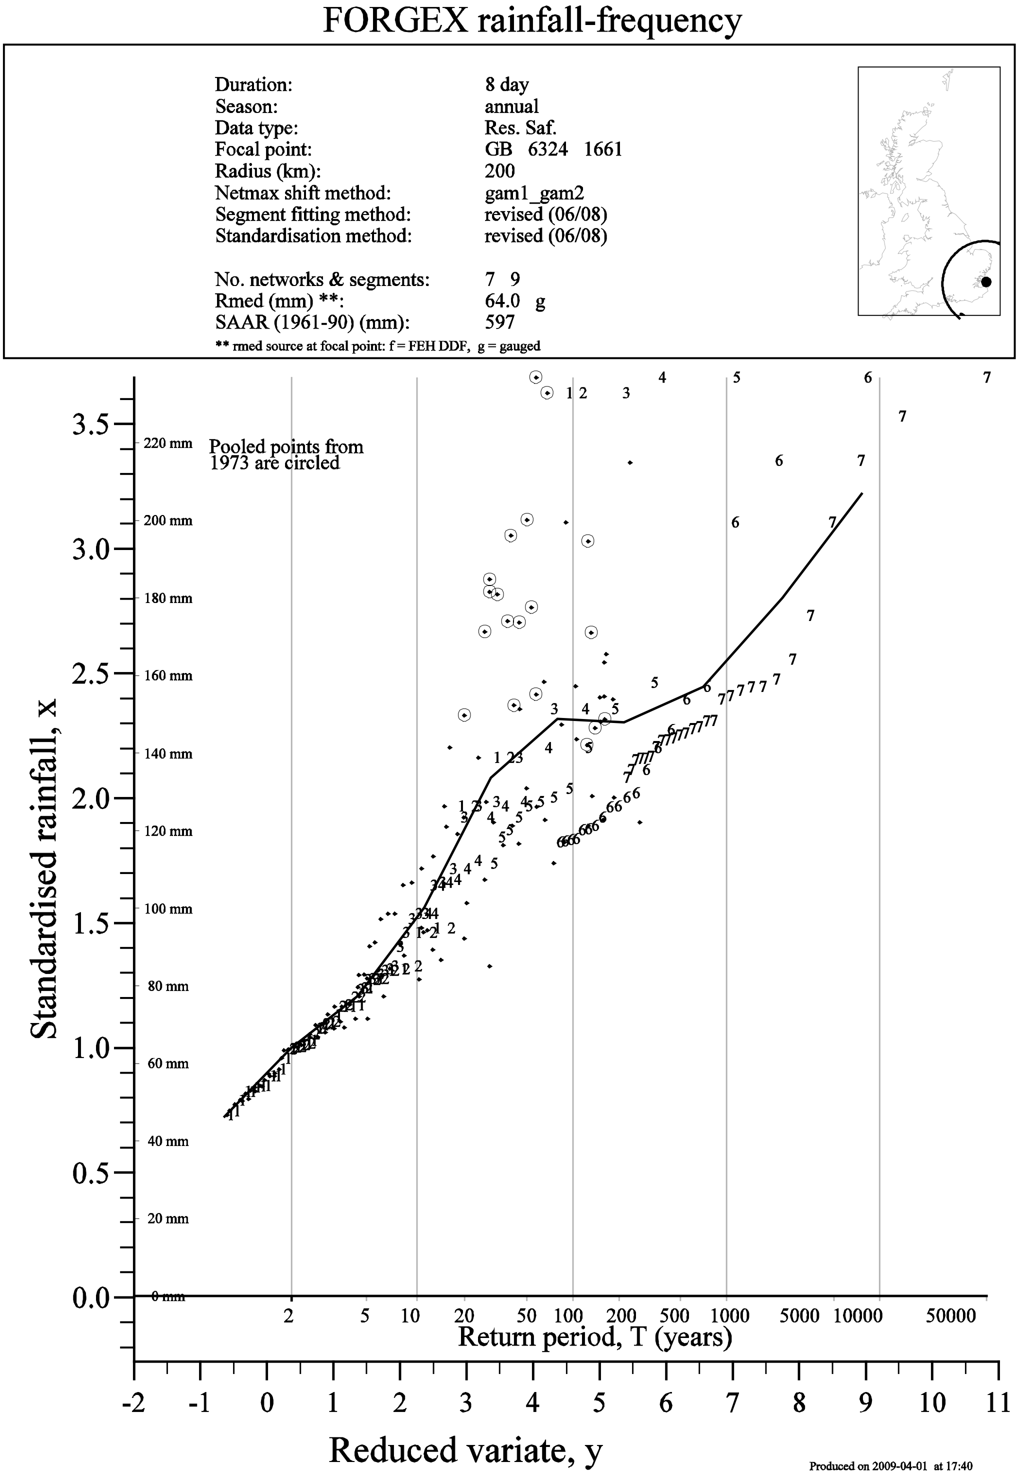

Figure 7.8 8-day rainfall frequency curve for a location close to the 1973 West Stourmouth event, produced using the FEH network definition and segment fitting method. Note the effect of the same event being recorded by many gauges (circled points).

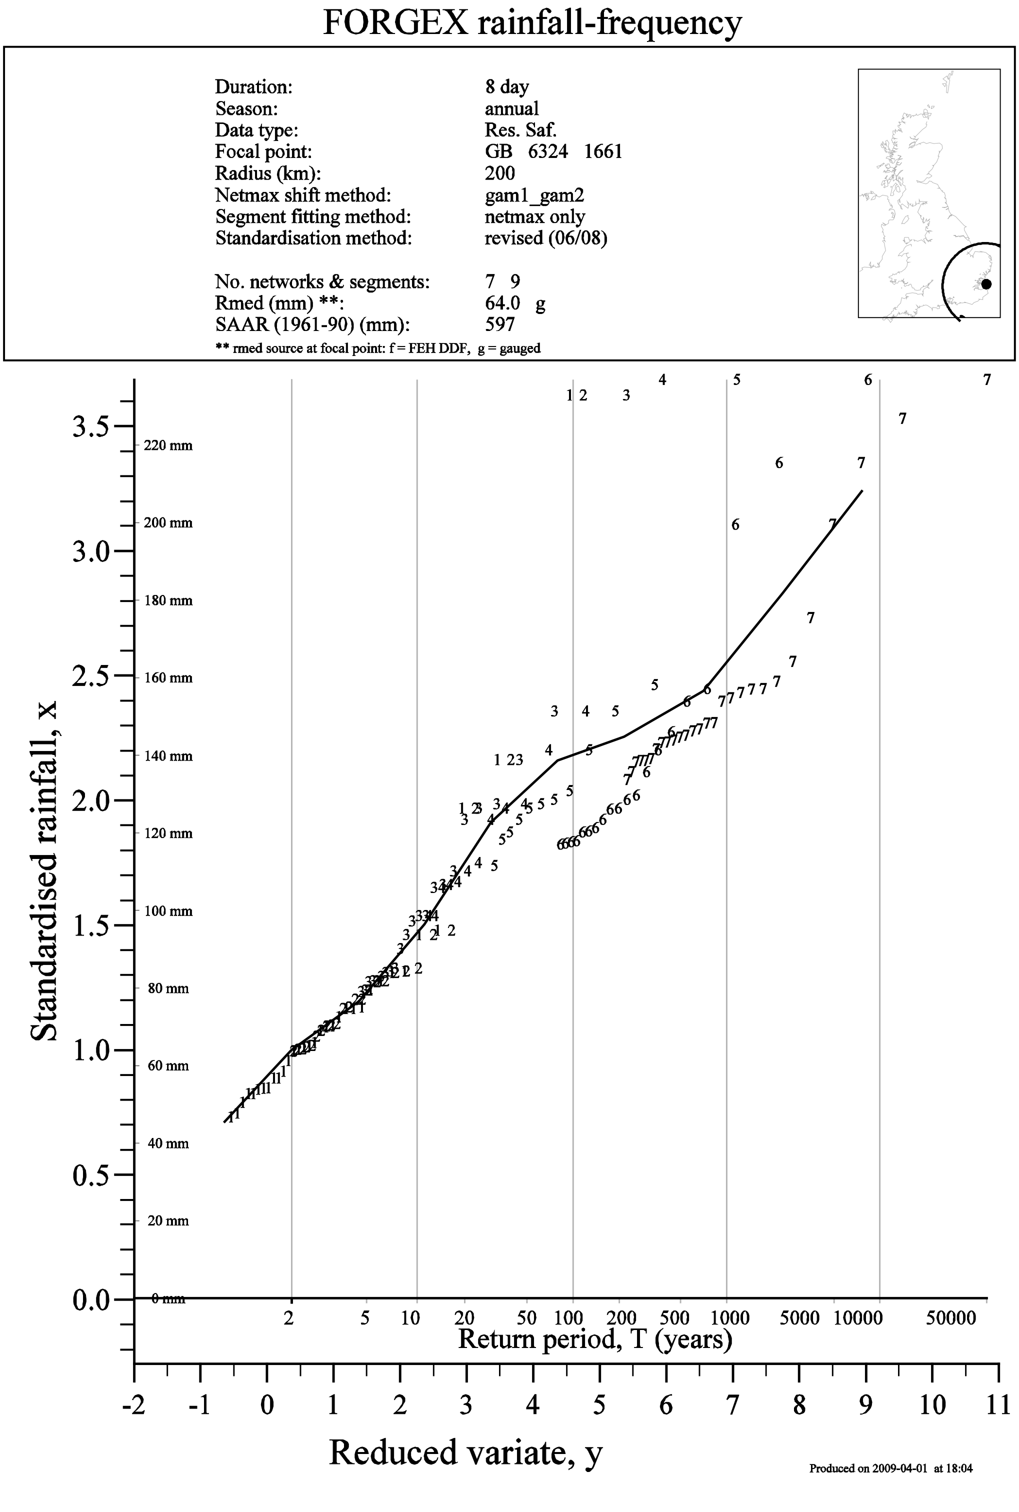

Figure 7.9 8-day rainfall frequency curve for a location close to the 1973 West Stourmouth event, fitted to network maxima only. Compare with Figure 7.8.

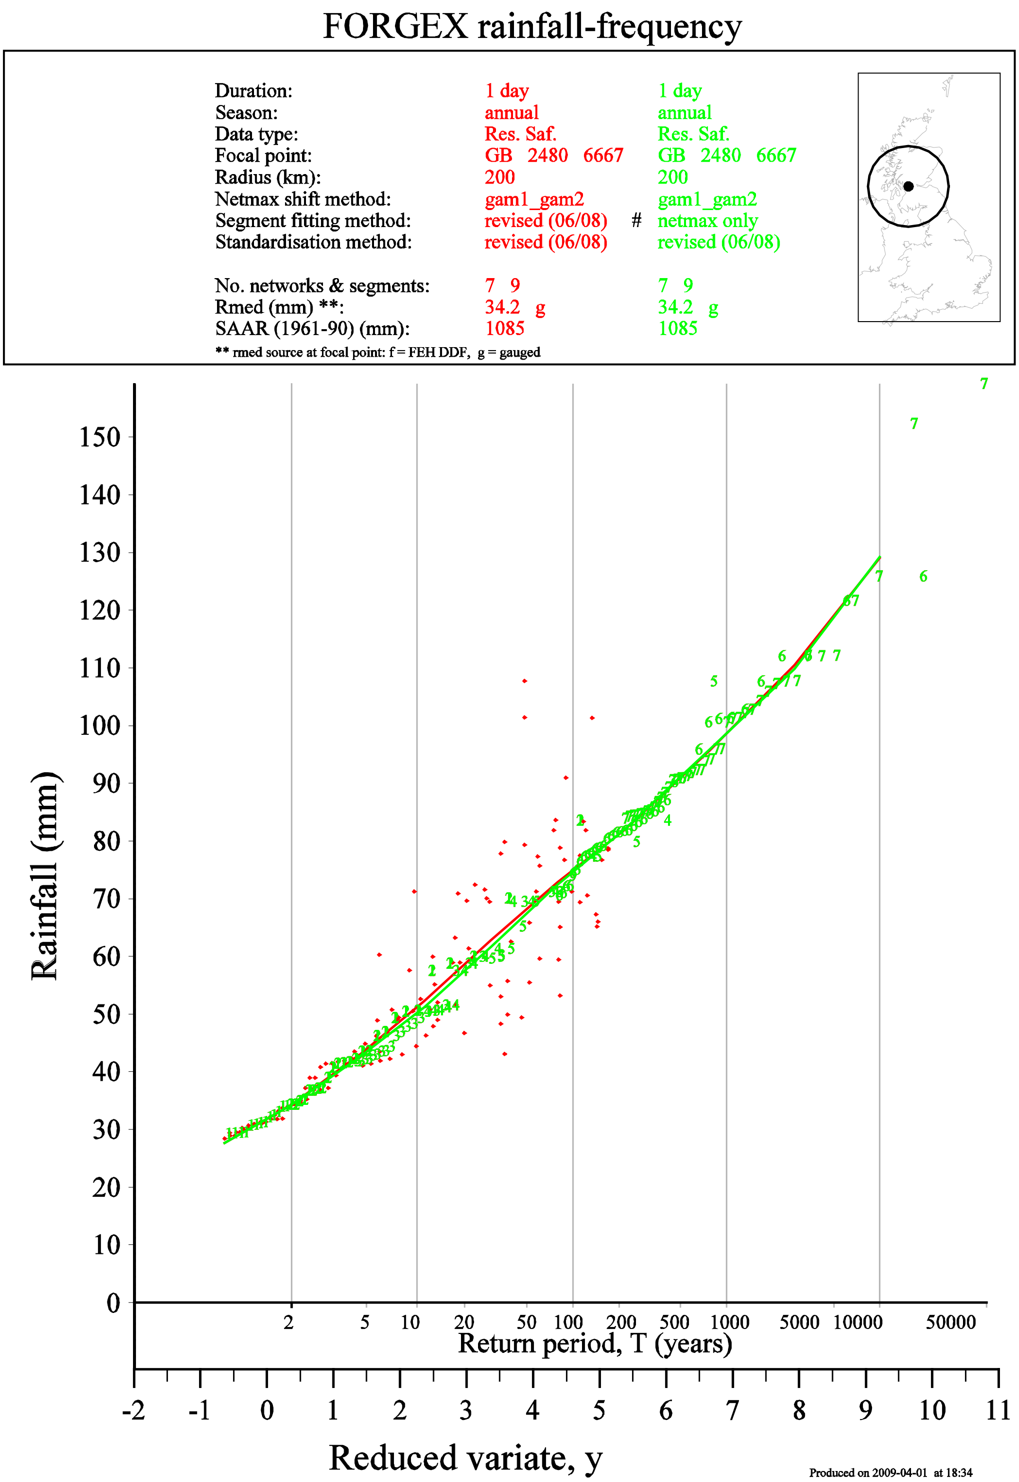

Figure 7.10 Typical example of the effect of excluding pooled points from rainfall frequency curve fitting (the green curve is fitted to network maxima only)

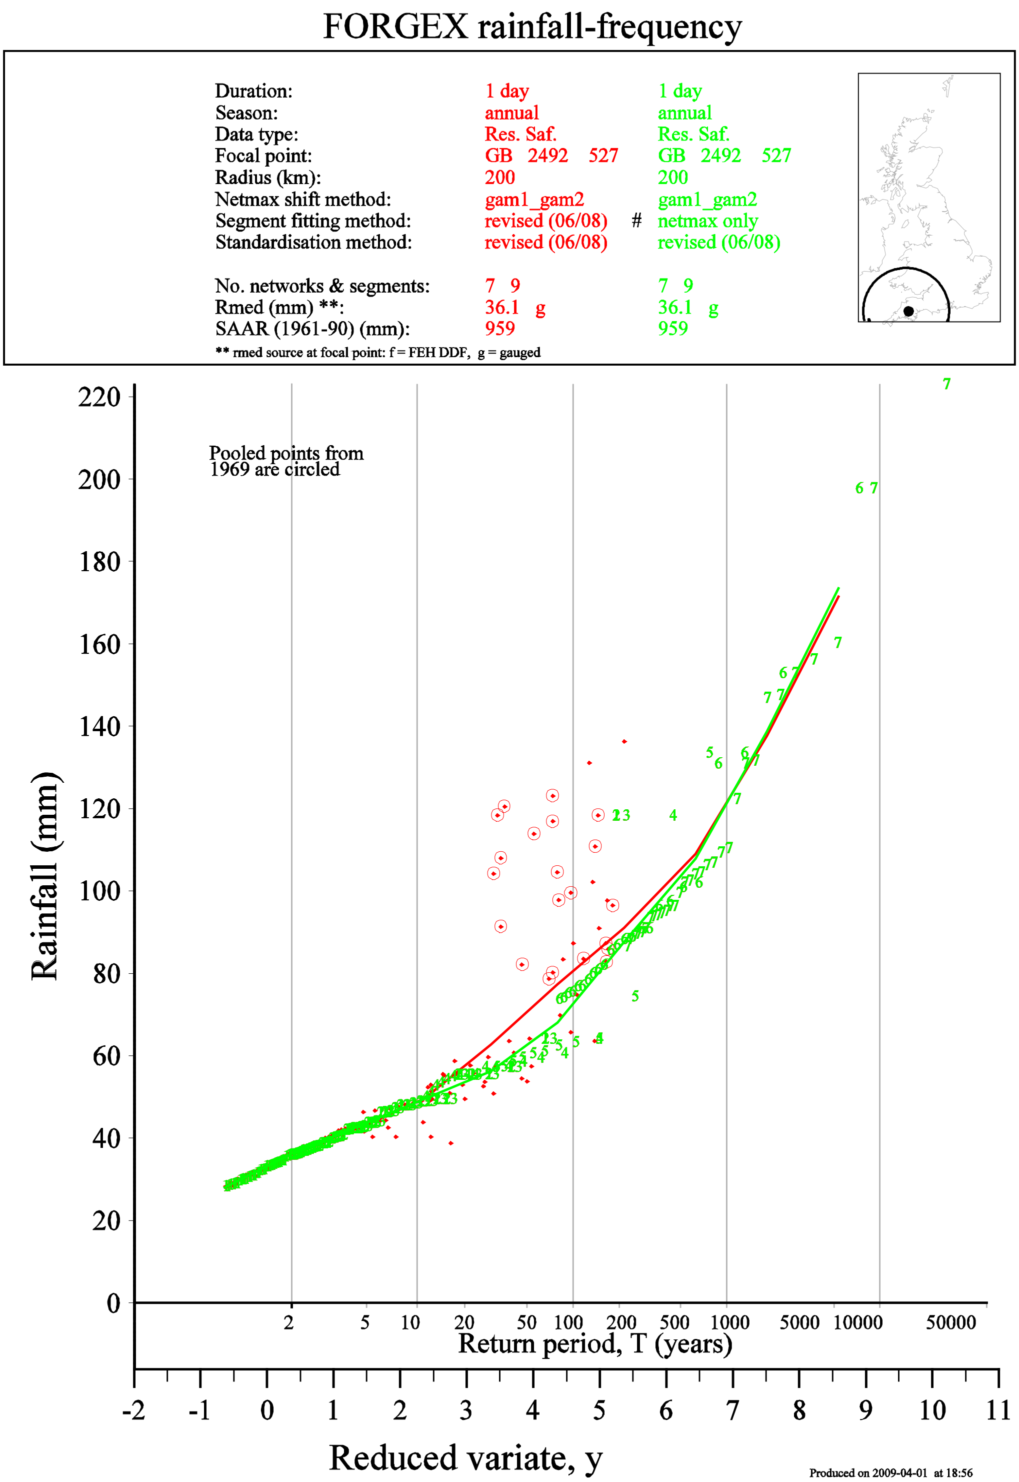

Figure 7.11 The change in the 1-day rainfall frequency curve for the location where excluding pooled points made the second biggest difference (the green curve is fitted to network maxima only)

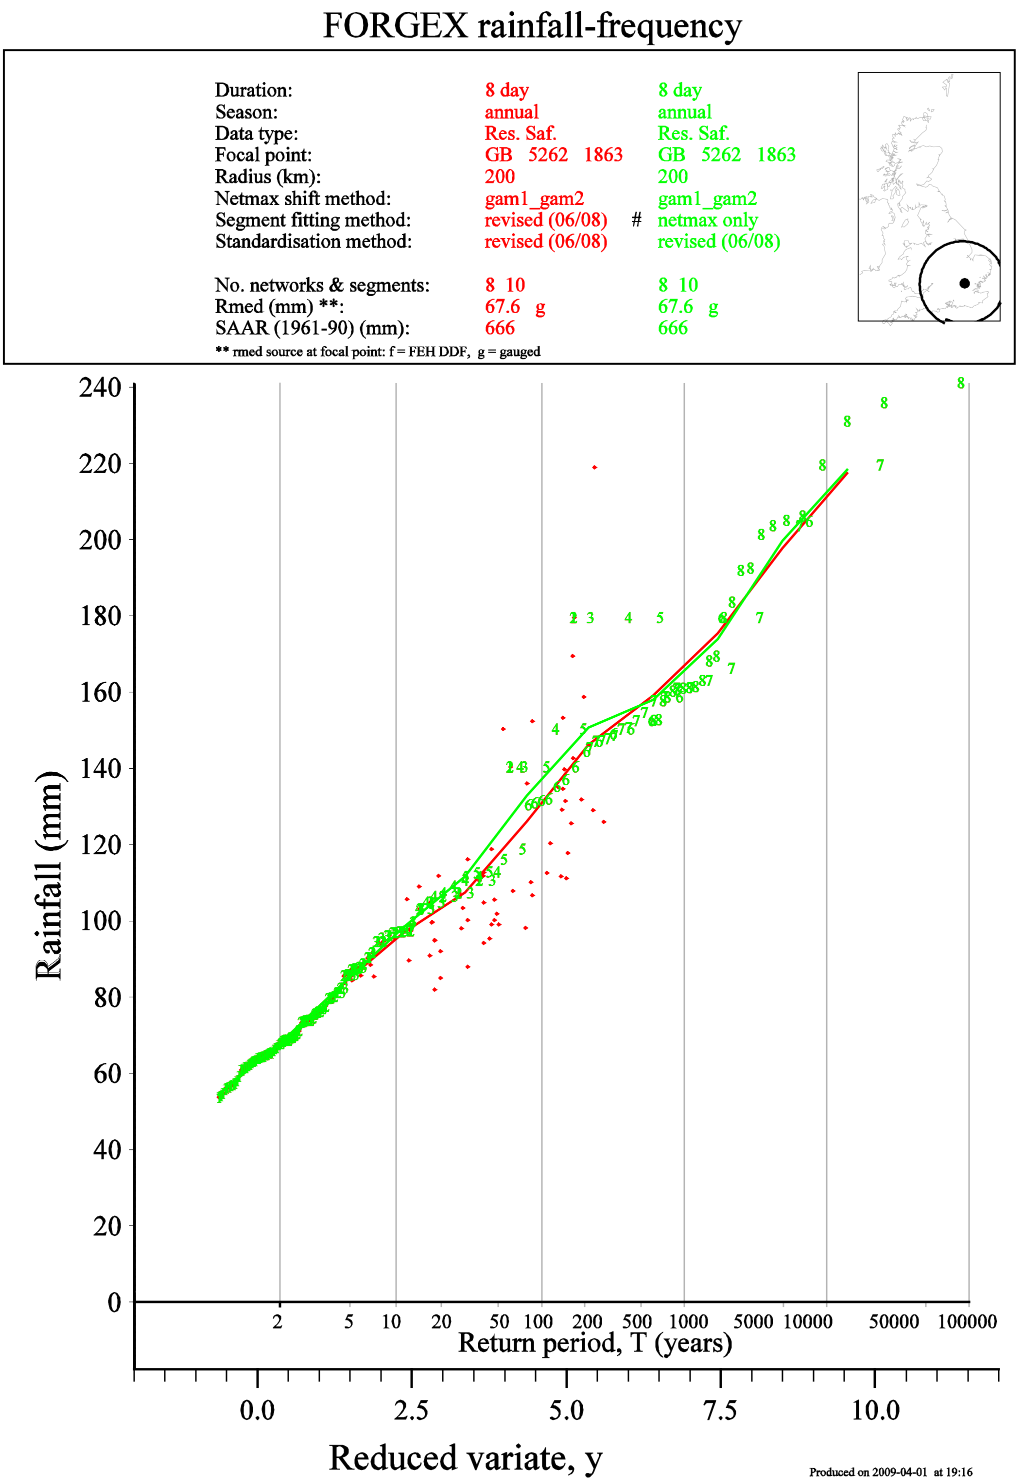

Figure 7.12 The change in the 8-day rainfall frequency curve for the location where excluding pooled points caused the biggest increase (the green curve is fitted to network maxima only)

7.7.2 New rules for the definition of network radii

As described above, the network radii in the FEH procedures were determined primarily by the disposition of the pooled points. With pooled points no longer being used for fitting, it was decided to develop new network definition rules.

The radius of the first network (network 1) is selected to give the smallest network with a total number of effective gauge-years of at least 40. The total number of effective gauge-years is the sum of the effective number of independent gauges in each year, according to the spatial dependence model, for all the years with data. Thus wherever the closest gauge has at least 40 years of data, it will be the sole member of the first network.

It had been observed that, under the FEH rules, the step in radius to the final network was often large (for example, Figure 7.13, which corresponds to the rainfall frequency curve in Figure 7.8) and it was considered that there was scope for obtaining a more locally focused result and a spatially smoother final product (as reflected in the variation of results between neighbouring locations in the final national 1-km grid) if there was a more gradual increase in network radius up to the final network.

Under the new rules, with the exception of the first network, networks are defined on the basis of the number of gauges. The number of gauges in each network is defined to be a specified proportion of the number of gauges in the next higher network, in the following way:

Terminology

Primary radius – 200 km (see also ‘secondary radius’, in Section 7.7.3)
Primary network – the network whose radius is the primary radius
G_p – the number of gauges in the primary network
$G_{p\text{-}1}$– the number of gauges in the network immediately below the primary network
$G_{p\text{-}m}$ – the number of gauges in the network m networks below the primary network
G_1 = the number of gauges in the first network

Radii are selected such that (truncating to integer)

$G_{p\text{-}1} = 0.8736\times G_p$
$G_{p\text{-}2} = 0.8736\times G_{p\text{-}1}$
$G_{p\text{-}3} = 0.8736\times G_{p\text{-}2}$
$G_{p\text{-}4} = 0.6667\times G_{p\text{-}3}$
$G_{p\text{-}5} = 0.6667\times G_{p\text{-}4}$
…

and so on, using 0.6667, until

$G_{p\text{-}m} \leq 1.5\times G_1$.

Note that the final member of the above list is not used; network 2 is then defined by the penultimate member, so ensuring it contains at least 50 per cent more gauges than network 1.

This ensures that, stepping up through the networks from network 2 onwards, each contains 50 per cent more gauges than the preceding one, with the exceptions of the final three steps, which have been made smaller ($0.8736^3 = 0.6667$) to promote spatial smoothness (described below).

7.7.3 Use of additional networks, up to a radius of 300 km

This change has been introduced for two reasons. Firstly, the change promotes spatial smoothness, so that the final national grid of results does not exhibit ridge effects at radii of 200 km from major storms: this is achieved, in combination with a suitable weighting system, by providing a 100 km transition zone. Secondly, it provides larger networks which enable higher return periods to be estimated and these guide the upper sections of the rainfall frequency curve. The outer radius used here is referred to as the 'secondary radius'.

The rules for defining network radii between the primary and secondary radii are similar to those described above: increments in the numbers of gauges are equal to 50 per cent of the gauges in the previous network, with the exception of three smaller increments immediately following the primary network, again to promote smoothness, and a rule to avoid a very small final increment.

Spatial smoothness is promoted in two ways:

- by applying a reduced weight, termed the '*distance reduction weight',* to network maxima from gauges located between the primary and secondary radii according to weight = $1 - 0.09999d^{0.5}$, where d is the distance (in km) of the gauge beyond the primary radius; this gives a weight that tails off to 0.0001 at the secondary radius; and
- by applying a reduced weight, termed the '*higher-network reduction weight',* to networks above the primary, equal to 0.75^{n-p}, where n is the network number and p is the number of the primary network; thus the network after the primary is given a weight of 0.75, the next a weight of 0.5625, and so on.

The second of these weights was introduced because a top-ranked network maximum at the distance of the secondary radius, even though lowly weighted, could otherwise have a marked indirect effect on the segment fitting because of its effect on the estimated return periods of the other maxima in the network.

For the same location and duration as in the previous example, Figure 7.14 shows the network limits defined by the new rules.

7.7.4 New method of segment fitting

In the new method, there is no longer a relationship between segments and networks. As previously, the first segment extends horizontally from y = -0.6335 to 0.3665, and its upper end has a fixed value on the vertical axis of 1.0 for consistency with the standardisation model: *i.e.* the standardised rainfall of 1 for a return period of 2 years. The y range from 0.3665 to the third highest point in the primary network (for consistency with the FEH method) is divided into the smallest possible number of equal-width segments of y width ≤ 1.0. The y range from the third highest point in the primary network to the highest point in the secondary network is divided into the shortest possible equal-width segments of y width ≥ 1.0. The fitted values corresponding to these upper segments do not directly provide data for the DDF model, but, because of the penalty placed on the angle between adjacent segments by the fitting procedure, they do influence the position of the other segments.

All of the points from network 1 are used in segment fitting. All higher networks have a lower cut-off from below which points are not used. This is the first point in the network that has a return period at or below one third of the total number of effective gauge-years in the previous network. Table 7.3, together with Figure 7.15, shows what this means in practice for the 8-day Kent example. Note that network 18 (denoted by ‘H’) is the primary network, and network maxima from higher networks are shown in a smaller font.

It had initially been intended to base the lower cut-off on half, rather than a third, of the total number of effective gauge-years in the previous network, echoing the FEH flood frequency guideline that an *M*-year record may be used to estimate up to an *M*/2-year flood. However, this was found to produce less smooth curves due to the smaller overlap between networks.

Before the segments are fitted, further weights are applied to the data, and these are described in the following section.

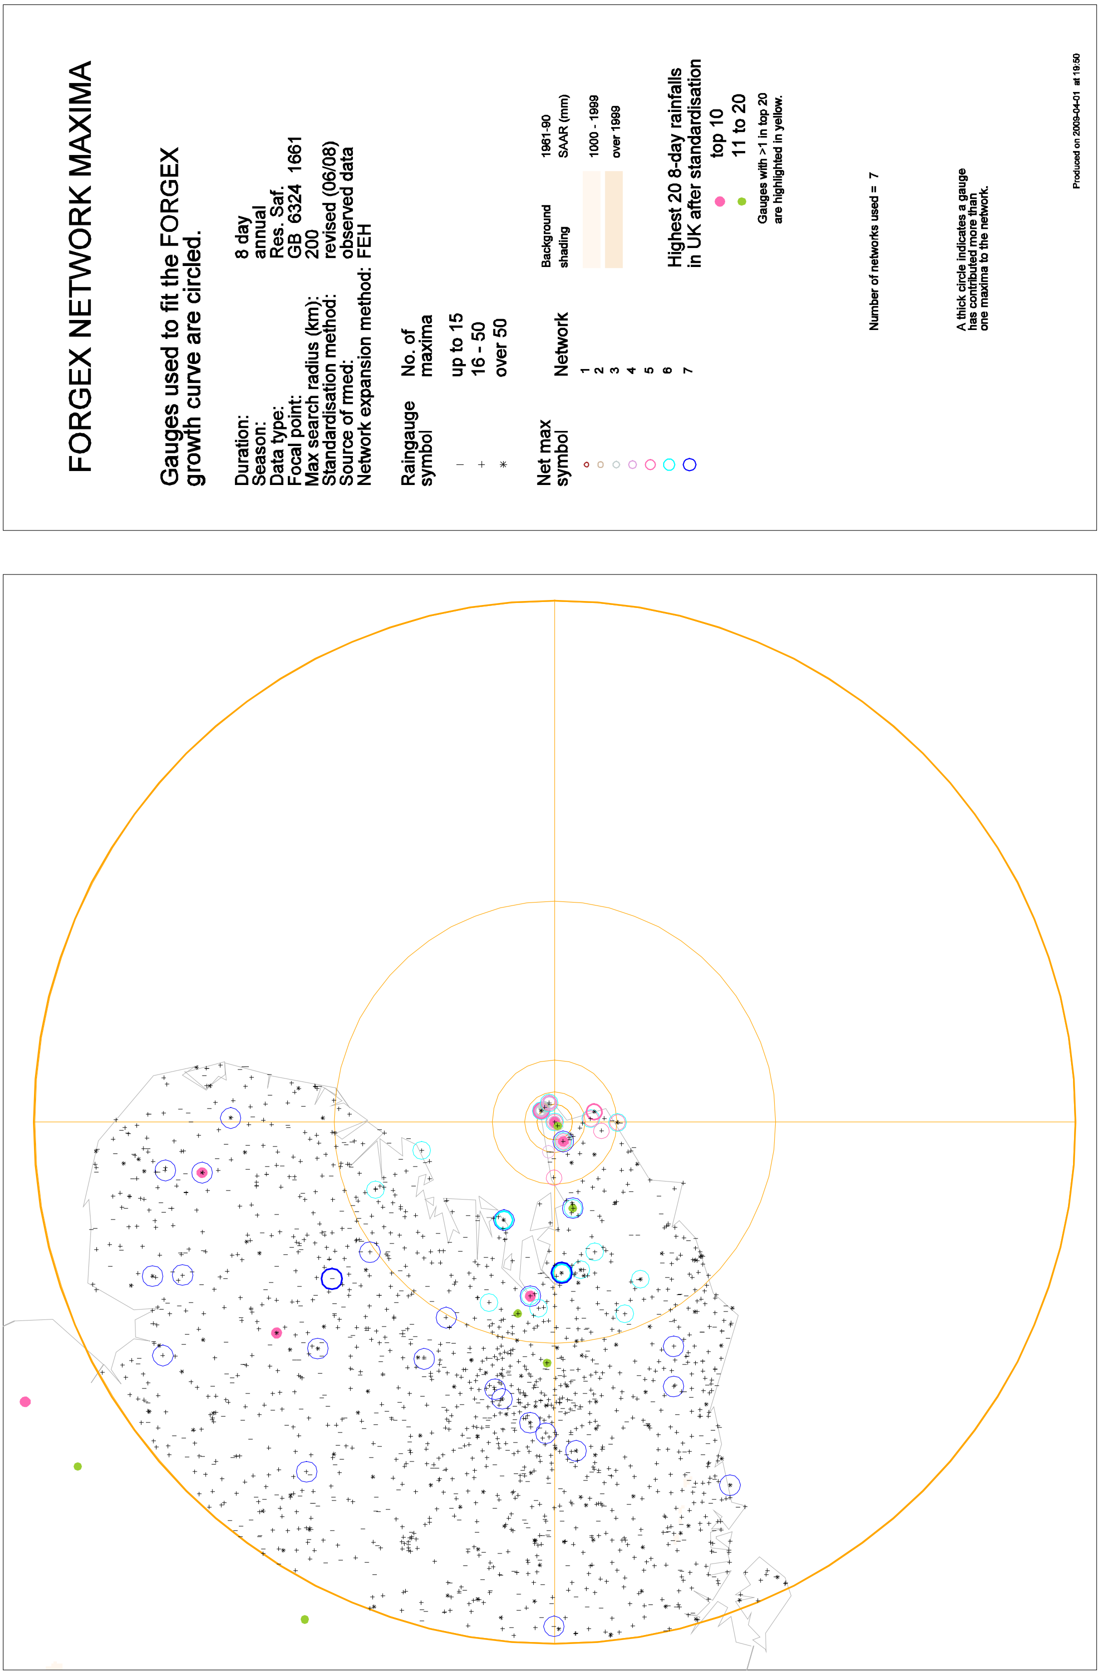

Figure 7.13 Example of the location of network limits and network maxima, using FEH rules

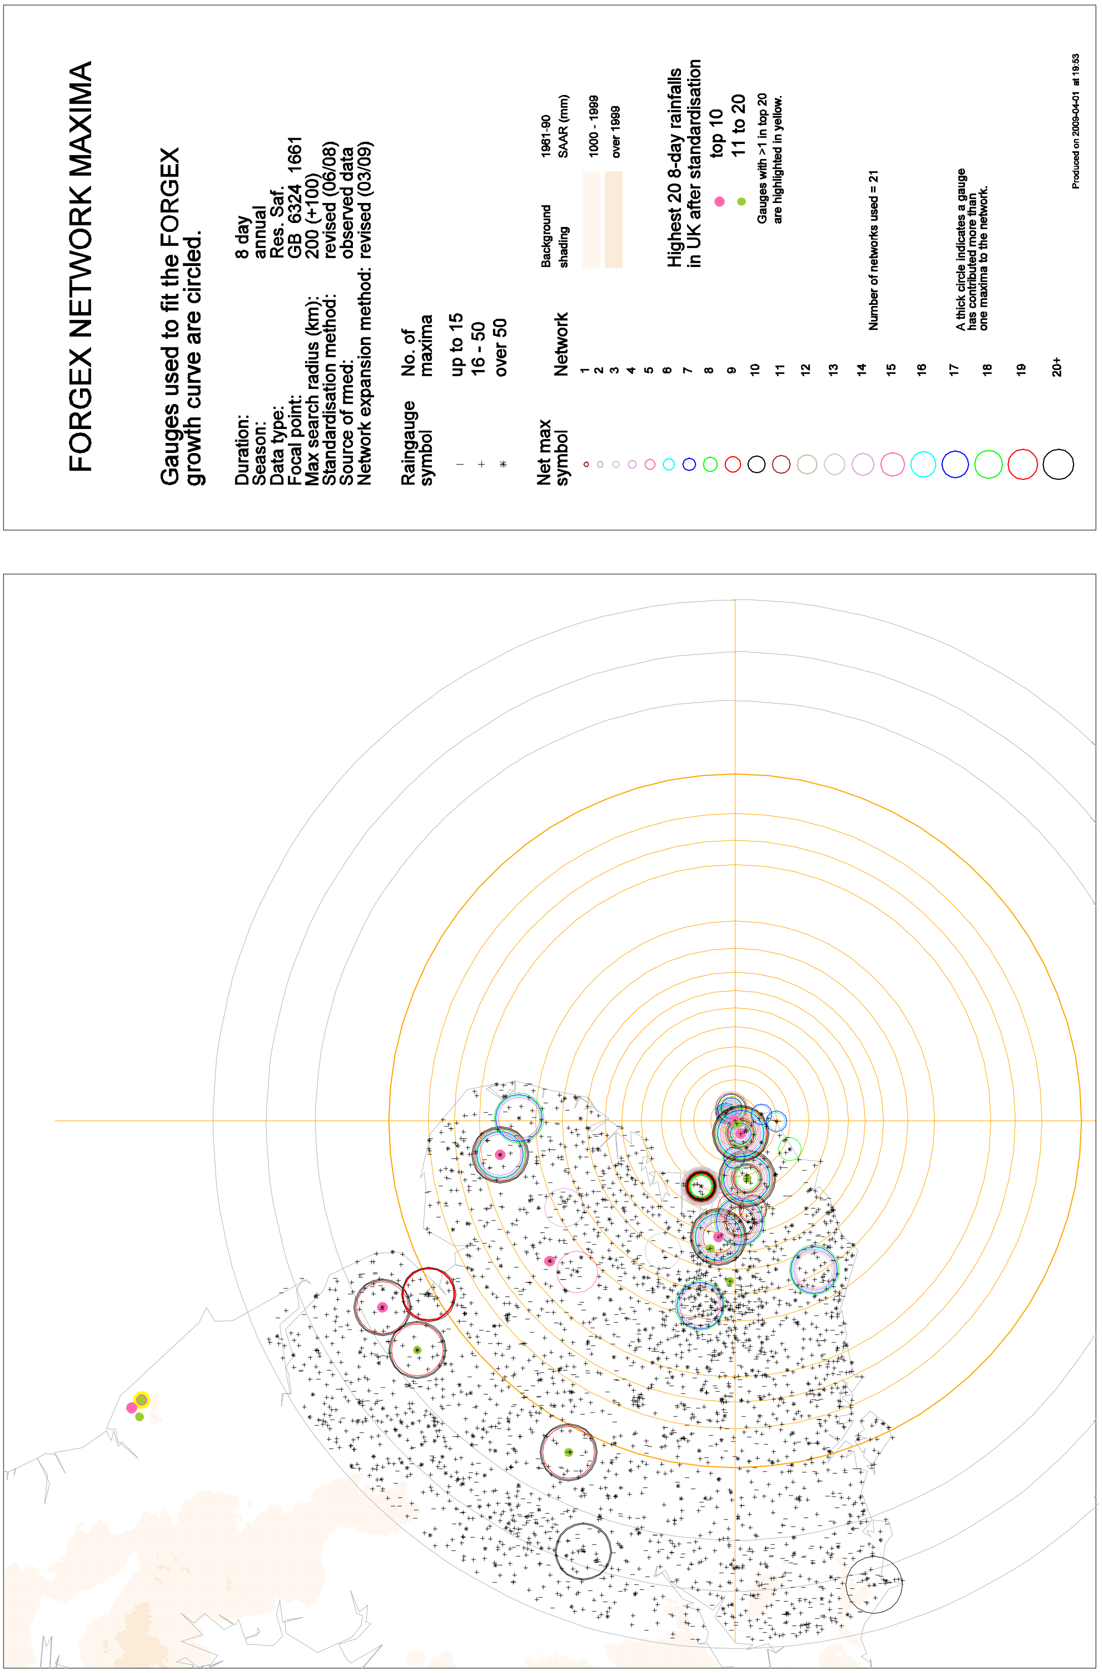

Figure 7.14 Example of the location of network limits and network maxima, using the revised rules

Table 7.3 Network statistics associated with Figure 7.15

Network	No. of gauges	Radius (km)	Total number of effective gauge-years	Return period used to define the lower cut-off of the next network (years)	Previous column expressed in terms of Gumbel reduced variate, *y*	Network width weight (see Section 7.7.5)
1	2	1.9	45	15	2.68	1.00
2	5	7.0	111	37	3.59	0.91
3	8	8.1	129	43	3.75	0.16
4	12	9.8	139	47	3.83	0.08
5	19	14.1	161	54	3.97	0.14
6	29	24.0	317	106	4.65	0.68
7	44	32.0	377	126	4.83	0.18
8	66	42.7	522	174	5.16	0.33
9	99	54.3	659	220	5.39	0.23
10	149	65.4	870	290	5.67	0.28
11	224	75.2	1127	376	5.93	0.26
12	336	85.6	1522	507	6.23	0.30
13	504	99.8	2104	701	6.55	0.32
14	756	115.3	2722	907	6.81	0.26
15	1135	147.5	3883	1294	7.17	0.36
16	1300	161.8	4473	1491	7.31	0.14
17	1489	177.1	5249	1750	7.47	0.16
18	1705	199.8	6319	2106	7.65	0.18
19	2111	242.2	7877	2626	7.87	0.22
20	2516	269.8	10166	3389	8.13	0.26
21	2922	300.0	12244	n/a	n/a	1.00

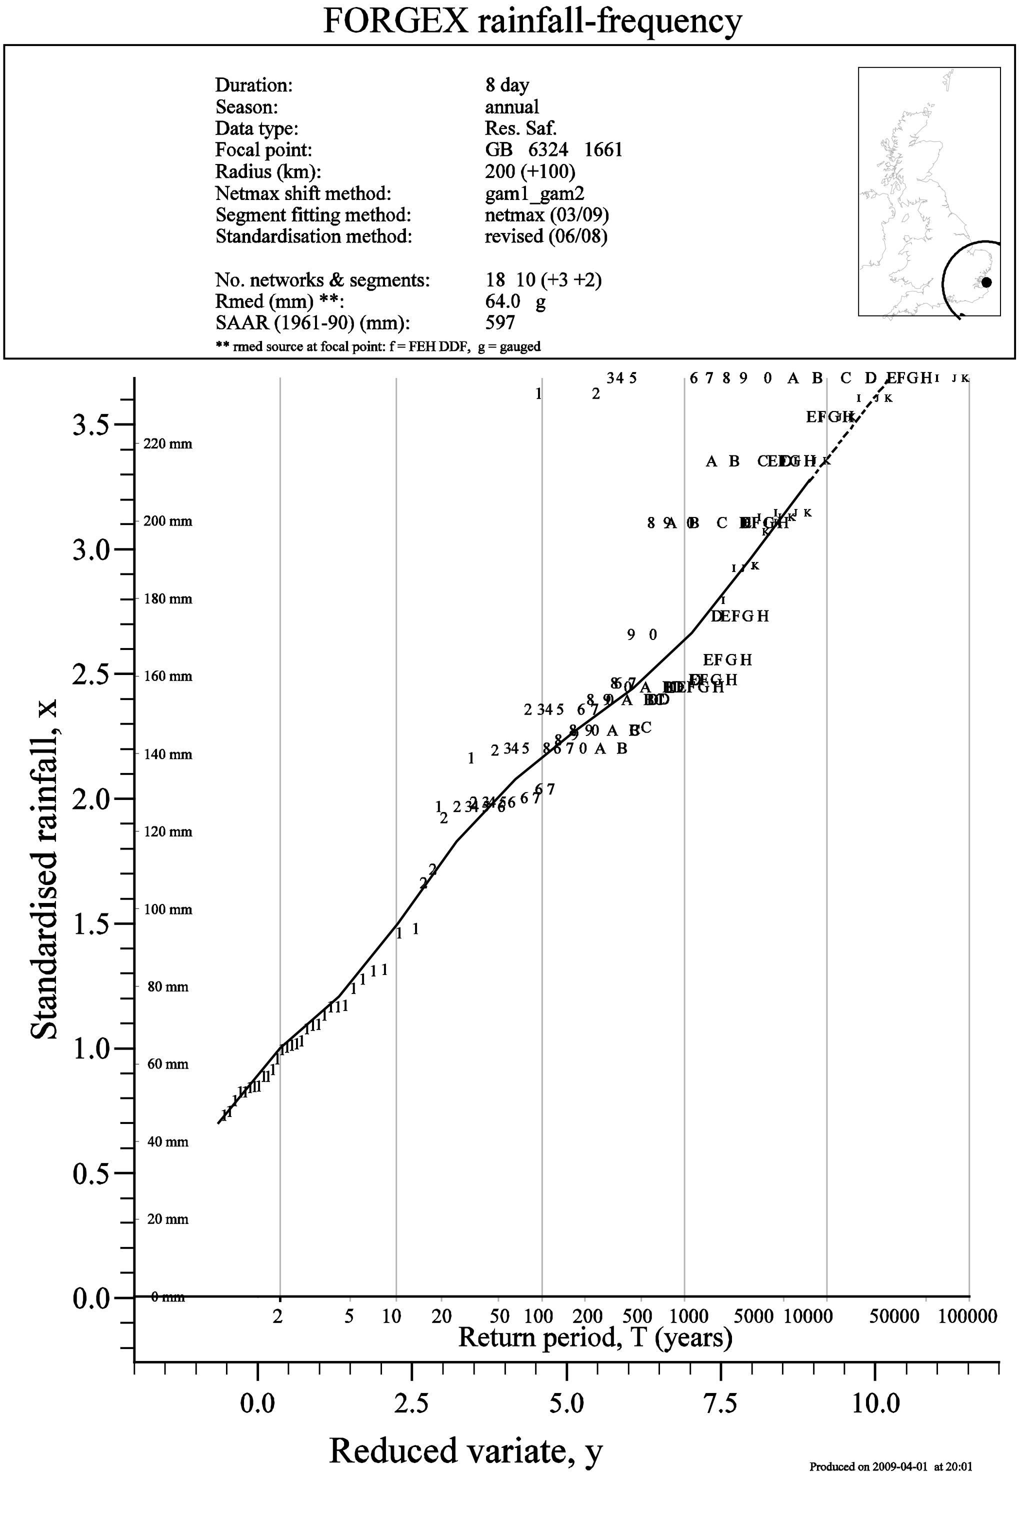

Figure 7.15 8-day rainfall frequency curve for a location close to the 1973 West Stourmouth event, produced using the revised network definition and segment fitting method.

7.7.5 Weighting of network maxima

In addition to the distance reduction weight and the higher-network reduction weight, which were introduced to promote spatial smoothness (Section 7.7.3), the network maxima are subject to up to three additional weights.

Uncertainty weight
The uncertainty in the estimate of the return period (or Gumbel *y* value) for a particular member of a ranked series of annual maxima depends on its position in the ranking, the number of maxima and the properties of the distribution of the underlying population. Generally, the uncertainty is greater towards the two extremes. When the return period is being estimated from network maxima, the uncertainty is also affected by the number of effective gauges in each year. On the assumption that the data can be represented by GEV distributions with a shape parameter, *k,* of -0.1 for hourly durations and -0.05 for daily durations, the uncertainty in the *y* value of each of the network maxima has been estimated and used to weight the values prior to segment fitting. The main effect of this is that when a value appears in several networks, its weight in the higher networks tends to be higher due to a combination of the higher number of effective gauge-years, often the larger number of years represented, and a fall in the rank of the value. Figure 7.16 shows the influence of this weight on the eight-day Kent example (note that on both lines, all the other weights are turned on).

Network width weight
As explained above, six of the networks (from primary-2 to primary+3) have been defined using increments in gauge numbers that are effectively one third of the usual increment. The purpose of this is to promote gradual change in the vicinity of the primary radius, the distance at which full distance and network weighting cease to apply. However, this measure alone would result in an over-representation of information from networks in that radius range. To counter this, each network is weighted according to its 'width', where width is defined as the y-range from the point at its lower cut-off to the point at the lower cut-off of the next higher network. The weights of the first and last networks are set to 1, as the above definition of width is not applicable in these cases, and a ceiling of 1 is applied to the weight of all other networks (although it is expected that this will not be reached in most instances). Figure 7.17 shows the influence of this weight on the eight-day Kent example (note that, for both lines, all the other weights are turned on). The network width weights are shown in the final column of Table 7.3.

Higher-network priority weight
Testing of the new method revealed a number of instances that produced a curve considered to be too low at the top end, and where the results from the FEH fitting method appeared to better reflect the data. This problem was countered by applying a reduction weight of 1/*h* to a network maximum if a higher network contained a maximum that had the same or higher rainfall and the same or lower *y* value. Where the higher network was at or below the primary network, *h* was set to 10. For higher networks *h* was reduced according to the higher-network reduction weight, described in Section 7.7.3, in order to

moderate the effects of data from beyond the primary radius. For example, it is shown in Section 7.7.3 that network primary+2 has a higher-network reduction weight of 0.5625. Therefore, if a point in a lower network has a higher (or equal) y value and a lower (or equal) rainfall than a point in network primary+2, it will receive a higher-network priority weight of 1/(10 × 0.5625). Of the sites and durations tested, the London area for one-hour duration is most affected by this weight (Figure 7.18: note that on both lines, all the other weights are turned on).

7.7.6 Effects of the revisions on spatial smoothness

The two largest standardised one-day annual maximum rainfalls on the project database are the July 1955 event at Martinstown in Dorset and the September 1973 event at West Stourmouth in Kent. The two locations NGR 4532, 2654 and NGR 4533, 2653, near Daventry in Northamptonshire, are separated by only 141 m, and would be expected to have very similar rainfall regimes. However, they are respectively just over and just under 200 km from both Martinstown and West Stourmouth. Figure 7.19 shows how this affects the rainfall frequency curves derived using the FEH FORGEX method. Figure 7.20 shows the same locations fitted using the new method. The difference has been considerably reduced, with the discrepancy at a return period of 22,000 years falling from 17 mm to 3 mm. Due to the introduction of the secondary radius, it is possible that the new method could have shifted the problem by 100 km. Figure 7.21 applies the same test to the new method at the equivalent 300 km points (near Hathersage in the Peak District). This shows that while the two large maxima have a strong influence on the dashed section of the line, their effect on the sections that are used (which in this case go up to 16,545 year return period) is minor. It is reassuring to note that at the 10,000 year return period (the primary objective of this project) using the new method there is negligible difference between the two curves at either location.

Given the way that these sites have been selected, it is unlikely that there will be circumstances where neighbouring points differ by much more than this. So, it is reasonable to expect that a national DDF model built on the output from this method will be spatially smooth.

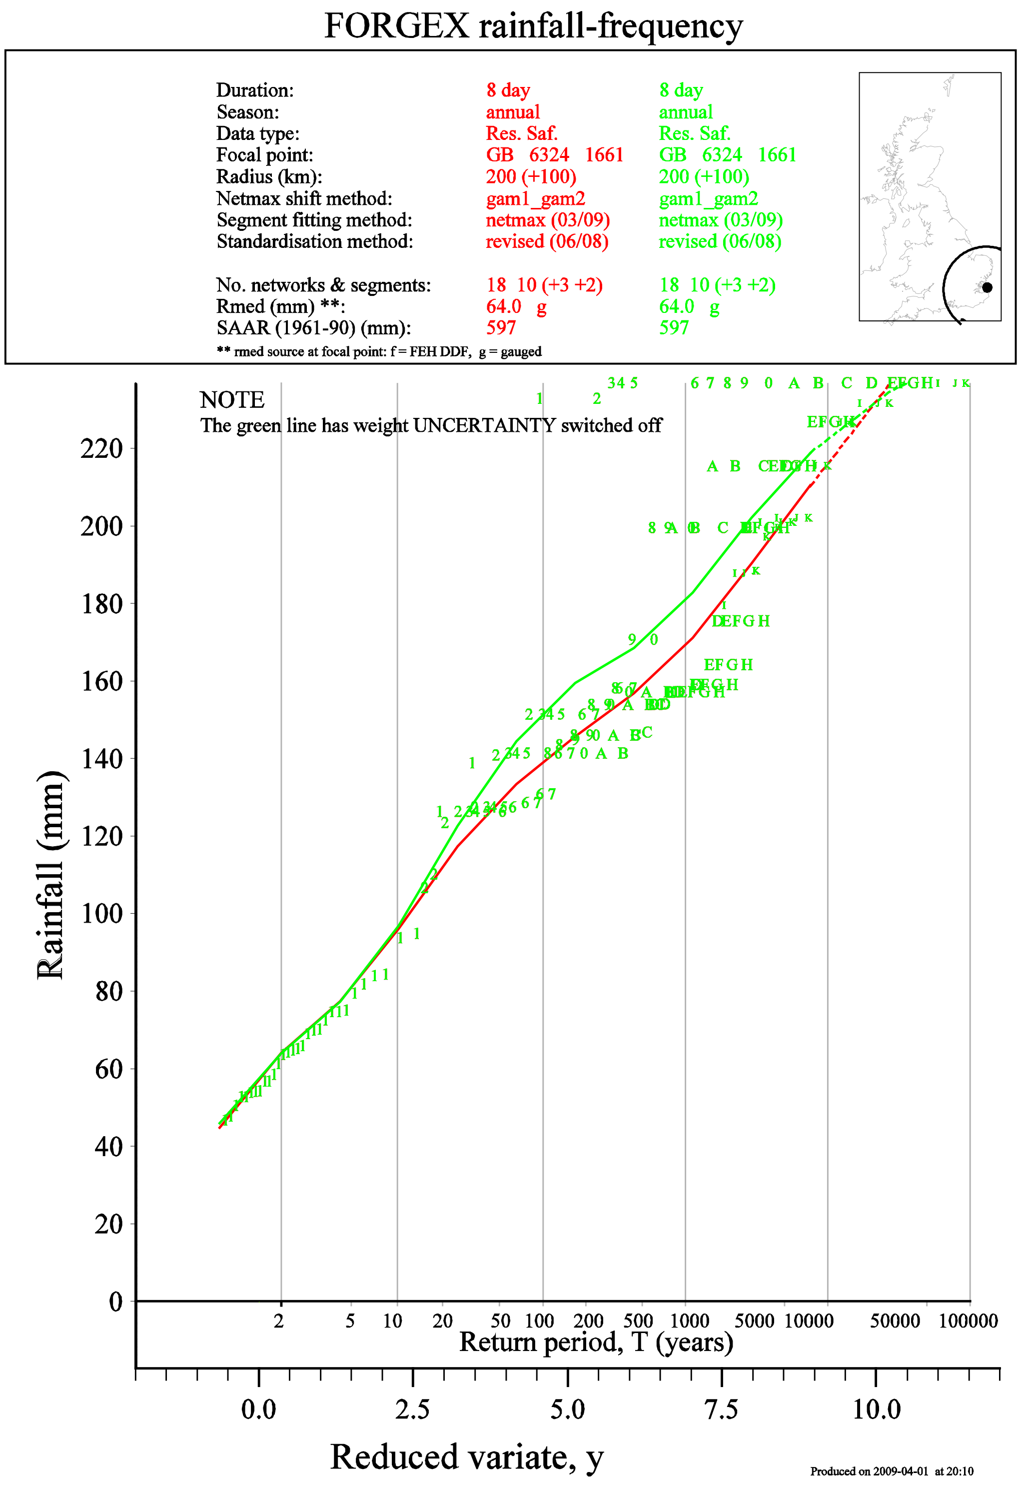

Figure 7.16 New segment fitting method: example of the effect of the uncertainty weight (the red line uses the uncertainty weight)

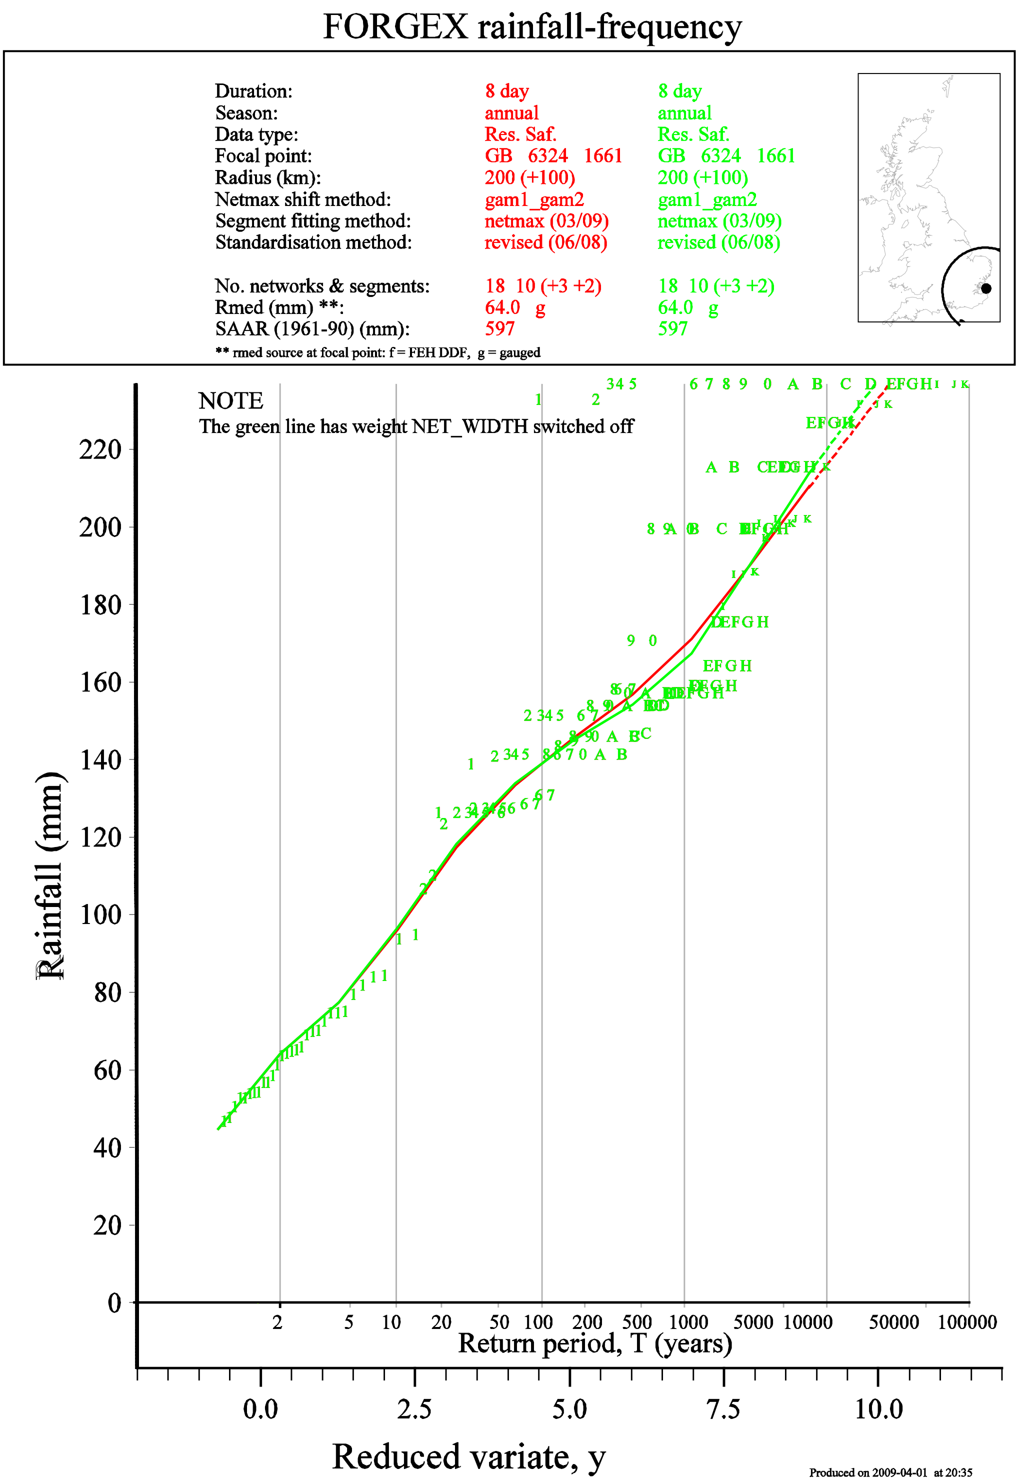

Figure 7.17 New segment fitting method: example of the effect of the network width weight (the red line uses the network width weight)

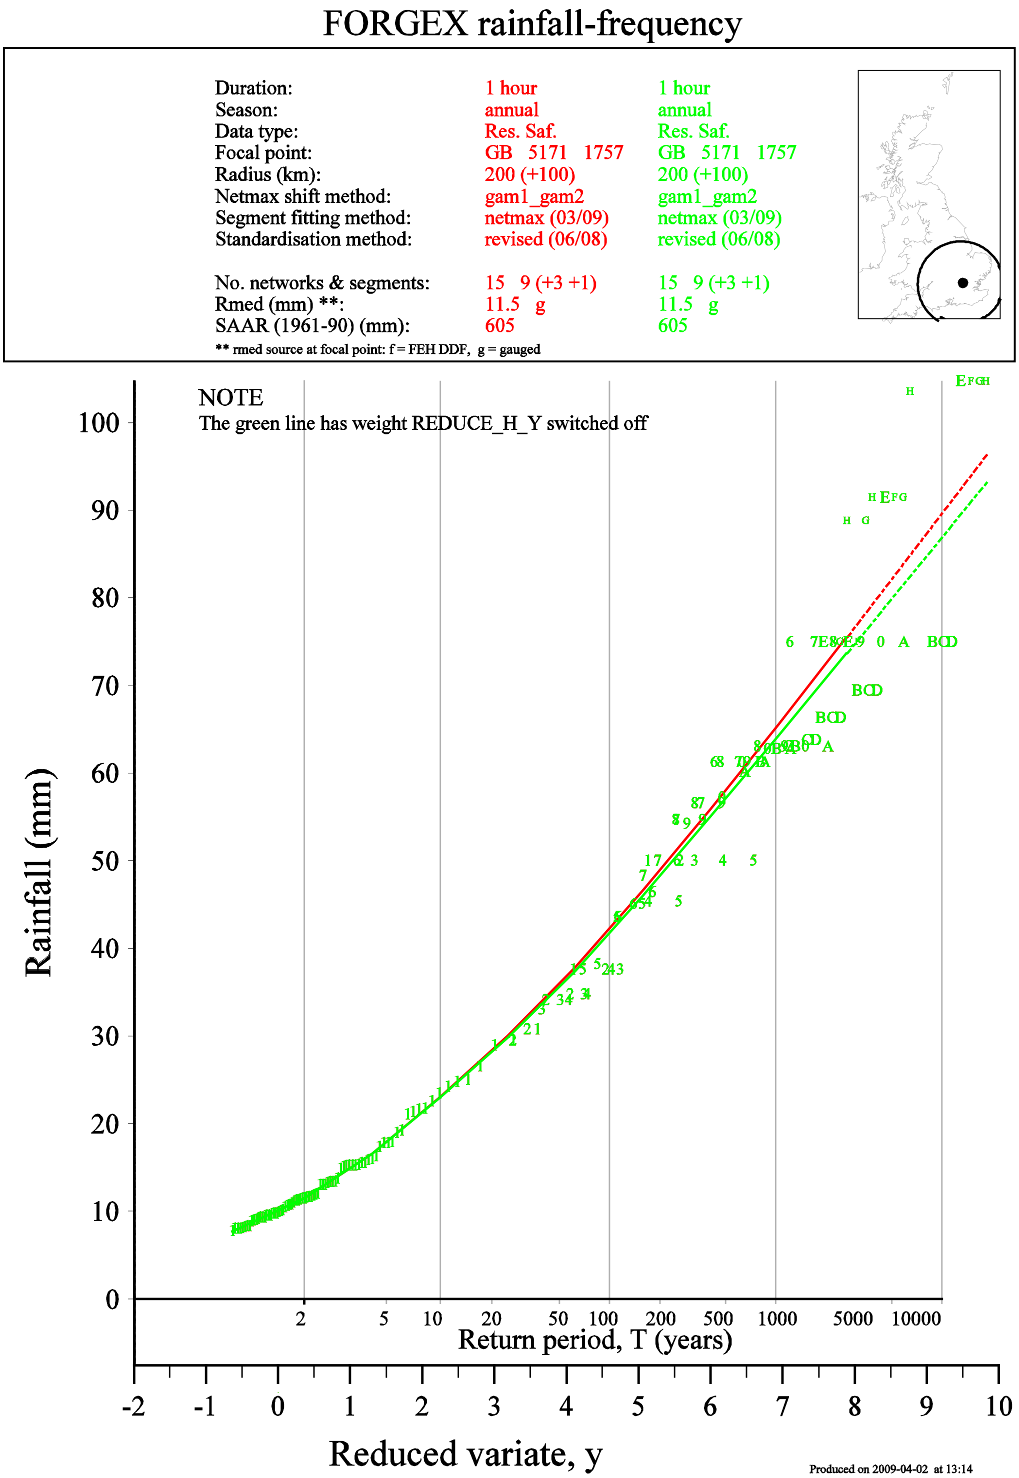

Figure 7.18 New segment fitting method: example of the effect of the higher-network priority weight (the red line uses the higher-network priority weight)

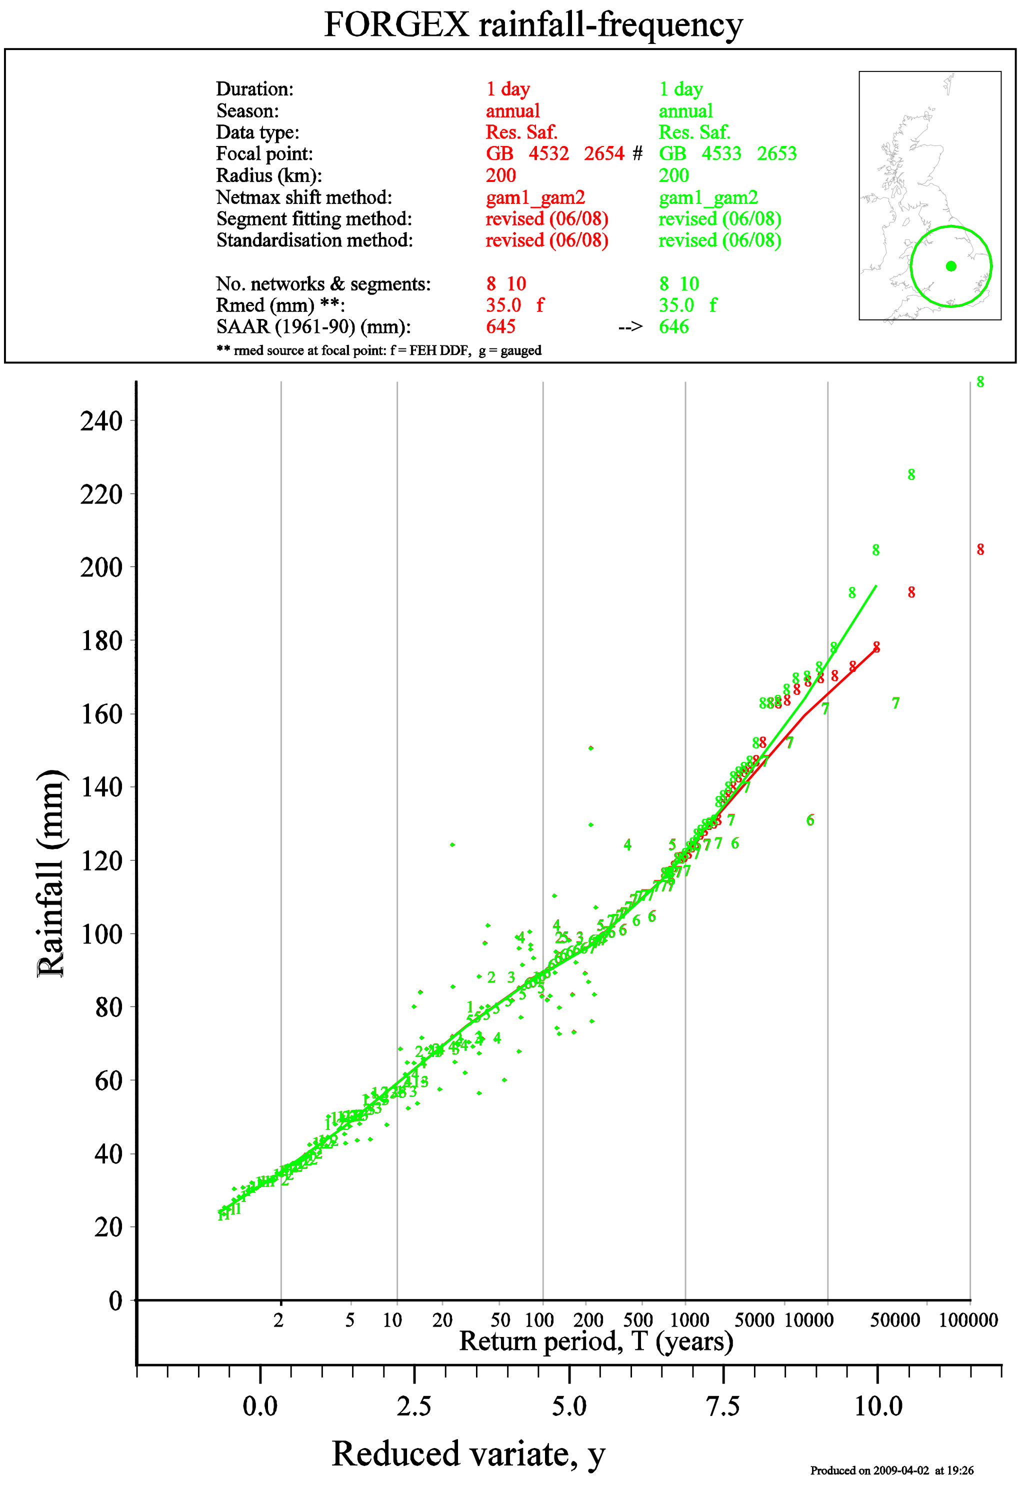

Figure 7.19 Spatial smoothness: FORGEX rainfall frequency curves using the FEH network selection and segment fitting method for two locations separated by 141 m (the green curve is for a site just under 200 km from the two largest 1-day events)

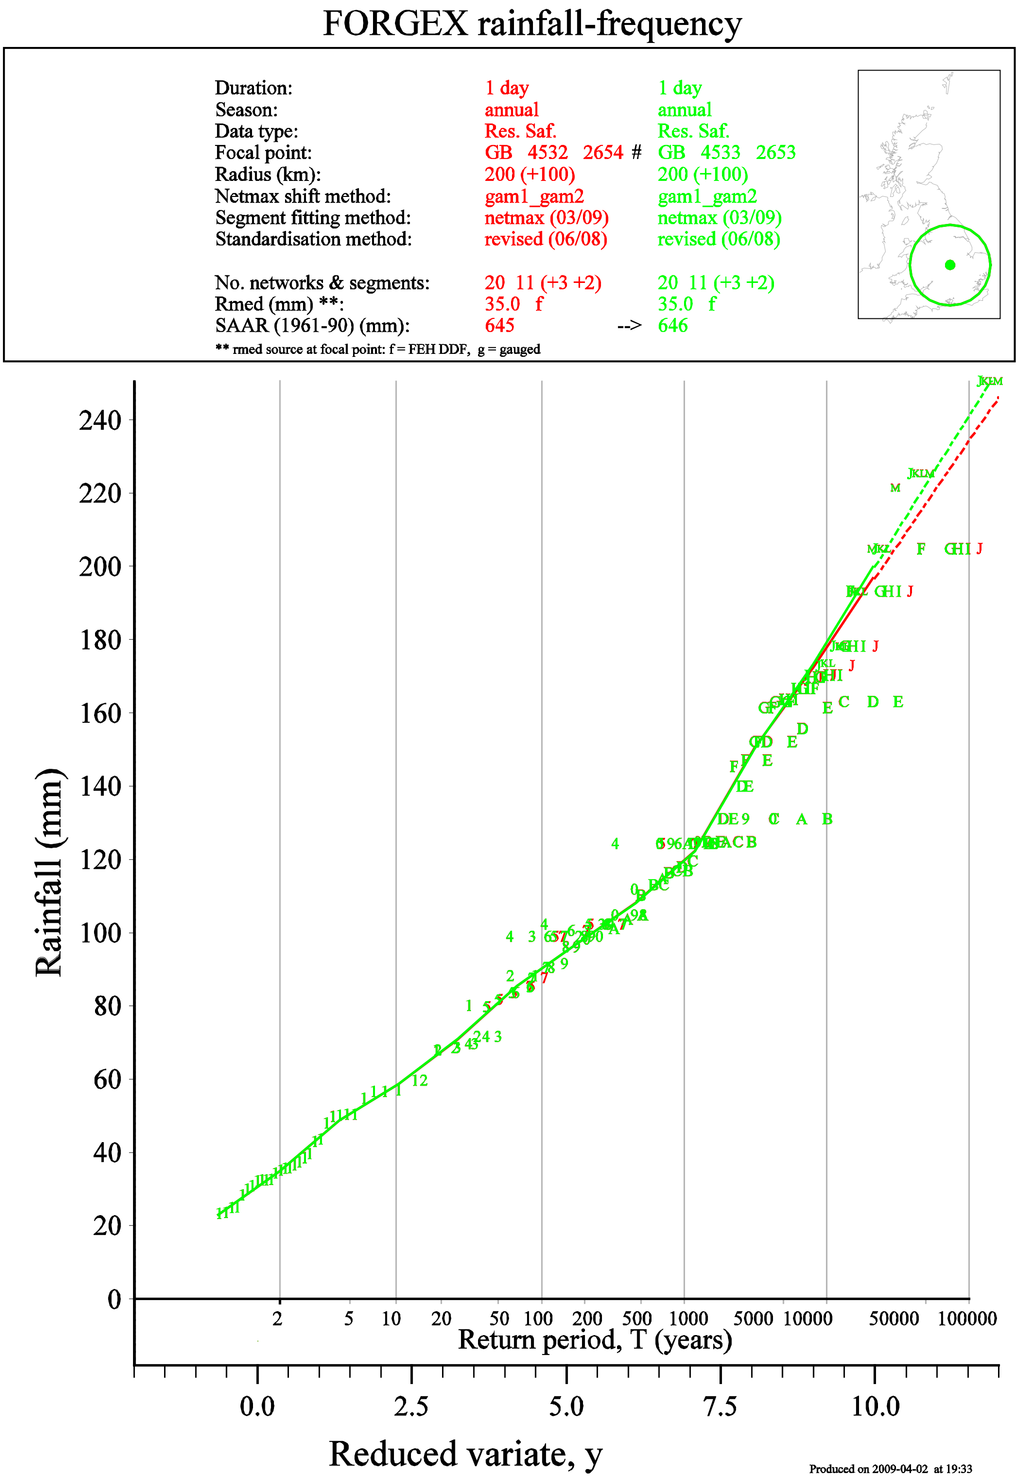

Figure 7.20 Spatial smoothness: FORGEX rainfall frequency curves using the revised network selection and segment fitting method for two locations separated by 141 m (the green curve is for a site just under 200 km from the two largest 1-day events)

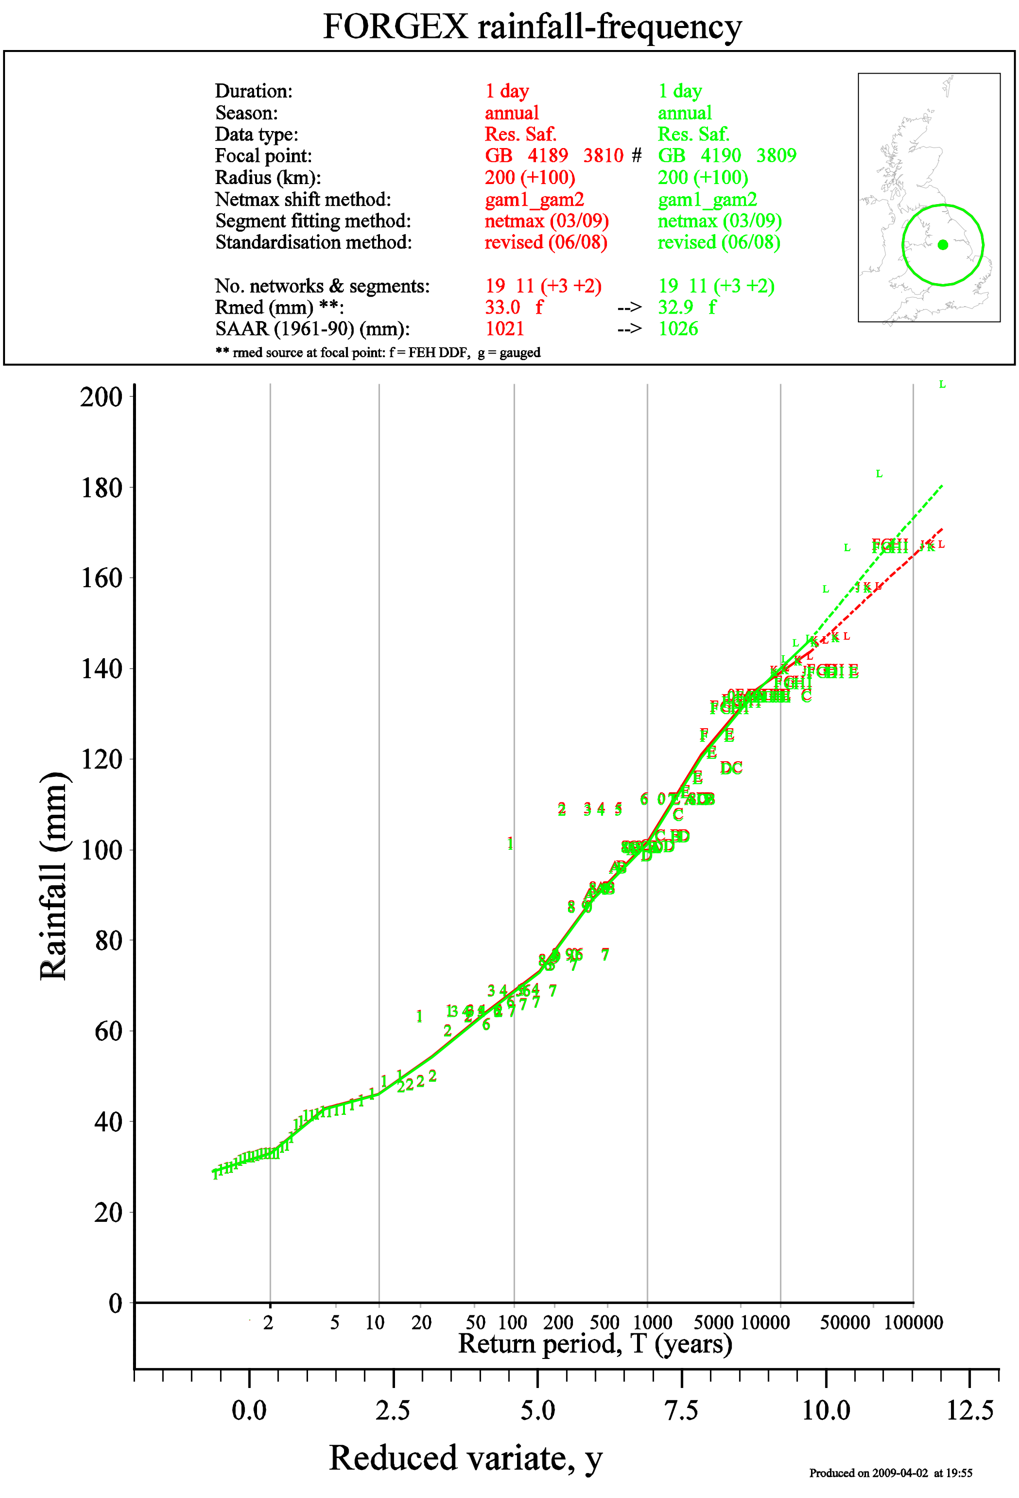

Figure 7.21 Spatial smoothness: FORGEX rainfall frequency curves using the revised network selection and segment fitting method for two locations separated by 141 m (the green curve is for a site just under 300 km from the two largest 1-day events)

7.7.8 Effects of the revisions on high return period rainfall estimation

To test the effects of the revisions, comparative plots were produced showing the rainfall frequency curve for the original and revised FORGEX methods (both using the new standardisation and spatial dependence model) at each of the 35 test locations listed in Appendix G for each of the durations one hour, one day and eight days. The changes at the top end of the curve are summarised in Table 7.4.

Table 7.4 Summary of the effects of FORGEX methodology revisions

Duration	Number where new method is higher	Number where there is little difference	Number where new method is lower
1 hour	15	9	11
1 day	14	10	11
8 days	20	10	5
Total	49	29	27

This suggests that these revisions are more likely to cause an increase than a decrease in rainfall estimates at high return periods.

The performance of the new method was then assessed qualitatively by visually examining all the plots where the curves differed at high return periods and judging which curve best represented the data points. Whilst this is a somewhat informal approach, it should give a reasonable comparison between the methods. Aggregated over the three durations, the results of this exercise are shown in Table 7.5, and show that the new method appears, in general, to be preferable.

Table 7.5 Summary of the performance of the revised FORGEX

New curve relative to old	New considered to be a better representation	No clear best performer	Old considered to be a better representation
Lower	5	16	6
Higher	34	14	1
Total	39	30	7

For each duration, plots for the locations with the greatest positive and negative differences have been included as Figures 7.22 to 7.27.

7.8 The net effect of all the changes

Again using the 35 test sites at one hour, one day and eight day durations, and the latest dataset, the net effect of the new standardisation, spatial dependence model and FORGEX revisions has been assessed with comparative plots. The changes at the top end of the curves are summarised in Table 7.6.

Table 7.6 Summary of the net effects of changes to standardisation, spatial dependence model and FORGEX methodology.

Duration	Number where new method is higher	Number where there is little difference	Number where new method is lower
1 hour	4	3	28
1 day	1	1	33
8 days	0	0	35
Total	5	4	96

The following examples have been included:

- 1-hour duration: the case where the new exceeds the old by the most (Fig. 7.28); a typical example, which is the central England site used previously (Fig. 7.29); and the case where the new is below the old by the greatest amount (Fig. 7.30). Note that the latter is the test site with the highest SAAR (Honister) and is strongly influenced by the change in the standardisation procedure.

- 1-day duration: the only case where the new exceeds the old (Fig. 7.31); a typical example, again the central England site used for hourly durations (Fig. 7.32); and the case where the new is below the old by the greatest amount (Fig. 7.33). As with the one hour duration, the greatest reduction occurs at the Honister site, down by 37 per cent at the 10,000 year return period. But, even in low SAAR locations the reduction can be substantial; for example, over 25 per cent in the London area.

- 8-day duration: the case where the new is closest to the old (Fig. 7.34); a typical example, again the central England site used for hourly durations (Fig. 7.35); and the case where the new is below the old by the greatest amount (Fig. 7.36). As for one hour and one day duration, the greatest reduction occurs at the Honister site, down by 36 per cent at the 10,000 year return period. But, in the London area the reduction is not much less than this.

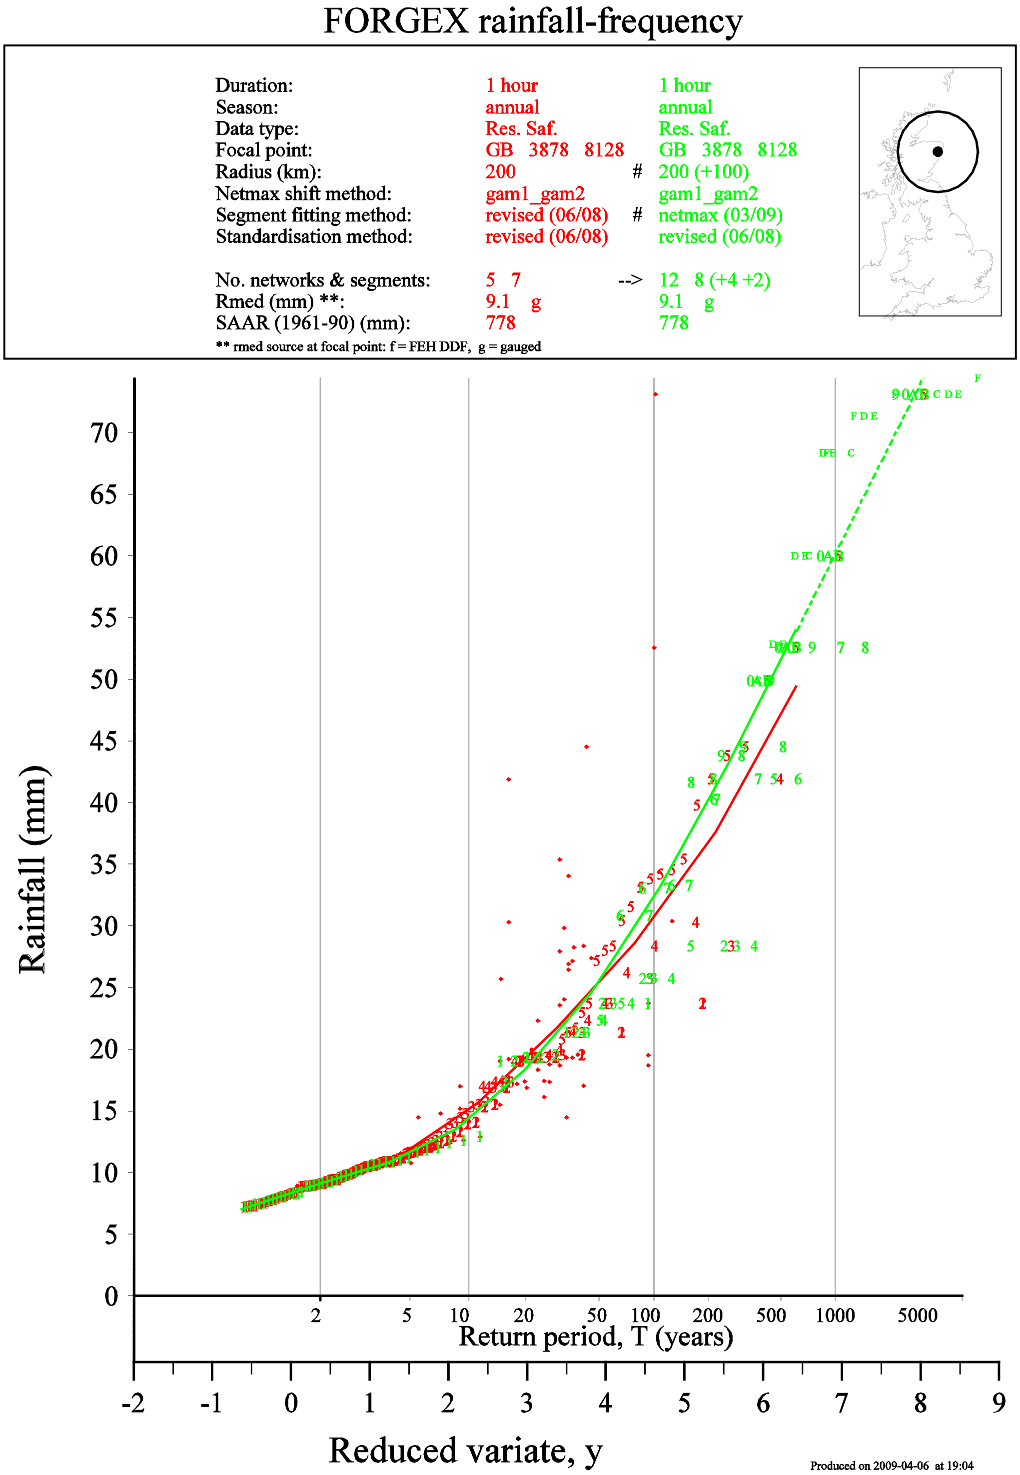

Figure 7.22 The largest positive effect at 1-hour duration of the changes to FORGEX network selection and segment fitting (the green curve is from the new method)

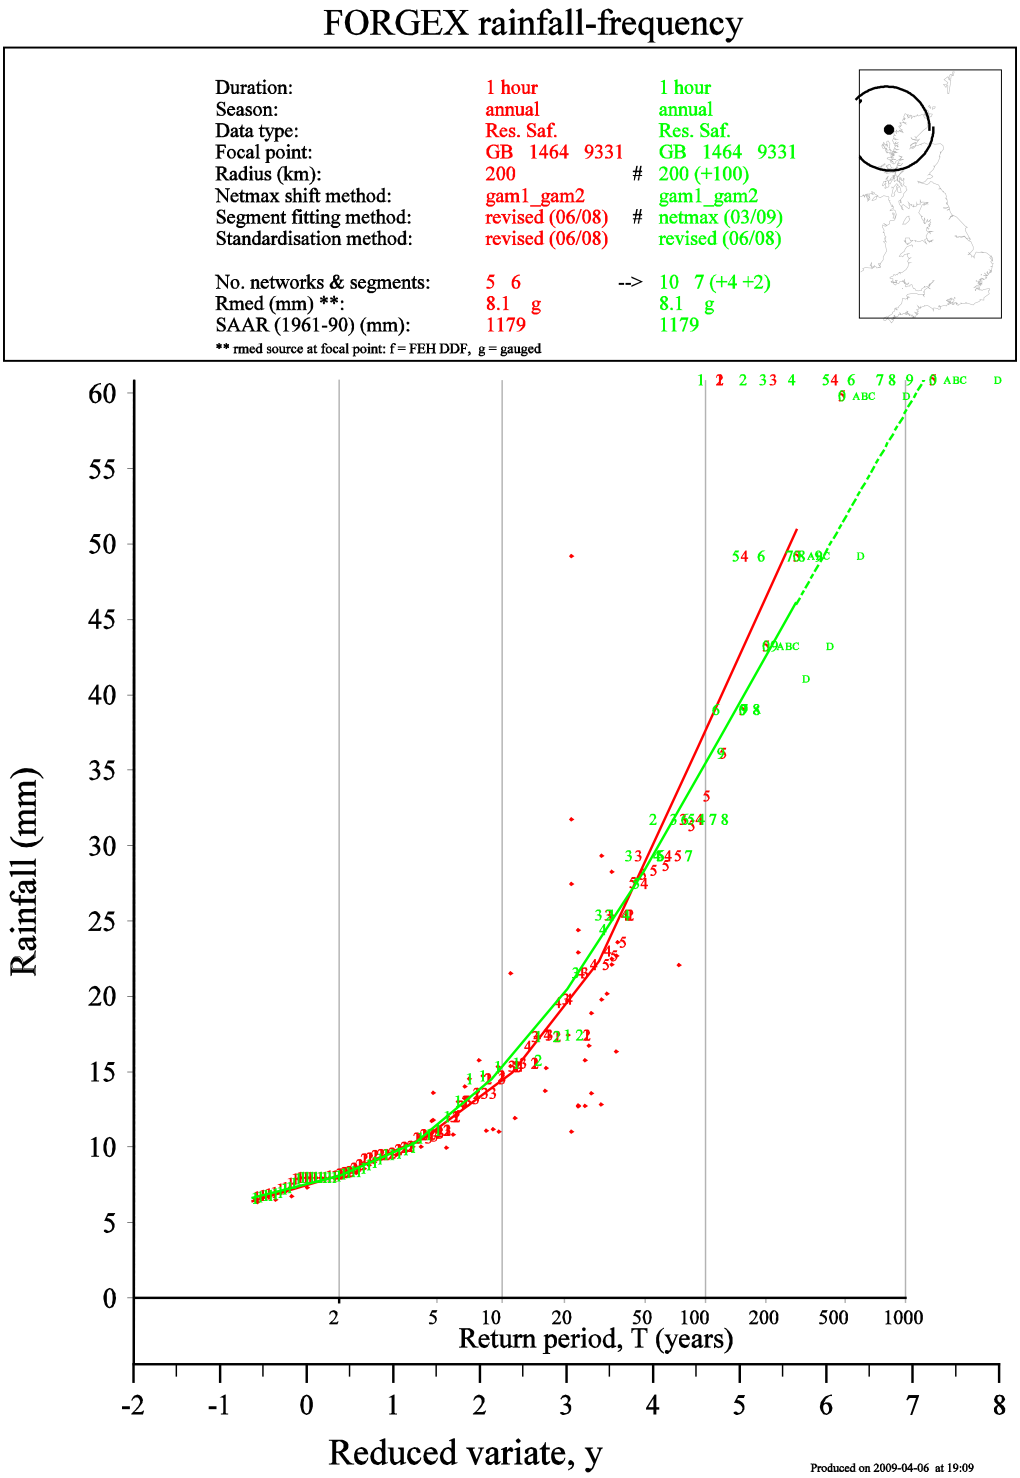

Figure 7.23 The largest negative effect at 1-hour duration of the changes to FORGEX network selection and segment fitting (the green curve is from the new method)

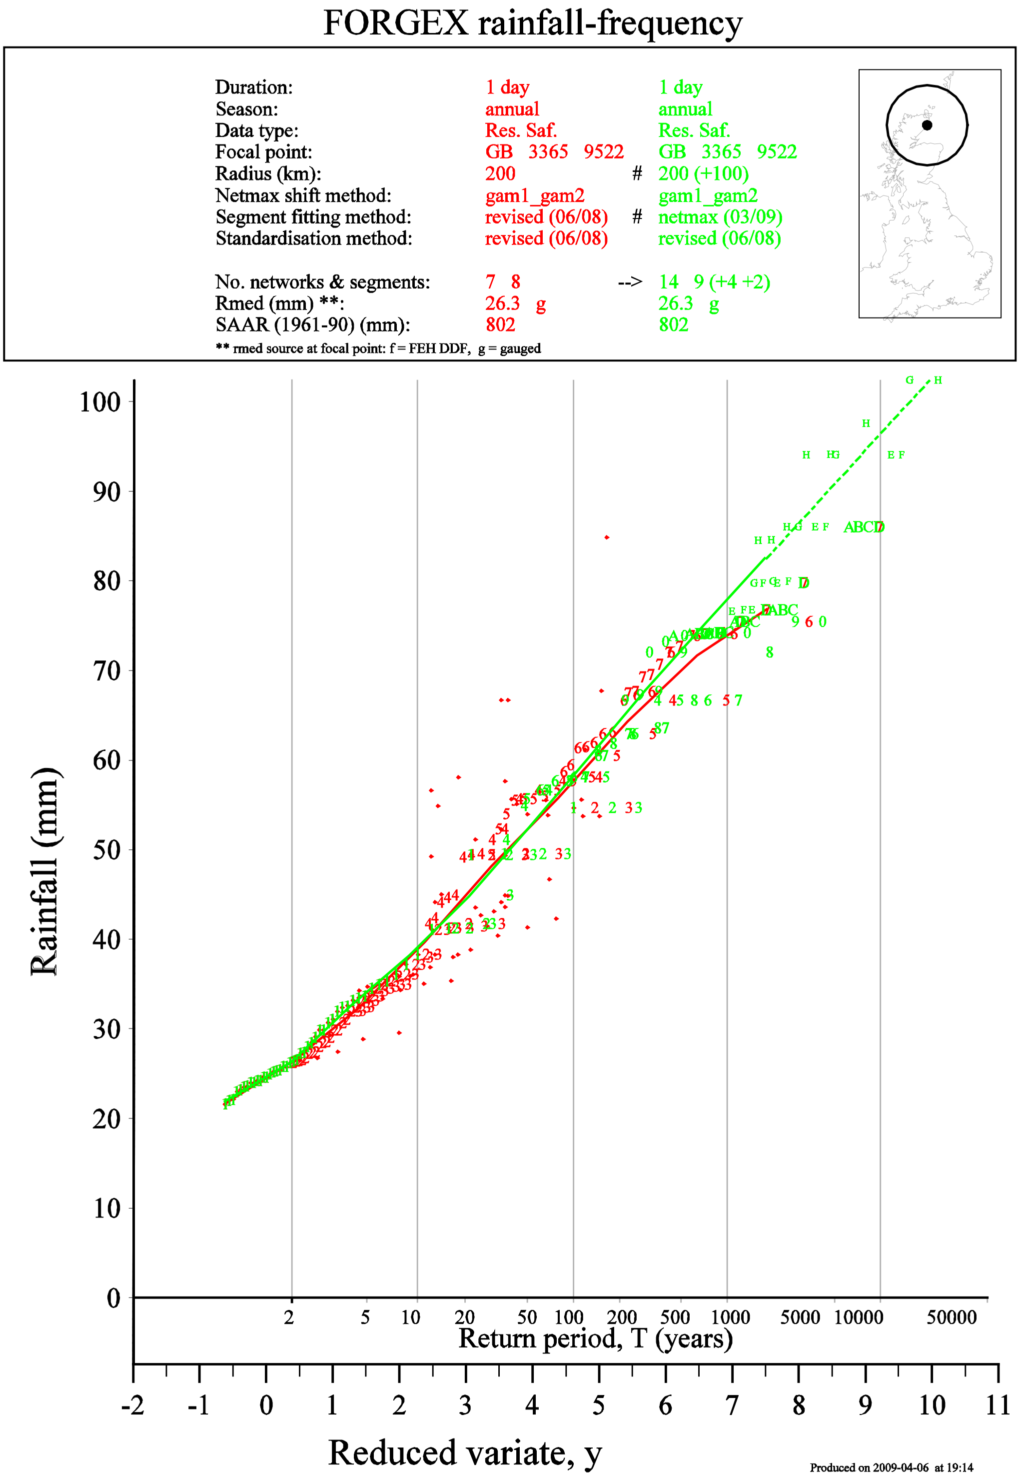

Figure 7.24 The largest positive effect at 1-day duration of the changes to FORGEX network selection and segment fitting (the green curve is from the new method)

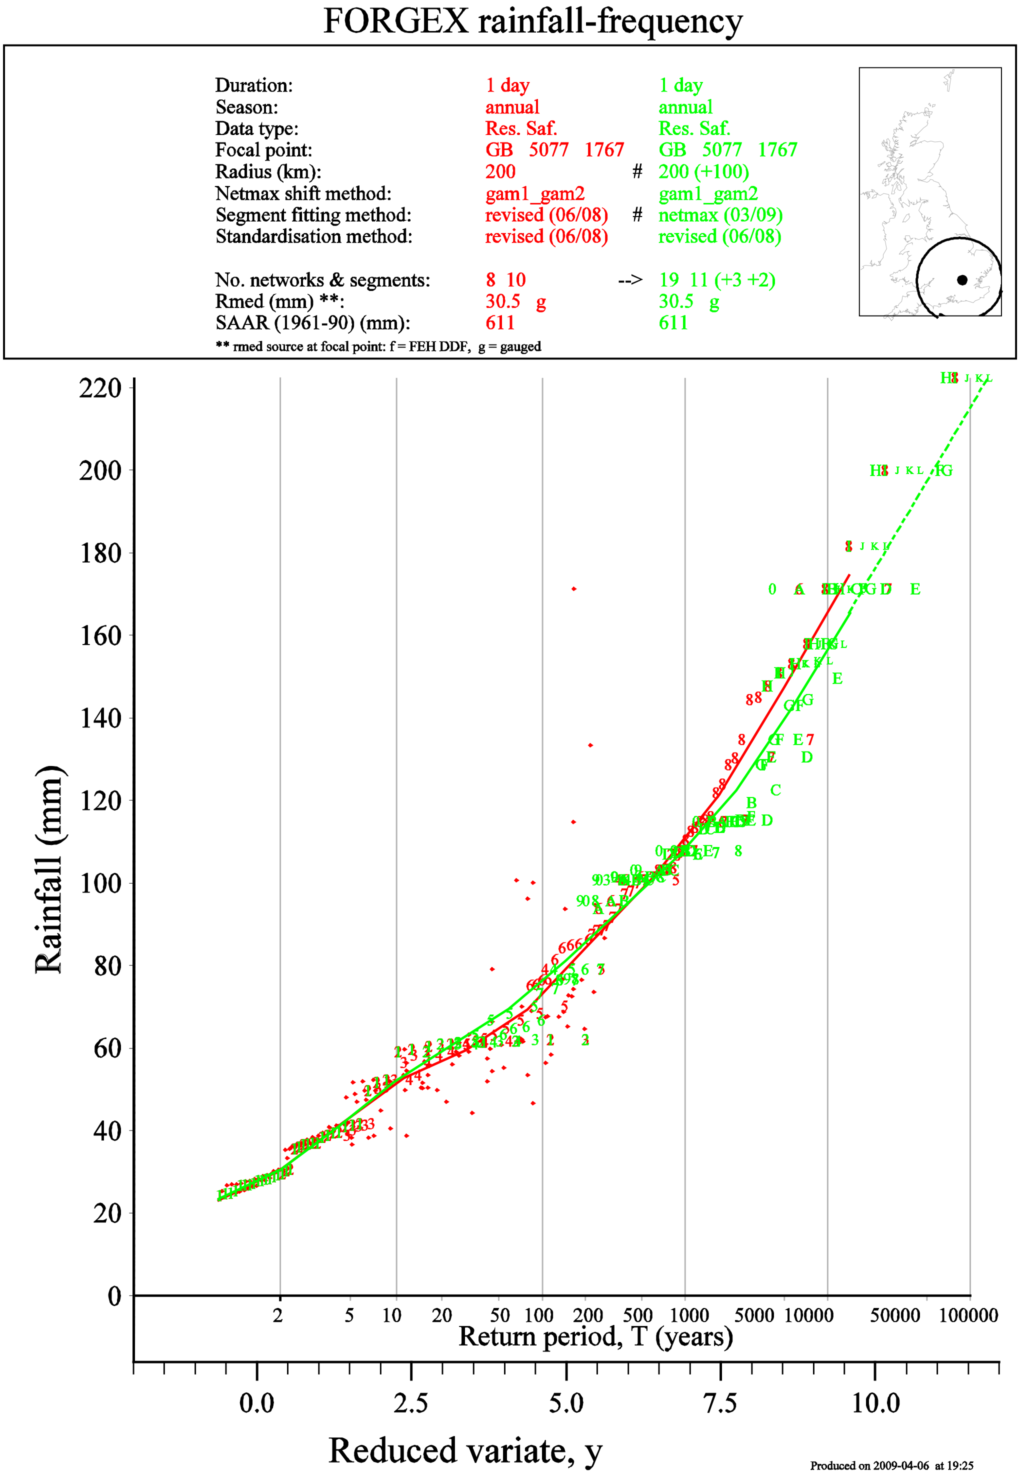

Figure 7.25 The largest negative effect at 1-day duration of the changes to FORGEX network selection and segment fitting (the green curve is from the new method)

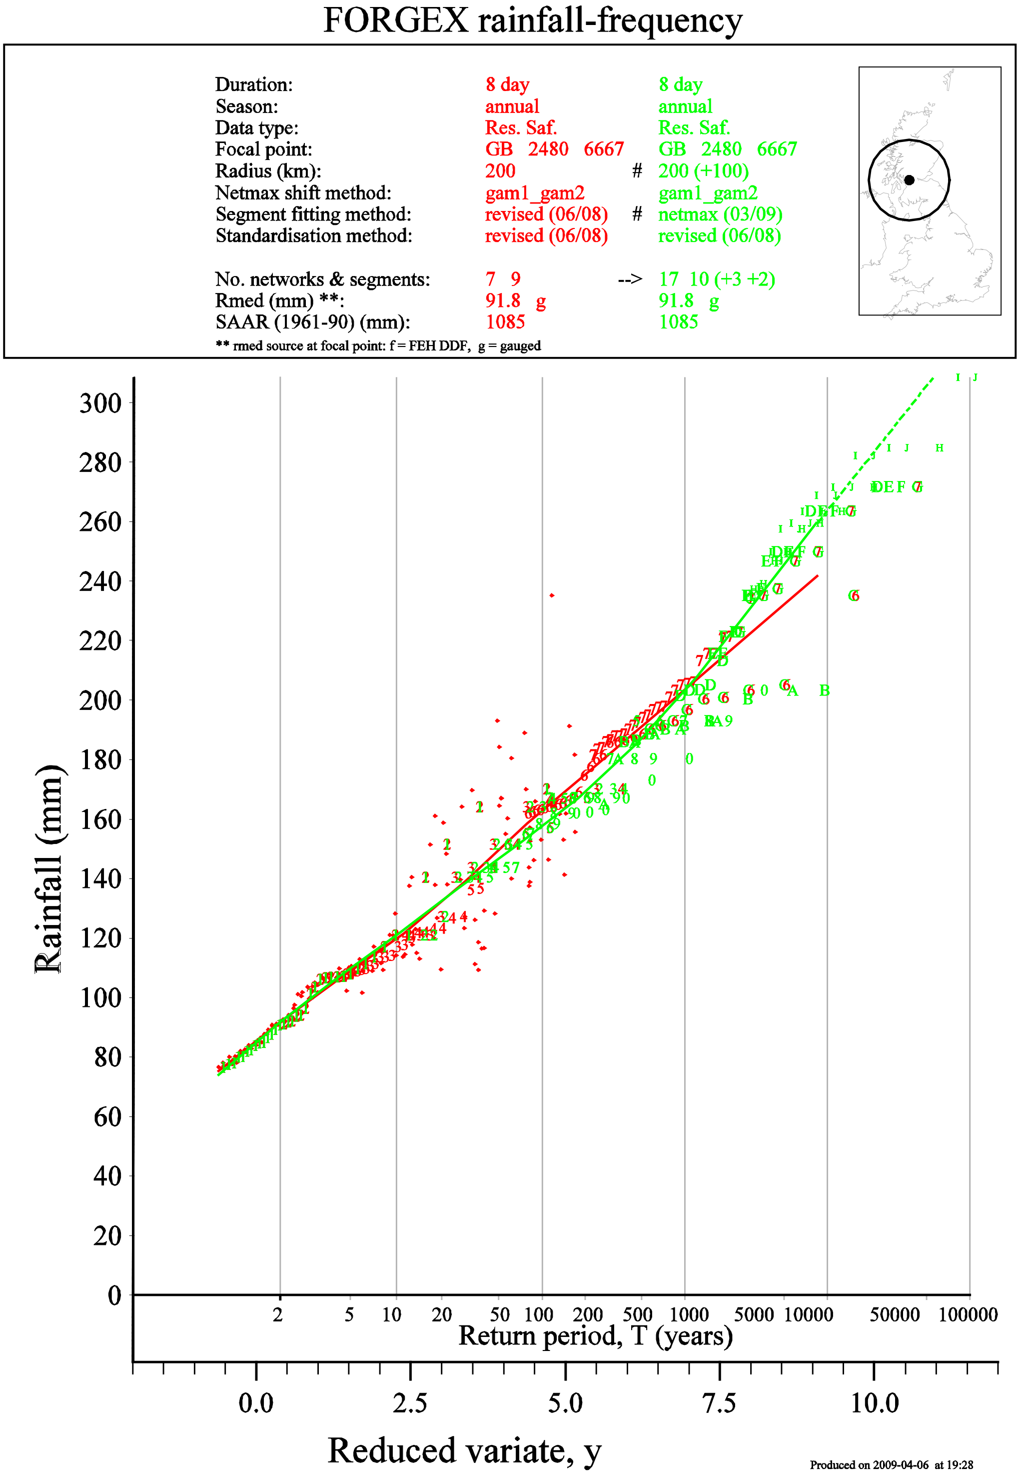

Figure 7.26 The largest positive effect at 8-day duration of the changes to FORGEX network selection and segment fitting (the green curve is from the new method)

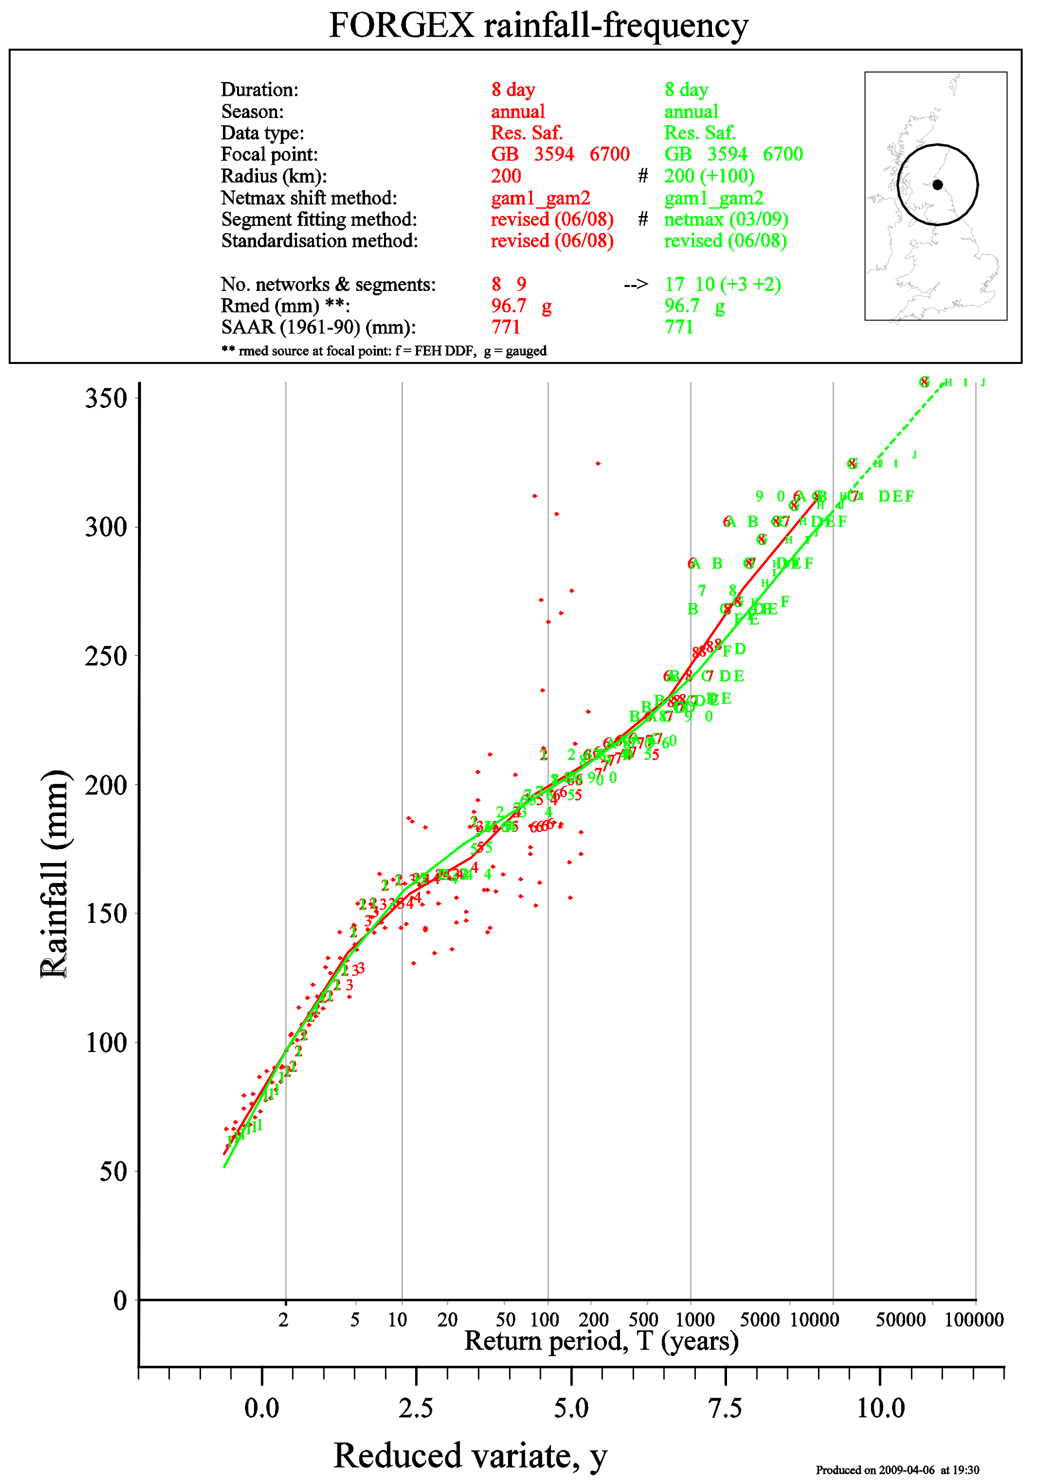

Figure 7.27 The largest negative effect at 8-day duration of the changes to FORGEX network selection and segment fitting (the green curve is from the new method)

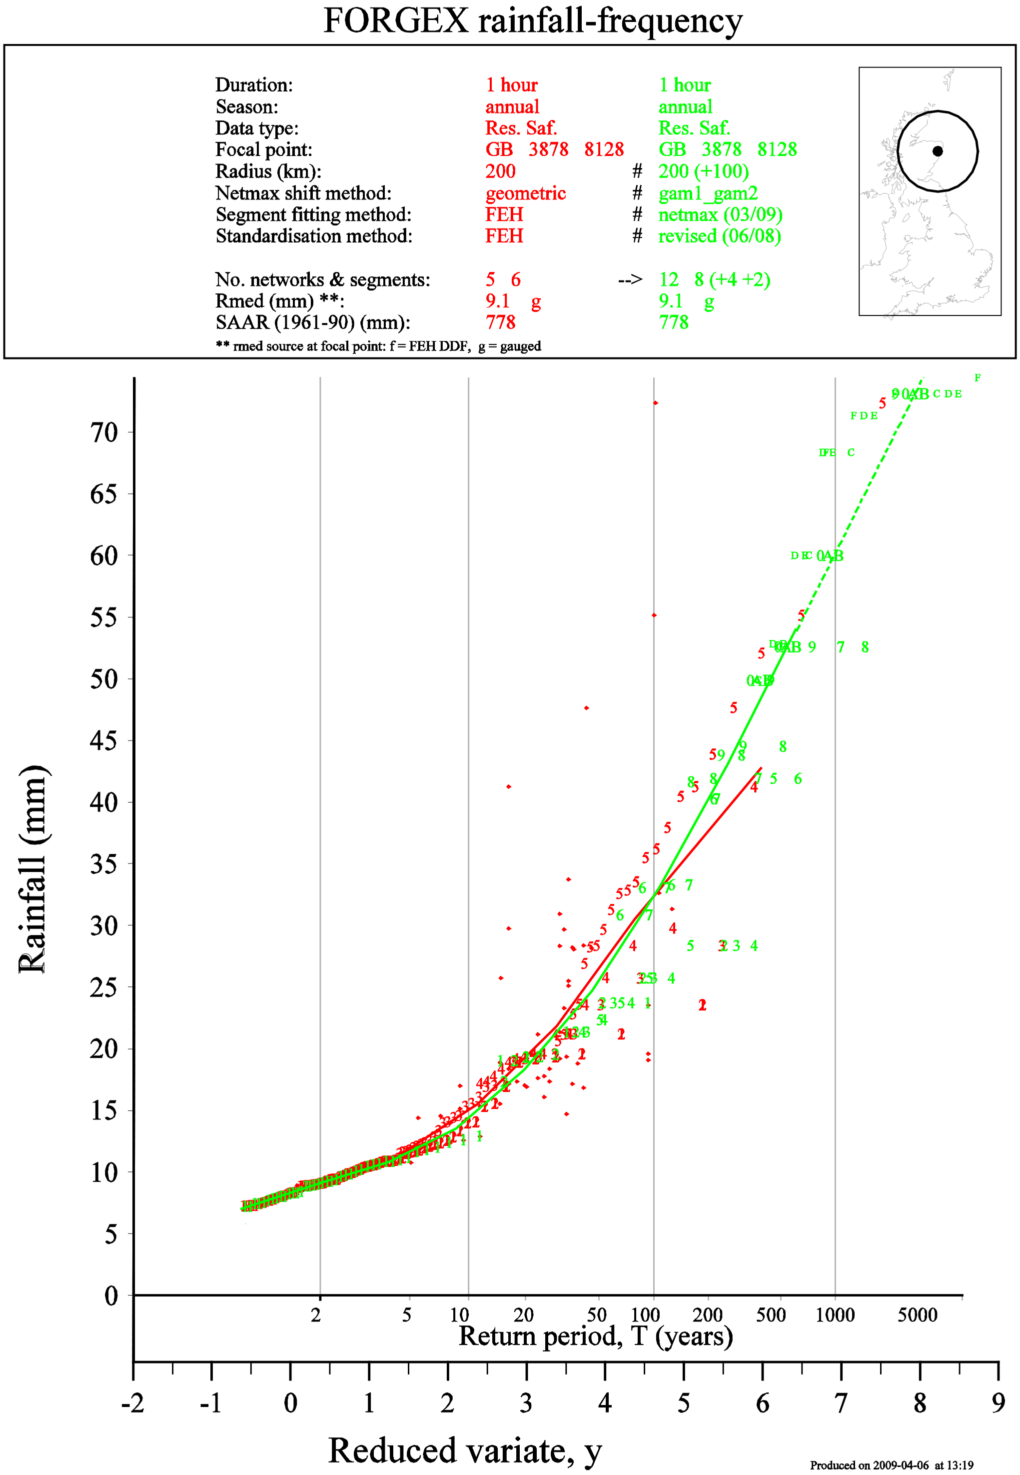

Figure 7.28 The largest positive effect at 1-hour duration of all of the changes to FORGEX (the green curve is from the new method)

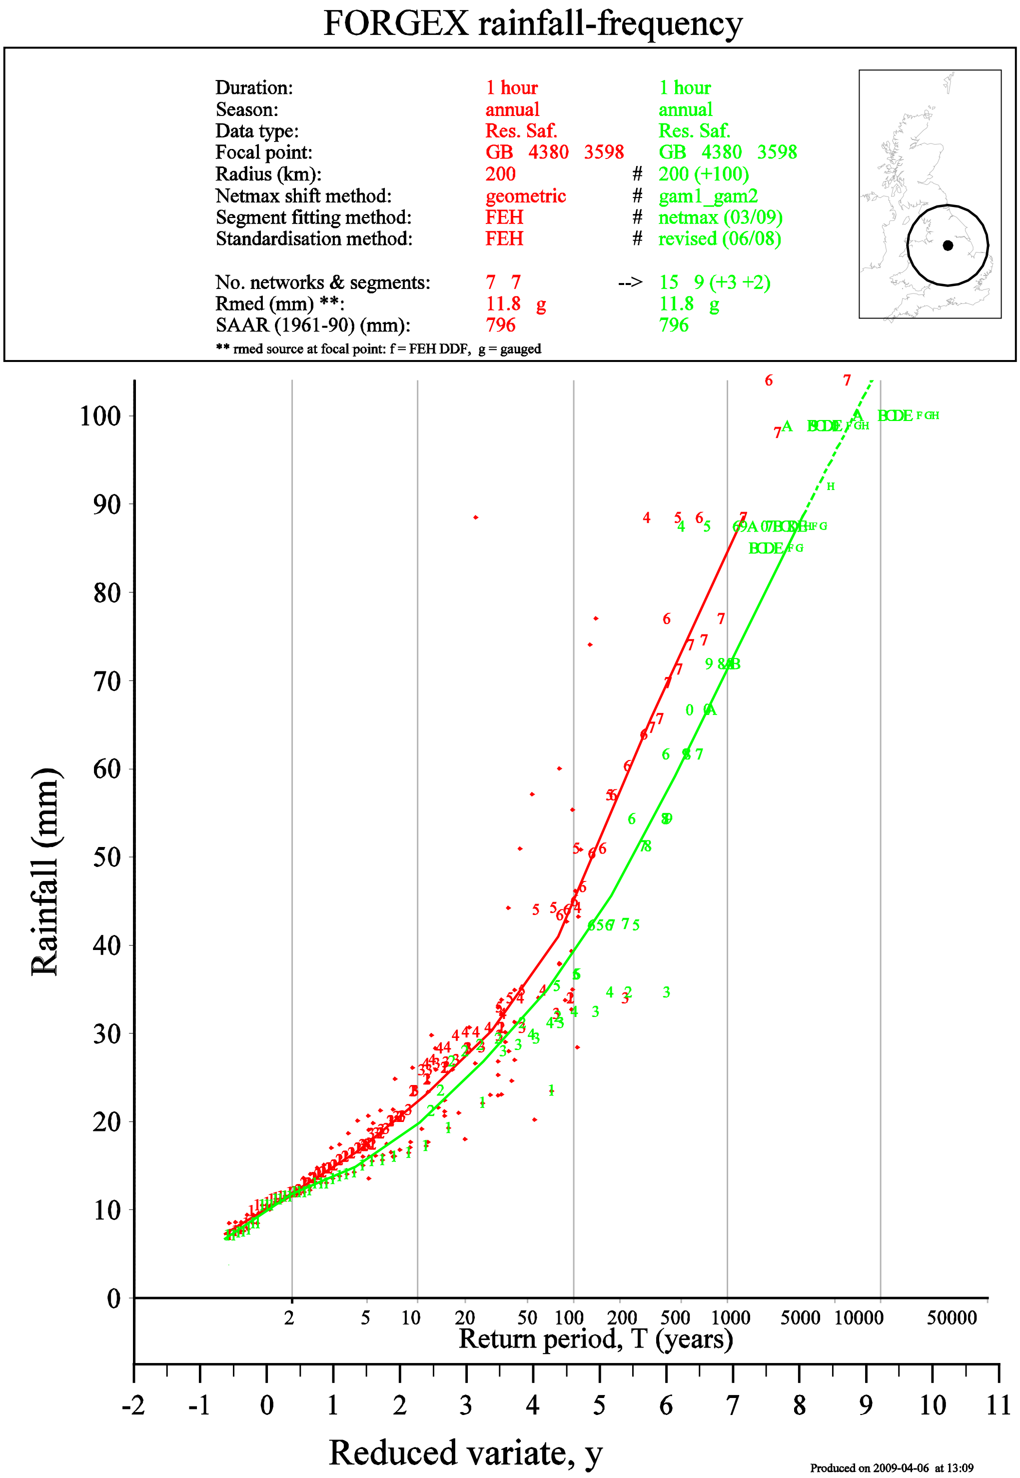

Figure 7.29 A typical example at 1-hour duration of the effect of all of the changes to FORGEX (the green curve is from the new method)

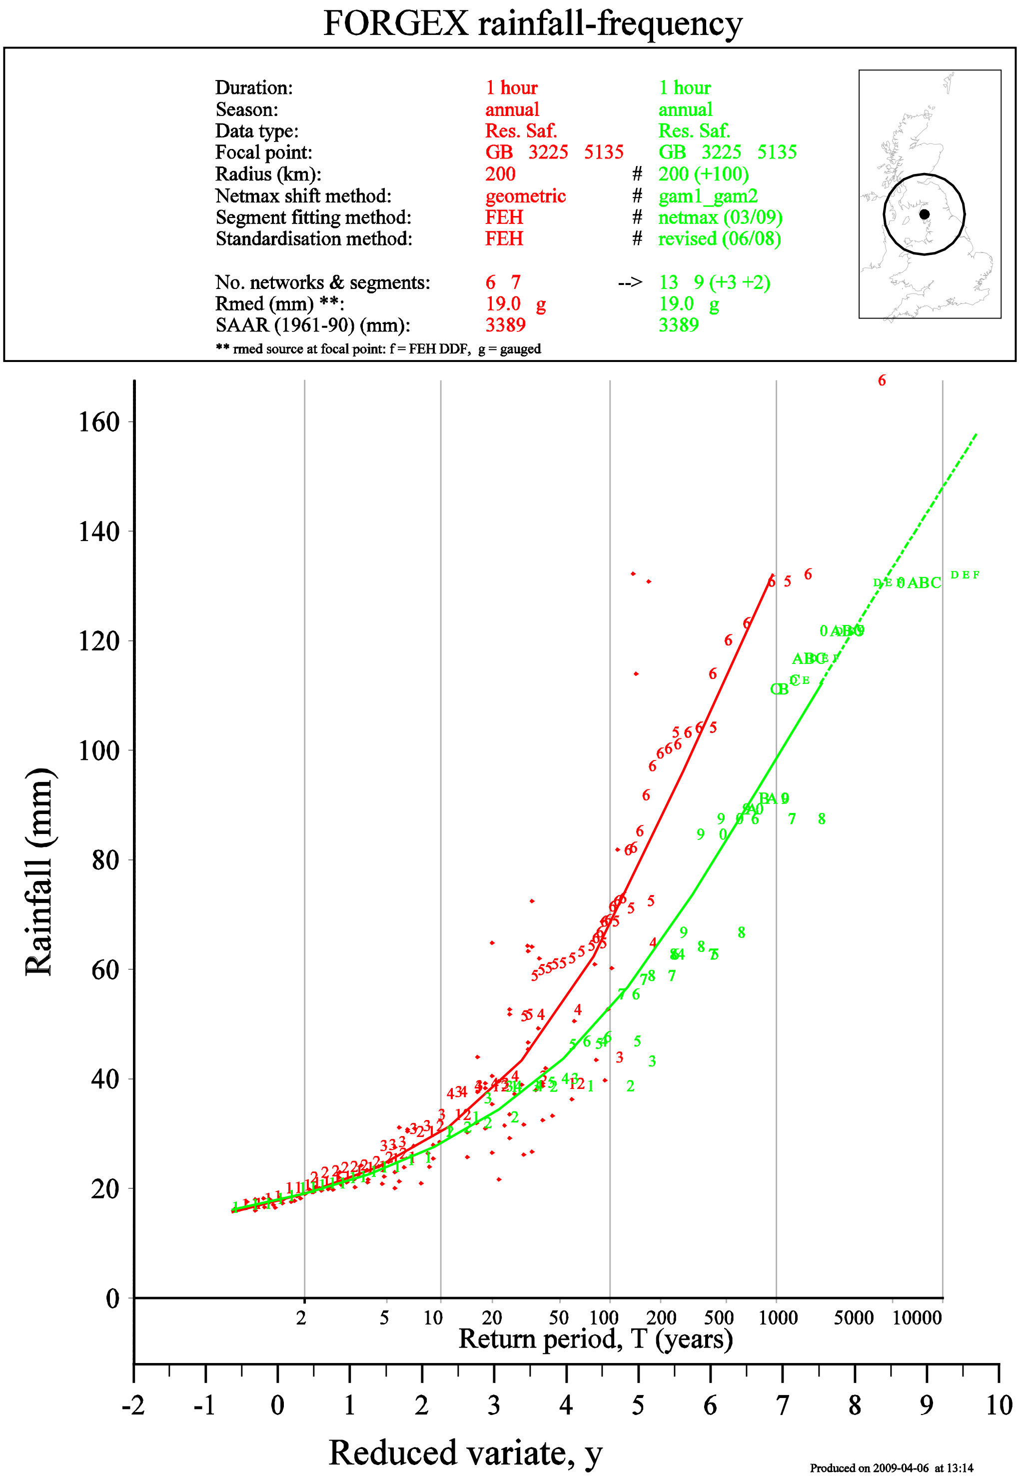

Figure 7.30 The largest negative effect at 1-hour duration of all of the changes to FORGEX (the green curve is from the new method)

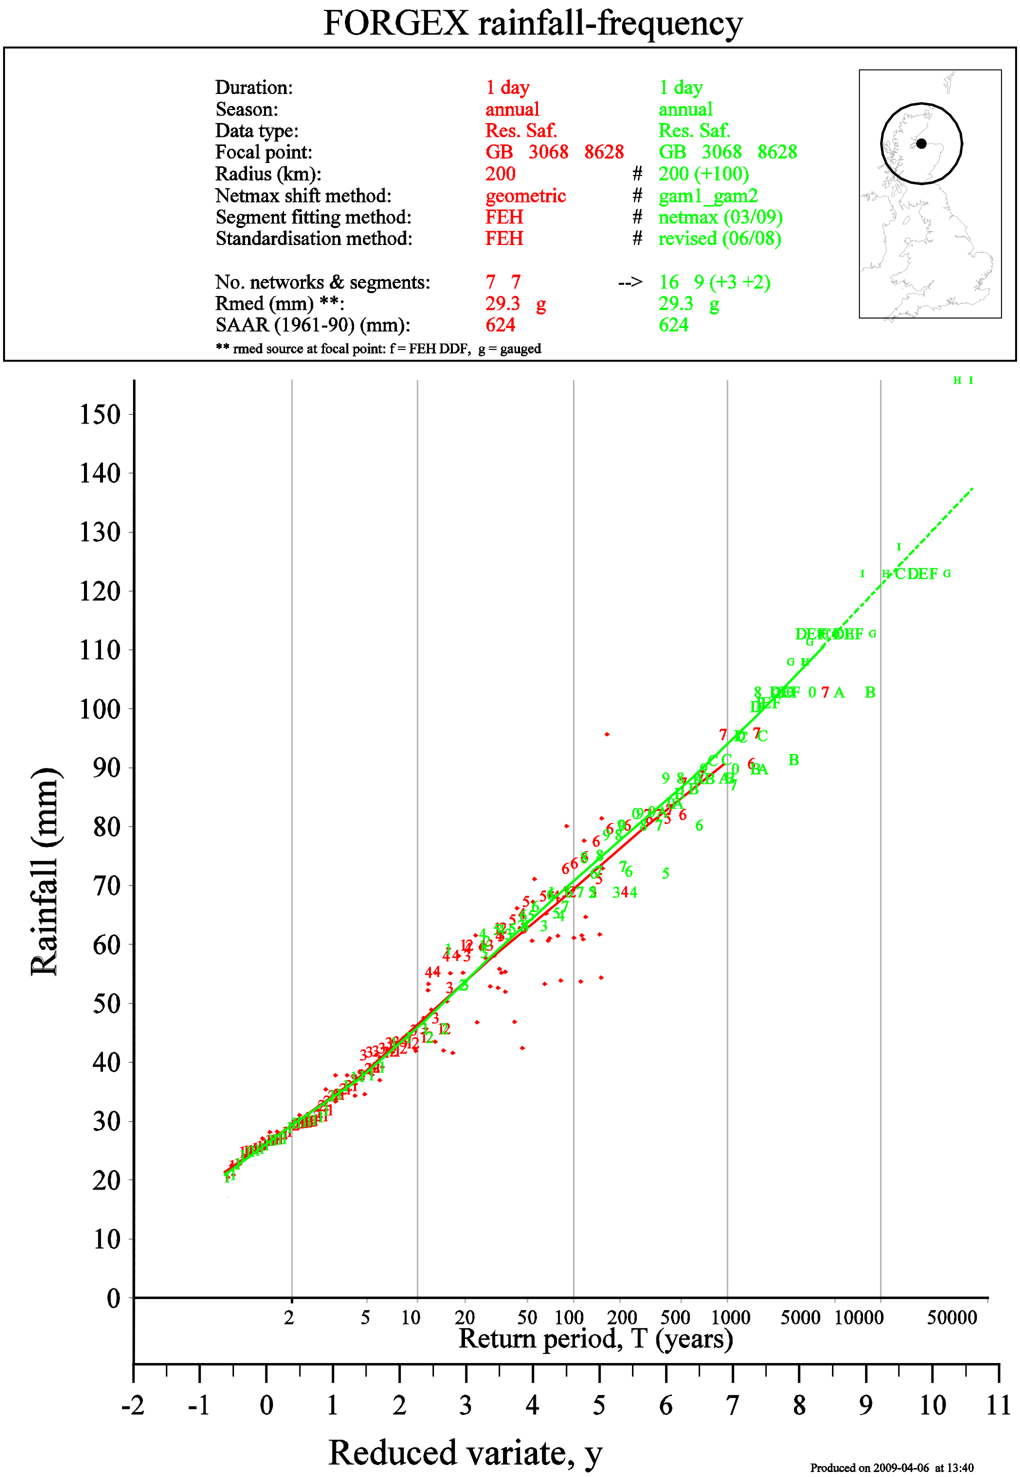

Figure 7.31 The largest positive effect at 1-day duration of all of the changes to FORGEX (the green curve is from the new method)

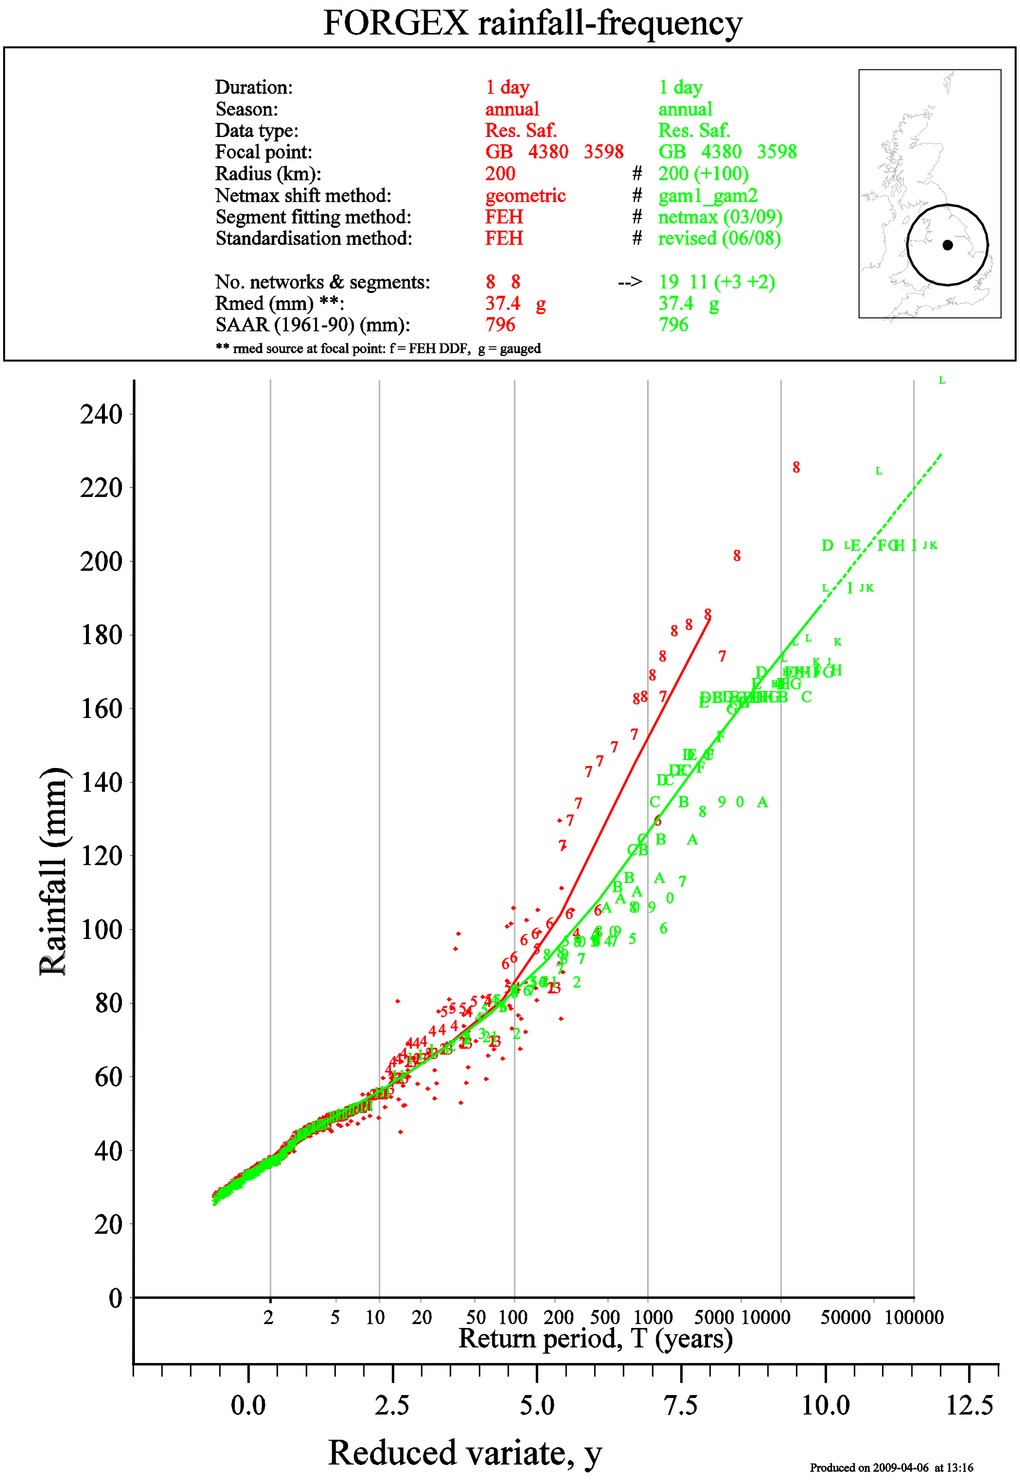

Figure 7.32 A typical example at 1-day duration of the effect of all of the changes to FORGEX (the green curve is from the new method)

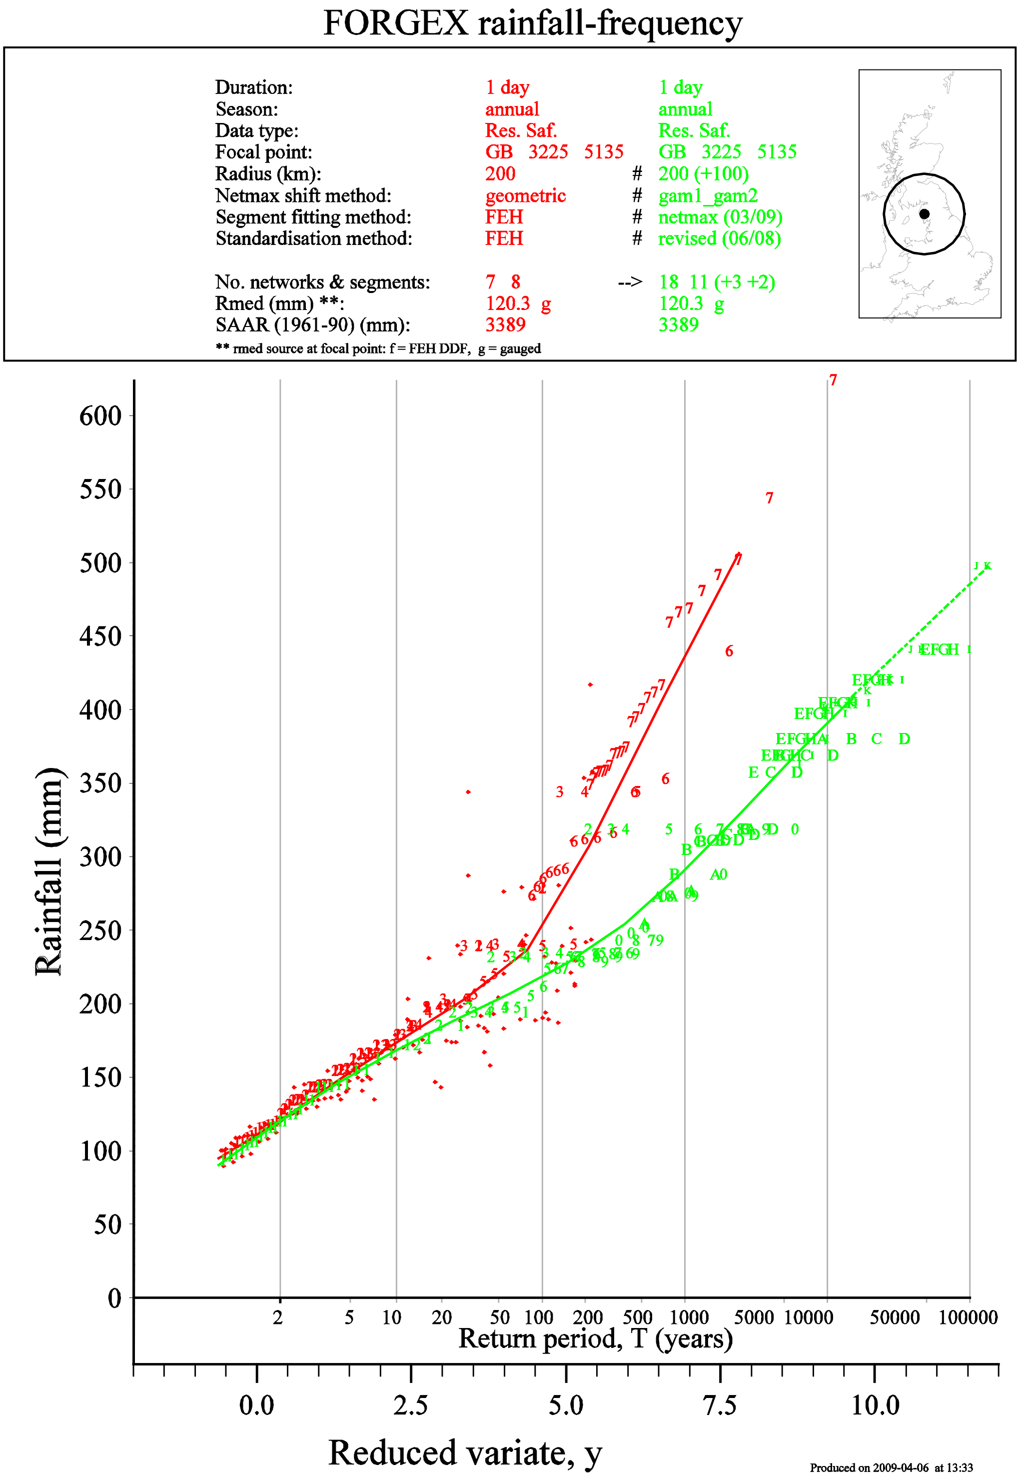

Figure 7.33 The largest negative effect at 1-day duration of all of the changes to FORGEX (the green curve is from the new method)

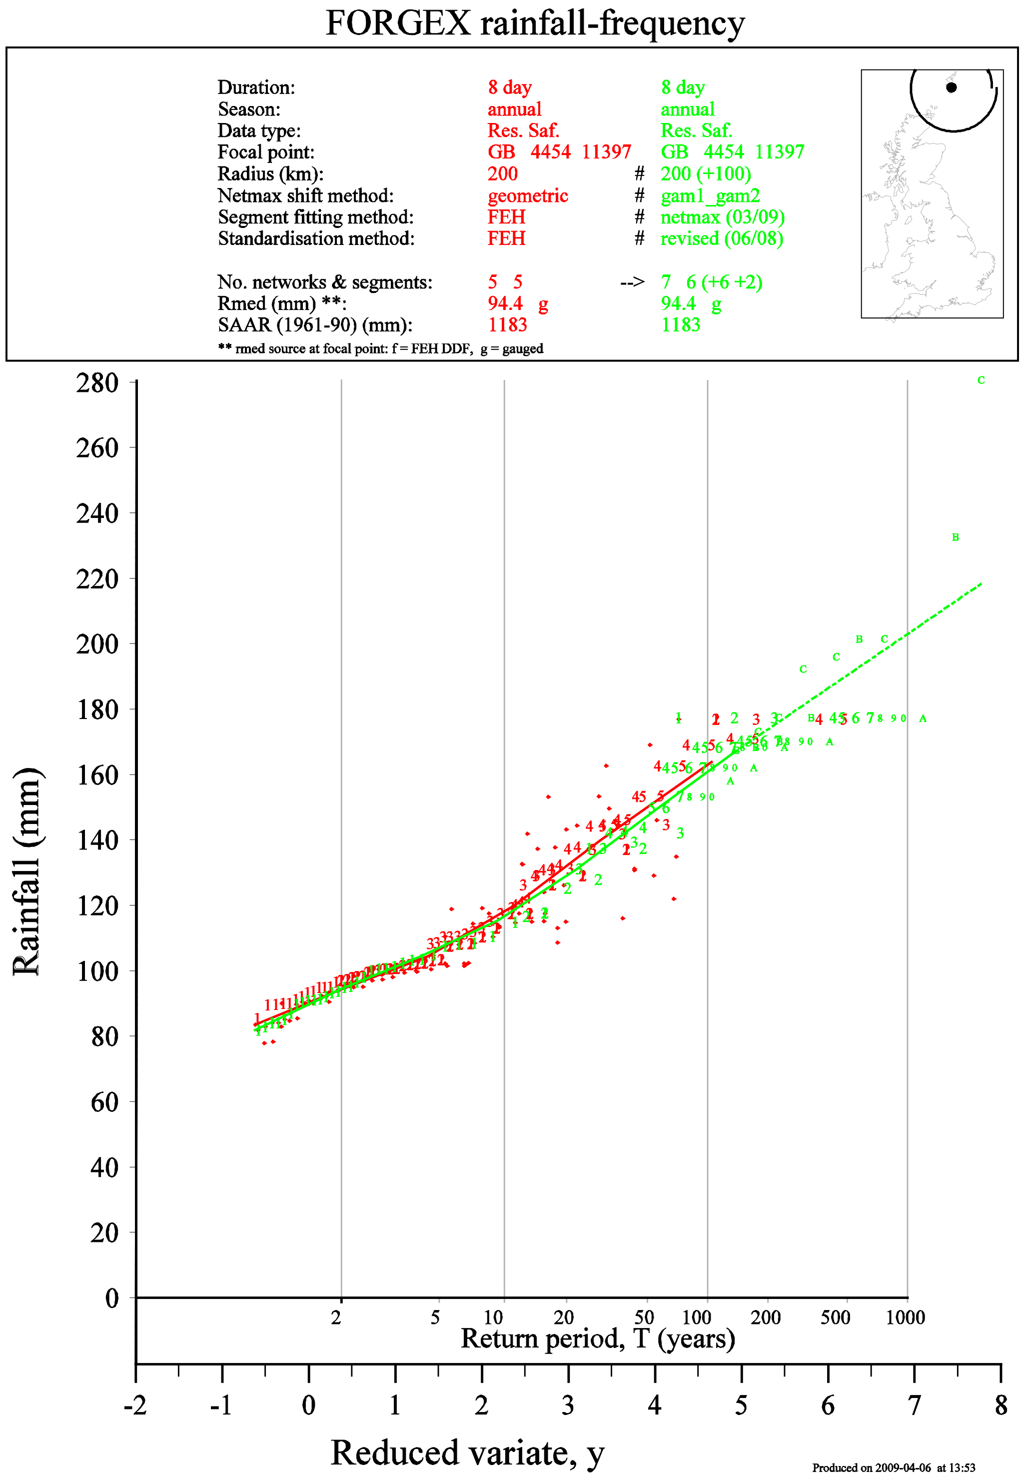

Figure 7.34 The smallest effect at 8-day duration of all of the changes to FORGEX (the green curve is from the new method)

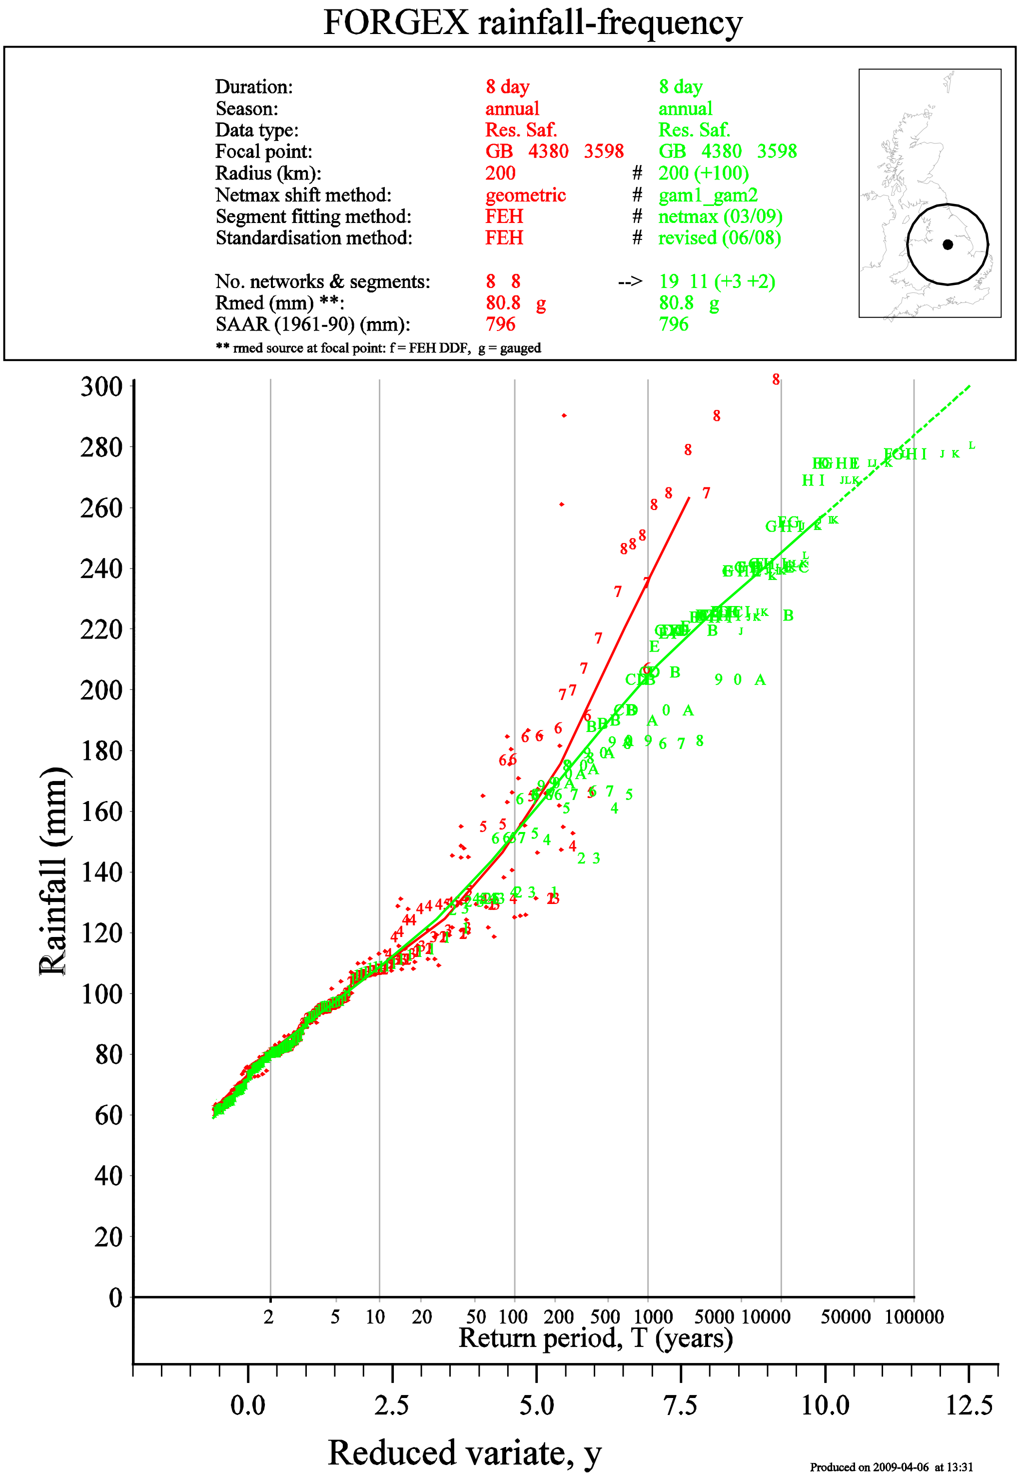

Figure 7.35 A typical example at 8-day duration of the effect of all of the changes to FORGEX (the green curve is from the new method)

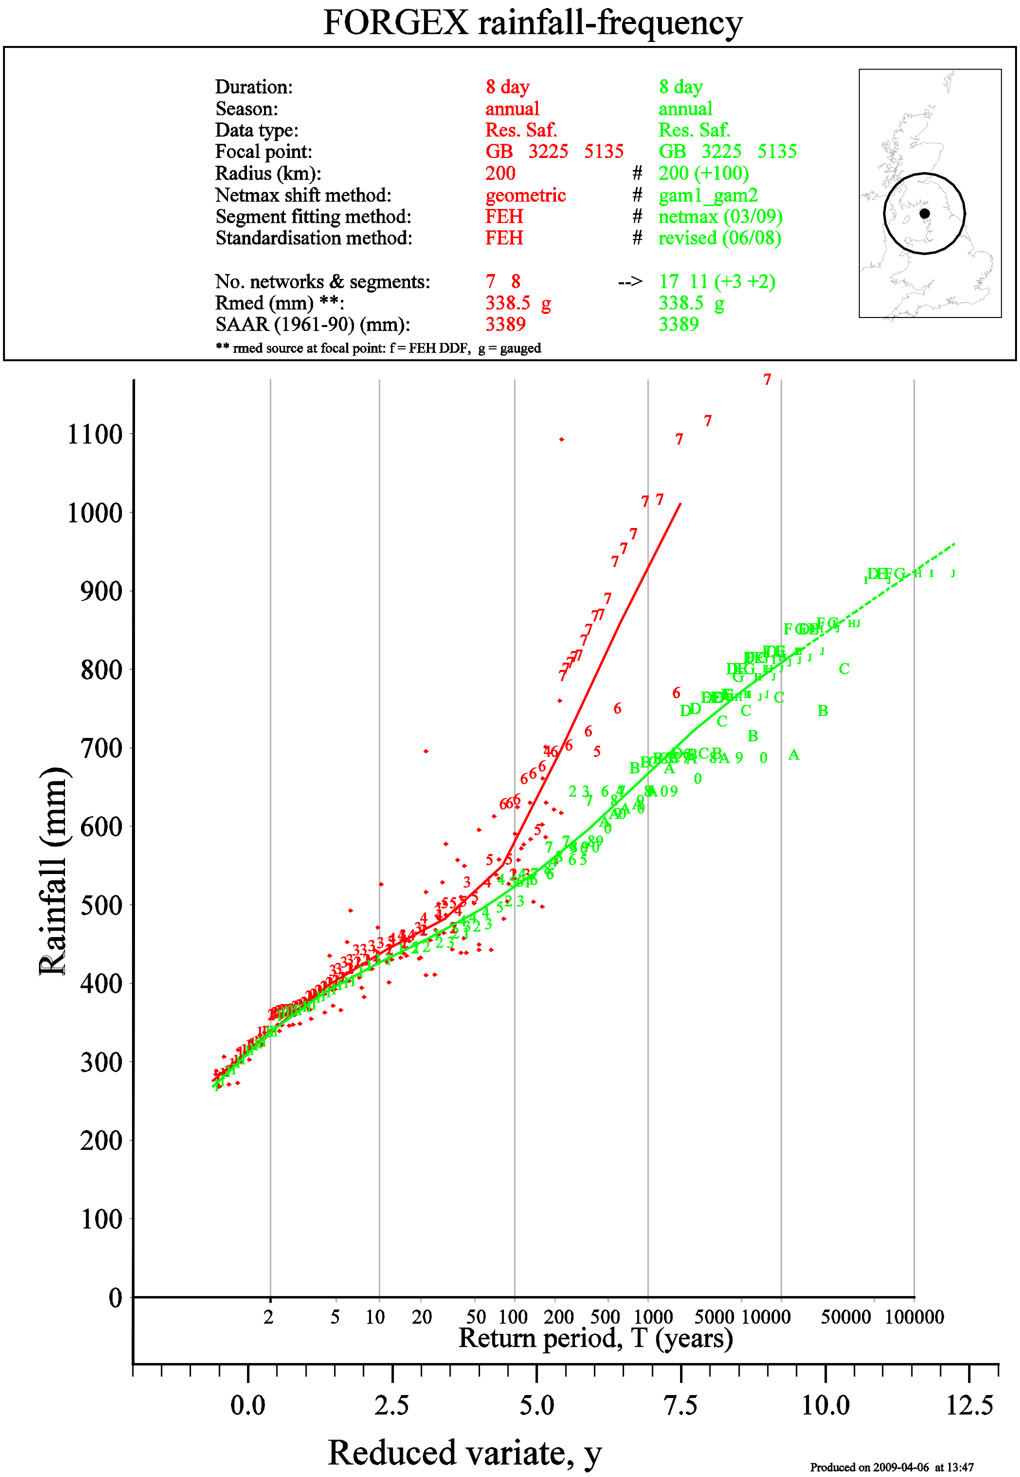

Figure 7.36 The largest negative effect at 8-day duration of all of the changes to FORGEX (the green curve is from the new method)

Finally, for four varied UK locations, Figures 7.37 to 7.40 comprise pairs of multi-duration (one-hour to eight-day) plots, with the FEH methods on the left (a), faced by the new method on the right (b). The latter represent the data that are taken forward to construct the DDF model.

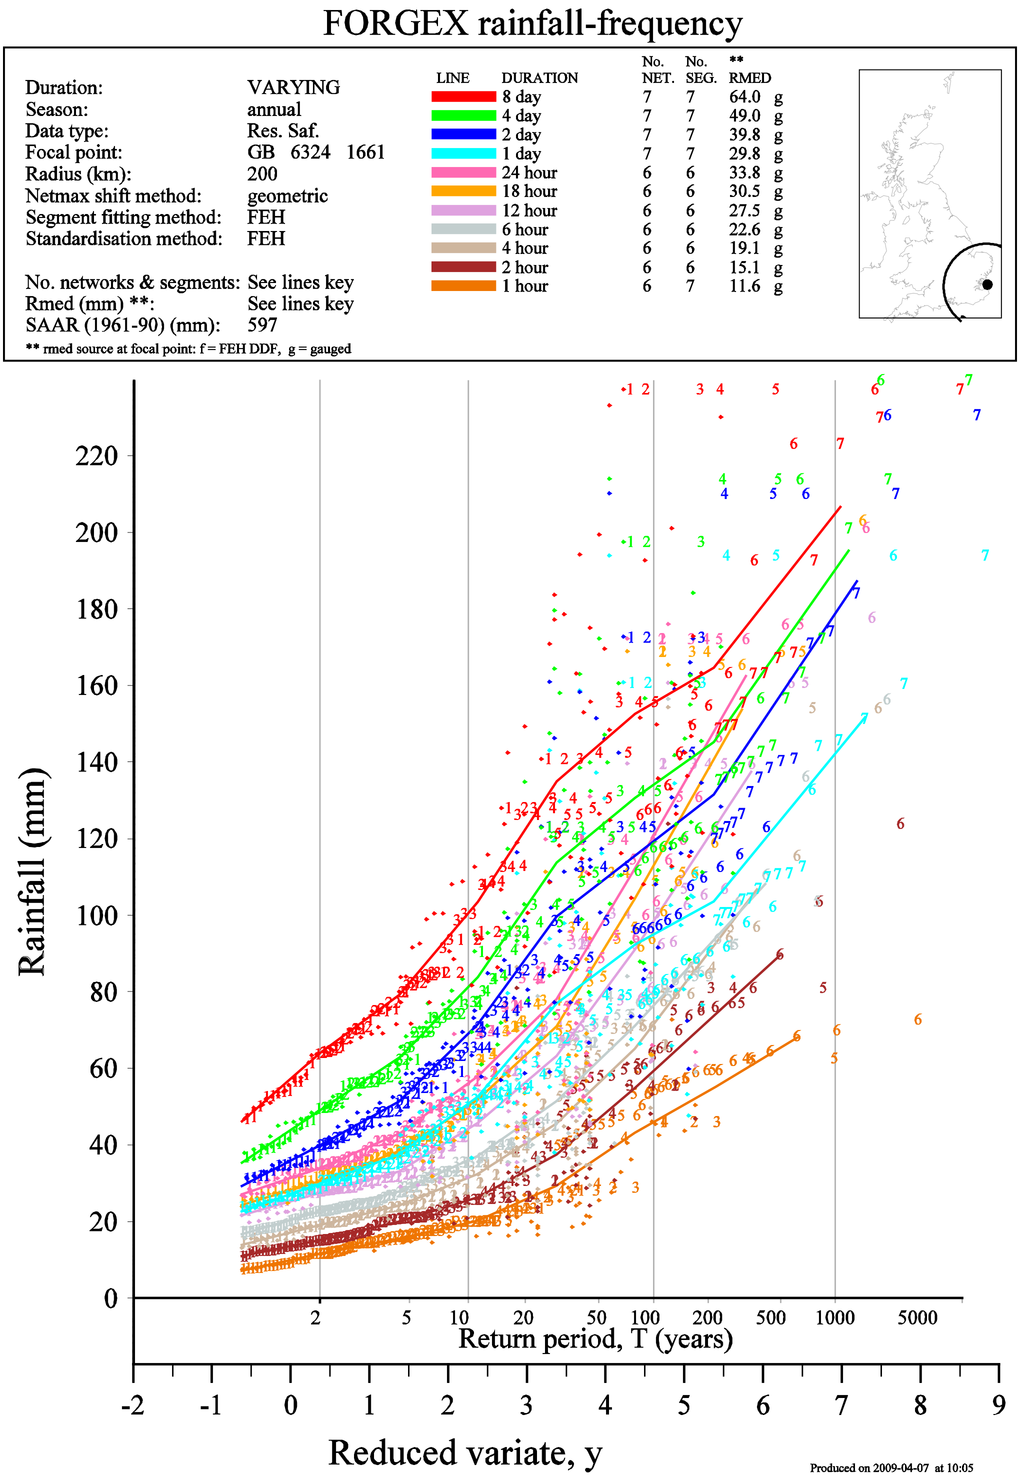

Figure 7.37(a) Rainfall frequency curves for durations ranging from 1 hour to 8 days for a location in south-east England. Produced using the FEH FORGEX procedures.

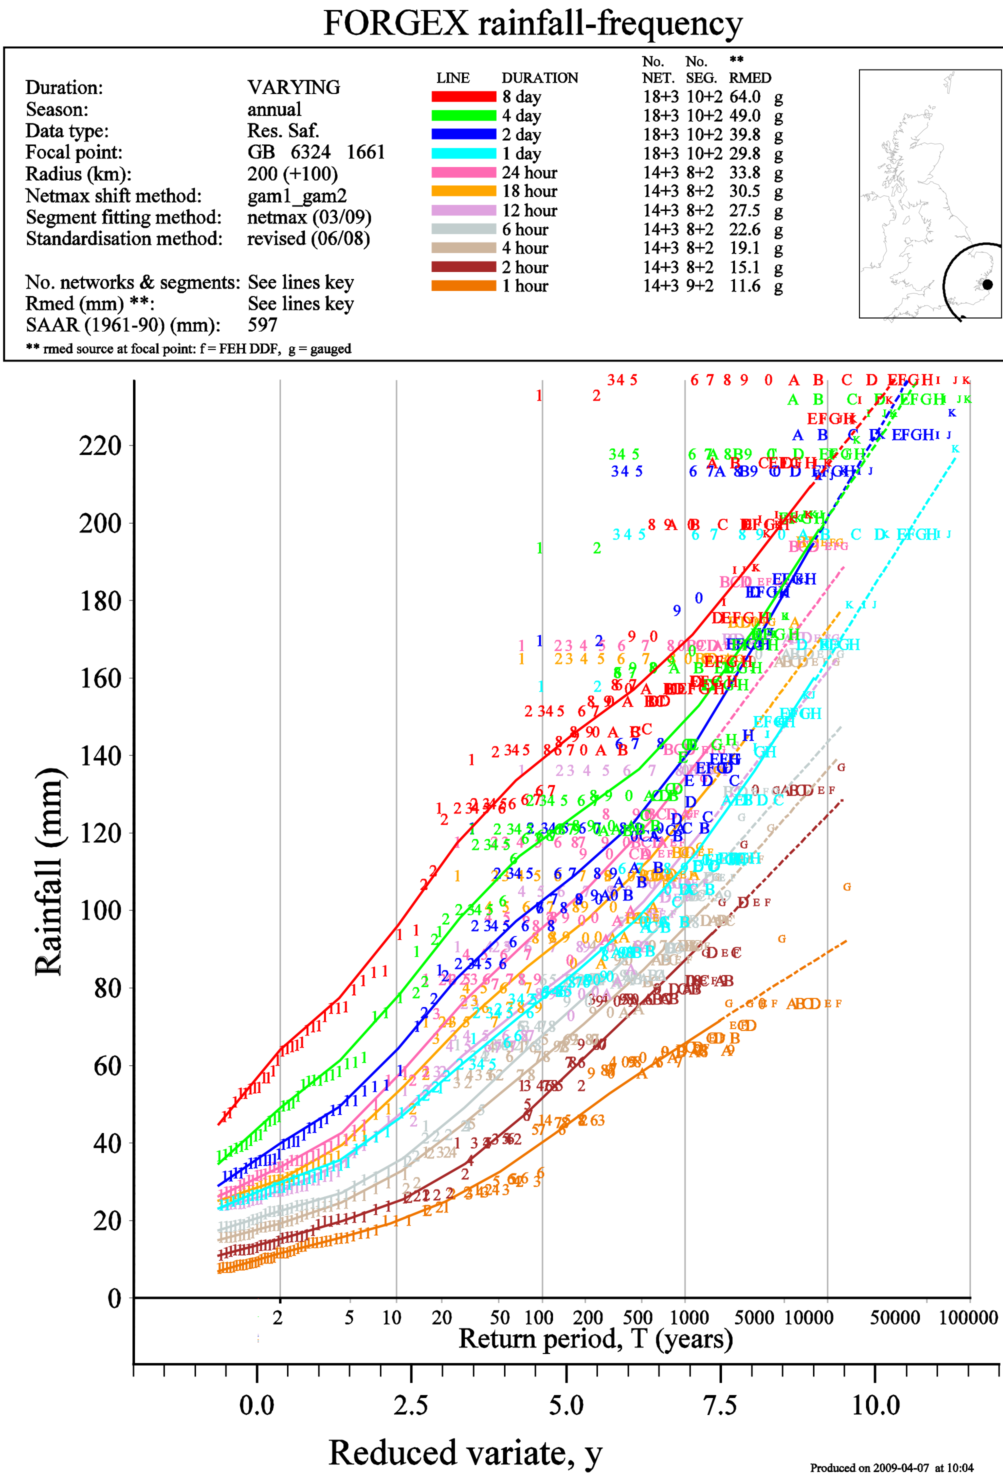

Figure 7.37(b) Rainfall frequency curves for durations ranging from 1 hour to 8 days for a location in south-east England. Produced using the new FORGEX procedures.

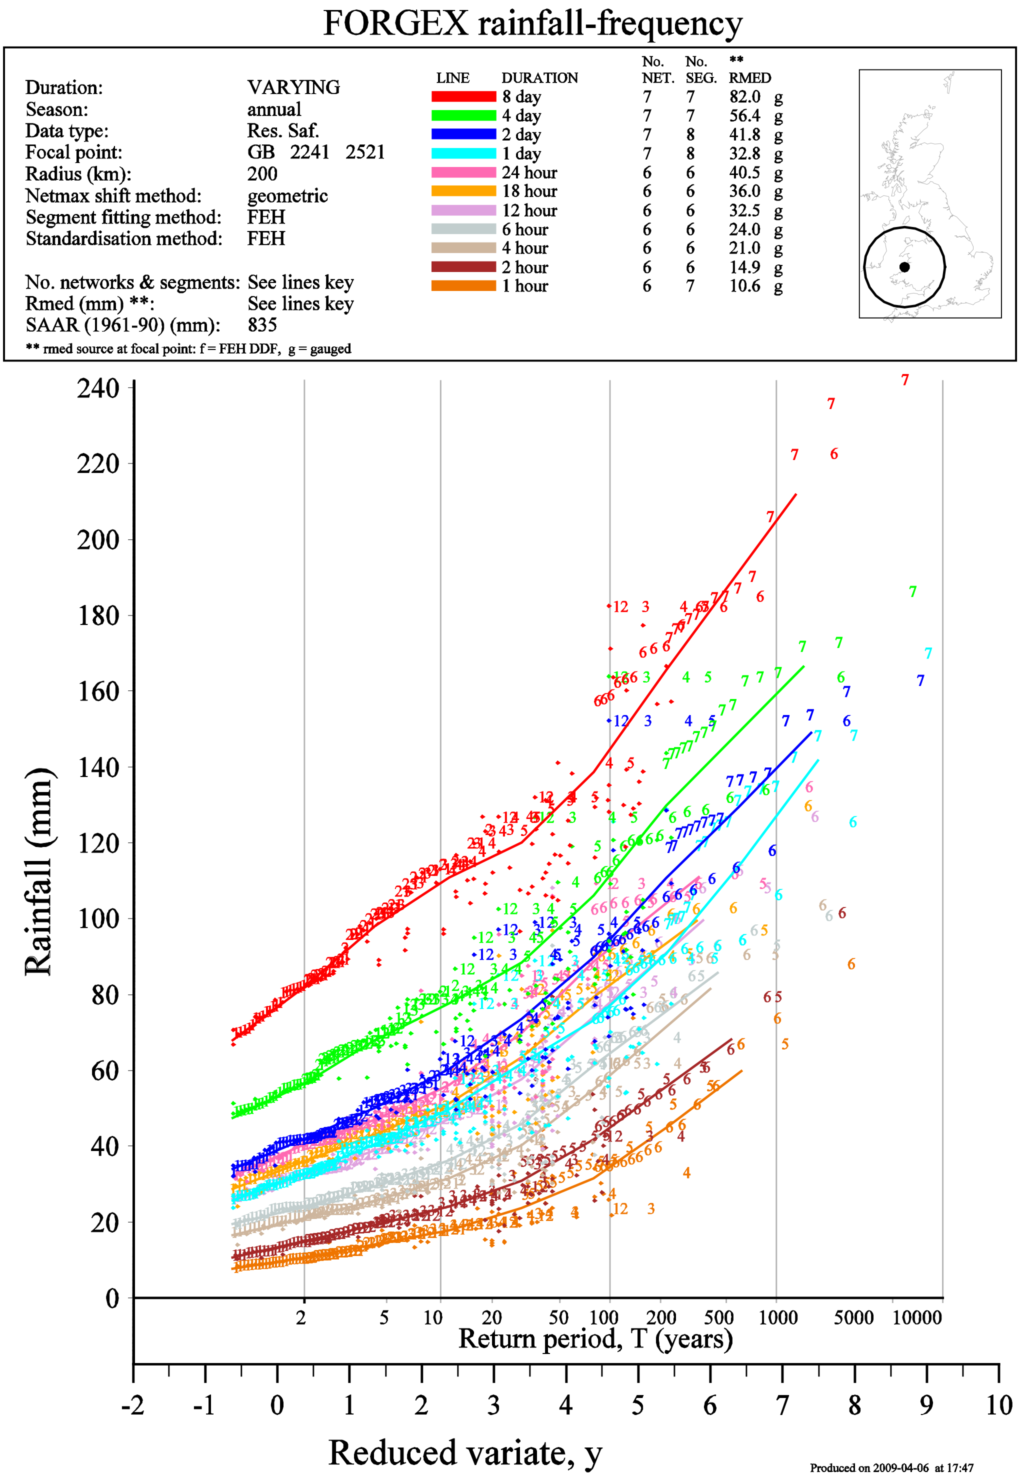

Figure 7.38(a) Rainfall frequency curves for durations ranging from 1 hour to 8 days for a location in Wales. Produced using the FEH FORGEX procedures.

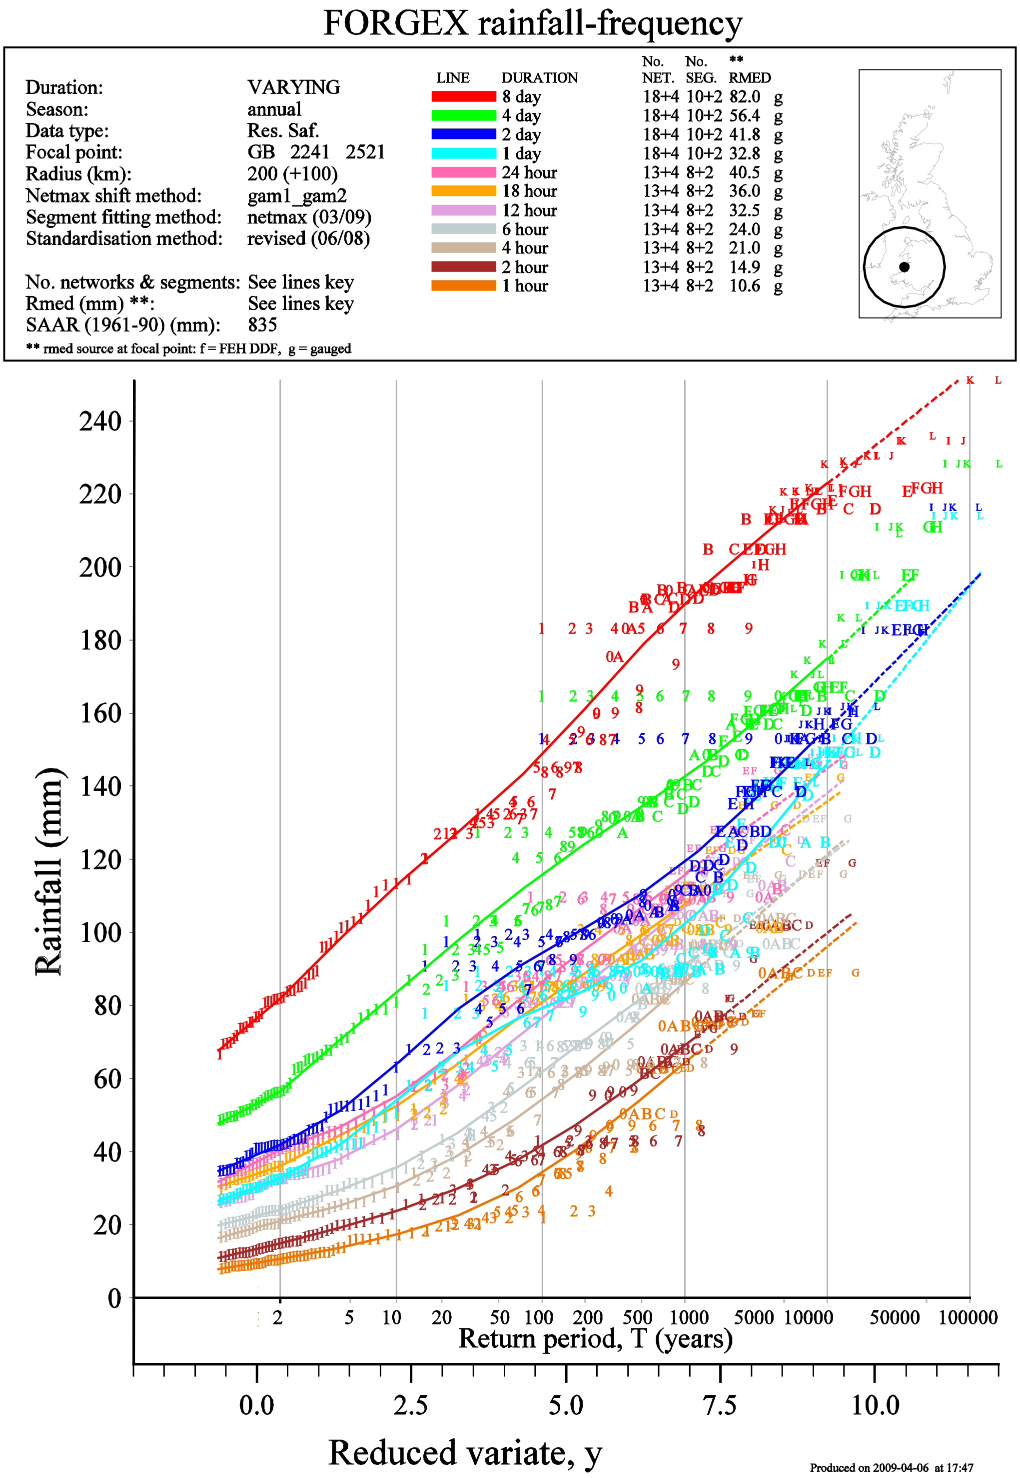

Figure 7.38(b) Rainfall frequency curves for durations ranging from 1 hour to 8 days for a location in Wales. Produced using the new FORGEX procedures.

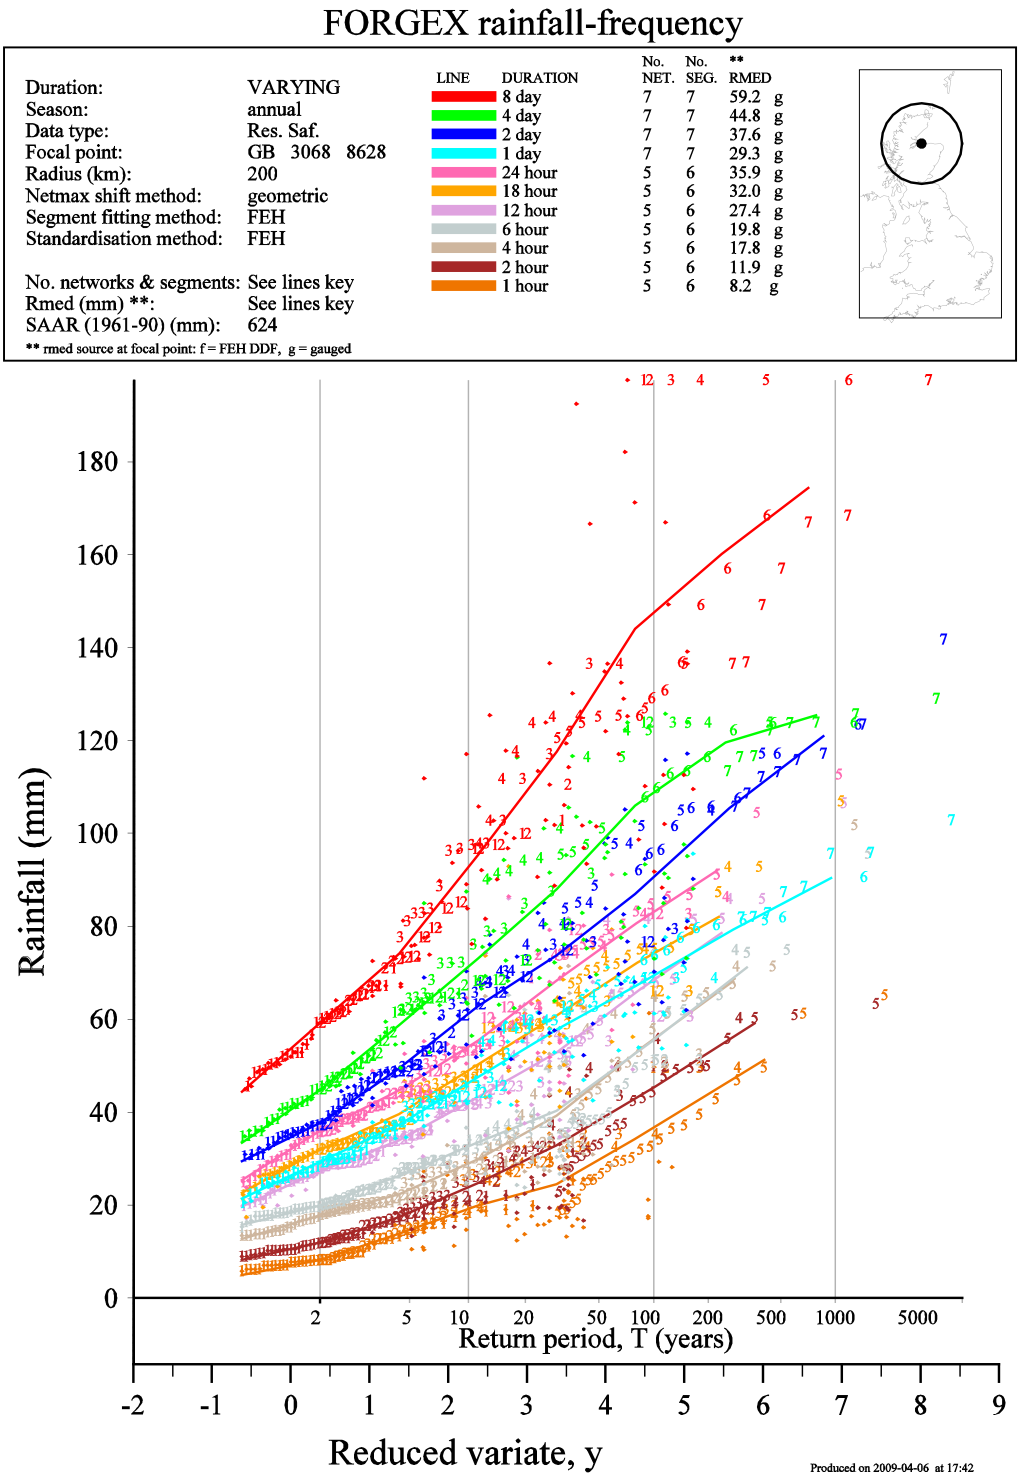

Figure 7.39(a) Rainfall frequency curves for durations ranging from 1 hour to 8 days for a location in north-east Scotland. Produced using the FEH FORGEX procedures.

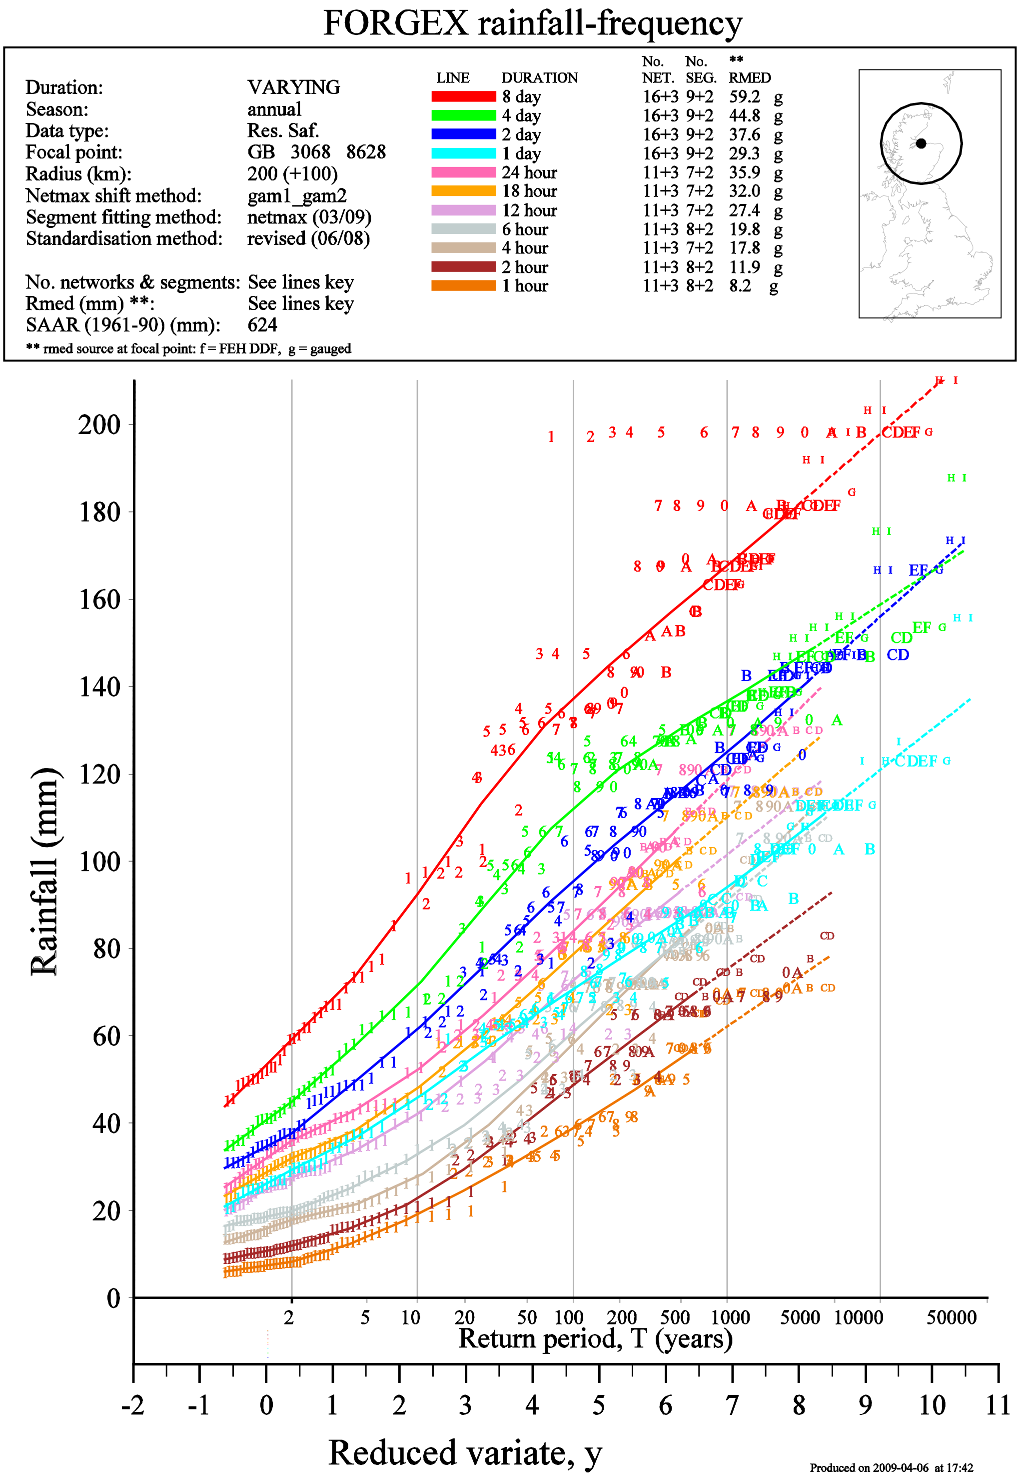

Figure 7.39(b) Rainfall frequency curves for durations ranging from 1 hour to 8 days for a location in north-east Scotland. Produced using the new FORGEX procedures.

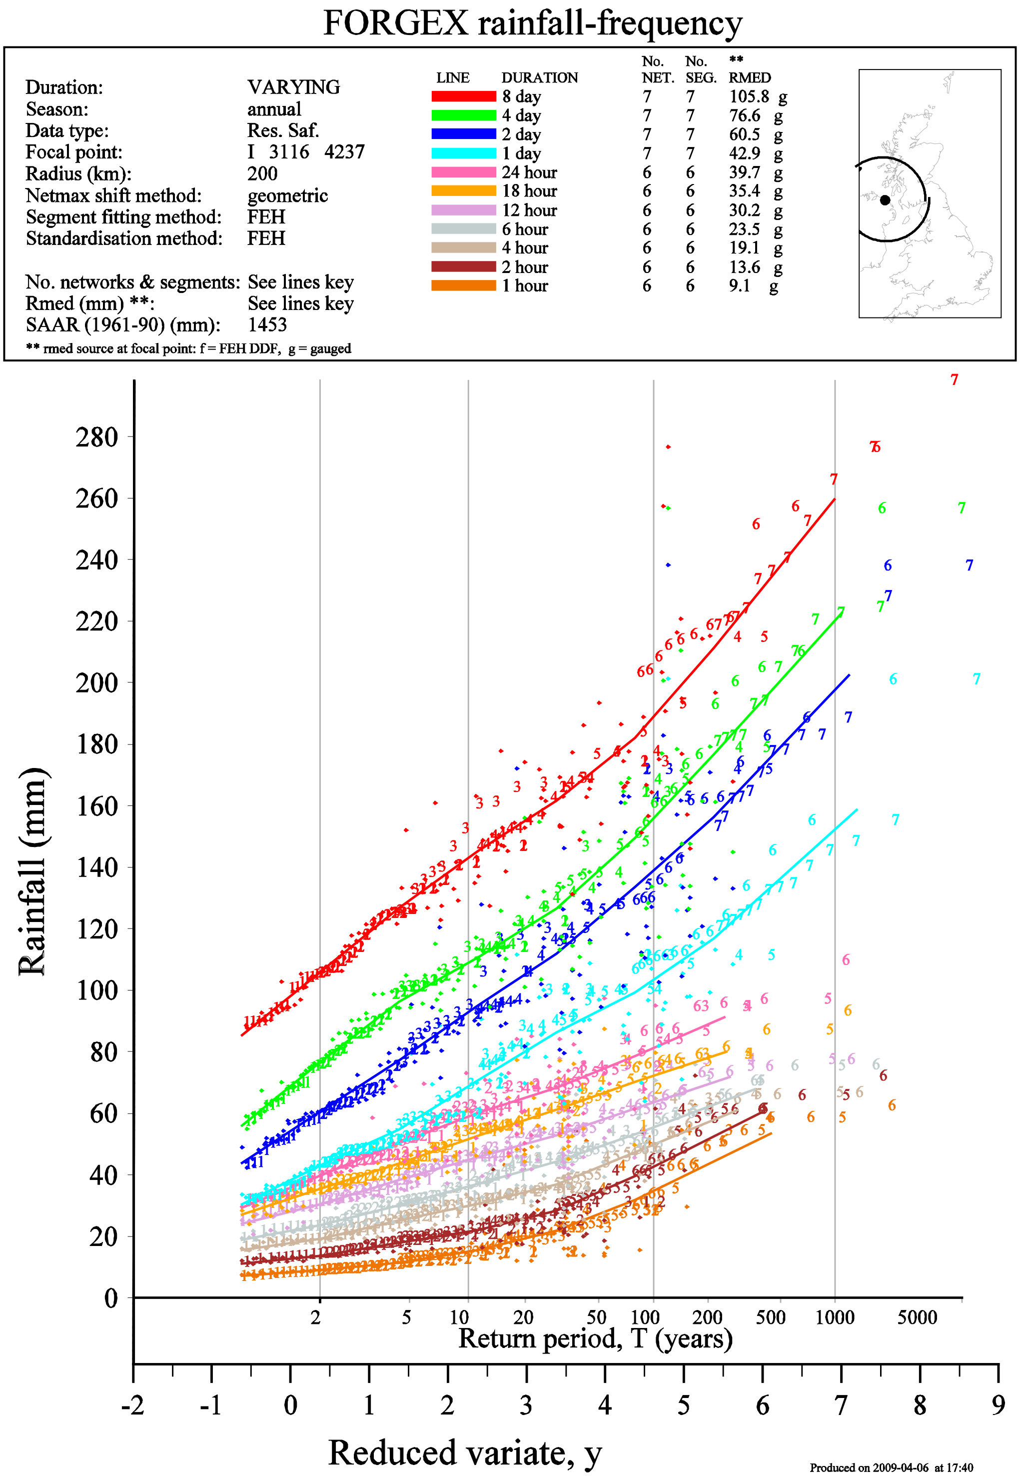

Figure 7.40(a) Rainfall frequency curves for durations ranging from 1 hour to 8 days for a location in Northern Ireland. Produced using the FEH FORGEX procedures.

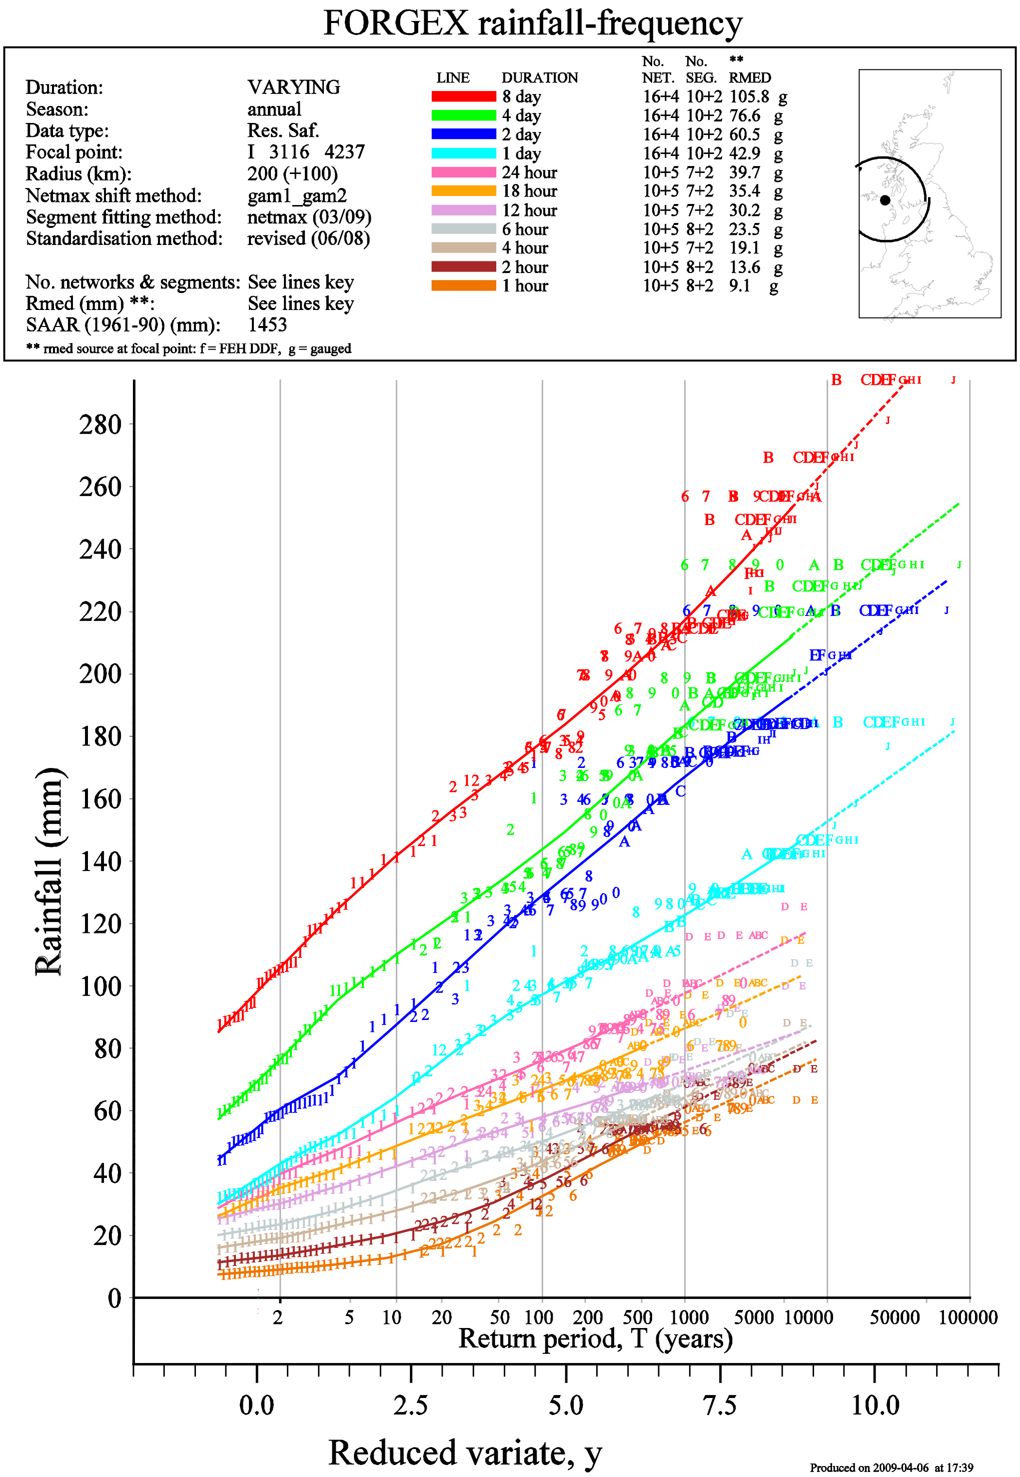

Figure 7.40(b) Rainfall frequency curves for durations ranging from 1 hour to 8 days for a location in Northern Ireland. Produced using the new FORGEX procedures.

7.9 Summary

This section has described a number of changes made to the FEH FORGEX procedure. The main purpose of these has been to remove some anomalous behaviour in the frequency curves inferred for a few locations and to ensure that there will be only small changes between the frequency curves for immediately neighbouring locations. The examples presented have shown that the most noticeable changes to the outcome of the FORGEX methodology since the FEH project arise because of the following two factors:

- the improved availability of data for the hourly-based durations, particularly in Scotland;
- the use of the revised model for spatial dependence described in Section 6.

The revised FORGEX procedure can provide a rainfall frequency curve (represented by a set of fitted line segments) for any duration from one hour to eight days and for any site of interest (focal point) in the UK. In the FEH analysis, the next step taken was to fit a depth-duration-frequency (DDF) model to the FORGEX curves at each focal point for the full range of rainfall durations. Within the current study, it was decided that it would be preferable to avoid this intermediate step of model fitting, and to fit the DDF model to the underlying information used within the FORGEX procedure, rather than to the lines fitted by FORGEX. In the approach adopted it has been possible to make use of the sets of weights attributed to the network maximum points within the revised FORGEX procedure which have been developed and described within this section. These weights are attached to each of the network maximum points, which themselves consist of a value of rainfall and a plotting position calculated according to the underlying spatial dependence model.

Overall, the revised FORGEX procedure provides a means of synthesising the information about annual maximum rainfalls that is available for a target location, taking into account records at immediately neighbouring and more distant locations, in preparation for the fitting of a DDF model. The DDF model is described in Section 8.