Rainfall Frequency Estimation in England and Wales
Phase 2: Production

Technical Report
W183

Rainfall Frequency Estimation in England and Wales
Phase 2: Production

R&D Technical Report W183

D S Faulkner and C Prudhomme

Research Contractor:
Institute of Hydrology

Further copies of this report are available from:
Environment Agency R&D Dissemination Centre, c/o
WRc, Frankland Road, Swindon, Wilts SN5 8YF

tel: 01793-865000 fax: 01793-514562 e-mail: publications@wrcplc.co.uk

Publishing Organisation:
Environment Agency
Rio House
Waterside Drive
Aztec West
Almondsbury
Bristol BS32 4UD

Tel: 01454 624400 Fax: 01454 624409

ISBN: 1 85705 031 2

Dissemination status
Internal: Released to Regions
External: Released to Public Domain

Statement of use
This report presents the results of the anlaysis undertaken during Phase 2 of the project Rainfall Frequency Studies: England and Wales. Phase 3 of the project will involve implementation of the work and wide dissemination of the results through Volume 2 of the Flood Estimation Handbook.

Research contractor
This document was produced under R&D Project W5A-008 by:

Institute of Hydrology
Wallingford
Oxfordshire
OX10 8BB

Tel: 01491 838800 Fax: 01491 692424

Environment Agency's Project Manager
The Environment Agency's Project Manager for R&D Project W5A-008 was:
Ms Linda Aucott, South West Region.

ACKNOWLEDGEMENTS

This report presents the reults of the analysis undertaken during Phase 2 of the project "Rainfall frequency studies: England and Wales" commissioned by the Environment Agency under R&D Project W5A-008. The project leader at IH was Nick Reynard. Thanks are due to the Agency's project leader, Linda Aucott (South West Region), who also co-ordinated the gathering of information on sub-daily rainfall data from the Regions.

The work of staff who provided rainfall data is gratefully acknowledged, particularly Andrew Pimperton of the Midlands Region and Philip Procter and Ann Ruane of the Yorkshire Area of the North East Region who spent much time arranging the digitizing of rainfall charts. Viv Turner of the Northumbria Area, North East Region, is thanked for undertaking the difficult task of gathering and checking data in several different formats. Brian May of the Met. Office gave permission to use the May & Hitch dataset, and these and the continuous data were supplied by George Anderson and Paul Harker.

Helpful advice and comments were provided by Duncan Reed and David Jones, who is also thanked for suggesting the methods given in Appendices 1, 2 and 3. Computational support was provided by David Hewertson and Philip Davies, sandwich course students at IH.

CONTENTS

EXECUTIVE SUMMARY

Background

New methods of rainfall frequency estimation have been developed during the project "Rainfall Frequency Studies: England and Wales", which is part of the Flood Estimation Handbook, a project collaboratively funded by MAFF and the Environment Agency. This report on Phase 2 of the project describes the methods and their results, which will be incorporated into the new Flood Estimation Handbook. Design rainfall estimates are compared with the results of the Flood Studies Report, which has been the source of rainfall frequency estimates in the UK for the past 22 years.

Objectives

A principal objective of Phase 2 was to extend the methods developed in the previous phase for a pilot study region in the East Midlands to the whole of England and Wales. This has involved collecting data to build up a comprehensive database of annual maximum rainfalls, estimating rainfall growth curves for sites throughout England and Wales, and producing maps of an index variable. Other objectives included extending growth curves to return periods longer than 200 years, deriving confidence limits for growth curves, and ensuring the results are consistent, both across different rainfall durations and between different locations.

Results

The achievements of the project include the compilation of a UK database of annual maximum rainfalls with durations from 1 hour to 8 days, from 375 recording raingauges and 6078 daily gauges. Records of sub-daily rainfall totals were gathered from many sources, and received in 19 different formats.

The approach to estimating design rainfalls has a number of key elements:

- The median of annual maxima at a site (RMED) is used as an index variable;
- Rainfall growth curves and maps of RMED are produced for key durations;
- Growth curves can be focused on any location;
- A hierarchical approach is used so that the degree of pooling of data is related to the rarity of the growth rate sought;
- The effect of spatial dependence is allowed for in the fitting of the growth curve;
- Estimates of growth rate and RMED are multiplied to give design rainfalls, which are combined for different durations in a model of depth-duration-frequency.

A new growth curve estimation method (FORGEX) has been developed. FORGEX plots network maximum rainfalls which extend growth curves to return periods as long as 200 years at some sites, accounting for spatial dependence between annual maximum rainfalls. The method has been tested to examine its sensitivity to features such as particularly large rainfall totals, and confidence limits have been derived using a bootstrapping technique.

To improve the mapping of RMED, particularly in sparsely-gauged upland areas, the method of georegression has been adopted. This incorporates specially-developed topographic variables, such as the angle subtended by a topographic barrier in the west direction, which are related to the processes governing extreme rainfall. The growth curves and mapping results have been combined into a model of rainfall depth-duration-frequency, which enables the estimation of rainfall for any duration and return period, ensuring that the resulting design rainfalls are consistent across durations.

Conclusions

The results of the new analysis confirm that the rainfall frequency estimates given by the Flood Studies Report (FSR) do not sufficiently allow for regional and local variations. It is concluded that continued use of the FSR results would lead to under-design of flood protection in some areas and over-design in other areas. Plans for implementing the results, principally in the Flood Estimation Handbook, are outlined at the end of the report.

KEY WORDS

rainfall frequency estimation; Flood Estimation Handbook; annual maxima; growth curves; FORGEX; spatial dependence; georegression; depth-duration-frequency.

1. INTRODUCTION

1.1 The project

This report describes the work carried out during Phase 2 of the project "Rainfall Frequency Studies: England and Wales", funded as part of the collaborative Flood Estimation Handbook, by the Environment Agency under R&D Project W5A-008. The main objectives of the project as a whole are to review the methods of rainfall frequency estimation in England and Wales, and to develop new procedures where the current methods give unsatisfactory results. Details of the new UK-wide methods of rainfall frequency estimation resulting from this project will be published as one volume of the Flood Estimation Handbook which is currently under development at the Institute of Hydrology.

The project is divided into three phases: feasibility, production and implementation. Phase 1 was subdivided into two parts. Phase 1a was a survey of the scope of the study which reviewed methods of rainfall frequency estimation and the availability of rainfall data, the results of which are given in NRA R&D Note 175 (Stewart and Reynard, 1994). Phase 1b was a pilot study which developed a new method of rainfall growth estimation for return periods up to 200 years in the East Midlands. Phase 1b also included a study of mapping an index of extreme rainfall in the Severn catchment. The results are given in NRA R&D Note 478 (Stewart et al, 1995).

1.2 Objectives of Phase 2

The detailed terms of reference for Phase 2 are given in Appendix 4. A principal aim has been extending the techniques developed in Phase 1b to the whole of England and Wales, which has required the collection of many records from recording raingauges. Phase 2 has also involved the development of new methods, for example to extend growth curves to return periods beyond 200 years, to map rainfall statistics in mountainous areas, and to estimate confidence limits for growth curves.

1.3 Approach and outline of report

The two-stage approach to rainfall frequency estimation used in Phase 1b has been retained. This involves estimating rainfall growth curves which are standardized by an index variable. Design rainfall estimates are obtained by multiplying the rainfall growth factor for the required return period by the index variable. As in Phase 1b, the index variable is the median of at-site annual maximum rainfalls, RMED. The structure of this report reflects the two stages of rainfall frequency estimation.

The report describes the collection of rainfall data in Chapter 2. The following four chapters are concerned with rainfall growth curves. The details of developments during Phase 2 are given in Chapter 3. Chapter 4 is a brief explanation of the final FORGEX method of rainfall growth estimation. The results are presented in Chapter 5, mainly in the form of maps of growth rates for the whole of the UK. Chapter 6 describes the estimation of confidence limits for growth curves, and various tests of the FORGEX method.

Chapters 7 and 8 are concerned with the mapping of the index variable, RMED. The development of new mapping techniques, incorporating topographic information, is described in Chapter 7, and the resulting maps of RMED are presented in Chapter 8.

The two stages are brought together in Chapter 9, which presents maps of design rainfall and compares the results with those of the Flood Studies Report, and Chapter 10, which introduces a model of depth-duration-frequency, reconciling the rainfall estimates for different durations.

The report avoids detailed derivations and mathematical descriptions of methods. Some derivations are given in Appendices 1 to 3, and others are given in publications cited in the report.

2. DATA COLLECTION

2.1 Requirement

Series of annual maximum rainfalls for durations between 1 hour and 8 days are required for sites across England and Wales. The analysis is restricted to records at least ten years long, as for Phase 1b. However, records lasting for 8 or 9 years have also been collected and added to the database. Because data for all of the UK are held on one database, the summary statistics in this chapter include records from Scotland and Northern Ireland.

2.2 Sources

2.2.1 Hourly data

Hourly rainfall measurements are made and stored by many different organisations in the UK, and much time was spent contacting these bodies, sorting through lists of gauges, and processing data received in 19 different formats. Some records were provided as continuous hourly data, others as times of tip from tipping bucket recorders, and others as tabulations of annual maxima. Software written at IH was used to convert times of tip into hourly totals.

The sources of data are listed in Table 2.1, together with the number of records received with at least ten annual maxima. In addition to these records, series shorter than ten years were in some cases combined with records from other sources so that the same gauge produced a series of at least ten annual maxima.

The organisations providing data for England and Wales are listed below. A comprehensive survey of rainfall data in England and Wales is given in the Phase 1a report (Stewart & Reynard, 1994).

The *Met. Office* provided computerized hourly data (most records starting in the 1980s), computerized 1-hour annual maxima (May & Hitch, 1989) and manuscript tables of annual maxima, which were computerized at IH. The sets of annual maxima include the data used for the Flood Studies Report analysis. The May & Hitch dataset provides the largest number of records, including several very long ones.

The *Environment Agency* provided computerized hourly data in various formats, although the majority of records (from the Midlands, Anglian, Thames and NW regions and the Yorkshire area of the NE Region) were in Rainark format and were easily copied to the Institute of Hydrology's Rainark system which was purchased as part of this project. Most records are short and recent, apart from several extending back into the 1970s which were digitized from charts by the Midlands region.

The *PEPR database* at IH contains hourly rainfall records which were digitized from charts by a Precision Encoder and Pattern Recognition (PEPR) machine by the Met. Office and the Greater London Council. Most records are from gauges operated by the GLC and cease in the mid-1970s.

Some extend as far back as the 1920s. The data are comprehensively described by Moore *et al.* (1993).

Table 2.1 Sources of hourly rainfall data

Organisation	Region	Type of data	Number of records with at least ten annual maxima (+ number with 9 AMs)
Met. Office	Great Britain	Continuous	99
	N. Ireland	Continuous, some digitized from charts	6 (+2)
Met. Office (May & Hitch)	UK	Tabulated annual maxima for 1 hour	159
Met. Office	UK	Tabulated annual maxima mostly for 2-24 hours	99
Environment Agency	Midlands	Continuous	67
	Anglian	Continuous	16
	Thames	Continuous	16 (+ 16 to extend PEPR data)
	Southern	Continuous	4
	NW	Continuous	0 (7 to extend May & Hitch data)
	NE	Continuous	8 (+1 to extend May & Hitch data)
Institute of Hydrology	Greater London	Continuous PEPR data	27
	Elsewhere	Continuous PEPR data	7
	Experimental catchments	Continuous AWS data	8
SEPA	Scotland	Continuous	12 (+3)
DANI	N. Ireland	Continuous on paper	8
Irish Met. Service	Rep. of Ireland near the border	Continuous	2

Automatic weather stations (AWSs) installed by *IH* in experimental catchments added eight records, six of which are from gauges close together at Plynlimon, mid-Wales.

The total number of records received from recording raingauges is 567. Note that the total number of sites is rather smaller, because some sites have more than one record from different sources. Additional hourly data were provided by the Environmental Change Network for Moor House in Teesdale: unfortunately the record was not complete enough to be included in the database.

Figure 2.1 is a map of all recording raingauges providing at least ten annual maxima. The density of gauges is highest in central England and Greater London. The largest gaps in the network are in inland Devon and Cornwall, the north Pennines and North Wales. A search through catalogues of recording raingauges indicated that there have been gauges in these areas, for example at the Dinorwic pumped storage power station in Snowdonia. Several organisations in North Wales were approached for data, but the records from those gauges have not been found. We still hope to receive annual maxima abstracted from charts for several extra sites in south-west England. Records from several more EA gauges in Yorkshire are also expected.

2.2.2 Daily data

The network of daily gauges is much more dense than that of recording gauges across all of England and Wales. The vast majority of daily annual maxima were abstracted from computerised records provided by the Met. Office. The number of gauges with daily data on the ORACLE database at IH is 12913. Of these, approximately 6100 have records for at least ten years. This total includes a set of 561 long-term records with data from before 1961.

In addition, a few records of daily annual maxima were provided by the Climatological Observers Link.

The hourly and long-term daily data from the Met. Office and the Environment Agency were provided under the terms of the Memorandum of Understanding agreed between the data suppliers and IH (see Appendix B).

2.3 Abstracting annual maxima

Annual maxima were abstracted from the continuous hourly and daily data using the algorithm described in the Phase 1b report (Stewart *et al.*, 1995). Years were rejected only if more than 25% of the data were missing. The effect of this assumption was tested (see Section 6.4). Dates and times of annual maxima were also recorded and stored in the database.

Some of the tabulated annual maxima were rainfall totals for a duration of 120 minutes rather than 2 clock hours. These were converted to 2-hour maxima by dividing by the discretization factor 1.08 (Dwyer & Reed, 1995). Likewise, some May & Hitch annual maxima were converted from 60 minutes to 1 clock hour, using the factor 1.16.

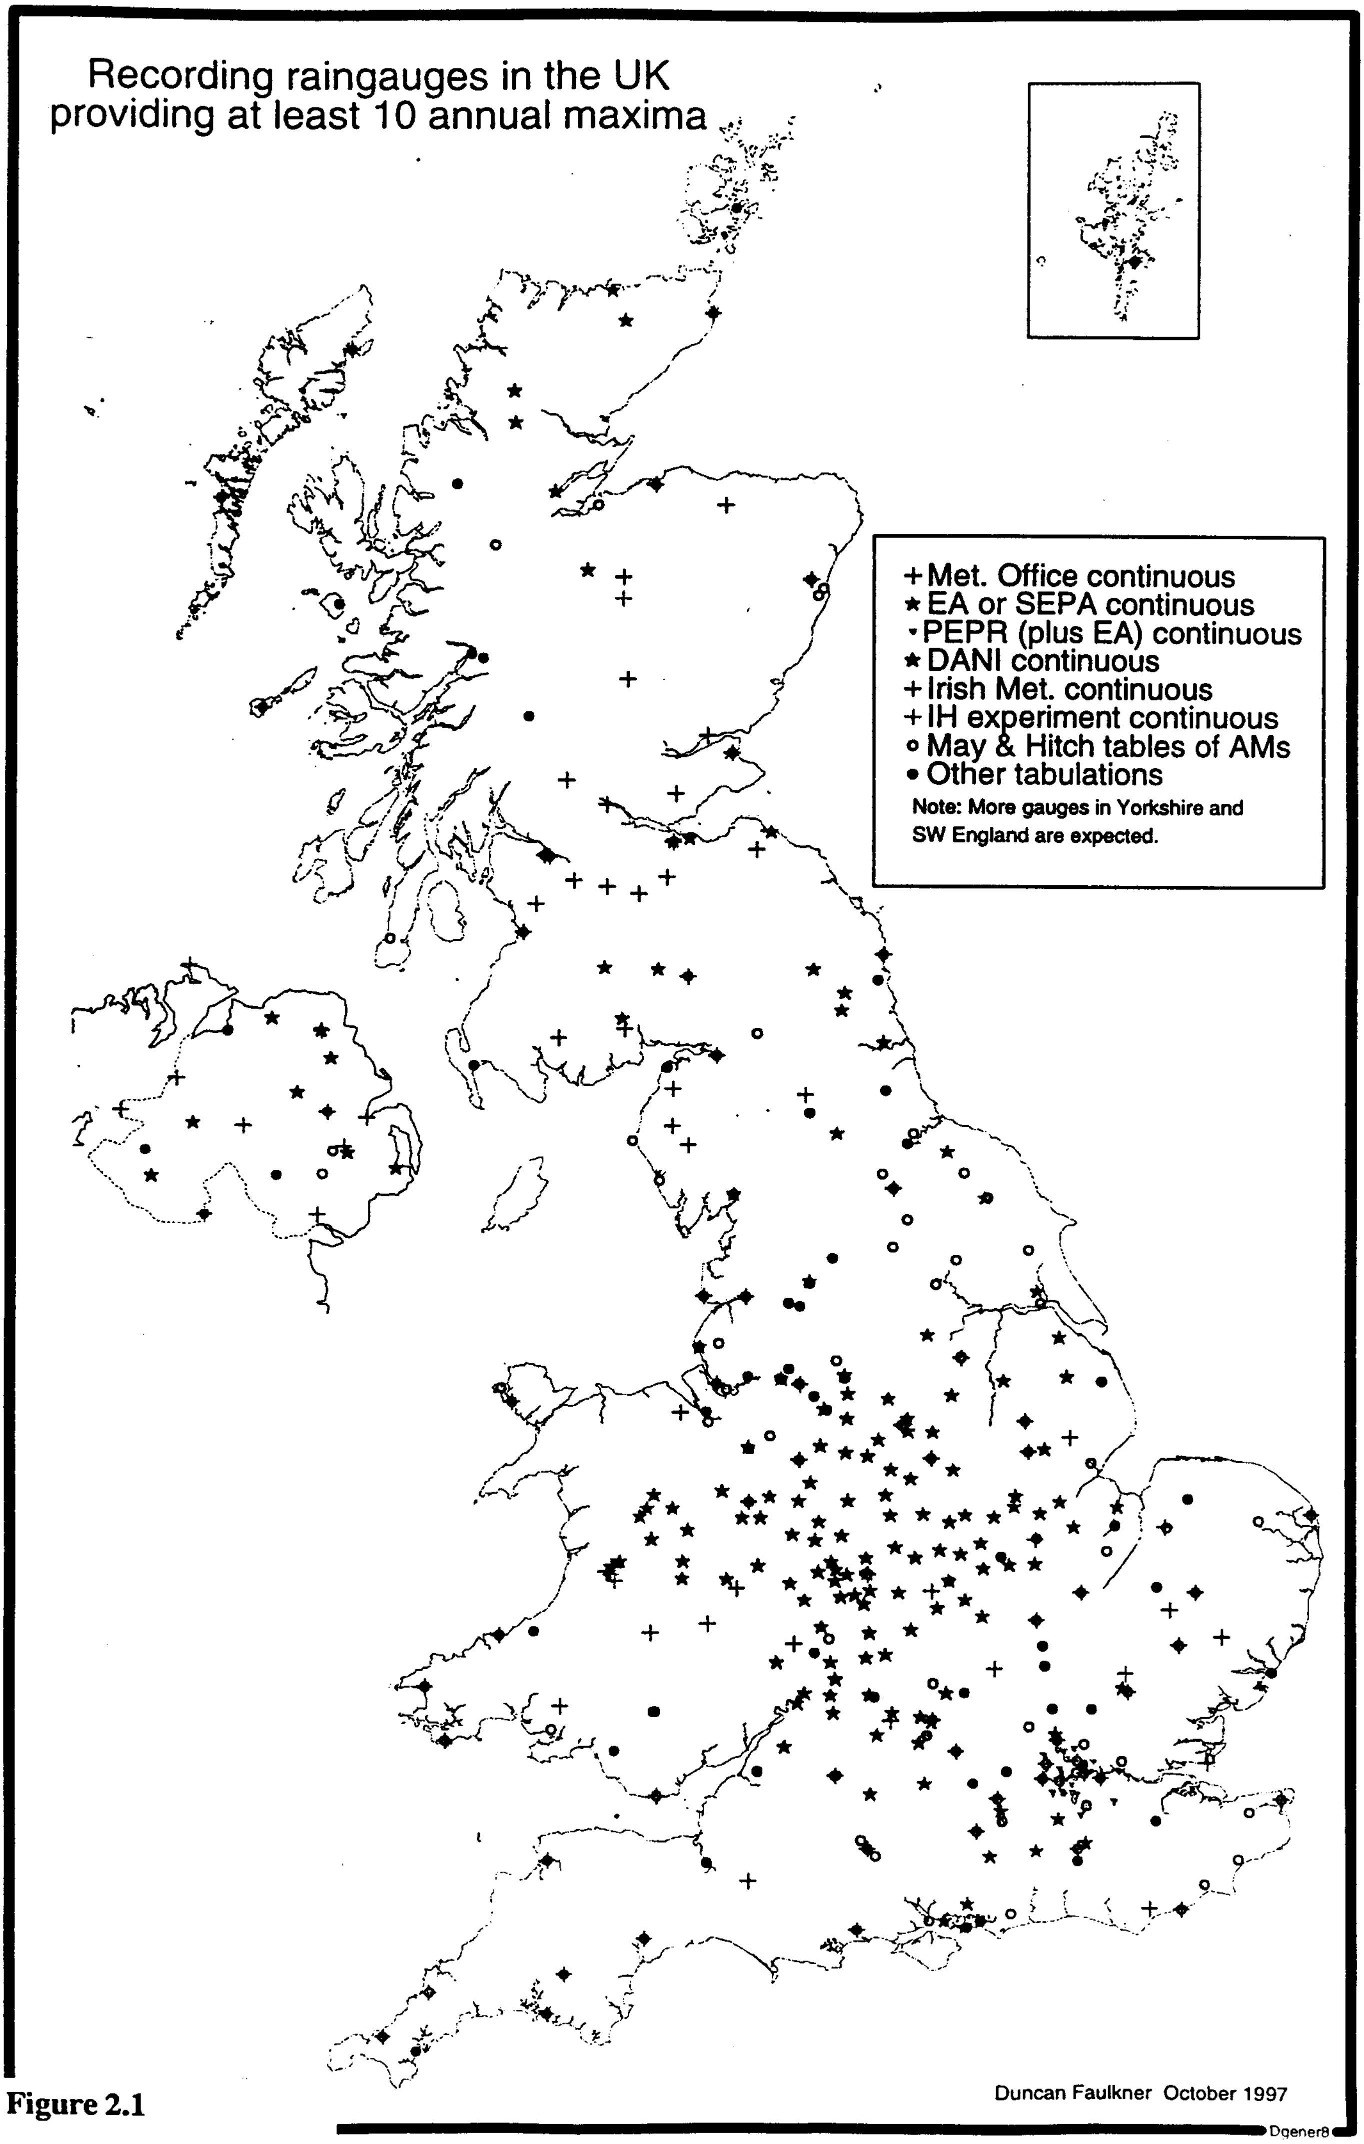

Recording raingauges in the UK
providing at least 10 annual maxima
+ Met. Office continuous
★ EA or SEPA continuous
▾ PEPR (plus EA) continuous
★ DANI continuous
+ Irish Met. continuous
+ IH experiment continuous
○ May & Hitch tables of AMs
● Other tabulations
Note: More gauges in Yorkshire and SW England are expected.
Duncan Faulkner October 1997
Dgener8

Figure 2.1

2.3.1 Hourly annual maxima

Annual maxima abstracted from continuous hourly data were checked with nearby daily totals. The hourly annual maxima were compared with totals for the three days surrounding the day on which the maximum was recorded, from the three nearest daily gauges.

Any suspicious values were investigated further by looking at the continuous data. This visual check highlighted errors such as accumulations, where a large hourly total follows a period of missing data. Records from tipping bucket recorders (TBRs) sometimes showed large numbers of simultaneous tips. Some suspiciously large hourly totals were further investigated with more daily data and information from British Rainfall. After consultation with the suppliers of data, erroneous values were removed and the annual maxima re-abstracted. The total number of suspect annual maxima investigated was approximately 290.

Quality control was not confined to looking for suspiciously high annual maxima. Annual maxima can be under-estimated if a gauge is out of action for a significant period, or if it under-reads and is not scaled by totals from a daily check gauge. It was not always possible to distinguish dry periods from periods of missing record in TBR data. As a basic check, the depths recorded by TBRs for each year were accumulated to spot any years with long periods of missing data. It was thought possible that TBRs in the Hampshire area were under-reading, since they gave very low median annual maxima. The daily TBR totals were compared with check gauges, and in the end only one TBR record required adjustment.

Tabulated annual maxima and the May & Hitch dataset were cross-checked with each other and with annual maxima abstracted from continuous records, where they overlapped, to spot any inconsistencies. The previously abstracted annual maxima were held to be of good quality, and few errors were found.

As a final check, all annual maxima loaded into the ORACLE database were scanned for unusually large or small values. Only one change was made as a result.

2.3.2 Daily annual maxima

The daily data had already been subject to extensive quality control at the Met. Office. This was confirmed by searching the annual maximum database for the presence of suspiciously large or small values. A few large annual maxima were compared with nearby records, but only one change was made as a result: the 1 to 8-day maxima for 1963 at gauge 297785 were deleted.

One other change was made, undoing a change made by the Met. Office quality control procedure. Gauge 77255, Walshaw Dean Lodge, recorded the Calderdale storm on 19th May 1989. The recorded total of 193.2 mm is generally accepted as valid (Acreman & Collinge, 1991), but the rainfall for the whole of May had been set to missing. The annual maxima for 1, 2, 4 and 8 days were set to 193.2 mm.

2.4 The database of annual maxima

Annual maximum rainfalls are now held in two tables (hourly and daily) on the ORACLE database at IH. Some summary statistics of the database are given in Table 2.2. Note that these statistics refer only to gauges used in the analysis ie those with sufficiently long records. The total number of station-years is compared with the approximate number used in the Flood Studies Report (FSR) rainfall analysis (reported by Stewart and Reynard, 1994). The increase in data quantity is striking for the 1-hour duration, where the number of station-years has more than trebled.

Table 2.2 Summary statistics of the annual maximum database

Duration	Number of gauges	Number of station-years	Approx. number of station-years used in FSR analysis.
1 hour	375	7400	2300
1 day	6078	150000	96000

Figures 2.2 and 2.3 show the distributions of record length in the hourly and daily tables respectively. Figure 2.2 reveals that most 1-hour series are short, with relatively few longer than 20 years. The longest records are around 100 years. The difference in numbers of short-term and long-term daily gauges is reflected in Figure 2.3, which has a sharp drop in the number of series longer than 35-40 years. Short-term gauges starting in 1961 will have produced 35 annual maxima by 1995, which is the most recent year of annual maxima currently stored in the database.

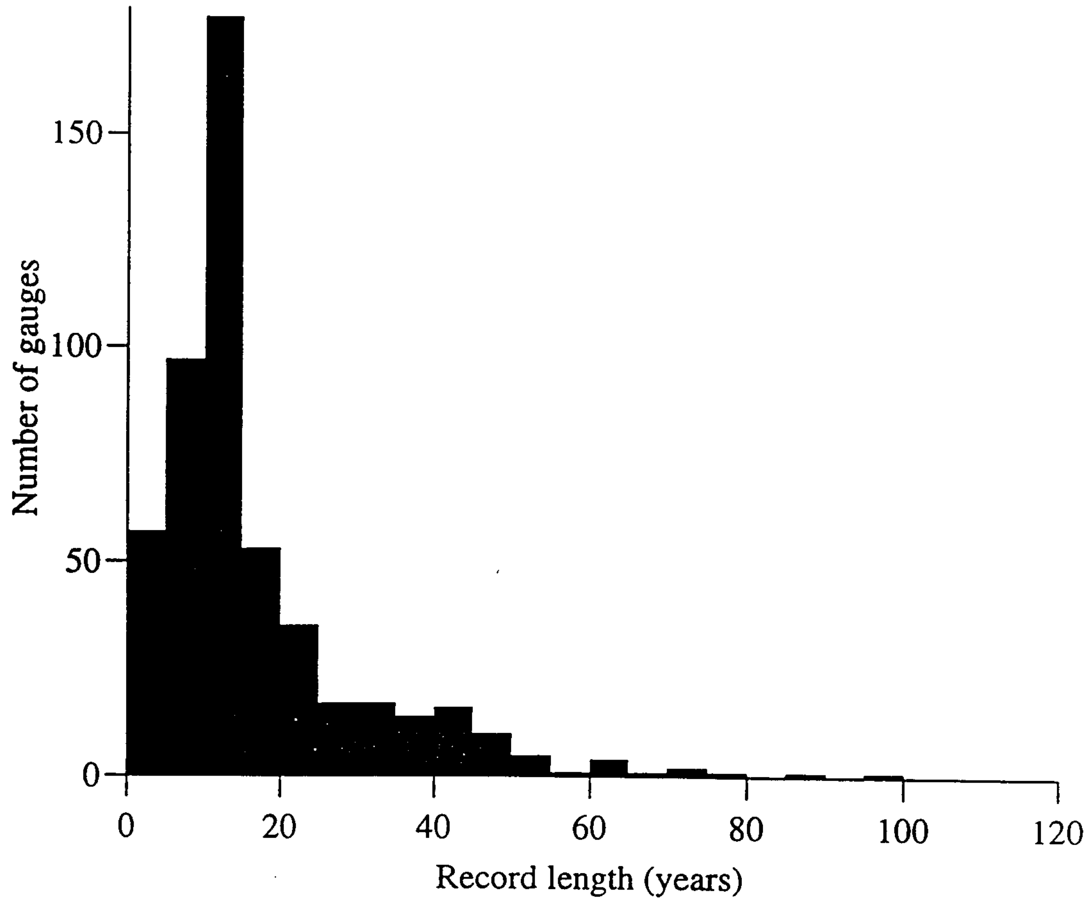

Figure 2.2 Distribution of record length for hourly data

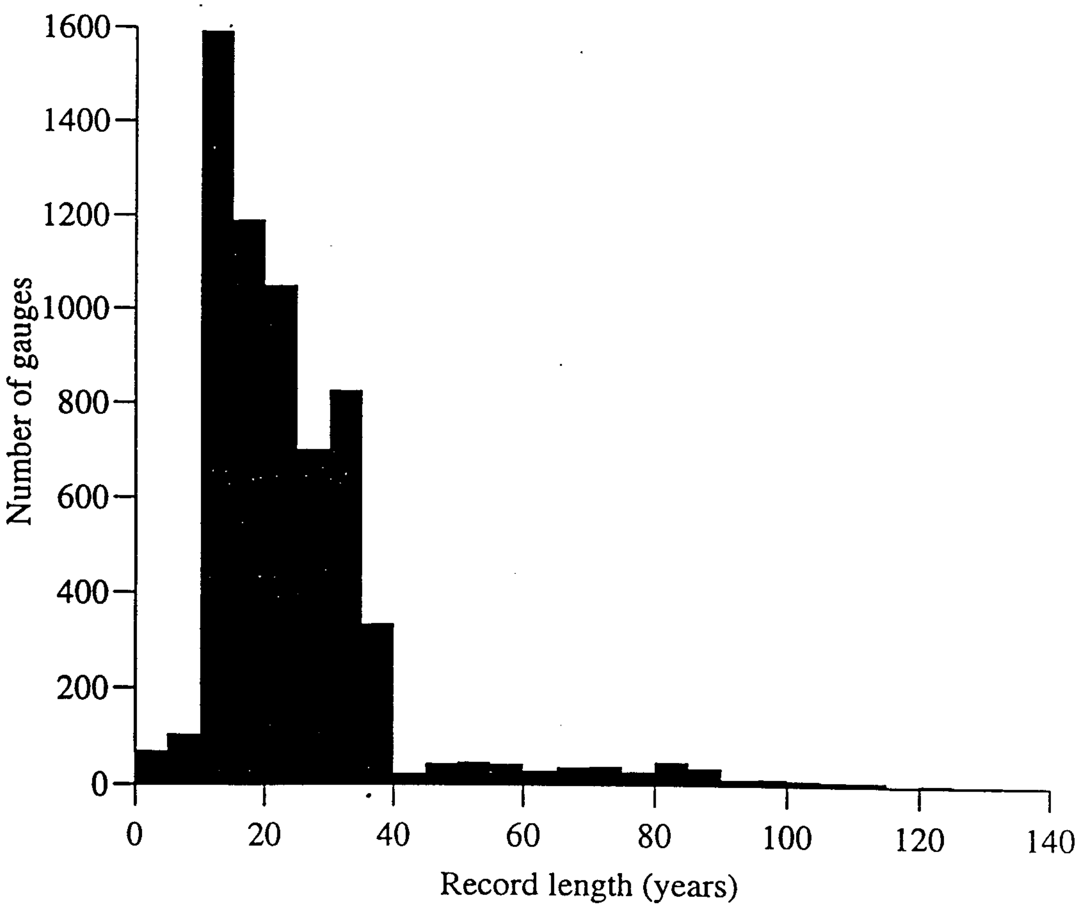

Figure 2.3 Distribution of record length for daily data

3. DEVELOPMENT OF GROWTH CURVES

3.1 Introduction

The main elements of the growth estimation method developed in the pilot study (Phase 1b) were retained in Phase 2. However, the method has been substantially modified, and in particular it has been extended to produce growth rates for return periods longer than 200 years. This chapter describes development of the method since Phase 1b and experiments *en route*. The final method is summarised in Chapter 4.

First, it is helpful to summarize key features of the Phase 1b focused growth curve estimation method:

- rainfall frequency estimation is based on the analysis of annual maxima;
- the median of at-site annual maxima, RMED, is used as the index variable;
- individual durations are treated separately in the construction of growth curves;
- growth curves are focused on the site of interest, rather than derived for fixed regions;
- annual maxima are pooled from a network of gauges which expands with return period, each network providing points which plot in a given "y-slice" on the Gumbel plotting scale;
- a network maximum growth curve is fitted jointly with the pooled curve, to allow for spatial dependence;
- the resulting segmented growth curve is fitted using least squares.

3.2 Extending growth curves beyond 200 years

3.2.1 The modified station-year method

Phase 1b growth curves extended to return periods of 100-200 years, the return period corresponding to the maximum raingauge record length. It was initially thought that they would be extended to longer return periods using a modified station-year approach. This approach was eventually rejected in favour of an extension of the network maximum plotting method.

The modified station-year method was introduced as part of the FORGE methodology (Reed and Stewart, 1989). The basic station-year method combines records of length M years from N sites with no consideration of dependence between them. It plots the largest annual maxima in the combined dataset as if they came from a single record of length MN years. The modified station-year method accounts for spatial dependence in the network by replacing MN with MN_e, the equivalent number of independent station-years. N_e (the effective number of independent sites) is estimated from the characteristics of the raingauge network using the Dales & Reed model (Dales and Reed, 1989). The largest independent standardized rainfalls are plotted as the largest events in a record of length MN_e years.

Most concerns about the modified station-year method sprung from the way it selects the largest independent annual maxima. Events occurring on the same day or on neighbouring days are deemed to be dependent (ie. arising from the same weather system), and only the largest of the group is plotted and used to fit the growth curve. Three possible problems with this are:

- What if a pair of annual maxima occur on the same day at widely-separated locations? They would be treated as dependent, when in fact they may not be due to the same event. This possibility was tested when FORGE was developed, and was further investigated in the study. It was thought to be unlikely to happen, and was in fact found not to have a major effect on growth curves.

- Keeping only the largest of a group of dependent annual maxima may bias the result. It may be better to include the other annual maxima with suitable weights (Jones, D.A., personal communication).

- There is no date information available for many sub-daily data. Tabulated annual maxima and the May & Hitch dataset only give the month of occurrence, and some gauges do not even have that. This prevents effective testing for dependent values.

Growth curves were derived using the modified station-year method, and it was altered to produce a growth curve for large networks which matches up with the segmented growth curve. In the end it was decided to use a network maximum approach instead, as described below. This does not require date information, and it gives a better idea of the dependence in the network, not suffering from the problems listed above. Furthermore, it is consistent with the method used to construct the lower portion of the growth curve.

3.2.2 The extended network maximum method

This is an extension of the Phase 1b method, but the Dales & Reed model is applied rather differently. The network maximum series consists of the highest standardized annual maximum rainfall observed across a network of gauges in each year of record. A distribution can be fitted to network maximum (netmax) points. In the previous method, netmax and pooled growth curve segments are fitted jointly, with the horizontal offset between them obtained from the Dales & Reed model. The new method first shifts the netmax points to the right, using the offset given by the model, and then fits a single growth curve to all points (pooled and network maximum), by least squares.

Figures 3.1 and 3.2 compare growth curves focused on Leicester fitted by the old and new method. Note that netmax points are represented by numbers, corresponding to the number of the network (1 is the smallest network). Pooled points are represented by dots. As in the previous method, only a few netmax points are shown on the diagram. Netmax points for networks 1 to j which plot within the jth y-slice are plotted on Figure 3.2, and only they are used to fit the growth curve. The points from smaller networks are more valuable, because they are more local to the focal point. This highlights one advantage of the new method: it overcomes the previously rather arbitrary assignment of netmax points to y-slices.

This amendment enables an extension to higher return periods. Once the pooled data have run out (because there are no records longer than around 150 years), the expansion of the gauge network

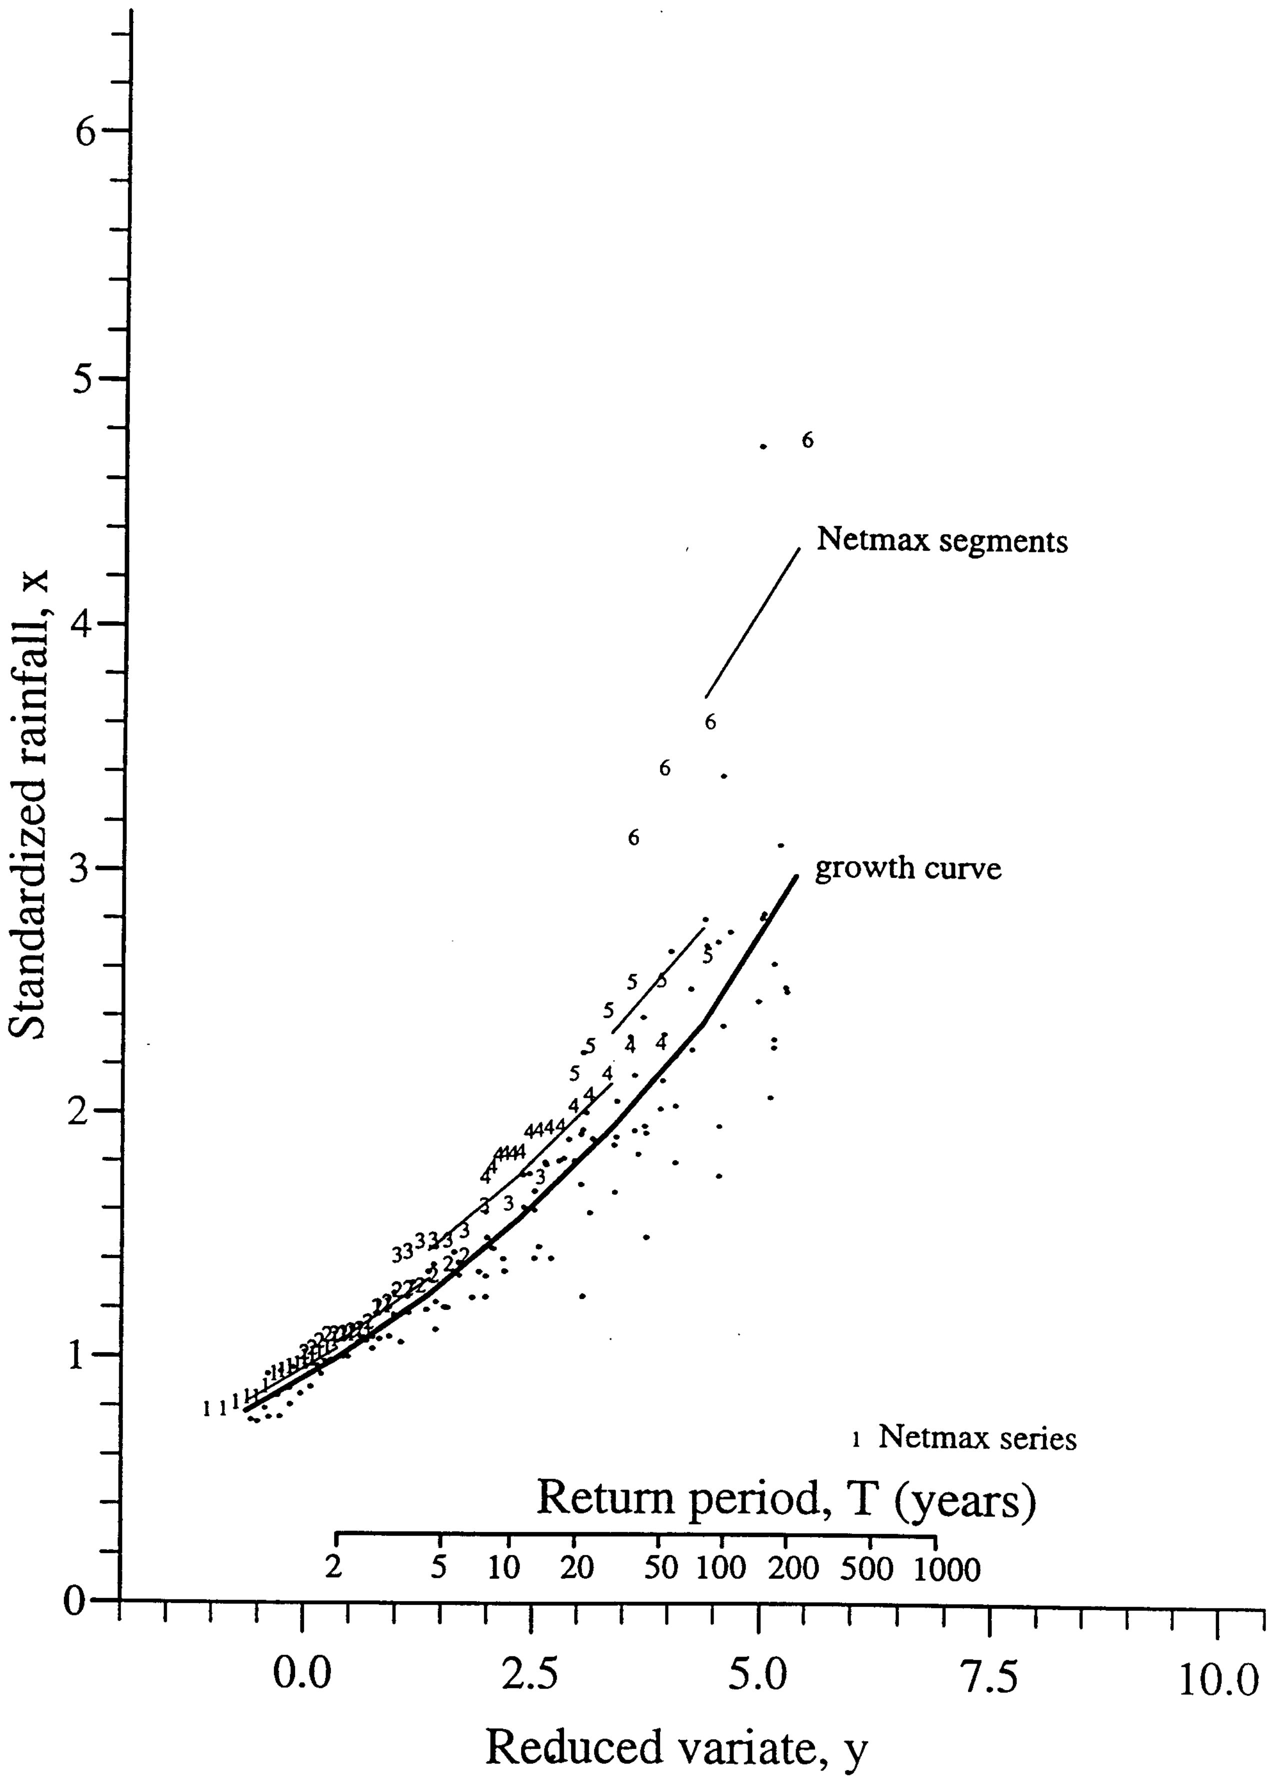

Figure 3.1 Growth curve focused on Leicester fitted by Phase 1b method

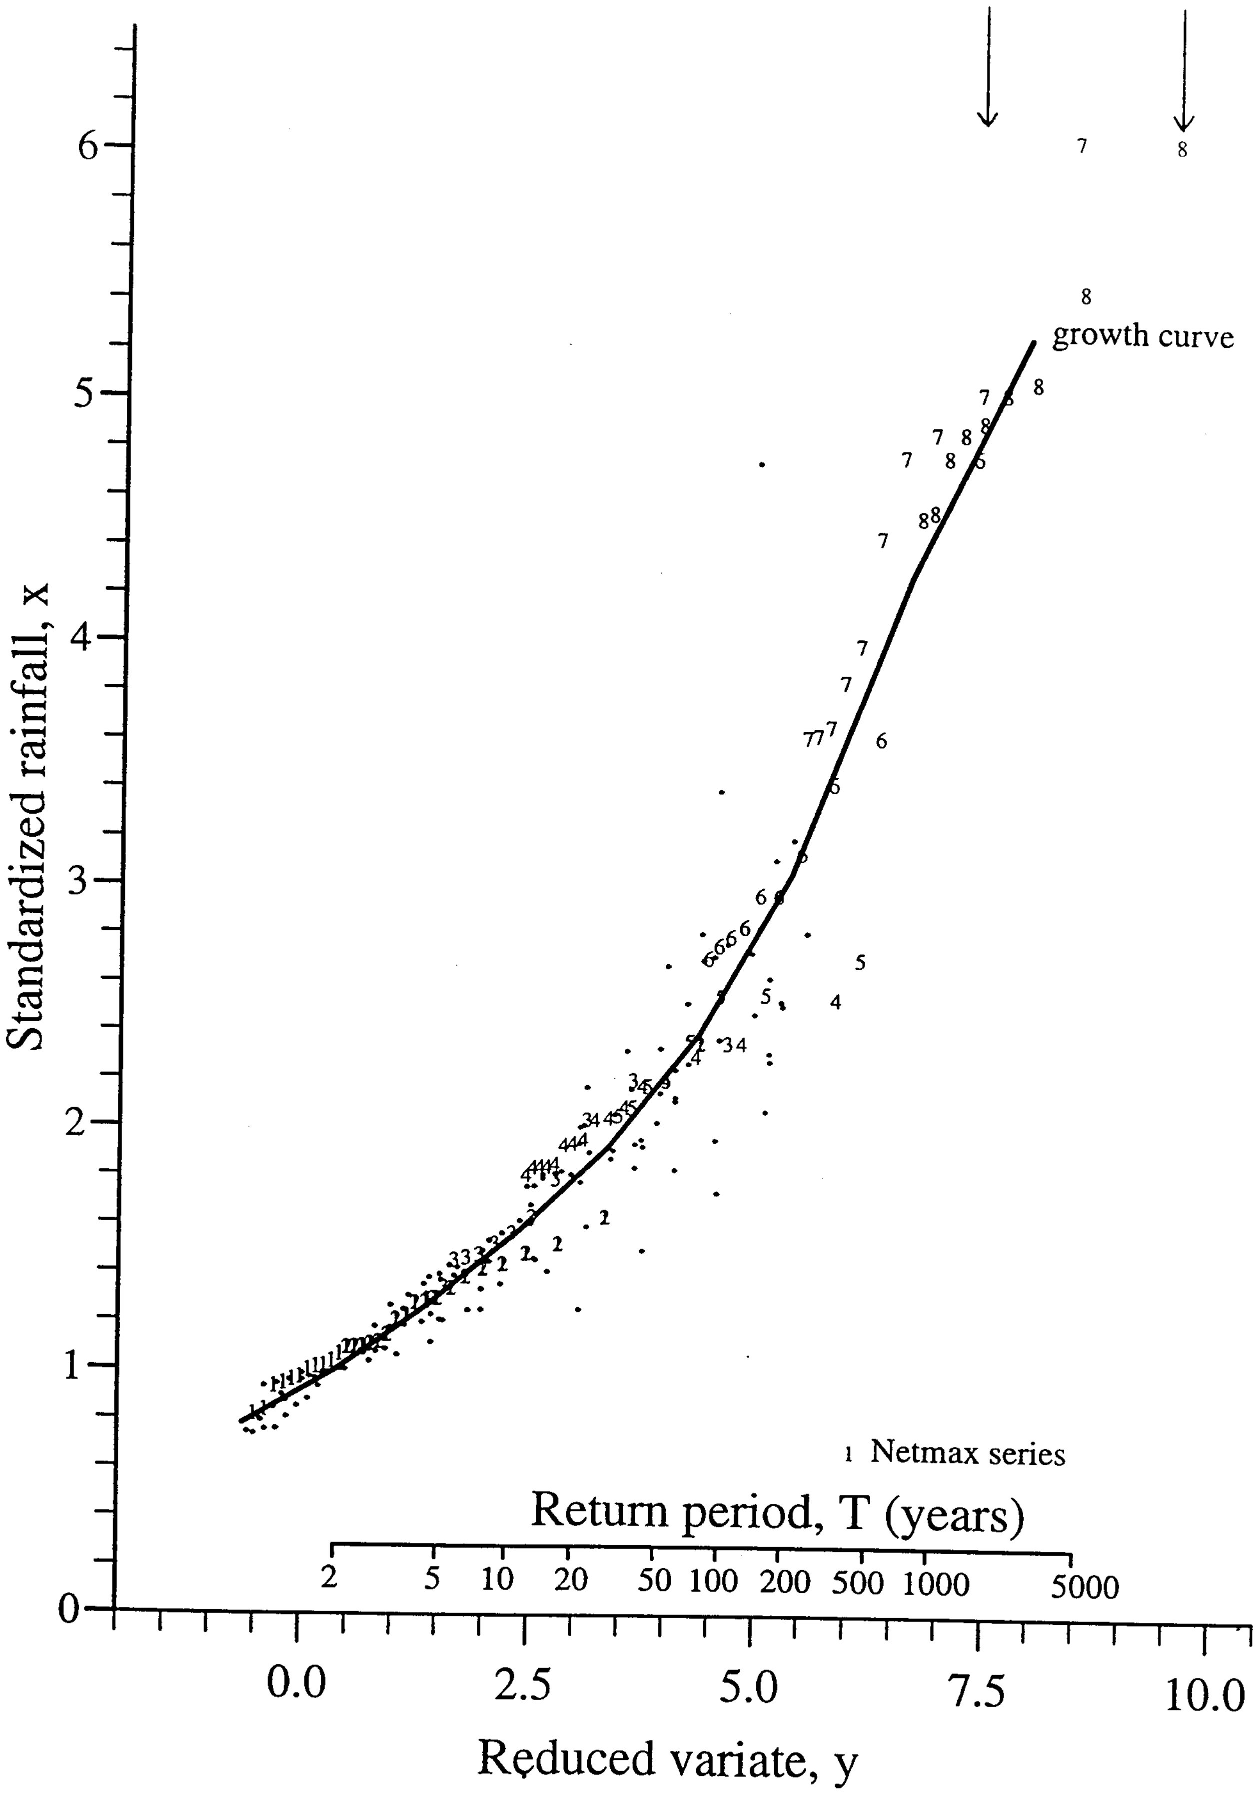

Figure 3.2 Growth curve focused on Leicester fitted by shifting netmax points

is continued, and only the netmax points are plotted for each new network. These netmax points are shifted to the right and define the growth curve for high return periods. As well as being more consistent, this new method is simpler.

During the development of this method, it was noticed that the shifted network maximum series do not all pass through the point (2,1) which defines the position of the median of the distribution. They should pass through this point because the median is the standardizing variable. See for example Figure 3.3 (netmax series for six networks focused on Preston), where the netmax series from larger networks do not pass through the median point. The effect on growth curves is not major, because they are fitted mainly to the highest netmax points in each series. However, the information given by the position of the median is worth using. The next section outlines a new way of assigning plotting positions to the netmax series, which overcomes this problem.

Another reason for needing new plotting positions is that the gauge network varies through time. As the number of gauges changes, the spatial dependence in the network changes. In the Phase 1b method, this was accounted for by averaging the characteristics of the network only over those years which provided the network maxima which were used to fit the relevant segment of the growth curve. This is not possible (without iteration, at least) when shifting the netmax series, because it is not known which netmax points will be used to fit the growth curve until the series has been shifted.

3.3 New plotting positions for netmax points

Because the network characteristics change from year to year, each netmax point is shifted individually by an amount appropriate to the network operating in the year concerned. The plotting positions are derived by a maximum likelihood method, which is explained in Appendix 1.

The formula for the plotting positions is solved by iteration, and it includes a summation such that the plotting position of a given netmax point depends on the plotting positions of all higher-ranked netmax points.

One advantage of these new plotting positions, and the individual shifting of netmax points can be seen in Figure 3.4 (focused on Preston as for Figure 3.3), where netmax points from all networks now pass very close to the median point.

3.4 Trials on weighting points and smoothing growth curves

During the development of the new method, several experiments were carried out to find the best method for fitting growth curve segments to the pooled and netmax points. Differential weighting of pooled and netmax points was tried, both to ensure that netmax points were given sufficient influence in the fitting, and also to help fit smoother growth curves. Other methods of smoothing were also investigated.

In the end it was decided to give all plotted points equal weight in fitting the growth curve. The

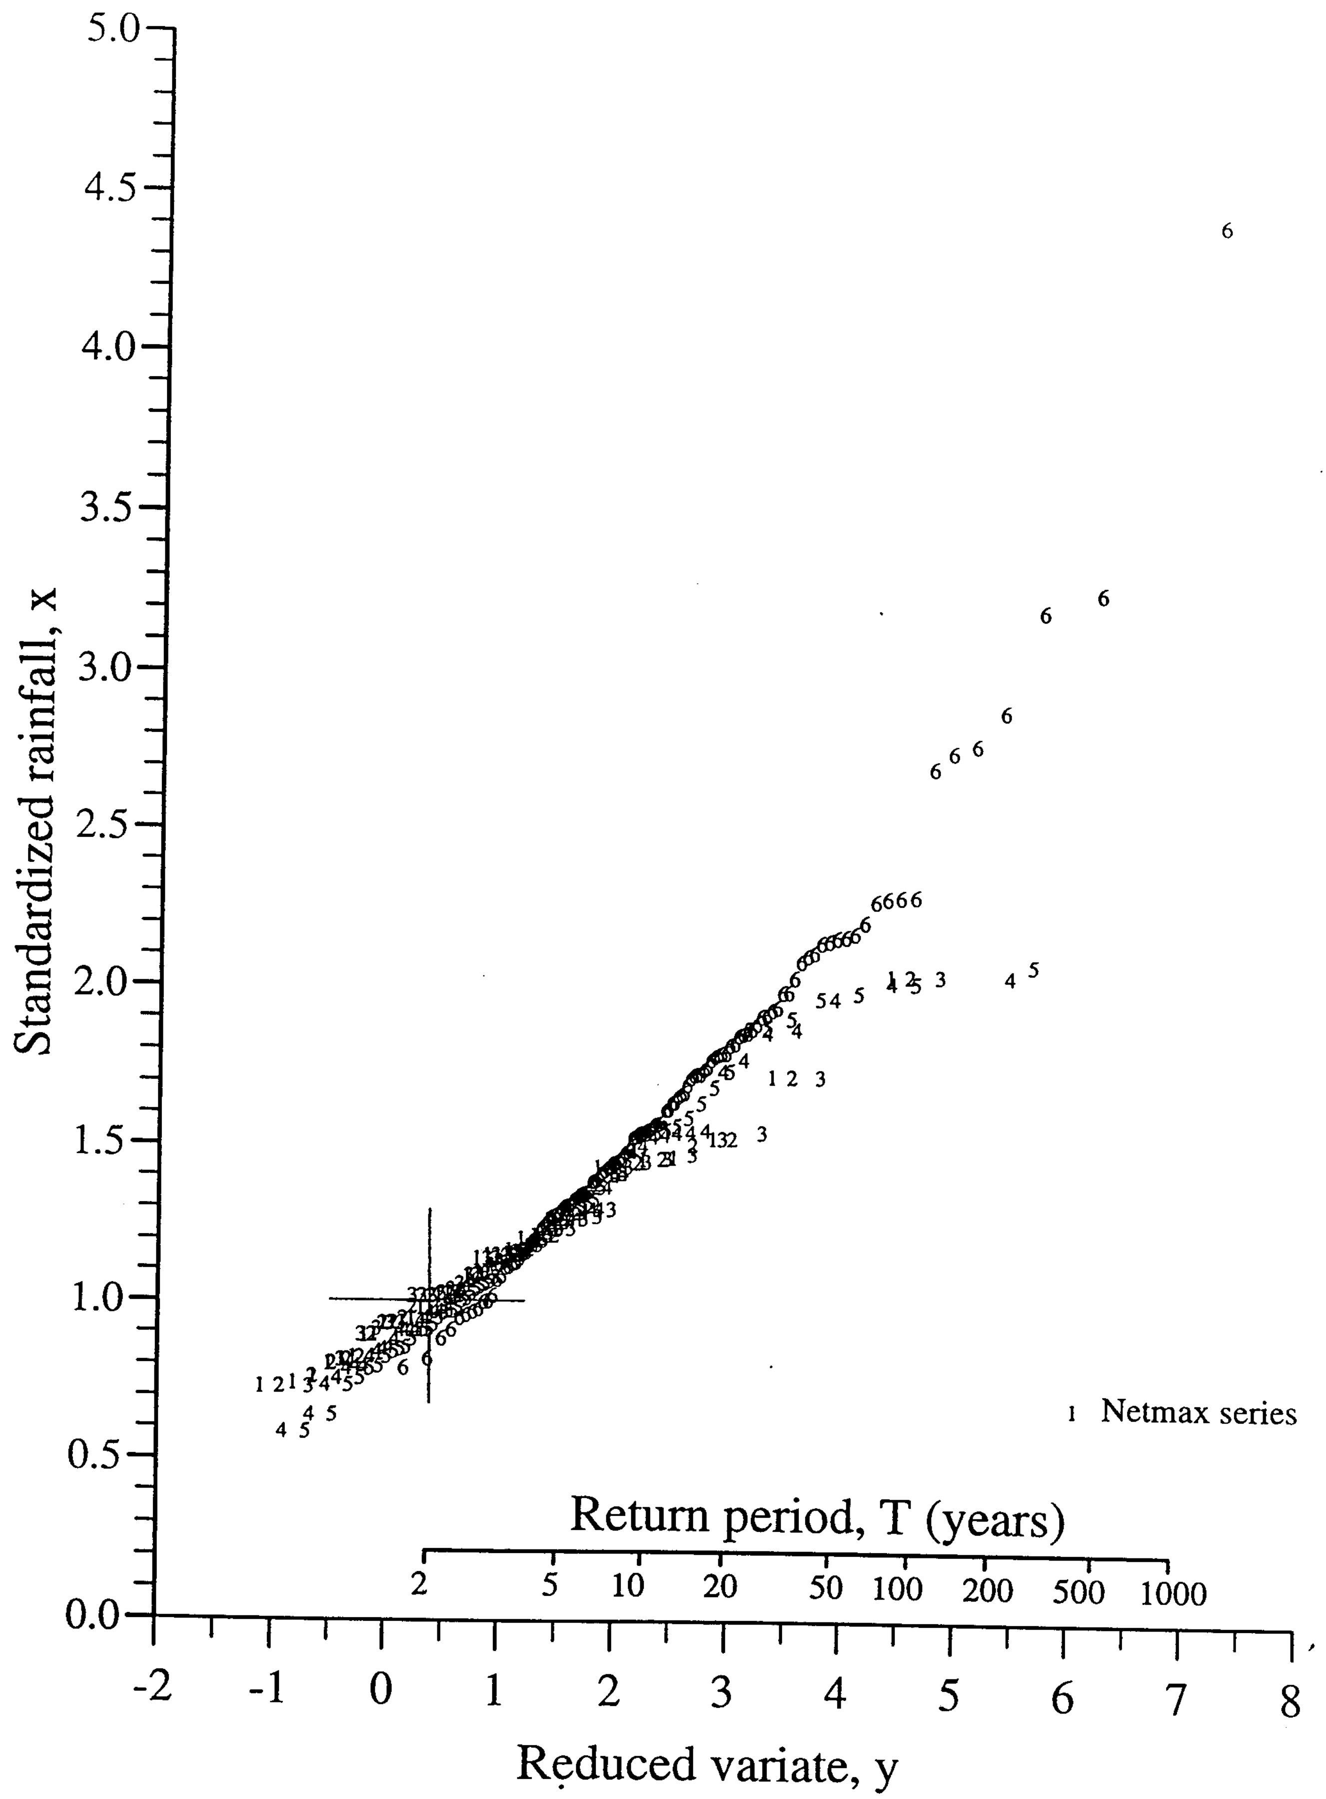

Figure 3.3 **Shifted netmax points focused on Preston: larger networks fail to pass through median**

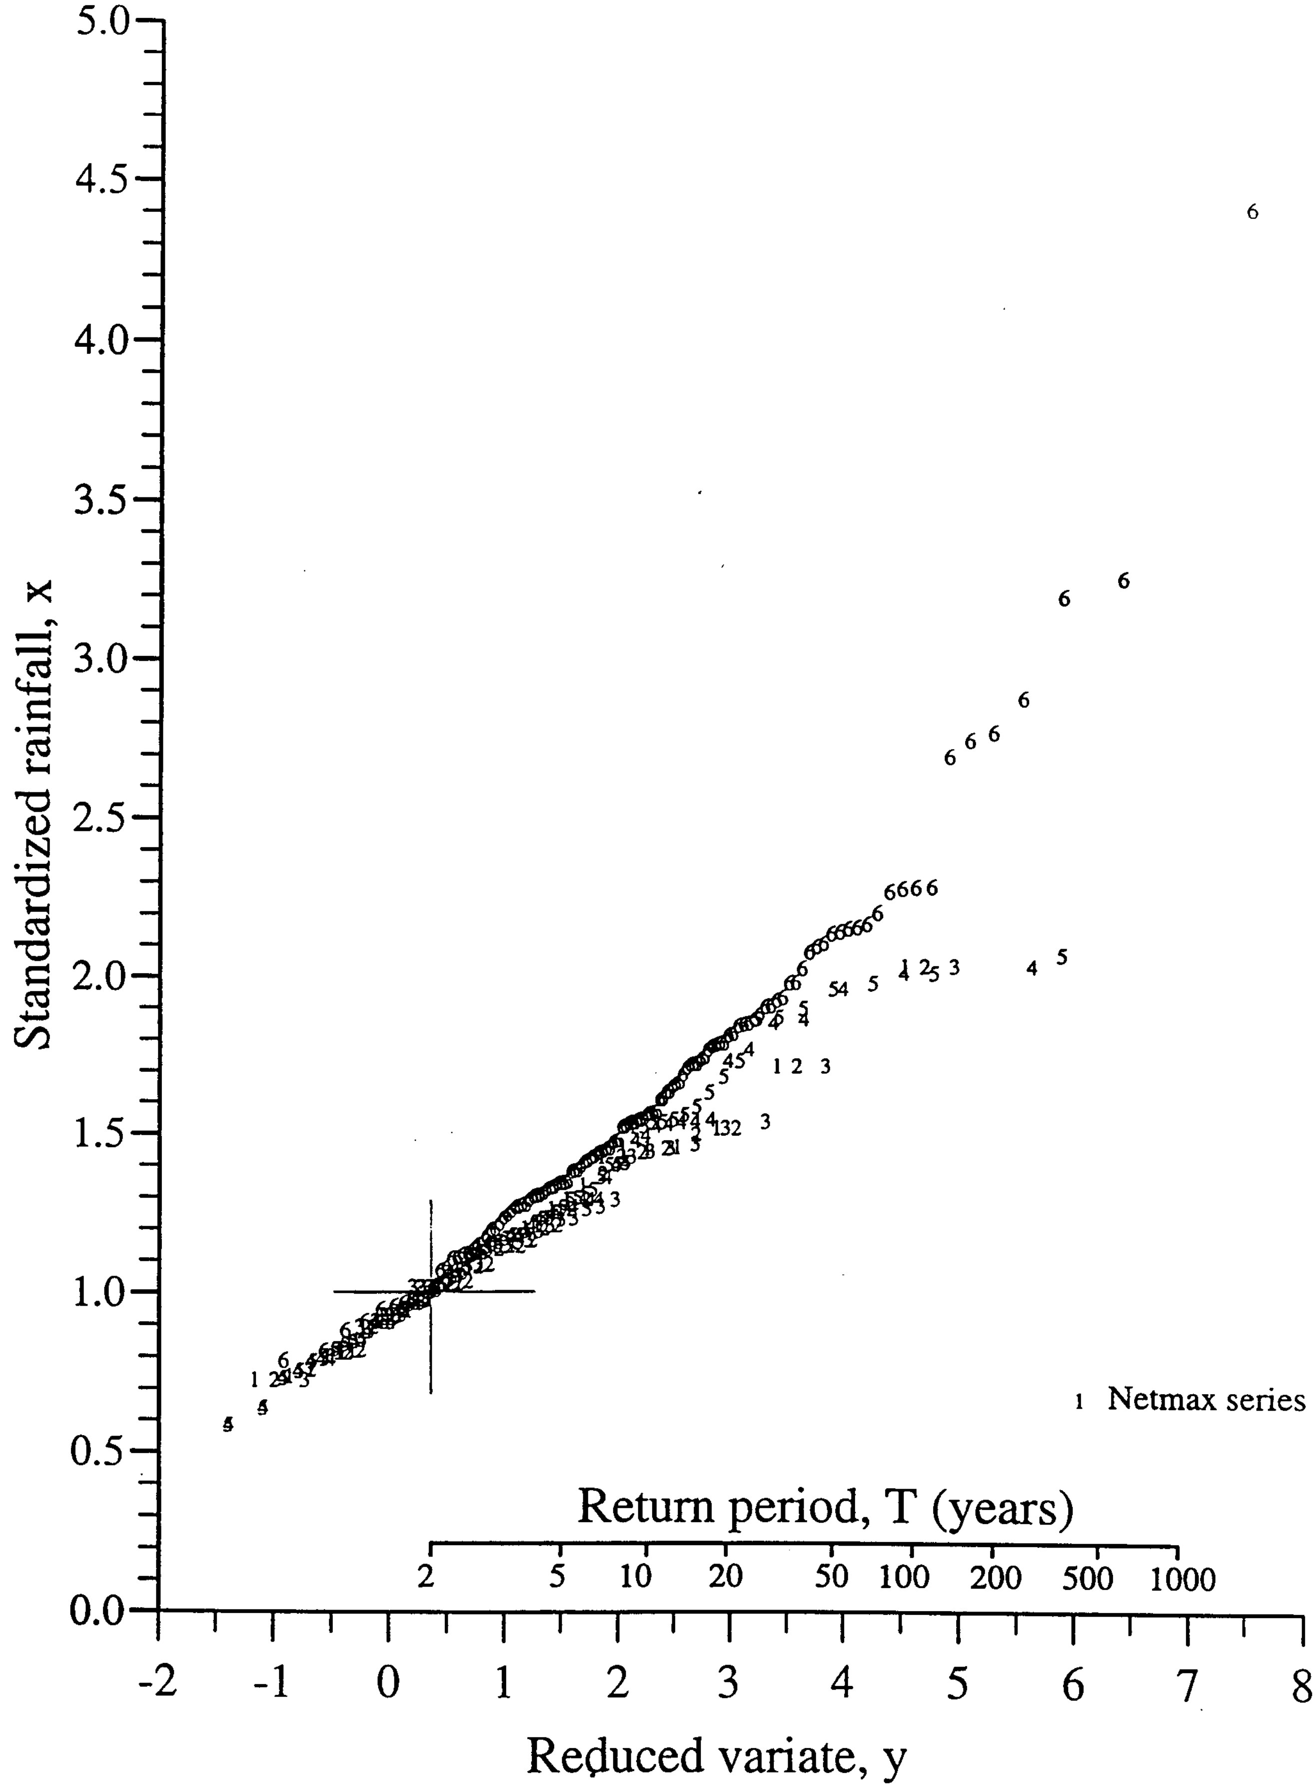

Figure 3.4 **Shifted netmax points focused on Preston with new plotting positions: all series pass close to the median**

relative influence of pooled and netmax points is thus determined by the number of each type of point which is plotted. The rules which determine how many netmax points are plotted are discussed in the following Section.

Large changes of gradient between adjacent growth curve segments are undesirable, particularly when they are due to one segment having negative gradient, a feature which occurred in a small number of growth curves before the introduction of smoothing. Various options for smoothing growth curves were explored, for example changing the width of y-slices, changing the relative weights of pooled and netmax points (see above), using netmax points from adjacent y-slices, or giving different weights to the various points according to their rank.

The chosen method of smoothing involves moderating the least squares fit by a type of penalty function (see for example Adby and Dempster, 1974). The penalty function is added to the sum of squared differences (between observed and modelled standardized rainfall) to give the following objective function by which the n unknown parameters (gradients of the n segments) are determined:

$$\sum_{1}^{M}(observed-modelled)^2 + 0.05\frac{M}{n-1}\sum_{j=2}^{n}(a_j-a_{j-1})^2$$

where M is the total number of pooled and netmax points and a_j is the gradient of the jth segment. The factor 0.05 was chosen experimentally; it provides a degree of smoothing while sustaining characteristic features of the growth curves focused on particular sites.

The effect of the penalty function is demonstrated for a rather extreme example in Figures 3.5 and 3.6. The fourth segment of the unsmoothed growth curve in Figure 3.5 is pulled up by several large local annual maxima. This produces a kink which is primarily a consequence of having a segmented growth curve. The kink is smoothed out by adding the penalty function to the least squares fit (Figure 3.6), while the locally high growth rates for return periods around 50 years are still evident. The effect of smoothing is smaller at most other sites.

3.5 Rules for network sizes, y-slices and growth curve segments

3.5.1 Expanding the gauge network

A key requirement is that the degree of pooling be appropriate to the rarity of event being estimated, with greater pooling at long return periods. Local data should be used where they are adequate for defining the growth rate. For low to moderate return periods, this requirement is met by regulating the network expansion so as to obtain a near-uniform density of pooled data points along the reduced variate axis. For longer return periods, a similar rule is applied to netmax points.

The expansion of the networks which provide pooled points is unchanged from the method published in the Phase 1b report. The target is that there should be no fewer than 20 pooled data points within each y-slice of width 1.0 on the reduced variate scale. Thus, in a sparsely gauged region, the circle defining the jth gauge network will be larger than elsewhere, as the network

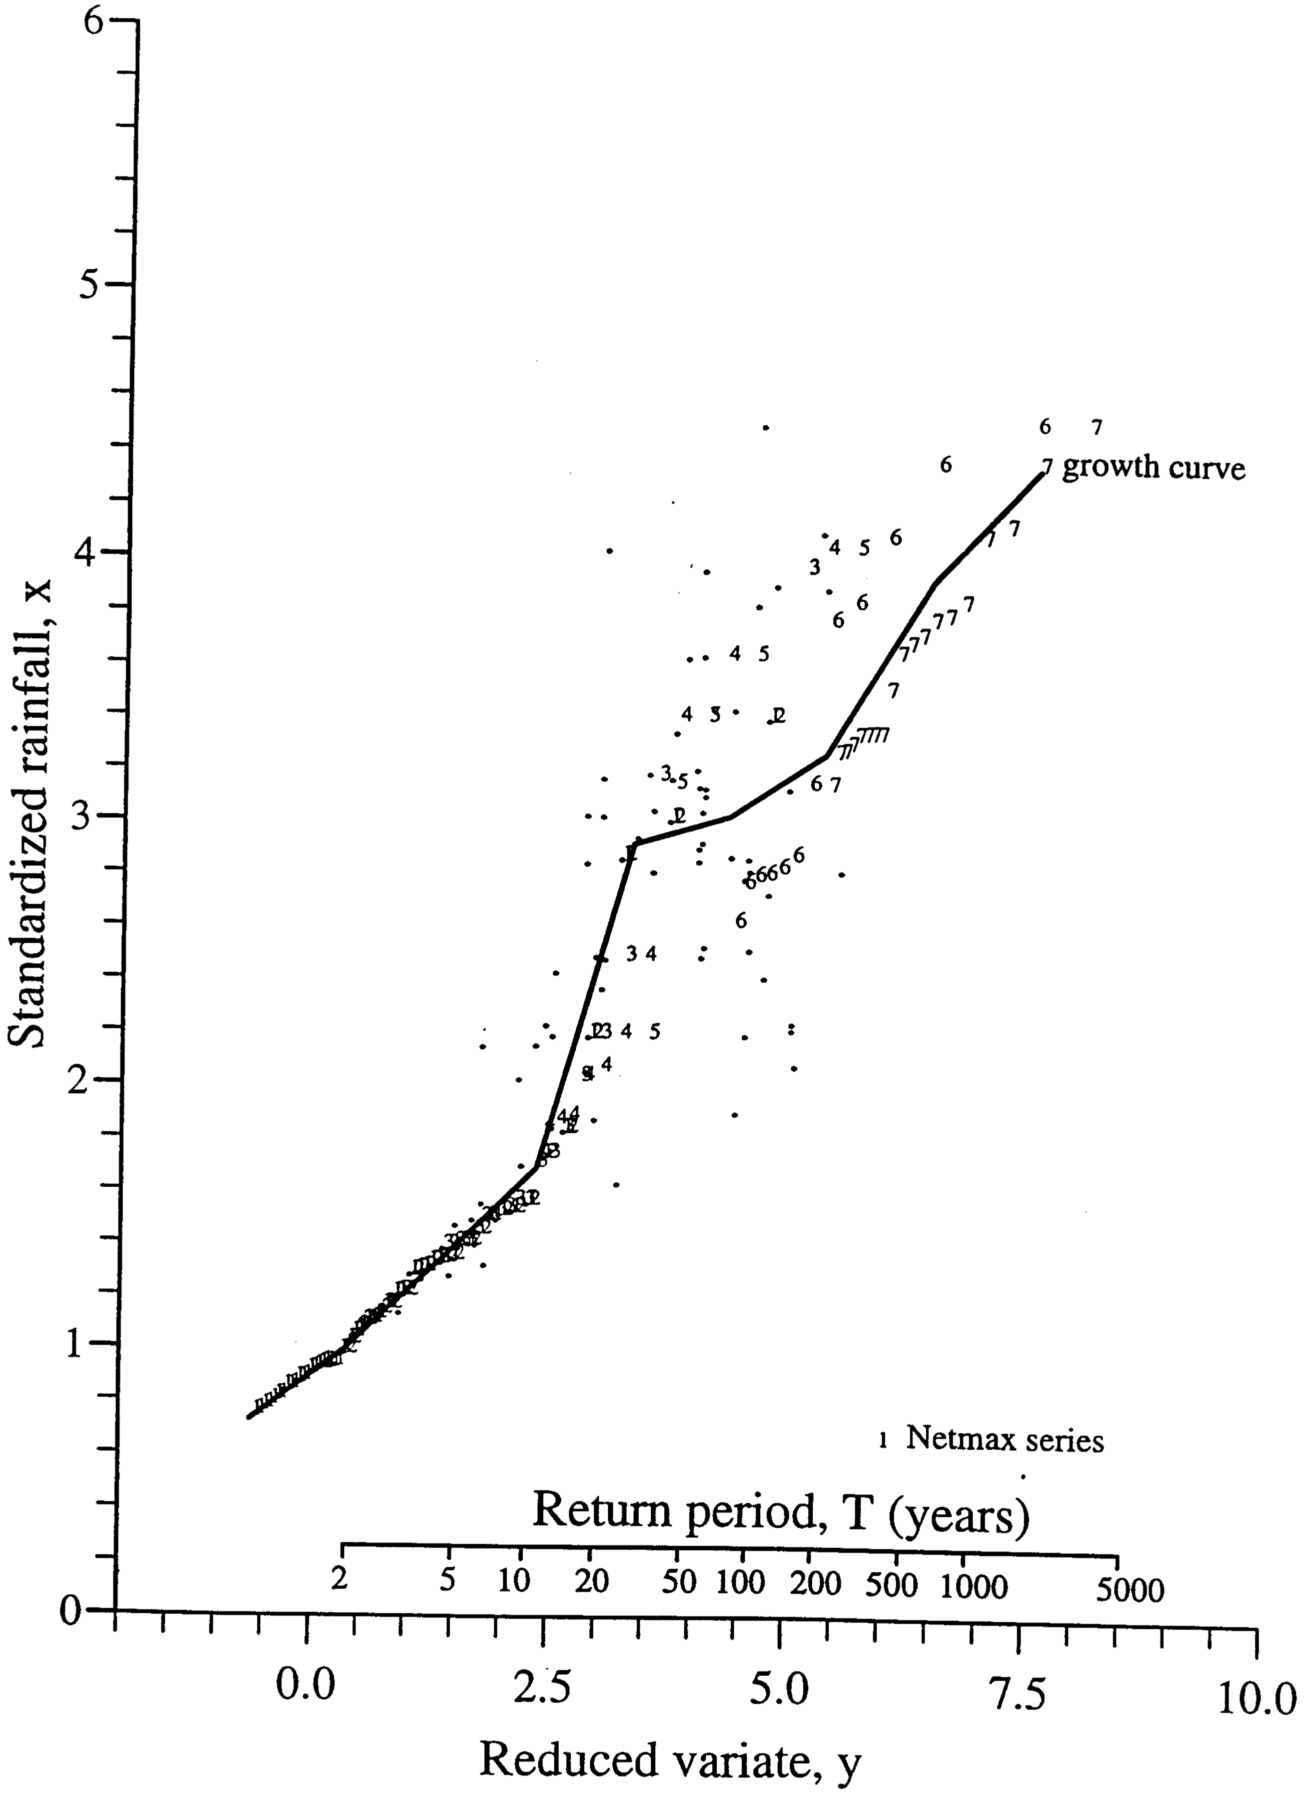

Figure 3.5 Unsmoothed 1-day growth curve focused on Bridgwater

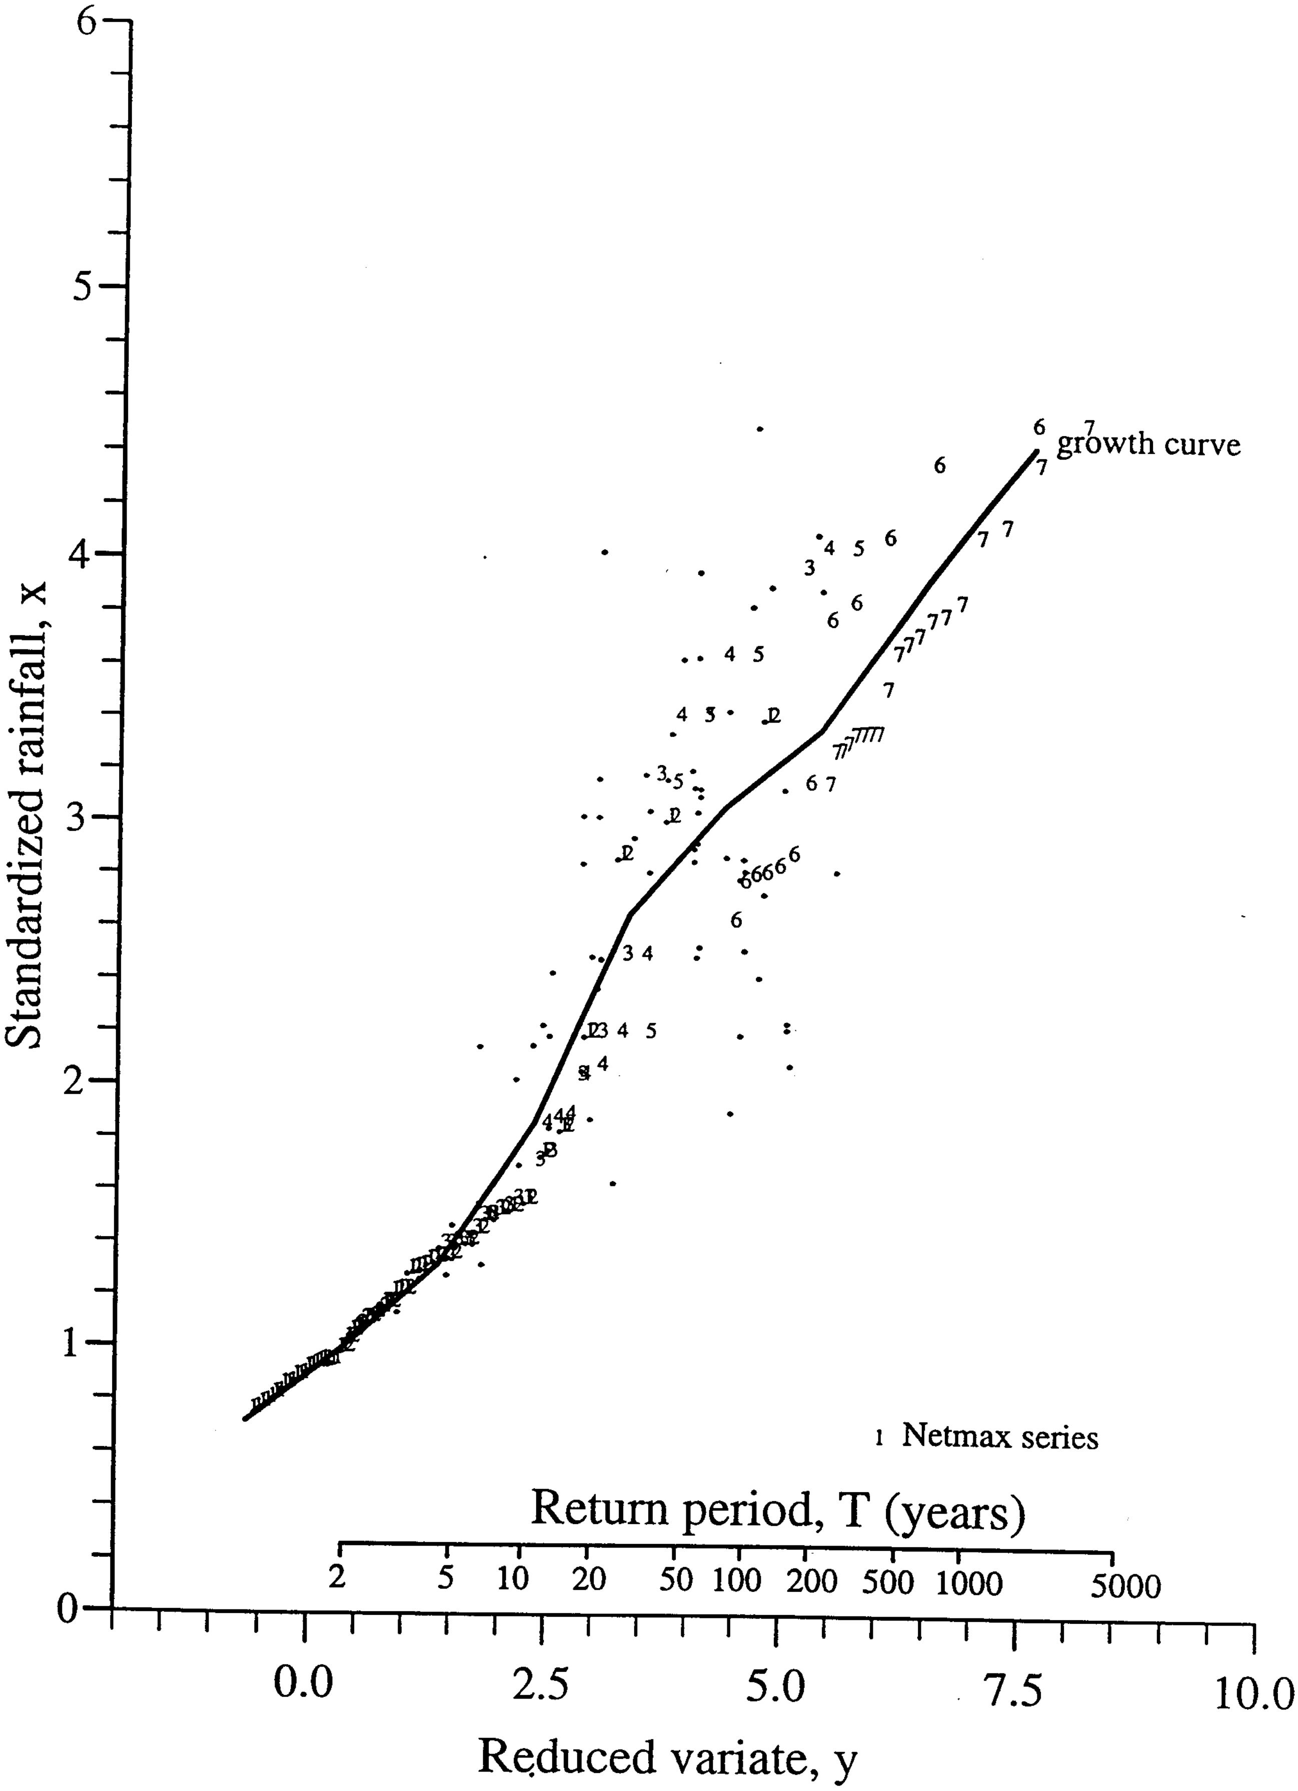

Figure 3.6 **Smoothed 1-day growth curve focused on Bridgwater**

expands to attain the required number of data points in the jth y-slice. This rule defines the first P networks. Typically, depending on the availability of long annual maximum rainfall records in the area, P is five or six, for 1-day rainfall growth analysis.

New rules for defining larger networks to provide only netmax points for higher y-slices (see Section 3.2.2) have been developed. The largest (Lth) network is defined first: this is determined by a limit set for the maximum range over which rainfall data are used, $dist_{max}$. Exceptionally, if the highest shifted netmax point in the Lth network is less than 0.2 units to the right of the highest point in the Pth network, only P networks are used.

An auxiliary rule is required to define any intermediate network between the Pth and the Lth. The procedure followed is to note the plotting positions of the highest shifted netmax points in the Pth and Lth networks. These are shown by downwards pointing arrows in Figure 3.2. If these plotting positions are more than 2.0 units apart, one or more intermediate networks are inserted. In Figure 3.2, the highest netmax points in the 6th and Lth networks are 2.2 units apart on the reduced variate axis, and there is room for one intermediate network (the 7th). The largest network is thus determined to be the 8th.

In choosing how many intermediate networks can be accommodated, the rule is followed that the highest shifted netmax points in consecutive networks should have plotting positions separated by no less than 1.0 and no more than 2.0 units.

The spatial extent of any intermediate network is determined by progressively increasing the defining radius to accommodate additional gauges until the plotting position of the highest shifted netmax value is the required proportion of the distance between the plotting positions of the highest shifted netmax values for the Pth and Lth networks. When using a maximum pooling distance of 200 km, there will rarely be more than one intermediate network in typical UK applications.

3.5.2 Setting the maximum radius

Setting a maximum distance over which data can contribute to growth curve estimation guards against the pooling of data from sites of radically different climate. The distance, d_{max}, is a compromise between this requirement for restrained pooling and the need to gain access to enough observations to support the definition of the growth curve at long return periods. After some experimentation, d_{max} was set to 200 km for UK applications.

3.5.3 Defining the growth curve segmentation by reference to y-slices

For the lower part of the growth curve, the y-slices are defined to be of width 1.0. In the Fig. 3.6 example, the six y-slices based on pooled data points extend the growth curve as far as y=5.3665. For the upper part of the growth curve, somewhat coarser y-slices are used. In order to avoid excessive reliance on the very highest netmax points, the growth curve is taken to extend only as far as the plotting position of the third highest netmax point in the largest network. The upper part of the growth curve is then divided into segments of equal width, the width being no less than 1.0 nor greater than 2.0 units.

For the Fig. 3.6 example, the procedure defines eight segments. Thus the eight hierarchical networks support a growth curve with eight segments. While it is typical to have equal numbers of networks and segments, this need not always be the case. At some sites this apparent discrepancy is necessary to avoid wasting additional netmax points which would otherwise plot to the right of the end of the last segment.

Although the segmented growth curve in Fig. 3.6 is not explicitly defined beyond the eighth y-slice, the uppermost segment of the growth curve is assumed to extend further, and the netmax points that plot to the right of this y-slice are included in the least-squares summation.

4. SUMMARY OF THE FORGEX METHOD

4.1 Introduction

This chapter is a brief summary of the final growth curve estimation method, developed as described in the Phase 1b report and the previous chapter. Detailed rules are not given in this chapter. The acronym FORGEX stands for FOcused Rainfall Growth curve EXtension.

The key features of FORGEX are:

Rainfall frequency estimation is based on the analysis of annual maxima;
- the median of at-site annual maxima, RMED, is used as the index variable;
- individual durations are treated separately in the construction of growth curves;
- growth curves are focused on the site of interest, rather than applying to rigid regions;
- annual maxima are pooled from a network of gauges which expands with return period, giving precedence to the use of local data;
- shifted network maximum rainfalls account for inter-site dependence in rainfall extremes;
- the growth curve is seamlessly extended to long return periods;
- the growth curve is made up of linear segments on a Gumbel scale, avoiding an explicit distributional assumption.

4.2 A focused method of growth curve estimation

4.2.1 Annual maxima and RMED

The raw materials for growth curve estimation are series of annual maximum rainfalls of a certain duration, for example 1 day. Data from several sites are used because record lengths are seldom long enough to estimate growth curves directly from annual maxima observed at the subject site.

Growth curve construction is concerned with estimating the T-year extreme rainfall relative to an index extreme rainfall. The use of an index, or standardizing, variable insulates the regional analysis from site-specific values which influence the typical magnitude of extreme rainfalls but may have little systematic effect on their variability and skewness. The annual maximum series is standardized by dividing by the at-site median, RMED. The median has the advantage over the mean of being less sensitive to individual large sample values. Also the median corresponds to a fixed return period of two years on the annual maximum scale. This point of reference simplifies and unifies growth curve construction.

A required minimum record length of ten years ensures that RMED is reasonably estimated.

4.2.2 Focusing

The FORGEX method pools data from a hierarchy of raingauge networks. Data are pooled according to distance from the subject site, i.e. in circular regions. As successive networks are constructed, the radius of the circle is increased and annual maximum rainfalls are drawn together from larger numbers of sites. The name of the method is taken from this idea of focusing the analysis on the site for which the growth curve is required. An upper limit of 200 km (for the UK) is set for the network radius. Figure 4.1 shows the networks used for deriving a 1-day growth curve focused on Leicester.

4.2.3 Plotting positions

Annual maxima from individual records are ranked and allocated plotting positions on a Gumbel reduced variate scale. Following established practice, the Gringorten plotting position formula is used:

$$F(i) = (i-0.44) / (N+0.12)$$

where F(i) is the non-exceedance probability, i the rank in increasing order, and N the number of annual maxima. The Gumbel reduced variate is defined by:

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

The pooling of standardized annual maxima from different stations is presented graphically by superposing extreme value plots for each site. However, not all points are plotted. Note that only points which are plotted are used to fit the growth curve.

4.2.4 Definition of y-slices

Each raingauge network in the hierarchy is associated with the definition of the growth curve within a particular y-slice. The y-slices have width 1.0 on the reduced variate scale, and the first one ends at y = 0.3665 which is the position of the median. Pooled data points which are plotted within the jth y-slice come from gauges within the jth network. This ensures that local data are used in preference to data from further afield.

Larger networks include more long-record stations, and thus provide pooled data points that plot in y-slices that correspond to rarer events. However, there are few sites in the UK with records longer than about 100 years. This means that pooled points alone cannot define the growth curve beyond about the fifth y-slice.

4.2.5 Network maximum points

The *network maximum* (netmax) series is defined as the annual maximum series of the largest standardized value recorded by the network of raingauges. Dales and Reed (1989) show that the distribution of the network maximum from N independent and identically distributed GEV distributions lies exactly ln N to the left of the regional growth curve, and Reed and 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 and Reed label this distance ln N_e, terming N_e the effective number of independent gauges.

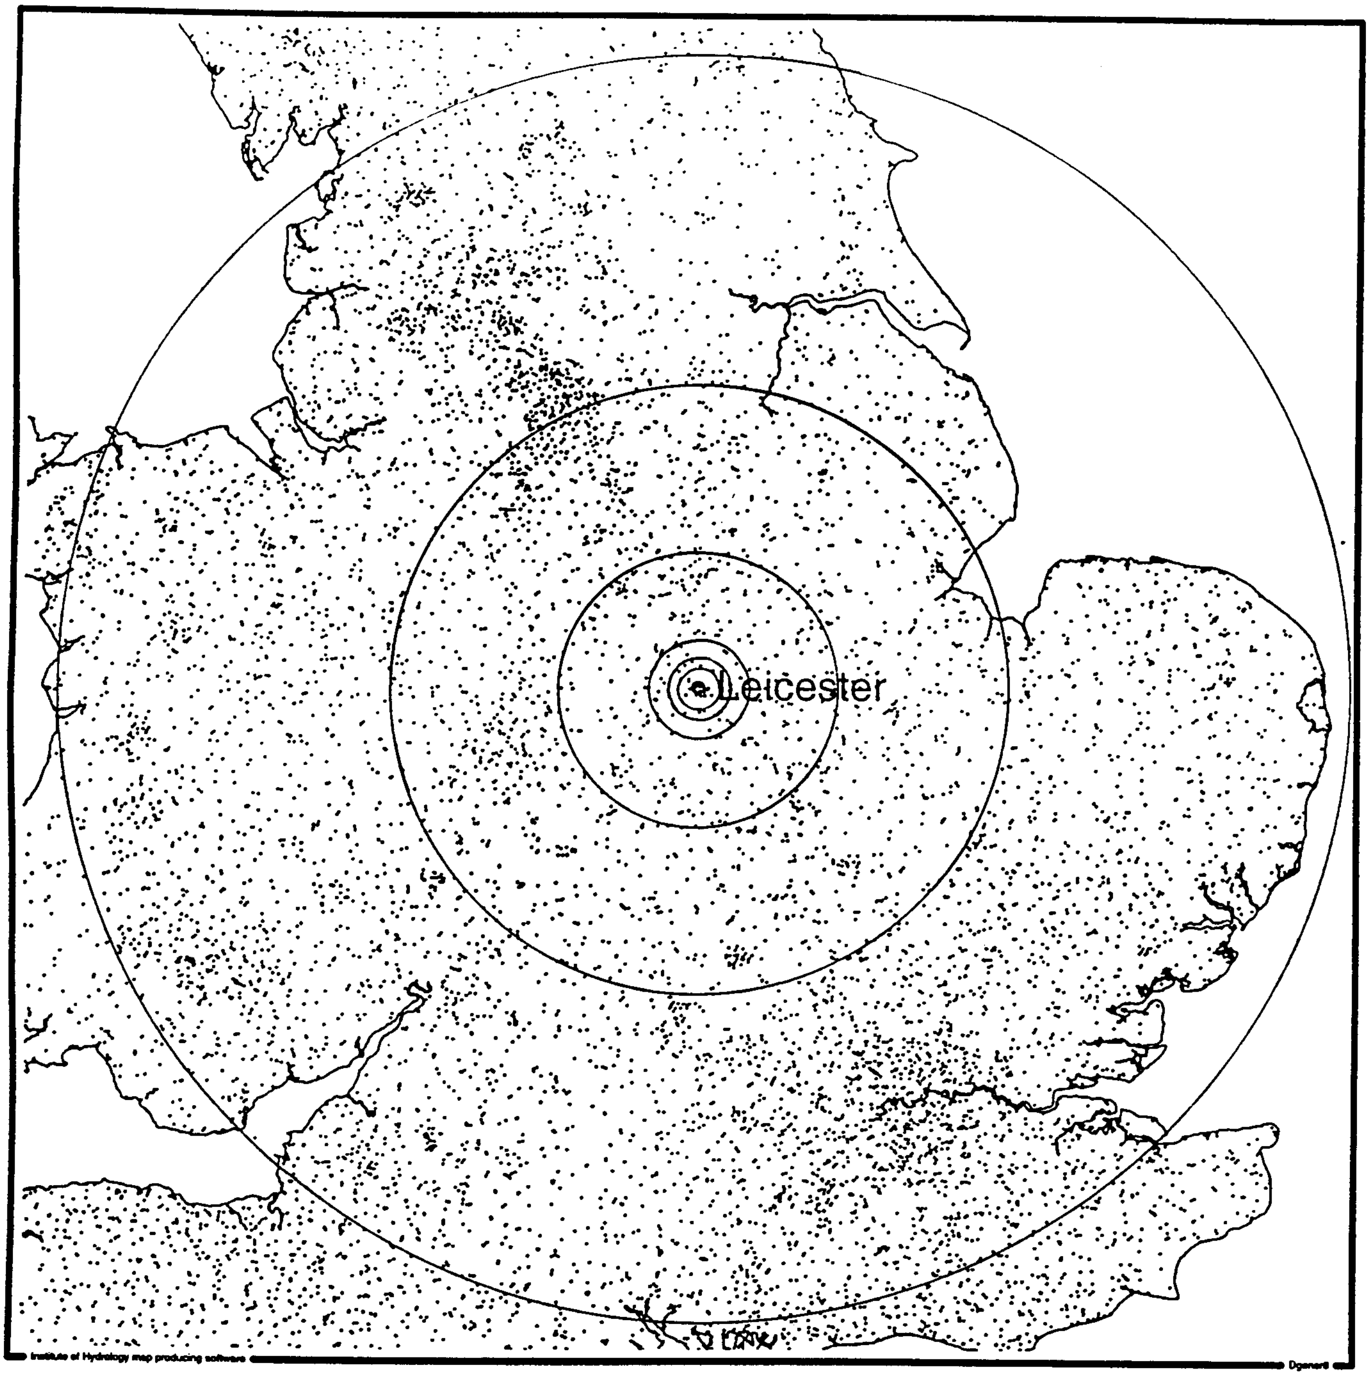

Figure 4.1 **Networks used for deriving a 1-day growth curve focused on Leicester. The largest network has radius 200 km.**

Dales and Reed developed a model of spatial dependence to estimate the offset distance, ln N_e, for any raingauge network. The UK model is:

$$\ln N_e / \ln N = 0.081 + 0.085 \ln AREA - 0.051 \ln N - 0.027 \ln D$$

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

$$AREA = 2.5(\bar{d})^2$$

where denotes the mean inter-gauge distance evaluated pairwise across all combinations.

The FORGEX method supplies netmax points for use in extending the growth curve to long return periods in a way which parallels the derivation of the spatial dependence model. For each network in turn, the netmax series is constructed and the values ranked. The plotting position of each point includes a shift to the right by ln N_e where N_e is the effective number of independent stations estimated using the Dales and Reed model. N_e is estimated separately for each point from the characteristics of the gauge network operational in the year concerned.

4.2.6 Fitting the growth curve

The rainfall growth curve is represented as a concatenation of linear segments. Because of the standardization by the median, the growth curve is constrained to take the value 1.0 at a return period of 2 years. Thus fitting the growth curve involves only determining the gradient of each segment.

The growth curve is fitted to pooled and network maximum points by a least-squares routine, and smoothed to avoid large changes in gradient between adjacent segments.

5. FORGEX RESULTS

5.1 Introduction

This chapter discusses the rainfall frequency estimates produced by the FORGEX method. The method produces growth curves, examples of which are shown for focal points in the East Midlands and south-west England. To give a nationwide picture of the results, maps of growth rates for selected durations and return periods have been compiled. These are displayed and discussed in detail. The spatial consistency of growth rates over short distances is also considered.

5.2 Growth curves

5.2.1 East Midlands

Figure 5.1 is a collection of growth curves for 1-hour rainfall at sites in the East Midlands, the region which was used for the pilot study in Phase 1b of the project. Note that all growth curves pass through the point where T = 2 years and x = 1, a consequence of standardizing by the median, RMED. The growth curves extend to return periods of 700 to 800 years. Growth rates do not vary greatly between the sites, although there is some evidence of a regional trend which was also noticed in the rather shorter growth curves shown in the Phase 1b report. Sites in the north and east of the region (for example Lincoln and Boston) tend to have rather higher growth factors than sites to the south and west.

One-day growth curves for the same focal points are shown in Figure 5.2. These extend to a return period of around 2000 years in this region, thanks to the plentiful data within the maximum range of 200 km. There is more variation between sites at this duration, particularly in the middle of the return-period scale. Growth curves vary in shape as well as scale, so that no single site has the highest growth rates over the complete range. However, the regional trend mentioned above is evident for 1-day rainfall as well, the highest growth curves being for sites in the east of the region, notably Peterborough and Boston. The lowest growth curves are those for Northampton and Leicester in the west. The 200-year growth rate at Boston is 28% higher than that at Northampton.

5.2.2 South-west England

Rainfall frequency estimation in south-west England has received much attention since the publication of the Flood Studies Report. Focused rainfall growth curves for ten sites in south-west England are shown in Figure 5.3. At moderate return periods (8 to 150 years), the curve for Bridgwater is the highest by some margin. At longer return periods, other sites in the east of the region have the greatest growth rates. Growth rates in the west, near the coast, at sites like Plymouth, are lower.

These results reflect the occurrence of unusually high rainfall totals in the counties of Somerset, Dorset and Devon, some of which are described by Bootman and Willis (1977). This anomaly in the national rainfall pattern has been attributed to the geography of the area, in particular the central basin in Somerset surrounded by hills, well suited to the development of thunderstorms. The occurrence of infrequent but severe and prolonged convective storms is the sort of phenomenon which could give rise to a high 1-day rainfall growth rate, by producing a few very large annual maxima without increasing the median.

The shape of the Bridgwater growth curve, with a "bump" between the 10 and 50-year return periods, occurs because this is the closest focal point to several of the historic notable rainfalls. These rainfall totals are recorded by gauges which form a relatively small network around the focal point, and so they plot at lower return periods than they do for other focal points.

Bootman and Willis (1977) suggest that the FSR underestimates 1 and 2-day extreme rainfalls in Somerset and Dorset. FORGEX growth curves for 1 and 8-day rainfall at Bridgwater in Somerset are compared with the results of the FSR in Figure 5.4. The FSR growth curves are standardized by the 2-year rainfall so that they can be plotted on the same scale as the focused growth curves. Bootman and Willis's suggestion is borne out: for a return period of 100 years, the FORGEX 1-day growth rate is 35% greater than the FSR estimate. Put another way, a 1-day rainfall of three times the median annual maximum is five times more likely according to FORGEX than according to the FSR. The 8-day growth curve is much closer to the FSR result. Bridgwater is an extreme example of the difference between FSR and FORGEX results. However, further comparison with FSR results is kept until Chapter 9, where design rainfalls in millimetres are compared nationwide.

5.3 Maps of growth rates

These maps were constructed by running the FORGEX algorithms for a given rainfall duration at every point on a 10×10 km grid across the UK. Growth rates for several return periods were estimated for each site. The software was streamlined for this purpose, but it is not yet quick enough to produce a map with a 1 km resolution within a reasonable time (a few days). The 10 km maps do however reveal regional and local variations in growth rates. Maps for a selection of durations and return periods are displayed here.

Note that results are shown for some squares over the sea, to aid the identification of peaks and troughs in growth rates. Growth rates are not calculated for squares far from land, to speed up the software. The limiting distance from the coast is defined by distance to the closest raingauge, and it is longer for hourly durations and in northern Scotland to ensure that no grid squares over the land are missed. Some land areas are missed off maps for longer return periods such as Figure 5.6 (the Western and Northern Isles) and Figure 5.10 (more of the western and northern parts of the UK). This is because FORGEX growth curves for these remoter areas do not extend to such long return periods.

Hourly durations

Figures 5.5 and 5.6 show 1-hour growth rates for return periods of 20 and 100 years. There is an overall decrease of growth rates from south-east to northern Britain, but this is obscured in some regions by spatially extensive peaks in growth factors. The largest peaks are in Greater London, Lancashire, north Wales and Herefordshire. It is difficult to find a single characteristic that these areas all have in common; the peaks are probably due to several factors. There are several unusually high standardized annual maximum 1-hour rainfalls (greater than 5.0) in these areas, which pull up the growth curves.

Although the peaks may look dramatic, the nationwide variation in growth rates (for all durations and return periods) is much smaller than the variation in RMED (see the maps in chapter 8). For example, the 1-hour 100-year growth rate varies over the UK from 2.7 to 4.2, whereas the 1-hour RMED (Figure 8.4) varies from 7 to 17 mm. Thus the final combined maps of design rainfall in chapter 9 resemble the RMED maps more closely than the growth rate maps.

Figure 5.7 shows 6-hour growth rates for the 100-year return period. There is a clearer national trend towards higher growth rates in the south-east. Six-hour growth rates are lower than those for 1 hour, as expected. There are a few areas where the 6-hour growth curves do not extend to 100 years. The anticipated acquisition of more data for south-west England and Yorkshire would help to fill in two of the gaps.

Daily durations

One-day growth rates for return periods of 20, 100 and 1000 years are mapped in Figures 5.8 to 5.10. The pattern of growth rates changes with return period, illustrating the greater scope for different shapes of 1-day growth curves at different focal points. For T=20 years, the highest growth rates are confined to four small areas in Somerset, the Fens, east Norfolk and north Kent. In the latter two cases, the influence of notable storms can be easily traced, at Ditchingham, Norfolk, in 1994, and West Stourmouth, Kent, in 1973. The lowest growth rates are found in western Scotland, north-west England, south Wales and the south coast of England.

For T=100 years (Figure 5.9) the pattern is similar but rather more blurred, with peaks in growth rate merging into each other. There is a distinct concentration of the lowest growth rates in north-west England, as hinted at by Dales and Reed (1989). For the 1000-year return period, the pattern changes. Growth rates appear to decrease almost radially from East Anglia, with a clear east-west trend. The effect appears to mirror the "continentality factor" introduced in the FSR (Volume 2, section 3.3.3). There is a local increase in growth rates in Dorset due to data from the Martinstown storm of 1955. Not unexpectedly, growth rates tend to be largest in areas where RMED is smallest (see Figure 8.3). There is therefore a cancelling-out effect when growth rates are multiplied by RMED to give the required design rainfall estimates (see Figure 9.8).

The large 1000-year growth rates in East Anglia were investigated thoroughly to ensure that they were not influenced by characteristics of the raingauge network (such as the obvious but unfortunate fact that there are no gauges at sea). Two additional maps help to explain the pattern of 1-day growth rates for long return periods. Figure 5.11 shows the mean of the highest ten network maximum 1-day rainfalls from the largest gauge network around each site. The portion of the growth curve for return periods close to 1000 years is generally fitted to these network maximum points. It can be seen that the mean is highest over a broad area from the southern Midlands to East Anglia.

This is not entirely due to meteorological factors, as shown by Figure 5.12 which maps the equivalent number of independent station-years in the largest gauge network. The number is greatest in the centre of England, because a 200-km network centred at that point includes more gauges then anywhere else (because the point is furthest from the coast, and also the density of gauges is high in central England). In the southern Midlands, it is not surprising that the mean of the top ten network maxima from the largest network is large, because that network includes a greater number of station-years than elsewhere. These high network maxima are therefore plotted at long return periods and so the 1000-year growth rate is not particularly large. Only in East Anglia are large network maxima co-located with a relatively small number of station-years. These network maxima are plotted at rather shorter return periods, and thus the growth curve is justifiably highest in East Anglia.

Finally, Figure 5.13 shows 8-day growth rates for T=100 years. This map has some features in common with the 1-day map, but the peak in growth rate over Somerset is much less marked, and there is a new peak over NE Scotland, which contributes to a clear east-west tend in growth rates for 8-day rainfall.

5.3 Spatial consistency of growth rates

The maps of growth rates enable a brief look at spatial consistency. What are the largest differences between growth rates at adjacent focal points (separated by 10 km)? Are these justified?

The local variation appears to be strongest for 1-day rainfall at the 20-year return period. For example, in Essex (at the western extremity of the Thames estuary as shown on Figure 5.8), there is a grid point with growth rate greater than 2.4. The point to the north (near Basildon) has a growth rate of 2.0 - 2.1. The plotting of growth curves for these two points reveals that the high 20-year growth rate is influenced by a considerable number of local rainfall events, which appear in networks 1, 2 and 3 (Figure 5.14). At the focal point 10 km further north, these events do not appear until the 4th or 5th networks (Figure 5.15), and thus do not exert much influence in fitting the fourth segment of the growth curve (which contains the 20-year return period). The fact that the difference in growth rates is caused by at least five separate rainfall events is encouraging.

Section 6.3 also mentions the spatial consistency of the results, measuring the difference in growth rates over a much shorter distance.

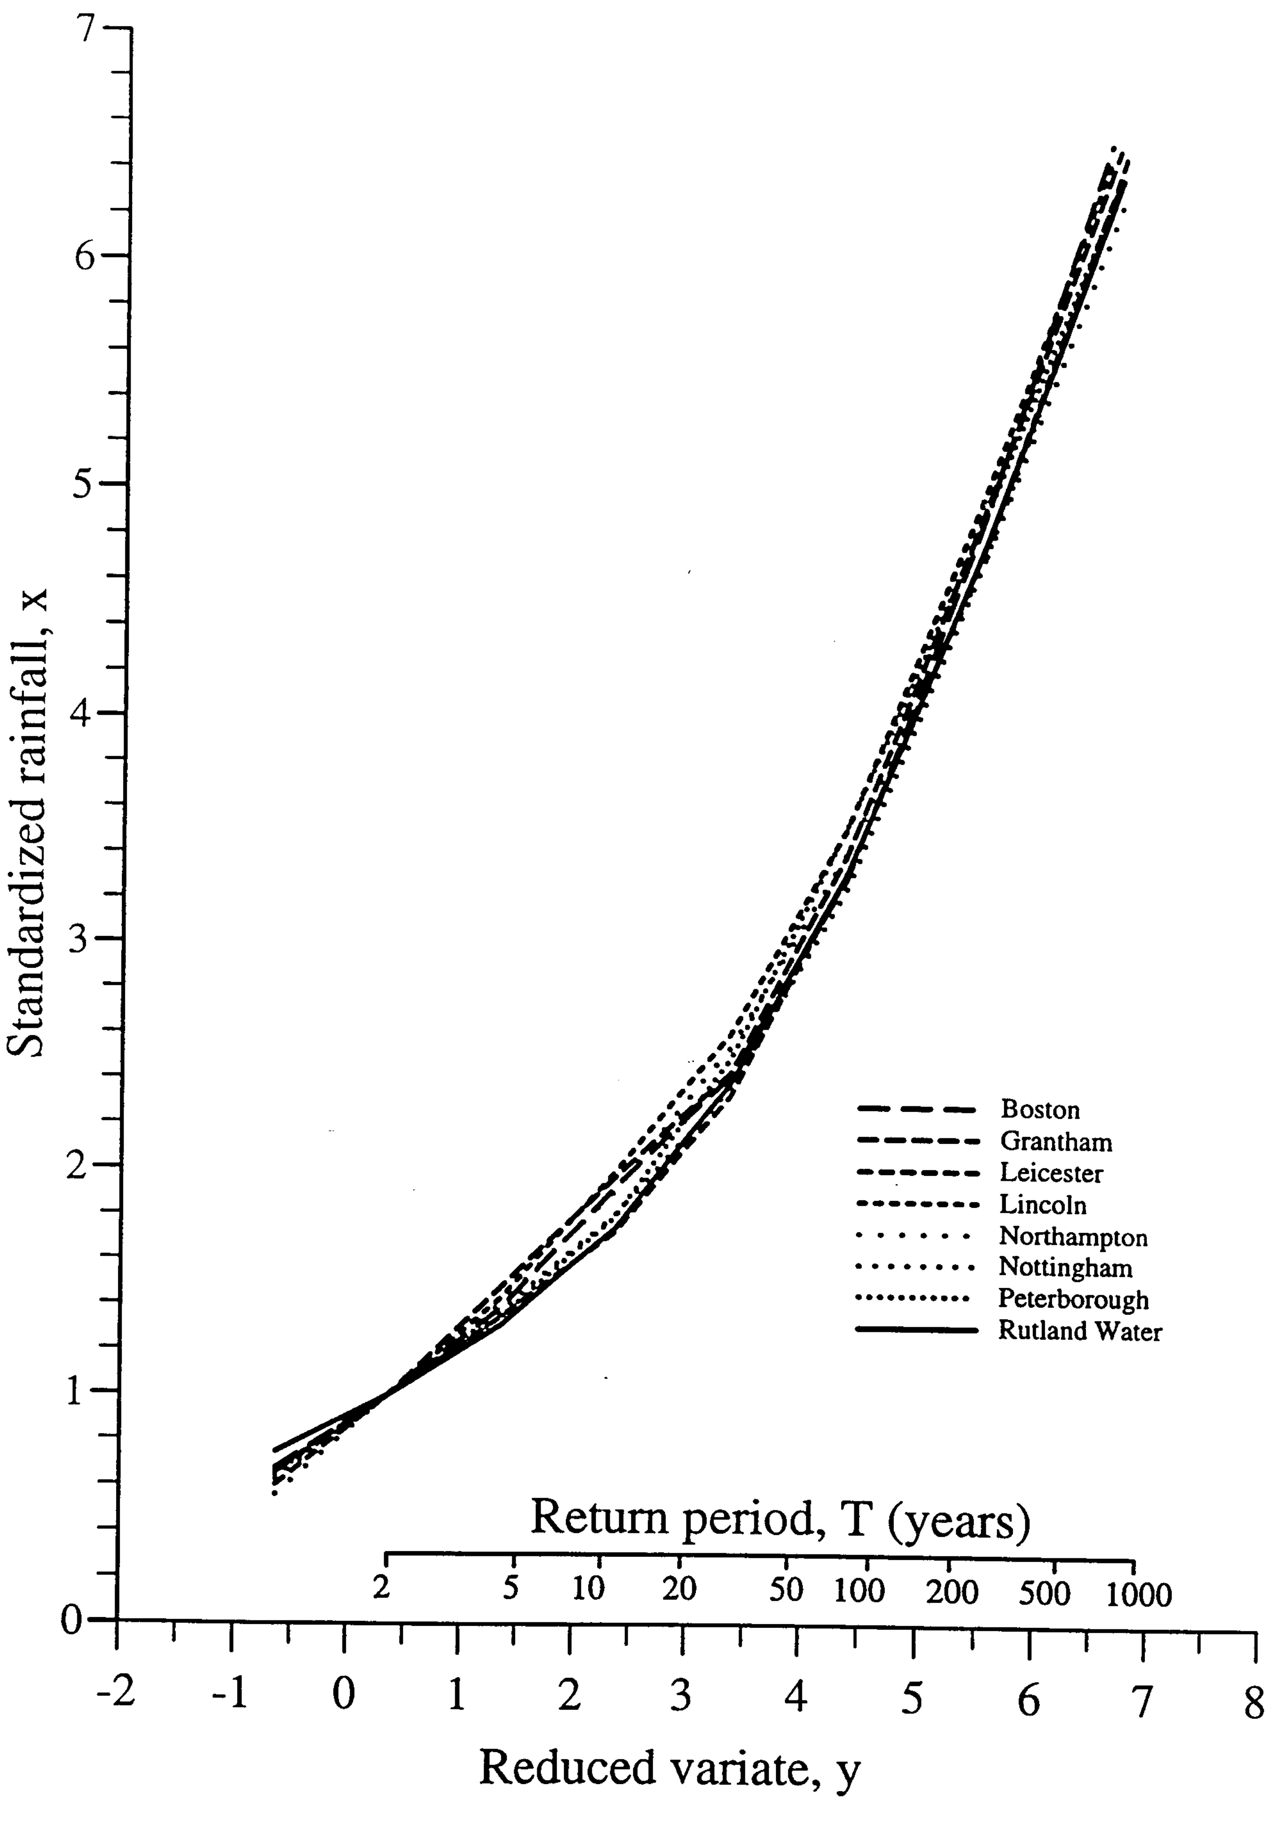

Figure 5.1 1-hour growth curves for the East Midlands

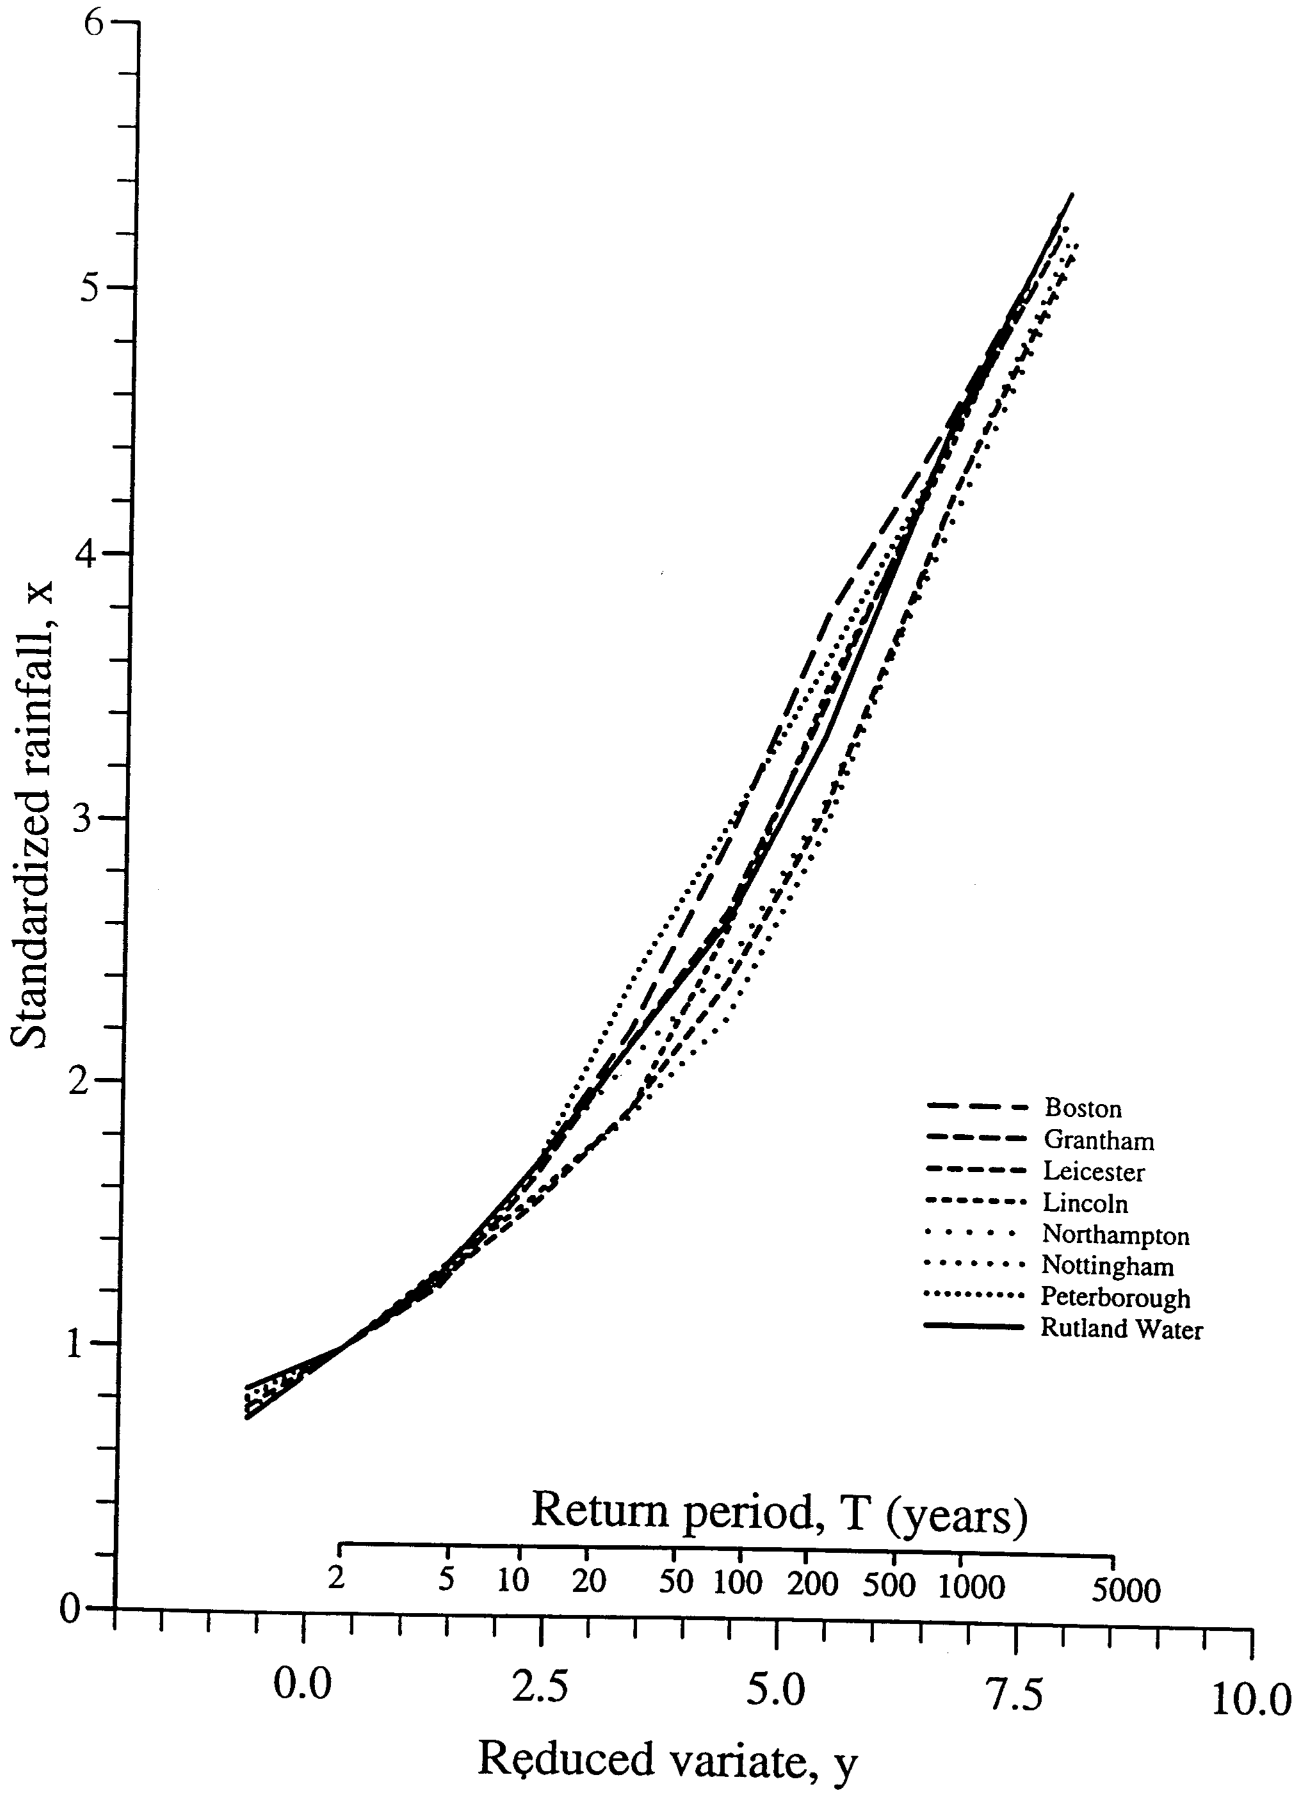

Figure 5.2 1-day growth curves for the East Midlands

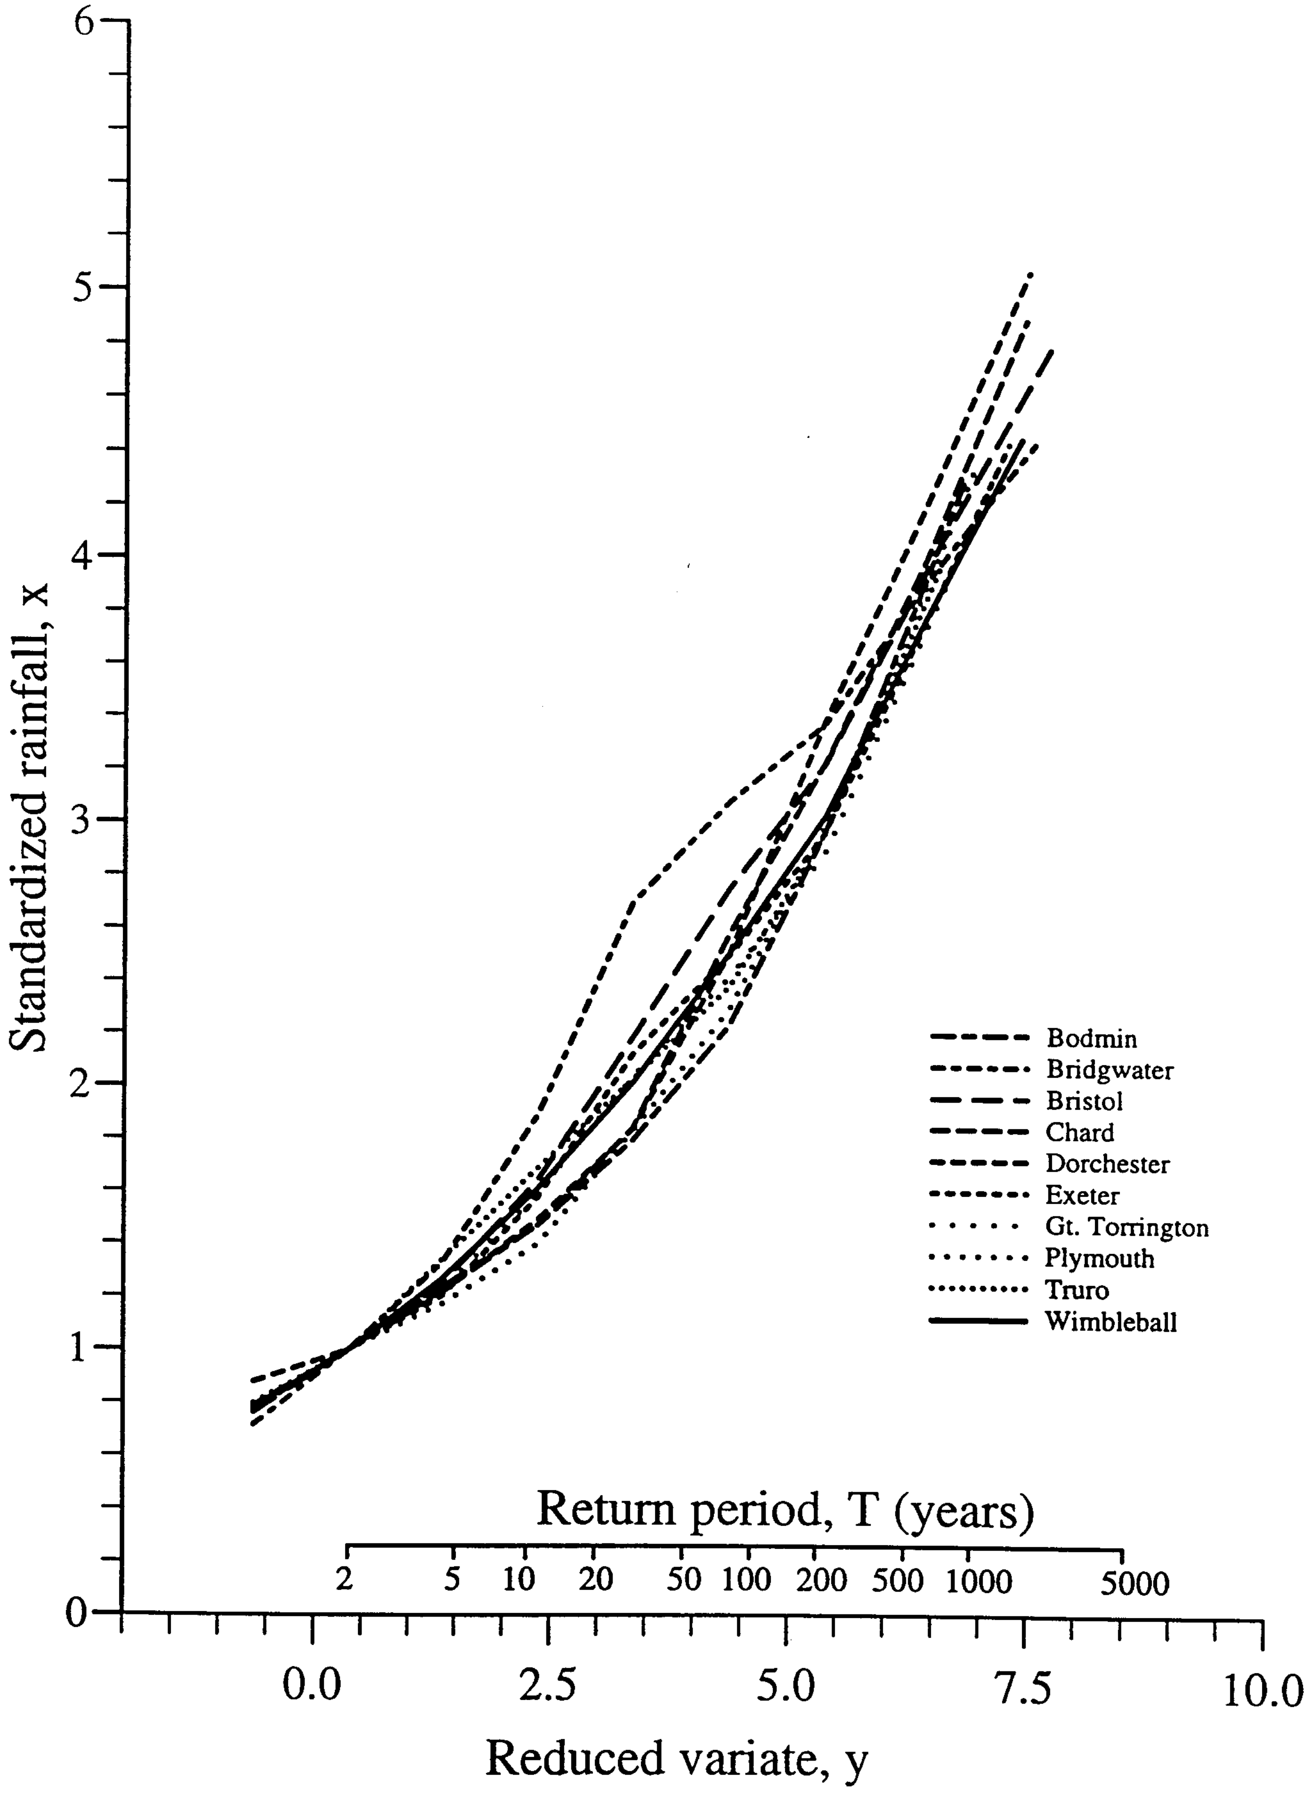

Figure 5.3 1-day growth curves for south-west England

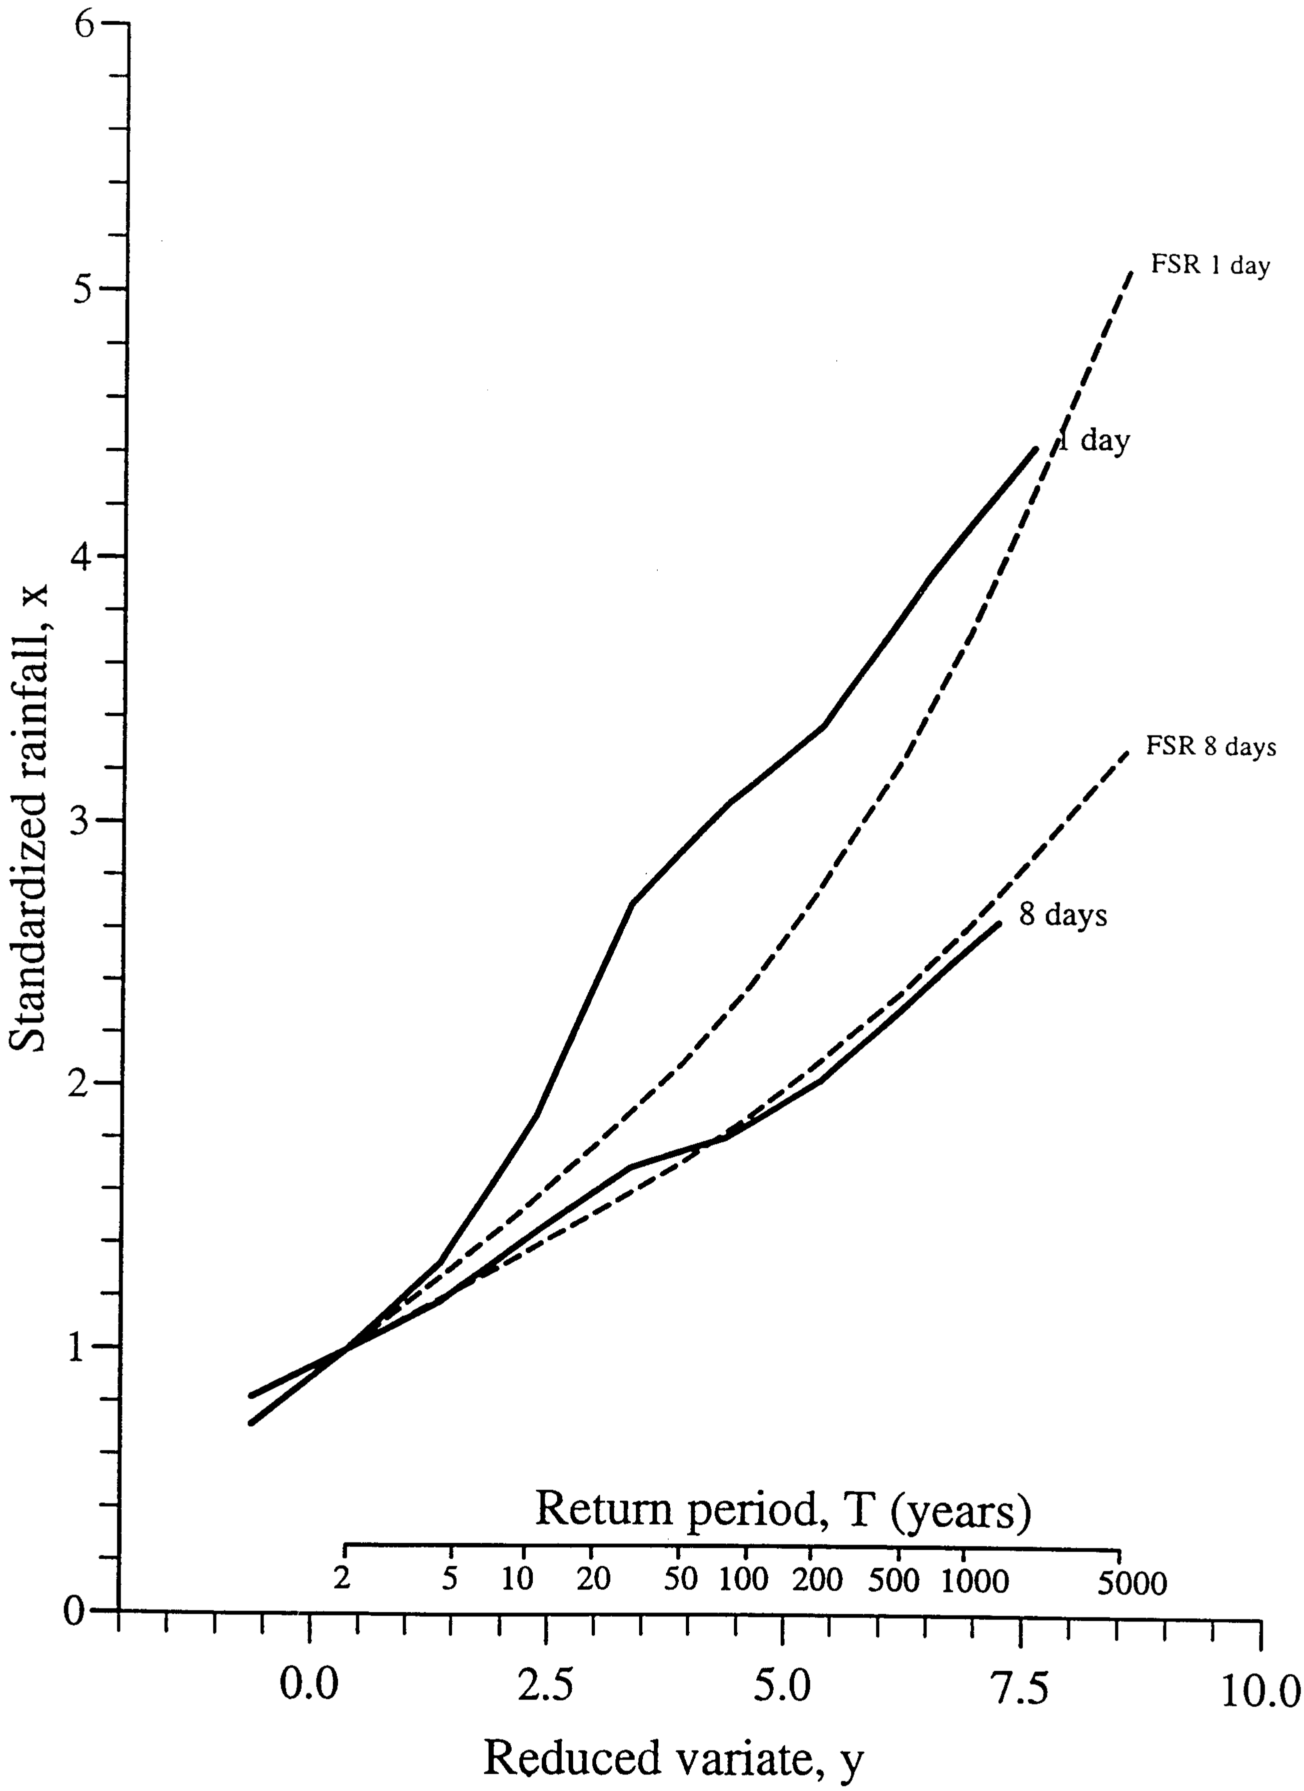

Figure 5.4 FORGEX and FSR growth curves for 1 and 8-day rainfall at Bridgwater

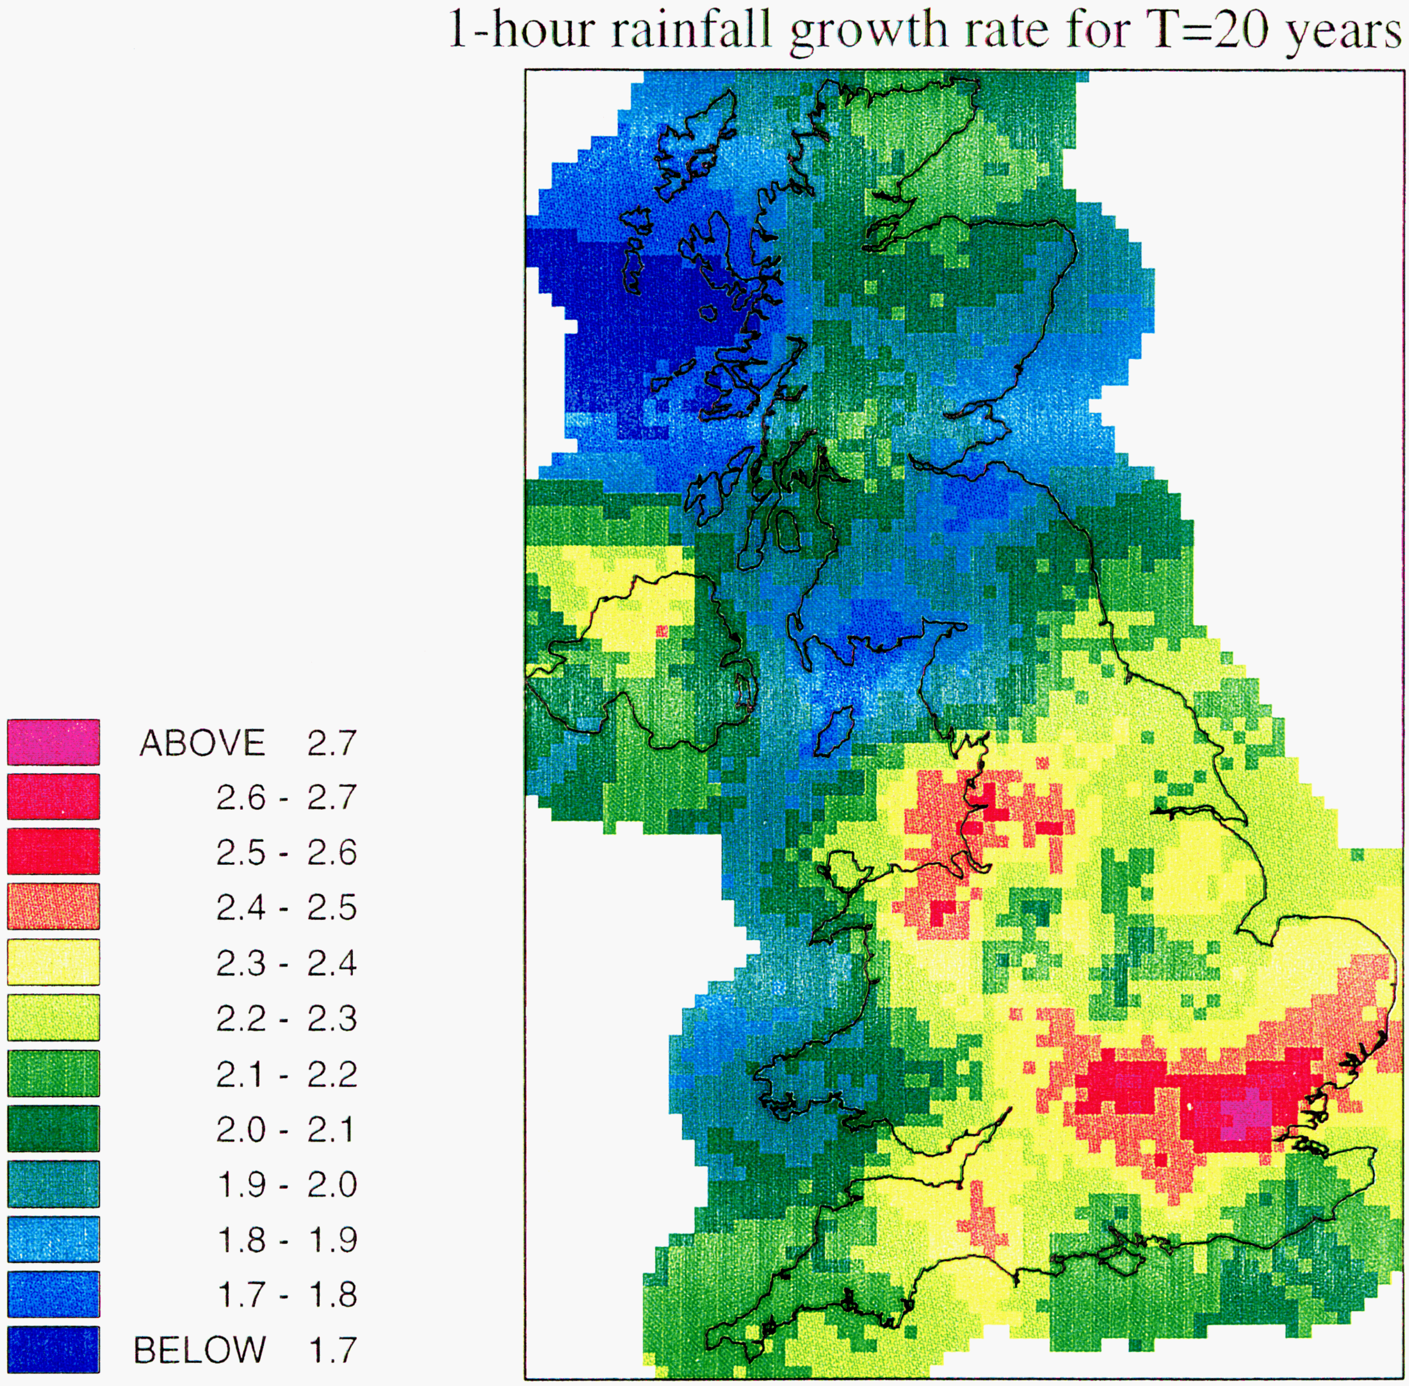
1-hour rainfall growth rate for T=20 years
ABOVE 2.7
2.6 - 2.7
2.5 - 2.6
2.4 - 2.5
2.3 - 2.4
2.2 - 2.3
2.1 - 2.2
2.0 - 2.1
1.9 - 2.0
1.8 - 1.9
1.7 - 1.8
BELOW 1.7

Figure 5.5

1-hour growth rate for T=100 years

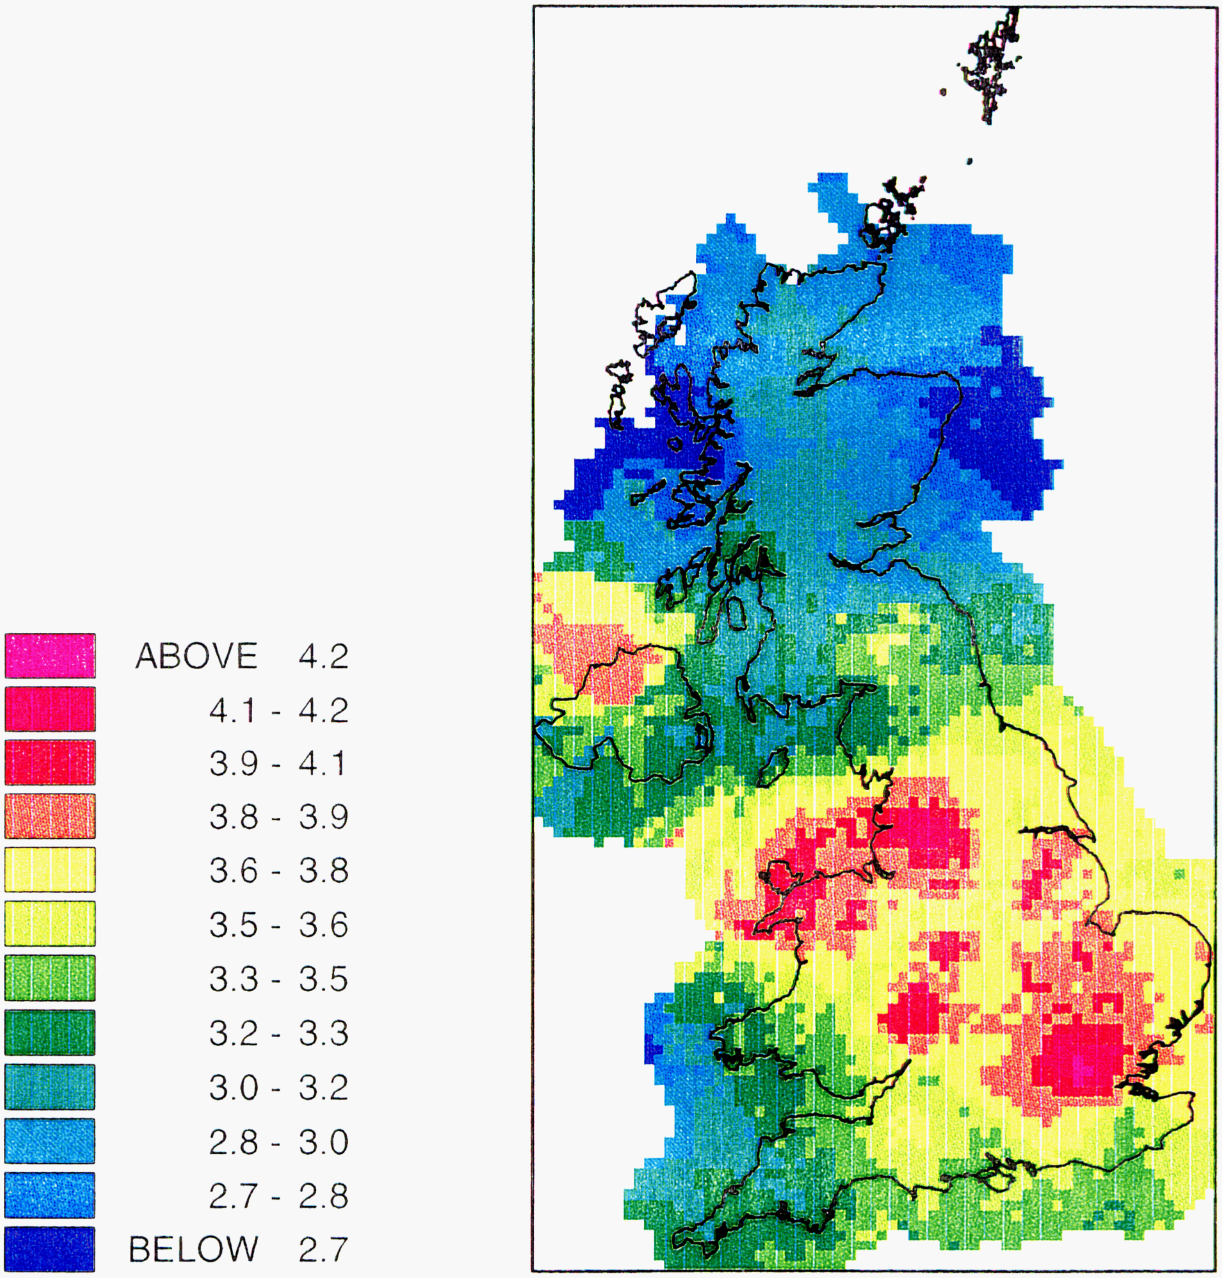

Figure 5.6

6-hour rainfall growth rate for T=100 years

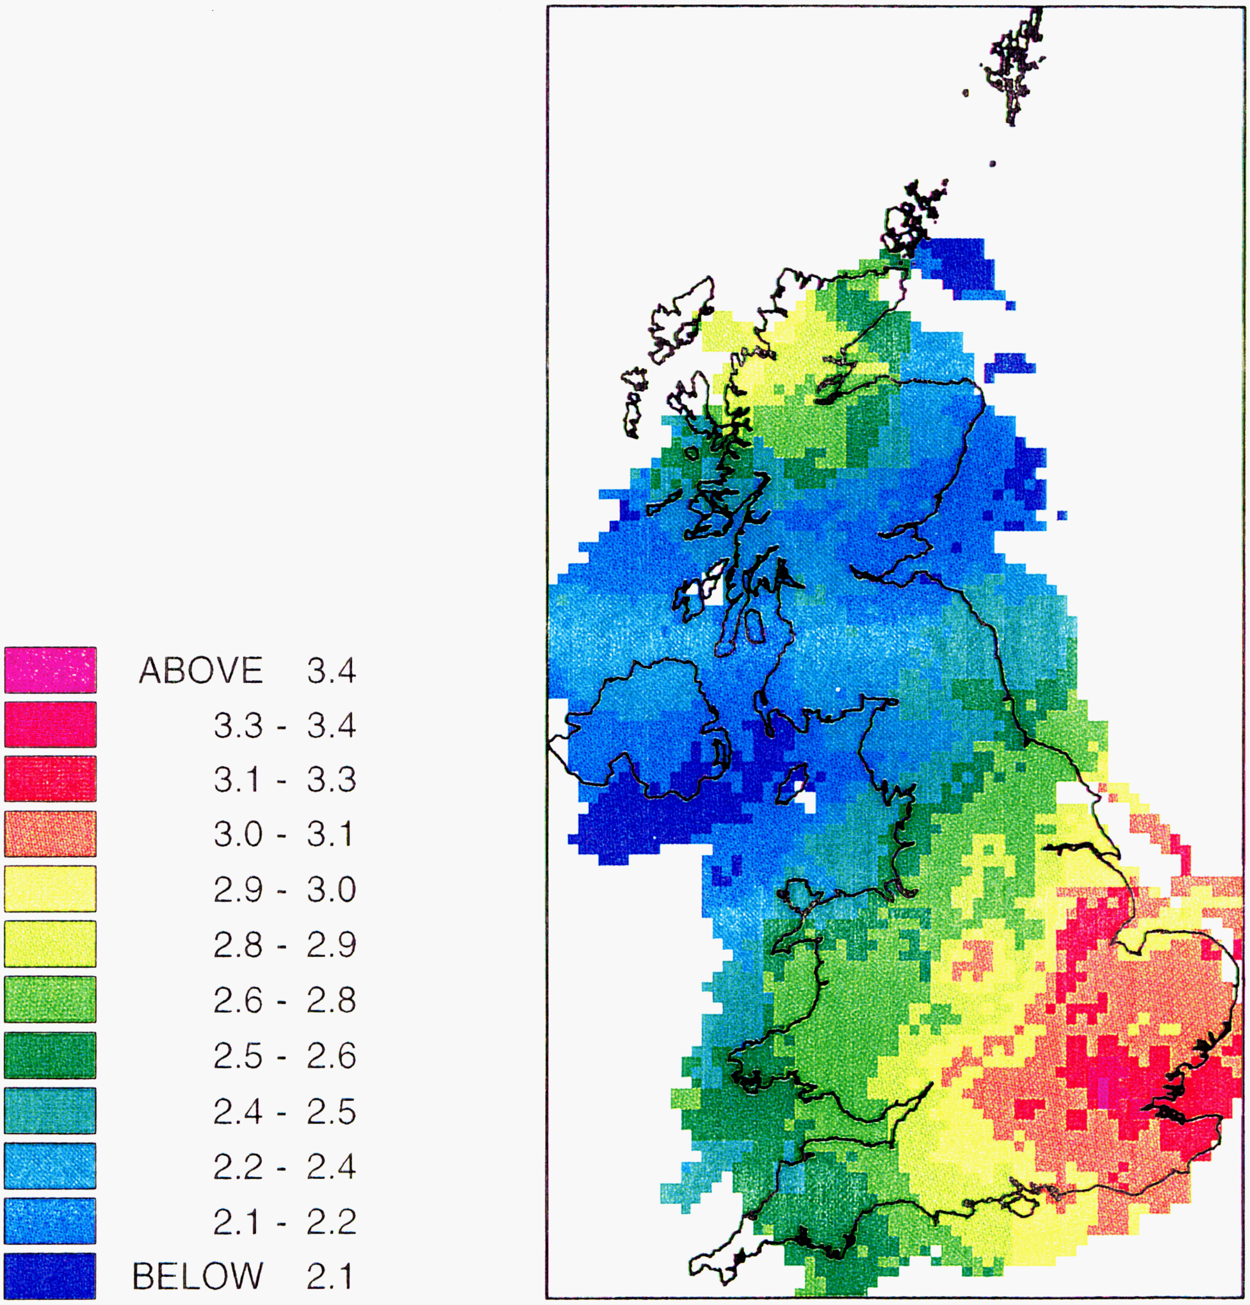

Figure 5.7

1-day rainfall growth rate for T=20 years

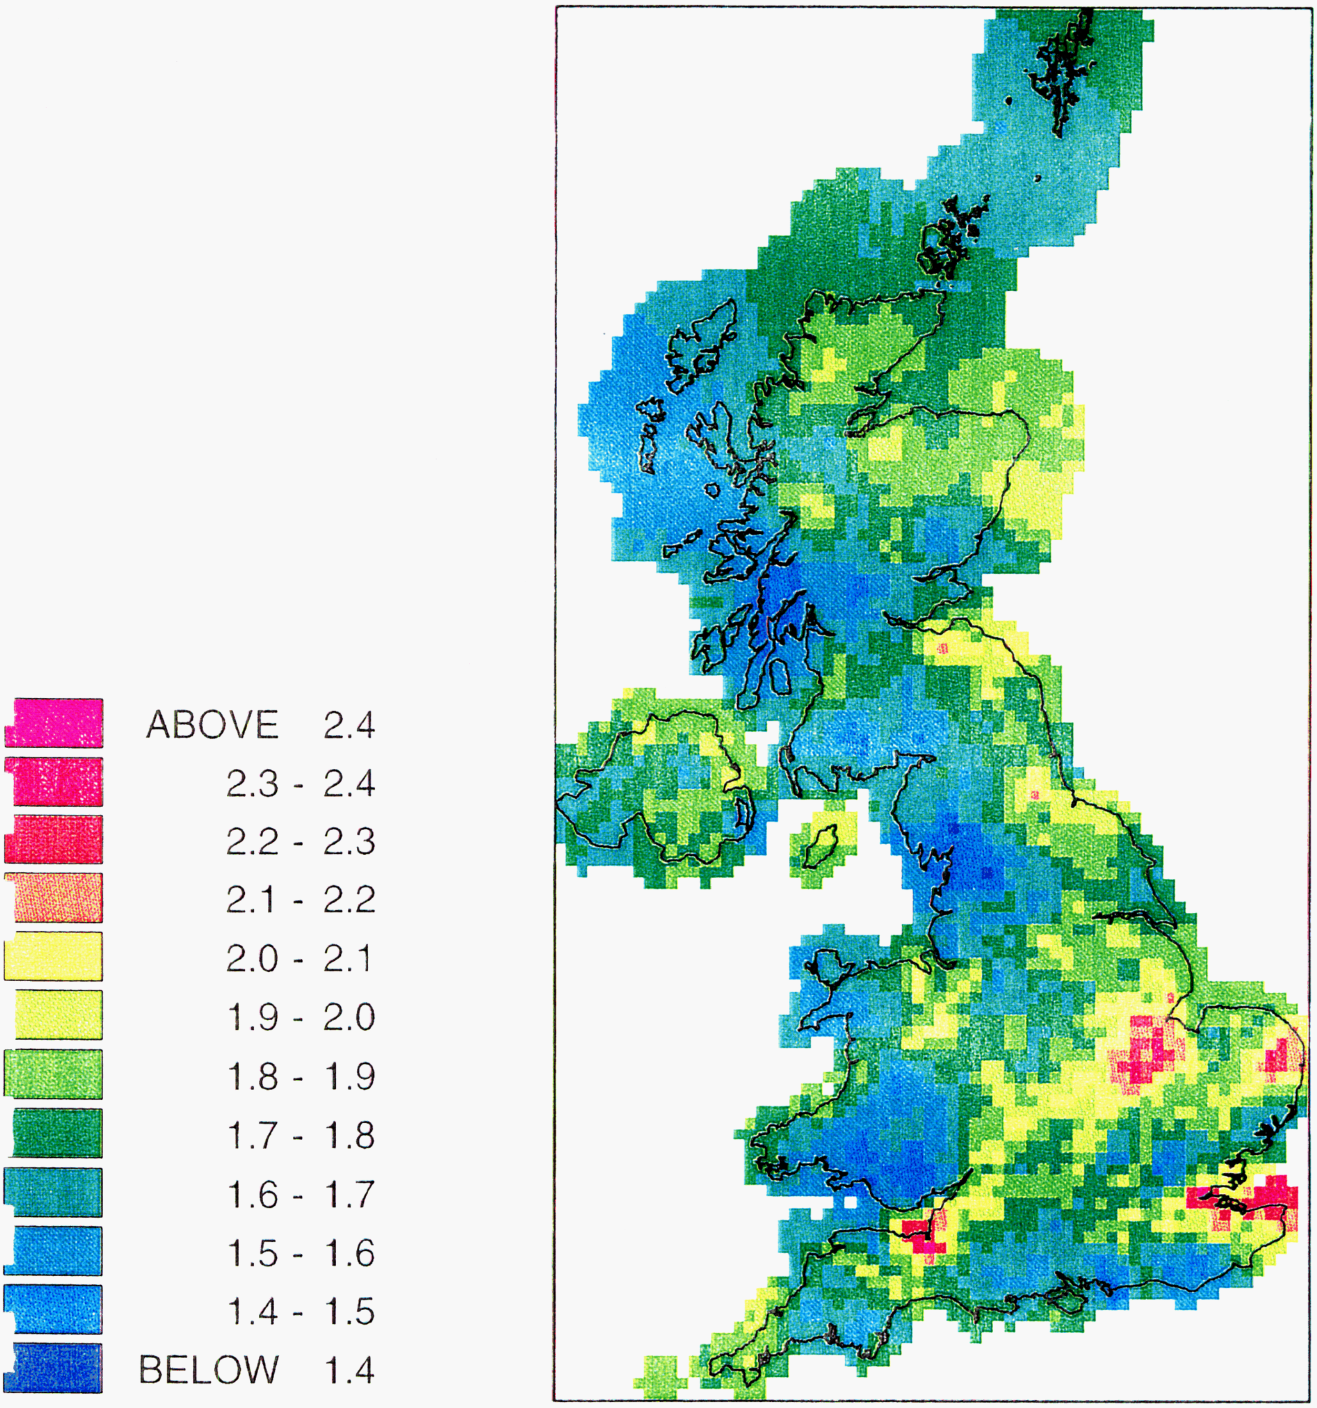

Figure 5.8

1-day rainfall growth rate for T=100 years

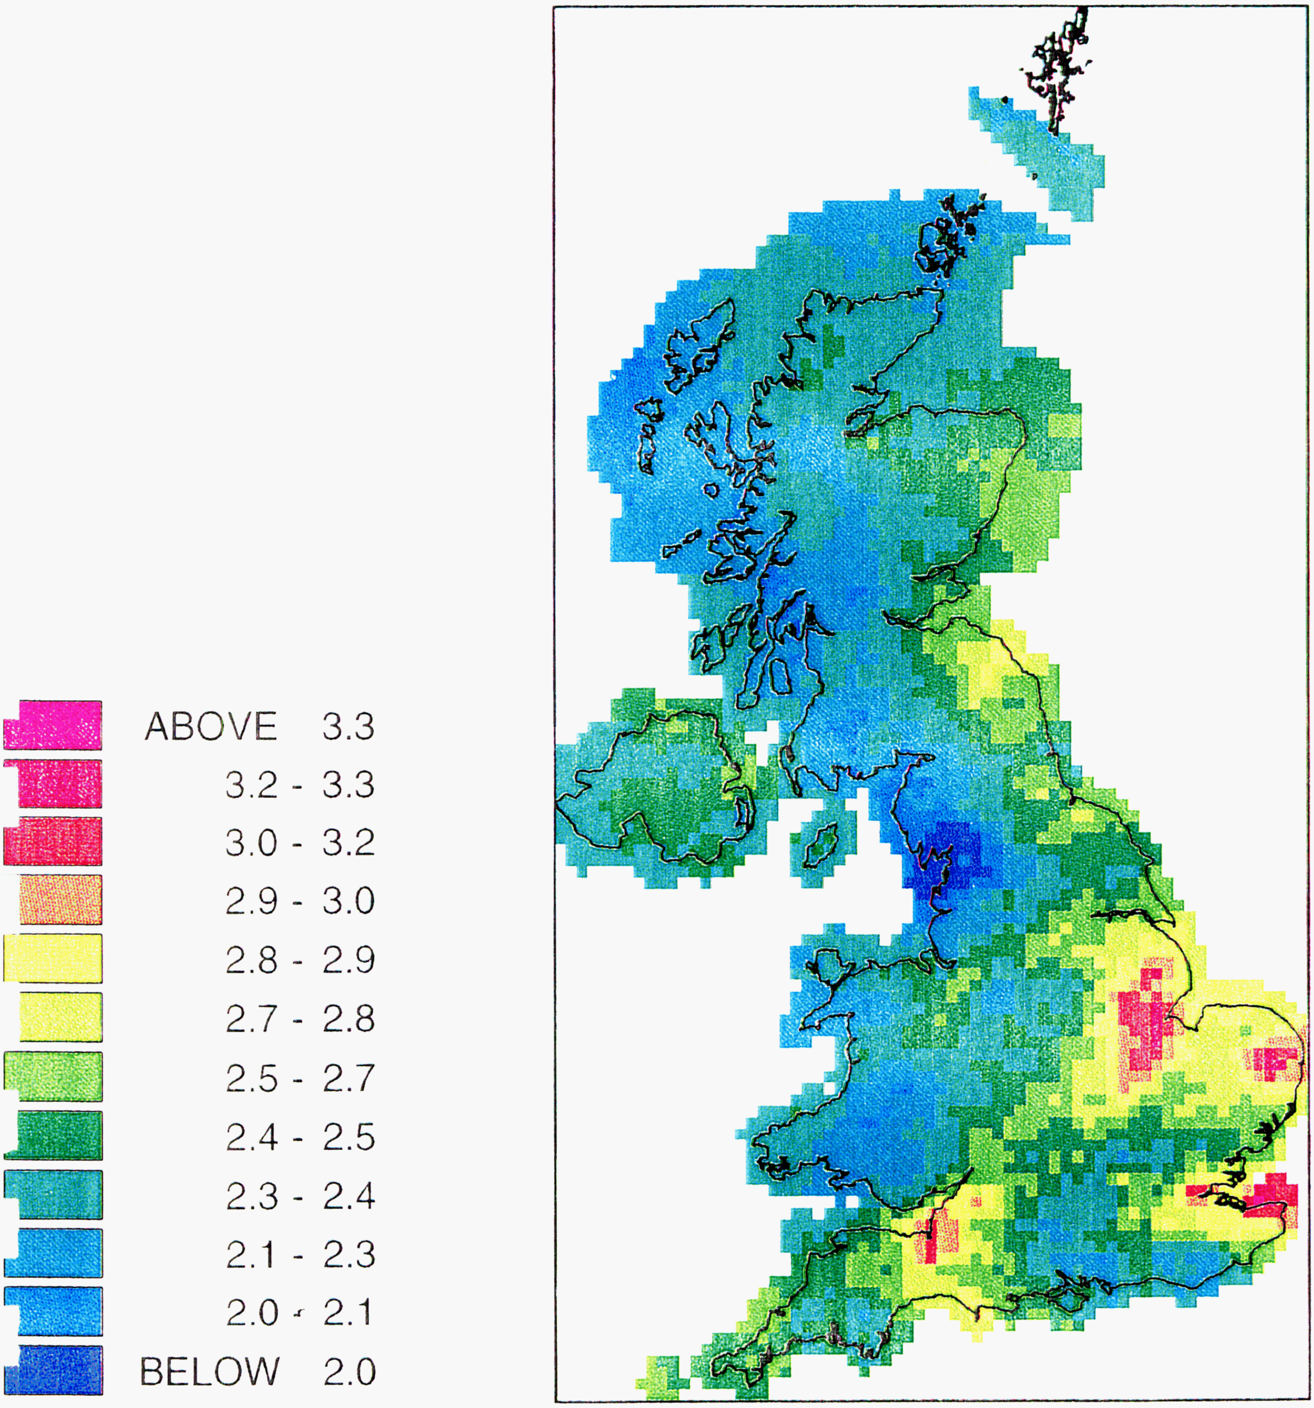

Figure 5.9

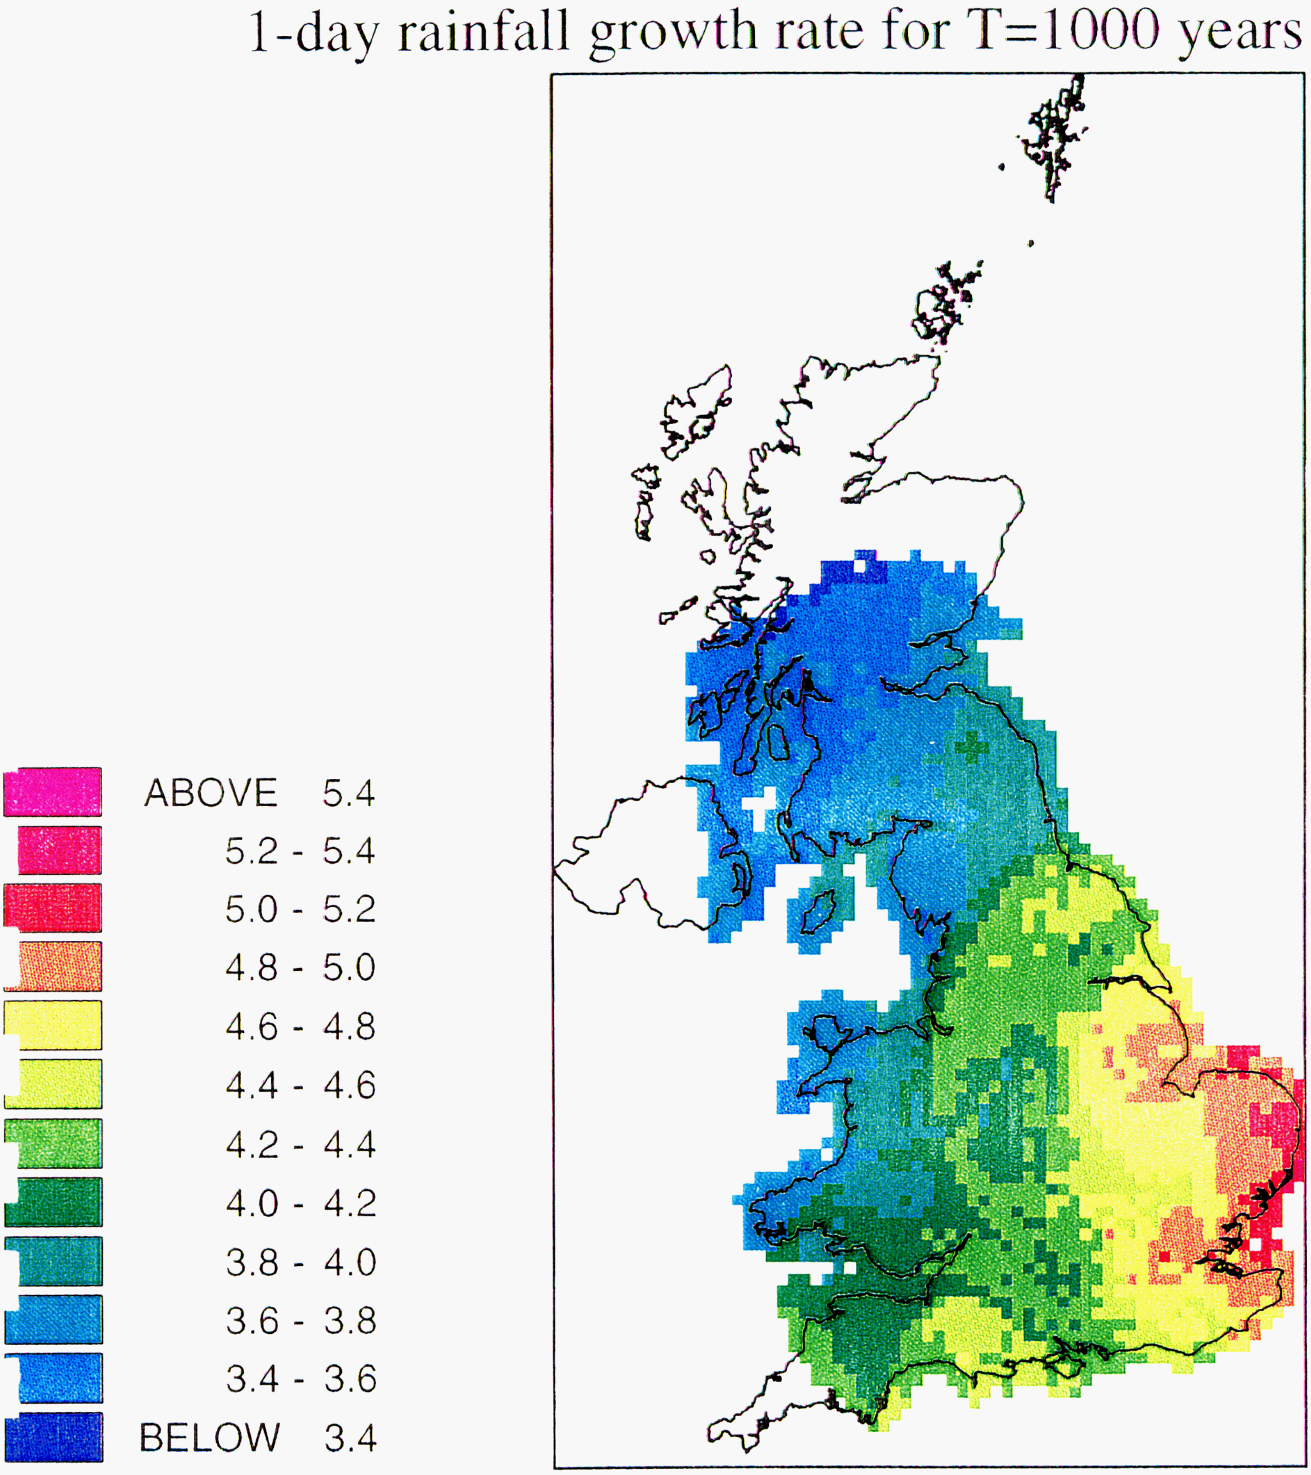
1-day rainfall growth rate for T=1000 years
ABOVE 5.4
5.2 - 5.4
5.0 - 5.2
4.8 - 5.0
4.6 - 4.8
4.4 - 4.6
4.2 - 4.4
4.0 - 4.2
3.8 - 4.0
3.6 - 3.8
3.4 - 3.6
BELOW 3.4

Figure 5.10

Mean of top 10 network maxima from largest network

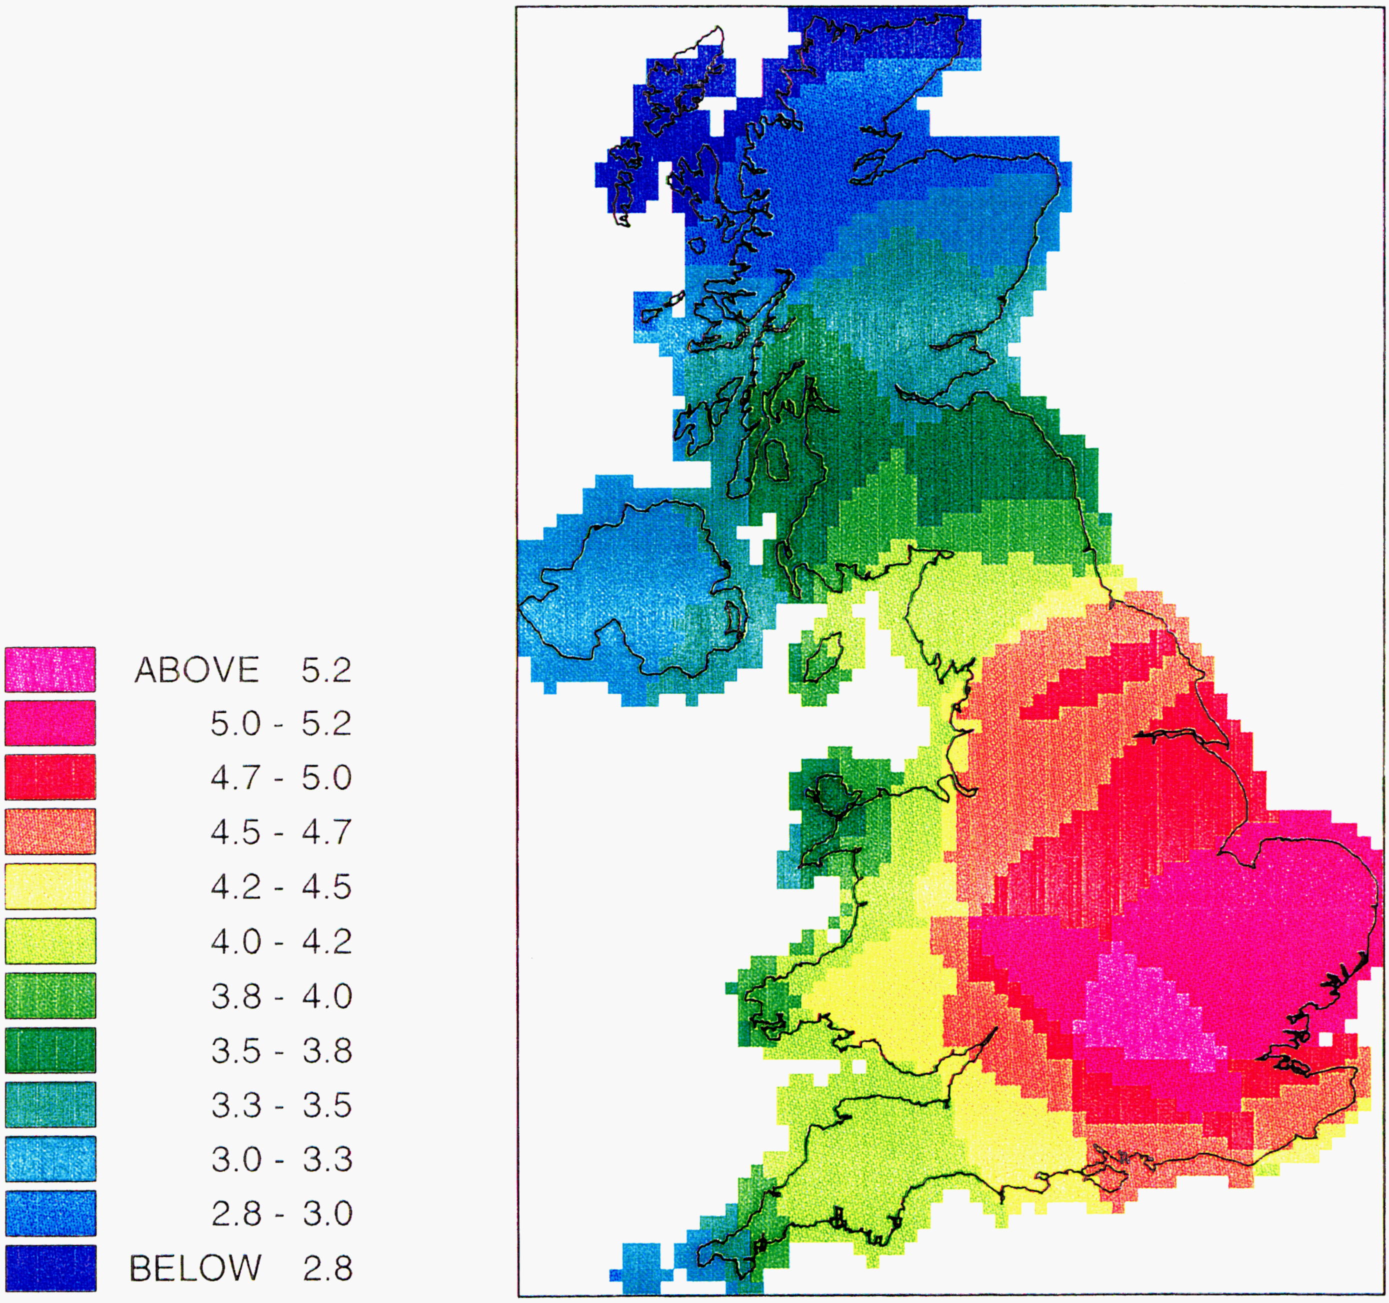

Figure 5.11

Equiv. no. of indep. station-years in top network

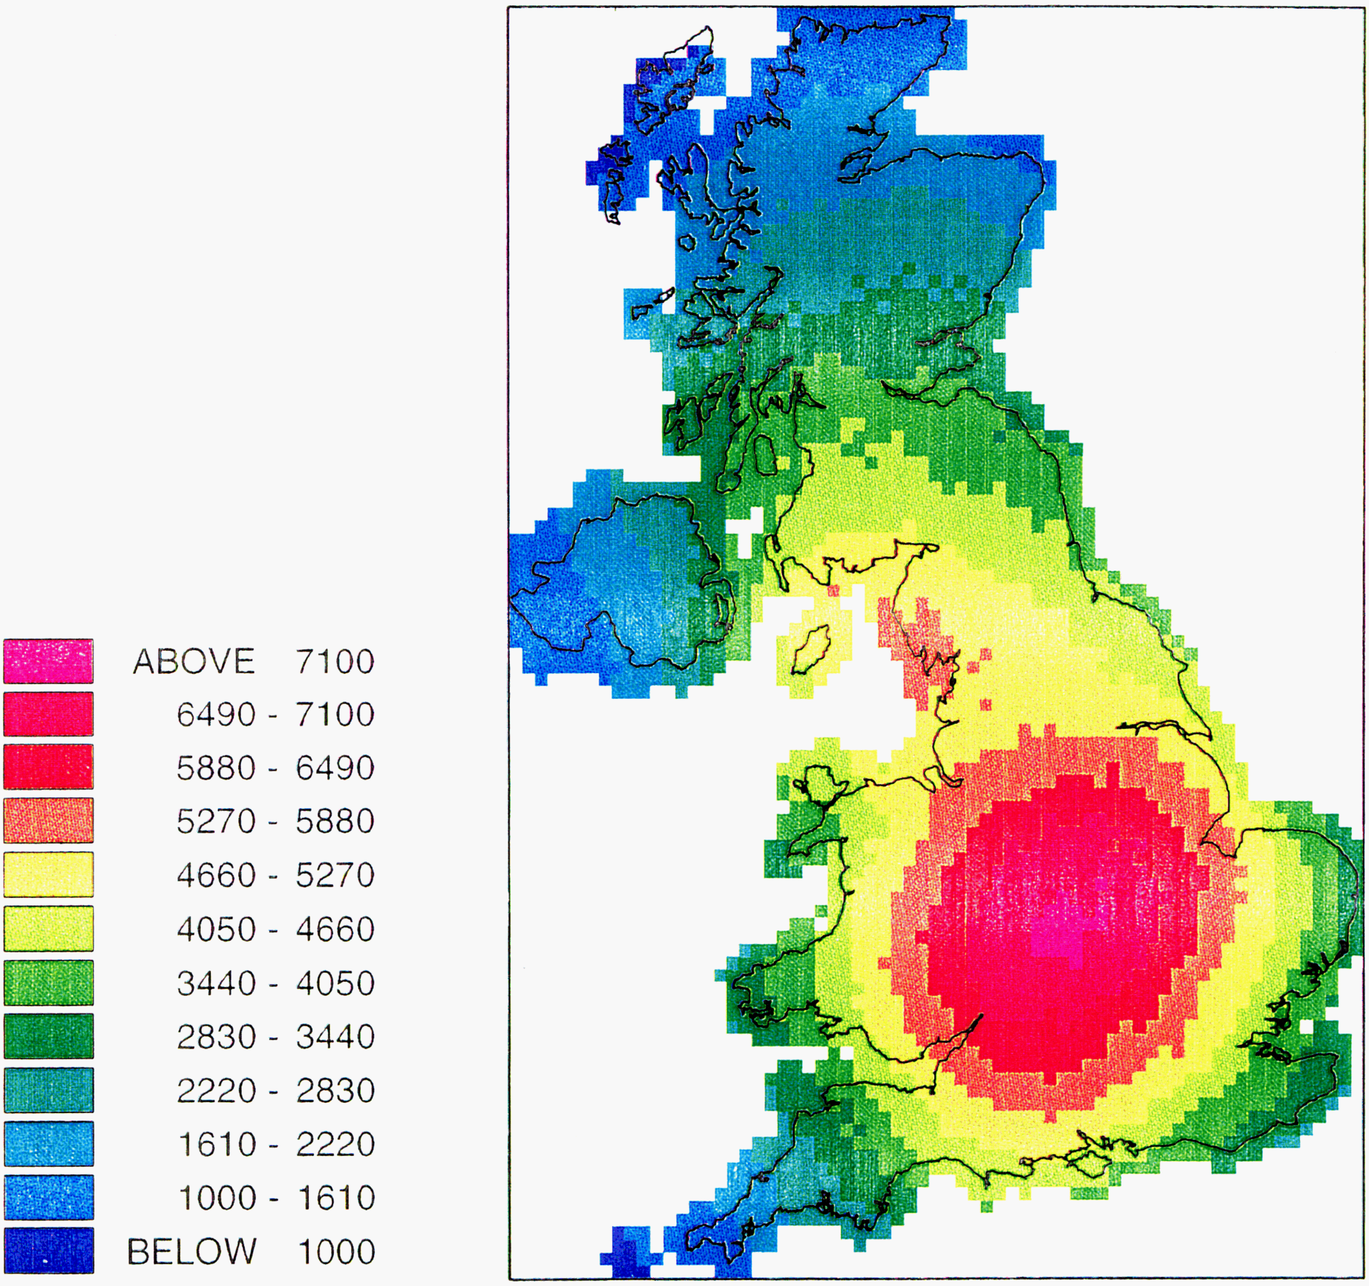

Figure 5.12

8-day rainfall growth rate for T=100 years

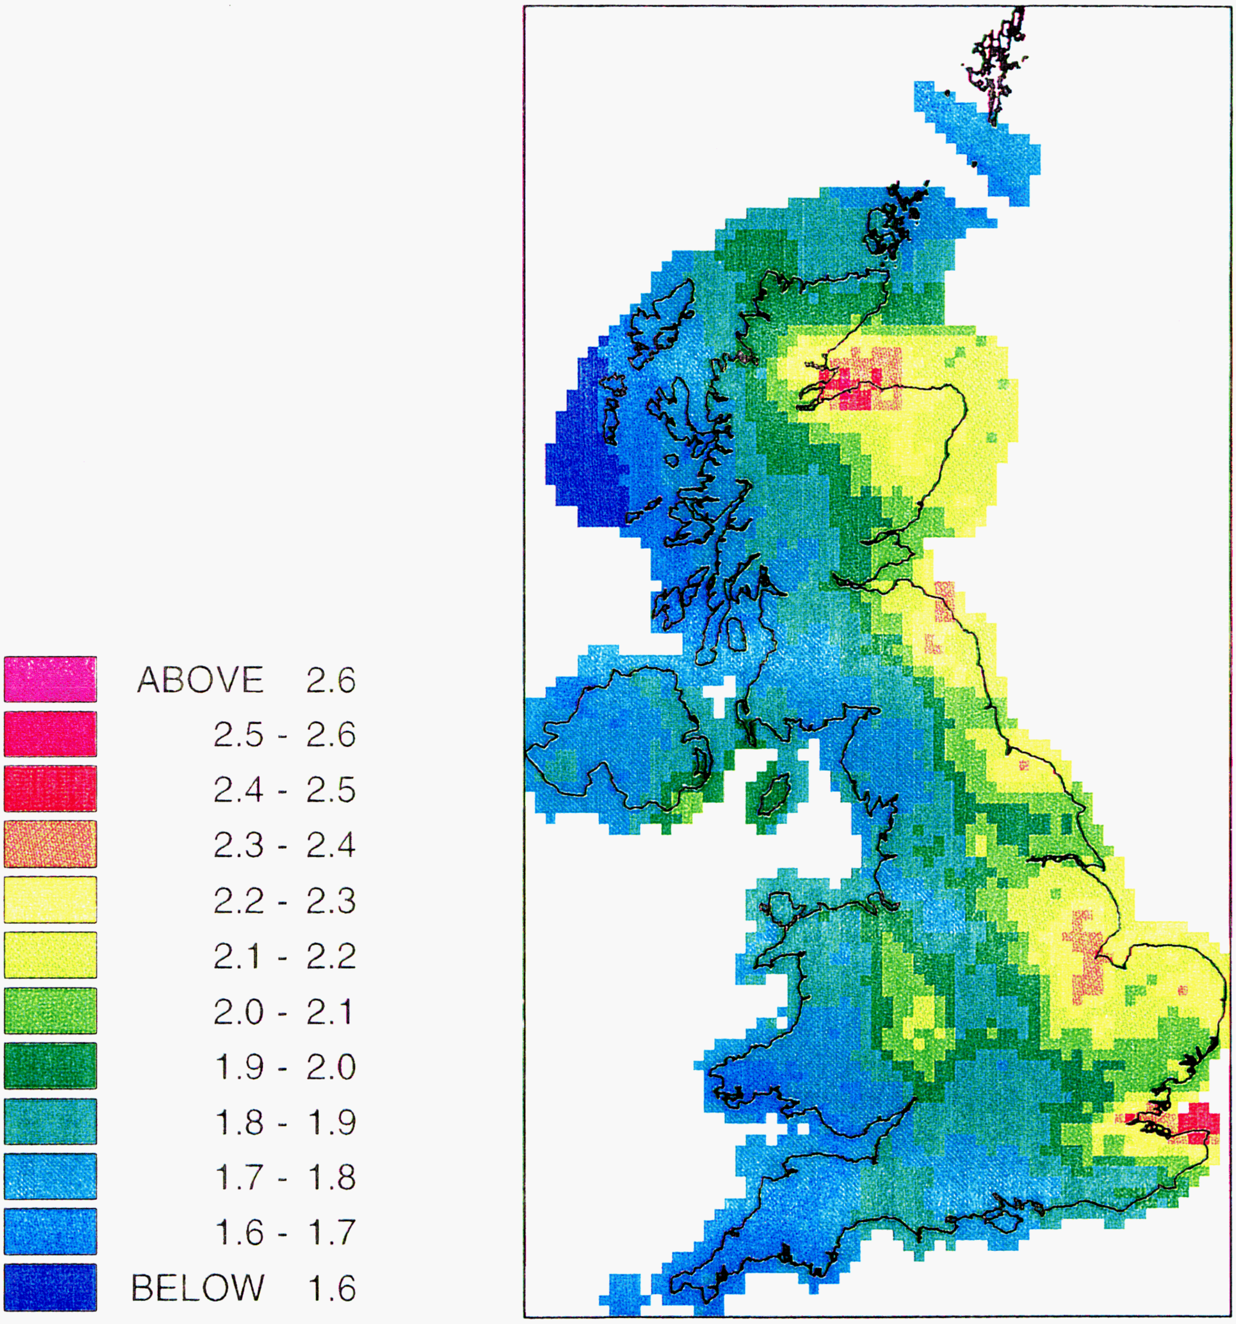

Figure 5.13

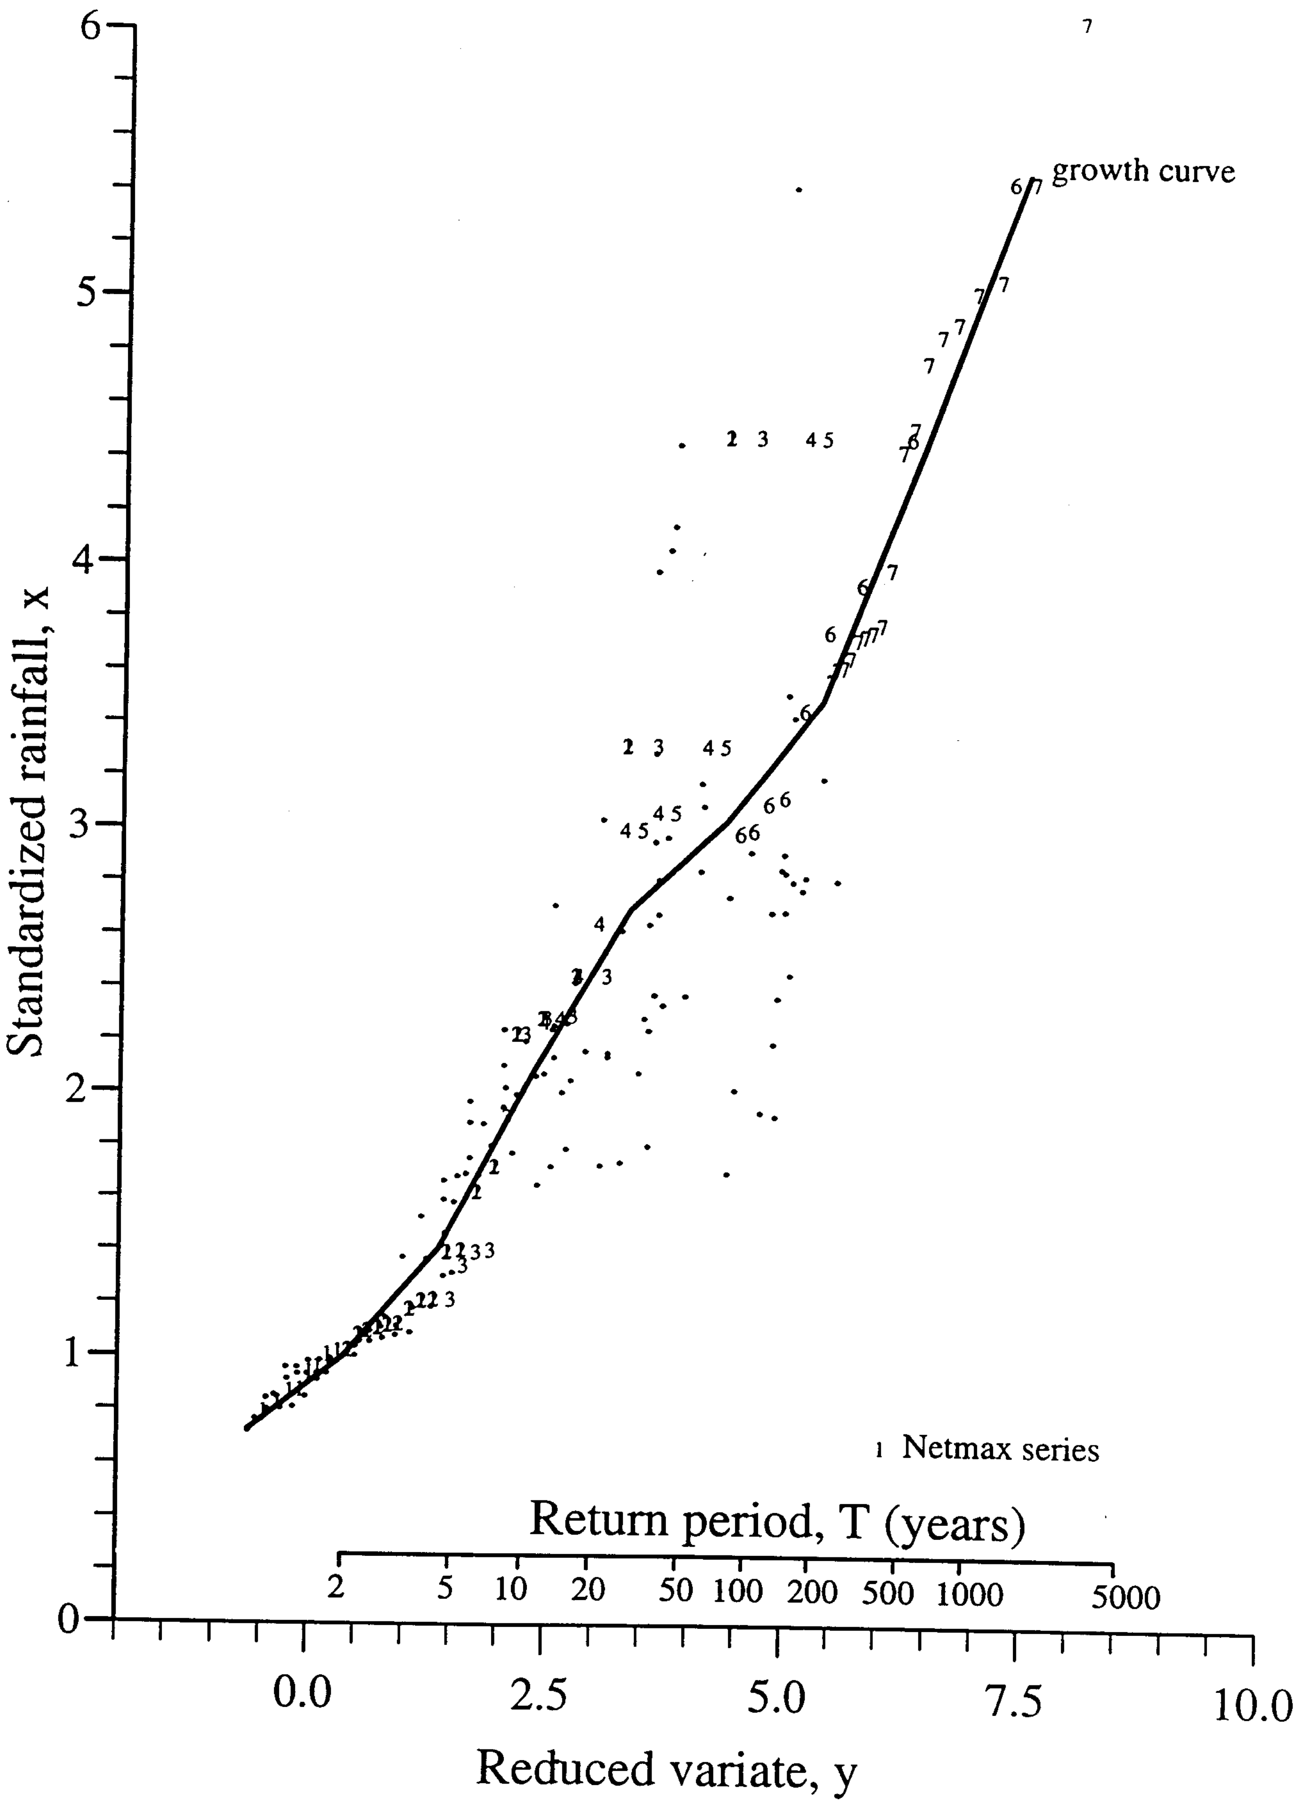

Figure 5.14 1-day growth curve for western end of Thames estuary: high netmax points in networks as local as number 1.

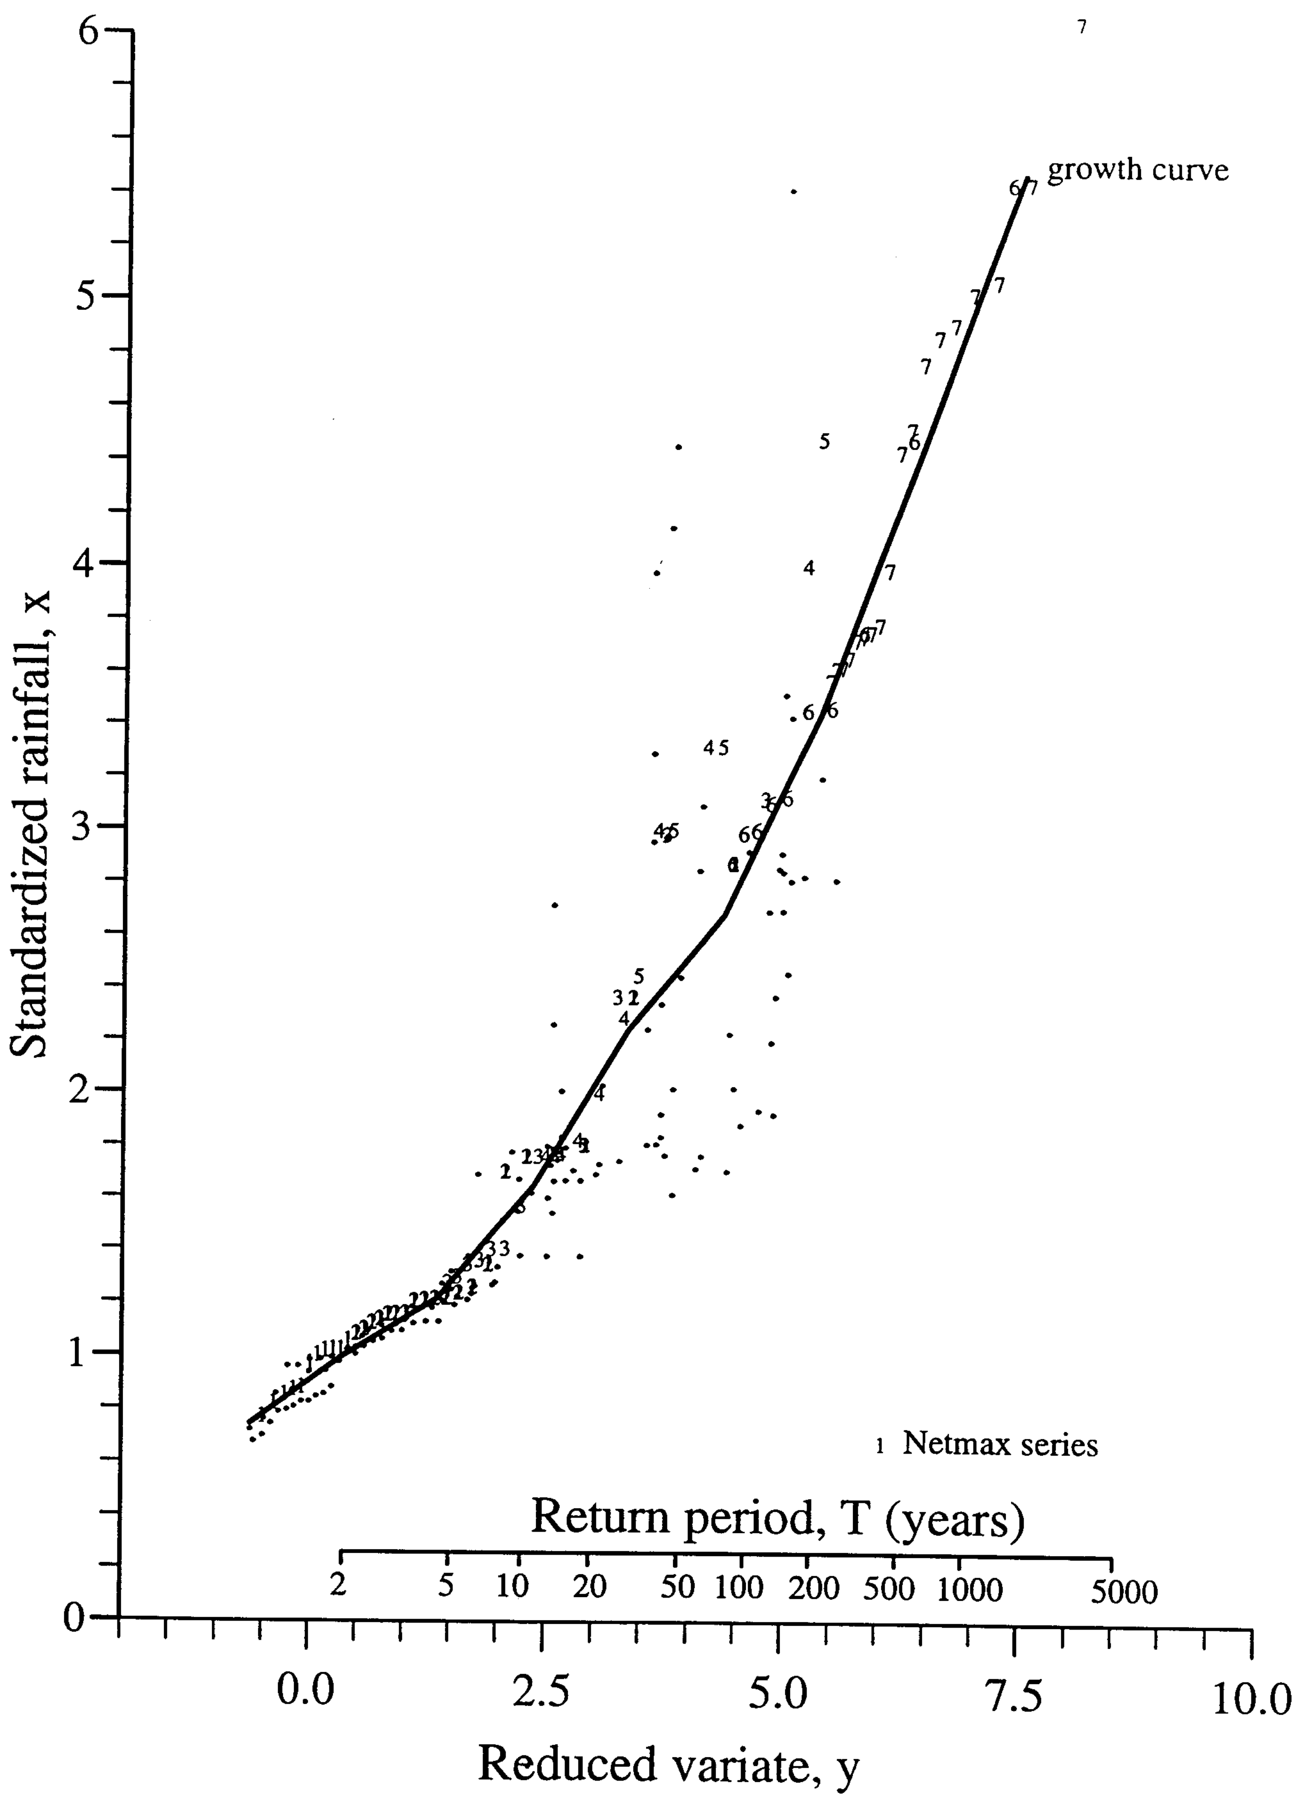

Figure 5.15 1-day growth curve for near Basildon: no high local netmax points

6. CONFIDENCE LIMITS AND TESTS

6.1 Confidence limits on growth curves

6.1.1 Background

There is an increasing demand for measures of confidence to be supplied along with frequency estimates. Confidence limits for rainfall frequency estimates are difficult to quantify because of the number of sources of uncertainty. For example, the FSR volume on rainfall gives no measure of confidence, apart from noting that the station-year method gives estimates at longer return periods "with lesser confidence" (FSR Volume 2, Section 2.3). The terms of reference for this project (Appendix 4) include a requirement to determine confidence limits for the rainfall growth curves.

The confidence intervals described here measure the uncertainty in growth rates due to the limitations of the sample size. They do not attempt to account for sources of error such as gauging inaccuracies. Confidence limits give an indication of the range of values in which we expect the true growth rate to lie. The true growth rate could only be known if we had an infinitely long record of rainfall, in which case we could derive the underlying population of annual maxima (assuming no climate change).

Bootstrap methods have been developed relatively recently (Efron, 1977). They enable the derivation of confidence limits and the use of significance tests in situations where the underlying statistical population is unknown. Bootstrapping is based on the generation of many resamples, which are selected from the original sample. This sample is used as the distribution from which the resamples are chosen randomly, with each value being returned to the original sample after it has been chosen, so that it may be chosen again.

In regional rainfall frequency estimation there are two alternative samples to work with: the years for which we have data or the sites at which we have data. We tried both alternatives, but eventually chose the time domain because there is more scope for variation between annual maxima in different years than at different sites, due to spatial dependence. It is easy to imagine an immense rainfall occurring in a year which is just before records start, or just after they end. It is less likely that an isolated storm would be totally missed by the dense (daily) raingauge network in the UK.

6.1.2 Method

The elements used for resampling are the individual years of data across all sites. Resampling must take place across the entire network rather than at individual gauges, to preserve the spatial dependence which is an inherent and important feature of extreme rainfall which the FORGEX method exploits. The basic method to find the $100(1-2\alpha)\%$ confidence interval is as follows.

(i) Draw a resample from the M years providing rainfall data

(ii) Use the resulting at-site records to derive a growth curve with the FORGEX method

(iii) Repeat the above steps B times (B=199, see Appendix 3) to give B growth curves.

(iv) For various return periods, find the bootstrap residuals E_i which are the deviations of each new growth rate G_i from the original sample growth rate G_{sam}, i.e. $E_i = G_i - G_{sam}$.

(v) Rank the residuals in ascending order and find E_m and E_n where $m = \alpha B$ and $n = (1-\alpha)B+1$. To obtain 95% confidence limits, the 5th and 195th values are used out of an ordered sample of 199 (see Appendix 3).

(vi) The confidence interval for the unknown growth rate is $(G_{sam}-E_n$, $G_{sam}-E_m)$.

The assumptions and theory behind this method are treated more fully in Appendix 2.

In practice, the more efficient method of balanced resampling was used. The principle of this method is to ensure that each year occurs equally often overall among the B bootstrap samples. This is implemented by creating a series of length BM consisting of the M years of record repeated B times. This series is then randomly re-ordered, and divided into slices of length M, to obtain B bootstrap samples.

6.1.3 Tests and modifications

Most published applications of bootstrap methods involve a relatively simple statistic such as the mean of a sample. The use of resampled years of record to find confidence limits for growth rates has more layers of complexity, so the method was extensively tested and several modifications were introduced.

The tests included an investigation of the influence of the mean record length in a resample. These mean record lengths differed considerably between resamples, although it was not seen to cause any systematic variation in growth curves. One reason for the variation in mean record length was the small proportion of daily raingauges supplying records from before 1961. These few records are much longer than the others, and so there are many years for which relatively few gauges provide data. This feature of the dataset was preserved by resampling the periods before and after 1961 separately.

Another test examined the minimum record length in the resamples. To avoid estimating the standardizing variable (RMED, the median annual maximum) from records shorter than ten years, it was decided to standardize each record before resampling. This is not an ideal solution, because the median of each resampled series of standardized annual maxima will not necessarily be exactly equal to 1. It could, however, be argued that, if it is acceptable to multiply growth rates by a value of RMED taken from a map, it is also acceptable to standardize records by a pre-determined value of RMED.

6.1.4 Results and discussion

Figure 6.1 shows 199 growth curves for 1-day rainfall focused on Leicester in the East Midlands, together with the original growth curve estimated from all the data. The resulting 95% confidence limits, obtained by ranking the bootstrap residuals as explained in Appendix 2, are shown in Fig. 6.2. The dashed lines connect the results for different return periods. The upper

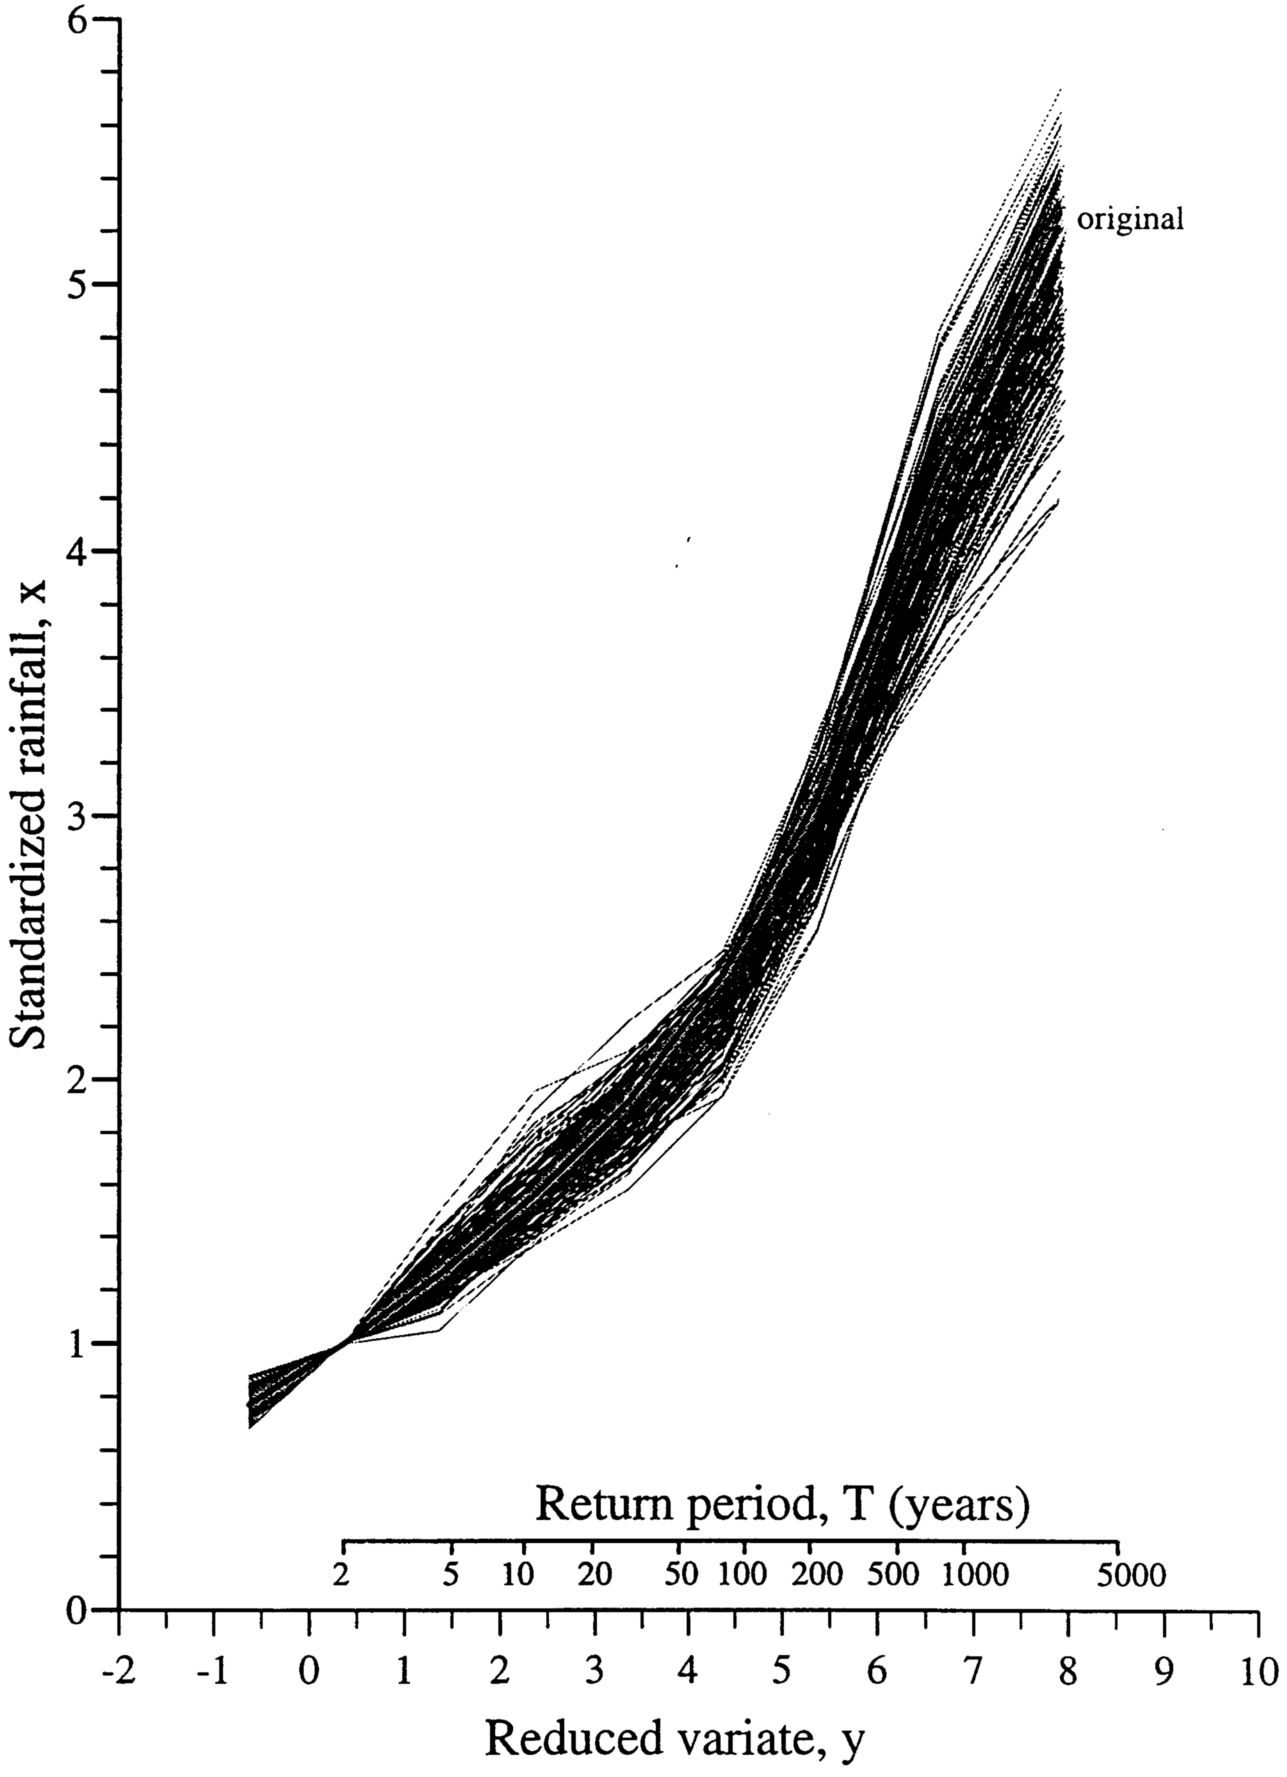

Figure 6.1 1-day rainfall growth curve and 199 resamples focused on Leicester

limit is typically 0.4 to 0.6 above the growth curve, and the lower limit varies from very close to 0.3 below. The asymmetry is in part due to the upward curvature of the growth curve. This was confirmed by some tests on single-site flood frequency curves with different shapes. It only takes one large event to pull up the growth curve significantly, and so the upper confidence interval is rather large.

Another example, for 1-day rainfall at Kendal in the Lake District, is shown in Figure 6.3. Growth rates at Kendal (for moderate return periods) are among the lowest in the country, and the confidence limits are narrow, only broadening when T exceeds 1000 years. The position of the growth curve within the 95% confidence interval follows the same pattern as for Leicester. The confidence limits are asymmetrical around halfway along the growth curve, where the curvature is strongest.

If Figures 6.2 and 6.3 are superimposed, it can be seen that the confidence intervals for growth curves at Leicester and Kendal do not overlap (for T>10 years). Note that this is not a formal test for a significant difference between the growth curves. However, the technique of bootstrapping is a valuable tool for assessing the effect of a limited period of record on growth curve estimation.

6.2 Gauge density test

The method of resampling from years of data as described above gives an idea of the variation between possible growth curves that might have arisen from the real world, given data from different periods. Another way to test the FORGEX method is to examine its sensitivity to the density of raingauges. A simple test was carried out at a few sites to test whether growth curves can be biased due to gauge density. For example, would estimated growth rates in central England be lower if there were fewer raingauges?

The daily gauges in the UK were divided into three subsets, so that the resulting three networks covered the country with a third of the density of the entire network. Allocations were made by ordering gauges by Met. Office reference number and numbering them 1, 2, 3, 1,... At four focal points, growth curves for 1-day rainfall were estimated from each subset of gauges, as well as from the complete dataset.

The results for one site (Wallingford) are shown in Figure 6.4. The growth curves estimated from subsets of the whole dataset are numbered 1 to 3. They extend to a slightly shorter return period than the growth curve fitted to data from all gauges. The relative positions of the growth curves vary with return period, but it can be seen that growth rates estimated from a sparser network can be similar to, higher than or lower than those estimated from a dense network. Although not shown, the overall picture is similar at the other trial sites: for example, at Bodmin in Cornwall, all three re-estimated growth curves lie slightly above the original one, whereas at Glen Nevis in the Highlands, they all lie slightly below the original growth curve.

At the fourth site (Waddington in Lincolnshire), two of the subsets of gauges missed out a nearby gauge which recorded a very high standardized annual maximum ($x = 6.0$). The growth curves from these subsets (marked 2 and 3 on Figure 6.5) lie significantly below the original estimate and the curve fitted to subset 1 gauges. This illustrates that a single very large event which appears at

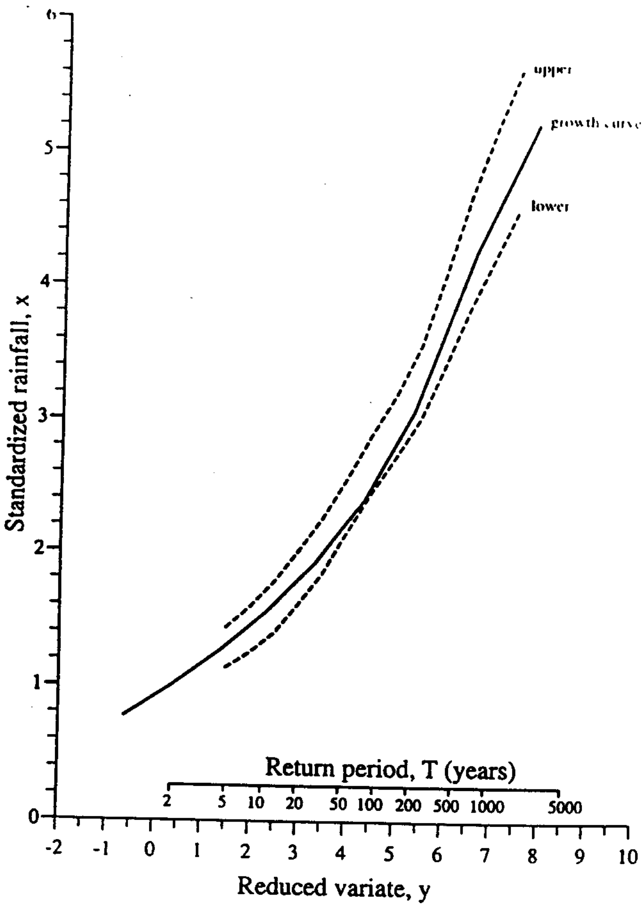

Figure 6.2 Confidence limits for 1-day growth rates at Leicester

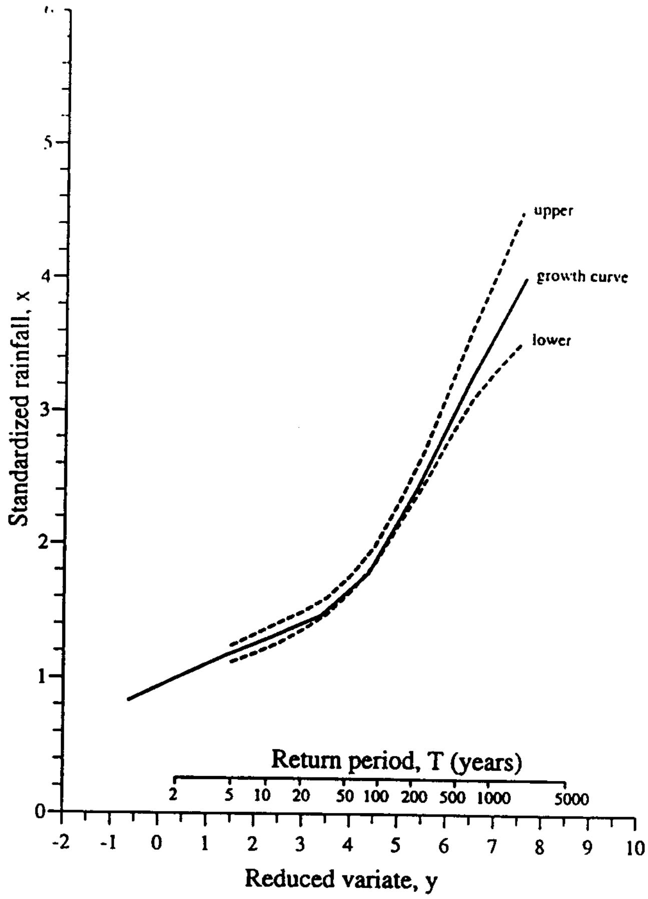

Figure 6.3 Confidence limits for 1-day growth rates at Kendal

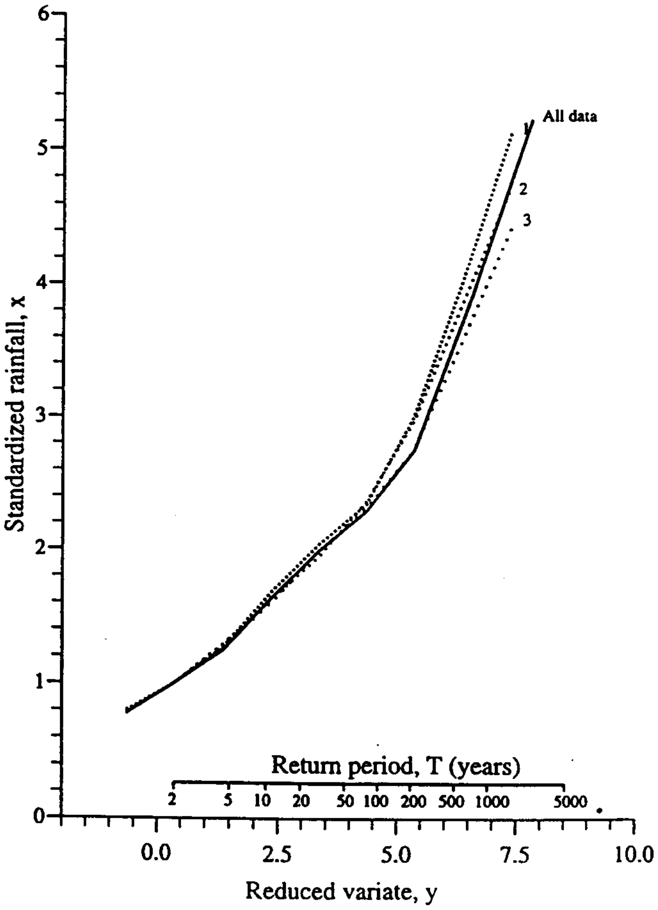

Figure 6.4 Gauge density test at Wallingford

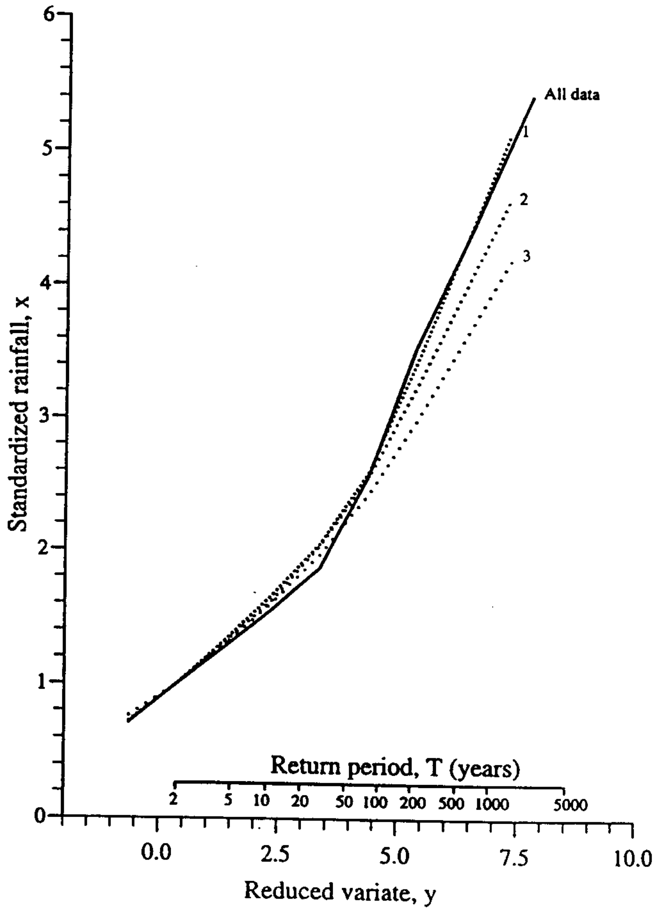

Figure 6.5 Gauge density test at Waddington

the top of the network maximum series for several networks (and also as a pooled point) can have a large influence on growth curves estimated at nearby focal points. However, in general the results do not reveal any bias due to gauge network density.

6.3 Test of sensitivity to large rainfall events

FORGEX bases growth estimates for long return periods on networks of gauges with diameters as large as 200 km. The largest recorded daily rainfall total in the UK was 241.3 mm at Upwey during the Martinstown storm of 18th July 1955, which translates to a standardized value of 6.3. The sensitivity of growth curves to this event was examined by deriving growth curves for pairs of adjacent sites just within and just outside the range of Upwey. For example, Figure 6.6 shows two growth curves for sites in Sussex 200 m apart. One site is inside the 200 km range of Upwey and the other is outside. The growth curve for the inside point is slightly higher, by 2%. The largest recorded rainfall is thus seen to have a discernable but small influence on growth rates 200 km away. This result is pleasing, avoiding the extremes which are, on the one hand, no sensitivity to large events and, on the other, large discontinuities in growth rates.

6.4 Tests of effect of including annual maxima from incomplete years

6.4.1 Background

When annual maximum rainfalls were abstracted from continuous data, as described in Section 2.2, years were regarded as providing acceptable annual maxima if at least 75% of the data were present. In other words, gaps of up to 3 months were allowed. Annual maxima from years with fewer data were not accepted. The motive in setting this condition was to include as many annual maxima as possible (because many station-years have at least some missing data), without compromising the analysis.

A question was raised after the publication of the Phase 1b report (Richard Chandler, personal communication) about the effect of including annual maxima from incomplete years. It was thought possible that if there were many station-years with substantial amounts of missing data, rainfall extremes could be underestimated. The seasonal distribution of gaps in data was also of interest, because of its effect on the likelihood of missing annual maxima.

6.4.2 Extent of gaps in daily rainfall data

For the purposes of this investigation, a "gap" is a period of missing data in a year for which at least 75% of data are present. Any missing period longer than a gap results in the year being ignored. Gaps in daily data on the Met. Office archives occur in the form of missing months.

A sample of 50 daily gauges from all over the UK was investigated. The extent of missing months is shown in Figure 6.7. The large majority of station-years do not contain any gaps. Less than 1% of station-years have more than 20% of the data missing. The seasonal distribution of gaps (Figure 6.8) has no strong bias toward any season. The months with most missing data (July and August)

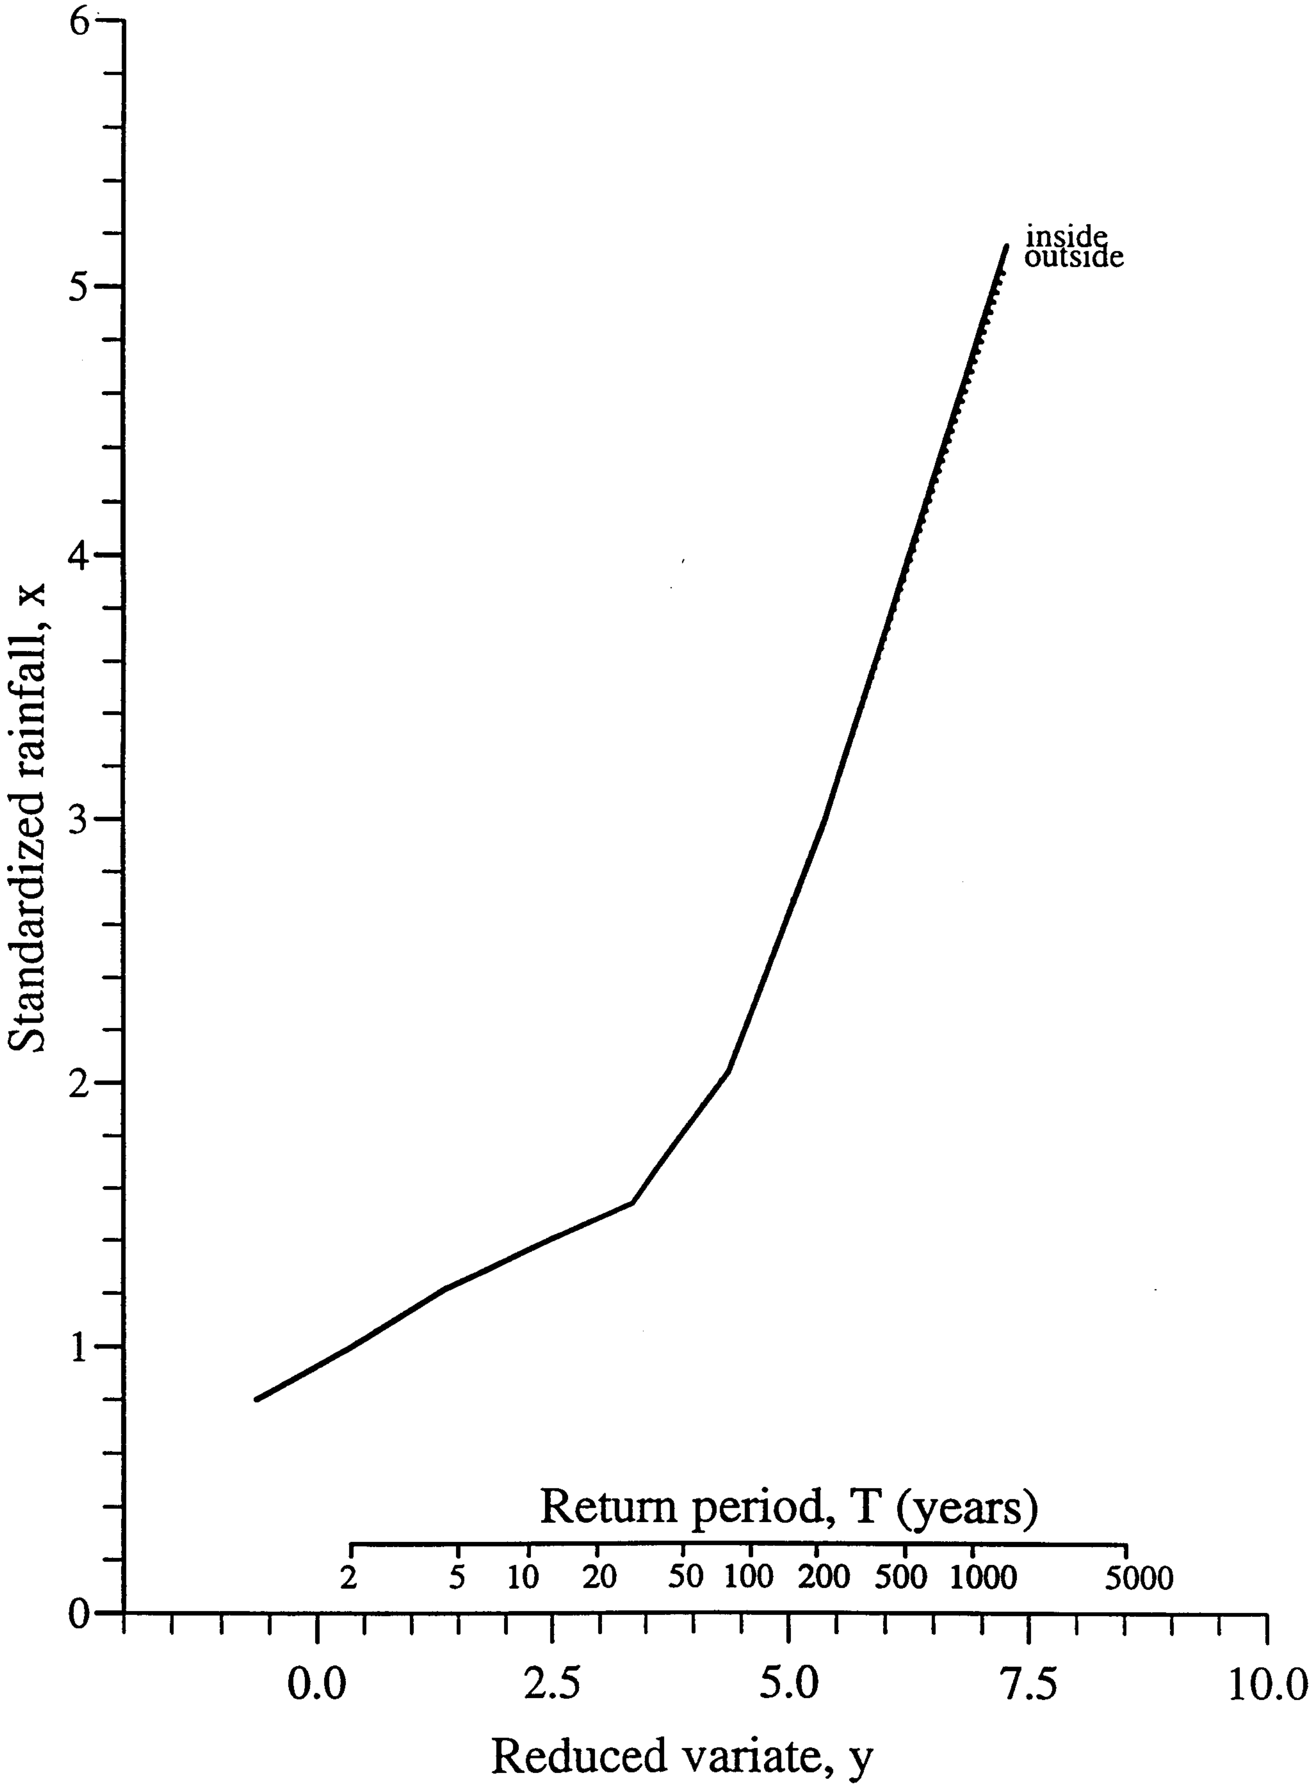

Figure 6.6 **1-day rainfall growth curves for sites 200 m apart, just inside and just outside range of Martinstown storm**

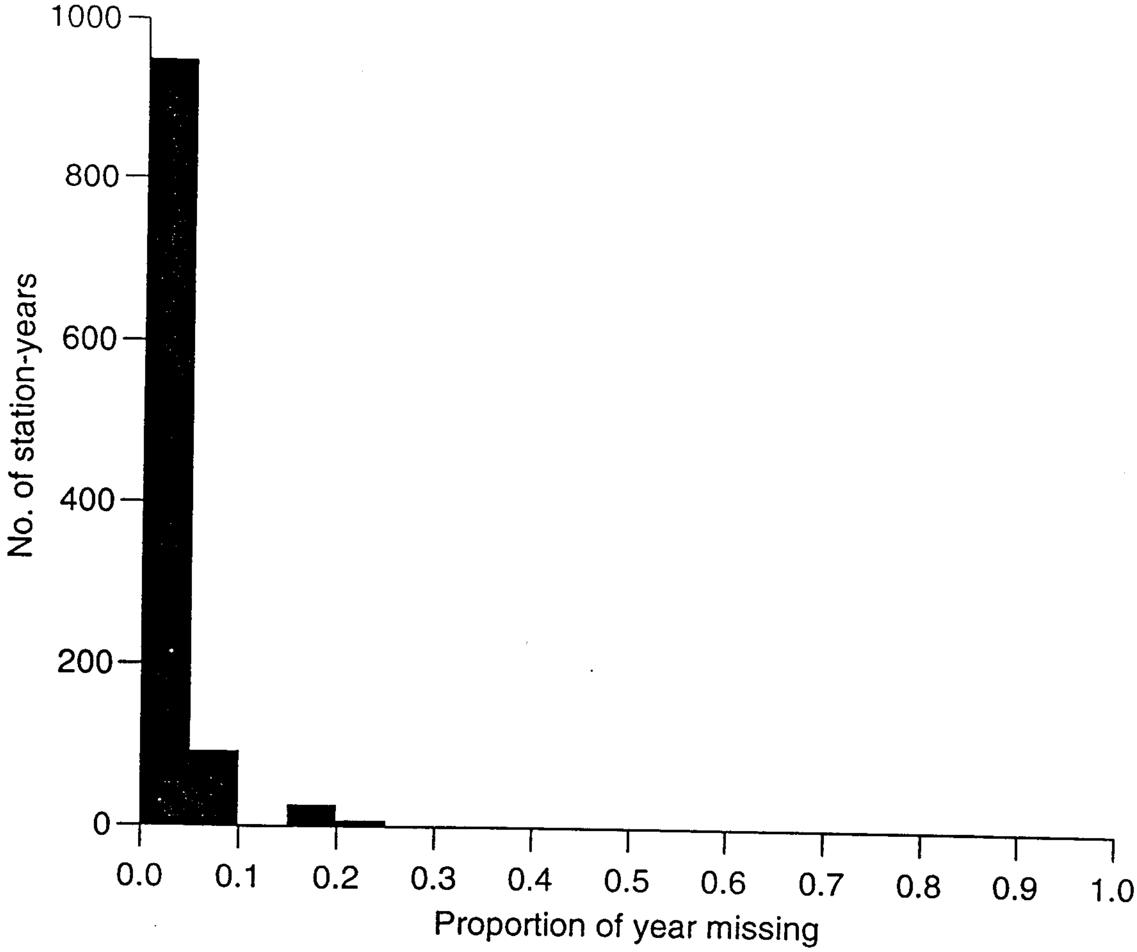

Figure 6.7 Distribution of missing data for sample of 50 daily gauges

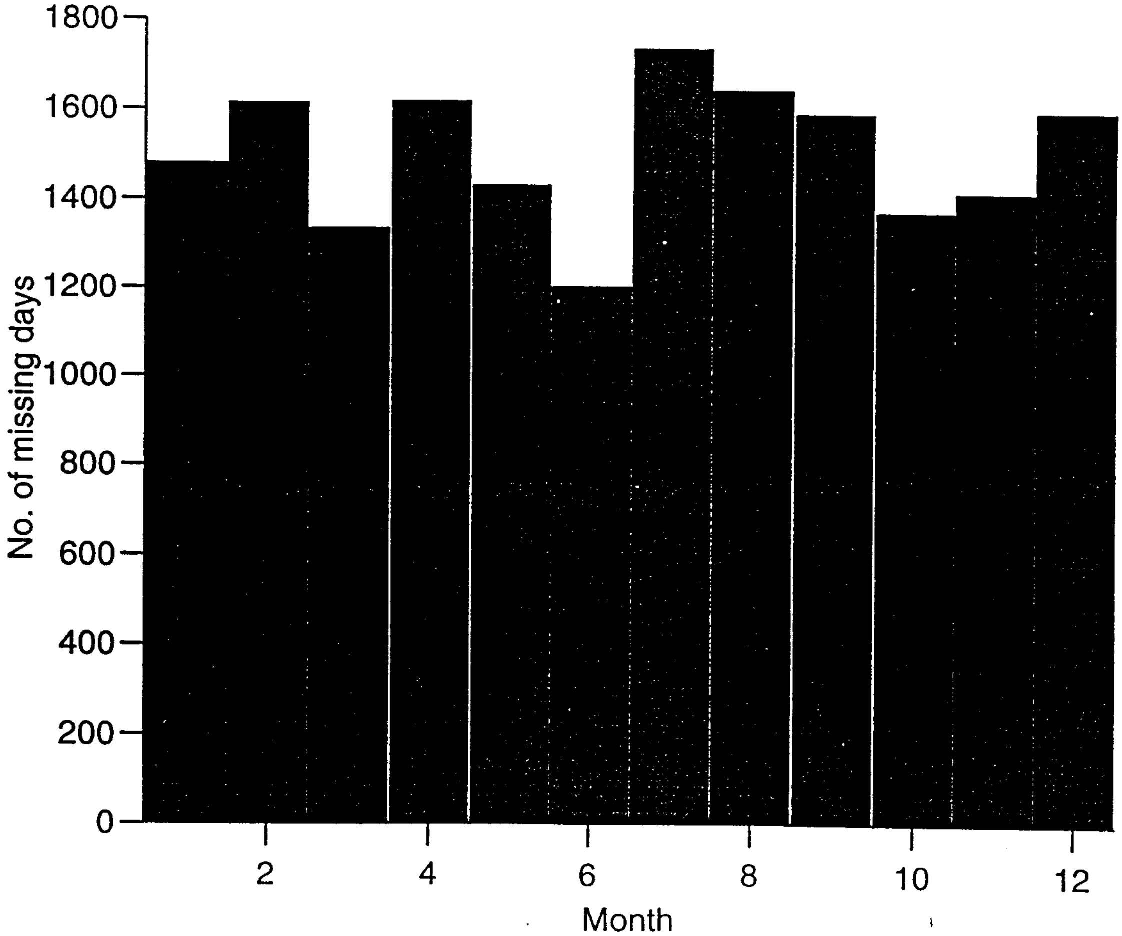

Figure 6.8 Seasonal distribution of missing data in 50 daily gauges

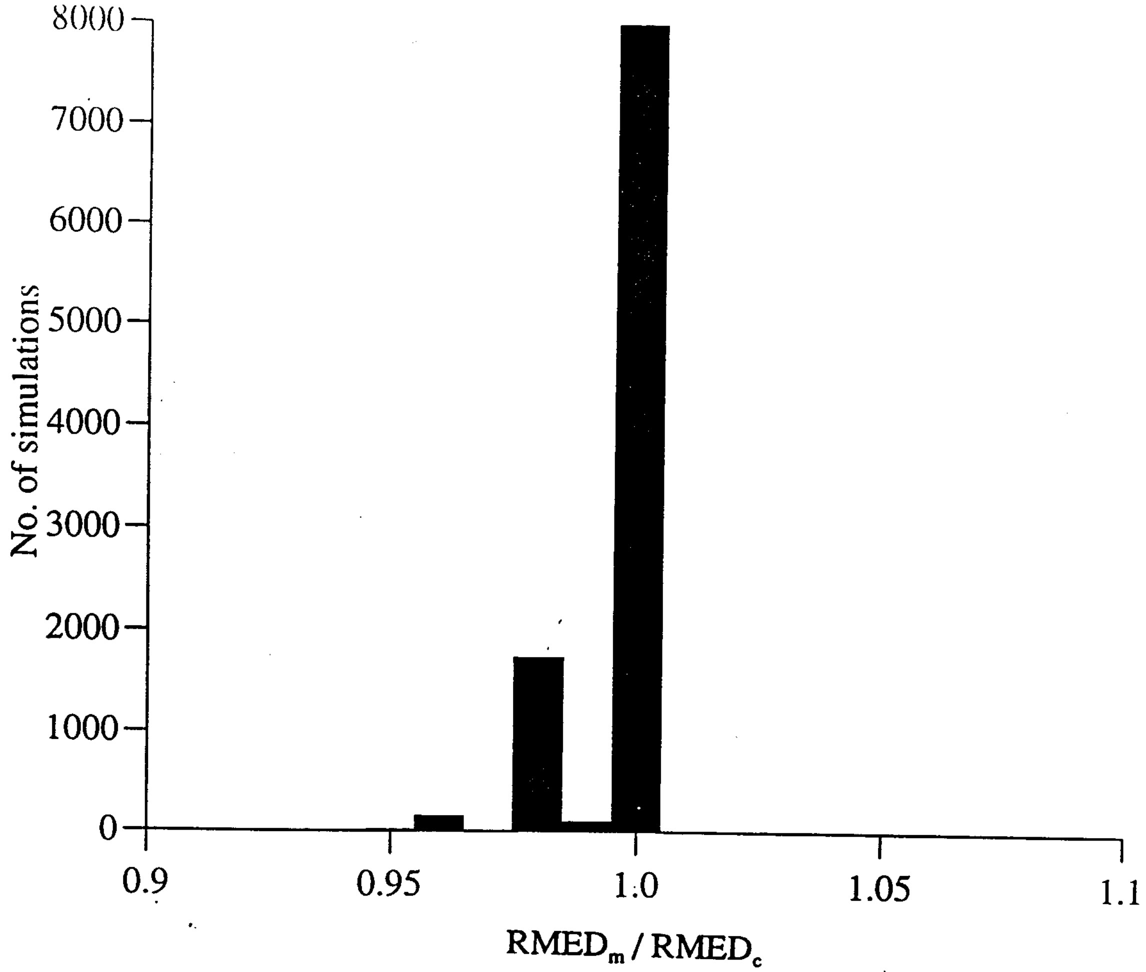

Figure 6.9 Effect of missing data on RMED: distribution of RMED$_m$/RMED$_c$ at 92233

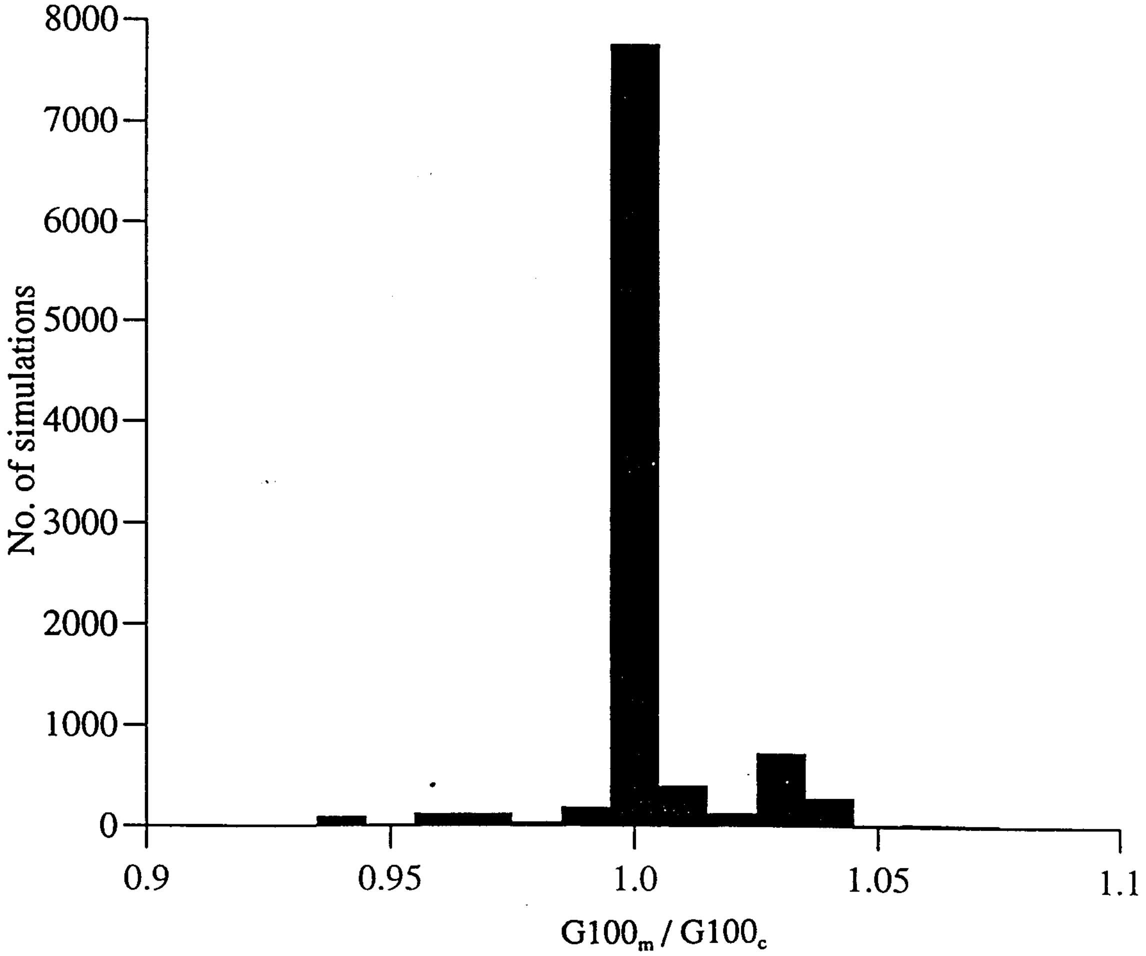

Figure 6.10 Effect of missing data on 100-year growth rate, G100 at gauge 92233

coincide with holidays. There is no indication that the many annual maximum 1-day rainfalls which occur in the winter are more likely to be missed.

6.4.3 Experiment: method

An experiment was devised, with the aim of testing the effect on growth curves and RMED estimates of including some years with missing data. Rainfall frequency estimates from complete records were compared from results from the same records, with some months set to missing.

The starting point was a complete record from a daily gauge, with no gaps. From this, many simulated series of daily rainfall were produced, with different patterns of missing months. The number of missing months in each year was selected randomly from the distribution shown in Figure 6.7. Within each year, the missing months were allocated randomly, bearing in mind the lack of clear seasonal variation evident in Figure 6.8. The procedure for abstracting annual maxima was altered slightly, so that annual maxima from years with three months missing were allowed, even if the three months were longer ones, so that slightly less than 75% of the days of the year were present. If these annual maxima had been ignored, the effect of missing some years would have confused the results.

To keep the experiment simple (avoiding regional analysis), a single-site General Extreme Value (GEV) distribution was fitted to each set of annual maxima. The estimates of RMED and G100 (the 100-year growth rate) were compared with their equivalents taken from the complete record.

10000 simulations were made, and the entire experiment was run twice, using data from gauges 92233 (Penkridge, Stafford) and 5349 (Cockle Park, Northumberland), both of which had complete records spanning 34 years.

6.4.4 Results and discussion

The results are given in terms of the ratio $RMED_m$, found from the series with missing data, to $RMED_c$, from the complete series, and the ratio $G100_m$ to $G100_c$. The distributions of the two ratios resulting from 10000 simulations are shown in Figures 6.9 and 6.10. These distributions are for data from gauge 92233; the other gauge produced similar results.

The large majority of rainfall frequency distributions are unaffected by the presence of missing months, because the annual maxima are not changed. The maximum change in the estimated RMED is a drop of approximately 5%, and the only frequently-observed change in RMED is a decrease by 2%. The changes in RMED on Figure 6.9 are "lumpy" because there is a limited number of values the median is likely to take. The estimated 100-year growth rate has more freedom to vary. It is possible to envisage an increase in G100 if some low-ranking annual maxima are reduced, and higher ones are unchanged, bearing in mind that G100 is an extrapolation beyond the highest-plotting point in a 34-year series. The maximum change in G100 (Figure 6.10) is an increase of 7%, which occurs in a very small number of simulations. The most common change in G100 is a 3% increase.

It is likely that the general decrease in estimated RMED could compensate for the general increase in estimated G100 in some cases. It seems wiser to let this happen than to try to adjust rainfall frequency results to take account of the effect of missing data. This adjustment would be possible only if the presence of gaps in the record had a consistent effect on the results. This is not so, as illustrated by Figure 6.11 which plots the relationship between the proportion of months missing in the 34-year record (taken from gauge 5349) and the change in G100. The diagram indicates that it is possible to have a relatively large proportion of months missing and still estimate the growth curve accurately. Conversely, it is possible to see a large change in the estimated growth curve due to a small amount of missing data.

Overall, the approach taken to handling years with incomplete data does not appear to be causing a problem.

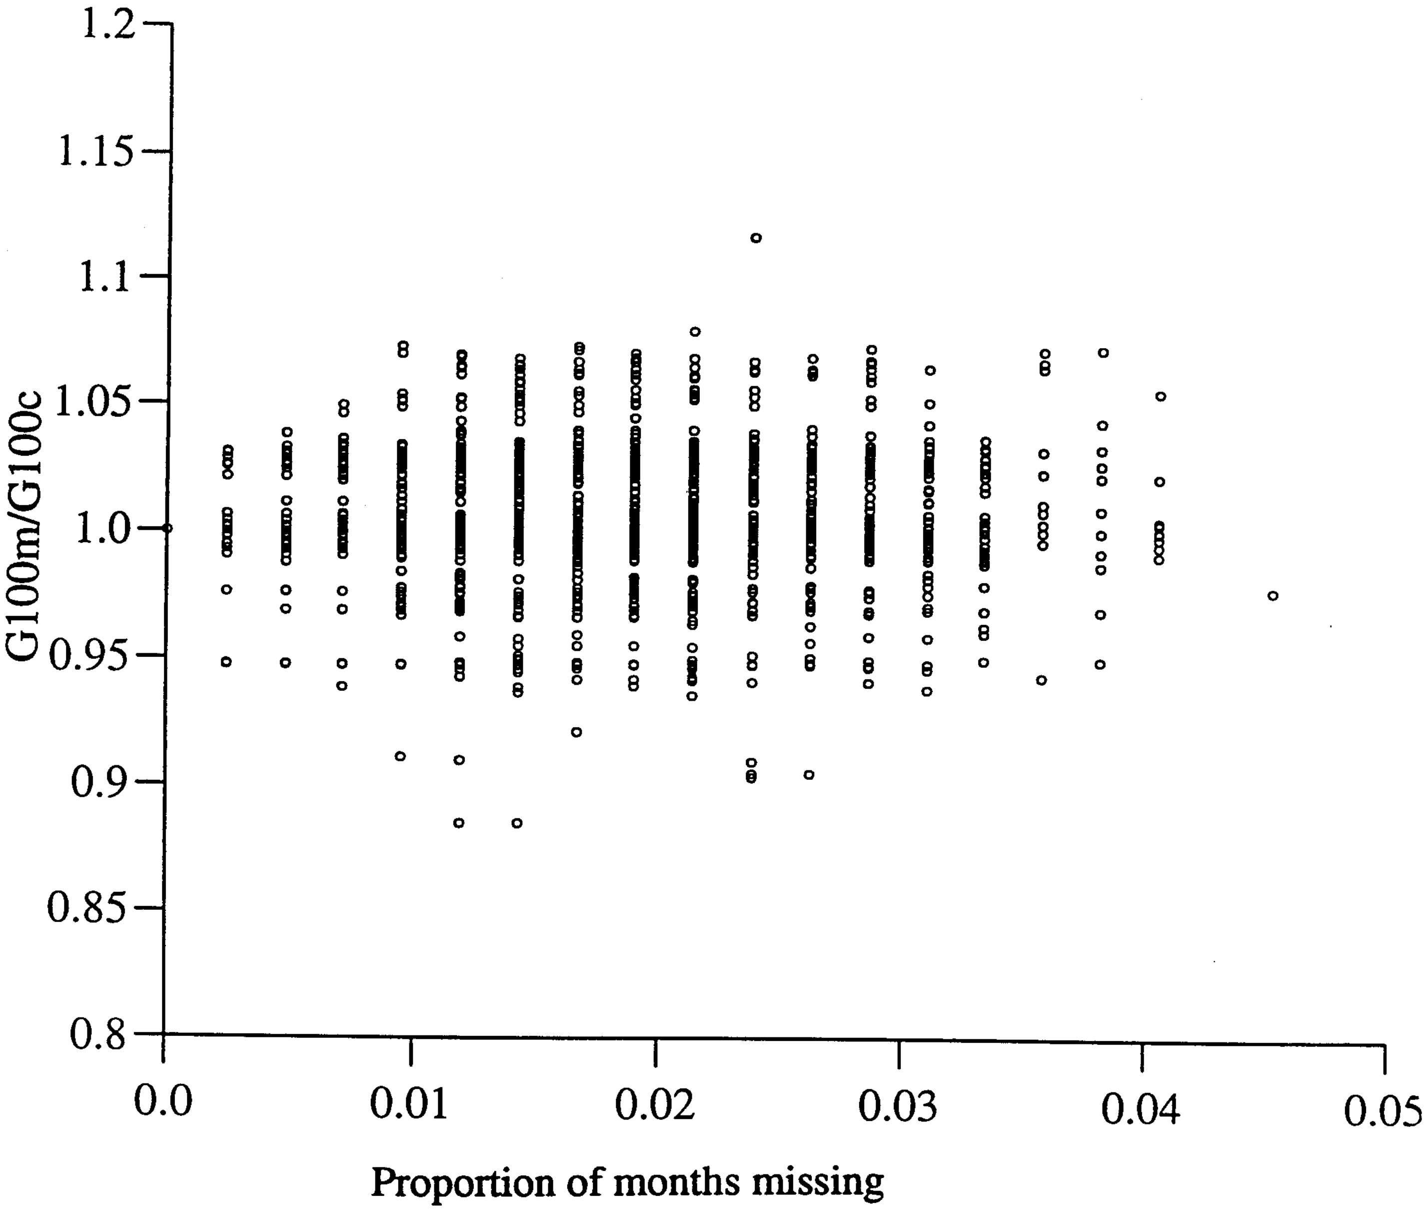

Figure 6.11 Relationship between amount of missing data and G100 at gauge 5349

7. DEVELOPMENT OF MAPPING TECHNIQUES

7.1 Introduction and background

The objective of the work described in this chapter is to produce maps of an index of extreme rainfall. Rainfall frequency estimation is based on the analysis of annual maximum rainfall totals of various durations (from 1 hour to 8 days). The index to be mapped across the UK is the median of annual maximum rainfalls at a site, RMED. This index is used together with the FORGEX rainfall growth method to provide design rainfall estimates at any location in the UK.

The report on Phase 1b (Stewart *et al.*, 1995) describes the mapping of 1-day RMED in a study region (the Severn basin) using the geostatistical interpolation technique of kriging. In addition, altitude data were incorporated via cokriging and the rather simpler method of georegression. It was concluded that kriging was adequate for mapping 1-day RMED in England and Wales, at least if the only alternative was to use site altitude as a covariate. The correlation between 1-day RMED and altitude was found to be poor.

In extending the mapping to all of the UK, the analysis has had to cope with more varied topography and poorer coverage by raingauges, for example in North Wales and the Scottish Highlands. A feature of mountainous regions is that raingauges tend to be sited close to settlements in the valleys. Furthermore, Phase 2 has extended the mapping to 1-hour RMED values, for which data are very much sparser than in the case of 1-day values. This chapter describes how topographic information has been used as a covariate for mapping RMED of various durations, using the georegression approach. The project has made much use of topographic variables developed in a study of RMED mapping in Scotland (Prudhomme and Reed, 1997a). The development of these variables is summarized in the following section.

7.2 Development of topographic variables: Scottish study

7.2.1 Approach

Relationships between precipitation and topography, especially elevation, have been studied for a long time. For example, Bleasdale and Chan (1972) review research in Great Britain and Ireland from the late 19th century. One of the best known relationships is the orographic effect: rainfall increases with elevation. This simple relation has to be moderated by a secondary phenomenon, the shadow effect (Flohn, 1969). In addition to the elevation and the shadow effects on rainfall, "other factors, such as the direction and the distance to moisture sources, as well as the synoptic climatology of a given region, considerably complicate and, hence, weaken precipitation-elevation relationships" (Konrad, 1996).

A study of extreme rainfall in Scotland (Prudhomme and Reed, 1997a) aimed to account for the influence of these factors on the distribution of RMED. Scotland, and in particular the Highlands region, was chosen as a challenge. It was expected that a method capable of mapping RMED in this mountainous and sparsely-gauged region would be applicable in other remote parts of the UK.

The method of georegression was adopted, involving the following steps:

(i) 1-day RMED was extracted from the daily annual maximum records.
(ii) Relationships between RMED and topography were investigated, and a 4-parameter linear equation fitted to the calibration sample.
(iii) Differences between the observed and the estimated RMED (i.e. the residuals) were calculated at all sites.
(iv) An ordinary kriging method was applied to the residuals, giving a map of correction factors to apply to the regression estimates for improving the estimation.
(v) Finally, maps of regression estimates and of correction factors were combined to give the final RMED estimate.

7.2.2 The topographic variables

This section describes the topographic variables developed in the study, which were later used for mapping RMED across the UK. Other aspects of the study can be found in Prudhomme and Reed (1997a). The variables were derived from a 1-km Digital Terrain Model (DTM), provided by the Met. Office.

Elevation is the topographical variable most often used in rainfall studies. ELEV is defined as the elevation of the nearest grid point of the DTM to the gauge. Hereafter, this position is referred to as the gauge grid point, ggp. The difference between the "true" elevation (altitude given by the rainfall database) and ELEV can sometimes be up to 100 m, due to the particularly steep relief encountered in Scotland and the relatively large-scale DTM used. However, the error is not considered too significant and is ignored in the rest of the study.

Average elevation is an alternative elevation suggested by Konrad (1996). Two average elevations are calculated: (i) ELEV4, arithmetic mean elevation of the typically 25 DTM-grid-points in a 4x4km4 x 4 kmre centred on ggp; (ii) ELEV10, arithmetic mean elevation of the typically 121 DTM-grid-points in a 10x10 km square centred on ggp.

Geographical position is represented by EASTING and NORTHING (in km units) of ggp on the national grid reference system.

SLOPEi is the *mean slope* between ggp and a point i km away in each of the eight cardinal directions. In this study, SLOPE2, SLOPE5 and SLOPE10 are calculated.
Two *average slopes* are evaluated in each of the cardinal directions:

(i) arithmetic slope, $$ASLOPE=\frac{SLOPE2+SLOPE5+SLOPE10}{3};$$

(ii) weighted slope, $$WSLOPE=\frac{2\times SLOPE2+5\times SLOPE5+10\times SLOPE10}{17}.$$

Unlike in most previous studies, the slope variable was considered at several distances and in several directions.

Average distance to the sea (SEA) is the average distance from the sea in a 90° sector centred on

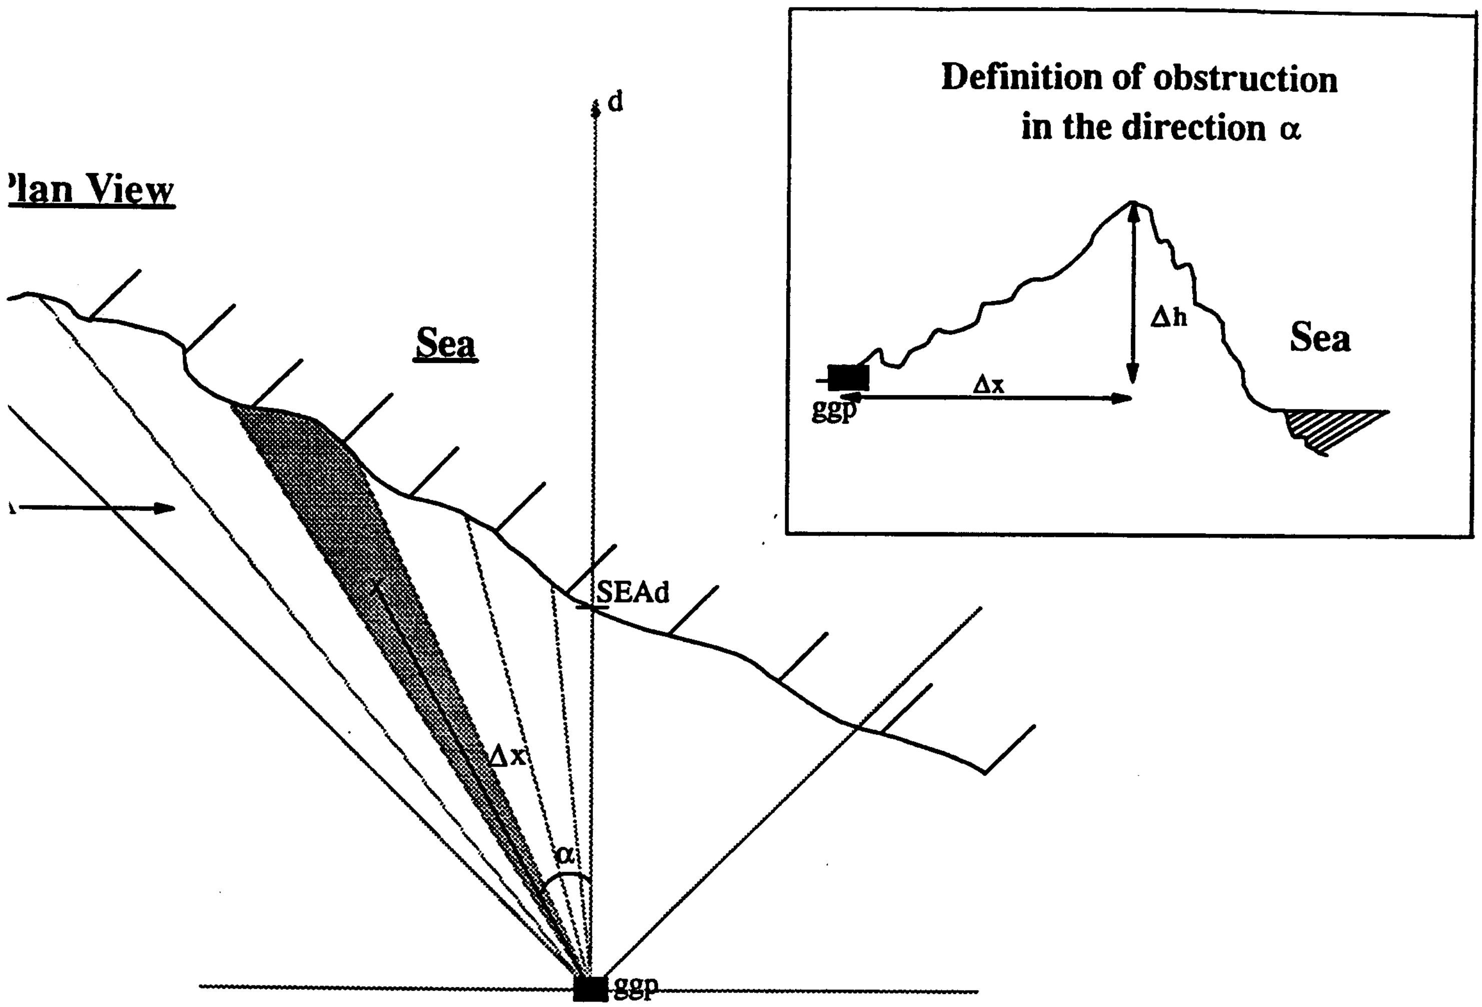

Figure 7.1 Distances used in definition of the variables SEA, OBST and BAR

the direction of interest d. More weight is given to directions close to d.

Obstruction (OBST) represents the average angle subtended by the highest barrier in a sector centred on the direction of interest, d. Again, more weight is given to directions close to d. BAR is the distance to that barrier.

Figure 7.1 illustrates the definition of SEA, OBST and BAR. The plan view shows the 90 degree sector A and some of the eleven sub-directions radiating from the ggp. The variables are averaged over these eleven directions, each value being weighted by cos α to give more weight to values close to the direction d. The longitudinal section shows the gauge and the horizon in the direction inclined at angle α to the main direction, d. The variable BAR is the average of the eleven distances to the horizon, Δx. The variable OBST is the average of the eleven angles subtended by the horizon, arctan ($\Delta h / \Delta x$).

SHIELD represents the *roughness index* for a sector centred on the direction d. It is the sum of the "ups and downs" in elevation, each elevation difference being divided by its distance from the gauge, so that roughness closer to the gauge is given more weight.

The use of averages over sectors centred on the direction of interest ensures that the variables are not unduly influenced by a particular, possibly very small, feature (for example a hill or a sea loch) which happens to lie in the direction of interest. Full details of the calculation of these variables can be found in Prudhomme and Reed (1997b).

7.2.3 Regression with RMED

A standard statistical package (SAS Institute Inc., 1989) was used for the analysis. Single variable regressions between 1-day RMED and the various topographical variables were first carried out. Then multiple regression analysis was performed using the above topographical variables to explore variation in RMED. The criteria for choosing variables were that the explanatory variables should not be highly correlated and, as much as possible, they should respect physical explanations for RMED variations. For example, considering that the prevailing winds in the study region come from the west and the southwest, explanatory variables oriented W-E or SW-NE are more likely to be physically meaningful than variables oriented N-S or NW-SE.

An initial multiple regression analysis on RMED (Prudhomme and Reed, 1997a) revealed systematic structure in the residual error in estimating RMED. The undesired structure in the residuals diminished when RMED was replaced by 1000/RMED. Only the results for the transformed variable are presented here.

The model-building technique used was a stepwise multiple regression: one after the other, variables are added to the regression. To improve the regression performance, each new variable should explain a part of the variance of RMED which is not explained by the previous variables. The preferred model for 1000/(1-day RMED) in Scotland was:

$$1000/RMED = 35.69 - 0.103 OBST_{NE} - 0.023 SHIELD_{SW} + 0.019 SEA_{SW} - 0.040 SEA_{NE}$$

This model explains 53.4% of the variance of 1000/RMED in the Highlands (fitting area) and

42.8% of the variance in the whole of Scotland. It is a mixture of geographical parameters (average distance from the sea in opposing directions) and of topographical parameters (obstruction against the prevailing wind, and roughness between the main moisture source and the gauge).

7.3 Mapping hourly RMED across the UK

RMED was mapped across the UK using the georegression procedure outlined in Section 7.2.1.

7.3.1 Data and covariates

The challenge of mapping RMED was greatest for hourly durations. As is evident from Figure 2.1, parts of the UK are very sparsely covered with recording raingauges. Because several key gauges in remote locations (parts of Northern Ireland and south-west Scotland) had only nine years of records, the minimum record length was reduced from ten to nine years. A total of 375 gauges with 1-hour RMED data were available for the UK, including 23 gauges in Northern Ireland. The amount of data available in Northern Ireland was not sufficient to fit a separate regression model. Therefore, the data from the whole of the UK were considered simultaneously, with Northern Ireland grid references transformed into the national grid system for the regression.

Maps of RMED were produced for durations of 1, 2, 6 and 12 hours. The numbers of gauges with sufficient data for each duration were: 375 for 1 hour, 334 for 2 hours, 300 for 6 hours and 272 for 12 hours. The reason for the differences between durations is that tabulated annual maxima were not available for all durations at all gauges.

RMED was particularly difficult to map for the 1 hour duration, owing to large variations in the sample values over short distances, especially in eastern England. It was thought that data connected with thunderstorm activity might help explain these variations, and so lightning strike data were analysed. The data were obtained from EA Technology and consist of the average strike density within 10x10 km "bins" over the UK for the years 1988 to 1996. The average density is defined by counting the number of strikes falling within a bin and divided by 100. The mean (AVLIGHT) and median (MEDLIGHT) strike density for the period 1988-1996 were calculated, and added to the topographic variables for the multiple regression analysis.

In addition to the topographic variables developed in the Scottish study (Section 7.2.2) and the lightning data, two other covariates were thought to be worth trying. The first was the interpolated 1-day RMED (RMEDDY) at the site of the recording raingauge. The interpolated value (mapped as described in Section 7.4) was used because not all recording raingauges have a nearby daily gauge with sufficient record length. The second new variable was an index of continentality, which was taken as being the distance from Lille in north-west France, DLILLE. The latter index has a similar effect to the "continentality factor" which was included in the model for M5 rainfall given in FSR volume 2 section 3.3.3 (NERC, 1975). DLILLE is small in south-east England where thunderstorms are more frequent and short duration falls tend to be heavy.

7.3.2 Regression

Single variables

Figure 7.2 shows the sample values of 1-hour RMED at gauge locations. The largest values are on Ben Nevis, in the Lake District, and at two sites (Bedford and Tugby) in the East Midlands. The smallest are in Hampshire, eastern Scotland and Northern Ireland. The large local variations in the East Midlands appear to be unrelated to topographic features, and are discussed further in the next section.

The relationships between x-hour RMED and the topographic and lightning data were first explored by single regressions. Selected results are displayed in Table 7.1.

For all of the UK taken together, the two most influential variables on 1-hour RMED are DLILLE and NORTH. For longer durations, as convective rainfall becomes less important and frontal rainfall more so, RMEDDY has a rapidly increasing influence. Note that over 70% of the variation in 12-hour RMED is explained by interpolated 1-day RMED.

Table 7.1 Percentage of variance in RMED explained (R^2) for selected covariates

Variable	1 hour RMED	2 hour RMED	6 hour RMED	12 hour RMED
DLILLE	14.4	1.9	3.2	5.3
NORTH	13.9	3.1	1.2	2.3
RMEDDY	5.1	21.6	62.9	70.3
AVLIGHT	4.1	0.0	4.7	6.5
MEDLIGHT	5.6	0.4	3.9	6.5

The lightning strike variables (AVLIGHT and MEDLIGHT) explain surprisingly little of the variance of RMED for any duration. In addition to analysis of the relation between the lightning data and RMED in the whole of the UK, different sub-regions were then considered. In none of the sub-regions was a strong relationship between x-hour RMED and lightning strike density found.

Multiple regression

To build a model suitable for georegression, a stepwise multiple regression was performed between x-hour RMED and topographical variables, lightning variables and RMEDDY. The best equations were chosen on the basis of physical realism as well as the degree of variance explained. The chosen model parameters are shown in Table 7.2.

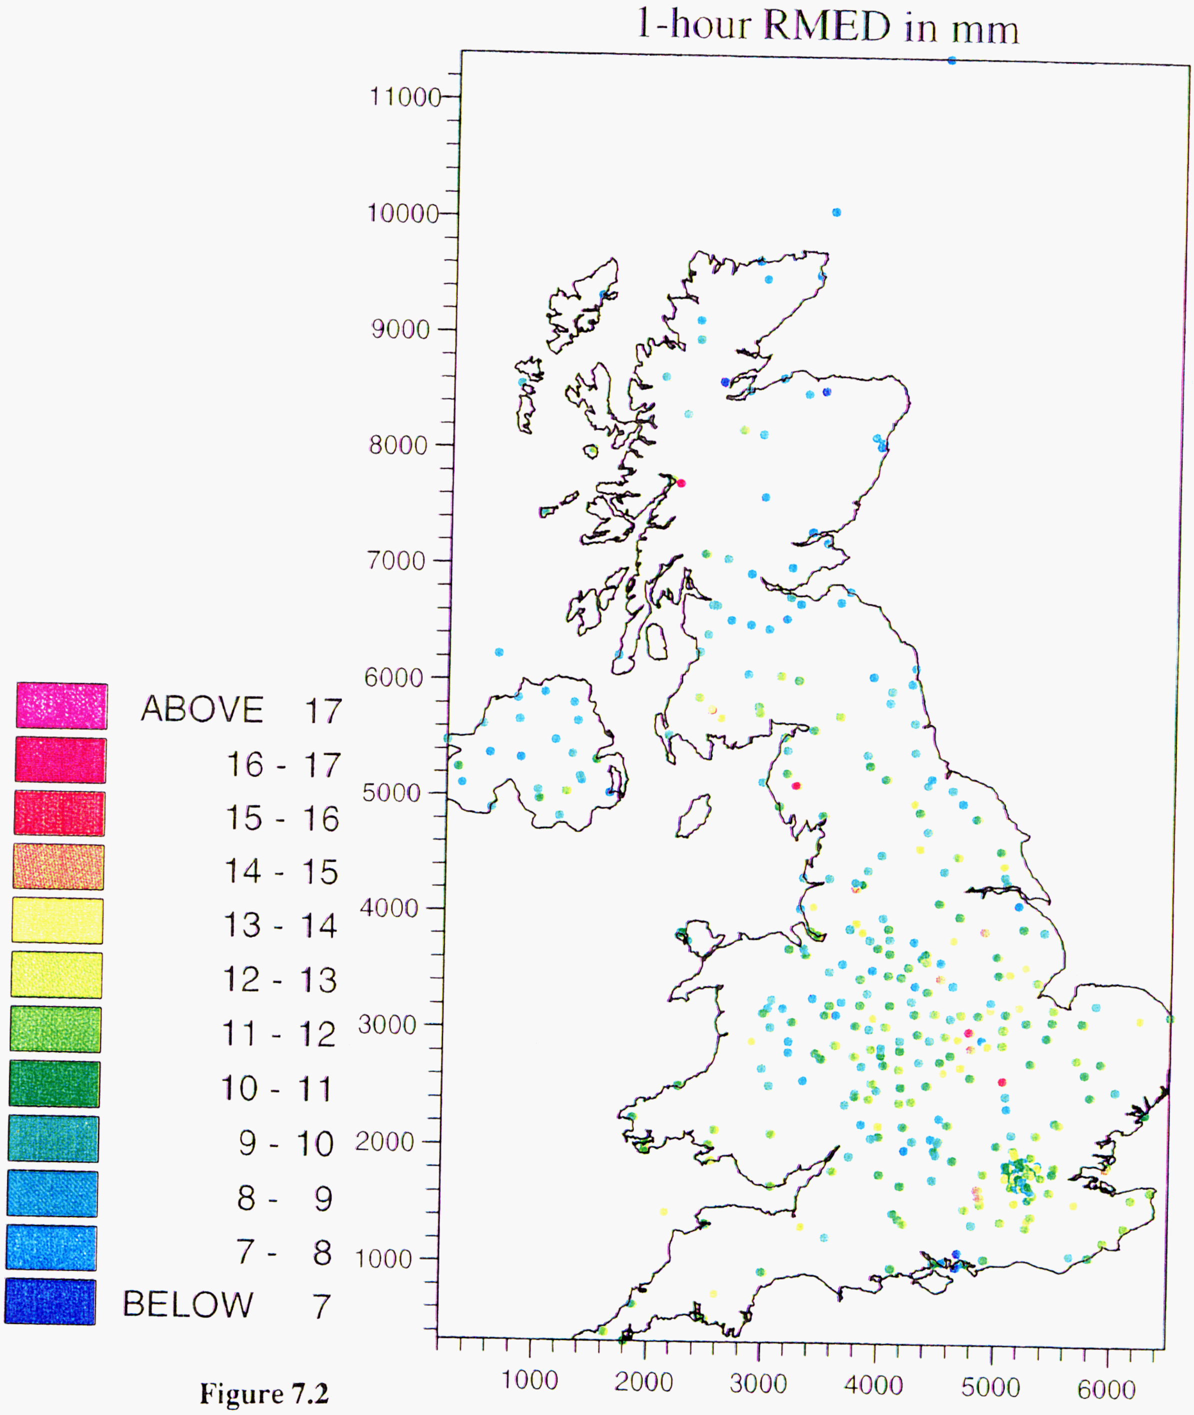
1-hour RMED in mm
ABOVE 17
16 - 17
15 - 16
14 - 15
13 - 14
12 - 13
11 - 12
10 - 11
9 - 10
8 - 9
7 - 8
BELOW 7
11000
10000
9000
8000
7000
6000
5000
4000
3000
2000
1000
1000
2000
3000
4000
5000
6000

Figure 7.2

Table 7.2 **Coefficients of the variables in regression models for *x*-hour RMED**

Duration	Constant	$OBST_{NW}$	BAR_{SE}	DLILLE	RMEDDY	R^2 (%)
1 hour	8.636	-0.030	0.012	-0.00036	0.100	36
2 hours	8.868	-0.049	0.015	-0.00035	0.216	46
6 hours	6.756	-	0.014	-0.00033	0.540	70
12 hours	3.515	-	-	-	0.793	79

DLILLE expresses the increase in short-duration extreme rainfalls towards the south-eastern corner of the UK. The presence of $OBST_{NW}$ and BAR_{SE} in the models supports some findings by May and Hitch (1989). They mapped 1-hour M5 (i.e. 5-year return period) rainfall in the UK and looked at station values along a swath running approximately SE-NW from Hastings to Watnall. The maxima in 1-hour M5 rainfall appeared to be co-located with south-east facing ground slopes. May and Hitch suggested that this was an effect of heavy rainfalls produced by orographic uplifting associated with summer thunderstorms which develop over France, cross the English channel and travel north-westwards into the Midlands. The present model, with a positive coefficient for BAR_{SE} and a negative coefficient for $OBST_{NW}$, also produces higher values for 1 and 2-hour RMED on south-east facing slopes.

The number of parameters required is smaller for 6 and 12-hour RMED, because RMEDDY explains so much of the variance on its own (Table 7.1), and some topographic variables are no longer significant terms in the regression model.

7.3.3 Further comments on the 1-hour RMED model and data

Several attempts were made to improve the model for 1-hour RMED, which explains only 36% of the variance in the sample data. The data were explored by splitting the country into regions and fitting separate models. This exercise indicated that the poorest model performance is in the east and south of England, where a four-parameter model could explain only 23% of the variance. Models for the regions contained different topographic variables, which suggests that different factors influence 1-hour RMED in various parts of the UK. However, it was decided to retain a nationwide model to avoid the problem of discontinuities at regional boundaries.

Another idea was to map the ratio of 1-hour RMED to 1-day RMED. It was found that this ratio cannot be estimated at ungauged locations significantly better than 1-hour RMED on its own. The two models were compared on the basis of a regression between modelled and observed 1-hour RMED. On this basis, the non-ratio model had a similar R^2 value and a regression line closer to 1:1.

The "awkward" 1-hour RMED sample values in the East Midlands were thoroughly investigated for effects such as RMED being large for gauges with short records covering particularly wet years. No such effects were found. Annual maxima from sites in this region a few kilometres apart can be remarkably different, and in some cases these differences are consistent throughout the record, giving rise to significant local differences in RMED. One of the longest records, at Watnall, covers 46 years and has a 1-hour RMED of 14.2 mm. This is close to three gauges with RMED smaller than 9.0 mm. In the end, all these data were retained, although the local variations were smoothed due to the nature of the kriging applied to the residuals (section 7.3.4).

7.3.4 Interpolating the residuals

The residuals from the models for each duration were evaluated at the gauge locations. They were defined by:

$$res = RMED_{modelled} - RMED_{observed}$$

The residuals are not very consistent from site to site, particularly in the East Midlands. Over short distances there appears (from semivariogram plots) to be no spatial structure in the residuals. However, it was possible to fit a semivariogram model over fairly long inter-site distances and interpolate the residuals on a 1 km grid using the method of ordinary kriging.

The theory of geostatistics and kriging is explained in the Phase 1b report (Stewart *et al.*, 1995). For more detail, the book by Journel and Huijbregts (1978) is recommended. The calculation of semivariograms and kriging were carried out using a library of Fortran subroutines known as the Geostatistical Software Library or GSLIB (Deutsch and Journel, 1992).

Figure 7.3 shows the semivariogram model used for residuals of 1-hour RMED. Note that the nugget effect is large, reflecting the large variability over short distances. All semivariogram models used were isotropic, i.e. the same in all directions. The semivariogram model in Figure 7.3 does not fit the data very well. However, it was found that the exact form of the model has little influence on the results (due in part to the large nugget effect).

The kriging technique was set up to cope with the sparsity of gauges in parts of the UK, with a minimum of one gauge required within the search radius, and a maximum of ten. The maximum search radius was set to 100 km. In other words, there must be at least one sample value within 100 km for a kriged estimate to be produced. Because of the large nugget effect, even sample values very close to a grid point do not have an overwhelming influence on the estimate at that grid point. This acts to smooth the local variations in sample values. The software was modified to ensure that no sample value could be located exactly on a 1-km grid point. If this had occurred, the estimate would have been set equal to that sample value, and the smoothing would have broken down.

7.3.5 Running the models and producing final maps of *x*-hour RMED

A model estimate of *x*-hour RMED was evaluated from the topographic variables and 1-day RMED at every 1-km grid node in the UK. The interpolated residuals were then added to these estimates to give final combined maps of RMED, which are shown and discussed in the following chapter. This method of georegression thus accounts explicitly for local variations in the performance of the regression model.

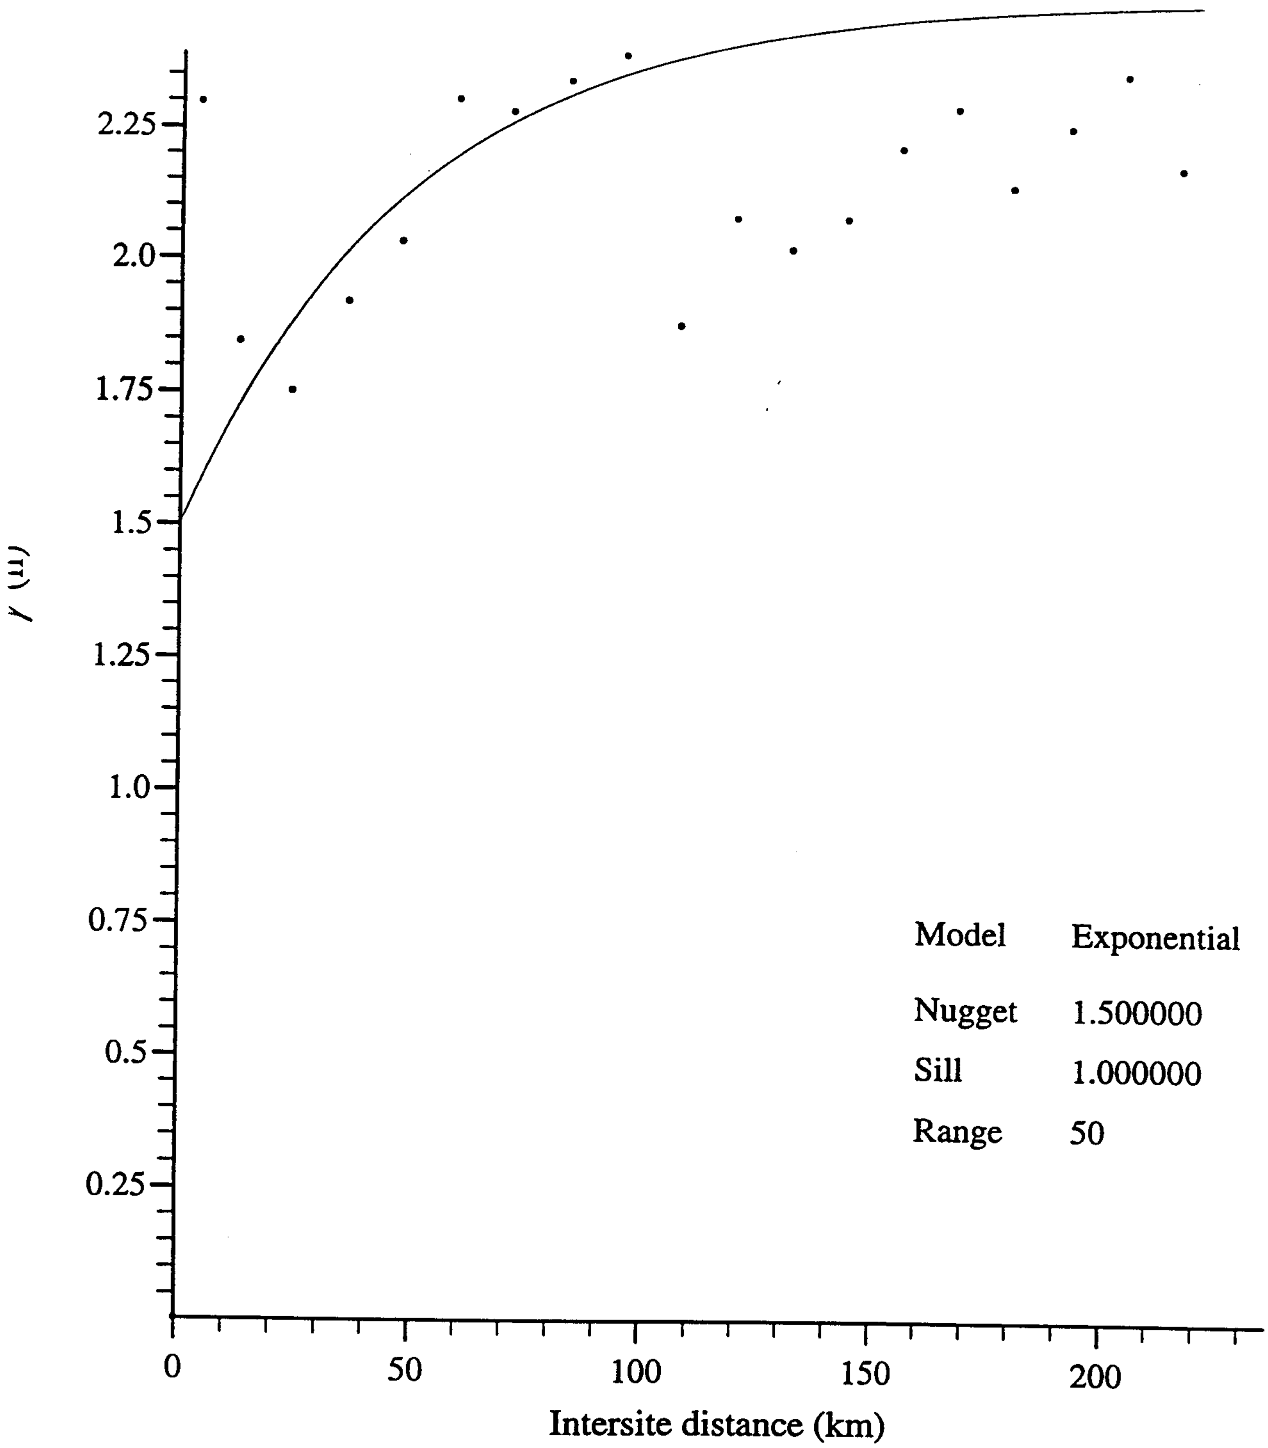

Figure 7.3 Semivariogram for residuals of 1-hour RMED

7.4 Mapping daily RMED across the UK

7.4.1 Data

Daily data are available for many more sites than hourly data, however, as was noted in Section 7.1, regions where topography (and hence rainfall) is most varied tend to be the most sparsely gauged. A total of 5843 gauges with 1-day RMED estimates were available for Great Britain, using a minimum record length of ten years. Since ample data were available for Northern Ireland, the province was mapped separately using a different regression model. This was not attempted for other parts of the UK because of the problem of discontinuities at regional boundaries.

RMED for durations 1, 2, 4 and 8 days was mapped across Great Britain by georegression. Figure 7.4 shows the sample values of 1-day RMED at gauge locations. RMED is obviously higher in mountainous areas, particularly in the west.

7.4.2 Regression

The regression model used in Scotland was found to be inappropriate for mountainous parts of Great Britain outside Scotland, possibly because of the use of the variables $OBST_{NE}$ and SEA_{NE} which can measure much longer distances in areas such as Wales. The Scottish model explained just 15% of the variance of 1-day RMED across GB, compared with 53% in the Scottish Highlands. Interestingly, the elevation variable ELEV10 (mean elevation of a 10×10 km grid square) is correlated with *x*-day RMED when considering the whole of Great Britain, which was not true in Scotland. One reason for the lack of correlation in Scotland is the rain shadow effect which reduces RMED over the Grampian mountains, a large upland area towards central and eastern Scotland.

It was decided to focus the regression analysis on the mountainous regions of the UK. Topography and rainfall are expected to be more clearly correlated here than in regions with less complex relief (as for instance in the Severn). The priority in regression was to find a model which fitted the data well in these sparsely-gauged regions. Elsewhere, it was expected that the good density of gauges would enable the kriged residuals to account for any variations in the fit of the model.

After several alternative groups of gauges were tried, the model was fitted to data from the "Celtic" region, comprising the Scottish Highlands, north-west England, Wales (but not the Severn region), and south-west England. The region includes 1589 gauges with at least ten years of daily annual maximum. The boundary of the region is marked on Figure 7.4.

Following the Scottish work (Section 7.2.3), the transformed variable 1000/RMED was used. After a stepwise regression for 1000/(1-dayRMED) on topography, the following model was chosen:

$$1000/RMED = 26.74 - 0.0899 OBST_W - 0.0277 ELEV10 + 0.0270 SEA_{SW} + 0.0202 SEA_W$$

The model explains 55% of the variance in the Celtic region, and 57% of the variance in the whole of GB.

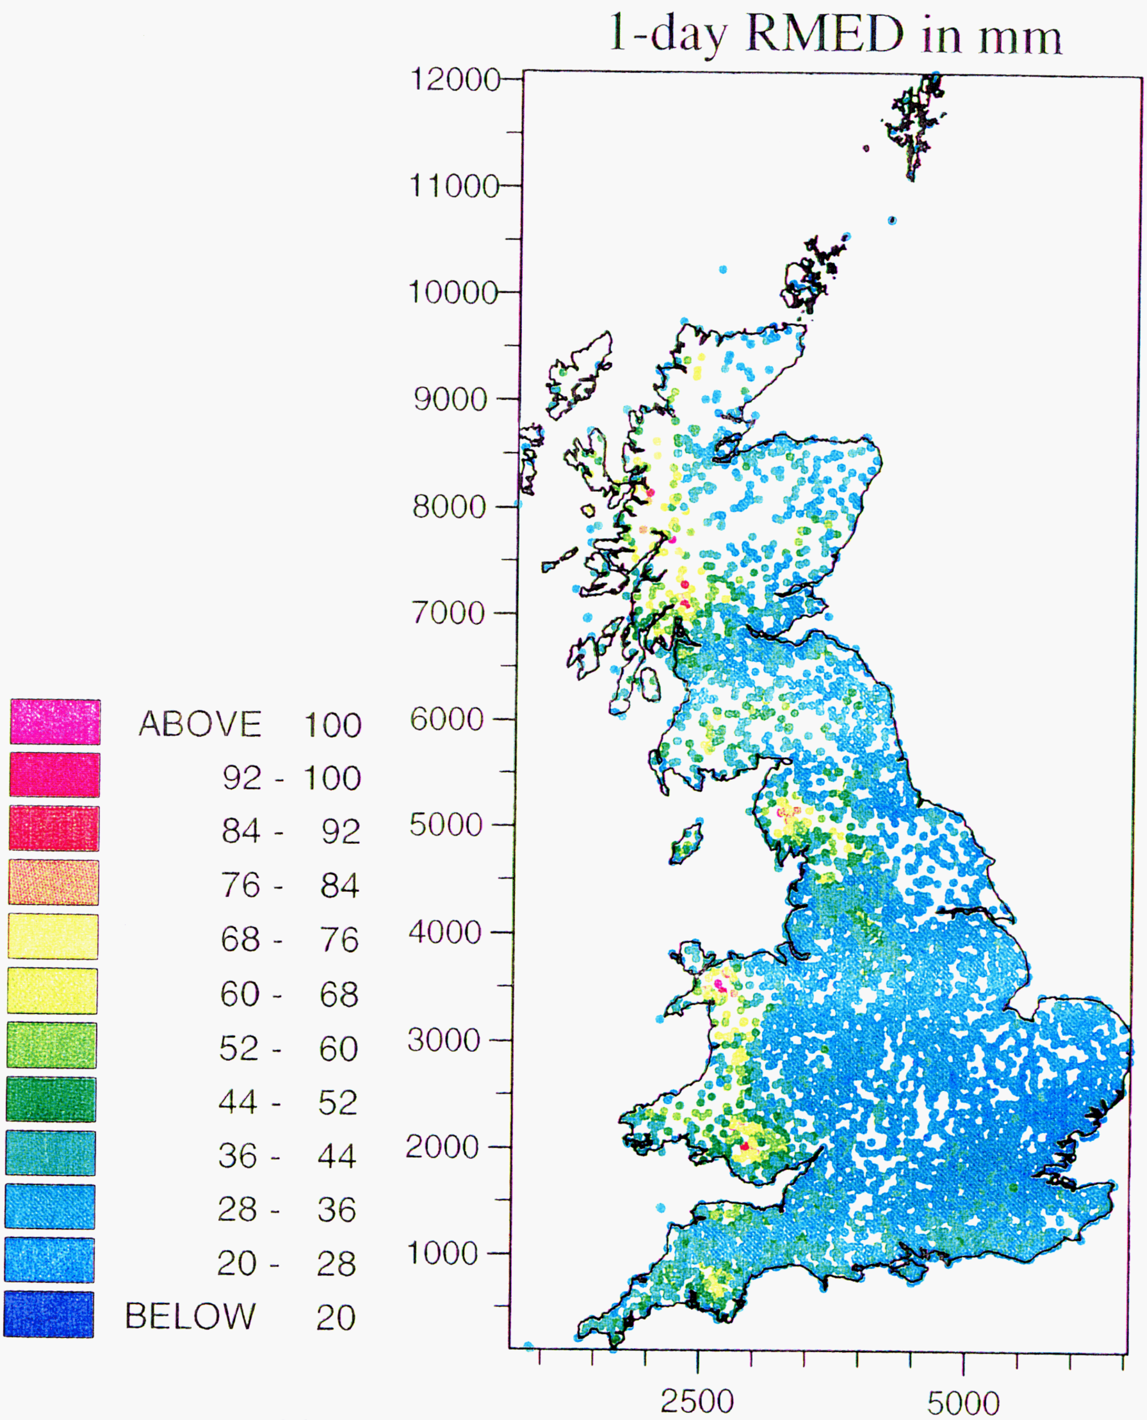
1-day RMED in mm
12000
11000
10000
9000
8000
7000
6000
5000
4000
3000
2000
1000
2500
5000
ABOVE 100
92 - 100
84 - 92
76 - 84
68 - 76
60 - 68
52 - 60
44 - 52
36 - 44
28 - 36
20 - 28
BELOW 20

Figure 7.4

This model makes physical sense: the rainfall increases with elevation and with the obstruction to the west, which reflects orographic enhancement of rainfall moving from the west. RMED decreases with distance to the sea in the west and southwest directions, in accordance with the prevailing wind directions which bring frontal rainfall, particularly to the westerly parts of the country where the model is fitted. Even though SEA_W and SEA_{SW} are highly correlated (correlation coefficient of 0.77), the part of the variance they both explain and their meteorological sense are good reasons to keep them both together in the model. *NB. We may consider trying a variable SEA_{WSW} instead.*

For longer durations, it was decided to use the same quartet of topographic variables, since similar processes give rise to extreme rainfall for durations between 1 and 8 days. Table 7.3 gives the equations for durations 1, 2, 4, and 8 days.

Table 7.3 Coefficients of the variables in regression models for *1000/(x*-day RMED)

Duration	Constant	$OBST_W$	ELEV10	SEA_{SW}	SEA_W	R^2 (Celtic)	R^2 (GB)
1 day	26.74	-0.090	-0.028	0.027	0.020	55%	57%
2 days	20.50	-0.078	-0.023	0.017	0.021	58%	64%
4 days	15.15	-0.063	-0.019	0.018	0.015	62%	73%
8 days	10.59	-0.049	-0.014	0.015	0.012	63%	75%

The coefficients vary smoothly from one duration to the other. For longer durations, more of the variance is explained by the topographical model. This is partly explained by a reduction in the geographical variation of RMED for longer durations.

The regression models perform in general fairly well everywhere, with more of the variance explained in the whole of Britain than in the fitting region. This is because the Celtic region was deliberately chosen to include the largest variations in topography and rainfall. In much of eastern and southern England RMED varies over a very small range (see Figure 7.4).

7.4.3 Comparison between georegression and ordinary kriging for 1-day RMED

Comparison of the performances of ordinary kriging (using no topographic data) and georegression was done through a cross validation. It is not sufficient simply to compare by cross-validation the performance of kriging RMED and kriging the residuals, because the two methods are estimating different variables. Instead, the accuracy of the final estimates of RMED were compared at the sites of a set of validation gauges which were not used at all in the kriging or georegression. Five different ways of splitting the gauges into two groups (calibration/validation) were selected randomly. The procedures reported below were thus repeated five times.

Ordinary kriging

Ninety percent of the total gauges were selected to constitute the calibration samples. Theoretical variograms were fitted to the calibration samples, and then used for the complete kriging procedure.

Georegression

The gauges situated in the fitting area (Celtic region) were reselected from the calibration sample, and the regression equation using the variables ELEV10, SEA_W, $OBST_W$ and SEA_{SW} re-fitted.

The residuals between the observations and the new equation were calculated for all the gauges. Theoretical variograms were fitted to the residuals for the calibration sample, and used for kriging the residuals of the calibration gauges.

Final estimates of 1-day RMED were calculated for the sites of the validation gauges, combining the regression estimates and the kriged residuals.

Comparison of the results

Linear regression between estimates and observations was performed for each of the five validation samples. The bias in the estimation is judged by the deviation of the model from a perfect model with slope 1 and intercept 0. The accuracy of the estimation is measured by the percentage of variance explained by the model (R^2), which in the perfect case would be 100%.

For all the samples, the estimation of RMED by georegression is less biased than when using ordinary kriging (table 7.4). For all samples but two, the estimation of RMED is more accurate when using georegression, more variance of RMED being explained than with ordinary kriging.

Table 7.4 Comparison of the estimates of RMED for the 5 validation samples using ordinary kriging and georegression

Sample	R^2 (%)	Slope	Intercept	Method
1	85.3	0.772	8.076	ordinary kriging
	86.2	0.871	4.644	georegression
2	83.6	0.810	6.841	ordinary kriging
	80.2	0.892	3.924	georegression
3	80.9	0.745	8.839	ordinary kriging
	84.0	0.798	7.98	georegression
4	86.6	0.818	6.471	ordinary kriging
	87.4	0.856	5.617	georegression
5	85.4	0.802	6.722	ordinary kriging

Sample	R^2 (%)	Slope	Intercept	Method
	83.3	0.857	4.841	georegression

This nationwide comparison of the two methods underplays the more substantial increase in performance of georegression over ordinary kriging in sparsely gauged areas where the regression model plays a large part in providing the georegression estimate.

7.4.4 Interpolating the residuals

The structure of the maps of residuals is quite similar for all durations. Residuals of 1000/RMED are generally positive in the eastern part of Britain, particularly in Lincolnshire and Yorkshire where SEA_W and SEA_{SW} are large but RMED is not substantially smaller than elsewhere in England. Residuals are most negative over the Grampian region, where the model cannot account for the rainshadow effect evident in the data (Figure 7.4). Figure 8.1 shows the interpolated 1-day residuals. Note that residuals are defined by:

$$res = \frac{1000}{RMED_{modelled}} - \frac{1000}{RMED_{observed}}$$

In contrast to hourly durations, the residuals are consistent between nearby gauge locations, and semivariogram models can be fitted easily to the spatial structure of the residuals. An example (for 1 day residuals) is shown in Figure 7.5.

The minimum number of points used to estimate RMED at each grid point was 1 and the maximum 12. The maximum search radius was defined as 80 km.

7.4.5 Problems in estimating long-duration rainfall in mountainous areas

In some regions where the relief is very steep, the regression model tends to give inconsistent results in several grid cells for durations longer than 1 day. This is notably the case near Ben Nevis in Scotland, and near the Cuillins in the isle of Skye, and it also occurs at a few grid cells in Snowdonia.

The problems are caused by the inverse transformation of RMED, which produces small figures in mountainous regions where $RMED_{modelled}$ is large. At some grid squares, the interpolated residual is based on data from gauges at some distance, and is not appropriate for the local conditions. In these cases, the residual may be similar to or larger than $RMED_{modelled}$. When the residual is subtracted from the model estimate, the final result (in mm, after back-transformation) is near-infinite, either large and positive or large and negative.

Several solutions to the problem were investigated. It was decided to define problem cells as those where the combined estimate of RMED was negative or exceeded the following limits:

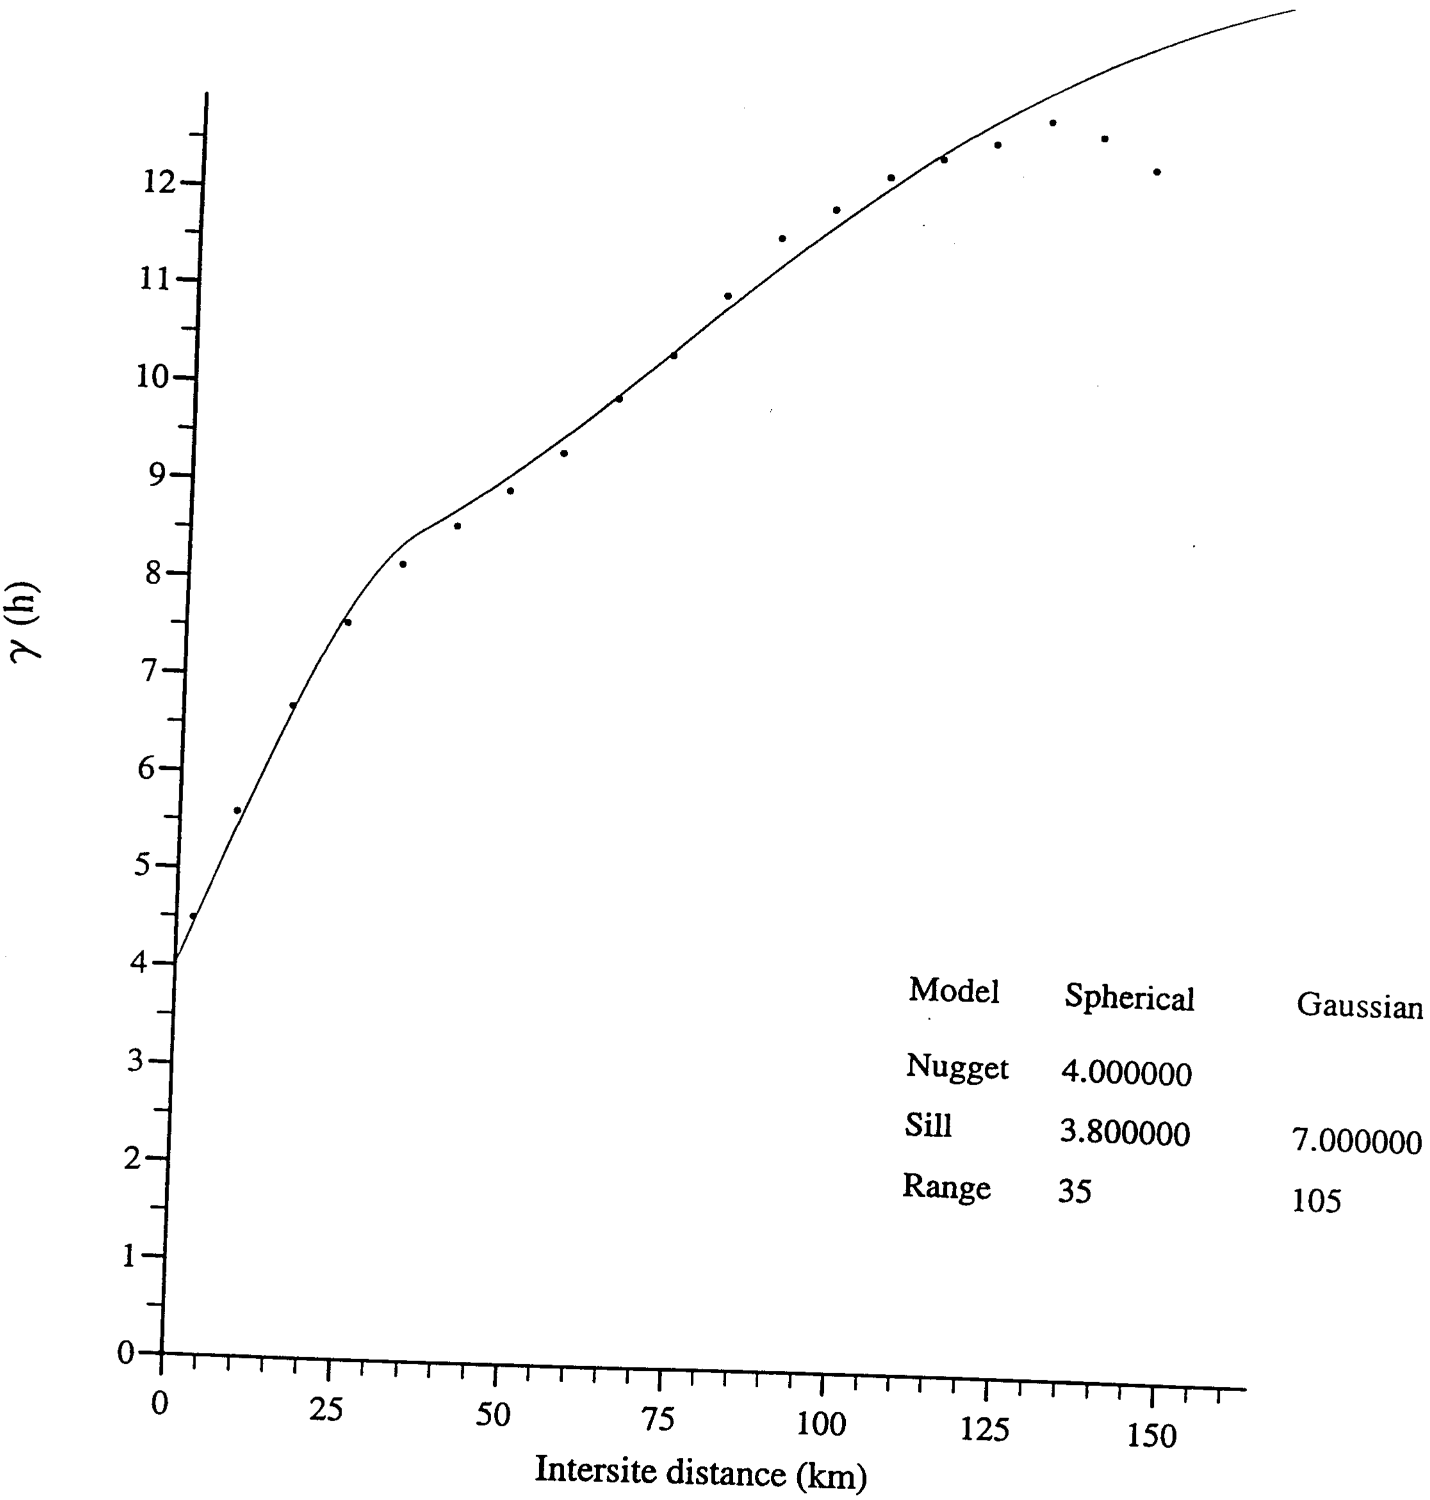

Figure 7.5 **Semivariogram for residuals of 1-day RMED**

Duration (days)	*Upper limit on RMED (mm)*
2	250
4	350
8	500

The preferred solution is to use the estimates of 1-day RMED (which do not suffer from the problem) as a base to estimate RMED for longer durations in the cells where the problem occurs. The ratio $\frac{x\text{-}day\ RMED}{1\text{-}day\ RMED}$ was mapped and found to vary more smoothly in space than x-day RMED. A typical value of the ratio for each duration was found for cells surrounding the problem areas. This value was then used to estimate x-day RMED in the problem cells.

8. MAPPING RESULTS

8.1 Illustration of georegression for mapping 1-day RMED

The previous chapter explains the method of georegression which was used to map RMED on a 1-km grid across the UK. This chapter discusses the resulting maps, particularly in the context of rainfall frequency estimation for England and Wales.

Figures 8.1 to 8.3 illustrate the stages involved in georegression. Figure 8.1 is a map of the kriged residuals from the model which estimates 1000/(1-day RMED). The map shows how well the regression model is performing across the country. The overestimation of RMED (positive residuals) is mainly concentrated in the Grampians, mid-Wales, the north Pennines and the north part of the Southern Uplands. RMED is underestimated in eastern England and also the western Highlands of Scotland.

Figure 8.2 shows the estimates of 1-day RMED produced by the regression model (before correction). The map is very detailed, following the relief closely. The largest area with high modelled RMED is the Grampian mountains, where the high altitude and steep slopes give rise to estimates of RMED which are significantly larger than in reality (Figure 7.4). In north-east England, RMED is underestimated.

The map of modelled RMED is combined with the map of kriged residuals to give a final corrected map showing the combined estimate of 1-day RMED, Figure 8.3. The overestimation in the Grampians and underestimation in north-east England have been corrected. The final map matches the data values (Figure 7.4) closely, while retaining local details which are not obtainable by using ordinary kriging on its own.

8.3 Maps of RMED for hourly durations

Figures 8.4 and 8.5 are the final maps of RMED for 1 and 6 hours. Figure 8.4 is the most striking, demonstrating that georegression can account for increases in 1-hour RMED in eastern and southern England, due to convective storm activity, as well as a rise in western mountainous regions due to orographic effects. The lowest values are on the east coast of Scotland. Note that some of the extreme local variations in sample values in the East Midlands (seen in Figure 7.2) have been smoothed out, as described in Section 7.3.4.

A rise in RMED towards south-east England is also evident in Figure 8.5, but the effect diminishes with increasing duration. The pattern of 12-hour RMED is similar to that of 1-day RMED (Figure 8.3).

8.4 Maps of RMED for daily durations

Figures 8.3 and 8.6 are the maps of 1 and 8-day RMED. The structures of these maps are similar, as are those for 2 days and 4 days (not shown). The highest RMED values are found typically in mountainous regions, and especially in the western parts of the mountains. The south coast has generally higher RMED than more inland areas, and the lowest RMED can be found at the east coast, and especially in East Anglia.

Kriged residuals of Celt3 model for 1000/RMED

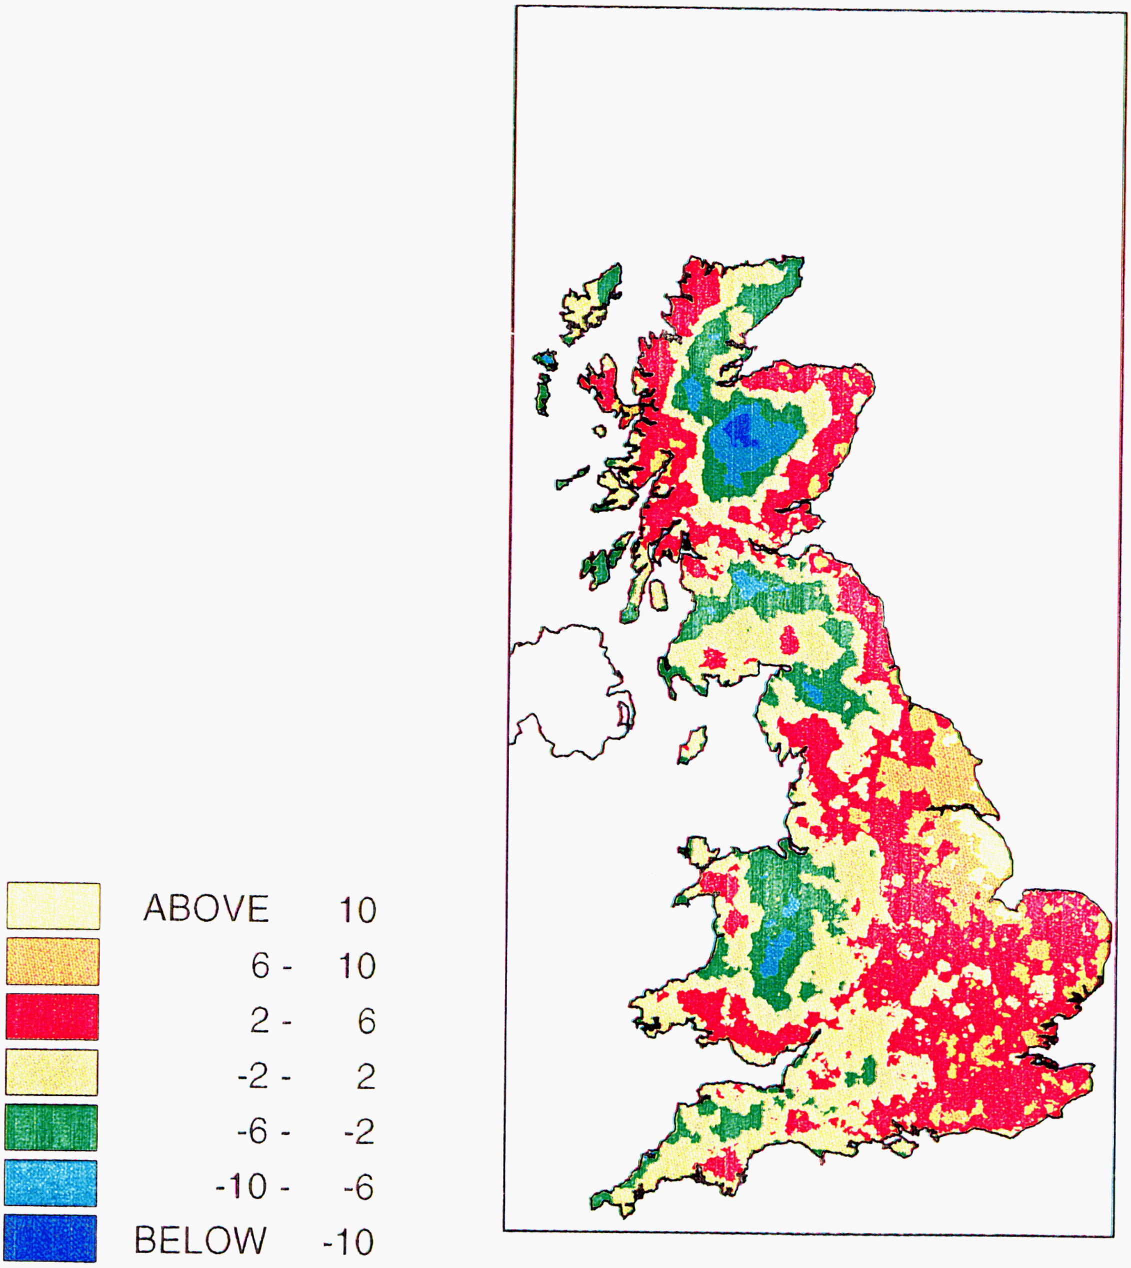

Figure 8.1

1-day RMED (mm) from Celt3 regression

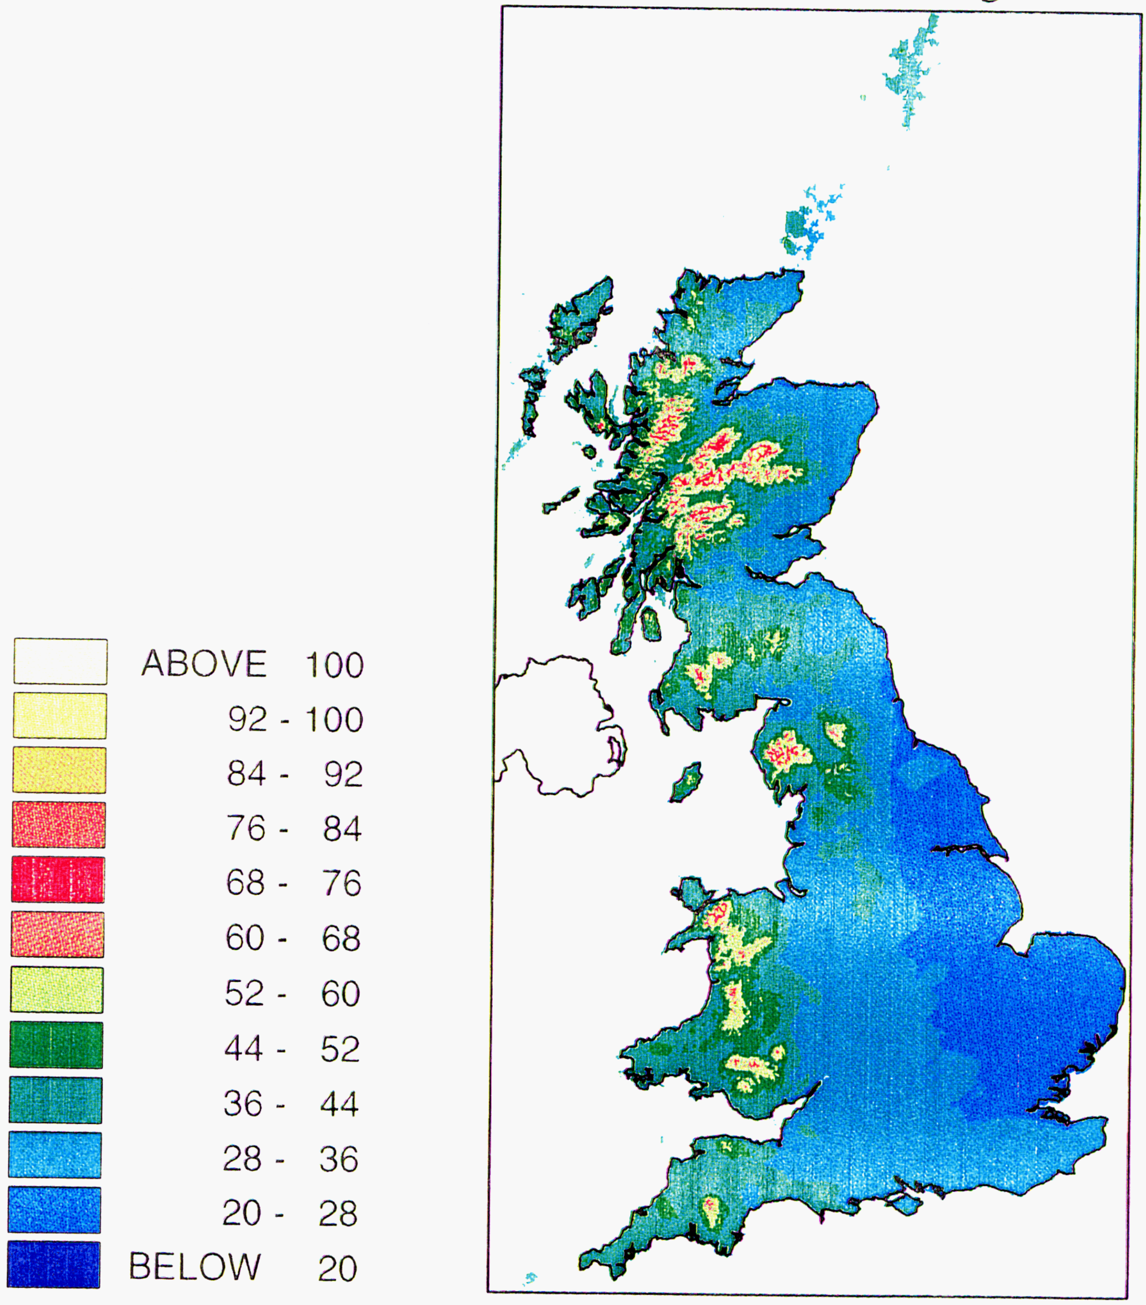

Figure 8.2

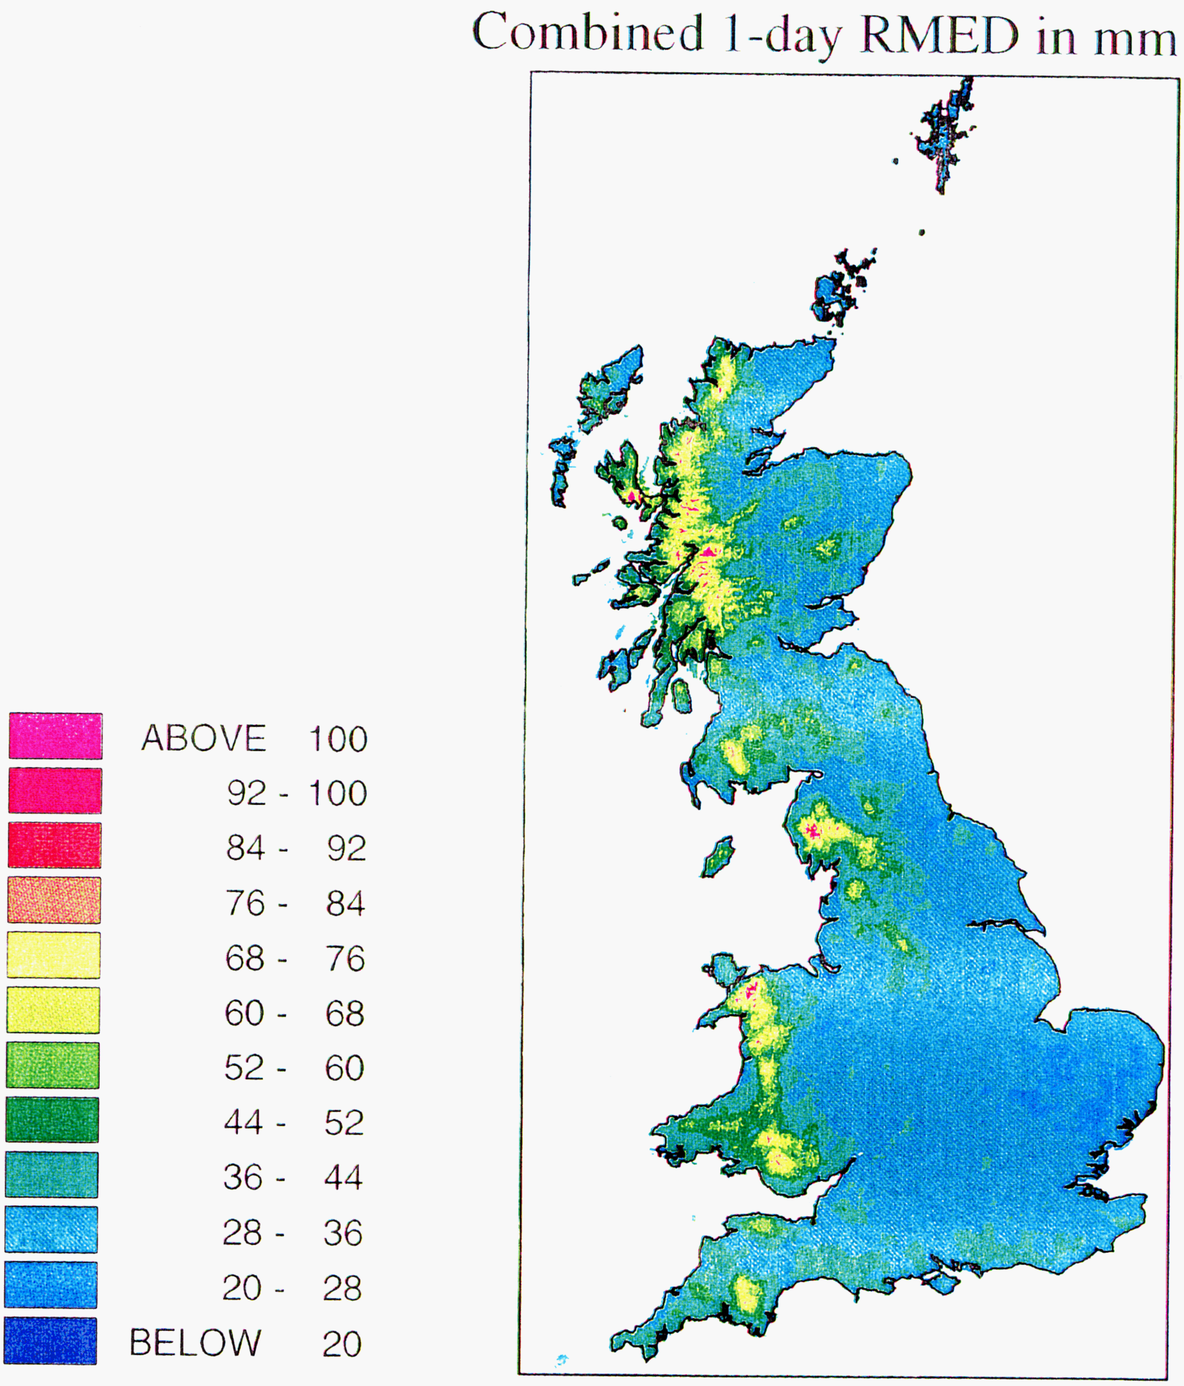
Combined 1-day RMED in mm
ABOVE 100
92 - 100
84 - 92
76 - 84
68 - 76
60 - 68
52 - 60
44 - 52
36 - 44
28 - 36
20 - 28
BELOW 20

Figure 8.3

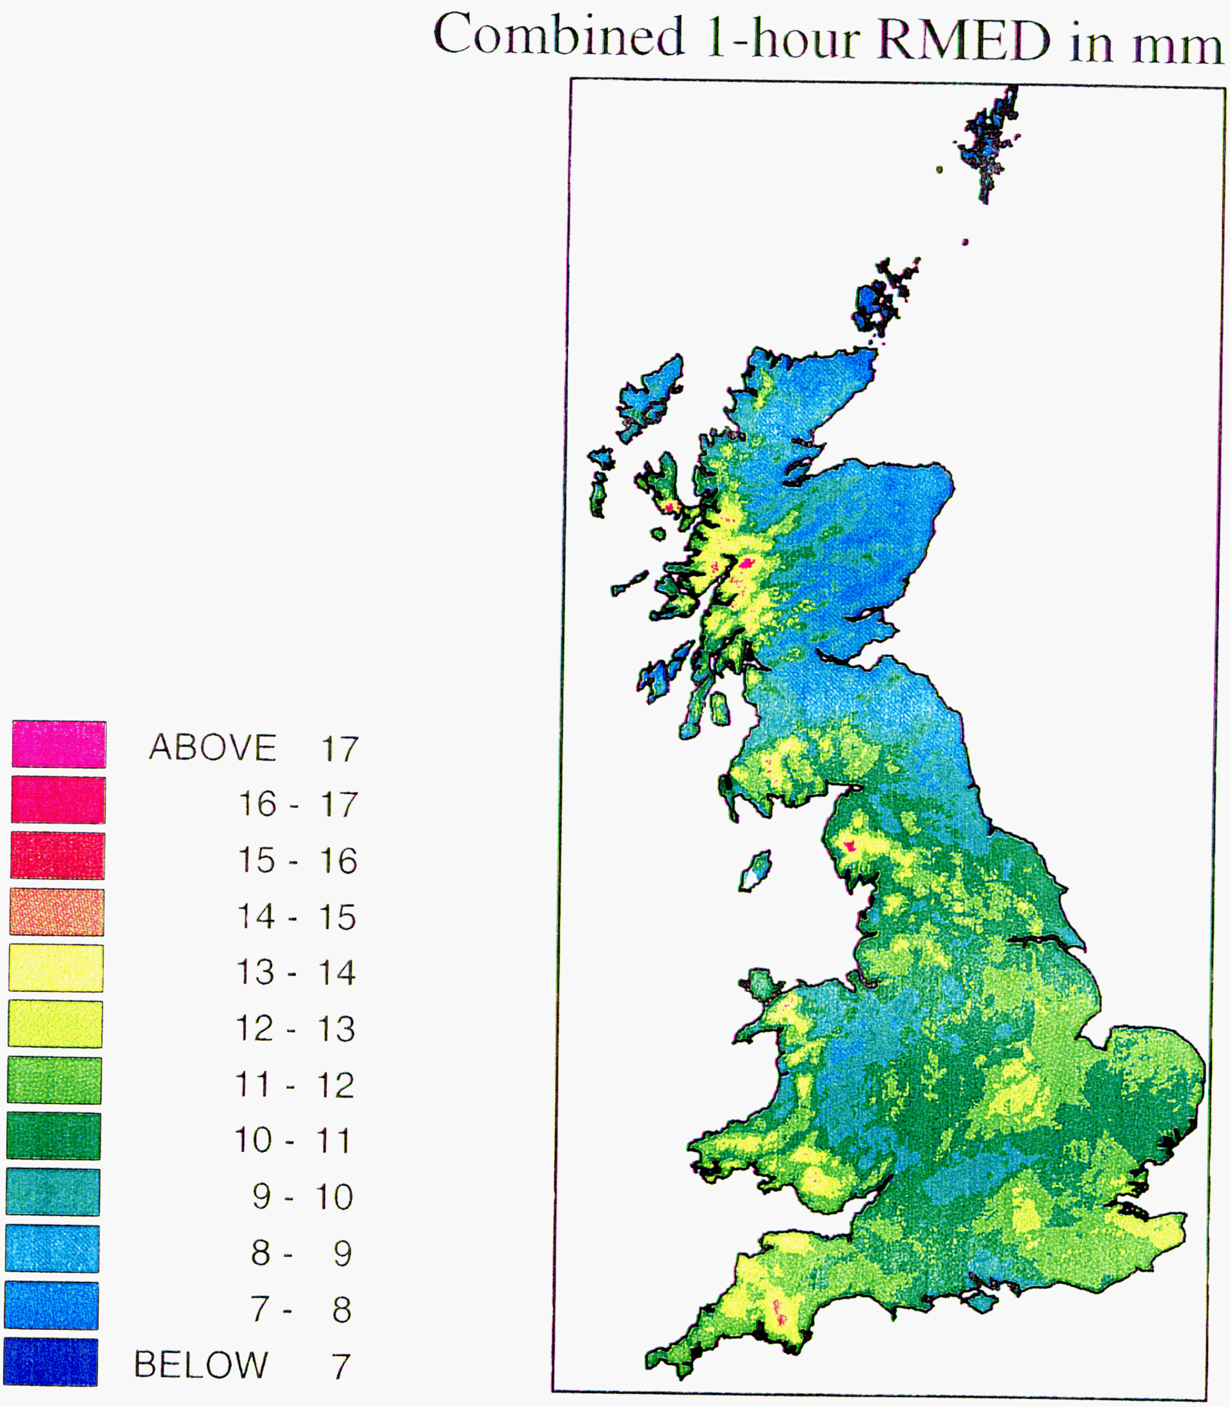
Combined 1-hour RMED in mm
ABOVE 17
16 - 17
15 - 16
14 - 15
13 - 14
12 - 13
11 - 12
10 - 11
9 - 10
8 - 9
7 - 8
BELOW 7

Figure 8.4

Combined 6-hour RMED in mm

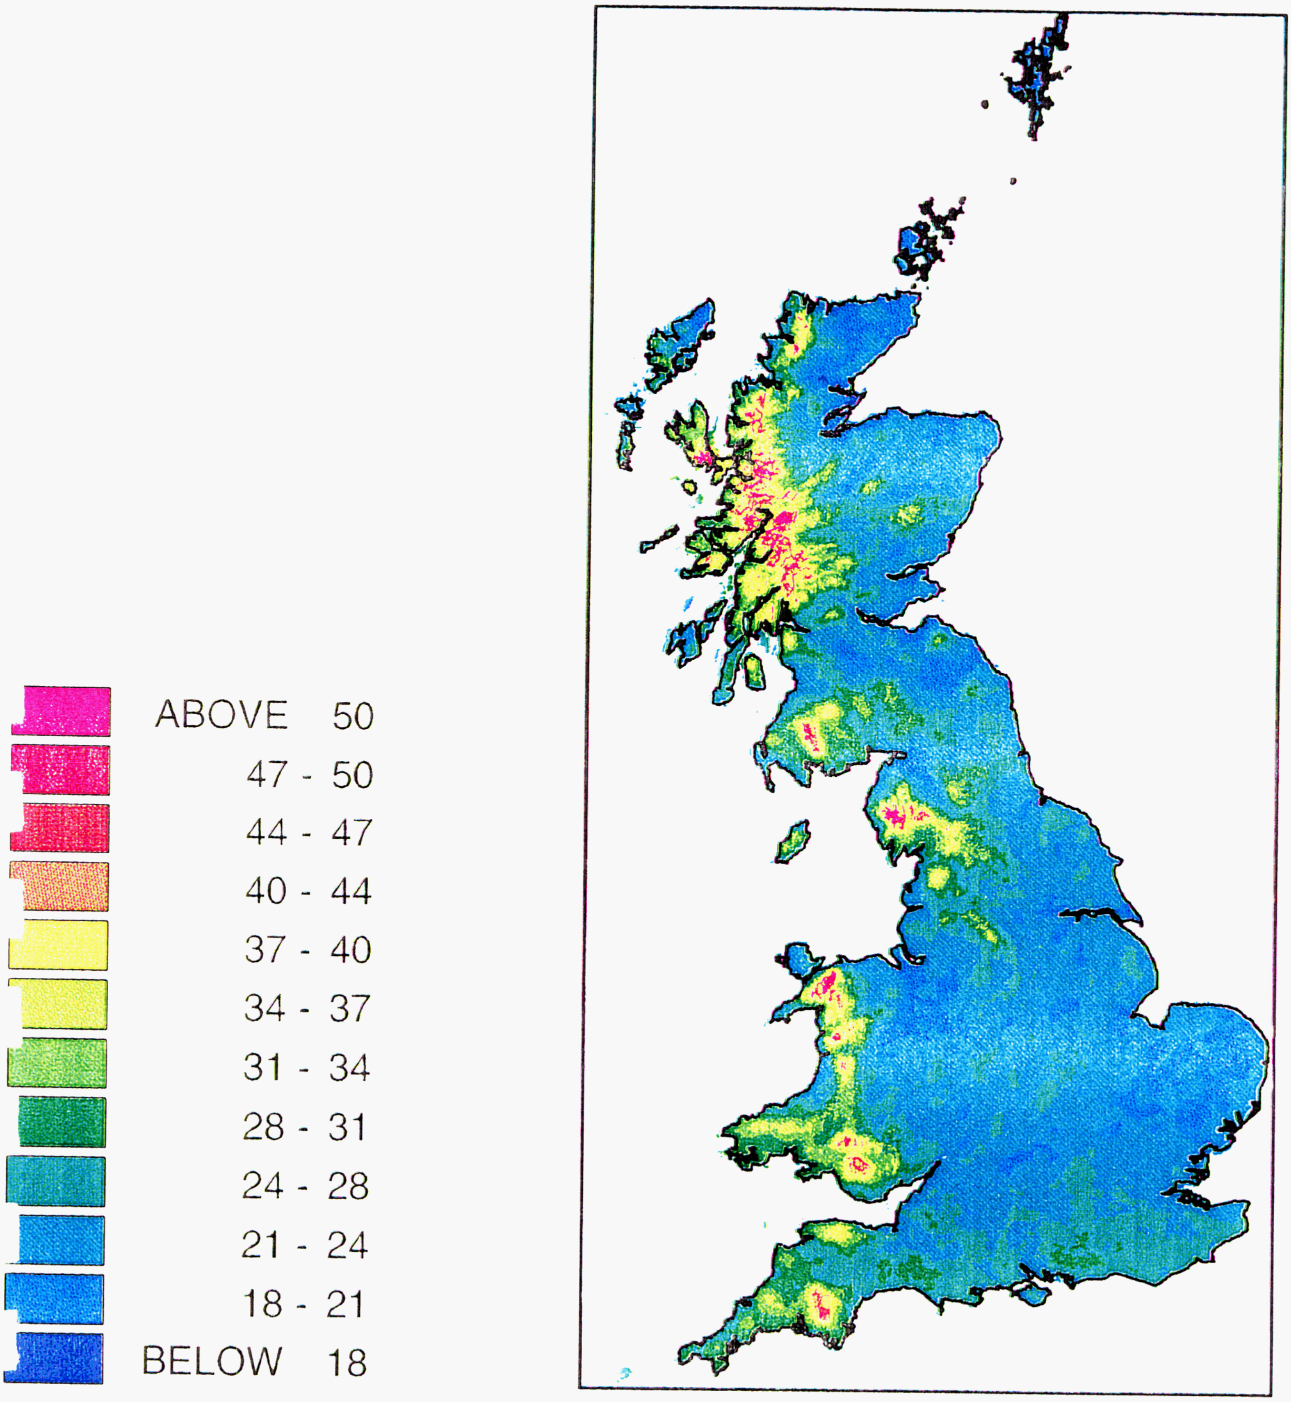

Figure 8.5

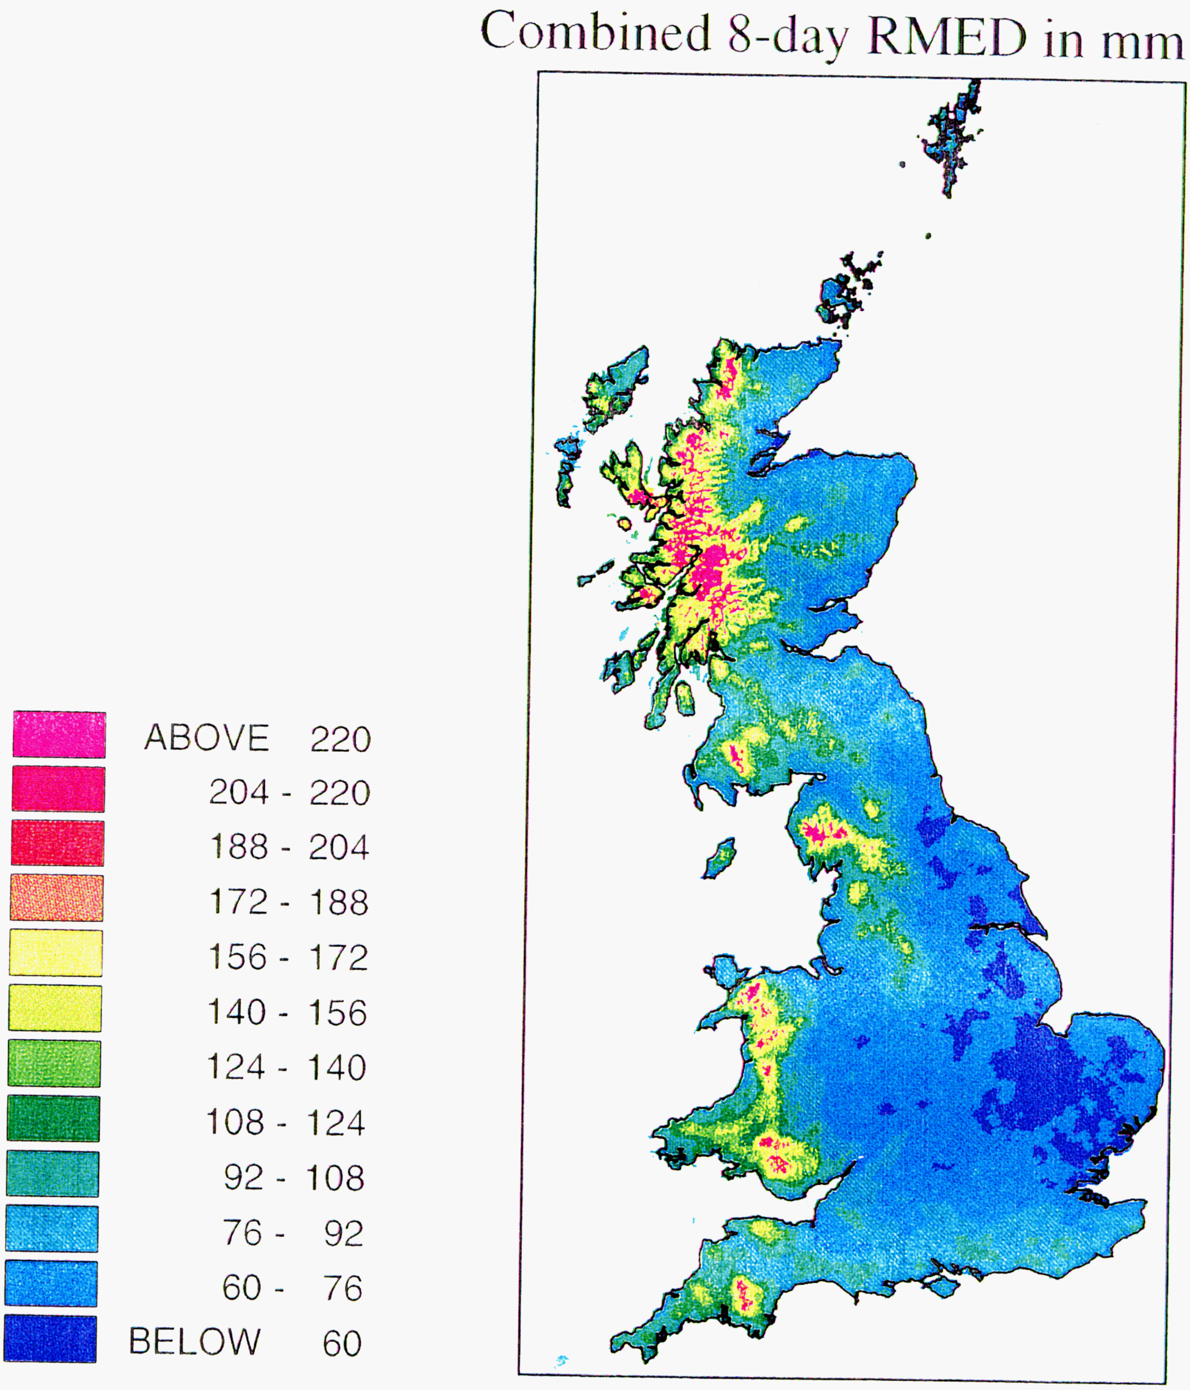
Combined 8-day RMED in mm
ABOVE 220
204 - 220
188 - 204
172 - 188
156 - 172
140 - 156
124 - 140
108 - 124
92 - 108
76 - 92
60 - 76
BELOW 60

Figure 8.6

9. DESIGN RAINFALL MAPS AND FSR COMPARISONS

9.1 Introduction

This chapter discusses the results produced by combining the maps of growth rates (Chapter 5) and RMED (Chapter 8). Maps of design rainfall with durations 1 hour and 1 day for various return periods are commented on and compared with the equivalent results from the Flood Studies Report rainfall analysis.

9.2 Duration one hour

Maps of 1-hour design rainfall for return periods 20 and 100 years are shown in Figures 9.1 and 9.2. The layout of the two maps is similar, with the largest rainfalls in mountainous regions and in parts of England where growth rates are high such as London, Lancashire and the East Midlands (see Figures 5.5 and 5.6). The smallest rainfalls are in eastern Scotland, and the smallest estimates for England and Wales are in Hampshire and the Welsh borders.

Figures 9.3 and 9.4 are maps of the FSR results for the same return periods. These are produced by running the automated FSR rainfall calculations for every 10-km grid square, using digitized versions of the key FSR rainfall statistics. The software used in the calculations is the same as that incorporated in the Micro-FSR package. The FSR estimates obviously vary less across the country. There is little evidence of the patches of high rainfall in parts of England which exist in the new results. The highest values are in upland areas such as the Lake District, the Pennines and Snowdonia.

To aid comparisons, Figure 9.5 shows the ratio of the FORGEX estimates to the FSR estimates for 1-hour 100-year rainfall. Over most of England and Wales, the new results are 10 to 30% greater. The difference rises to over 30% in Greater Manchester, Lancashire, parts of the East Midlands, Essex and Berkshire. Further west and south in England and Wales, there are areas where the new estimates are similar to the FSR ones. The only large area where the new results are lower is eastern Scotland.

The principal reason for the differences with the FSR results is the over-generalization of the FSR growth rates for 1-hour rainfall. For example, across all of England and Wales, the FSR 20-year growth rate varies by less than 5%, whereas the FORGEX results vary by nearly 50%.

9.3 Duration one day

Maps of 1-day design rainfall for return periods 20, 100 and 1000 years are shown in Figures 9.6 to 9.8. Once again, the pattern of the results does not vary much with return period. The wettest areas are, as expected, in the western uplands, and the driest in East Anglia and parts of Shropshire, Staffordshire and, interestingly, northern Cumbria. Southern counties of England have rather higher estimates of the 1000-year rainfall than elsewhere in lowland England.

Note that the dramatic-looking patches of high growth rates over parts of England (Figures 5.8 to 5.10) have relatively little influence on the maps of design rainfall, because the proportional variation in RMED across the country is significantly greater.

The equivalent FSR results are shown in Figures 9.9 to 9.11. They vary less across lowland England, and do not reach such high peaks in upland areas. At return periods of 20 and 100 years, the new results are close to the FSR estimates for much of England and Wales. Figure 9.12 aids the comparison, showing the ratio between the results for the 100-year return period. There appears to be a band running from Somerset to Lincolnshire where the new results significantly exceed the FSR ones. The difference exceeds 30% in small areas of SW England (particularly Somerset, as illustrated for Bridgwater in Figure 5.4), Lincolnshire, Kent, and the highest parts of most mountainous areas. There may of course be other localised examples of such exceedances which have not been resolved by the 10-km grid.

There are several grid points where the new 100-year estimates are significantly smaller, for example in Hampshire, the Welsh borders, East Anglia, NW England and central Scotland. In a few cases, the difference is as large as 25%. For the 20-year return period the FORGEX results are smaller than the FSR ones over rather larger areas of central southern England, the West Midlands, mid-Wales and NW England. The difference exceeds 30% in some locations.

For the 1000-year return period, FORGEX gives greater design rainfalls than the FSR over much of England and Wales, more than 50% greater for a few grid points in mountainous areas. However, the methods give similar results in the West Midlands, parts of southern England and the far north of England. There are several grid points in these areas where the new estimates are slightly smaller.

9.4 Discussion of FEH/FSR comparisons

The rainfall growth factors in the FSR appear to be over-general, masking important local and regional variations in rainfall growth and, hence, rainfall frequency. The new method takes more account of local data, both in constructing the focused growth curves, and in mapping the standardizing variable, RMED.

The results presented in this chapter suggest that continued use of the FSR rainfall frequency method for hydrological design will lead to under-design for some catchments and over-design in others. The errors in flood estimates are not likely to be quite as large as the maximum point differences in design rainfalls, because design flood estimates are based on catchment-average rainfall statistics. Underestimation of floods using the FSR rainfall analysis is likely for many small or urban catchments where the characteristic rainfall duration is closer to 1 hour than 1 day. For catchments which respond more slowly to rainfall, under-design is most likely in those areas where the FSR underestimates 1-day rainfall, i.e. in Somerset, the East Midlands and around the Thames estuary. Over-design is likely in several locations, for example parts of NW England, particularly for short to moderate return periods.

Although the new estimates of design rainfalls are significantly higher in upland locations, they can vary a great deal over short distances in mountainous areas because of the greater influence of topography in mapping RMED. Thus the effects on flood estimation for catchments including mountainous areas are less easy to anticipate.

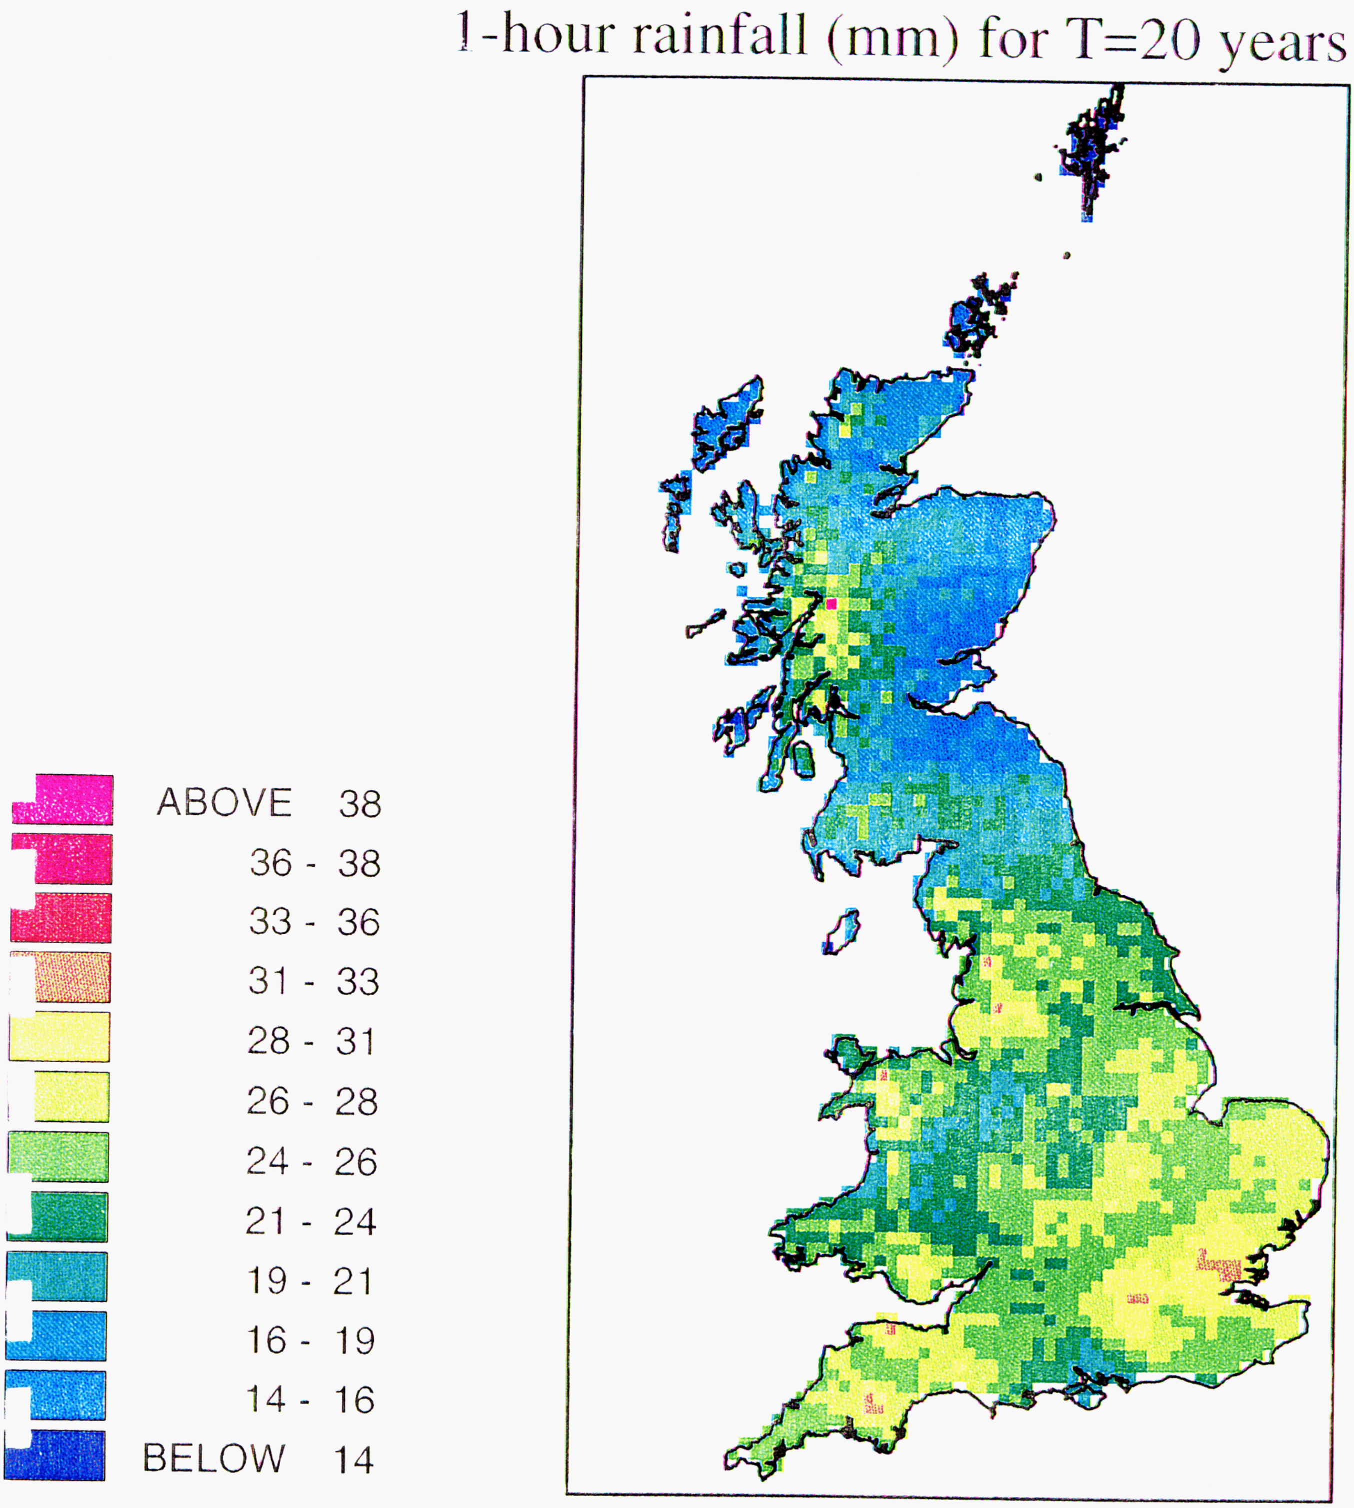
1-hour rainfall (mm) for T=20 years
ABOVE 38
36 - 38
33 - 36
31 - 33
28 - 31
26 - 28
24 - 26
21 - 24
19 - 21
16 - 19
14 - 16
BELOW 14

Figure 9.1

1-hour rainfall (mm) for T=100 years

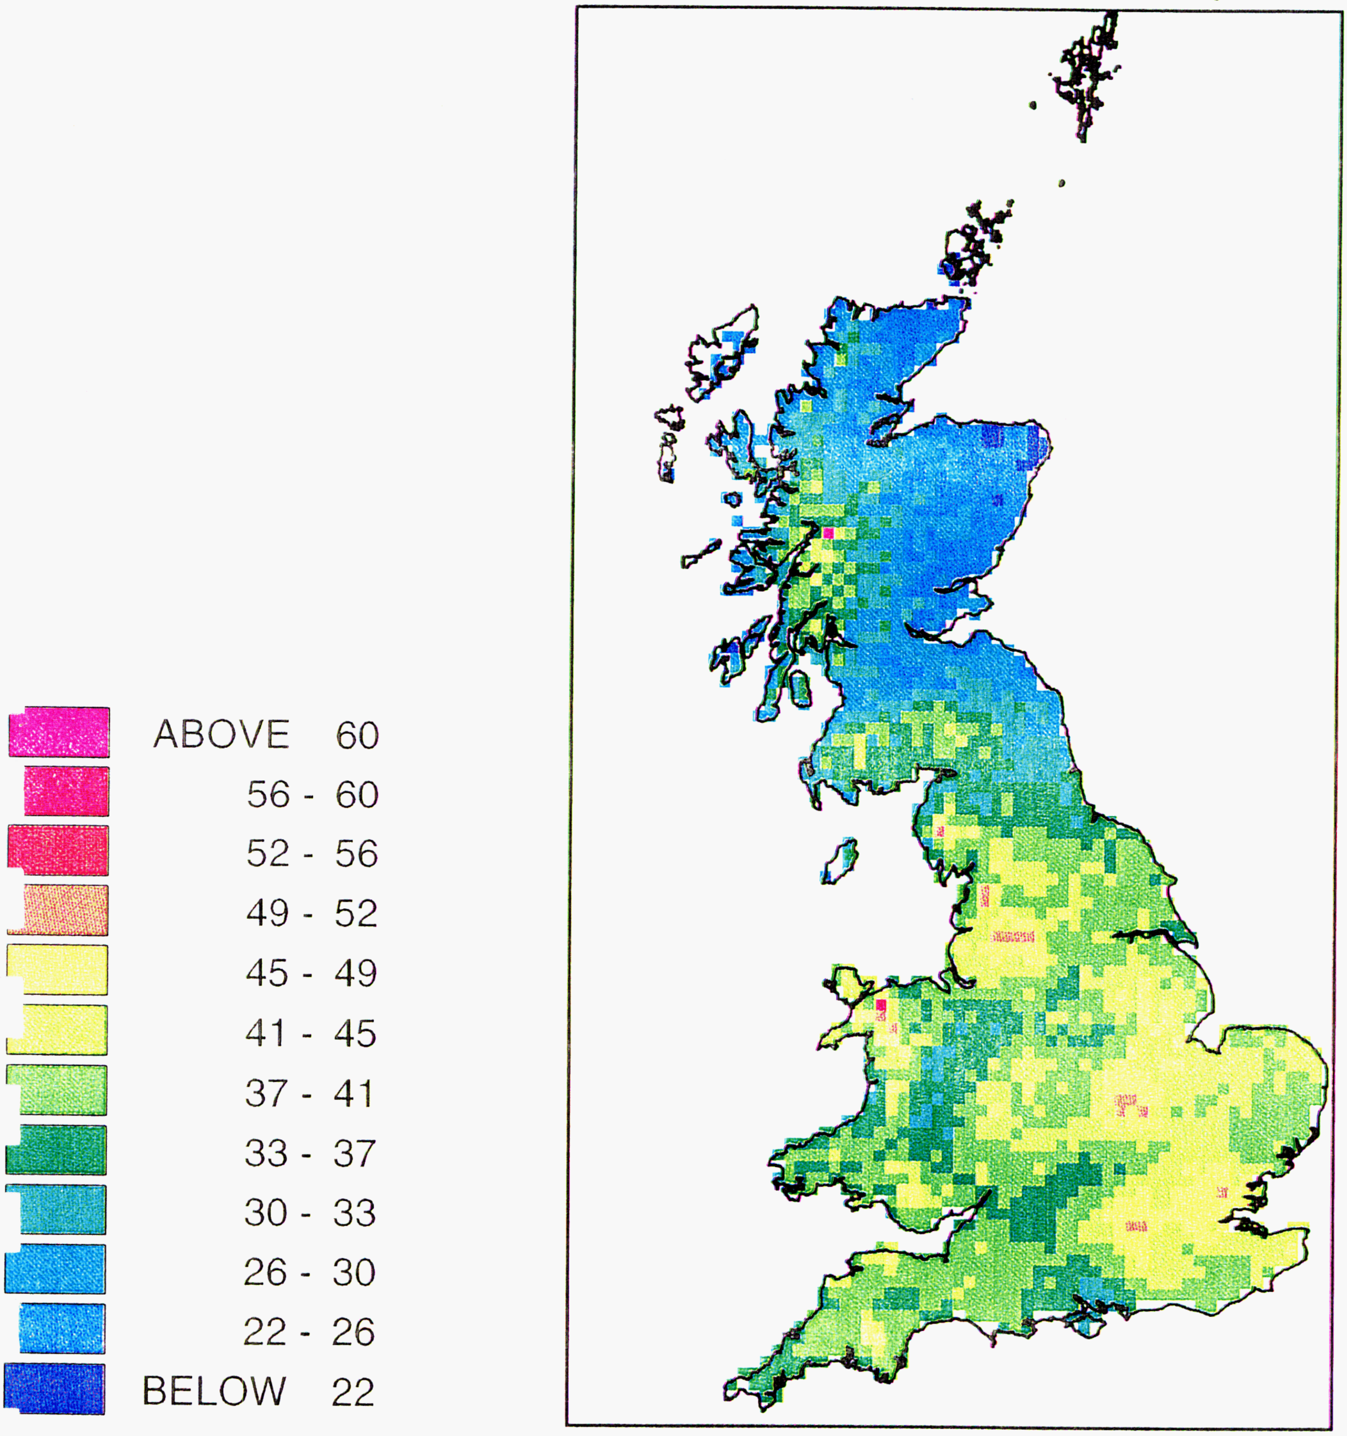

Figure 9.2

FSR 1-hour rainfall (mm) for T=20 years

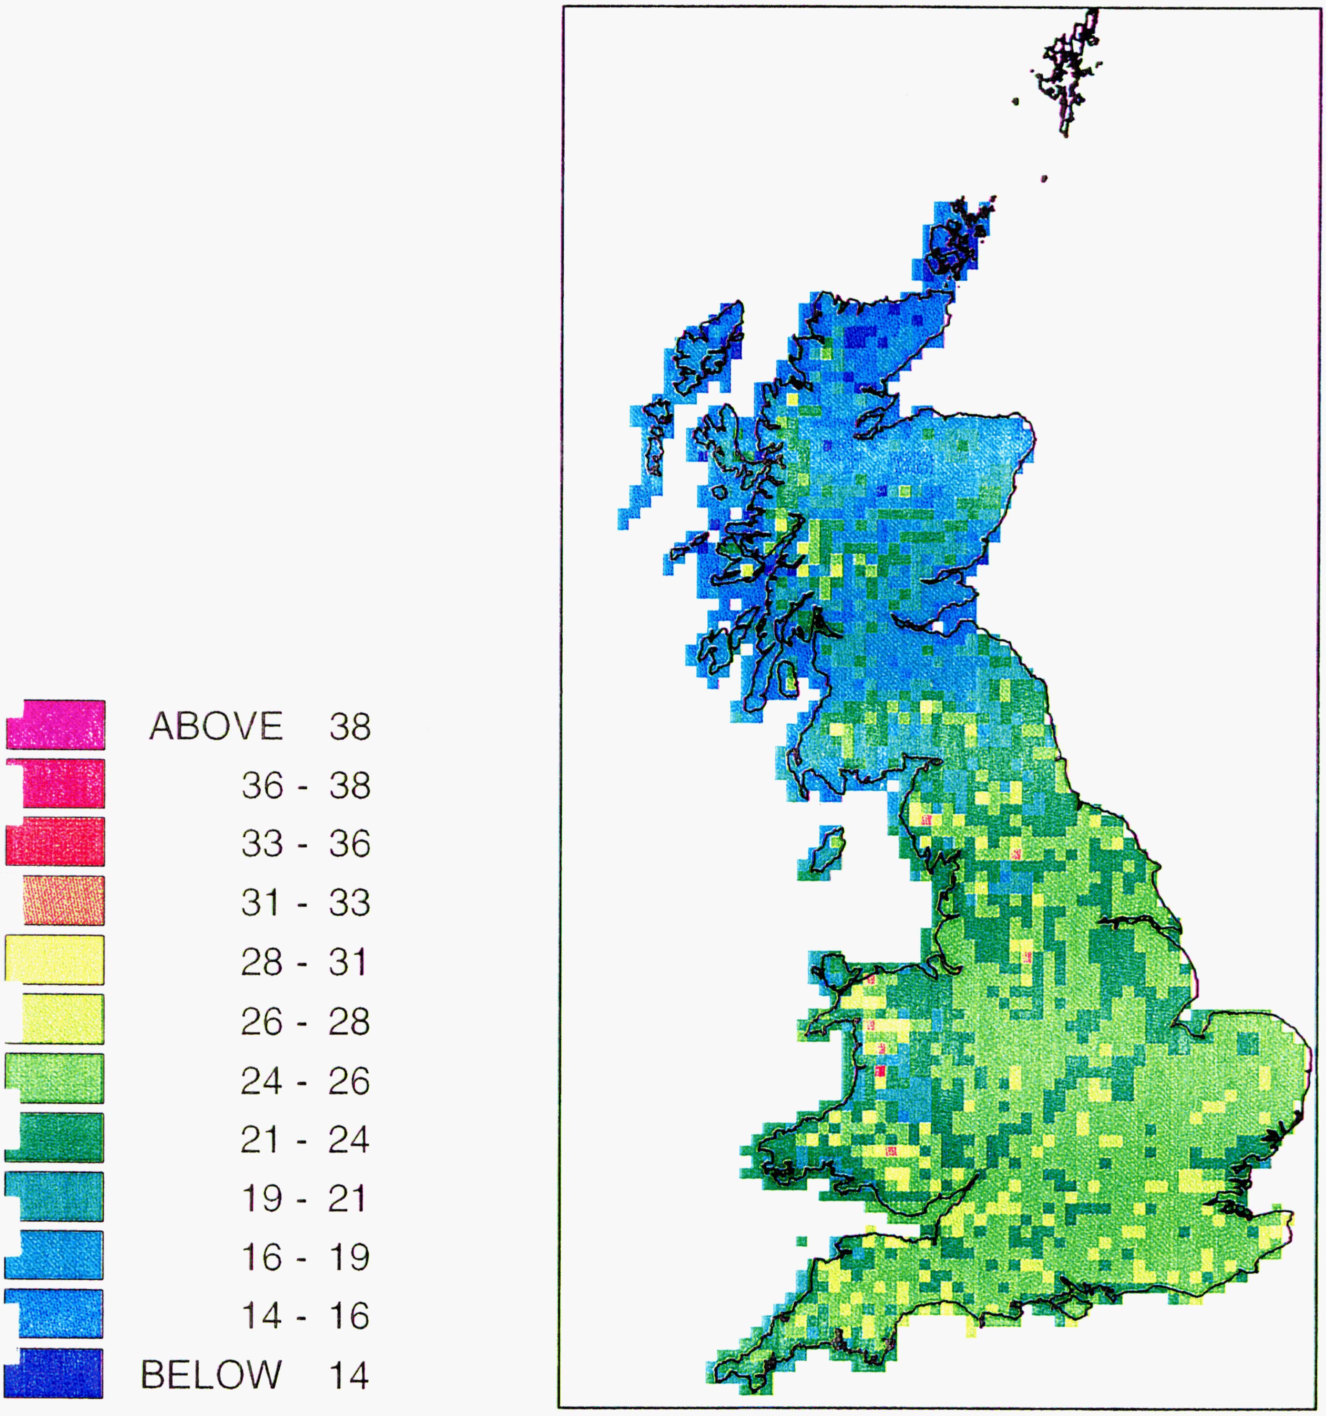

Figure 9.3

FSR 1-hour rainfall (mm) for T=100 years

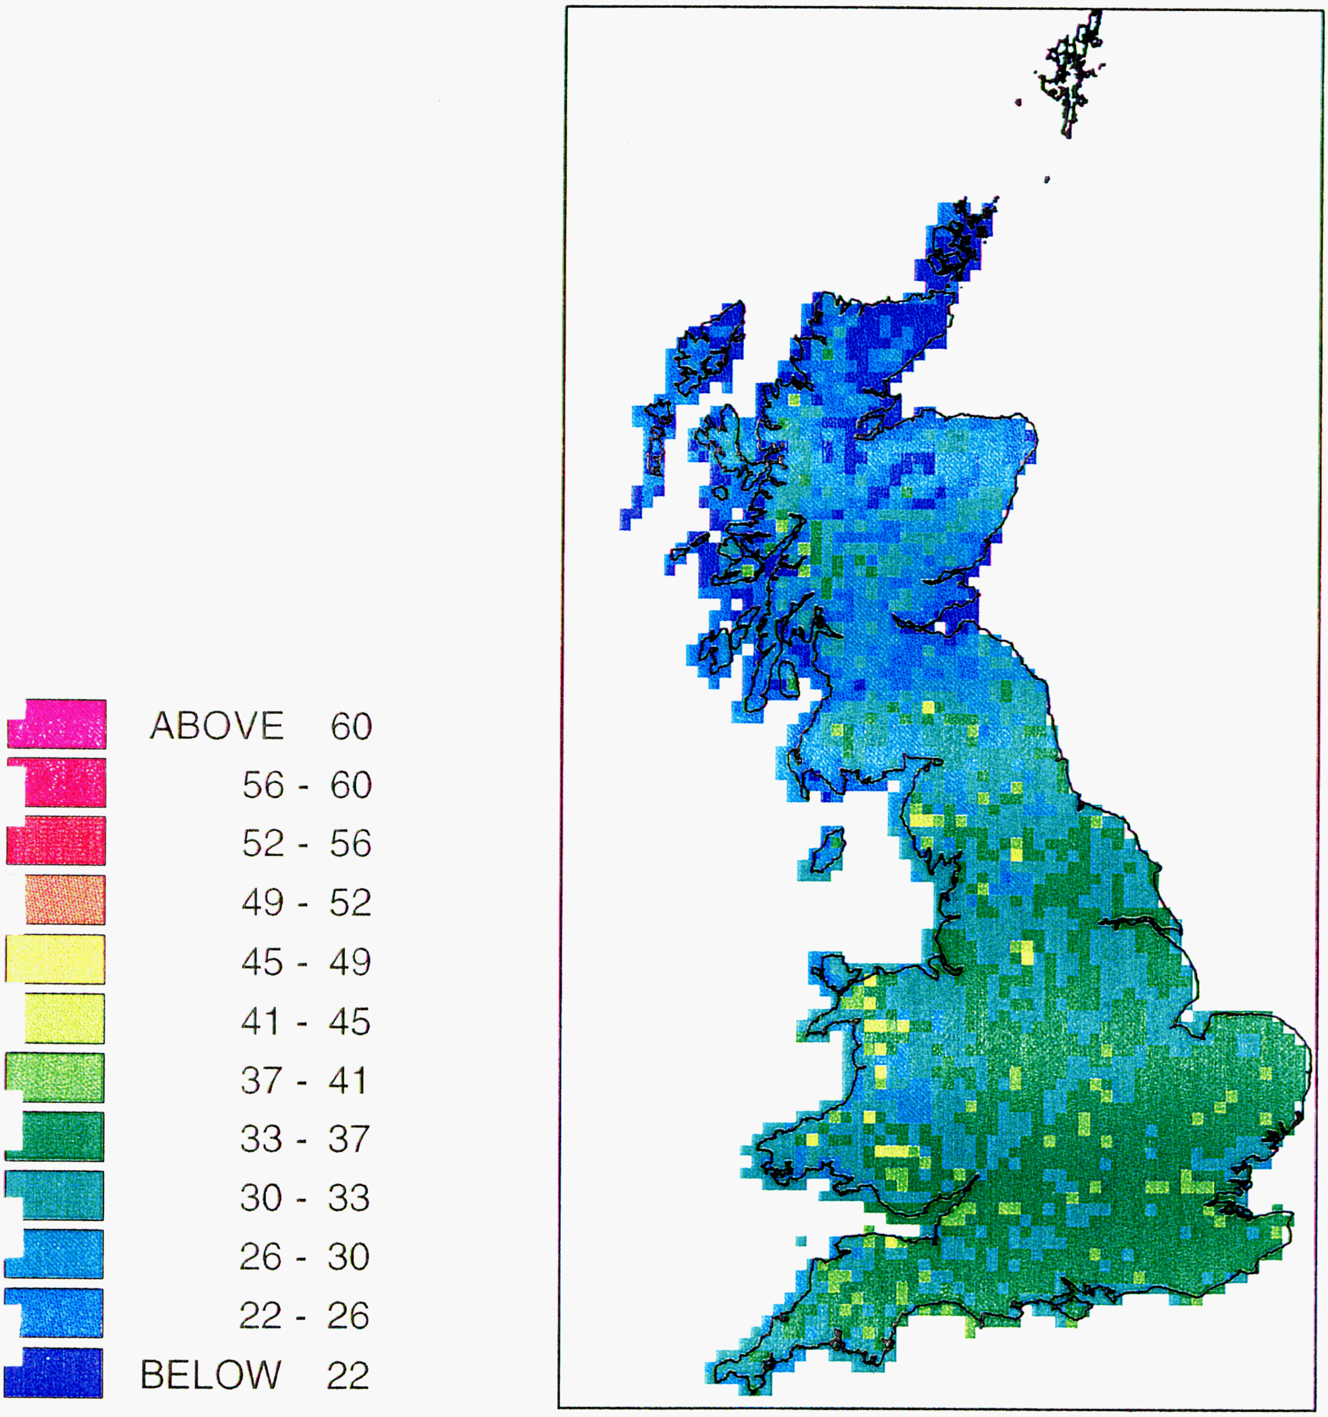

Figure 9.4

Ratio of FORGEX to FSR estimate of 1-hour 100-year rainfall

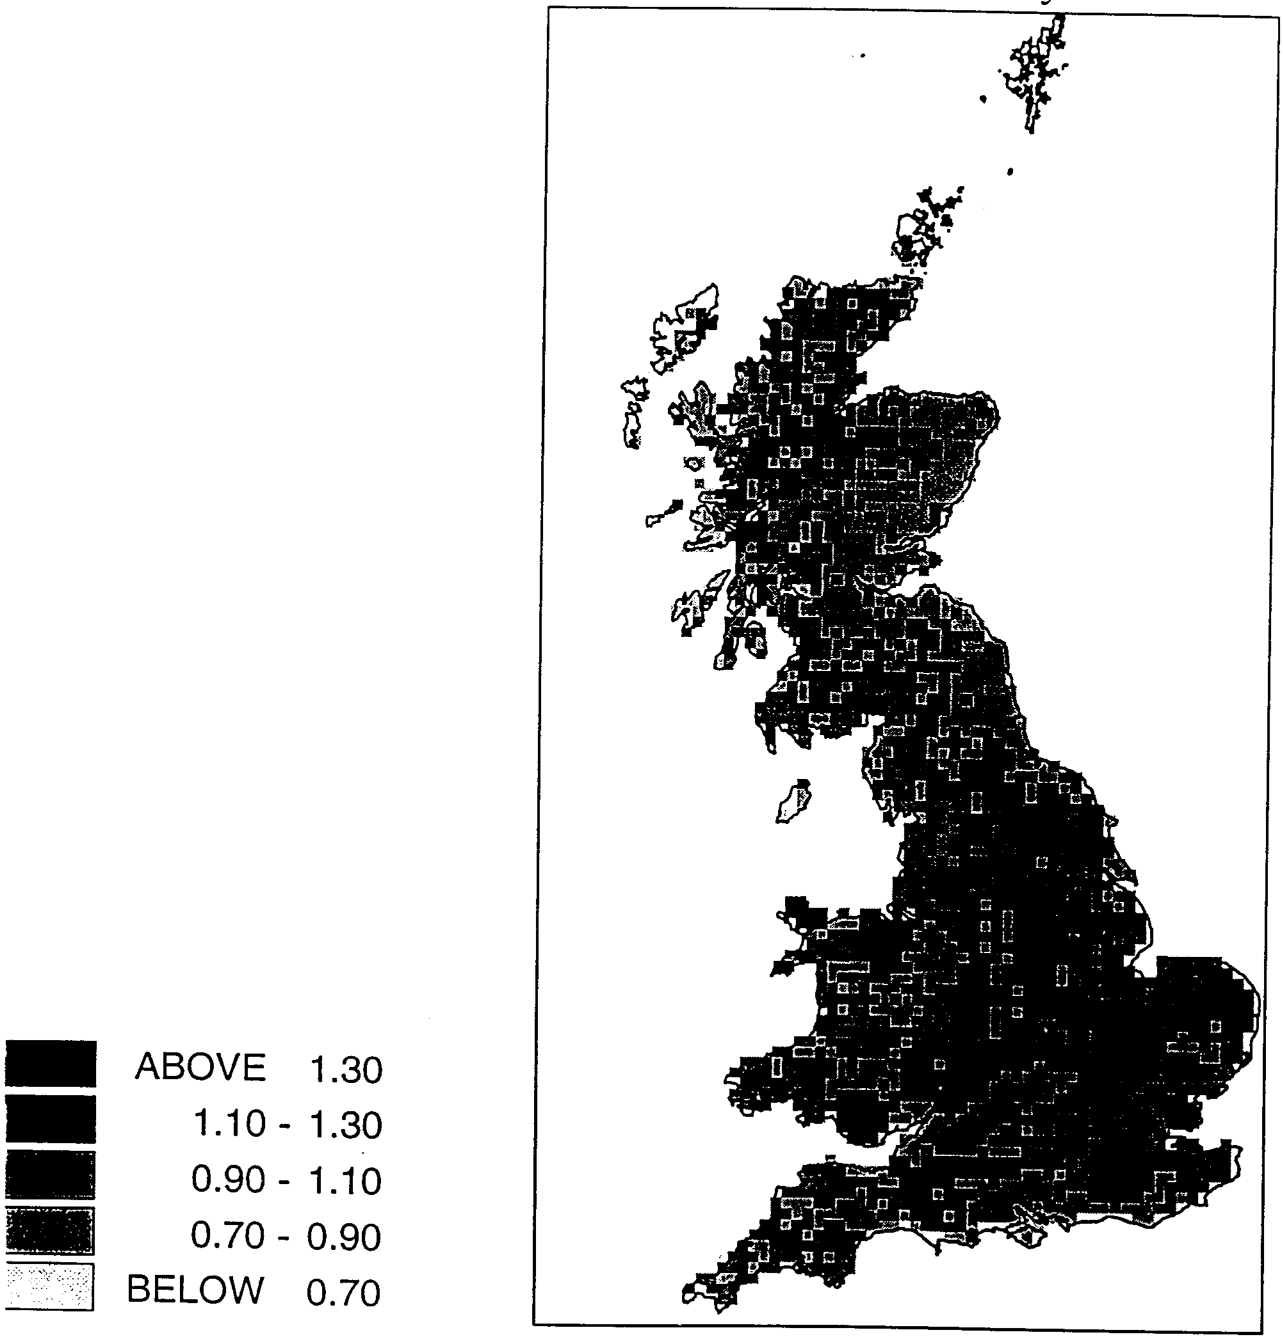

Figure 9.5 **Ratio of FORGEX to FSR estimate of 1-hour 100-year rainfall**

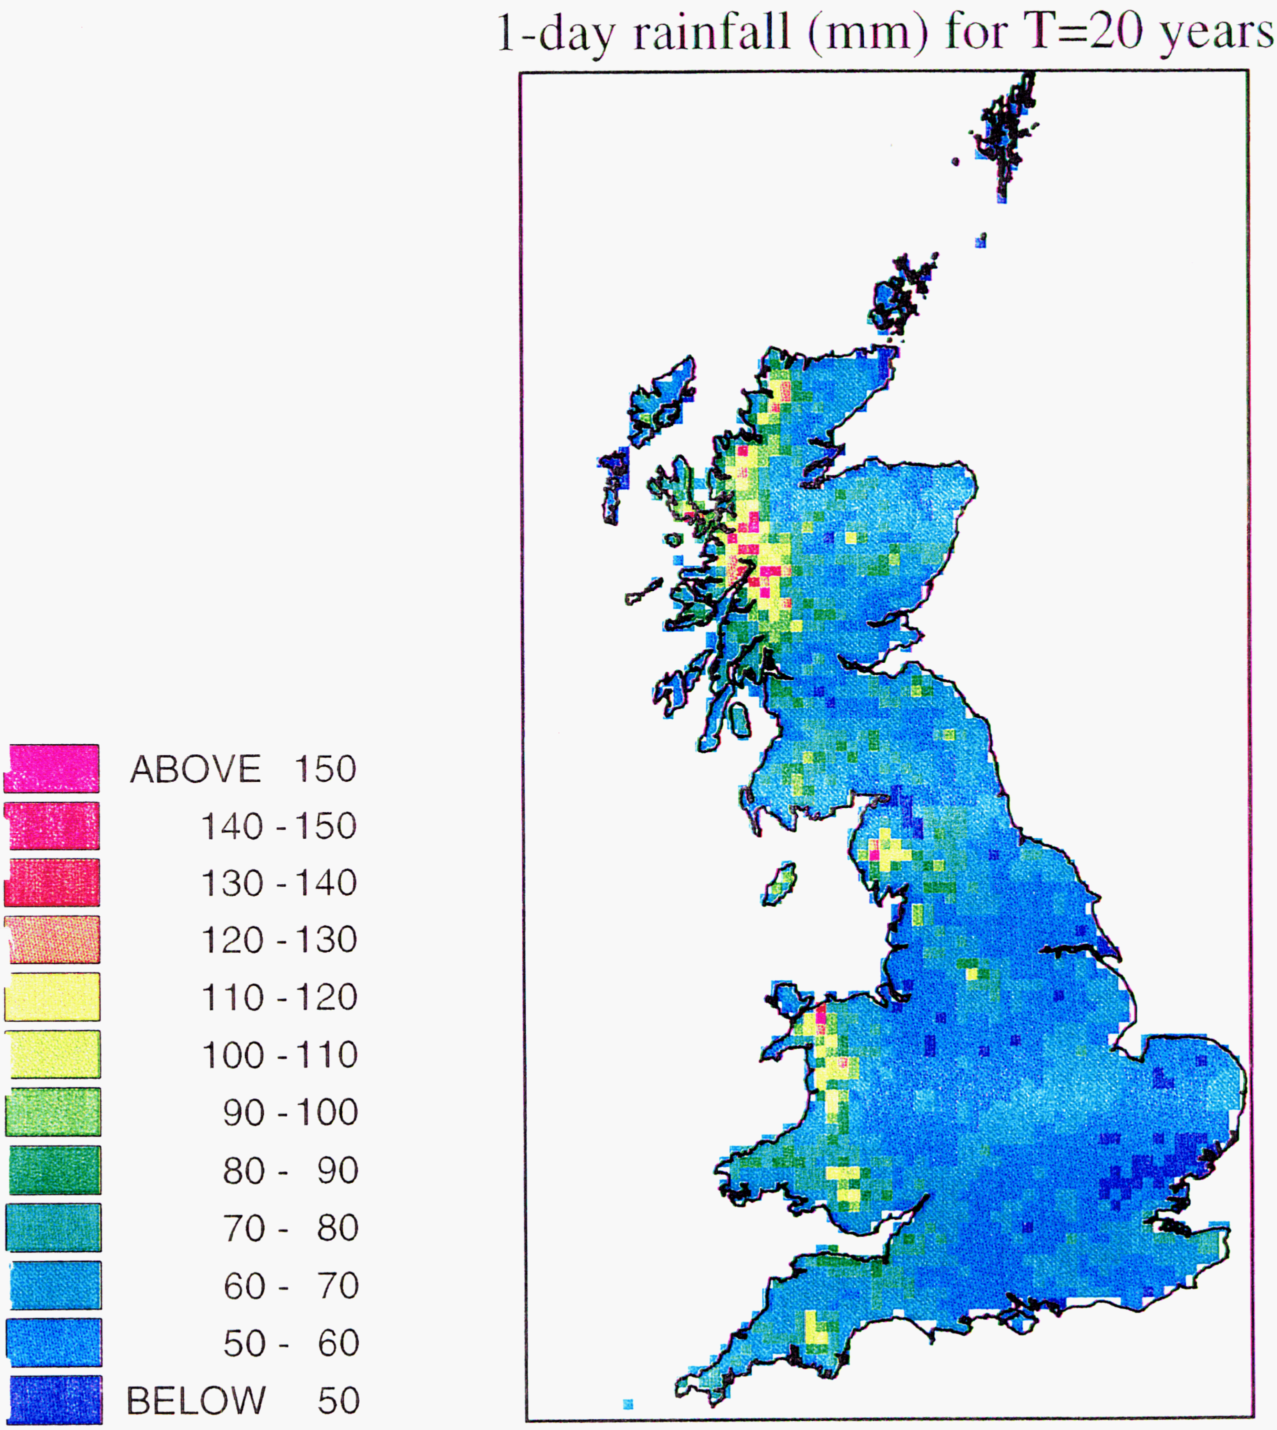
1-day rainfall (mm) for T=20 years
ABOVE 150
140 - 150
130 - 140
120 - 130
110 - 120
100 - 110
90 - 100
80 - 90
70 - 80
60 - 70
50 - 60
BELOW 50

Figure 9.6

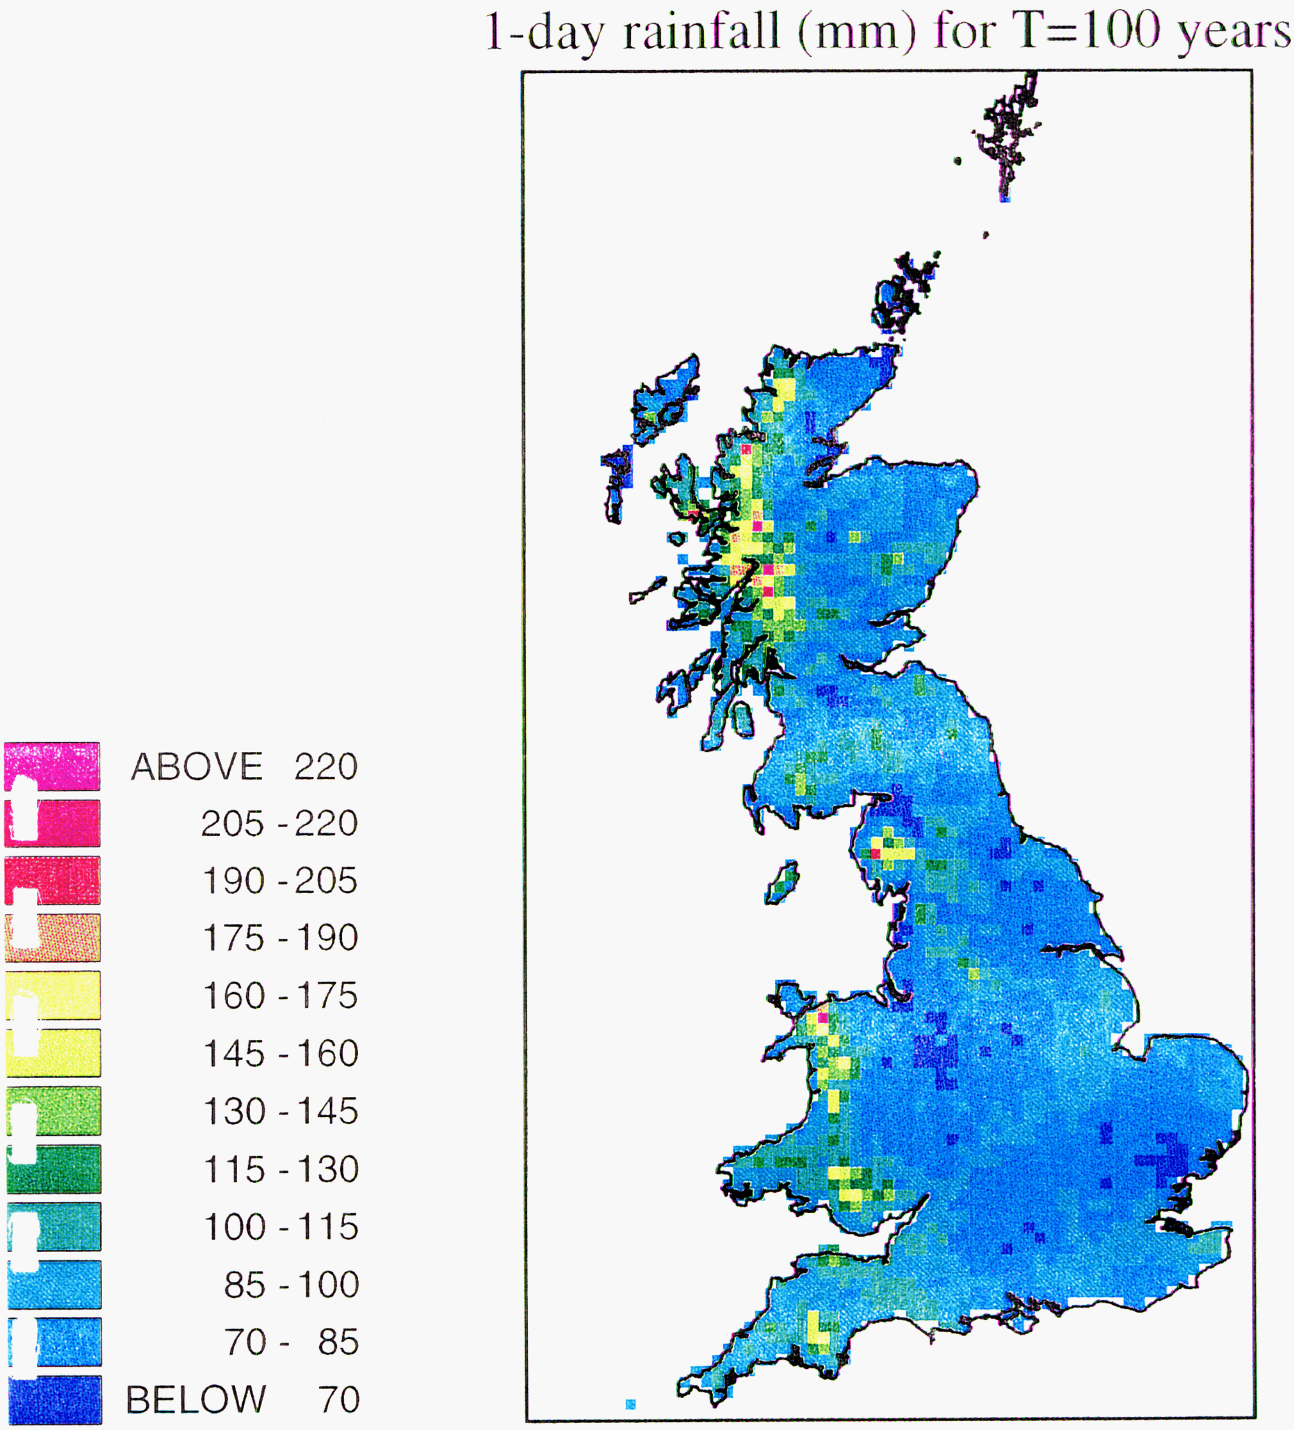
1-day rainfall (mm) for T=100 years
ABOVE 220
205 - 220
190 - 205
175 - 190
160 - 175
145 - 160
130 - 145
115 - 130
100 - 115
85 - 100
70 - 85
BELOW 70

Figure 9.7

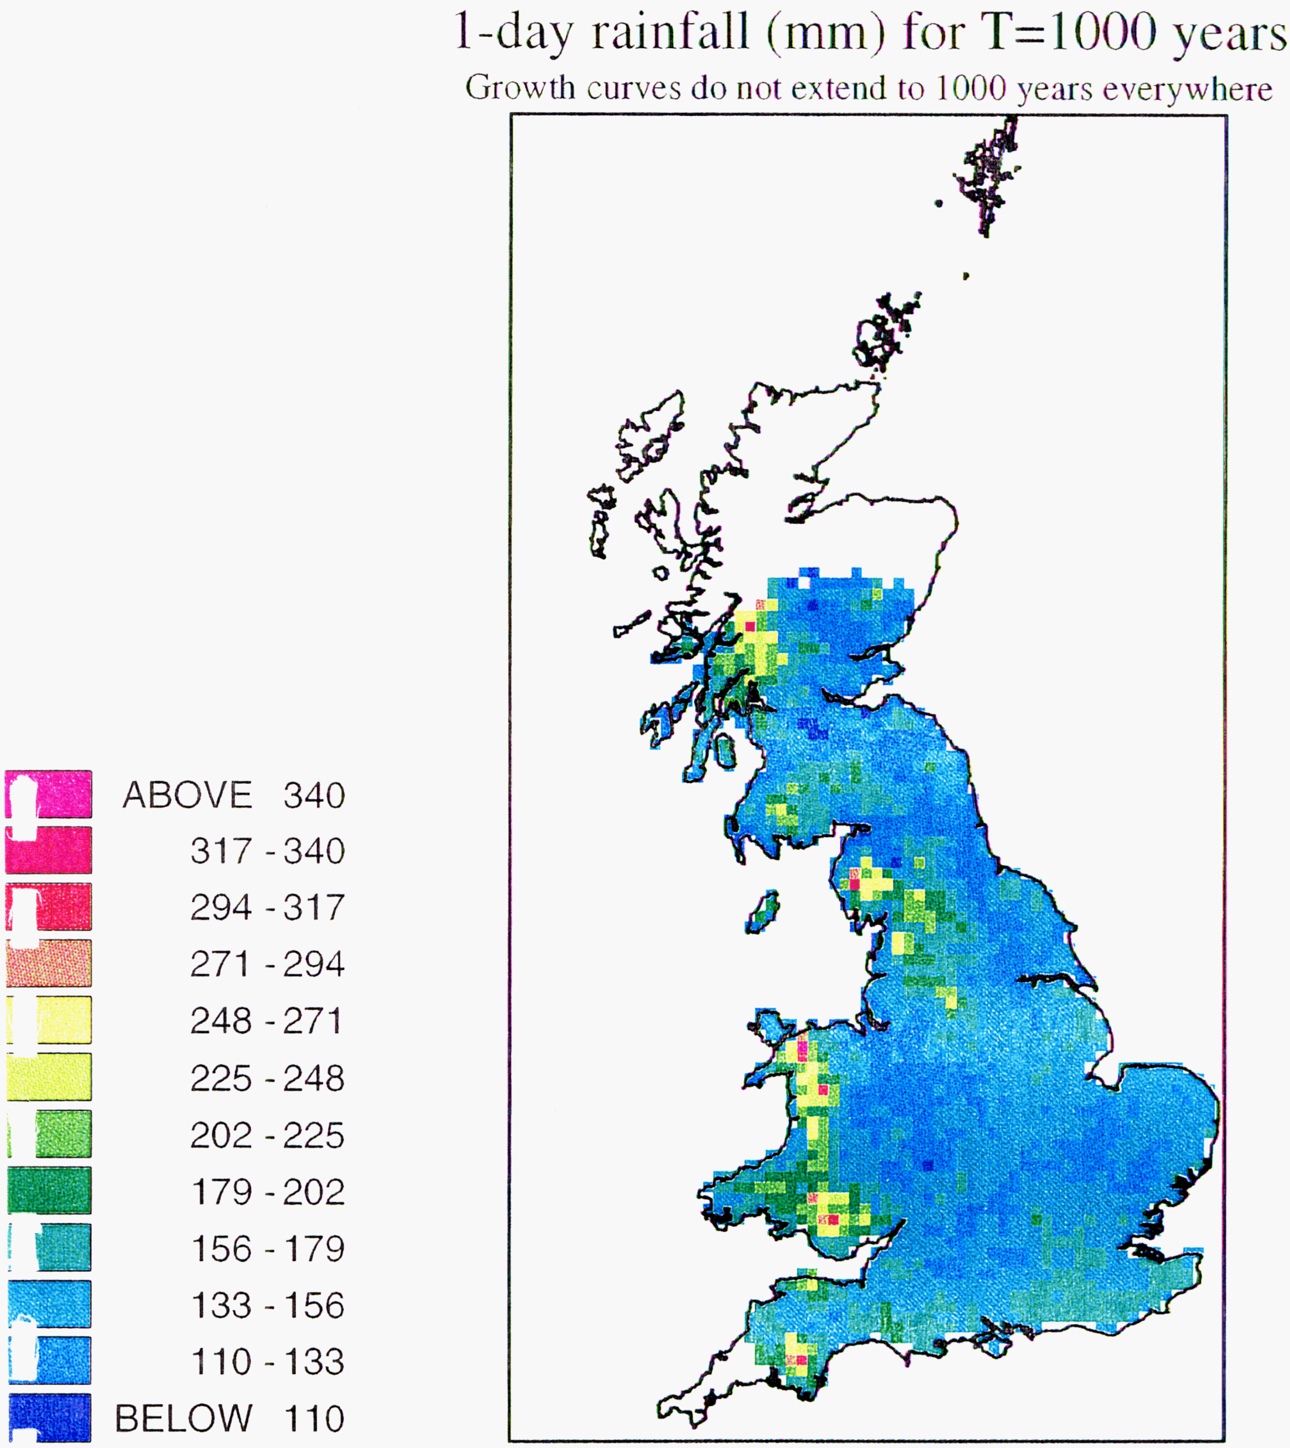
1-day rainfall (mm) for T=1000 years
Growth curves do not extend to 1000 years everywhere
ABOVE 340
317 - 340
294 - 317
271 - 294
248 - 271
225 - 248
202 - 225
179 - 202
156 - 179
133 - 156
110 - 133
BELOW 110

Figure 9.8

FSR 1-day rainfall (mm) for T=20 years

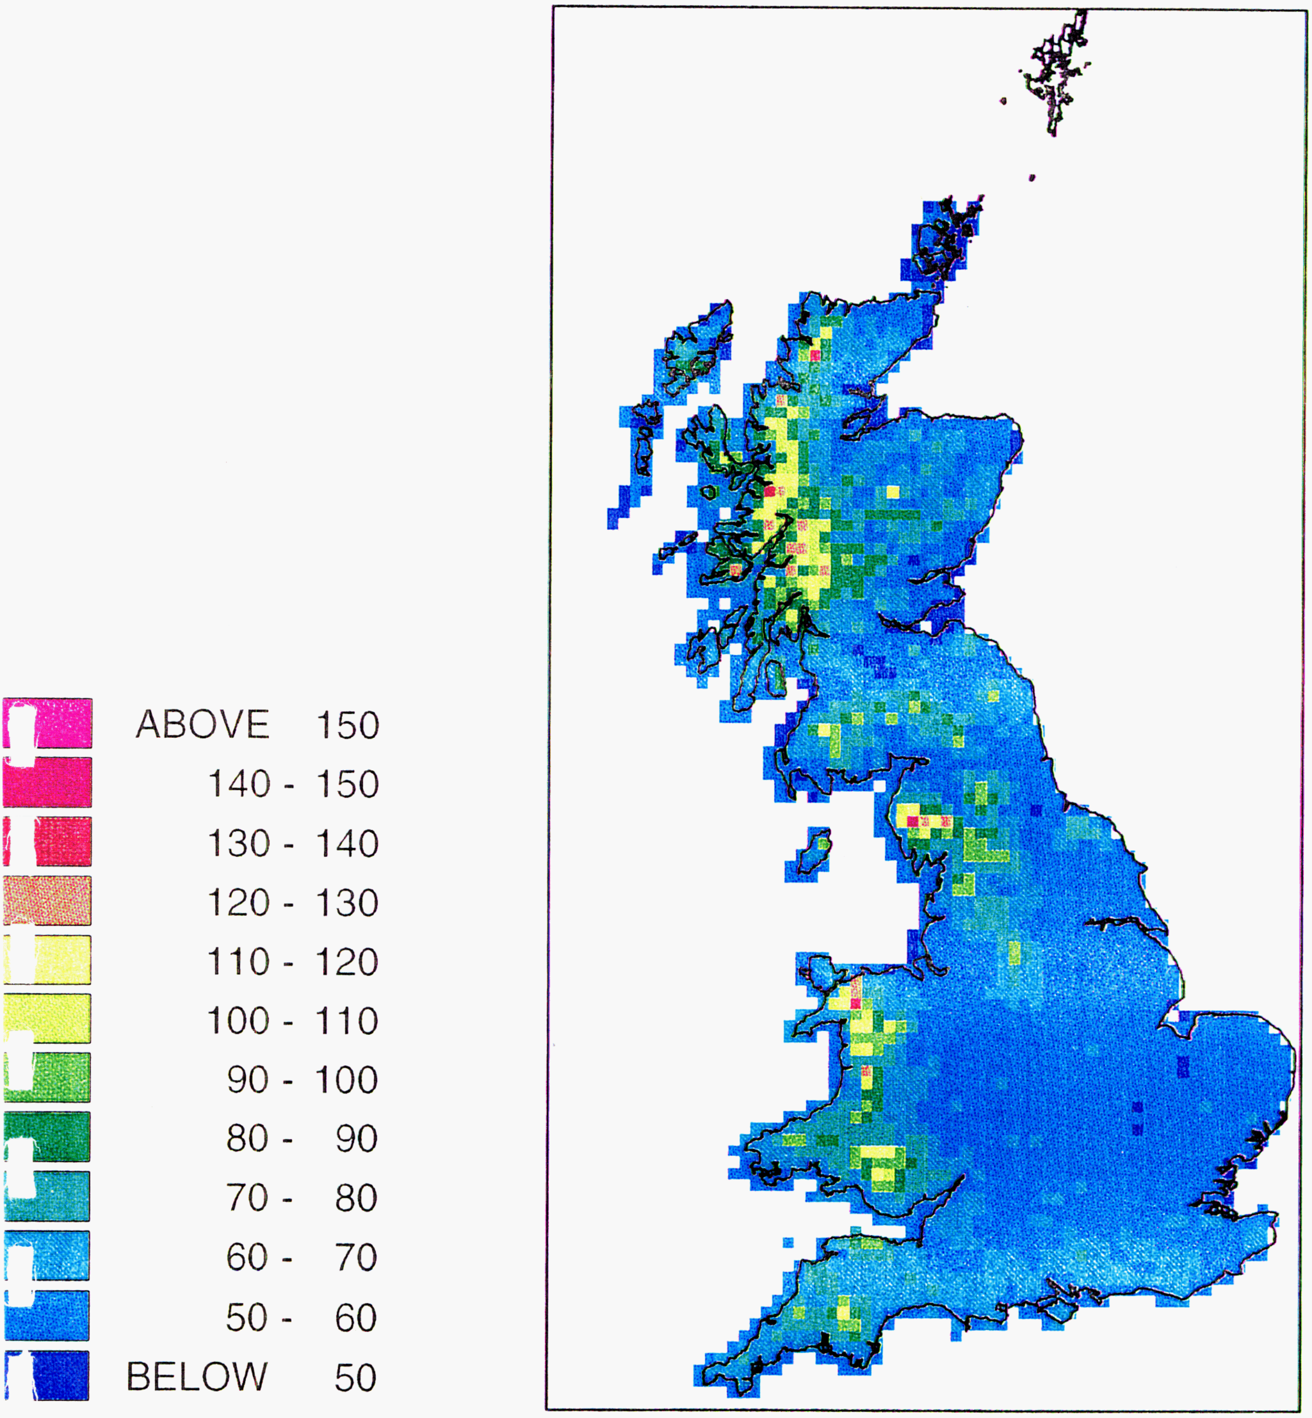

Figure 9.9

FSR 1-day rainfall (mm) for T=100 years

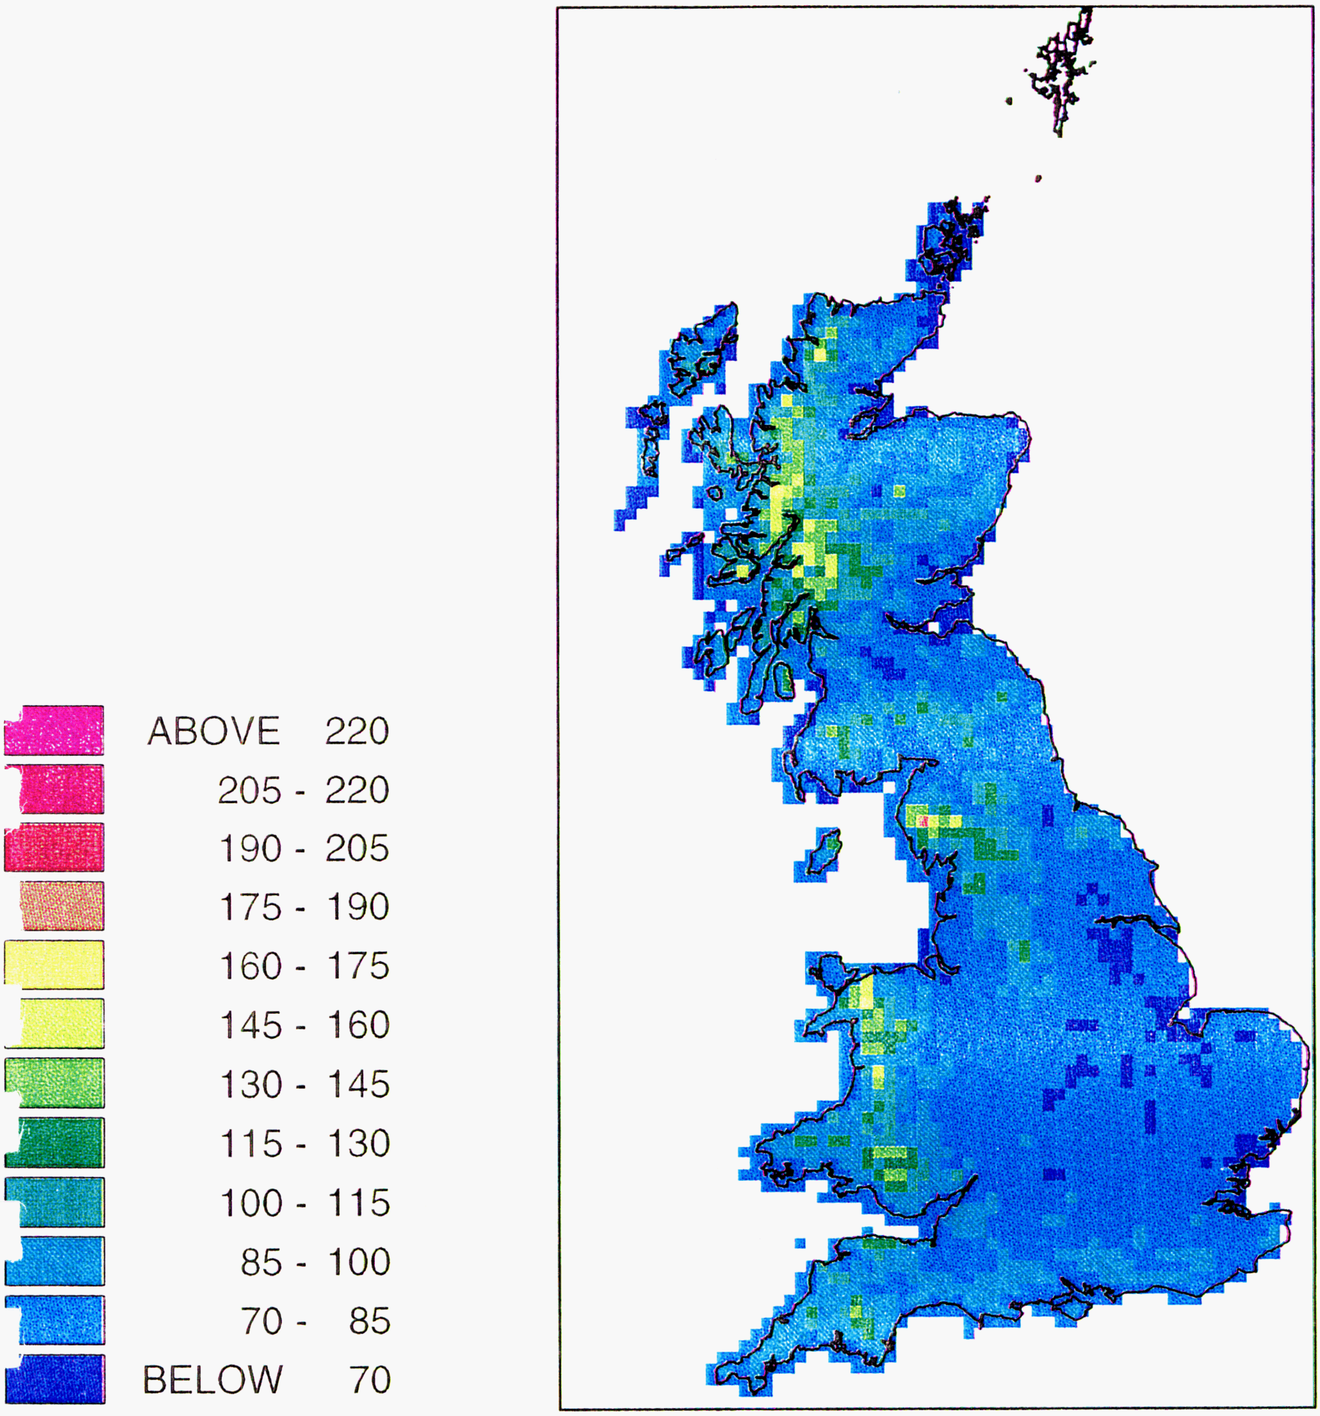

Figure 9.10

FSR 1-day rainfall (mm) for T=1000 years

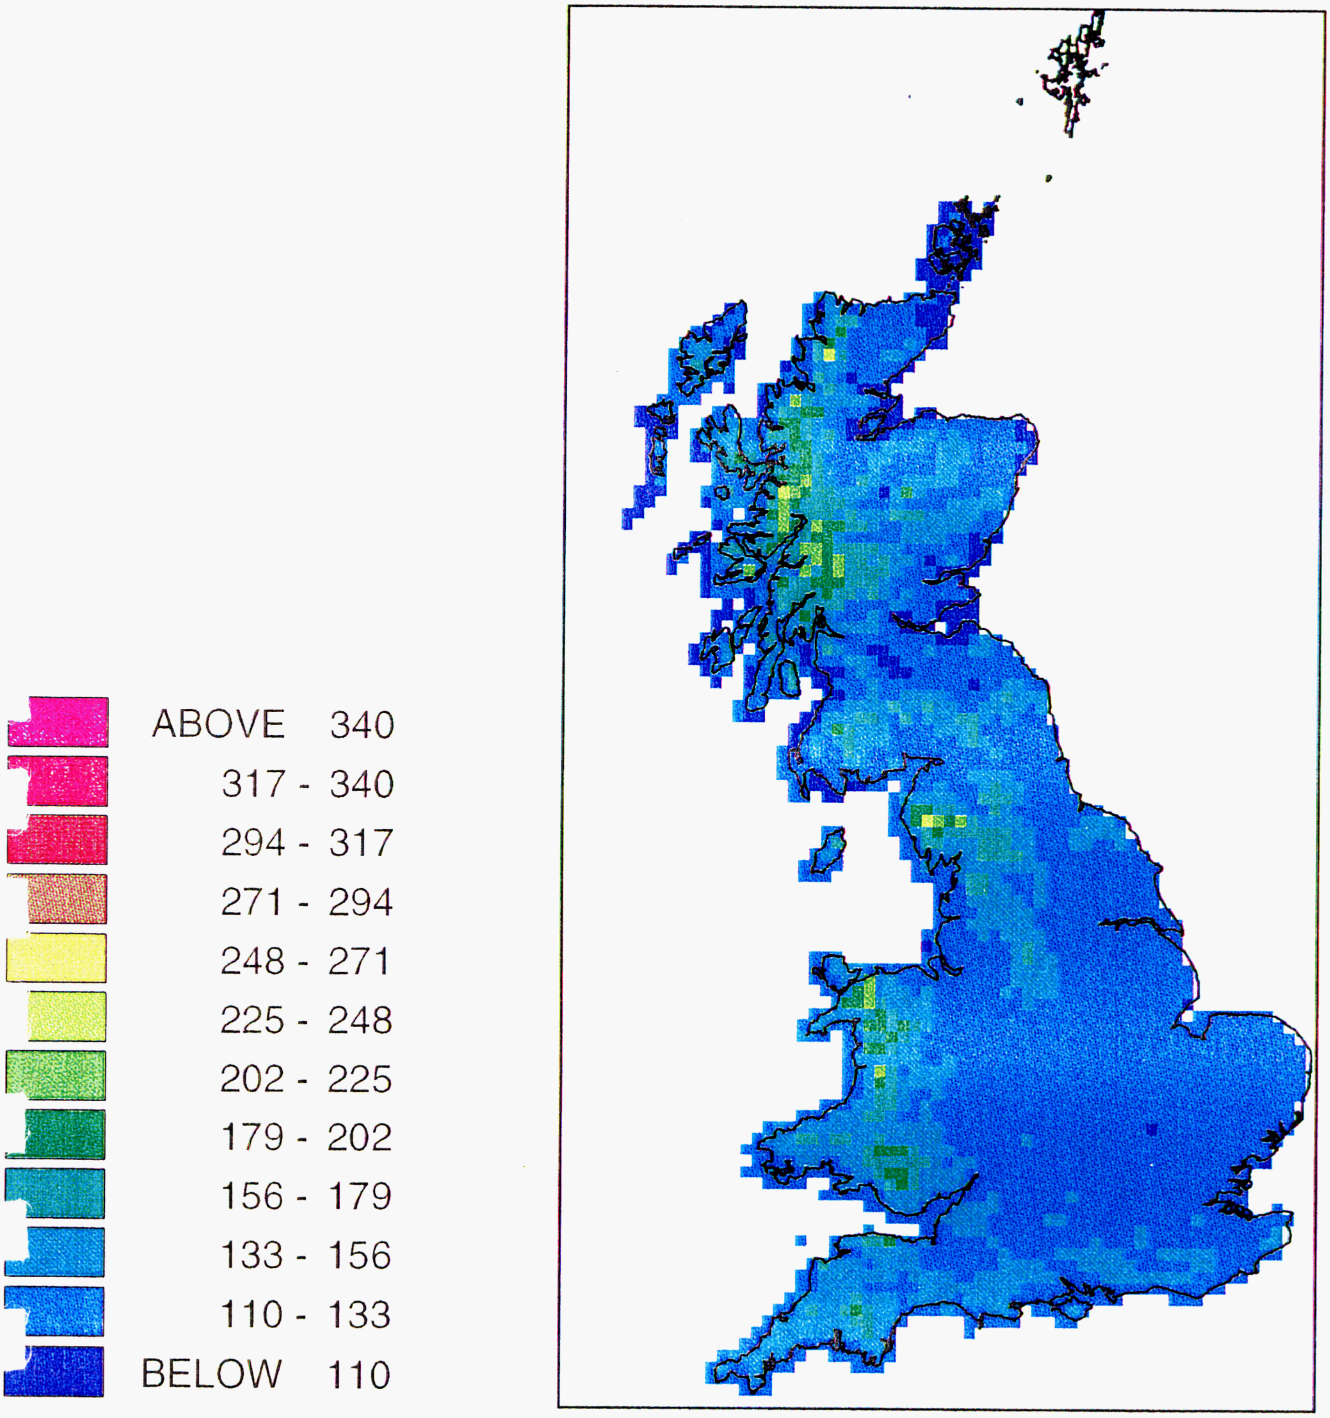

Figure 9.11

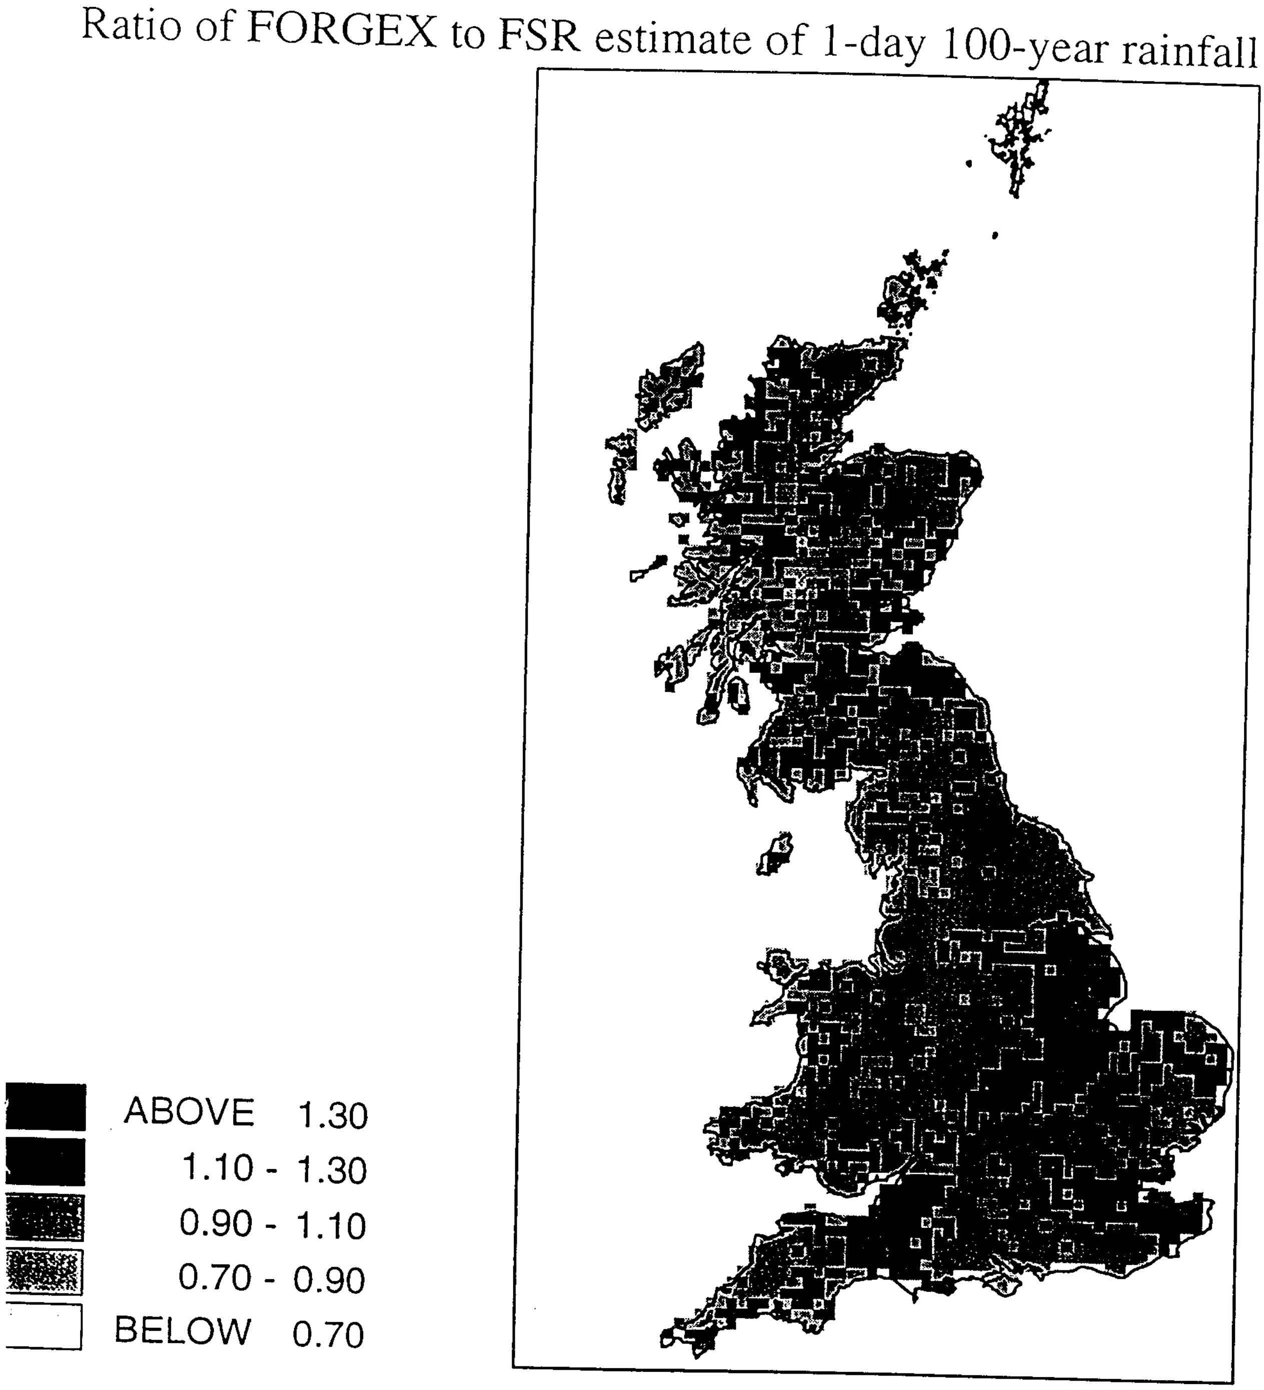

Figure 9.12 Ratio of FORGEX to FSR estimate of 1-day 100-year rainfall

10. DEPTH-DURATION-FREQUENCY MODELLING

10.1 Aims and approach

During the construction of growth curves, the mapping of RMED and the production of maps of design rainfall, each rainfall duration has been treated separately. Growth curves can be estimated for any duration with annual maxima available in the database, i.e. for 1, 2, 4, 6, 12, 18 and 24 hours and for 1, 2, 4 and 8 days. Maps of RMED have been produced for durations 1, 2, 6 and 12 hours and for 1, 2, 4 and 8 days. Other durations of RMED could be mapped if necessary. However, the final required output of the project is a method of estimating design rainfalls for any duration and return period. This chapter describes the incorporation of the results into a model linking rainfall depth, duration and frequency (a DDF model).

Another reason for requiring a DDF model is that rainfall estimates for different durations need to be reconciled. If each duration is treated separately, it is possible that contradictions between durations could exist, for example if the estimated 2-day 100-year rainfall were smaller than the 1-day 100-year rainfall at the same site. In addition, a model enables extrapolation of rainfall frequency estimates for durations shorter than 1 day, for which growth curves do not extend to return periods such as 1000 years. Incorporating information from longer durations can guide what might otherwise be an uncertain extrapolation.

The requirements for a DDF model are as follows:

- it should provide design rainfall estimates for any duration and return period;
- it should follow local and regional variations in design rainfalls where possible;
- it must avoid any contradictions between durations or return periods.

Note that a DDF model should not be expected to improve the accuracy of estimating rainfalls for the primary durations at which data are available. Although a model incorporates information from other durations, Buishand (1993) showed that the dependence between annual maxima over different durations is such that little or no improvement can be obtained.

This chapter describes the development of a suitable model and compares it with a model for M5 rainfall given in the FSR.

10.2 Developing a DDF model

10.2.1 Approaching the problem

It was decided to build a model which would be fitted separately at each site, possibly on a 1 km grid eventually. There are various ways of viewing the rainfall frequency results at a site. One example is shown in Figure 10.1, where growth curves for six durations are multiplied by the appropriate duration of RMED and plotted in the depth - frequency (i.e. return period) domain. The data are for Waddington in Lincolnshire. There is one line for each duration. Note the dependence between growth curves for different durations. Because there is already a continuous relationship between rainfall, R, and return period, T, one approach would be to work in this domain and fit a model to explain the variation with duration, D.

This approach was not taken, because of the large number of parameters which would have to be related to D. Each line on Figure 10.1 has seven or eight parameters, and, for the highest segments, the segment boundaries are at different return periods, which effectively adds further parameters. It would not be easy to relate all these parameters to D, and in particular to ensure that curves for different durations never crossed.

The chosen approach plots the results in the depth - duration domain, as shown in Figure 10.2 for Waddington. Rainfall frequency estimates are shown for selected return periods and durations. A log-log scale is used so that, if the depth-duration model is of the form $R = aD^b$, it will plot as a straight line for each return period. Power laws of this form have been used previously, for instance by Ferreri and Ferro (1990) and Reed and Stewart (1994). Chen (1983) proposed a model with an extra parameter, $R = aD(D+b)^{-c}$, but this is also linear on a log-log scale for all but the shortest durations, because $b \ll D$ for durations longer than half an hour. The earlier model developed for the FSR was essentially similar (see Section 10.3 below).

Note that these formulae are often expressed in terms of rainfall intensity rather than rainfall depth over a certain duration. The latter has been chosen for use in the present study because aggregated rainfall depth, rather than instantaneous intensity, is the quantity which is actually measured and which is of importance for flood estimation.

10.2.2 Adjusting for discretization

In reconciling rainfall estimates for different durations, adjustments are required to take account of discretization. For example, design rainfalls based on "fixed" 1-day annual maxima are not expected to be as large as those based on "sliding" 24-hour annual maxima, where the period of 24 hours can start at any hour. All design rainfalls were converted from "fixed" or "sliding" to "true" estimates, using standard factors taken from Dwyer and Reed (1995). The multiplication factors are shown in Table 10.1.

Table 10.1 Factors used to correct design rainfalls for discretization

Duration (days)	Multiply by
1	1.16
2	1.11
4	1.05
8	1.01

Duration (hours)	Multiply by
1	1.16
2	1.08
4	1.03
6	1.01
≥12	1.00

10.2.3 Building a model

The points in Figure 10.2 do not lie in straight lines for each return period, although fitting a straight line might be a first approximation. It was decided to fit a more flexible function, using data for sufficient durations and return periods to ensure that the model would represent the variations in growth curves for different durations that can be seen in Figure 10.1. The model is based on several concatenated straight line segments, which can have different slopes for different return periods. The final form of the model is illustrated in Figure 10.3.

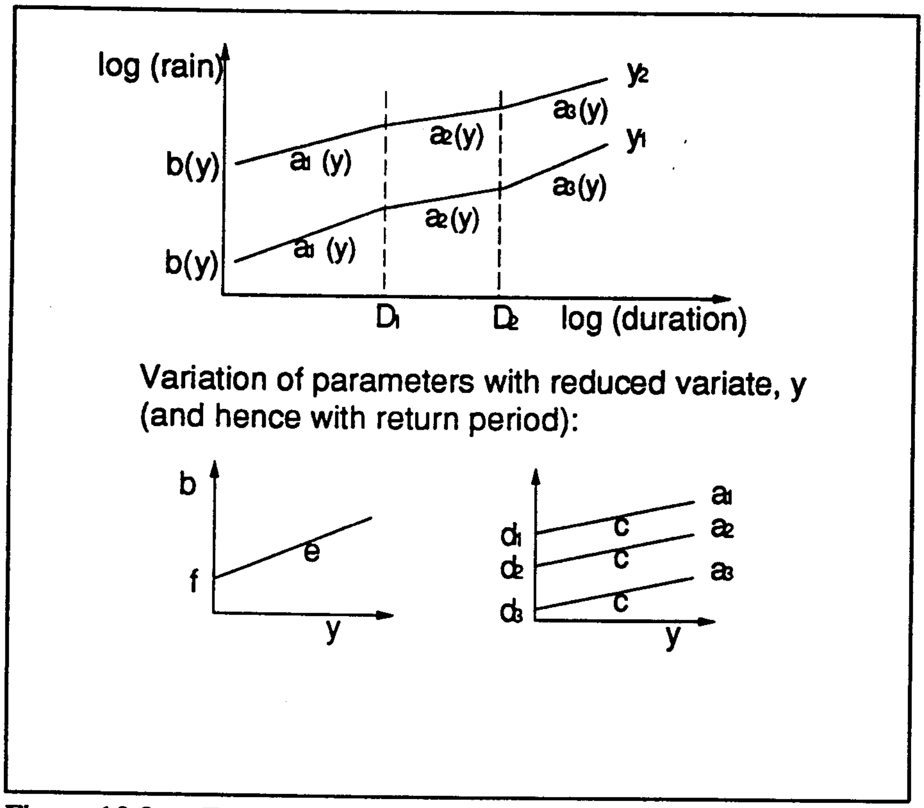

Figure 10.3 Form of the DDF model

There are three line segments which represent the increase of rainfall with duration. The lines change slope at the durations D_1 and D_2. The parameters of the function change linearly with reduced variate y, which is related to T thus:

$$y = -\log\ (-\log\ (\frac{T}{T-1})).$$

The gradient a_i of segment i varies linearly with y, having gradient c and intercept d_i. The intercept b of the first segment varies with y, having slope e and intercept f. For the three-segment model there are six parameters altogether: c, d_1, d_2, d_3, e and f.

Thus, for the first segment ($D<D_1$), the formula is:

$$\log R = (cy + d_1) \log D + ey + f.$$

The model is fitted jointly to data for all durations and return periods at once, using a least-squares criterion.

Variations on this model were investigated, including simpler versions with only one or two line segments. A more complex version was also tried, with more than one c parameter, allowing the slopes a_i to vary independently with y. By comparing the mean squared error from fitting the models to several sites, the above 6-parameter model was chosen.

The choice of the two durations D_1 and D_2 at which the slope should change is not obvious. By looking at DDF plots for several sites, two candidate pairs of durations were chosen:

A 12 hours and 2 days;
B 1 day and 3 days.

Figures 10.4 and 10.5 show models A and B fitted to data from Lowestoft in Suffolk. At this location, model A fits the data rather better. The change of slope at 2 days is particularly strong, as long-duration rainfall increases more slowly with duration at this site. Unusually, short durations (1 to 6 hours) are not very well fitted at this site.

Table 10.2 Mean squared error from two versions of the 6-parameter DDF model. The smaller MSE for each site is printed in bold.

Site	Mean squared error ($/10^{-4}$)	
	Model A (breaks in slope at 12 and 48 hours)	Model B (breaks in slope at 1 and 3 days)
Waddington	17.9	**12.2**
Wallingford	20.1	**19.0**
Bodmin	12.9	**11.3**
Lowestoft	**9.8**	11.5
Snowdonia	**9.5**	9.8
Glen Nevis	**29.5**	30.0
John O'Groats	12.6	**8.9**

Table 10.2 compares the mean squared error (MSE) between modelled and observed rainfall at various sites, for models A and B. The models are fitted to data for seven durations and seven return periods. Although model B gives a smaller MSE at more sites, neither model is a clear favourite. To keep the number of parameters reasonably small (at six), it is necessary to choose one version of the model. Model A, with slope breaks at 12 hours and 2 days, is easier to explain on meteorological grounds. For durations shorter then 12 hours, it is likely that many rainfall extremes are due to convective rainfall. Frontal rainfall is likely to give rise to most extremes for

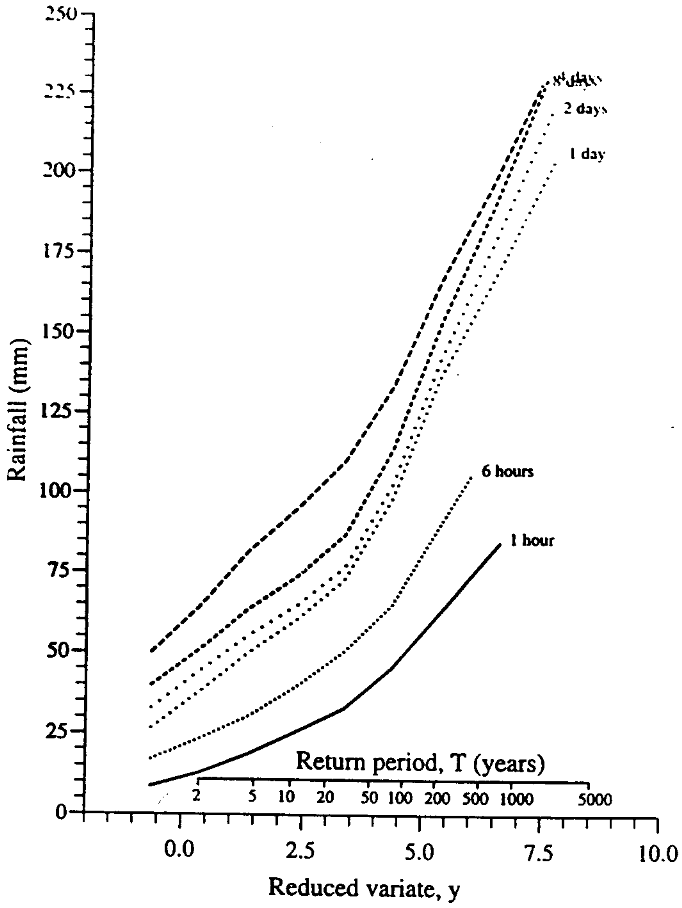

Figure 10.1 Rainfall frequency for Waddington

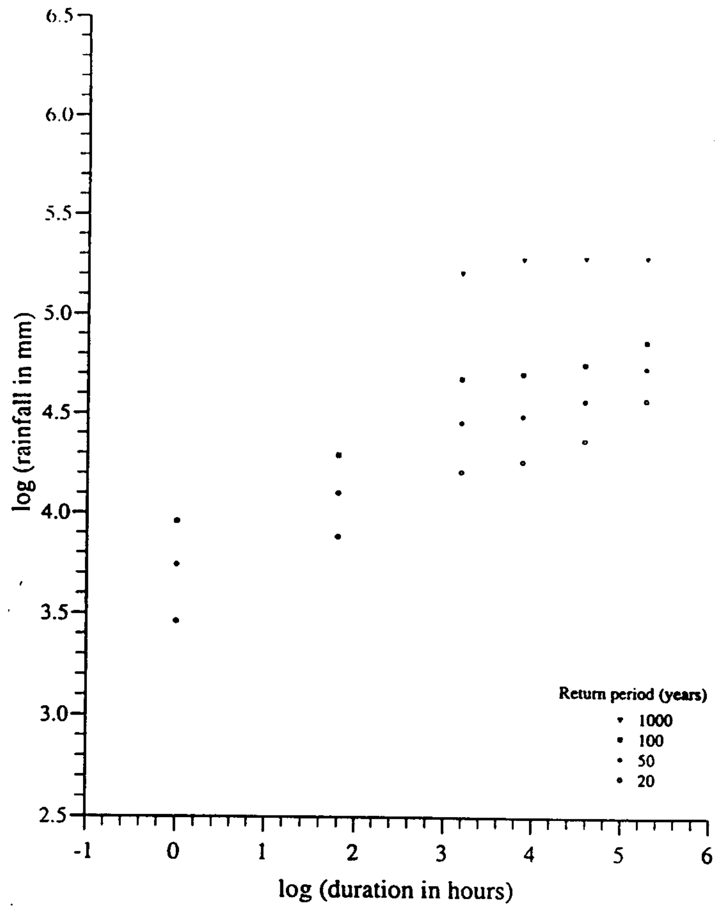

Figure 10.2 DDF plot for Waddington

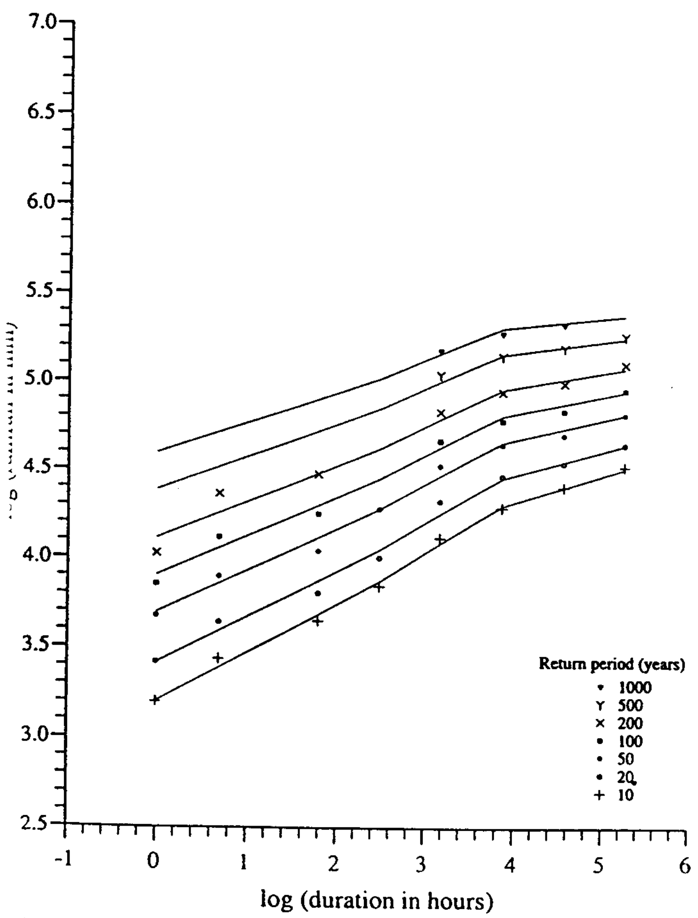

Figure 10.4 DDF model A for Lowestoft

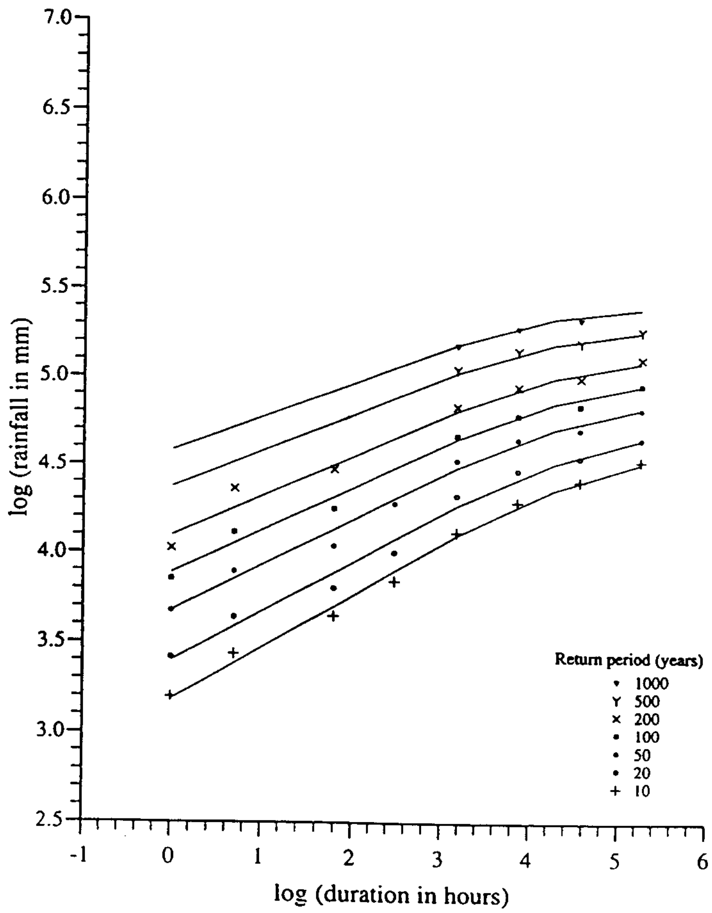

Figure 10.5 DDF model B for Lowestoft

longer durations. Depressions typically pass over Britain in less than two days (Chandler and Gregory, 1976), and so it is reasonable to expect a change in the nature of rainfall extremes at about this duration. Model A is used in the following examples.

It is possible that in reality the durations D_1 and D_2 at which the nature of extreme rainfall changes may vary with return period. For example, a set of convective storm cells is unlikely to give heavy rainfall persisting over very many hours, but could do so in an extreme case, which would produce a long return period rainfall. However, introducing this extra degree of freedom would increase the number of parameters in the model, and also make contradictions between return periods harder to avoid (see below).

The shape of the DDF plot can be rather different from that for Lowestoft. DDF points and fitted models for Waddington and a location in northern Snowdonia are shown in Figures 10.6 and 10.7. The model fits very well for short durations at Waddington, because the data lie in approximately straight lines. Longer durations are not so well fitted, as the empirical relationship changes strongly with return period. For Snowdonia, the results for 1 hour are similar to elsewhere, but the gradients are much steeper, reflecting the importance of frontal rainfall as the source of high rainfall totals for longer durations. The model fits well at all durations apart from 4 days, which is underestimated at moderate and long return periods.

10.2.4 Avoiding contradictions

The DDF model must avoid any contradiction between durations, which would occur if a line on the depth-duration scale had negative gradient, or between return periods, which would occur if two lines for different return periods intersected.

No constraints are imposed on the parameters as they are estimated, but their values can be checked against the following conditions.

To avoid contradictions between durations, the gradient $a_i > 0$. (1)

To avoid contradictions between return periods, R must increase with reduced variate, y:

$$\frac{\partial(\log R)}{\partial y} > 0.$$

For the first segment ($D<D_1$), this yields $c \log D + e > 0$. (2)

For the other segments, the terms in D_1 and D_2 cancel out, so that the above condition (2) applies for $0 \leq D \leq 192$ hours. Because this must be satisfied over the whole range of durations at once, it must apply when $D = 1$ hour, so we know that e must be positive. Therefore, if $c \log D + e > 0$ for $D = 192$ hours, then it is true for all shorter durations as well. Thus it is sufficient to verify that condition (2) is satisfied at the two durations of 1 and 192 hours.

10.3 Comparison with the FSR model for M5 rainfall

A model for M5 (5-year return period) rainfall was used to smooth the values of M5 given in the FSR (Volume 2, section 3.3.3). The model, which applies for durations up to 2 days, is:

$$I = \frac{I_0}{(1 + BD)^n}$$

where I is the rainfall intensity (i.e. R/D), D the duration in hours and I_0, b and n are parameters indexed by the annual average rainfall, SAAR. Because the value of B ranges from 15 to 45, for durations longer than 1 hour, BD>>1. Thus, apart from a gentle inflection at short durations, the model is essentially a straight line on a log-log plot, with gradient 1-n:

$$\log R = \log I_0 - n\log B + (1-n)\log D \quad \textit{for } D>1.$$

Values of M5 for durations longer than 2 days are found by interpolation of the results in Table 3.2 of Volume 2 of the FSR.

The FSR model for M5 has a similar form to the DDF model described in the previous section. However, it applies only to 5-year rainfalls, up to a duration of 2 days. Design rainfalls taken from the FSR for longer return periods and longer durations show relationships with duration similar to those found for the new results. Figure 10.8 is an example, showing FSR results for Waddington plotted on a log-log scale. For durations up to 1 day, the points lie very close to straight lines. For longer durations, the lines decrease in slope and then begin to bend upwards, similarly to the new results for Waddington shown in Figure 10.2 (for return periods of 20, 50 and 100 years at least).

This similarity between FSR results and FEH results indicates that the form of the depth-duration relationship is not a spurious consequence of some aspect of the new rainfall frequency analysis, but perhaps reflects genuine meteorological factors.

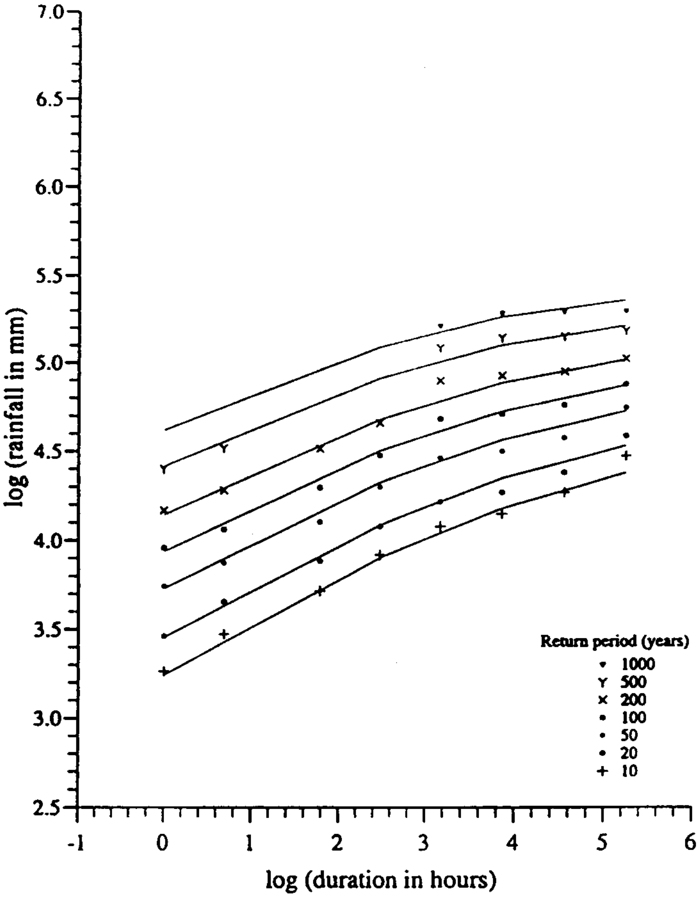

Figure 10.6 DDF model A for Waddington

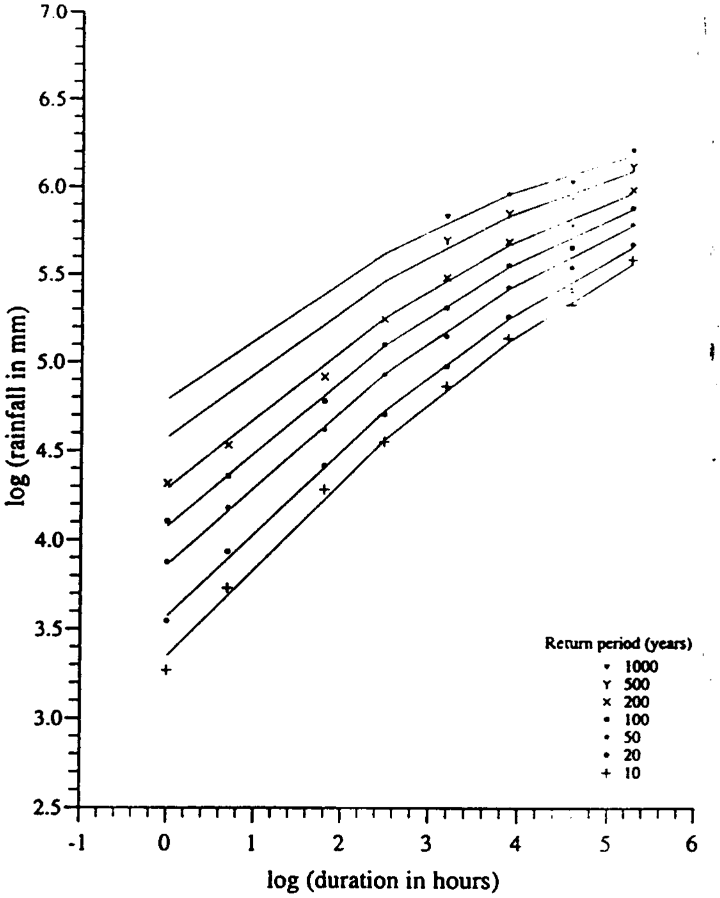

Figure 10.7 DDF model A for Snowdonia

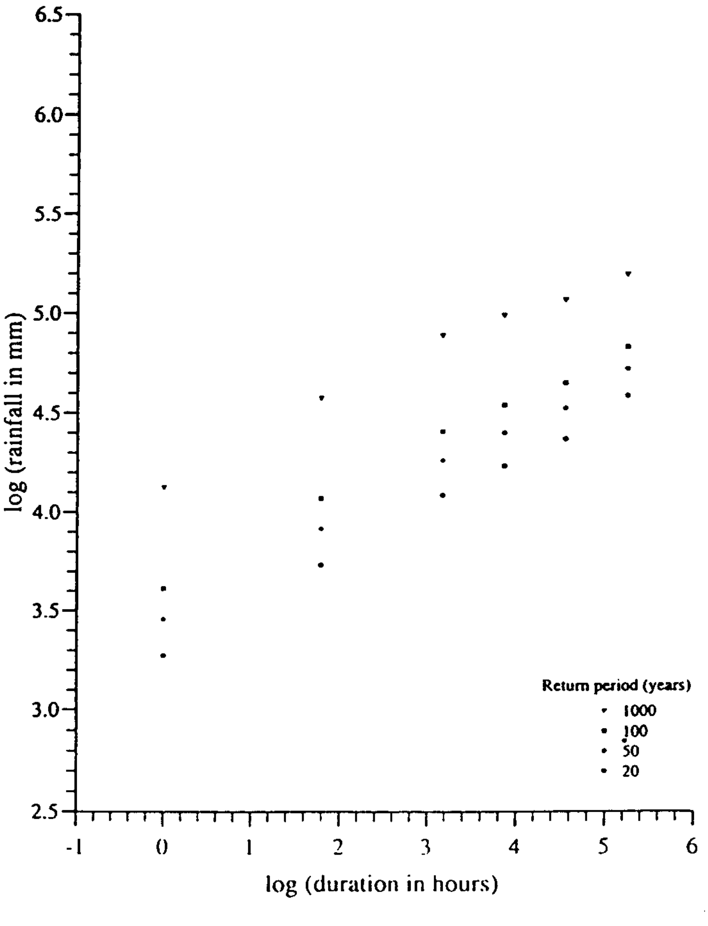

Figure 10.8 DDF plot of FSR results for Waddington

11. PLANS FOR IMPLEMENTING THE RESULTS

Phase 3 of the Rainfall Frequency Studies: England and Wales project is planned to implement the results, making available to users a method of rainfall frequency estimation that is easy to apply and comprehensively documented.

Since the project started, it has become clear that a major implementation will be in the form of a rainfall volume in the new Flood Estimation Handbook (FEH), which is being co-ordinated under funding by MAFF. The completion of FEH Volume 2, Rainfall frequency estimation, is planned for December 1998, and the FEH is due to be published in early 1999. FEH Volume 2 will include a description of the methods contained in this Technical Report, and worked examples. Users will be able to estimate design rainfalls for any point in the UK, using derived data which will accompany the FEH, possibly on a CD-ROM.

The nature of the derived data which will accompany the FEH is a matter for negotiation between IH, the Met. Office (as data suppliers) and the Environment Agency (as funders and data suppliers), as stated in the Memorandum of Understanding (MOU, Appendix 5). The views of future users of the FEH will be sought via the FEH Advisory Group. One possibility for implementation is a model of rainfall depth-duration-frequency with parameters evaluated on a 1-km grid across the UK. Negotiations with the Met. Office have been started.

Other implementations are likely through updates to existing software packages such as the successor to ITED (Met. Office), Micro-FSR (IH) and RAINARK (Hydro-Logic Ltd.). Such implementations will be agreed between the three parties to the MOU. As required in the MOU, IH will make available to the Environment Agency and the Met. Office the model of depth-duration-frequency, copies of the database of annual maximum rainfalls, maps of RMED and software for the FORGEX method.

12. CONCLUSIONS

Following Phase 2 of the project "Rainfall Frequency Studies: England and Wales", new estimates of design rainfalls are available across the UK, for the durations between 1 hour and 8 days, and return periods as long as 2000 years. A comprehensive and up-to-date database of annual maximum rainfalls has been created, comprising records from over 6000 daily gauges and 375 recording raingauges.

The conclusions of the work described in this report are summarised here.

- The FORGEX method produces rainfall growth curves for return periods up to 2000 years for daily data, and up to 800 years for hourly data.
- The plotting of network maximum rainfalls enables estimation of growth rates for long return periods in a way which accounts for spatial dependence among annual maximum rainfalls.
- FORGEX uses rainfall data from large regions, while giving precedence to local data where they are available.
- Bootstrapping is a useful technique for estimating confidence limits for growth curves.
- Maps of growth rates reveal local and regional variations due to meteorological and geographical features. The results are not unduly sensitive to gauge density or individual large rainfall events.
- Georegression is an effective method for using topographical information to assist the mapping of the index variable, RMED, particularly in sparsely-gauged upland areas. Maps of RMED for durations between 1 hour and 8 days have been produced on a 1-km grid.
- Maps of design rainfall, combining growth rates and RMED, show that the largest estimates for long durations are in upland areas to the west. For 1-hour rainfall, there are also relatively large estimates in the East Midlands, London and Lancashire.
- Compared with the Flood Studies Report rainfall frequency estimates, the new results are more spatially variable, tending to be greater in most areas for 1-hour rainfall. For the 1-day duration, design rainfalls are greater than the corresponding FSR results in some areas, and smaller in others.
- A model of depth-duration-frequency has been developed which provides estimates at any duration and reconciles estimates at different durations.

The final phase of the project will involve implementing the work and making the results widely available, principally through Volume 2 of the new Flood Estimation Handbook.

References

Acreman, M.C. and Collinge, V.K. (1991) The Calderdale storm revisited: an assessment of the evidence. In: *Proc. Third National Hydrology Symposium* (Southampton 1991), 4.11-4.16.

Adby, P.R. and Dempster, M.A.H. (1974) *Introduction to Optimization Methods*. London: Chapman and Hall.

Bleasdale A. and Chan Y.K. (1972) Orographic influences on the distribution of precipitation. In *Proc. Distribution of precipitation in mountainous areas* (Norway 1972) **326** (II), 140-146. World Meteorological Organization, Geneva.

Bootman, A.P. and Willis, A. (1977) Extreme two day rainfall in Somerset. Wessex Water Authority.

Buishand, T.A. (1993) Rainfall Depth-Duration-Frequency Curves; a Problem of Dependent Extremes. In *Statistics for the Environment* (ed. V. Barnett and K.F. Turkman), pp. 183-197. Wiley.

Chandler, T.J. and Gregory, S. (1976) Editors. *The Climate of the British Isles*. London: Longman.

Chen, C. (1983) Rainfall intensity-duration-frequency formulas. *Journal of Hydraulic Engineering, ASCE,* **109**, No. 12, 1603-1621.

Dales, M.Y. and Reed, D.W. (1989) Regional flood and storm hazard assessment. *Report no.102*, Institute of Hydrology, Wallingford.

Deutsch, C.V. and Journel, A.G. (1992) *GSLIB Geostatistical Software Library and User's Guide*. Oxford University Press.

Dwyer, I.J. and Reed, D.W. (1995) Allowance for discretization in hydrological and environmental risk estimation (ADHERE). *Report no. 123*, Institute of Hydrology, Wallingford.

Efron, B. (1979) Bootstrap methods: another look at the jackknife. *Annals of statistics*, **7**, 1-26.

Ferreri, G.B. and Ferro, V. (1990) Short-duration rainfalls in Sicily. *Journal of Hydraulic Engineering, ASCE,* **116**, No. 3, 430-435.

Flohn H. (1969) *Climate & weather*. World University Library, 253pp.

Jones, D.A. (1997) Plotting positions via maximum-likelihood for some non-standard situations. *To be submitted?*

Journel A.G. and Huijbregts C.J. (1978) *Mining geostatistics*. London: Academic Press

Konrad C. (1996) Relationships between precipitation event types and topography in the southern Blue Ridge mountains of the southeastern USA. *International Journal of Climatology*, **16**, 49-62.

Prudhomme C. and Reed D.W. (1997a) Mapping extreme rainfall of mountainous regions: a case study in Scotland. *Internal IH review*, May 1997.

Prudhomme and Reed (1997b) Relationships between extreme daily precipitation and topography in a mountainous region, A case study in Scotland. *Submitted to International Journal of Climatology*, September 1997.

Reed, D.W., Faulkner, D.S. and Stewart, E.J. (1997) The FORGEX method of rainfall growth estimation II: Description. *Awaiting submission to HESS*.

Reed, D.W. and Stewart, E.J. (1989) Focus on rainfall growth estimation. In: *Proc. Second National Hydrology Symposium* (Sheffield 1989), 3.57-3.65.

Reed, D.W. and Stewart, E.J. (1994) Inter-site and inter-duration dependence in rainfall extremes. In *Statistics for the Environment 2: Water-related Issues* (ed. V. Barnett and K.F. Turkman), pp. 125-143. Wiley.

SAS Institute Inc. (1989) *SAS/STAT® User's guide*. Version 6, 4th edition, Cary, NC, SAS Institute Inc.

Shao, J. and Tu, D. (1995) *The Jackknife and Bootstrap*. Springer Series in Statistics.

Stewart, E.J., Faulkner, D.S. and Reynard, N.S. (1995) Rainfall Frequency Estimation in England and Wales. Phase 1b: Pilot Study. R&D Note 478, National Rivers Authority, Bristol.

Stewart, E.J. and Reynard, N.S. (1994) Rainfall Frequency Estimation in England and Wales. Phase 1a: Survey. R&D Note 175, National Rivers Authority, Bristol.

Appendix 1 Derivation of network maximum plotting positions

The plotting positions for network maximum points which are individually shifted (Section 3.3) are derived as follows. Full mathematical details can be found in Jones (1997).

For a simple case where observations (annual maxima) have a common distribution function F, it is possible to find a maximum likelihood estimator p^* of $p=F(x)$:

$$p^*=r/N$$

where: r is the number of observations less than or equal to x;
N is the total number of observations.

This gives an estimated non-exceedance probability for any x-value. We require a plotting position for an observation. Unfortunately, r changes as x moves from just below to just above an actual observation, so the value of p^* for the observation is not clear.

One way forward is to use a (weighted) average of two log-likelihood functions, one for an x-value immediately below the rth observation, and one for a value immediately above it. The first value is weighted by b and the second by (1-b). In this simple case, this yields a plotting position:

$$p^*=(r-b)\,/\,N.$$

For a netmax series, we want to shift each point by an amount $\ln(N_e(i))$ where $N_e(i)$ is the effective number of independent sites in year i. This is equivalent to raising F(x) for the ith value of x to the power $N_e(i)$. The above average likelihood procedure is used again, with a weighted average of two log-likelihood functions. After several changes of variables, the likelihood function for the plotting position p can be written down and then maximised by differentiation to give a formula which can be solved by the Newton-Raphson approximation. The formula includes a summation such that the plotting position of the rth netmax point depends on the plotting positions of all higher netmax points.

The specification for the weighting constant b to produce a Gringorten plotting position is not immediately obvious in the case where the numbers of equivalent sites vary. For the largest observation, it can be shown that b = 0.56 produces a Gringorten plotting position (Jones, D.A., personal communication). More generally, for the observation with rank r from the bottom,

$$b = \frac{0.44\sum_{i=1}^{N} N_e(i) + 0.12\sum_{i=1}^{r} N_e(i)}{\sum_{i=1}^{N} N_e(i) + 0.12}$$

This has the effect of decreasing the weight b for lower-ranked netmax points.

Appendix 2 Derivation of confidence intervals by bootstrapping

The key assumption for deriving bootstrap confidence intervals is that the bootstrap *residuals*

$$E_i = G_i - G_{sam}$$

(where G_i is the estimate from bootstrap sample i and G_{sam} is the basic estimate from the sample data) are assumed to be representative of values drawn from the same distribution as the actual unknown residual $E = G_{sam} - G_{true}$. Shao and Tu (1995) call this the hybrid bootstrap.

If the distribution is asymmetric, the correct assumption gives results which are the reverse of what the user may have expected, based on the wrong assumption that the resamples are simply alternative realisations of the original sample. Using the correct assumption, if most of the range of resampled growth curves lies below the estimated growth curve, then most of the range of possible estimated growth curves lies below the true growth curve.

Mathematically speaking, if E_L and E_U are appropriate lower and upper percentage points of the unknown distribution of residuals, such that the probability

$$Pr\ (E_L \le E \le E_U) = 1\text{-}2\alpha$$

then $$Pr\ (E_L \le G_{sam} - G_{true} \le E_U) = 1\text{-}2\alpha$$

or $$Pr\ (G_{sam} - E_U \le G_{true} \le G_{sam} - E_L) = 1\text{-}2\alpha.$$

So $(G_{sam} - E_U\ ,\ G_{sam} - E_L)$ is a confidence interval with coverage $1\text{-}2\alpha$.

The percentage points E_L and E_U are estimated by the empirical percentage points of the ordered bootstrap residuals, E_m and E_n:

$$(G_{sam} - E_n\ ,\ G_{sam} - E_m).$$

This is the formula given in the main part of the paper. If we go on to substitute $G_m - G_{sam}$ for E_m where G_m is the mth smallest resampled growth rate, and similarly for E_n, we obtain:

$$(2G_{sam} - G_n\ ,\ 2G_{sam} - G_m),$$

so that the upper percentage point is used to find the lower confidence limit. The incorrect approach yields the interval (G_m, G_n).

Note that it is possible for this formula to give a lower confidence limit which dips below 1.0, which is unrealistic for growth rates at return periods over two years. However, this is unlikely to be a problem for higher return periods. An alternative strategy would involve defining the residual as the ratio of growth rates. This would permit the use of a transformation to ensure that the lower confidence limit could not drop below 1.0. This alternative method can give results which are rather different to those of the above method, when the confidence limits are asymmetrical.

Note also that it would be possible to use the adjusted estimator G_{sam}-E_{med}, where E_{med} is the median of the bootstrap residuals, as a bias-adjusted estimate, where the bias is estimated via the bootstrap.

Appendix 3 Reason for using 199 resamples in bootstrapping

It is not immediately clear how many resamples to use. Some authors recommend the use of the 6th highest and lowest values out of 200. We used the 5th highest and lowest out of 199 for the following reason.

The 95% confidence region consists of all those growth rates g for which there is not enough evidence to reject the null hypothesis that the variables $E = (G_{sam}-g)$ and the bootstrap residuals $E_i = (G_i - G_{sam})$ have the same distribution, i.e. the null hypothesis is $g=G_{true}$. The hypothesis is rejected if E is in one of the tails of the distribution of the E_i's.

Suppose there are B random samples, and they are ordered $E_1, E_2, \ldots, E_B$. Under the assumption of a common distribution for E and E_i (i.e. the null hypothesis), E is equally likely to appear at any point in the ordered set of E_i's. Each of the following occurrences is equally likely:

$$E < E_1,$$
$$E_1 < E < E_2,$$
$$\ldots$$
$$E_{B-1} < E < E_B,$$
$$E_B < E.$$

Each has probability $1/(B+1)$. Therefore the probability of E occurring in one of the 2m most extreme regions (i.e. $E < E_m$ or $E > E_{B-m+1}$) is exactly $2m/(B+1)$. This is the probability of rejecting the null hypothesis when it is actually true. When this is subtracted from one, it yields the probability that the rest of the distribution of the E_i's contains the true value. Hence the coverage probability of the confidence interval is $1 - 2m/(B+1)$.

When B=199 and m=5, this yields a coverage probability of 0.95. If resources permit the extra computation, values of B=1999 and m=50 could be used.

Appendix 4 Terms of reference for Phase 2

These terms of reference are taken from the R&D project initiation document.

Overall objectives

To extend the work of the pilot study to all areas of England and Wales, to develop rainfall growth rates for return periods in excess of 200 years, and for durations between 1 hour and 8 days.

Specific objectives

(a) Collect rainfall data from a variety of sources to derive annual maxima for durations of 1 hour to 8 days.

(b) Develop focused growth curves for the range of durations above, for return periods ip to 200 years, using the methodology derived in Phase 1b.

(c) Develop a modified station-year method to allow these growth curves to be extended to return periods in excess of 200 years.

(d) Compare the results with those derived from FSR Vol II and FORGE where appropriate.

(e) Analyse the spatial variability of growth rates and determine the optimum number of focal points.

(f) Consider the consistency between growth curves of different durations.

(g) Determine the confidence limits of the new growth curves.

(h) Produce a report in the form of an R&D Technical Report summarising the work undertaken in Phase 2.

(i) Produce quarterly progress reports detailing the progress made against planned including any variation in planned expenditure

Appendix 5 Memorandum of Understanding (1992)

Introduction

This Memorandum of Understanding sets out to define the responsibilities of the principal parties involved in the project "Rainfall Frequency Studies: England and Wales". These are the National Rivers Authority (NRA), the Institute of Hydrology (IH), and the Meteorological Office (Met. Office). The Project is defined in Para. 3 and, more fully, as part of the Project Investment Appraisal prepared by the NRA. The Memorandum of Understanding acknowledges the close links between the three parties and identifies their respective contributions to rainfall frequency estimation in the UK. It also acknowledges that the NRA contributes generally to the work of the Met. Office provision of rainfall data from NRA - operated sites.

The Memorandum of understanding is separate from, but complementary to, the Phase 1 contract between the NRA and IH.

It is agreed that no party to this agreement will seek to profit unduly from the inputs (as listed in Section 4) of any other party.

It is agreed that the Memorandum of Understanding may be altered during the course of the project by agreement of all three parties.

1. Background

Procedures for rainfall frequency estimation were set down in Volume II of the UK Flood Studies Report (Natural Environment Research Council, 1975), based on research at the Met. Office. Subsequent to particular studies in South West England, IH was asked by the NRA to investigate the feasibility of reworking rainfall frequency estimation procedures for England and Wales. The Project requires sub-daily rainfall data held by the Met. Office as custodians of the national rainfall archive.

2. Objectives

i) Review the current status of research in methods of rainfall frequency estimation for flood studies;

ii) Where appropriate, develop new methods of determining rainfall frequency in England and Wales;

iii) Produce documents, algorithms and software to allow the procedures to be implemented with clarity and ease, to the satisfaction of all parties.

3. Project

The terms and scope of the project agreed between the NRA and IH are attached as Annex to this Memorandum of Understanding. The project will be in three main phases.

a) Phase 1 will be a feasibility study comprising a "scoping survey" (in which requirements, methods and data are appraised) and a pilot study in an agreed region. If Phase 1a (scoping survey) is successful then Phase 1b (a pilot study) will be undertaken. The outcome of Phase 1a will be a Project Report. The outcome of Phase 1b will be a Project Report which will include a detailed proposal for Phase 2. The proposal will take account of the market study undertaken by the Met. Office and IH, identifying the potential beneficiaries of Phases 2 and 3.

b) If the feasibility study (Phase 1a and 1b) is successful, Phase 2 will extend the study to the whole of England and Wales. It will develop a series of algorithms and related documentation.

c) Phase 3 will see the conversion or adaptation of the algorithms derived in Phase 2 into agreed forms suitable to each party.

4. Inputs

a) Inputs from the Met. Office

   i) The Met. Office will supply IH with sample hourly raingauge data for the scoping survey (Phase 1a). If the Project moves on to Phase 1b the Met. Office will supply hourly raingauge data for an agreed pilot region (probably East Midlands). If the Project moves on to Phase 2, the Met. Office will supply further hourly rainfall data for gauges throughout England and Wales (and in the Scottish borders). These data, taken from the national rainfall archive, will include some stations operated by the NRA, and stations for which data is collated by NRA.

   ii) The Met. Office will supply IH with such hourly statistical data (May and Hitch, 1989) as are available: for sample gauges in Phase 1a, for gauges in the pilot region if Phase 1b proceeds and, if the Project proceeds to Phase 2, for gauges throughout England and Wales (and in the Scottish borders).

   iii) Met. Office charges for the supply of these data will be confined to staff, computing and administration. No charge will be made in respect of costs accrued in observing, processing, quality control and analysis of the data, or in instrument inspection or database maintenance.

b) Input from NRA

i) The NRA will supply IH with hourly data for sample gauges in Phase 1a. Subject to the evaluation of Phase 1a, the NRA will supply additional raingauges in the agreed pilot region for Phase 1b and throughout England and Wales for Phase 2.

ii) The NRA will fund Phase 1a work at IH, and also Phase 1b should it proceed. It will also give strong support to Phases 2 and 3 if feasibility is proved in Phase 1.

c) Input from IH

i) IH will contribute expertise in rainfall frequency analysis derived in projects funded by DOE (from 1985), South West Water (1988), MAFF (1989) and NERC (1990).

ii) IH shall carry out its study for the NRA as described under "objectives" above and subject to the NRA/IH contractual agreement for "Rainfall Frequency Studies: England and Wales".

5. Outputs

a) Derived statistical data and algorithms

IH will provide copies of derived statistical data sets and algorithms (and with any subsequent amendments thereto) in agreed forms compatible with NRA and Met. Office requirements. The derived statistical data will comprise three main types of information:

i) rainfall statistics - deriving reasonably directly from the basic data,

ii) rainfall maps - coming from a fairly advanced analysis using digital terrain data,

iii) rainfall growth curves - deriving from an advanced analysis of extreme values.

b) Software implementations

Various implementations of the derived statistical data and algorithms for rainfall frequency estimation at any site in England and Wales are anticipated with existing packages such as Micro-FSR and RAINARK (both of which are heavily used by the NRA, either directly or by delegation), and within the successor to the Met. Office's ITED rainfall frequency estimation program. The form of implementation will reflect the characteristics of the particular package, and may be carried out (or contracted out) by any of the parties to the Memorandum of Understanding, subject to agreement. IH will supply copies of the derived statistical data and algorithms in computer-compatible form to assist in any approved implementation.

c) If all parties are in agreement that the revised methods supersede previous FSR methods for rainfall frequency estimation, IH will produce a Flood Studies Supplementary Report.

6. **Rights and Restrictions**

a) The restriction imposed on IH concerning use of pre-1961 daily rainfall data for some 400 long-term stations (held at IH under a previous project) will be lifted. This relaxation applies only to the "Rainfall Frequency Studies: England and Wales" project, and any other use must (as at present) be with the prior approval of the Met. Office.

b) IH will use the statistical data supplied by the Met. Office (Para. 4aii) and the basic hourly rainfall data supplied by the Met. Office (Para. 4ai) and NRA (Para. 4bi) only in the specified project. Any further use of these data by IH or their contractors or any transfer outside IH, would require further agreement with the supplier (ie. Met. Office or NRA); the Met. Office or NRA reserve the right to refuse a request, to apply commercial charging rates, or to enter a further special agreement.

c) The NRA and its delegated representatives (as agreed by the other two parties) will be free to use the derived statistical data and algorithms (Para. 5a) in any application or further research.

d) IH will be free to use the derived statistical data and algorithms in any application or further research, and to supply (but not sell) these to any bona fide researcher.

e) The intellectual property rights for the outputs of the research will be shared according to the relative value of inputs to this, and related earlier projects. Subject to the feasibility study not demonstrating undue profit, the suppliers of data (Met. Office and NRA) shall be granted sole rights to sell a general rainfall frequency estimation service based on the derived statistical data and algorithms. Software implementations in commercial packages such as Micro-FSR, RAINARK and WIS will be consistent with their general roles and will not seek to displace the successor to the Met. Office's ITED package.

This paragraph will be subject to review following Phase 1 in the light of more detailed economic data.

f) The parties will openly advertise this study and their collaborative activities and responsibilities. Any publication package or sale arising out of the project will acknowledge the three principal parties, and may quote the Memorandum of Understanding.