UNIVERSITY COLLEGE CARDIFF

DEPARTMENT OF

MECHANICAL ENGINEERING AND ENERGY STUDIES

Angular distribution of diffuse radiation: data processing

Peter B. Lloyd

Departmental report no. 1373 (SEU 562) vol. 1
May 1987

ABSTRACT

A description is given of software that has been developed for accessing and processing the large set of irradiance measurements taken at UCC with scanning pyranometers.

The pre-processing of the data is described, with the correction and rehabilitation of the data being treated in detail.

Some basic results are presented in the form of polar plots of the solar irradiation on tilted surfaces for the twelve months of the year, and for representative types of sky condition (cloudless, sparsely clouded, densely clouded, and overcast skies). An initial assessment of the anisotropy of the diffuse radiation field is given.

This report was written between March and May 1987, as the final part of the NERC-funded project, "The distribution of the diffuse radiation field and its dependence on atmospheric conditions and solar elevation" (Grant GB3/5061).

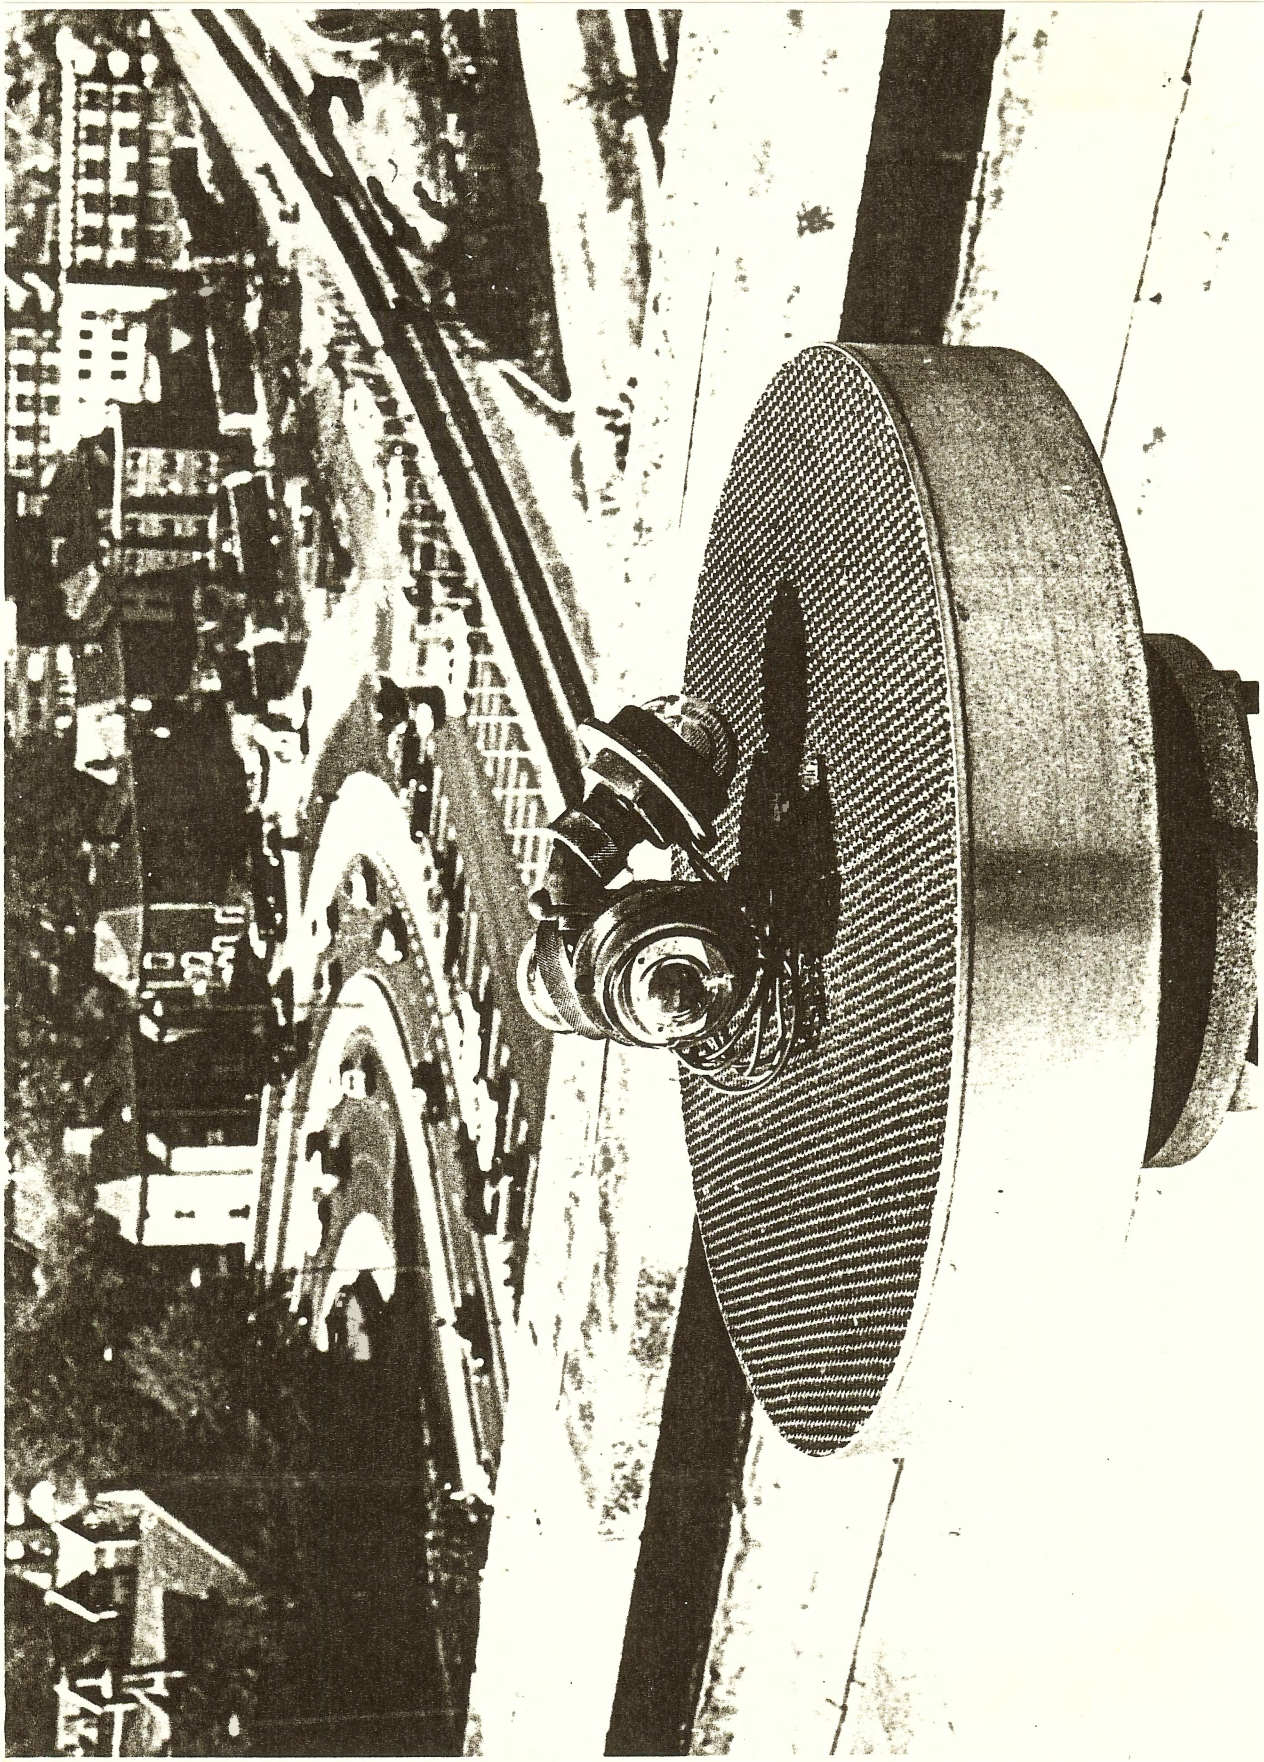

Photograph 1: Scanning pyranometers (D.A. Svendsen, 1978)

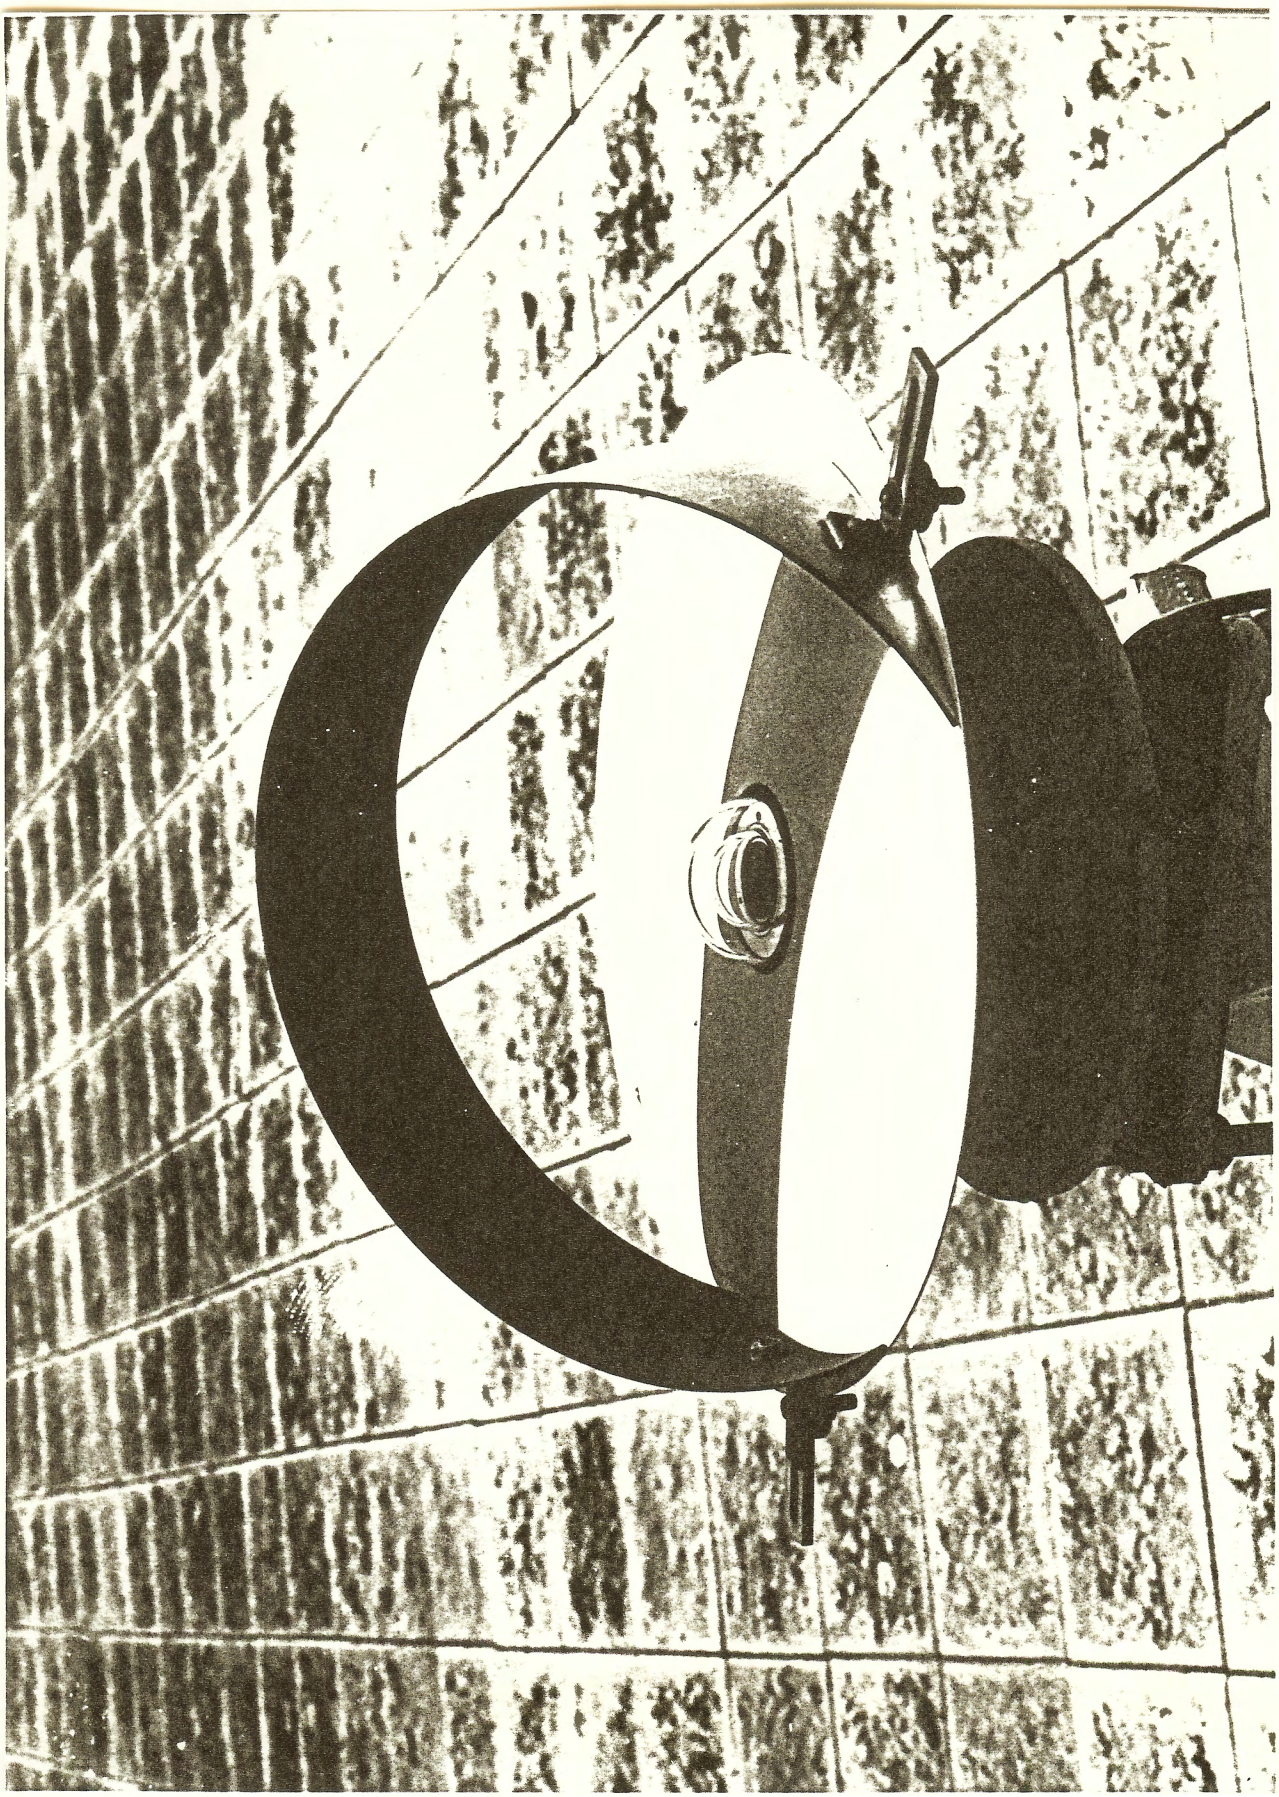

Photograph 2: Summer shade-ring (D.A. Svendsen, 1978)

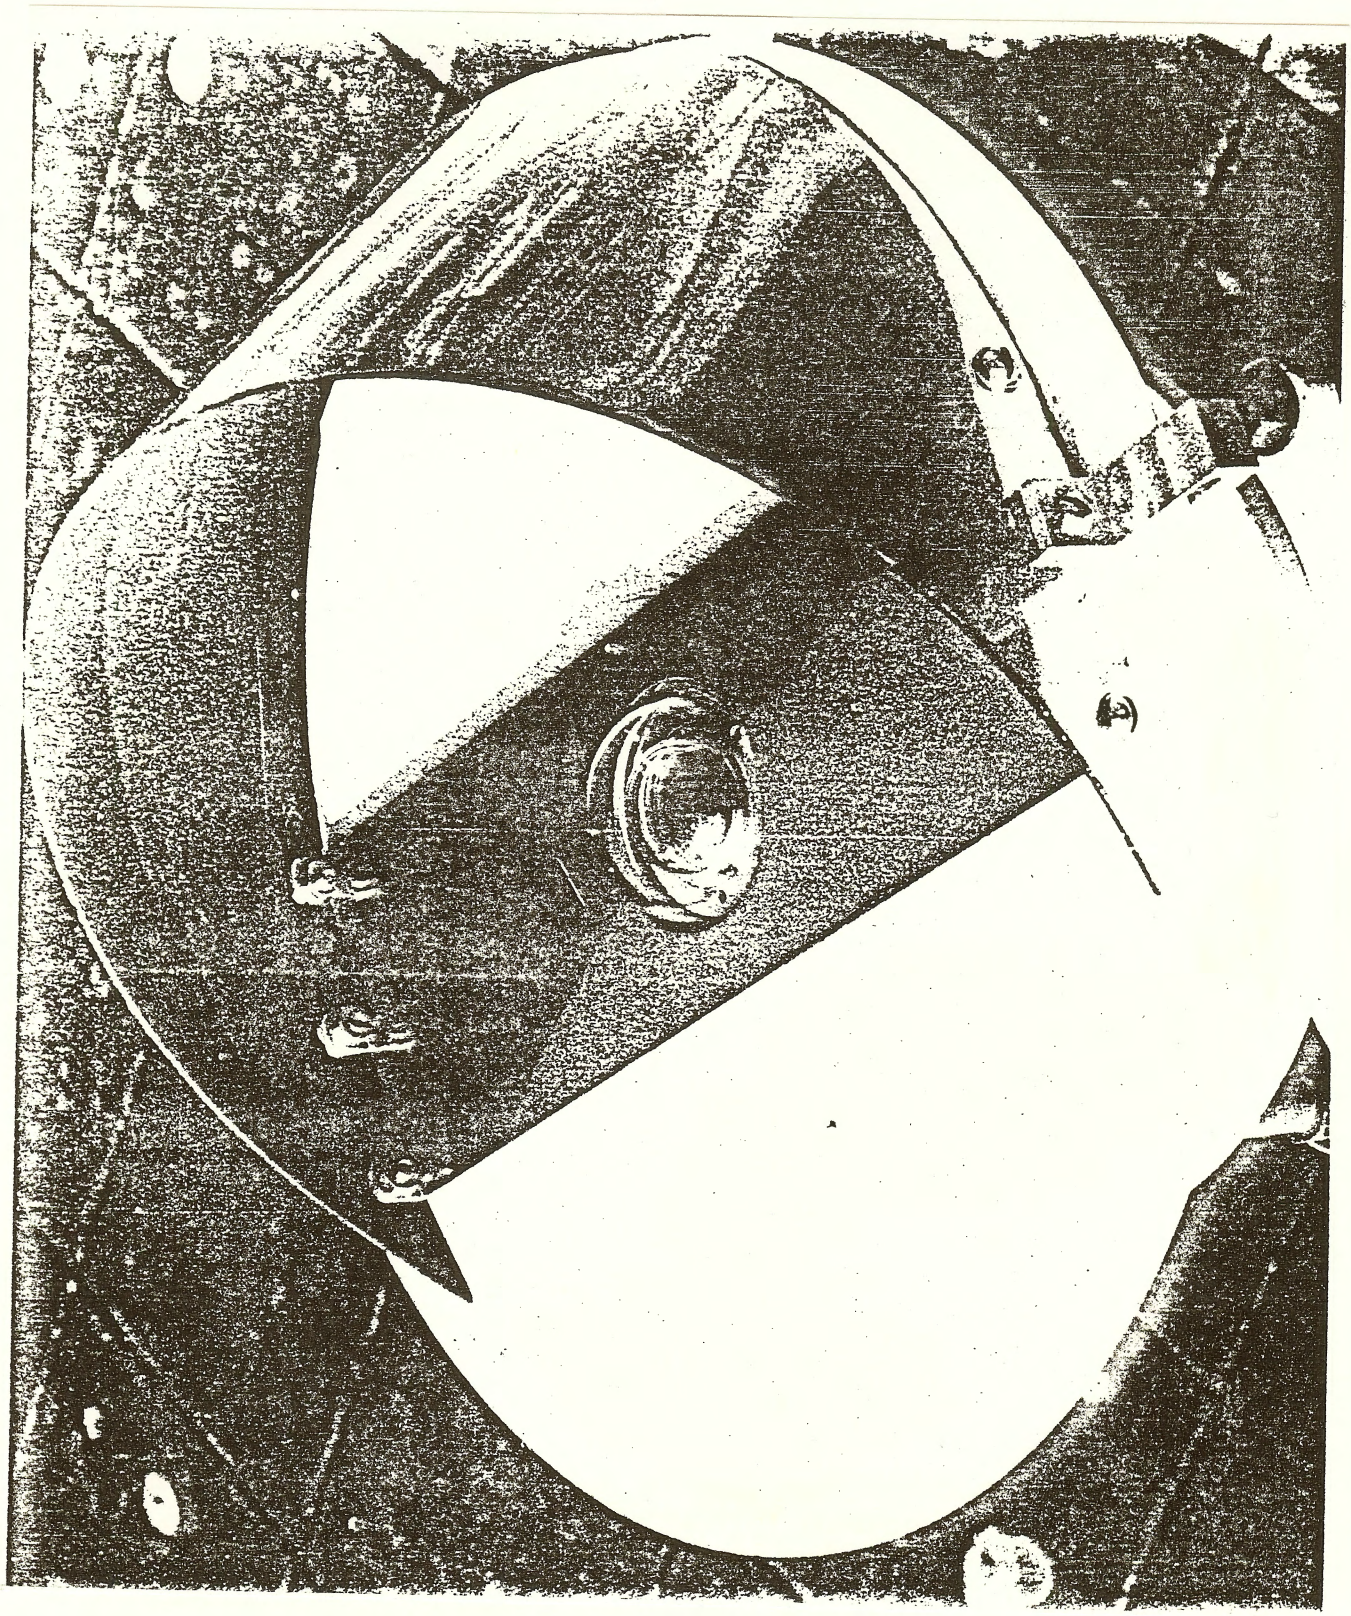

Photograph 3: Winter shade-ring (D.A. Svendsen, 1978)

CONTENTS

Volume 1

Volume 3

Annexes: listing of computer programs

(Programs written by Lloyd, except where D.A.Svendsen is indicated.)

A. Program for restructuring DEC files into GEC files
- A.1 TOPV master program
- A.2 STRUV common core of subroutines
- A.3 RTGCV subroutines specific to RT-11 and GEC-4090 operating systems
- A.4 MTFIV subroutines for reading magnetic-tape sources
- A.5 DCFI3 subroutines for reading disc sources
- A.6 COPYLIBV subroutines for copying files from one structure to another
- A.7 AUXV miscellaneous auxiliary subroutines

B. Software for plotting a chronology of radiation monitoring
- B.1 CHROPROG program for plotting a chronology of radiation monitoring
- B.2 CHRODATA information specifying the chronology

C. Programs for processing pyranometric time responses
- C.1 LOGTAU program for measuring the time response of a pyranometer (DAS)
- C.2 GENTAU program for computing pyranometer time-constants from measured time response (DAS)
- C.3 TSTEP program for measuring the time response of a pyranometer & finding time constants (DAS)
- C.4 TRAMP ditto, but with ramp input (DAS)
- C.5 TRAMP1 ditto (DAS)
- C.6 CHKTAU program for computing sensitivity of corrected irradiances to errors in time constants (DAS)
- C.7 CHKOLD ditto (DAS)
- C.8 PBLTAU program for computing time constants for, and generating plots of, Svendsen's data (Lloyd)
- C.9 LOOKTAU program for plotting candidate curves for fitting to pyranometric time responses (Lloyd)

D. Program for RRR-to-PPP conversion
- D.1 RTOP program for converting RRR files to PPP files (DAS)

E. Programs for analysing Svendsen data
- E.1 DAS3LIB library of core subroutines
- E.2 TIMEPLOT program for plotting irradiance against time
- E.3 POLEPLOT program for plotting polar graphs of irradiances and errors in irradiance estimates
- E.4 PPPEDIT program for editing PPP files

F. Programs for shade-ring computations
- F.1 ANICHECK program for checking implementation of Steven's anisotropic model
- F.2 SHADES program for computing proportion of irradiance obscured by shade rings

G. Programs for geometric modelling
- G.1 UNISCAN program for scanning a Meteorological Office data file

1. Archive tapes of radiation monitoring at UCC

During the period 1976 to 1982, a large set of measurements of solar irradiance were recorded by D.A. Svendsen and J. McGregor using radiometers situated on the roof of the Mining building at UCC's Newport Road site.

The radiation data were recorded in the first instance on domestic cassette tapes, and then archived on standard half-inch tapes. More than half of the archived data were taken with an apparatus that was set up by Svendsen in 1978. This consisted of a system of pyranometers designed to allow comprehensive monitoring of the angular distribution of irrradiance. These are the data with which we are concerned in the present project.

The contents of the data tapes left by Svendsen and McGregor, after their successive departures from UCC, were not fully documented. Some discussion is therefore required of the problem of locating and extracting the information that is on the tapes. For certain sorts of file, this problem is discussed in the sub-sections from 1.2 onwards. It is stressed that none of the data files or program files that have been retrieved from the archive tapes would now be available without the file-restructuring program, STRUC, which is described in the next sub-section.

1.1 Software for accessing the archive tapes

The Svendsen data were written onto magnetic tape with a DEC PDP-11 computer running an RT-11 operating system. For financial reasons the tape drive was sold before the beginning of the present programme of work. It was therefore necessary to read and process the data on the SERC's GEC 4090 located at UCC. This computer employs a different operating system and a different file structure from the PDP-11. In order to read the data into 4090, some software was written to restructure DEC files into a form that could be used on the GEC machine. This file-restructuring software (called STRUC) is detailed in Appendix A.1, and briefly outlined here.

Another option that was considered was to find some other academic institution that possessed an under-used PDP-11 computer with a 1600 bpi tape drive, take the tapes there, and re-write them all in some format that would be mutually comprehensible to both the GEC and DEC machines. There were four reasons for not taking that option. First, there are the time and financial costs of making the arrangements with the other institution, and of transport and accommodation for the week or more that would be needed. Second, re-writing the tapes in a special format would itself require some software development, albeit not as great as that of producing STRUC. Third, we would have to either over-write the tapes, which would involve a risk of losing files if the re-writing process were faulty, or buy a new set of about thirty tapes at a cost of about $300, for which funds were unavailable. Third, the re-structuring program would have other uses beyond the present project: Microdata loggers are still in use in the SEU, and QQQ files are still being created. The STRUC package offers a means of processing those files on the GEC 4090,

which is more powerful than the DEC PDP-11.

In the remainder of this sub-section, some of the functions that STRUC must perform are described in outline.

(a) Bytes and words.

Archived radiation data are normally stored as integers. In the data-handling procedure that was set up by Svendsen, the raw data are stored as 8-bit integers, as their values lie in the restricted range of 0-255 logger units (LU). Calibrated data are stored as 16-bit integers, as their values lie in the range 0-1400 Wm^{-2}.

The FORTRAN system on the GEC computer, unlike that on the DEC machine, cannot handle single bytes as integers. It is therefore necessary for the STRUC software to expand 8-bit integers to 16-bit words, whenever it is copying a raw-data file into the GEC computer.

In the DEC system, the architecture is set up to handle 16-bit integers in which the two component bytes are in the reverse of the natural order. The bytes of each 16-bit integer must therefore be reversed by the program STRUC before the data are output to a file on the 4090.

(b) Records and blocks.

The organisation of records inside blocks is completely different between the DEC and GEC computers. In addition, there are several quite different block structures in use on the DEC machine; some are standard DEC structures, others were devised by Svendsen or McGregor. STRUC must therefore recognise which internal structure a file has, unpack the logical records in accordance with that structure, and output the records to the GEC file.

Lloyd (1985, Rept 1204) has collated full details of the internal file structures that occur on the archive tapes.

(c) Tape blocks.

Individual blocks on the DEC discs are quite small, only 512 bytes. If blocks of this size were written onto tape, then the obligatory inter-block gaps would take up a relatively large proportion of the physical space that is available on the tape. To avoid this wastage, some tapes (those written with the ROLLIN package) lump together eight disc blocks at a time to form one tape block. When reading tapes that have been written in this way, the program STRUC must unpack the tape blocks to get the original disc blocks.

(d) Device images.

The operating system RT-11 has a facility for putting a complete byte-for-byte copy of a device into a file situated on another device. This facility was used by McGregor to put binary images of his floppy discs onto the archive tapes. In order to get any

information out of these 'floppy images', the program STRUC must be able to read and interpret the directory that is held inside the image. Using that directory, it can then locate and open other files within the image.

(e) File names and sub-directories.

In the DEC system, a file has a six-character filename, followed by a three-character 'filename extension'. There is no hierarchy of directories or sub-directories; there is only a single directory for a disc (and no directory for a tape). In the GEC system, a file has an eight-character filename, and there is a hierarchy of directories (called 'catalogues' in GEC terminology). To convert filenames from the DEC regime to the GEC one, the following convention has been adopted: if the full DEC name of the file is "ffffff.eee" then the GEC filename is ".ffffff" and it is placed in a catalogue called ".eee", so that its full GEC name is ".eee.ffffff". The catalogue ".eee" is created if it does not already exist.

(f) Directory listings.

Some archive tapes were externally labelled to indicate their contents. Most of them, though, were unlabelled and their contents unknown. It was therefore essential that STRUC should be able to produce a listing of the contents of a tape. In the case of binary images of whole discs, STRUC is able to read and interpret the existing directory and print out its contents. But in the case of tapes that were written with the RT-11 command COPY, there is no resident and STRUC must create its own by passing along the tape and opening each file in turn. In both cases, a directory is printed out in standard RT-11 format. A directory listing has now been produced for each archive tape.

1.2 Radiation data

1.2.1 Introduction

Three data-sets were identified in the Work Plan as being able to serve as a basis for establishing an empirical model of the diffuse field.

The most important of these is the set of scanning-pyranometer recordings that was initiated by D.A. Svendsen at UCC. These are measurements of global irradiances in 25 planes taken about every minute. The data-set spans a total of almost 24 months, predominantly in the Summer. Simultaneous measurements of beam irradiance were also taken for about half of these months, although the data for most of that period have not been found. Contemporary observations of the duration of sunshine and the amount and type of cloud were recorded hourly at Rhoose airport about eighteen kilometres away.

1.2.2 Nature of the data

In 1978, Svendsen set up an apparatus consisting of three thermopile pyranometers, which were tilted at angles of 45°, 90°, and 135° from the zenith and mounted on a shaft driven

by a stepper motor. The cluster of three pyranometers moved azimuthally through eight positions every fifty seconds. At each position, the cluster would stand still for almost five seconds and the irradiance reading would be logged automatically. This apparatus was situated on the roof of the tall Mining Building at the UCC site on Newport Road. Photographs 1 to 3 illustrate the apparatus: Photo. 1 shows the scanning pyranometers; Photos. 2 and 3 show respectively the fixed Summer and Winter shade rings.

On the same roof, two stationary pyranometers recorded the global irradiances in the horizontal plane and in the South-facing 45°-tilt plane. (The latter instrument provided an essential check on the accuracy of the scanning pyranometers, as its reading should always match that of the scanning 45° pyranometer when in the the South orientation.) A sixth pyranometer, also stationary, registered the diffuse irradiance in the horizontal plane, using a broad shade ring of Svendsen's own design. Logging of global irradiance in the horizontal plane started in November 1976, followed shortly afterwards, in March 1977, by the logging of global irradiance in the 45°-tilt plane and of horizontal diffuse irradiance.

Svendsen made two shading-rings. Each ring obscured the Sun for an entire season - one ring for Summer, one for Winter. His intention was to obviate the need for the almost daily adjustments that conventional narrow shade-rings need. Inevitably, each ring will also obscure a part of the diffuse irradiance. A correction factor must therefore be applied to give an better value of the diffuse irradiance. Svendsen evaluated this correction factor on the assumption of an isotropic diffuse field; the value was 1.38 for both rings (Svendsen 1977). He incorporated the correction factor for the two rings into the program RTOP, which generates PPP files from the intermediate RRR files. Therefore, all the values of diffuse irradiance contained in the PPP files will have been corrected, albeit on the assumption of isotropy (which will have caused the corrected values to be slightly lower than the true values).

In the later months of the project, beam irradiance in the plane normal to the Sun's beam was also recorded on the roof, using an Eppley pyrheliometer. Provided that the tracking is maintained accurately, the pyrheliometer offers the most accurate means of determining the beam and diffuse composition of the global irradiance.

Svendsen left UCC a few months after the scanning pyranometers came into operation. The monitoring was then taken over by A.A. Green and later by J. McGregor, but the archived measurements will be referred to as the "Svendsen data". Details of the apparatus and the basic pre-processing of the logged data are given in Svendsen's eleven Departmental Reports which are listed in his Final Report (Svendsen 1978, Rept 431). A chronological account of this monitoring programme, including technical information omitted from Svendsen's reports, is given by Lloyd (1985, Rept 1204). The latter report has proven to be highly useful in interpreting the contents of the data-files. Until it is replaced with a more formal manual, it must be regarded as an essential adjunct to the data files.

The Svendsen data-files are made up of records, each record containing the values of all the monitored quantities for a single 50-second scanning cycle. Each such record comprises two components: first, 24 values of global irradiance from pyranometers that scanned through planes at 3 tilts and 8 azimuths; second, 3 values of irradiances (two global, one diffuse) from the stationary sensors. When they are first recorded, the irradiances are expressed in 'logger counts', or 'logger units'. These are the uncalibrated output from the analogue-to-digital converter, which is fed with the EMF signal of each pyranometer in turn. Files of these raw data were given the arbitrary name "QQQ". We shall not need to use the QQQ files, unless there is some reason to doubt the calibration that has been applied to them and to re-calibrate the raw data. The data pass through an intermediate file format called "RRR", which we shall not consider either. Finally, the finished data (as far as Svendsen was concerned) were put into "PPP" files - again, an arbitrary designation.

In passing from the RRR stage to the final PPP stage, the irradiance values were converted to engineering units (Wm^{-2}) and subjected to two of the 'engineering' corrections. First, the effect of the pyranometers' slow time-response was removed on the assumption of a first-order linear response, using previously measured time constants. Second, the effect of the pyranometers' temperature response was partially removed by applying a fixed correction factor that was presumably based on a long-term average value of ambient temperature.

1.2.3 File names

(a) Svendsen's file names

Svendsen gave each data-file a name that indicated the date on which the contents of the file began (Svendsen 1977, Rept 357, p 4). The basic form of this filename syntax is "MMMddy", in which: "MMM" stands for a three-letter abbreviation of a month name, e.g. "JUL"; "dd" is the two-digit day number within the month, and "y" stands for a single-digit year number, e.g. "8" for 1978. For example, the first file has the name "NOV106", which indicates that it holds sequential data that began on 10 November 1976. On some occasions, however, Svendsen found it necessary to create more than one file on a given day. To distinguish these files, he apparently interposed digits within the filename. Sometimes, he seems to have followed this scheme:

first file:	mmmddy.QQQ	e.g. NOV106.QQQ
second file:	mmmydd.QQQ	e.g. NOV610.QQQ
third file:	mmmdyd.QQQ	e.g. NOV160.QQQ

But at other times, he seems to have interposed digits at random, so that one must refer to his chronological list of filenames (Svendsen 1978, Rept 430). File names conforming to Svendsen's convention are found for dates up to 13 September 1979. It is therefore inferred that Green continued using this format when he took over in August 1978, and that McGregor did so when he took over in July 1979. A break occurred in the irradiance monitoring after 10 September 1979. When the monitoring restarted, McGregor

used a different filename format. His reasons for doing so are not known.

Data were recorded onto cassettes, and each cassette was read into a single, separate QQQ file on a disc in the DEC PDP-11 computer. Each cassette normally held roughly one day's data, although sometimes a cassette would be put into the logger at sunset of one day, ready to take the next day's data. By looking at the names of the data-files, one can ascertain almost precisely which days are included in the extant set of raw data.

The QQQ files were converted one-by-one into RRR files, retaining their file names. Hence one can ascertain precisely which days are included in the extant sets of intermediate processed data.

The PPP files were concatenated into large files that covered several days, making it much easier to handle the data. It appears to have been McGregor who put together files containing concatenated data from whole months. These files had names of the form "mmmPPP.PPP" (for the period July 1978 to May 1979) or "PPPmmm.PPP" (for the period July and August 1979), in which "mmm" stands for the abbreviated month name, e.g. "JULPPP.PPP". (No PPP files have been found for periods after August 1979.)

For three months, shorter PPP files have been found: JULYA and JULYB for July 1978; AUGA, AUGB, and AUGC for August 1978; SEPPP1, SEPPP2, SEPPP3, SEPPP4 for September 1978. It would that seem that these were left over when McGregor had finished creating the monthly PPP files.

(b) McGregor's first file names

At first, McGregor followed Svendsen's conventions for naming files. In September 1979, however, there was a break in the running of the scanning pyranometers and Svendsen's convention was not used again. (The last file in the old convention was SEP109, for 10 September 1979.)

No written log has been found for the period 11 September 1979 to 1 May 1981. Some references in McGregor's monthly SEU reports suggest that radiation monitoring continued for at least some part of this period. Moreover there is a set of data files with names of the form "DYdddd.DAT", which seem to have been created at some time during January to April 1981. The "dddd" is almost certainly a day number, ranging from 0 to 81; if so, then each DY file spans about one week, and a total of about 88 days are covered. Although the filename extension is "DAT", the files have been examined and found to be QQQ files. They contain only three streams of irradiance data, and must therefore be assumed to be records of the three fixed pyranometers. It is not certain when day number 0 was, but it was probably in the second week of February 1981.

Since the DY files do not contain readings from the scanning pyranometers, we shall not consider them further.

(c) McGregor's second file-names

On 2 May 1981, the scanning pyranometers were restarted, and the QQQ data from the Svendsen apparatus were put into files with names of the format "MOddhh.QQQ", in which "dd" stands for the day number (ranging from 0 to 56) and "hh" stands for the hour number at which the data cassette was loaded. At this time, a second logger was recording irradiance values from an Eppley pyrheliometer and an Eppley PSP pyranometer. Files resulting from this second logger were given names of the form "STddhh.QQQ", in which "dd" and "hh" have the same significance as before. (It is thought that the codes "MO" and "ST" stand for "MOving" and "STationary".) Data files with names conforming to this convention are found for the period from 2 May 1981 to 15 June 1981. Only QQQ files have been found for those two months: no RRR or PPP files have been seen complying with either the "MO" or "ST" conventions.

The data from these two months are important, as they are the only simultaneous recordings from the scanning pyranometers and the tracking pyrheliometer.

Monitoring of the scanning pyranometers continued until the end of July 1981, but no files have been found for the period from 16 June until the of July. Then, the Svendsen apparatus was dismantled and never ran again.

(d) McGregor's third file-name format

On 10 August 1981, a new monitoring programme began, with a new file format: "ssdddd.QQQ", in which: "ss" indicates which of two data loggers was used - "L1" and "L2"; and "dddd" is the day number.

The new monitoring programme included measurements of global irradiance in the horizontal plane and in the South-facing 45-degree tilt plane, and of the beam irradiance at normal incidence. These data carry only limited information about the diffuse field and therefore will not be considered further.

1.2.4 Extant data files

The data tapes have been systematically examined and a record made of what files of radiation data are present. Figure 1.1 shows the results of this exercise in a bar chart. Solid boxes represent data whose files have been found and identified. Dashed boxes represent data which might exist but which have not been found or have not been identified. E.g. the almost 90 days of data in the DY files are marked in dashed boxes, because their precise dates are unknown. (The program and data file that were used to create the graph are listed in Annexe B.)

It will be seen that there is an almost continuous record of data from the scanning pyranometers over the fifteen months from April 1978 to September 1979.

A comparison has been made between the data files found on the tape and the written logs for the almost three years from 10

November 1976 up to 21 August 1979. It has been established that all the data which were recorded during that period are still in existence on the tapes. Unfortunately, no such comparison can be made for the two years from 22 August 1979 to 9 August 1981 because no written log has been found. The scanning apparatus was finally shut down on or just before 9 August 1981.

At the top of the diagram is an axis showing the existence of PPP files for the fourteen months from July 1978 to August 1979 inclusive, and the existence of DDD files for the ten months from July 1978 to April 1979 inclusive.

In the first instance, analysis must be restricted to the PPP or DDD data, because only they have been corrected and calibrated. Hence we have: (a) fourteen months of immediately usable data from the scanning pyranometers; (b) about a further month of scanning-pyranometer that must be pre-processed before use; and (c) about twelve months of data from the fixed pyranometers, which again must be pre-processed.

Table 1.I lists the PPP files that were found on the archive tapes, with the starting time ofthe first and last continuous segment of data in each file. These original PPP files do not have a one-to-one correspondence with calender months. Some months (e.g. September 1978) have several PPP files. Some PPP files contain odd days from adjacent months (e.g. MARPPP, besides containing March 1979, includes one and a half days from February 1979 and part of a day from April). The program PPPEDIT has been written to tidy up the file. The result is seen in the right-hand half of Table 1.I, which lists the new PPP files and their contents. These files are all held on disc at present, but will be copied to archive tapes in the future.

1.3 Pyranometric time-response data

One of the archive tapes held a series of files called RUN1.DAT to RUN13.DAT, each file comprising a short header followed by two columns of numbers. On inspecting the files, it was found that they were the output from Svendsen's program GENTAU.FOR, which was also found on the tape. This program was used to used to control and monitor a laboratory experiment to determine the time response of the pyranometers. The files RUN1 to RUN2 therefore contained the measured responses of pyranometers to ramp- and step-function changes in the incident irradiance. In each file, the first column of numbers is the time since the start of the run, and the second is the millivolt reading from the pyranometer.

These data files are potentially very useful, especially as Svendsen did not write up this part of his work in a Departmental Report. The time responses of pyranometers are discussed in the next section.

1.4 Software written by Svendsen

Svendsen wrote numerous computer programs for processing the radiation data on the DEC PDP-11. His reports, however, contain very few program listings and no detailed documentation.

What appears to be a full complement of Svendsen's programs has been retrieved from the archive tapes using the file-restructuring package, STRUC. This set comprises 31 FORTRAN programs and 2 programs written in the assembly language, MACRO-11.

One program was RTOP, which converts data files from the intermediate RRR stage to the usable PPP stage. This program file is vitally important, as it allows us to ascertain precisely what corrections were applied to the radiation data. Svendsen left only very brief documentation of the program. For example, he did not state in his reports whether or not the correction for the pyranometric time response had been applied to the data in the PPP files. But it can be seen from an inspection of the program RTOP that it has. The main functions of the program RTOP are: to filter out the time responses of the pyranometers; apply the temperature correction; and convert the irradiance values from crude logger counts to units of Wm^{-2}. Because of its importance, the program is listed in Annexe D.

Other program files that are informative are: QTOR.FOR, a program for converting QQQ files to RRR files; CASS.MAC, a program for reading cassette tapes; SOUTH.FOR, a program for determing due South; GENTAU.FOR, a program for determining pyranometric time responses.

2. Pyranometric time responses

2.1 Introduction

In the scanning apparatus, each pyranometer remained in a particular position for only five seconds. This is much less than the settling time of the sensor. If the irradiance on the sensor changes appreciably just before the reading is taken, then the irradiance reading may have a non-negligible error. There are two ways in which the irradiance on the scanning pyranometers can change: first, when the pyranometer moves from one direction to another, the irradiance on its sensor will usually change appreciably; second, when the Sun becomes obscured by a cloud (or ceases to be obscured), the irradiance may fall greatly (or rise greatly).

In either of these two situations, the irradiance reading will differ appreciably from the true irradiance at that moment. The action to be taken to correct the data is different for the two situations.

2.2 Irradiance changes due to pyranometer motion

If the Sun is visible and beam radiation is incident on the pyranometer cluster, then most of the movements of the cluster will yield large, systematic changes in the irradiance between azimuths, due to changes in the incidence angle of the solar beam. The irradiance changes will be systematic, in the sense that at a particular time of day the irradiance in a particular plane will be consistently higher or lower than that in another plane. Therefore the error in the irradiance reading, which is due to the slow response of the sensor, will be consistently

negative or positive. Hence these appreciable errors will not average out.

The time response of the pyranometer is usually modelled by a linear differential equation, so that the response curve is a simple exponential curve which is completely determined by the sensor's time constant. In principle (barring numerical errors), therefore, we can remove the error due to the slow response by applying the inverse of the exponential response function.

In order to be able to do this, Svendsen carried out laboratory tests on the pyranometers to determine their time constants. Those values were then fed into the computer program RTOP and applied to the irradiance data when converting RRR files to PPP files.

2.3 Irradiance changes due to solar obscuration

Within the five seconds for which the pyranometer cluster is stationary, the Sun may become obscured by a cloud, or cease to be obscured by a cloud. In either situation, the irradiance reading will have an appreciable error, owing to the sensor's slow response. Since a pyranometer will remain stationary in a particular plane for 5 seconds in every 50, it follows that one tenth of all the solar obscurationas and cessations of solar obscuration will occur while the irradiance in that particular plane is being measured. Therefore, this source of error cannot be dismissed as being infrequent. It will be quite frequent when the sky has a large amount of broken cloud.

The significance of this error depends on what the data are to be used for.

(a) Integrated irradiations

If the irradiances are integrated over whole days, or hours (or, probably, even for shorter periods), then the errors will balance and the resulting error will be negligible. This is because the negative error occurring when the irradiance rises is matched by a positive error when the irradiance falls. Since the upward and downward exponential curves are approximately the same, the match between negative and positive errors will be fairly good.

(b) Frequency distributions of irradiance

The irradiance values can be analysed into frequency distributions by counting how many times the irradiance falls in each of a number of predefined intervals. The errors due to the slow time response will cause some data points to be counted in the wrong bins. These will, however, be matched by approximately equal numbers of data points being mis-counted in the opposite direction. Therefore, we may expect that a frequency distribution of irradiance formed over a period of a day or more will have a negligible error due to the slow pyranometric response.

(c) Comparisons of contemporaneous irradiances

For some purposes, we may wish to compare irradiance values that were measured 'at the same time'. An example of this is the comparison of readings from (i) the pyranometer fixed at a tilt of 45°, facing South, and (ii) readings from the scanning pyranometer at the same tilt angle which, once a minute, faces South. Ideally, the readings should be the same, and discrepancies might suggest an inaccuracy in one or both of the pyranometers. (The error is more likely to be in the reading from the scanning pyranometer because of the greater degree of processing to which its readings are subjected.)

2.4 Determination of time constants

It was mentioned above that Svendsen determined the time constants of the pyranometers in a laboratory experiment. The experiment consisted in applying beam irradiance to a pyranometer and monitoring its irradiance reading, the irradiance being made to follow both step- and ramp-functions. It is not clear why the ramp-function was used, as the data generated with the step function are sufficient to characterise the instrument's time response, and are simpler to process. The irradiance changes and the logging of the readings were both under the control of a PDP-11 computer, so that precise timing of the data could be obtained.

Unfortunately, Svendsen did not write up this work in a formal report, and no informal notes have been found on it. In the examination of the data tapes, however, a number of files were found that relate to this work. These are of two sorts: files of irradiance measurements, which were logged with the PDP-11 computer; and FORTRAN programs for processing the data. The programs are listed in Annexe C. On this basis, we can repeat Svendsen's curve fitting and see how well the fitted curves match the observed time responses, and in particular whether the time constants are the same for both increases and decreases in irradiance. (We cannot readily reproduce the experiment itself, as the apparatus has long since been dismantled and the computer sold. Another apparatus would have to be set up, using a different computer, with new software.)

The control and logging of the experiment was done with the program LOGTAU, which produced files of irradiance values. Another program, GENTAU, fitted a simple exponential curve to those data points. (Judging from the dates of the files (between 10 July and 8 August 1978), Svendsen carried out this work shortly before he left in late August 1978. He therefore probably did not have enough time to report the work in detail.)

Svendsen wrote the program GENTAU in the RT-11 version of FORTRAN and it cannot run on the GEC computer. Therefore, his algorithm has been coded into the program PBLTAU in the GEC version of FORTRAN-77. In addition, this new program plots the measured data, with the fitted curves super-imposed. It also applies a different curve-fitting method, printing out the results of both methods so that they can be compared. The programs LOGTAU, GENTAU, and PBLTAU are listed in Annexe C.1, C.2

and C.8 respectively. The thirteen basic graphs generated by PBLTAU are shown in Figures 2.1 to 2.13.

Both programs, GENTAU and PBLTAU, are designed to use the same files of input data, which were produced by Svendsen's proram LOGTAU in August 1978. There are thirteen known data files, which were named RUN1 to RUN13. In runs 1 to 9, the irradiance followed a ramp function, and in runs 10 to 13 it followed a step function.

For runs 1 to 9 (with the ramp input), the program PBLTAU applies two different methods to fitting the time-response curve. First, it applies the method used by Svendsen, which is to (a) iterate through several possible time-lags from 0.46 to 0.56 seconds, and (b) for each value of the time-lag, minimise the error in the end-point of the exponential curve. (The user then selects whichever of the time-lags gives the lowest error.) Second, the program fits the exponential curve by the method of least-squares.

The purpose of this exercise was to ascertain whether a lower RMS error could be got by explicitly minimising this error statistic. It was found that the difference in RMS error between the two methods is negligible.

None of the data files contains any identification of the pyranometer. There is therefore a problem with deciding which time constants belong to which pyranometers. The time constants computed by the program PBLTAU for runs 1 to 9 are as follows:

	End-point method	RMSE method
Run 01:	5.2360	5.2192
Run 02:	5.2215	5.2105
Run 03:	5.3296	5.3501
Run 04:	5.3371	5.3405
Run 05:	5.2289	5.2233
Run 06:	5.2098	5.2054
Run 07:	5.3078	5.2909
Run 08:	5.2751	5.3510
Run 09:	5.3433	5.4113

The values all lie in the approximate range 5.2 to 5.3, with a mean of 5.28, which suggests that the data relate to a single pyranometer. Svendsen (1978) reports that older Kipp CM5 pyranometers (e.g. those manufactured in 1974) had time constants around 5.4 seconds, whereas later ones (e.g. those of 1977) had smaller time constants around 2.4 seconds. Extant records (collated by Lloyd 1985, Rept 1204) indicate that at the time of the experiments (August 1978), the SEU possessed only seven pyranometers: one dated 1974, three dated 1975, and three dated 1977. Of these, only the first would be expected to have a time constant as large as 5.28 seconds. It is therefore likely that runs 1 to 9 were carried out with the 1974 pyranometer.

In Runs 10 to 13, only the RMSE method is used. Separate fits are made of the exponential curves on the upward and downward steps of the irradiance. This allows us to see whether the upward and downward time constants are different, and it is

found that there are slight differences. To judge the significance of these differences, we can compute the magnitude of the sensor's response at a point in time five seconds after a step. Recall that the scanning pyranometers remain in each particular position for only five seconds. Therefore, the values tabulated below for the response at five seconds give an indication of how much systematic error might be introduced into the corrected data (i.e. PPP data) by using a single value of the time constant, say the mean.

	Time constants (sec)		Response after five seconds		
	Upward	Downward	Upward	Downward	Difference
Run 10:	5.1258	5.0634	62.3%	62.7%	0.4 pp
Run 11:	2.1194	1.9365	90.6%	92.4%	1.8 pp
Run 12:	2.3695	2.2062	87.9%	89.6%	1.7 pp
Run 13:	2.5088	2.3108	86.4%	88.5%	2.1 pp

These figures can be interpreted in the following manner. Consider the pyranometer used in Run 13, and let us suppose that the downward time constant (2.3108) is used in correcting its readings. If the irradiance on this pyranometer suddenly falls from 1200 Wm^{-2} to 200 Wm^{-2} at the start of its five-second monitored period, then the logged reading will be 315 Wm^{-2}, which will be corrected to the true figure of 200 Wm^{-2}. If, on the other hand, the irradiance suddenly rises by 200 Wm^{-2} to 1200 Wm^{-2} at the start of the monitored five seconds, then the logged reading will be 1064 Wm^{-2}. This will be 'corrected' (using the downward time constant) to 1176 Wm^{-2}. Thus the 'correction' leaves an error of -24 Wm^{-2}, which is a relative error of -2%.

Three points should be noted about this error. (a) The magnitude of the error, 24 Wm^{-2}, is comparable to the anisotropic part of the diffuse irradiance. (b) The error will be systematic, so that it will not cancel out in integrated irradiations. (c) The worst case has deliberately been chosen in the worked example given above. The actual errors will be lower than this, for two reasons. First, it would be reasonable to choose the average of the two time constants, rather than one of them - this at once halves the error. Second, a change of irradiance of 1000 Wm^{-2} would be infrequent. We may therefore expect that, in practice, the error due to the asymmetry of time constants will be rather less than 1% of the true irradiance, and will rarely exceed an upper bound of about 3 Wm^{-2}.

When running the program RTOP to correct the irradiance data for the time response, the user must type in the time constants at the terminal. Consequently, these constants are not incorporated into the program. There are, however, three lines in the program that give default values of 2.40, 2.16, and 2.50 seconds to the time constants for the pyranometers tilted at 90°, 45°, and 135°, respectively. These default value are used if the user offers no values for the time constants. In addition, a comment statement in RTOP that gives an example of the command for running the program, and this includes example values of the time constants. These are 2.40, 2.50, and 2.40 seconds.

The only conclusion that can be drawn from this is that we cannot determine the time constants of the scanning pyranometers, but they all seem to have been in the approximate range 2.4 to 2.5 s.

2.5 Non-simple exponential time responses

In the plots of Runs 1 to 9 (Figures 2.1 to 2.9), the predicted pyranometer response fits the measurements very well. In the more stringent case of the step input (Runs 10 to 13, Figures 2.10 to 2.13), however, there is a visible discrepancy between the predicted and measured responses.

The shape of the discrepancy suggests the presence of a time lag between the irradiance step and the start of the pyranometer's response. Using the least-squares criterion, the best value of this hypothetical lag was determined for each run. Figures 2.14 to 2.17 show the predicted pyranometer response with this timelag incorporated. There is little or no improvement in the fit of the curve to the measurements.

Two non-simple exponential functions (a sum of exponentials, and nested exponentials) were tried as possible response functions, but the results were not altogether satisfactory. A major constraint on fitting a usable function to these measurements is that the function must be the solution of a differential equation describing the behaviour of the sensor. Without the differential eqiation, we cannot predict the response of the pyranometer to the ramp-shaped change in irradiance that occurs when the pyranometer moves.

For the present, therefore, we must accept the simple exponential response curves that Svendsen fited.

3. Basic subroutines for analysing Svendsen data

In the course of procesing the archived irradiance data that were recorded with the Svendsen apparatus, certain computations need to be done repeatedly. These have been put into a set of subroutines call collectively DAS3LIB. (The code number 3 is the standard GEC-4090 code for FORTRAN-77.) They will form a core of subroutines for the growing suite of programs that carry out particular analyses of the data. Documentation is given in Appendix A.2.

4. Diurnal profiles of irradiance

4.1 Introduction

A program has been written to generate graphs of irradiance versus time, superimposed on curves representing the physical upper bound of the irradiance, for each of the planes scanned by the Svendsen apparatus. These graphs serve four important purposes. (a) They provide a check on the integrity of the irradiance measurements. It is much quicker to make a visual inspection of the data, and thereby discover any major faults, than it is to write a computer program for applying numerical criteria of validity. (b) They provide a check on the validity

of the computer subroutines that evaluate the terms in the geometry of the irradiances received by the pyranometers. This is consistent with checking the integrity of the measurements, because it is highly unlikely that the measurements and the computations will suffer from an identical fault. (c) The collection of graphs can be used as a catalogue for the data set. If one requires a day having particular weather conditions, e.g. a cloudless day, then such a day can found by looking through the catalogue. Conversely, the weather conditions on a given day can be ascertained from the catalogue. (d) Selected graphs can serve to illustrate qualitative features of the radiation field. This is especially useful when one considers features for which there is not yet any quantitive parameterisation.

4.2 Upper bound for irradiance

An upper bound is required for the physically possible ranges of the irradiances that may be encountered in processing the Svendsen data. One possible choice for an upper bound would be the irradiance at the top of the atmosphere, i.e. the extra-terrestrial irradiance. The extra-terrestrial irradiance in the horizontal plane, G_o, is widely used as an upper bound when dealing with the ground-level irradiance in the same plane, G.

The Svendsen apparatus, however, employed various planes throughout the hemisphere. In some planes, such as the 90° plane facing approximately West, the solar beam is at normal incidence on the pyranometer when the Sun is setting at the natural horizon. In these planes, the extra-terrestrial irradiance is very much higher than the terrestrial irradiance could be in the same plane. This is because the mass of air through which the light is travelling is large at this low solar altitude, and greatly attentuates the beam that is incident on the pyranometer. Therefore, the extra-terrestrial irradiance is not a suitable upper bound for use with the Svendsen data.

An alternative upper bound would be an estimate of the beam irradiance under a cloudless atmosphere with high visibility. Hottel (1976) has presented a simple formula for computing the beam irradiance under a model atmosphere having a visibility of 23 km. This formula is recommended by Duffie and Beckman (1980) and so may be assumed to be widely used. We should prefer to use a simplified formula rather than a more detailed one for two reasons: first, we require only a rough upper bound, not a precise estimate; second, we do not have detailed information about the atmospheric conditions over the Mining Building at the time of the irradiance measurements.

Hottel's (1976) formula is as follows.

$$G_{bn} = G_{on} K_{b1}$$

$$K_{b1} = a_o + a_1 \exp(-km)$$

in which:

G_{bn} = normal-incidence beam irradiance

G_{on} = normal-incidence extra-terrestrial irradiance

K_{b1} = beam transmission through a cloudless atmosphere

m = relative air mass, approximated as $\sec(Z_o)$

Z_o = solar zenith angle

a_o, a_1, and k are Hottel's parameters

The parameters in this formula are evaluated as follows.

$$a_o = r_o \, (0.4237 - 0.00821 \, (6.0 - alt)^2 \,)$$

$$a_1 = r_1 \, (0.5055 + 0.00595 \, (6.5 - alt)^2 \,)$$

$$k = r_k \, (0.2711 + 0.01858 \, (2.5 - alt)^2 \,)$$

in which 'alt' is the altitude of the site above sea level, in kilometres. The 'r' factors are given in the following table

climate type	r_o	r_1	r_k
tropical	0.95	0.98	1.02
mid-latitude summer	0.97	0.99	1.02
sub-arctic summer	0.99	0.99	1.01
mid-latitude	1.03	1.01	1.00

To allow comparison with values of global irradiance, the Hottel beam irradiance must be supplemented with some estimate of the diffuse irradiance. For this purpose, it was simply assumed that the diffuse irradiance in the horizontal plane was 15% of the beam irradiance (which might be expected on a very clear day), and that the anisotropy was such that all the diffuse energy was received at normal incidence in the plane of the pyranometer. Such a high degree of anisotropy is physically impossible, but it gives an upper bound for cloudless skies.

The resulting sum, of Hottel beam and estimated diffuse irradiances, generally follows the profile of the measured irradiances. The match is especially close with cloudless skies. In cloudy skies, however, cloud masses often reflect irradiance into a pyranometer's field of view and thereby push the actual irradiance above the nominal upper bound. Consequently, the sum is not a strict upper bound, but it nonetheless serves as a valid nominal curve for comparing with the measured profile.

4.4 Catalogue of irradiance profiles

A complete catalogue of irradiance profiles from the Svendsen data set has been created as a separate document (Lloyd 1987). This spans the fourteen months from July 1978 to August 1979. For ten of those months, from July 1978 to April 1979, the graphs are complemented by McGregor's tables of hourly and daily irradiations, together with daily values of other meteorological variables recorded at Rhoose airport (McGregor 1980).

The program, TIMEPLOT, which was used to create the profile graphs, is listed in Annexe E.2. Full documentation will be written when time is available.

5. Determination of the beam irradiance

5.1 Introduction

The only major deficiency in the design of Svendsen's apparatus was that the beam irradiance was not measured. Instead, the diffuse irradiance was measured with a shade-ring. Svendsen recognised that a more accurate separation of the beam and diffuse components could be obtained from measurements taken with a tracking device - either a pyrheliometer or a shading ball. But he regarded the financial and time costs of that approach as prohibitive: a tracking system would be expensive to buy; it would need more maintenance than a shade-ring; and most importantly it would require inspection at least once a day in order to ensure correct alignment.

In June 1980, McGregor started running an Eppley pyrheliometer on the roof of the Mining Building. No logged data, however, have been found for the period from June 1980 until the end of April 1981. It is known that there were problems with the scanning pyranometers for at least part of this period, and it is possible that no irradiance measurements were recorded. Data from the scanning pyranometers have been found for the two months May and June 1981, and pyrheliometer data have been found for May 1981. The two sets of data are in separate computer files, as different Microdata loggers were used, and they exist only in the QQQ form and have not passed through the normal pre-processing. For the present study, it was decided not to try to use the data from this period because there was expected to be difficulty in synchronising the two data-sets, and in checking the alignment of the pyrheliometer, besides the work of the normal pre-processing. Moreover, the software that would need to be developed for the single month of May 1981 would of no use in study the twelve months of July 1978 to June 1979. It is intended to process the measurements that were recorded in May 1981 when time is available.

In order to carry out any worth-while analysis of the 1978/79 data, an estimate must be made of the beam irradiance for each data record. The numerical value of the normal-incidence beam irradiance, G_{bn}, must be derived from measurements of global irradiance, G, in twenty-five planes and of the diffuse irradiance, G_d, in the horizontal plane. There are several ways in which these streams of information can be combined to yield an estimate of G_{bn}. The methods that were considered will be described below, including the method that has been adopted. First, however, the nature of the 'diffuse irradiance' measurements will be discussed.

5.2 Measurement of the diffuse irradiance

Svendsen used a pair of shade-rings of his own design and construction (see photograph 2). They are based on the same principle as the rings used by other practitioners, such as the

UK Meteorological Office (Met.O. 1982), but they are much broader than usual. Geometrically, a shade ring is a section of a cylinder whose axis is parallel to the rotational axis of the Earth and passes through the pyranometer. The ring is positioned so that the pyranometer is in the shadow of the ring throughout the day. As the solar declination changes from -23.45° in mid-winter to +23.45° in mid-summer, the position of the Sun's apparent trajectory moves upward across the sky. In order to ensure that the ring keeps the pyranometer shaded, it must be moved from time to time to keep pace with the Sun, the frequency of moving depending on the width of the ring and the time of year.

In order to avoid the need for frequent adjustments to the shade ring, Svendsen employed one ring that was fixed throughout the Winter and another ring that was fixed throughout the Summer. There was sufficient overlap between the two rings to ensure that the pyranometer remained shaded throughout the transition periods when the rings were interchanged. These transition points were 20 April and 23 August. (Nevertheless, there is one occasion (13 May 1979) on which Green reported beam irradiance on the supposedly shaded pyranometer. This might be attributable to a misalignment error; but if that were so then we should expect more frequent occurrences.)

An inherent problem that arises with any shading ring is that part of the diffuse irradiance is excluded, besides the solar beam. Under an imaginary isotropic sky, the narrow shade ring used by the Meteorological Office obscures about 5% of the diffuse irradiance; Svendsen's rings would obscure about 27% of it. (In reality, the diffuse radiance is heightened in the region near the Sun, so these fractions would actually be slightly higher.) In principle, the measurements that are taken with a shade ring can be corrected with a factor that is determined thus: if a shade ring obscures a part G_{dr} of the hemispherical diffuse irradiance G_d, then the correction factor is $G_d/(G_d-G_{dr})$, or $1/(1-G_{dr}/G_d)$.

Note that the proportion of the diffuse irradiance that a shade ring obscures is not simply the relative angular area of the ring as seen by the sensor. It is, rather, the relative angular area weighted by the cosine of the incidence angle of the diffuse radiation. In this connection, it is rather alarming to read that Svendsen (1977, Rept 340) derived the correction for his shade ring from what he calls the "proportion of sky obscured". When one analyses his undocumented computer program, which is listed in that report, however, one sees that what he has mistakenly called the "proportion of sky obscured" is really the correctly calculated proportion of diffuse irradiance osbcured.

Svendsen (1977, Rept 340) expressed uncertainty about his method as follows: "... a static method using rings of greater than usual dimensions has been chosen. This means that it will not need to be adjusted regularly, but will have greater than usual correction factors. This latter point may make the data worthless, but the alternative would be to delay measurements for many months and it is possible that usable data may nevertheless

be obtained." As will be seen below, Svendsen's shade rings do offer a valid means of determining the beam/diffuse composition, but only through rather elaborate data processing which involves referring to measurements of the global irradiance in the equator-facing vertical planes.

5.3 Isotropic shade-ring correction

The isotropic correction factor is the value of $1/(1-G_{dr}/G_d)$ when G_{dr} and G_d are computed from an assumption of an isotropic diffuse field.

Since the method of evaluating these quantities is not detailed in the literature that I have found on the subject, some details will be given here.

The two irradiance terms in this expression are evaluated by integrating the diffuse radiance, weighted by the cosine of the incidence angle, over the area of sky in question. For G_d, the integral is carried out over the entire hemisphere of view; for G_{dr}, the integral is carried out over the region of sky that is obscured by the shade ring. For thin shade rings, such as that used by the Meteorological Office use, it is possible to regard the ring as being infinitesimally wide and to approximate the area integral with a line integral along the centre of the ring. But Schmidt (1976, cited by Steven & Unsworth 1980) has suggested that this approximation is valid only for shade rings in which the ratio of width to radius does not exceed 0.2. But both of Svendsen's rings are much wider than this limit. Therefore, the irradiance G_{dr} must be computed with a true area intregral. This is done by breaking up the area integral into a line integral along the the length of the ring, which in turn is integrated across the width of the ring. For each such line integral, the 'imaginary declination' is what the solar declination would be if the Sun were to follow the path of that line integral. Within the line integral, the position of the integration element is parameterised with an 'imaginary hour angle', which is what the solar hour angle would be if the Sun were in that position and the solar declination were equal to the imaginary declination. The apparent angular area occupied by shade ring is then

$$A_r = \int \int \cos(\mathrm{dec})\ d(\mathrm{dec})\ d(\mathrm{han})$$

in which

A_r = apparent angular area of shade ring
dec = imaginary declination
han = imnaginary hour angle

and in which the integral is carried out over the range of hour angles from dawn to dusk (or from dawn to midday, and multiplied by 2.0), and over the range of declinations implied by the width of the ring. The isotropic irradiance is given by

$$G_{dr} = \int \int \cos(Z(\mathrm{dec},\mathrm{han})).\cos(\mathrm{dec})\ d(\mathrm{dec})\ d(\mathrm{han})$$

in which

Z = zenith angle, computed as
cos(Z) = sin(lat).sin(dec) + cos(lat).cos(dec).cos(han)

in which

lat = latitude of site

A program, called SHADING, has been written to get these quantities. The actual integrations are done by subroutines that have been put into the library DAS3LIB: subroutines HEMINI and HEMINT for the hemispherical integrals, and subroutines RININI and RININT for the 'ring integrals', that is, the integrals over the area of sky obscured by the shade ring.

For Svendsen's rings, the fraction of sky occupied by the ring, A_r, and the relative 'ring irradiance' from an isotropic field, G_{dr}/G_d, are as follows.

	Summer	Winter
A_r:	0.2710	0.3759
G_{dr}/G_d:	0.2745	0.2740

Following Svendsen, the relative 'ring irradiance' is rounded to 0.274 for both rings. This entails a correction factor of 1/(1-0.274) = 1.377, which Svendsen rounded to 1.38: this is the value that is employed in the program RTOP and hence the factor that has been applied to all measurements of diffuse irradiance stored in the PPP files.

In its defence, it must be said that the isotropic correction factor is widely used. Not only Svendsen (1977, Rept 340), but after him McGregor (1980, Rept 616) and Muneer (1987) have all used the isotropic correction factor with the Svendsen data. The isotropic correction is an established practice that was advocated by Blackwell (1954) and Drummond (1956). It is the standard practice described in the Meteorological Office's manual (Met.O. 1982).

Nevertheless, it has been known for many years that the isotropic correction factor systematically underpredicts G_d and hence overpredicts G_b. Painter (1979) and Steven & Unsworth (1980) have investigated the daily average correction factor by different means and found that the true value of the daily average of G_{dr} may be up to 2.3 times the isotropic value. For Svendsen's rings, the isotropic value of G_{dr}/G_d is 0.274, so in this worst case we would have:

isotropic correction factor = 1/(1-0.274) = 1.38
true correction factor = 1/(1-2.3*0.274) = 2.7
relative error = (2.7-1.38)/2.7 = 0.49

which implies that G_d could be under-estimated by up to 49%. To give some typical numerical values to these terms, let us suppose that the measured values of the global and diffuse irradiances in the horizontal plane are G = 800 and G_d 160 Wm^{-2} respectively, and that the solar elevation is 45°. The

value of the diffuse irradiance as 'corrected' by the isotropic assumption is

$$G_d(\text{corr}) = 1.38*160 = 220 \text{ Wm}^{-2}$$

so the computed value of the normal-incidence beam irradiance is

$$G_{bn} = (G - G_d(\text{corr})) / \cos(Z_o)$$
$$= (800-220)/0.7071 = 820 \text{ Wm}^{-2}$$

On the other hand, if we feed into this formula the worst-case true value of the daily shade-ring correction factor, 2.7, then we get a true value of the G_{bn} = 520 Wm^{-2}.

It is clear from this worked example that using the isotropic correction factor can lead to very substantial over-prediction of the beam irradiance. Now, we shall be obliged to use computed values of G_{bn} in determining the anisotropic part of the diffuse irradiance in arbitrary planes. These determinations would be completely swamped by errors of the magnitude that can arise from using the isotropic shade-ring correction. To illustrate the rubbish that can be produced with this simple correction factor of 1.38, predictions were made of the isotropic irradiance in the 24 planes of the scanning pyranometers, using the value of G_{bn} got with the correction factor. The hour around solar noon was chosen, for two cloudless days, so that we shall have a conservative estimate of the spurious values that are generated by the isotropic correction factor. Figures 5.1 and 5.2 show the measured irradiances during the noon hours of 23 March 1979 and 14 July 1978. Figures 5.3 and 5.4 show the discrepancies between the measured irradiances and the isotropic predictions (in which G_{bn} has been derived with the correction factor of 1.38). On the first of the two days (23 March, Figure 5.3), the results are obviously unphysical: the isotropic model substantially over-predicts the irradiance in the sunward direction - which would imply that the radiance around the Sun is appreciably less than the mean radiance over the sky. On the second day (14 July, Figure 5.4), the results are credible. These observations can readily be explained. As was mentioned above, the isotropic correction over-predicts G_{bn}, and the magnitude of the over-prediction is inversely proportional to cos(Z), the cosine of the incidence angle on the horizontal plane. This over-prediction of G_{bn} carries through to an over-prediction of G_t, and the magnitude of the over-prediction is greater on 23 March than on 14 July because the Sun is at a lower elevation on the former date.

5.4 Anisotropic shade-ring correction

Steven (1977) has carried out a thorough study of the diffuse field in cloudless skies, and fitted a formula for computing the radiance at any point in the cloudless sky. (The formula is merely a means of interpolating the mean measured values that are tabulated by Steven (1977); it is not intended to have any physical significance.) In view of the stability of the cloudless-sky radiance field within the range of turbidities encountered in Britain, it might be reasonable to use Steven's radiance function to compute correction factors for Svendsen's

shade rings.

Unsworth & Steven (1980) have set out some shade-ring correction factors based on the above-mentioned function, but their correction factors are daily means, whereas we require appropriate correction factors throughout the day. Moreover, Unsworth & Steven give their correction factors only for much narrower shade rings than Svendsen's.

The anisotropic correction factor can be computed in much the same way that the isotropic factor was computed, but with Steven's relative radiance inserted into the integrand:

$$G_{dr} = \int \int N(Z(han),A(han)).\cos(Z(han)).\cos(dec)\ d(dec)\ d(han)$$

Again, the integration is carried out as a line integral, which is then integrated with respect to the 'imaginary declination'. In this case, however, the line integral must be done numerically rather than analytically.

Three possible methods for computing the radiance, N, present themselves. First, one could set up an array holding the mean measured values of N, and then employ some simple procedure to interpolate between them. Second, one could employ the shorter of the two functions fitted by Steven (1977), which is a functional form that Steven attributes to Dogniaux. Third, one could employ the longer of Steven's two formulae, which is based on orthogonal spherical harmonics. Regarding the short formula, Steven (1977) states that "the maximum error in any individual value was 10%"; and, regarding the long formula, he states that "the maximum error of any individual fitted value was 8%, which is well within the uncertainties of the measured data". As will be seen below, these statements are slightly misleading.

On a first attempt to use Steven's formulae, spurious results were obtained. Therefore, the program ANICHECK was written to investigate the accuracy of the short and long formulae. This program evaluates both formulae for each of the conditions for which Steven tabulated mean measured radiances, and prints out a table comparing the three sets of numbers. It was found that the formulae as given in the Quarterly Journal of the Royal Meteorological Society (Steven & Unsworth 1979) contain four significant misprints. Two were later corrected in the Journal (Steven & Unsworth 1980), but corrections of the other two misprints have apparently not been published. They were found by comparing the Journal paper with Steven's thesis. The formulae that have been used are coded into the FORTRAN subroutines STEVE1 and STEVE2 in the library DAS3LIB. For ease of reference, however, they are also listed below. (The language is FORTRAN 77 as implemented on the GEC 4090 computer.)

As will be seen, the subroutines have been written for ease of reading rather than for computational efficiency. They may be made to run faster in the future, if this is thought to be necessary.

The formula for the angular distance between an arbitrary point and the Sun is not given in Steven's writings. Following

```
C==== Steven's anisotropic model ================================================
C     Using the fitted curves given by M.D. Steven (PhD thesis 1977)
C     The value returned by this function is the normalised ratio:
C        (radiance) / (hemispherical diffuse irradiance in horizontal plane)
C     expressed in units of 1/(pi.ster)
C==== The shorter of the two functions ==========================================
      FUNCTION STEVE1(AZI,ZEN,AZISOL,ZENSOL)
C     AZI and ZEN are the azimuth & zenith angles of evaluation point
C     AZISOL and ZENSOL are the azimuth & zenith angles of the Sun
      IMPLICIT DOUBLE PRECISION (A-H,O-Z)
      DIMENSION D1(4),D2(4),D3(4),D4(4),D5(4)
      DATA
     :D1/ 0.61D0,  0.65D0,  0.73D0,  0.76D0/
     :D2/11.90D0, 10.70D0, 11.10D0, 13.00D0/
     :D3/-2.97D0, -2.82D0, -2.97D0, -3.09D0/
     :D4/-0.12D0, -0.02D0, -0.07D0, -0.17D0/
     :D5/-0.45D0, -0.45D0, -0.48D0, -0.42D0/
     :PI/3.14159265D0/
      DEGFAC=180/PI
C---- Determine the range of solar zenith angles
      IF (ZENSOL.LT.40/DEGFAC) THEN
         IC=1
      ELSE IF (ZENSOL.LT.50/DEGFAC) THEN
         IC=2
      ELSE IF (ZENSOL.LT.60/DEGFAC) THEN
         IC=3
      ELSE
         IC=4
      ENDIF
C---- Determine azimuthal distance of specified point from the Sun
      AZIDIF=ABS(AZI-AZISOL)
      IF (AZIDIF.GT.PI) AZIDIF=2*PI-AZIDIF
C---- Compute angular distance of specified point from the Sun
      SOLDIS=ACOS(COS(ZENSOL)*COS(ZEN)+SIN(ZENSOL)*SIN(ZEN)*COS(AZIDIF))
C---- Evaluate Steven's formula, which comprises three factors:
      STEVE1=
C     The normalisation constant
     :(1/PI)*
C     The circumsolar component
     :( D1(IC)+D2(IC)*EXP(D3(IC)*SOLDIS)+D4(IC)*COS(SOLDIS)**2 )*
C     The circumhorizontal component
     :( 1-EXP(D5(IC)/COS(ZEN)) )
      RETURN
      END
```

```
C==== The longer & more accurate function ====================================
      FUNCTION STEVE2(AZI,ZEN,AZISOL,ZENSOL)
C     AZI and ZEN are the azimuth & zenith angles of evaluation point
C     AZISOL and ZENSOL are the azimuth & zenith angles of the Sun
      IMPLICIT DOUBLE PRECISION (A-H,O-Z)
      DIMENSION C(4,12)
      DATA  C
     :/ 1.10D0,  1.00D0,  0.40D0,  1.60D0,
     :  0.12D0,  0.70D0,  0.90D0,  1.50D0,
     : -2.60D0, -0.70D0,  0.70D0, -1.60D0,
     :  0.04D0,  0.26D0,  0.38D0,  0.42D0,
     : -0.04D0, -0.50D0, -0.70D0, -2.00D0,
     :  1.40D0, -0.04D0, -1.30D0,  0.20D0,
     :  0.01D0,  0.04D0,  0.09D0,  0.12D0,
     : -0.10D0, -0.30D0, -0.30D0, -0.32D0,
     : -0.16D0,  0.01D0,  0.006D0, 0.19D0,
     : -0.60D0,  0.06D0,  0.60D0,  0.10D0,
     :  5.60D0,  5.80D0,  7.00D0,  9.80D0,
     :  2.30D0,  3.90D0,  4.40D0,  5.20D0/
     :PI/3.14159265D0/
      DEGFAC=180/PI
C---- Determine the range of solar zenith angles
      IF (ZENSOL.LT.40/DEGFAC) THEN
         IS=1
      ELSE IF (ZENSOL.LT.50/DEGFAC) THEN
         IS=2
      ELSE IF (ZENSOL.LT.60/DEGFAC) THEN
         IS=3
      ELSE
         IS=4
      ENDIF
C---- Determine azimuthal distance of specified point from the Sun
      AZIDIF=ABS(AZI-AZISOL)
      IF (AZIDIF.GT.PI) AZIDIF=2*PI-AZIDIF
C---- Compute angular distance of specified point from the Sun
      SOLDIS=ACOS(COS(ZENSOL)*COS(ZEN)+SIN(ZENSOL)*SIN(ZEN)*COS(AZIDIF))
C---- Precompute trigonometric functions
      SINZ=SIN(ZEN)
      COSZ=COS(ZEN)
      COS2Z=2*SINZ*COSZ
      COSA=COS(AZIDIF)
      SINA=SIN(AZIDIF)
      SIN2A=2*SINA*COSA
      COS2A=COSA*COSA-SINA*SINA
      COS3A=COSA*COS2A-SINA*SIN2A
C---- Evaluate Steven's long formula
      STEVE2=(1/PI)*
     :( C(IS,1)
     : +C(IS,2)*SINZ*COSA
     : +C(IS,3)*COSZ
     : +C(IS,4)*SINZ*SINZ*COS2A
     : +C(IS,5)*SINZ*COSZ*COSA
     : +C(IS,6)*(3*COSZ*COSZ-1)/2
     : +C(IS,7)*SINZ*SINZ*SINZ*COS3A
     : +C(IS,8)*SINZ*SINZ*COSZ*COS2A
     : +C(IS,9)*SINZ*COSA*(5*COSZ*COSZ-1)
     : +C(IS,10)*(5*COSZ*COSZ*COSZ-3*COSZ)/2
     : +C(IS,11)*EXP(-C(IS,12)*SIN(SOLDIS))
     : )
 100  RETURN
      END
```

Robinson (1966), though, one can derive it by using the following theorem of sph erical geometry: consider a spherical triangle ABC with sides a,b,c; then

$$\cos(a) = \cos(b).\cos(c) + \sin(b).\sin(c).\cos(A)$$

In our case, b,c are the zenith angles of the point and of the Sun, and A is the azimuthal distance between them; a is the angular distance of the point from the Sun.

The output produced by the program ANICHECK is reproduced in Table 5.I. Two forms of error statistic have been computed and are included in the output. First, the maximum absolute error is found for both the short and the long formula. Those maxima are also expressed as relative errors (the 'associated' relative errors). It is seen that the maximum absolute errors, expressed as relative errors, are 13.9% and 8.6%, which are slightly higher than Steven's claimed maxima of 10% and 8%. The reason for these small discrepancies is unknown. Second, the maximum relative errors are also found and printed out; these are 24.4% and 13.3% respectively for the short and long formulae. These are much higher than the errors quoted by Steven. In this respect, Steven's statements (quoted above) might be misleading, for they could easily be read as meaning that the maximum relative errors were 10% and 8%. (Almost certainly, Steven meant that the maximum absolute errors, when expressed as relative errors, were 10% and 8%.)

Using the subroutines STEVE1 and STEVE2, a program called SHADING was written to compute and tabulate the anisotropic shade-ring correction factors.

Because Steven's anisotropic model depends on the position of the Sun, there is no single unique value of the anisotropic shade-ring correction factor. Strictly speaking, the correction factor should be evaluated afresh for each time point. That, however, would be grossly inefficient. Therefore, an array of correction factors was created, covering the full range of possible azimuthal and zenithal positions in steps of about 6^{o}. In order to use these correction factors, a program will discretize the Sun's position to get a pair of integer indices, which are then employed to read off the required correction factor from the array.

When the program SHADING is run, the user can select which model is to be used to compute the correction factors (isotropic, or Steven's anisotropic model computed with the short formula, or ditto computed with the long formula). Tables 5.II to 5.IX are the output from this program.

5.II	Obscured irradiance: Steven's short formula, Summer ring
5.III	Correction factors: Steven's short formula, Summer ring
5.IV	Obscured irradiance: Steven's short formula, Winter ring
5.V	Correction factors: Steven's short formula, Winter ring
5.VI	Obscured irradiance: Steven's long formula, Summer ring
5.VII	Correction factors: Steven's long formula, Summer ring
5.VIII	Obscured irradiance: Steven's long formula, Winter ring
5.IX	Correction factors: Steven's long formula, Winter ring

Differences between the correction factors obtained with the short formula and those obtained with the long formula are small but not negligible.

Figures 5.5 amd 5.6 show the discrepancies between the measured irradiances and those predicted with the isotropic model, using values of G_{bn} got with the anisotropic correction factors; again, these are for the noon hours on 23 March 1979 and 24 July 1978. The graph for 23 March is less unrealistic than the one based on the isotropic correction (Figure 5.3), but it is still not credible.

Steven's anisotropic correction factor has a fundamental disadvantage in the present context, as it is approriate only to cloudless skies, whereas the overwhelming majority of the Svendsen data were recorded under cloudy skies. A further disadvantage is that it presumes a particular model of the diffuse field, whereas it would be preferable to avoid any such assumptions when determining G_{bn}, because the final objective is to employ G_{bn} in determining the anisotropy of the diffuse field.

5.5 Method of simultaneous equations, isotropic

The isotropic model implies the following equations:

$$G = G_{bn}.\cos(Z) + G_d$$

$$G_t = G_{bn}.\cos(Z_t) + R_d.G_d + R_g.alb.G$$

in which

G = global irradiance in horizontal plane
G_t = global irradiance in tilted plane
G_d = diffuse irradiance in horizontal plane
G_{bn} = normal-incidence beam irradiance
Z = solar zenith angle wrt horizontal plane
Z_t = solar zenith angle wrt tilted plane
R_d = tilt-conversion factor for diffuse irradiance
R_g = tilt-conversion factor for ground-reflected irradiance
alb = albedo

Of these quantities, the following are known: the global irradiances G and G_t are measured; Z, Z_t, R_d, and R_g are computable from the geometry of the measuring planes; alb is assumed to be zero for the Svendsen apparatus. This leaves G_d and G_{bn} as the two unknowns. Since we have two simultaneous equations, we can get estimates of G_{bn} and G_d by algebraic manipulation.

This method of splitting up the global irradiance into beam and diffuse components was tried, but was found to be highly unreliable.

5.6 Method of simultaneous equations, bi-isotropic

The approach described above in sub-section 5.5 was refined by splitting the diffuse field into two sectors, within each of

which the radiance was assumed to be isotropic, although no assumption was made about the relative magnitudes of the radiances of the two sectors. We shall refer to this as a "bi-isotropic" model. One sector was defined to be the part of the sky that is obstructed by the shade ring, and the other sector was the remainder of the sky. This bi-isotropic model implies the following equations.

$$G = G_{bn}.\cos(Z) + G_{dr} + G_{du}$$

$$G_t = G_{bn}.\cos(Z_t) + R_{dr}G_{dr} + R_{du}.G_{du} + R_g.alb.G$$

in which

G_{dr} = irradiance from part of sky obscured by shade ring
G_{du} = irradiance from unobscured part of sky
R_{dr} = tilt-conversion factor for ring irradiance
R_{du} = tilt-conversion factor for unobscured irradiance

The quantity G_{du} is measured directly, for it is the diffuse irradiance received by the shaded pyranometer. In this pair of simultaneous equations, therefore, the two unknowns are G_{bn} and G_{dr}. As before, estimates of the two unknowns can be got by algebraic manipulation. The resulting estimates of G_{bn} and G_{dr} are more stable than those obtained with the pure isotropic model, but they are still unsatisfactory.

Note that the geometric factors R_{dr} and R_{du} are not equal to R_d. In the software, they are computed as follows. Integration of the assumed radiance distribution (isotropic in this case) over the obstructed and unobstructed sectors of the sky is precomputed, and the results are stored in a file. This file contains the values of the irradiances from the obstructed and unobstructed sectors, normalised to G_d. (The choice of normalising to G_d was made following Steven's convention.) These values are generated, and stored in the file, for the horizontal plane and for the 45°-tilt South-facing plane. A subsequent program reads in the values of G_{dr}/G_d and G_{tdr}/G_d. For each tilted plane, it can determine the tilt-conversion factors thus:

$$R_{dr} = (G_{tdr}/G_d) / (G_{dr}/G_d)$$

$$R_{du} = (G_{tdu}/G_d) / (G_{du}/G_d)$$

The 45°-tilt South-facing plane was chosen because the fixed pyranometer in this plane is polled by the data logger immediately after the horizontal pyranometer is polled. There would therefore be a minimum likelihood of the Sun becoming hidden by a cloud between the two readings.

5.7 Correlation method

The method of simultaneous equations can be formulated for any pair of planes, not just the pair that was selected above in sub-section 5.6. A number of different combinations were tried, but none of them provided consistently reliable results.

Since each pair of the 24 measurements of G_t yields an estimate G and G_{du}, it is appropriate to use some statistical method to get a 'best' estimate. After some experimenting with this approach without success, a suggestion by R.H. Marshall was found to give promising results. Marshall's suggestion was to regard the problem as one of fitting a straight line by least-squares regression, with G_{bn} and G_{dr} being regarded as the parameters of the line.

Marshall's idea can be formalised as follows. The basic equation of the bi-isotropic model is this:

$$G_t = G_{bn}.\cos(Z_t) + R_{dr}.G_{dr} + R_{du}.G_{du} + R_g.alb.G$$

Now, let X and Y be defined thus:

$$Y = (G_t - R_{du}.G_{du} - R_g.alb.G) / \cos(Z_t)$$

$$X = R_{dr} / \cos(Z_t)$$

then the equation of the line is

$$Y = G_{dr}\ X + G_{bn}$$

in which G_{bn} and G_{dr} are respectively the gradient and intercept of the regression line. Note that the values of R_{dr} and R_{du} must now be computed for each direction; their values are given in Table 5.X.

For each sweep of the scanning pyranometers, a set of 24 pairs of <X,Y> are determined, and a least-squares fit of G_{bn} and G_{dr} is made.

It is found that there is a different regression line for each pyranomter tilt: that is, the 90°-tilt pyranometers yield one gradient and intercept; the 45°-tilt pyranometers yield another; and the 135°-tilt pyranometers another. This is probably caused by the pyranometers seeing different amounts of the darker parts of G_{du}. We must therefore decide which pyranometers to select for inclusion in the regression.

Five vertical pyranometers facing the southern hemisphere were selected for forming the basis of the regression, on the grounds that (a) the shade rings occupy more of the hemispheres of view in these planes than in the other planes, as is shown by the Table 5.X, and (b) the 45°-tilt pyranometers yield unrealistically high values of G_{bn}.

Having such a small basis (only five points), makes the regression vulnerable to disruption by rogue data. Therefore, numerous checks and vetting procedures have been built into the subroutine BEAM, which determines the values of G_{bn} and G_{dr}. This accounts for the large size of the subroutine (about five pages).

Obviously, surfaces must be excluded from the regression if the Sun is not above the surface (i.e. if $\cos(Z_t)$ is not positive). The surfaces that are rejected for this reason can

still be used to get estimates of G_{dr}. The basic equation in this sunless case is

$$G_{dt} = R_{du} \cdot G_{du} + R_{dr} \cdot G_{dr}$$

This can be rearranged to be explicit in G_{dr}, thus:

$$G_{dr} = (G_{dt} - R_{du} \cdot G_{du}) / R_{dr}$$

Hence each sunless plane among the selected five yields an estimate of G_{dr}. These are averaged, and a final value of G_{dr} is computed by interpolating between the sunless-plane estimate and sunlit-plane estimate of G_{dr}.

Figures 5.7 and 5.8 show the angular distributions of the isotropic prediction errors for 23 March and 14 July, respectively, in which the value of G_{bn} has been determined by the regression method that has been described here.

This regression method works well most of the time: the regression line fits the points with little scatter, and the values of G_{bn} and G_{dr} are credible. From time to time, however, the procedure seems to be unable to cope, and errors up to 30 Wm^{-2} have been observed in the back-predicted values of G_t in the vertical planes. There is therefore some further development to be done on this procedure. But, as it stands, the procedure yields values of G_{bn} and G_{dr}, and hence of G_t, whose errors are rather less than the magnitude of the anisotropic part of the diffuse irradiance.

A definitive validation of this method cannot be made until the radiation data from May 1981 have been pre-processed and collated with the simultaneous measurements of G_{bn}.

5.8 A potential application

The method that has been described for estimating G_{bn} and G_{dr} from a set of pyranometer readings is potentially useful elsewhere. There is a continuing need for measurements of the beam / diffuse composition of global irradiance, and this need is normally met with equipment that needs frequent attention (either shade rings or tracking devices). Svendsen's two shade rings, however, can remain unattended for half a year each - and, as has been reported here, their readings can be processed to yield fairly reliable values of G_{bn}. The Svendsen rings would therefore be suitable for use in isolated locations, such as remote villages in non-industrialised countries.

6. Model estimates of irradiation

A central part of the present programme of work is to determine the extent to which published models of the diffuse radiation field can predict the solar irradiation on arbitrarily oriented surfaces under arbitrary weather conditions. To carry out comparisons between the measurements and the model predictions, the program POLEPLOT has been written, which is listed in Annexe E.3. Full documentation will be written when time is available.

The program is designed to produce irradiation estimates of two sorts: first, estimates based on the isotropic model, computed analytically; second, estimates based on arbitrary models, computed numerically. It includes a subroutine HEMINT, which performs a hemispherical integration of the irradiance. Into subroutine HEMINT can be passed the name of any subroutine for evaluating the diffuse radiance at an arbitrary specified point in the sky. At present, three such subroutines have been written - for the isotropic field (subroutine ISONUM), for Steven's anisotropic field using his short formula (STEVE1) and using his long formula (STEVE2). Further such subroutines can readily be written for other models, such as Pokrwoski's. By comparing the irradiation values yielded by the analytic subroutine for the isotropic model (subroutine ISOANA) with those produced by the combination of HEMINT and ISONUM, it has been established that HEMINT works correctly and accurately. The program is therefore ready for including other subroutines for each of the existing correlation formulae for the diffuse field.

The program POLEPLOT evaluates and accumulates the errors in the irradiation estimates, as compared with the measured values. These are then plotted on a polar diagram, in which the angle represents the azimuth, the radial distance represents the magnitude the error (in Wh/m^2), and three plots are given for the three tilt angles of the pyranometer.

Unfortunately, the integration subroutine HEMINT in its present form is prohibitively slow when run on the GEC 4090. The processing of a day's data takes about an hour. It is believed that this is due largely to the unnecessarily high density of integration elements around the zenith of the sky (and the nadir of the ground view). Therefore, another integration scheme has been conceived, in which the integration elements are forced to have approximately constant angular areas. In addition, this integration scheme has its apex coincident with the Sun, rather than with the zenith so that a higher angular resolution will be possible when integrating the irradiance from a strongly anisotropic cirumsolar region (such as that in Pokrowski's model). The equations for this scheme have been derived, and they will be coded into FORTRAN when time is available.

Meanwhile, important results can be extracted from the program with just the analytical isotropic model, ISOANA. A number of error curves produced by the program are discussed in sections 11 and 12 of the present report.

The program also generates plots of the measured irradiation, using the the same graphics subroutines.

7. Alignment error of the pyranometer cluster

Polar graphs of the irradiance were generated using the program POLEPLOT and it was found that the irradiance distributions on apparently cloudless days were asymmetric about the north-south axis. At first, it was thought that the asymmetry could be explained by assuming that there were clouds to the south-east, reflecting appreciable amounts of light. The reflected light would give rise to a higher level of irradiance in the

south-eastern quadrant of the sky. It had to be further assumed that the clouds happened not to pass in front of the Sun, as the irradiance profiles show that the Sun was not obscured. This explanation was regarded as plausible because there is a land mass to the south-east of Cardiff, but a water mass (the Severn Estuary) to the south-west. One would therefore expected a higher degree of cloudiness due to convective Cumuli visible in the south-western direction.

Further study of the few clear days that occurred during the monitoring period (16 July 1978, 2 February 1979, 23 February 1979, and 23 March 1979), however, showed that this hypothesis was not consistent with the irradiance distributions in the early morning and late afternoon. It was found that the irradiance was higher in the period up to 9:00 than in the comparable period after 15:00. Also, the irradiance was concentrated in the south-east sector during the morning and in the south-west quadrant in the afternoon. If the 'cloud hypothesis' were true, then the afternoon irradiance would have been higher than that in the morning, and the morning distribution should have included a large component in the south-west direction.

It was therefore postulated that the cause might have been an error in the alignment of the scanning apparatus. The program POLEPLOT was adapted to allow the user to specify hypothetical errors in the azimuth and tilt of the apparatus, together with the azimuthal direction of the tilt error. These three parameters were varied in an attempt to get the irradiance plots to look symmetric. No perfectly satisfactory combination of values was found by this means.

Svendsen's reports were re-examined, and the following sentence was found: "Recently, a check was made on the azimuth positioning and it was discovered that an error had been made in the estimation of the direction of due South of about 6°" (Svendsen 1978, Rept 431).

It was therefore decided to determine the best value of a hypothetical azimuthal error by an objective method. There was assumed to be no error in the tilt of the apparatus. To this end, the program POLEPLOT was adapted to compute the RMS discrepancy between the measured irradiance and the irradiance computed with a model, for five directions facing nearly South. These directions were selected because they would have the greatest beam irradiances. The supposed error in the azimuth was systematically varied so as to minimise this RMS discrepancy. Underlying this approach is the fact that the values of global irradiance generated by the model (which is the isotropic model of subroutine ISOANA) include a beam component which is computed as

$$G_b = G_{bn}\cos(Z_{to})$$

in which

G_{bn} = the normal-incidence beam irradiance
Z_{to} = incidence angle on the plane

An error in the azimuth angle would thus yield errors in the predicted irradiances. By minimising those errors, we hope to eliminate the error in the azimuth. A source of inaccuracy in this method is the presence of diffuse irradiance anisotropy, but the inaccuracy is expected to be small because (a) the irradiance data have been selected so as to minimise the diffuse irradiance, and (b) there are physical grounds for expecting the diffuse radiance to be symmetric about the North-South axis.

The two days 16 July 1978 and 23 February 1979 were selected. The variation in the RMS errors was examined and found to be consistent with an azimuthal misalignment of approximately 6^{o}. A correction of 6^{o} has therefore been incorporated into the subroutine INIGEO of the library DAS3LIB.

8. Rehabilitation of vertical pyranometer data

8.1 Description of the problem

As was mentioned above, one of the purposes of plotting the irradiance profile for each day was to check for visible errors. An error of this sort was found in the irradiance readings that were logged from the scanning pyranometer tilted at 90^{o}. It was noticed that these irradiance values were smaller than one would expect, by a factor of roughly 2.5, during the periods

(a) From 21 August to 3 December 1978
(excluding the following short periods:
10:30 GMT to 10:42 GMT on 6 October
12:44 GMT to sunset on 7 October
14:39 GMT to sunset on 8 October
sunrise to sunset on 11 to 13 October inclusive
12:20 GMT to sunset on 27 October).
(b) From 14 to 30 January 1979.
(c) From 9 April to 22 June 1979.

No data were logged between 22 June and 25 June 1979. Thereafter, the irradiance readings for this channel continue to be smaller than one would expect, but by a factor of roughly 1.5, during the period

(d) From 25 June to 20 July 1979.

Data that were logged after the latter date have yet to be examined. (Logging continued until September 1979, but measurements for August and September 1979 are available only as raw data in QQQ files; they have yet to be converted to PPP data.)

The problem can be illustrated with irradiance profiles from comparable days before and after one of the dates when the fault appears. Figures 8.1 and 8.2 show the irradiance profiles for 8 and 28 April 1979 respectively. That these two days have similar weather conditions is evident from the shapes of their irradiance profiles. Moreover, the irradiance levels logged from pyranometers other than the 90^{o}-tilt one are of similar magnitudes on the two days. Nevertheless, the irradiance readings logged from the 90^{o}-tilt pyranometer are much smaller

on 28 April than on 8 April.

The range of possible 'rehabilitation factors', which might correct the logged data to their true values, can be guessed at from Figures 8.3 to 8.5. These show the irradiance profile for 28 April 1979 with possible 'rehabilitation factors' of 2.0, 2.5, and 3.0 multiplying the logged irradiance values. Comparing them with Figure 8.1, we see that the magnitude of the reduction must lie between about 2.5 and 3.0.

According to the notes collated by Lloyd (1985, Rept 1204), McGregor noticed the occurrence of this fault during the period from 21 August to 3 December 1978, but was unable to ascertain its cause. He therefore applied a correction factor of 2.5 to the affected data, which factor he acknowledged to be a subjective guess based on a visual inspection of the irradiance profiles. The irradiation values tabulated by McGregor (1980, Rept 616) include this correction for the above-mentioned periods up to 3 December 1978.

It seems that McGregor did not notice the return of the fault in April 1979. If we select from his tables of hourly irradiations the data for 8 and 28 April 1979 (which are the two days that were considered above) and compare them, we see that his irradiation values for 28 April are much less than the corresponding ones for 8 April.

Obviously it is unsatisfactory to apply subjective corrections of such a large magnitude. This is especially so when the radiation quantity that is of interest to us, namely the anisotropic part of the diffuse irradiance, is much smaller than this correction. An error analysis is given below, to ascertain the magnitude of the problem.

Let us suppose that the irradiance values have been reduced by a factor of X, and let us suppose that we adopt a slightly inaccurate rehabilitation factor of $(X+E_x)$. For example, X might be 2.6 and E_x might be -0.1. So

$$G_t = X\ G_{tl}$$
$$G_{tr} = (X + E)\ G_{tl}$$
$$= G_t + E\ G_{tl}$$

in which

G_t = true value of the irradiance
G_{tl} = logged value of the irradiance
G_{tr} = rehabilitated value of the irradiance

The anisotropic part of the diffuse irradiance is defined thus:

$$G_{ta} = G_t - G_{ti}$$

in which

G_{ti} = estimate of G_t by the assumption of isotropy
G_{ta} = anisotropic part of the diffuse irradiance

In practice, however, it will be computed thus:

$$\begin{aligned} G_{tac} &= G_{tr} - G_{ti} \\ &= (G_t + E_x\, G_{tl}) - G_{ti} \\ &= (G_t - G_{ti}) + E_x\, G_{tl} \\ &= G_{ta} + E_x\, G_{tl} \\ &= G_{ta} + (E_x/X)\, G_t \end{aligned}$$

Hence the error in the computed value of the anisotropic part of the diffuse irradiance is $+(E_x/X).G_t$. Some typical values can be given to these terms, as follows. The 90°-tilt irradiance G_t might be about 900 Wm^{-2} in a South-facing direction; the reduction factor X is probably 2.6; the value of E_x might be -0.1. Hence the error in G_{tac} will be -36 Wm^{-2}, which is -4% of G_t.

In contrast, the true value of G_{tac} would probably be about 5% of G_t in the sunward direction and about -20% of G_t in the anti-sunward direction. So the error in G_{tac} due to this problem could be substantial, although it would not swamp the value of G_{tac}.

It is clear from this brief analysis of the errors that the accuracy of the rehabilitation is important. It would be unscientific to derive results from the faulty data (which constitute almost half of the data set) without first obtaining an objective rehabilitation factor.

8.2 Possible causes of the problem

The following hypotheses have been considered as possible explanations of the anomalous reduction of the irradiance values recorded from the 90°-tilt pyranometer.

(a) Fault in the pyranometer

The sensor of a pyranometer becomes degraded with time, both through oxidation of the surface material and though damage caused by mechanical shocks. As the scanning pyranometers were being subjected to mechanical impulses produced by a stepper motor throughout each day, it might not be surprising if the delicate black paint on the thermopile were to become detached. If this had happened, then the irradiance reading would have been reduced approximately proportionally to the area of the detached paint. Hence about two thirds of the paint would have had to have fallen off. Since this pyranometer was tilted at 90°, the detached paint would easily fall away from the thermopile.

Although this hypothesis seems plausible, there are several facts that make it improbable. (i) The fault was intermittent: it disappeared abruptly on several occasions and throughout the period 4 December 1978 to 8 April 1979. This excludes chemical degradation of the sensor, which would be irreversible. The intermittency can be reconciled with the 'flaking paint' hypothesis only by supposing that the flake remained hinged onto the sensor and occasionally swung back into place. (ii) A physical defect of this magnitude would have been noticed by Green, who inspected the apparatus every day during this period.

It would also have been noticed by McGregor during the course of a subsequent laboratory investigation which included this pyranometer, CM5 77-3538 (McGregor 1983, Rept 1027). (iii) An identical Kipp CM5 pyranometer (which had become defunct) has been dismantled by Lloyd and Marshall, who found that the black paint adheres strongly to the individual strips of the thermopile and cannot be expected to come off in large flakes. The 'flaking paint' hypothesis has therefore been rejected.

One of the two electrical wires that connect the pyranometer to the data logger could have been knocked loose, e.g. by the mechanical shocks of the stepper motor. This could produce an intermittent fault, with the wires occasionally being knocked back into contact. It is very difficult to see how this could reduce the signal strength by a fixed factor of about 2.5. If the electrical circuit was broken, or if the detached wire somehow produced a short ciruit, the signal would be zero as there would be no potential difference across the wires. Each of the two wires was made up of several strands, and some of the strands could have become disconnected. If that had happened, the signal would have been unaffected because each strand can carry the very small current involved. Therefore the 'detached wire' hypothesis has been rejected.

Water could have penetrated the interior space of the pyranometer and corroded the wires, thereby increasing the contact resistance and conceivably reducing the signal. This hypothesis has been rejected because (i) the corrosion would have been irreversible and therefore not capable of causing an intermittent error; (ii) the water would have yielded condensation on the outer glass dome (which encloses a space that is continuous with that containing the wires), and Green would have observed it - but Green made no mention of condensation in this pyranometer.

Concern was expressed by McGregor that solar heating of the backs of the pyranometers would introduce errors in the irradiance readings. This source of error cannot, however, explain the observed reduction in the irradiance readings because: (i) the reduction occurs during overcasts when the pyranometer body would not have been heated by sunlight; and (ii) the other scanning pyranometers would have been subject to a comparable degree of solar heating, but they did not exhibit a spurious depression of readings of the magnitude that we are considering.

The presence of an obstruction on the roof would diminish the solar irradiance, but it would have had an all-or-nothing effect on the solar beam - either casting its shadow on the pyranometer or else not affecting the beam at all. An optical barrier on the pyranometer's glass dome would reduce the detected irradiance, e.g. bird excrement, or condensation on the inside. Two facts make this hypothesis improbable: (i) the contamination would not be uniformly spread over the surface of the glass dome, and hence would not give rise to the constant reduction factor which is observed for all angles of solar incidence; (ii) Green would have noticed, recorded, and removed any such contamination (as he did with the 45^{o}-tilt pyranometer, which occasionally

suffered from condensation). Therefore the 'obstruction hypothesis' and the 'optical barrier hypothesis' have been rejected.

(b) Operator error during pre-processing

Uncalibrated data in RRR files were converted from logger units (LU) into engineering units (Wm^{-2}) with the program RTOP. The program also corrected the data for the slow time responses of the pyranometers. When running this program, the user must type in at the terminal the time constant of each of the scanning pyranometers. If the operator were to type in a grossly wrong time constant for the pyranometer tilted at 90^{o}, then the resulting irradiance values could have large errors.

This hypothesis can be tested by comparing the raw data (in the QQQ files) for the two representative days, 8 and 28 April 1979. Figure 8.6 shows occurrence frequency distributions for the raw data of the two days. It is evident that the values for 28 April are systematically lower, by roughly two thirds, than those for 8 April. Hence we can exclude the possibility that the reduction of the irradiance values might have arisen in the pre-processing of the data on the computer.

(c) Fault in the logger

The mechanism of the data logger can be regarded as comprising two parts: first, the circuitry that is common to all of the data channels (which includes most of the logger electronics, and covers the digitisation of analogue signals and the writing of data onto the tape); second, the amplifier cards, each of which is an independent circuit applying to only one channel. The error that we are considering here occurred in only one channel (that of the 90^{o}-tilt pyranometer), so it is more likely to have arisen in the amplifier card corresponding to that channel (namely DAS1).

The logger does, however, process the channels sequentially and it is conceivable that it could malprocess the 90^{o}-tilt channel owing to this channel's being the first of the pyranometer channels. There is thus a possibility, albeit unlikely, that the fault could have been in the common circuitry rather than in the channel-specific circuitry.

The action of the amplifier cards that are used with the pyranometers is simply to magnify the voltage signal. The pyranometer produces an EMF of up to about 15 mV, whereas the A-to-D converter in the logger operates on voltages in the range up to 3.333 V (= 200 LU) (or up to 4.0 V (= 240 LU) if the normal voltage clamping is switched off).

(c-1) Digital fault in the logger

The digitised signal exists only in the logger circuitry that is common to all the channels, and does not exist in the amplifier cards. This at once makes the hypothesis of a digital fault unlikely, as it would have to be somehow selective of just the one channel.

In 1983, a fault was found in one of the clock cards in the logger, in which individual binary digits ('bits') in the digital register would intermittently cease to function. This fault had the effect of reducing the logged value by an amount equal to some power of two. For instance, if the most significant bit of the 8-bit register were to become dead, then a subtrahend of 2^7 = 128 would be taken from all values that were above 128. This, however, would leave a discontinuous jump in the time profile of the logged quantity. (Imagine a quantity increasing gradually from 120 to 140 logger units. Upon reaching 128, the logged value would suddenly drop to zero and then resume its continuous increase until it reached 12 (=140-128).) No such discontinuity is present in the 90°-tilt irradiance data. Therefore, this hypothesis has been discarded.

Another possibility is that a fault might cause the logged readings to be shifted by one or two binary digits. A rightward shift of one bit would reduce the logged reading by a factor of 2; a shift of two bits would reduce it by a factor of 4. It was inferred above (from Figures 8.1 and 8.3 to 8.5), that the rehabilitation factor must lie in the approximate range 2.5 to 3.0. We can hence be sure that the error is not a reduction by a factor of 2 or 4. Therefore the hypothesis that the fault was caused by a binary shift has been rejected.

(c-2) Analogue fault in the logger

The amplifier cards that were used with the scanning pyranometers were those built by Svendsen himself and labelled DAS1 to DAS3. The generic circuit diagram is given as Drawing 19 of Svendsen (1978, Rept 427). (This is slightly different from the equivalent circuit used by Microdata.) The fixed resistors R1 and R2 of this circuit were selected to produce the required gain in the amplifier circuit.

In several of the signal conditioning cards that have been used during the course of the SEU's radiation monitoring, the resistances that determine the amplifier gain have been modified by the addition of extra resistors soldered together in parallel or in series. Consequently, this element is sometimes physically fragile. It is the only element of the signal conditioning cards that has suffered visible damage through normal use, and damage in this element has been noticed in several cards.

The cards DAS1 to DAS3 have been visually inspected (in 1987) and the resistors in elements R1 and R2 were found to be securely mounted. Indeed, all the components in DAS1 were found to be intact and an electrical test showed that its calibration is still 13 LU/mV, which is Svendsen's specification. (The card DAS3 was also fully intact and correctly calibrated. But DAS2 has been damaged by a mechanical impact, and no longer functions correctly.) All the resistors in the DAS cards are a high quality and unlikely to become 'open circuit' spontaneously.

The hypothesis of a fault in the amplifier card, DAS1, has therefore been rejected.

(d) Operator error with the data logger

There is a possibility that a different amplifier card was inserted into the slot for the 90°-tilt pyranometer, without this fact being recorded in Green's notes.

First, let us try to determine which amplifier should have been in the slot for the 90°-tilt pyranometer.

A brief digression is necessary. Cassette tapes were read on the PDP-11 computer with a program called CASS. This program allowed the operator to insert a code number into the QQQ data file. Svendsen used this code number to indicate the configuration of his apparatus, and in particular which set of signal-conditioning cards were in use. Code numbers 10 onwards were used with the data that were logged from the scanning-pyranometer apparatus. For the three code numbers 10, 11, and 12, the configuration of signal conditioning cards remained unchanged. Code number 13 was skipped. With code number 14, the calibrations of the three custom-built cards, DAS1, DAS2, and DAS3 were changed (the amplification being reduced, presumably to diminish the risk of the irradiance readings going out of range). The two configurations are listed below.

		Code number: 10-12	14
		LU/mV	LU/mV
Channel 1: Gts(90°)	card DAS1	18	13
Channel 2: Gts(45°)	card DAS2	18	13
Channel 3: Gts(135°)	card DAS3	80	40
Channel 4: G	card 0576B	10	10
Channel 5: Gd	card 0576A	10	10
Channel 6: Gt(45°)	card 0567C	10	10

Green's notes inform us that code 14 was in effect up to and including 29 August 1978, after which code 15 was used; code 16 was used from 3 May to 4 June 1979, after which code 15 returned. His notes do not, however, say what changes in the configuration were signified by codes 15 and 16. But McGregor was still using in September 1979 the calibration factors associated with code number 14. Hence we may infer that the correct amplifier card for the 90°-tilt pyranometer was the card DAS1, with a calibration factor of 13 LU/mV, throughout the period that we are considering.

Let us suppose that the card for the 90°-tilt pyranometer (13 LU/mV) was replaced with another card with a different calibration factor. This change would have to reduce the irradiance data by a factor between 2.5 and 3.0 (as was seen above from Figures 8.1 and 8.3 to 8.5). Hence the new card would need to have a calibration between 13/3.0 = 4.33 and 13/2.5 = 5.2. An inspection of the signal conditioning cards left by McGregor reveals only one card with a calibration factor in this range: 0576C (later assigned the reference number 13), which has a calibration factor of 5.0 LU/mV. If this card had been inserted in the 90° channel, then the logged irradiance readings would have been reduced by a factor of 2.6.

There are three weaknesses in this hypothesis: (i) It implies that Green did not make a note of this significant change in the logger configuration, even though he regularly noted less significant things, such as changes in the electrical battery. This weakness could be removed by assuming that another member of the SEU changed the card, without telling Green. Microdata loggers were being used in two other projects in the SEU at this time, and another user of Microdata equipment might have swapped cards. (ii) According to Svendsen (1978, Rept 430), when Green took over the running of the apparatus on 18 August 1978, card 0576C was already in use with a stationary pyranometer measuring diffuse irradiance. Its calibration was then 10 LU/mV. At some later time, someone soldered an extra resistor onto the amplifier circuit, thereby changing the card's calibration to 5 LU/mV. The hypothesis that we are considering therefore requires one of two further complications: first, we could assume that card 13 was modified between 18 August and 20 August and moved into the channel for the 90°-tilt pyranometer, while at the same time another card (with a calibration of 10 LU/mV) replaced card 13 in the diffuse-irradiance channel; or, secondly, we could assume that there was another card (besides no. 13) calibrated at 5 LU/mV, which was inserted into the 90°-tilt channel, but which card has subsequently either been lost or modified. (iii) There would have been no logical reason for wanting to reduce the irradiance readings during the period from late August to early December, whilst not doing so for the further period until April. The only legitimate reason for wanting to reduce the logged readings is to prevent the readings going outside the range that can be handled by the logger circuitry. If Svendsen or Green had expected this to occur on the 90°-tilt channel during the Autumn and early Winter, then the same expectation would have been applicable to the late Winter and early Spring, as the solar incidence angles on the pyranometers would have been in the same range. Moreover, a simple calculation would have shown that there was no risk of the irradiance readings going out of range. An upper bound for the irradiance in a vertical plane at Cardiff's latitude can be taken as 1000 Wm^{-2} (see Figure 8.1); this would generate an EMF of 12.8x1.0 = 12.8 mV in the pyranometer CM5 77-3538; this in turn would generate 13*12.8 = 166 logger units via the amplifier card DAS1. Such a logged reading is well within the limit of 200 LU below which Svendsen's cards were linear.

In conclusion, the 'swapped card' hypothesis is technically capable of explaining the reduction in the 90°-tilt irradiances, but circumstantial evidence makes it unlikely. The hypothesis would become much more probable, however, if there was evidence that some person other than Green had changed the card.

As a postscript, it might be relevant to mention something noted by McGregor at SEU meetings on 31 January and 14 March 1980. He said that he had carried out a calibration test on the card that was then being used with one of the 45°-tilt pyranometers. (He neglects to say whether it is the scanning or the stationary pyranometer.) The test yielded a calibration of 10 LU/mV, instead of the value of 13 LU/mV that Svendsen had specified. McGregor concluded that somebody else had changed the card and asked whoever did it to notify him so that he could use

the appropriate calibration factor in processing the data. Nobody admitted changing it. At the meeting on 14 March, McGregor said that he had found the original card that should have been in this channel. Again, he asked for someone to admit changing it.

The scanning 45°-tilt pyranometer had a signal conditioning card calibrated at 13 LU/mV, whereas the stationary instrument at the same tilt had a card calibrated at 10 LU/mV. From these two facts, it should follow that McGregor is referring to the scanning pyranometer. Nevertheless, the three scanning-pyranometer channels had Svendsen's custom-made cards (labelled DAS1 to DAS3), and these were quite different in appearance from the other signal conditioning cards, which were all Microdata cards. Therefore, McGregor would have been able to tell if the card had been changed simply by looking at it - without carrying out a calibration test.

McGregor hoped to be able to locate the date at which the card was changed in the 45°-tilt pyranometer channel, by inspecting the irradiance profiles. In March 1980, he reported that he had found no evidence of a change before January 1979. We do not know whether he ever found the date of the change. I have examined the data from both of the 45°-tilt pyranometers, and have found no evidence of a change in calibration, up to August 1979.

What happened to the 45°-tilt channel remains uncertain, but McGregor's statements do indicate that an unknown person was tampering with the logging system, and this makes it more plausible that the fault in the 90°-tilt channel was due to the card being swapped.

8.3 Empirical rehabilitation factors

In order to be able to use the data that have been affected by the spurious change in the apparatus, a 'rehabilitation factor' must be applied. (The term "rehabilitation factor" is used, rather than "correction factor", as the cause of the fault has not been ascertained and the rehabilitation factor might not fully correct the data.)

We can find a rehabilitation factor by comparing irradiance readings recorded on two similar days before and after the change. If $G_t(1)$ and $G_t(2)$ are two readings from the 90°-tilt pyranometer, taken under identical sky conditions, with $G_t(1)$ taken before the reduction and $G_t(2)$ taken after, then the rehabilitation factor will be computed as

$$X = G_t(1) / G_t(2)$$

It is impracticable, however, to use days having identical sky conditions, as they are too rare. A practicable method has been adopted as follows: (a) ensure that the solar position is approximately the same; (b) select days which have very similar angular distributions of radiance, which really implies either an overcast sky or a cloudless one; (c) compute the rehabilitation factor as

$$X = (G_t(1) / G(1)) / (G_t(2) / G(2))$$

in which G(1) and G(2) are the global irradiances in the

horizontal plane on the two occasions. The reason for normalising G_t to G is to eliminate the effect of changes in the turbidity. (As Steven has shown, the variations in turbidity that are normally enountered in this country reduce the value of G without altering the angular distribution of radiance in a cloudless sky.)

Only overcast skies or cloudless (or nearly cloudless) skies will be sufficiently well defined to allow similar conditions to be found on two different days. Of these, the overcast conditions will be the more stable; whereas on a nearly cloudless day the solar beam irradiance will fluctuate with passing dust clouds and the diffuse field is likely to be perturbed by reflections by the few clouds present. Therefore, let us first consider two overcast periods. Figures 8.7 and 8.8 show irradiance profiles for two overcast afternoons: 27 March (before the change) and 9 April 1979 (after the change). Figures 8.9 and 8.10 are polar diagrams showing the angular distributions of irradiation on the two days. It is immediately seen from the two polar diagrams (and their accompanying tables) that the irradiances recorded from the 90°-tilt pyranometer are very low on 9 April, whereas the irradiances recorded in the other channels are much the same on the two days. Under the columns labelled "45°, %" and "135°, %" are listed the ratios of $G_{ts}(45^o)$ and $G_{ts}(135^o)$ to the horizontal irradiance G. It is seen that they are almost the same on the two days, which implies that the distribution of irradiations in these directions is the same on both days. Therefore we might also expect that the ratio of $G_{ts}(90^o)$ to G would be the same on the two days. Assuming that to be so, we get eight rehabilitation factors for the eight azimuthal directions:

Azimuth	Rehabilitation
189°:	4.3
144°:	4.0
99°:	3.9
54°:	3.9
9°:	3.9
324°:	3.9
279°:	3.9
234°:	3.9
Mean:	3.96

This gives a factor of approximately 4.0, which is much higher than the range of 2.5 to 3.0 that was determined earlier for the possible values of the rehabilitation factor. It seems likely therefore that the diffuse radiance had a different anisotropy on the two days, and that the rehabilitation factor has been affected accordingly.

If we consider two cloudless periods instead of two overcast periods, then differences in the diffuse anisotropy should be much smaller. Unfortunately, truly cloudless skies are quite rare in Britain (especially in coastal areas in Spring) and therefore, instead of cloudless days, we must select two days on which the cloud cover is small and clouds do not obscure the Sun.

After the appearance of the fault, the morning of 15 April 1979 is the earliest occasion on which the sky is fairly cloudless (Figure 8.11). The nearest period before the appearance of the fault when there is a comparably clear sky is the morning of 4 April (Figure 8.12). It is evident, however, that there is some cloud in the sky on that morning. Further back in time, we find that 23 March has much clearer skies (Figure 8.13). This date, however, is about three weeks away from 15 April, and the Sun's elevation will be somewhat different during the corresponding hours of the day. Thus it is not obvious which of the two days, 23 March and 4 April, should be used.

Figures 8.14 to 8.16 show the distributions of irradiation between the times 9:00 and 11:00 (solar time) on the mornings of the three days (23 March, 4 April, and 15 April). It can be seen that the distribution of irradiations in the planes tilted at 45° and 135° are quite similar on all three days. We may therefore proceed on the working assumption that the true values of the 90°-tilt irradiations are also very similar. For each of the two days 23 March and 4 April, we get one value of the rehabilitation factor, X, for each of the eight azimuthal directions.

Rehabilitation factor, X
(in relation to 15 April 1979)

Azimuth	23 March	4 April 1979
189°:	2.87	2.53
144°:	2.98	2.50
99°:	2.93	2.57
54°:	2.30	3.20
9°:	2.67	4.17
324°:	2.33	3.67
279°:	2.60	4.20
234°:	2.33	3.22
Mean:	2.63	3.26

It is seen that using 4 April gives a greater scatter of rehabilitation factors (X ranges from 2.50 to 4.17) than using 23 March does (X ranges from 2.30 to 2.98). Moreover, many of the rehabilitation factors on 4 April are too high to be credible. It would seem therefore that the anisotropy of the diffuse radiation on 4 April is perturbing the rehabilitation factors. That would be consistent with the fact that there is some cloud near the Sun at times between 9:00 and 11:00 on that day. In view of these two defects, we shall not use the rehabilitation factors from 4 April.

The rehabilitation factors for 23 March are fairly self-consistent, although the factors for azimuths near the Sun (from 105° to 195°) are rather high. Plausible explanations for this can be conceived, as the irradiances in the 45°-tilt plane in the near-solar azimuths are higher on 23 March than on 15 April. It is inappropriate, however, to spend further time at present on investigating the rehabilitation factor. The proposed

factor of 2.63 seems satisfactory, although further investigation might be carried out when time is available.

A further depression of the irradiance values occurred during the period 25 June 1979 to 20 July 1979, having a different magnitude from the reduction by 2.63 that has been discussed above. The numerical value of this second reduction can be determined by the method that has already been employe. Two fairly cloudless days were selected, 19 June before the start of the new reduction, and 3 and 4 July afterwards (Figures 8.17 to 8.19 respectively). Figures 8.20 to 8.22 are polar diagrams showing the distributions of irradiation for the three days, for the period 9:00 to 15:00 in each day. There was a problem with the 45°-tilt channel during this period, and the readings from this channel may not be altogether reliable. For the present purpose, however, we are interested only in the 90°-tilt channel, which seems to be producing valid data.

Proceeding as before, we define the rehabilitation factors as one of the following two.

$$X = (G_t(19\text{-Jun})/G(19\text{-Jun})) / (G_t(3\text{-Jul})/G(3\text{-Jul}))$$

$$X = (G_t(19\text{-Jun})/G(19\text{-Jun})) / (G_t(4\text{-Jul})/G(4\text{-Jul}))$$

These are evaluated with the irradiances recorded from the 90°-tilt pyranometer for each of the eight azimuthal positions. The resulting values for each day are tabulated below.

Rehabilitation factor, X
(in relation to 19 June 1979)

Azimuth	3 July	4 July 1979
189°:	1.47	1.38
144°:	1.54	1.74
99°:	1.56	1.75
54°:	1.60	1.45
9°:	1.63	1.44
324°:	1.30	1.30
279°:	1.29	1.05
234°:	1.36	1.10
Mean:	1.47	1.40 Overall mean: 1.44

The scatter of these X factors is rather greater on 3 July than on 4 July (the ranges of X being 0.34 and 0.70 respectively), but the mean of X is very close. We may therefore take the overall mean, which is 1.44. One likely source of the scatter in the X values are the numerical errors in the small irradiances in the North-facing surfaces. For example, between the azimuths 54° and 324°, the irradiance in the 90°-tilt surface on 19 June does not exceed 10 Wm^{-2}. If we exclude the mean irradiances that do not exceed 10 Wm^{-2}, then we get mean values for X of 1.44 and 1.41.

If the earlier depression of the irradiance data (by the factor of 2.63) was due to the insertion of an incorrect

amplifier card, then it would be reasonable to suppose that this second depression is due to the insertion of another card. In order for that to be so, the replacement card would need to have a calibration of 13/1.44 = 9.0 LU/mV, but there is no evidence of such a calibration being used by Svendsen or others in the SEU. The closest known calibration of any of the cards is 10 LU/mV, which would imply a rehabilitation factor of X = 1.30 rather than 1.44. Nevertheless, we cannot rule out the possibility of such a card having existed and subsequently been changed to a different calibration.

8.4 Checks on the rehabilitation factors

Having arrived at rehabilitation factors by the method described above, we must check the validity of the factor by comparing several irradiance distributions for similar days before and after the appearance of the fault.

(a) Cloudless sky

Figure 8.23 shows the irradiance profile for 15 April; and Figure 8.24 shows the irradiation distribution for 9:00 to 11:00 on the same day, when the sky was almost cloudless. (This is the period used in setting up the rehabilitation factor.) In this and all subsequent plots, the rehabilitation factors have been applied to the 90°-tilt data. It should be compared with the equivalent graphs for 23 March and 4 April, Figures 8.14 and 8.15 respectively. Clearly, the shapes of the curves are very similar, indicating that the rehabilitation has been correctly implemented.

(b) Overcast sky

Figure 8.25 shows the irradiance profile, and Figure 8.26 the angular distribution of irradiation, for the rehabilitated data from 9 April, which was overcast all day. This polar diagram should be compared with the equivalent graph for 27 March, Figure 8.9. Clearly, there is an appreciable difference between the two graphs. In one respect, this is not surprising, since implausible rehabilitation factors were derived - and rejected - earlier from this pair of days. The question arises of whether the obvious differences between the curves are due to real differences in the radiance distributions, or due to some deficiency in the rehabilitation of the data.

Irradiation distributions were plotted for a number of other overcast days dated before the appearance of the fault, and one such is shown in Figure 8.27, for 15 March 1979. All these distributions showed that the 90°-tilt irradiances were equal to about 40% of the horizontal global irradiance, G. Yet, the plot for 9 April (Figure 8.26) shows that the 90°-tilt irradiances were only about 27% of G. To investigate this further, irradiance distributions were plotted for a number of other overcast days after the appearance of the fault, and one such is shown in Figure 8.28, for 19 April. This is almost identical to Figure 8.27 (for 15 March). We are led to the conclusion that 9 April was an unusual day, in which the irradiances in the vertical planes were much lower than usual.

It is noted that measurements are missing for two hours from 12:00 to 14:00, out of a day length of twelve hours. This gap will bias the data toward the non-noon hours of the day, and could conceivably have caused the vertical-surface irradiances to be lower than usual. Such an explanation, though, seems unlikely because of the shortness of the gap.

Digression:-

> Under an isotropic sky, a vertical surface would get precisely 50% of the irradiance of a horizontal surface, provided that the albedo of the neighbouring ground can be neglected. In the Svendsen data, it has been found that the 90°-tilt irradiance under an overcast is usually only 40% of the horizontal irradiance - which implies that the overcast sky is darker at the horizon than at the zenith. This is consistent with conclusions drawn from a literature survey (Lloyd 1985, Rept 1204). An unusual condition seems to have arisen on 9 April, however. On that day, the 90°-tilt irradiance is merely 27% of the horizontal irradiance. Hence the relative brightness at the horizon was much less on this day than usual. From the absolute values of the irradiances, it can be seen that the sky was much darker on 9 April than on the other overcast days (15 March, 27 March, and 19 April). We are thus led to the tentative conclusion that the anisotropy of the diffuse radiance becomes more accentuated as the overcast becomes thicker.

The excellent agreement of the irradiance distributions for 15 March and 19 April (Figures 8.27 and 8.28 respectively) is a good corroboration of the rehabilitation factor, and permits confidence in the rehabilitation.

(c) Partly-clouded skies

Comparing irradiance distributions for partly clouded skies, before and after the appearance of the fault, is a less stringent test than comparing distributions for cloudless or overcast periods. This is because it is practically impossible to ensure that the weather conditions are the same when the sky is partly clouded. Nevertheless, these skies are important because they are more relevant to average conditions.

Sky conditions were very similar on 8 April and 28 April, with considerable broken cloud (Figures 8.1 and 8.29 respectively). The distributions of irradiation for the two days are given in Figures 8.30 and 8.31. The shapes of the curves are quite similar, which indicates the validity of the rehabilitation factor. The absolute irradiances on 28 April are lower than those on 8 April, but the distribution of irradiances is almost the same on the two days.

8.5 Occasional relapses

There are several occasions in October 1978 when the depression of the 90°-tilt irradiance values ceases for short periods. These have been listed above in sub-section 8.1. Most of the

periods start either at dawn or shortly after midday, and apparently stop either at dusk or at dawn the next morning. These events are consistent with somebody experimenting by changing the amplifier card for either whole days or half-days. (The only exception is the first such period, which lasted for only about 12 minutes on the morning of 6 October.) On four of the days, these changes occur within the daylight hours. These are especially informative, and the irradiance profiles for these days are shown in Figures 8.32 to 8.35, and in greater detail in Figures 8.36 to 8.39. (Note that the graphs are in solar time, while the periods listed above are in GMT; also note that the graph time is expressed in hours and decimal fractions of hours, so 12.50 is 12:30.)

For a few minutes around the moments when the fault ceases, the irradiance values deviate from what one would expect. Because we cannot tell precisely when the fault stops, we cannot know whether these deviations are immediately before the cessation, or immediately after; or whether they represent some intermediate state. The deviations can be seen more easily in the large-scale plots of the un-rehabilitated irradiance profiles shown in Figures 8.36 to 8.39.

If the hypothesis were true that the irradiance depression was caused by an incorrect amplifier card being inserted, then it is difficult to explain any such deviations. When the (hypothetical) card was inserted, the logged readings should have immediately ceased to be depressed, without any gradient or fluctuation. Instead, we see a gradual change of the reduction factor from 2.63 to 1.0 over a period of about 3 minutes, apparently followed by erratic fluctuations. Two explanations are conceivable: (a) if the card was not changed very quickly, there would have been a gap in the recording while the card was out of its slot, and Green or McGregor might have doctored the data to insert interpolated values; (b) the card that was inserted (DAS1) might have had some defect that prevented it from operating normally for a few minutes after insertion (a 'warm-up' period). Since these transitional data are very few in number, they have simply been deleted. On 6 October, the period during which the depression ceased was very short, and it has been deleted entirely.

A small error will be introduced immediately after the end of the data depression, by the numerical process of removing the time response of the pyranometer. This is because the correction of a given datum involves referring to the immediately previous irradiance value. Since the latter will have been recorded as a depressed value, the numerical procedure will over-estimate the gradient of the pyranometer's response curve, and hence over-estimate the irradiance. In order to prevent this error, the rehabilitation factor of 2.63 should strictly be applied to the data before the time reponse is filtered out. This might be carried out in the future, but time is not available for it it now. The error introduced in this way is small and cannot account for the deviations that have been observed. Moreover, these deviations are also present in the RRR files, which have not been corrected for the pyranometric time-response.

Further investigation of the deviations will be made at a later date, and it might be possible to establish a proper correction. Of course, the transitional data are not important in themselves, as they are so few in number. But they might yield some useful insight into the nature of the fault that afflicted the 90°-tilt data. For this reason, they deserve more attention than can be afforded at present.

Because the fault ceases at points within each of the three days 7, 8, and 27 October, we are able to compare irradiance profiles for very similar weather conditions before and after the cessation of the depression. Figures 8.40 to 8.42 show the irradiance profiles for these three days after rehabilitation (and after the transitional points have been deleted). It is seen that the rehabilitated and un-rehabilitated sections are consistent. This further corroborates the numerical value of the rehabilitation factor of 2.63.

8.6 Conclusion

It is concluded that the irradiance readings from the 90°-tilt scanning pyranometer should be multiplied by a rehabilitation factor of X = 2.63 for the periods listed below under (a), (b), and (c):

(a) From 21 August 1978 to 3 December 1978
(excluding the following short periods:
10:30 GMT to 10:42 GMT on 6 October
12:44 GMT to sunset on 7 October
14:39 GMT to sunset on 8 October
sunrise to sunset on 11 to 13 October inclusive
12:20 GMT to sunset on 27 October).

(b) From 14 January 1979 to 30 January 1979.

(c) From 9 April 1979 to 22 June 1979.

and multiplied by a factor of X = 1.44 for the following period:

(c) From 25 June 1979 to 20 July 1979.

Irradiance measurements from this channel should be deleted for the periods listed below.

10:29 - 10:46 GMT on 6 October 1978
12:42 - 12:47 GMT on 7 October 1978
14:39 - 14:47 GMT on 8 October 1978
12:19 - 12:23 GMT on 27 October 1978

It is noted that the first rehabilitation factor (2.63) is consistent with the hypothesis that the channel's amplifier card (calibrated at 13 LU/mV) was swapped with another card (calibrated at 5 LU/mV), which would imply a reduction by a factor of 2.6. The discrepancy of 0.03 might be due either to inaccuracy in the calibration of the amplifier card (which has a numerical resolution of only 1 in 256) or to inaccuracy in the empirical determination of the rehabilitation factor X. The second rehabilitation factor is consistent with the hypothesis

that the channel's amplifier card was later replaced with one calibrated at 9 LU/mV.

A photograph of the Microdata interface unit has been found, which was taken by Green at some time between August 1978 and July 1979. It shows that an amplifier card manufactured by Microdata was applied to the 90°-tilt channel at that time, and not the card DAS1. This proves that the amplifier card was changed at some time. Of course, the photograph gives no indication of the calibration of the replacement card.

No rehabilitation factor has been determined for the period after 20 July 1979. This does not affect the present work. No measurements were logged from the scanning pyranometers during the period 21 July to 14 August, and the data after that date are not yet available for processing on the GEC 4090 computer.

9. Loss of data in May 1979

Throughout the period from the start of July 1978 to the end of April 1979, there are occasional gaps in the monitoring, which arise from various causes, but they never last more than a few days. In May 1979, however, there was a massive loss of data which makes this month almost useless for studying the anisotropic diffuse field. This was probably one of the reasons why McGregor apparently made a special effort to take measurements from the scanning pyranometers during May 1981, in order to ensure that the data-set covered a complete year of twelve months. The data that were logged during May 1981 are extant on the archive tapes, but only in the form of raw data held in QQQ files. In view of the further work that would be needed to pre-process these May data, they have been omitted from the present study.

The first two days of May 1979 exhibit no problems. On 3 May, all channels disappear in the morning (see Figures 9.1 to 9.4). About three hours later, 16 of the 27 channels reappear, while the remaining 11 channels give no signal. This pattern persists throughout May and up to midday on 12 June 1979: the sixteen live channels have a higher than normal incidence of small gaps, while the eleven other channels remain dead. The dead channels are distributed as follows:

Azimuth	90°-tilt	45°-tilt	135°-tilt
189°	dead	dead	dead
144°	live	dead	live
99°	live	dead	live
54°	live	dead	live
9°	live	dead	live
324°	live	dead	live
279°	live	dead	live
234°	live	dead	live

Stationary pyranometers:-
measuring horizontal global: live
measuring 45°-tilt global: dead
measuring horizontal diffuse: live

During most of May, therefore, it is difficult to determine reliably the angular distribution of irradiation. And it seems impossible to determine reliably the beam irradiance, because that would need all four of the South-facing channels of the vertical pyranometer. Without reliable values of the beam irradiance, it is impossible to assess the validity of any models of the diffuse field, so no graphs have been prepared for errors in the isotropic predictions in May (beyond 3 May).

10. Notice of other data faults

In this section, notice is given of faults in the Svendsen data that have not been corrected owing to a lack of time. It is believed that these will have a very small effect on results obtained from the data-set.

10.1 Loss of data due to fractured leads

As has been mentioned above, several channels went dead from 3 May to 12 June 1979, including all of the channels from the 45°-tilt pyranometer. After 12 June 1979, these channels became live again, but several channels developed an intermittent loss of signal. At first, the fault was restricted to the scanning 45°-tilt scanning pyranometer in azimuthal positions 39° and 84°. First instances of the fault occurred on 13 June (Figure 10.1), and the fault seems to have reached a saturation level by 16 June (Figure 10.2) - on which date it spread to another 45°-tilt channel, in azimuthal position 354°. As the fault mimics the obscuration of the Sun, it is difficult to see its effect on a partially clouded day; but the cloudless day of 19 June shows up the fault especially clearly (Figure 10.3). Also on 19 June, the fault spread to the 129°-azimuth channel. At times, the fault seems to have faded away, such as on 27 June (Figure 10.4), before gradually returning again as it had done by 12 July (Figure 10.5).

On 3 July, the fault spread to the 135°-tilt pyranometer (Figure 10.6), although it never reached a very high frequency in that pyranometer's channels. On 9 July, it finally spread to the 90°-tilt pyranometer (Figure 10.7), and by 15 July reached its highest frequency (Figure 10.8) in that pyranometer's channels.

McGregor traced the cause of the fault to fractures in the 'flying leads' that connect the moving pyranometer to the logger. This might also have been the cause of the extensive loss of data in May of the same year.

This fault persisted in the channels of all of the three scanning pyranometers until the afternoon of 22 July, when all channels went dead. On this date, McGregor replaced the leads connecting to the pyranometers, but incorrectly wired them up, so that three days' data were lost from all channels (Lloyd 1985, Rept 1204). The stationary pyranometer channels restarted on 26 July, but logging from the scanning pyranometers did not resume until 15 August. It is not known why logging of the scanning pyranometers did not resume on 26 July. McGregor's notes contain the simple entry "Moving pyranometers connected" for 15 August, but no reason has been found for why they were disconnected

during the 23 days from 23 July to 14 August. When they were restarted, the scanning pyranometers no longer exhibited the signal loss that has been described above.

The signal loss is manifested in the raw data by the occurrence of zero data values instead of the actual readings. In the processed (PPP) data, however, there are also spikes in the irradiance profile caused by the algorithm that filters out the time response of the pyranometer. The reason for the spikes is as follows. When the signal from a pyranometer returns to its proper value after being zero, the filtering algorithm 'assumes' that there is a very steep gradient in the irradiance, and that the true irradiance is therefore appreciably greater than the actual reading. Consequently it over-predicts the irradiance. This can be seen particularly clearly in Figure 10.3, with the cloudless day of 19 June. Whenever the signal from the 84°-azimuth channel becomes zero, there is a corresponding spike in the next channel, which is for the 39° channel - except when the latter channel's signal has also become zero.

At present, the program TIMEPLOT treats these spurious data as valid and plots them along with the valid data points. The program POLEPLOT, on the other hand, excludes irradiance values that are zero, although it does not vet out the other spurious values (i.e. the spikes). Consequently, this fault will affect to some extent the polar graphs of the irradiation distribution and of the isotropic prediction error. This can be seen in Figures 10.9 and 10.10, for 19 June, which is a cloudless day. The irradiation value for the 39°-azimuth pyranometer is slightly too high, while that for the 84°-azimuth pyranometer is slightly too low (so the isotropic prediction errors are respectively too high and too low). Nevertheless, the magnitude of this error is plainly quite small.

The correction of this fault is fairly straightforward in principle and will be carried out when time is available. It will be more accurately carried out on the RRR data, rather than the PPP data, in order to eliminate spurious spikes through the algorithm for filtering out the pyranometric time response.

10.2 Obscuration of the Sun by buildings

Whenever solar irradiance is monitored in a city, there is a risk of the Sun being obscured by tall buildings. In order to assess the risk of this, Svendsen surveyed the horizon as seen from the roof of the Mining Building, using a theodolite (Svendsen 1977, Rept 357). His diagram of the nearby tall buildings is reproduced as Figure 10.11. It is seen that there are two main constructions that will block the sunlight at certain times of the year, both early in the morning: one is the lift shaft, situated on roof of the Mining Building at an azimuth of about 50°; the other is the Julian Hodge building (not labelled in Svendsen's diagram) at an azimuth of about 105°. There are also several buildings that will briefly obscure the Sun when it is very low, if it is not obscured already by the natural skyline.

Obviously, the sunlight can be blocked by these buildings only if the Sun is not already obscured by cloud. As cloudless skies at dawn are fairly uncommon in Cardiff's climate, the frequency of solar obscuration by the buildings will be low.

Figures 10.12 to 10.14 show irradiance profiles from the 90°- and 45°-tilt pyranometers for 26, 27, and 29 June 1979. These dates occur shortly after the Summer solstice. Early in the morning (between 4:00 and 5:00), it possible to discern what is obviously the obscuration of the Sun by a pair of buildings. This is inferred from two facts: first, the disappearance of the beam irradiance happens at the same time on each of the days, which is too great a concidence for clouds; second, there is no spike in the irradiance at the beginning or end of the obscuration, which almost always appears when the Sun is obscured by a cloud.

It is not clear which two buildings have produced these shadows. Judging from the curve for the Summer solstice in Svendsen's diagram, there should be only a single obscuration of the Sun, which should start at dawn. But if we introduce the known 6° error in Svendsen's determination of due South, then it becomes possible that the two shadows are those of the lift shift and of the broader maintenance room at the base of the lift shaft.

Only the pyranometers tilted at 45° and 90° and facing roughly eastward will register the building shadows appreciably. The building shadows are only barely visible in the profile of the global irradiance in the horizontal plane.

The correction of the irradiance for obscuration by buildings is straightforward in principle, and will be carried out when time is available.

11. Monthly graphs

In many applications of information about solar irradiation, monthly totals are an adequate form in which to use the data. Often, this is because the data are to be fed into correlation models which have been set up to run on monthly data only. It is therefore desirable to have a set of reference curves of monthly irradiations in arbitrary planes.

Ideally, the reference curves should be based on many years' data. In the present author's experience, ten years's data are generally necessary to ensure that average monthly irradiations are within about 5% of the long-term averages. Meteorological textbooks, however, normally recommend thirty-year averages as the basis for determining the climate of a place. Contrasted with this ideal requirement of at least ten years, the Svendsen set of thirteen months' data, with gaps, might seem useless. Nevertheless, reference curves derived from the Svendsen data are valuable in two respects: first, they indicate at least the rough value of the anisotropic diffuse irradiations in arbitrary planes - which is better than having no guide at all; second, they indicate the relative magnitudes of the anisotropic diffuse irradiations, which are probably less changeable than the

absolute magnitudes of these irradiations.

The greatest cause for concern about the monthly graphs is their sensitivity to the number of overcast days in a month. The angular distribution of radiance is completely different between overcast and sunny days. Small changes in the number of overcast days in a month may therefore be expected to have a large effect on the monthly angular distribution of irradiation. This suggests that a fruitful approach might be to generate two sets of reference curves: one for overcast days, the other for non-overcast days. A user of the results would then take an average of the two curves weighted by the proportion of overcast days in the month in question. In the present project, however, there is insufficient time to pursue this approach.

Figures 11.1 to 11.15 are polar diagrams showing the angular distribution of monthly irradiations, for the thirteen months from July 1978 to July 1979. As has been discussed above, the data from May 1979 have large gaps. An attempt has been made to salvage as much as possible from this month, and the results are shown in three graphs: Figure 11.11 is the graph for 1 to 2 May (with no missing data); Figure 11.12 is the graph for 4 to 31 May (only the 90^{o}-tilt and 135^{o}-tilt irradiances, and including many gaps); Figure 11.13 is the combined graph for 1 to 31 May. The use of Figure 11.11 is preferred, although none of these May graphs is altogether satisfactory. (A replacement based on data from May 1981 will be prepared in the future.)

Figures 11.16 to 11.28 are polar diagrams showing the angular distribution of the errors in irradiation estimates predicted by the isotropic model. They can also be interpreted as distributions of the anisotropic part of the diffuse irradiation. Again, these are for the thirteen months from July 1978 to July 1979. Here, there is only one May graph, which is for 1 to 2 May; this is because it has not been possible to separate the beam and diffuse components of the global irradiance on the days from 4 to 31 May, which would be a prerequisite for deriving any model predictions.

12. Representative daily graphs

It is of theoretical interest, and potentially of practical interest, to have reference curves for the angular distribution of irradiation on recognisable types of day.

12.1 Classification of days

Two obvious categories for breaking down the irradiance data are cloudless and overcast days. Cloudless days, however, are so rare that we must use instead days that are 'almost cloudless', a condition that corresponds to the standard meteorological class of 'zero okta'.

Partially clouded skies make up the third broad category, but this must be further sub-divided. The purpose of the sub-division is to arrive at sky conditions that are sufficiently well-defined to determine the leading aspects of the diffuse radiance field, as nearly uniquely as possible.

There are three dimensions in which these cloudy days can be classified: first, by some measure of the horizontal extent of the cloud field; second, by some measure of the vertical thickness of the clouds; third, by some measure of the extent to which the clouds are broken up within the horizontal plane. Since we are working with just irradiance measurements (without the aid of the crude and infrequent sky observations taken at Rhoose airport 18 km away), convenient numerical indices of these three dimensions are, respectively: the duration of solar obscuration (or its complement, the duration of bright sunshine); the height of the irradiance spike that immediately precedes or succeeds the obscuration of the Sun by a cloud; and the frequency with which the Sun becomes obscured.

The significance of the first index (duration of solar obscuration) is obvious, as the Sun is the dominant point-source of radiant energy in the sky. The significance of the other two indices lies in their relation to the amount of radiant energy reflected from the sides of clouds. Immediately before and immediately after the Sun is obscured by a cloud, the beam irradiance will be reflected by the cloud's side to an appreciable extent - this is the cause of the spikes in the irradiance profile (see e.g. Figure 12.1) on each side of the troughs. Obviously the spikes themselves do not carry much irradiation because of the shortness of their duration. But the clouds do reflect appreciable amounts of irradiation even when they are not near the Sun's disc, and the magnitude of the spikes is probably correlated with the magnitude of that irradiation. How much energy is reflected by the clouds' sides is influenced by how many cloud sides are visible from the ground. That depends partly on how broken the cloud field is in the horizontal plane, and that can be indexed by the frequency with which the Sun is obscured by clouds.

Time is not available to carry out a regression analysis of the irradiation distributions against these indices. Instead, we can only examine distributions that are typical of some sky conditions.

12.2 Cloudless skies

Figures 12.2 to 12.5 show the irradiance profiles for four 'almost cloudless' days: 30 January, 23 March, 14 July, and 20 September. The angular distributions of irradiation for these days are given in Figures 12.6 to 12.9, and the angular distributions of the isotropic prediction errors are given in Figures 12.10 to 12.13.

The error plots for cloudless skies mostly confirm what has already been established from the literature regarding the effect of the high levels of diffuse radiance around the horizon ('circumhorizontal') and around the Sun ('circumsolar'). Let us consider the effect on vertical planes first. Circumhorizontal radiance causes vertical surfaces to receive more energy than the isotropic model would predict; and the circumsolar radiance adds to this effect if the surface faces the equator. Consequently, for a vertical surface under a cloudless sky, the isotropic-prediction errors will almost always be negative, and

the magnitude of the error will be greatest toward the equator. (Poleward vertical planes, however, may receive an irradiance that is less than the isotropic prediction in an unusually clear sky, when the circumhorizontal radiance is weak.) Tabulated below are some illustrative results selected from Figures 12.10 to 12.13. For the 90°-tilt pyranometer, these results are fully consistent with the conclusions that were drawn from the literature survey.

Table 12.I

Daily error in isotropic prediction of irradiation, Wh/m^2

	90°-tilt pyranometer		45°-tilt pyranometer	
	South	North	South	North
30 Jan:	-186 (-4%)	-73 (-16%)	-257 (-6%)	+152 (+30%)
23 Mar:	-256 (-5%)	-188 (-25%)	-770 (-12%)	+36 (+4%)
14 Jul:	-224 (-5%)	-261 (-13%)	+524 (+7%)	+440 (+9%)
20 Sep:	-124 (-3%)	-27 (-3%)	-847 (-15%)	+21 (+1%)

We shall now consider 45°-tilt planes. The circumsolar irradiance, rather than the circumhorizontal, will dominate the anisotropic diffuse irradiance in these tilted planes, because the circumhorizontal energy will be received at a large incidence angle. From the literature, it was concluded that 45°-tilt surfaces facing the equator should receive more energy than that predicted by the isotropic model, and those facing the pole would receive less than the isotropic prediction. Therefore we should expect the isotropic prediction errors to be negative for the equator-facing surfaces and positive for the pole-facing surfaces. In Table 12.I, it is seen that this is true for most of the year, but not for the middle of Summer. On 14 July the isotropic model over-predicts the daily irradiation on all 45°-tilt surfaces, even those facing South - which is somewhat surprising. A partial explanation can be found in the long duration of the Summer day: for a large fraction of the day, the Sun will be either at a low elevation angle, or behind, any given 45°-tilt surface. Consequently the surface will not receive much circumsolar irradiance, but will instead see the darker sky that lies toward the anti-solar point. This can be illustrated by breaking up 14 July by the time of day: Figures 12.14 and 12.15 show the isotropic error distributions for the two hours centred on 6:00 and 12:00 LAT. In the six o'clock distribution, it is clear that the South-facing 45°-tilt pyranometer sees irradiance from the darker part of the sky and the isotropic model over-predicts by 7%; in the midday distribution, the same pyranometer sees the circumsolar irradiance at nearly normal incidence, and the isotropic model over-predicts by only 2% - which is within the likely error of these results. It would seem reasonable to assume, therefore, that the diffuse irradiance in the 45° planes takes values greater than the isotropic prediction in sunward directions, but that the magnitude of this anisotropy is too small to be determined by the methods employed here.

12.3 Sparse cloud fields

Figures 12.16 and 12.17 show the irradiance profiles for two days on which there are many small clouds scudding rapidly past the Sun: 25 April and 15 August. The cloud field on these days may be described as a 'sparse field'. The angular distributions of irradiation for these days are given in Figures 12.18 and 12.19, and the angular distributions of the isotropic prediction errors are given in Figures 12.20 and 12.21.

It is desirable to isolate the contribution that the clouds make to the irradiation distribution. This can be done by comparing the polar graphs for the cloudy day with the corresponding polar graphs for a cloudless day.

Solar declination takes approximately the same value on 25 April and 14 July, so we may compare the two days' distributions of irradiation (Figures 12.18 and 12.22 respectively) and their distributions of the isotropic prediction error (Figures 12.20 and 12.23 respectively). Overall irradiation levels on 25 April are inevitably lower than those on 14 July because the Sun is obscured for a large part of the former day. The angular distributions of irradiation are also different: the sparse field of clouds apparently causes the relative daily irradiation (H_t/H) to be greater in the equator-facing planes and smaller in the pole-facing planes. For example, the 45^{o}-tilt South-facing irradiation, H_t, when expressed as a percentage of the horizontal irradiation H, is about 110% on 25 April, whereas it is only about 99% on 14 July.

Two possible explanations can be considered for the change in the irradiation distribution that is produced by the presence of the sparse cloud field, as follows.

(a) Geometric hypothesis

The Sun is more likely to be obscured by clouds when it is low in the sky, and the Sun will be lower in the sky the further it is from due South. Therefore, the presence of clouds will have least impact on the mean global irradiance in the direction of due South, and increasingly greater impact in directions away from due South.

The strength of this effect, and hence the practical validity of the hypothesis, depends on the vertical thickness of the clouds in relation to their horizontal width and on the width of the horizontal gap between neighbouring clouds. If the clouds are shallow and have large gaps between them (e.g. lenticular Cumuli), the attenuation of the mean irradiance in off-South directions will be only slightly greater than that of the mean irradiance in the southward direction. If, on the other hand, the clouds are tall (e.g. Cumuli congesti), the effect will be stronger.

This hypothesis can explain the observed increase of the daily southward irradiation relative to the daily irradiations in other directions. But it also implies two other predictions: first, that the angular distribution of irradiation at noon should be very similar between the cloudless day and the day with a sparse cloud field; second, that the irradiation distributions

should be most different between the two sorts of day in the periods just after sunrise and just before sunset.

These extra predictions can be tested by plotting the irradiation distributions for the two time periods 9:00 to 10:00 and 11:30 12:30, for each of the two days, 25 April (Figures 12.24 and 12.25) and 14 July (Figures 12.26 and 12.27). (We are obliged to use the hour 9:00 to 10:00 rather than an earlier one because the earlier data are missing from 25 April.) It is seen that the irradiation distributions are actually most similar at 9 o'clock and most different at noon. Thus the present hypothesis must be rejected. Presumably, the clouds on 25 April were shallow.

(b) Turbidity hypothesis

Another hypothesis is that the air was dustier and more humid on 14 July than on 25 April, so that the turbidity would be greater on the former date. The higher turbidity would act primarily to reduced the beam irradiance, although some of the irradiance lost from the beam would be spread around the sky as diffuse irradiance.

Since the beam irradiance forms a greater part of the global irradiance when the Sun is high, and therefore near South, the impact of the turbidity on global irradiance may be expected to be greatest in the direction of due South. Thus the hypothesis of higher turbidity on 14 July can explain the difference between the two days' angular irradiation distributions.

According to McGregor's tables (McGregor 1979, Rept 616), the meteorological station at Rhoose airport recorded a visibility ranging from 30 to 45 km during 25 April, but a much lower visibility of only about 7 km on 14 July. In addition, the temperature on 25 April never exceeded 12°C, but on 14 July it was up to 22°C; similarly, the relative humidity was around 45% on 25 April but around 70% on 14 July. Thus 14 July was humid, hazy day whereas 25 April was a cold, clear day. (The information available for the three days 23 to 25 April indicate that a cold air mass passed over Cardiff on the morning of 24 April, bringing both fair-weather Cumulus and clearer air.)

It therefore seems that a difference of turbidity between 25 April and 14 July is the most likely explanation for the difference between the two days' angular distributions of irradiation.

12.4 Dense cloud fields

Figures 12.28 to 12.30 show the irradiance profiles for three days on which the sky was mostly filled with large clouds: 21 July, 3 November, and 14 December 1978. The cloud field on these days may be described as a 'dense field'. The angular distributions of irradiation for these days are given in Figures 12.31 to 12.33, and the angular distributions of the isotropic prediction errors are given in Figures 12.34 to 12.36.

12.5 Overcasts

Figures 12.37 to 12.39 show the irradiance profiles of three overcast days: 1 July 1978, 2 July 1978, and p.m. on 27 March 1979. The angular distributions of irradiation for these days are given in Figures 12.40 to 12.42, and the angular distributions of the isotropic prediction errors are given in Figures 12.43 to 12.45.

Overcasts can be differentiated into three sorts: first, layer clouds that are of sufficient optical depth to make the Sun's disc invisible (e.g. Stratus); second, fields of broken cloud that are sufficiently extensive to block out the Sun and the blue sky for most of the time (e.g. Stratocumulus); third, layer cloudds that allow the Sun to be seen (e.g. Cirrostratus); . We shall not consider the third type (thin layer clouds) as it seems to be impossible to distinguish it from haze on the basis merely of the irradiance profile. It would probably be necessary to refer to the hourly observations recorded at Rhoose to identify the presence of cirriform clouds.

The first two sorts of overcast (the smooth and the non-smooth) are differentiated because we would expect to see two differences in their radiation regimes. (a) There will normally be no azimuthal variation in the diffuse field under a thick layer overcast, whereas there could be an azimuthal variation under an overcast that is due to broken cloud. This would be due to the reflection of sunlight by the sides of clouds. Results that have so far been derived from the scanning-pyranometer data, however, show no appreciable azimuthal variation in irradiance in overcasts - except at very low solar elevations, when the irradiance in the direction of the Sun is lower than the mean irradiance for each tilt. (b) Under the non-smooth overcast, there should be temporal changes in the irradiance level as the brighter parts of the cloud field pass overhead. In Figures 12.37 and 12.38, one can discern a fairly regular fluctuation in the horizontal irradiance, which is suggestive of sinusoidal waves. It might be fruitful to carry out a Fourier analysis of the irradiance in overcasts of this sort. Such an analysis might yield a useful index for identifying the type of overcast.

Let us briefly compare the irradiation distributions on the three overcast days. The absolute irradiance levels on 1 July are appreciably higher than on 2 July, yet the angular distributions of irradiation are very similar on the two days. This is consistent with Steven's (1977) finding that the overcast radiance distribution was uncorrelated with the absolute level of irradiance. On both of these days (and numerous other overcast days in the data set), the irradiances fall into the following rough pattern; the isotropic prediction is given on the right-hand side, for reference:

45° tilt - 80% of horizontal irradiance (isotropic 85.4%)
90° tilt - 40% of horizontal irradiance (isotropic 50.0%)
135° tilt - 15% of horizontal irradiance (isotropic 14.6%)

On 27 March, the irradiance in the 135°-tilt has shrunk to almost nothing (2% of G), while the irradiances in the 45°-tilt and 90°-tilt planes maintain their usual values of about 80% and 40% respectively of the horizontal irradiance G. In this

connection, we may also consider 9 April 1979, which has been discussed at length above (see Figures 8.25 and 8.26). On that overcast day, the 135°-tilt irradiance is again a mere 2% of G, but the 90°-tilt irradiance has also fallen - to aboyt 26% of G, whereas the irradiance in the 45°-tilt remain unaltered at 80% of G.

Taken together, these results suggest that the idea of a single radiance distribution for all overcast skies is unrealistic. Steven (1977) followed the tradition that had been laid down by Moon and Spencer (1942), of seeking a 'standard overcast sky' distribution. Nevertheless, Steven found that there was a high degree of scatter of the actual distributions about his fitted 'standard curve'. Perhaps it might be fruitful to try to correlate the overcast irradiances in the three tilts against some readily determinable parameter. Two plausible candidates that come to mind are (a) the atmospheric transmisison (= global irradiance, G, divided by extraterrestrial irradiance, G_o, in the horizontal plane); (b) some measure of the fluctuation in the irradiance level under overcasts of the non-layer sort, e.g. the mean amplitude of the peaks or the mean time between peaks.

Figures 12.46 to 12.52 show the angular distribution of irradiation at different times throughout the overcast day of 1 July 1978. The periods represented are symmetric about noon:

Fig. 12.46 - 4:00 to 5:00
Fig. 12.47 - 5:30 to 6:30
Fig. 12.48 - 8:30 to 9:30
Fig. 12.49 - 11:30 to 12:30 (noon hour)
Fig. 12.50 - 14:30 to 15:30
Fig. 12.51 - 17:30 to 18:30
Fig. 12.52 - 19:00 to 20:00

It seems to be widely assumed that the radiance distribution under an overcast is independent of the solar elevation. The results presented here cast some doubt on this. From about 9:00 to about 15:00 there is no appreciable change in the radiance distribution, but outside these times the radiance above the horizon drops markedly. This is manifested in the fall of the irradiances measured by the vertical pyranometer: throughout the central part of the day, the irradiance in any of the vertical planes is about 40% of the horizontal irradiance, G; but in the morning at six hours before noon, or in the evening at six hours after noon, the vertical irradiance is only about 30% of G; and in the first hour after dawn or the last hour before dusk, the vertical irradiance is just 20% of G. Irradiances in the 45°-tilt planes, on the other hand, remain fairly stable at about 77% of G. In the 135°-tilt planes, the irradiance is stable at 15% of G during the central part of the day, but rises slightly to 17% of G just after dawn and just before dusk. This implies that the relative radiance immediately above the horizon rises at dawn and dusk, whereas the relative radiance at a slightly higher angle falls at those times.

To place this discussion in perspective, however, it must be noted that the energy quantities incident in the 90°-tilt and

135°-tilt planes are inconsequential before about 9:00 and after 15:00. In fact, the irradiances in these planes between 4:00 and 5:00 are barely even measurable.

13. Toward a geometric model

13.1 Introduction

The problem of formulating a physical model of the radiation field in a partly clouded sky has been discussed at length in the Work Plan. It was noted that a fully detailed ray-tracing model would be impractical, because (a) it would require computer speeds much greater than those available to us, and (b) it would require information about the atmosphere that is not routinely available and is certainly not available in conjunction with the Svendsen data. A compromise approach was therefore outlined. This consisted in tracing rays through a field of 'grey boxes' that represented clouds. Rays were to be traced only between the clouds, not within the clouds, thereby greatly saving computer time, as well as obviating the need to specify the water density and the droplet-size probability distributions inside the clouds.

It has not been possible, however, to devise a scheme for programming the grey-box model in such a way that it can be run within practical periods of time.

A more simplified model has therefore been formulated. Even this simpler model is at the limit of what can be done on a practical time scale on the GEC 4090 computer. At its present stage of development, the model takes between 40 minutes and 5 hours to run on the 4090 (depending on the make-up of the cloud field). In order to investigate the effects of parameter variations, days of computer runs are needed. It is clear that if any headway is to be made with this model then it must be done on a more powerful computer.

The implementation of the model involves sequences of computations that are highly independent of one another. This suggests that a parallel processor would a suitable means of running the model quickly. It is therefore hoped to run the model on one of the Cray computers that are accessible to the academic community.

13.2 Outline of the model

The basic ideas underlying the model can be stated simply. First, at any given elevation angle, the fractional area of sky obscured by cloud is the complement of the relative duration of sunshine. Second, when the Sun is low in the sky, it is more likely to become obscured by cloud than when it is high in the sky. We may call this the 'elevation effect'.

For a given field of clouds, this elevation effect can be represented by a zenithal distribution of the sunshine duration - a curve showing the mean duration of bright sunshine versus the Sun's elevation angle. Examples of such curves can be seen in Figures 13.1 to 13.42 (in which the 'relative sunshine duration' is the ratio of actual sunshine duration to the maximum

possible).

A profile of this sort indirectly conveys information about the zenithal distribution of cloud coverage. Indeed, we can regard the curves as a crude means of characterising the cloud field. These profiles can therefore be used to derive information about the angular distribution of diffuse radiance, which is usually dominated by the angular distribution of clouds. So far, though, attention has been focussed on the first step, that of predicting the zenithal distribution of the cloud coverage.

The starting point of the model is a horizontal grid of identical clouds, each cloud being the same distance from each of its neighbours. We then assume that each cloud moves purely randomly within its local 'cell'. If (x,y) is the nominal position of a cloud, and w_g is the width of the gap between neighbouring clouds, then the position of each cloud is assumed to have a uniform probability distribution, with the X-coordinate varying from $x-w_g$ to $x+w_g$ and the Y-coordinate varying from $y-w_g$ to $y+w_g$.

It is next assumed that the radiance field under this cloud field is approximately the same as that under a similar field in which the clouds are arranged in circles concentric around the observer. In this 'concentralised' field, each concentric row is at a distance w_g from its two neighbouring rows, and each cloud is at a mean distance w_g from its two neighbours in the same row. Thus the mean distance between clouds is the same in both the rectangular grid and in the concentric grid. Again, each cloud is assumed to move purely randomly within its local cell. If a cloud's nominal position is (r,c), with r being the radial distance from the observer's overhead point and c being the arc distance from North along that cloud's row, then the actual position of the cloud has a uniform probability distribution, with the R-coordinate ranging from $r-w_g$ to $r+w_g$, and the C-coordinate ranging from $c-w_g$ to $c+w_g$ on average.

The final assumption is that each cloud is curved in line with the row in which it lies. Hence each point on the cloud side facing the observer is perpendicular to the radial line drawn to the cloud from the observer's overhead point.

The purpose of the two main assumptions - the concentralisation of both the grid on which the clouds are placed and of the shapes of the individual clouds - is to eliminate the azimuthal variation in the apparent cloud coverage. This is realistic, because one expects the average distribution of clouds in the sky to be azimuthally uniform, unless there is a strong meteorological influence such as a steep-sided mountain near the observer.

Eliminating the azimuthal dimension from the computations reduces the running time of the model very greatly. A precise figure cannot be given, but it seems likely that running the model with a concentric field instead of a rectangular field decreases the running time by at least two orders of magnitude.

Given the assumptions described above, one can determine the proportion of sky obscured by cloud at each elevation angle, using elementary trigonometry. The actual computation, though, is complicated in three ways: (a) the sides of distant clouds may be hidden by nearer clouds; (b) each cloud is moving randomly in both the radial and azimuthal directions, so the computer program must iterate through the possible positions of each cloud; (c) very distant clouds will occupy a negligible area of sky, so the program must exclude clouds that are beyond some limiting distance. Care must be taken in point (c) to deal correctly with clouds whose bases are beyond the horizontal range of the model but whose bodies stretch upwards and occupy non-negligible amounts of sky.

As was mentioned above, the fractional area of sky obscured by cloud is the complement of the relative duration of sunshine. Therefore the model described above predicts a relationship between the solar elevation and the duration of sunshine, for a given cloud amount. Thereby, the model can be partially validated: one can determine the empirical relationship between sunshine duration and solar elevation, and compare it with that predicted by the model. This is done in the next two sub-sections.

13.3 Empirical data

An extensive set of meteorological data were available from six stations in the UK, which were recorded by H.M. Meteorological Office. These data sets include hourly measurements of sunshine duration, S (measured in tenths of an hour, but usually expressed as a fraction of an hour), and visual observations of cloud amount, C (recorded in okta, but usually expressed as a fraction of the sky). The six stations are as follows.

Station	Latitude	Longitude	Years
Aberporth	52.133°N	4.567°W	1959-79
Aldergrove	54.650°N	6.217°W	1969-79
Cardington	52.100°N	0.417°W	1972-79
Eskdalemuir	55.317°N	3.200°W	1959-79
Kew	51.467°N	0.317°W	1959-79
Lerwick	50.133°N	1.183°W	1959-79

For each station, curves of the elevation effect were determined, which are shown in Figures 13.1 to 13.42, in the following sequence.

Figs. 13.1 - 13.6: C = 1 okta, for Aberporth to Lerwick
Figs. 13.7 - 13.12: C = 2 okta, for Aberporth to Lerwick
Figs. 13.13 - 13.18: C = 3 okta, for Aberporth to Lerwick
Figs. 13.19 - 13.24: C = 4 okta, for Aberporth to Lerwick
Figs. 13.25 - 13.30: C = 5 okta, for Aberporth to Lerwick
Figs. 13.31 - 13.36: C = 6 okta, for Aberporth to Lerwick
Figs. 13.37 - 13.42: C = 7 okta, for Aberporth to Lerwick

The curves do not change by much for 1 to 4 okta, beyond which they change fairly rapidly, with the sunsine duration falling appreciably.

These curves were constructed in the following manner. The range of possible solar elevation angles (in theory 0^{o} to 75^{o}, but in practice 0^{o} to about 60^{o}) was divided into 25 intervals, and within each interval the percentiles were determined of the values of S. This was done separately for each possible value of C other than 0 okta (cloudless sky) and 8 okta (overcast sky). Graphs were then plotted of the median (50% percentile), the quartiles (25% and 75% percentiles) and the 5% and 95% percentiles.

The values of S are found to have highly skewed distributions, which can be seen from compression of the upper percentiles relative to the lower percentiles. Therefore it would be inappropriate to fit a curve by the usual technique of minimising the sum of squared errors between the curve and the individual data. The best way to fit a curve to these data would be to fit it to the medians of the data. (Note that it would be quite wrong to fit it to the values of the mean in each interval.) At present, however, it is impracticable to fit the model by any method, because the model is too slow to run.

The empirical curves were produced with two programs: UNISCAN, a general program for scanning through the Meteorological Office data and compiling the basic statistics that are needed for a univariate regression; and UNIPLOT, which reads the statistics compiled by UNISCAN and plots the percentiles of the dependent variable versus the independent variable (with, optionally, the means and the linear and quadratic regression curves). These programs are general in the sense that they can operate on any channels in the meteorological data base. In the present case, the user must specify that the dependent variable is channel 4 (sunshine duration), the independent variable is channel 43 (computed solar elevation), and the data are to be partitioned according to the value of channel 5 (the total cloud amount).

13.4 Implementation of the model

As the model is still in a tentative stage of development, it would be inappropriate to present a full exposition of it here. The model has, however, been encoded into a FORTRAN program, called ZENDMODL. A listing of the program is given in the separately bound annexes (Annexe G.3), but it is also reproduced as Appendix B of the present volume for easier reference. The algorithm that effects the model can be followed in the program, which is internally documented.

The model has four physical parameters: the height of the cloud bases above ground level, h_c; the vertical depth of the cloud masses, d_c; the width of the cloud in the radial direction w_c (which is assumed equal to the mean lateral width of the cloud); and the relative areal cloud coverage, CA, in the horizontal plane. The first three parameters are specified by the user when the program ZENDMODL is run, and the cloud coverage is automatically varied by the program.

A brief study of sensitivity to parameter changes has been carried out. Some of the results are shown in Figures 13.43 to

13.50. Note that when the Sun is at the zenith (elevation angle 90°), the relative sunshine duration must be numerically equal to CA, and cannot vary with h_c, d_c, or w_c. For a given value of CA, therefore, the end-point of the curve at the right-hand margin of the graph must be the same in all the Figures. This has been verified by superimposing the graphs. (The superpositioning is not shown here.)

(a) Effect of cloud height.

Figures 13.43 to 13.45 show the graphs for the following cases:

Fig. 13.43:	h_c = 2 km	d_c = 4 km	w_c = 4 km
Fig. 13.44:	h_c = 6 km	d_c = 4 km	w_c = 4 km
Fig. 13.45:	h_c = 12 km	d_c = 4 km	w_c = 4 km

It is seen that varying the height of the clouds above ground level has a completely negligible effect on the profile of the elevation effect. The only visible difference between the three Figures is that, in the lower left-hand corner, the curves start at progressively higher elevation angles. This is because the program ignores clouds more than 120 km from the observer, irrespective of the height of the clouds.

(b) Effect of cloud width.

Figures 13.46 to 13.50 show graphs for the following cases:

Fig. 13.46:	h_c = 6 km	d_c = 4 km	w_c = 16 km
Fig. 13.47:	h_c = 6 km	d_c = 4 km	w_c = 12 km
Fig. 13.48:	h_c = 6 km	d_c = 4 km	w_c = 8 km
Fig. 13.49:	h_c = 6 km	d_c = 4 km	w_c = 2 km
Fig. 13.50:	h_c = 6 km	d_c = 4 km	w_c = 1 km

It is seen that varying the mean width of the individual clouds has a large effect on the curves. (Note that Figure 13.47 is incomplete, owing to the computer job exceeding the limit of 300 minutes running time.) The effect can be summarised by saying that, the smaller the 'footprint' of the cloud, the less will be the sunshine duration at elevations away from the zenith.

(c) Effect of cloud depth.

Variation of the cloud depth, d_c, was found to give the same results as the inverse variation of w_c. That is to say, if the ratio w_c/d_c is held constant, then the individual parameters d_c and w_c can be varied arbitrarily without affecting the sunshine profiles. The graphs that resulted from varying d_c are not reproduced here as they are indistinguishable from Figures 13.46 to 13.50.

13.5 Discussion

It is too early to draw definite conclusions about the geometric model that has been described above. But it seems that the model can generate functions of sunshine duration versus solar elevation that resemble the empirical relationships derived from Meteorological Office records.

The most important cause of systematic discrepancy between the empirical and theoretical curves is expected to be the non-zero threshold of the Campbell-Stokes sunshine recorders, with which the Meteorological Office measure sunshine duration. The sunshine recorder fails to detect sunshine if the beam irradiance at normal incidence is below a threshold of about 200 Wm^{-2}. (The actual value of the threshold increases with humidity.) Therefore we should expect that, for a given low solar elevation, the empirical sunshine duration will be less than that predicted by the geometric model. This is consistent with the observed differences in shape between the empirical and theoretical curves. Consequently, we cannot judge the validity of the model until the effect of the Campbell-Stokes cut-off has been either incorporated into the model, or obviated by the use of pyrheliometric measurements of beam irradiance.

The model itself seems to be sensitive to only two parameters: the relative areal cloud coverage, CA, and the ratio of width to depth of the cloud, w_c/d_c. This result might lead to a further simplifying of the model.

References

Hottel, H.C. (1976), "A simple model for estimating the transmittance of direct solar radiation through a clear atmosphere", Solar Energy 18, 129.

Lloyd, P.B. (January 1985), "Solar radiation monitoring at the UCC Newport Road site", UCC Mech. Eng. Dept Rept 1204.

McGregor, J. (April 1980), "Daily tabulations of meteorological and tilted surface irradiance data at Cardiff for the period July 1978 - April 1979", UCC Mech. Eng. Dept Rept 616.

- (June 1983), "Studies on the measurement of solar radiation: final report on CEC ESF-206-80-UK(H) and ETSU E/5A/CON/1040/032", UCC Mech. Eng. Dept Rept 1027.

Met.O. (1982), "Handbook of Meteorological Instruments. 6: Measurement of Solar and Terrestrial Radiation", Met.O. 919f, 2nd ed., H.M. Meteorological Office, London.

Moon, P., and D.E. Spencer (1942), "Illumination from a non-uniform sky", Trans. Illum. Engng. Soc. N.Y., 37, 457-465.

Robinson, G.D. (ed.) (1966), "Solar Radiation", Elsevier Publishing Company.

Schmid, W. (1976), "Aufbereitung und Qualitatskontrolle langjahriger Messunterlagen der Globalstrahlung und Himmelstrahlung", Working Reports of the Swiss Met. Inst. No. 59. Cited by Steven & Unsworth, 1980.

Steven, M.D. (October 1977), "Angular distribution and interception of diffuse solar radiation", PhD thesis, Dept Physiology & Environmental Studies, University of Nottingham School of Agriculture.

- (1977), "Standard distributions of clear sky radiance", Q.J. Roy. Met. Soc., 103, 457-465.

Steven, M.D., and M.H. Unsworth (1979), "The diffuse solar irradiance of slopes under cloudless skies", Q.J. Roy. Met. Soc. 105, 593-602.

- (1980), "The angular distribution and interception of diffuse solar radiation below overcast skies", Q.J. Roy. Met. Soc. 106, 57-61.

- (1980), "Shade-ring corrections for pyranometer measurements of diffuse solar radiation from cloudless skies", Q.J. Roy. Met. Soc. 106, 865-872.

Svendsen, D.A. (February 1977), "The measurement of diffuse short-wave radiation", UCC Mech. Eng. Dept Rept 340.

- (August 1978), "Inclined surface irradiance measurement: instrument operation and circuit diagrams", UCC Mech. Eng.

Dept Rept 427.

- (August 1978), "Inclined surface irradiance measurement: measured data, January 1976 - August 1978", UCC Mech. Eng. Dept Rept 428.

- (August 1978), "Inclined surface irradiance measurement: final report", UCC Mech. Eng. Dept Rept 431.

APPENDIX A: Software documentation

This appendix contains documentation of the two largest blocks of software written for the NERC project: the file-restructuring program STRUC; and the library of subroutines for handling the Svendsen data, DAS3LIB. Documentation for the other software will be written when time is available, and archived as internal reports; that task is less urgent because the other software is made up of relatively small units which can be fairly easily understood by reading the source code. All of the software contains extensive internal documentation in the form of COMMENT statements.

The purposes of the documentation are to enable other people to use or modify the software, and to provide a record of the processing and pre-processing that has been applied to the radiation data.

The documentation consists of annotated lists of the subroutines and common blocks, and of the variables inside the common blocks, together with a list of the commands that can be given to the program STRUC. Local variables of the subroutines are not described because their function can generally be seen fairly easily by reading the source code.

A.1 STRUC: program for accessing the archive tapes

The program STRUC is written in FORTRAN and, where possible, FORTRAN-77 was used. Unfortunately, the system subroutines for accessing the tape drive exist in FORTRAN-66 only, and the linkage editor on the GEC-4090 does not permit the mixing of subroutines from the two forms of FORTRAN. Therefore, those parts of STRUC that deal with tapes are necessarily written in FORTRAN-66. The files whose names end in "V" are written in FORTRAN-66, and can be compiled with both the FORTRAN-66 and FORTRAN-77 compilers (since the former language is a subset of the latter language); these are TOPV, STRUV, MTFIV, and RTGCV. Files whose names end in "3" are written in FORTRAN-77 and can be compiled only with the FORTRAN-77 compiler; at present, this is only DCFI3. (The choice of the codes "V" and "3" flows from the names of the respective compilers on the GEC-4090, namely "FORTV" and "FORT3".)

The program normally reads files directly from tape. Some files on the tape, however, are actually binary images of entire floppy discs. It is easier to handle these by first copying them onto a GEC-4090 disc and then breaking them up into their constituent files. In addition, files of logged data are often transferred by land-line from PDP-11 computers on the UCC site to the GEC-4090. If these files comprise binary data, then they are transferred as binary images and must then be restructured using the STRUC program. For generality, therefore, the term "source" will refer either to a tape or to a binary image of a PDP-11 file (or whole disc) held on a GEC 4090 disc.

The STRUC source code is stored in seven files, as follows.

TOPV

The file TOPV contains the main program, which merely calls the user-interface subroutine, USERIN. That, in turn, calls whatever subroutines are required to execute the user's commands. The purpose of this arrangement (as opposed to having USERIN as the main program) is to make it possible to incorporate USERIN in other programs.

STRUV

The file STRUV contains the subroutines that are in both the configuration of STRUC that reads from a tape source and in the configuration that reads from a disc source. It contains the following subroutines.

USERIN provides an 'interface' with the user, that is, it reads commands from the user's terminal, syntactically analyses the commands and calls the subroutines that are necessary to execute them. This is the largest subroutine in STRUV.

OPRQIN opens the source file that has been requested by the user. This subroutine calls lower-level subroutines to locate and open the source file, the appropriate routines being selected according as the source is on tape or disc.

OPRQOT opens a requested disc file for output.

COPYFI copies a file from the source to the output stream (which can be attached to a disc file, or the lineprinter, or the user's terminal). It calls an appropriate copying subroutine according to the type of the file that is being copied. The copying routines are in COPYLIBV (see below).

IDFORM is a function that identifies a file-type on the basis of the 'extension' of the file's name (in the RT-11 syntax for file names). For example, the "PPP" in "JUN108.PPP" identifies the file as having the internal structure that is designated PPP.

TABCON converts tab characters into appropriate numbers of spaces. (The tab character is interpreted differently on the PDP-11 and GEC-4090 machines. The correct interpretation is essential for the valid copying of FORTRAN source-code files.)

FLINIT opens the source file.

LSTDIR lists an RT-11 directory (either of a tape, or of a binary image of a PDP-11 floppy disc).

The following are auxiliary subroutines:

INBYTE initialises the getting of bytes from the input buffer.
NXBYTE gets the next byte from the input buffer.
NXNBYT gets the next N bytes from the input buffer.
NXZBYT gets subsequent bytes until the end of the input buffer.
NXI2RE gets the next two-byte word and reverses it.

ERROR handles errors when bad blocks are detected on the tape.

MTFIV

The file MTFIV contains fairly low-level subroutines for handling the tape. It is written in FORTRAN-66, and makes extensive use of the lower-level system subroutines in the library MTVLIB. Note that an individual physical block on the tape may hold more than one logical block packed into it, and some of the MTFIV subroutines are necessary for unpacking the blocks. This file contains the following subroutines:

OPENIN opens the tape stream for reading.

INIPAK initialises the information that governs the unpacking of logical blocks from physical tape blocks.

NXBLOK gets the next logical block from the tape. (This may involve getting the next physical block from the tape.)

GOBLOK go to a specified logical block in the tape.

GOMARK find the next tape mark.

BACKMK move the tape back to the nearest past tape-mark.

REWIND rewind the tape back to the load point.

CLOSIN close the tape stream.

RTGCV

The file RTGCV contains subroutines that carry out processing which is specific to the RT-11 and GEC-4090 operating systems. (The other subroutines, in MTFIV, DCFI3, STRUV, etc., are independent of the choice of operating system.) It contains the following subroutines.

NORNAM normalises RT-11 file names, to facilitate subsequent processing.

MAKNAM constructs a suitable filename on the GEC-4090 for a file that has been read from an RT-11 source.

INITIN identifies the source device and carries out initialisations.

OPOTFI opens a GEC-4090 file for output. (This is at a lower level that the subroutine OPRQOT, which was mentioned above.)

OPRTMT opens a standard RT-11 tape and checks the header.

OPROMT opens a ROLLIN-format RT-11 tape and checks the headers.

OPRTNF opens the next RT-11 file in the source.

GETDIR extracts an RT-11 directory from the source. In the case of a tape, the directory is constructed by opening each file successively and recording its name.

RDDIR reads an RT-11 directory from the source.

NXTENT gets the next entry from an RT-11 directory.

SKIPNF skips over a file in the source.

RESET rewinds the tape and goes to logical block 0.

CLOSEF closes the tape stream.

DCFI3

The file DCFI3 contains subroutines that are used to access a source that is held on disc. These subroutines largely mirror those in MTFIV. The need for these subroutines can arise in two ways. First, it will arise when a binary image of a floppy disc is held on tape, and is read onto the GEC-4090 disc. The physical tape blocks will then be stored on the GEC disc as individual FORTRAN records; but they will still be referred to as 'physical tape blocks'. Second, a file containing binary data may be transferred to the GEC computer via land-line.

OPENIN opens the source.

INIPAK initialises information about the unpacking of logical blocks from physical tape blocks.

NXBLOK gets the next logical block from the source.

GOBLOK skips to a specified logical block in the source.

GOMARK goes to the next tape mark in the source.

REWIND restarts the reading of the source from its beginning.

CLOSIN closes the source stream.

COPYLIBV

The file COPYLIBV contains a library of subroutines for copying files from their original PDP-11 formats to appropriate formats on the GEC-4090. Each subroutine consists of a basic loop which involves extracting a logical record from a source block (which may be situated either on tape or on disc) and writing it to the output file, which is held on disc. Each subroutine includes a FORTRAN ENTRY statement with an entry-name of the form "COPYnn", so that the subroutine can be called by its reference number. For instance, the first subroutine is called COPIMA but has an ENTRY name of COPY01, so it can be called by either "CALL COPIMA" or "CALL COPY01". In the following list, the reference number of each subroutine is given after the name.

COPIMA 01 copies files as binary images, without any restructuring.

COPQQQ 02 copies QQQ files, which are Svendsen's files of raw logger data.

COPTXT 03 copies files as text; this includes files containing the source code of programs.

COPRRR 04 copies RRR files, which are Svendsen's files of partially processed logger data.

COPPPP 05 copies PPP files, which are Svendsen's files of fully processed logger data.

COPRNV 06 copies files comprising binary data, which are in the standard structure that RT-11 FORTRAN normally uses for binary data.

COPSCR 07 copies Scriven's files of raw logger data.

COPDIR 08 copies DIR files, which are Svendsen's directories to RRR files.

COPMOR 09 copies files from McGregor's Meteorological Office Rhoose tape.

COPMET 10 copies McGregor's METDAT file.

COPTO3 11 copies files from the data tapes received from Toronto.

COPDDD 12 copies DDD files, which are McGregor's files of processed logger data.

AUXV

The file AUXV contains a small library of minor auxiliary subroutines for data handling. They are taken from other libraries written by the present author, PTEXT and PFILE.

I1TOI2 converts a one-byte integer into a two-byte integer.

TESTON tests whether or not a specified stream is online (to the user's terminal).

PFGARG gets an argument from the RUN command that started the program running.

YESNO gets the user's answer to a yes/no question.

Common blocks

The subroutines communicate with one another via a number of common blocks, which are as follows.

/GOTDIR/GOTDIR is set to TRUE if the directory of the source has been obtained.

/SESION/OPENED is set to TRUE if the source has been opened.

FORMDF is the index to the default copying subroutine; it is initialised to zero to indicate that no default has been explicitly set by the user; it is used in the subroutine COPYFI.

/DEVTYP/DEVTYP is reserved for future use (when the tape handling can be done in FORTRAN 77).

EXIDIR is set to TRUE if the source contains a directory.

EXIHED is set to TRUE if the files on the source have file headers.

/LUN100/LUN100 is an internal buffer, used with logical unit 100.

/ONLINE/ONLINE is set to TRUE if the program is being run online to a terminal, as opposed to being run in a background job.

/RECDAT/INGERS is set to TRUE if the user has declared that a file to be copied contains only 16-bit integers.

FIXLEN is set to TRUE if the user has declared that a file contains only records of a fixed length.

/PTAPEF/IFILE is the index number to the currently open file on the tape; it is initialised to zero to indicate that no file has yet been opened.

FNAME is the name of the currently open file.

/IENTRY/IENTRY is the index number to the directory entry (if there is one) of the currently open file.

/TOSPEC/TOSPEC is the user's specification of the output destination, which may be the name of a file or the name of a GEC device.

/FFORMS/NFORM is the number of filename templates that have been defined.

KFORMS is the list of indices to copying subroutines that corresponds to the filename templates.

SFORMS is the array of filename templates.

/TABDAT/TABFLG is set to TRUE if any tab stops are currently set.

NTAB is the number of tab stops that have been set.

KTABS is the array of positions of the tab stops

/OBLOCK/OBLOCK is the block number that has been designated as a 'false origin' for relativising the indexing of the block numbers; it is initialised to zero, which is the true origin.

/BYTBUF/IBYTE is the index to the next byte to be processed in the the block of data that is currently being transcribed.

I0BYTE is the index to the first byte to be processed in the current block.

NRBYTE is the number of bytes to be processed in the current block.

BLOCK is a buffer that holds the current block.

/IGNORE/NIG is the number of bad blocks to be ignored before abandoning the attempt to read the tape.

/NOCATS/NOCATS is set to TRUE if no implicit catalogue name has been set for the GEC output file.

/MTNXEN/IWORD0 is the index to the first word in the RT-11 directory.

IWORD is the index to the next word in the RT-11 directory.

ISEG is the index to the next segment in the RT-11

directory.

/PTAPEB/IUNIT is the logical unit number of the tape source.
IBLOCK is the index to the current block in the tape.
I0PACK is the index to the current logical block within the current physical block.
NPACK is the number of logical blocks in a physical block (1 for ordinary tapes, 8 for ROLLIN tapes).
KREPLY is the error code that is returned by the most recently called system subroutine; a non-negative value indicates that there was no error.
WRMODE is set to TRUE if the source has been opened in write mode; this is not normally used.

/PTAPEP/ISUB is the index to the next logical block within the current physical block (1 for ordinary tapes, from 1 to 8 for ROLLIN tapes).
PAKBUF is a buffer that holds the contents of the physical block while it is being unpacked.

/PDTOGC/IUNIT is the logical unit number for the disc source.
IBLOCK is the index to the current logical block within the current physical block.
NPACK is the number of logical blocks packed into a physical block.
TXFLAG is set to TRUE if the output is to be sent to a text file.
EOF is set to TRUE if the end of the source has been reached.

How to use the program STRUC

When the program has been compiled, execution is begun in the normal way, by issuing the %B command to the operating system. The general proforma is

```
RUN %B FROM  = <source> -
       IN    = <command stream> (*) -
       OUT   = <advice stream>  (*) -
       LIST  = <directory stream> (SINK) -
       STR10 = <log stream> (SINK)
```

in which the usual OS4000 conventions apply: chevrons "<" and ">" denote items specified by the user; the default of each argument is given in parentheses; the keywords "RUN" and "FROM" etc. are optional; the asterisk "*" signifies the user's terminal; and the hyphen "-" is a line continuation mark.

The FROM argument (stream 3) must be the name of the source: if the source is a tape, then the FROM argument must be "MTD0"; if the source is a disc file, then the FROM argument is the name of the file. The IN and OUT streams are by default connected to the user's terminal, although they can be attached to files if the user so wishes. Additional information about the tape is sent to stream 10. By default, this is sent to SINK; but the user can receive this information at his terminal or in a file if he so wishes. Examples:

```
%B MTD0
%B MTD0 STR10=*
%B .FLOPPY
```

The program will then prompt for the user's instructions by displaying

```
Command:
```

The possible commands are as follows. A command may be in upper or lower case, and only the first four characters are significant.

OPEN

This opens the input stream, and must precede any commands that access the source. The first few blocks are examined to ascertain what kind of source it is.

WHAT

This command tells the program to report to the user whatever information it gleaned from reading the first few blocks of the file. The command will be accepted only after OPEN has been issued.

DIRECTORY

The command DIRECTORY causes a directory of the files held by the source to be listed. Three possibilities can be distinguished. First, if the source is a binary image of a disc, then it will contain an existing directory, which the program will read and display. Second, if the source is a tape containing files which have recognisable header records, then the program will read through the whole tape, making a note of each file name and thereby construct a directory which it then lists. Third, if the source is a tape whose files have no headers, then it is impossible to produce a directory.

The directory listing gives one line per file, containing the filename, the creation data, the size (in blocks of 0.5 kilobytes), and the position within the source of the first block of the file.

The program will prompt for a stream number for the listing. This may be stream 2 for the terminal, or stream 6 for a file or the lineprinter. If stream 6 is to be used, then the user must have attached stream 6 to the required destination when the "RUN %B" command was issued.

COPY

The command COPY copies a file from the source to a named destination file on disc. The program will prompt for the name of the file to be copied and the name of the file to which it is to be copied. Note that the file to be copied must have a name in RT-11 format and the output file must have a name in OS4000 format. The output can be sent to the user's terminal by typing

"*" in response to the prompt. For example (with the user's response underlined):

```
From file: JULPPP.PPP
To file: .JULPPP/LSB

From file: TAUFIT.FOR
To file: .TAUFIT

From file: TAUFIT.FOR
To file: *
```

If the file contains binary data, as opposed to text, then "/LSB" must be put at the end of the destination filename. In filenames of the RT-11 format, the characters after the dot (".") constitute the 'filename extension', and this normally denotes the type of file. E.g. file types FOR and PPP are FORTRAN source-code files and Svendsen's PPP files, respectively. For the commonest file types, special subroutines have been written for translating the file into a format that is usable on the GEC-4090. Some file formats are recognisable only by the main part of the filename, rather than by the filename extension. In order to recognise these file types, there is an array of template strings set up by the subroutine FLINIT. Actual filenames are compared with each of these strings and if a match is found, then the program uses the appropriate copying routine in COPYLIBV. The current set of templates are as follows, with the corresponding copying subroutine named in parentheses. As is normal in RT-11, the character "%" is wild and can stand for any other character.

```
%%%%%%.QQQ  (COPQQQ)  Svendsen's QQQ logged-data files
%%%%%%.QQD  (COPQQQ)  Svendsen's QQQ logged-data files
%%%%%%.PPP  (COPPPP)  Svendsen's PPP logged-data files
%%%%%%.RRR  (COPRRR)  Svendsen's RRR logged-data files
%%%%%%.QSI  (COPRNV)  Taylor's QSI logged-data files
FTN20 .DAT  (COPSCR)  Scriven's logged-data files
%%%%%%.DDD  (COPDDD)  McGregor's DDD logged-data files
METDAT.     (COPMET)  McGregor's meteorological data files
%%%%%%.MAC  (COPTXT)  programs written in MACRO-11 assembly language
%%%%%%.TXT  (COPTXT)  documents
%%%%%%.FOR  (COPTXT)  FORTRAN source-code files
%%%%%%.COM  (COPTXT)  RT-11 command files
%%%%%%.DAT  (COPTXT)  ASCII-coded data files
%%%%%%.HLP  (COPTXT)  help text files
%%%%%%.DIR  (COPDIR)  directories to Svendsen's RRR files
```

The complete file name in the RT-11 system can be regarded as consisting of nine characters (the dot being a punctuation sign). But the maximum length of a filename on the GEC system is only eight characters. To get around this problem, the file-name extension (e.g. "PPP") is automatically regarded as a catalogue name on the GEC. So, the file JUN068.PPP would be copied into the file .PPP.JUN068. The catalogue that is thus implied is automatically created (if it does not already exist) and used by the STRUC program; there is no need for the user to specify it. Unfortunately there seems to be a bug in the program which prevents the user from over-riding STRUC's choice of catalogue.

When the copying has been completed, the user is again given the "Command:" prompt.

TEMPLATE

The list of filename templates given above can be increased permanently by editing the subroutine FLINIT. Alternatively, the user may temporarily introduce new templates with the TEMPLATE command. This prompts the user for a template string, and then for the subroutine number. For example,

```
Filename template? %%%%%%.XYZ
Reference number of copying routine? 3
```

Thereafter (for the duration of the program run), files whose names match the specified template will be copied with the specified subroutine. If the specified template is already present in the list, then the specified routine number will replace that given in the list.

DEFAULT

The user can nominate a default copying subroutine, which is to be used if the program cannot identify a file type from the list of templates. When the DEFAULT command is issued, it prompts "Default format?". The user should then type in the index number of the subroutine number that he wants to use.

DEFINE

This command allows several parameters to be defined at one time. At present, only three parameters are settable with the DEFINE command, but there is scope for including more. Each parameter is prompted for, as follows.

(a) "Reference number of default copying subroutine? "

This is the same as in the DEFAULT command.

(b) "INTEGER*2 fields only? (Y/N) "

The answer "Y" should be given if the file to be copied contains only binary data in the form of two-byte integers. In the PDP-11, two-byte integers contain the bytes in the reverse of the natural order (with the low-order byte first), and must be reversed before the data can be used on the GEC-4090. If the file contains a mixture of two-byte integers and other data fields of other sorts, then the user must write his own program to do the conversion (or add a new copying routine to STRUC to handle it).

(c) "Fixed-length records? (Y/N) "

The answer "Y" should be given if all the records are of the same length, or if the file contains text rather than binary data. If the records are of variable length, and if the file holds binary data, and is to be copied with the general routine COPRNV, then each output record will begin with a two-byte integer containing

a count of the number of bytes in the record. Obviously, this would be unnecessary if all the records are the same length. (Note that the FORTRAN READ statement cannot handle binary records whose lengths are unknown beforehand.)

TABS

There are two differences in the interpretation of tab characters between the PDP-11 and the GEC-4090: first, most of the software on the PDP-11 is designed to accept tab characters in place of spaces (for instance, the FORTRAN compiler allows tab characters to be used in columns 1 to 6, whereas the GEC FORTRAN does not); second, the default positions of tab stops are different between the two machines. Therefore it is often necessary to convert tab characters into an appropriate number of spaces when taking files from the PDP-11 to the GEC-4090. The TABS command allows the user to specify the tab stops that are to be used when copying text files. When the TABS command is issued, it prompts the user for a sub-command, "Tab command: ". The permissible sub-commands are as follows.

ON switches on tab conversion
OFF switches off tab conversion
AT sets a tab stop at a specified position, for which the user is prompted "Character position?"
NOT unsets a tab stop, again the user is prompted for a position
LIST lists the currently defined tab positions
HELP lists the tab sub-commands
QUIT leaves the tab controller and returns to the main user interface

The tab settings remain in force until the end of the program run.

TITLE

The user may wish to have a title affixed to a directory listing, especially if this is to be sent to the lineprinter. The TITLE command prompts for such a title, which is stored and used until it is reset or until the end of the program run.

CLOSE

The command CLOSE closes the source stream.

QUIT

The command QUIT stops the program run.

A.2 DAS3LIB: Basic subroutines for analysing Svendsen data

A number of computations and actions are required repeatedly during the processing of the Svendsen data. The more important of these have been coded into subroutines and placed in a library that has been called DAS3LIB. The central routines are INITS, INITRS, and READS, which allow Svendsen files to be read. The data files may be either RRR or PPP, but not (yet) QQQ. It is essential to include the appropriate sequence of letters ("RRR" or "PPP") in the filename, as it is by this means that the program ascertains of what type the file is.

Throughout these subroutines: stream 3 is used for reading data from the Svendsen file; stream 10 is used to report technical information about the data as each sub-file is opened; stream 2 is used for reporting errors and other matters to the user at his terminal.

The current set of subroutines comprised by DAS3LIB are as follows.

INITS opens a Svendsen data-file and initiates the process of reading from it. It calls the subroutine INITRS, which initiates the reading of a sub-file.

INITRS initiates the reading of a sub-file: it reads the header record, which includes the date; converts the date into a number of computationally more efficient forms, such as the century day; reports the new sub-file on stream 10; initialises certain flags and counters.

READS reads the next record from the Svendsen file; it also updates an internal clock, and delivers to the calling program both the irradiance data and the time and date; upon finding a new sub-file, it calls INITRS to initiate the reading of it.

READSN reads the next n records from the Svendsen file, and delivers to the calling program the average of each channel over those n time-steps. It functions by calling the subroutine READS repeatedly, and hence the time will be set to that of the last of the n records.

SETPER sets the period within which data are to be supplied from the Svendsen file. If the subroutine SETPER is not called, then the subroutine READS will supply to the calling program all the data records that it finds in the file. If, however, SETPER is called, then the subroutine READS will simply ignore all data records that fall outside the specified period and supply only those records that lie within that period.

YERDAY is a function that converts a calendar date (expressed as year, month, and day) into a year day-number (range 0-365).

INIGEO initialises geometric variables that do not change with time, such as trigonometric functions of the latitude.

DAYGEO performs geometric calculations that need to be computed only once a day, e.g. the evaluation of the solar declination, which varies so slowly that a single, midday value can be employed for each day.

SCAGEO performs geometric calculations that need to be computed for each scan of the pyranometer cluster.

TIMGEO performs geometric calculations for a specified moment in time; this is an entry point within SCAGEO.

TOPATM computes the solar irradiance at the top of the atmosphere (the 'extra-terrestrial' irradiance).

HOTTEL computes the solar irradiance under a cloudless and very clear sky, according to Hottel's simple model.

HEMINI carries out precomputation of trigonometric quantities that are to be used by subroutine HEMINT.

HEMINT integrates a radiance function over the full hemisphere of view, as seen by a surface of arbitrary tilt and azimuth.

RININI carries out precomputation of quantities that are to be used by the subroutine RININT.

RININT integrates a radiance function over the sector of the sky that is obscured by a shade ring.

ISONUM is a function that returns the value of radiance at a specified point in the sky according to the isotropic model.

STEVE1 is a function that returns the value of radiance at a specified point in the sky according to Steven's anisotropic model, computed with the short formula.

STEVE2 is a function that returns the value of radiance at a specified point in the sky according to Steven's anisotropic model, computed with the long formula.

The subroutines communicate with one another via a number of common blocks, which are as follows.

/ISUNIT/ISUNIT is the logical unit number (i.e. stream number) for the input from the Svendsen file; this will normally be 3, but can be set to any value when the subroutine INITS is called.

/EOFSVE/EOFSVE is set to TRUE if the end of the Svendsen file has been encountered.

/SITDEF/LATIT latitude of the site, degrees.
LONGIT longitude of the site, degrees.
ALTIT altitude of the site above mean sea level, metres.

/SURDEF/AZIMS is an array of 24 azimuth values corresponding to the measurement planes of the scanning

pyranometers.

TILTS is an array of 24 tilt values corresponding to the measurement planes of the scanning pyranometers.

/SDATE/ YEAR year number (range 76 to 81).

YRDAY year day-number (range 0-364, or 0-365 in a leap year).

MONTH month number (range 1-12).

DAY day number (range 1-31).

/QEDAY/ QEDAY is the quadrennium day-number, i.e. the day number with a four-year leap cycle (range 0 to 4*365).

/HEADAY/CENDAY is the century day-number as read from the Svendsen file; i.e. the day number within the year.

DECLIN is the value of the declination as computed by Svendsen.

/PTIME/ HOUR hour number (range 0-23).

MINUTE minute number (range 0.0 to 59.9999).

SOLTIM true solar time (hours, range 0.0 to 23.9999).

/REQPER/REQPER is set to TRUE if a 'requested period' has been set up with the subroutine SETPER.

/PERIOD/YRDAY1 is the year day-number of the start of the 'requested period'.

HOUR1 is the hour number of the start of the 'requested period'.

MIN1 is the minute number of the start of the 'requested period'.

YRDAY2 is the year day-number of the end of the 'requested period'.

HOUR2 is the hour number of the end of the 'requested period'.

MIN2 is the minute number of the end of the 'requested period'.

/FOUND/ FOUND is set to TRUE if a specified period of data has been requested and the start of the period has been found.

/SEASON/SEASON is the index of the shade-ring season; this is 1 for Summer and 2 for Winter.

/SOLCON/SOLCON is the solar constant, 1367 W/m2.

/GROUND/ALBEDO is the assumed albedo of the terrestrial surface surrounding the scanning pyranometers.

GLOBAL is the current value of the global irradiance, used in computing the ground-reflected diffuse irradiance.

/RSHADE/NZO is the number of intervals into which the range of solar zenith angles has been split, for each of which the shade-ring fractions have been computed.

NAO is the number of intervals into which the range of

solar azimuth angles has been split, for each of which the shade-ring fractions have been computed.

Z0STEP is a two-element array holding the seasonal sizes of the step in zenith angles that was used in constructing the array of shade-ring factors; in this and the subsequent arrays in this common block, the two seasons are those of the Summer and Winter shade rings, respectively.

A0STEP is a two-element array holding the seasonal sizes of the step in azimuth angles that was used in constructing the array of shade-ring factors.

Z0NOON is a two-element array holding the seasonal maximum noon values of the solar zenith angle.

A0RISE is a two-element array holding the seasonal solar minimum value of solar azimuth (as measured from North) at sunrise.

RINRAD is an array holding the values of the predicted irradiance received from the sector of sky that is obscured by the shade ring, relative to the horizontal diffuse irradiance; if the isotropic model was used, then there will be a single value for each season, but if an anisotropic model is used then then RINRAD will contain two seasonal values for each solar position.

UNORAD is like RINRAD, except that it holds the values of the irradiance received from the unobstructed sector of the sky.

KSCHAN is an array of the channel numbers that are to be used in the regression to get the beam irradiance.

/SITDAT/SINLAT is the sine of the site's latitude.

COSLAT is the cosine of the site's latitude.

TANLAT is the tangent of the site's latitude.

SINTIL is a 24-element array of the sines of the tilts of the measurement planes of the scanning pyranometers.

COSTIL is a 24-element array of the cosines of the tilts of the measurement planes of the scanning pyranometers.

TANTIL is a 24-element array of the tangents of the tilts of the measurement planes of the scanning pyranometers.

SINAZI is a 24-element array of the sines of the azimuths of the measurement planes of the scanning pyranometers.

COSAZI is a 24-element array of the cosines of the azimuths of the measurement planes of the scanning pyranometers.

/DAYDAT/SINDEC is the sine of the solar declination.

COSDEC is the cosine of the solar declination.

TANDEC is the tangent of the solar declination.

EQTIME is the equation of time.

ESDIST is the factor for modifying the beam irradiance to account for the annual variation in the Earth-Sun distance.

/SOLPOS/AZISOL is the solar azimuth angle.
ZENSOL is the solar zenith angle.

/HEMDAT/NASTEP is the number of azimuthal steps in the hemispheric integration of a radiance function.
NZSTEP is the number of zenithal steps.
ASTEP is the size of the azimuthal step.
ZSTEP is the size of the zenithal step.
SINAH is an array of the values of the sine of the azimuth angle for each integration element.
COSAH is an array of the values of the cosine of the azimuth angle for each integration element.
SINZH is an array of the values of the sine of the azimuth angle for each integration element.
COSZH is an array of the values of the cosine of the azimuth angle for each integration element.

/INTDAT/ANGTOT is the integrated total angular area.
ISOTOT is the integrated total irradiance from an isotropic field.

/RINDAT/NHAN is the number of steps in the imaginary hour angle to be used in the integration of a radiance function over the sector of sky obscured by a shade ring.
NDEC is the number of steps in imaginary declination to be used.
DECSTP is the step size of the imaginary declination in the ring integral.
DECLO is the minimum declination value to be used in the ring integral; this, of course, corresponds to solar noon.

/FILTYP/PPP is set to TRUE if the Svendsen file is a PPP file
RRR is set to TRUE if the Svendsen file is an RRR file.

/ECHO/ ECHO is set to TRUE if the the user has requested that all the data read from the Svendsen file is to be echoed to the echoing stream.
ECUNIT is the logical unit number of the echoing stream.

/EQTIMX/EQTIMX the equation of time, as evaluated by Svendsen's program RTOP and inserted into the PPP file. It is distinct from the equation of time computed in DAYGEO and stored in the common block DAYDAT.

/CODE/ CODE the code number assigned to the data file by the operator when running the program CASS. The meanings of codes 1 to 14 are given by Svendsen (1978, Rept 430); the meanings of codes 15 and 16 are unknown. The subroutine INITRS merely checks that the code number is 10 or higher, i.e. that it is reading a valid Svendsen file with data from the scanning pyranometers.

/NEWSEG/NEWSEG is set to TRUE at the start of a new segment (or 'sub-file') of the Svendsen data file.

The subroutines in DAS3LIB do not interact directly with the user. All such interaction is carried out by the calling program.

APPENDIX B: Listing of program for geometric model, ZENDMODL

```
C==== Program for predicting the zenithal distribution of the relative
C     duration of bright sunshine in a partly clouded sky
      PROGRAM ZENDIST
      PARAMETER (NGAP=3,DSTEPN=0.6)
      LOGICAL FIRST
      CHARACTER CH,FTEXT*5
      DOUBLE PRECISION PI,RADFAC,Z,ZBOT,ZTOP
      COMMON /GEOM/HEIGHT,DEPTH,WIDTHC,WIDTHG
      DATA PI/3.14159265D0/
      RADFAC=PI/180
C---- Enable long lines in the output
      CALL RECLEN(6,132)
C---- Initialise graphics
      CALL GRINIT('SCALEPLOT; LINES; A4')
      CALL XAXIS(10,0.0,90.0,'Solar elevation angle')
      CALL YAXIS(10,0.0,1.0,'Predicted sunshine duration')
C---- Get the user-specified parameters
      CALL TXTOUT(2,42,"Cloud geometry: height,depth,width/depth? ")
      READ (1,*) HEIGHT,DEPTH,WRATIO
      WIDTHC=WRATIO*DEPTH
C     Echo the parameter values to the printed-output stream
      WRITE (6,'('' Height of cloud base: '',F4.1,'' km'')') HEIGHT
      WRITE (6,'('' Depth of cloud body: '',F4.1,'' km'')') DEPTH
      WRITE (6,'('' Width of cloud body: '',F4.1,'' km'')') WIDTHC
C     Set graphics title
      CALL TITLE7
     :(  'Height ',FTEXT(HEIGHT,'(F4.1)'),
     : ',  Depth ',FTEXT(DEPTH, '(F4.1)'),
     : ',  Width ',FTEXT(WIDTHC,'(F4.1)'),' (km)')
      CALL SUBTIT('[CA = areal cloud cover, C = apparent cloud cover]')
      CALL ANNO
      CALL CHARSZ(2.0)
C---- Loop for different cloud amounts starts here ----------------------------
C---- Get the required cloud amount from the user
5     CALL TXTOUT(2,22,"Areal cloud coverage? ")
      READ (1,*) CAREAL
      WRITE (6,'('' Areal cloud coverage: '',F4.2)') CAREAL
      WIDTHG=WIDTHC*(1/SQRT(CAREAL)-1)
C---- Initialise integration of total apparent cloud cover
      CTOTAL=0
      FIRST=.TRUE.
C---- Iterate through successively more distant viewed points
      DO 100 IVIEW=1,210
      IF (IVIEW.LE.10) THEN
         DSTEP=DSTEPN/10
         DVIEW=(IVIEW-1)*DSTEPN/10
      ELSE
         DSTEP=DSTEPN
         DVIEW=(IVIEW-10)*DSTEPN
      ENDIF
      ZVIEW=Z(DVIEW,HEIGHT)
C---- Evaluate contribution that cloud at viewpoint makes to local cloudiness
C--   Contribution due to cloud bases (CBASE):
      CBASE=CAREAL
C--   Contribution due to cloud sides (CSIDE):-
C     Position of fore edge of cloud:
```

```
      DFORE=DVIEW+WIDTHC
C     Position of aft edge of cloud:
      DAFT=DFORE-WIDTHC
C     Position of end of gap (on aft side of cloud):
      DGAP=DFORE-WIDTHC-WIDTHG
C     Initialise integration over possible positions of next row of cloud:
      PSUM=0
C     Iterate through possible positions of next row of clouds:
      DO 15 IGAP=-NGAP,NGAP-1
C     'Perturbed' width of gap between neighbouring clouds:
      WGPERT=WIDTHG*(IGAP/REAL(NGAP))
C     Add the projection to the accumulator:
15    PSUM=PSUM+PROJEC(DFORE,DAFT,DGAP+WGPERT,DVIEW)
C     Average the sum of projections over the set of possible positions:
      PMEAN=PSUM/(2*NGAP)
C     Finally, compute the apparent area of the sky occupied by the cloud side:
      CSIDE=SQRT(CAREAL)*PMEAN*(DAFT-DGAP)/(DFORE-DGAP)
C--   Total local cloud coverage = sum of base and side contributions:
      CLOCAL=CBASE+CSIDE
C---- Evaluate contributions from clouds further away than the viewpoint
C     Position of end of gap (on aft side of cloud):
      DGAP=DVIEW-WIDTHC-WIDTHG
C     Position of fore edge of cloud at viewpoint:
      DFORE=DVIEW
C     Amount of sky not yet obscured by calculations performed so far:
      CFREE=1-SQRT(CAREAL)
C     Position of fore edge of cloud in next row to be considered:
20    DFORE=DFORE+WIDTHC+WIDTHG
C     ... and position of aft edge of same:
      DAFT=DFORE-WIDTHC
C     Initialise integration over possible positions of next row of cloud:
      PSUM=0
C     Iterate through possible positions of next row of cloud:
      DO 30 IGAP=-NGAP+1,NGAP
C     Perturbation in position of cloud row:
      PERTUR=WIDTHG*(IGAP/REAL(NGAP))
C     Perturbed position of fore edge of cloud row:
      PDFORE=DFORE+PERTUR
C     Perturbed position of aft edge of cloud row:
      PDAFT=DAFT+PERTUR
C     Projection of cloud side into field of view:
      P=PROJEC(PDFORE,PDAFT,DGAP,DVIEW)
C     If the cloud row is too far away, abandon integration:
      IF (P.EQ.0) GOTO 35
C     ... otherwise add projection to accumulator and continue iteration:
30    PSUM=PSUM+P
C     Average the accumulated projections over all perturbed positions:
35    PMEAN=PSUM/(2*NGAP)
C     Evaluate contribution to local cloudiness
      CADD=CFREE*SQRT(CAREAL)*PMEAN
      CLOCAL=CLOCAL+CADD
C     If contribution to local cloudiness is still significantly large,
C     then reduce remaining free space and repeat loop for next row of cloud
      IF (CADD.GE.0.01) THEN
         CFREE=CFREE*(1-SQRT(CAREAL))
         GOTO 20
      ENDIF
C---- Compute & output coordinates of data point
```

```
C     Elevation angle:
      ELEVAT=90-ZVIEW/RADFAC
C     Sunshine duration associated with current elevation angle:
      SLOCAL=1-CLOCAL
      CALL DATUM(ELEVAT,SLOCAL)
C---- Add local cloud amount to total cloud amount
      ZBOT=Z(DVIEW+DSTEP,HEIGHT)
      ZTOP=Z(DVIEW,HEIGHT)
      AWIDTH=ZBOT-ZTOP
      ACIRCU=2*PI*SIN((ZBOT+ZTOP)/2)
      CTOTAL=CTOTAL+CLOCAL*(AWIDTH*ACIRCU)
C     Make a note of topmost Z value
      IF (FIRST) ZTOPMO=ZTOP
C---- Make a note of the coordinates of the starting point
      IF (FIRST) THEN
         STARTX=ELEVAT
         STARTY=SLOCAL
         FIRST=.FALSE.
      ENDIF
100   CONTINUE
C---- Normalise value of apparent cloud cover
C     Divide the angular area of the clouded part of the sky by
C     the angular area of the modelled part of the sky (max 2*PI)
C     [because a small patch of sky near the horizon may be unmodelled]
      ZBOTMO=ZBOT
      RANGE=2*PI*(COS(ZTOPMO)-COS(ZBOTMO))
      CTOTAL=CTOTAL/RANGE
C---- Print out the accumulated total cloud amount
      WRITE (6,102) CAREAL,CTOTAL,CTOTAL*8
      WRITE (2,102) CAREAL,CTOTAL,CTOTAL*8
102   FORMAT ('Areal cloud amount: ',F4.2,
     :'   Apparent cloud amount: ',F5.3,' ( = ',F5.3,' eighths)')
      CALL TEXT4(STARTX,STARTY,
     :' CA=',FTEXT(CAREAL,'(F4.2)'),' C=',FTEXT(CTOTAL,'(F4.2)'))
C---- Ascertain whether user wishes to run program for more cloud configurations
      CALL TXTOUT(2,13,"More? (Y/N)  ")
      READ (1,'(A1)') CH
      IF (CH.EQ.'Y'.OR.CH.EQ.'y') THEN
         CALL FRESH
         GOTO 5
      ENDIF
      CALL GRSTOP
      STOP
      END

      FUNCTION PROJEC(DFORE,DAFT,DGAP,DVIEW)
      PARAMETER (TSTEP=1,SPEED=0.01)
      DOUBLE PRECISION PI,Z,ZVIEW,ZTOP,ZGAP,ZTPRE,ZGPRE
      COMMON /GEOM/HEIGHT,DEPTH,WIDTHC,WIDTHG
      DATA PI/3.14159265D0/
      RADFAC=PI/180
      ZVIEW=Z(DVIEW,HEIGHT)
      ZTOP=Z(DAFT,HEIGHT+DEPTH)
      ZGAP=Z(DGAP,HEIGHT)
      IF (ZTOP.GE.ZVIEW) THEN
         PROJEC=0
         RETURN
      ENDIF
```

```
      TIME=0
      TDAFT=DAFT
      TDGAP=DGAP
10    TDAFT=TDAFT+TSTEP*SPEED
      TDGAP=TDGAP+TSTEP*SPEED
      ZTPRE=ZTOP
      ZGPRE=ZGAP
      ZTOP=Z(TDAFT,HEIGHT+DEPTH)
      ZGAP=Z(TDGAP,HEIGHT)
      IF (ZGAP.LT.ZTOP.AND.ZTOP.LT.ZVIEW) THEN
         TIME=TIME+TSTEP
         GOTO 10
      ENDIF
      TIME=TIME+TSTEP*(ZVIEW-ZTPRE)/(ZTOP-ZTPRE)
      TFREE=(DVIEW-DGAP)/SPEED
      PROJEC=MIN(1.0,TIME/TFREE)
C      WRITE (2,*) DAFT,DGAP,TIMGAP,PROJEC
      RETURN
      END

      DOUBLE PRECISION FUNCTION Z(DISTAN,HEIGHT)
      PARAMETER (RADIUS=12500)
      DOUBLE PRECISION T
      T=ABS(DISTAN)/(RADIUS+HEIGHT)
      Z=ATAN(SIN(T)/(COS(T)-RADIUS/(RADIUS+HEIGHT)))
      IF (DISTAN.LT.0) Z=-Z
      RETURN
      END
```