Supercomputers
Radar rainfall — PARAGON
Noctilucent clouds from Ben Nevis
Diagnosis of NWP products

THE METEOROLOGICAL MAGAZINE

HER MAJESTY'S STATIONERY OFFICE

March 1988

Met.O.982 No. 1388 Vol. 117

METEOROLOGICAL OFFICE Met.O.982

THE METEOROLOGICAL MAGAZINE

No. 1388, March 1988, Vol. 117

551.509.313:681.3

The application of supercomputers to weather forecasting*

K.M. Rogers
Meteorological Office, Bracknell

Summary

The aim of this paper is to examine the architectures of large conventional computers and of the new generation of supercomputers, in relation to using these designs in operational weather forecasting and other meteorological applications. The CYBER 205 and the CRAY X-MP range of computers are covered in more detail, since they are the current operational computer systems at the Meteorological Office and at the European Centre for Medium-range Weather Forecasts.

1. Conventional computers

1.1 *Basic concepts*

(i) *Design.* A typical conventional computer has the general structure and constituent parts illustrated in Fig. 1. The basic idea of this design is over 40 years old and was first put forward by Von-Neumann. The control unit co-ordinates the flow of data and instructions between the various parts of the system (for example, any transfers of data between main and secondary memory) thus leaving the processor free to perform all the arithmetic and logical calculations. The transfer of information to and from secondary memory or the input/output (I/O) devices is much slower than that to and from main memory. The term 'hardware' applied to a computer system refers to the actual physical components of which it is comprised, whereas its 'software' consists of the programs which can be processed on the computer.

In the context of this discussion, the most important aspect of this design is the sequential execution of instructions, with the processor continually performing the following cycle of actions:

— fetch next instruction from memory
— decode instruction
— fetch operands from memory
— perform operation on operands
— store result

This type of computer can be classified as a 'Single-Instruction stream Single-Data stream' (SISD) computer, since one instruction operates on a single set of operands, and only one operation may be

* This article is based on an investigation undertaken as part of the 1986/87 Scientific Officers' Course held at the Meteorological Office College, Shinfield Park, Reading.

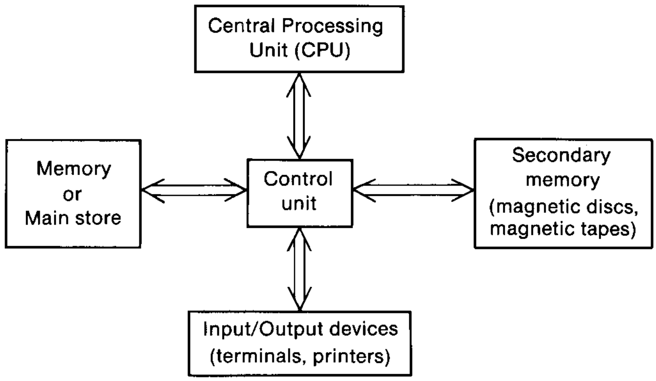

Figure 1. A simplified diagram of the basic units of a conventional computer.

performed at one time. An example of an SISD computer is the IBM 360/195, the main computer used at the Meteorological Office from 1971 to 1982. The scheme of classification used here divides computers into four general classes and is described by Hockney and Jesshope (1983). Computers in some of the other classes will be referred to later.

(ii) *The compiler.* The instructions will initially be written in a 'high-level' language such as Fortran. High-level languages have 'statements' to be interpreted by the computer, consisting of a restricted set of English words and symbols, each given a precise meaning in the language. They are designed to be hardware independent and are much easier for the programmer to understand and to use than the 'low-level' language which corresponds directly to the machine's instruction set and is therefore unique to a given type of computer. A high-level program is transformed ('compiled') to low-level machine code by a special program known as a compiler. The machine code must then be 'run' for the instructions to be actually executed. A program such as the operational weather forecast model only needs to be recompiled every few months, when changes are made to the program. The resulting program is then run twice a day (at the Meteorological Office) from its machine code form.

(iii) *The Operating System.* Every general purpose computer has an Operating System (OS), which is a set of interrelated programs. A major function of the OS is to simplify using the machine's hardware. On larger computers it provides a multi-programming environment for users and programs, in which a large number of users may appear to have 'simultaneous' access to the processor, each program being allowed a certain share of time in control of the Central Processing Unit (CPU), according to its priority. The OS also allocates resources, aids efficiency, and protects users and programs from each other.

(iv) *Operations and data.* The basic unit of data for numeric processing on a computer is the 'word' which is a consecutive sequence of binary digits, usually known as bits, of a fixed length for any one computer. A 64-bit word length is used in a number of present-day supercomputers. Operands and instructions are often an integral number of words in length. A sequence of 8 bits is normally referred to as a byte and these are typically used to store characters, one to a byte.

The main memory of modern computers can typically hold a few million words of data, stored in 'interleaved' memory banks. Accessing any one bank successively is slow, so generally consecutive words are stored in different banks to speed up their retrieval. Although secondary memory is usually provided in the form of magnetic discs and tapes, semiconductor or solid state memory, such as the

Solid-state Storage Device attached to the CRAY X-MP at the European Centre for Medium-range Weather Forecasts (ECMWF), has recently become sufficiently inexpensive to be used in this way. The semiconductor memory is much faster than the other types of secondary memory, thus going some way towards reducing the large difference between the CPU and I/O speeds of more conventional systems. Such a reduction is important because the large speed gap results in the system 'throughput' (the amount of work that can be processed by the system) being considerably less than that to be expected from the CPU processing rate.

All activities in the CPU are synchronized by a 'clock', with a clock cycle time in the range 5–50 ns (1 ns = 10^{-9} seconds) for the computers considered here. This gives the basic indivisible unit of time for the shortest operation, and gives a rough guide to the speed of the machine since the smaller the clock cycle time the faster operations can take place. If the clock cycle time is reduced to achieve greater speed, this must be backed up by correspondingly faster facilities and speeds on the rest of the system. However, one must also take account of the fact that one machine may do more or less work per clock cycle than another. The memory access time, which is usually a few clock cycle periods, gives the time to transfer one word of stored information from a memory bank to the CPU.

(v) *Comparison of performance.* There is no absolute standard for comparison of computer performance, but computer speeds are commonly measured in:

(a) MFLOPS (millions of floating-point operations per second) or
(b) MIPS (millions of instructions per second).

In comparing similar computer systems, it is standard practice to take a set of test programs, called bench-marks, typical of the work-load of an establishment, and run them on the different machines. It is particularly important to apply bench-marks to supercomputers since their performance varies considerably depending on the application, and may be much less than the 'peak rate' given by the manufacturer. (The performance is affected by factors such as the quality of compiler, the OS environment, the efficiency of mathematical library routines and the precision of floating-point arithmetic, and not just of the hardware.)

In scientific applications, the Fast Fourier Transform is often used to compare the performance of different computers, since it is an integral part of many scientific programs. There are also well known suites of programs which have become 'standard' bench-marks, e.g. Livermore Loops.

1.2 *Pipelined vector processors*

The computers outlined above can be described as scalar processors because one instruction is needed for each operation (e.g. adding a pair of operands).

In contrast, vector processors (e.g. CRAY-1, CYBER 205), a specialized form of the conventional computer designed to process large quantities of scientific data efficiently, operate on long vectors of data. One instruction now performs the same operation on each element of a vector in turn. In this context, a vector is just a sequence of numbers all of the same type, such as, for example, temperature values at a series of forecast model grid points. The method by which vector data are processed is known as 'pipelining'. Successful pipelining depends on being able to break down an instruction, such as an addition, into small segments, each of which can be performed several times faster than a complete operation.

Pipelining is analogous to processing an 'assembly line' of data. One pair of operands is added to the beginning of the assembly line each clock cycle and it takes a while before the line is full (the start-up time). The start-up time may be affected by the time taken to get the first operands from memory as well as by the length of the pipeline. Once the line is full, new operands are taken in, and one result is produced each clock cycle, with intermediate operations being performed on elements 'in the middle of' the assembly line. Obviously this is much more efficient than putting one element onto the assembly line and waiting until it emerges at the other end before putting the next element on the line, which is generally what would happen in a computer that did not have the pipeline facility.

The compiler on a vector processor automatically vectorizes sections of code (i.e. it enables the data to be pipelined) where this is possible. The only type of construct that can be vectorized is a loop of some sort, such as the 'DO-loop' in Fortran, and even this is vectorizable only under certain conditions. However, a good programmer or compiler will spot ways of formulating vectorizable code which may be far from obvious in a program written orginally for a conventional (scalar) computer.

A simple example of a DO-loop that vectorizes is given below. A, B and C are vectors, with A(I) being the Ith element of vector A, etc. and I takes on each of the values 1, 2, 3, ... up to a fixed integer N, in turn. For each value of I, the Ith elements of vectors B and C are added together, and the Ith element of A is given the value of this sum.

```
      DO 10 I = 1, N
         A(I) = B(I) + C(I)
   10 CONTINUE
```

When compiled, this Fortran loop would become a single instruction operating on a vector of length N. On the first clock cycle, B(1) and C(1) would be fed into the first part of the floating-point add pipeline. On the second clock cycle B(2) and C(2) would be fed in to the first part and the first operands would be moved to the second part of the pipeline. Similarly B(3) and C(3) are fed in on the third clock cycle, etc. After a few cycles the results A(1), A(2), etc. would start emerging from the add pipeline.

This is possible because the add operation can be broken down into several smaller operations. Firstly, the two numbers (in floating-point notation) to be added must be adjusted so that they have the same exponent. They are then added and finally the result is put back in floating-point notation before it leaves the pipeline.

Most vector processors have several such pipelined units. The main difficulty is in getting the correct instructions and data in the right place at the right time to keep the processor working efficiently. The machine is at its most efficient when the pipeline unit is kept full, i.e. for long vectors.

Numerical weather prediction models can be formulated to work with long vectors, particularly for the equations representing the dynamical processes of the atmosphere, where the same operations are generally applied at every point in the forecast domain. This is why the Meteorological Office purchased a CDC CYBER 205, and ECMWF a CRAY-1A, which has since been upgraded by stages to a CRAY X-MP/48.

Most vector processors are designed to be run in conjunction with a computer which acts as a front-end processor. In general, the front-end processor controls the preparation of input data, presentation of output data (a large task in weather forecasting systems), job preparation and other data management tasks. The vector processor is then free to spend all its processing time on processing numerically intensive work, the type of computation at which it is most efficient.

2. Supercomputers

2.1 *Introduction*

The term 'Supercomputer' is used by different people in various ways. It is sometimes used to denote the fastest computers currently available; this would include the CYBER 205 computer, manufactured by CDC, and the CRAY range of computers, manufactured by CRAY Research Inc. Alternatively, the term can be used to denote the advanced systems about to emerge in the next few years. These new computers typically have multiple processors and many have a completely different structure from the conventional computer, enabling them to exploit the data structures of the problems they are designed to solve.

In some of these computers, two or more processors execute the same instruction simultaneously on different items of data 'in parallel'. The array processor discussed in section 2.2, operates in this way, an example of which is ICL's Distributed Array Processor (DAP).

In other computers, multiple processors again work together, but execute *different* functions at the same time. For example, the CRAY X-MP has up to four CPUs as its processing elements. The meteorological applications of the CRAY X-MP at ECMWF are discussed more fully in section 4, and the general characteristics of this class of computers, called multiprocessors, are covered in section 2.3.

2.2 *Array processors*

The array processor is designed to operate on arrays of data. It consists of a large array of processors, each with its own memory and connections to neighbouring processors. A single instruction issued by a separate control processor operates simultaneously on multiple pieces of data (i.e. each element of the array), one in each processor. Therefore this can be classified as a 'Single-Instruction stream Multiple-Data stream' (SIMD) computer.

If a scalar problem is presented, and all real problems have parts which are irreducibly scalar, then only one of the many processors can be used; this is extremely inefficient. For this reason the array processor would normally be attached to a front-end processor which would perform these scalar tasks.

In theory this type of processor is well suited to performing meteorological work because the forecast models are based on data at a grid of points. It is easy to visualize a system being applied to numerical weather prediction in which there was one processor for each grid point, with communication to nearest neighbour processors. With such an arrangement, SIMD systems are ideal for dynamical calculations where each calculation is the same at every point. However, in dealing with physical calculations, such as those due to radiation or rainfall, certain computations will only be performed at some grid points and therefore many processors will be switched off for most calculations (see section 3.4). This reduces the efficiency of SIMD machines.

Some Meteorological Centres use array processors as a cost-effective way of running dynamically based forecast models. They are particularly applicable in situations where the forecasts of physical processes, such as precipitation, are less important than those of variables calculated from dynamical processes, such as wind. However, at the Meteorological Office, this type of computer is not general purpose enough or powerful enough to handle the many models and varied applications for which it is required. In particular, there would be major problems in achieving the fast input and output of data required for the running of the operational forecast models.

2.3 *Multiprocessors*

These may be classified as 'Multiple-Instruction stream Multiple-Data stream' (MIMD) computers, since each processor is working on a different task using separate data.

When a computer has many users, such a multiprocessing environment can be organized so that different programs are processed by different processors, as they become free, thus increasing the throughput. The OS organizes the allocation of programs to processors and so the programmer is not concerned with this aspect of the processing, in that the work is set up as for a single-processor system. Machines of this type are quite common (an example being the two IBM 3081D computers in use at the Meteorological Office) and such computers will not be covered further in this section.

Alternatively, the processors may all be used to process one program, with the processors executing different sections of the program concurrently. This can significantly decrease the elapsed time needed to execute a program such as an operational forecast model, but the program becomes more difficult to organize and overheads arise due to the need to synchronize the processors in order to ensure that each section is executed in the correct order. This whole area is subject to much current research.

In running an operational forecast model, the programmer can make the execution more efficient by setting up algorithms and strategies that are easy to execute as multiple processes. These processes would be large, to avoid overheads, and well balanced so that processors are idle for the minimum amount of time. If the program has been split into self-contained tasks and placed in a 'task bin', then as each processor becomes free it takes the task at the top of the 'bin' to process next, subject to task dependencies, such as one task having to complete before another one can start. In this type of multiprocessing, a 'task' is defined by Dent (1985) as, 'any block of executable code which can run independently of other blocks'. On the CRAY X-MP, a task could be a subroutine or may even be part of a DO-loop, the total number of passes through the loop being divided into several tasks.

The speed-up factor compares how much faster a multiprocessor environment is relative to a single-processor environment:

$$\text{speed-up factor} = \frac{\text{time taken for task in single-processing mode}}{\text{time taken for task in multiprocessing mode}} .$$

A speed-up factor of greater than one is only achieved if, as expressed by Gibson (1985), there is:

> sufficient parallelism to utilize more than one CPU to perform computations at a faster rate than could be achieved by a single processor, allowing for the extra overheads necessarily incurred by the multiprocessing control mechanism. The full benefit of multiprocessing can only be realised when the multiprocessing job has sole command of the computer. Multiprocessing should be reserved for time critical jobs which are likely to execute on the computer either in isolation, or at high priority.

It can be proved that for n processors, the speed-up factor must be less than n, since the overheads prevent the possibility of achieving a speed-up of n.

3. The CYBER 205 at the Meteorological Office

3.1 *The computer installation*

The present operational supercomputer at the Meteorological Office is the CDC CYBER 205, but this will be replaced during the next year. The CYBER, together with its two front-end processors, both IBM 3081D machines, form a very powerful processing environment. The IBM machines are general purpose machines designed to be efficient in a multi-programming environment, so they are ideal for providing the user interface to the CYBER as well as meeting much of the processing requirements in their own right. The two 3081s are identical and interchangeable but, at any one time, one of them is the control processor for both the other 3081 and the CYBER. This means that if one 3081 'goes down' it can always be replaced by the other 3081 and if the CYBER fails the combined 3081s can temporarily take over some of its processing, although global models are never run on the 3081s. The jobs given to the CYBER are those with a large amount of data which can be vectorized, since it performs this kind of processing efficiently, as discussed earlier.

Both the CYBER and the 3081s have disc storage which holds the operational forecast and climate models, as well as being used for temporary storage during the execution of programs. Magnetic tapes are used for archiving data and for back-up purposes.

The CYBER compiler has an automatic vectorizer which substitutes vector instructions for DO-loops, where such replacement cannot alter the program logic. In addition, there is an optimizer which reschedules the order of scalar instructions to make optimum use of scalar registers and pipelined scalar units without needing user intervention.

3.2 *Design features and limitations of the CYBER 205*

A vector on the CYBER is defined to be up to 65 535 contiguous elements (i.e. elements are stored in successive memory locations). This definition of a vector is rather limiting in that the elements have to be contiguous, but there are instructions to form vectors in consecutive storage locations from random elements of other vectors, and to reverse the process. This facility can be used to good effect when the required elements are defined by other data values, as in the modelling of physical processes (see section 3.4).

Execution of a vector operation consists of two phases. Firstly there is the start-up phase, which is independent of vector length and involves filling the general purpose pipelines. Secondly there is the stream phase during which the vector elements are processed. Here the time consumed is proportional to the vector length, where the constant of proportionality is the average number of clock cycles needed to produce one result. Table I, calculated from values given in Dickinson (1982), shows that for small vector lengths the start-up time dominates the MFLOPS rates. Note that the Meteorological Office's CYBER 205 has two pipes (i.e. two independent pipelines) so the figures given in Table I are directly applicable to the performance achievable on this system.

Table I. *Performance for 32-bit arithmetic operations on a two-pipe CYBER 205*

Vector length	Speed, in MFLOPS, for addition or multiplication
25	21.8
50	39.4
100	65.8
500	142.1
1 000	166.1
10 000	196.0
50 000	199.2
asymptotic performance	200.0

The use of 32-bit arithmetic (see section 3.3) gives a peak rate twice that of the standard 64-bit arithmetic on a CYBER 205. In addition, the use of a facility known as a linked triad, whereby the output of one vector unit is fed directly into the input of a second vector unit, can, where applicable, further increase speeds by a factor of two. This level of performance can be achieved only by use of careful programming, in which the code is tailored to the particular computer configuration. Such techniques can increase performance by an order of magnitude. By comparison, in scalar arithmetic program optimization can only increase performance by, say, a factor of two or three.

Once accessed, a memory bank on the CYBER cannot process another memory request for the same bank for four clock cycles. This can lead to very inefficient code, for example when there are successive accesses of the same array element or of array elements which are a multiple of the number of memory banks. There are 256 memory banks on the CYBER 205 when using 32-bit words.

3.3 *Use of the CYBER for numerical weather prediction and climate modelling at the Meteorological Office*

The time available on the CYBER is used primarily for numerical weather prediction and general circulation studies. In addition to supporting the Office's operational forecasting effort, research activities cover a wide range of topics which include such subjects as the study of climate change, the development and testing of a mesoscale model, experiments on the observing system such as data impact studies, long-range forecasting and research leading to a better understanding of the physical processes in the atmosphere.

All the above research activities require the use of atmospheric models that can resolve the physical and dynamical processes in time and three-dimensional space. The designs of the computer programs which run two such models are discussed in more detail in the following sections. These are the global version of the operational forecast model and the climate model. Between them these models account for probably 60% of the time used on the CYBER. Descriptions of the formulations of these models are given by Bell and Dickinson (1987) in the case of the forecast model and Slingo (1985) in the case of the climate model.

Since the operational forecast models are run each day to a tight deadline, they must be efficient and the programs must be reliable, so changes to the models are thoroughly tested before introducing them. The emphasis is on producing the most accurate forecast in the time available. In practice this means running models with the highest possible spatial resolutions. In contrast, an experimental run of the climate model generally takes much longer to execute than a typical run of the forecast model, since it is simulating atmospheric events on a time-scale of years rather than days. These models, by necessity, are designed to use much lower spatial resolutions than the one used in the forecast model. Program efficiency is again important because small increases in efficiency can save large amounts of computer time, although the necessity to use modified equations to ensure the simulated atmospheric dynamics is correct in the long-term means that these equations are somewhat less efficient to solve. Also there must be provision for the amount of data processing needed in the very long program runs and for the storing of large quantities of intermediate results. The program must, in addition, have the property of being restartable in the case of a computer malfunction.

The implementation of both the forecast and climate models on the CYBER uses 32-bit words instead of the full word length of 64 bits. This gives, in effect, twice as many words of store and also has the advantage that 32-bit vector arithmetic operates at speeds up to twice that of 64-bit arithmetic. In addition, it has been shown that using only half the precision has a negligible effect on the accuracy of the weather prediction and climate models.

The CYBER is operated in 'stand alone mode' for runs of large research tasks and during operational weather forecast model runs, that is, there is only one user on the machine at once. Not only does this make the entire memory and processing power available to a single program, but it also means that no CPU time is lost in changing over to another user and bringing other files into main memory. Hence the 'wall clock' time, the time taken for the program to run, is reduced even more. Some sections of the model are 30–50 times quicker than the equivalent best coding on the previous computer, an IBM 360/195, thus demonstrating how successful the vectorization and other improvements have been.

3.4 *The treatment of dynamical and physical processes*

From a vector programming point of view, large-scale models of the atmosphere consist of two quite distinct parts; the dynamical and the physical processes. The dynamical processes are represented by the momentum, thermodynamic, humidity, continuity and hydrostatic equations. Both the climate and forecast models use explicit finite-difference techniques to integrate this governing set of fluid dynamic equations. The physical processes, on the other hand, occur on length scales much less than that of the grid-point system and their statistical effects are represented in terms of the prognostic variables. The schemes for modelling these processes are characterized by sets of calculations that are both conditional and intermittent. The physical processes include such processes as convection, rainfall, radiation, surface exchanges and turbulent mixing.

The integration of the dynamical equations by finite-difference techniques requires each level of the forecast area to be covered by a regular grid of points at which the current values of the prognostic variables are stored. A set of linear equations, derived from the finite-difference approximations to the governing equations, are then used to step the forecast forwards one time step at a time by modifying the values stored at each grid point and at each level of the integration domain. This means that in general the same computations are done at every point on the grid, so the solutions to the dynamical equations vectorize well and thus are processed efficiently on the CYBER.

The physical processes, however, are much more difficult to vectorize efficiently. The most important aspect of these processes, in terms of modelling, is that they only apply 'randomly' at grid points; that is, they only apply at a varying non-contiguous subset of the forecasting area. For example, areas of rainfall would be present at some grid points but not at others, and in all likelihood those points forming a developing area of rainfall would not be a contiguous set and so would not form a vector for processing on the CYBER (see section 3.2(a)). It would be inefficient to process the whole vector, since unnecessary calculations will be performed at grid points which have no rain. On the CYBER 205 there are two ways of reducing this overhead once the number of active elements in a vector becomes small. These methods use the 'compress and expand' and 'gather and scatter' instructions. Instead of suppressing the storing of the results of a calculation at unwanted points, these instructions allow shorter vectors to be first constructed from the active elements of the long vector. Once this has been done, subsequent calculations may be carried out at the shorter, faster vector length. The circumstances of any given case will determine which method is most efficient. Dickinson (1982) discusses the application of these techniques in more detail.

3.5 *Data organization for vector computation*

The current sizes of the grids used by two of the Meteorological Office's main models are given in Table II. Other climate models are also used, with both greater and lesser horizontal resolutions.

Table II. *The dimensions of two of the models in use at the Meteorological Office*

	Global forecast model	Climate model
No. of levels	15	11
No. of points round each line of latitude	192	96
No. of points between North and South Poles	121	72

On the CYBER 205 vector processor, long vectors are required for efficient running, so the global operational weather forecast model uses 'horizontal slices' to give the maximum vector length of $192 \times 121 = 23\,232$. The timings given in Table I indicate that this vector length, used for dynamical processes, is 99% efficient. This division of the grid also gives the longest possible vector length for physical processes, after taking into account the removal of elements that do not require processing. The physical calculations tend, in practice, to have a typical vector length of around 2000, which still gives a processing rate that is about 90% of the peak rate.

In the climate model 'vertical slices' are used to produce vector lengths of $11 \times 96 = 1056$. The flexibility given by using a smaller-sized slice compensates for an efficiency of only 84%. At present the climate model can reside entirely in main memory, but if the model were to increase in size so that it used more than this amount of store, the smaller slices, when inactive, would be more easily swapped in and out of main memory. The computations for some grid-scale physical processes are done in blocks of a number of vertical slices placed together. Since many of the variables for these processes are derived from quantities at other vertical levels, the natural way to solve the physics is by layers over a small horizontal area.

4. CRAY computers

4.1 *The CRAY-1A at ECMWF*

At the commencement of its operational activities ECMWF purchased a CRAY-1A computer, a vector processor with one CPU, that was roughly equivalent in power to the CYBER 205 when performing 64-bit arithmetic.

The distinctive cabinet shape shown in Fig. 2 is characteristic both of the CRAY-1A and the CRAY X-MP series, with the vinyl-padded seating disguising power supply units. In this generation of supercomputers, finding the correct shape for the cabinets was important, because the wiring lengths, and hence the correct timing of the electronic signals, depended on it. Now that a whole CPU can be implemented on a single circuit board, this is becoming less of a problem.

Figure 2. The CRAY-1A situated at the European Centre for Medium-range Weather Forecasts.

4.2 *The CRAY X-MP/22 and the X-MP/48 at ECMWF*

In March 1984 a CRAY X-MP/22 was installed at ECMWF. It had two CPUs and two million words of memory, the two quantities to which the '22' in its name refers. This was double the number of processors and twice the amount of memory than there was previously on the CRAY-1A. Both the Solid-state Storage Device (SSD) and the Input Output Subsystem (IOS) were new fast devices making the running of the whole system much more efficient and the clock cycle period had also been reduced to 9.5 ns from 12 ns. It was estimated that there was an increase in throughput over the CRAY-1A of a factor of 3.

The CRAY X-MP/48, installed in January 1986, has four CPUs and eight million words of memory, another substantial increase in resources. In addition, the number of memory banks was increased, and the amount of SSD went up from 16 million words to 32 million words. It is estimated that the performance is on average 2–2.3 times better than the X-MP/22. Table III gives the speed-up of the two- and four-processor versions of the X-MP, over the one-processor version, when processing ECMWF's spectral forecast model.

Table III. *Multiprocessing timings for the spectral model on the X-MP/48*

No. of processors	1	2	4
Seconds per time step	19.3	10.3	5.4
Speed-up ratio	(1)	1.87	3.6

Since the ECMWF operational weather forecast model already contains the code to run it on a variable number of processors, it is possible to estimate the performance that would be given by greater numbers of processors, using the present multiprocessing strategy (see Dent 1985). From this, it is clear that, as in any multiprocessor system, the ratio of the speed-up against the number of processors decreases as the number of processors increases, so eventually adding more processors would make no significant difference or would actually be detrimental to the speed of the forecast model. This problem of diminishing returns is well known and is caused by communications overheads between processors and the need to synchronize tasks.

4.3 *Strategies for multi-programming numerical weather prediction models*

Consider a latitude–longitude–height grid as used in an operational forecast model. Suppose that the processing over this grid is to be spread over several processors. Then the grid could be split into latitude sections as shown in Fig. 3(a) with each processor being allocated its own section.

Firstly, there will need to be some transfer of boundary conditions across the dividing boundaries, in addition to the actual processing.

Secondly, there will be the same number of calculations due to dynamical processes to do at each grid point, but different amounts due to the physical processes. For example, there is more convective activity in the equatorial regions than in the other regions. Also there is more land in some latitude sections than others which means more calculations of the computationally expensive land surface processes will be necessary. This type of consideration will lead to unequal amounts of work for the various processors so one or more processor(s) may be idle while the others are still working, thus achieving less than the potential efficiency. Therefore, although the processors will keep in step overall through synchronization, there will always be inefficiencies arising from some processors being idle for part of the time.

On the two-processor CRAY X-MP, the ECMWF model is used to process one latitude row in the northern hemisphere at the same time as its corresponding one in the southern hemisphere (Fig. 3(b)). Consequently, for the two rows being processed simultaneously, effects due to latitude are more or less equal, although other more randomly distributed effects cannot be made equal. This method of processing is convenient in terms of the Fourier components and Legendre transforms appearing in the equations of the ECMWF spectral model, since they have corresponding northern and southern terms, which can be calculated simultaneously (see Dent 1985).

On the four-processor CRAY X-MP, originally two northern rows were processed at the same time as their equivalent southern rows. Although each north–south pair still needed to be synchronized when they completed a process, there was no need to synchronize between the two pairs of processors until the

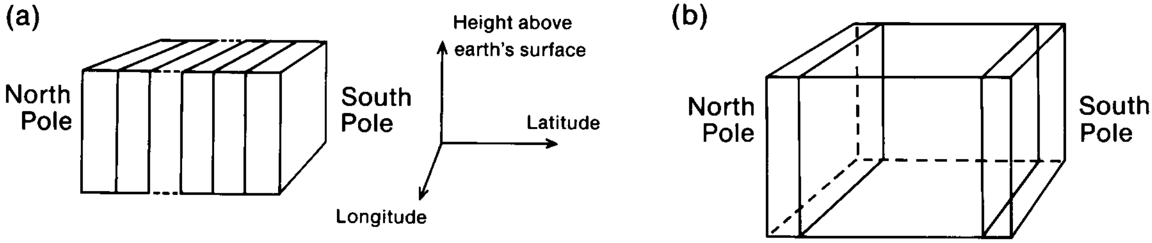

Figure 3. (a) The division of an observational grid into N sets of grid points in order to be processed by N processors and (b) the two corresponding northern and southern hemisphere latitude slices, that can be processed simultaneously in the two-CPU version of the CRAY X-MP.

complete global grid had been processed. This used a 'static' strategy in which most of the decisions concerning which task to execute in which order, the synchronization of the subprocesses and the handling of boundary transfers between subprocesses, were taken by the programmer at coding time.

Now a 'dynamic' strategy is being introduced (see Dent 1987) in which these decisions can be made at run time, processors taking on tasks as they become free. Each task still deals with a north-south pair of latitude rows but the need for these to synchronize (i.e. wait for each other) after each process is avoided by having each row updating different copies of the same files. This requires significantly more memory but reduces the time wasted when processors are idle. The dynamic method becomes more efficient than the static when using large numbers of processors because of these savings in time.

The speed-up over the single-tasked version of the dynamic strategy is an impressive 3.6. For this, the model executes 99% of the floating-point calculations in vector mode and has an execution rate of some 335 MFLOPS.

5. Future supercomputers and their use at the Meteorological Office

Both the speed and main memory of the most powerful computers have been increasing by approximately a factor of 10 every 5 years, a trend demonstrated by the main computers used at the Meteorological Office over the past 30 years. Present indications are that this general pattern of exponential growth will continue, at least for the next few years.

The latest generation of multiprocessing supercomputers includes the CRAY 2, from CRAY Research, of which several have already been sold, as well as the ETA10, manufactured by a subsidiary of CDC, and the CRAY Y-MP range, also from CRAY Research, both of which are now on the market. Each of these computers is capable of giving a processing rate an order of magnitude faster than the CDC CYBER 205. The precise speed-up will be dependent on the problem being solved, the numerical precision used and the number of processors. For example, an ETA10 with eight processors is projected to have a maximum performance of around 20 times that of the CYBER 205.

Although these are multiprocessor systems, they use only a relatively small number of CPUs, in which each CPU is essentially a vector processor 'supercomputer' such as the CYBER 205. If massive parallelism can be fully exploited, larger numbers of simpler processors could lead to cheaper and more powerful systems. Some manufacturers have already built, or are developing, systems with many processors. For example, the FPS T-series comprises, in principle, up to several thousand Inmos Transputer processor chips, which have been specifically designed to be used in multiprocessor systems.

Since an entire supercomputer CPU can now be located on just one circuit board, the physical cabinets of the newest supercomputers are correspondingly much smaller, compared to those of the early 1980s. One example of these compact machines is the CRAY 2, its main cabinet being less than a cubic metre in volume. It is possible to envisage the 'desk-top supercomputer' in the not too distant future. However, the greater heat dissipation per unit volume produced by this reduction in size has presented engineers with major design problems concerned with the development of increasingly sophisticated cooling systems which use cooling agents such as liquid nitrogen. An efficient cooling system is essential because the processors must be maintained at a low temperature to give the best performance.

Within the Meteorological Office, the supercomputer is a tool used not only to run the operational numerical weather prediction models but also to aid the research into producing even more accurate forecasts. The limited amounts of computational power and main memory available at each stage have constrained the accuracy to which weather systems and physical processes could be represented. Indications are that significant improvements in the forecasts would be gained by doubling the resolutions of the global, limited-area and mesoscale models. This would mean that the global 'coarse-mesh' model would have a resolution equivalent to that of the limited-area 'fine-mesh' model at present, with twice the number of grid points to be processed in both the latitude and longitude directions. If this was done, the model time step would also have to be halved to maintain computational stability, thus doubling again the number of necessary calculations. Hence, altogether eight times the computational power used at present would be needed to run this model within the same deadlines. Even more power and memory would be required if more vertical levels were included and if there were more runs per day.

Further planned development of the mesoscale model, which covers the British Isles, will lead to an increase in resolution, more levels in the vertical, larger areal coverage, more runs per day and runs for a longer period ahead. It has been estimated that this demand would require 50 times the power used for the present model and over 30 times the memory.

The computer is also used to support the research and development of enhancements to the operational forecast models, such as improved physical parametrizations, as well as for thoroughly testing them before implementation. Again in the research area, the climate model, which is very important for investigating long-term changes, will require increases in computer speed and memory of two orders of magnitude within the next decade. Long-range forecasting, which goes up to 30 days ahead, and the development of coupled ocean–atmosphere prediction models will both also make large demands on future supercomputers.

It has been estimated that the increase in processing power necessary for all these applications in the early 1990s would require a computer with at least 50 times the power of the CYBER 205 and about 500 million words of memory. This performance target will not be met by a single processor so, at least for the major models, it would be necessary to have a multiprocessor with probably 8–16 processors. There are several parts of the numerical forecast suite that can be made independent of each other. For example, the mesoscale and global forecasts could be run simultaneously with each integration being shared by several processors. There appears to be no intrinsic difficulty in adapting both the atmospheric and ocean models to be used on a multiprocessor, as the experience of ECMWF has demonstrated.

It would have to be possible to run the models on a variable number of processors to cater for periods when processors are not working or new processors are added to the system, and so the execution of the subprocesses needs to be independent of particular processors. Software tools are available to help users of multiprocessor systems, but even so it is likely that the Meteorological Office would still need to develop some software itself to implement a multiprocessor system fully.

During 1987 the Meteorological Office obtained approval to replace the CYBER 205 and, as a result of a competitive procurement exercise, a contract was awarded to CDC in December 1987 for supply of an ETA10 system, with June 1988 as the target date for its acceptance. The configuration is based on four processors, each of which is about twice as powerful as the current CYBER 205, and has twice the local memory. These processors also have access to 64 million words of shared memory and the overall system is expected to have a peak performance of about 3000 MFLOPS. This represents an important step towards meeting the computational requirements of the 1990s and there will also be the option of upgrading the system in the future.

In conclusion, it is clear that the accuracy of forecasts, achieved using predictive models, has progressed in stages with the available state-of-the-art computing power and this close liaison between numerical weather prediction and supercomputers is likely to be important for the foreseeable future.

Acknowledgement

I would like to thank Dr A. Dickinson for providing papers and guidance at the outset of the project, and for his continued help during the preparation of this paper.

References and bibliography

Bell, R.S. and Dickinson, A. 1987 The Meteorological Office operational numerical weather prediction system. *Sci Pap Meteorol Off*, No. 41.

Dent, D. 1985 The multitasking spectral model at ECMWF. *ECMWF Newsl*, **29**, 21–26.

1986 The multitasking model — Part 2. *ECMWF Newsl*, **33**, 9–12.

1987 The multitasking model — Part 3. *ECMWF Newsl*, **37**, 5–9.

Dickinson, A. 1982 Optimizing numerical weather forecast models for the CRAY-1 and CYBER 205 computers. *Comput Phys Commun*, **26**, 459–468.

Gibson, R. 1984 Using the CRAY X-MP. *ECMWF Newsl*, **26**, 10–12.

Gibson, J.K. 1985 A production multi-tasking numerical weather prediction model. *Comput Phys Commun*, **37**, 317–327.

Hockney, R.W. and Jesshope, C.R. 1983 Parallel computers. Bristol, Adam Hilger.

Kascic, M.J. 1979 Vector processing on the CYBER 200. Infotech state of the art report 'Supercomputers'. Maidenhead, Infotech International.

Simons, G.L. 1983 Towards fifth-generation computers. Manchester, The National Computing Centre.

Slingo, A. 1985 Handbook of the Meteorological Office 11-layer atmospheric general circulation model. (Unpublished, copy available in the National Meteorological Library, Bracknell.)

Storer, N. 1985 The ECMWF CRAY X-MP/48. *ECMWF Newsl*, **30**, 15–18.

White, P.W. 1985 Vectorization of weather and climate models for the CYBER 205. *In* R.W. Numrich (ed.); Supercomputer Applications. New York, London, Plenum Press.

551.501.777:551.501.81

Progress in the development of PARAGON

B.R. May
Meteorological Office, Bracknell

Summary
A description is given of PARAGON — a system developed in the Meteorological Office for the routine archiving and adjustment by gauge observations of daily rainfall totals measured by radar.

1. Introduction

The Meteorological Office requires estimates of surface rainfall amount at ungauged locations for a range of periods from five minutes to one year. The data are required for enquiry bureau use, investigations and research.

For historical studies of daily or longer rainfall totals, observations from the climatological network of rain-gauges are available. The gauge spacing in this network ranges from 3 km to over 30 km depending upon locality, with an average of about 8 km. For more urgent requirements for daily rainfalls the only observations reported in near real-time are those from the synoptic network with a gauge spacing ranging from 20 km to 100 km and an average of 40 km.

Quantitative measurements of rainfall from radar observations are now available. Because of their regular and dense spatial coverage (5 km spacing) they can be used to interpolate between gauge observations especially from the more widely spaced gauges in the synoptic network. Their operational availability every 15 minutes increases their potential for a range of advisory uses especially when prompt or detailed information is required.

On the basis of considerable experience in the non-routine processing of radar rainfall observations for quantitative studies (Palmer *et al.* 1983), proposals were made for the design of a fully automated comprehensive system for the routine processing, adjustment by gauge observations and storing of radar data (Smith *et al.* 1984). The system was planned to produce daily rainfall totals, with the computer software being developed in two phases — phase 1, off-line (non real-time) data and phase 2, on-line (real-time) data. The development of phase 1 is finished, but it now seems unlikely that phase 2, as it was originally conceived, can be made operational before 1990 because of difficulties in transferring to the Meteorological Office central computer (COSMOS) in real time the large amounts of single-site radar data involved.

This paper describes progress so far on the development of this system which has been given the name PARAGON (Processing and Archiving of RAdar and Gauge data Off-line and in Near real-time). An outline block diagram of the PARAGON processing system as originally conceived is shown in Fig. 1.

2. Radar observations

Radar waves are emitted by a radar transmitter and reflected by raindrops or snow and ice particles. The strengths of these reflected signals are converted by a simple formula into rainfall rates (intensities) and averaged over contiguous 5 km × 5 km squares partly to reduce their random errors and also to reduce the data to manageable amounts.

The basic rainfall intensities are measured virtually instantaneously by each radar every 5 minutes, accurately synchronized, and given an automatic real-time on-site calibration (Collier *et al.* 1983) using,

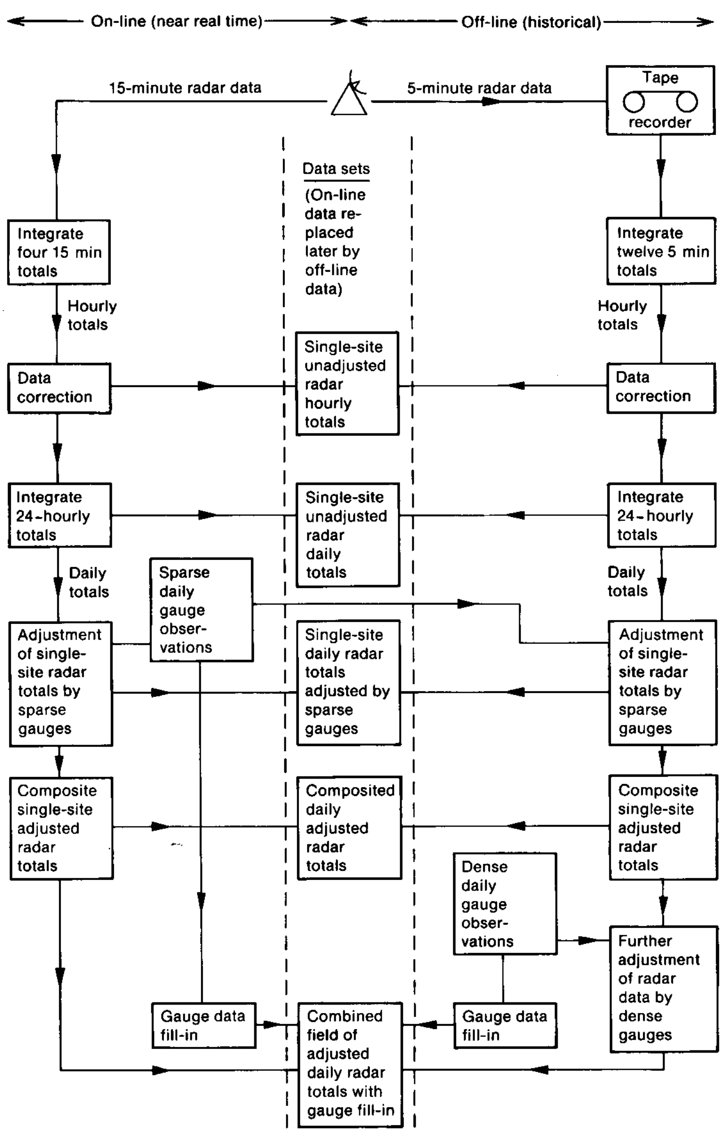

Figure 1. The PARAGON processing system.

as ground truth, observations from three–five dedicated gauges within 75 km of each radar. Areas of permanent clutter (reflections from hills, buildings, etc.) are erased on site and a correction is applied to compensate for loss of radar sensitivity at longer ranges. Currently the presence of bright band, the enhanced signal from melting snow, is detected on site but its effect on apparent rainfall rates is not yet removed operationally. The resulting rainfall intensities are regarded as the 'raw' radar data in the context of PARAGON.

Although the intensities or rainfall totals are areal averages they need to be ascribed to specific locations for the purposes of adjustment by gauge observations, or simply for plotting. These locations are chosen to be the centre of the squares and the spatial radar rainfall distributions are represented by a regular array of grid-point values with a 5 km spacing. This determines the way in which all radar data (and some associated gauge data) are processed, stored and displayed by PARAGON. From these grid-point values, estimation of radar rainfall at any location, for instance of a gauge, can easily be carried out by two-dimensional interpolation.

At present, radar observations are being made at the five installations shown in Fig. 2. The coverage of single-site data is a circle of radius 210 km contained in squares of 84 × 84 grid points. Data from two new radars, at Castor Bay in Northern Ireland and Ingham, near Lincoln, will be available within the next year.

From the beginning of 1984 the PARAGON data sets are as complete as they can be, though data from the radar at Chenies did not start until the beginning of 1985. The data for May 1981 to the end of 1983 have also been processed but are less complete than for later years.

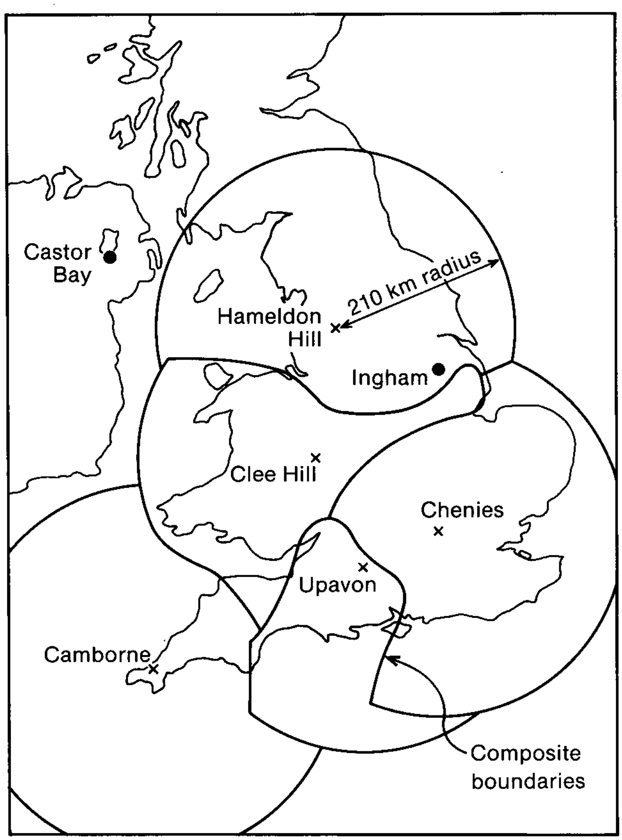

Figure 2. Radar coverages with existing operational radars. The sites of operational radars are indicated by X whereas those expected to become operational in the next years are denoted by ●.

3. PARAGON input

The basic observations required for input into PARAGON are hourly radar rainfall totals for each grid point as observed by each radar separately (single-site data), and obtained on-line or off-line as described in section 5. Much of the subsequent processing, consisting of quality control, integration to daily totals, adjustment by gauge observations and compositing, is common to on-line and off-line working and uses the same computer programs. The main difference in the two versions was planned to be the number of gauges used for adjustment of the radar data as ground truth. The on-line version would use only the (sparse) synoptic gauge observations while the off-line version would further adjust the radar data by dense gauge observations. The timetable for the processing would be within a few hours, after the end of each day, for on-line working but after a few weeks for off-line working. The on-line products, assumed to be less accurate, are required for urgent advisory use until the more accurate off-line products become available. At present only the off-line system is operational.

4. The PARAGON processing system

4.1 *Quality control and production of daily totals*

Quality control of hourly radar data is carried out within PARAGON to supplement that already applied on site. Areas of partial or complete occultation (again caused by hills, buildings, etc.) are deleted completely. Areas of 'anaprop', the enhanced clutter occurring during particular weather conditions, are corrected after detection by distinguishing between the usually stationary anaprop and moving rainfalls. The data are then held as hourly totals from each radar separately, unadjusted by gauge observations, other than those used for the on-site calibration.

The 24-hourly totals for 09 GMT to 09 GMT the next day are formed for the grid points and stored as a set of single-site unadjusted radar daily totals. This particular 24-hour interval is chosen to conform with the standard rainfall day for gauge observations, which are required for the following adjustment process.

4.2 *Adjustment of daily radar data by sparse gauge observations*

Previous experience indicates that on-site calibrated data are still not sufficiently accurate for certain quantitative uses of surface rainfall information such as determining monthly river-catchment rainfalls. As a consequence the radar data are modified by gauge observations using a form of objective adjustment which is distinct from the more meteorologically based on-site calibration. This adjustment is a type of quality control and so is consistent with the long-standing commitment of the Meteorological Office to improve the accuracy of archived weather data.

The objective adjustment involves the calculation of daily adjustment factor g/r (g and r are daily totals measured by gauge and radar respectively) at the location of the gauge observations, interpolating between these by surface fitting to the location of each 5 km grid point and applying the estimated adjustment factors to the grid-point radar values. This alters the radar data on a scale comparable with the spacing between the gauges, but leaves the shape of smaller features between gauges unaltered. Adjustment factors outside the range 0.1 to 10.0 and ones based on too small values of r and g are rejected. The adjustment is applied to the data from each radar separately.

For both off- and on-line processing the sparse daily gauge observations used for this adjustment are those from the synoptic gauge network received a few hours after 09 GMT each day.

The result of the adjustment process is a set of single-site daily radar totals adjusted by sparse gauge observations.

4.3 *Compositing*

In order to reduce data storage requirements, and to facilitate the use of the data generally, a composited version is now produced to give a single field of rainfall totals by eliminating overlaps between the coverages of adjacent radars. This is done using a previously specified set of boundaries, but with alternatives if the data from a particular radar are missing altogether, to maximize radar coverage. This product is a set of composited daily adjusted radar totals.

If compositing is carried out on unadjusted data (and the raw data from two adjacent radars can be very different in the overlap area) then sharp boundaries can be produced which cannot be removed by subsequent gauge adjustment because they are small-scale features. Adjustment before compositing suppresses these boundaries, in contrast with the compositing after on-site calibration (but without any further gauge adjustment) which takes place in the production of the Network composited radar picture sent to forecasters, and is an important feature of PARAGON.

4.4 *Further adjustment by dense gauge observations*

Off-line hourly radar totals are received after a delay of a few weeks by which time the daily gauge observations from the dense climatological network are available. These daily gauge observations arrive at the Office by post or magnetic tape and are then subjected to a lengthy quality control. This further adjustment of the composited off-line radar data by the dense gauges is carried out exactly as described previously for sparse gauges, the intention being that the increased amount of gauge data should improve the accuracy of the final product.

4.5 *Gauge-only data fill-in*

Areal averages for 5 km × 5 km squares, in areas not yet covered by radar data, or where the data have been deleted because of contamination by occultation or clutter, are estimated directly from sparse or dense gauge observations as appropriate to give completeness of data coverage. In this process the spatial distribution of average annual rainfall (AAR) is used in the adjustment, in effect, as a substitute for specific radar data for the day. Since gauge observations are always used with interpolation between them by radar data or AAR, a uniformity of accuracy of the two resulting types of fields interspersed over large areas can be obtained.

The product is stored in the final data set of combined fields of adjusted daily radar totals with gauge fill-in; gauge-only derived amounts are distinguished from gauge-plus-radar derived amounts by a negative sign.

5. Sources of hourly radar rainfalls and storage of data

Hourly radar rainfalls would be available to PARAGON from two sources for use in the two time-scales.

For off-line processing 5-minute intensity measurements are recorded on magnetic tape at each radar site independently for each grid point. These tapes are retrieved within 2–3 weeks of recording and processed into hourly totals by integrating 12 successive intensities. These totals should be the more accurate ones available from radars because this frequent sampling rate can account for more rapid changes in intensity.

For on-line processing every third set of single-site 5-minute data is also transmitted to Bracknell from each radar via direct telecommunication links to the RADARNET computer before compositing and transmission to forecasters as the Network composited pictures. Originally it was planned to transfer these 15-minute single-site observations to COSMOS in real time for processing to hourly totals at each grid point, by integration of four successive intensities. These hourly totals would be less accurate as radar data than the ones derived from the recorded data because of the lower sampling rate, but this is unlikely to be a serious source of error apart from occasions when the rainfall intensity is changing rapidly.

The general strategy of data storage is that on-line data would be automatically stored to be replaced later by the off-line data, which are expected to be of higher quality. There would be exceptions to this if for instance off-line data were partially or completely missing due to on-site recorder malfunction, in

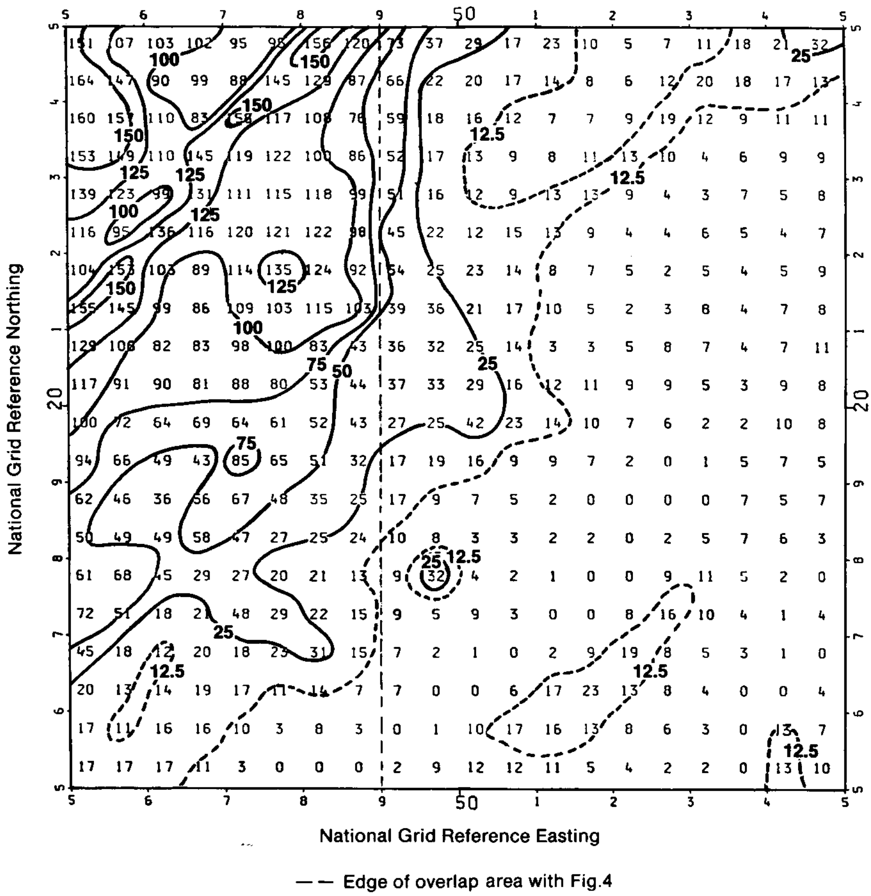

Figure 3. Unadjusted rainfall totals (units of 0.1 mm) measured by Chenies radar for the rainfall day commencing 09 GMT on 7 April 1987.

which case the on-line data would be retained for completeness. Until the on-line system is fully operational a procedure exists whereby the processing of off-line data for periods of outstanding interest can be accelerated, but this causes disruption of the routine processing of the radar-site magnetic tapes.

6. Output from PARAGON

6.1 *Data output*

A block diagram of the complete PARAGON system is shown in Fig. 1 with the data sets after each processing stage arranged in order down the centre. The contents of these data sets can be printed as fields of grid-point rainfall totals on fiche, page print-out or film, or produced in contoured forms. Examples of grid-point fields of radar observations unadjusted, and adjusted by sparse gauge observations with gauge fill-in (all 5 km square areas average), are shown in Figs 3 and 4. Two overlapping 100 km squares are shown such that the left-hand portion of the square in Fig. 3 coincides with the right-hand portion of the square in Fig. 4, the overlapping area being bounded by dashed lines. Fig. 3 shows the unadjusted radar data from the Chenies radar which is located at the centre of the

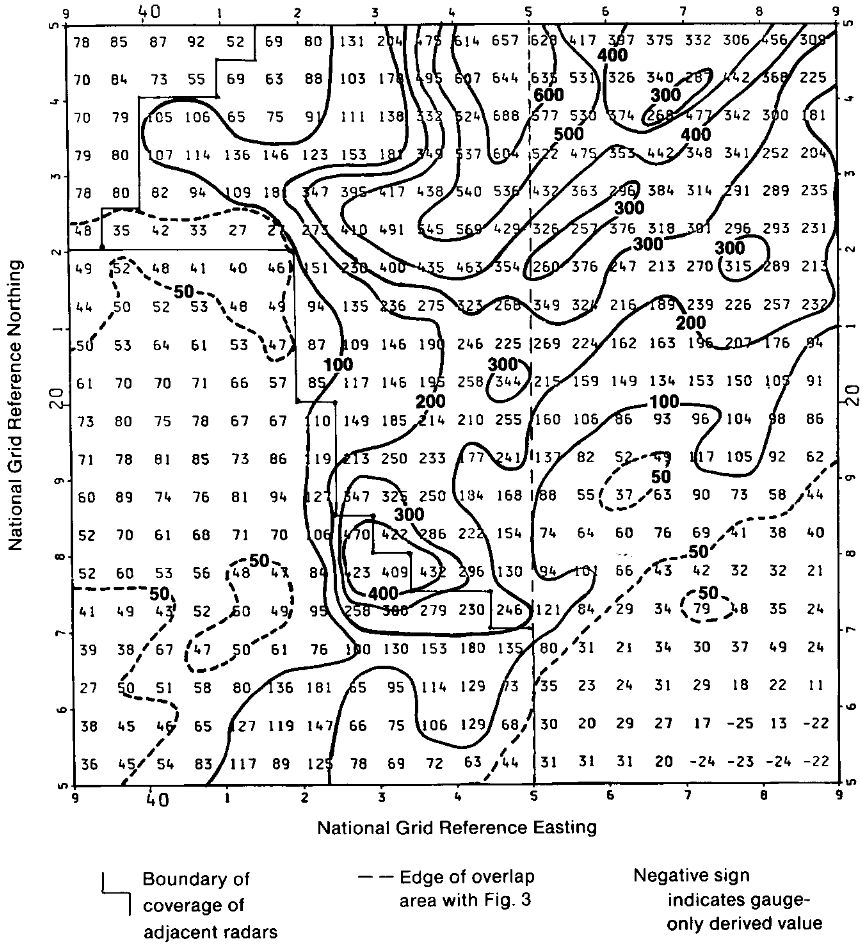

Figure 4. Radar rainfall totals adjusted by observations from sparse gauges (units of 0.1 mm) for the rainfall day commencing 09 GMT on 7 April 1987.

square. Fig. 4 shows the radar data adjusted by the sparse gauges and composited, according to the indicated boundaries, with the Chenies data to the right, Upavon data to the lower left and Clee Hill data to the upper left of the square.

In the overlap area (all Chenies data) it can be seen that the effect of adjustment is to leave the smaller-scale detail of the rainfall field unchanged while altering the absolute values on the larger scale, in this case increasing the rainfalls by a factor of about 2 to the south but nearer 3.5 to the north. The peak of rainfall to the south-west of the centre of the square in Fig. 4 shows no influence on the isohyets of the composite boundary between the Chenies and Upavon adjusted radar data. A small area in the south-east corner of the square in Fig. 4 shows some gauge-only values which are consistent with neighbouring radar-derived values.

6.2 *Diagnostic output*

Information regarding the data, such as availability, quality, changes to radar-site hardware and software and its off-line or on-line status is contained in MINILOG (Banks 1985, Crummay 1984).

Quarterly Radar Data Assessment Reports are produced giving up-to-date summaries of this information (Banks and Crummay 1985). A full technical description of data sets, processes and computer programs comprising the system is contained in the *PARAGON users' handbook.*

7. Future developments

7.1 *Selective use of radar data*

At present, the adjustment procedure using sparse gauge observations is applied to all radar values within the area of coverage of valid adjustment factors. No consideration is made as to whether or not a better estimate of the surface rainfall at a grid point could be obtained by direct interpolation between the gauge observations without the use of radar data at all, as described in section 4.5. The choice of which is the better of the two sparse gauge-derived estimates is a matter of assessment of probability which depends on the quality of radar data and the nature of the rainfall fields in the vicinity (May 1986).

Algorithms to determine the choice at each grid point and their success are being investigated in the Advisory Services Branch of the Meteorological Office.

Investigations have been made of the differences between fields derived from dense gauges only and from radar observations adjusted by dense gauges. First indications are that the small inter-gauge spacing eliminates the contribution of radar data to a great extent and so it is possible that in future the routine adjustment of radar data by dense gauges will be omitted from PARAGON.

7.2 *Use of Network composite data for on-line operation*

It now appears unlikely that a system to transfer single-site radar data in near real-time to COSMOS can be implemented before at least 1990 because of data volume restrictions. Therefore on-line PARAGON processing, as originally conceived, cannot proceed before then.

The requirement for improved near real-time daily and sub-daily rainfall information is becoming more urgent, for use in MORECS (the Meteorological Office rainfall and evaporation calculation system) and in general advisory services. Consequently it is now planned, as an interim measure, to use the 15-minute Network composited radar data already available in COSMOS as the basic data for PARAGON. The benefits of the 'adjustment before compositing' feature of PARAGON would have to be sacrificed but adjustment by gauges will be retained as a means of quality controlling the radar data.

References

Banks, C.J. 1985 MINILOG — a catalogue of information on the radar data archive. (Unpublished, copy available in the Meteorological Office Advisory Services Branch.)

Banks, C.J. and Crummay, G.P. 1985 The production of radar hardware and software reports. (Unpublished, copy available in the Meteorological Office Advisory Services Branch.)

Collier, C.G., Larke, P.R. and May, B.R. 1983 A weather radar correction procedure for real-time estimation of surface rainfall. *Q J R Meteorol Soc*, **109**, 589–608.

Crummay, G.P. 1984 A user's guide to browsing PARAGON radar archive using ISPF panels. (Unpublished, copy available in the Meteorological Office Advisory Services Branch.)

May, B.R. 1986 Discrimination in the use of radar data adjusted by sparse gauge observations for determining surface rainfall. *Meteorol Mag*, **115**, 101–115.

Palmer, S.G., Nicholass, C.A., Lee, M.J. and Bader, M.J. 1983 The use of rainfall data from radar for hydrometeorological services. *Meteorol Mag*, **112**, 333–346.

Smith, C.J., Spalding, T.R. and Palmer, S.G. 1984 The acquisition and use of radar data by Met O 8 in real-time and offline. (Unpublished, copy available in the Meteorological Office Advisory Services Branch.)

551.593.653 (412.1)

Observations of noctilucent clouds from Ben Nevis Observatory

D. McConnell
Meteorological Office, RAF Scampton

Summary

Noctilucent clouds were regularly seen from Ben Nevis Observatory during part of its lifetime. Notes about the clouds were entered in the logbooks but have remained unpublicized, apart from having appeared within the log entries reprinted in the *Transactions of the Royal Society of Edinburgh*. This article brings to light the occasions when the clouds were observed from the Ben, together with descriptions, largely as the observers wrote them, of the more significant sightings.

1. Introduction

Noctilucent clouds (NLC) are tenuous, brilliant, silvery clouds visible against the twilight sky because they are being illuminated by sunlight when the sun is below the observer's horizon. At a mean altitude of 82 km in the coldest part of the atmosphere (called the mesopause) where temperatures can reach −140 °C, these luminous 'night' clouds, at ten times the average height of cirrus, are the earth's highest clouds. They are believed to consist of meteor particles, or possibly ions, around which are deposited minute crystals of ice, accounting for the high reflectivity. Although NLC are regularly seen by those who enthusiastically scan the summer twilight skies, one surprising feature of the clouds' history is that they were virtually unknown before 1885. It was only in that year that they were widely recognized when

The Ben Nevis Observatory in summer.

there were extensive and brilliant displays visible from the British Isles, other parts of Europe, and Russia. But why astronomers and meteorologists did not record them with certainty beforehand is still a mystery.

One place from where NLC were 'discovered' in 1885 was the Ben Nevis Observatory. Being at a latitude of 56.8° N (almost midway within the best range of 50–65° N) it was well situated for observing them. They were recorded in subsequent years by the skilled observers of the Ben's essentially meteorological establishment which was in existence from 1883 to 1904. The observatory's altitude, 1343 m (4406 ft) at the mountain's summit, provided an excellent vantage for observation. These sightings, including any considered to be possibilities, are listed in Table I. However, in order to place the first Ben Nevis observation of NLC into perspective, it is necessary to outline when the clouds were first seen elsewhere in 1885.

Table I. *Dates of noctilucent cloud observations with notes, including possible sightings, from Ben Nevis Observatory 1884–96*

Date — night of			Notes	Date — night of			Notes
1884	July	1/2	Possible only	1888	June	30/1	Very bright
1885	June	6/7	Possible only		July	7/8	
	July	1/2	First 'pearly Ci'	1889	May	26/27	
		29/30	Possible only		June	6/7	
1886	June	2/3	Almost certainly			7/8	
		17/18	Almost certainly			9/10	
1887	June	13/14	Confirmed as NLC			10/11	
		24/25				17/18	
		25/26				21/22	To 70° altitude
		28/29	To 30° altitude		July	2/3	
		29/30	To 20° altitude			3/4	
	July	5/6				4/5	
		17/18	Not following sun	1890	May	26/27	Notes taken
		20/21			June	28/29	
1888	June	4/5			July	15/16	
		5/6				16/17	
		15/16	To 80° altitude	1891	July	9/10	Notes taken
		17/18	To 85° altitude			13/14	Notes taken
		19/20				14/15	
		21/22		1892	July	13/14	NLC and aurora
		22/23				23/24	
		23/24				24/25	
		24/25		1895	July	16/17	
		25/26	Very bright	1896	May	30/31	NLC or aurora
		26/27					

Total: 44 positive, 5 possible

The first officially recorded sighting of NLC is credited to T.W. Backhouse, a keen skywatcher from Sunderland, who saw the clouds from Kissengen in southern Germany on 8 and 10 June, and again from there on 7 July (Backhouse 1885). They were also seen on 10 June from Prague, and on 12 June the Russian astronomer V.K. Tseraskii observed them from Moscow (Bronshten and Grishin 1976). On 23 June they were noticed by various other observers, including the German astronomer O. Jesse who saw them from near Berlin (Jesse 1890). Jesse, like Tseraskii, was a pioneer in measuring the heights of

NLC. Also, on several unspecified dates during June, the clouds were seen from Dublin, and on 6 July from Southampton and London. There were numerous other occasions throughout June and July 1885 when they were observed from Europe, particularly Germany (Vestine 1934).

2. The early sightings from Ben Nevis

NLC may have been discerned from Ben Nevis in 1884 and also in 1885, earlier than Backhouse. The possible occasion in 1884 was on the night of 1/2 July when the logbook contained this entry:

> At midnight long streamers radiating from the zenith were observed. They seemed to be composed of filmy cirrus, but looked almost like an auroral arch. They were about 45° long, and reached from NE to W.

And then on 7 June 1885:

> At 1 h and 2 h the sky to NNE above where the sun was, was covered with very thin cirrus clouds, coloured pearly grey and pale green. Dull orange at horizon.

Was either of these two instances NLC?

The most common general description in English referred to luminous cirrus clouds; the name noctilucent clouds had not at that time come into existence. The Ben Nevis observers had their own unique name for the clouds, and they very aptly called them pearly cirrus. Thus, on the only definite occasion of their observation from the Ben in 1885, on the night of 1/2 July, the logbook records:

> Beautiful pearly coloured cirrus to northward at 23 h and midnight.

No further mention was made of the clouds that night, even though the sky remained clear, but did the observer(s) wonder about them?

There was another possible sighting on 30 July 1885:

> At 1 h the whole sky was dull blue above the thick haze that covered the horizon, except to the N and NNE where a distinct path of light, coloured red and pearly grey, and bounded by a sharp well-marked line, was observed. This light was probably due to the sun, and was not auroral.

In 1886 there were only two reports but these were almost certainly of NLC. On the night of 2/3 June:

> At midnight the higher cirrus clouds to northward were bright, apparently with sunlight; but the lower ones were quite dark;

and on 17 June:

> At 23 h and at midnight the cirri clouds to northward shone with a bright silvery light (not auroral); the brightest parts were about 15° above the horizon.

During 1886 the clouds were well observed elsewhere in the British Isles, from such widely separated places as the Isle of Wight (as early as 28 May), Sunderland, Belfast, Dublin, Bideford, Southampton, Cumberland and Edinburgh, with the last reported sighting that year from Sunderland again (as late as 11 August).

3. Sightings in 1887–89

The next three years continued to produce many reports of the clouds, both from the British Isles generally — even as far south as the island of Sark (49.5° N) — and from Ben Nevis where the weather observers saw them on 8 occasions in 1887 and on 13 in 1888.

In 1887 they were first seen from the Ben on 13 June at 23 h when filmy cirrus, which was confirmed as being NLC, was recorded to the north. Thereafter, during June and July, the words 'pearly-white cirrus' becoming 'pearly cirrus' began to be commonly used; the observers no doubt by this time realizing that these 'cirrus' clouds were something different from the familiar variety. On the night of 28/29 June 1887, the angular elevation of the clouds was noted:

> Pearly-white cirrus reaching to about 30° altitude at midnight to northward.

The following night, they were seen to about 20° altitude. An interesting remark appeared with regard to the display on the night of 17/18 July:

> At night bright pearly cirrus to northward, height about 12°, highest to NNW all night, not following the sun round.

NLC do not always 'follow' the sun round; their position in the sky, as seen by the observer, depends on their relative extent and actual movement, and the sun's depression angle, as well as its azimuth.

On two nights in 1888, the clouds were observed very high in the sky. On 15 June:

> Pearly cirrus seen to N at night. At 23 h it reached to almost overhead, not more than 10° from zenith, and even at midnight was about 45° above horizon.

And two nights later, it was slightly higher:

> At 23 h some of it was not more than 5° from zenith.

During 1888 the clouds were observed almost exclusively in June, with only two sightings in July. Having been seen both on the night and morning of 30 June/1 July, they were thus recorded on twelve occasions in June, from which there were sightings on the six successive nights from 21/22 to 26/27 June; quite a sequence. The displays of 26/27 June and 30 June/1 July were described as being 'very bright'.

Their appearances in 1887 and 1888 were summed up by the observatory's superintendent, Robert Traill Omond, in a letter dated 2 July 1888 to the British scientific journal *Nature* which, since 1885, had carried numerous reports of the clouds. This letter was published in the issue of 5 July 1888 under the heading 'Sky-coloured Clouds at Night':

> In *Nature*, June 28 (p. 196), Mr Backhouse notes the appearance of illuminated clouds to northward at night. Similar clouds are seen from here on almost every clear night near the summer solstice. For the last two years special note has been taken of them. In 1887 they were first seen at midnight on June 13, and last seen on July 20; this year their first appearance at midnight was on June 4, and they are still visible every clear night. The clouds are not, as far as I have observed, coloured, but shine with a pearly or silvery lustre. I have seen them at midnight as high as 30° altitude, but they are generally confined to the first 10° or so above the northern horizon. The facts that they vary greatly from night to night in appearance, being sometimes almost absent, and that one or two photographs that have been taken of them show them simply as ordinary cirrus clouds, all seem to indicate that they really are very high cirrus lighted by the sun.

Ironically, after this letter, the clouds were seen only once more in 1888, on 7/8 July. Omond, a careful observer, had noted their 'non-colour' and he was not misled, as was easily done, by the influence of the twilight sky background which would give the impression that the clouds were sky-coloured. Even today, noctilucent clouds are usually described as a blue-silvery colour, and occasionally as yellowish when very near the horizon due to more haze lower down, but it is nevertheless the case that the cloud particles reflect blue wavelengths more strongly than any others.

The logbook in 1888 also contained two 'negative' reports of the clouds, confirming that the observers did keep a look-out for them. Thus, on the night of 16/17 June:

> N horizon dusky red at midnight, but no pearly cirrus all night;

and on the night of 11/12 July:

> At midnight N horizon was faintly red, but no pearly cirrus was seen.

In 1889 the clouds were observed on ten occasions, beginning on the night of 26 May at 23 h which had become the earliest date of the year that they were recorded from the Ben. On the night of 21/22 June, they were again seen high in the sky:

> Pearly cirrus was seen to 70° altitude at 23 h, and a small patch to the NNW at midnight.

Six occasions were recorded in June and three in July.

4. Sightings in 1890–92

For the sightings of the clouds during 1890 and 1891 — numbering only four and three respectively — the observers supplied, in contrast to all previous years, detailed notes relating to the clouds' position, movement and structure.

In 1890 the clouds were also recorded as early as 26 May, when before midnight two photographs of them were taken. This was the first instance that the logbook recorded any photographs having been attempted, although Omond, in his letter to *Nature* in 1888, had mentioned that 'one or two' had been taken. This display was described as having a 'ribbed' and 'wavy' appearance through which stars could be seen. The movement of the cloud mass was slowly from the east. On 27 May, another photograph was taken, but only ten minutes later:

> 3 h, no pearly cirrus or any kind of cloud in sky, only faint bands and streaks of haze near horizon, and pink 'foreglow' streamers.

NLC are nowadays classed into four main types — veils, bands, waves and whirls — and from the description of the display on the night of 26/27 May 1890, the Ben Nevis observer would have seen the wave or band type or a combination of both. NLC nearly always have a motion generally from east to west across the sky, which is what the observer had noted.

After that display, the clouds were not observed for over a month, until 29 June, and thereafter they were noticed only twice more, in mid-July.

The clouds were observed for the first time in 1891 as late as the night of 9/10 July, and exact angular measurements were made:

> At 23 h 5 m the highest point of the pearly cirrus was 5° 18′ and the lowest 3° 14′ above the horizon.

And on the night of 13/14 July:

> Pearly cirrus seen to N at midnight, stretching from NW to NE, its highest point being then 14° 16′ above horizon, and its lowest limit 7° 8′ above same.

These measurements were made with an instrument called a stephanome which the observers frequently used for obtaining the angular radii of various sections of coronae, glories, haloes, etc. More photographs of the clouds were also taken.

No increase in the frequency of the cloud sightings occurred in 1892. The first that year, however, on 13 July, contributed to a beautiful and unusual conjunction because an aurora was observed at the same time:

> At 23 h and midnight pearly cirrus was seen about 10° above NE horizon. At 23 h 30 m while watching the pearly cirrus, streaks of aurora were seen, and the cirrus seemed behind these. Spectrum bands seen in N horizon at night. Photograph taken of these and of the pearly cirrus.

The reference to the pearly cirrus being behind the aurora may imply that the aurora was nearer and thus lower in altitude, but this was only an optical illusion in the same way that altocumulus sometimes appears higher than cirrus in the same part of the sky. Auroras occur at a minimum altitude of around 90–100 km. The reference to 'spectrum bands' was not to the actual spectrum of the aurora but to a description of the form of the aurora, i.e. that it was of the 'rayed-band' variety, like curtains. However, this Ben Nevis observation was not the first of aurora and NLC together. On 27 July 1886, Charles Piazzi Smyth, the Astronomer Royal for Scotland, had recorded the two phenomena at the same time from Edinburgh, although he was confused at first — unlike the observer at Ben Nevis who knew immediately of the simultaneous occurrence in 1892.

5. Height determinations elsewhere

While the Ben Nevis observers knew that their pearly cirrus clouds were at a very high altitude, they were not able to determine what that height was. In any case, the observers would not, until some years after, if at all, have been aware that such measurements had already been made as early as the clouds' year of discovery, 1885, and even before they had first been seen with certainty from the Ben. It was in late June 1885 that Vitol'd Karlovich Tseraskii in Russia had successfully obtained a range of values from simultaneous observations, resulting in an average height of 79 km (Bronshten and Grishin 1976). Two years later, Otto Jesse in Germany, using baseline photography, gave a height of 75 km, but in 1889 he was able to refine his accuracy to values of 81 and 82 km. The Ben Nevis observers may have read Jesse's account in *Nature* (Jesse 1890), but they would not have been aware of Tseraskii's work due to its lack of publicity because it was in Russian.

Photographs of noctilucent cloud taken from Ben Nevis Observatory.

6. Observation trends

It is strange that, after 1892, the clouds were reported with certainty from the Ben only once more — in 1895 — until the observatory's closure in 1904. But there was a possible occurrence on 30 May 1896:

> Some brightness in northern sky at 23 h, was not distinct enough to make it certain whether it was due to aurora or pearly cirrus.

However, this dearth generally corresponds with the low frequency of sightings of NLC from elsewhere after 1895, and the reason can only be surmised in connection with the trend as a whole. 'The period 1885–94 was particularly rich in displays of noctilucent clouds which were brightest in 1885–86 and diminished in brilliance and frequency to 1894' wrote E.H. Vestine of Canada in his historical survey of 1934. Although the greatest frequency of the clouds as seen from Ben Nevis was not from 1885 to 1886 but from 1887 to 1889 inclusive, the number of sightings from an individual station in relation to others depends on the extremely variable factor of the frequency of tropospheric clouds obscuring the noctilucent variety.

Nevertheless, with the distribution of the Ben Nevis sightings fitting the overall trend, one possible reason for the high frequency from 1885 until 1894, and the low frequency thereafter, is the suspected correlation — an inverse one — between NLC and sunspot maxima. After the 1883–84 sunspot maximum, the clouds increased in frequency until after the sunspot minimum of 1889–90, but by the succeeding maximum of 1894–95, there were few reports. However, although similar inverse occurrences have been noted from then until the present time, there is still no agreement that such a correlation can be termed conclusive. Researchers believe that the heating effect of an aurora, whose presence is proportional to solar activity, may be sufficient to either prevent the formation or cause the dispersal of NLC so that the two phenomena are rarely seen together for long. It was, appropriately enough, between a departing solar minimum and an approaching solar maximum, in 1892, that the instance of coexistent NLC and aurora was recorded from Ben Nevis.

One final point of interest was a logbook entry on 29 December 1898:

> Clouds like 'pearly cirrus' seen in E sky at 7 h.

From the time of year, this was probably an observation of nacreous or mother-of-pearl clouds, which form in the stratosphere at an average height of around 30 km. They occur far less frequently than NLC.

7. Photographs

The photographs of NLC reproduced here are two of several kept in the archives of the Meteorological Office in Edinburgh. However, these photographs are generally not in good condition; there is some fading and some lack of clarity. Unfortunately, the photograph of NLC with aurora has not been found.

References

Backhouse, T.W.	1885	The luminous cirrus clouds of June and July. *Symons's Mon Meteorol Mag*, **20**, 133.
Bronshten, V.A. and Grishin, N.I.	1976	Noctilucent clouds. Jerusalem, Keter (translated from the Russian 1970 publication).
Jesse, O.	1890	Luminous clouds. *Nature*, **43**, 59–61.
Vestine, E.H.	1934	Noctilucent clouds. University of Toronto Press.

Conference report

Summer School on the Diagnosis of NWP Products, Meteorological Office College, Shinfield Park, England, 6–10 July 1987

This was the second such summer school to be held at the Meteorological Office College, following the successful one on mesoscale meteorology held in 1985. The principle behind the summer school is to bring together forecasters and active research scientists, from both inside and outside the Meteorological Office, in an environment where they may share and exchange ideas, thereby enriching each other's understanding of meteorology. Over lectures, in classrooms and in the bar or out on the lawns, there was ample opportunity for meteorologists from different disciplines to get together.

Over 60 people attended the summer school which lived up to its name, as may be discerned from the group photograph. The universities were well represented, particularly at post-graduate student level, and there were four participants from European organizations. Most of the Meteorological Office representation came from the research branches. There were about ten practising forecasters and a handful of others with some forecasting experience. It was disappointing to see so few outstation forecasters; those unable to attend have missed a valuable opportunity to widen their knowledge of numerical weather prediction (NWP) products.

Participants in the Summer School on the Diagnosis of NWP Products.

As in the previous summer school, lectures were held in the mornings followed by case-study sessions in the afternoons and some evenings. The speakers were drawn from the European Centre for Medium-range Weather Forecasts (ECMWF), the Meteorological Office and the University of Reading. The scene was set by Dave Burridge (ECMWF) who reviewed the progress made in NWP in recent years and the expected developments in the near future. Alan Gadd (Meteorological Office) compared the design and formulation of the global forecast models used by ECMWF and the Meteorological Office, and Stuart Bell (Meteorological Office) compared the analysis schemes at the two centres. The foundations had thus been laid for the case-study sessions in which the working groups of seven or eight people were involved in analysing and understanding the performance of the forecast models.

The first case study was one dealing with explosive cyclogenesis where the aim was to identify the main processes taking place and the ability of the models to predict such events. It also afforded the opportunity to look at some other model output diagnostics apart from those used routinely; in particular, potential vorticity and Q-vectors as advocated by Brian Hoskins (University of Reading) in his lectures on the theory of rapidly developing systems. This case also showed the differences in performance of the analysis schemes so there was plenty of material for the working groups to analyse. The results of using the newer diagnostics were rather inconclusive. As with most methods of diagnosis, practice and familiarity are essential for effective application so that the interpretation of potential vorticity and Q-vectors was handled better by the research workers. However, the forecasters were just as effective in their use of the more familiar NWP products. This case showed the importance of higher resolution in the models for the successful prediction of explosive cyclogenesis, a factor borne out by the Meteorological Office's experience with the fine-mesh model.

The second case study was concerned with the variability of forecasts and compared the performance of the global models from the Meteorological Office and ECMWF over an 8-day period. The background to forecast variability was provided by Tony Hollingsworth (ECMWF) in his lecture on error growth and propagation. He also demonstrated how analysis differences between forecast centres could sometimes explain a major part of forecast differences. Charts of differences between models for various fields, e.g. 500 mb height difference, can be used to trace the differences back through a forecast to find the area of origin. However, the differences at analysis time are often very small and the largest analysis differences do not necessarily become the largest forecast differences. This makes the use of difference charts to estimate forecast confidence in real time rather difficult and inexact, as the working groups were to find in the case study. The case study was useful in demonstrating that the later forecast is not necessarily the most accurate. The problem of predicting the likely skill of a particular forecast was examined by Tim Palmer (ECMWF) who demonstrated three ways in which this might be achieved. The first is to identify certain regimes where forecast skill has been found to be highest, e.g. stronger-than-average Rockies ridge. Alternative ways of predicting forecast skill involve running an ensemble of forecasts from slightly perturbed initial states or using successive forecasts to estimate variability. Elements of each of these methods can be found in current medium-range forecast practice in a subjective way, but a more rigorous formalization using a combination of three techniques is eagerly awaited. Lennart Bengtsson (ECMWF) charted the improvements that have been made in medium-range forecasts over the last ten years and suggested that the gains achieved through resolution increases are becoming less marked. This also demonstrates the potential benefit in predicting the skill of a forecast.

Forecasts of a different variety were attempted by Andrew Lorenc and Mike Cullen (both from the Meteorological Office) in their assessment of prospects in data assimilation and NWP over the next ten years. These talks generated a lively exchange about the merits of automation and manual intervention. Another area of concern is that of data and in particular the likely demise of weather ships. Peter Ryder (Meteorological Office) gave a vivid account of the kinds of decision regarding cost and impact of observing systems that are having to be made now.

Reading through the questionaires filled in by many of the participants, it appears that most found the lectures interesting and the week stimulating and enjoyable. There were some reservations concerning the amount of material in the case studies, but even so most found the exercise worthwhile and we all appreciated the amount of effort that had gone into setting these up, mostly by Will Hand at the Meteorological Office College. The summer school is certainly a forum worth continuing. Perhaps forecasters on roster duties should pencil in July 1988 in their diaries and try to get along.

T. Davies

Satellite photograph — 9 January 1988 at 0911 GMT

This infra-red picture taken from NOAA-10, displayed on the Meteorological Office's HERMES system, has been archived as part of the Anglo-French Mesoscale Frontal Dynamics Project or FRONTS 87 experiment*. The cloud spiral near 60° N, 15° W shows an intense low that had deepened explosively during the preceding 24 hours.

At the time of the picture, the Hercules C130 aircraft of the Meteorological Research Flight was completing the second of four runs across the front (near latitude 48° N) during which dropsondes were released at intervals of 30–50 km. The measurements showed the front to be a clearly defined feature only above 800 mb. At the surface, although radar observations indicated intermittent line convection, the wind veer was only about 20° and the temperature fell by only 2 °C.

The front was the seventh of eight that were observed as part of the FRONTS 87 experiment. During the Intensive Observational Periods, considerable amounts of data were archived including measurements from three aircraft, many routine and special radiosonde soundings over the United Kingdom and France, and several high-power Doppler radars.

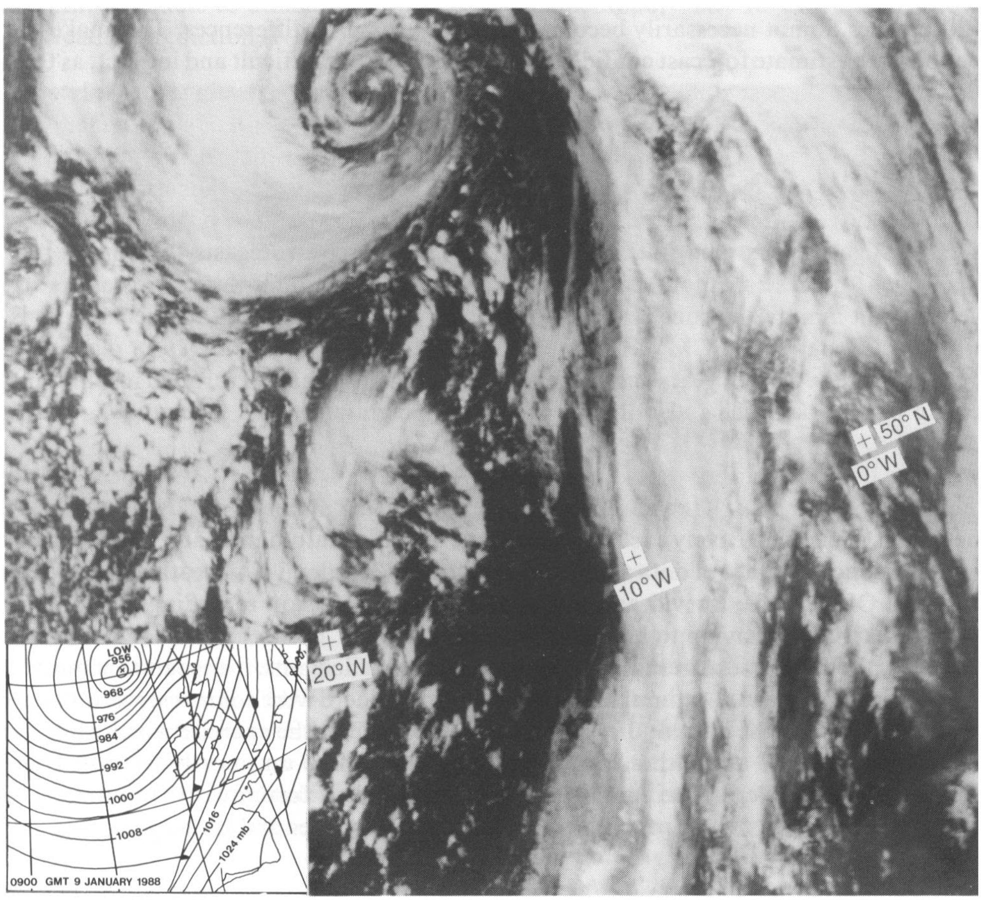

* Clough, S.A.; The mesoscale frontal dynamics project, *Meteorol Mag*, **116**, 1987, 32–42.

Meteorological Magazine

GUIDE TO AUTHORS

Content

Articles on all aspects of meteorology are welcomed, particularly those which describe the results of research in applied meteorology or the development of practical forecasting techniques.

Preparation and submission of articles

Articles for publication and all other communications for the Editor should be addressed to the Director-General, Meteorological Office, London Road, Bracknell, Berkshire RG12 2SZ and marked 'For *Meteorological Magazine*'.

Articles, which must be in English, should be typed, double-spaced with wide margins, on one side only of A4-size paper. Tables, references and figure captions should be typed separately.

Spelling should conform to the preferred spelling in the *Concise Oxford Dictionary*.

References should be made using the Harvard system (author, date) and full details should be given at the end of the text. If a document referred to is unpublished, details must be given of the library where it may be seen. Documents which are not available to enquirers must not be referred to.

Tables should be numbered using roman numerals and provided with headings. We consider vertical and horizontal rules to be unnecessary in a well-designed table; spaces should be used instead.

Mathematical notation should be written with extreme care. Particular care should be taken to differentiate between Greek letters and Roman letters for which they could be mistaken. Double subscripts and superscripts should be avoided, as they are difficult to typeset and difficult to read. Keep notation as simple as possible; this makes typesetting quicker and therefore cheaper, and reduces the possibility of error. Further guidance is given in BS1991: Part 1: 1976 and *Quantities, Units and Symbols* published by the Royal Society.

Illustrations

Diagrams must be supplied either drawn to professional standards or drawn clearly, preferably in ink. They should be about 1½ to 3 times the final printed size and should not contain any unnecessary or irrelevant details. Any symbols and lettering must be large enough to remain legible after reduction. Explanatory text should not appear on the diagram itself but in the caption. Captions should be typed on a separate sheet of paper and should, as far as possible, explain the meanings of the diagrams without the reader having to refer to the text.

Sharp monochrome photographs on glossy paper are preferred: colour prints are acceptable but the use of colour within the magazine is at the Editor's discretion. In either case contrast should be sufficient to ensure satisfactory reproduction.

Units

SI units, or units approved by WMO, should be used.

Copyright

Authors wishing to retain copyright for themselves or for their sponsors should inform the Editor when they submit contributions which will otherwise become UK Crown copyright by right of first publication.

It is the responsibility of authors to obtain clearance for any copyright material they wish to use before submitting it for publication.

Free copies

Three free copies of the magazine are provided for authors of articles published in it. Separate offprints for each article are not provided.

Meteorological Magazine

May 1988

Vol. 117
No. 1390

CONTENTS

Contributions: It is requested that all communications to the Editor and books for review be addressed to the Director-General, Meteorological Office, London Road, Bracknell, Berkshire RG12 2SZ, and marked 'For *Meteorological Magazine*'. Contributors are asked to comply with the guidelines given in the *Guide to authors* which appears on the inside back cover. The responsibility for facts and opinions expressed in the signed articles and letters published in *Meteorological Magazine* rests with their respective authors. Authors wishing to retain copyright for themselves or for their sponsors should inform the Editor when submitting contributions which will otherwise become UK Crown copyright by right of first publication.

Subscriptions: Annual subscription £27.00 including postage; individual copies £2.30 including postage. Applications for postal subscriptions should be made to HMSO, PO Box 276, London SW8 5DT; subscription enquiries 01–211 8667.

Back numbers: Full-size reprints of Vols 1–75 (1866–1940) are available from Johnson Reprint Co. Ltd, 24–28 Oval Road, London NW1 7DX. Complete volumes of *Meteorological Magazine* commencing with volume 54 are available on microfilm from University Microfilms International, 18 Bedford Row, London WC1R 4EJ. Information on microfiche issues is available from Kraus Microfiche, Rte 100, Milwood, NY 10546, USA.

ISBN 0 11 728084 4

ISSN 0026–1149

Printed in the United Kingdom for Her Majesty's Stationery Office.
Dd 0240021 C13 4/88 3936 12521.