European Centre
for Medium Range Weather Forecasts

COMPUTER NEWSLETTER

Fitzwilliam House Skimped Hill Bracknell Berkshire England

Bracknell (0344) 24777 Telex 847908

In this issue:

Number 9 - July 1978

Edited and Produced by User Support Section, Brandon House Ext. 286

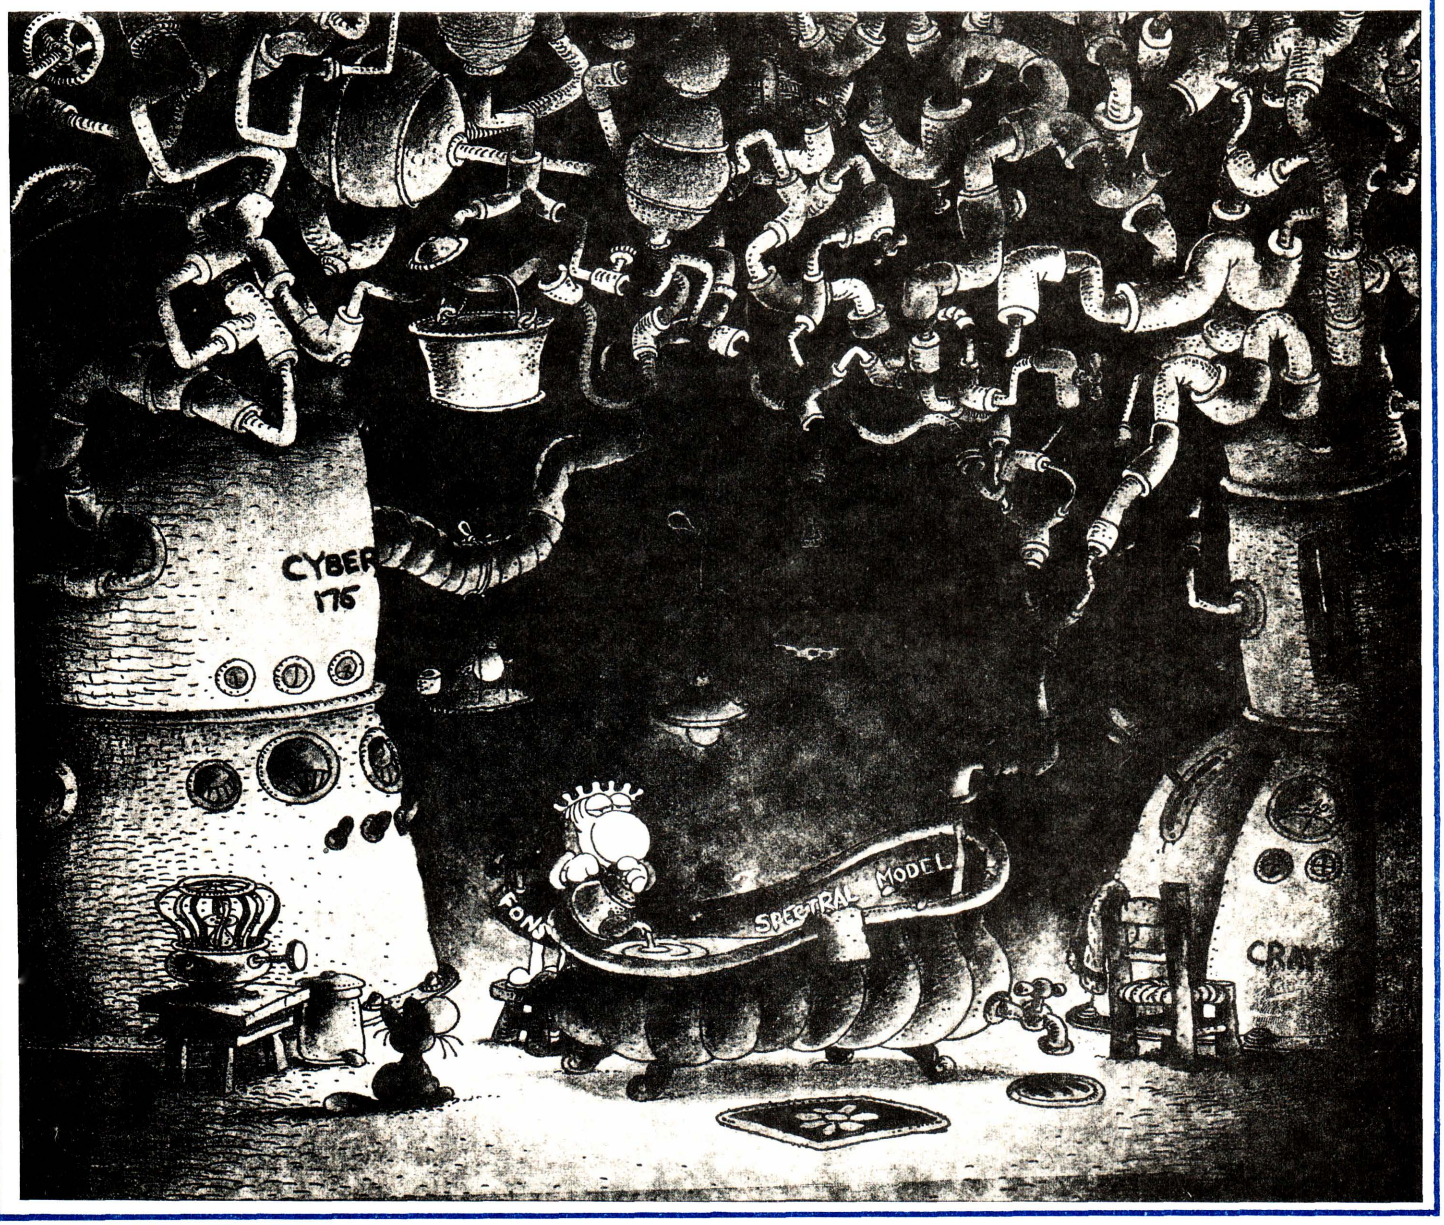

The Spectral Model

Many people at the Centre are now familiar with the concept of a grid point model. However, the representation in grid points of the different quantities, such as wind and temperature, is not the only numerical technique available to us to solve the non-linear partial differential equations governing the behaviour of the atmosphere. There is another group of methods, whereby the fields are represented by analytic functions. One of these methods is used in the so-called spectral model. Here the horizontal fields are represented by a finite sum of spherical harmonics $Y_{m,n}$, analytic functions defined on the sphere:

$$x(\lambda,\mu,t) = \sum_{m=-M}^{+M} \sum_{n=|m|}^{N} X_{m,n}(t)\, Y_{m,n}(\lambda,\mu)$$

where λ,μ are the spherical co-ordinates on the sphere and $X_{m,n}(t)$ are the complex spectral coefficients. The character of the functions $Y_{m,n}(\lambda,\mu)$ is such that, the larger M and N are chosen, the smaller the details of the horizontal fields which can be represented.

Instead of the set of partial differential equations for the set of quantities $X(\lambda,\mu)$, we now obtain a set of ordinary (but still non-linear of course) differential equations for $X_{m,n}(t)$ which then can solved by the same finite difference time stepping techniques as in the grid point model.

In the vertical direction several techniques can be used, but in most existing spectral models the atmosphere is divided into a finite number of layers. In this respect, the spectral model is identical to the grid point model.

This spectral technique has certain advantages, but also some disadvantages. An advantage is that the space derivatives of $X(\lambda,\mu,t)$ can be computed exactly. Another important advantage is that the so-called polar problem, inherent to the grid point model, is avoided. It seems to me that the spectral method is mathematically and numerically more elegant than the grid point method, although I shall be the first one to admit that this judgement may suffer from a certain professional bias.

A disadvantage is that the equations cannot be solved on an arbitrary part of the globe (a "limited area" model) although hemisphere integrations are possible. The grid point method is more flexible in this respect. Moreover spectral models are very much core demanding. Much attention was paid to this aspect in the development of the Centre's spectral model.

A remarkable aspect of the spectral model is that is still requires a network of gridpoints on the globe for an efficient solution of the equations! This has to do with the fact that the equations are non-linear. This grid makes it moreover possible to introduce in the model the effect of small-scale physical processes, such as the formation of clouds, the heating of the atmosphere due to solar radiation etc., just in the same way as in the gridpoint model.

A spectal model along these lines was completed recently on the CYBER and has also been run successfully on the CRAY. The external organisation of the initial data, the history files, etc. is kept as close as possible to that of the grid point model. This will simplify the combination with the physics package and will make it possible to incorporate the spectral model in the operational system in the same way as the grid point model. Ideally, the operational system should not have to know whether the model is spectral or grid point!

There was no previous experience with the vectorisation of spectral models. It seems that we succeeded in this respect, also thanks to the vectorised Fast Fourier Transform written by Clive Temperton, which plays such a prominent role in spectral models.

When the physics package has been built and the model runs fool proof on the CRAY, we shall undertake a set of integrations to compare the spectral model with the grid point model, but with respect to forecast results and efficiency.

- Fons Baede

Hardware Design Fault Uncovered

Congratulations to Clive Temperton on being the first ECMWF staff member to discover a hardware design fault on the CRAY-1. The problem became apparent when he was testing a vectorised version of the FFT package, written entirely in Fortran. Under certain circumstances, the program gave incorrect results, while the scalar version always worked. The suspect area was reduced to a single loop containing 4 statements and it became clear that the special circumstances were related to memory bank conflicts, that is array subscript increments of 16.

No fault could be found in the compiler generated code and we began to suspect a hardware problem. Cray analysts and engineers took over the investigation and reduced the test code further until the fault could be reproduced with a few machine instructions. Three conditions were necessary to force the error.

1. Illustrated by the following CAL instructions:

```
,A0,A1   V0
 V0      V1*RV2
```

   i.e. a vector store into memory is immediately followed by a vector computation where the result register is the same as that involved in the store instruction. In this circumstance, a reservation is placed by the hardware logic on the register V0 to prevent subsequent instructions using it until the store is complete.

2. The vector length is 3, 4 or 5.

3. The memory increment for the store, i.e. the content of address register A1, is 16.

Condition 3 implies that the vector store will take 4 times as long to execute because of the memory bank conflict. This in turn requires the reservation on V0 to be held for a proportionally longer period. This last requirement was not being satisfied. The effect was that V0 was refilled before the store completed. No error occurred for vector lengths of 1 or 2 because the store had completed before any overwriting took place. For vector lengths longer than 5, a fortuitous 'accident' meant that the results were correct. Clive's program was the first (to our knowledge) where all three conditions coincided and the problem became apparent.

Cray design engineers at Chippewa Falls produced a wiring modification and this was implemented at Rutherford within 8 weeks of the date when CRI were originally notified of this problem.

Congratulations to Cray staff on a speedy diagnosis and solution. Any more bugs lurking in there?

- David Dent

Graphics Progress

Two projects are under way:

i) creation of a library of routines for the Tektronix 4014 terminals.
ii) conversion of the Varian basic software for the new Versatec plotter.

The Tektronix 4014 is a large-screen graphic display terminal, measuring 48cm (19") across the screen diagonally. It is a high resolution device, offering addressability up to 4096 in X and 3120 in Y. In addition to a 96 character ASCII keyboard, there is a crosshair controlled by two thumbwheels. The 4014 is a storage tube; that is to say that once a vector is drawn it remains on the screen until the whole picture is erased. On the other hand, because it is not necessary to refresh the picture continuously, complex pictures (such as weather maps) can be drawn without needing large amounts of memory to store the complete picture.

In addition to drawing solid lines, the 4014 has the ability to draw dotted, dashed lines or certain combinations of these. This facility is available in the hardware of the 4014 and so is fast (about 127 metres a second). Characters can either be generated by software and defined by line segments, or by the hardware. Software characters are more flexible because any size or orientation is possible, but they are slower. Hardware characters are faster to draw, but are limited to 4 character sizes and cannot be rotated.

At the moment, there is only one Tektronix 4014 situated in Brandon House. More will be provided in the future for Met. Operations and for general use.

Initially, a library of Fortran subroutines will be created with similar functions to those of the Varian basic software. This is still under development. A future note will be published when this library can be tried by users.

The Versatec plotter is model 8122 with 100 dots per inch and 21.12 inch wide paper just as with the existing Varian plotter. It will be connected on-line to the Cyber and not off-line as for the Varian. As far as users are concerned, there should be no changes required to existing programs when the conversion from Varian to Versatec is made. It will, however, be necessary to re-link existing programs before the conversion takes place. This is necessary because some internal changes are needed to the Varian basic software. A note will be published on this in the future.

- Howard Watkins

TEKTRONIX

4014

Graphic Display Terminal

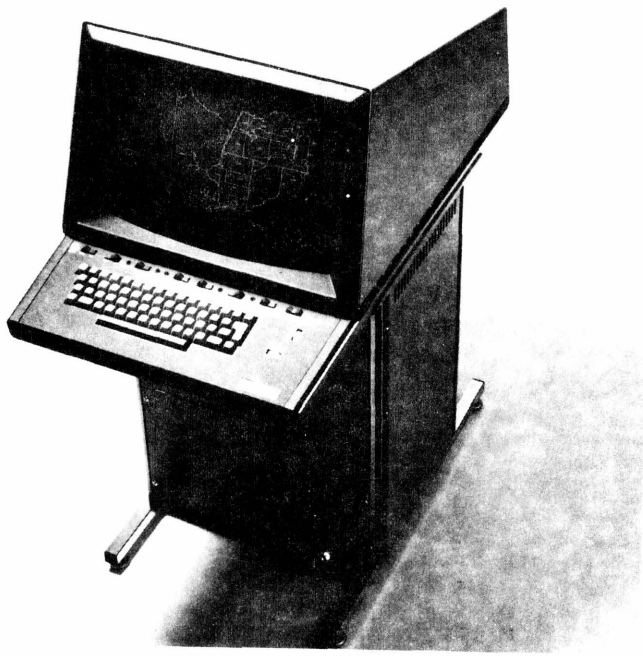

I/O Transfer Rates on CRAY-1

Introduction

The mass storage subsystem of the CRAY-1 consists of DD-19 moving head non-dismountable disc packs connected to the mainframe through controllers. At the Rutherford preliminary installation, the CRAY-1 currently has 2 controllers and 3 discs on line. The final machine at Shinfield will have 4 controllers and 8 discs. A controller can handle up to 4 discs but contains only one pair of buffers for handling data transfer. Each disc can hold approximately 38 million words of data, 64 bits per word. A disc/controller combination is capable of streaming data to or from the mainframe at a sustained transfer rate of about 500KW/sec, allowing for head movement across neighbouring cylinders. For comparison, the equivalent rate on the Cyber using 844 discs and 'full tracking' is 0.925 Mcharacters/sec (87KW/s).

Within the Cray operating system, the data flows through circular buffers which are filled or emptied by the system routines as required. The user may decide on the size of these buffers as well as on the device where a dataset is to reside. A key factor in achieving a high transfer rate, for example on input, is the mechanism by which sectors of data are transmitted to the system buffers in advance of a user request. This requires that there be sufficient room in the buffer, i.e. that the user program reads from the buffer at as fast a rate as the new data arrives. This mechanism is known as streaming. The size of the buffer in relation to the user's record length is important for efficiency.

Although the Cray Fortran compiler supports the BUFFER IN statement, there is no system software comparable to the CYBER buffered I/O whereby data is passed directly between the controller and the user area, without involving a system buffer. This means that we are dependant on efficient operation of the circular buffering scheme in order to achieve high transfer rates.

Performance Measurements

Initial tests with one I/O stream were encouraging. Using long records and a large buffer, sustained rates of 500 Kwords/sec were achieved for output. However, for input under the same conditions, the performance was noticeably degraded. The diagram illustrates results of a test program which measures transfer rate under varying levels of computing load. While the test remains I/O bound, a transfer rate of 500KW/sec is measured. When the test becomes CP bound, the transfer rate is reduced, but it is interesting to observe that almost total overlap is achieved between computing and I/O (95%). It is simple to show that the CP cost of a Fortran READ or WRITE is less than 2%.

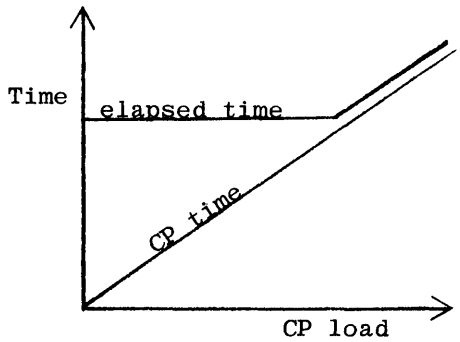

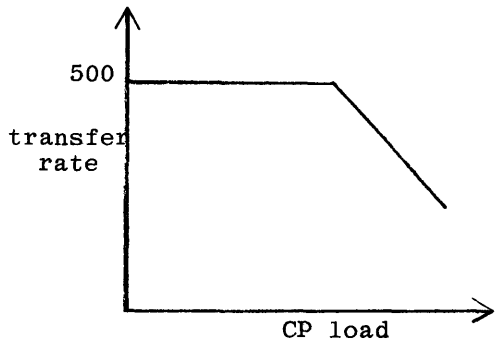

.../cont.

When performance measurements were made using the ECMWF forecast model, a less satisfactory situation emerged. The model is dominated by the I/O requirements of its work datasets. These may be organised so that there are either 2 or 4 active streams during the sweep through the data for each timestep. The 4 stream version is designed to take advantage of 4 disc controllers and trial runs demonstrate that it is unprofitable if there are only 1 or 2 available. The table shows that the 2 stream model achieved 352 Kwords/sec (70% throughput), using only one controller. However, when a second controller became available at Rutherford, this sustained transfer rate (averaged over the complete execution of the model) rose to a disappointing 521 Kwords/sec (52%). These early tests were carried out on a coarse N16/24 grid. The runs include some radiation calculations and during this phase the program becomes CP bound. Additionally, there is some lost performance at the end of each time step due to buffer flushing and filling requirements when switching between read and write mode, unavoidable when using sequential I/O.

	Grid	Record Length	Controllers	Streams	Buffer Size	Time Steps	Elapsed time(s)	Transfer rate(KW/s)	Efficiency %
Model	N16/24	3828	1	2	44	288	613	352	70
		3828	1	4	44	288	688	314	63
		3828	2	2	44	288	415	521	52
		3828	2	2	110	288	379	570	57
		3828	2	4	44	288	464	466	47
Simulation	N48	17072	4	4	44	18	68	1753	88
		17072	1	2	44	18	256	466	93
	N60	21120	4	4	44	18	102	1804	90

Consequently, a simulation model was written to exactly reproduce the work dataset requirements of the explicit model. This enabled more precise measurements of the I/O performance to be made including estimates of elapsed times for any desired grid.

At the same time, Cray system analysts identified several areas of inefficiency in the handling of logical I/O and the circular buffers. In particular, one modification dealt with differences between reading and writing while another improved the triggering of a physical I/O. A new operating system was built incorporating these improvements, and was tested using the simulation program. The second part of the table indicates the improvements achieved. The best case, on N60 simulation, returned a very acceptable 1804 Kwords/sec using 4 disc controllers. This new system is not yet available operationally on the Cray-1 at Rutherford. It has therefore not been possible to carry out further tests yet. Note that even with this high transfer rate, an N60 global explicit model would still require 12 hours for a 10 day forecast.

Random Input/Output

All discussion so far has been confined to sequential I/O using READ/WRITE or BUFFER IN/OUT statements. However, in the semi-implicit formulation of the model, there is an important requirement to read the records of a work dataset in reverse order. Currently, this is achieved by coding

```
BACKSPACE
BACKSPACE
READ( )...
```

An alternative is to make use of the Cray random I/O facility. This allows the user to position randomly to the beginning of any record within a sequential dataset. Sequential I/O statements may then be used to read or write.

Both methods suffer from the disadvantage that the streaming mechanism is destroyed. Clearly, it is impossible for the system to anticipate the users next (non-sequential) request. Work remains to be done to establish the I/O performance for a reverse sweep, but preliminary measurements indicate that the transfer rate will be significantly reduced. The explicit simulation model will be converted to reproduce the more complicated I/O of the semi-implicit model, in order to estimate forecast elapsed times under operational conditions.

.../cont.

There remains one further alternative which may have to be considered in the future if the available software proves inadequate. This is to provide the option of performing I/O directly between the disc controller buffers and the user's work area.

There are several disadvantages:

a) The project to implement this facility is non-trivial and as such would require an (as yet) unknown quantity of effort from either Cray analysts or ECMWF systems staff.

b) The structure of the dataset would have to be unblocked which probably means that there would be no record structure and no security against corruption (as provided by the blocked format). It would therefore not be such a general purpose package as the standard logical I/O facility.

On the other hand, the advantages for the forecast model are clear.

a) The model is already designed around this approach and can adequately overlap the computing load with I/O by initiating I/O requests well in advance.

b) The record structure is well defined and would not suffer from the necessary use of an unblocked format.

c) With careful selection of record lengths, it should be possible to achieve maximum transfer rates on 4 active streams working through 4 controllers for forward AND reverse sweeps.

d) The inefficiencies due to switching from output to input with sequential I/O are avoided.

e) A significant quantity of main memory is freed for use by the model, since the (large) circular buffers are not required.

- David Dent

ECMWF's Class Scheduling System - Cyber 175

This system which was described in a previous News Letter, and whose implementation was recently announced in a News Sheet, will have to be temporarily withdrawn. Certain intermittant problems have been found and hopefully will soon be cured. A News Sheet will announce its re-commencement.

- Neil Storer

Graphics Files

It is no longer necessary to use the "MR=1" parameter on the ATTACH card for the following files:

VARIANLIB,ID=EWPLOT	Varian basic software library
NEWCONTLIB,ID=EWPLOT	Contour package library
BACKBASE,ID=EWPLOT	Coastline data

Although no longer necessary (because the files have the multi-read attribute by default), the presence of "MR=1" does no harm and does not have to be removed.

- Howard Watkins

Recent COS Permanent File Problems

Last month a user ran a job to delete editions 3 and 4 of a permanent file. The CRAY crashed just after doing this and as the user had not received any output, he re-input his job. However, editions 3 and 4 of this file no longer existed, but editions 5 and 6 did, so COS mistakenly deleted these editions instead. THIS SHOULD NOT HAVE HAPPENED.

The problem cannot be reproduced but Cray Research have been made aware of the seriousness of this. Anyone who experiences this problem (i.e. ACCESSing a specific edition of a file, one that does not exist, and being given a different edition) should contact User Support.

Another problem has recently been encountered on the CRAY which is of relevance to users who may produce extremely large permanent files. The problem is this:

> if a file is produced which is greater than about 1600 tracks on the CRAY disc (about 15,000,000 words) and this is SAVEd, then the file may be used only until a deadstart or warmstart is performed. At deadstart, the file will become corrupt in such a way that the first 1600 tracks (approx) will contain valid data, while the rest will become overwritten.

This problem is not as serious as it at first seems. 1600 tracks is a large amount of disc space (about 2/5 of a DD-19). The problem is fixed in COS 1.0.1

- Neil Storer

Clarification of the Eclipse Back-up of Cray Permanent Files

As you will no doubt be aware, there are 6 permanent file IDs reserved for secure files. The disc space occupied by these files has been restricted to 15 Megabytes, this limit being set by the capacity of the Eclipse discs to which the files are dumped. There are 18 such discs grouped into 3 sets (A,B and C) of six (for DUMP00 to DUMP50). These sets are used in the following way:

'A' on Monday and Wednesday

'B' on Tuesday and Thursday

'C' on Friday

As you can see, this ensures that those files dumped on Friday will remain available on set 'C' for 7 days. It was hoped that set 'C' would be backed-up to 3600 foot tape, before being overwritten, thus giving a secure period of 14 days, however, it was found that the maximum size such a tape could accommodate was 12 to 13 megabytes. At user liaison meetings, it was made clear by the users that the provision of 15 megabytes of space with a 1 week security period was preferable to a 13 megabyte limit with a 2 week security period, so that is the situation, as it stands.

- Neil Storer

NOS/BE Changes

The following changes to the NOS/BE system have been made since the last News Letter:

1) Fix to cure 'HUNG IN AUTO RECALL' problem.
2) QCRAYIN speeded up.
3) Plotter spooling utility driven from the NOS/BE sequencer.
4) Permanent file catalog increased in size.

- Neil Storer

CYBER Upgrade Plans

An article in the May News Letter (pages 2 and 3) described the installation plans for Shinfield and their impact on computer users. The following additional information is now available:

1) The additional Cyber PPU's will be installed on the 175 at Rutherford, probably during the first weekend in August. Consequently, there is unlikely to be any CYBER service on that weekend (5/6 August).

2) The version of NOS/BE required at Shinfield (release 4) will be made available for special test sessions during September. It is important that all users who have working programs on the CYBER should take this opportunity to try the new release of the operating system. Early experience of conversion difficulties bugs, etc. will be valuable to all.

For further information on NOS/BE release 4 changes, see the following article by Luigi Bertuzzi.

- Andrew Lea

NOS/BE (Release 4) Record Manager

A number of new features will be introduced by the NOS/BE system release available at Shinfield. Many of these affect the Record Manager.

This article is an attempt to provide information on what is going to change in the Record Manager. A few general aspects will be addressed in greater detail by future documentation.

The new Record Manager

The new Record Manager will comprise two separate file managers, each of them being described by an appropriate Reference Manual.

From Shinfield onwards, the term CRM (Cyber Record Manager) will be a generic substitute for either:

BAM (Basic Access Methods) or
AAM (Advanced Access Methods),

or for them both, as the case will dictate.

The BAM file manager deals with SQ (Sequential) and WA (Word Addressable) file organisations.

The AAM file manager deals with IS (Indexed Sequential), DA (Direct Access) and AK (Actual Key) file organisations. It includes a process called MIP (Multiple Index Processor), allowing AAM file access by alternate keys.

Main Characteristics

The function organisation and the loading scheme of the BAM/AAM structure allows field length reduction at execution time and improved performance.

A permanently loaded subset of BAM will dynamically load all necessary CRM object modules using the FDL (Fast Dynamic Loader).

Many of these modules will be grouped into transient capsules, which remain in memory until another capsule must be loaded or a buffer must be allocated.

Other modules are grouped into permanent capsules, which must be in core as long as any activity is taking place on any file. They are unloaded when no longer in use by any file.

To avoid loading and unloading transient capsules many times, wasting FDL CP time, users will be recommended to open and close all their files together.

.../cont.

Capsule loading and buffer allocation will start at HHA (Highest High Address), which is the end of a basic load or the end of the longest overlay.

Memory management will be controlled by the CMM (Common Memory Manager), except in the case when BAM is loaded statically, thus allowing execution of programs doing their own memory management (memory requests).

One instance of the above is where a Fortran user declares a blank common area at the end of the loading region and uses this area for overindexing of arrays. Such programs must be recompiled invoking the STATIC parameter on the FTN call statement, at the expense of somewhat more memory space.

The FIT (File Information Table) has been reformatted and combined with the FET (File Environment Table), but upward compatability of binaries generated at or after FTN level 446 is supported (i.e. the FIT/FET structures of these binaries are dynamically converted to the new format). Code generated by FTN prior to level 446 must be regenerated by recompiling the source programs. Our current FTN level is 452.

Most commonly used file structures remain unchanged; e.g. the default structures produced by Fortran WRITE or BUFFER OUT. Some of the more unusual file types require conversion. To ease compatability problems the FORM utility has been revised and is being released as FORM 1.1. It will read files written by CRM, BAM and AAM, with the exception of files using K or E type blocks with W type records, and K type blocks on PRU devices. These excepted, file structures should be converted prior to the installation of the new system.

Error processing will improve. The two FIT fields related to error processing (SDS and EXD) have been replaced by the two new fields DFC (Dayfile Control) and EFC (Error File Control), but the old ones will still be accepted, producing slightly different results.

The error information is going to reside on file ZZZZZEG (instead of ZZZZZEF) in binary format, and a CRMEP control statement must then be used to write selected meaningful diagnostics to the output files.

The RA+2 list of files will no longer be supported. A pointer to a table of the names and FIT addresses of all active files will be maintained instead. Details on use of this new feature will be found in the NOS/BE Reference Manual (Rev.E), and in Appendix F of the new BAM Reference Manual (to be distributed shortly).

LDSET(FILES=) control statements will no longer be needed following a FILE control statement. Their presence will be ignored.

Coming Soon

More information, as made available by Control Data literature, will be given shortly.

Expected changes introduced by BAM and AAM will be covered as exhaustively as possible, but validation of reported information can only be made when user test time is available with the new system.

- Luigi Bertuzzi

Data Management on CRAY-1

The owners of many large data sets on the Cray cannot be identified at present. This causes problems when the discs overflow as the owners of the old unused data sets cannot be found.

For the Spring Experiment to run successfully on the Cray, about 40 million words of disc space are required. Recent experience suggests this is not always available. Therefore, ALL users are urged to adhere to the following conventions in order to improve data management:-

a) The current system of DUMP IDs continues for all data that needs to be backed up, and which is short enough to be dumped to eclipse disc.

b) All other data sets created by ECMWF staff to be saved with IDs conforming to the current CYBER standard.

c) Data owned by groups of users to be saved with special IDs identifying the group. Information concerning such IDs and nominating a person responsible for their management to be addressed to my office at Rutherford.

- Eric Walton

CRAY-1 Test Compiler

The Fortran compiler and library changes described in the last News Letter (May 1978 - page 7) are now available. Nearly all of the compiler bugs discovered by ECMWF staff have been fixed in this new release. The CYBER file

CRAYBUGS,ID=EWDD1 (see News Letter No. 7)

has been updated to contain this information.

Note that the new compiler directives

CODE,NOCODE,VECTOR,NOVECTOR

apply to the complete program unit (i.e. main program or subroutine). It is not possible to switch on and off code generation or vectorising attempts within such a unit.

- David Dent

CRAY-1 Second Disk Controller

The second CRAY disc controller was introduced successfully on 19 June. User programs which perform significant amounts of I/O may well benefit from careful positioning of important datasets. Where I/O is active on 2 datasets simultaneously, the datasets should be placed on discs connected to separate controllers:

```
e.g.    ASSIGN(DN=FT01,BS=22,DV=DD-19-20)
        ASSIGN(DN=FT02,BS=22,DV=DD-19-22)
```

The use of large system buffers (BS=22) is usually beneficial for record lengths greater than about 1000 words. The default buffer size is 4 (i.e. buffer holds 2048 words).

Disc units DD-19-20 and DD-19-21 reside on one controller while DD-19-22 resides on the other.

The forecast model being used for the Spring Experiments initially required about 30 hours elapsed time for a 10 day forecast. By directing the work datasets so that they use separate controllers, this elapsed time has been reduced to about 20 hours.

- David Dent

Quality Assurance Test Jobs

It is extremely important that we continue to build up a set of working CRAY-1 programs in order that new compiler, operating systems, etc. can be adequately tested. Would all who have such programs and who are willing to allow us to use them for QA work, please contact User Support with the information necessary to run them.

- Frank Hollmann

Q and A

Question : My program intermittently issues the dayfile message
"BUFFER ARGUMENT ERROR"
The program executed without error. Why?

Answer : The program does not terminate the last BUFFER I/O statement with an IF(UNIT) statement. Hence, at job termination time, there may be some uncompleted I/O in progress. A PSR has been issued to CDC making Record Manager check I/O complete before attempting to close the file.

Interdisciplinary Laws

We propose to publish in the News Letter on a regular basis, some of the fundamental 'laws' of modern life. We begin this month by reproducing the well known Murphy's Law.

Murphy's First Law

Nothing is as easy as it looks.

Murphy's Second Law

Everything takes longer than you think.

Murphy's Third Law

In any field of scientific endeavour, anything that can go wrong will go wrong.

Murphy's Fourth Law

If there is a possibility of several things going wrong, the one that will cause the most damage will be the one to go wrong.

Murphy's Fifth Law

If anything just cannot go wrong, it will anyway.

Murphy's Sixth Law

If you perceive that there are four possible ways in which a procedure can go wrong, and circumvent these, then a fifth way, unprepared for, will promptly develop.

Murphy's Seventh Law

Left to themselves, things tend to go from bad to worse.

Murphy's Eighth Law

If everything seems to be going well, you have obviously overlooked something.

Murphy's Ninth Law

Nature always sides with the hidden flaw.

.../cont.

Murphy's Tenth Law

Mother Nature is a bitch.

Murphy's Eleventh Law

It is impossible to make anything foolproof, because fools are so ingenious.

Murphy's Law of Thermodynamics

Things get worse under pressure.

O'Toole's Commentary on Murphy's Laws

Murphy was an optimist.

Previous Newsletter Articles

Listed below are the major articles published in previous News Letters which are still valid. However, as one goes back in time, some points in these articles may no longer be accurate. When in doubt, contact the author, or User Support (ext. 286).

	NEWSLETTER		
CRAY 1	No.	Date	Page
Bug list for users to inspect	7	April 78	7
Computer configuration	5	Jan. 78	1
Fast Fourier Transforms	8	May 78	1
Libraries (NAG, NCAR)	8	May 78	8
Service Schedule	6	March 78	3
Software - next release of COS/CFT etc.	8	May 78	7
Transfer of CB files to the Cyber (DSOUT)	7	April 78	6
UPDATE - advice for users	5	Jan. 78	5
CYBER 175			
Computer configuration	5	Jan. 78	1
Disk packs (844-41, double density)	8	May 78	5
Fortran - Ill conditioning in programs	4	Nov. 77	4
- MNF compiler for debugging	1	Aug. 77	1
Libraries (NAG)	8	May 78	8
Job Scheduling	7	April 78	3
Routing printed output	7	April 78	5
Service schedule	6	March 78	3
Terminals - Newbury VDU problems	8	May 78	4

.../cont.

GENERAL

Documentation - Computer Bulletins	6	March 78	3
- Computer Manuals	8	May 78	6
- News Sheets	4	Nov. 77	2
Fortran 77	2	Sept. 77	4
Plotter plans	6	Mar. 78	2
Shinfield moves - Oct. to Dec. 78	8	May 78	2
Telecommunications project	8	May 78	4
Transport schedule to Rutherford	6	Mar. 78	3

METEOROLOGY

Operational forecast suite	7	April 78	1
Spring Experiments (1978)	6	March 78	1

USEFUL NAMES AND PHONE NUMBERS

		Room	Ext.
ADVISORY - Cray	David Dent	BH102	286
- Cyber 175	Luigi Bertuzzi	BH103	284
Computer Division Head	Rob Brinkhuysen	FH509	210
Computer Terminal Rooms	Brandon House	BH107	296
	Fitzwilliam House	FH401/2	225/229
	John Scott House	JS Central Room	256
Disk space and permanent file problems		As for ADVISORY	
DOCUMENTATION	David Dent	BH102	286
INTERCOM - registering new users	Jean-Luc Pepin	BH111	294
OPERATIONS - Computer Room Rutherford - Graham Holt	}	(0235) 838099 or (0235) 21900 x229	
- Eric Walton		RL2-40	(0235) 83288
Research Dept. Computer Co-ordinator	Rex Gibson	FH411	248
Tape requests	Pauline Litchfield	As for OPERATIONS	
Terminal problems		As for ADVISORY	
User Support Section Head	Andrew Lea	BH101	289

CERN Program Library

The CERN program library is now available on the CYBER. It contains a number of very useful utility routines. A Bulletin is being prepared describing their use. To reference the library:

```
ATTACH,LIB,CERNLIB.
LIBRARY,LIB.
```

Corrections to the May News Letter (No. 8)

On page 7, last line:

The CFT compiler directive to specify a 24 bit integer type is-

```
CDIR$ INT24   a,b,c
```

The example on the same page should read-

```
A(K+I) = B(2*K+J)
         (K+I) and (2*K+J) calculated in A register.
```

Coming Soon

Meteorological Satellites.

NOS/BE Release 4 - further details

News Letter - Next Issue

The next issue will be in September. PLEASE contribute.
