User Requirements for the Data Handling Facilities at ECMWF

J.W. Münch

University of Siegen

1. Introduction

On account of its objectives, ECMWF stores and processes large quantities of meteorological data mostly on magnetic tape. Data volume and number of magnetic tapes are now so great that access to the various data by means of the existing data processing system is becoming difficult. Thought must therefore be given to a new data handling system.

User requirements with regard to this system are based on the following tasks:

- Collection and storage of meteorological observation data,
- production of operational forecasts,
- improvements of the forecast model, and the research work related to it,
- various services for the Member States.

All four of these activities need data processing capacity. The following review is restricted to

- data storage demands,
- data retrieval demands,
- and software demands concerning data storage and retrieval.

The growth of the volumes of the data sets results from the

- accumulation of observational and analysis data with time,
- refinements of the used model of the atmosphere in space and time,
- and increasing use of ECMWF by the Member States.

The corresponding user requirements were compiled from various memos of the Research Department written by A. SIMMONS and from user interviews.

2. Description of the Current Data Sets, their Contents, Usage and Structure

All datastreams archived are listed in Table 1 at the end of this section. They may be subdivided into operational and research data. Only the more important kind of data of a data set is listed under the heading 'contents'. The operational data streams DASU, DAOB, and FUSG containing, besides observations, analysis and forecast data stored directly in the coordinate system of the model. They are the basis for other datastreams, and consist of sets of data defining the state

of the model atmosphere at a particular time. Such a set called 'model file' consists of a sequence of records each containing the complete data for one line of latitude. These files comprise a sequence of descriptor records followed by data records. The descriptor records contain also information characterizing the model used. The data on the tapes may be packed, i.e. three or four values per 60 bit Cyber word.

These datastreams are used both by the Operations and the Research Departments. They are kept for approximately one year, and the tapes are then recycled. Henceforth, datastreams stored only for a year or less belong to the "short term archive", while data stored for more than one and less or up to five years belonging to "medium term archive".

For the verification of operational and research forecasts, pressure-level forecast and analysis data from the datastreams,AIPG, FUPG and FSUR are used. These forecast and analysis data are postprocessed data of the above mentioned model files. The basic element of such a dataset is a "field", namely one model variable at one pressure level (or on the earth's surface) at a particular time.

These fields are stored as coefficients of a special harmonic expansion for the upper air or in grid point form for the surface fields. Two or less commonly four fields, normally at eight pressure levels, are used for verfication. A subset of AIPG and FUPG is called VERIFY archive.

All sorts of verification and evaluation of forecasts use pressure level forecast and analysis fields for plotting. These data are kept for five years before their tapes are recycled.

The pressure-level data will also be used in applications involving access not only to a few fields but to most fields at most levels. Such are case studies and budget studies. The latter ones may include all analysis and all forecasts at selected time ranges for a particular season or several seasons.

On the contrary,applications involving long time series make only use of a few fields such as the 500mb hight, mean sea level pressure, vorticity and divergence at 200mb and 850mb, and 850mb pressure.

Besides the pressure-level data, i.e. fields, analysis and forecast data in model file form are also used for the similar applications like production of monthly mean files, comparison with research forecasts or diagnostic studies.

The observational data archive RDBG is kept permanently. Its unit of retrieval

and storage is a complete report, containing data from one source e.g. a radiosonde report. Regular access to this datastream may be anticipated for the calculation of monthly mean rejection and flagging rates. It is also accessed for the diagnosis of the analysis and initialization of the analysis , for which the data streams AUPG and FGPG are used, too. The latter two contain pressure-level data.

For quick access to the archived data there exist three European subsets of global datastreams, namely AUPE, FAPE, and RDBE which are kept permanently and contain data in a grid-point coordinate system. There exist four other data sets not directly related to the afore mentioned ones:

The Experimental forecast archive records the results e.g. for higher resolution model experiments or other types of new forecasting systems. It is unlikely that many forecast results will be retained for longer than five years.

In the same way the Experimental Data assimilation archive is used for data assimilation experiments. The majority of these data are stored in a format called "PT", which is likely to be changed into one of those already mentioned.

The observational and analysis data of the Special Observing Period are contained in the FGGE-Reports and FGGE data assimilation archive. They are often used to compare the results of model changes. Part of the FGGE-data assimilation archive has the same data structure as the Experimental Data assimilation archive and will also be changed; the FGGE observational data stream has a format which does not agree with RDBG but has been drawn up to fit best into the assimilation process.

There exists also a dataset containing analysis from other weather centres and various climatological data; these data are stored in various formats.

Rationalization of the data streams may be achieved by making all the Experimental Data assimilation archive and the FGGE data assimilation archive compatible with the DASU and DAOB data streams.

There exists a strong user requirement to create a new operational stream containing routine diagnostic data. Significant demand for times series individual fields is anticipated both from within and without the Centre. It would therefore be advisable to maintain such series well documented as part of the operational data stream.

NAME	CONTENTS	STRUCTURE	PACKED	ARCHIVE
DASU	ANALYSES	MODEL FILE	NO	SHORT TERM
DAOB	OBSERVATIONS and ANALYSES	MODEL FILE	NO	SHORT TERM
FUSG	FORECASTS ANALYSES	MODEL FILE	YES	SHORT TERM
AIPG	ANALYSES	FIELDS	YES	MEDIUM TERM
FUPG	FORECASTS	FIELDS	YES	MEDIUM TERM
FSUR	ANALYSES and FORECASTS (SURFACE VAL.)	MODEL FILE	YES	MEDIUM TERM
VERIFY	ANALYSES and FORECASTS	FIELDS	YES	MEDIUM TERM
RDBG	OBSERVATIONS	REPORT	YES	PERMANENT
AUPG	ANALYSES	FIELDS	YES	PERMANENT
FGPG	FIRST GUESS	FIELDS	YES	MEDIUM TERM
AUPE	ANALYSES	GRID POINT	YES	PERMANENT
FAPE	FORECASTS	GRID POINT	YES	PERMANENT
RDBE	OBSERVATIONS	REPORT	YES	PERMANENT
FVSG	FORECASTS and ANALYSES	FILE	YES	MEDIUM TERM
EXPERIMENTAL FORECASTS	FORECASTS	FIELDS	YES	MEDIUM TERM
EXP. DATA ASSIM.	DATA ASSIM.	PT,FIELDS	YES	MEDIUM TERM
FGGE REPORT	OBSERVATIONS		YES	PERMANENT
FGGE ASSIM.	DATA ASSIM.	PT, FIELDS	YES	MEDIUM TERM
DIAGNOSTIC and CLIMATOLOGICAL ARCHIVE	ANALYSES and CLIMATOL. DATA			MEDIUM TERM

Table 1 List of the archived datasets

3. Required Facilities for Storage and Retrieval of Data

An essential requirement for a new data handling system is that it should support general data base activities. The requirements concerning storage and retrieval of data cover the following activities:

Storage of data should be able to be carried out from any computer within ECMWF's network.
Exact recovery of data must always be guaranteed unless the user agrees to a packing density which may lead to less precision.
Data retrieved should be available on any computer regardless of where they were originally written.
Unpacking, data conversion and data transformation should be performed on request for retrieved data, especially transformation from spherical-harmonic to grid-point form. The user should have the option to specify a different format or structure in which retrieved data are loaded.

During the retrieval process an optional operation should be possible, enabling averaging or subtraction of fields or files.

Field descriptors in standard format should be provided with retrieved fields. The user should be able to indicate retrieval priority.

All these functions and options should be able to be used interactively by VDUs as well as from JCL and by FORTRAN-callable subroutines.

Any planned data handling system should at least be able to manipulate and transfer data at rates achieved at present.

4. Current and Expected Volume of Data Sets

The requirements concerning the time periods for which data have to be stored lead to three different types of archives:

permanent archive	(PA)
medium term archive	(MA)
short term archive	(SA)

The increasing data storage requirements for the individual years are shown in the following.

To a great extent, they depend heavily on analyses and models used. Models using higher resolution in time and space need more storage space than those using lower resolution.

Models stored as spherical coefficients need less space than those in grid format. Storing several variable values in one computer word results in a corresponding reduction of space required.

Use of spherical coefficients in a packed format would therefore require minimal storage space.

If, however, data stored in such a way are often used for weather maps, computer load will be heavy because of unpacking and conversion into grid format. The same considerations apply to representation of redundant data in field or model file form. With regard to each archive used, it will have to be considered which type of representation is suitable for a particular purpose.

Fig. 1 shows the number of tapes accumulated in the different archives as well as the number of all tapes to be stored at ECMWF. To calculate the number of tapes for individual archives and years, the following has been taken into account:

From 1985 onwards, there will be a new model using a greater number of pressure levels, and three hourly values of the analyses will be stored. The number of tapes of the permanent archives (PA) also includes the number of back-up tapes required. The medium term archive (MA) is subdivided into tapes used for operational MA (OP) and research MA (RS) data streams. Growth of the MA (RS) archive is partly due to the higher frequency of analyses in future and a repeat of the analyses of the FGGE archive. The number of tapes required by the Member States (MS) is assumed to increase by 20% per year, while the number of pool tapes, i.e. tapes for temporary storage, is assumed to increase by 10% per year.

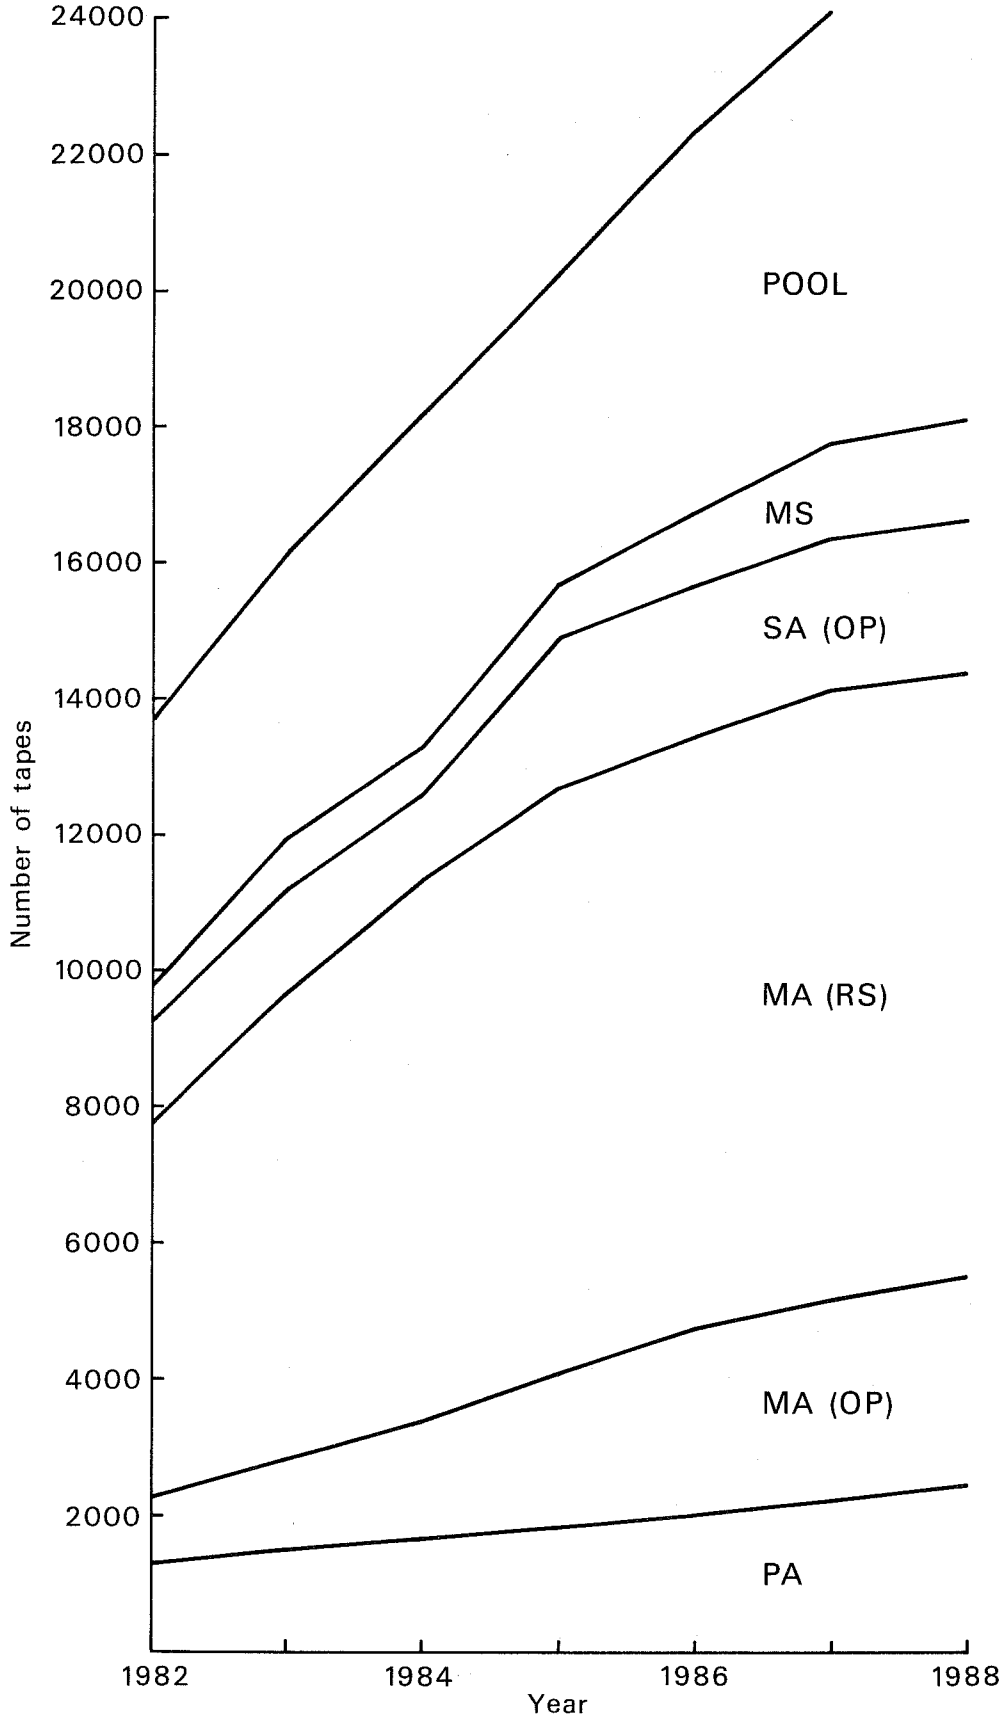

Fig. 1 Number of tapes accumulated. (For explanation, see text)

5. Access Times and Frequency of Requests

The essential requirement with regard to access times to data is that, at worst, they will not be longer with a new data handling system than at present.

One improvement is greatly needed: maps of individual fields from at least the most recent ten days of operational analyses and forecasts should almost instantly be available. For evaluating analyses, observational data are also required.

Table 2 shows the required access times and frequency of requests as well as volumes of data to be handled.

Data Use	Machine on which data usually required	Data volume	Required net access and transfer times	Estimated frequency of type of request
Maps	CY/CR	10^3-10^4 W	0-5 mins	several 1000's/day
Initial data for forecast	CR	2×10^6 W	$\lesssim$ 5 mins	several 10's/day
Initial files for data assimilation	CR	4×10^6 W	$\lesssim$ 5 mins	several per day
Reports for a data analysis	CR	5×10^5 W	$\lesssim$ 5 mins	10 per day
Verification of an individual forecast	CY	10^7 W	$\lesssim$ 30 mins	10 per day
Monthly-mean verification and diagnostics	CY/CR	2×10^8 W	24 h	several per month
Time series	CY	10^6-10^7 W	5 mins - 24 h	variable, but up to several per week
Personal user archives	CY/CR	ANY	$\lesssim$ 5 mins for $\lesssim 10^6$ W	several 10's/day

Table 2: Required access times and frequency of requests

6. Present frequency of tape usage

One of the problems connected with the growth of the number of tapes stored at ECMWF is the steadily increasing number of tape mounts per week. On the one hand, this results in longer access times for users and, on the other hand, it leads to a situation where the physical limits of the tape drives are reached. The following considerations deal with tape mounts irrespective of tapes for input or output. The amount of data retrieved from or stored on a tape has therefore not

been taken into account either. For four weeks in April and in September 1982, the number of tape mounts was analysed, evaluating which tapes and how often they had been mounted.

Table 3 shows, for example, the number of tapes used per week and the number of tapes mounted per week, i.e. one week in April and one week in September. (Tapes used for system analysis and related activities have been subtracted.) The figures studied for the eight weeks showed the same tendencies:

	Week 18.4 - 25.4	Week 19.9 - 26.9
Tapes used	970	720
Tapes mounted	2270	1904

Table 3: Number of tapes used and number of tapes mounted per week

When we analyse the number of tapes used and the number of tapes mounted for the different projects, we can easily distinguish two different types of tape usage: projects with a large number of mounts of an equally large number of tapes, and projects with a large number of mounts using only a few tapes. For the latter projects, the amount of tape mounts could be reduced by keeping multiple mounted tapes in a directly accessible storage for the user.

When we study the mounts of tapes for the former projects, we can determine for how long - in terms of consecutive weeks - a specific tape was used and how often it was mounted during consecutive weeks. If the data of such a tape were stored in a directly accessibly storage and kept there for the weeks during which they are requested, all the mounts of tape could be saved. Table 4 shows the number of tapes which should be kept in a directly accessible storage - following the arguments given above - and how many tape mounts could be saved by these means:

	Week 18.4 - 25.4	Week 19.9 - 26.9
Tapes to be kept in storage	248	241
Mounts saved	1275	1029
Mounts remaining	44%	48%

Table 4: Number of tape mounts to be saved per week

It was furthermore investigated from which archives the 241 tapes listed in Table 4 for the week 19.9 - 26.9 originated. It was found that 51 tapes belonged to the operational archives and 190 to the research archives, 22% and 38% being the corresponding percentages of tapes used. 338 or 62% mounts of tapes from operational archives and 691 or 51% mounts of tapes from research archives would have been saved. These figures indicate that access to the data stored could be very much improved and the workload for mounting the tapes could be much reduced if the new data handling system would provide enough online storage.

7. Concluding remarks

The existing archives of ECMWF contain many different data streams of different structure, some of which can contain redundant information.

With the introduction of a new data handling system, variety and redundancy should be reduced as much as possible without giving up user friendliness. This concerns especially data substreams which may be kept for this reason. A new system should by no means be slower than the existing one. It should provide better facilities for data storage, data retrieval and data manipulation by means of a general data base concept and making it usable interactively or via JCL or in FORTRAN callable subroutines.