Workstation MARS - data retrieval and manipulation

B. Raoult
ECMWF
Reading, UK

1. Introduction

MARS (Meteorological Archival and Retrieval System) is both one of the largest meteorological archive and a simple user interface to this data. This archive has been available for years on the mainframes. With the arrival of more and more powerful workstations at ECMWF, there was a need to access MARS from any machine within the centre. One of the main design issue of MARS on the workstation was the development of the Metview project, allowing interactive access to any fields, field manipulations and plots. Providing data for direct manipulation should be as fast as possible, but easy on the network.

2. One of the largest meteorological archive

The MARS archive represents more than 2.5 TeraBytes of Operational data and the same amount of Research data.

The volume of data archived is growing all the time. In 1985, each day 20 MegaBytes of analysis, 20 MegaBytes of initialization, 26 MegaBytes of first guess and 200 MegaBytes of forecast were archived. In 1991 the figures were respectively 138 MegaBytes, 124 MegaBytes, 353 MegaBytes and 1165 MegaBytes per day!

The fields are saved using the GRIB format. Observations, satellite images and feed back are packed in BUFR. The archive has also some time series and means.

Some other projects are handled by the system, such as the Ensemble Prediction System or the Wave project. Other centers products are save for comparison purposes.

3. A simple user interface

The MARS interface is based on a very simple language. It is batch oriented and the requests are defined in meteorological terms, not in computer terms.

The following request will retrieve 12 fields from the 10th forecast day into the file named `file`:

```
RETRIEVE,
   TYPE= FORECAST,
   LEVTYPE= PRESSURE LEVELS,
   LEVELIST= 1000/850/700/500/400/300,
   REPRES= SH,
   PARAM= Z/T,
   DATE= 931121,
   TIME= 1200,
```

```
   STEP= 240,
   TARGET= "file"
```

This request retrieve some observations:

```
RETRIEVE,
   TYPE= OBSERVATION,
   AREA= EUROPE,
   OBSTYPE= SYNOP SHIP,
   DATE= 931122,
   TIME= 03/TO/09
```

4. Fully networked, distributed databases

The main archive is sitting on an IBM 9000. The access to this server is done using TCP/ip. Because of the networking aspect of the system, its design must take care of the speed and amount of data transfers over the network. To minimize the traffic, a multiple cache system was implemented. The more often the data is requested, the "closest" it comes to the user. These caches are shared by all users. Figure 1 shows the data path between the archive and the user's target file.

Figure 1 MARS Caches

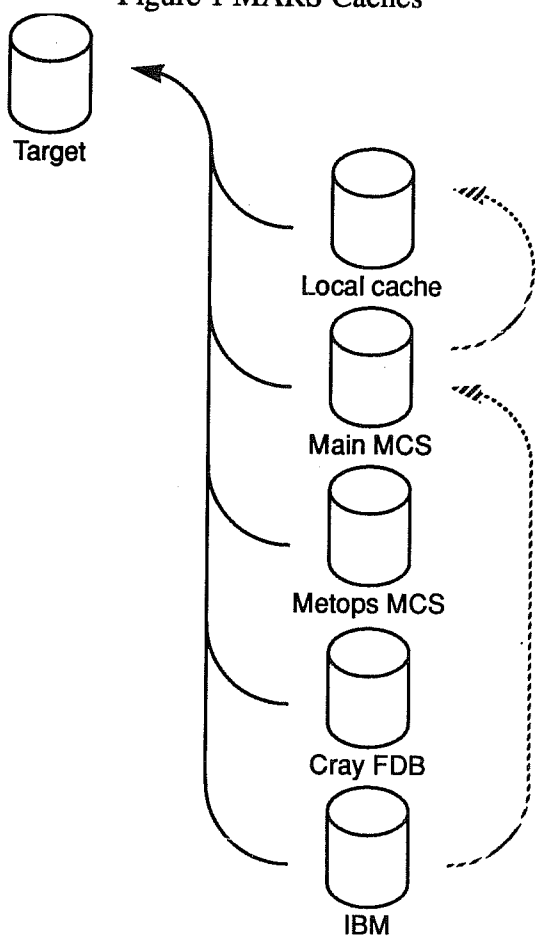

The request goes down from the Local cache to the IBM. When the data is found, it is returned to the user. Any data retrieved from IBM will be copied to Main MCS. Any data retrieved from Main MCS is copied to Local cache. Local cache is a set of files on the users own workstation. Main MCS is 4 GigaBytes in a shared Empress/SQL database.
The system can be configured, and as many databases as needed can be added.

To retrieve a few fields from the system, it can take 1 mn 30s if the data is on tape on the IBM, 40s if it is on a disk, and 10s if it is cached.

5. Data processing abilities

MARS is not only a retrieval system. It can also manipulate fields.

- Built-in manipulations
MARS can change the representation of the fields e.g. from spherical harmonics to Gaussian grid. It can also change the resolutions of the fields, perform sub-area extraction and derive the wind component from the vorticity and the divergence.

- Free form computation
MARS can also perform formulae oriented computations. The formulae can be any mathematical expression. The next example shows how to compute the mean layer temperature.

```
retrieve,
   date       = -1,
   param      = z,
   fieldset   = z1000,
   levelist   = 1000,
   grid       = 3/3

retrieve,
   fieldset   = z850,
   levelist   = 850

compute,
   fieldset   = mean layer t,
   formula    = "-(9.8/8.3)*(log(z850)-log(z1000))/(log(850)-
log(1000))"
```

Or, retrieving LSP (large scale precipitations) and CP (convective precipitations), which are cumulative precipitation, the following computes the total rain between to time steps.

```
retrieve,
   type     = fc,
   param    = lsp,
   fieldset= lsp,
   step     = 12/to/240/by/12,
```

```
   grid    = 3/3
retrieve,
   param   = cp,
   fieldset= cp
compute,
   fieldset= precip,
   formula = "cp + lsp"
compute,
   fieldset= rain,
   formula = "precip[2,count(precip)]-precip[1,count(precip)-1]"
```

The formulae are expressed using operators such as + - / * ^ or = < > >= <= <>, and a set of mathematical functions (log, sin, cos,...) or statistical functions (mean, variance, covariance, rms,...)