UNIX CONVERSION AT THE CANADIAN METEOROLOGICAL CENTRE

Richard Jones
Canadian Meteorological Centre
Dorval, P.Q.
Canada

ABSTRACT

This paper will describe the major activities of the Unix conversion of the operational runs at the Canadian Meteorological Centre (CMC). This includes the conversion of operational runs from a XMP416 to a NEC SX3-44, and front end jobs from a CYBER 830 to a MIPS 4680. The Unix environment and workstations will be briefly described as well as changes to the main sets of software.

1. INTRODUCTION

Unix conversion began in March, 1990, when the decision to convert the front end computer activities to Unix-based machines was taken. In the Spring of 1991, the decision to replace the XMP416 with a NEC SX3-44 was taken.

The decision to convert to Unix involved a considerable amount of discussion within the CMC and RPN (Recherche en Prévisions Numériques). The principle reasons offered to convert were: to minimize the costs of future conversions; better transportability of applications from machine to machine; not being tied to one vendor; the capacity and connectivity of Unix, the availability of software under Unix, the reduced cost and the increased efficiency of learning one operating system instead of many.

CMC operations expect to be fully converted on the front end computers as of April 1, 1992. In the summer of 1992, CMC operations expect to be converted on the NEC.

At this point, the total human resource cost of converting CMC operations as well as related development activities to Unix is estimated at **60 person years.**

2. HARDWARE

The present operational computers include a dual CPU Cyber 830, and a 930. The backend

machine is an XMP416. As backups and for the use of development and research there are another 830 and a 930.

On the Unix side, there are presently four Cyber 920's each with 4 processors (30MIPS/processor) and 4GBytes (GB) of disk. One 920 is operational, and the others are used for research and development.

There is also a MIPS 4680 with 16GB of space. This is expected to be upgraded to a two-processor MIPS with 32GB of disk space early in 1992 (possibly four processors with 48GB of disk space). On April 1, 1992 CMC will be operational on this MIPS with the removal of the operational Cyber 830 machine. Another identical MIPS will be used as backup and for the use of development and research.

The power of one processor of a MIPS 4680 is in the order of 55MIPS or 10Mflops. This is roughly the scalar processing speed of one processor of the XMP. This is the reason that much of the graphical alphanumeric production is being shifted from the backend to the front end.

As of November, 1991, the NEC SX3-44 is being installed at the CMC. The NEC will be installed with four processors each with four functional units with four vector pipelines, for a total of 16 pipes/processor. It will be installed initially with 1GB of central memory and 2GB of XMU (a SSD equivalent), and about 100GB of disk. The CPU clock speed is 2.9ns (6ns for the YMP8, 8.5 for the XMP416). Peak speed is 22Gflops (2.7Gflops for the YMP8, and .85Gflop for an XMP416). Matrix product have attained 4.8Gflops on 1 NEC processor (2.4Gflops for a YMP8) and 17-18 Gflops with 4 processors. Similarly, Livermore loops show that one 1 NEC CPU (2950Mflops maximum) performs at twice the rate of a YMP8 (1308Mflops maximum).

These performance numbers apply to certain numerical kernels. It is not clear at this point how the NEC SX3-44 will be able to perform with the numerical models at CMC.
It is estimated at this point that at least a ten-fold increase in performance over the XMP416 will be possible.

Networks: Mainframes are connected though an FDDI network. Shortly two FDDI rings will be installed through CISCO routers to allow interconnectivity with operational and non-operational machines.

At present at CMC/RPN, there are 55 Cyber 910's (Silicon Graphics) with 35 X-terminals and 20 Sun workstations. There are also several HP9000 workstations as well as Microvaxes.

There is a fast link that connects Toronto computers and the local area networks. Through CISCO routers, the Wide Area Network (WAN) connects many weather centres across Canada as well as military sites. Locally these CISCO routers are connected to their LAN's and workstations. Presently, alphanumeric traffic a well as interoffice communications use the WAN.

3. MAJOR CHANGES IN THE OPERATIONAL RUNS UNDER UNIX:

3.1 Adoption of standards

Several standards are being adopted in the conversion to Unix necessitating considerable effort in conversion. These are:

- X-windows and Motif in order that CMC be able to run windowing applications across many platforms. Older software written in SunViews are being converted to X.
- Bourne shell for scripts, in order that they will able to run under Korn shell when System V release 4 is readily available. Csh is preferred for interactive sessions.
- ANSI Fortran and standard C. Code that used the "grey areas" of CDC and Cray Fortran Eg. OPENMS files and Hollerith will be converted. This adds considerably to the amount of conversion. Under Unix an attempt will be made not to use these "gray" areas.
- Network protocols such as X.25, TCP/IP, and NFS and FTP for files are being adopted.
- Internal files. Grid data will be kept in RPN standard files and observational data will be kept in Binary Universal Reports Protocol (BURP) files also developed in collaboration with RPN. A great amount of effort is being spent in the conversion to BURP.
- External distribution of grid data will be made in GRIB format.
- Graphic production uses a graphics utility called SIGMA developed by RPN from a modified NCAR graphics package. This utility produces METACODE which can be sent a utility called TRAMES, which, in turn, produces a device-specific raster image. CMC charts for outside consumption are converted into a compressed DIFAX code and sent on the Weather Service's satellite communication system for charts, METSIS. The conversion to SIGMA has also added considerably to the conversion load, since the present graphics production uses a different graphics utility.
- Reel Librarian for tape management. Early in 1992 an automated cartridge (3480)

library will be installed to manage the tape libraries.
- extended NQS (Network Queuing System) will be used on all machines for batch.

3.2 Graphics and alphanumeric production

The front end computers are extremely powerful and are well suited for the scalar processing of data, such as graphics. For this reason, operational graphics and alphanumeric bulletin production will be placed on the front end, in addition to the decoding of observational data. This is a major shift in the operational production. The supercomputer will be left to what it does best: highly vectorized compute-intensive code; i.e analyses and numerical models in operational and parallel runs, and research models.

3.3 The grid point data base

As a network resource, a large grid point data base on the front end or on the NEC will supply the needs of production, verification, case studies, and visualization graphics. A subset of this will be kept in the CMC Archives on tape.

The grid data will be kept in RPN standard file format, which uses a tagged record for a field. Fields can be searched by type, by level, and time with secondary search keys available. Subroutines (in C and Fortran) and utilities exist to select and process the standard files.

DB Structure: Model output is kept in files containing many fields. These files are stored by origin time and by a 6 hour forecast slice; analyses are kept by origin time. These are kept in "families" which use Unix pathnames to identify them. In general, one week of analyses will be kept online as well as all forecasts until they can be verified.

Utilities exist to identify what model output is current and to extract selections of the forecasts for further processing in other utilities. Forecasts received in GRIB from other centres (ECMF,EGRR,KWBC for now) are also decoded in RPN standard files and placed in the GRID Point DB.

Subsets of these families are archived. Once archived, files are removed and pared down in size as they become older.

The main production of graphics and alphanumeric data will operate from this data base. As a 6-hour forecast slice of the model arrives, jobs accessing this output will be initiated. This method prevents file contention and reduces the delays in production.

3.4 ADE and BURP

The automatic data extraction (ADE) system decodes incoming observational data and places them in BURP format. These are flexible, extendable word addressable files that are well suited in storing data at one point. They are efficient in the use of disk space as well as in the use of CPU. Conventions used are those of table B of BUFR.

The development of BURP is in-house (RPN/CMC), so that modifications to the code in the future will be done relatively easily. The conversion to BURP will prepare the infrastructure for real time data quality control.

The ADE under Unix has been completely overhauled, with many decoders having been re-written. Raw observational data as well as the decoded data (in BURP) will be kept on-line by the hour of the observation rather than the arrival time. Data quality control checks from the derivate (which prepares data for the O/A) and from the O/A itself will be kept. Trial field values for the observation can also be kept in the ADE DB.

At least a week of raw, decoded and treated (quality control) observational data will be kept online. They will also be kept in the CMC Archives.

3.5 GRIB messages

GRIB produced from the gridpoint data base will be sent to the different weather centres and military offices across Canada using the Wide Area Network (WAN). Jobs to produce these GRIB messages and to send them out via FTP to the regions are initiated as soon as the forecast data arrives in the data base. Two sets will be sent per region: a national GRIB from the Regional Finite Element (RFE) at lower resolution and a regional GRIB at the resolution of the model, over a much smaller window.

GRIB will also allow tailor-made model output messages to be sent to the regions as soon as they are available, since CMC will send the Grid Definition Block (GDB) as a rule.

CMC will also send GRIB of direct model output on the GTS. CMC will discontinue the production of its GRID and other alphanumeric messages containing fields.

3.6 Standardized charts

Most of the METSIS charts from CMC will have a standard window and legend that covers N. America, eastern Pacific and the N. Atlantic. The conversion to a standard window and legends will make future chart development easier.

4. MODEL CHANGES

4.1 Regional analysis

An analysis of the atmosphere will be made on the grid of the RFE with the RFE as input trial field. Presently the RFE uses the analyses from the global data assimilation cycle. Much detail in the vertical and the horizontal are lost through such a system. The regional analysis with greater horizontal and vertical resolution will allow the preservation of features and vertical structures of the RFE.

The regional analysis system will spin up from the global analysis system at T-12 (T=0,12 UTC). A first 6-hour RFE forecast will be starting at T-12 to be used by the T-6 analysis on the RFE grid. Another 6-hour RFE forecast from this analysis completes the spin-up. This trial field is then used by the analysis at time T. This approach incorporates all the features (i.e. long waves) of the global analysis system, and allows for higher resolution data in the core regions.

4.2 Geophysical fields analyses

New analyses of ice, snow depth and albedo will be produced. The ice analysis is based on satellite (SSM/I) data while the snow depth analysis uses conventional observations. Both these analyses are global and relax toward climatology in the absence of data. The albedo is calculated from climatology and from the snow depth analysis.

4.3 Climatological geophysical fields

Improvements in the quality and resolution of the geophysical fields will be implemented. The new climatology is a merge of fields obtained from the ECMWF, Bracknell and the USAF Environmental Technical Applications Centre.

4.4 Global assimilation system

Several changes will possibly be made to the global data assimilation system (use of satellite data for humidity, variational normal mode initialization). The resolution of the analysis will be increased as well as the number of observations used for the analysis at each grid point.

4.5 RFE

The regional model will be coupled with the same physics as the global model and will be run with the regional analyses. If possible, model resolution will be increased to 25-50km, with an increased vertical resolution.

4.6 SEF

The global model will be tuned (diffusion, upper boundary conditions) and the performance enhanced (reduced Gaussian grid toward the poles, efficient Legendre transforms, semi-Lagrangian optimizations) and have its resolution increased (T160?) with a corresponding increase in the vertical.

5. GRAPHICAL APPLICATIONS

Several workstation applications (developed by RPN and CMC) are of note:

XREC: An X-window application to view grid point fields. Fields are contoured and coloured on the fly, with zoom and pan and animation capabilities. Fields can be superposed on each other and on satellite images.

XMETAVIEW: Metacode generated by SIGMA can be viewed by this application. Zoom, pan and animation are possible.

ANIM/Carousel: Allows the display of raster images and the setting up of a Carousel of slides. Panning and animation is possible.

TEPHIGRAM: For the display of soundings.

METEOGRAM: For the display of meteograms.

FLOWCHART: An application to monitor the runs; jobs that bomb send messages to the screen.

DATAMON: An application to view the amount of observational data as it arrives.

Near Future:

XREC+: A new application that has the features of XREC will also have Operations, BOGUS and plotting added to it. Operations will allow addition, subtraction, scalar multiplication, and derivatives to be calculated on fields on the fly. BOGUS operations will allow the entry (through graphical screens) of bogus points into the data assimilation cycle. Plotted observational data will be displayed on request.

EDIGRAF: This is used for the graphical editing and production of weather depiction images on a workstation.

CHARTFLOW: All chart products from the operational runs will light up boxes as they are produced.