Workshop on System Collaboration

Météo-France, Toulouse, 13 to 15 May 2002

Report

Gerard Cats

1. Introduction

In the Memorandum of Understanding for the fifth HIRLAM project, signed in 1999, one of the instructions for the HIRLAM system manager reads:

> *During the course of the Project, the System Manager should undertake a comprehensive reappraisal of the entire suite of system codes and scripts (...) and devise a strategy for overall restructuring of the system.*

It was felt that an important target, when restructuring the system, should be to make the system better accessible to users, in particular those novel to HIRLAM. Conformance to international standards, conventions, and procedures would be a great advantage in this. One of the actions taken in the context of the quoted instruction to the HIRLAM system manager was to seek contact with sister organisations in Europe (through the SRNWP network), with ECMWF, and with the WRF project (references to web sites are in the references section). The aim of this action was to find out what standards *etc*. those organisations are using, and whether there are options to join forces in further system developments. This resulted in the organisation of a 'Workshop on System Collaboration', hosted by Météo-France from 13 to 15 May 2002. The purposes of the workshop were to investigate what subjects would be suitable for collaboration, to establish expert groups for those, and to investigate what external resources would be available for common actions. This document is a report of that workshop.

The workshop was in the (usual) format: a number of presentations, followed by discussions in working groups and in a plenary session. The programme of the workshop is in an appendix. In the presentations the organisations mentioned above gave an overview of their current system, and their plans for the near future. Furthermore, there are several projects that aim at facilitating collaboration between NWP-related systems. A number of these projects were also presented. The workshop programme includes a list of presentations.

In section 2 the presentations are summarised. The section consists of subsections, each addressing one issue within the context of system collaboration. A major motivation to collaborate is the sharing of knowledge and experience among the European (and world-wide) NWP experts. So this is the subject of the first subsection (2.1). When working on common codes, a good system for version control is required (2.2). Exchanging codes requires cross-database access methods (2.3). Many presentations mentioned the control of the complex suite of jobs in the NWP system; hence it is addressed in a separate subsection (2.4) event though it does not seem to be a real issue for cooperation. Several presentations addressed how to couple models or model components. A good mechanism for such coupling would greatly simplify the exchange of *e.g.* physics modules, and

hence this is a real issue for collaboration (2.5). GRID computing would further facilitate the exchange of software and data considerably: 2.6. A user sees a system through its user interface; sharing of knowledge is much supported if there would be some uniformity in user interfaces (2.7). In 2.8 there is a number of loose notes from the presentations.

The second part of the workshop consisted of discussions in working groups, to identify issues for collaboration. There were three working groups: on internal and on external code matters, and on user interfaces. The working group on external code matters reformulated its assignment to 'Modularity and interfaces to facilitate collaboration and exchange'. After the working group discussions, the participants convened in plenary discussions. During these, a number of conclusions of the workshop were reached, based on the discussions of the working groups. The working group and plenary discussions are summarised in section 3.

The fourth section gives my conclusions on the prospects of cooperation.

2. Presentations

2.1 The knowledge data base

An important consideration to try and act together is to share the knowledge that is present amongst the European researchers (the 'knowledge data base'). One of the main conclusions of the Aladin presentation was that Aladin users have to be taught how others use the system. One of the key drivers for UKMO to change their system is collaboration with universities and other national institutes. WRF has a similar motivation: it aims at establishing closer ties between operational and research institutes. ECMWF strongly support remote access to their systems, which contributes to 'knowledge sharing'.

2.2 Version control

For any collaborative work on software packages, a good version control system is required. This administration of software versions was presented as a serious issue in the Arpège/IFS cooperation. The merging of new versions costs 1 person month per half year. Problems of portability of Aladin are mainly in the maintenance of the variety of national versions, in particular when it comes to keep in phase with Arpège/IFS. (Portability across platforms is not a serious issue). The OLIVE project at Météo-France hopes to produce a better version control system to facilitate this 'back phasing': keeping local Arpège and Aladin options alive across new releases of Arpège/IFS.

Version control is based on the free *rcs* or on the commercial ClearCase. The latter is still often used for version control, but its users apparently tend to change to *rcs*, which is deemed good enough. (DWD is using VCS, a shell around *sccs*, but a switch to *rcs* is possible).

2.3 Data bases

DWD built their system successfully around Oracle; they expect some (trivial) porting work for IBM; overall they are very happy with Oracle. ECMWF has *mars*, which is now

accessible over the web. The WRF model assumes a local IO library is plugged in, to connect to GRIB, NetCDF, *etc*., to interface to local databases. UKMO uses a 'Met Database', in BUFR, with transparent decoding at retrieval. There are some drawbacks, *e.g*.: BUFR handling is inefficient; the format is less adequate for feedback during data assimilation. ECMWF's ODB is being considered for replacement.

2.4 Suite of jobs

NWP systems, be it operational or research, describe a complicated suite of jobs, to do various tasks as: setting up the run; compilations; collection and pre-processing of climate and model data and observations; data assimilation; forecast; dissemination; archiving; and validation. Several systems exist to control these suites. Almost all main European centres use the ECMWF Supervisor Monitor Scheduler SMS for operations. The (only?) exception is UKMO; they use their home-built Suite Control System SCS. Quite a few centres (Météo-France, DWD) maintain different suite control systems for operations (usually SMS) and for experimentation. One of the aims of the OLIVE project at Météo-France is to integrate the two suite control systems.

HIRLAM developed a subset of SMS (mini-SMS). It provides the facilities for job scheduling and monitoring as required by HIRLAM (and a few more). It has the advantage over ECMWF's SMS of being available for free, but it does not support full SMS functionality. Within the context of this document, SMS and mini-SMS are exchangeable.

2.5 Model coupling

PALM, PRISM: These projects try to develop the infrastructure to couple two or more models. Issues addressed are how to handle different horizontal, vertical, and temporal resolution of the models; whether to run the models as separate executables or by alternating invocation as subroutines. Performance is a serious consideration.

PALM should become a coupler that is able to run the models both as one, or as separate executables. An important feature of the PALM system is the ability to launch models dynamically, i.e. during the run. PALM offers an alternative approach to the classical monolithic structure of NWP code. The PRISM project aims to automate the coupling of models of different resolutions, based on meta-data describing the spatial and temporal layout of the coupling data. It will not provide dynamical model launching.

A major component of the UKMO project FLUME (to develop a *fl*exible *u*nified *m*odel *e*nvironment) is the design of the coupling to other models. It has been proposed to build upon PALM.

The WRF model has defined an internal data/memory layout that allows the coupling of a variety of dynamics, physics, and IO modules. This layout allows efficient execution on vector and parallel machines. The layout has been enforced on all contributed codes; hence it would be a major task to implement a different one if need ever would arise. As it is now, this strategy of requiring a suitable data layout has led to a system that supports modelling scales from LES to synoptic-scale, two-way nesting, Eulerian and semi-

Lagrangian schemes, IO in GRIB or NetCDF, *etc*. Direct maintenance of the data layout scheme would be an horrendous task, but it becomes manageable thanks to the use of a data registry that is converted into Fortran by Computer Aided Software Engineering (CASE) tools. IO can be done asynchronously (*i.e*. parallel to calculations); this gives very good performance improvements on multi-processor systems.

2.6 GRID computing

There were two presentations on GRID computing, each describing a EU-funded project. The DataGRID project aims at providing an infrastructure for large-scale computing, handling huge amounts of data. It has a work package on earth-observation. Some applications are now being tested using basic DataGRID services. DataGRID builds on the GLOBUS grid environment.

MeteoGRID (a work package of the EU EuroGRID project) aims at developing a version of DWD's LAM with adjustable domain, and provide 'weather prediction on demand' as an ASP solution. A GUI helps the user choose domain, resolution, forecast date, range, and products. Global topographical data are pre-processed at DWD, and DWD provides observational and lateral boundary data. The hardware to run the forecast on is provided by the 'EuroGRID HPC', which is a GRID compute environment based on UNICORE.

2.7 User interfaces

Many current projects have a strong activity in the development of (graphical) user interfaces. There are two main streams, one based on X (in particular Tcl/Tk), and the other on web technology. Those based on X in principle allow more direct interaction between the user and a process, and thus are more flexible than the web-based ones. But the advantage of web technology is its universal portability.

The user interface of (mini-)SMS is based on X. UKMO's SCS is based on Tcl/Tk. So virtually all centres in Europe use an X based system to monitor and control the suite of jobs. One of the targets of the OLIVE project is to develop a graphical user interface. To make sure it can be applied from any computer, in particular also from Windows systems, it will be based on web technology.

The same choice has been made rather earlier already by ECMWF. They now provide an extended service over the web, including access to their data archives and even to their compute resources.

Presentation of results is generally with ECMWF's Metview. ECMWF have integrated this presentation package with their web services, which is a particularly user-friendly method. However, most centres experienced problems with Metview installation.

2.8 Some remarks

The WRF project shows that is it possible to combine developments from many institutions into one. Operational implementation is targeted for 2005 (Michalakes, 2002). Considerable resources have been put into WRF to come here.

Météo-France has started the AROME project, which aims at the integration of their meso-scale model groups, for operational deployment within comparable time frames.

3. Working groups

After the presentations the participants split up into 3 'working groups', to discuss scope for collaboration on the following subjects:

1. Codes (Chair: John Michalakes; rapporteur: Kristian Mogensen)
2. Modularity and interfaces to facilitate collaboration and exchange (Chair: Stuart Bell; rapporteur: Nils Wedi)
3. User interfaces (Chair: Gerard Cats; rapporteur: Philippe Marguinaud)

The working group discussions were concluded with a plenary session, chaired by John Michalakes.

3.1 Codes

In this working group there were representatives from Aladin, HIRLAM, and UKMO. These consortia expect to (further) develop a non-hydrostatic model, with variational analysis (Aladin in parallel with AROME), and to invest heavily in system rewrites.

Portability is the primary requirement for collaboration. Scientific codes will be in Fortran-90. For portability, Fortran-90 code rules should be imposed (probably a revised form of the existing SRNWP rules); but enforcement at the level we see in WRF is beyond our reach. C may be used for IO related routines. There was no recommendation concerning the 'system language' (Perl, Python, …?).

Standardisation of namelists would much simplify the exchange and/or coupling of software modules. Due to model differences this would require some (modest) effort. Standardisation of ASCII 'codes' like namelists and other model specification data could perhaps be based on XML, but no Fortran-90 interface to XML is known to exist already.

For modularity, physics and dynamics should be independent. Memory should be managed dynamically, which would allow shorter argument lists. Loop lengths should be tuneable through namelists.

For version control, CVS is recommended because it is powerful, and free (CVS is a shell around *rcs* allowing concurrent developments).

Couplers should be considered, and so should the use of CASE tools, following the example set by WRF.
The following proposals for collaboration were made:

- Set up a common software structure and interface group within the SRNWP context, and have it meet regularly
- Standardise namelists, in particular to couple models
- Set up a newsgroup or web-based database with information about vendor tools, coupling tools, compilers, and so on.

3.2 Modularity and interfaces to facilitate collaboration and exchange

To identify possible subjects for collaboration the group drew the diagram on the right. It shows the processes in an NWP system and data streams between them. Two separate issues, job scheduling and version management, were also identified. The conclusions of the workgroup are based on an analysis of this diagram. There is no conceptual difference between the data flow for observations and climate data. Hence, in the following, only the former will be discussed; climate data are treated analogously.

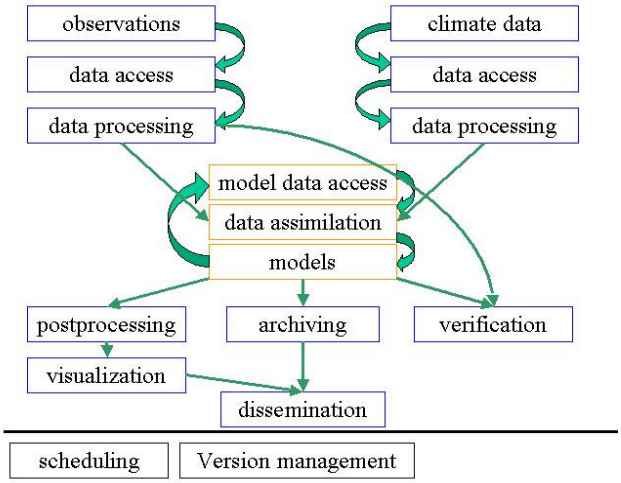

The data flowing between the various processes can be categorised in observation type and field type file; equally important are the meta-data associated with each file. Those meta-data describe data layout, model parameters like resolution, *etc*. In GRIB and BUFR messages some of the meta-data is part of the message itself but other meta-data are to be provided separately, *e.g.* in the form of GRIB and BUFR tables. Almost each NWP system has its own flavour of GRIB, with meta-data that are not always explicit.

The group made recommendations for collaboration on 6 subjects, with a possible time frame within which probably fruitful collaboration could become established:

i) Model data flow

To be able to exchange data, we need transformation tools for the various versions of GRIB, GRIB2, NetCDF. It is proposed that the varieties of versions, and the tools to process them, be collected. These should be made available through a common website and this information should be used to adopt local tools (*e.g.* Metview, IDL), to be able to process any of them. This can be done immediately, and it would facilitate model data access, visualization, post-processing, archiving, verification, and the development of models and data assimilation.

ii) Tools

Useful tools (scripts, makefiles, macros) should be made available, through a common website and/or a possibly restricted newsgroup. It can be done immediately and its goal is to create a general platform for collaboration and exchange.

iii) Observational data

Harmonisation of the data and access to data catalogues is considered useful. DataGRID may in the medium to long term facilitate universal access.

iv) Model configuration interfaces

Automatic interpretation of meta-data would be useful to be able to run another model with ones own set-up. PRISM interfaces may be built into the systems to achieve this, in the medium to long term.

v) Version management

There is no striking single software for this. Fortunately, the need for harmonisation is not strong because version management is considered useful only within the consortia or when coupling different model components. PRISM may help here, again in the medium to long term.

vi) Scheduling

Again no strong need for a common system was felt. Subgroups usually are happy with what they use now (even though some feel that SMS is the best currently available). GRID computing may, in the longer term, provide us with more user-friendly methods for scheduling.

3.3 User interfaces

There are essentially two streams of Graphical User Interface (GUI) developments: Tcl/Tk (SCS, PALM) and web-based (OLIVE, webmars, PrepIFS). There was not much incentive to make these streams converge. In particular UKMO (with SCS) and DWD (with NUMEX) seem to be happy with what they have now. DWD only missed the possibility to submit ECMWF jobs in full batch mode. Any future system should allow the users to do what they can do now. The first action in an initiative to harmonise user interfaces would be to make a list of user requirements. For this, it was suggested that use-case scenarios be developed to find out what functionality of the interfaces users really use. But the GUI developers of some consortia were so sure that they knew the requirements of their users already that they did not even want to compile a list of those requirements by development and analysis of use-case scenarios.

The existing GUIs may provide building blocks for a future common system, but probably only in bi-lateral exchanges, given the reluctance of some institutes to venture into European wide harmonisation. There exist methods to write GUIs (CASE tools), which may prove useful for (bi-lateral?) collaborative developments. The representatives from Météo-France and HIRLAM showed keen interest to cooperate (there was no representative from ECMWF, but later, in the plenary discussions, ECMWF also joined this interest group).

3.4 Plenary

Based on the summaries of the working groups, a number of fields for collaboration were identified:

- Develop and exchange utilities like GRIB to NetCDF converters, supporting all flavours of file formats that exist within the SRWNP countries. Adapt tools like Metview, IDL, to accept all of those. Make use of meta-data.
- Develop and exchange utilities like common scripts, makefiles, macros.
- Harmonise observational data access.

- Make codes as modular as possible, use CASE tools, and follow software conventions. Fortran-90 is the standard, but update the SRNWP coding rules.
- Standardise namelists.
- Use couplers (in particular, PRISM) and follow GRID developments.

Bi-lateral collaboration may be a good strategy where pan-European is too ambitious. This was most explicitly stated in the field of user interfaces. The WRF project proves that it is possible that diverse institutions join forces to develop a community weather model, but in Europe hardly anybody appeared ready for initiatives into this direction.

The meeting also identified how collaboration could be organised. The first suggestion is to set up a group of experts, consisting of one or two representatives of each consortium, which should meet regularly, to further work out the collaboration. A second suggestion was to set up a web site and a (possibly restricted) news group, to exchange software, data, meta-data, knowledge, *etc*.

Because I was worried about finding resources for these actions, in particular human resources, I proposed to seek EU funding, by having paragraphs on collaborative system developments in possible project proposals within the Framework 6 programme. For a start, it was suggested to ask the SRNWP coordinator to include such a paragraph in the Expression of Interest on Very High Resolution Modelling that he was about to write (in the mean time, this has been effectuated).

4. Conclusions

There are many developments, worldwide, that eventually could lead to considerable savings on the maintenance of the non-meteorological components of NWP systems. These include automatic model couplers (PALM, PRISM), automatic hardware selection (GRID computing), and standardisation of interfaces between software modules (*e.g.* XML for meta-data); but also the development of CASE tools, as *e.g.* used in WRF. Most of the system experts within the European consortia, however, do not believe that these developments will change their ways of working within the foreseeable future. In fact, quite a few believe that the current ways need not be abandoned soon, and are far from ready to cooperate. It is interesting to see that the consortia that have a longer history of international cooperation show a stronger incentive to further increase cooperation: ECMWF, HIRLAM, and Météo-France (the last with Aladin and HIRLAM experience). This may be an indication that the virtues of collaboration only become clear by doing it. Another reason for the two other consortia (COSMO, UKMO) to be apprehensive of increased collaborative activity is in their current staffing (DWD) and housing (UKMO) problems, which, they expect, will occupy their staff for some time. It will not be easy to discontinue the long history of concurrent, competitive developments in European NWP, and replace them by concert.

Yet, the workshop identified a number of fields where collaboration on systems could be fruitful. They are listed above, in the summary of the plenary discussions. In general, they appear to aim at facilitating exchange of data and software components. Only a few consortia (in particular those that believe in cooperation in general) believe that software components can be developed in common, so as to save on system development and

maintenance costs. Usually, savings need initial investments. It has been proposed to try and find the resources for those investments in the EU framework 6 programme.

Acknowledgments

In the report of the working group sessions, I made use of the reports written by the rapporteurs (Kristian Mogensen, Nils Wedi, and Philippe Marguinaud) and by Per Undén. The report of the plenary discussions is also based on mission reports of Stuart Bell, Filip Vana and Claude Fischer, and Emanuele Zala. I thank these people for letting me use their notes. I am grateful to John Michalakes and Sophie Valcke for their contributions through discussions during and in particular after the workshop.

References

DataGRID: http://www.eu-datagrid.org

MeteoGRID: http://www.eurogrid.org

Michalakes, John, 2002: Software Infrastructure for the Weather Research and Forecast Model. Presented at the U.S. Dept. of Defence HPC Users Group workshop, Austin, Texas, 10-14 June 2002.

PALM: http://www.cerfacs.fr/~palm

PRISM: http://prism.enes.org

SRNWP: http://srnwp.sma.ch

UNICORE: http://www.unicore.de and http://www.fz-juelich.de/unicoreplus. Developments started by Pallas GmbH, http://www.pallas.com

WRF: http://www.wrf-model.org

Appendix

Workshop Program

Monday 13 May

Session: Running Projects - Chairperson: Stuart Bell

start	*end*	*title*	*by*
09:45	10:00	Welcome to the workshop	Gerard Cats (KNMI)
10:00	10:30	The PRISM Project	Sophie Valcke (CERFACS)
10:30	11:00	DataGRID: a collaborative environment for scientific data analysis	John Van de Vegte (KNMI)
11:15	11:45	Overview and Status of the Weather Research and Forecast Model (WRF) Project	John Michalakes (NCAR)
11:45	12:15	AROME, an integrated mesoscale modelling environment	François Bouttier (Météo-France)

Session: Current achievements at operational centres - Chairperson: Per Undén

start	*end*	*title*	*by*
13:30	14:00	ECMWF, a web services provider?	Nils Wedi (ECMWF)
14:00	14:20	The massively nested Aladin configuration at DMN Casablanca	Radi Ajjaji and Zahra Sahlaoui (DMN - Casablanca)
14:20	14:40	The Met Office NWP Suite	Stuart Bell (Met Office, UK)
14:40	15:00	Operational suite and tools at MeteoSwiss	Emanuele Zala (MeteoSwiss)

Session: Distributed Projects - Chairperson: Filip Vana			
start	*end*	*title*	*by*
15:30	15:50	The HIRLAM Project	Gerard Cats (KNMI)
15:50	16:20	Operational model suite of DWD	Thomas Hanisch (DWD)
16:20	16:40	Review of ALADIN/LACE operational application	Filip Vana (CHMI)
16:40	17:00	ALADIN portability issues	Claude Fischer, Jean- François Geleyn, and Aladin team

Tuesday 14 May

Session: Design and interfaces - Chairperson: Nils Wedi			
start	*end*	*title*	*by*
09:30	10:00	Lessons from the design of a dynamic MPMD coupler	Samuel Buis & Etienne Gondet
10:00	10:30	Design and Implementation Issues and Challenges in WRF	John Michalakes (NCAR)
10:30	10:50	On the parallelization of NCEP's spectral model on IBM-SP3	Jean-François Estrade (Météo-France)

Session: Design and interfaces (continued) - Chairperson: François Bouttier			
start	*end*	*title*	*by*
11:10	11:30	Met Office Suite Control System	Jim Bolton
11:30	11:50	The OLIVE project in the SWAPP (*s*hared *w*eb *app*lications) framework	Eric Sevault
11:50	12:20	Meteo-GRID - Local Weather Forecast Using GRID-Computing	Claus-Jürgen Lenz
12:20	12:30	Introduction to Working Groups sessions	Gerard Cats (KNM)

Session: Working groups		
start	*end*	
13:30	15:30	WG1, WG2, WG3
16:00		WG1, WG2, WG3

Wednesday 15 May

Session: Outcomes from working groups - Chairperson: John Michalakes		
start	*end*	
09:00	09:20	Report of WG1
09:20	09:50	WG1 discussion
09:50	10:10	Report of WG2
10:10	10:40	WG2 discussion
11:00	11:20	Report of WG3
11:20	11:50	WG3 discussion
11:50	12:30	Discussion: follow-up of this workshop
12:30	12:45	Closure of the workshop