

The operational HIRLAM at the Finnish Meteorological Institute

Kalle Eerola
Finnish Meteorological Institute

1 Introduction

Significant changes have taken place in the operational HIRLAM system at the Finnish Meteorological Institute (FMI) since the previous status report (Eerola, 2002). The changes concern both the computational and the meteorological environment. On the computational side the co-operation with CSC - Scientific Computing Ltd. (CSC) has continued during the year, allowing FMI to use the most powerful computers in Finland. In December 2002 the Cray T3E supercomputer was replaced with a new supercomputer based on IBM technology. On meteorological side the new operational HIRLAM system, based on the HIRLAM version 5.1.4 and 3DVAR, has been tested thoroughly during the year and as a result of this work the new ATX suite was introduced into operational use in March 2003.

This report describes the operational HIRLAM environment at FMI. In chapter two, the new computational environment at CSC, relevant to HIRLAM, is described and the performance of HIRLAM in the IBM environment is discussed. Next, chapter three describes the current operational HIRLAM system at FMI from the meteorological point of view. Both the old system, based on HIRLAM 4.6.2, and the new system, based on HIRLAM 5.1.4, is described. Also, the local modifications are discussed. Chapter four describes some tools and products developed at FMI. Finally, in chapter five, the conclusions are drawn and some prospects of the future are given.

2 New technical environment

The replacement of the T3E parallel supercomputer at CSC with an IBM-based cluster was a major event, which affected very much all the HIRLAM-related work at FMI in the year 2002. The switch-off of the T3E was planned to take place in the beginning of December and by that time we should have our HIRLAM running on the new computer.

The new supercomputer at CSC is an IBM eServer Cluster 1600, which is a distributed memory supercomputer. It consists of 16 IBM p690 nodes, each equipped with 32 Power4 processors. So there are 512 processors in total. Nodes are connected to each other with a high speed Colony PCI switch providing the possibility to run parallel applications, which need fast communications between nodes. The system has a theoretical peak performance of 2.2 Tflop/s. The clock frequency of the Power4 processor is 1.1 GHz and the amount of memory in a node is 32 GB.

One node (32 processors and 32 GB of shared memory) is totally dedicated for FMI purposes, mainly for running the operational and experimental HIRLAM runs. This solution fulfils in the best way the reliability requirements and the need of top priority of the operational HIRLAM runs. The dedicated node with its own local disks can be seen as a separate computer and is independent of the other nodes. Also the batch job system has been optimized for HIRLAM environment allowing the top priority for the operational runs. One of the "normal" nodes is nominated as a backup node, where operational HIRLAM can be run in case the dedicated FMI node is down.

The operational HIRLAM version uses 30 processors in the 3DVAR analysis and in the forecast model inside one node. In both parts the MPI parallelization strategy is used. In co-operation with IBM we have optimized some of the parallelization routines in the forecast model and this has improved the efficiency by about 10%. These optimizations are not IBM-specific and they have been tested to be beneficial also on some other platforms. Also the IBM-specific mathematical libraries are used to improve the efficiency.

Concerning the parallelization efficiency of the forecast model, Figure 1 shows the scaling of the operational ATX forecast model (256x186x40 gridpoints), when the number of processors is increased from one to 30. The right bar shows the idealized perfect scaling, while the middle shows the scaling of the computations only, excluding the input/output. Finally, the left bar shows the scaling of the operational configuration, where a lot of output is produced with a frequency of one hour. The scaling of the computations is excellent. Even when increasing the number of processors from 24 to 30 the speed-up is about 97% of the perfect speed-up. The bottleneck is the input/output, which reduces the scalability: almost 40% of the time of the forecast model is taken by the input/output. In the current configuration the input/output is not parallelized. So a remarkable improvement in the scalability can be expected when applying the parallel output strategy suggested by Heikonen and Eerola (2003).

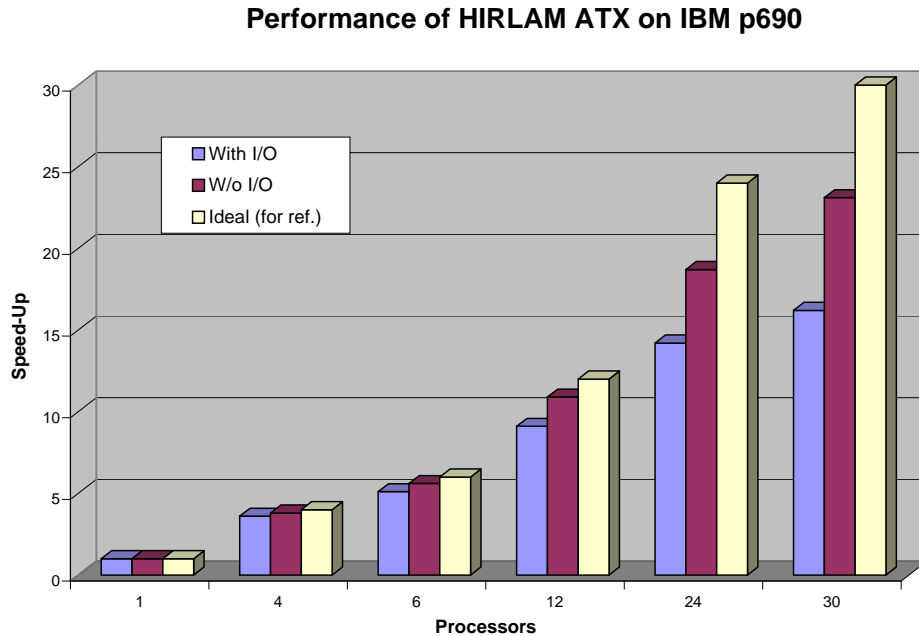


Figure 1: The scaling of the HIRLAM forecast model on IBM p690 with different number of processors. The right bar: perfect scaling, shown for reference. Middle bar: measured speed-up for computations only. Left bar: measured speed-up for computation and hourly I/O. The size of the problem is 256x186x40 gridpoints.

At the moment the whole operational ATX suite takes about half an hour of wall clock time. Of this time, about 20 minutes is taken by the forecast model. As stated earlier the bottleneck is the input/output. Another slow part of the HIRLAM system is the successive correction part of the surface analysis, which is not parallelized at all at the moment.

3 The meteorological environment

3.1 The ATA/ENO suite

The HIRLAM ATA/ENO suite, based on version 4.6.2, has been our main operational NWP routine since November 1999 until the mid-March 2003. It has been described in earlier Newsletters (Eerola, 2002, 2001) and therefore only the main points will be repeated here.

Shortly, the integration area of the Atlantic suite (ATA) contains Europe, the North Atlantic and parts of North America and the horizontal resolution is 0.4° . The number of vertical levels is 31. The European suite (ENO) is run on a smaller area containing mainly Northern Europe and the resolution is 0.2° . The horizontal areas of the ATA and ENO suites are shown with dashed lines in Figure 2 (outer and inner area). Both suites consist of OI analysis system based on HIRLAM version 5.0.0 and the forecast model based on version 4.6.2 with later updates.

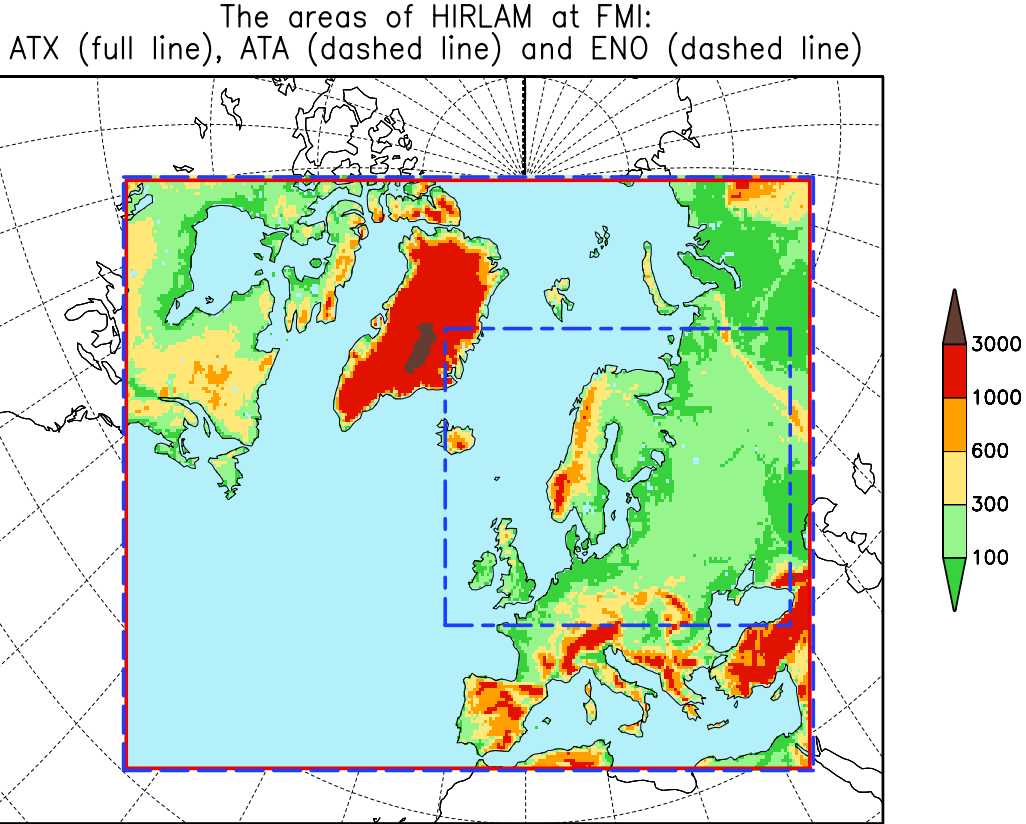


Figure 2: The areas of the different operational suites at FMI. ATA: outer dashed area, ATX: solid line, ENO: inner dashed line. The topography is from the ATX suite.

Both suites are run four times a day and the forecast length is 54 hours. This was the users' request to get an effective forecast length of two days in the end-products. The forecasts are available, both on model and pressure levels, at the frequency of one hour.

In summer 2002 it was decided that the ANA/ENO suite will be implemented on the IBM system, although the original idea was that the HIRLAM 5.1.4 will be the first model to be run on IBM. This decision allowed us a smooth transition from T3E to IBM and ensured that we have an operational HIRLAM system after the switch-off of the T3E even in the case IBM is not reliable enough for operational use, because ATA/ENO suite had already been installed on the SGI system at CSC. On the other hand, this gave us more time to test the meteorological behavior of the new HIRLAM 5.1.4. The implementation of the ATA/ENO suite was rather straight forward, and even the OI analysis system worked without modifications on IBM.

It appeared that implementation of the ATA/ENO suite on IBM was a good decision, because in December 2002 it was not possible from the technical point of view to implement the HIRLAM 5.1.4

system operationally on IBM, because there were problems in the local real-time data base at FMI and also problems in arranging the backup on the IBM system.

3.2 Testing the HIRLAM version 5.1.4

Tests with the HIRLAM 5.1.4 and Hirvda 5.0.3 started in spring 2002. The IBM computer was still at that time in the configuration phase and the progress in the beginning was rather slow.

The original idea was to run the new HIRLAM on an area equivalent the ATA area, but with the resolution of 0.18° in the horizontal and 40 levels in the vertical. After the first tests it appeared that the system was not stable, i.e. there were several crashes, which we did not understand. Also the meteorological fields (analysis and forecasts) were noisy for the human eye. The reasons for these problems have been partly understood later: errors in the implementation of the 3DVAR and problems in the boundary relaxation scheme. The latter topic is discussed by Järvenoja (2003b) in this volume. However, the noise, discovered by human eye, could not be understood or accepted. For this reason, the horizontal resolution was changed to 0.3° and the parallel runs with this system started in September 2002.

In January 2003 it was concluded that the new suite, now after some modifications called ATX, was from the meteorological point of view ready for operational implementation. Especially, the two-metre temperature forecasts were much more realistic compared to the ATA/ENO suite. However, there are still some problems left, especially the negative bias in the surface pressure forecasts. The results of the parallel runs are discussed by Järvenoja (2003a) in this volume. Due to technical problems on the backup arrangements at CSC and on the local on-line data-base at FMI, the new ATX suite became operational two months later, in mid-March 2003.

3.3 Features of the operational ATX suite

The operational ATX suite is based on HIRLAM version 5.1.4 and HIRLAM 3DVAR version 5.0.3, but there are some modifications to the reference system which are described later. The main features of the system are listed in Table 1. The horizontal resolution is 0.3° . The area together with the current topography is shown in Figure 2. It is very close to the ATA area, which allows the use of the same "frame" lateral boundaries from ECMWF in both suites. At the moment the frames are received in the resolution of 0.4° , but after the ATA suite will be discontinued, the resolution of the frames will be changed to 0.3° . The number of vertical levels is 40 and there are nine levels in the lowest one kilometer layer. The lowest level is about 30 m above the ground. The semi-Lagrangian advection and semi-implicit time integration allow a time step of 450 s. The forecasts are available both on the model and constant pressure levels to the forecasters and applications with the frequency of one hour. The forecast length is 54 hours and the data assimilation cycle six hours.

In data assimilation only conventional observations are used at the moment. The cut-off time of observations is at the moment 2h 30min. Digital filter is used for initialization.

3.4 Modifications to the reference HIRLAM 5.1.4

There are a few modifications to the reference HIRLAM system, which are worth mentioning. First of all, we use the Canadian MC2 boundary scheme, implemented in HIRLAM by McDonald (2003) and tested by Järvenoja (2003b). The basic difference is that in the original HIRLAM scheme the boundary relaxation is the last modification to the fields, i.e. the boundary relaxation is done in one step after the physical calculations. In the MC2 treatment the fields are relaxed towards the host model first time after the dynamics and then after the physics the physical tendencies are relaxed towards zero in the boundary zone. The experiments confirm that this prevents unrealistically large precipitation amounts close to the boundaries and therefore also prevents the model to blow up in some extreme cases.

In the climatological fields the fraction of sea ice was corrected by removing the artificial ice zone close to the coastlines. There are also small changes in the topography.

In the post processing we apply the HIRLAM 6.0.0 convention, where the near-surface variables are computed separately as a grid average and as an average over the land fractions. The user can choose, which one better suit to any specific application. The error in soil freezing over low vegetation was corrected.

In the data assimilation of sea surface temperature and ice fraction extra observations from the Baltic Sea are used and over Finland we apply climatological lake surface temperatures. These features have been installed also in the ATA and ENO suites at FMI.

Concerning technical differences, full SMS is used instead of miniSMS to control and schedule HIRLAM. The advantage of this is that the operators handle HIRLAM in the same way as all other routines, because SMS is used in all operational runs. We have re-written the SMS definition file to be suitable in our environment and written a few new scripts to control the different tasks. Figure 3 shows, how the ATX HIRLAM suite looks like in the XCDP window. A basic idea is that SMS knows very little of HIRLAM. It only knows, what the current date is, what the time schedule is, how to start different jobs of HIRLAM and what the relations between jobs are.

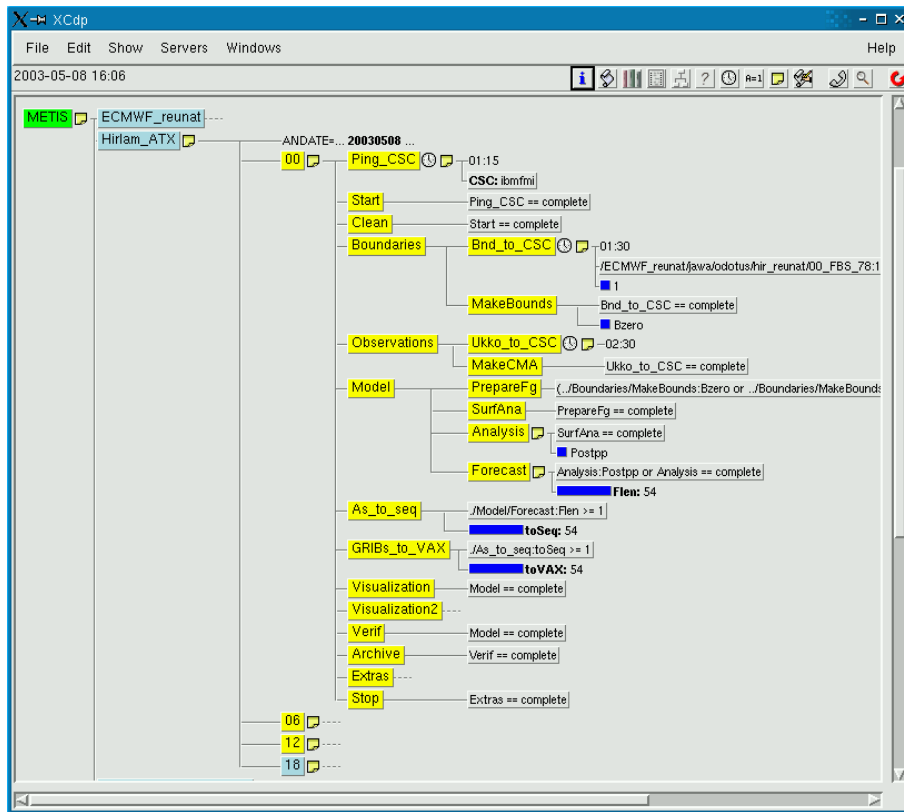


Figure 3: SMS structure of the HIRLAM ATX suite as seen by the XCDP program.

The more optimized parallelization routines for IBM were already discussed in chapter two. In the file naming, date, hour and forecast length are added in the analysis and forecast filenames. For example, the six hour forecast from the day 17.4.2003 at 00UTC can be found in the directory and file \$HL_DATA/20030417_00/fc20030417_00+006. This improves the readability and decreases the possibility of errors because of redundancy in the naming convention.

4 Tools and products developed at FMI

In the research environment the Grads graphical package, which is free software, is used extensively. Around it we have built quite an extensive set of tools, both the basic tools like to "mapping" the HIRLAM Asimof fields into Grads environment and application tools to draw different kind of products.

An example of an interesting application is shown in Figure 4. The predicted outgoing longwave radiation at the top of the atmosphere is shown with different shades of gray, representing cloud imagery. Superposed on it, the predicted surface pressure and the predicted precipitation intensity are shown. In this figure a beautiful family of cyclones can be seen. In the FMI intranet these predicted "pseudo satellite pictures" are shown as an animation of the predicted cloud, precipitation and surface pressure fields with the time resolution of one hour. This is one example of products, which originally was developed for research purposes, but can be used also in everyday forecast work.




Figure 4: An example of HIRLAM products: a predicted pseudo satellite picture. The outgoing longwave radiation at the top of the atmosphere (different gray shades) is shown together with the surface pressure and precipitation intensity.

5 Conclusions and prospects

During the latest year there have been big changes both in the technical environment and in the meteorological content of the operational HIRLAM system at FMI. On the technical side, the era of the T3E supercomputer ended in December 2002 and all HIRLAM forecasts are now run in the IBM environment. In the meteorological sense the new HIRLAM, based on version 5.1.4, was implemented as an operational system, containing many new features, like the variational data assimilation, digital filter initialization, semi-Lagrangian advection and ISBA surface scheme. Together they form the biggest change, which have taken place during one year in the history of the operational HIRLAM at FMI. All this work was harder and required more manpower than expected before, both in the technical and meteorological sense. However, we now have a new system running operationally on a new computer platform. The biggest improvements in the forecasts can be seen in the 2-metre temperatures forecasts.

The horizontal resolution of the ATX suite is 0.3° . At the moment there is an urgent task to get a nested higher resolution suite, of the order $0.1^\circ - 0.2^\circ$, inside the ATX suite. The first version of this higher resolution suite will be based on the HIRLAM 5.1.4. This work has the highest priority in the operational side at the moment.

In longer perspective, the higher resolution suite on the Atlantic scale will be studied. An important task in the data assimilation is to look at the new observation types. FMI has proposed to possibility to run the Reference Cycle Runs (RCR) as an operational system at FMI. This question is still open, but it gives new challenges to the FMI HIRLAM activities.

References

- Eerola, K., 2001: The operational HIRLAM at the Finnish Meteorological Institute. *Hirlam Newsletter*, **38**, 22–28.
- Eerola, K., 2002: The operational HIRLAM at the Finnish Meteorological Institute. *Hirlam Newsletter*, **41**, 19–24.
- Heikonen, J., and K. Eerola, 2003: Improving load balance in a weather code: Asynchronous output in HIRLAM with MPI. *Applied Parallel Computing: advanced scientific computing; 6th International Conference*, J. Fagerholm, J. Haataja, J. Järvinen, M. Lyly, P. Röback, and V. Savolainen, Eds., Springer, 567–577. Proceedings of PARA 2002.
- Järvenoja, S., 2003a: Experiences from pre-operational HIRLAM 5.1.4 test runs at FMI. *Hirlam Newsletter*, **43**, xx–xx.
- Järvenoja, S., 2003b: Testing of the MC2 boundary treatment in HIRLAM. *Hirlam Newsletter*, **43**, xx–xx.
- McDonald, A., 2003: Re-arranging the HIRLAM boundary relaxation treatment. *Hirlam Newsletter*, **43**, xx–xx.

Data assimilation	
Upper air analysis	3-dimensional variational data assimilation
Version	HIRVDA 5.0.3
Parameters	Surface pressure, temperature, wind components, humidity
Surface analysis	Separate analysis, consistent with the mosaic approach of surface/soil processes, for: <ul style="list-style-type: none"> • SST, fraction of ice • snow depth • screen level temperature and humidity • soil temperature and humidity at two layers
Grid length	0.3° in the horizontal
Integration domain	256 x 186 gridpoints in rotated lat/lon grid
Levels	40 hybrid levels defined by A and B.
Observation types	TEMP, PILOT, SYNOP, SHIP, BUOY, AIREP, AMDAR
First guess	6 hour forecast, six hour cycle
Initialization	Digital filter
Cut-off time	2 h 30 min.
Assimilation cycle	6 hour cycle

Forecast model	
Forecast model	Limited area gridpoint model
Version	Hirlam 5.1.4
Basic equations	Primitive equations
Independent variables	λ , θ (transformed lat-lon coordinates), η , t
Dependent variables	T , u , v , q , p_s , cloud water, turbulent kinetic energy
Discretization	Arakawa C grid
Grid length	0.3 ° in the horizontal
Integration domain	256 x 186 gridpoints in rotated lat/lon grid
Levels	40 hybrid levels defined by A and B
Integration scheme	Semi-Lagrangian, semi-implicit, $\tau=450$ s.
Orography	Hirlam physiographic data base, unsmoothed
Physical parameterization	<ul style="list-style-type: none"> • Savijärvi radiation scheme • Turbulence based on turbulent kinetic energy • STRACO condensation scheme • Surface fluxes using the drag formulation • Surface and soil processes using mosaic tiles • No gravity wave drag
Horizontal diffusion	Implicit fourth order
Forecast length	54 hours
Output frequency	1 hour
Boundaries	<ul style="list-style-type: none"> • “frame” boundaries from ECMWF • received four times a day • forecast interval six hours • linear interpolation between six hours

Table 1: Characteristic features of the ATX suite.