Report of DTC Activities from November 2010 - October 2011

Don Morton
Arctic Region Supercomputing Center
Universtiy of Alaska Fairbanks

Introduction

The focus of the 2010-2011 visitor's program project was to enhance the capabilities of the Alaska High Resolution Rapid Refresh prototype being run in one form or another at the Arctic Region Supercomputing Center since mid-2010. Forecasters in the Alaska Region NWS, as well as personnel in the RR / HRRR group at NOAA GSD/ESRL/AMD have been interested in this work, but up until this project, we had received no funding for these efforts. We saw this as a wonderful opportunity to pursue a few improvements in the modeling efforts, namely the incorporation of automated "verification" products using MET, which would be used by forecasters as a first step in detecting strengths and weaknesses of the model runs on a long-term basis. Our plan, following the deployment of verification tools, was to begin making changes to the basic model configurations, and to experiment with data assimilation with GSI in an effort to determine whether assimilating at the 3km resolution (given that our initialisation data was the heavily-assimilated 13km Rapid Refresh) would yield improved forecasts.

Although we had the opportunity to deploy MET operationally, as planned, and to work with model improvements and GSI as planned, radical shifts in ARSC infrastructure resulted in many unanticipated hours porting the entire modeling systems to three different supercomputers in less than a year, and we are only now just beginning to stabilise on one platform again. Additionally, with the new funding structure of ARSC, research faculty are now expected to support over 80% of their time, and this obviously decreases time spent in important, but non-funded activities. This unfortunately, severely limited the applicability of the new capabilities to our modeling environment, but we did succeed in building the infrastructure and continually seek new opportunities to improve this work on a funded basis.

The good news is that for the time being, ARSC and UAF see our weather modeling activities as being a potentially useful investment, and it has succeeded in bringing in a small pool of funding (DTC and Alaska NASA EPSCoR), and there is some potential for more. The funding provided in this project did help to leverage a small pool of funding from Alaska NASA EPSCoR which will help us to carry out our operational GSI aspirations in the next few months, and continues to open a number of doors, some of which will hopefully lead to sustained funding for these activities.

Operational "Verification" Products

The mandatory first step in the project, before implementing other model changes and experimenting with GSI on an operational basis, was to implement a "hands off" automated verification system. The design goal was to set up the infrastructure that would collect and process all of the raw observations and forecasts, as they became available, making them accessible for a wide range of analyses, over varying time-, space- and domain-scales.

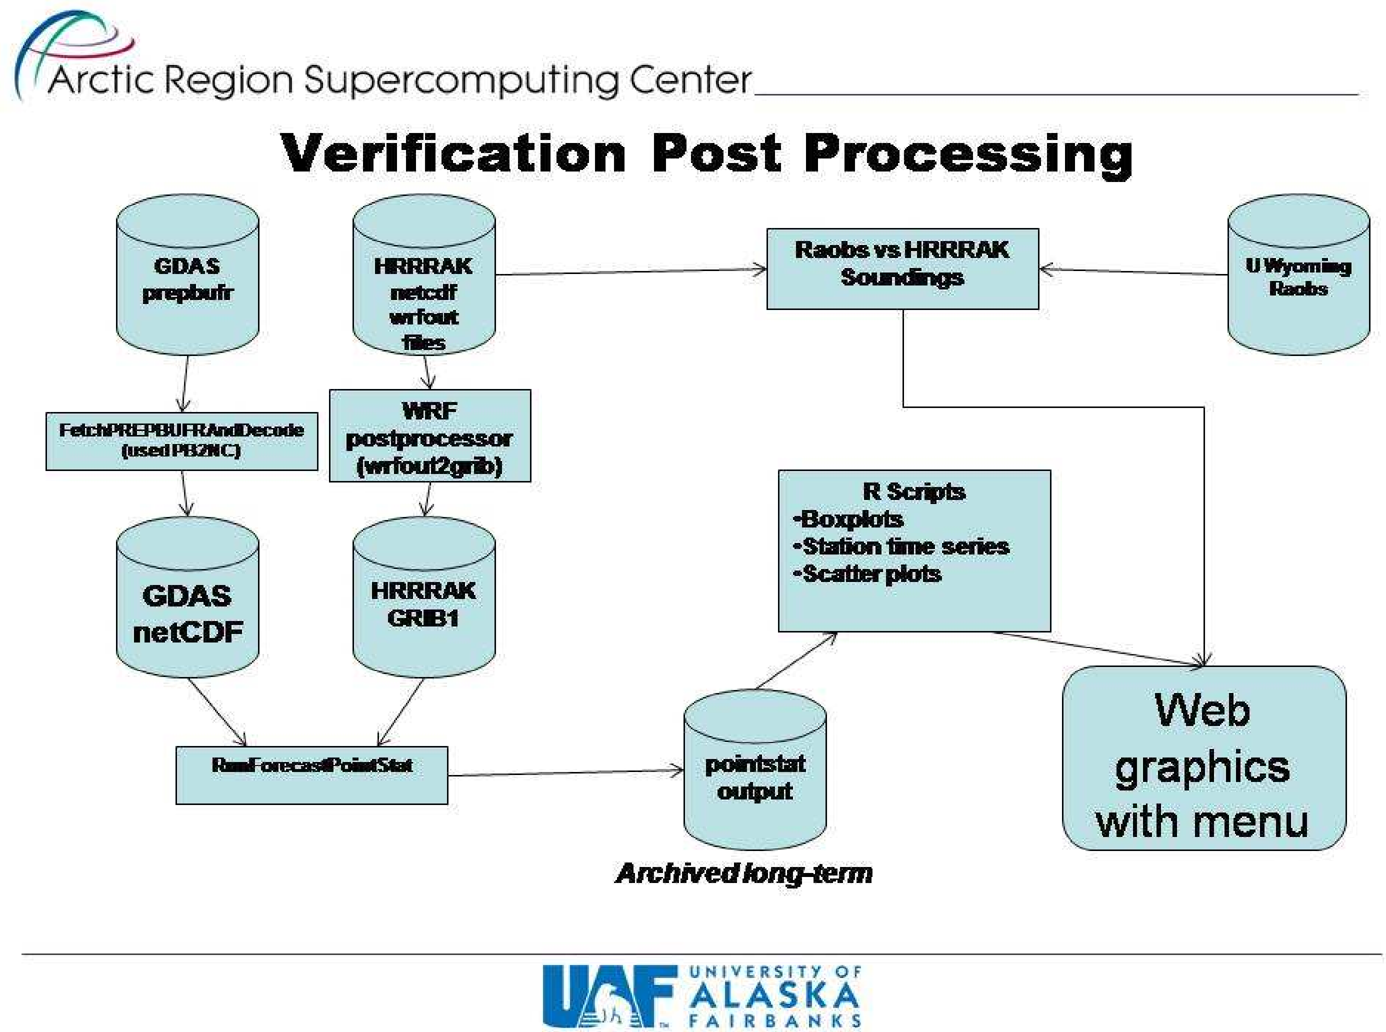

The above schematic crudely represents the operational system implemented during this period. Most of the components are driven by crontab entries. Every six hours GDAS PREPBUFR observations are downloaded and converted to netCDF with the MET pb2nc tool. Likewise, every six hours, as HRRRAK forecasts are produced the output is converted to GRIB1 (this is a separate process from our GRIB2 conversion for AWIPS) and stored. At a later time (fired by crontab), when it is assured that forecast and observations are all available, the MET tool point_stat is invoked to create files of matched pair forecast/observations, which are retained indefinitely[1]. These files, then, form the basis of future “verification” products. Although we have only created a few prototypes, it is our intent to retain these for long term, varied, rigorous analyses of the HRRRAK forecasts.

[1] See http://weather.arsc.edu/Operational/HRRR-AK/Products/MET/point_stat/

On a regular, post-forecast basis (when all observations and forecasts have become available) the data in the archived point_stat output is used to create simple graphics for providing a cursory overview of forecast skill. Though a number of graphical packages exist, after a lot of research, it was felt that R provided the most utility in producing the comparisons (box/whisker plots, scatter plots, etc.) that we were interested in. Hence, driven by crontab entries, plots are produced for each forecast, generated from the point_stat output files, to create box/whisker plots, scatter plots, time series for individual stations and, not directly related to MET, comparison of forecast versus actual vertical soundings. One area of weakness, as pointed out by one NWS SOO, is that the GDAS PREPBUFR files don't contain many of the local observation points of interest, so it will be necessary for us to obtains there elsewhere - MADIS, for example - and bring into the verification system via ASCII2NC.

The graphical products are presented through a simple interface[2] that allows users to look at a specified forecast for domain-wide box/whisker plot error display, domain-wide scatter plots, and time series observations versus forecasts at specified stations, in addition to the separately derived sounding comparisons[3].

And, the text point_stat products are available indefinitely, through a well-organized directory structure, making it possible to perform complex analyses over many forecasts with specific criteria, such as station locations, elevations and observed or forecast values (or even differences between such values). On a weekly basis, we assemble time series observation versus forecast at specified forecast hours for specified stations[4]. With simple scripts that access the point_stat files, we can also produce seasonal comparisons. An important area of future consideration is the ability to pick subsets of the point_stat output. For example, one of the Alaska SOO's has suggested the need for regime-based verification rather than domain-wide. A classic case arises in the deep winter inversions experienced in Interior Alaska. Rather than getting an overall feel for forecast temperature skill over the entire domain, it would be much more useful to make these comparisons for valley bottoms, whether defined by low elevations or a set of specific stations. With all of this information available in the archived point_stat output files, it's relatively trivial to make these kinds of analyses. Where we would really like to go is the user-driven specification of comparison parameters. The

---

[2] See http://weather.arsc.edu/Operational/HRRR-AK/Products/HTML/IndividualForecastError.html Please note that recent forecasts are missing, and right now only 00Z and 12Z forecasts are available

[3] See, for example, http://weather.arsc.edu/Operational/HRRR-AK/Products/SoundingVerify/2011/11/2011-11-09_12/PAFA/

[4] See, for example http://weather.arsc.edu/Operational/HRRR-AK/Products/HTML/MonthlyStationErrors.html

data is already there, and perhaps it is possible without too much investment to allow users to specify time periods and parameters to choose customised analyses.

MET is a very useful system of tools, designed in the true Unix spirit of building a collection of simple, single-purpose tools, which can then be combined in a myriad of ways to suit the users' tastes. Of course, the MET tools are far from simple, and there is a bit of a learning curve, but frankly, this probably can't be avoided. The tutorials are well-written and, as they should, take the novice user to a point where (s)he can begin to explore in more depth. Of greatest value to us is the ability to chain our own data-gathering tools and the MET preliminary conversion and processing tools to create a detailed set of matched pair records for each forecast. This, coupled with a little knowledge of R and HTML allows for a wide variety of "verification" products that we have only begun to explore. Of immediate interest is the ability to use the already-archived point_stat files to facilitate a web-based user-driven query over multiple forecasts within specified parameters. We had initially hoped that METViewer would facilitate this, but an admittedly brief introduction to it convinced me that it probably wasn't what we were looking for. Further discussion with MET developers, however, has convinced me that I should take a second look.

The other primary component in this project was to experiment with GSI in order to determine if assimilation of observations on our 3km HRRRAK initial conditions would be advantageous. As I look through my comprehensive research notes with GSI from the past year, I see a huge amount of time spent in just "getting it running" on large problems needing a large number of processors for memory and speedup. A problem common to both versions 2 and 3 has been error messages of the form

```
PGFIO-F-217/list-directed read/unit=999/attempt to read past end of file.
 File name = fsize_prepbufr0007    formatted, sequential access record = 1
 In source file read_obs.f90, at line number 22
```

in which, it turns out, the stated file is supposed to have an integer denoting the filesize, but in these error cases the value is 0. Unfortunately, this problem has often been coupled with large-domain problems in which it sometimes looks like memory allocation issues are the real culprit and, as a result, many hours have been spent on trial-and-error testing in order to get things finally working (or at least to understand the constraints under which the cases "will" run). For serial cases, I think it's safe to say that GSI runs "out of the box," but for parallel cases this hasn't been my experience. I do note that GSI user support has generally been very responsive to queries, but it's my opinion that the GSI community would benefit by having more support directed at making more accessible reports of potential problems, FAQ's, etc. I note that the current GSI support page says nothing about “Known Issues” in version 3.0, and only one in version 2.0. I fully recognize that everybody’s time is limited, but I do know that I have spent a great deal of time just trying to get GSI to work on systems, and others at other institutions share my sentiments, and maybe this is a useful problem to target?

Despite the problems, once solved, GSI seems to have an amazing amount of support for a variety of observations, and has allowed us to perform a couple of interesting, useful experiments. One experiment, presented at the AMS 2011 Annual Meeting (Petrescu 2011; Morton 2011) tried to determine (admittedly, in a single case) whether it might be more beneficial to cold start from the highly-assimilated Rapid Refresh runs, assimilate observations into the 3km HRRRAK initial conditions, or to initialise HRRRAK from previous HRRRAK runs, with data assimilation. In most cases it appeared that initialising directly from the Rapid Refresh produced the best results, but further investigation suggested that “maybe” in regions of highly complex terrain, such as Skagway, the model might intialise much better from its own previous forecasts (at 3km resolution) with new observations assimilated. More study needs to be performed here.

Finally, in collaboration with both DTC and the NASA Short-term Prediction and Research Transition (SPoRT) Center, we investigated the assimilation of NASA Atmospheric Infrared Sounder (AIRS) data (Divacarla 2006), representing hundreds of vertical profiles, into the HRRRAK. The sounding data comes in two forms

1. The raw, radiance data. GSI is set up to ingest this as-is, and some have recommended using this, due to its “original” non-processed format. Others, however, suggest that GSI will throw away any profiles with underlying clouds. In the Alaska region, these cloudy regions can fill a large part of the problem domain.
2. Processed data in the form of vertical temperature and moisture profiles. This data is processed by JPL and in many cases, in conjunction with the Advanced Micrcowave Sounding Unit (AMSU), can derive vertical soundings above the cloud layers, providing a more complete picture. In its preliminary stages, NASA SPoRT is making the data available by (and, recognizing the shortcomings of the approach!) treating it as radiosondes, and simply appending to existing GDAS and NDAS PREPBUFR files. We tried both approaches (Morton and Harrison 2011), and overall findings are summarized in the following graphic. In the first row, we are looking at differences (in temperature) between a non-assimilated model and a model assimilated with basic surface observations, at Forecast Hours 0, 12 and 24. In the second row we are looking at the temperature differences resulting from the assimilation of processed AIRS profiles, and in the third row we are looking at the temperature differences resulting from the assimilation of the raw radiance data. At first glance it appears that the use of the processed AIRS profiles might be adding great value to the forecasts, but deeper analysis revealed degraded forecasts. We had an opportunity to visit the NASA SPoRT office in Huntsville while attending

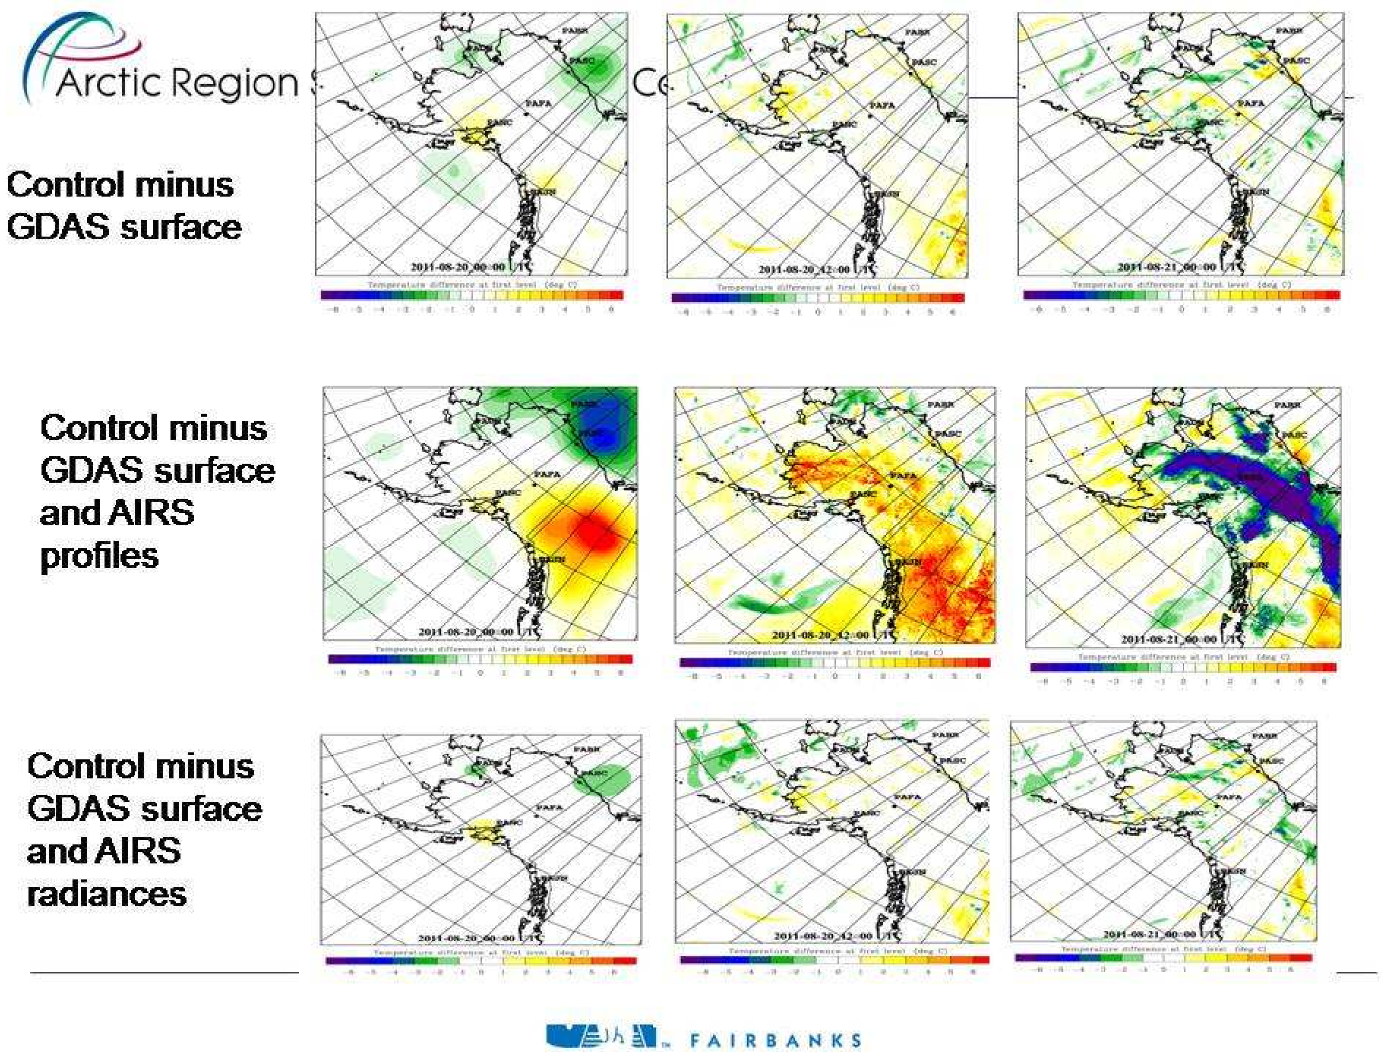

NWA in October 2011, and found that they had been experiencing some of the same problems in other cases. A number of ideas sprung from the meeting, and will be pursued in the upcoming months.

Finally, several very fruitful meetings with the NOAA ESRL/GDS/AMD RR and HRRR team members resulted in our making a number of incremental modifications to the HRRRAK over the course of the year. Although a rigorous analysis of the effects of these changes has not yet been pursued, we fortunately have all of the point_stat files produced from each forecast and, in the future, will be able to go back and see if we can detect any interesting changes appearing at the time of model changes.

<u>Closing</u>

Time spent at DTC was a truly rewarding experience and one of the highlights of my professional career. As an interdisciplinary scientist with one small set of skills, and lacking in so many others, it was enlightening to be exposed to the NCAR environment where similar people are working hard together on important problems in the atmospheric sciences. If I wasn't so in love with the Interior Alaska lifestyle, I would be applying for full-time work at NCAR!

I had the opportunity to meet a number of wonderful people at DTC, NOAA, and elsewhere in the NCAR system to discuss various issues in WRF, trying to arrange international collaborations, etc.. Since this report will be public I hesitate to mention specifics, but there are some very interesting collaborative possibilities, some leading to possible funding, and some which are fascinating and I regret that I simply didn't have the time to pursue further.

References

Divakarla, M.G., C.D. Barnet, M.D. Goldberg, L.M. McMillin, E. Maddy, W. Wolf, L. Zhou, and X. Liu, 2006: Validation of Atmospheric Infrared Sounder temperature and water vapor retrievals with matched radiosonde measurements and forecasts. *Journal of Geophysical Research*, **111**, D09S15, doi:10.1029/2005JD006116.

Morton, D., K. Harrison, B. Zavodsky, G. Jedlovec, "Assimilation of Atmospheric Infrared Sounder (AIRS) Data into the Prototype High Resolution Rapid Refresh for Alaska (HRRRAK)" poster presentation, *National Weather Association 36th Annual Meeting*, 15-20 October, 2011.

Morton, D., E. Petrescu, C. Dierking, "The High Resolution Rapid Refresh for Alaska (HRRRAK)" *The Alaska Weather Symposium*, Fairbanks, AK, 15-16 March 2011.

Petrescu, E., C. Dierking, D. Morton, "A prototype implementantion of the high resolution rapid refresh for Alaska," 91st American Meteorological Society Annual Meeting, Seattle, WA, 23-27 January 2011.