BC Clean Air Research Fund

BlueSky Forest-Fire Smoke Forecast Enhancements

Final Report

April 1, 2011 to March 31, 2012

April 15, 2012

Principal Investigator: Prof. Roland Stull
Department of Earth, Ocean & Atmos. Sciences
University of British Columbia
6339 Stores Rd., Vancouver, BC, V6T 1Z4
604-822-5901, FAX: 604-822-6088, rstull@eos.ubc.ca

Report written by George Hicks II
604-822-4760, ghicks@eos.ubc.ca

PROJECT OVERVIEW

Abstract

BlueSky predicts the spread and concentration of smoke (as PM2.5) from wildfires in Western Canada and the bounding regions. Five goals were outlined in this project to improve performance, reliability, and forecast quality.

1) Purchase, deploy, and test a new high performance computer server at UBC.
2) Devise a method for handling smoke carryover from previous runs or creating extended forecasts.
3) Test persistence of wildfire burns to improve latter forecast particulate levels.
4) Implement and test integration of the HYSPLIT model v4.9 into the framework. Determine if it affects the 'bulls-eye' problem.
   a) Extend Domain for the Trans-border Effects
   b) Enhancements: Incorporate and test new nested grid codes provided by STI
5) Conduct daily real-time tests of the enhancements during the 2011 fire season.

Results from this project included the following: The new computing server and enhancements were successfully implemented for the 2011 fire season. Fire persistence showed improvement for longer term fires allowing for more smoke presence in the latter portion of the forecast. A smoke carryover approach was eschewed in favor of secondary daily forecasts allowing for a daytime forecast taking advantage of up-to-date fire information for rapidly changing conditions. Finally, HYSPLIT v4.9 was integrated into the framework and initial testing conducted. Results require further evaluation but it is clear that v4.9 does not solve but did reduce the bulls-eye problem.

The success of the forecasts in the 2011 fire season drew further attention from the eastern provinces and interest in expanding the forecast to cover more of Canada. As of Spring 2012, provincial governments from AB to ON are working together to provide funding for operational runs of BlueSky during the 2012 fire season.

Accommodating a larger domain will require improvement of the forecast performance and expansion of the meteorological domain. From a performance perspective, development work is being conducted on integrating SmartFire2 into a framework that would allow for improved fire information with the ability to cluster and reduce fire inputs (hopefully speeding up HYSPLIT). In addition, work will continue with HYSPLIT v4.9 to determine if it can provide stable and reasonable operating forecasts under heavy fire conditions.

FINANCIAL OVERVIEW

Revenue Description

Table 1 Projected Total Project Revenue (cash and in-kind)

Organization	2010/11		2011/12		Total
	Cash	In-kind	Cash	In-kind	
BC CLEAR - Fraser Basin Council	$0	$0	$25,000	$0	$25,000
BC Ministry of Health Services	$19,955	$0	$3,445	$0	$23,400
BC Ministry of Environment	$0	$18,750	$0	$0	$18,750
AB Department of Environment	$0	$25,000	$0	$0	$25,000
University of British Columbia	$0	$0	$0	$7,500	$7,500
TOTAL	$19,955	$43,750	$28,445	$7,500	$99,650

Table 2 Actual Revenue for Reporting Period (cash and in-kind)

Organization	2010/11		2011/12		Total
	Cash	In-kind	Cash	In-kind	
BC CLEAR - Fraser Basin Council	$0	$0	$25,000	$0	$25,000
BC Ministry of Health Services	$12,000 (a)	$0	(b)	$0	$12,000
BC Ministry of Environment	$0	(c)	$3,445	$0	$3,445
AB Department of Environment	$0	(d)	$0	$0	
University of British Columbia	$0	$0	$0	$7,500	$7,500
TOTAL	$12,000		$28,445	$7,500	$47,945

Note: Please attach copies of letters or agreements confirming additional funds.

Please explain revenue discrepancies (if any)

(a) **Cash** budget for contributions from the BC Ministry of Environment (formally Health Services).

Item	**Planned Funds to UBC**	**Planned Funds for subcontracts**	**Actual Funds to UBC**
1) Labour for George Hicks II to install new computer at UBC	10h @ $53/h = $530		10h @ $53/h = $530
2) Enhance BlueSky to include carryover "old" smoke from previous days. Labour by George Hicks II (See Results Overview)	115h @ $53/h = $6,095		106.42h @ $53/h = $5,640
3) Test burn persistence option with new computer, explore cause of "bulls-eye" artifacts in the forecast. Labour by George Hicks II.	10h @ $53/h = $530	$7,500 contract from BC MoE to Sonoma Tech Inc. (STI)	10h @ $53/h = $530 (e)
4) Enhance BlueSky to include effects of smoke blowing into the BC/AB domain from outside. Labour by George Hicks II.	100h @ $53/h = $5,300		100h @ $53/h = $5,300
7) Terminate daily runs. Archive results. Write Users Manual & Final Report. Report on reliability. Start implement-ing future improvements.	65h @ 53/h = $3,445.		65h @ 53/h = $3,445.
TOTAL	**$15,900**	**$7,500**	**$15,445**

(b) BC MoHS funding of $3,445 for fiscal 2011-2012 arrived from BC MoE instead.
(c) Computer purchased by BC MoE and delivered to UBC. Cost unknown.
(d) Smoke blow-in work done by STI as contracted. Cost unknown.
(e) BC MoE modified their agreement due to reduced funding. The support to UBC for a subcontract to STI was canceled, hence the "bulls-eye" reduction work was not done.

Expenses Description

Table 3 Projected Expenses for Reporting Period (cash and in-kind)

Project Costs	**Expenses**		
	All Sources		
	Cash	*In-kind*	*Total*
Salaries and fees	$23,400	$0	$23,400
Travel and accommodation	$0	$0	$0
Equipment and supplies	$0	$18,750	$18,750
Communications and outreach	$0	$0	$0
Analysis	$25,000	$32,500	$57,500
TOTAL PROJECT COSTS	$48,400	$51,250	$99,650

Table 4 Actual Expenses for Reporting Period (cash and in-kind)

Project Costs	**Expenses**		
	All Sources		
	Cash	*In-kind*	*Total*
Salaries and fees	$12,356 salary + $5,965.10 fees (h)	$7,500	$25,821.10
Travel and accommodation	$0	$0	$0
Equipment and supplies			(unknown, see note c above)
Communications and outreach	$0	$0	$0
Analysis	$22,123.90	$0	$22,123.90
TOTAL PROJECT COSTS	$40,445(i)	$7,500	$47,945

Please explain expense discrepancies (if any)

(g) Totals are for UBC only. The other revenues and costs from BC MoE and Alberta Environment are unknown. We recommend you contact Steve Sakiyama of BC MoE (Steve.Sakiyama@gov.bc.ca) for these numbers.
(h) Fees are the UBC overhead (13% for CLEAR, 25% for BC MoE)
(i) The only cash received at UBC for BlueSky was $25,000 from CLEAR and $15,445 from BC MoE for a total of $40,445.

RESULTS OVERVIEW

Activity Description

Table 5 Summary of Activities for the Reporting Period

Activity*	Completion Date	Description of Results
1. Install new computer at UBC	Feb 27, 2011	Computer installed with Ubuntu 10.10 operating system.
2. Experiment with extended 2-3 day BlueSky forecasts, smoke carryover.	Incomplete	See item 2 below.
3. Test burn persistence.	April 28, 2011	See item 3 below.
4. Modify the MM5 numerical weather prediction model domains.	Dec 31, 2010	See item 4 below.
5. Conduct real time tests of BlueSky enhancements.	Nov 14, 2011	See item 5 below.
6. Archive forecasts, terminate runs, and report reliability.	Nov 14, 2011	See item 6 below.
7. Explore implementation of HYSPLIT v4.9	Dec 31, 2011	See item 7 below.

*As outlined in the project contribution agreement or contract.

Please explain activity discrepancies (if any)

2. As the hardware and software changes to HYSPLIT were not complete until late April, there was insufficient time prior to the fire season to test smoke carryover. Stakeholders determined that an alternative to smoke carryover would be to do a secondary run everyday at noon PDT with the latest fire information and extended original meteorological forecast. This was successfully implemented. Evaluation of the system shows that the secondary forecast can improve the latter portion of the forecast in cases with rapidly changing fire conditions giving a kind of carryover improvement.

3. Persistence burning was tested and incorporated into the operational framework. It did not solve the 'bulls-eye' problem as anticipated, but did show improvement of the forecast quality in cases with larger persistent fire events. The computational forecast performance was balanced by hardware improvements.

   Burn persistence was tested for the period of August 1-29, 2010. Persistence showed improved forecast output for the latter 24 hours of the forecast period while doubling computational expense. It was determined that with the newer hardware, the HYSPLIT model could still be run fast enough to utilize burn persistence for an operational forecast. As such, it was included in the operational runs for the 2011 fire season.

4. In the 2010 fire season, an MM5 weather forecast domain with 4km grid spacing was used to drive the HYSPLIT dispersion model. This domain covered most of BC and

Alberta. Due to software constraints, it was not possible to extend this domain to cover more area. To accommodate additional interest in smoke forecasting outside the 4km grid-spacing domain, it was determined that the MM5 domain with 12km grid spacing could be used. This domain was increased and extended north into the Yukon and Northwest Territories, south into the northwestern United States, and east into Ontario.

MM5 forecasts weather using a nested grid approach. It begins with a coarse grid that is fed on analysis and observations. To achieve high resolution forecasts, subsequent finer spaced grids are nested within the coarse grid and each other to achieve a spatially fine forecast. All these grids are computed together and feed information back and forth between each other. As a result, to enlarge the 12km domain as outlined, it was necessary to increase the 36km domain. Additionally, a change was made to use a richer initialization data set to begin the 36km domain. All changes were made operational starting December 31, 2010.

5. Daily BlueSky runs were conducted operationally each day from April 29 to November 14, 2011. Output was made available on an ftp server for Dr. Michael Brauer's team and the BC Ministry of Environment. Documentation in the form of a users guide was completed February 28, 2012 (see appendix 1).

6. BlueSky forecast output was archived locally on two computer servers for the duration of the April 29 to November 14 season. In addition, output for testing runs was also archived.

   The season comprising April 29 to November 14 had 338 possible forecasts (accounting for the daily forecast plus secondary forecasts beginning May 31). Of these, 21 forecasts were not completed, resulting in 9 days without at least one forecast. This yielded a 94% reliability for total forecasts and 95% reliability for days with at least one forecast.

   Of the nine days, seven occurred in May (5, 6, 9, 10, 13, 16, 17) due to a bug in the BlueSky framework resulting from the upgrades put in place prior to April 29. A fix was put in place on May 19 after which, 14 forecasts were missed. Nine were due to problems occurring in the Canadian Wildfire Information System (CWFIS) which we rely upon for fire information. Three were pre-empted due to testing. The remaining two were days with no fire activity.

   If we discount the first seven missed forecasts due to the bug, and the two missed forecasts due to a lack of fire activity, then the system only missed 12 forecasts and only 1 day without any forecast giving a 96% reliability for total forecasts and 99% reliability for days with at least one forecast. This information was communicated to stakeholders via email and conference call.

   In addition to system reliability, Dr. Brauer's group and the Alberta Ministry of Environment have been conducting separate verification of the forecasts using a variety of quantitative and qualitative methods. We have assisted Angela Yao (UBC) and Valerie Klikach (ABDoE) by providing data, performing additional testing, and assisting in interpreting results. Ms. Yao's results were presented at the Pacific NorthWest International Section of the Air and Waste Management Association 2011 Conference. Ms. Klikach's report is forthcoming.

7. The more recent version of HYSPLIT (v4.9) includes significant model changes over the version utilized in the BlueSky Framework (v4.8) with more sophisticated puff and particle dispersion algorithms. A copy of a v4.9 executable compiled on older Ubuntu v8 Intel architecture machines was supplied by STI. A couple of tests were conducted for Aug 18, 2010 using this executable. Results were significantly different, somewhat better than v4.8 when compared with MODIS data, but runtimes were at least an order of magnitude longer. It did appear to reduce 'bulls-eyes'. Initial attempts to compile the code locally with an older code repository provided by STI were unsuccessful.

   In August 2011, we were able to obtain the full v4.9 source code from NOAA Atmospheric Research Laboratory (ARL) and began having successful local compilations. However, initial tests were not satisfactory with very low emissions output. Initial investigation showed that the internal dynamics and settings for the model had changed and the BlueSky Framework would have to be modified to accommodate. After much research and testing, it turned out that the problem did not lie with HYSPLIT but rather a fix that had been put in place to correct erroneous emissions in v4.8. Once corrected, emission levels appeared closer to expected levels, and output could be evaluated. The correction was put in place for testing beginning Sept 14, 2011. Ex post facto testing has continued using previous high fire dates (July 24, 2010, Aug 18-20, 2010) and current daily tests (Sept 14, 2011).

   Overall runtimes using the locally compiled model are better in performance than the previously supplied binary from STI, however, they were still many times longer than runs with v4.8. As it turned out, more recent revisions to the v4.9 code allow for parallel processing using MPI (message passing interface). By taking advantage of the MPI compiled code, we were able to run the model using 10 processor cores instead of 1. As is typical with parallelized code, this did not result in linear scaling (HYSPLIT was not 10 times faster) but it was fast enough to make the overall runtimes close to the original v4.8 runtimes.

   As such, we have also begun examining increasing the number of vertical layers relative to the number that was required to run v4.8. Tests using 20 vertical levels show increased complexity in model output, but increase overall runtime by a factor of 3. This has also shown a reduction in near surface particulate levels which may be realistic but could be hard to verify using MODIS data. Output has been passed along to stakeholders for verification. Evaluation is ongoing.

Deliverable Description

Please include copies of all deliverables with the final report (e.g. publications, presentations, research reports, etc.). The final report will be considered incomplete without copies of the project deliverables.

Table 6 Summary of Key Deliverable Accomplishments for the Reporting period

Deliverable*	Description	Description of Results
1. Complete real-time tests of BlueSky enhancements	Daily forecast runs performed from April 1, 2011 to Sept 20, 2011.	See (1) below
2. Summarize outcome of results.	Archive results. Create a users manual. Report on forecast reliability.	See (2) below.
3. Final Report	Final Report to BC CLEAR	This document.

*As outlined in the project contribution agreement or contract.

Please explain deliverables discrepancies (if any)

1. Daily forecast runs including the enhancements began on April 29, 2011. This was the earliest day we could get complete successful runs due to the late installation of the new hardware and software improvements from STI. The forecasts continued through November 14, 2011, well past the designated fire season. This was due to a combination of an unusually inactive fire season and late fire activity in October.

   During April 29 to November 14, 2011, a total of 317 successful forecast runs were conducted. The addition of the second daily forecast run each day (to handle extended smoke forecasts) began on May 31, 2011 continuing through to November 14, 2011. Forecasts did not complete for May 5, May 9, May 10, June 4, June 5, June 19, June 20, August 16, August 31, and September 13. With the exception of the difficulties in May, the other forecasts were interrupted by unavoidable problems obtaining data from CWFIS. Overall, 95% of all forecasts successfully completed.

2. Results of the daily forecasts were archived each day onto two different computer servers. In addition, these results were made available for download via ftp to all stakeholders. A summary of the fire season results and forecast reliability was presented to the stakeholders during a conference call on November 3, 2011 and made available via email.

   A users' manual was completed on February 28, 2012. A copy was previously forwarded to BC CLEAR but is also attached to this document.

DELIVERABLES

a) Real-time smoke concentration forecasts were sent via ftp for display on the BC MoE public website, http://www.bcairquality.ca/bluesky

b) Also delivered and archived at UBC, are kmz files to display smoke concentration forecasts using Google Earth and similar GIS programs. A sample from one of the forecasts this summer is attached here.

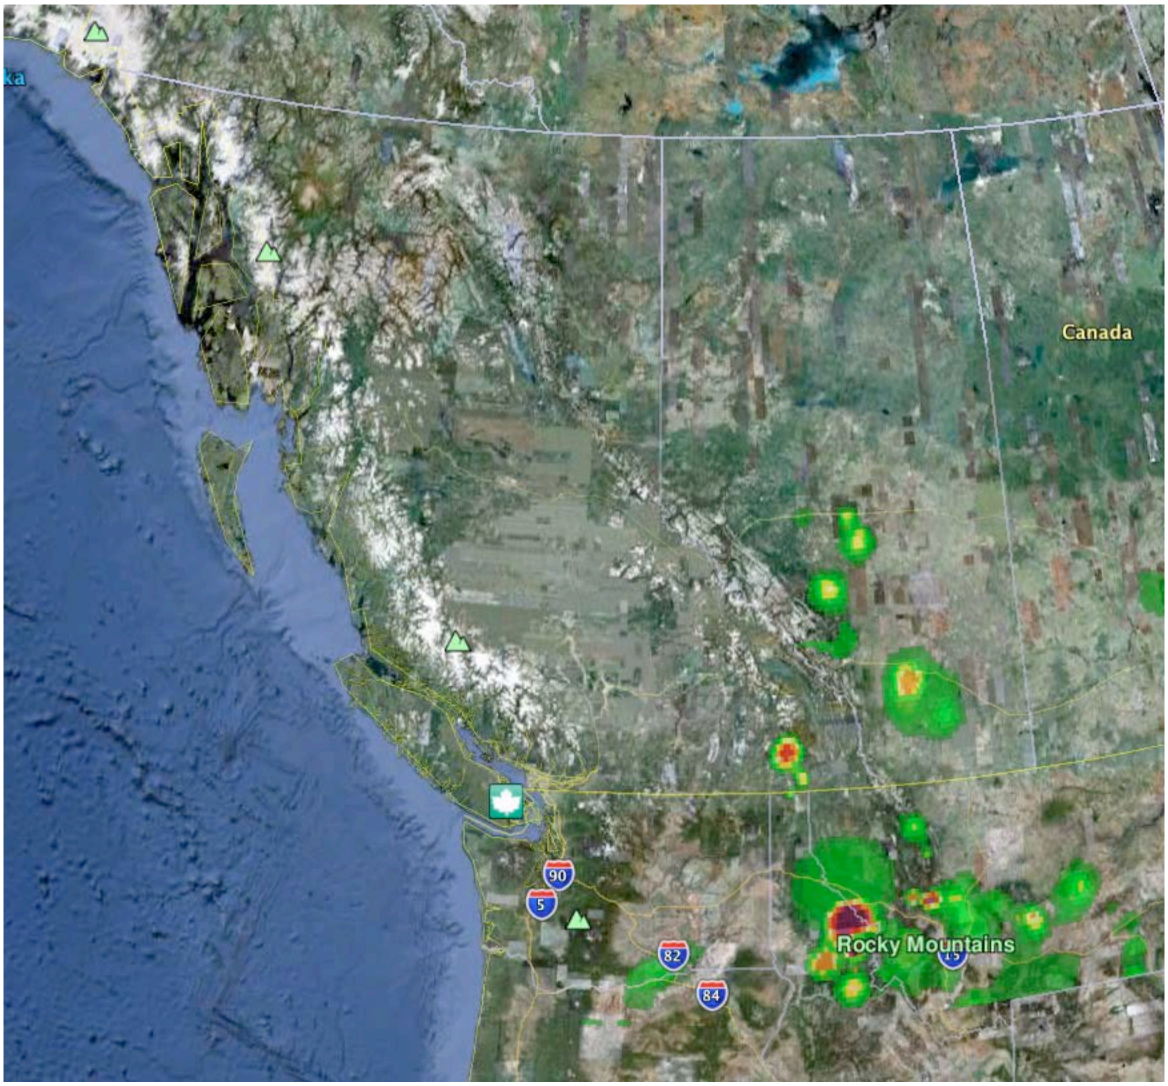

Fig. 1. Google Earth display of kmz smoke concentration forecast valid at 8am on September 15, 2011, as produced by the BlueSky project at UBC.

• **Appendix 1:** ***UBC BlueSky Users' Manual.*** *(see separate attachment)*
• **Appendix 2:** ***Pacific NorthWest International Section of the Air and Waste Management Association 2011 Conference Presentation.*** *(see separate attachment)*

Canadian BlueSky Framework Users' Guide

prepared by George Hicks II
University of British Columbia
ghicks@eos.ubc.ca
604-822-4760
February 28, 2012

Director: Professor Roland Stull
UBC Geophysical Disaster Computational Fluid Dynamics Center
rstull@eos.ubc.ca
604-822-5901

1. The BlueSky Framework

BlueSky is not a single model, but a system designed to run and interconnect a variety of models to produce dispersion and trajectory forecasts of smoke. BlueSky takes care of how the individual models communicate with each other. It creates a directed acyclic graph (DAG) of the steps necessary to complete all model operations. Here is the DAG for the Canadian installation.

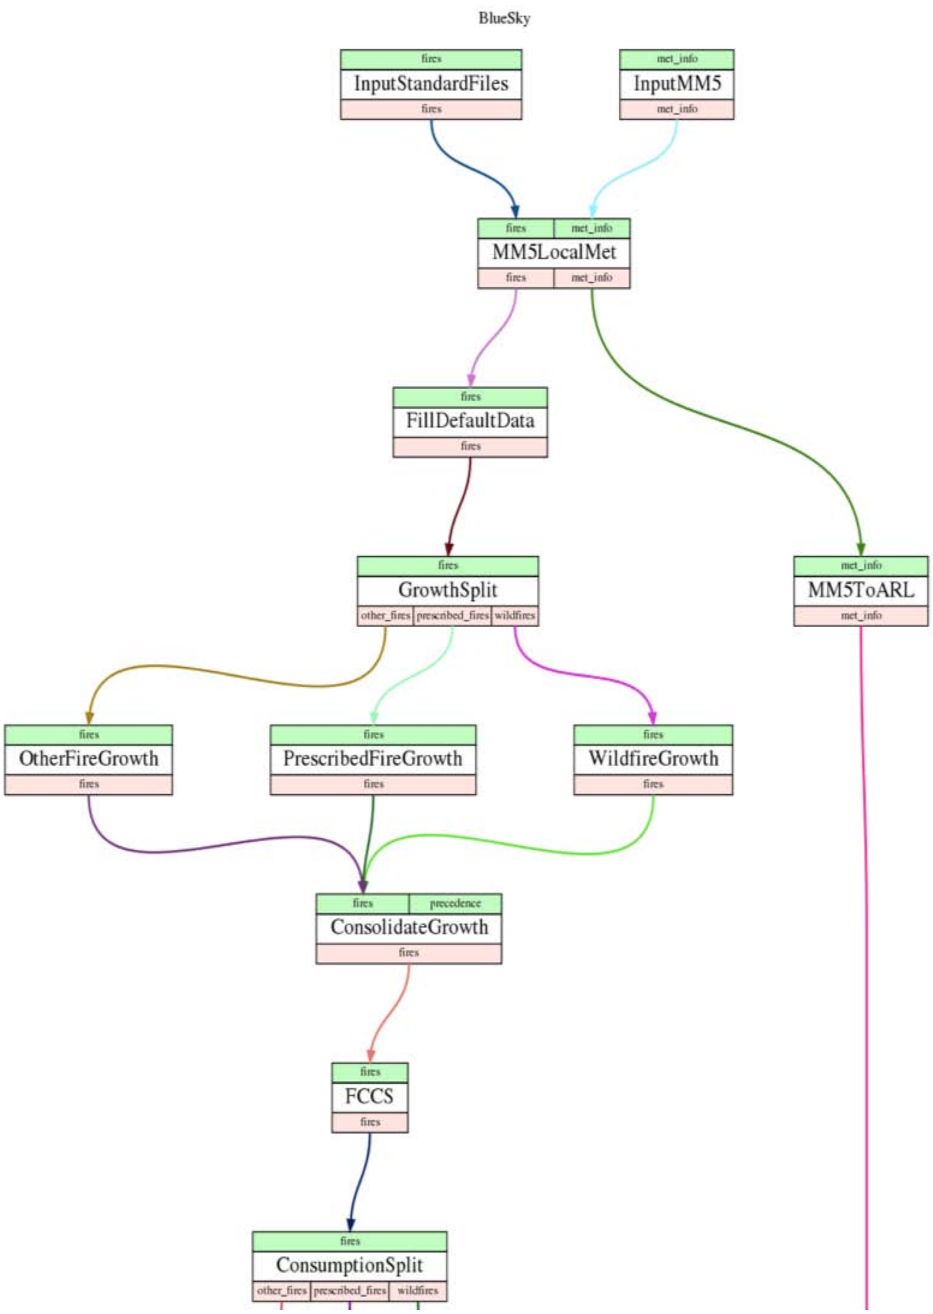

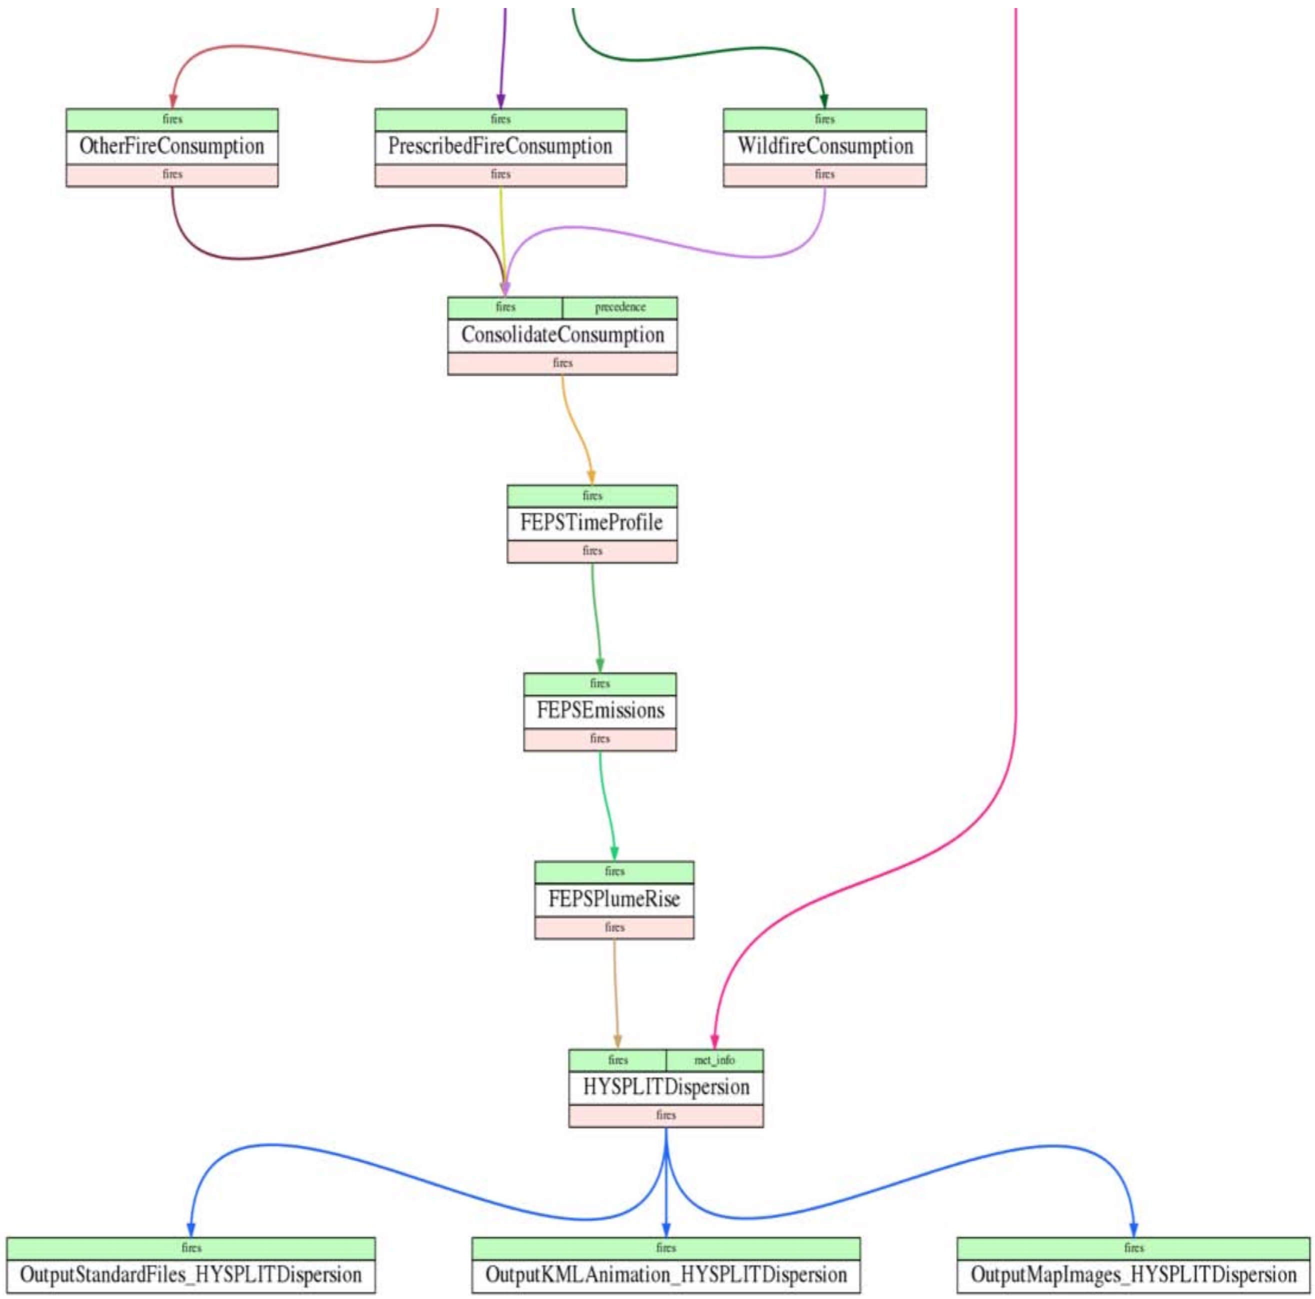

In short, the Canadian BlueSky installation takes Canadian Wildfire Information System (CWFIS) consumption data and runs it through various stages of the Fire Emission Production Simulator (FEPS) model to produce emissions profiles. Coincidentally, it takes raw MM5 weather forecast data and uses mm52arl to convert it to a format useable by the dispersion model HYSPLIT. The emissions and Air Resources Laboratory (ARL) data are fed into HYSPLIT and the resulting dispersion data is then converted into output files for the display on the web or use in analysis.

2. Canadian Wildfire Information System (CWFIS)

CWFIS tracks fires in Canada using information from satellite imaging. Based on the satellite data, fire intensity is determined. Combining that information with the Canadian fuel inventories allows for CWFIS to determine fire consumption. The specifics of the system can be found in the paper "An approach to operational forest fire growth predictions for Canada".[1]

Once this information is calculated, CWFIS produces a database file called hotspots_ec.csv. This file is available for download from the CWFIS ftp site ftp://ftp.nofc.cfs.nrcan.gc.ca/pub/fire/maps. The file typically accumulates data throughout the year (becoming large by the end of the year) so part of the setup phase involves culling older fires from the available data.

A special initialization file (download_cwfis.ini) instructs BlueSky to download the hotspots_ec.csv file and process it into fire_locations.csv and fire_events.csv files with only recent fire data (from the previous 24 hours). These files are used as input to the FEPS model to perform the smoke emissions and time profile calculations. They mimic the types of fire data that would be made available from the American SMARTFire system. Password and user information for downloading the file (in case you need to do so manually) can be found inside the download_cwfis.ini file.

The fire_locations.csv and fire_events.csv files are stored in /bluesky/archive/CWFIS as fire_locations_YYYYMMDDHH.csv and fire_events_YYYYMMDDHH.csv where YYYY represents the year, MM represents the month (00 .. 12), DD represents the day (01 .. 31) and HH represents the initialization hour (typically 00 or 12) of the date of the fire data. This allows for additional runs to be made in future using past fire events. However, to use the files within the system, they must be copied to the /bluesky/bluesky_nested/input/fires directory and renamed fire_locations_YYYYMMDD.csv and fire_events_YYYYMMDD.csv. Typically you will also need to run the bounding script (CWFIS_Bound_Polar.pl) once you have copied the files into place and renamed them to cull any fires outside the meteorological domain. The typical call to that script would look like:

```
/usr/bin/perl /bluesky/bin/CWFIS_Bound_Polar.pl 55.0 -120.0 12.0 335 285
```

The five numerical arguments indicate the polar stereographic center of the domain (latitude and longitude), the resolution of the domain, and the number of x and y points for the domain (reduced slightly to exclude the boundary).

3. Mesoscale Model v5 (MM5)

The University of British Columbia runs the MM5 weather forecast model to produce meteorological data for the BlueSky system. A set of nested domains run together feeding data back and forth as the forecast is produced. The coarsest domain has 36km resolution and covers most of Canada. Within this domain, 12km and 4km resolution domains are produced. The 12km domain extends over British Columbia, Alberta, Saskatchewan, Manitoba and the western portion of Ontario. It also covers the southern portions of the Yukon, Northwest Territories, and Nunavut as well as the northern portion of the western United States. The 4km domain covers British Columbia and Alberta as well as a portion of southern Yukon and Northwest Territories

---

[1] Authors: K.R. Anderson, P. Englefield, J.M. Little, Gerhard Reuter, International Journal of Wildland Fire, 2009, v18, pg 893-905

and a portion of the northwestern United States. Both forecasts are initialized at 0 UTC each day and produce 84 hours of forecast although only the first 60 hours are used in the smoke forecast.

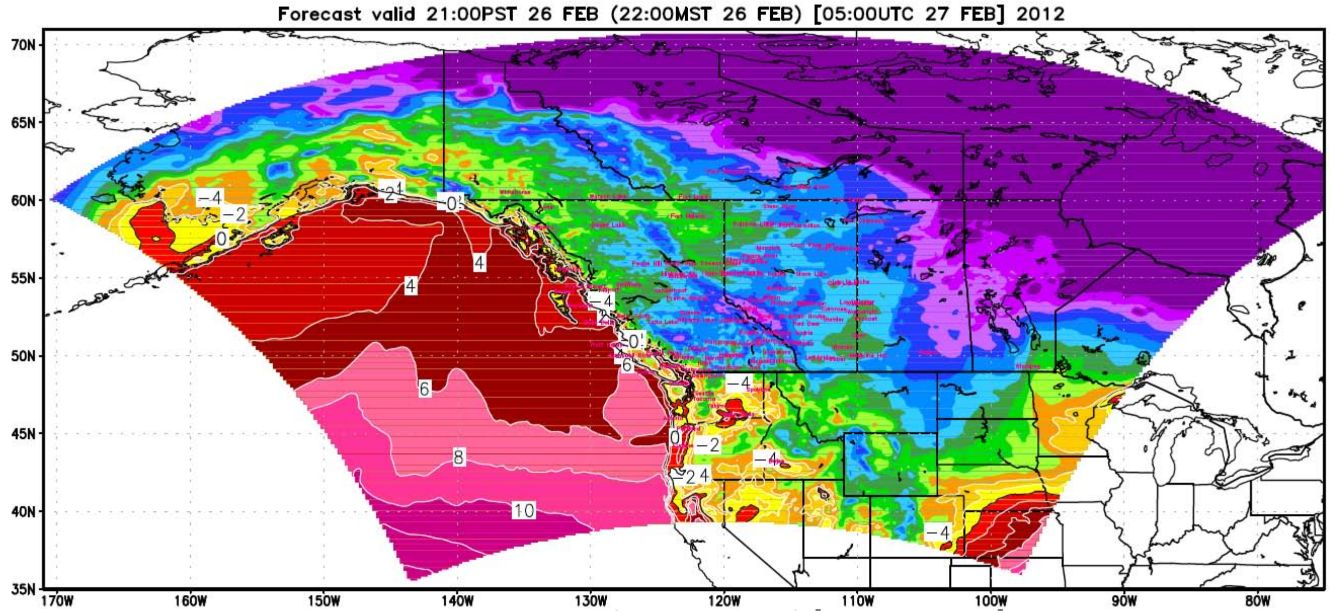

MM5 12km domain surface temperature forecast

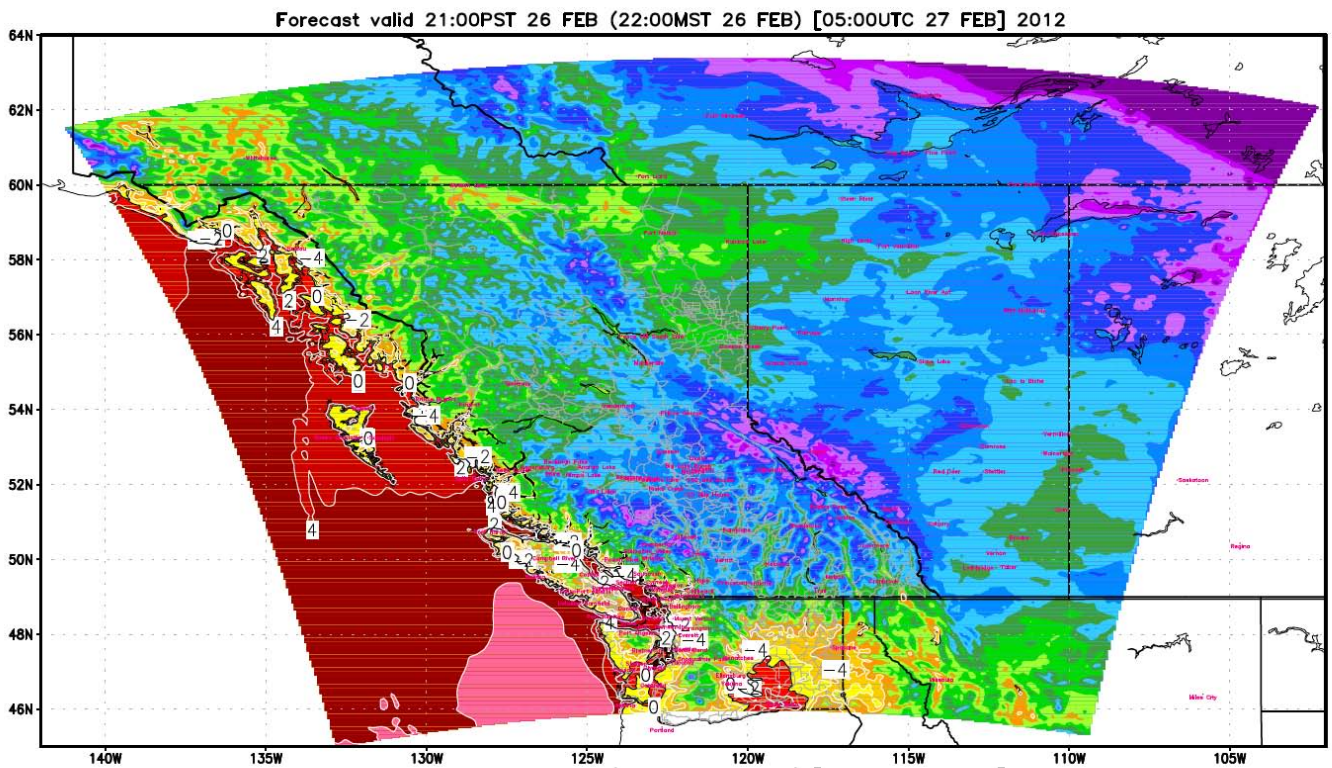

MM5 4km domain surface temperature forecast

Raw forecast data is produced in files MMOUT_DOMAINx_yyy where x is the domain nest ID (the 12km domain is nest 3, and the 4km domain is nest 4) and yyy is the file number (incrementing by 1 for each output time). The MM5 data is output every 30 minutes, so MMOUT_DOMAIN4_002 would be the 4km domain output for the first forecast hour (1 UTC). Because the output is for every 30 minutes, but BlueSky only makes use of the hourly forecasts, the raw output needs to be linked or renamed. (ie MMOUT_DOMAIN4_002 should be MMOUT_DOMAIN4_001 if we only count hourly output).

The ARL converter for MM5 takes this raw data and converts it into a format that can be used by HYSPLIT. The Bluesky_nested.pl script can be used to link the MM5 data so that it can be converted. To do this, the script can be invoked with the -l flag, or it can be run normally (in which case it will link the files as part of the overall process). However, the original data must be available on the savash cluster in /temp/mm5_op/MM5_RAW for it to be linked (locally available in /nfs/savash/data). The data can be made directly available in the /bluesky/bluesky_nested/input/met directory instead by copying it there. Linking across filesystems is preferred due to the large size of the daily output (84GB for 60 hours) which takes considerable time to transfer over the network.

4. HYSPLIT Dispersion Model

The dispersion modeling of the smoke is conducted by the HYSPLIT model, which is available from the ARL at http://www.arl.noaa.gov/HYSPLIT_info.php. For the 2011 fire season, we used a precompiled v4.8 executable. A newer version of HYSPLIT (v4.9) is available, but it was not successfully compiled and tested for the 2011 fire season. Continued work has shown progress, but evaluation is needed to determine if output is reasonable.

HYSPLIT is run as a single process. The newer v4.9 has support for MPI directives which allows it to be run as parallel processes. The mode in which HYSPLIT is normally run (horizontal gaussian, vertical top-hat) is not supported in MPI but can still be run this way (does not appear to fail or show any negative side effects). Running a more complex model (like 3D particle) can be done in parallel but is significantly slower due to increased computation.

In v4.8 we had significant computational problems running the model with a full complement (20) of vertical levels. STI introduced a variable for reducing the number of vertical levels by a set factor. We have typically run v4.8 with a vertical reduction of 20, which means that we have run HYSPLIT with only 1 vertical level. In v4.9, this reduction still works, however, with improved performance in parallel runs, we have been able to run with more levels. This seems to create a complication since the output only tracks pollutants near the surface (more pollutants seem to be transported/trapped in the upper levels). This gives the appearance that there is less pollution in the output since more pollutants are contained in higher levels. Since the output cannot capture this (and is not intended to), the result is an apparent decrease in pollutants.

5. Model Output and Storage

BlueSky takes the HYSPLIT output and first produces netcdf files of the data. From these netcdf files, HYSPLIT produces .jpeg files of each hour of forecast (or less if you change the configuration). It also produces kml files which it compiles into a kmz file for use in Google Earth. Lastly, it can produce SMOKE ready files for CMAQ in the pre-nested framework, but this is not currently implemented in the nested version.

All of the model output is stored in /bluesky/archive/FULL. The visualized output, netcdf, and, when available, SMOKE-ready files are automatically tarballed together and put into

/bluesky/archive/FULL/output-YYYYMMDDHH.tar.gz files. Operational files (intermediate data, logs, configuration files) are stored in a file called /bluesky/archive/FULL/archive-YYYYMMDDHH.tar.gz. The run.log file is stored in the archive tarball and has a lot of information about the complete run.

The daily CWFIS download is stored in /bluesky/archive/CWFIS. This includes the fire_locations and fire_events files.

6. Running the Framework

The current version of the framework is installed in the /bluesky/bluesky_nested directory. There is an older non-nested version of the framework installed in /bluesky/bluesky. Direct control of the components of the framework is handled by an initialization file called UBCFULL_NEST.ini which is located in /bluesky/bluesky_nested/setup. CWFIS downloading has a special initialization file which can also be found in setup called download_cwfis.ini. Inputs, such as fire and met data, are located within the /bluesky/bluesky_nested/input/fires and /bluesky/bluesky_nested/input/met directories. Temporary files and logs generated during the forecast are stored within the /bluesky/bluesky_nested/working/YYYYMMDDHH.x directory (the .x is a number, typically 1, which signifies which run or re-run is covered). Output generated is contained within /bluesky/bluesky_nested/output/YYYYMMDDHH.x.

To run the framework, you can simply invoke the executable

/bluesky/bluesky_nested/bluesky

This will run the framework with the configuration in the default initialization files. To run with specific initialization files, use the -d flag with the filename.

/bluesky/bluesky_nested/bluesky -d UBCFULL_NEST.ini

If the initialization file is located in the setup directory, then you can use just the filename. If it is located somewhere else, you will need to specify the full path.

The framework has additional flags that can be specified to affect how it is run. For a full listing, invoke the help flag

/bluesky/bluesky_nested/bluesky -h

The -P flag is good for getting information about individual modules. The -G and -L options will show you how a particular .ini file will execute (or create the graph). There are many other information flags as well which have descriptions in help.

The CWFIS download and model runs are handled separately via different initialization files, thus a typical forecast will require multiple invocations of the framework with each .ini file. In addition, although the framework will run the models, other work has to be done to archive and ftp output, cleanup files, and link data. To handle the overall creation of a forecast, a script called

Bluesky_nested.pl was created. This script can be invoked in several ways to handle different aspects of a forecast, but in essence it:

- cleans up the input, output, and working directories of the framework
- runs the framework to download the CWFIS data, stores a copy of this data in /bluesky/archive/CWFIS and invokes the CWFIS_Bound_Polar.pl script to remove fires outside the meteorological domain from the working copies in the /bluesky/bluesky_nested/input/fires directory
- links the meteorological data into the /bluesky/bluesky_nested/input/met directory
- runs the framework (with the UBCFULL_NEST.ini) to invoke all other stages and create a complete dispersion forecast
- tars and copies the output and archive to /bluesky/archive/FULL, also copies the output to the ftp server for download

The protocol for invoking the script is:

perl /bluesky/bin/Bluesky_nested.pl [-d | -l | -m] [YYYYMMDD] [HH]

where

- (optional) only one of the -d, -l, or -m flags can be invoked, if no flag is supplied, all stages are run
- -d flag indicates that only the CWFIS download stage should be done
- -l flag indicates that only the linking of the MM5 data stage should be done
- -m flag indicates that only the running of the framework for the dispersion forecast stage should be done
- (optional) the YYYYMMDD indicates the date for which the framework should be run, if not supplied, the current date is used
- (optional) the HH indicates the initialization hour for which the framework was run (typically 00 or 12 for 0Z and 12Z), if not supplied, 00 is used.

The daily runs are cron'd (ie scheduled) under the blueop user so that both runs occur automatically. Emails are sent to the blueop user regarding the success or failure of the cron'd runs. Normally a complete daily run can be handled simply by executing

perl /bluesky/bin/Bluesky_nested.pl

To do a non-daily run (such as a test run or re-run of a previous day), you typically cannot run the framework by the Bluesky_nested.pl script alone. The issue tends to be with getting CWFIS data from the appropriate period. We have archived daily files of processed CWFIS data (the files you would normally get from running the framework to do the CWFIS download) in /bluesky/archive/CWFIS. To use these files, you first need to copy the appropriate fire_locations...csv and fire_events...csv files to the /bluesky/bluesky_nested/input/fires directory. These files have a date field in them that can be YYYYMMDD or YYYYMMDDHH (ie 20110705 or 2011070500). The framework does not handle the HH field, so any fire files containing it need to be renamed without it (ie fire_locations.2011070500.csv should become fire_locations.20110705.csv).

Next you need to bound the fires in the files to the meteorological domain. You can do this by running

/usr/bin/perl /bluesky/bin/CWFIS_Bound_Polar.pl 55.0 -120.0 12.0 335 285

Now that the CWFIS data is set, you need to link the meteorological data. Do this by running

/usr/bin/perl /bluesky/bin/Bluesky_nested.pl -l YYYYMMDD

for the date you want. This will only work if the MM5 data is available on the savash cluster in /temp/mm5_op/MM5_RAW/YYMMDD00 (locally mounted as /nfs/savash/data).

Finally, complete the dispersion forecast by running

/usr/bin/perl /bluesky/bin/Bluesky_nested.pl -m YYYYMMDD

for the date you want. The output will be tarballed by the script and stored in /bluesky/archive/FULL/output-YYYYMMDDHH.tar.gz and /bluesky/archive/FULL/archive-YYYYMMDDHH.tar.gz. This could overwrite existing tarballs so it is best to check and rename any existing tarballs for that date beforehand.

7. The BlueSky Machine

The machine running BlueSky has the network name bluesky.eos.ubc.ca. It is a dual hex-core 2.6GHz Xeon processor machine with 24GB of RAM and 1.6TB of RAID disk, owned by the BC Ministry of Environment. The operating system is Ubuntu v10.10 (maverick). The original BlueSky computer supplied by the USFS was a dual opteron running Ubuntu v8. The changes in operating system and architecture required some minor changes to pre-compiled binaries, but for the most part BlueSky moved well from the older system to the newer one. STI still only supports BlueSky on the older Ubuntu v8 operating system so changes to the framework need to be tested locally before being made live.

8. Configuration and Framework Specifics

Most of BlueSky is written as python modules directed by configuration files. The configuration files are .ini files and can be found in /bluesky/bluesky_nested/setup and /bluesky/bluesky_nested/base/etc. Copies of the local .ini files (which are available in /bluesky/bluesky_nested/setup) can also be found in /bluesky/ini.

The .ini files can be configured to alter any part of BlueSky's operation by changing the standard configuration. The .ini files in /bluesky/bluesky_nested/base/etc provide default values for most every part of the framework. The .ini files in /bluesky/bluesky_nested/setup contain operational configurations that override the defaults. Our local file UBCFULL_NEST.ini can be found there. There is also a download_cwfis.ini file which is used to download the CWFIS data. Only these two files should be altered if changes are to be made to the way the framework runs. The ini files in /bluesky/bluesky_nested/base/etc should only be referenced, especially as these files will change whenever BlueSky is updated.

To change the operation of BlueSky within an .ini file, you need to go to the section of the framework within the .ini that you want to alter. Sections are denoted by titles in square brackets corresponding to particular operations in the DAG at the start. For example, to alter HYSPLIT, we would go to the [HYSPLITDispersion] section. There are also [DEFAULT] and [META] sections for determining what modules will run within the framework (you can limit the framework to only running one module, for example see download_cwfis.ini).

Under a given section, you can set variables that would be accepted for configuration. Usually checking the default.ini file is a good reference. Sometimes individual modules may also have .ini files (which can be found in /bluesky/bluesky_nested/modules/... under the appropriate module name). These files can be referenced for specific module variables. Here is an example of the HYSPLIT section of the UBCFULL_NEST.ini file.

```
[HYSPLITDispersion]
USER_DEFINED_GRID = true
CENTER_LATITUDE = 55.0
CENTER_LONGITUDE = -117.5
WIDTH_LONGITUDE = 55.0
HEIGHT_LATITUDE = 20.0
SPACING_LONGITUDE = 0.10
SPACING_LATITUDE = 0.10
VERTICAL_EMISLEVELS_REDUCTION_FACTOR = 20
OPTIMIZE_GRID_RESOLUTION = false
FIRE_INTERVALS = 0 50 100 500 1000
MAX_SPACING_LONGITUDE = 0.10
MAX_SPACING_LATITUDE = 0.10
```

This configuration sets the grid HYSPLIT will work on and the vertical level reduction.

Initialization files have the same formatting. Headers are contained within square brackets. Variables are usually all in capitals and set equal to some value. Any variables appearing after a header belong to that header section. Comments begin with '#' and continue to the end of a line. Complete listings of headers and variables appear in the default initialization files located in /bluesky/bluesky_nested/base/etc or individual module directories.

a. download_cwfis.ini

This initialization file downloads and preps CWFIS data into a format useable by FEPS then quits the framework. The file is broken down by section where the .ini file code appears in bold.

[META]
includes=${GRAPH_DIR}/noop.graph
inputs=CWFIS
outputs=StandardFiles

The META section helps the framework determine how to setup the DAG and know what modules to run. It expects CWFIS inputs and it outputs 'StandardFiles' which is framework code for intermediate files useable by other modules.

[DEFAULT]
OUTPUT_DIR=${INPUT_DIR}
SPIN_UP_EMISSIONS=true
EMISSIONS_OFFSET=-24
HOURS_TO_RUN = 84

The DEFAULT section can contain any variables and can overrule variables in other sections, but typically only contains more general runtime variables. OUTPUT_DIR is where the outputs should end up (in this case it is the fires directory under inputs). SPIN_UP_EMISSIONS determines if emissions need to be converted. EMISSIONS_OFFSET determines how many hours from start time to include fires (negative is prior to start time, positive is after start time). HOURS_TO_RUN is the duration of emissions.

```
[InputCWFIS]
# Add the following 2 lines if you use a local file
USE_LOCAL_CWFIS_FILE = false
#LOCAL_CWFIS_FILE = /bluesky/archive/CWFIS/bluesky20100816-19.csv
#LOCAL_CWFIS_FILE = /bluesky/archive/CWFIS/hotspots2011.csv
LOCAL_CWFIS_FILE = /bluesky/bluesky_nested/hotspots.csv
HOSTNAME=ftp.nofc.cfs.nrcan.gc.ca
# Note: You can edit the next two lines to specify a username/password to use
# when connecting to the CWFIS FTP server.  Alternatively, you can specify
# an [InputCWFIS] section with USERNAME= and PASSWORD= lines in your
# user-specific config file: ~/.bluesky
USERNAME=anonymous
PASSWORD=blueskyframework
#FTPPATH=/pub/fire/maps/hotspots_archive/
#FTPFILE=hotspots%Y.dbf
FTPPATH=/pub/fire/maps/
#FTPFILE=hotspots_ec.dbf
HOTSPOT_RADIUS=350.0
DUFF_CONSUMPTION_RATIO=0.5
FLAMING_SMOLDERING_CONSUMPTION_RATIO=0.5
INCLUDE_ALL_HOTSPOTS=false
ERROR_ON_ZERO_FIRES=true
WARN_ON_ZERO_FIRES=true
```

This section handles the CWFIS module configuration. USE_LOCAL_CWFIS_FILE and LOCAL_CWFIS_FILE allow for prep to be done from a CWFIS data file that is located on the server as opposed to downloading the data file. This was intended for testing purposes and is not typically used.

HOSTNAME is the ftp server hosting the CWFIS data. USERNAME and PASSWORD are provided for secure login (when necessary).

FTPPATH is the location on the ftp server where the data is located. FTPFILE is the name of the data file (it may include special flags like %Y for year, %m for month, %d for day).

HOTSPOT_RADIUS is the number of meters used to define an area for a fire event (all fires within 350m are inside a single event). DUFF_CONSUMPTION_RATIO and FLAME_SMOLDERING_CONSUMPTION_RATIO determine how much to count these values in consumption. INCLUDE_ALL_HOTSPOTS is supposed to determine whether or not to limit hotspots to the dispersion and meteorological domains or simply include them all. However, it is not currently implemented properly for polar stereographic domains and thus does not work with MM5. ERROR_ON_ZERO_FIRES and WARN_ON_ZERO_FIRES let the user know that there are no valid fires within the data and prevents a forecast run from starting with no fires.

[OutputStandardFiles]
SPIN_UP_EMISSIONS=true
EMISSIONS_OFFSET=-24
HOURS_TO_RUN = 84
EVENTS_FILE=fire_events_%Y%m%d.csv
LOCATIONS_FILE=fire_locations_%Y%m%d.csv
HOURLY_FILE=fire_emissions_%Y%m%d.csv

The OutputStandardFiles section determines how output should be handled. We have the same emissions variables here (we could leave them out since they appear in DEFAULT). Converting the CWFIS data produces two (or if it becomes supported, three) standard files: fire_events.csv, fire_locations.csv, and fire_emissions.csv. The events, locations, and hourly emissions filenames are set using EVENTS_FILE, LOCATIONS_FILE, and HOURLY_FILE respectively.

b. UBCFULL_NEST.ini

This initialization file covers the running of the framework from FEPS through to final output. This initialization file requires fire input to be pre-processed and meteorological input to be copied or linked into the appropriate input directory.

[META]
BS_DIR=/bluesky/bluesky_nested
 # The graph file that describes how the models interconnect
includes=${GRAPH_DIR}/default.graph
 # The input tasks we're going to run to read in our input data
inputs=$MET StandardFiles
 # What step accepts thet input data from the input(s) above
start=$EXTRACT_LOCAL_MET
 # The step(s) we are trying to accomplish -- the "stopping point" on the graph
targets=$DISPERSION
 # What output tasks to run once we've reached our stopping point
#outputs=StandardFiles SMOKEReadyFiles MapImages KMLAnimation
outputs=StandardFiles MapImages KMLAnimation

The META section contains the typical includes, inputs, and outputs variables which help determine what portions of the framework run. In this case, we have a default.graph for the overall run. Inputs include StandardFiles and meteorological data ($MET). Outputs include StandardFiles (which includes the netcdf files), MapImages (which are the jpegs), and KMLAnimation (which are the kmz files). The BS_DIR variable sets the root directory for the framework (which defaults to /bluesky). Since we have put the framework in a non-standard location, we can use this variable to inform the framework where to find itself.

```
[DEFAULT]
# Analysis time range (override with -d and -H command-line switches)
#DATE = 2008050100Z
HOURS_TO_RUN = 60
# What models are we going to run to get those outputs?
#MET=CALMET
MET=MM5
#EXTRACT_LOCAL_MET=CALMETLocalMet
EXTRACT_LOCAL_MET=MM5LocalMet
FUEL_LOAD=FCCS
TIME_PROFILE=FEPSTimeProfile
EMISSIONS=FEPSEmissions
PLUME_RISE=FEPSPlumeRise
#DISPERSION_MET=CALMET
DISPERSION_MET=MM5ToARL
#DISPERSION=CALPUFF
DISPERSION=HYSPLITDispersion
TRAJECTORY_MET=MM5ToARL
TRAJECTORY=HYSPLITTrajectory
# Growth models (by fire type)
WILDFIRE_GROWTH=Persistence
PRESCRIBED_GROWTH=NoGrowth
OTHER_GROWTH=Persistence
#OTHER_GROWTH=NoGrowth
# Consumption models (by fire type)
WILDFIRE_CONSUMPTION=CONSUME
PRESCRIBED_CONSUMPTION=CONSUME
OTHER_CONSUMPTION=CONSUME
# Consumption canopy fraction (by fire type)
WILDFIRE_CANOPY_FRACTION = auto
PRESCRIBED_CANOPY_FRACTION = 0.0
OTHER_CANOPY_FRACTION = auto
# Other run-specific options (overrides $BS_DIR/base/etc/defaults.ini)
#HOURS_TO_RUN_TRAJECTORY = 12
HOURS_TO_RUN_TRAJECTORY = 60
SPIN_UP_EMISSIONS = true
EMISSIONS_OFFSET = -24
DISPERSION_OFFSET = 0
STOP_IF_NO_BURNS = true
STOP_IF_NO_MET = true
```

MM5_NEST = true
 # UBC input Met Dir (if different than default location)
#MET_DIR = /bluesky/bluesky/input/met
 # Stop BlueSky from removing fire locations outside the domain
 # This algorithm currently fails for Polar Stereographic MM5 domains
REMOVE_INVALID_LOCATIONS=False
 # Location of the map image shape files
BC_MAPS_DIR = /bluesky/bc_maps
SHAPE_PATH = ${BC_MAPS_DIR}/shapefiles

The DEFAULT section includes a lot of new variables. DATE can be used to specify a particular date (otherwise the current date is used). HOURS_TO_RUN specifies how long the forecast sections of the framework should run. In the CWFIS data, we started 24 hours in the past and extended 60 hours into the future for a total of 84 hours. For the forecast section though, we are only going 60 hours into the future.

The MET and EXTRACT_LOCAL_MET variables determine what type of meteorological data to use and how to interpret it. FUEL_LOAD, TIME_PROFILE, EMISSIONS, PLUME_RISE, DISPERSION_MET, DISPERSION, TRAJECTORY_MET, and TRAJECTORY determine which models to run for each section. Although we specify models for trajectories, we do not actually run that portion as specified by the targets variable in META.

WILDFIRE_GROWTH, PRESCRIBED_GROWTH, and OTHER_GROWTH determine what type of fire growth model should be used to grow fires of specific types. NoGrowth will not grow the fires. Persistence will assume the fire characteristics do not change (gets no larger or smaller, just keeps burning the same throughout). There is also the option Growth which will attempt to grow the fires, but given our lack of fuel loading information in Canada, this option is not possible. Normally, there would be variables for fuel loading, but CWFIS negates the need for this providing the consumption data directly.

WILDFIRE_CONSUMPTION, PRESCRIBED_CONSUMPTION, and OTHER_CONSUMPTION determine the consumption model to use. Possibilities include CONSUME, FEPSConsumption, and NoConsumption (to skip that phase). Since CWFIS provides this data, we can select NoConsumption, but just to be safe we select the CONSUME model which just passes the data directly through to StandardFiles.

 # Define what the MM5 filename structure will look like
[InputMM5]
MM5_PATTERN = ${MET_DIR}/MMOUT_DOMAIN3_*
MM5_NEST_PATTERN = ${MET_DIR}/MMOUT_DOMAIN4_*

The InputMM5 section sets up where the MM5 raw input will be and what the filenames will look like.

 # Look for file patterns in the input fire files
[InputStandardFiles]
USE_DAILY_FILE_PATTERNS=true
#LOCATIONS_FILE = ${INPUT_DIR}/fire_locations_20100805.csv

#LOCATIONS_FILE = /bluesky/data/fires/fire_locations_20090815.csv

InputStandardFiles handles where and how most input files appear (particularly fire_locations and fire_events). Usually default values should be used, although this does create an option for using specific files or files outside of the usual input directory. USE_DAILY_FILE_PATTERNS indicates whether or not the file should contain a YYYYMMDD field in the filename.

 # User-defined sampling grid option for Hysplit
[HYSPLITDispersion]
USER_DEFINED_GRID = true
CENTER_LATITUDE = 55.0
#CENTER_LONGITUDE = -115.0
CENTER_LONGITUDE = -117.5
#WIDTH_LONGITUDE = 60.0
WIDTH_LONGITUDE = 55.0
HEIGHT_LATITUDE = 20.0
SPACING_LONGITUDE = 0.10
SPACING_LATITUDE = 0.10
VERTICAL_EMISLEVELS_REDUCTION_FACTOR = 20
#VERTICAL_EMISLEVELS_REDUCTION_FACTOR = 1
OPTIMIZE_GRID_RESOLUTION = false
FIRE_INTERVALS = 0 50 100 500 1000
MAX_SPACING_LONGITUDE = 0.10
MAX_SPACING_LATITUDE = 0.10

This section covers the configuration of HYSPLIT. Most of the variables listed need to be defined here (defaults are not applicable to a normal setup).

WIDTH_LONGITUDE defines a total width of the HYSPLIT domain (half the width on each side of CENTER_LONGITUDE). In the same way, HEIGHT_LATITUDE AND CENTER_LATITUDE define a height for the domain. Together, these four variables define a latitude/longitude box. The SPACING_LONGITUDE and SPACING_LATITUDE define the distance between points within the domain box in degrees. MAX_SPACING_LATITUDE and MAX_SPACING_LONGITUDE allow for a bit of wiggle room around the points if specified to larger values that the spacing variables. This is only applicable if OPTIMIZE_GRID_RESOLUTION is true (which will allow HYSPLIT to attempt to make a 'best fit' to the fire information by relaxing the distance between grid points based on the number of fires in an area). While HYSPLIT sets its domain as a latitude/longitude box, the meteorological domains of MM5 are polar stereographic. As a result, for HYSPLIT to run properly, its domain must be completely contained within the MM5 domain. If it is not, the model may crash, or smoke may be removed from the model at the meteorological domain boundaries (which leads to strange artifacts).

By default, HYSPLIT is designed to run with 20 vertical levels. This creates the largest computational strain on the system. VERTICAL_EMISLEVELS_REDUCTION_FCATOR provides for a way to reduce the number of levels. This variable must be a factor of 20. Typically it is set to 20, meaning that HYSPLIT's vertical levels are reduced by a factor of 20, or to 1 level.

FIRE_INTERVALS is only used if OPTIMIZE_GRID_RESOLUTION is set. It sets an number of thresholds (total fires) around grid points. As the number of fires increases, the threshold determines how far apart grid points can be (up to the MAX settings). This allows the grid to become coarser around heavy fire areas, treating them more like single events. Although this was tested, by default it is not used and can lead to changes in fire resolution.

 # Enable SMOKE ready output
#[OutputSMOKEReadyFiles]
#COUNTRY_TO_PROCESS = CANADA

If SMOKE output is set, you can determine the countries to output files for (if you wish to limit the output). The nested version of BlueSky does not support SMOKE output at this time.

 # Enable Canada map images
[OutputMapImages]
MAP_TEMPLATE = ${BC_MAPS_DIR}/MapTemplate_Canada.map
USE_DATA_PROJECTION = false
 # Define the name of the output images
IMAGE_FILE_PATTERN=dispersion_%Y%m%d%H.jpg

The OutputMapImages section allows for setting related to the production of the jpegs and kmz files to be set. MAP_TEMPLATE defines the map template file to use (the default is in /bluesky/bluesky_nested/base). We use a specially designed map template. USE_DATA_PROJECTION will attempt to override the setting of the map template to make the output appear on the same grid projection as the dispersion output. We do not use this. IMAGE_FILE_PATTERN shows what the jpeg filenames should appear as. It is important not to change this unless the change is communicated to other users who display the images on websites, etc.

9. Modules

In the /bluesky/bluesky_nested/modules directory, you can find the various modules installed as part of the framework. Modules can be individual models (such as HYSPLIT) or processing components (such as MM5data). Each module might have subdirectories with different versions of the module. BlueSky will usually make use of the most recent module unless you have specified not to in one of the .ini files. You can get more general information about the module with

/bluesky/bluesky_nested/bluesky -P <module>

Inside each module directory you will find executables, python code, and possibly .ini files. The executables are pre-compiled and distributed as part of the framework so you cannot change these. If you are having trouble with a specific executable, you will likely have to go to STI to get a replacement.

The python files (.py for the scripts and .pyc for the compiled scripts) contain the code that runs the executables for this specific module within the framework. It is important to note that these files get changed when the framework is updated so if you make changes to the python itself, you should either communicate the changes back to STI or make notes so that you can implement those changes on future frameworks.

Generally, it is not a good idea to try to change the running of a particular module outside of what can be configured via .ini files. In your running .ini file (ie UBCFULL_NEST.ini) you can add a section for changing a given module and overriding the variables you find within its own default .ini file.

10. Output Generation

The output images make use of the mapping capability of the framework. This is controlled by the /bluesky/bluesky_nested/base/modules/maps.py python script and the map template files located in/bluesky/bc_maps. The python script itself should not require changing. The MapTemplate_Canada.map handles most of the map configuration (background, cities, colour scheme) and is the only element that should require modification to change the appearance of output.

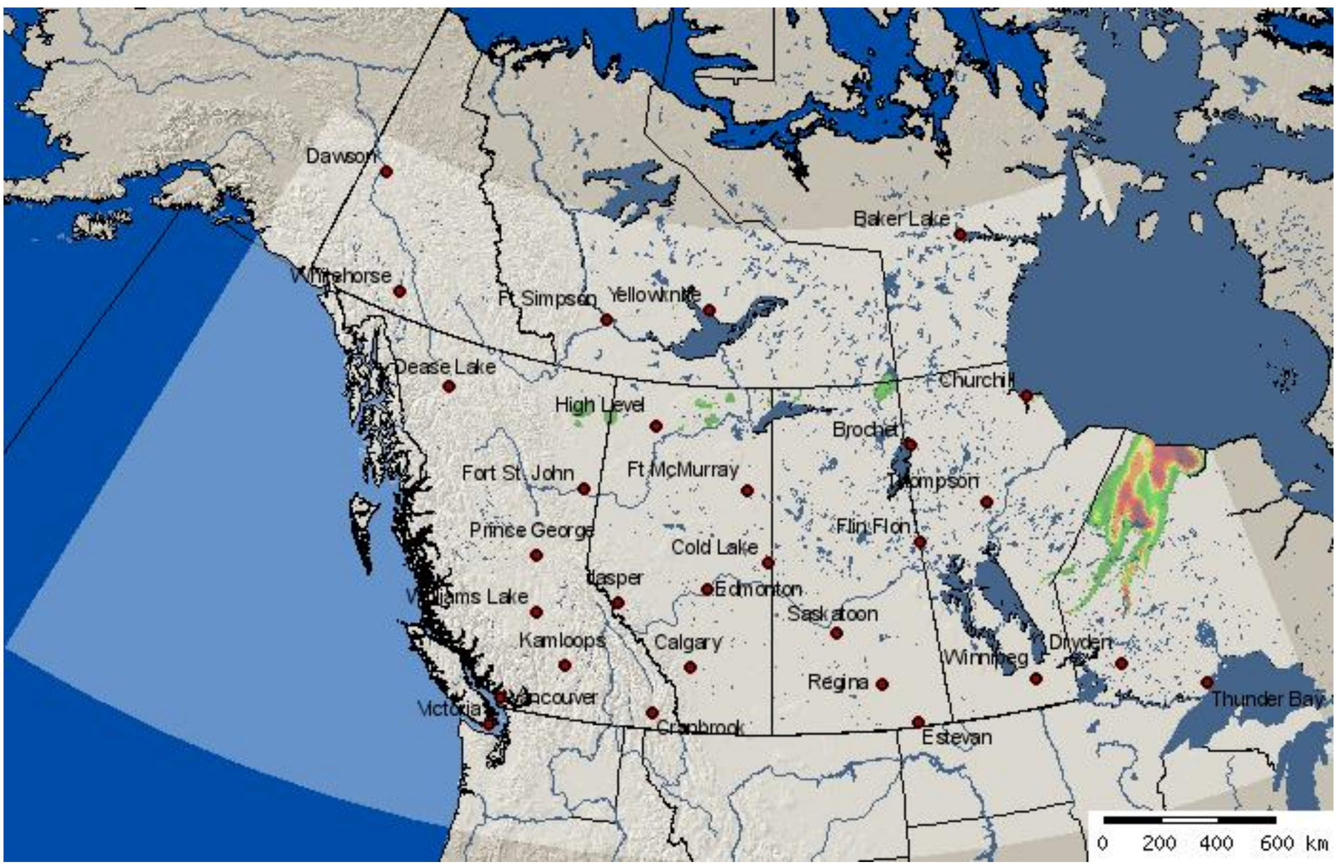

Example of jpeg output, July 16, 2011

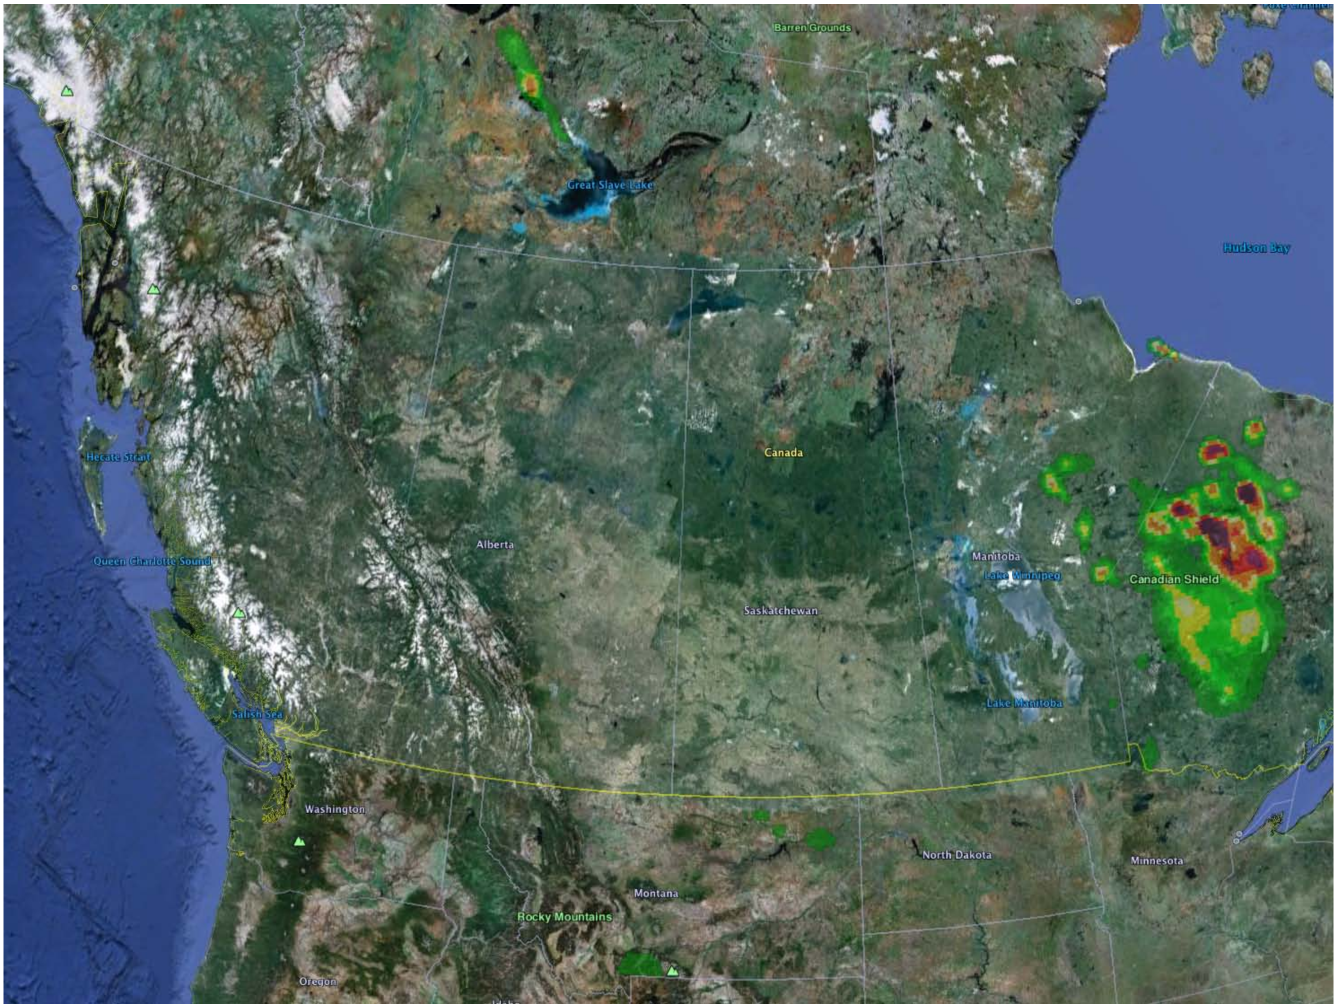

Example of KMZ output, July 19, 2011

11. Problem Solving/Missing Forecasts

On occasion, the framework may not complete either due to missing fire data, missing meteorological data, a lack of fires within the dispersion domain, hardware failures, or some other cause. The result of this is that the output tarball normally sent to crypt.eos.ubc.ca will not be created or will be empty. This is the best indicator of a problem.

To diagnose the problem, it is best to first look for the file run.log which can be located within the working directory (if the forecast run has not been deleted) or the archive-YYYYMMDDHH.tar.gz file in /bluesky/archive/FULL. The run.log file is usually very large, but if you scan it for 'ERROR' or 'Error', it should bring you to occurrences where the framework had trouble. Often the first occurrence is what caused your problem.

Lack of fires

If you find that there is an error running download_cwfis.ini resulting in fire_locations and fire_events files not being created, then there will either be a problem with not being able to connect to the CWFIS ftp server or that the file will be missing or corrupt. In this case, you need to notify CWFIS of the problem and see what they can do to work towards a resolution.

In other cases, download_cwfis.ini will show no errors but HYSPLIT or FEPS will complain there are no fires. In this case, check the fire_locations and fire_events files in /bluesky/bluesky_nested/input/fires. Does the information match the information in the file on the ftp server? Do the files contain fires which are recent (within the past 24 hours) and within the dispersion domain? Usually, there are no fires because there are no recent fires within the dispersion domain. In this case, there is nothing to do. The framework will not run an empty forecast. Generally, these cases result in an empty tarball being uploaded to the ftp server which serves as an indicator to the web manager that no forecast is present.

Missing meteorological data

If MM5data complains of missing MM5 data in run.log, the first thing to check is that the /bluesky/bluesky_nested/input/met directory contains links to the appropriate MM5 data files. If the links are missing, they should be made (run Bluesky_nested.pl -l) and re-run the framework. If the links are broken, the data is not located on the cluster fileserver. In that case, you need to contact the systems administrator at UBC to arrange for data to be made available. If the links appear fine but things still do not work, it is likely due to either a network file service error or corruption in the MM5 data files. In either case, contact the UBC systems administrator for assistance.

Hardware failure

This will not be directly detectable by BlueSky. If no output or archive tarballs exist, nothing has been uploaded to the ftp server, and the framework does not appear to be running on the computer, then you can restart the framework and observe it as it runs. If it appears to quit for no reason, and the run.log just cuts off without the appropriate completion notices, chances are there is some sort of hardware problem. By the same token, if you are unable to access the computer or it suddenly becomes non-responsive, hardware problems are likely occurring. In this case, contact the UBC systems administrator for assistance.

12. Support Information

For additional help on specifc components of the BlueSky framework, contact Sean Raffuse at Sonoma Technologies, sraffuse@sonomatech.com.

The 2011 Western Canada BlueSky System: A Framework for Forecasting Wildfire Smoke

PNWIS Conference Nov. 10, 2011

Steve Sakiyama
British Columbia Ministry of Environment
Victoria, BC Canada
Steve.Sakiyama@gov.bc.ca

George Hicks
University of British Columbia
Vancouver, BC Canada
ghicks@eos.ubc.ca

Partners and Acknowledgements

BC Ministry of Environment – C. Jenkins, S. Sakiyama, G. Okrainetz
Alberta Environment - L. Cheng, D. Lyder
Fraser Basin Council – CLEAR fund
University of British Columbia – R. Stull, M. Brauer
Natural Resources Canada – K. Anderson, P. Englefield, S. Taylor
Sonoma Technologies Inc. – S. Raffuse, K. Craig
US Forest Service – S. Larkin, R. Solomon
BC Ministry of Forests and Range – E. Meyer
Alberta Sustainable Resource Development - C. Tymstra
Environment Canada - A. Pankratz, B. Wiens, R. Vingarzan

Why do we need wildfire smoke forecasts?

- Western Canada can experience thousands of wildfires.
- Wildfire smoke covers large areas and affects distant cities and people.

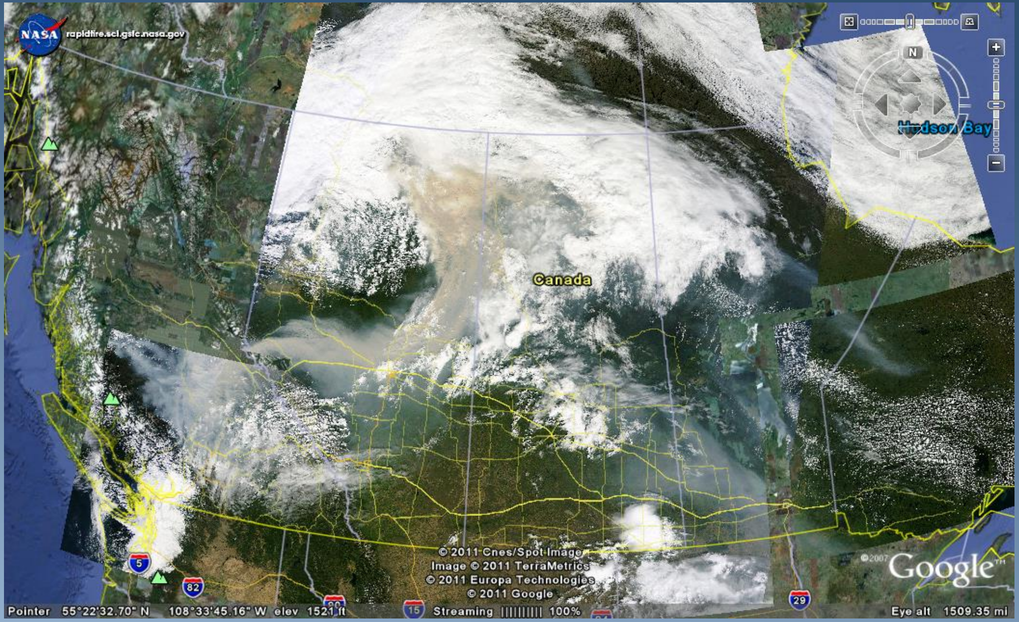

August 19, 2010 Satellite image showing extensive smoke coverage from wildfires – affecting the lives of millions.

Image courtesy: NASA MODIS Image/Google Earth

Why do we need wildfire smoke forecasts?

- Public safety
  (health alerts/advisories, evacuations)
- Health Authorities
  (advisories, evacuations)
- Transportation (visibility)
- Tourism (health, visibility, nuisance)
- Fire Management
  (downstream effects, air tanker safety)
- Weather Forecasters
  (current/future location, enquiries)
- Air Quality Management Agencies
  (public alerts, concentrations)

California wildfire, Sept 5, 2008

What is BlueSky?

- BlueSky : a software framework developed by the U.S. Forest Service that links different models and components to produce wildfire and controlled burn smoke forecasts.
- Integrates fire information with models for fuel loading , emissions and combustion, along with a dispersion/ trajectory model and weather forecast model
- Western Canada implementation for 2010 (BC and Alberta)
  - Fire Characteristics: fire locations (via satellite detection), fuels - Canadian Wildfire Information System (CWFIS)
  - Transport and dispersion - HYSPLIT
  - Weather forecast model - MM5 (4 km grid spacing)

Western Canada BlueSky Operational for the 2010 Wildfire Season

Hourly Meteorological Forecast: MM5 Weather Forecast Model, University of British Columbia (UBC) in Vancouver, B.C.

Daily Wildfire Location and Fuel Consumption: Canadian Wildland Fire Information System (CWFIS), Northern Forestry Research Centre in Edmonton, Alberta

HYSPLIT and BlueSky Framework: Framework links pieces and produces smoke transport and dispersion forecast (UBC)

Web Output: Animations up to 60 hours into the future of forecast hourly $PM_{2.5}$, *ground-level concentrations* for BC/Alberta available from www.bcairquality/bluesky

Improving BlueSky for 2011

In 2010, Western Canada BlueSky produced daily wildfire forecasts for BC and Alberta - wide public/stakeholder interest

Research support for improvements for the 2011 fire season:

- Increase the forecasting domain to cover more of the Yukon, NWT, Western and Central Canada, and Northern U.S.
- Increase persistence of fire emissions (just 24 hours in 2010)
- Increase frequency of forecasts to leverage up-to-the-minute fire data and better understand the impacts of smoke carryover from previous forecast.

Increasing the Domain

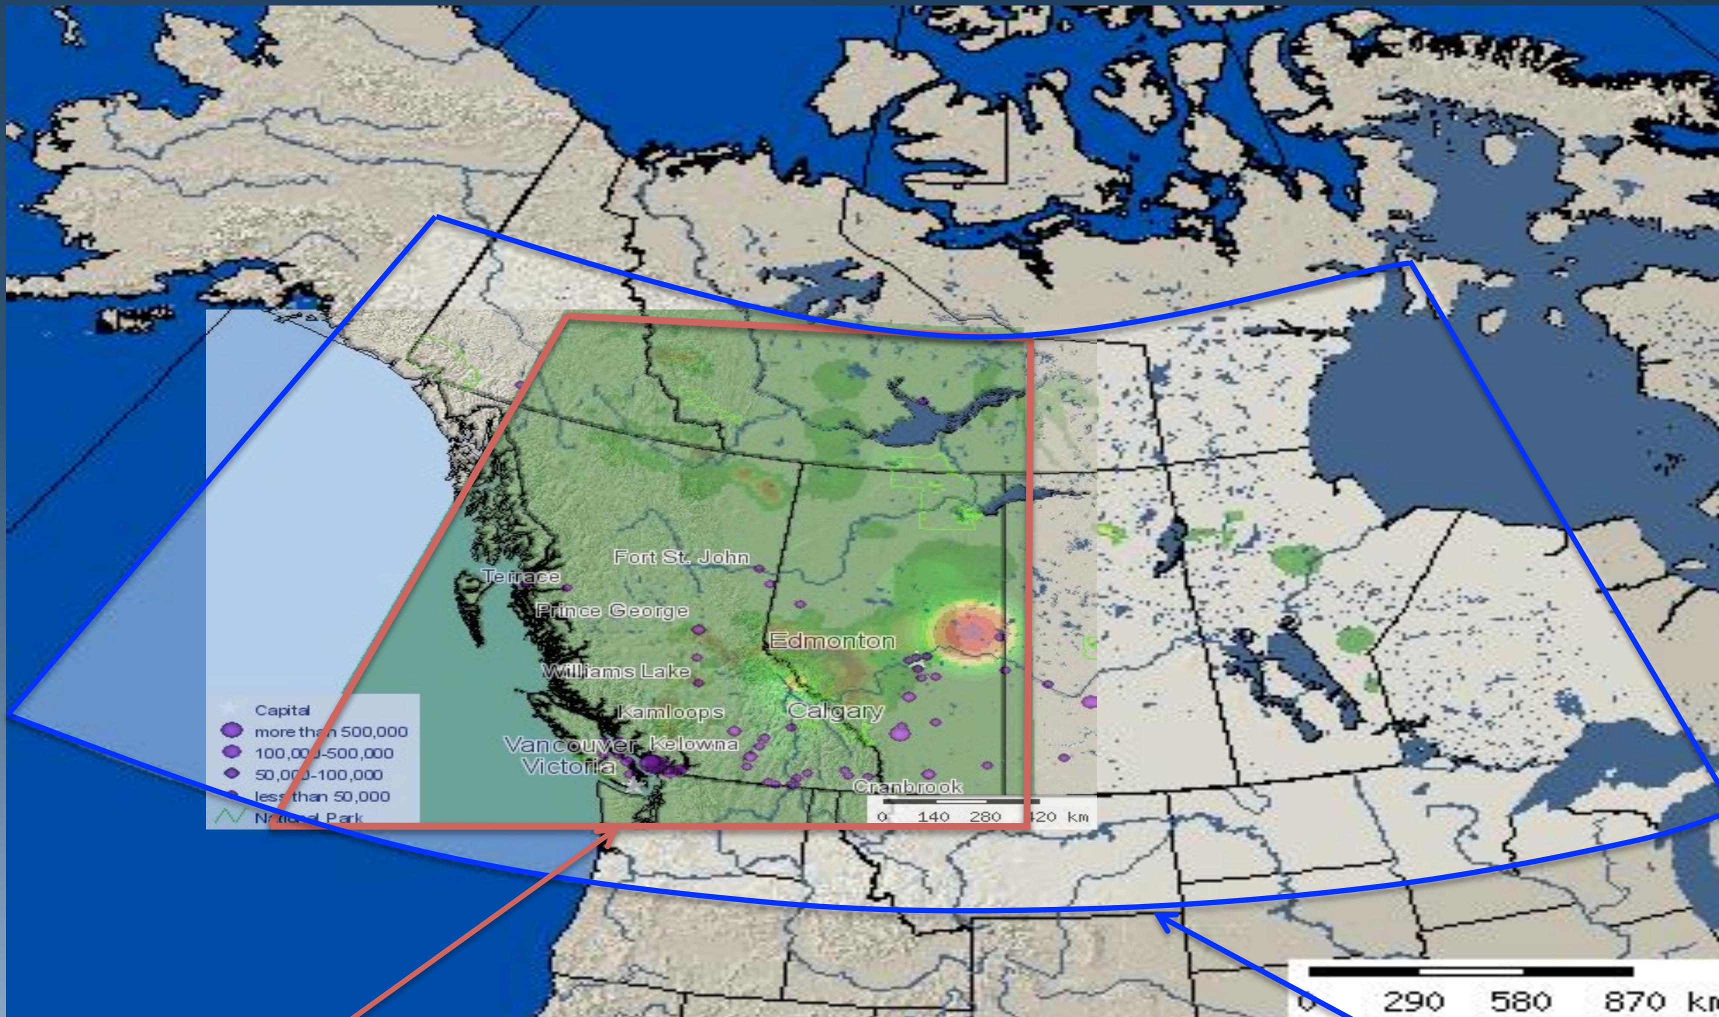

HYSPLIT Domain 2010

HYSPLIT Domain 2011

Accommodating the Increased Spatial Domain

- Required increasing the MM5 meteorological domain, the HYSPLIT dispersion model domain, and adding CWFIS (fire information) coverage for the Northern U.S.
- Nested Grid approach: Meteorological 4km grid for BC and Alberta nested within a larger 12km grid mother domain to cover the remainder .
- Framework adjusted to meld meteorological data from nested domains to act as input for the HYSPLIT model.
- Doubling the HYSPLIT domain more than doubled runtime and computational cost (due to increased fire events).
- New and improved computational hardware installed: offset increased computational cost.

Increasing Fire Persistence

- Wildfires now burn throughout the 60 hour forecast period.
- Adding persistence increased realism for most fire events.
- Increased computational cost in the dispersion model was offset by improved hardware.

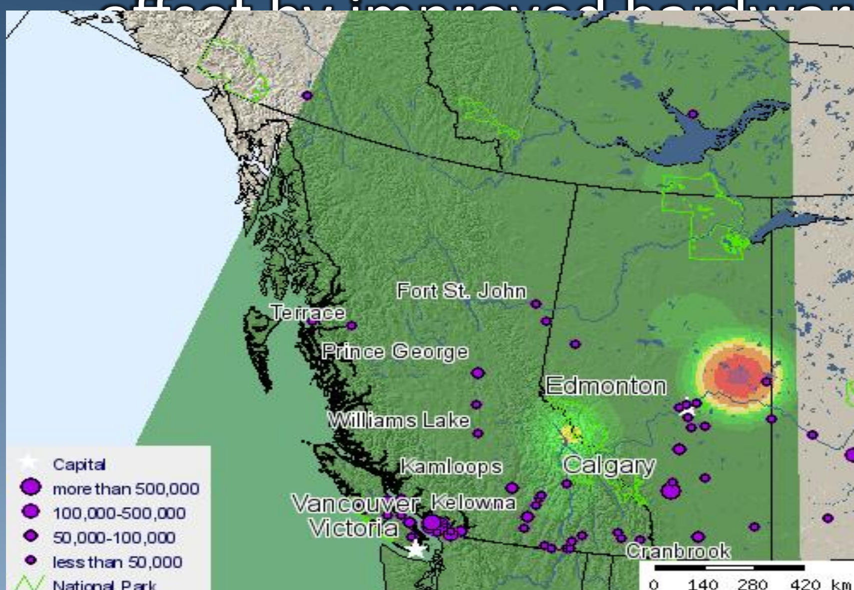

Aug 5, 2010 forecast (59th hour) without persistence.

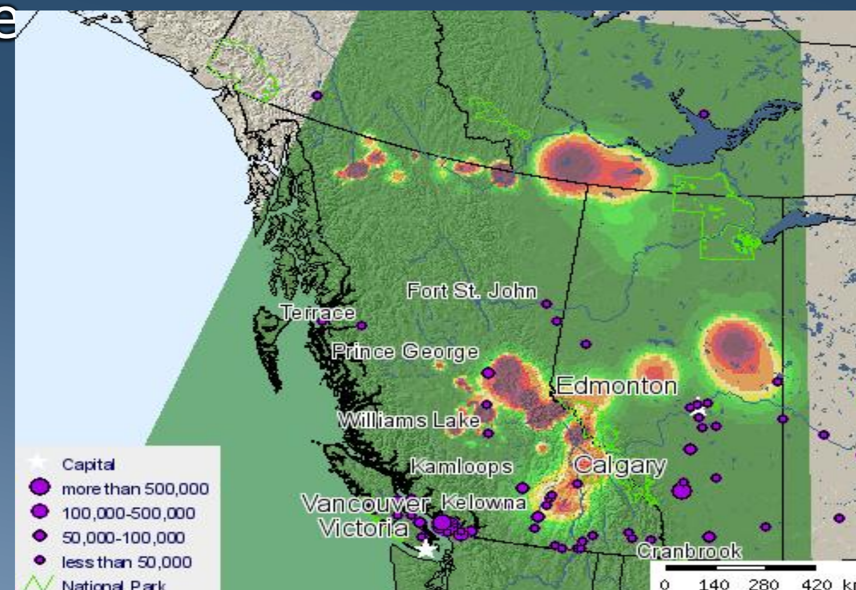

Aug 5, 2010 forecast (59th hour) with persistence.

More Frequent Forecasts

- Original aim to have forecasts available by early morning Pacific Time.
- Required BlueSky start at midnight Pacific Time.
- Leaves available time each day for a secondary run.
- Second run begins at noon and utilizes most recent CWFIS (satellite detection) data.

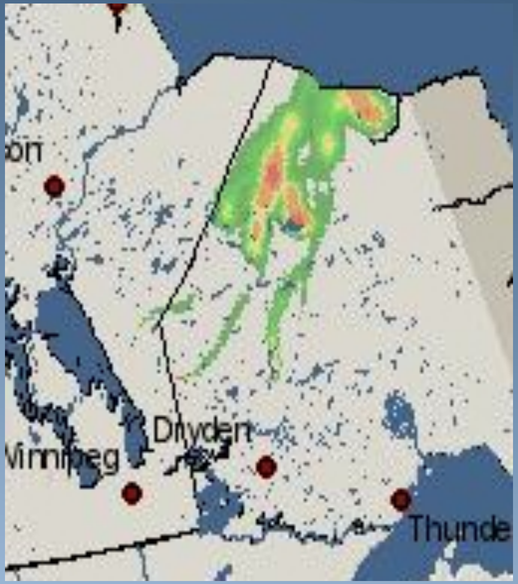

Forecasts for 5am July 16, 2011 (midnight July 15 forecast left, noon July 15 forecast right). Increased fire activity in Ontario caught by second forecast.

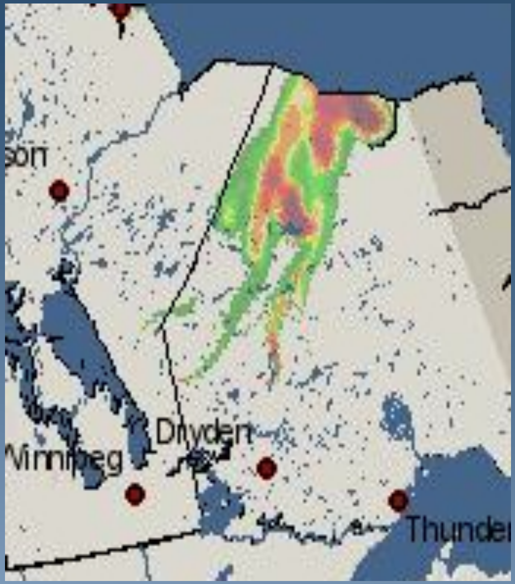

Qualitative Evaluation: Aug 19th 2010
(100's of fires, huge areas impacted by smoke)
Edmonton > 500 ug/m_3 hourly $PM_{2.5}$

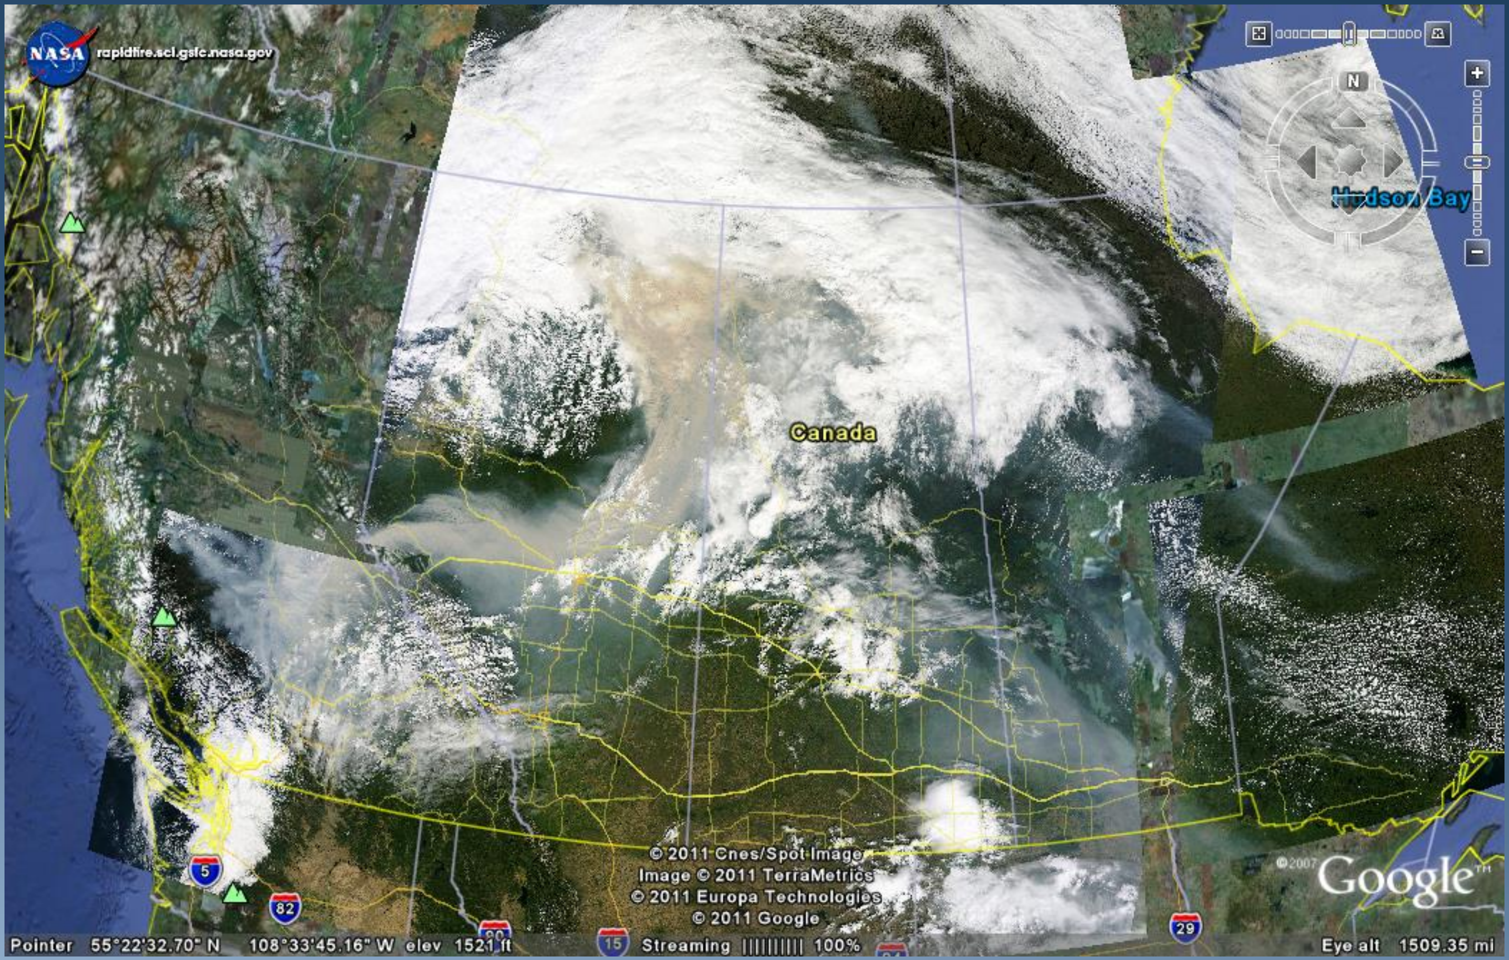

Image courtesy: NASA MODIS Image/Google Earth

Evaluation (Preliminary): Comparison to Measured PM2.5 Observations

- Aug 19th Smoke Event - Time and Space comparison for Maximum Concentration:
- For the corresponding times when the observed maximum occurred at Calgary and Edmonton, BlueSky under-predicted the peak by a factor of 3 (Edmonton) and 10 (Calgary) from fires 500 km away
- However, BlueSky was able to forecast peaks of this magnitude:
  - within a few hours of the actual peak, and
  - within a radius of 100 km of the specific location
- Magnitude predicted but not exactly at the right time and location.

Smoke Forecast issued Aug 18th, showing ground level PM2.5 concentrations for Aug 19th, 1100 PDT
Comparison to MODIS Satellite Image

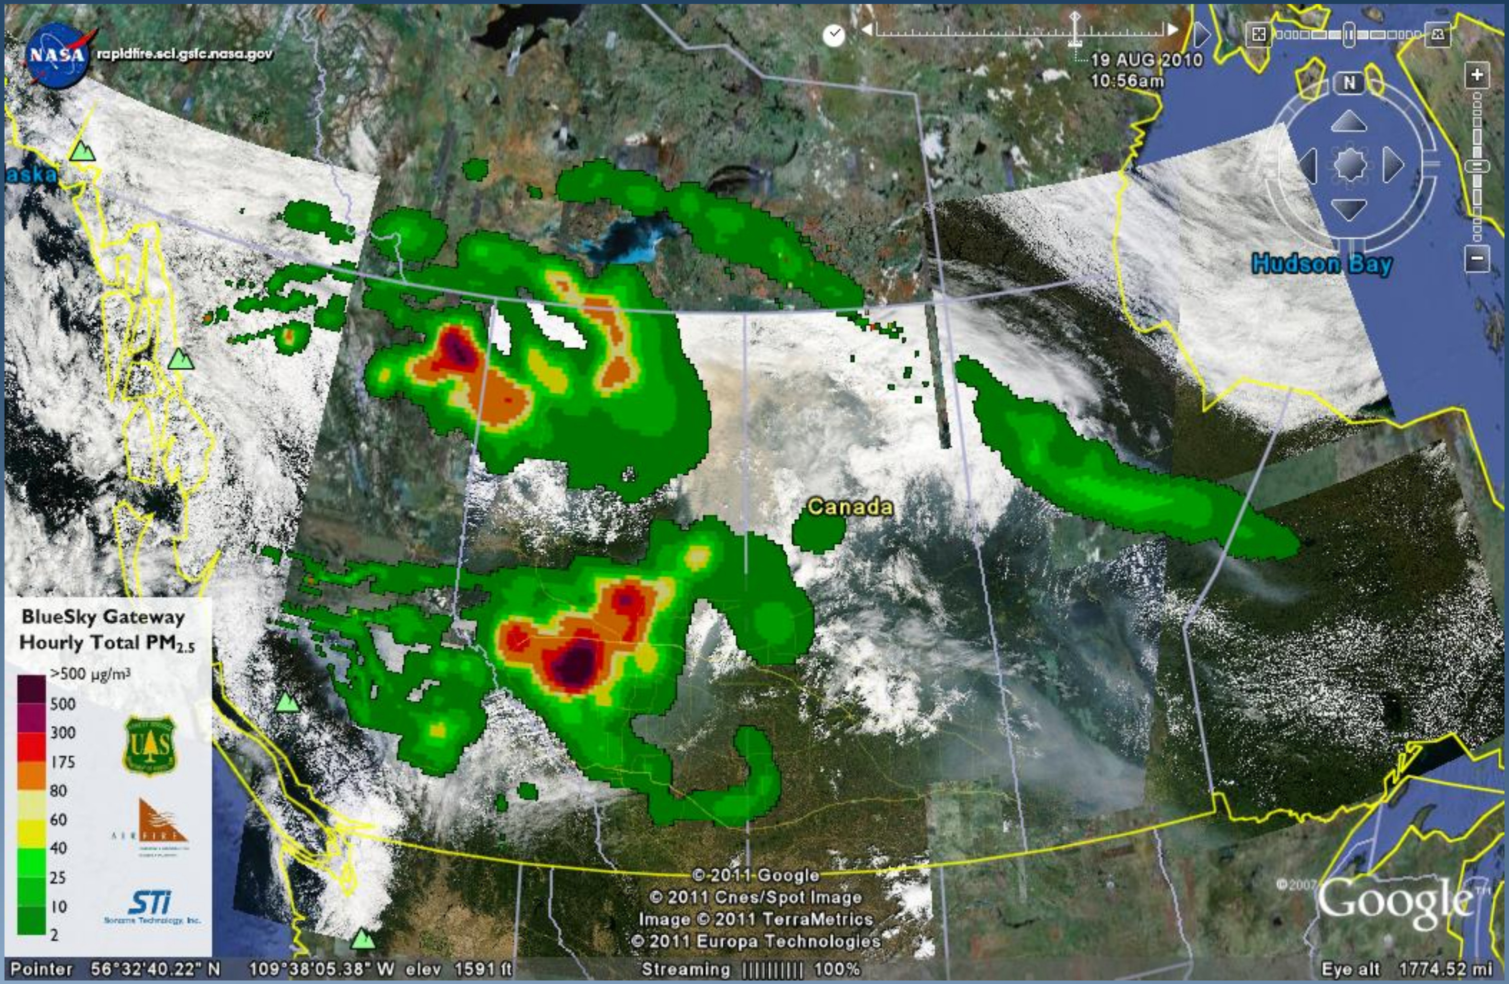

<u>Qualitative:</u> *Forecast smoke distribution* are reasonable

Preliminary Evaluation

- System is producing output consistent with the science and capability of the components.
- Limited test cases with large fire complexes, forecast patterns of smoke impacted areas (smoke footprint) consistent with satellite images.
- Timing and location of peak hourly PM2.5: “Close but no Prize” - can’t expect precision in timing and location of peaks at these distances.
- Evaluation continues.

Possible Improvements for 2012

- Extend domain further east, capture Ontario entirely (interest from Ontario for a similar tool)
- Group nearby wildfires to reduce computational load on dispersion model
- Incorporate carryover smoke from previous forecast period (evaluation indicates that this will improve predictions)
- Speed up run times and remove occasional "bulls-eye" effect.
- HYSPLIT is the most computationally expensive part of the framework. Can the latest version of HYSPLIT (version 4.9) solve some of these problems?

Evaluating HYSPLIT v4.9

Forecasts for 5am Aug 18, 2010

HYSPLIT v4.8

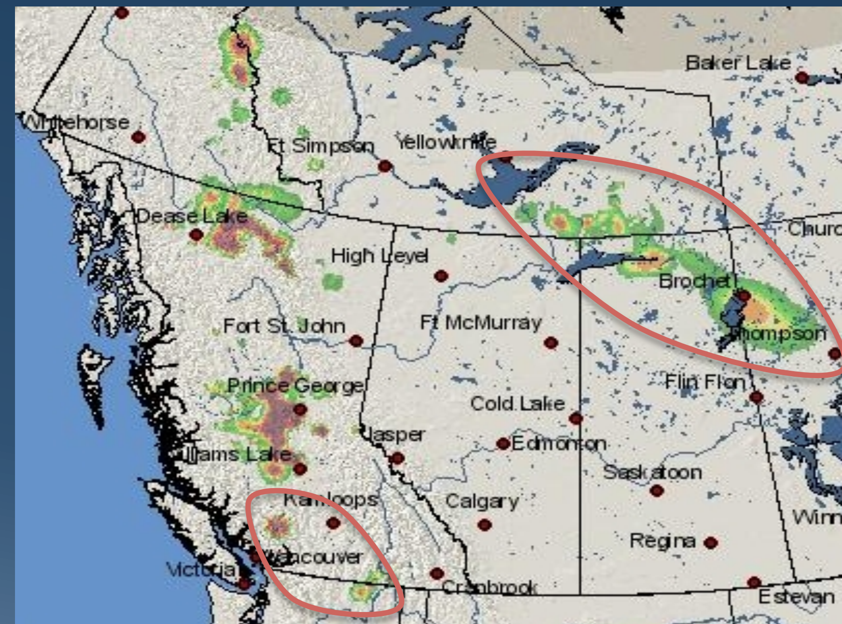

HYSPLIT v4.9

Both forecast were initialized from the same fire and meteorological data. Dispersion is greater in v4.9 but peak concentrations are lower (not apparent in image). Some additional fire activity is caught by v4.9 in southwest B.C. not originally caught by v4.8.

Evaluating HYSPLIT v4.9

- Significantly more complex schemes.
- Dispersion is more diffuse and peak concentrations are lower.
- “Bulls-eyes” mostly removed but ephemeral nature suggests a problem with emissions profiling.
- Computational cost is 5-10 times greater than v4.8, increasing exponentially with additional fire events.
- Parallelization of portions using multi-core processing allows for computational cost to be offset.

Summary: Western Canada BlueSky 2011 Version and Next Steps

- Expanded domain facilitated by nested meteorological grid and expanded fire spatial information
- Increased source persistence provides more realism.
- Twice a day forecasts for up to date wildfire information – useful for rapidly developing fire situations.
- New hardware to offset increased computational costs
- Forecasts consistent with the science built into the models , and producing smoke footprints consistent with satellite images
- Peak magnitudes are consistent with observations but details (exact timing and location) are not.
- Plans to continue evaluation, expand domain, include previous day(s) smoke and improve efficiencies in run-times with source grouping and the parallelization capabilities in HYSPLIT 4.9.

Web Links

1. BlueSky Smoke Forecasts (Updated Daily)
   – http://www.bcairquality.ca/bluesky

2. Canadian Wildfire Information System (CWFIS) – fire location, fire danger rating
   – http://cwfis.cfs.nrcan.gc.ca/en_CA/index

3. Current Wildfires (Status, Location) in British Columbia
   – http://bcwildfire.ca/hprScripts/WildfireNews/Fires.asp

Steve Sakiyama
British Columbia Ministry of Environment
Victoria, BC Canada
Steve.Sakiyama@gov.bc.ca

George Hicks
University of British Columbia
Vancouver, BC Canada
ghicks@eos.ubc.ca