

Running Puffin on Vhub

Introduction

Puffin is a coupling of the programs Puffin (also known as Bent) and Puff. Puffin is a theoretical model of a volcanic plume based on applying the equations of motion in a plume-centered coordinate system. Puff simulates the transport, dispersion and sedimentation of volcanic ash. Puff also prepares some scripts to turn the cdf files generated into images. [ImageMagick](#) then creates an animated gif using these images, so we can see how ash spreads over time.

Required Preparations

The tools we use require large databases, too large to be stored on Vhub. For this reason, you will need to mount onto the irods database in order to be able to run Puffin. Additionally, for large runs you have the option of sending larger jobs to the cluster.

First of all, you need to join the **network** group. By default, the hub does not allow outgoing network access. You will need to request an administrator to be added, and then it may take an hour for access to be granted.

To set up irods, look up the topics page for [installing irods](#) to get details on how to set up irods. For the purposes of running Puffin, I recommend having the irods directory be in your home directory and to stay logged in as the public user. All necessary databases are in the public folder.

Example Run

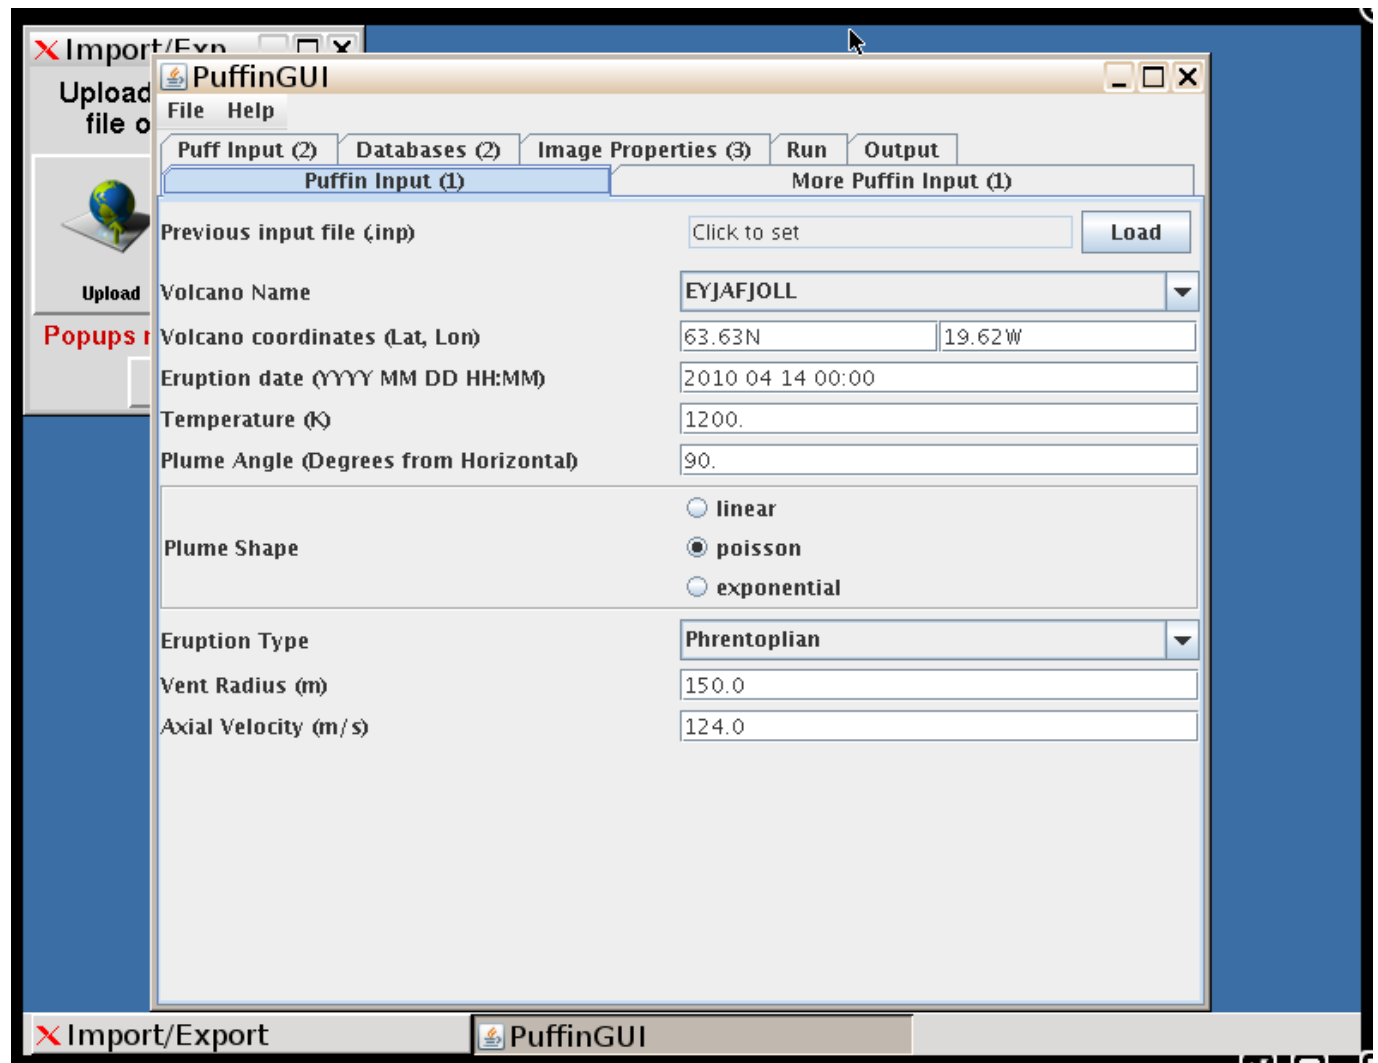
When you start Puffin, most of the entries will start out as filled. The only entries you may need to change for a test run are the locations of the reanalysis databases on the Database tab. If you set up irods to be in your home directory, it will automatically set up the locations in the gui.

In order to run the test run, you just need to press the three buttons on the **run** tab in order. The first button runs puffin, which fills out some information about the plume on the puff input tab. The second button runs puff, which will create the cdf files representing the ashes. The third button creates the animations, which will then display on the **output** tab. You can view these various animations by clicking the buttons on the bottom of the output tab.

The prepared run simulates a volcanic plume in Iceland. The sample runs only 10000 particles over a 72 hour period, so it runs rather quickly. Running puffin should take less than a minute. Running puff should take less than five minutes. Creating the animations should take less than a minute.

Discussing each tab in more detail:

Puffin Input



PuffinGUI

File Help

Puff Input (2) Databases (2) Image Properties (3) Run Output

Puffin Input (1) More Puffin Input (1)

Previous input file (.inp) Click to set Load

Volcano Name EYJAFJOLL

Volcano coordinates (Lat, Lon) 63.63N 19.62W

Eruption date (YYYY MM DD HH:MM) 2010 04 14 00:00

Temperature (K) 1200.

Plume Angle (Degrees from Horizontal) 90.

Plume Shape

- ☐ linear
- ☒ poisson
- ☐ exponential

Eruption Type Phrentoplion

Vent Radius (m) 150.0


Axial Velocity (m/s) 124.0

Import/Export PuffinGUI

All the entries here are pretty self explanatory.

If you have the .inp from a previous run, you can load the information from a previous run. Just select the .inp file, and press load. All relevant entries will be updated.

More Puffin Input



The screenshot shows the PuffinGUI application window. The title bar is 'PuffinGUI'. The menu bar has 'File' and 'Help'. The tab bar has 'Databases {2}', 'Image Properties {3}', 'Run', and 'Output'. The 'More Puffin Input {1}' tab is selected. The 'Puff Input {2}' tab is also visible. The 'Radiosonde File' section has a 'File Location' text box with the path '/home/vhub/alejandro/puffin/data/Apr17_12Zcop'. The 'Radiosonde Data Type' section has two radio buttons: 'Height (m) Temp (C) Windspeed (m/s) Density (kg/m^3)' (selected) and 'Height (m) Temp (C) Windspeed (knots) Pressure (hPa)'. The 'Column Numbers' section has four text boxes with values 4, 5, 7, and 10. The 'Generate Radiosonde from NWP' checkbox is unchecked. The 'Water Fraction (% by weight)' text box has the value 0.017. The 'Particle Density (kg/m^3)' text box has the value 1000. The 'Particle Shape (1=circular, <1 = oval)' text box has the value 0.8. The 'Random Grain Mean' checkbox is unchecked. The 'Custom Grain: Overrides volcano type defaults' section has two text boxes: 'Grain Size Mean (phi scale)' and 'Grain Size Standard Deviation (phi scale)'. The 'Display Plume' checkbox is unchecked.

In order to simulate how the ash will move up into the atmosphere, you need to give information about the air conditions, given by a weather balloon file. If you have one already, just mark in where the file is, what are the relevant columns, and what the relevant columns represent. If you do not have one, select generate weather file in order to automatically generate one from the databases. Note that this requires all 3 databases from the database tab to be filled out.

Additionally, you have the choice to specify the mean and standard deviation, or just select grain random.

Puff Input

The screenshot shows the PuffinGUI application window. It has a menu bar with 'File' and 'Help'. Below the menu bar are several tabs: 'Databases {2}', 'Image Properties {3}', 'Run', 'Output', 'Puffin Input {1}', 'More Puffin Input {1}', and 'Puff Input {2}'. The 'Puff Input {2}' tab is currently selected. The main area contains several input fields and checkboxes. 'Ash Output' and 'Grid Output' are checked. 'Run Hours' is set to 72.0, 'Eruption Duration' to 3.0, and 'Number of Particles' to 10000. 'Grid Latitude Range' is set to -14.8819 and 32.8058, 'Grid Longitude Range' to 37.3622 and 68.3775, and 'Grid Height Range' to 0 and 16000. There is a 'Previous Puff run' section with a 'Click to set' button and a 'Load' button. At the bottom, there is a 'Puffin Output {Run Puffin to set}' section with several input fields for 'Plume Max', 'Plume Min', 'Plume Horizontal Width', 'Plume Verticle Width', 'Ash Log Mean', and 'Ash Log Standard Diviation'.

Parameter	Value
Ash Output	<input checked="" type="checkbox"/>
Run Hours	72.0
Eruption Duration	3.0
Number of Particles	10000
Grid Output	<input checked="" type="checkbox"/>
Grid Latitude Range	-14.8819 32.8058
Grid Longitude Range	37.3622 68.3775
Grid Height Range	0 16000
Previous Puff run	Click to set Load
Puffin Output {Run Puffin to set}	
Plume Max	16908.8
Plume Min	12824.5
Plume Horizontal Width	18.514
Plume Verticle Width	2.042
Ash Log Mean	-4.20302
Ash Log Standard Diviation	0.90163

First of all, having ash output selected is very important. With no intermediary ash outputs, no images can be made for the animation.

Second, the more particles you simulate, the more computation time the program will take. Because of this, you are limited in how many particles you can simulation. If running locally, you are capped at 100,000. If you submit to the u2-grid, you are capped at 10,000,000.

Also, note that the bottom six entries will be filled out by running puffin. You do not need to fill them out yourself if you are doing a full run through.

Databases

The screenshot shows the PuffinGUI application window. It has a menu bar with 'File' and 'Help'. Below the menu bar are tabs: 'Databases {2}', 'Image Properties {3}', 'Run', and 'Output'. Under the 'Databases {2}' tab, there are three sub-tabs: 'Puffin Input {1}', 'More Puffin Input {1}', and 'Puff Input {2}'. The main area contains a list of input fields for different databases. The fields are organized into two columns: the left column lists the database names, and the right column contains the corresponding values entered in the text boxes.

Database Name	Value
UWND Model	reanalysis
UWND Directory	/home/vhub/alejandro/irods/puff
UWND Mask {Filename Template}	uwnd.YYYY.nc
UWND VarUname	uwnd
VWND Model	reanalysis
VWND Directory	/home/vhub/alejandro/irods/puff
VWND Mask {Filename Template}	vwnd.YYYY.nc
VWND VarUname	vwnd
Temperature Model	reanalysis
Temperature Directory	/home/vhub/alejandro/irods/temperature_data
Temperature Mask {Filename Template}	air.YYYY.nc
Temperature VarUname	air
Irods Name	public

There are three databases that may need to be set up.

The wind u and wind v databases (which is in miles per hour) are needed to run puff.

The temperature database (which is in Celcius) is only needed for generating a weather balloon file for puffin.

For each database you need to set:

- * MODEL- The name of the model you are using.

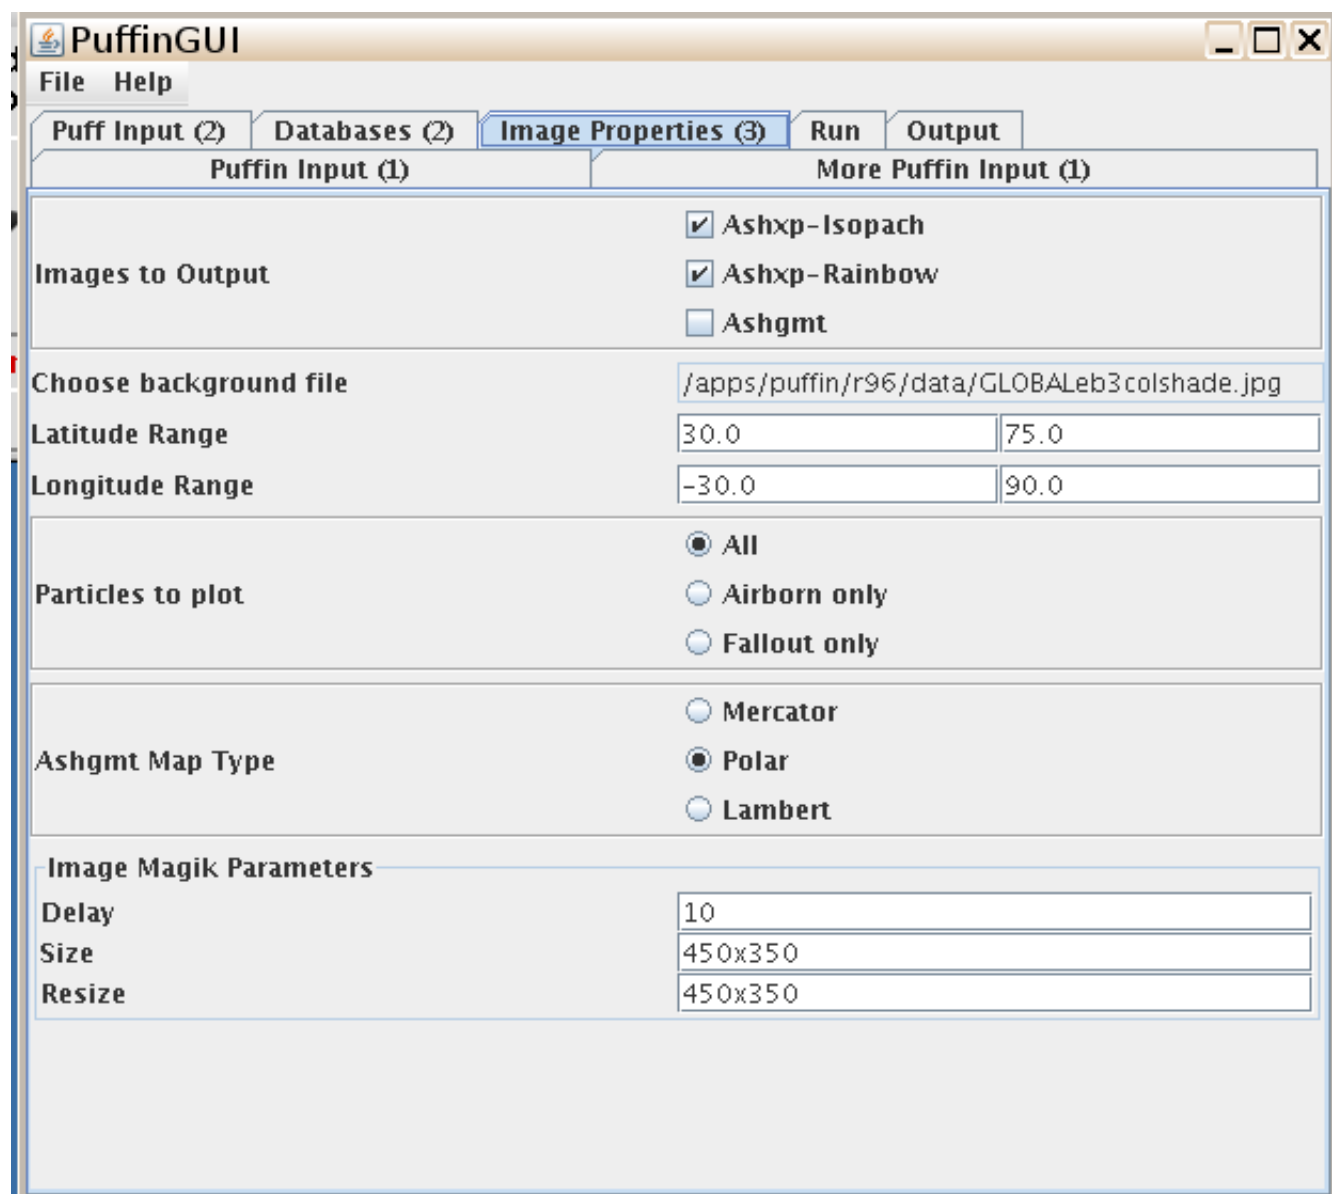
- * DIR- The directory where your nc file is.

* MASK- The filenames of the data files. YYYY represents where the year would go.

* VARUNAME- The name of variable within the nc file that we use. This can be found by running `ncdump -h` on the file.

The irods username is only needed for remote runs. The Irods username is used to convert the directory paths of an irods mount into the directory paths on the cluster.

Image Properties

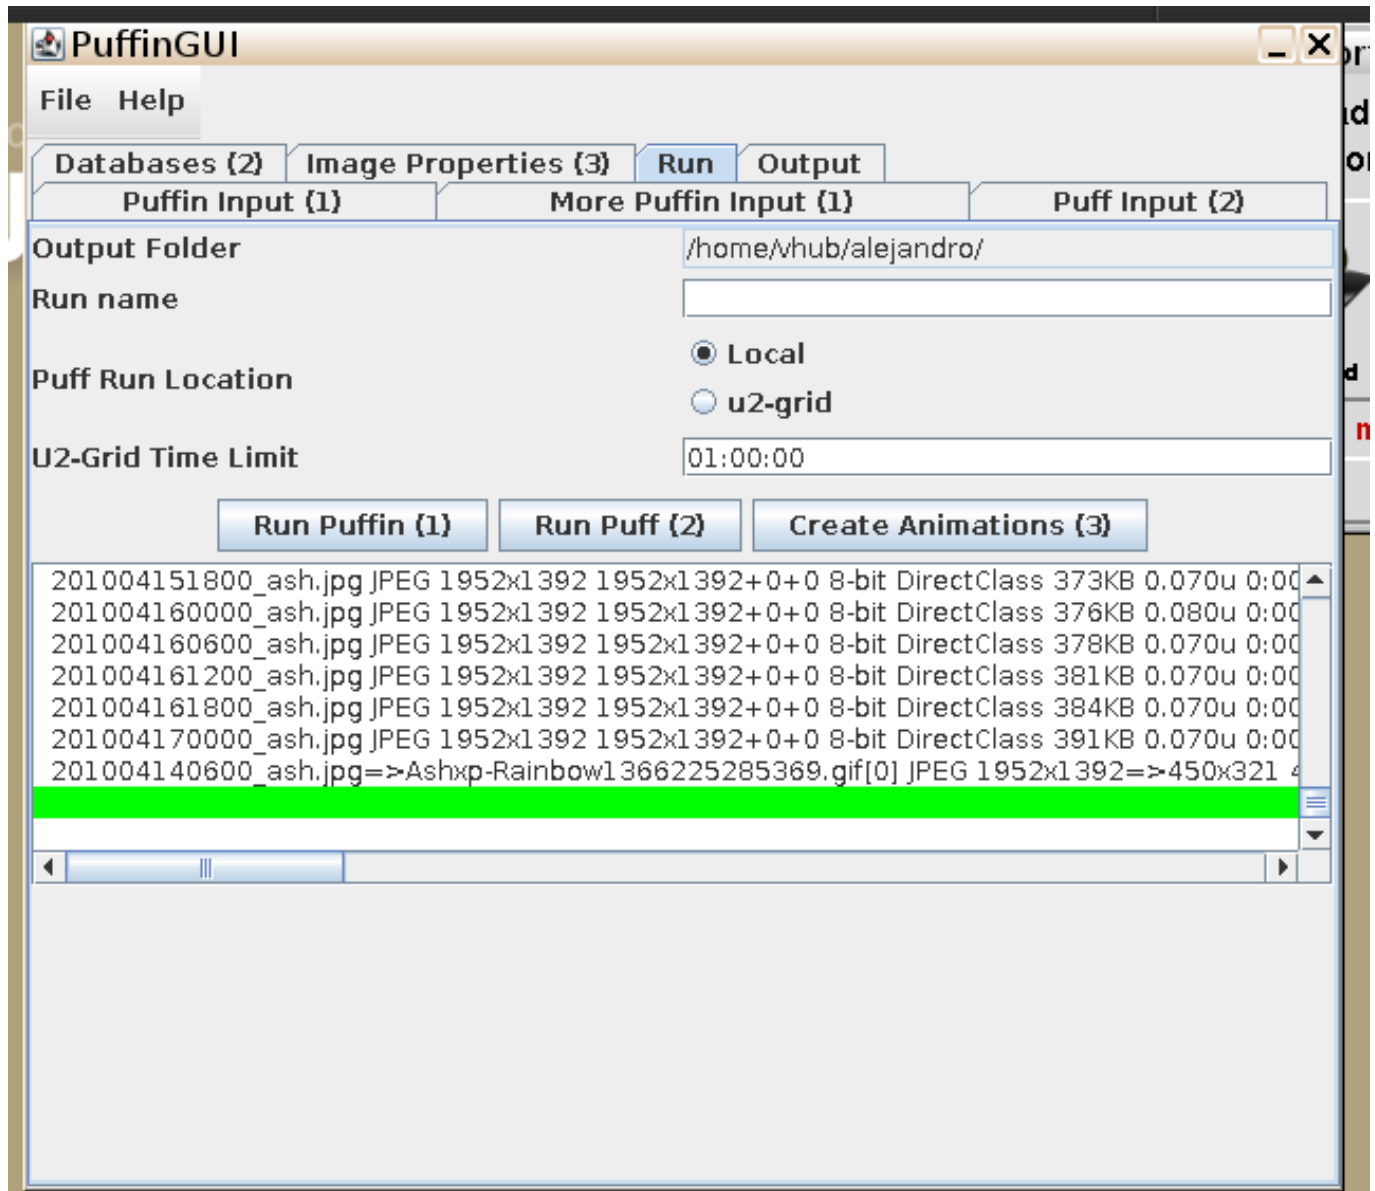


The screenshot shows the PuffinGUI application window with the 'Image Properties' tab selected. The window has a menu bar with 'File' and 'Help'. Below the menu bar are tabs for 'Puff Input (2)', 'Databases (2)', 'Image Properties (3)', 'Run', and 'Output'. The 'Image Properties' tab is active and contains two main sections: 'Puffin Input (1)' and 'More Puffin Input (1)'. The 'Images to Output' section has three checkboxes: 'Ashxp-Isopach' (checked), 'Ashxp-Rainbow' (checked), and 'Ashgmt' (unchecked). The 'Choose background file' section has a text box containing '/apps/puffin/r96/data/GLOBALeb3colshade.jpg'. The 'Latitude Range' section has two text boxes: '30.0' and '75.0'. The 'Longitude Range' section has two text boxes: '-30.0' and '90.0'. The 'Particles to plot' section has three radio buttons: 'All' (selected), 'Airborn only', and 'Fallout only'. The 'Ashgmt Map Type' section has three radio buttons: 'Mercator', 'Polar' (selected), and 'Lambert'. The 'Image Magik Parameters' section has three text boxes: 'Delay' (10), 'Size' (450x350), and 'Resize' (450x350).

Puffin Input (1)		More Puffin Input (1)	
Images to Output		<input checked="" type="checkbox"/> Ashxp-Isopach <input checked="" type="checkbox"/> Ashxp-Rainbow <input type="checkbox"/> Ashgmt	
Choose background file	/apps/puffin/r96/data/GLOBALeb3colshade.jpg		
Latitude Range	30.0	75.0	
Longitude Range	-30.0	90.0	
Particles to plot	<input checked="" type="radio"/> All <input type="radio"/> Airborn only <input type="radio"/> Fallout only		
Ashgmt Map Type	<input type="radio"/> Mercator <input checked="" type="radio"/> Polar <input type="radio"/> Lambert		
Image Magik Parameters			
Delay	10		
Size	450x350		
Resize	450x350		

The images to output button selects which images to create. Warning: Ashgmt currently requires some files that are not on Vhub. Ashgmt will not be able to create an image.

Run

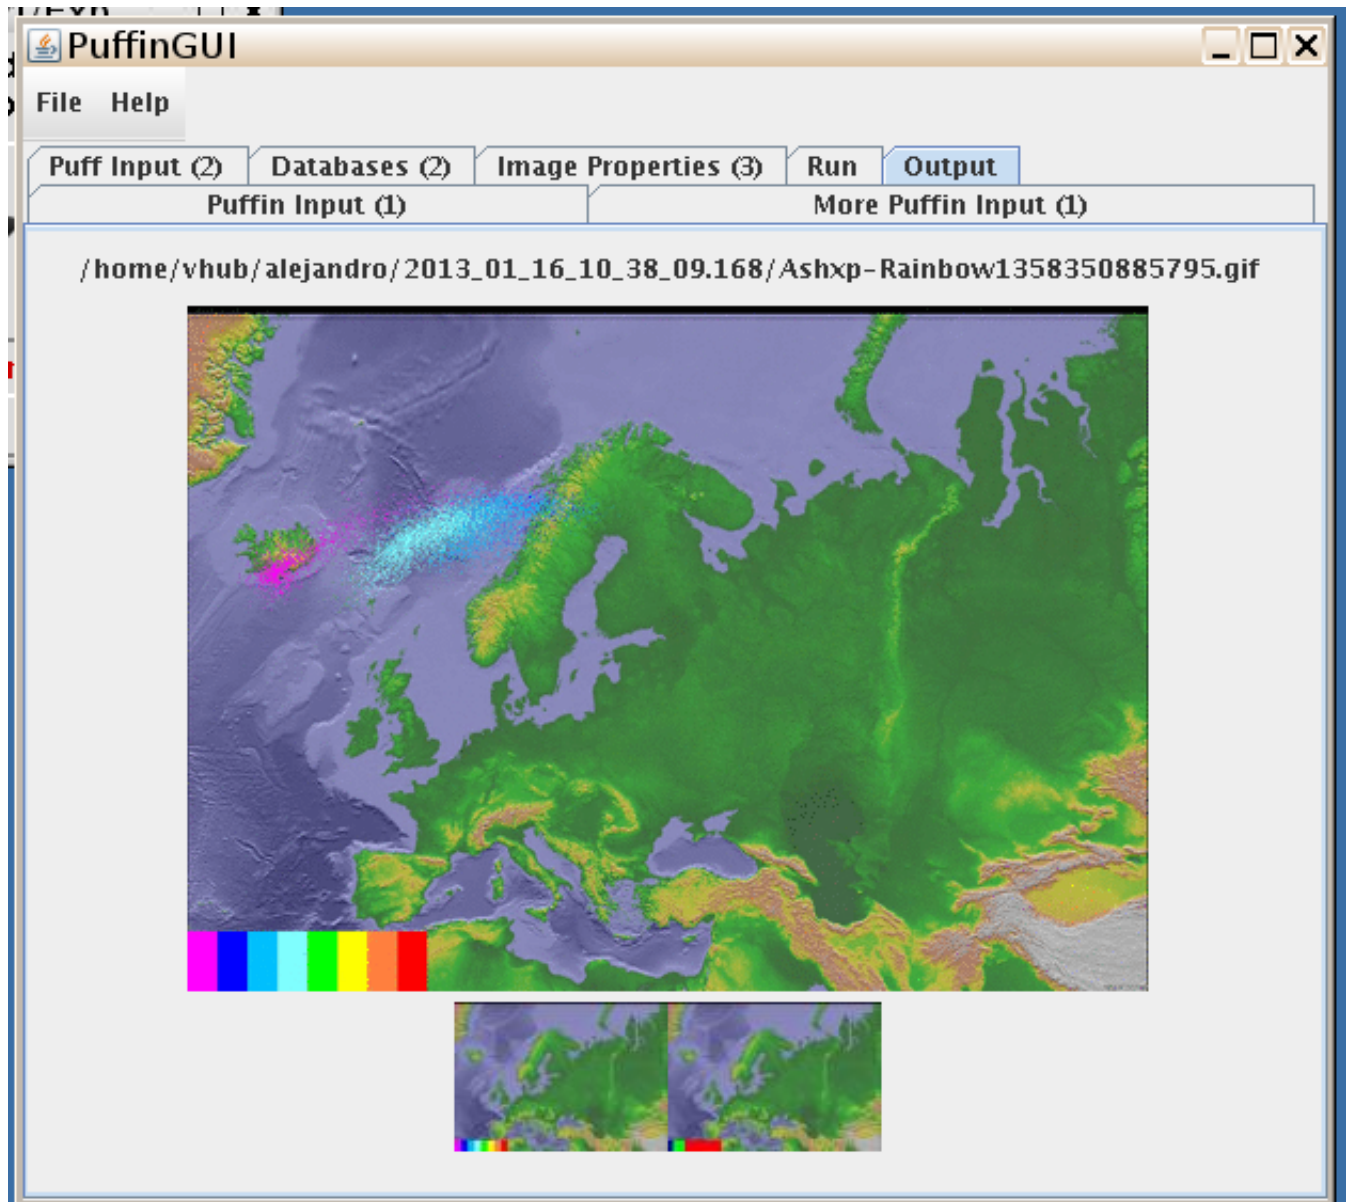


First of all, you specify where you want to output in the output folder, and whether you want to create a new folder to keep the information in with run name. Every time puffin run, a new folder will be created based on the current time.

Normally, you will want to press run puffin, then run puff, then make animations.

You have the option to send the job to the cluster. Puffin runs (when you hit the first button) will always be done locally. You can send Puff runs (second button) and the animation runs (third button) to the cluster.

Output



If successful, there will be images on the output tab after make animations is run. Click the thumbnails on the bottom in order to cycle through the images made.

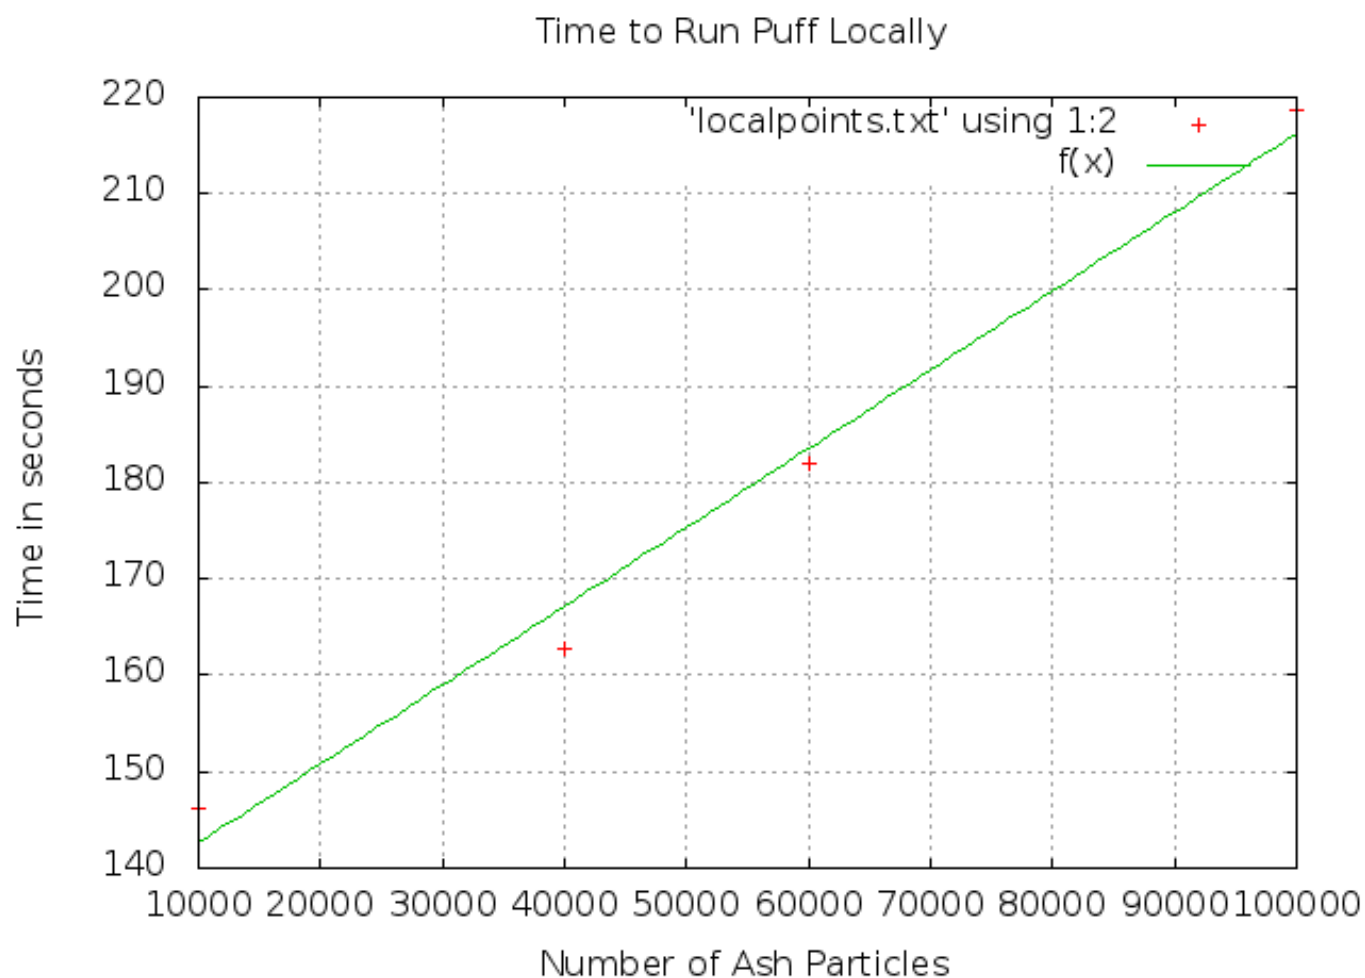
Run Time

Running Puffin should take seconds, even if you are outputting the plume information.

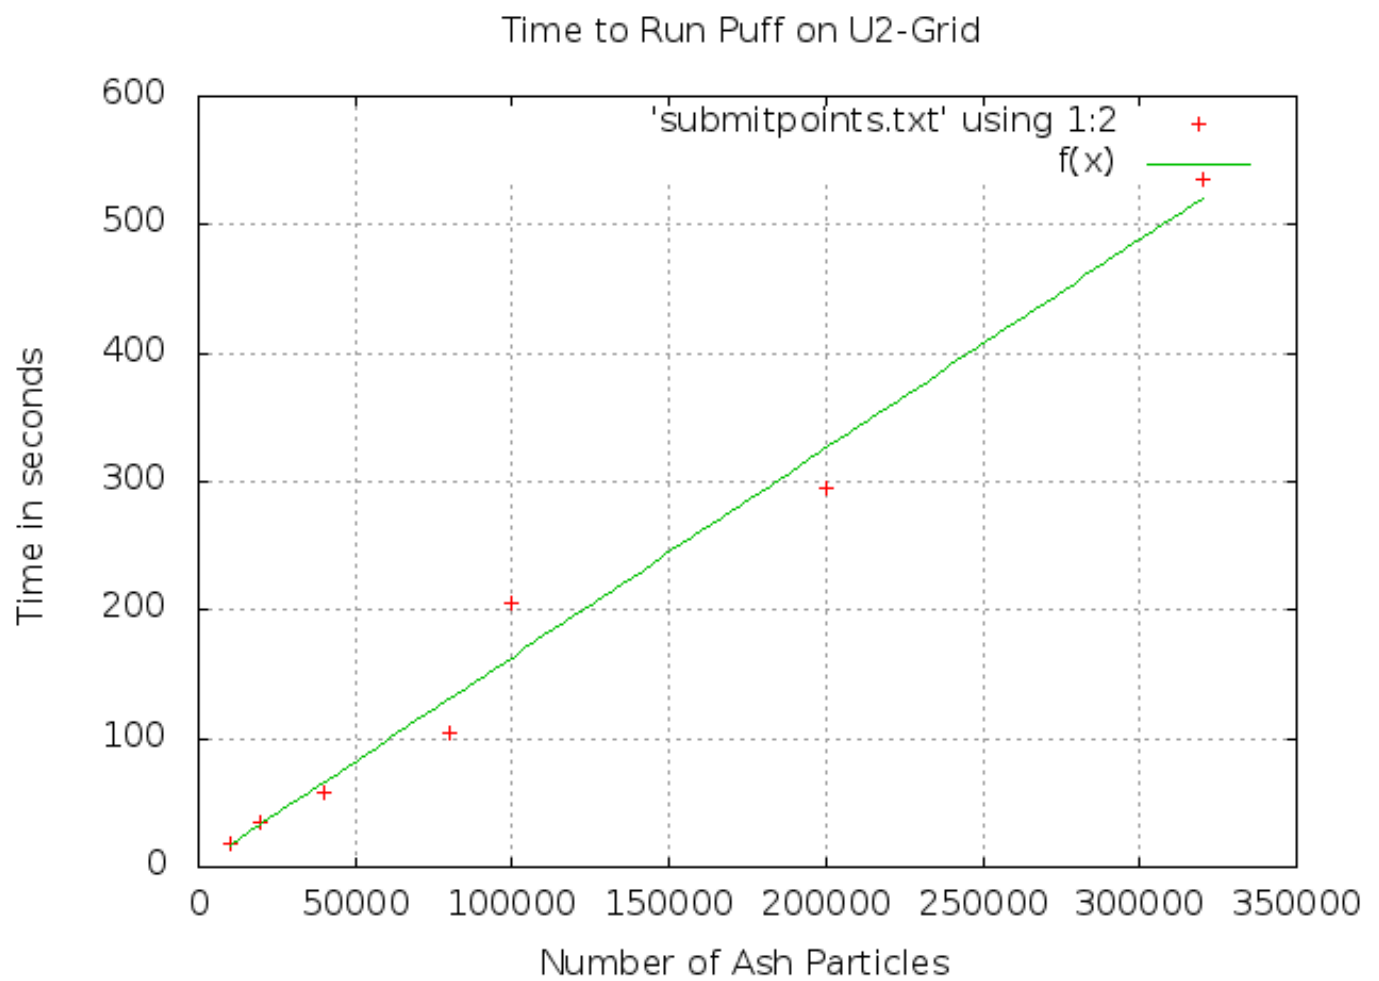
If you generate a weather balloon file, it will take considerably longer because it has to load the databases. The entire calculation should take on the order of 4 minutes.

Running Puff is much more computationally intensive, and takes time dependent on how many particles you want to simulate over how much time. You are restricted to only using 100,000 particles locally and 10,000,000 on the hub. Running locally, loading the databases takes about 2 minutes and most computations will take at most 2 minutes. From running several runs, here

is a graph of how much time you can expect your runs to take.



When you submit puff to u2, since databases don't have to be loaded the run will be a lot faster. However, often you will have to wait a couple minutes before your job is queued onto the server. As a result, runs on u2 tend to take just as long as local runs. Here is a graph of the computation time for some larger runs (but note that the time to queue is not included).



Creating the animations should also take on the order of a minute.