U. S. Department of Commerce
National Oceanic and Atmospheric Administration
National Weather Service
National Centers for Environmental Prediction
5830 University Research Ct.
College Park, MD 20740

Technical Note

Automatic generation of land masks and bathymetry for satellite analysis and model use[†].

Robert Grumbine [‡]
Environmental Modeling Center
Marine Modeling and Analysis Branch

November 8, 2012

[†] MMAB Contribution No. 301.

[‡] e-mail: Robert.Grumbine@NOAA.gov

This page is intentionally left blank.

1 Abstract

This document describes the programs used for geographic grids and processing used by sea ice and sea surface temperature analyses, and for high resolution ocean modeling, at NCEP. This includes generating land masks, bathymetries, identifying lakes for further more detailed analysis of satellite data where gridding is insufficient, and determining distance from land for use in, so far, conducting satellite sea surface temperature analysis.

2 Introduction

Information about the earth's geography is important in a number of ways when conducting satellite analyses or modeling. One routine situation is in the construction of land masks. This is complicated by the fact that the earth does not reside on a grid, so that grid cells can be partly land and partly water. Further, when we turn to constructing a bathymetry for an ocean model, the depths of water, which are agreed to be water by both the bathymetry and coastline information, maybe be insufficient for the numerical model to treat it as water. Again, it can also happen that the two data sets are in conflict and need to be reconciled.

The programs and methods described here started from a desire to be able to construct a bathymetric field and land mask for a high resolution ocean model on a curvilinear grid. Further, it was deemed desirable that the land mask and bathymetry be in mutual accord, and that the land mask be as high resolution as possible. Since the bathymetric data sets were pre-gridded at 2' resolution, and the model's resolution was as fine as this in some areas – but on a curvilinear grid not necessarily aligned with the cells the bathymetry were on – interpolations of the bathymetry could easily create or destroy water areas in ways inconsistent with a more faithful knowledge of where the coastlines are. Hence this approach, to ensure that land was masked where the highest resolution information (as fine as 200 m) said it should be, and ocean was placed where ocean should be. The original programs for this were released in August, 2003. The RTOFS-Atlantic model using this approach was implemented operationally in December, 2005.

A related approach was taken by Chawla and Tolman [2007] to construct grids for wave modelling purposes, and to determine partial obstruction by islands.

This basic program, minus bathymetric concerns, has also been applied to generating land masks on latitude-longitude and polar stereographic grids for satellite sea surface temperature and sea ice analysis and the results will be discussed here. Such masks have been used in NCEP operations since August 2006.

Finally, two derivative program bases have been developed for addressing additional satellite concerns. Satellite sensors have a 'footprint' over which they

observe the earth. This is an area to which they are sensitive to terrestrial conditions. The size of the footprint depends on the sensor and wavelength used, and varies from 1.1 km for AVHRR [Goodrum et al., 2000] in local area resolution to about 50 km for SSMI in the 19 GHz channel [Raytheon, 2000]. There are further complexities arising from the fact that the footprint is typically not a perfect circle. But our initial question is: given a satellite observation with a known center location and known maximum distance from center over which the observation responds to terrestrial conditions, is the observation far enough from land (speaking as oceanographers, generalize to 'different surface type' for others) that our observation is not contaminated? It is entirely possible that even though an observation falls inside a grid cell that is labelled 'water', it is too close to the shoreline for the information to be a straightforward representation of the water conditions.

A more complex issue arose with the Real Time Global Sea Surface Temperature analysis [Gemmill et al., 2007], where the physical retrieval process near land was returning poor quality sea surface temperatures. This is currently considered to be a result of the atmospheric model being much coarser than the sea surface temperature analysis grid, and therefore considering as land points that the analysis considers to be water. The method used to address this was to use different retrievals for points within 35 km (approximate resolution of the global atmospheric model at the time) of land. And, since the distances were taken from a grid, and because it's desirable to have the analysis vary smoothly, the other retrievals and physical retrievals are blended with linearly varying weight from all navy retrieval at 35 km to all physical retrieval for points 200 km or greater from shore. This was implemented in NCEP operations in February, 2010 [MMAB, 2010].

3 Satellite Land Mask Generation

This is the simpler situation though it occurred later in history – constructing a land mask file for satellite product use. For this, we can examine only the coastline curves and decide at each grid point whether the point is ocean, coast (boundary), land, or inland water.

Information about the location of coastlines is taken from the GSHHS data set [Wessel and Smith, 1996]. The version used, to date, is the version 1.2 release of 18 May 1999, retrieved in February, 2003. The data set includes the latitude-longitude locations of points on polygons enclosing surface types. The types are 1: land, 2: water, 3: land that is inside a type 2 polygon, 4: water that is inside land that is inside a type 2 polygon. The global ocean is those points which are not enclosed by any of these types. The polygons are available on a range of pre-selected resolutions, from the full resolution (down to 200 m between points) to crude resolution (approximately 25 km).

A winding number algorithm [Sunday, 2001] is used to determine whether a given point is inside or outside the polygon at hand. With large polygons, such as the Old World (Europe, Africa, and Asia are contained in a single polygon), at high resolutions, this can be very time consuming as there are over a million points to consider for each grid point that we wish to determine whether it is inside or outside.

To construct the first-guess land mask, the program *refill* reads each polygon and then tests each point on the target grid for whether it is inside the polygon or not. As some polygons are quite small, the extreme north, south, east, west points on the curve are used to pre-screen points for consideration. As the target grid may not be a simple latitude-longitude grid, the actual extreme values are increased by 2 degrees (northmost point move 2 degrees further north, eastmost moved 2 degrees further east, etc.) to ensure that map projection issues do not cause a point which, in (i,j) space, is inside the bounding curve to be skipped. Figures 1 and 2 illustrate the process graphically and in flow chart, respectively.

Cells that the bounding curve itself passes through are flagged as 'coastal' (equivalently, 'boundary') on the grid. For satellite products, such points likely should be treated as water. For open ocean models, they likely should be treated as land – the specific decision will be made by reference to the bathymetry. When a point strictly inside the curve (bounding curve does not pass through the cell, and the point is inside the bounding curve) is found, a flood algorithm is used to insert the appropriate flag type for all grid points which are connected (by motion either in i, or in j) to the current point – out to the limit when the boundary curve is reached. The flood fill algorithm used is from MIT [2001]. This works, but encounters stack overflow problems when the area to fill has too many points. The practical limit seems to be a grid with about 10 million points, making a global grid at 5' resolution (9.3 million points) about the limit.

This hierarchical data set ensures that no type 2 curves (inland lakes and seas) are encountered before the type 1 curve (continents and oceanic islands) which encloses it. This gives us a 'painter's algorithm' process of successively updating the overall flag grid.

Program *refill* produces a flag file named fout (if you use the standard script new_build.sh) and an .xpm graphic for a quick look at the results of the land mask generation. The format of fout is unsigned characters varying first in longitude and then latitude (as specified by the increments for latitude and longitude in the grid's class).

Flag values are:
1 Boundary
3 Undefined
5 Land
15 Ocean

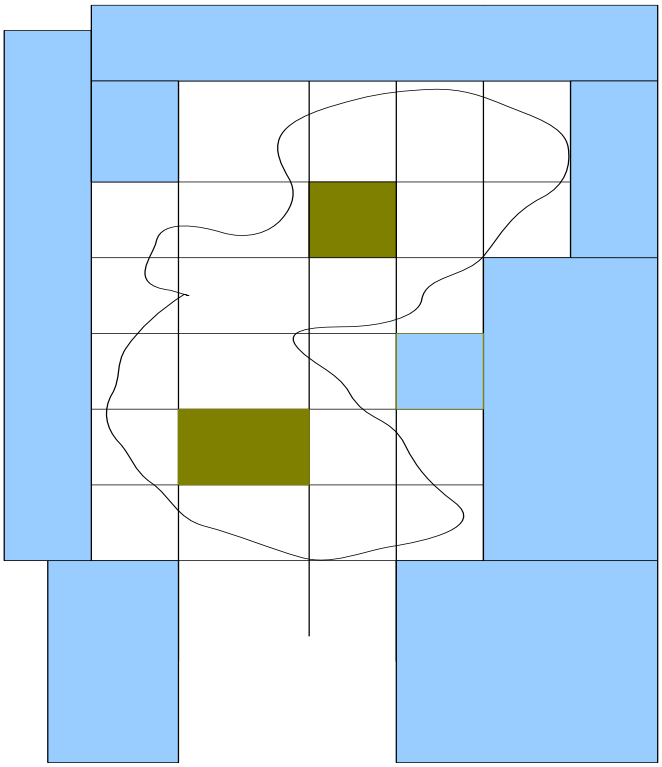

Fig. 3.1 : Boundary curve passing through grid space. Points strictly inside curve or strictly outside curve are flagged appropriately to the curve type, and those the curve passed through are flagged 'coastal'.

Fig. 3.2 : Flow chart of first guess land mask

17 Water (outside of ocean)

4 Reconciler for RTOFS-Atlantic Use

The satellite use is content with having points flagged 'boundary', as this can be interpreted as being points which required additional or different processing. Or simplifying assumptions may be made. For ocean modeling purposes, it is necessary to resolve whether the point should, or should not, be included inside the model domain. To decide this, I use the bathymetry. This is program *bathy*.

All high resolution bathymetry points which fall inside a given target model grid cell are used. Points which are above sea level are summed and counted separately from those which are sea level and below. Whichever has the higher count is used to provide the bathymetry. If more points are above sea level than at or below, then the point is considered land and an above sea level elevation is returned. If there is a tie on counts, then whichever is farther from sea level is used. That is, if there are equal grid points above and below sea level, but the average depth of the water is 200 meters, while the average elevation of land is 20 meters, the point will be considered water with an average depth of 200 meters.

Then a check is made for points which are flagged as ocean from the coastline file, but which the bathymetry has said are above sea level. If the number of bathymetry points was equal between above and below sea level, the bathymetry is reset to be below sea level, reversing the prior tiebreaker. This mostly happens near Greenland.

The final step in program *bathy* is a scan applying minimum depth limits. Any point shallower than a specified minimum depth and deeper than a specified maximum – used to ensure that all land points aren't reflagged – are flagged as coast. The algorithm flowchart is given in figure 3.

The final mask results from program *paving*. This first flags as 'boundary' points lying along straits that the user wishes to close for modeling purposes. It then takes a location specified by the user as being a point of interest to the model and performs a flood fill, flagging as final_ocean all points which are reachable from that initial point. In this way, the Great Lakes and Pacific Ocean are excluded from the mask for RTOFS-Atlantic. Straits specified for Gibraltar and Odense close off the Mediterranean Sea and Baltic Sea, respectively. Figure 4 provides the flow chart for this procedure.

RTOFS-Atlantic demonstration:
∘ read in bathy and drop in bucket average values, separately for positive and negative
∘ count depth = 0 as ocean (negative)
∘ in final pass, apply a critical depth for 'ocean' status – avoid going too shallow for intended usage
∘ Remove lakes and ocean points that are disconnected from a 'point of interest'

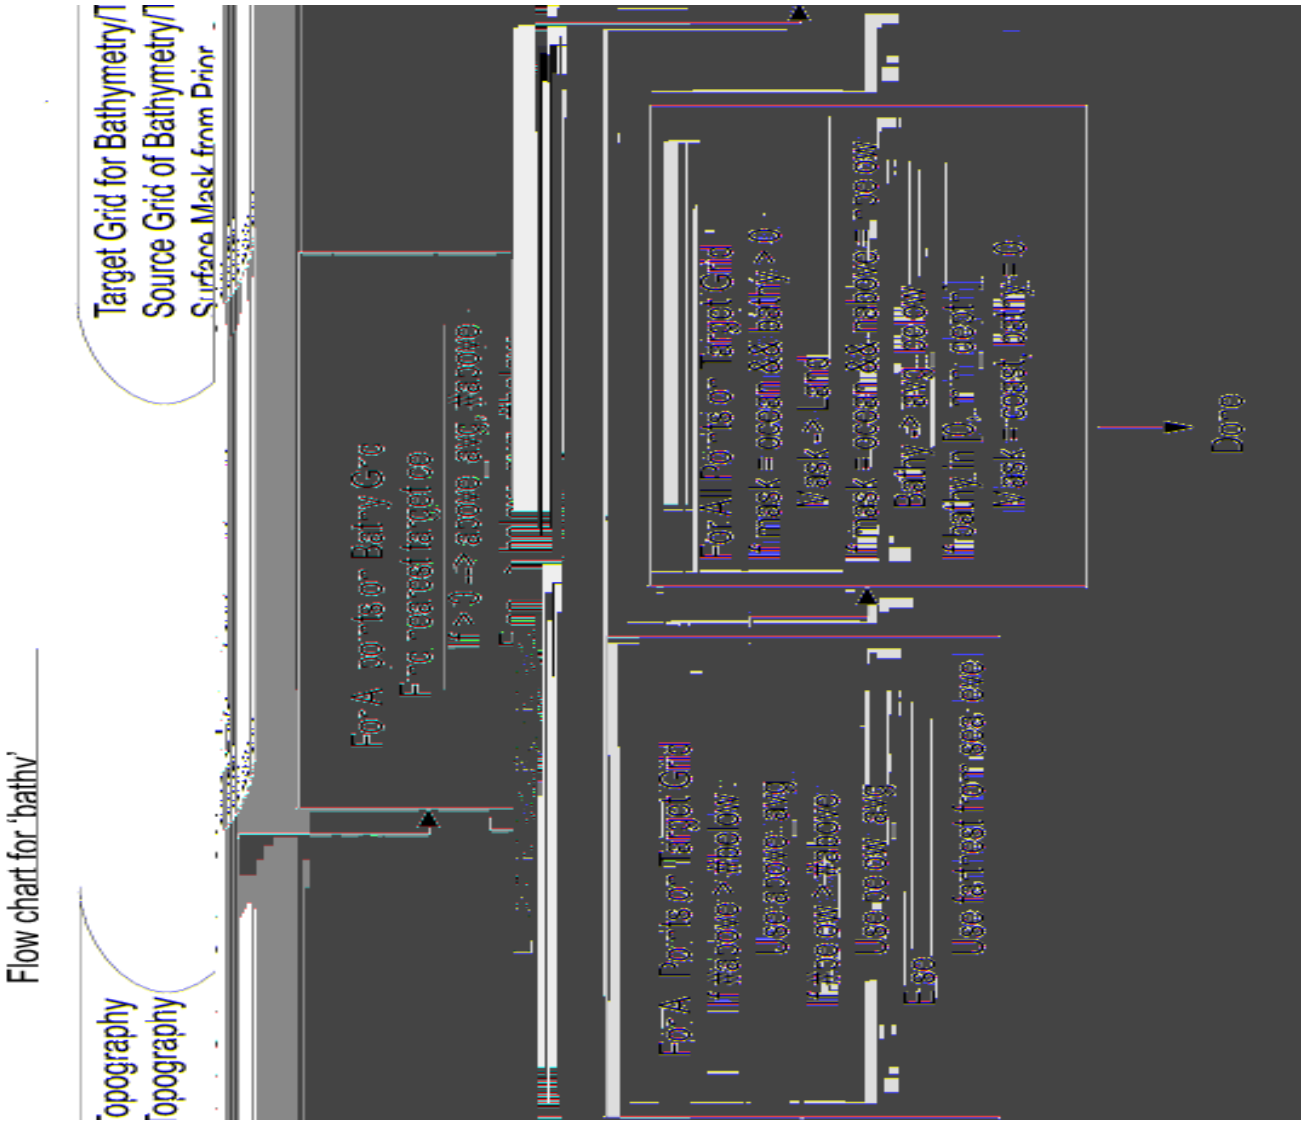

Fig. 4.1 : Flow chart of bathymetry analysis

Fig. 4.2 : Flow chart of paving in undesired points

– i.e., in doing regional model, even one which crosses into another ocean, skip those points that are not in the ocean of interest.
◦ Add in straits manually as desired (Gibralter, ...)

The bathymetry used is the ETOPO2 [NGDC, 2006]. DBDB2 [NRL, 2006] was also examined and found to differ substantially from ETOPO2. As collaborators were using ETOPO2, we continued with that data set.

5 Inland bodies observation processing

For satellite processing, it becomes important to recognize that lakes and inland seas (and embayments, wide parts of rivers) do not orient themselves on grids. When an observational footprint includes both the water surface of interest, and the land around it, the data are, from a water standpoint, corrupted. It would be ideal to work with full knowledge of just how much of the footprint was how heavily influenced by each surface type. Until that time, however, it is useful to consider the inland bodies separately from the open ocean problem. Then, to check how far from the bounding curve the center of the footprint is. If the observation is too near the boundary curve, given the observation system, then such observations should be ignored. Or, as this becomes possible, handed to a processing system that can manage mixed types.

In order to carry out such analyses, a reduced data set is extracted from the GSHHS [Wessel and Smith, 1996]. It contains only the bounding curves for types 2, 3, 4 – lakes, islands in lakes, and lakes on islands in lakes. This data set, at 'fine', resolution, is only 6 Mb, to the 87 of the full data set.

It is further useful to consider *a priori* how far from the bounding curve it is possible to get. A nearly circular lake, like Lake Victoria, makes it possible to get quite far from the edge (105 km). Lakes with convoluted boundaries, like Great Slave Lake (half the area of Lake Victoria, so 70% the radius if equally circular), are more difficult to get away from the coast (45 km maximum, vs. the 75). Far fewer observations in this case than suggested by the areas. Great Slave will have only $(45/105)^2$, 18% as many observations even though its area suggests it should have 42.4%.

Program *polygon* runs through a file of gshhs polygons looking for the point on the given grid which is a) inside the polygon, and b) farthest from the polygon. It then lists out this maximal distance, its location in ij space and latitude, longitude, the gshhs bounds for the polygon, its area in km^2 and the filling fraction (comparing the given area to the area of a circle with radius equal to the maximal distance). Figures 5 and 6 provide text and graphical illustration of the process.

Notes: A final pass to look through points within 0.5 i,j of the center point found by the above method may make significant difference for the results of smaller curves. Many of the polygons (from the 1/16th km file) wind up with

Fig. 5.1 : Flow chart for analyzing distance to land (edge of polygon)

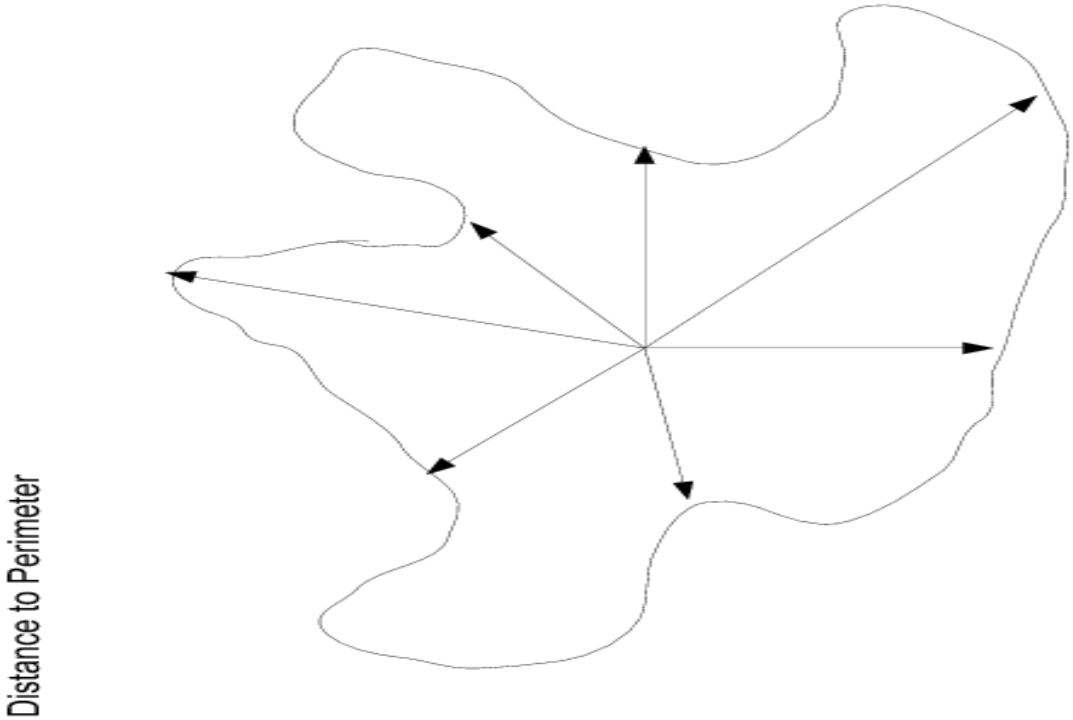

Fig. 5.2 : Graphic of distance to edge of polygon algorithm

maximal distances less than 1/2 of a grid point (tests were done with 5 minute global grid as the target).

Maximal distance depends on resolution of gshhs files – coarser grids also include greater spans between points. Fine resolution was used in these.

Maximum inland water body distance to coast is 170 km for Caspian Sea, giving it a 23% fill. Lakes Superior-Huron-Michigan combine for 104 km maximum, and fill of 16%. Aral Sea gives the highest fill fraction, but this is for the erroneous old coastline, 111 km to shore and 58% fill.

The fill fraction gives a sense of the probability of getting an observation which falls inside the lake but far enough away from the coast. This isn't rigorous, but it should be possible to make a statistic which would be. **

With 50 km as a cutoff (19 GHz SSMI plus sidelobe issues – note that the point is 50 km away from anything in all directions, so 100 km footprint would be safe), only 5 bodies are nominally resolvable:

Name	Distance	Latitude	Longitude	Area (km^2)	Fill fraction
Caspian Sea	169.9	38.708	51.292	396480.6	0.229
Aral Sea	111.1	44.792	59.708	67392.7	0.575
Lake Victoria	105.1	-0.875	32.875	69367.1	0.500
Lake Superior (+Huron/Michigan)	104.1	47.792	272.625	207700.4	0.164
Lake Ladoga	55.5	60.792	31.625	17765.3	0.544
Great Slave Lake	44.5	61.375	245.375	29411.9	0.212
Lake Erie	43.5	42.208	278.792	26997.2	0.220
Lake Winnipeg	41.3	52.708	262.042	25268.1	0.212
Great Bear Lake	37.0	66.042	239.625	30826.3	0.140
Lake Baikal	36.0	53.042	107.542	32131.8	0.127
Lake Ontario	35.5	43.625	282.458	19678.4	0.201
Lake Nicaragua	34.8	11.542	274.542	8185.1	0.465
Lake Nyasa	33.3	-11.792	34.625	29056.3	0.120
Lake Tanganyika	32.6	-6.792	30.042	32874.6	0.102
Manicouagan Lake*	29.4	51.375	291.292	4176.9	0.649
Lake Onega	27.2	61.792	35.292	10113.9	0.230
Koko Nor	26.7	36.875	100.125	4448.8	0.503
Lake Chad	25.7	12.958	14.125	12022.3	0.172
Lake Khanka	25.5	45.042	132.458	4060.1	0.502

Down to 25.4 km (SSMI 19 GHz without sidelobes, 37 GHz with them, and approximate reporting spacing) adds 14, giving a total of 19 bodies:
12.7 km (SSMI 37 GHz clean ignoring sidelobes) takes us to 60 bodies
6.3 km (SSMI 85 GHz, AMSR-E 37 GHz) 192 bodies
3.1 km (AMSR-E 91 GHz) 629 bodies

2.2 km (AVHRR GAC half footprint) 1050 bodies
0.55 km (AVHRR LAC half footprint) 3267 bodies

Name	Distance	Latitude	Longitude	Area (km^2)	Fill fraction
Lake Uvs	24.8	50.292	92.708	3317.0	0.584
Lake Balkhash	24.8	45.708	73.875	17571.7	0.110
Issyk-Kul	24.1	42.375	77.125	6272.3	0.291
Lake Titicaca	24.1	-15.792	290.542	8145.2	0.224
Lake Nipigon	22.9	49.875	271.458	4713.9	0.349
Lake Athabasca	21.3	59.375	250.625	8100.4	0.176
Lake Alakol	20.9	46.208	81.708	2843.4	0.482
Lake Vanern	20.5	59.125	13.542	5856.3	0.226
Lake Mweru	20.5	-9.125	28.625	5059.3	0.261
Lake Rudolph	20.4	3.792	36.042	7766.9	0.168
Lake Urmia	19.8	37.458	45.542	4407.3	0.280
Lake Tana	19.8	12.042	37.375	3165.0	0.387
Lake Dubawnt	19.4	63.125	258.458	3823.3	0.309
Lake Albert	19.4	1.792	31.125	5442.4	0.216
Lake of the Woods	19.0	49.125	265.292	4632.7	0.245
	18.6	31.208	120.125	2521.5	0.433
Lake Peipus	18.6	58.792	27.458	3578.5	0.304
Lake Kivu	18.2	-1.958	29.125	2674.8	0.387
Rybinsk Reservoir	18.1	58.458	38.375	4699.6	0.218
	17.9	61.875	29.208	11121.9	0.091
Reindeer Lake	17.8	57.208	257.625	6885.8	0.145
Lake Edward	17.4	-0.292	29.625	2219.6	0.429
Lake Van	17.2	38.625	42.708	3511.5	0.265
Lake Okeechobee	16.9	26.958	279.208	1604.6	0.556
Lake Manitoba	16.6	50.458	261.708	4833.0	0.179
Amadjuak Lake	16.4	64.958	288.958	3120.2	0.270
Mar Chiquita	16.2	-30.708	297.458	2097.9	0.392
Lake Manicuoagan*	15.7	51.375	291.208	2004.7	0.388
Lago Mirim	15.6	-32.625	307.208	3859.2	0.199
Lake Winnipegosis	15.3	53.375	259.792	4522.3	0.162
Iliamna Lake	15.0	59.458	204.375	2696.1	0.262
Selwyn Lake	14.9	58.292	256.708	2471.7	0.284
Lake Beloye	14.6	60.208	37.708	1139.6	0.585
	13.8	56.208	285.458	1339.5	0.446
	13.7	13.042	103.958	2565.7	0.228
	13.6	51.208	100.542	2758.5	0.211
	13.3	48.875	117.292	2349.5	0.235
	13.2	-11.208	29.792	2132.6	0.258
	13.0	55.125	255.042	1445.3	0.369
	12.9	30.708	90.458	1919.9	0.270
	12.8	-7.792	31.958	1955.2	0.263

Some filling fractions (not the above) can be below 1%. These are typically rivers with wide areas which get included in the GSHHS data set. Particularly large area inland water bodies which are nevertheless always within 12.7 km of a shore are:

Distance	Latitude	Longitude	Area (km^2)	Fill fraction
10.4	-2.542	294.542	28347.0	0.012
5.8	65.125	124.542	12659.0	0.008
8.1	1.708	19.458	11925.2	0.017
12.5	7.625	0.125	8281.1	0.060
2.5	66.708	68.542	7850.3	0.002
9.5	-27.458	303.625	7479.2	0.038
10.6	11.375	105.125	6130.7	0.058
4.4	49.875	136.458	5735.2	0.010
10.7	54.542	48.625	5721.4	0.062
11.7	54.042	296.042	5644.8	0.076
5.8	69.708	85.542	5576.1	0.019
11.3	52.708	260.125	5487.6	0.073
12.2	49.208	32.875	5309.2	0.088
5.4	9.125	297.792	4928.7	0.019
8.1	29.125	116.208	4921.3	0.042
9.1	74.542	101.625	4644.2	0.056
8.3	-17.042	27.792	4622.7	0.047
11.4	48.125	83.625	4181.5	0.098
6.9	61.458	25.375	4153.0	0.036
6.6	55.958	101.875	4062.7	0.033
12.5	45.542	286.292	3982.1	0.124
5.6	23.208	32.792	3391.8	0.029
9.1	47.708	42.375	3057.8	0.085

We see that at over 3000 km^2 each, these 23 bodies contribute significant areas of water for numerical weather prediction models, but in such a way that they cannot be analyzed by coarse resolution satellite instruments.

6 Distance to Land

The RTGSST implementation of February 2010 [MMAB, 2010] also made use of a distance to land grid. This was generated from the high resolution land mask, rather than directly from the bounding curves. The algorithm was simply to start from each water point. Then examine the points within 1 point in i,j space to see if any were land. If any are, the distance from grid center to grid center was computed, and the minimum value retained. If no points were available within 1, then points 2 grid spaces away were checked, and so on, spiralling outward until

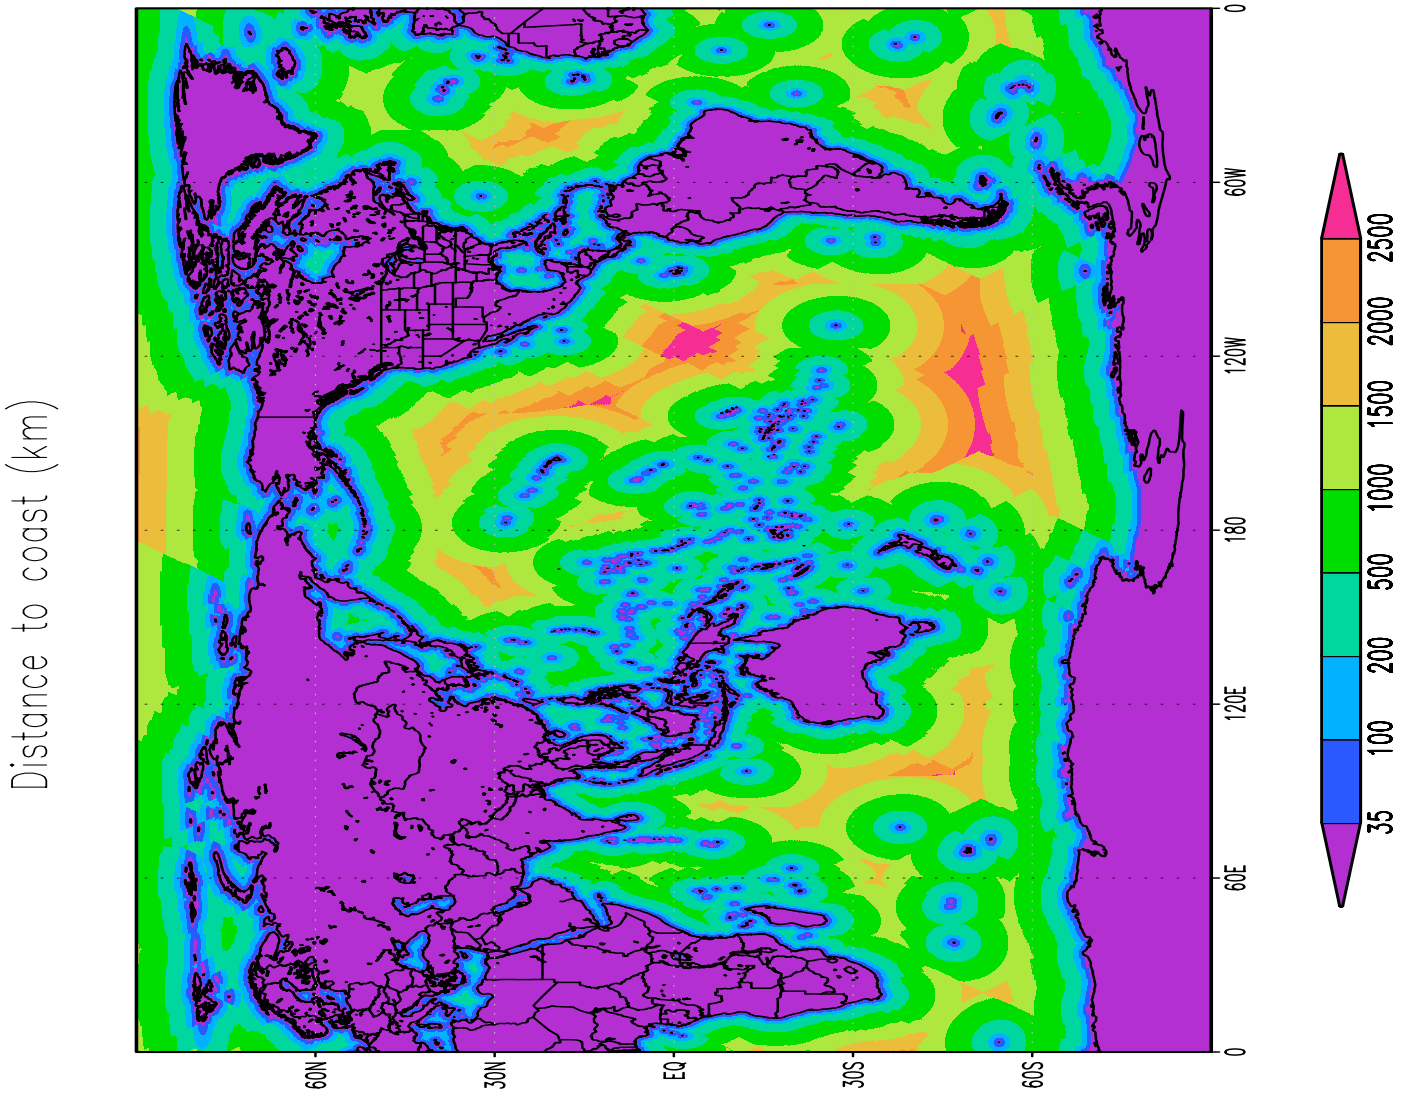

Fig. 6.1 : Distance to land (km)

either a limit range (which is an argument to the program) was reached or a land point was found. The point farthest from land was found to be 3077 km away. Current results are shown in figure 7.

This algorithm has the virtue of simplicity and speed, taking about 10 minutes on a 3.3 GHz pentium. It has the drawback of working in grid space, which ignores the convergence of meridians towards the poles – such that the nearest point geographically may be farther away in i,j space. This is more a problem for large distances from land. As the interest was for points within 200 km of land, this was deemed not a problem. It also has imprecision due to working with pre-gridded masks rather than directly with the bounding curves.

7 Conclusions

This suite of programs has evolved over a period of years, and will continue to do so. The main goal, which has been largely achieved at this point, was to construct a system which was capable of producing by consistent methods all geographically-related fixed files needed by satellite analysis systems and ocean models, while at the same time being indifferent to the nature of the target grid. Bathymetry, masks, distance to land grids have been generated on latitude-longitude grids, polar stereographic grids, and curvilinear orthogonal grids.

Improvements to come include: Moving to the most current version of GSHHS data, improving the distance to bounding curve algorithm, improve the distance to land grid's algorithm, use a flood fill algorithm that avoids the stack overflow issue of the current, and to develop a means of constructing 'coastal' polygons – such that polygons for coastal Florida, for instance, will be available to use analogously to the polygons for lakes.

Current versions of the software will be available in a tar file under
http://polar.ncep.noaa.gov/mmab/papers/tn2NN/

8 Bibliography

AMSR-E, AMSR-E Instrument Description,
http://www.ghcc.msfc.nasa.gov/AMSR/instrument_descrip.html, Last Accessed 1 February 2011.

Chawla, Arun, Hendrik L. Tolman, Automated grid generation for WAVEWATCH III, MMAB Technical Note #254, 71 pp., 2007.

Gemmill, W., B. Katz, and X. Li, : 2007: Daily Real-Time, Global Sea Surface Temperature – High-Resolution Analysis: RTG_SST_HR,
NOAA/NWS/NCEP/EMC/MMAB, Science Application International Corporation, and Joint Center for Satellite Data Assimilation Technical Note Nr. 260, 22 pp., 2007.

Goodrum, Geoffrey, Katherine B. Kidwell, Wayne Winston, eds., NOAA KLM User's Guide, 2000, amended 2004. http://www2.ncdc.noaa.gov/docs/klm/cover.htm last accessed 1 February 2011.

Grumbine, Robert W. Automated Passive Microwave Sea Ice Concentration Analysis at NCEP OMB Tech. Note 120, 13 pp., 1996.

MMAB, http://polar.ncep.noaa.gov/sst/ophi/Changed_RTG_SST_HR.html, last accessed 1 February 2011.

Mehra, Avichal, Ilya Rivin: 2010: A Real Time Ocean Forecast System for the North Atlantic OceanSysTerr. Atmos. Ocean. Sci., Vol. 21, No. 1, 211-228, February 2010 doi: 10.3319/TAO.2009.04.16.01(IWNOP)

MIT 2001
http://groups.csail.mit.edu/graphics/classes/6.837/F01/Lecture03/lecture03.pdf last access 12/18/2009.

NGDC, 2006: 2-minute gridded global relief data (ETOPO2).
http://www.ngdc.noaa.gov/mgg/fliers/01mgg04.html.

NRL, 2006: Digital bathymetry data base 2-minute resolution v. 3 (DBDB2).
http://www7320.nrlssc.navy.mil/DBDB2_WWW/NRLCOM_dbdb2.html.

Raytheon, SSM/I User's Interpretation Guide, 104 pp, 2000.

Suranjana Saha, Shrinivas Moorthi, Hua-Lu Pan, Xingren Wu, Jiande Wang, Sudhir Nadiga, Patrick Tripp, Robert Kistler, John Woollen, David Behringer, Haixia Liu, Diane Stokes, Robert Grumbine, George Gayno, Jun Wang, Yu-Tai Hou, Hui-Ya Chuang, Hann-Ming H. Juang, Joe Sela, Mark Iredell, Russ Treadon, Daryl Kleist, Paul Van Delst, Dennis Keyser, John Derber, Michael Ek, Jesse Meng, Helin Wei, Rongqian Yang, Stephen Lord, Huug Van Den Dool, Arun Kumar, Wanqiu Wang, Craig Long, Muthuvel Chelliah, Yan Xue, Boyin Huang, Jae-Kyung Schemm, Wesley Ebisuzaki, Roger Lin, Pingping Xie, Mingyue Chen, Shuntai Zhou, Wayne Higgins, Cheng-Zhi Zou, Quanhua Liu, Yong Chen, Yong Han, Lidia Cucurull, Richard W. Reynolds, Glenn Rutledge, Mitch Goldberg
The NCEP Climate Forecast System Reanalysis, Bulletin of the American Meteorological Society, Volume 91, Issue 8 (August 2010) pp. 1015-1057, doi: 10.1175/2010BAMS3001.1

Sunday, Dan, 2001 Fast Winding Number Inclusion of a Point in a Polygon, http://softsurfer.com/Archive/algorithm/algorithm_0103.htm, last access 12/18/2009.

Wessel, P. and W. Smith, 1996: A global self-consistent hierarchical high-resolution shoreline database., *J. Geophys. Res.*, **101**(B4), 8741-8743.

Data Source: Global Land Cover Facility (GLCF) at www.landcover.org, last access