A Tile-Based System for the Rapid Display of Symbolized Digital Nautical Chart Data

Norman Schoenhardt, Dr. John T. Sample, Frank McCreedy
Naval Research Laboratory
Stennis Space Center, MS 39529

***Abstract*- For efficient retrieval and display of Digital Nautical Chart (DNC®) imagery, a tiled-based approach is presented. By preprocessing the data into persistent image tiles, the presented method can minimize or eliminate image processing steps during the retrieval process. Tile-based approaches define constant image sizes and scales for geographic regions, therefore tiles can be pre-generated and stored in a format needed for transmission to users. This greatly improves the performance of applications that visualize DNC® data and allows for advanced caching of DNC® imagery tiles. In addition, the evenly-spaced and commonly-sized images are stored in files that are easily indexed allowing for simple and fast lookup of DNC® imagery tiles for a given geographic area.**

I. Introduction

Digital Nautical Chart (DNC®) is a vector-based digital database containing maritime significant features that are essential for safe marine navigation. Examples of such features are buoy locations and descriptions, maritime shipping routes, and underwater hazards. It is produced by the National Geospatial-Intelligence Agency (NGA) and was designed to support applications such as navigation, mission planning, command and control, and situational awareness. In order to provide DNC® data as imagery to users in a rapid, yet easy to use way, a tile-based delivery system was developed. The tiling of DNC® data divides geographic regions into evenly-spaced and commonly-sized images containing DNC® symbolized vector data. Since the tile size is constant, DNC® imagery tiles can be pre-generated, thus preventing any time consuming image processing to take place upon request for the tile image. Also, the tiled-based delivery system is well suited for online distribution of imagery. Tiled images can be easily sent over the internet via http/https requests and responses allowing for broad access to DNC® imagery.

As previously stated the native format of DNC® data is vector-based, therefore to create tiled images of DNC® datasets a method of symbolization is required. The complexity of the DNC® database structure presents an additional challenge when displaying the symbolized vector data as imagery. The DNC® features are contained in overlapping layers and coverage areas that when displayed digitally can dramatically reduce the readability of DNC® imagery. A layering structure is needed to prevent the final imagery from containing duplicate data as well as appearing overly cluttered.

II. DNC® Symbolization

DNC® imagery is produced in the Vector Product Format (VPF). VPF is a standard format, structure, and organization for large geographic databases that are based on a georelational data model and are intended for direct use [1]. In order to display DNC® data as imagery, a method of symbolization is needed. The symbolization method used within the DNC® tiling effort is the Geospatial Symbols for Digital Displays (GeoSym). GeoSym is a comprehensive collection of symbology produced in the Computer Graphics Metadata (CGM) format developed specifically to support the display of VPF data; therefore it was a natural fit for the symbolizing of DNC® data[2].

In order to produce tiles using the DNC® data, a conversion would need to take place transforming the raw VPF data into imagery data using GeoSym symbology. A software module named the “GeoSym Visualization plugin”, or “GeoSym module” for short, was developed by NRL as part of the DNC® tiling system to do just this. The module ingests the raw DNC® VPF data and produces DNC® feature images, in the PNG image format. The GeoSym module extracts all the attributes found within the VPF data and processes those attributes through a series of assignment tables located within the GeoSym package. The GeoSym module uses the assignment tables to produce a series of equations which when evaluated link the attributes to one or more GeoSym symbols.

For some DNC® features up to six symbols may be assigned to represent various attributes of the feature [2]. Table 1 shows an example of a DNC® Feature and Attribute Coding Catalouge (FACC) Buoy feature (BC020) whose attributes match multiple GeoSym symbols.

0-933957-38-1 ©2009 MTS

Report Documentation Page

Form Approved
OMB No. 0704-0188

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED
JUN 2010	**N/A**	**-**

4. TITLE AND SUBTITLE
A Tile-Based System for the Rapid Display of Symbolized Digital Nautical Chart Data

5a. CONTRACT NUMBER

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Naval Research Laboratory Code 7180 Stennis Space Center, MS 39529 USA

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

10. SPONSOR/MONITOR'S ACRONYM(S)

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release, distribution unlimited

13. SUPPLEMENTARY NOTES
See also ADM202806. Proceedings of the Oceans 2009 MTS/IEEE Conference held in Biloxi, Mississippi on 26-29 October 2009. U.S. Government or Federal Purpose Rights License, The original document contains color images.

14. ABSTRACT
For efficient retrieval and display of Digital Nautical Chart (DNC®) imagery, a tiled-based approach is presented. By preprocessing the data into persistent image tiles, the presented method can minimize or eliminate image processing steps during the retrieval process. Tile-based approaches define constant image sizes and scales for geographic regions, therefore tiles can be pregenerated and stored in a format needed for transmission to users. This greatly improves the performance of applications that visualize DNC® data and allows for advanced caching of DNC® imagery tiles. In addition, the evenly-spaced and commonly-sized images are stored in files that are easily indexed allowing for simple and fast lookup of DNC® imagery tiles for a given geographic area.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			
unclassified	**unclassified**	**unclassified**	**SAR**	**6**	

Table 1 **Example DNC® FACC Buoy Feature Attributes and Matching GeoSym Symbols**		
Feature Attribute Type	**Example Attribute Value**	**Matching GeoSym Image**
BTC – Buoy Type Category	Unknown and Structured Shape Category equals Pillar/Spindle	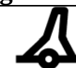
TMC – Top Mark Characteristic	Can(Open), Can(Filled) or Can Over Ball(Open)	
COL- Character of Light	Red	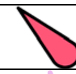
SST- Sound Signal Type	Horn, whistle, etc	
NST – Navigation System Type	Racon or Ramark	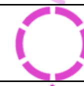
REF – Radar Reflector Attribute	Radar Reflector Present	

In this case the GeoSym module would take the matching GeoSym symbols and combine them to produce a completely symbolized portrayal of the DNC® feature.

Open Can Buoy with Red Light, Fog Signal and RADAR Reflector	

Figure 1 - Example FACC Buoy feature (BC020) final image

III. LAYERING SCHEME

The DNC® database consists of 29 geographic regions or databases which span an area between 84° North latitude and 81° South latitude. These regions are comprised of four overlapping library categories of data coverage based on scale including Harbor, Approach, Coastal and General, from largest to smallest scale respectively. Within each scale, the DNC® imagery is further divided into feature types that are thematically organized into coverages including: Cultural Landmarks, Earth Cover, Environment, Hydrography, Inland Waterways, Land Cover, Limits, Aids to Navigation, Obstructions, Port Facilities, Relief, and Data Quality. These features all cover the same geographic areas so when all feature types are displayed they can produce maps that are extremely cluttered and difficult to read. In order to display the DNC® data using tiles, without a tremendous amount of data overlap and clutter, a layering scheme needed to be developed.

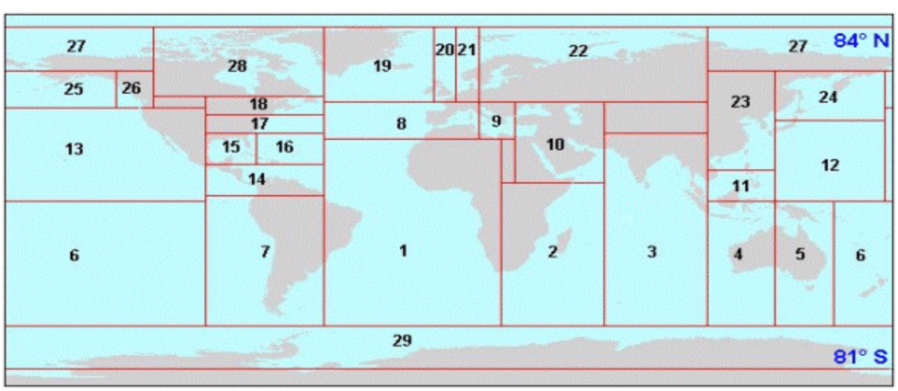

Figure 2 – DNC® Database Regions, Worldwide Coverage [3]

In a normal tiling scheme a single tile would represent image data of a single geographic region at a particular scale; however, due to the complexity of the DNC® data set, this scheme would have to be adapted. In order to display the DNC® imagery in a clear and concise way are layered approach was adopted. A single geographic region of a particular scale would now be comprised of multiple tiles, a background layer and a single layer for each individual feature type contained within the DNC® database. By separating the layers in this fashion the user can select the types of data to be displayed by simply turning on or selecting the desired layers. Once the layers are selected the tiles can be created accordingly and combined together to produce the final image. Unlike the old paper charts DNC® was intended to replace, this layering scheme allows the user to add and remove different types of DNC® imagery data that are viewable by simply adding or removing tile layers.

The background layer or "base" layer serves as the backdrop for all DNC® imagery and is comprised of three DNC® feature types: FORESHORE_AREA, DEPTH_CURVE_AREA, and ISLAND_AREA. As mentioned earlier the DNC® coverage area only spans from 84° North latitude and 81° South latitude so there would be missing data in a worldwide view. To help fill in these missing areas the World Vector Shoreline (WVS) dataset was used as a foundation layer onto which the three above mentioned feature types are drawn. WVS is an imagery data set provided in Vector Product Format which is the highest resolution demarcation of coastline globally available [4]. With the addition of WVS data, a complete worldwide background layer for DNC® was created.

To visualize the rest of the DNC® features, individual layers would be created for each DNC® feature type for each scale. These layers would be comprised of tiles that contain symbolized DNC® features for a particular type and scale and each tile would have a transparent background so that the tiles could be layered upon one another.

IV. Tiling of DNC® Data

The tiling system works by dividing the world into fixed equal sized geographic regions. For this application, each tile image was created of size 512x512 and each tile represents one particular geographic region. An image tile may be any size, however, a 512x512 image tile provides a compromise between limiting image size and limiting the number of images required to fill a user's map. By limiting the number of images, the number of requests that have to be made for image tiles is reduced, decreasing the time it would take to completely fill a map with imagery data. In addition, limiting the size of the images limits the amount of data to be sent over a network connection, therefore improving the speed of tile delivery.

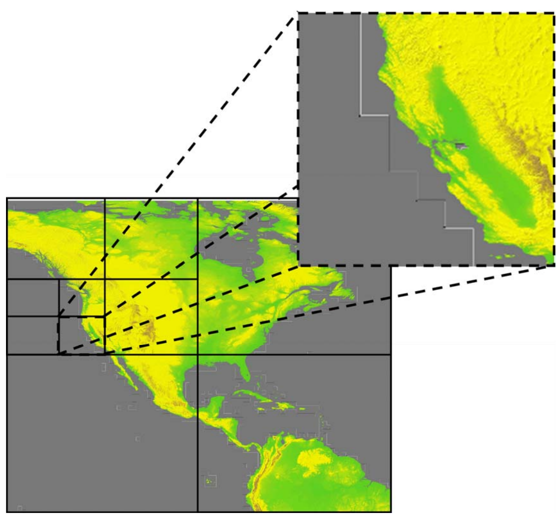

Figure 3 – An Example World Image Broken Into Successively Higher Resolution Tiles.

This system necessitates predefined map scales. The lowest map scale, having a value of "1", partitions the world into two tiles; a tile matrix of one row and two columns. All whole number (i.e. 1, 2, 3, … n, n+1) increasing consecutive scales have double the number of rows and columns as the previous scale[5]. The following formulas can be used to calculate the number of rows and columns for an arbitrary scale. Let *NR* be the number of rows for a given scale *s*, then the formula for computing the number of rows for an arbitrary scale is as follows:

$$NR = (2)^{(s-1)}$$

Now, let *NC* represent the number of columns for a given scale *s*, then the formula for computing the number of columns for an arbitrary scale is

$$NC = (2)^{s}$$

The tiling system allows for arbitrarily high map scale, though it is rarely necessary to have more than twenty scales.

An advantage of using a tiling scheme for the distribution of DNC® imagery is that tiling is well suited for online distribution. Tiled images can be easily sent over the internet via http/https requests and responses in their native PNG image format. A tile server was developed to provide direct access to tiled DNC® imagery over the internet. Each tile image is accessible from the tile server via a unique Uniform Resource Locator (URL) containing the scale, row, column, and DNC® layer name of the desired tile. This URL allows both web browser clients as well as stand alone clients to retrieve tiles in an effortless and timely manner. When making tile requests, the client is to determine which tiles are necessary for a particular geographic region and will manage the map view themselves. Clients will usually display the tiles at their native resolution, therefore, no costly image processing needs to be performed by the client [5]. Tile based clients such as Google Earth and OpenLayers would be an example of clients that are compatible with this methodology. Also, by using the tile server as a means to access tiles, the client can eliminate the need for cumbersome and large imagery data storage on their local machines.

The tile server was developed for the dissemination of DNC® imagery over the internet. The tile server receives a request from a client and then extracts the raw vector feature data from the DNC® database using the requested row, column, scale and layer name. The latitude and longitude bounds for the requested tile are then computed based on the row, column, and scale. The tile bounds data along with the vector feature attributes are passed to the GeoSym module which is located internally to the tile server. The server then returns the tiled image produced by the GeoSym module via an http response.

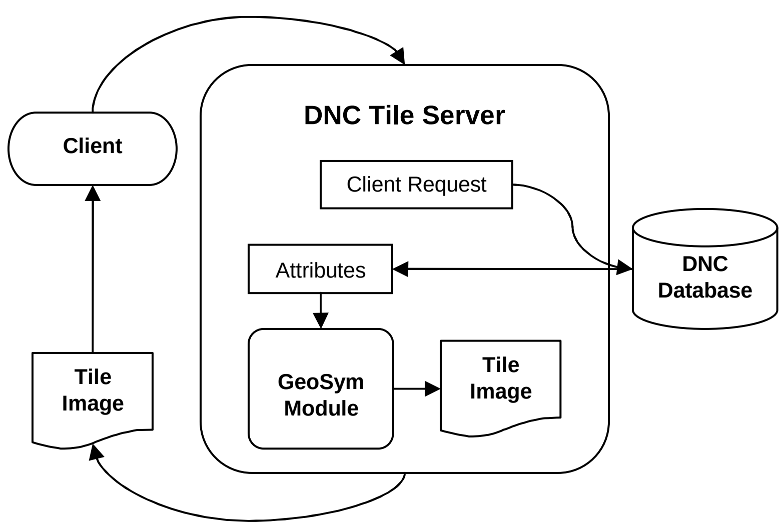

Figure 4 – DNC® Tile Server Model.

V. TILE CACHING

When a client makes a request for a DNC® tile image using the DNC® tile server model, time-consuming data extraction and image processing must take place in order to generate that tile image. As the number of tile requests increases the amount of image processing and data extraction can start to affect performance of the tile server. In order to increase the speed and efficiency of tile delivery to clients, a two-tiered caching system was developed. The idea behind the caching system is to pre-generate tiles ahead of time and store them in a way in which they can be retrieved rapidly. The tiling scheme is well suited for this application in that the tile format provides for evenly-spaced and commonly-sized images which can be stored in files that are easily indexed allowing for simple and fast lookup. By using a caching system, when a client makes a request the tile can be retrieved from the cache and no image processing is required.

The first tier of the caching system is called the “packed cache” system. In the packed cache system, tiles are stored in files in a binary format. Each file is designed to store all the tiles that make up one particular DNC® layer and scale. Each binary file can contain full or partial tile sets and all tiles that comprise the binary files are created and stored before the caching system is in use. Accompanying each binary file is an index file that contains the tile offsets of the tiles in the binary file. Each tile is indexed based on its row and column, representing the tile’s location in the particular scale’s tile set. A tile can be retrieved from the packed cache by retrieving its file position from the index file using the row and column number of the tile. Using random access, the tile is retrieved from its location in the binary file and returned.

The second tier of the caching system is referred to as the “image cache” system. Due to the potential size of the files required to store complete larger scale DNC® imagery layers, it may not be practical to generate all DNC® tiled images within the packed cache system. In the event that a tile request is received and the tile is not located within the packed cache, that tile must be generated by the DNC® tile server. To prevent that tile from being recreated each time it is requested, the tile must be added to the tile cache. In order to prevent collision with other tile requests accessing the packed cache system, it is not ideal to write the tile to the binary file in which it belongs while the caching system is in use. Therefore the tile is stored in the image cache for later retrieval.

The image cache is composed of a series of directories, located on the file system, based on DNC® layer names, scale and row. Tiles are stored in the PNG image format in a hieratical directory structure starting with the layer name directory, followed by the scale directory, followed by the row directory in which the tile belongs to. The tile is saved using the column number it represents as a filename. By using this directory scheme and naming convention, tiles can be quickly located and retrieved by simply parsing through the directory structure using the tile’s layer name, scale, row and column.

The image cache is designed so that at any point the packed cache system is not in use, the tile image files can be extracted from the image cache, combined and stored into the packed cache system. Ideally, this would take place during planned server outages.

A tile cache server was created to accept tile requests and act as a buffer between clients and the tile server. Upon receiving a tile request the tile cache server will check the tile cache for the existence of the tile. If present, the tile is extracted from the cache and returned to the requesting client. If the tile does not exist in the tile cache, the tile cache server will make a request to the tile server for the particular tile. Once the tile is received from the tile server, the tile cache server will store the tile in the cache system before the tile is returned to the requesting client.

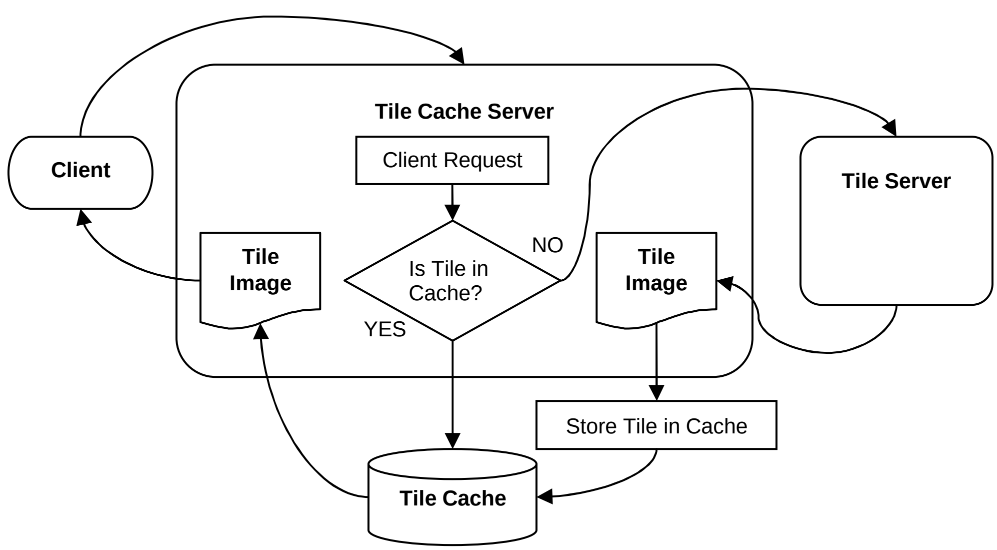

Figure 5 – DNC® Tile Cache Server Model.

VI. SUMMARY

The tile based system to display DNC® data was designed to provide rapid dissemination of symbolized DNC® imagery. Through the use of tiles, DNC® imagery can be produced ahead of time and stored in way that allows for little to no image processing. This allows for DNC® imagery to be retrieved and visualized in a quick and concise way. The imagery can be served out via the internet to provide broad access to clients and users of the DNC® product. The DNC® tiling system provides a sophisticated, yet easy to use imagery management system which removes the need for clients to store large amounts of imagery on their local machines as well as removes the need for clients to perform complex and time consuming image processing.

ACKNOWLEDGMENT

This work was sponsored by the U.S. Naval Research Laboratory's Base Program, Program Element No. 0602435N.

REFERENCES

[1] "Department of Defense Interface Standard for the Vector Product Format," MIL-STD-2407, 28 June 1996, p.1

[2] "Detailed Specification Geospatial Symbols for Digital Displays (Geosym ®)," MIL-DTL-8904A, 15 February 2007, pp. 1-7

[3] National Geospatial-Intelligence Agency. "Digital Nautical Chart (DNC®)." http://www.nga.mil/portal/site/dnc/. (July 23, 2009).

[4] National Oceanic and Atmospheric Administration. "World Vector Shorelines." http://shoreline.noaa.gov/data/datasheets/wvs.html. (July 25, 2009).

[5] Ioup, Elias, Sample, John and McCreedy, Frank, "Tiled Image Archival and Distribution for Seafloor and Terrestrial Imagery," NRL Review, 2009.