Building Intermap's Data Management System Interface Using ArcPublisher and ArcReader

Richard Asselin

Intermap Technologies Corp.
2 Gurdwara Rd
Suite 200
Ottawa, ON K2E 1A2
(613) 226-5442

www.intermaptechnologies.com

Abstract

Intermap is in the business of creating digital descriptions of the Earth's surface. Our unique data acquisition technology enables Intermap to quickly and affordably acquire large and continuous datasets. Through our NEXTMap program, we're acquiring complete and accurate map databases of entire countries - something never before undertaken on a commercial basis.

Managing such a large undertaking demands a thorough understanding of what's happening in production, and where the work is heading. A tool called the "Data Management System – Explorer" (DMSe) was developed to aid production staff and management to tackle this daunting task.

This paper focuses on the design, development, and implementation of the DMSe. Details are given regarding the programming of the ArcReader / Microsoft SQL Server production database interface, as well as the various DMSe database views and stored procedures. Timesaving tips are also discussed, allowing for speedier changes to the DMSe with minimal or no downtime.

Introduction

Intermap and the NEXTMap Program

Less than a decade later after its establishment in 1996, Intermap Technologies has employees and operations worldwide that serve a growing digital elevation model (DEM) and geographic information system (GIS) marketplace.

Intermap is in the business of creating and licensing accurate digital descriptions of the surface of the earth through the use of their STAR mapping technology. They have acquired, and continue to build, a library of elevation data and imagery products that provide a precise and accurate digital model of the shape of the earth's surface at an unprecedented accuracy. These products serve a wide range of commercial, government, military and consumer applications.

The demand for detailed, accurate and low-cost digital topographic maps of the world has never been greater. Intermap answered that demand with NEXTMap.

Their unique data acquisition technology enables them to quickly and affordably acquire large and continuous datasets, at night or in cloudy conditions. Through the NEXTMap program, they're acquiring complete and accurate map databases of entire countries - something never before undertaken on a commercial basis. Intermap's data library currently features the national coverage of Great Britain, Jamaica, Puerto Rico, Indonesia, Solomon Islands and the island nation of Vanuatu. The data collection of the United States is currently underway, with initial collections in Mississippi, West Virginia, California and Florida.

Data Management System

In order to process such a large dataset in a timely and efficient manner, great efforts have gone towards automating parts of production and simplifying the management of the production flow. One such effort was the development of the Data Management System, or DMS. The DMS, as its main role, aids managers and employees control the flow of data through the production process. The DMS consists of a set of tools accessible through one location. Some of these tools interact with the production database and the editing software. The other part of the DMS, the "DMS – explorer", serves as the eyes into the production process. The DMSe, created using ArcPublisher and ArcReader, allows users to, at a glance, see the exact location of a tile in production process. The user can then drill in for more details on that specific tile. The DMSe will be the focus of this paper.

Outline

This paper will begin by giving a general description of the production flow from the initial acquisition planning right through to the final delivery of data to Intermap's online store. The second section will delve further into the actual data editing, focusing on the software used to create Intermap's products, and specifically on the Data Management System. The final section will focus entirely on the DMSe and how it was developed to interface with the production database. This paper will conclude with lessons learned, as well as what the future holds for the DMSe.

IFSAR Data Processing

Overview

Intermap is in the business of collecting Interferometric Synthetic Aperture Radar (IFSAR) data for creating highly accurate yet affordable Digital Elevation Models (DEMs) and Orthorectified Radar Images (ORIs) of the Earth. Intermap delivers three core products to its clients: Digital Surface Models (DSMs), Digital Terrain Models (DTMs), and the aforementioned ORIs. ORIs look somewhat like monochromatic aerial photos and are always generated in pairs for any particular scene. The pairing of the ORIs allows Intermap to create DSMs. ORIs and DSMs display the first surface on the ground that the radar strikes, including terrain features, buildings, power lines, and vegetation such as large trees or forests. Through a combination of automated tools and manual editing, a DTM is created by digitally removing cultural features, as well as treed areas, from the DSM.

Intermap's data acquisition / editing processes follow an ISO framework, meaning that the activities have been structured in terms of well-defined and documented processes. As such, core products are produced to tightly controlled specifications. The general flow of data from the initial contract award right to the final delivery occurs as follows:

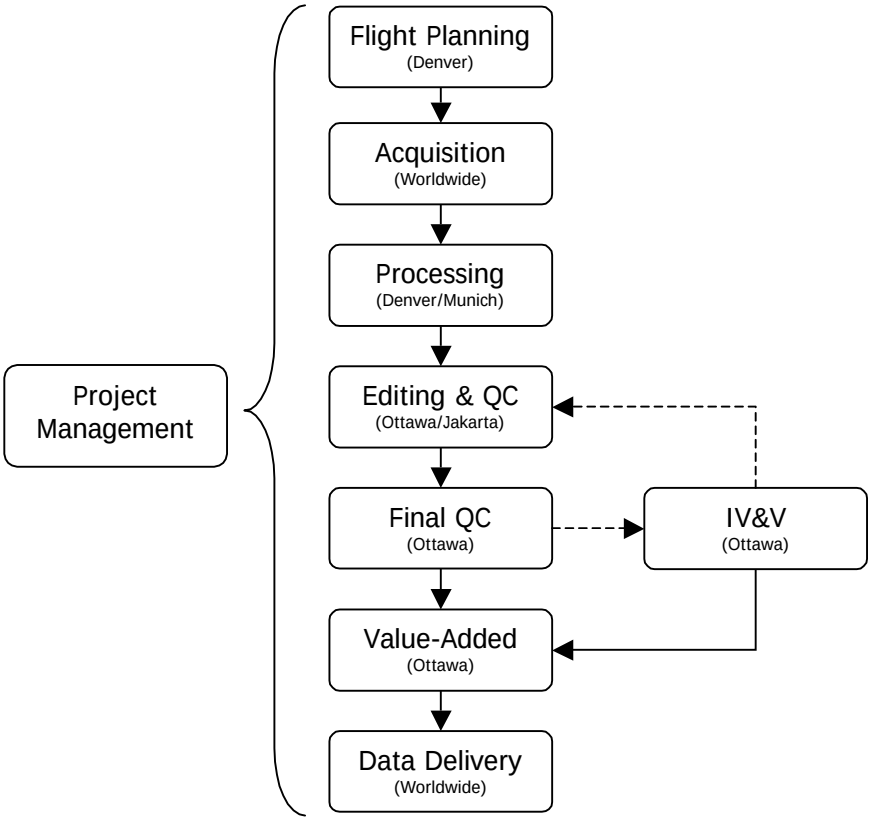

Figure 1: Production Flowchart.

Primarily each of the departments responsible for steps in Intermap's data acquisition & processing system have their own methods for controlling and tracking their production. At the same time, it is of vital importance that all departments work synchronously in order to get data out in a timely and organised manner, all the while avoiding bottlenecks. This requires teamwork on the part of many individuals from different departments, and clear and concise methods for communication.

Data Processing Tools

Interferometric Editing Software

The greatly increased production flow of the NEXTMap programs demanded that better methodologies were embraced for editing. As such, in 2002 efforts began towards streamlining the editing process used by the production department. The Boeing Company (and their affiliate Autometric) had already developed software that the Boeing – Intermap team used for editing data from the Shuttle Radar Topography Mission in February 2000. After modifying the software for

use on Intermap's Star data, it was renamed IES ("Interferometric Editing Software"). The SRTM & IES software included one very important advantage over most commercial off-the-shelf software packages: it offered the ability to control production flow using a relational database management system. Tiles passing through the editing process were controlled by the database, and the history of the tile was tracked in the database. Additional functionality such as an operator's access level for the software, or a tile's difficulty rating were also included in the database design. The software also had the ability to automatically launch certain processes based on the tile's state. For example, once interactive editing was completed, the software would automatically check the edges for any errors.

Data Management System

The software – database interface, although a great improvement over the old manual editing process, introduced a new challenge to the production environment. Management had to manually edit fields in the database in order to get data through the production process. This manual editing of tables was, for obvious reasons, a big concern for the database administration team. The probability of introducing erroneous data into the database was rather high. Accordingly, a new toolset was introduced into the production environment: the Data Management System (DMS). The major goal of the DMS was to provide to the end user a simplified interface into the production database. Tools to change the editing state of tiles, as well as perform and track Quality Control of tiles were built into the DMS. The other major goal was to provide a one-stop shopping location for all DMS-related tasks. Users logged in to a secure web-based environment and were, based on their access level, presented with a list of available tools. The DMS in a sense became the control interface for the IES software and the other production editing processes.

Data Management System "Explorer"

Overview

While the web-based DMS solved the issue of getting information *into* the production database, there was still the challenge of getting information *out* to the user. With all the information stored in the database, there was still no way to view what was going on in production other than viewing a tabular report of all tiles and their associated editing information. While useful, tabular reports have one major shortcoming, in that they don't allow easy "at a glance" viewing of adjacent tiles or editing areas. They also make it more difficult to interpret information from a multitude of tables. Intermap needed a way to get that information out to the user without having to train all people on the intricacies of

the database tables and structure. From this need came the birth of the Data Management System "Explorer" (DMSe). The DMSe provided a graphical interface, through ESRI's ArcReader, to the production database allowing users the ability to quickly see a tile's editing status, while still giving the option of drilling down further into the details of the tile.

DMSe Design Goals

One of the major goals of the latest version of the DMSe was to facilitate the "DMSe experience." Since the DMSe was intended to be the gateway to the production environment, ease of use was of the utmost importance. In order to meet this goal, focus was placed on four main areas: a live join to the database, better integration with the DMS tools, an easier at-a-glance understanding of where tiles were in the editing process, and more information basemap layers.

Basemap Layers

Only basemap layers that would be of use in the editing process were included, such as the various hydrology, transportation, and urban area layers. While not directly used in the production process, these layers do provide a general idea of the degree of difficulty that an operator will encounter during editing. This in turn allows the production coordinator to more easily assign work to operators based on their experience levels or editing specialties. Additional layers such as political boundaries provide information more generally used by project managers during the initial stages of the production process.

Another basemap layer in the DMSe is the ORI (orthorectified raster image) raster catalogue. Before the editing process is launched, a MrSID mosaic is made from the ORIs for a particular mission. These MrSID files are then added to an unmanaged raster catalogue stored in an ESRI personal geodatabase. This ORI raster catalogue layer facilitates management of the DMSe, since as soon as a new MrSID file is added to the raster catalogue, it is instantly available in the DMSe without the need to republish the PMF file (see Figure 2). Another use of the raster catalogue technology is in displaying Intermap's targeted editing masks. These masks are Boolean rasters used to aid operators in editing. Again, following the same steps as for the ORIs, the targeted editing masks can easily be added to an existing raster catalogue and become immediately available to the DMSe users.

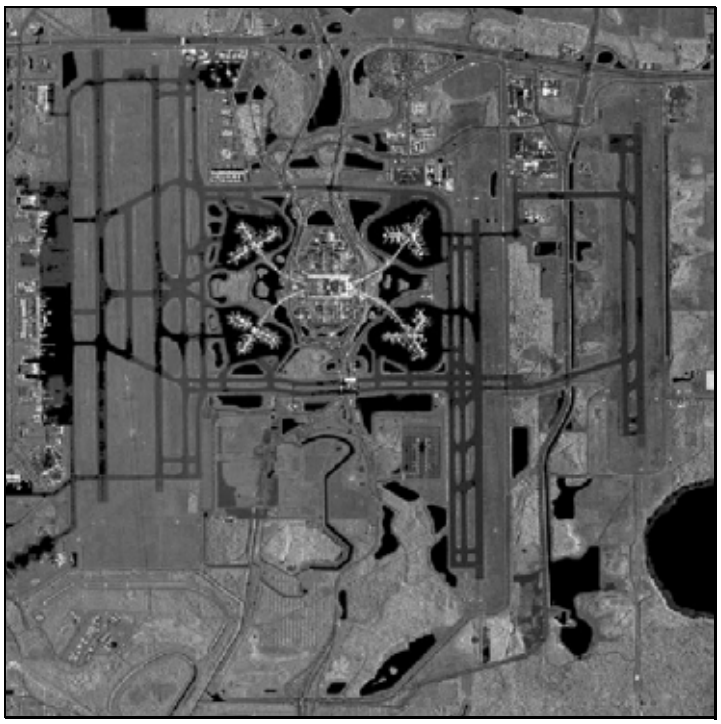

Figure 2: Orthorectified Raster Image catalogue.

Live Database Join

The biggest improvement of the current DMSe version was the addition of a live connection or join to the production database. In the past, the SQL Server exported data out to a flat file on a 5-minute basis, and ArcReader connected to this file to fetch the tile's production attribute information. This was a limitation of the stand-alone version of ArcReader 8.x. While joins could be made directly to a database in ArcGIS and then published to an ArcReader PMF file, there were issues when it came time to view that joined information in ArcReader. With the release of ArcGIS 9.0 came the ability to save the database connection info (i.e. username and password) with the PMF file (see Figure 3). Thus, the current version of the DMSe replaced the flat file join with a direct join to the database. Now changes made to the database by the IES software or the DMS tools could be immediately viewed in the DMSe.

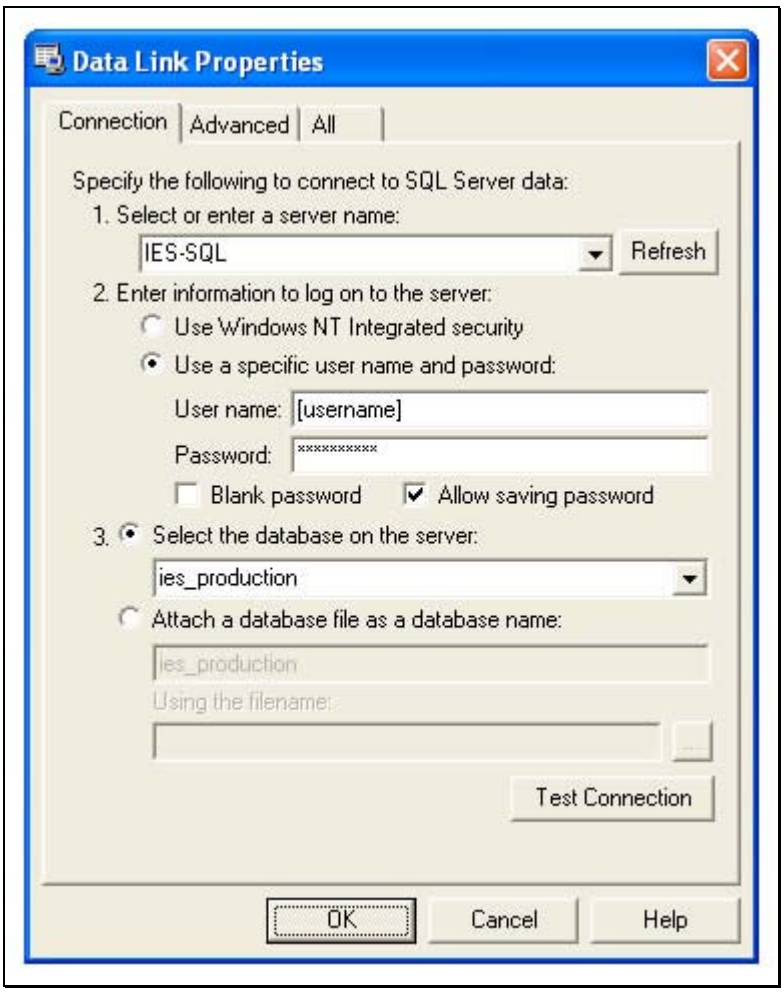

Figure 3: Database Connection Properties.

Since the production database stored important tile information in a multitude of tables (see Figure 4), creating more than one join in ArcMap significantly increased the amount of time required to redraw a layer. Additionally, there was also an increased demand on ArcReader to symbolise tiles based on information stored in different joined tables. For this reason the ArcMap layer was joined to a view created on the SQL Server. A view can be thought of as a virtual table or a stored query. The data accessed through a view is not stored in the database as a distinct table, but rather as a SELECT statement. When a call is made to the view, the data from the tables in the view's SELECT statement is assembled on the fly. This moves the data-crunching burden from the machine running ArcReader to the database server, which dramatically increases the redraw speed.

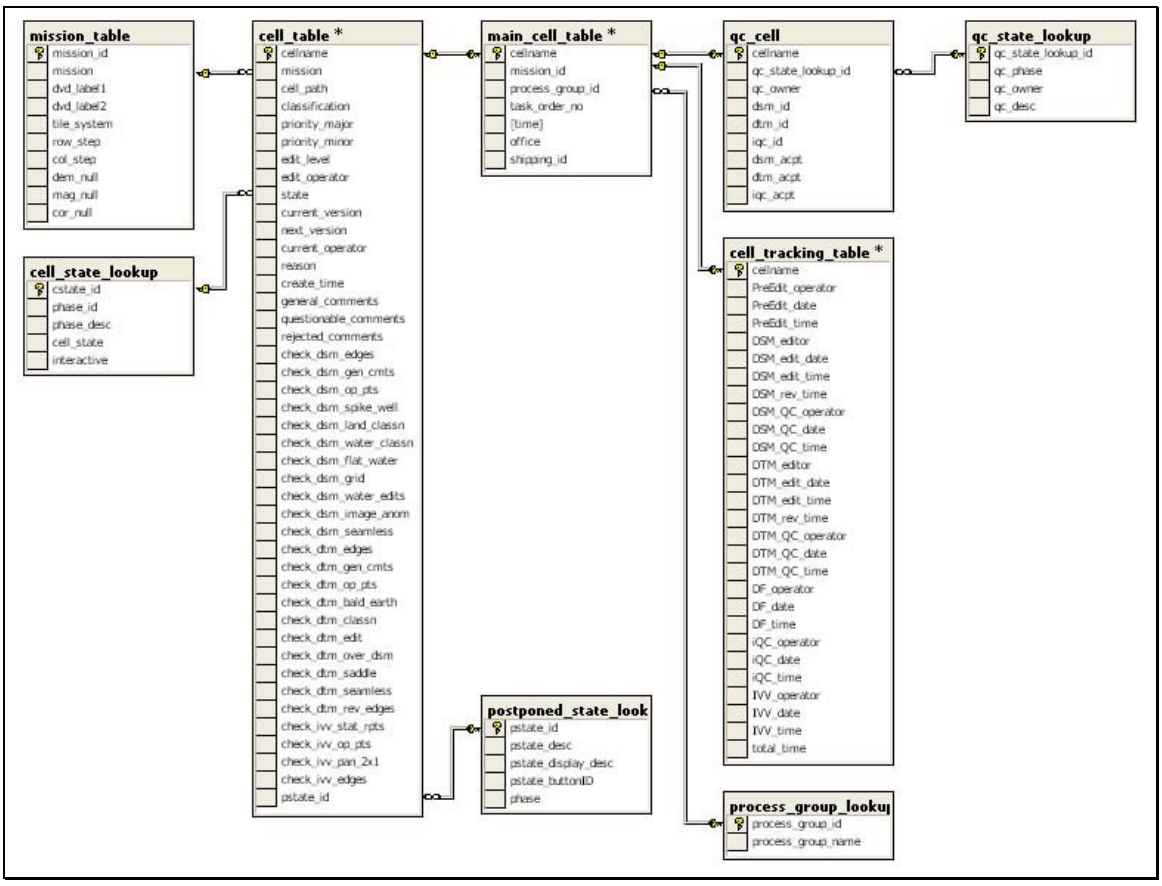

Figure 4: DMSe Database Table Relationships.

A further step was taken with the SQL Server view: creating symbolization classes on the server side. CASE statements are special Transact-SQL expressions that allow the user to return alternate values when a list of conditions is met (see Figure 5). CASE statements allowed the SQL Server to determine under what symbolization class tiles fell (i.e. “Ready for DSM Edit”) depending on the values stored in the various database tables. The symbolization class was then returned in an integer “DMSe_id” column. This noticeably reduced drawing time, as ArcReader no longer had to look at multiple columns to determine how to symbolize a tile.

```
-- DMSe id
CASE    -- Error and Reject states ------------------------------------------------------------------------------------
        WHEN (ct.state IN (0,5,10,15,23,32,38,43,45,49)) OR
             (qc.qc_state_lookup_id=999)                    THEN 9000       -- Error
        WHEN (ct.state IN (11,24,33,39,54))                 THEN 8000       -- Reject

        -- Ingest states -----------------------------------------------------------------------------------------------
        WHEN (ct.state BETWEEN 0 AND 15)                    THEN 100        -- Ingest/Screen/DSM Auto

        -- DSM Processing states ---------------------------------------------------------------------------------------
        WHEN (ct.state IN (16,17))                          THEN 200        -- DSM: Ready for Edit
        WHEN (ct.state=19) AND (ct.pstate_id IN (0,1))      THEN 201        -- DSM: Edit Started/Incomplete
        WHEN (ct.state=19) AND (ct.pstate_id=2)             THEN 203        -- DSM: Ready for QC
        WHEN (ct.state=19) AND (ct.pstate_id=8)             THEN 206        -- DSM: QC Completed

        -- DTM Processing states ---------------------------------------------------------------------------------------
        WHEN (ct.state IN (25,26))                          THEN 300        -- DTM: Ready for Edit
        WHEN (ct.state=20) AND (ct.pstate_id IN (0,1))      THEN 301        -- DTM: Edit Started/Incomplete
        WHEN (ct.state=28) AND (ct.pstate_id=2)             THEN 303        -- DTM: Ready for QC
        WHEN (ct.state=28) AND (ct.pstate_id=8)             THEN 306        -- DTM: QC Completed

        -- Data Finish states ------------------------------------------------------------------------------------------
        WHEN (ct.state IN (31,34,35,36))                    THEN 400        -- DF/DF Edge Started
        WHEN (ct.state=37)                                  THEN 401        -- DF Edge Hold

        -- Data Delivery states ----------------------------------------------------------------------------------------
        WHEN (ct.state IN (40,41,42))                       THEN 500        -- Ready for Data Delivery
        WHEN (ct.state IN (44))                             THEN 501        -- Data Delivery Done

        -- Auto iQC states ---------------------------------------------------------------------------------------------
        WHEN (ct.state IN (46,47,48))                       THEN 600        -- iQC: Ingest/Auto Proc

        -- Inter iQC states --------------------------------------------------------------------------------------------
        WHEN (ct.state IN (50,51))                          THEN 700        -- iQC: Ready for iQC
        WHEN (ct.state IN (52,53))                          THEN 704        -- iQC: iQC in Progress
        WHEN (ct.state=55)                                  THEN 706        -- iQC: iQC Completed

END AS DMSe_id
```

Figure 5: Transact-SQL Query for DMSe View.

Another advantage of using CASE statements is the ability to quickly adapt to new tile states or DMS tools. If the new state should fall under one of the existing symbolization classes, the specific condition in the CASE statement just has to be modified to include the new state. The ArcReader PMF project does not need to be republished. Should a new symbol class be needed, another condition can be added to the CASE statement and the new colouration scheme added to the ArcMap document. In this case, the PMF file will need to be republished.

In order to quickly show which missions are currently being worked on, another view was added to the DMSe project: the aptly named Active Missions view (see Figure 6). This view dynamically calculates the number of tiles in a mission, and how many of those tiles are currently in the editing process. The view returns a table of only the active missions to the DMSe. The mission shapefile is joined to this view, but with the “Keep only matching records” option under the Advanced menu. As new missions come online or old ones go offline, this Active Missions layer redraws accordingly.

Figure 6: Active Mission Layer in the DMSe.

Multi-Site Functionality

Given the nature of business today, often times companies will be faced with the task of managing and controlling production at a multitude of different production centres. Intermap is no different from other companies in this respect, as they are currently faced with maintaining three production centres. The IES editing software, being controlled by a database, and given that a disconnection of one production centre from a localized database could be costly, has gone the route of giving each production centre their own database to work from. As such, managing production over three databases poses its own limitations. The DMSe had to connect to these three databases, of which one lies on a separate network. As such, three separate versions of the DMSe were created, with each office receiving its own copy. For users with VPN access between the two networks, a full version was published, allowing all production centre layers to be displayed simultaneously.

Comprehensive Symbolization

One of the major goals of the DMSe was to facilitate understanding the location of editing for all users. Since the roles of the DMSe users are pretty varied, this posed a challenge. Upper level managers do not need to see as much detail about the interactive editing as, for example, an editor. For this reason a 2-tiered approach to symbolization was embraced. The first tier was the base colour of the tile. The three major steps in the editing process that some users would want to drill down in more detail (DSM edit, DTM edit, and the Independent QC, or iQC, step) were given each their own colour: DSM = green, DTM = blue, and iQC = purple. The lighter shades of each colour signified that

the tile was ready for that step. The medium shade indicated that the work was in progress, and the darker shade indicated that the work was complete. All other processes outside of the interactive editing steps were given distinct colours outside of the DSM / DTM / iQC choices.

Additionally, since there are many processes in the “in progress” step, the second symbolization tier comprised using different hatching levels. No hatching signified editing was in progress; checkmarks meant the tile was ready for QC, triangles signified that QC was in progress, and diagonal orange hatching signified revisions were needed for the tile. Additionally, any tiles postponed as “Guidance Required” in the IES editing software were marked with question marks, thus allowing production coordinators or Team Leads to quickly identify tiles needing attention.

Figure 7: Tile Symbolization Using Comprehensive Colours and Hatching.

Operators requiring more information on a tile can drill down further and, using the Identify tool, call up the production information. Attributes such as the general editing phase (i.e. DSM or DTM) or state (i.e. “DSM QC in Progress”), the production centre responsible for the tile, the current operator, or any comments

written on the tile can all be viewed from here. Operators' names and dates are also gathered for a tile when that specific process is finished. In Figure 8, the "dsm_edit" and "dsm_qc" fields correspond to the operators responsible for that step, as well as the dates they finished the work. This also allows operators to query tiles beside the ones they are working on to see how the edit is progressing, as well as the person to contact should they need to talk to them for one reason or another.

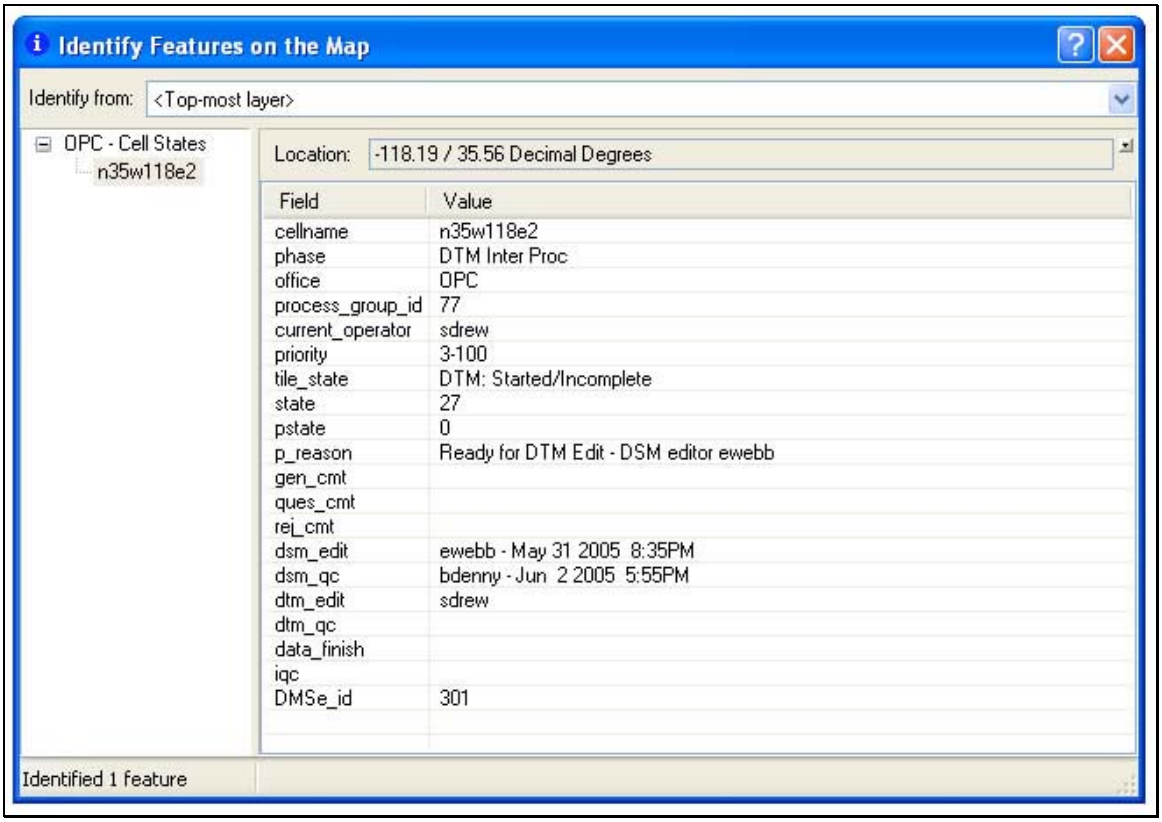

Figure 8: Addition Tile Information.

DMS Integration

The final big requirement of the latest DMSe was to provide a better integration with the various DMS tools. While most of the DMS tools interacted solely with the IES software tables in the database, new developments in the DMS toolset brought new tables in the database that were not tapped by the DMSe. One such example is the QC Control tool, which is used during the QC process to track error locations, as well as the type of errors committed. This tool facilitates Quality Control metrics gathering, as it incorporates into the IES editing software platform. QC Control allows operators to enter operator points in the IES software with an associated QC error code. These operator points are then uploaded into the database, where additional information from the QC Control

tool is stored. These points are immediately viewable in the DMSe, and are symbolized by their rework status: “Rework Needed” = star, “Rework Completed – Awaiting Approval” = diamond, and “Rework Approved” = circle. Using the Identify tool in the DMSe brings up supplementary information such as the latitude and longitude of the point, to aid locating it in the IES software, as well as QC error descriptions, usernames of QC and rework operators, and the upload and rework dates. Furthermore, QC points are displayed using the same colouration scheme as the editing phases: DSM = green, DTM = blue, and iQC = purple.

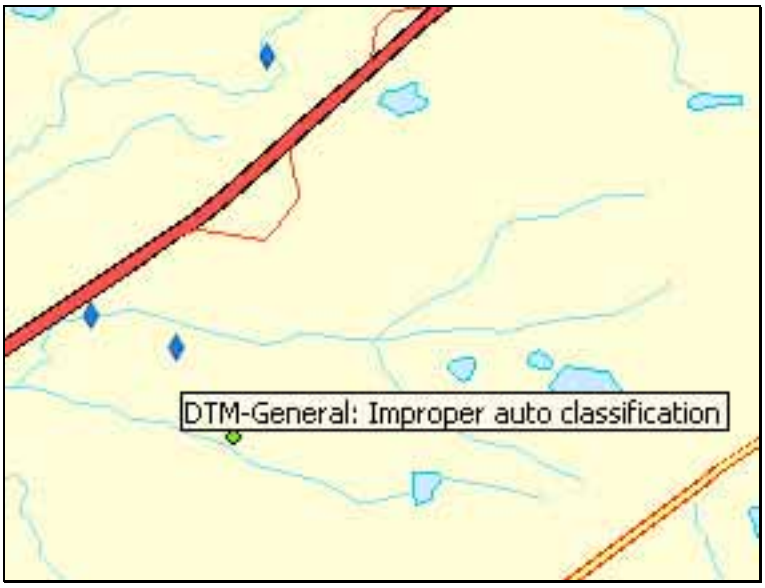

Figure 9: QC Control Error Points in DMSe.

While the IES editing software stores the operator points’ latitude and longitude in degrees – minutes – seconds, QC Control adds additional fields to the database table for storing the lat / long in decimal degrees, therefore facilitating getting the points into other software packages. The QC Error points were added to the DMSe by adding an X-Y table, and pointing ArcMap to the proper fields in the QC Control table in the database. As soon as new QC Error points are uploaded into the database through QC Control, they are immediately viewable in the DMSe. When an editor receives notice that they have a revision to complete, by viewing the QC Error points first in the DMSe, along with the ORI layer, they can quickly gauge how much time to budget for revisions. Managers can, at a glance, view how the editing for an area is coming along, as well as the general types of errors being committed by operators, and adjust accordingly.

Conclusion

Lessons Learned

As previously stated, with the release of ArcGIS 9.0, improved database connection parameters greatly improved the DMSe functionality by allowing a live connection to the database, rather than the 5-minute delay of the previous version. Moving the data-crunching strain to the database further improved the performance of the DMSe, as it was no longer necessary for ArcReader to deal with determining symbolization from multiple fields. The use of dynamic views also facilitates adjusting to new states or procedures in the IES editing software. Using a comprehensive symbolization plan also makes it easier for a multitude of user-access levels to effortlessly, at a glance, determine the level of information they need.

Limitations

The IES editing software is a memory intensive program, for just reasons. ArcReader itself is another burden on the system. Certain workstations that don't perform as well as others, for one reason or another, will occasionally work better if ArcReader is not used by the operator. Whether a problem resides in the installation of ArcReader or the computer itself, in either case the operator of that computer is presented with an issue that hinders his or her performance. For these reasons, the development of an ArcIMS solution for tracking production flow would take the processing burden off the editing workstations and put it in the hands of a dedicated web-server. Should time and budgets allow, Intermap might very well embrace a web-based version of the DMSe.

The default legend size as well caused some issues with properly displaying the symbolization pattern chosen for this project. Due to the small size of the legend boxes, certain hatching patterns were not visible. Since the DMSe uses the same base colouration for tiles that are considered "in progress", this made it difficult at first to understand what each hatching pattern referred to. Over time, this problem has diminished, as users become more familiar with the patterns. Perhaps a future modified version of the DMSe will contain a better reference for users.

Advantages

Certainly, the incorporation of many different levels of information into one all-inclusive package has greatly facilitated the task of managing production. All levels of users, from VPs to editors can open the DMSe, and quickly get the information they need about an area. The integration of DMS tool data into the DMSe has also led to a more thorough understanding of the production flow.

Whereas in the past certain steps in the process may not have been visible through the DMSe, the new version completely amalgamates all DMS – tile specific information into its symbolization.

The ability for managers and coordinators to view information from different production centres facilitates the ability to easily and quickly track what's happening in production, and to adjust accordingly. Whereas in the past it was commonplace for discussions of production progress to occur over email, now with the click of a button information can be viewed by any party with less chance for confusion.

Intermap has embraced the challenge of providing first-class, countrywide datasets to its clients in a timely manner. Modifications were made to the editing software to better handle the dramatically increased data flow through production. Whenever the flow of data through a production centre increases, it becomes all the more important that management and coordinators have an easy way to better control that data flow. The DMSe has fulfilled those requirements completely, by providing a leading edge tool for the leading edge NEXTMap production environment.

Author Information

Richard Asselin

Geospatial Database Applications Developer
Intermap Technologies Corp.
2 Gurdwara Rd, Suite 200
Ottawa, ON K2E 1A2
Phone: (613) 226-5442
Fax: (613) 226-5529
Email: rasselin@intermaptechnologies.com