Microsoft TerraServer: A Spatial Data Warehouse

Tom Barclay
Jim Gray
Don Slutz

June 1999

Technical Report
MS-TR-99-29

Microsoft Research
Advanced Technology Division
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052

Microsoft TerraServer: A Spatial Data Warehouse

Tom Barclay. Jim Gray, Don Slutz
{TBarclay, Gray, Dslutz}@microsoft.com

Microsoft Research, 301 Howard St., Suite 830,San Francisco, CA 94105
http://research.microsoft.com/barc

Abstract

Microsoft® TerraServer stores aerial, satellite, and topographic images of the earth in a SQL database available via the Internet. It is the world's largest online atlas, combining five terabytes of image data from the United States Geological Survey (USGS) and SPIN-2. Internet browsers provide intuitive spatial and text interfaces to the data. Users need no special hardware, software, or knowledge to locate and browse imagery. This paper describes how terabytes of "Internet unfriendly" geo-spatial images were scrubbed and edited into hundreds of millions of "Internet friendly" image tiles and loaded into a SQL data warehouse. Microsoft TerraServer demonstrates that general-purpose relational database technology can manage large scale image repositories, and shows that web browsers can be a good geospatial image presentation system.

1. Overview

The TerraServer is the world's largest public repository of high-resolution aerial, satellite, and topographic data. It is designed to be accessed by thousands of simultaneous users using Internet protocols via standard web browsers.

The TerraServer is a multi-media data warehouse. It differs from a traditional data warehouse in several ways: (1) it is accessed by millions of users, (2) the users extract relatively few records (thousands) in a particular session and, (3) the records are relatively large (10 kilobytes). By contrast, classic data warehouses are (1) accessed by a few hundred users via proprietary interfaces, (2) queries examine millions of records, to discover trends or anomalies, (3) the records themselves are generally less than a kilobyte. In addition, classic data warehouse queries may run for days before delivering results. Initial results typically cause users to modify and re-run queries to further refine results.

One thing the TerraServer has in common with classic data warehouses is that both manage huge databases: several terabytes of data. Terraserver's topographic maps cover all of the United States at 2 meter resolution 10 million square kilometers), the aerial photos cover 30% of the United States today (3 million square kilometers at one-meter resolution, and 1% of the urban areas outside the United States (1 million square kilometers) at 1.56 meter resolution.

This report describes the design of the TerraServer and its operation over the last year. It also summarizes what we have learned from building and operating the TerraServer.

Our research group explores scaleable servers. We wanted first-hand experience building and operating a large Internet server with a large database and heavy web traffic. To generate the traffic we needed to build an application that would be interesting to millions of web users. To have a huge database, we needed a huge data source: trillions of bytes that are relatively inexpensive to acquire and process.

Based on our exposure to the EOS/DIS project [Davis94], we settled on building a web site that serves aerial, satellite, and topographic imagery. We picked this application for three reasons:

1. The web is inherently a graphical environment, and these images of neighborhoods are recognizable and interesting throughout the world. We believed this application would generate the billions of web hits needed to test our scaleabilty ideas.
2. The data was available. The USGS was cooperative, and since the cold war had ended, other agencies were more able to share satellite image data. The thaw relaxed regulations that had previously limited the access to high-resolution imagery on a global basis.
3. The solution as we defined it – a wide-area, client/server imagery database application stored in a commercially available SQL database system – had not been attempted before. Indeed, many people felt it was impossible without using an object-oriented or object-relational system.

This paper describes the application design, database design, hardware architecture, and operational experience of the TerraServer. The TerraServer has been operating for a year now. We are just deploying the third redesign of the database, user interface, and online image loading system.

2. Application Design

Microsoft TerraServer is accessed via the Internet through any graphical web browser. Users can zoom and pan across a mosaic of tiles within a TerraServer scene. The user interface is designed to function adequately over low-speed (28.8kbps) connections. Any modern PC, MAC, or UNIX workstation can access the TerraServer in this way. If you have never used it, look at the TerraServer web site at http://terraserver.microsoft.com/.

Imagery is categorized into "themes" by data source, projection system, and image type. Currently, there are three data themes:

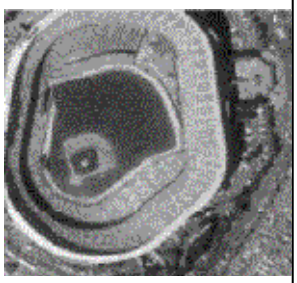

Figure 1. A USGS DOQ Image of 3Com Park near San Francisco

USGS Digital Ortho-Quadrangles (DOQ) are gray-scale, 1-meter resolution aerial photos. Cars can be seen, but 1-meter resolution is too coarse to show people. Imagery is orthorectified to 1-meter square pixels. Approximately 50% of the U.S. has been digitized. The entire conterminous U.S. is expected to be completed by the end of 2001.

Figure 2. A USGS DRG 2-meter resolution image

USGS Digital Raster Graphics (DRG) are 13-color digitized topographic maps, with scales varying from 2.4 meter resolution to 25.6 meter resolution. DRGs are the digitized versions of the popular USGS topographic maps. The complete set of USGS topographic maps have been scanned including Alaska, Hawaii, and several territories such as Guam and Puerto Rico.

Aerial Images SPIN-2™ are grayscale 1.56 meter resolution Russian satellite images. The images are re-sampled to 1-meter resolution. Microsoft TerraServer contains SPIN-2 images of Western Europe, the United States, and the Far East. Unfortunately, there is little coverage of Canada, South America, Africa, and Southeast Asia.

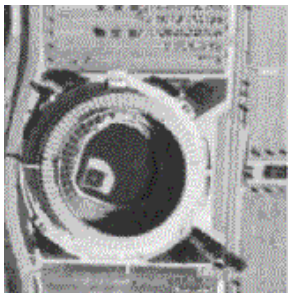

Figure 3. a SPIN-2: 1.56-meter image of Atlanta's Fulton County Stadium.

2.1 Projection Systems and Scenes

The earth is a bumpy ellipsoid. Maps and computer monitors are flat. It is impossible to accurately present a spherical object on a flat surface.

Cartographers have addressed this issue by developing projections of the geoid onto flat surfaces [Robinson95]. There are many projection systems, each suited to present certain regions or properties. Multiple images in a projection system can often be joined together to form a seamless mosaic within certain boundary conditions. These mosaics either have extreme distortion as they scale out, or they introduce seams.

DOQ and DRG data are projected by the USGS into Universal Transverse Mercator (UTM) projection using the North American Datum (NAD) ellipsoid created in 1983 [Snyder89]. UTM is a projection system that divides the earth into 60 wedge shaped *zones* numbered 1 thru 60 beginning at the International Date Line. Each zone is 6 degrees wide and goes from the equator to the poles. UTM grid coordinates are specified as zone number, then meters from the equator and from the zone meridian[1].

The conterminous United States is divided into 10 zones (see Figure 4). Each of these UTM zones is a *scene*. The TerraServer mosaics each scene, but two adjacent scenes are not mosaiced together. Users can pan and zoom within a scene, and can jump from one scene to another.

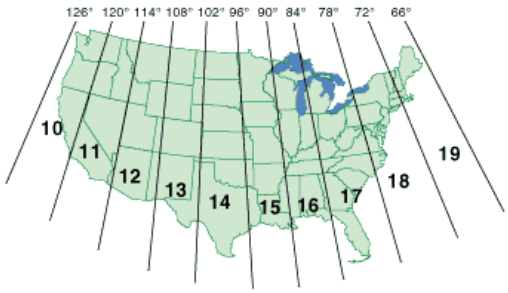

Figure 4: The ten UTM zones in the continental United States.

The Russian SPIN-2 imagery is digitized from Russian satellite photographs. The Russian satellite captures 160km wide by 40km high areas in a single image. The satellite takes one image and then begins the adjacent image, overlapping the last image. The overlap is variable, and when digitized does not line up on a pixel boundary.

To create a seamless mosaic of SPIN-2 imagery, all SPIN-2 imagery would have to be orthorectified. This requires precise geo-location of each image, which was not possible due to security concerns. Without rectification, if tiles extracted from separate SPIN-2 satellite images are mosaiced, the tile edges are misaligned. Roads, rivers, and other geographic features do not line up. While GIS experts may tolerate this, it is disorienting and unacceptable to novice users.

[1] Actually, UTM grid units can be in inches, feet, meters, or kilometers. The USGS chose meters for most of their assets in the UTM projection. UTM is not used above 80N or 70S [Robinson95].

Consequently, the TerraServer treats each 160km x 40km SPIN2 image as a separate scene. These scenes are not mosaiced together. Users can pan and zoom within a scene, and can jump from one scene to another.

2.2. TerraServer Grid System

Users can zoom and pan across a mosaic of tiles within a TerraServer scene. The tiles are organized in the database by theme, resolution, scene, and location within a scene.

TerraServer is designed to support a fixed number of resolutions in powers of 2 from 1/1024 meters per pixel (scale 0) through 4096 meter (scale 22). The scale is related to resolution in meters per pixel by

$$\text{Scale} = \log_2(\text{resolution}) + 10$$

The highest resolution images currently in the database are one meter per pixel, which is scale 10. Coarser resolutions are derived by sub-sampling fine-resolution images.

For UTM projection data-sets, the SceneID is the UTM zone assigned to the original image a tile's pixels were extracted from. For SPIN2 data-sets, a unique SceneID is assigned by TerraServer for as each scene is loaded.

Each TerraServer scene is planar. A tile can be identified by its position in the scene. The tile loading program assigns a relative X and Y tile identifier to each tile as it is loaded.

For UTM projected data, the X and Y tile address is the UTM coordinate of the top-left pixel in the tile divided by the tile image size in UTM units in meters. The following are the formulas:

X = TopLleftUTM_X / (TilePixWidth • Resolution)

Y = TopLeftUTM_Y / (TilePixHeight • Resolution)

For SPIN2 scenes, the X and Y tile addresses are relative to the upper left corner of the scene.

The six fields – Resolution, Theme, SceneID, X, and, Y - form the unique key by which any TerraServer image tile can be directly addressed. Each TerraServer web page contains image tiles from a single Theme, Scale, and SceneID combination. For example, our building in USGS DOQ theme (T=1), has scene UTM zone 10 (S=10), at scale 1 meter (Z=10) with X=2766 and Y=20913. The URL is: http://terraserverv.microsoft.com/tile.asp?S=10&T=1&Z=10&X=2766&Y=20913.

The TerraServer search system performs the conversion from geographic coordinate systems to the TerraServer coordinate system. The TerraServer image display system uses TerraServer grid system coordinates to pan and zoom between tiles and resolutions of the same theme and scene.

2.3. Imagery Database Schema

Each theme has an ***OriginalMeta*** table. This table has a row for each image that is tiled and loaded into the TerraServer database. The *OrigMetaTag* field is the primary key. The meta-fields vary widely from theme to theme. Some of the meta fields are displayed by the Image Info Active Server Page (for example, see http://terraserver.microsoft.com/GetOrigMeta.asp?OrigMetaId=104578&SrcId=1&Width=225&Height=150&ImgSize=0&DSize=0)

All the image tiles and their metadata are stored in an SQL database. A separate table is maintained for each (theme, resolution) pair so that tiles are clustered together for better locality. USGS DOQs have resolutions from 1-meter through 64-meter. USGS DRG data supports 2-meter through 128-meter resolution. SPIN supports resolutions from 1-meter to 64-meter.

Each theme table has the same five-part primary key:

- *SceneID* –individual scene identifier
- *X* – tile's relative position on the X-axis
- *Y* – tile's relative position on the Y-axis
- *DisplayStatus* – Controls display of an image tile
- *OrigMetaTag* – image the tile was extracted from

There are 28 other fields that describe the geo-spatial coordinates for the image and other properties. One field is a large "blob" type that contains the compressed image.

These tile blobs are chosen to be about ten kilobytes so that they can be quickly downloaded via a standard modem (within three seconds via a 28.8 modem).

2.4. Gazetteer Database Schema

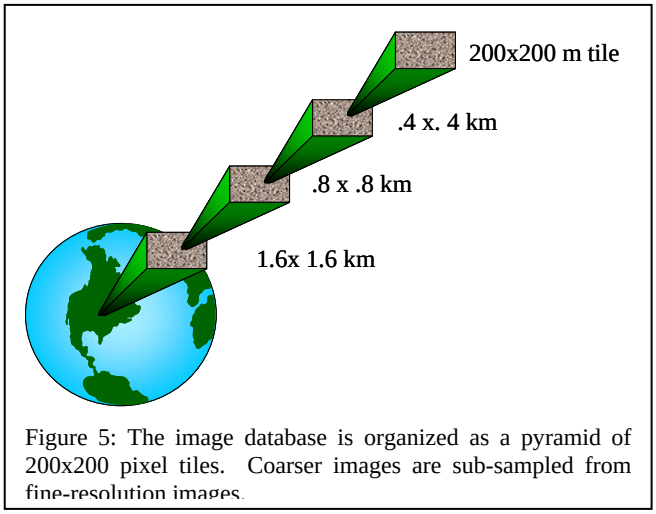

Figure 5: The image database is organized as a pyramid of 200x200 pixel tiles. Coarser images are sub-sampled from fine-resolution images.

The Gazetteer lets users find images by name. It contains the names for about 1.5 million places, with many alternate spellings. It is a simplified version of the Gazetteer found in the Encarta Virtual Globe™ and Microsoft Streets™ products.

The Gazetteer Schema has a snowflake structure. *Place* is the center table. It contains the formal name for a unique place on earth and maps the uniquely named location to the TerraServer Grid System. The *AltPlace* table contains all the synonyms of a unique place. The *State* and *Country* parent tables identify a place's state/province and country. The *AltState* and *AltCountry* tables contain the state/province and country synonyms.

DOCKET ALARM

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts

Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research

With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips

Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS
Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS
Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS
Sync your system to PACER to automate legal marketing.