SPATIAL DATA STORAGE AND TOPOLOGY IN THE REDESIGNED MAF/TIGER SYSTEM

David Galdi

Abstract: *Since the late 1980's, the Geography Division has utilized the TIGER® system to provide geographic support for Census Bureau surveys, censuses, estimates, and partnership programs. While this system has served its purpose well, changing technology and requirements are dictating its replacement. The new system will utilize commercial off-the-shelf software, including Oracle Spatial and Oracle Spatial Topology Data Model. As with its predecessor, the new system will utilize persistent topology. It will also merge the largely non-spatial Master Address File with the spatial TIGER database to form a new database. This paper examines issues related to topology and selection of a spatial data storage mechanism for the new database, as well as issues regarding integration of spatial and non-spatial data. It also addresses some of the early design implementation decisions for the new system.*

BACKGROUND

The Census Bureau has a long history of innovation in data collection and processing techniques. The Geography Division of the Census Bureau has pioneered efforts in the field of automated cartography, including development of the Address Coding Guide in the 1960's, the GBF/DIME files in the 1970's, and the Topologically Integrated Geographic Encoding and Referencing (TIGER®) system in the 1980's (Marx 1986).

The TIGER system, still in use 20 years after its inception, utilizes topological data structures based on the landmark work by Corbett and White (1979, 1981), which relates the mathematics of graph theory and topology to the storage of geographic and cartographic data. A significant feature of the TIGER system was the ability to automate production of all geographic support products from a single integrated digital database (LaMacchia 1990). The TIGER system was built using a database management system (DBMS) called TIGERdb. Geography Division staff developed the TIGERdb DBMS, along with associated software to support spatial functionality and indexing, and automated maintenance of the topological data (Boudriault 1987). The use of topology, stored persistently in the database, improved the efficiency of spatial data retrieval and data storage and helped enforce implementation of data integrity and consistency rules.

MAF/TIGER REDESIGN

While the TIGER system has served its purpose well, the time has come for a change. The homegrown database system does not integrate well with current commercial off-the-shelf (COTS) tools and Web technology. It is cumbersome to change, difficult to learn for new developers, does not allow multi-user access, and is not accessible via a standard query language.

The MAF/TIGER Redesign project, which is one of the five components of the MAF/TIGER Enhancement Program (MTEP), will result in a seamless national integrated database that will utilize commercial database software in place of TIGERdb. To the extent possible, homegrown applications will be replaced by COTS software.

Oracle's relational database management system will be used to store all data in the new database. Spatial data management, including spatial operations and indexing, will utilize Oracle Spatial. The Oracle Spatial Topology Data Model will store and manage geographic features and the topological data structures on which they are built.

MAF/TIGER Functionality and Content

The MAF/TIGER system provides geographic services in support of Census Bureau surveys, censuses, estimates, and partnership programs. These services include:

- Delineation and Maintenance of Geographic Areas
- Mapping
- Spatial Data Exchange
- Address Geocoding
- Address Matching
- Creation of Geographic Data Extracts
- Creation of Geographic Comparability and Equivalency Files
- Creation of Address File Extracts
- Spatial Analysis

Provision of these services requires ongoing maintenance of the MAF and TIGER databases, which are being combined as a part of the redesign. The content of these databases is described in the following sections.

TIGER

The TIGER database includes geographic features such as roads, railroads, geographic areas, landmarks, waterways, and other geographic information that is needed to support the programs of the Census Bureau. The database serves as the repository for all of the geographic information needed for census and survey data collection, data tabulation, data dissemination, geocoding services, geographic and statistical analysis, and production of maps.

MAF

The MAF is designed to be an accurate, national residential address inventory and supports data collection efforts and questionnaire delivery to each residence. It provides for storage of a mailing address to support questionnaire mail-out, as well as location information to support personal interviews and address canvassing operations. Location information might consist of a city-style address, a latitude/longitude coordinate, a census block number, an E-911 address, and/or a textual location description.

While the existing TIGER system contains primarily spatial data, the integrated MAF/TIGER system will combine TIGER with the largely non-spatial MAF database. This mirrors a common trend in the Information Technology (IT) and Geographic Information System (GIS) industries: the integration of spatial and non-spatial data into a single enterprise data set.

Spatial Data in a Relational Database Management System

Historically, spatial data and GIS applications have been developed as stand-alone systems for spatial analysis and map production. Often, the spatial data in these "stovepipe" systems are stored in a separate, proprietary database, on separate hardware, and require proprietary data update and access tools, as well as separate archiving, system maintenance, and tuning. For many organizations, GIS projects represent a major integration challenge consuming a large portion of the IT budget. Much of this is directed towards integration between spatial and non-spatial systems (Batty 2004).

Over the last decade, however, isolated "stovepipe" GIS systems have begun to be integrated with mainstream IT and used to support location awareness in additional business applications. Database vendors have partnered with GIS vendors to make a concerted effort to ensure that spatial data can be blended seamlessly into the enterprise database (Gonzales 2000). When the relational database management system (RDBMS) is expanded to handle spatial data, core database capabilities, such as scalability, security, versioning, and replication, can be extended to spatial datasets. An open spatial database also allows for spatial enabling of many enterprise applications with associated improved functionality (Weinberger 2002).

One of the key components of incorporating spatial functionality into an RDBMS is implementation of spatial data types. Standard RDBMS attributes support data types such as character, date, and integer. In a spatial database, additional data types are required that can be used to represent point, line, and area features. In addition to data types, spatial databases must provide support for spatial indexing and clustering, as well as spatial operators (Van Oosterom 2002, Batty 2004). Spatial operators include:

- Predicates that return Boolean values (e.g., in_window, equal, adjacent, inside)
- Functions that return numbers (e.g., area, perimeter, centroid, distance)
- Operators that return new spatial objects (e.g., union, intersection, difference) (Schneider 2002).

The OpenGIS Consortium (OGC) recognized the importance of the integration of spatial data into the IT mainstream by standardizing the basic spatial data types and functions in the Simple Feature Specification (Van Oosterom 2002). The OGC Simple Feature Specification defines a "standard SQL schema language that supports storage, retrieval, query, and update of simple geospatial features via the ODBC API" (Open GIS Consortium 1999). Simple features are based on 2D geometry and have both spatial and non-spatial attributes.

Oracle, the database vendor selected for the MAF/TIGER Redesign, has been a leader in integration of spatial data and functionality into the RDBMS, maintaining a dominant share of the market for geospatial database management (IDC 2003). Utilization of Oracle and Oracle Spatial for the redesigned MAF/TIGER database will ensure a single database in which spatial data is stored seamlessly with the associated feature attributes and with non-spatial datasets. This will allow for an integrated and improved approach to scalability and data base management that includes replication, archiving, and tuning. It will also provide for spatial enabling of enterprise data that includes the MAF. Oracle Spatial adheres to the OpenGIS Simple Features Specification for SQL, Revision 1.0, Normalized Geometry (Oracle Corporation 2004).

WHAT IS TOPOLOGY?

Topology involves the mathematical study of spatial relationships. It describes the characteristics of a geometric figure that do not change under continuous transformation. In a graph, the number of line segments, intersection points, and polygons, and their relationship to each other, are constant as the plane in which they exist is stretched or distorted. In GIS applications, topology is the means to describe, manage, and retrieve these relationships explicitly without resorting to time-consuming spatial comparisons (Ramage and Woodsford 2002).

The principles of topology are utilized to implement a system that provides for:

- Rapid spatial data retrieval.
- Enhanced spatial data analysis.
- Spatial data editing and clean up.
- Improved data consistency.
- Management of shared geometry (Hoel 2003).
- Definition and enforcement of data integrity rules (Hoel 2003).

Some GIS applications use "persistent topology", that is, they structure the data according to topological principles so that the topological relationships are stored and available persistently in the database. In addition to the above benefits, this persistent topology approach also provides for:

- More efficient storage of spatial data.
- Reduction or elimination of redundancy of spatial data.
- Easier implementation of certain spatial business rules in the database.
- Improved management of hierarchical geographic relationships (Ramage and Woodsford 2002).

- A more straightforward way to address coordinate precision and tolerance issues that sometimes lead to gaps or slivers (Egenhofer et. al. 1989).

PERSISTENT VS. ON-THE-FLY TOPOLOGY

An alternative to persistent topology is called "on-the-fly" topology, in which intelligent client applications compute topology as needed for selected subsets of the database. This can potentially provide most of the benefits of persistent topology, although it does not address the redundant data storage issue. An additional drawback can be the time that it takes to determine the topological relationships, as they are needed. The choice of "on-the-fly" or persistent topology would appear to depend largely on the nature of the GIS system being utilized and maintained. The majority of processing for the MAF/TIGER system is spent on very large batch processes that run on the whole nation and utilize topology to improve performance. Continual "on-the-fly" recalculation of topology to support these programs could prove problematic. In addition, due to the high degree of interdependency between different feature types in the TIGER system, editing becomes more straightforward with a system that minimizes coordinate redundancy.

The redesigned MAF/TIGER system will use the persistent topology data structure that is part of Oracle Spatial, starting with the release of Oracle 10g. This system, called Oracle Spatial Topology Data Model, provides persistent topology to support batch or interactive applications. In addition, since it is implemented via a server-side topology engine, this solution becomes much more interoperable.

Thick, topologically aware clients are not required, so any client applications, even thin web-based ones, can be used to update the database, while the topology engine manages the spatial updates and maintains the topological data structure (Lessware 2004).

THE BUILDING BLOCKS OF TOPOLOGY[1]

Topology works by dividing spatial data into low-level primitives, which form the building blocks for spatial data. These building blocks include the following components:

- Edge
- Vertex
- Isolated Node
- Connecting Node
- Face

---

[1] In this paper, the terms node, edges, and face are used to describe point, line, and area topological primitives. This is the terminology used by the OpenGIS Consortium and Oracle's Topology Data Model. The legacy TIGER system uses the terms 0-cells, 1-cells, and 2-cells, but to avoid confusion, the newer terms are used, even when referring to the legacy system.

The edge is the central component of two-dimensional topology. It is a linear or one-dimensional construct that has a starting point and an ending point. The end points are referred to as connecting nodes. An arbitrary direction is assigned to each edge, allowing designation of one of the nodes as the Start Node, and the other as the End Node. The edge may be defined by the line segment connecting these two nodes or it may have intermediate points called vertices (Figure 1).
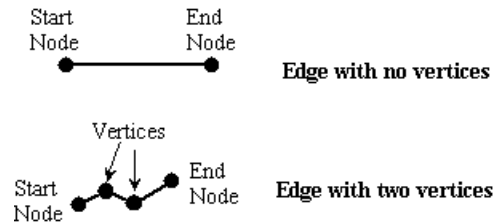
Figure 1. Edges

A face corresponds to a simple polygon bounded by edges, less any "holes" created by the formation of polygons within its boundaries. The formal definition of the face also includes any interior edges that have the face on both the left and right side (Figure 2).
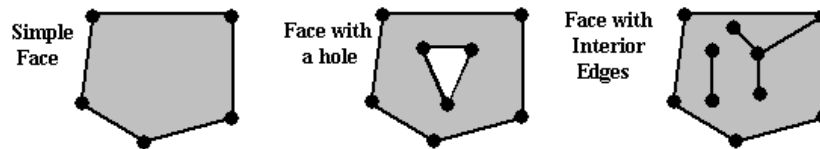
Figure 2. Faces

Because the face is built from edges, no coordinates are stored explicitly for the purpose of representing a face. Rather, the coordinates are stored at the edge level. In order to determine the coordinate geometry of a face, the geometries of the edges that bound the face must be retrieved.

Topologically consistent datasets have the following properties (Kainz 2004):

- Every edge must be bounded by two nodes (start and end nodes).
- Every edge has a left and right face.
- Every face has a closed boundary consisting of an alternating sequence of nodes and edges.
- There is an alternating closed sequence of edges and faces around every node.
- Edges do not intersect each other, except at nodes.

These rules apply to systems with full "polygonal" or "planar" two-dimensional topology, such as is used by the legacy TIGER system and the new MAF/TIGER system. Some applications use less rigorous types of topology. Examples are Network Topology, which consists of only

nodes and edges, but not faces (Ramage and Woodsford 2002) or Winged-Edge Topology, which utilizes only edges and faces (Baumgart 1975).

Another component of topology that is recognized by Oracle Spatial Topology Data Model, but is not covered in the formal topology rules described above, is the isolated node. This is a point that does not have any attached edges. An isolated node might be used to represent a housing unit, a nursing home, or a mountain peak. It is commonly required to identify the relationships between these isolated nodes and other spatial features. For example, one might want to identify the closest road to a structure, or all of the point features within a face or areal feature. If the isolated nodes are not integrated into the topology, these queries can require significant spatial searching. The Topology Data Model allows isolated nodes to be integrated into a topology layer and linked to faces. For the redesigned system, consideration was given to storing the isolated nodes in a separate topology layer, so that they would not interfere with updates to the linear network, especially in regards to snapping and tolerance rules. However, the decision was made to store isolated nodes in the same layer as the other topology components. In addition to the more onerous spatial searching that would be required if isolated nodes were in a separate layer, there could be other disadvantages related to the use of the geometric update methods provided with the Oracle Topology Data Model. When geometric features are moved or reshaped, it is useful to know if the features have "crossed over" housing units, causing a housing unit to end up in a different relative location to the feature. If housing units are stored as isolated nodes and integrated with the other spatial data, then the Oracle Topology Data Model will be aware of this change in the relationship of the road to the housing units, and can inform the calling application, providing the opportunity to abort the change. If the housing units are just stored as coordinates or are isolated nodes in a separate layer, the geometric methods will be unaware of their location or the change in relationship caused by the move. See Figure 3.
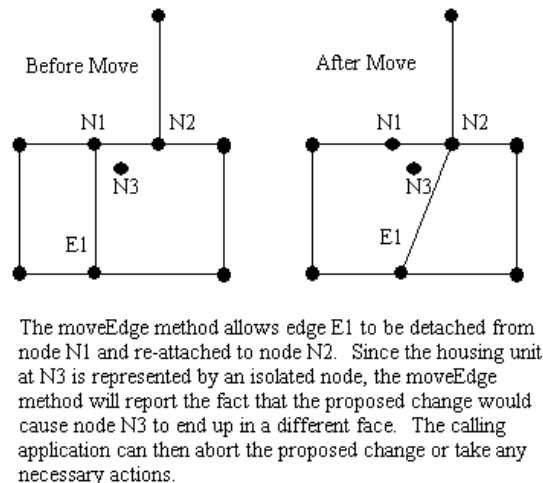

The moveEdge method allows edge E1 to be detached from node N1 and re-attached to node N2. Since the housing unit at N3 is represented by an isolated node, the moveEdge method will report the fact that the proposed change would cause node N3 to end up in a different face. The calling application can then abort the proposed change or take any necessary actions.

Figure 3. Move Edge

TOPOLOGY DIFFERENCES BETWEEN THE LEGACY TIGER SYSTEM AND ORACLE SPATIAL TOPOLOGY DATA MODEL

The storage of topology by Oracle's Topology Data Model is very similar to the legacy TIGER system. There are, however, a few differences:

- Isolated Nodes - As described above, Oracle's Topology Data Model allows isolated nodes to be integrated into the topology layer. The significant benefit of this is that the isolated nodes within a face can be identified through database relationships and without resorting to spatial searches. In addition, the face containing an isolated node can be determined via a database retrieval as well. In the legacy TIGER system, while isolated nodes can be linked to a face, they are not actually part of the topology, and so cannot be readily retrieved as if they were components of a face.
- Ordering of Edges belonging to a Face - In the legacy TIGER system, edges are ordered in a clockwise direction around the exterior boundary of a face. Oracle's Topology Data Model uses a counter-clockwise direction for the edges in the exterior boundary of a face, and a clockwise direction for the edges in the interior boundary of a face, as specified by the Geographic Information – Spatial Schema standard 19107 of ISO TC/211
- Coordinate Storage of Connecting Nodes - While the legacy TIGER topology has no redundant storage of coordinates; a small amount exists with Oracle's Topology Data Model. The coordinates for connecting nodes are stored twice. This is because the Topology Data Model stores each edge as an Oracle Spatial SDO_Geometry data type. This SDO_Geometry for the edge includes all of the coordinates, including the ones corresponding to the start and end nodes. So for these start and end nodes the latitude/longitude coordinate is stored twice; once as part of the SDO_Geometry of the edge and once on the connecting node record. The Oracle Topology Data Model API manages this data transparent to the user and keeps it consistent.

TOPOLOGY MECHANICS

A distinct advantage of using topology derives from the relationships that are maintained between the "topological primitives", i.e., nodes, edges, and faces. These relationships speed queries concerned with adjacency, connectivity, and containment of features or topological primitives. For example, from a connecting node, it is straightforward to retrieve all of the edges that connect to it, through a simple database query. Similarly, it is a simple query to retrieve the faces on either side of an edge as well as to retrieve all of the edges belonging to a face. More complex queries are also fairly straightforward, such as retrieving all the faces that share a given node or all the faces adjacent to a given face. Similar operations are available to determine adjacency, connectivity, and containment of features. Again, these are available as direct database queries, without having to resort to expensive spatial data searches.

Egenhofer, Frank, and Jackson (1989) enumerated all the possible topological relationships that can occur between two spatial entities by examining the combinations of the intersections of the boundaries and interiors of the two objects. The eight possible relationships for objects of the same dimension are:

- Disjoint
- Meet
- Overlap
- Inside
- Contains
- Covers
- Covered By
- Equal

Calculation of these relationships and other functionality is possible in a database that does not utilize topology and stores a complete set of coordinates explicitly to represent each feature. However, calculations in such a database require spatial searches and comparisons rather than more efficient database queries. With the refined use of spatial indices, such as R-trees, spatial searches are getting more efficient, but they still are significantly more time-consuming than other types of data access.

In addition to faster retrieval and more efficient data storage, topology also provides for more efficient and effective data cleansing, error detection, and data integrity. If the topology is stored persistently, storage of redundant data is greatly reduced or eliminated, which simplifies enforcement of data consistency. Topology allows for easy detection of gaps and overshoots, and helps prevent inadvertent overlap of areal features. It also facilitates implementation of a snapping and tolerance system to avoid slivers, arbitrarily close nodes, very small polygons, and so forth.

EXAMPLES OF TOPOLOGY USE AT THE CENSUS BUREAU

Topology is utilized to a significant degree by many of the applications that comprise the Census Bureau's Geographic Support System. The following are some examples.

Areal Delineation

Areal Delineation programs, such as Automated Block Numbering and Automated Assignment Area Delineation, make extensive use of topology, "walking around" nodes or faces, and examining adjacent or connected primitives to determine the prospects for adding additional territory to a candidate block or assignment area.

Interactive Update Software

A common function in interactive update software is for the operator to select a feature on the map. This requires an initial spatial search to determine the closest node or edge of the point

selected, or the containing face. However, once the appropriate topology primitive has been determined, all associated or adjacent features can be readily determined via topological relationships, without having to do additional spatial queries.

Mapping

- Boundary Symbolization
  Topology can facilitate boundary symbolization where multiple features share the same segment of boundary.
- Block Number Placement
  Each polygon on a Census Block map must be labeled with a block number or fish hooked to another polygon. Topology facilitates this determination.
- Thematic Maps
  Assignment of color on thematic maps to assure that adjacent features are not represented by the same color.
- Linear and Areal Feature Chaining
- Expansion of Subject Map Entity Extent
- Processing Feature Density
- Building Inset Area Polygons
- Error Detection and Clean-up During Line Generalization
- Clipping
- Identification and Symbolization of Internal and External Fringe
- Coincident Feature Identification

Spatial Data Matching

Spatial data matching, or conflation, makes extension use of topology. The topological primitives play a central role in feature match recognition (Saalfeld 1993). The Geography Division uses spatial data matching software to support digital exchange with local governments and other sources, and to upload data received as part of the MAF/TIGER Accuracy Improvement Project (MTAIP), which is designed to improve the accuracy of coordinate locations of the road centerline spatial features in TIGER. Permanent identifiers assigned to edges (TIGER/Line ID's) also play a critical role in MTAIP and spatial data exchange activities.

Quality Assurance

Topology can be utilized to identify and correct or report anomalies in the dataset based on entity specific rules for adjacency, coincidence, etc.

"I" of the TIGER (Integration)

The most straightforward way to represent spatial features in a GIS would be to independently store the geometry for each feature. "Geometry" in this context refers to the entire ordered set of latitude/longitude coordinates that represent the location and extent of the feature. In this scenario, features that share geometry would overlay each other resulting in redundant

coordinate storage. However, independent or layered representation of the geometry of spatial features is not optimal if there is a high degree of interrelationship among features and many features share underlying coordinate strings or polygons. For example, roads, rivers, and other linear geography in the TIGER database often also serve as boundaries for geographic areas, such as places or counties. In addition, the Census Bureau manages and maintains the boundaries of over 75 different types of tabulation and collection geographic areas. And for many of these, multiple vintages must be maintained simultaneously. These areas often share portions of their boundaries with each other and/or with linear features. For example, for a given section of a road that serves as a county boundary, it would not be unusual for it to also comprise a boundary for some or all of the following:

- Incorporated Place
- Census Tract
- Block Group
- Census Block
- County Subdivision
- Voting District
- Traffic Analysis Zone
- Public Use Microdata Area
- ZIP Code Tabulation Area
- Regional Office
- Local Census Office
- Collection Assignment Area
- Field Office Supervisor District

If a layered approach were used, the independent storage of geometry for each feature would result in the string of coordinates, representing the section of road, being stored multiple times. If the shape of the road needs to be adjusted or corrected, the geometries for the geographic areas that follow the road must be rebuilt to reflect the change. Alternatively, if topology is used to organize spatial data, the coordinates for a section of road that also bounds geographic areas are only stored once, and the correction becomes more straightforward, although it is still subject to business rules that govern feature consistency.

Features

Features in the new MAF/TIGER database will be built upon the topological primitives. The feature records will contain name, class code, and other attributes, while the location information will be managed via linkage to the topological primitives. All of this linkage occurs through a central relationship table managed by Oracle's Topology Data Model, which allows a many-to-many relationship between features and topological primitives. In general, the spatial features will be purely areal (comprised of faces), linear (comprised of edges), or point (comprised of nodes). However, complex features comprised of more than one primitive type are allowed. So, a river could be represented as a linear feature near its source, while changing to an areal feature as it widens.

In the legacy TIGER database, much of the attribution is assigned at the topological primitive level to nodes, edges, and faces. The redesigned MAF/TIGER makes greater use of higher-level features and assigns attributes to these features where possible. However, certain attributes vary over the extent of a feature and are more appropriately stored at the topological primitive level. Examples of such edge attributes are the Permanent Edge Identifier (TIGER/Line ID), Number of Lanes (for roads), and Track Type (for railroads). For faces, attributes at the primitive level include Internal Point (of the face), Permanent Face Identifier, and Land/Water Flag. Since Oracle's Topology Data Model manages the creation of edges and faces, these topological primitives are not readily modified by the user, in the event that it is necessary to store attributes at the primitive level. For the redesigned MAF/TIGER implementation, the Geography Division has utilized Oracle's Feature Management API to create auxiliary edge, node, and face "feature" tables that mirror the topology primitive tables. The staff also is developing software to manage the maintenance of the attributes on these mirrored node, edge, and face tables. This maintenance software is invoked whenever Oracle's Topology Data Model initiates updates to the node, edge, and face tables, thereby keeping the mirrored tables consistent with the base topology and assuring that the attributes on these mirrored tables are maintained in real time. Business rules manage the attribution on these mirrored tables, including assignment and tracking of permanent ID's (e.g., TIGER/Line ID) whenever edges are split or merged.

Hierarchical Features

Areal features are represented as groups of faces and the faces are, as stated previously, represented by their bounding edges. However, areal features are often hierarchical in nature. For example, states are comprised of counties, and counties are comprised of county subdivisions. In order to improve efficiency and performance for these hierarchical features, Oracle's Topology Data Model allows geographic areas to be defined in terms of other geographic areas, rather than defining them directly in terms of faces (Figure 4). For these hierarchically defined features, the same API is used to query and manipulate them that are used for the non-hierarchical features. The fact that they are stored differently is transparent to the developer.
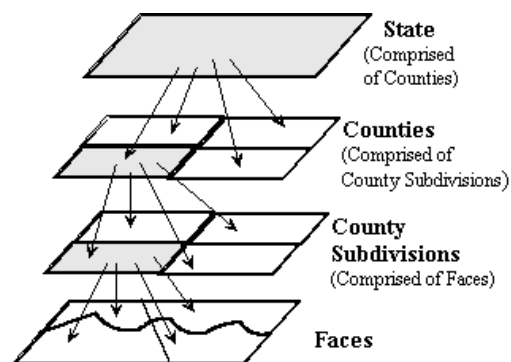

Figure 4. Geographic Hierarchies

Product Creation Database

The MAF/TIGER database, also known as the Transaction Database, will be optimized for batch and interactive updates, with a minimum amount of redundant data. However, such a database is not necessarily optimal for data extraction, mapping, and creation of other geographic products and services. For example, in the Transaction Database, consider a geographic area that is stored as a collection of faces. In order to determine the "geometry" of it's boundary, the application software must:

- Identify all the faces associated with the areal feature.
- Retrieve the edges associated with those faces (ignoring the "internal" ones).
- Order the external edges around the boundary.
- Retrieve the coordinates for the external edges.
- Return the ordered string of coordinates for the entire boundary.

From the developer's perspective, this is simple because Oracle's Topology Data Model API provides a single call to do it. However, the processing time to determine boundaries for large or complex geographic areas can be significant. It is inefficient if every application that needs this boundary has to re-calculate it, especially if the boundary has not changed. That is essentially how the current system works and it can be problematic, especially for mapping applications (Trainor 2003). For this reason, a Product Creation Database will be generated that contains the information from the MAF/TIGER Transaction Database, plus additional calculated information, including explicitly stored boundaries for geographic areas, as well as explicit geometric representations of linear features such as roads. The Product Creation Database will be replicated or created from the Transaction Database on a regular basis. The exact nature and format of the Product Creation Database has not been determined and it could depend on the selection of COTS application tools for mapping and other applications.

TIGER® is a registered trademark of the U.S. Census Bureau.

REFERENCES

Batty, P., 2004, Future Trends & the Spatial Industry, Part One, http://www.geospatial-online.com/geospatialsolutions/ article/articleDetail.jsp?id=101548.

Baumgart, B., 1975, Winged-Edge Polyhedron Representation for Computer Vision. *National Computer Conference, Stanford University*, Stanford, California. http://www.baumgart.org/winged-edge/winged-edge.html.

Boudriault, G., 1987, Topology in the TIGER File. *Proceedings of the 8th International Symposium on Computer Assisted Cartography (Auto-Carto 8)*, Baltimore, MD.

Corbett, J., 1979, Topological Principles in Cartography. *Technical Paper No. 48*, U.S. Bureau of the Census.

Egenhofer, M., A. Frank, and J. Jackson, 1989, A Topological Data Model for Spatial Databases. *Design and Implementation of Large Spatial, Lecture Notes in Computer Science*, 409, 271-286.

Gonzales, M., 2000, Seeking Spatial Intelligence, 3(2), http://www.intelligententerprise.com/000120/feat1.jhtml.

Hoel, E., S. Menon, and S. Morehouse, 2003, Building Robust Topologies. *In Advances in Spatial and Temporal Databases. Proceedings of the 8th International Symposium on Spatial and Temporal Databases. SSTD 2003*. Santorini Island, Greece, Springer-Verlag Lecture Notes in Computer Science 2750.

Kainz, W., 2004, Geographic Information Science, Version 2.0, http://www.geografie.webzdarma.cz/GIS-skriptum.pdf.

LaMacchia, R., 1990, The TIGER System. *1990 Exemplary Systems in Government Awards Competition, Urban and Regional Information Systems Association*, Edmonton, Alberta, Canada.

Lessware, S., 2004, Comparing Topology Management Techniques, http://www.geoplace.com/gw/2004/0401/0401tch.asp.

Marx, R., 1986, The TIGER System: Automating the Geographic Structure of the United States Census. *Government Publications Review*, 13: 181-201.

Open GIS Consortium, Inc., 1999, OpenGIS Simple Features Specification for SQL Revision 1.1.

Oracle Corporation, 2004, Oracle® Spatial Option and Oracle® Locator Datasheet - Location Features in Oracle Database 10g, http://www.oracle.com/technology/products/spatial/htdocs/ data_sheet_9i/10g_spatial_locator_ds.html.

Ramage, S. and P. Woodsford, 2002, The Benefits of Topology in the Database, http://spatialnews.geocomm.com/features/laserscan2/.

Saalfeld, A., 1993, Conflation: Automated Map Compilation. *Center for Automation Research, CAR-TR-670, (CS-TR-3066)*, University of Maryland, College Park.

Schneider, M., 2002, Spatial Data Types: Conceptual Foundation for the Design and Implementation of Spatial Database Systems and GIS, http://www.cise.ufl.edu/~mschneid/ Research/Tutorials/TutorialSDT.pdf.

Trainor, T., 2003, U.S. Census Bureau Geographic Support: A Response to Changing Technology and Improved Data. *Cartography and Geographic Information Science*, 30(2), 217-223.

Van Oosterom, P., J. Stoter, W. Quak, and S. Zlatanova, 2002, The Balance Between Topology and Geometry. *Symposium on Geospatial Theory, Processing and Applications*, Ottawa.

Weinberger, J., 2002, The Spatial RDBMS in the Enterprise, http://www.directionsmag.com/ article.asp?article_id=259, accessed 11 Oct. 2004.

White, M., 1984, Technical Requirements and Standards for a Multipurpose Geographic Data System. *The American Cartographer*, 11(1), 15-26.