Turning Transportation Planning Data Into Effective Web Sites

Giovanni Flammia and Andres Rabinowicz
giovanni@caliper.com *andres@caliper.com*

Caliper Corporation
1172 Beacon St, Newton, MA 02461
www.caliper.com

Abstract

The Web offers a unique opportunity for State Departments of Transportation and planning agencies to disseminate to large audience interactive results of transportation studies. This paper provides a technical backgrounder about the differences between GIS desktop applications and web mapping applications in the context of transportation planning. In a web mapping application the application server performs all of the computations and manages multiple user sessions. This technical constraint determines what functions are available and how they can be implemented effectively to display interactive maps and reports.

Introduction

In this paper we compare and contrast the functions of desktop geographical information systems (GIS) applications vs. web mapping applications in the specific context of transportation studies, and we discuss some of the benefits as well as the technical issues that arise when transferring features from a desktop application to a web application. We will use case studies about linear referencing, routing and demographic analysis to illustrate the points we will make.

Engineers use GIS and transportation planning applications to analyze road network data, solve routing problems and predict changes in travel patterns and the utilization of the transportation system in response to changes in demographics and transportation supply (see [1] for a review of some desktop GIS applications used for linear referencing and dynamic segmentation).

Usually, the functions offered by desktop applications are geared towards expert users. The applications allow engineers to create and edit geographic databases (e.g., route systems), and join them to tabular information such as pavement conditions, traffic accident reports along mileposts, and demographic tables. While desktop applications enable engineers to perform complex geographic editing functions and run interactive queries by location, they require technical knowledge and skills in order to be used effectively.

Until the advent of the World Wide Web, the most appropriate way of disseminating to a large audience the information gathered by transportation studies has been via static reports, tables and map images created with the GIS desktop applications. The use of interactive reports and interactive maps that allow users to perform queries by location has been limited to the users that have access to the GIS desktop applications and to the underlying data.

In the last few years, we have witnessed technical advances in web-based interactive mapping. The web sites listed in [2], [3] and [4] provide many significant examples of accessing and querying geographic data via a simple web browser interface connected to GIS application servers. The Web offers a unique opportunity for State Departments of Transportation and planning agencies to disseminate to large audience interactive results of transportation studies. Unlike a static report or map image, an interactive map allows anybody connected to the Internet to query the result of a study by location and perform drill down analysis, such as getting traffic accident reports for specific segments of a route system.

A web mapping application is a very effective way to visualize large quantities of data sets for a large audience. In many cases, it is desirable to disseminate the data very fast, as soon as it is available, without requiring users to download the data to their desktop and to use specialized software for viewing the data. For example, the Census 2000 has released the County-To-County Worker Flow Files [6]. These data contain the number of workers 16 years old and over who commute to work. The counts give the number of workers traveling within or between counties in the 50 states and the District of Columbia, plus travel to workplaces outside the U.S., such as Puerto Rico or a foreign country.

To quickly display and query this large data set (127,798 non-zero flows between pairs of 3,294 counties), we developed the web mapping site www.caliper.com/countytocounty. The map on that web site displays desire lines and tables for any selected county. To display the flows to and from any county in the United States, the user can simply click on a county area. The flows are displayed on the map using a scaled symbol theme (see Figure 1), and in a table below the map (see figure 2). To reduce clutter, only flows of 50 workers or more in or out of the county are shown on the interactive map.

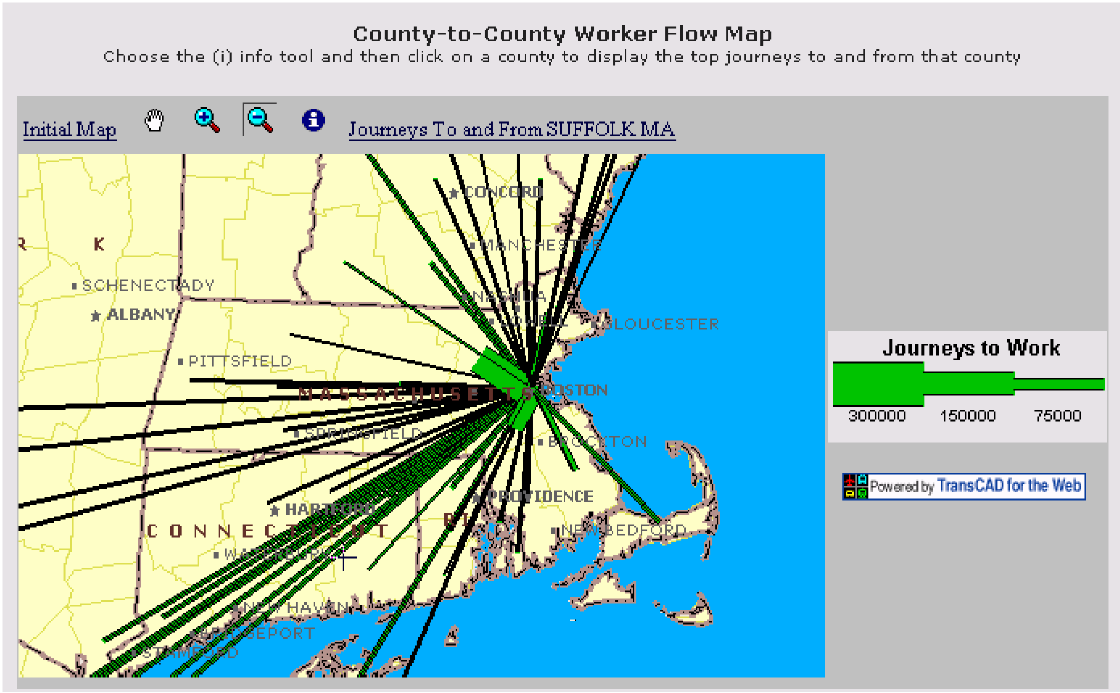

Figure 1: In the County-to-County worker flow web site, the user can choose the info (i) tool and then click anywhere on the map to display the home-to-work travels to and from any county in the United States.

Journeys To and From SUFFOLK MA (Threshold = 50)

County A	County B	Journeys from A to B	Journeys from B to A
ANDROSCOGGIN ME	SUFFOLK MA	86	18
BARNSTABLE MA	SUFFOLK MA	2,767	237
BELKNAP NH	SUFFOLK MA	156	--
BERKSHIRE MA	SUFFOLK MA	97	46
BRISTOL MA	SUFFOLK MA	11,103	1,409
BRISTOL RI	SUFFOLK MA	216	35
CARROLL NH	SUFFOLK MA	126	--
CHESHIRE NH	SUFFOLK MA	85	--
CHITTENDEN VT	SUFFOLK MA	76	28
COOK IL	SUFFOLK MA	110	186

Figure 2: Below the map, the worker flow data is also presented as a table.

Desktop Applications vs. Web Applications

From a software architecture perspective there are important technical differences between a desktop GIS application and a web mapping application. Knowledge about these differences will enable us to understand and overcome the technical challenges in designing effective web applications and decide which functions are appropriate to transfer from a desktop GIS application to a web mapping application.

In contrast, in a typical web application session many concurrent users interact with one application server via a sequence of web pages and map images that are dynamically generated. Users connect to the application server via a web browser such as Internet Explorer, and there is no GIS application or geographic data installed on each user personal computer or local area network. In the rest of the paper, we will limit our discussion only this "thin client" software architecture, where most if not all of the computations are executed on the application server (see figure 3).

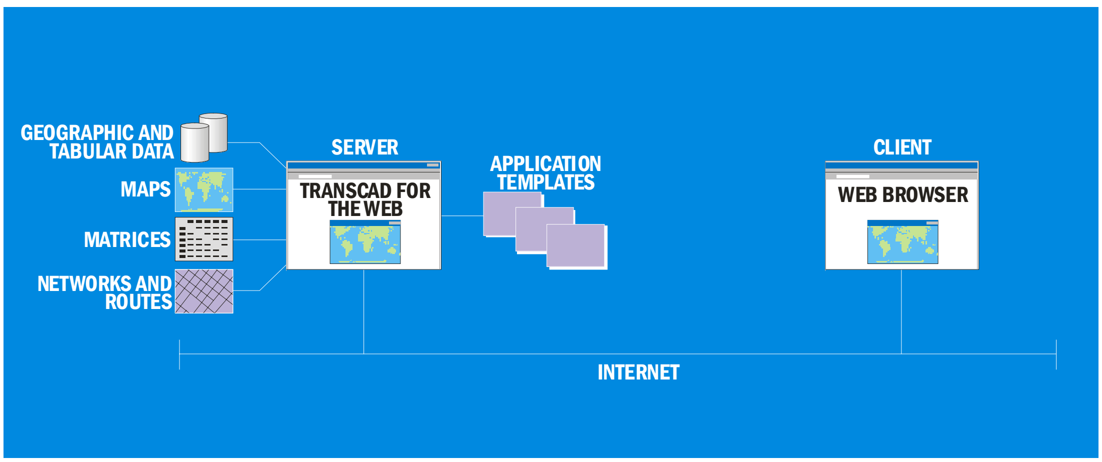

FIGURE 3: In typical web application architecture, all of the computations are performed by the GIS application server (to the left of the figure). Many users connect to the application server via the Internet and a thin client – e.g., Internet Explorer (to the right of the figure). The web application server sends to the browser all the information needed to display a map, including a map image file and one or more HTML tables and input forms.

The web application user interface provides a limited number of tools when compared to the desktop application. This is due to two constraints. Firstly, the target audience for the web application is as large as possible, and includes non-expert users. The web application user interface provides only the essential display and reporting tools that can be used with minimal training or knowledge of the underlying route system data structures employed by the software. Secondly, since the application server executes most of the computations the web application implements the set of functions of the desktop application that can be executed efficiently in a multi-user environment.

When a user clicks on the map or fills an input form in the web page, the application server performs the appropriate selection in the database stored in the server local area network and displays the results as HTML tables, images and maps. The web page displayed in the browser stores only the map image and the user session parameters that are needed by the application server to restore the user session when the next page is requested.

Managing Multiple User Sessions in a Web Application

In general, there are three approaches in storing the user session information in the web application server. The approaches can be combined and are presented here in an increasing level of technical complexity.

Approach 1 – Storing Session Information In a Web Page

The first and simplest approach is to store all of the information needed to create the map and the user selection in the web page. This is an appropriate solution if the information needed to restore the map is limited (e.g., map extent, origin, destination, and route name), and the amount of time it takes to restore the user selection in the memory of the application server is negligible. In this case, the page control flow for every page includes restoring the user selection and map extent, drawing the map or the chart, and releasing the user selection. This approach has the advantage of being scalable to an unlimited number of users. The application server does not maintain a session in memory for each user connected to the web application, and a web page can be served by any one of a number of servers that stores a copy of the web application. If the number of concurrent users increases, the increase in page requests can be handled by adding web application servers with a copy of the web application.

Approach 2 – Storing Session Information in Persistent Objects

The second approach is to compute some of the information needed to draw a map once and then store it a persistent fashion in the application server file system or as an object in the application server database system. This approach is being used by our direction-giving web applications. When a user asks for directions from an origin to a destination, the directions are computed once by the application server and then stored as an XML file in the server file system **(*figure 4*).** The web page stores information about the origin, destination and a unique identifier for the directions stored in XML format. When the user pans or zooms to a specific segment of the directions, the application server restores the directions by reading them from the XML file. In this second approach the application server does not maintain a specific session in memory for each user connected to the web application, and all session-specific information (e.g. directions) are stored in persistent objects or files by the application server. This approach scales well to multiple users because persistent objects (e.g. XML files or database objects) can be shared across a number of application servers; however it is appropriate only if the time it takes to store and reload objects from the local file system or database system is negligible.

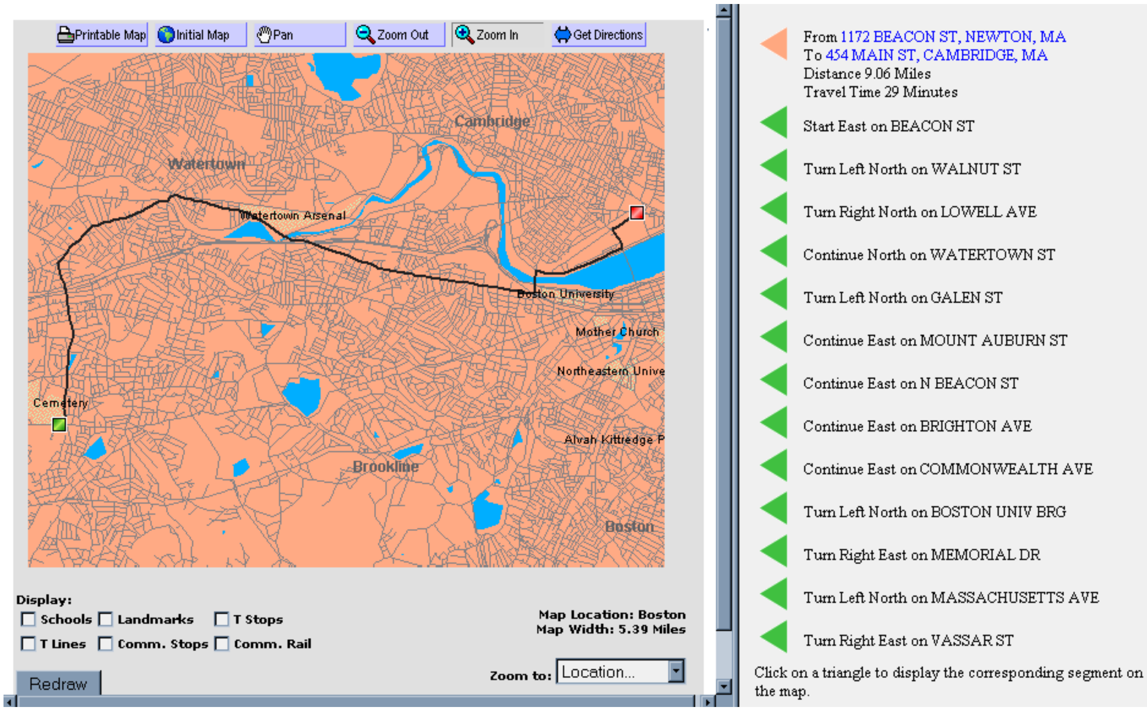

Figure 4: In the direction giving web application, the user can select an origin address and a destination address. The application server translates the input addresses to geographic locations and computes the directions a series of line segments in the road network. The directions are stored as an XML file and displayed as an HTML table to the right of the map image. The directions are restored from the XML file into the server memory when the user zooms or pans on the map image or zoom in to a specific segment by clicking any one of the triangles.

Approach 3 – Storing Session Information in Memory

The third approach is to maintain a pool of user sessions in the application server memory space. Using this approach, the web page stores the individual user unique session identification number, and the application server stores all the information needed to restore the map in memory. This approach is taken in one of our web applications to enable the user to add or remove ZIP Codes from a potentially very large selection set and then display aggregate information about the selection set specific to each user – e.g., total population in the selected area ***(figure 5).*** In this approach, the application server maintains a pool of selection sets with time stamps. The selection sets can be stored in the server memory for a limited amount of time. As soon as a user clears the selection set or when a time-out occurs the selection set can be used by another user. This is appropriate when it would take to much time to store and retrieve a potentially large selection set from a file or external database for every page, but it requires that a specific server is assigned to a user selection set until the user discard the selection set or a time-out occurs.

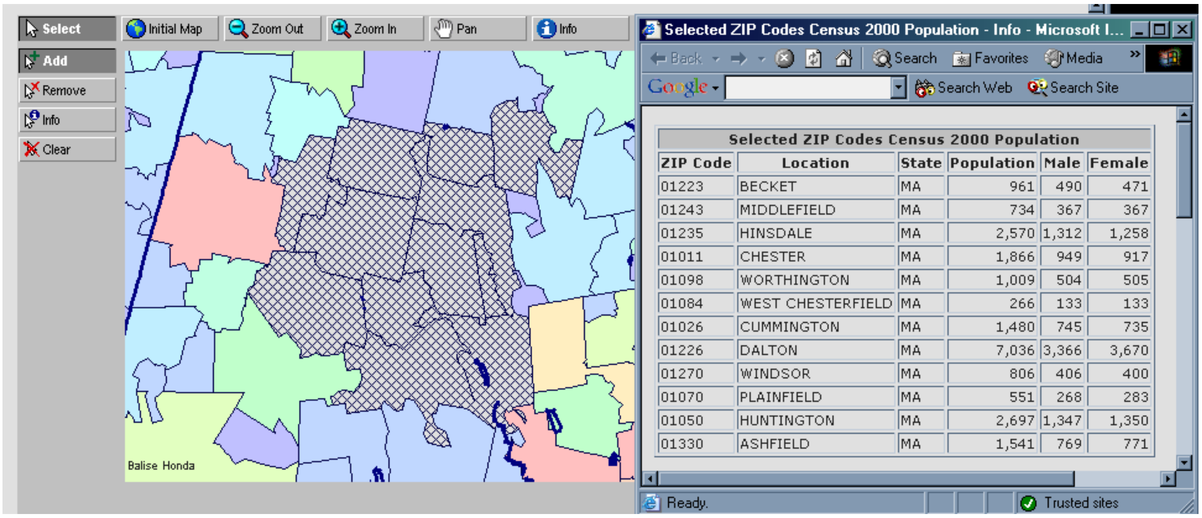

Figure 5: In this web application, the user can use the selection tool in the upper left corner to add or delete ZIP Codes from the current selection. The application server maintains a selection set of ZIP Codes for each specific user. When the user clicks on the Info button in the selecion toolbox, the application server collects the demographic information (Population, Male, Female) for the selection set and displays it in a pop-up window.

Editing and Personalizing Routes Online

The computational power, memory and disk space of mapping servers increases substantially every year. As a consequence, it is now practical to store and edit geographic data for hundreds of individual users on one mapping server. Authenticated users can log in to a server and set their own map display setting preferences, as well as review and edit their own personal geographic data.

The web mapping software must be carefully designed to manage multiple user data sets at the same time. In a load-balanced distributed architecture, individual user preferences and data will be stored in their own protected folders, and the folders may be physically stored on a separate machine from the web mapping servers, to increase performance and security.

The geographic data collected from each user can be aggregated and analyzed by transportation planning software, simplifying substantially the process of collecting data for transportation models. For example, travel surveys are necessary for building accurate transportation demand models. Typically, home-to-work travel data are collected on paper forms and booklets, the data is transferred to databases by hand, and eventually the database columns are associated to attributes on road network geographic layers that are used for estimating the models. This is a very time consuming process that provides limited data sets and that is prone to errors and omissions, especially when travel demand models need to be specified at the street address level.

An online travel survey service is a cost-effective alternative to paper surveys (see figure 6). Users can log in to a web site and fill a travel survey in stages. Typically, a user records their demographic profile and their home-to-work travels by activity. For each activity, the user enters all of the data that she would otherwise enter in a paper survey. Each activity correspond to a location (e.g., a street address), a purpose (e.g., at home, at work, school, shopping, etc) and a time of day.

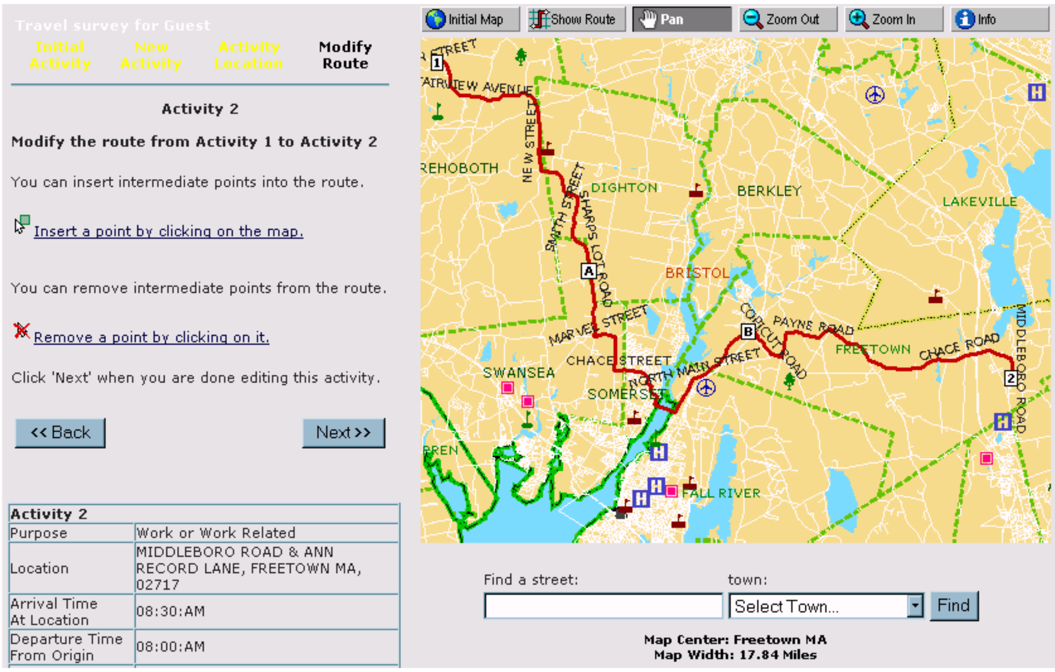

Figure 6: Unlike a paper travel survey, the on-line travel survey includes an interactive map. The map is used to specify the precise street address location of each activity, and to edit the route from one activity to the other. The screenshot displays how to edit a route between two activities. To the left, the user can choose an action (e.g., adding a route stop, modifying a route). If the user chooses to modify the route between two activities, she can either insert or remove intermediate points by clicking anywhere on the map to the right of the figure. The mapping server responds by redrawing the route to correspond to the user's choices.

Unlike a paper survey, users can specify the exact geographic location for each activity using an interactive map. The interactive map is also used to edit the route taken from one activity to the other. Initially, displays the route between activities based on the shortest path between them. While filling the survey, the user is always able to edit the route by inserting and deleting intermediate points, thus recording the actual route that was taken, rather than an ideal route.

An online travel survey offers two main advantages over paper data collection. Firstly, it is possible to collect data from a large number of users in a very short time and at a fraction of the cost of collecting data on paper. Secondly, since the interactive map includes the road network geographic layer that will be used to estimate the travel demand models, the data collected is immediately associated to the planning network nodes and segments. Thus data entry inaccuracies are minimized.

The ability of editing and personalizing routes for each individual user is a function that is not limited to collecting travel survey data. It is also a key feature that can be used by route delivery and fleet management systems, public transport route planning systems, and carpooling systems.

Extending the Functionality of Web Applications

Because the application server must maintain or restore each user specific session data structure with one of the three approaches described above, it is clearly beneficial that each page is processed as fast as possible. As a consequence, not all of the functions of the desktop application may be appropriate for being implemented in the web application. For example, we would consider appropriate to implement a function that always take a limited amount of time, between 1 and 5 seconds. If a reporting or display function takes more than 5 seconds, it might be better to compute the report offline, and deliver it to the user via email.

However, at least three considerations lead us to believe, as do many others, that the number of functions that can be implemented in a web application will increases significantly in the next few years. Firstly, the computational power of the hardware increases and the cost decreases substantially every year. The performance of a GIS application depends on many features including CPU clock speed, memory size, data bus size, and file input/output throughput. All of these features tend to improve substantially from one year to the next. Secondly, the generic user interface technology that can be embedded in the web browsers allows developing very rich graphical interfaces that go well beyond raster map images and HTML tables. Soon, users will be able to manipulate directly in the web browser vector images that overlay points, lines and area layers via the standard scalable vector graphics language (SVG) [5]. Thirdly, the number of people connected to the Internet with high-speed connections increases steadily every year, making it viable to transfer high-resolution images with printer-ready quality and large amount of geometric vector data between the web browser and the application server.

In summary, a web mapping application can implement a number of functions in order to present results of transportation studies in a very graphically rich and interactive presentation. In this paper we tried to illustrate how the web application control flow is substantially different from the control flow in a desktop application. As a consequence, the functions of a web mapping application must be designed carefully, sometimes from scratch, in order to support multiple user sessions in the most efficient manner.

Bios

Giovanni Flammia, PhD is Director of Web Development at Caliper Corporation. He is the lead software architect of TransCAD for The Web and Maptitude for the Web. Dr. Flammia holds a PhD in Computer Science from MIT. He can be reached by email at: *giovanni@caliper.com*

Andres Rabinowicz, PhD is Vice President of Caliper Corporation. He is an expert in transportation and GIS. He directs TransCAD product and application development. Dr Rabinowicz holds a PhD from TECHNION. He can be reached by email at: *andres@caliper.com*

References

[1] Jane E. Goodman "Maps in the Fast Lane – Linear Referencing and Dynamic Segmentation" *GIS Vision Magazine*, Dec 2001.

http://www.gisvisionmag.com/vision_printer_friendly.php?article=200112/review.html

[2] "GIS on the Internet: Web Mapping powers diverse applications" *Geo World Magazine*, September 2001. http://www.geoplace.com/gw/2001/0901/0901wb.asp

[3] W. Frederick Limp "Web Mapping 2002: Commercial Improvements, New Web Services and Interoperability Initiatives Make for Interesting Times" *GEO World Magazine,* March 2002. http://www.geoplace.com/gw/2002/0203/0203web.asp

[4] Rick Wayne "Cartographic Advances: Location, Location, Location" *Software Development Magazine*, March 2002, pp. 39-43. http://www.sdmagazine.com/documents/s=4077/sdm0203a/0203a.htm

[5] Scalable Vector Graphics Standard Specifications and Tutorials http://www.w3.org/Graphics/SVG/Overview.htm8

[6] U.S. Census 2000 Conty-to-County Worker Flow Files http://www.census.gov/population/www/cen2000/commuting.html