FUSING LIDAR DATA, PHOTOGRAPHS, AND OTHER DATA USING 2D AND 3D VISUALIZATION TECHNIQUES

Robert J. McGaughey, Research Forester
Ward W. Carson, Research Engineer
USDA Forest Service
Pacific Northwest Research Station
University of Washington
PO Box 352100
Seattle, WA 98195-2100
bmcgaughey@fs.fed.us
wcarson@fs.fed.us

ABSTRACT

The process of understanding, analyzing and communicating the large spatial datasets commonly output from LIDAR systems is difficult. Such data typically consist of several million 3D sample points distributed over several thousand hectares. Visualization techniques can help users better understand the raw data, formulate suitable analysis methods, and present results. Unfortunately, few visualization environments can represent entire LIDAR datasets in a way that promotes easy interaction and clear understanding. Even immersive techniques such as CAVE- and headset-based systems tend to overwhelm the viewer with the shear volume of data presented. This paper describes a prototype visualization framework that fuses LIDAR data, photographs, and GIS data using 2D displays and a 3D stereoscopic visualization environment. We have used the system for projects where the terrain surface is of interest and in the derivation of vegetation characteristics from LIDAR data. The primary interface consists of a 2D display of a geo-referenced orthophoto, various point, line, and polygon overlays, and contour lines derived from a digital elevation model. Users interact with the 2D display to specify a subset of LIDAR data for stereoscopic viewing. The 3D viewing tool displays the terrain represented as a shaded surface, the specified subset of the raw LIDAR data points, and a portion of the orthophoto image mapped onto a horizontal plane that can be positioned anywhere along the vertical axis of the scene. The system provides methods for organizing and indexing large datasets making it useful for a wide variety of LIDAR projects. This paper presents an overview of the data interface and visualization system along with several examples of its use.

INTRODUCTION

The Silviculture and Forest Models Team of the USDA Forest Service, Pacific Northwest Research Station based in Olympia, Washington and the Precision Forestry Cooperative located at the University of Washington in Seattle, Washington are evaluating the accuracy of LIDAR data and developing tools to describe vegetation characteristics such as tree density, tree size, canopy volume, and crown bulk density using LIDAR data. Efforts to date include the comparison of LIDAR-derived terrain models with models produced by the USGS in both 10- and 30-meter resolutions, accuracy assessments of LIDAR-derived terrain models using surveyed control points (Reutebuch and others, 2003), development of algorithms to locate individual trees and characterize their size and overall crown shape using raw LIDAR returns (Andersen and others, 2001), comparison of LIDAR-derived terrain and canopy models with photogrammetric measurements collected from 1:3000 color photography, and the comparison of terrain and canopy surface models derived from LIDAR data with models developed using interferometric radar data. These efforts have all included some type of data visualization to help communicate the wealth of information contained in the LIDAR datasets. Visualizations are commonly hill-shaded or perspective renderings of surface models or renderings of large point clouds using various color schemes to enhance particular attributes of the return data.

Project scientists have been frustrated by the difficulty inherent in manipulating the large datasets common in LIDAR analysis. Commercial geographic information systems (GIS) are often incapable of processing files of

several million points or at least suffer severe performance degradation when drawing or searching through such files. Programming products such as the Interactive Data Language[*] (IDL) from Research Systems, Inc. provide a more general solution for analysis and visualization of LIDAR data but generally require task-specific programming to produce analysis procedures and visualizations.

After working on several LIDAR analysis project, we have identified several capabilities that we feel are important for a visualization system. We think such a system should:

- use existing data files stored on networked file servers and not rely on local copies on individual computers,
- display a variety of data types and let users display or hide individual data layers,
- allow users to interactively select portions of large LIDAR datasets for detailed viewing and examination,
- clip all data layers to a specified sample area and make them available for detailed viewing,
- operate quickly enough to encourage use of the system (subsets need to be extracted in a few seconds rather than a few minutes),
- provide support for stereoscopic display hardware when viewing data subsets, and
- operate on readily available computer hardware and operating systems.

This paper describes the prototype system we are developing to test the usefulness of these capabilities and describes its application on one of our project areas.

BACKGROUND

Most researchers involved in the analysis of LIDAR data rely on some type of data visualization to help understand and explore specific features in the data. Visualization techniques include fully immersive display systems using specialized hardware (Eggleston and others, 2000), 2D and 3D plots of LIDAR data generated using IDL (Andersen and others, 2002), and custom developed applications. For many projects, visualizations show entire datasets or large subsets of the data. In our visualization system, we wanted to view relatively small subsets of large LIDAR datasets rather than entire datasets. A simple visualization tool that could display collections of 3D point data seemed to be the best choice.

Hoping to find a suitable visualization tool, we searched the world wide web for programs that could be used directly or modified to provide the needed display capabilities. Our search resulted in two promising choices: the Stereoscopic Field Analyzer (SFA) and a simple 3D scatter plot program. SFA allows the visualization of both regular and irregular grids of volumetric data (Ebert and others, 1996). The system provides a minimally immersive interactive visualization tool designed to facilitate the understanding of volumetric data and provide for the real-time manipulation of the data display. It uses a pair of 3D magnetic sensors to operate the user interface and manipulate the data volume. SFA required a Silicon Graphics computer and used the proprietary GL language.

A simple 3D scatter plot program was developed and documented[**] by William J. Heitler. It is a Microsoft Windows application that displays random XYZ data points and allows manipulation of the data volume using a mouse. Individual points can be assigned different colors and the point data are input from ASCII text files. Given the hardware requirements of SFA and difficulties associated with the GL programming language, we decided to start development using the simple scatter plot program.

In its original form, the scatter plot program could display subsets of LIDAR data using a pre-assigned color. However, the program did not provide many of the display features we desired, had some memory management problems, and did not provide stereoscopic display capabilities. We modified the program to provide stereoscopic and anaglyph capabilities, expanded its options for point markers and point coloring, added a virtual trackball control mode to make it easier to manipulate the data volume, provided support for shaded surface rendering and texture mapping images onto surfaces and arbitrary planes, and modified the data clipping capabilities. In the end,

---

[*] Use of trade or firm names in this publication is for reader information and does not imply endorsement by the U.S. Department of Agriculture of any product or service.
[**] Documentation is available at http://www.codeguru.com/opengl/scatter3d.html.

the new program, which we call LDATAVIEWER, uses very little of the original code. However, we are indebted to Mr. Heitler for posting his code and the accompanying article as we found his program an excellent prototype for developing the design of our LDATAVIEWER display program.

VISUALIZATION SYSTEM OVERVIEW

The visualization system consists of two programs, FUSION and LDATAVIEWER, and a set of utility functions to manipulate and reformat data files. All are implemented as Microsoft Windows applications coded in C++ using the Microsoft Foundation Classes. The primary interface, provided by FUSION, consists of a graphics display window and a control window. The FUSION display presents all project data using a 2D display typical of geographic information systems. It requires that all data be geo-referenced using the same projection system and units of measurement. LDATAVIEWER provides a 3D visualization environment, based on OpenGL, for the examination of data subsets. In operation, FUSION communicates with LDATAVIEWER by passing simple messages using the Windows message queue.

In FUSION, data layers are classified into four categories: images, raw data, points of interest, and surface models. Images can be any geo-referenced image but they are typically orthophotos. Raw data include LIDAR return data and simple XYZ point files. Points of interest (POI) can be any point, line, or polygon layer that provides useful visual information or sample point locations. Surface models are expected in a gridded format that represent either a ground surface or other surfaces of interest such as the top of a forest canopy. The current FUSION implementation limits the user to a single image and a single surface model, however, multiple raw data and POI layers can be specified.

The FUSION interface provides users with an easily understood display of all project data. Users can specify display attributes for all data and can toggle the display of the four data types. The entire raw data layer can be rendered, however, it generally requires excessive time and is often left as a hidden layer.

FUSION allows users to quickly and easily select and display subsets of large LIDAR data. Users specify the subset size, shape, and the rules used to assign colors to individual raw data points and then select sample locations in the graphical display. LDATAVIEWER presents the subset for the user to examine. Subsets can include not only raw data but also the related portion of the image and surface model for the area sampled. Subsets can be either circular, rectangular or square with a fixed or variable size. A corridor subset type allows extraction of rectangular subsets defined by the ends of a line segment and the width of the corridor. One can snap to a sample location at a POI to generate subsets centered on a particular location. Subsets can be made relative to a local surface model prior to display in LDATAVIEWER. This feature is especially useful when viewing data of forested regions in steep terrain as it is much easier to examine returns from vegetation after subtracting the ground elevation.

LDATAVIEWER presents raw data points as a variety of markers that range from single pixels, to simple geometric objects, to complex superquadric objects. Orthographic images are mapped onto a horizontal plane that can be positioned vertically within the cloud of raw data. Surface models are rendered as a shaded 3D surface. LDATAVIEWER operates in monoscopic, stereoscopic[*], and anaglyph display modes. To enhanced the 3D effect on monoscopic display systems, LDATAVIEWER provides a simple rotation feature that moves the data subset continuously through a simple pattern (usually circular). We have dubbed this technique "wiggle vision", and feel it provides a much better sense of the 3D spatial arrangement of points than does a static, fixed display.

Data Indexing

One of the primary goals in the development of our visualization system was to provide a fast way to browse large LIDAR datasets. Extracting subsets with purpose-written utility programs proved to be too slow. Typically, extractions in batch mode required several hours for a few hundred subsets. We wanted a system where the user could interactively extract and display subsets and easily change the coloring rules used to display the data. Obviously, reading entire datasets containing several million returns would not result in the desired performance.

---

[*] Stereoscopic display requires a graphics adapter that supports OpenGL quad-buffered stereo and display hardware such as active and passive systems available from StereoGraphics Corporation.

Our most commonly encountered format for LIDAR data has been ASCII, space-delimited text. Data describing each return typically include the pulse number, return number (if data were produced by a system capable of collecting multiple returns), X (often longitude), Y (often latitude), elevation, off-nadir angle, and the recorded intensity for the return. While such data files are easy to read and process, they are not suitable for fast, efficient access. For the prototype visualization system, we developed an indexing scheme to facilitate rapid examination of the data points located in specific areas within the data extents. In addition, we converted the ASCII text files into a binary format that preserves all of the original information contained in the ASCII files. The binary format requires less storage space (about 33% less) and provides a fixed length record for each LIDAR return. The current binary format is unique to this project but we expect to modify the classes that manipulate LIDAR files to use the newly adopted LAS format (supported by the American Society for Photogrammetry and Remote Sensing and described at http://www.lasformat.org) in a future version.

The indexing scheme is simple and can be applied to our binary files or the original ASCII text files. In fact, this indexing scheme could be used with data stored in a variety of formats given that the index is only used to jump to the start of a point record in the raw data file. The indexing procedure first scans a LIDAR data file to determine the extents of the data coverage. The area is then overlaid with a 256 by 256 grid. A new file, called the index, is created containing one record for each LIDAR return in the source file. The record contains the column and row for the cell containing the data point and an offset into the raw data file to the start of the point record. After completing the index, it is sorted using the column and row values and a second file, called the first point file, is created listing the offset into the index file to the start of the first index entry for each cell in the index grid. Using the index and first point file, we can quickly read all data points contained in a specific cell in the index grid.

EXAMPLE PROJECT

Our example focuses on a small project area located in western Washington State. Overall goals for the project include the evaluation of LIDAR accuracy for ground surface measurements, determination of vegetation characteristics using LIDAR data, and a general exploration of the utility of LIDAR data for detecting terrain and vegetation features. The area is part of Capitol State Forest and is managed by the Washington State Department of Natural Resources. The site is mountainous with elevations varying from 150-400 m and ground slopes from 0-45 degrees. Forest canopy within the study area is primarily coniferous and highly variable. It includes recent clearcuts, areas thinned to varying tree densities, and forest plantations ranging in age from recently planted to 70-year-old mature forests. As part of a forest management study (Curtis and others, 1996), the canopy of the 70-year-old forest stand was partially harvested in 1998, resulting in four different residual canopy density classes. In clearcut areas, the number of residual trees per hectare (TPH) is zero; in the heavily thinned areas, approximately 40 TPH remain; in lightly thinned areas, 175 TPH remain; in the uncut area, 280 TPH remain. The dominant tree height in the harvest area was approximately 50 m. Figure 1 is a screenshot of FUSION showing a 1999 orthophotograph of the study area (outlined) taken 4 months prior to the LIDAR flight, treatment unit boundaries, ground control points, and contour lines generated from a digital terrain model.

Several types of data are presented in this example:
- digital orthophoto,
- LIDAR data,
- digital terrain model,
- point overlay depicting ground control points, and
- polygon overlay delineating treatment area boundaries.

For this project, all horizontal data are referenced to the State Plane projection system, Washington south zone (4602), NAD 83, U.S. survey feet. Vertical data are referenced to NAVD88, U.S. survey feet.

Digital Orthophoto

The Washington Department of Natural Resources, Resource Mapping Section produced a digital orthophoto for the study area using a softcopy system (Socket Set). The source imagery was 1:12,000 color aerial photography. Orthorectification was accomplished using a canopy surface model developed using autocorrelation techniques. The final image used a 0.3m (1 foot) pixel.
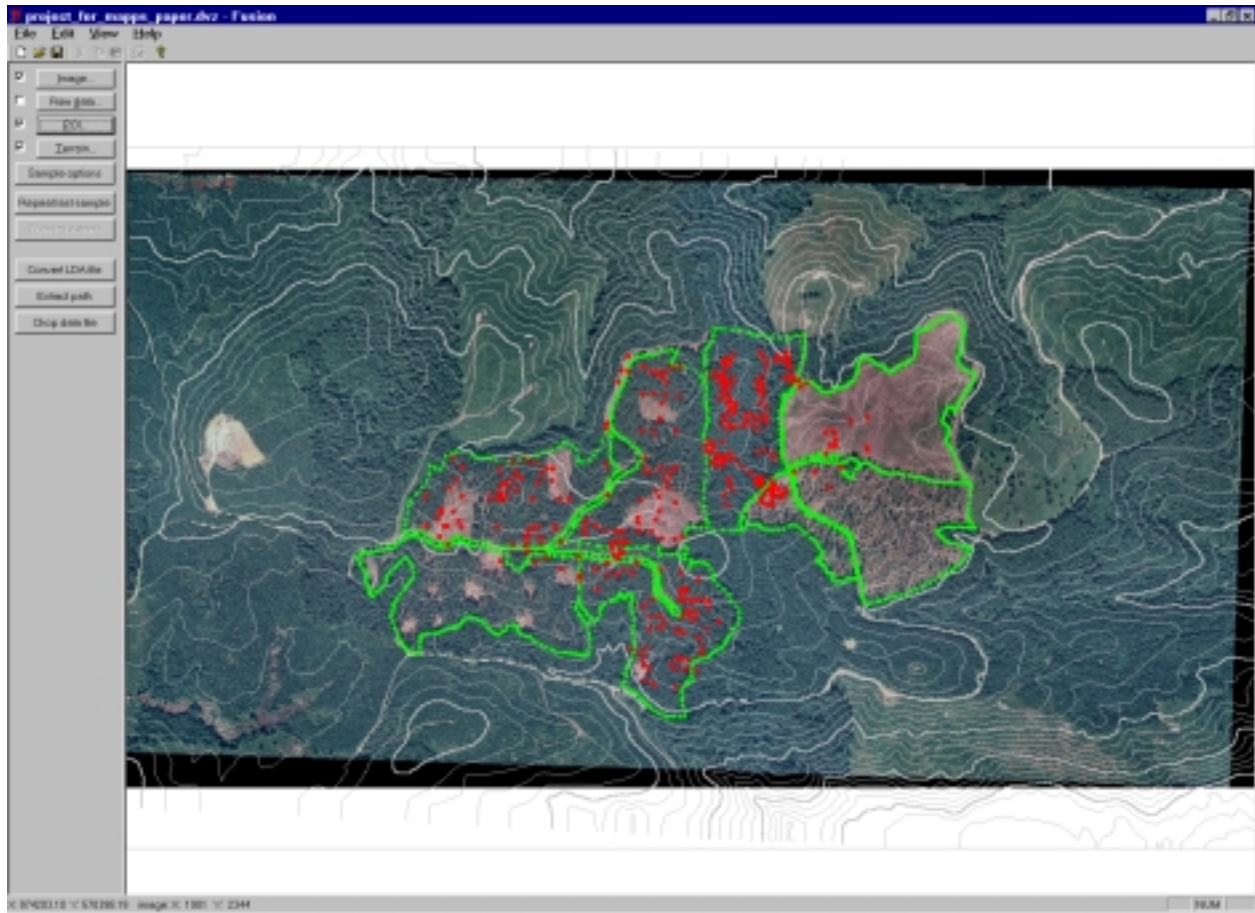
Figure 1. Screenshot of FUSION showing an orthophotograph of the Capitol Forest study site, 15-meter contours created from the digital terrain model, control points (red), and treatment area boundaries (green).

LIDAR Data

A small footprint, discrete return LIDAR system was used to map 5 km^2 of the study area in the spring of 1999. The contractor used a Saab TopEye LIDAR system mounted on a helicopter to collect data over the study site. Table 1 presents the flight parameters and instrument settings for the data acquisition. LIDAR data are organized into 15 separate files each containing 600,000 to 3,000,000 returns. There are a total of 37,077,317 returns in the project dataset. Each return includes the pulse number, return number for the pulse (up to four returns were recorded per pulse), X, Y, elevation, off-nadir angle and return intensity.

Table 1. Flight parameters and scanning system settings.

Flying height	200 m
Flying speed	25 m/sec
Scanning swath width	70 m
Forward tilt	8 degrees
Laser pulse density	4 pulses/m^2
Laser pulse rate	7,000 points/sec
Maximum returns per pulse	4
Footprint diameter	40 cm

Digital Terrain Model

For this study, the contractor sorted the data into "last returns"--approximately 6.5 million ground-point candidates in a rectangular, 5 square kilometer area clipped from the project area. They then used a proprietary routine (most probably based upon the removal of spikes) to further filter the data down to approximately 4 million points assumed to be on the ground. We used these filtered ground points to create a ground-surface digital terrain model (DTM) using a 1.524- by 1.524-meter (5 by 5 feet) grid resolution. The final filtering of last returns for the purpose of generating a "ground surface" is a commercially important and, thereby, inherently proprietary activity. Some discussions have been published in the open literature (Elmqvist, 2002; Haugerud and Harding, 2001), however, it is common for the consumer of LIDAR data to order the "ground-surface" coordinate data and accept it as a commercial product. In this case, we worked with the filtered last returns or "ground surface" coordinate dataset as delivered by the contractor.

Additional Data Layers

A Topcon ITS-1 total station surveying instrument was used to survey 347 checkpoint locations in the central portion of the area. Approximately 85 percent of the survey points were under forest canopy. The ground survey consisted of three closed traverses that originated at reference points established using survey-grade GPS. After adjustment of each traverse, the horizontal and vertical accuracy of ground points were approximately 15 cm and 3 cm, respectively. Reutebuch and others (2003) presents an evaluation of LIDAR-derived terrain model accuracy using these checkpoints and found an average LIDAR elevation error of 22 cm.

Treatment area boundaries were traversed using field GPS units. In areas with sparse vegetation, boundary points were occupied and measured directly using GPS. In areas of heavy vegetation, boundary points were measured using offsets from GPS stations in clearings using a Criterion laser rangefinder.

Example Visualizations

We used FUSION and LDATAVIEWER to produce several examples to illustrate some of our visualization system's capabilities. Unfortunately, it is difficult to demonstrate the interactive capabilities of the system and the ease with which the user can explore the project data in a printed document.

Figure 2 shows a vegetation profile created by extracting all returns within a 4- by 850-meter strip. The terrain model elevation was subtracted from each data point to better show the vegetation structure.

Figure 3 shows a 1-hectare portion of the terrain model rendered as a shaded surface. Figure 4 shows the same area with return data colored by intensity. The prominent feature in the surface model is a road. The area below the road is a recent clearcut covered by slash, shrubs, and small trees. The area above the road is a second growth, Douglas-fir stand.

Figure 5 shows a 1-hectare sample of return data. The terrain model elevation was subtracted from each data point and the markers colored according to their height above ground. Figure 6 shows the same area with a portion of the orthophoto texture mapped onto a flat plate positioned 3m above the lowest data point. The shadows in the orthophoto were cast from the trees evident in the LIDAR return data.


Figure 2. Vegetation profile showing all LIDAR data in a 4-meter wide strip colored by height above ground.

Figure 3. One-hectare patch of the LIDAR digital terrain model.

Figure 4. One-hectare patch of the LIDAR terrain model with LIDAR data super-imposed and colored by intensity.

Figure 5. One-hectare patch of LIDAR data colored by height relative to ground surface. The red markers are on or very close to the ground.


Figure 6. One-hectare patch of LIDAR data shown in Figure 5 with a portion of the orthophoto texture mapped onto a plate positioned 3m above the lowest data point.

CONCLUSIONS

Initial experiences with the FUSION and LDATAVIEWER prototypes have demonstrated the overall usefulness of a simple visual interface to access and examine large sets of LIDAR data. The ability to quickly extract subsets of return data and view the subsets interactively using a variety of coloring rules has led to both a better understanding of the ability of LIDAR to capture a high level of terrain and vegetation detail and a series of new questions regarding the exact tree features that are being measured by a LIDAR pulse.

Development of our visualization system is on going. We expect the system to change over the next year as we work with new data and identify new features that facilitate the understating and analysis of LIDAR data. One of our major goals is to add a new program that will display aerial photo pairs in stereo with the superimposition of LIDAR and other data on the stereo image. We hope that such a program will further enhance our ability to understand the specific tree features that are actually measured by a LIDAR pulse and to develop algorithms that can "see the trees" within the vast cloud of points.

REFERENCES

Andersen, Hans-Erik; Reutebuch, Stephen E.; Schreuder, Gerard F. (2002). Bayesian object recognition for the analysis of complex forest scenes in airborne laser scanner data. In: *International Archives of Photogrammetry and Remote Sensing*. Graz, Austria. Volume XXXIV. Part 3A. pp. 35-41.

Ebert, David S.; Shaw, Christopher D.; Zwa, Amen; Starr, Cindy. (1996). Two-handed interactive stereoscopic visualization. In: *Proceedings IEEE Visualization 96*, ACM Press, New York, pp. 205-210.

Eggleston, N.T.; M. Watson; D.L. Evans; R.J. Moorhead; J.W. McCombs. (2000). Visualization of airborne multiple-return LIDAR imagery from a forested landscape. In: *Proceedings of the 2nd International Conference on Geospatial Information in Agriculture and Forestry*, Lake Buena Vista, FL. January 10-12, 2000. ERIM International. Volume I. pp. 470-477.

Elmqvist, M. (2002). Ground surface estimation from airborne laser scanner data using active shape models. *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*. Volume XXXIV-3/W3. pp.114-118.

Haugerud, R.A.; D.J.Harding. (2001). Some algorithms for virtual deforestation (VDF) of LIDAR topographic survey data. *International Archives of Photogrammetry and Remote Sensing*, Volume XXXIV-3/W4. pp. 211-217.

Reutebuch, Stephen E.; McGaughey, Robert J.; Andersen, Hans-Erik; Carson, Ward W. (2003). Accuracy of a high resolution LIDAR terrain model under a conifer forest canopy. *Canadian Journal of Remote Sensing*. Volume 29, No. 5. (in press).

Curtis, R.O.; DeBell, D.S.; DeBell, Jeff; Shumway, John; Poch, Tom. (1996). Silvicultural options for harvesting young-growth production forests. Study plan on file with the Pacific Northwest Research Station, Westside Silviculture and Forest Models Team, 3625 93rd Ave SW, Olympia, WA 98512.