Extracting Points, Lines, Surfaces, Features & Models from Point Clouds

Geoff Jacobs

In an earlier, introductory article on high-definition surveying (3D laser scanning), it was pointed out that its single most important, distinguishing feature is its high point density (May 2004). Other key features—high-speed data capture, reflectorless measurement, 3D visualization, and colorful imagery—can be considered bonus by-products of the technology's focus on economically capturing a site or structure with high point density. The potential advantages of dense point data, such as creating more complete and accurate as-builts, have also been covered in earlier articles (May, July, and August 2004). High point density, however, raises some interesting new questions—and opportunities—regarding the task of extracting traditional deliverables.

What Do I Do with All Those Points?

Whether traditional surveying is done with a tape or chain, a total station, or GPS, the focus is on the collection of individual, significant points. In a day's work, hundreds or even thousands of points of interest are collected *intelligently*. A professional places a pole/prism on, or aims a reflectorless instrument at, *selected* spots to be measured by the instrument. Surveyors intelligently collect points (and sometimes add codes to identify associated features) at markers and monuments, corners, intersections, edges, selected features, etc., knowing that these selected points/features will form a basis for creating final deliverables.

In contrast, for the same site or structure, a high-definition or laser scan survey can literally blanket it with millions of closely spaced points. This article will describe various ways of extracting, from high-definition survey data, the desired points, lines, surfaces, symbols, models, and other elements needed for final deliverables. This article will also describe an increasingly popular new deliverable option—point clouds themselves—and how progressive clients today are effectively using this deliverable to better perform their engineering, planning, QA, analysis, archiving, and other tasks.

Loading Point Clouds for Processing

Data captured by laser scanners is stored on electronic media. Data is typically stored directly on a portable computer, along with software that operates the scanner and may also be used to view and process the data. Depending on the scanner, data may alternately be stored on a removable media/card that resides either on the scanner or in a handheld collector, but without 3D viewing and processing capabilities.

Some vendors provide both scanner hardware and point cloud processing software; others supply one or the other. Point cloud software is either file-based or utilizes a database structure, which can be beneficial for managing large data sets and for sharing scan data among multiple users. Most point cloud processing software can import

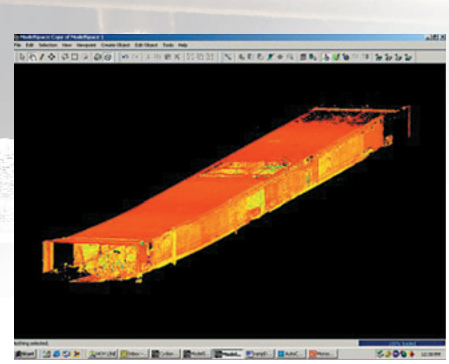

Multiple scans registered and geo-referenced. Image courtesy BSC/Cullinan, Inc.

standard file types, such as simple ASCII formats, but each software has an internal, native representation on which it operates most efficiently. Vendors that provide integrated scanning hardware and point cloud processing software typically make it such that no extra data conversion is needed to process the output from the scanner. That is, Vendor X's scanner outputs data in a format that is optimal (native) for Vendor X's software. When mixing and matching is done between point cloud software and scanning hardware from different vendors, an extra step of data import and data format conversion must be done. The conversion time depends on the size/format of the files and the horsepower of the processing PC.

Creating Registered, Geo-referenced Point Cloud Sets

The first step for creating deliverables from any high-definition survey is to register together and/or geo-reference point clouds captured from different scanner

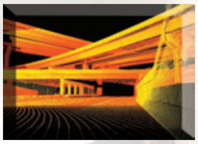

setups. Even for "single-scanner-position" projects, it may be important to tie scan data to site coordinates. Today, various ways are available to register and geo-reference point clouds. Some scanners can be directly geo-referenced over points, similar to the way a total station is used, or can utilize GPS. For high-accuracy work today, the common registration/geo-referencing methodology for scanning is analogous to that for photogrammetry: known points or targets within the scans are used for registration/geo-referencing. In addition, overlapping scans may also be aligned together. In many cases, both known points and overlapping scans are used to create a final "registered, geo-referenced" point cloud data set.

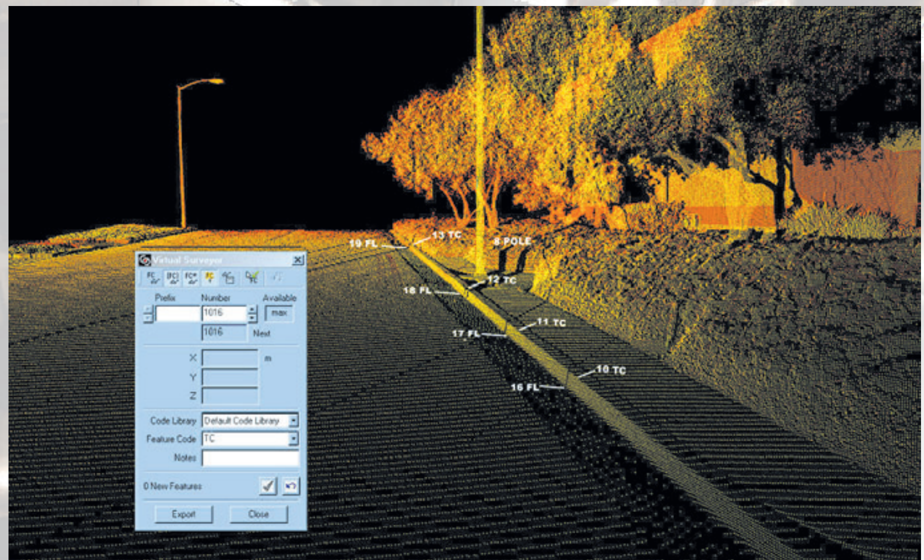

Virtual surveying: scan points and codes can be exported to mapping software.

Data Cleanup

A second step sometimes used on high-definition surveys is cleaning up the data, *i.e.*, removing unwanted data prior to creating final deliverables.

In high-definition surveys, unwanted data can come from various sources. For example, laser scanners can literally capture everything within range of the scanner, including people or vehicles passing through the scanner's line of sight. Likewise, not all static features collected (*e.g.*, vegetation, construction junk, tools, etc.) may be pertinent to the deliverable. Additionally, laser scanners can collect spurious data reflected off mirrored surfaces (including puddles of water). Yet another source is "edge effects," where a laser beam that partially reflects off the edge of an object may give a "false point" that appears by itself out in space. The edge effect phenomenon worsens with increasing spot size and is often more pronounced with phase-based scanners, which rely on integrating consecutive point samplings. Phase-based laser scanners are also very photon-sensitive and can record ambient light as random point noise within a scan. They are also subject to phase ambiguity resolution effects, which may require additional data clean-up (an explanation of this is considered beyond the scope of this article).

In practice, scan data from phase-based scanners is filtered in the office to remove most noise. Users report good results for this process, even though 25% or more of the data collected may be filtered out. For time-of-flight scanners, data may optionally be cleaned up before it is used, depending on required deliverables. I've spoken with experienced time-of-flight scanner users who used to clean up their data sets before processing them, but now don't bother with that step if the extraneous data does not negatively impact results or office efficiency.

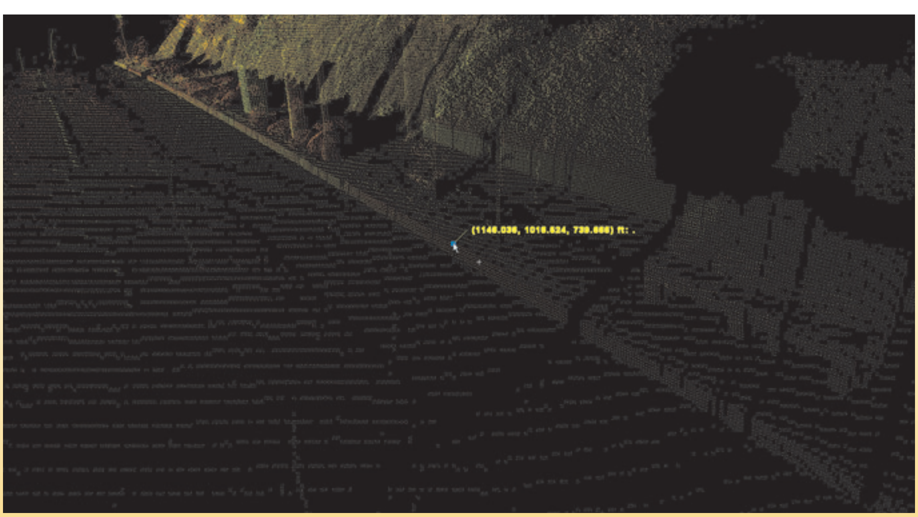

Users can click on any point within a point cloud for its X, Y, Z coordinates.

Extracting Single Points

Each point in a point cloud has unique X, Y, Z coordinates, so a user needs only to click on a point to observe its coordinates. This can be done in specialty point cloud software or in popular CAD applications such as AutoCAD or MicroStation (or in applications built on top of these platforms). For point extraction within CAD, there are two possible paths. One path is to import the point cloud or part of the point cloud into CAD. This path is efficient only with relatively small scan data sets (in the low tens of thousands of points). The other path is to use special point cloud "CAD-integration" software, much like a CAD plug-in, that allows large point clouds to be accessed from within CAD. Such software is commercially available for both AutoCAD and MicroStation and is being developed for other CAD applications.

As easy as it is to extract single points from point clouds, you can't be sure that a given scan point is exactly on a corner or an edge, or represents the

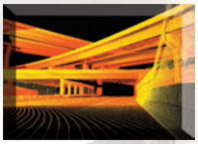

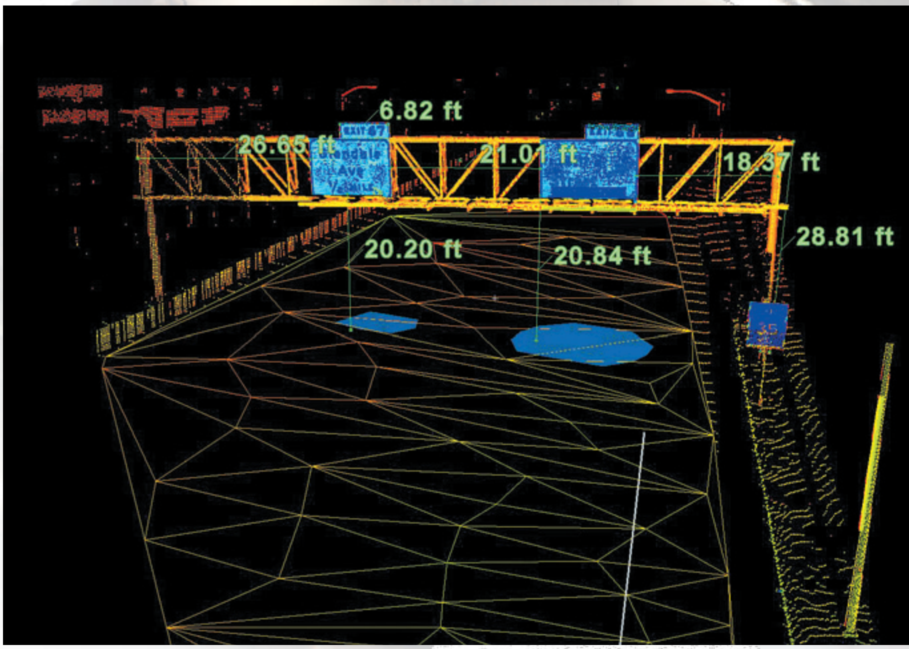

Distance measurements to points and surfaces

Roadway cross-sections at regular Intervals

lowest point of a curb, etc. Use of a single scan point in this way will depend on the accuracy required. If a user only needs point accuracy on the order of the scan spacing, say 2cm, then it may be acceptable to use single points from the scan data, depending on the actual single-point accuracy of the scanner. If higher accuracy were needed, then an alternative (again depending on the actual single point accuracy of the scanner) would be to execute a high-density scan of the area of interest (*e.g.*, with 1.5mm point-to-point spacing) and click on the point of specific interest.

If even higher point accuracy is needed, then it may be necessary to create a model of the area of interest and derive the point of interest from the model. This is exactly the process that is used to locate the center point of scan targets used for geo-referencing. A planar or spherical target is tightly scanned (*e.g.*, 1.5mm point-to-point spacing); fitting algorithms are then applied to fit the target, representing the center as a vertex. Target centers acquired this way can be derived to 1.5mm accuracy, depending on the type of scanner used. Most software allows target centers to be extracted automatically. A similar approach can be applied to other objects/features. This technique of doing fine scans to extract highly accurate, single point features is common practice today, provided that the scanner has this type of capability.

Section created by tracing over scan data using point cloud CAD-integration software and MicroStation. Image courtesy Washington Group International.

Extracting Point IDs, Features and Symbols

When professionals first see point cloud data, they may quickly note that a point cloud has unmistakably captured a traffic sign or light pole or other objects often mapped as symbols. The point cloud provides the highest level of confidence of what the object is and where it is; the question is, "How do I map a feature that I see in the point cloud as a symbol in my deliverable?" Today, software allows a user to select a point from the point cloud and attach a point ID with standard or custom feature codes to that point. This is essentially "virtual surveying," whereby point cloud software emulates the functionality of standard handheld data collectors, except the work can happen in the office rather than in the field. When the selected point and features codes are exported via ASCII format and Excel tables to standard mapping software, the symbols are likewise recognized and placed on the map deliverable. Some software also provides automatic sequential numbering capabilities.

Taking Distance Measurements

All commercially available point cloud software has the ability to extract slope distance measurements from any two selected points. This capability is very useful for quick, approximate distance measurements. Horizontal and vertical distances can also be readily extracted, provided the software has this capability. Again, the accuracy and utility of the distance measurements from the point cloud are constrained by the cloud's point-to-point spacing and the accuracy of each selected point. For more accurate distance measurements, point cloud modeling is needed. As noted above, such distance measurements can be made within specialty point cloud software or directly within CAD.

Extracting Lines

Creating line work is a mainstay of surveyors and other measurement professionals, and it is here that many wonder how you get from dense point clouds to traditional line work, especially 2D plans and elevations. In the earliest days of laser scanning, the only way to get to 2D plans and section line work was to first create a 3D model and then extract 2D plans and sections from the 3D model. Today, however, there are many ways to get to 2D (or 3D) line work without having to create any 3D models.

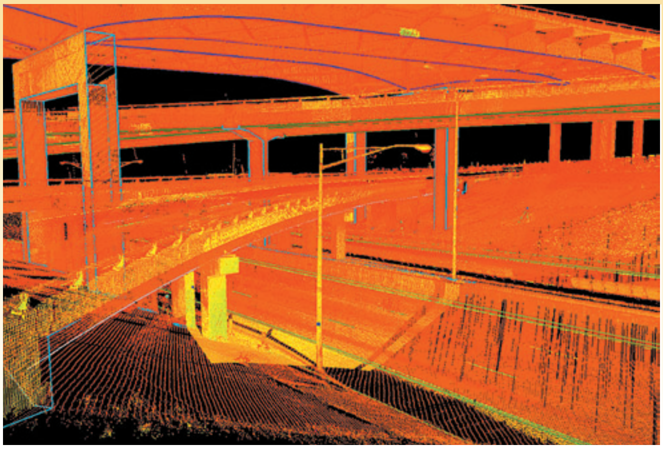

3D line work from point clouds. Image courtesy Bohannan Huston, Inc.

Choosing the best method for line generation is dependent on how accurate the line work needs to be and how efficiently it needs to be generated. One of the fastest ways to generate line work from point clouds is to simply trace over points within the cloud. "Digital tracing" is simply the act of moving the cursor along the point cloud and clicking on or near individual points in a sequence. Line work is then the trail of the digital tracing. If the operator wants to ensure that only certain points are selected, the operator can zoom in on the point cloud area of interest. Software tools are also available to isolate portions of point clouds for doing such work.

Using visualization aids and/or slicing tools, it is fairly easy to get a view of the point cloud, either in 3D or 2D, which clearly shows where to trace over the points. Some software tools automatically adjust the manual tracing to best fit a straight line, a common task for creating line work of buildings. Some software will also automatically create line work from point clouds without intervention from the operator, but this approach appears to have a long way to go before it is robust enough for creating final deliverables. The problematic issues here are that fully automatic line generation requires high point density and regular object geometry. In the real world of laser scanning, adequate point spacing and regular object geometry are not always simultaneously present. Many architectural facades, for example, have highly irregular geometry that make automatic line fitting quite difficult.

Another useful method for creating lines is "semi-automatic" line generation. In this case, the operator focuses on a specific region of the scan data with *a priori* knowledge about the types of geometry that are applicable. Good examples of this are the creation of line work for curbs and gutters and the creation of "line work"

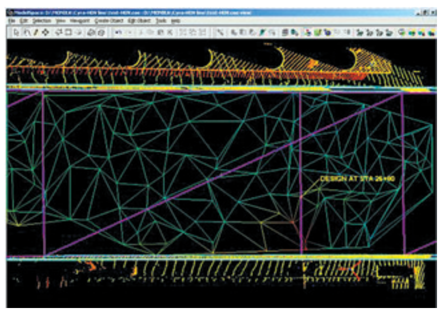

Surface created from high-definition survey. Image courtesy BSC/Cullinan, Inc.

representing road cross-sections at regular intervals. [Although road cross-sections are not commonly referred to as "line work," the sections are lines—as opposed to points or surfaces]. For curbs and gutters, a "cross-section template" can be created that best fits one section of the curb/gutter. Once this template is created, the operator can use this template to automatically follow curb and gutter points that are continuous within the scan data thus creating line work. In the case of road cross sections at regular intervals, the user can select large portions of the point cloud data set and then specify within the software the intervals at which sections are to be generated. The resulting sections are generated within seconds ready for export to CADD software.

Contour lines can also be generated automatically from surface meshes created from point cloud data. Data cleanup is important when generating the original surface meshes; otherwise contour lines may not follow the surface of interest.

Extracting Surfaces

Many types of point cloud software have the ability to generate surface models. These can be created from meshes or from CAD primitives, where appropriate. For civil applications such as road surfaces, initial work is done within the point cloud software to cleanup the data and to create TIN meshes. Meshes can be decimated intelligently to reduce file size while retaining critical changes in geometry. After this step, the points themselves are often exported to civil design software for creating final DTM deliverables.

Some software provides the capability to create NURBS surfaces, which are complex mathematical models of irregular surfaces. This type of software is commonly used in mechanical engineering (*e.g.*, automotive parts and reverse engineering) and in "organic modeling" domains for the entertainment industry

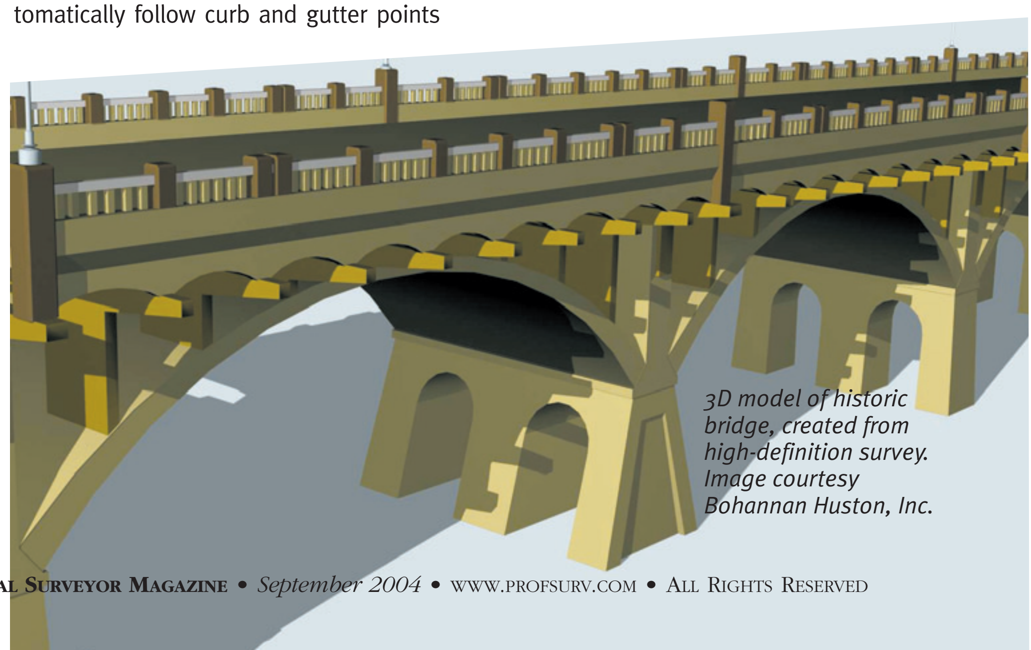

3D model of historic bridge, created from high-definition survey. Image courtesy Bohannan Huston, Inc.

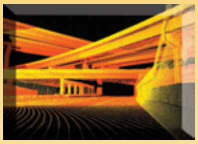

(*e.g.*, models of human faces), and it has also been used in architectural applications. This type of software can be fairly costly and typically requires advanced 3D modeling competency to effectively create accurate NURBS surfaces.

3D model of piping from scan data with true color overlay

Extracting Volumes

Extracting volumes from point cloud data can be a highly efficient process, which also yields more accurate volume calculations than traditional methods. Today, volumes can be extracted between surfaces (including TIN meshes) or between a surface and an assumed plane. The process involves creating surface geometry first and then extracting volumes and surface-to-surface deviations. This latter functionality can also be useful for excavation projects and for monitoring surface movements over time.

Extracting 3D Models

Although the need to create 3D object models for civil projects is uncommon, the extraction of 3D models from point clouds is quite well developed in advanced point cloud software, as this is where most point cloud processing software began. Over the years, there has been a steady stream of advancements in this area that has dramatically reduced the office time associated with creating 3D object models from point clouds, especially for plant and facility as-builts. Typically, object models generated from point clouds are surface models, *e.g.,* cylinders, planes, spheres, steel shapes, other geometric primitives, and arbitrary "catalog" shapes. Tools are also available in advanced point cloud software for applying standard spec tables, such as for piping and steel shapes.

Three-dimensional object models are generated "piece-by-piece" from point clouds. In some software, the user can click on a point that obviously lies on the surface of a pipe (for example) and instruct the software to fit a cylinder using all other appropriate points near the selected point. Similarly, the user may select a point on a wall or floor and instruct the software to create a plane, and so forth. Tools allow the management of CAD-style layers so that all piping can be shown in one layer, all equipment in another layer, etc. Three-dimensional models can be exported to intelligent 3D design software that enables the creation of "intelligent 3D models" that represent additional attributes of the object, such as the thickness and pressure rating of piping.

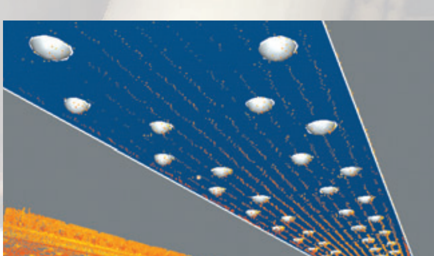

Models of bolt pattern on steel girder. Image courtesy Meridian Associates, Inc.

Three-dimensional models of irregular objects, such as statues, cars, boulders, etc. can also be created, but this involves more intensive office processing and the use of specialty software. Likewise, texture mapping can be applied to the point clouds or geometric surfaces for photo-realistic models, but this is also generally done outside of point cloud software using other specialty applications such as AutoDesk's 3D Studio, Bentley's TriForma, etc. As with complex surface modeling, creating photo-realistic 3D models may involve competencies not common to civil/survey organizations.

Using Point Clouds Directly – the New Survey Deliverable

As noted above, there are many ways to create a variety of deliverables from point cloud data sets. Perhaps the most interesting development, however, is the emerging use of point cloud data sets as a deliverable, *without* extracting anything from the data! Here's the scenario: A piping designer needs to check if the proposed routing of a new pipe through an existing plant will clash with any existing piping, structures, equipment, stairs, etc. Rather than develop an as-built of the existing plant and then import the as-built into the piping design model, the designer can simply integrate the proposed design with the point cloud. Software tools can automatically detect any interference between the proposed design and the point cloud. Measurement tools can be used to determine the geometric extent of any clash and the proposed design can be adjusted and re-checked.

The benefits and potential implications of this workflow are far-reaching. First, the office costs of creating as-built drawings are eliminated. More importantly, doing design work directly inside point clouds can dramatically reduce the risks and consequences of "omissions."

Many readers would be surprised at how extensively this work process is being used today. It is used often in plant and facilities retrofit design projects and has also been used in civil and architectural projects. This process is not only used in retrofit design, but also for comparing actual fabrication/construction against design. Clients can keep point clouds for archive and then simply view areas of concern themselves and decide whether or not to extract geometry from the data. Direct use of point clouds is a powerful, highly cost-effective work process that many predict will one day be a standard workflow.

Summary

In summary, point clouds from high-definition surveys represent a rich source of geometric information. Tools and work processes are available today that enable the extraction of traditional points, 2D and 3D lines, surfaces, volumes, features, and models from these rich point cloud data sets. Another new deliverable has also emerged: the registered, geo-referenced point cloud itself. Point clouds can be used directly for better retrofit design and fabrication/construction QA, without having to first extract other deliverables.

Geoff Jacobs *has been employed by Leica Geosystems HDS, LLC (formerly CyraTechnologies Inc.) since 1998. He currently acts as Senior Vice President, Strategic Marketing. He is also a Contributing Writer for the magazine.*

Design verification within point clouds