HIGH-DEFINITION SURVEYING (HDS): A NEW ERA IN REALITY CAPTURE

Erwin Frei, Erwin.Frei@lgshds.com
Jonathan Kung, Jonathan.Kung@lgshds.com
Richard Bukowski, Richard.Bukowski@lgshds.com

Leica Geosystems HDS, 4550 Norris Canyon Road, San Ramon, CA 94583, USA

KEY WORDS: Keynote, Proceedings, NATSCAN'04, High-Definition Surveying, Laser Scanning, Terrestrial LIDAR

ABSTRACT:

Terrestrial Laser Scanning, as a technology, has gone through an impressive development phase. Although laser scanners have been available for over ten years, their beneficial applicability as productive and competitive tools has just started to emerge in the last three years. On one hand, this may be attributed to large improvements in data acquisition rates, the photorealistic representation of reality, and a significant increase in accuracy, through the use of different base technologies (e.g., "Time of Flight", "Phase Based", etc.); on the other hand, and perhaps even more significant, there have been break-throughs in the handling, managing, and visualizing of large Clouds of Points (CoP). These advancements have opened up entirely new perspectives. The improvements in software applications to handle, visualize and mine these CoP's in an effective and efficient way have reached a stage where several industries and users have utilized this technology as their primary tool and workflow. These developments have created a new era in Surveying: the era of "High-Definition Surveying" (HDS). In applying these new tools and workflows, the physical reality can be captured in an efficient, accurate and reliable way, with tangible benefits for engineers and scientists. In addition, CoP's are now treated as a new form of data content, and their seamless integration into CAD, GIS, LIS and ERP applications have been successfully demonstrated and are the subject of many current developments.

The new capability to handle large CoP's is driving the demand for integrating data from different sensors into a single data model. This data fusion, where data from terrestrial systems (e.g., Laser Scanners, Total Stations and GPS) are combined with data from airborne LIDAR systems, photogrammetry, and remote sensing applications, is opening up new dimensions of data extraction, data interpretation and information gathering.

This keynote will briefly present the status of Terrestrial Laser Scanning, its potential, and its limitations, and will introduce the concepts and ideas behind "High-Definition Surveying". A more detailed look into the developments and the capabilities of a state-of-the-art Point Cloud Engine (*pc*E), which supports the interactive management, visualization, and processing of CoP's with billions of points, will lead into an outlook of the feasibility and current state of combining data from different sensors and systems into one seamless, high-definition model of physical reality. All of these latest developments, tools, workflows and methodologies will continue to change the way reality can be captured and utilized for the benefit of the broader engineering and scientific community.

1. INTRODUCTION

When just 10 years ago a few technology visionaries started to look into new ways of capturing as-built information, the motivation was clearly to fundamentally improve the quality and reliability with which engineers and designers can interact with an accurate geometrical and physical model of reality. The capturing was just one part of the obvious challenge to be solved; the other part was the seamless integration of the information in their native CAD and Engineering systems. Missing, inaccurate or even wrong as-built information was and still is a huge cost driver for any design, engineering and construction project. Although, there were some technologies available for the tasks at hand at that point in time, none of those technologies were able to deliver a cost effective, accurate and fast capturing of as-built information and the subsequent efficient utilization of the captured data for engineering tasks and purposes. The emerging terrestrial Laser Scanning technology seemed to possess all the basic ingredients to meet all of these requirements.

Subsequently, terrestrial Laser Scanning has gone through an impressive and rapid development. The first years were mainly devoted to develop the basic hardware technologies and the first prototype systems. Time of Flight (ToF) and Phase Shift (PS) techniques were utilized as the fundamental measuring technologies for the first scanning systems. With rapidly improving capabilities of these early systems in terms of the achievable measurement accuracy as well as the scanning speed, the development of software systems to handle the growing size of captured data sets got more attention and soon became an area of major developments by itself. The potential of this emerging technology could be successfully demonstrated in various applications and in many different vertical markets.

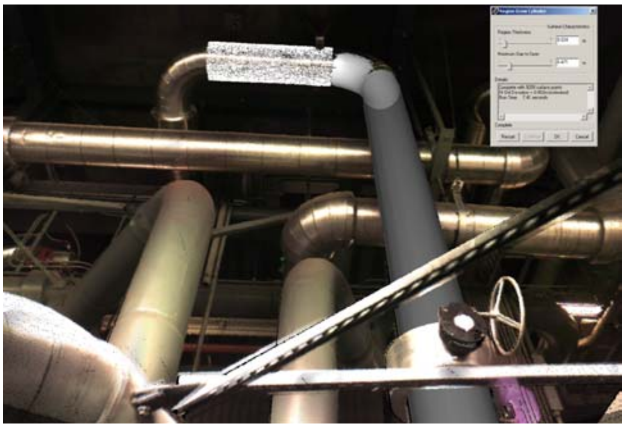

Figure 1. CoP's with photo overlay

Nevertheless, it has taken almost 6 years before terrestrial Laser Scanning was used for the first time in commercial projects. This has to be attributed to the fact that the technology is disruptive in its nature and calls for different workflows and approaches to fully exploit the rich and often overwhelming information content provided through these dense CoP's. Undoubtedly, terrestrial Laser Scanning has to be seen as a disruptive technology and its utilization to the full potential will take time. There are still ways to go before the apparently unlimited potential of this new technology will be fully exploited and becomes a truly sustaining mainstream technology.

Such a huge step forward in the successful deployment of terrestrial Laser Scanning was the introduction of High-Definition Surveying (HDS). The combination of the latest hardware and software technologies to provide seamless and efficient workflows as well as all the required tools to integrate the captured data with existing engineering tools goes a long way in improving the applicability of laser scanning for real jobs in a still early stage of a groundbreaking technology development.

State of the art laser scanners are capable to generate today in a very short time huge data sets. We are talking here 200 million points in just 15 minutes. In order to be able to manage, visualize and query these xx-large data sets totally new approaches and concepts had to be developed. This was the primary motivation for Leica Geosystems to develop the point cloud engine *pc*E, which can store and handle billions of points and hundreds of gigabytes of data, visualize the data quickly and allow examining and extracting information from the data in an interactive way.

The latest advancements in handling, managing and rendering large CoP's as well as in developing tools to efficiently extract information from these large data sets are truly impressive. These significant performance improvements have boosted the interest to combine data from different sensors and data sources to form even larger CoP's and hence even more detailed models of our physical reality. The obvious benefits of such a data fusion opens up a new and exciting chapter in managing, visualizing and mining large CoP's.

This paper will briefly summarize the developments of terrestrial Laser Scanning over the past ten years and will describe performance characteristics of state-of-the-art Laser Scanning systems. The latest developments in managing, visualizing and mining CoP's and the resulting capabilities will be discussed and explained. The paper will conclude with a reflection on the latest developments in the area of data fusion with an outlook for future opportunities in the field of Laser Scanning and point cloud management.

2. TERRESTRIAL LASER SCANNING

2.1 Base Technologies

Basically, there are just two fundamental technologies being utilized at this point in time for high precision terrestrial laser scanning, the time of flight technology and the phase-based continuous wave technology. Significant research and development work for both technologies started at about the same time and roughly ten years ago the first early prototype developments for real products were started. It has taken more than 5 years to get the first time of flight based scanner in 1998 into the market. Currently available products can be considered third generation products. It has taken the phase based versions of laser scanning a couple of years longer to get into the commercial market. However, the latest developments on the phase based side of the technology are impressive in terms of capabilities and performance characteristics and are truly supported by major advancements in handling and managing xx-large data sets. These two different technologies possess very distinctive characteristics and complement each other in an almost ideal way.

We will refrain from a detailed discussion in this paper of all the properties and the limitations of these technologies and would like to refer to some fundamental papers covering these aspects in a comprehensive way (Jon Hancock et al., 1998; Ryan M. Sullivan et al., 2000; Alonzo Kelly, 1997, W. Boehler et al., 2004)

2.1.1 Typical Performance Characteristics: Terrestrial scanners based on the time of flight or the phase based technologies have very distinctive characteristics. The following table characterizes the typical performance of these two technologies and does not represent an attempt to competitively compare the two technologies against each other.

Characteristics	Leica Geosystems Time of Flight	Leica Geosystems Phase Based
Applicable scan range	0.5 – 100 m	0 – {20 m, 50 m}
Applicable target range	Up to 100 m	< 10 m
Scan Density max.	1.2 mm @ 20 m	6.3 mm @ 20 m
Scan Density min.	selectable	101 mm @ 20 m
Scan Speed	1 – 3 kHz	200 – 600 kHz
Data Volume	moderate to xx-large	xx-large
Point Accuracy	3 – 6 mm	3 – 6 mm
Accuracy at Range	4 – 10 mm	n/a
Field of View	360x270	360x310
Spot Size	6 mm @ 50 m	5 mm @ 10 m

Table 1. Performance characteristics

In a more simplistic way, the main difference between the two technologies is the scanning speed and the applicable range.

3. HIGH-DEFINITION SURVEYING

3.1 The Idea and the Concept

Terrestrial Laser Scanning as a technology has gone in the first few years through an impressive and steep development curve. It reached a stage where individual components, namely scanners and cloud of point processing software modules were readily available for a broader deployment to real tasks in real projects. However, it was left to a great extent to the users to combine these individual components to design and establish efficient and effective workflows, which in turn were supposed to deliver the anticipated results. The main focus at this point in time was clearly on capturing and registering CoP's. This overemphasizing on the data acquisition side did not necessarily help to exploit the real value of the technology and to generate the expected returns of investments.

3.2 The Value is in the CoP's

The CoP's themselves do not necessarily represent the real value of the technology and are in most cases just an intermediate step to get to results and deliverables. The real value is in the information implicitly stored and contained in the CoP's (size, shape, form, location and surface characteristics of objects in the real world). The application specific extraction of this information and the subsequent management, analysis and documentation of this information generates the real value of the technology and hence delivers the anticipated results and the expected added value.

3.3 The Components of High-Definition Surveying

Based on these considerations and some fundamental new developments Leica Geosystems introduced in September 2003 a new series of products and solutions under the name of High-Definition Surveying (HDS), which are clearly supporting the customer in his/her task to generate value and results with solutions tailored specifically to his/her problems and needs.

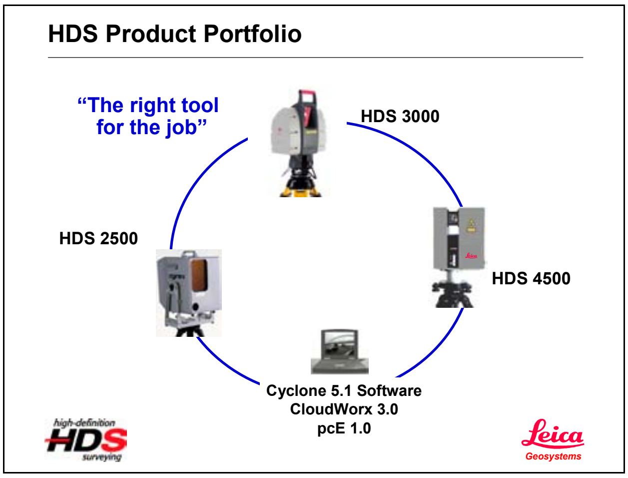

Figure 2. HDS Product Portfolio

HDS combines and integrates various scanning technologies (HDS3000 with a time of flight technology and the HDS45000 with a phase shift technology) and different software modules (*Cyclone*: Scan, Register, Model) with tools to seamlessly integrate the CoP's into CAD systems (CloudWorx: for AutoCAD, for MicroStation; *pc*E: *point cloud* Engine). This integration supports a data and work flow, which delivers results in an efficient and cost effective way for various applications and industries.

3.4 Features and Functionality

3.4.1 Data Management: The HDS software uses a robust object database to store HDS data. The data can be acquired directly from an HDS scanner via the HDS Scan Control interface, imported from a file from one of the supported formats, or generated as second-generation data by software tools. The database supports a hierarchical structure to help the user organize projects and models. Although a database is normally a single, self-contained file, databases can share data to reduce redundancy. Because of the volume of data acquired by HDS scanners, databases can grow in size to ten gigabytes or more. However, the database size does not affect stability or performance of the data.

HDS databases are served by a database server to HDS software clients. The client can be on the same computer as the database server. Multiple clients can access the same data on a given server. Modifications to the database are immediately observed by the other clients using that database. For performance reasons, the client only loads the parts of the database that it needs.

3.4.2 Visualization: HDS scanners acquire discrete samples of real surfaces. There may be over one hundred million samples acquired from a single position and billions of samples in a single project. Visualizing each sample as a 3D point is an obvious approach; rendering the collection of points in 3D displays a "CoP's".

In addition to the 3D coordinate of the sample, a point may also have other properties. For example, colouring each point based on the true colour from a digital image can aid visualization. HDS scanners also measure and record the "intensity" (strength of the return signal) of each sample. For example, surfaces that absorb the laser have lower intensities than surfaces that reflect the laser. Other factors that can affect the intensity include specularity, surface colour, and distance to the scanner; the ambient light in the scene does not affect the intensities.

The HDS software can optionally apply a false-colour mapping to the point cloud, based on the intensity of each point. *Greyscale* and *Hue-based* are the two colour maps most commonly used. Greyscale resembles a black-and-white photograph when the samples are dense enough, so it can be more familiar to the user. The hue-based (ranging from red, green, yellow, blue) colour map is not limited to the 256 shades of grey that modern video cards support and the higher contrast can reveal more subtle variations in intensity. Other colour maps include elevation-based maps that resemble topographic charts.

However, without context or other information, a static image of a CoP's can be difficult to interpret. By interactively changing the viewpoint, the user can understand the nature of the structures represented by the samples in the scene, based on how the 3D points shift in relation to each other on the video display. The HDS software must achieve a minimum frame-rate for the user to understand structure from motion, so rendering performance is critical.

To achieve the necessary interactive performance, the HDS software uses a dynamic level-of-detail (LOD) system to organize and display the points. Areas that are closer to the observer are drawn with more detail, and areas that are further away are drawn with less detail. Because the video display has a maximum resolution, the adaptive LOD needs only draw to the pixel level.

Modern computers with a 32-bit memory architecture have finite memory and may not be able to hold more than one hundred million points in memory. Correspondingly, the LOD determines how many and which points to be loaded and displayed. Points are loaded from the database and cached in system memory. Depending on the viewpoint and memory needs, points that are not displayed may be ejected from memory to make room for other objects. The HDS cloud engine manages this automatically for an optimal frame-rate.

Unlike other approaches, this system does not restrict the user to a pre-defined spatial resolution (e.g., no details finer than 1 centimetre) or require that the user pre-define areas of interest to be pre-processed and then loaded. The entire database of clouds is available dynamically, but is loaded intelligently based on the current viewpoint and task. A single point cloud may span hundreds of kilometres but still have points that are within millimetres of each other, where each point retains sub-millimetre precision.

The HDS software supports a combination of point clouds and geometric surfaces (e.g., boxes, cylinders, etc.). Geometric surface objects, which may be used to represent piping runs, building facades, etc., have a similar LOD system.

Because projects can be geo-referenced, registered coordinates can be large. The HDS software uses a custom 64-bit rendering engine that is capable of rendering points and objects millions of kilometres from the origin, and of scenes that span millions of kilometres, all with sub-millimetre precision. This 64-bit engine overcomes the 32-bit limitations in modern graphics cards, but simultaneously takes advantage of common graphics hardware acceleration for improved performance.

3.4.3 Data Processing: HDS scanner coordinates are not automatically referenced to a given coordinate system, they must be registered to the same coordinate system before scans from different locations can be combined into a single model. The HDS software provides robust registration tools. Registration can include: registering to surveyed coordinates using scans of those same targets; registering to other scans using common modelled geometry (e.g., targets, common surfaces, etc.); and/or registering to other scans using the surface geometry implied by the points in the clouds. Registration can be pair-wise (one position is registered to another) or global. Global registration is a sophisticated process that takes the pairs and simultaneously resolves the best arrangement for them all. The HDS software records statistics from the registration in the database for the user to inspect.

Once the point clouds are in a common coordinate system, a project consisting of multiple clouds can be conveniently visualized. However, although simply visualizing the point clouds may be sufficient for some deliverables, additional processing to extract non-visual information from the points is often required.

Because point clouds can be large and complex, the HDS software provides various tools to manage what is displayed. For example, the points outside a user-defined volume can be hidden, letting the user focus on the remaining points. Points can be selectively segmented from a cloud and deleted, placed on a layer (similar to CAD layers), or left segmented as a separate entity from the rest of the cloud. Because the original clouds are never deleted from the database, the points in the original clouds can always be recovered.

The linear distance between two points can be measured, although the accuracy of this measurement depends on the per-point accuracy of each sample as well as the scanning resolution. Potential interferences with proposed changes can be directly checked against the as-built conditions as represented by the point cloud. The HDS software supports the creation of Triangulated Irregular Network (TIN) meshes from points, with optional break lines. Meshes may be decimated intelligently, preserving triangles in regions with detail. The meshes support volumetric measurements, including surface-to-surface deviations.

Surface geometry can be best-fit to selected points. For example, cylinders represent pipes, planar surfaces represent flat walls, etc. The operations record statistics about the fit in the database, with the object. The user can query the object for its attributes and statistics. To supplement fitting, a variety of linear and surface geometry can be manually inserted into the model; 2D drawing tools support line-work.

3.5 Typical Applications and Workflows

HDS is applicable in many situations. Although there are substantial differences in the technologies, the HDS2500 and HDS3000 are both suited for medium-range civil/infrastructure projects (freeways, streets, terrain), as well as process plant and architectural projects. The HDS3000 has a high-resolution digital camera, and its 360°x270° field of view reduces the number of scanner setups; the HDS2500 can scan straight down and does not need to be leveled. The range of the HDS systems also make it safer for the operator and makes the measuring less disruptive, e.g., by reducing the need for lane closures on a busy freeway. With its high acquisition rate and 360°x310° field of view, the HDS4500 is efficient in process plant and architectural projects. The Leica HDS software suite supports these and other applications.

3.5.1 Scanning with HDS Scanner. After examining the site and planning the sequence of scanner setups, the operator places the HDS scanner at the first position and acquires the data available to that location, using *Cyclone*-SCAN software. If required for registration, targets are scanned and optionally surveyed. The HDS3000 supports use of station data to define the location of the scanner in an absolute or local coordinate system. If needed, the operator repeats the process at the next scanner location and/or orientation. The data are stored in a single database.

3.5.2 Registration: With the data captured, the *Cyclone*-REGISTER module is applied to combine the different scanner positions into a single coordinate system. Optionally, the user imports data from other devices (including TPS, GPS, laser scanners, etc.) to include in the registration. Although exclusive use of targets and exclusive use of cloud-registration are both supported, users have been combining both, supplementing target registration with cloud registration. While still at the field, the operator will often perform registration on the data already collected to verify that the scanner locations overlap sufficiently. Once registered, the data are ready for delivery to an end-user, who can use the data for visualization, as an archival record, for measurements, modelling, and more.

3.5.3 CloudWorx: By making the point cloud available to the CAD platform, the CAD user can make use of the point cloud within a familiar environment and leverage existing CAD tools that can use the points. CloudWorx supports a number of visualization, measurement, and reverse engineering tools that operate on the point cloud as a read-only object. The CloudWorx user typically receives a registered point cloud from a service provider and might never use the *Cyclone* software directly. The database is sometimes served from a central database server for multiple users to access concurrently.

3.5.4 Virtual Surveying: The data from HDS scans is often so complete that a user can take measurements based solely on the points, even without ever having been present at the real world site. This *virtual surveying* can be a cost-effective alternative to conventional surveying. Indeed, the same cloud data can be used for different surveys. The completeness of the data can reduce the need for site re-visits; the surveyor often finds that the already-captured data will also support a requirement newly-added by the client. Virtual surveying includes: acquisition of significant features such as lane stripes and gutter lines; 2D or 3D line-work; creation of cross-sections along a reference alignment; volumetric quantity surveys. *Cyclone*-SURVEY and *Cyclone*-MODEL support these workflows; CloudWorx currently supports a subset.

3.5.5 Reverse Engineering and Design Validation: The scanned points are samples on real world surfaces. A surface may be a wall, floor, pipe, street, terrain, etc. The deliverable often requires that these points be converted into some other form, e.g., for export as a solid or TIN mesh to CAD. By using *Cyclone*-MODEL, the user can transform the point data into other forms. CloudWorx supports a subset of these tools. Some deliverables require the entire project to be converted to surface geometry, whereas transforming selected portions is sufficient for others.

Points need not be necessarily modelled to support engineering deliverables. Representing the as-existing conditions, the point cloud can be used to quickly check for potential interferences of a proposed design. The design may be imported into *Cyclone*-MODEL, or may be native to the CAD platform on which CloudWorx is running.

4. POINT CLOUD ENGINE

4.1 The Challenge

With laser scanners that can capture hundreds of thousands of samples per second, management of the massive quantity of output data – enough to completely fill an entire state-of-the-art hard drive in a few days of normal use – becomes a serious challenge. Standard CAD programs are designed to work with datasets that can fit entirely in the computer's system memory; such datasets can be at most about 2 gigabytes, and typically less than half that. In contrast, current phase scanners such as the Leica HDS4500 can generate 1 gigabyte of compressed data in a single scan, representing about 15 minutes of real time and 200 million point samples. Early attempts to manage the data in the standard CAD fashion – loading each point as an entity from a file and working in system memory – failed spectacularly due to the fact that this method allows at most a few hundred thousand points to be manipulated at once, even on the largest computer you could buy.

Even if these memory constraints were lifted – if you were to imagine a 64-bit computer with a 64-bit CAD package and hundreds of gigabytes of system RAM – the massive, spatial nature of the data requires a different approach to visualization and manipulation. The fastest computers can only draw tens of millions of points per second; thus, the system needs to be intelligent about which points it needs to draw, because the traditional method of drawing everything visible would mean that it could take minutes to redraw the screen. Similarly, measurement and fitting methods often assume that the data is either pre-partitioned by the user into sets of interest, or that the dataset is small enough to exhaustively search in a reasonable amount of time. The former requires the ability to rapidly segment (cut into pieces) a large cloud, and the latter requires a sophisticated, database-style spatial query system. It is interesting to note that segmenting a huge cloud efficiently itself requires a spatial query system, so such a mechanism is clearly a basic requirement. Advanced databases such as Oracle 10g have recently started to provide such query systems, which are suitable for some measuring and extraction tasks, but still fall short of the functionality needed to support rapid 3D visualization of large clouds.

Thus, the technology to support our three basic operations on large point datasets – storing and retrieving the data, visualizing the data efficiently, and extracting information and measurements quickly and efficiently – are not supported by current CAD or database systems. This is the motivation behind the concept of the point cloud engine, or *pc*E, which should simultaneously provide the ability to store and index large cloud datasets (billions of points, hundreds of gigabytes of data, which cannot plausibly be loaded into memory at once), visualize the data from a certain viewpoint quickly (ideally in a matter of seconds), and allow measurement and extraction algorithms to perform spatial database queries that efficiently access localized parts of the data. Furthermore, the ideal *pc*E would interoperate seamlessly with other CAD systems to provide these functions within the user's desired visualization environment, be that a CAD package, a visualization and rendering package, a design review package, or on the web.

4.2 State of the Art Solutions

4.2.1 Managing Large Point Datasets: Managing a multi-gigabyte dataset efficiently requires a data engine that allows the system to dynamically load and unload sections of the data from hard disks into memory, such that the amount of data in memory at any given time is small enough for the system to handle. This is traditionally called a database, as opposed to a "flat-file" system which loads an entire dataset into memory, manipulates it, and then writes it back. Without a database, the system must either operate on a decimated representation of the dataset, which loses precision, or on a small section of the dataset, which creates problems in how to partition the data and which partitions to load to perform a particular task. Traditional databases such as Oracle operate in this fashion; access to the data is done through *queries*, where a portion of data is dynamically selected by sending criteria such as where in space you are interested in ("I want all of the data in the left side of this room") or the characteristics of the data ("All points which returned high intensity"), and the engine returns only those portions of the data that meet those requirements, and unloads that data when the function is finished executing. In this way, the entire dataset is available for viewing and computation, but the system memory requirements are reduced. A further benefit of this system is that once the database system is divorced from the client application doing the measurement, extraction, and visualization work, the database can now reside in a central location and be accessed via queries by many users, eliminating the headaches caused by replicating datasets on many client machines and tracking multiple versions of datasets.

There are many large database systems available; it would be convenient if existing relational database systems were directly applicable to this task. Unfortunately a traditional relational database has a number of limitations that make it unsuitable for the typical user of point cloud datasets. First, the types of geometrically and spatially based queries needed for efficient visualization and manipulation of 3D datasets are not available in standard database systems. These queries require special indexing structures that have not been integrated into current offerings. Second, these databases store geometric data in a "naïve" fashion that causes the same data to take up a huge amount of space relative to a more dedicated representation designed for point clouds. Finally, these systems are expensive and difficult for typical users to manage.

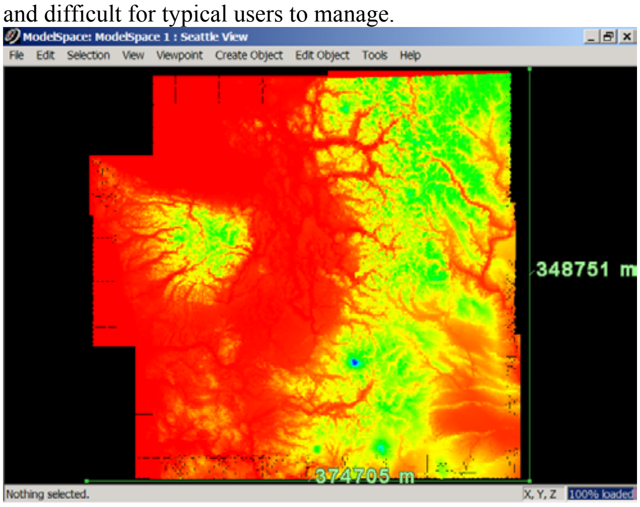

Figure 3. XX-Large data set (1.2 Billion points)

These shortcomings lead to the creation of custom database engines specifically designed for storing, viewing, and managing point cloud data to its full precision. The Leica HDS *Cyclone* family of software is based on such a customized database. With such a solution, the entire dataset is on-line at all times; it is seamless, without borders, so the user need not worry about whether the data they need is loaded or not; it can be visualized quickly with dedicated viewing queries that provide view updates in seconds, regardless of the data size; data can be represented in a compact form that consumes minimal disk space; and it can be measured and manipulated efficiently, with queries that process data quickly with a small memory footprint.

4.2.2 Visualizing and Extracting Information: In order to visualize a large point dataset efficiently in 3D, the database must support a specialized 3D *viewing query*. This query is essentially asking the database to load a particular number of point samples (say, 2 million) that best represent the scene from a given viewpoint. This allows the system to present a high-quality view of the data without requiring all of the data to be loaded into memory. To support interactive motion in the view, the database needs to be able to cache the shared information between nearby views so that it doesn't have to repeatedly load the same points. For example, if you are looking at a scene and turn your head slightly, most of what you see is the same from the last direction you looked. Similarly, if you turn your view slightly, most of the points in the last scene will be the same as the points in the new scene. The query system should support only loading those points which are different between the two scenes. This allows the graphics to update quickly as the viewpoint moves.

For data extraction, queries need to support specifying an area of interest (say, the area in which a pipe the user wants to model runs), as well as a resolution (say, points spaced at 1 cm along the pipe). Again, the resolution limit allows the system to conserve system memory while still providing the data needed to perform the requested operation (in this case, fitting a cylinder to the samples corresponding to the pipe).

Many data extraction queries can be imagined -- constraints on colour, return intensity of the laser, and so forth, in addition to the basic constraints of location and resolution, to support efficient extraction of data for multiple purposes.

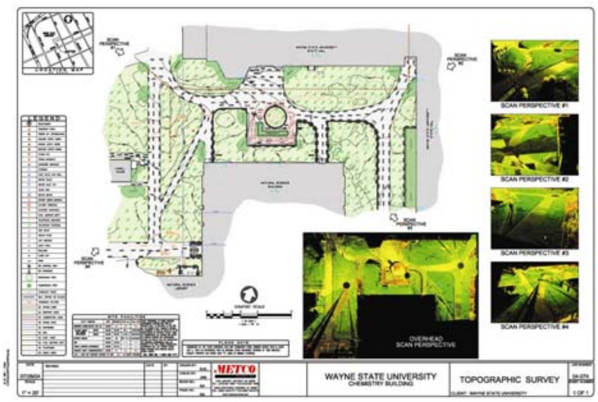

Figure 4. Topography mapping

4.2.3 Interoperating with Other Packages: Like any database system, a point cloud engine should be able to satisfy queries generated from multiple types of program clients, including native applications, CAD packages, or design review software. If this function extends to viewing queries, and the application allows plug-ins to access its drawing system, then the engine can take over the task of rendering point data in other environments. *Cyclone* CloudWorx is an example of this functionality in a CAD system; *Cyclone pc*E for SmartPlant Review is another example. In these examples, the point cloud engine not only allows large datasets to be viewed in native environments, but also interacted with – the user can pick points, measure, and extract data. These functions are implemented by queries to the engine just like a bank teller machine would query the bank's central database to get account information. The dataset never leaves the *pc*E database, but it appears in the client data space as if it was loaded normally.

4.3 Features and Functionality

With the support of a point cloud database, large datasets can be visualized and mined for data in real time in many ways. In a flat-file system, it can take minutes to open even small, single scan files; with *pc*E, a dataset can generate an overall thumbnail view in seconds, and pan and zoom in to detailed views in only seconds more. Once the desired data has been found, it can be visualized in many different ways; as a mesh, by laser intensities or colours, and with many different colour schemes. The engine supports fast collision and interference detection and fast selection of individual samples or groups of samples from the whole dataset. Local filters can use queries to remove noise or identify objects and geometries in the dataset; the identified points can then be used to perform best-fit operations with models or geometry to construct CAD or geometric models of the scene. The fast, resolution-controlled queries are also very useful to support high-speed and high-accuracy registration operations for integrating multiple scans from different sources; since large projects can involve billions of points, the registration operations need the database functionality to be able to register these datasets together, again without having to load more data than they need into system memory. Since one universal point data format is used, data from many different hardware scanners, with different resolutions and attributes, can be integrated into one unified database and operated on together, with the same set of tools. The point cloud engine supports end-to-end processing of huge datasets within a single environment; from acquisition, to registration, to visualization, modelling and fitting, all with a universal set of queries, visualization, and management functions.

5. DATA FUSION

5.1 The Value of Combining Data from Different Sources

Today, there are a number of ways to get accurate dimensional measurements of the world, including: TPS; GPS; HDS and other laser scanners; aerial LIDAR; terrestrial and satellite photogrammetry; and tape measure. Most of these require software to extract the most benefit from the measurements, but the software is tailored to best process that input only. An additional format is required to bridge the gap between different sources. Further, the scale and scope of data measured by the different technologies vary widely, from a sub-millimetre scan of a pencil, to a scan of the room that contains the pencil, to a scan of the building that contains the room, to a survey of the street that the building is on, to the neighbourhood, the city, the region, and so on.

Despite their differences, these methods have something in common: they can all produce 3D points, either indirectly (as in the case of tape measurements) or directly, to represent individual samples. Points have value as the common representation.

Although various approaches build upon point clouds (e.g., meshes, voxels, etc.) or use an alternate foundation (e.g., 2D images), these lack scalability, can be restricted to limited sources of input, or are not broadly applicable outside their targeted domain. In contrast, the Leica HDS Point Cloud Engine is designed to operate on massive point clouds with simultaneous extremes in scope and scale.

5.2 An Early Example

This is an example of a single model served through the Leica HDS *pc*E. The model has over one billion points imported from a U.S. Geological Survey data set (the underlying data were produced via satellite photogrammetry) integrated with a suburb (scanned by aerial LIDAR), a pedestrian bridge (scanned by a Leica HDS2500), an office building exterior (scanned by a Leica HDS3000), and a desk (scanned by a Leica HDS4500). The sequence follows the observer as she zooms into the desk from the initial viewpoint. The loading and viewing happens interactively in real-time on a desktop computer (2.6 GHz Pentium 4, 1 GB RAM, 7200 RPM SATA/100 hard disk).

(See picture series on the next page)

Data Fusion: An early example

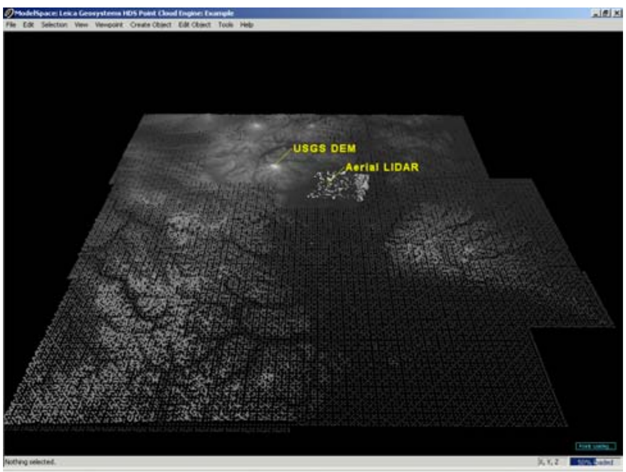

Picture 1. USGS DEM (350 x 375 km) initial loading 50%: +16 sec

Picture 2. USGS DEM loaded for this viewpoint (100%): +9 sec

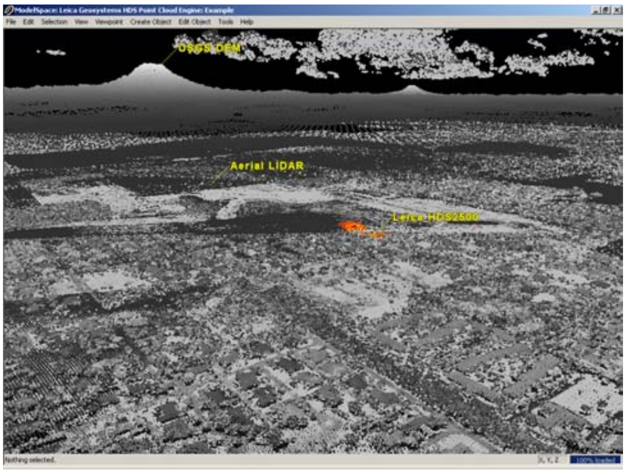

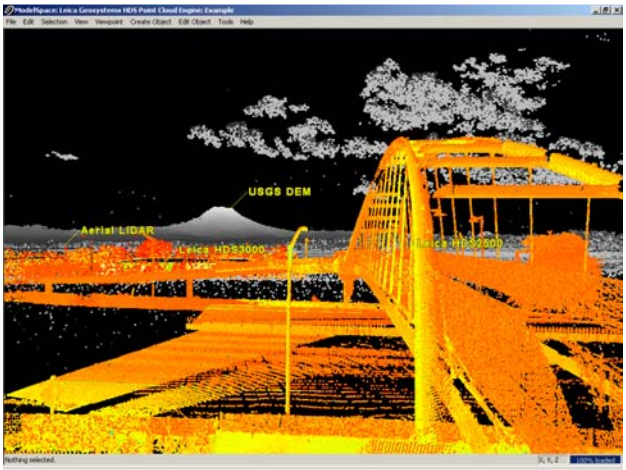

Picture 3. Aerial LIDAR (50 x 30 km) within DEM: +8 sec Note the street layout and the DEM mountain in the background

Picture 4. Bridge (HDS2500) with office and mountain in background: + 5 sec. The clouds in the sky are from the Aerial LIDAR

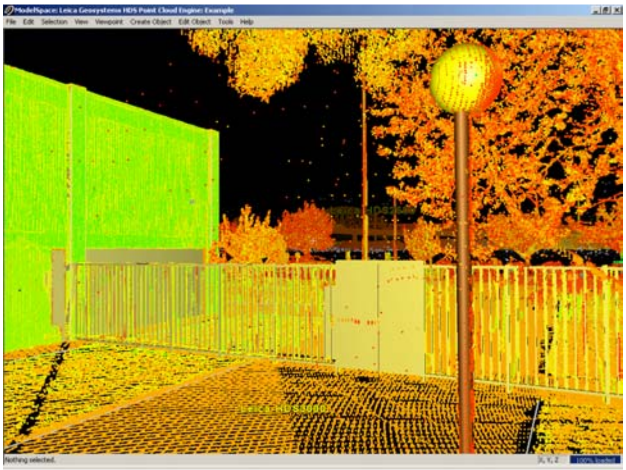

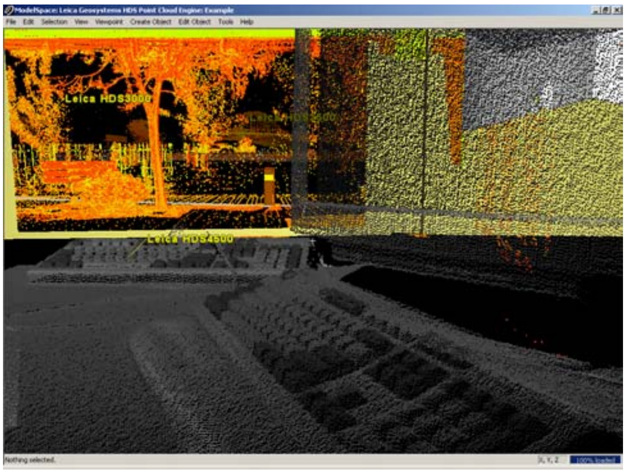

Picture 5. Office courtyard (HDS3000), looking in opposite direction, with bridge in background: +5 sec

Picture 6. Desk (HDS4500) inside office looking towards bridge: +6 sec

6. OUTLOOK

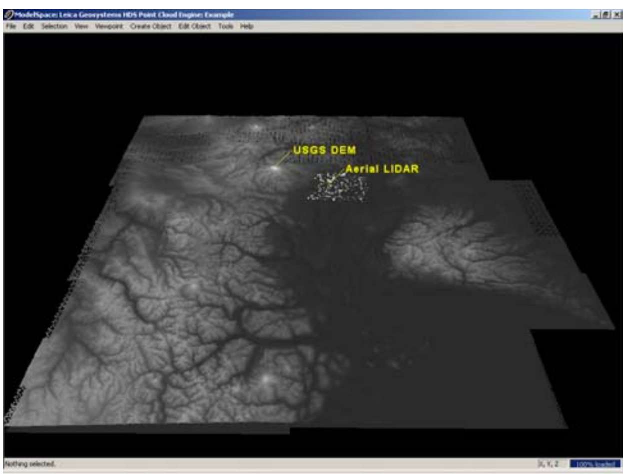

High-Definition Surveying is just in the beginning of a steady and steep development into a sustaining technology. Its deployment for different applications and for various engineering disciplines has just begun. The potential of this new technology to capture reality and its integration into different engineering workflows has seemingly no boundaries. Just the first few years of gathering experiences with this revolutionary way to cope with reality has unveiled a list of sheer endless opportunities to utilize this technology in new and innovative ways. Not just only the AECO industries (Architecture, Engineering, Construction and Owner/Operators) will benefit from these ground breaking developments, but also industries like asset management, GIS, heritage, deformation monitoring, agriculture, security, disaster prevention, forensic etc. just to name a few, will follow rapidly in utilizing the obvious benefits of HDS.

The hardware developments will follow the typical development path of a high-tech product. Smaller, lighter, faster, less power consumption and a significantly improved cost/benefit ratio will drive the HDS hardware components for the generations to come. Additional functionality and a significant improvement in the scanning speed will broaden the applicability of Laser Scanners and make it a true multi-purpose tool. Thanks to possible further increases in the scanning speed of the phase based scanners, measuring on-the-fly might become possible and Laser Scanners could therefore be used to capture reality in a kinematic mode. This opens up new and exciting applications for capturing highways, railway tracks in an efficient and cost effective way. These developments might even become a new basis for guiding and steering construction machinery.

Clouds of point will become a new data content, which will be fully supported in all engineering, ERP and asset management applications. The quality of the representation of objects and elements in these CoP's will increase due to the high-definition nature of laser scanning. The improved quality of texture mapping paired with powerful rendering and visualization capabilities will change the appearance of CoP's so dramatically that there is no need to model objects anymore for many engineering, inspection and managerial tasks. The CoP's itself becomes a photorealistic, geometrically correct representation of reality. This has truly the potential to change the way engineers, designers and managers are utilizing and working with "as-built" and "as-found" information on a daily basis. The "time to model" will be almost instantaneous and the cost for getting complete high quality representation of reality becomes affordable.

In the not so distant future laser scanners will not be the only source for populating large CoP's. Data from Photogrammetry, Remote Sensing, Aerial LIDAR, GPS and TPS applications will contribute to a unified, high definition representation of reality. These xx-large data sets will form the basis for city models, extensive multi-layer processing plants, an open pit mine with the entire infrastructure, a highway strip with bridges overpasses and tunnels, or even a gas-pipeline over several thousands of kilometres. Areas of low resolution can be easily combined with areas of highest detail. These models allow drilling down seamlessly into the smallest engineering details and representing any element of our physical reality in a geometrical correct and detailed way. These models can be queried and mined to unveil rich and complete information about every aspect of the model and the reality, which it represents.

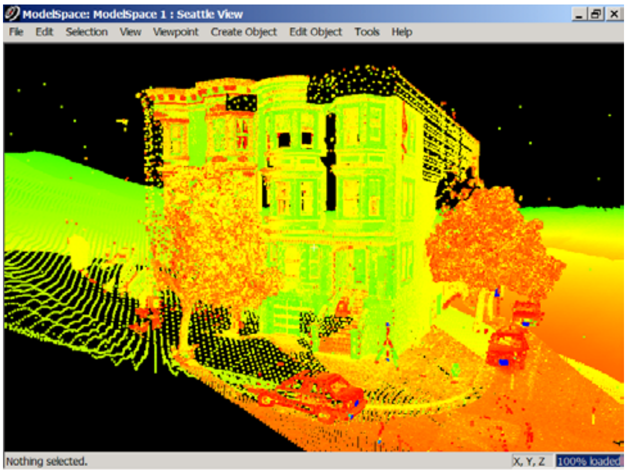

Figure 5. Data Fusion and xx-large CoP's

7. SUMMARY AND CONCLUSION

A fast and accurate capturing of the size, shape and form of our physical reality in an efficient, complete and cost effective way has become available with the combination of the latest hardware and software technologies combined in Leica Geosystems' HDS approach. Capturing data with the latest hardware generation leaves the option for an extremely fast and accurate capturing of data at short distances with a phase based scanner or for an accurate scanning at short to medium distances with a time of flight technology. The combination of these data sources in one unified data base combines the best of both worlds for the benefits of various applications and workflows. Recently developed software functionality allows managing and visualizing billions of data points in a fully interactive mode. Querying these large data sets unveils the rich and complete information for the creation of results and deliverables. HDS is unique in its approach and in its capabilities and has set new standards in the still young history of Laser Scanning.

Leica Geosystems supports with the new development of a point cloud Engine (*pc*E) the utilization of CoP's in engineering, asset management and ERP systems. The integration of the *pc*E allows to seamlessly integrating xx-large CoP's as new data content into these systems. This clearly adds a new quality for the design and the decision taking for engineering and management disciplines. Accurate and complete "as-built" or "as-found" information becomes readily available to support any engineering or management task.

The capability to handle xx-large data sets also provides for the first time the opportunity to combine data from different sources and measurement techniques in a single data set and makes the resulting CoP's even more valuable for many engineering, design and management tasks. Photogrammetry, Aerial LIDAR, Terrestrial LIDAR and even GPS and TPS data can be combined now in one unified CoP's. This new quality in representing reality is truly unique and will provide lots of valuable information and insights for the work of the engineering and the science community.

HDS has opened a new era in reality capturing and is about to fundamentally change the way we are interacting with a geometrically correct and complete representation of reality.

REFERENCES

Boehler W. et al., 2004. W. Boehler, M. Bordas Vincent, A. Marbs. Investigating Laser Scanner Accuracy. *XIX CIPA Symposium at Antalya, Turkey, 30 Sep – 4 Oct 2003*, Updated for Web presentation April 2004. i3Mainz, Institute for Spatial Information and Surveying Technology, FH Mainz, Holzstrasse 36, 55116 Mainz, Germany.

Hancock J. et al., 1998. John Hancock, Dirk Langer, Martial Hebert, Ryan Sullivan, Darin Ingimarson, Eric Hoffman, Markus Mettenleiter, Christoph Froehlich. Active Laser Radar for High Performance Measurements. *Proceedings of the 1998 IEEE International Conference on Robotics & Automation*, Leuven, Belgium, May 1998.

Kelly A., 1997. Alonzo Kelly. *Concept Design of a Scanning Laser Rangefinder for Autonomous Vehicles*. The Robotics Institute, Carbnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, May 8, 1997.

Funkhouser T. at al., 1992. Thomas A. Funkhouser , Carlo H. Séquin , Seth J. Teller, Management of large amounts of data in interactive building walkthroughs, *Proceedings of the 1992 symposium on Interactive 3D graphics*, p.11-20, June 1992, Cambridge, Massachusetts, United States

Funkhouser T. et al., 1993. Thomas A. Funkhouser , Carlo H. Séquin, Adaptive display algorithm for interactive frame rates during visualization of complex virtual environments, *Proceedings of the 20th annual conference on Computer graphics and interactive techniques*, p.247-254, September 1993

Levoy M. et al., 1985. Levoy, M. and Whitted, T. "The Use of Points as a Display Primitive," *Technical Report TR 85-022*, University of North Carolina at Chapel Hill, 1985.

Samet S., 1990. Hanan Samet, *Applications of spatial data structures: Computer graphics, image processing*, and GIS, Addison-Wesley Longman Publishing Co., Inc., Boston, MA, 1990

Levoy M. et al., 2000. Marc Levoy , Kari Pulli , Brian Curless , Szymon Rusinkiewicz , David Koller , Lucas Pereira , Matt Ginzton , Sean Anderson , James Davis , Jeremy Ginsberg , Jonathan Shade , Duane Fulk, The digital Michelangelo project: 3D scanning of large statues, *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, p.131-144, July 2000