

Leica P-Series Scanners: Details that Matter

Gregory Walsh, Ph.D.

March 11, 2015

1. Summary

Leica Geosystems, in creating the Leica Scanstation family of products, has designed and conducted a number of steps to ensure our accuracy and performance specifications are met. This document summarizes some of the steps taken in service and in manufacturing in order to ensure accuracy specifications are met, using the Leica Scanstation P40 as an example. Similar steps have been applied to the creation of all P-series and C-series scanners. Leica is in the business of selling survey instruments and the quality of the data produced by these instruments is central to the value the company delivers to the customers.



Figure 1: The Leica Scanstation P40, a survey grade laser scanner.

The specifications that a Leica Geosystems P and C series scanners are designed and built to satisfy hail from the Leica's surveying heritage. The first section of this document explains what these specifications mean and when they are important. The second section details Leica's manufacturing processes that are critical to support our specifications, and the third section explains some of the validation processes that scanners in the Leica Geosystems portfolio go through in order for us to release them as products.

2. Accuracy with Laser Scanners

Currently there are no standards facilitating the comparison of laser scanners. For this reason, the data sheets of P and C series scanners borrow terminology from surveying total stations. Indeed, the P and C series scanners are designed and built to be what is termed a survey grade instrument. The vocabulary on the data sheet, however, is typically not immediately accessible to those outside the surveying community. For this reason in this section we offer a short description of these terms and why they are chosen. The trouble with vocabulary starts right away with the term Measurement – this is an everyday word – but inside technical circles a slightly different meaning is attached to the idea of “Measuring” something and the resulting “Measurements”.

Measurement Uncertainty

Surveyors have long had to answer fairly practical questions – for example: where is the building located? How tall is the framing? Will the new addition block my neighbor’s view? How wide are the windows? Can I fit the new holding tank through this particular window? Knowing not just a number, but also the uncertainty in this number, is integral to providing a useful service for their customers. Determining that a window is 1.456 meters wide, plus or minus 6 millimeters, at its narrowest point, is far more useful than just the width by itself. The number without unknown uncertainty might possess many digits, say, 1.456234 meters, but if the uncertainty is plus or minus 0.5 meters, for example, then all of that extra detail is not practically useful.



Figure 2: Laser scan of a building provides a lot of detail. Literally this view shows in context millions of measurements.

All laser scanners provide a great deal of detail. Enormous detail does not mean the data is useful. Without a bounded uncertainty, there are many practical questions you simply cannot answer with confidence.

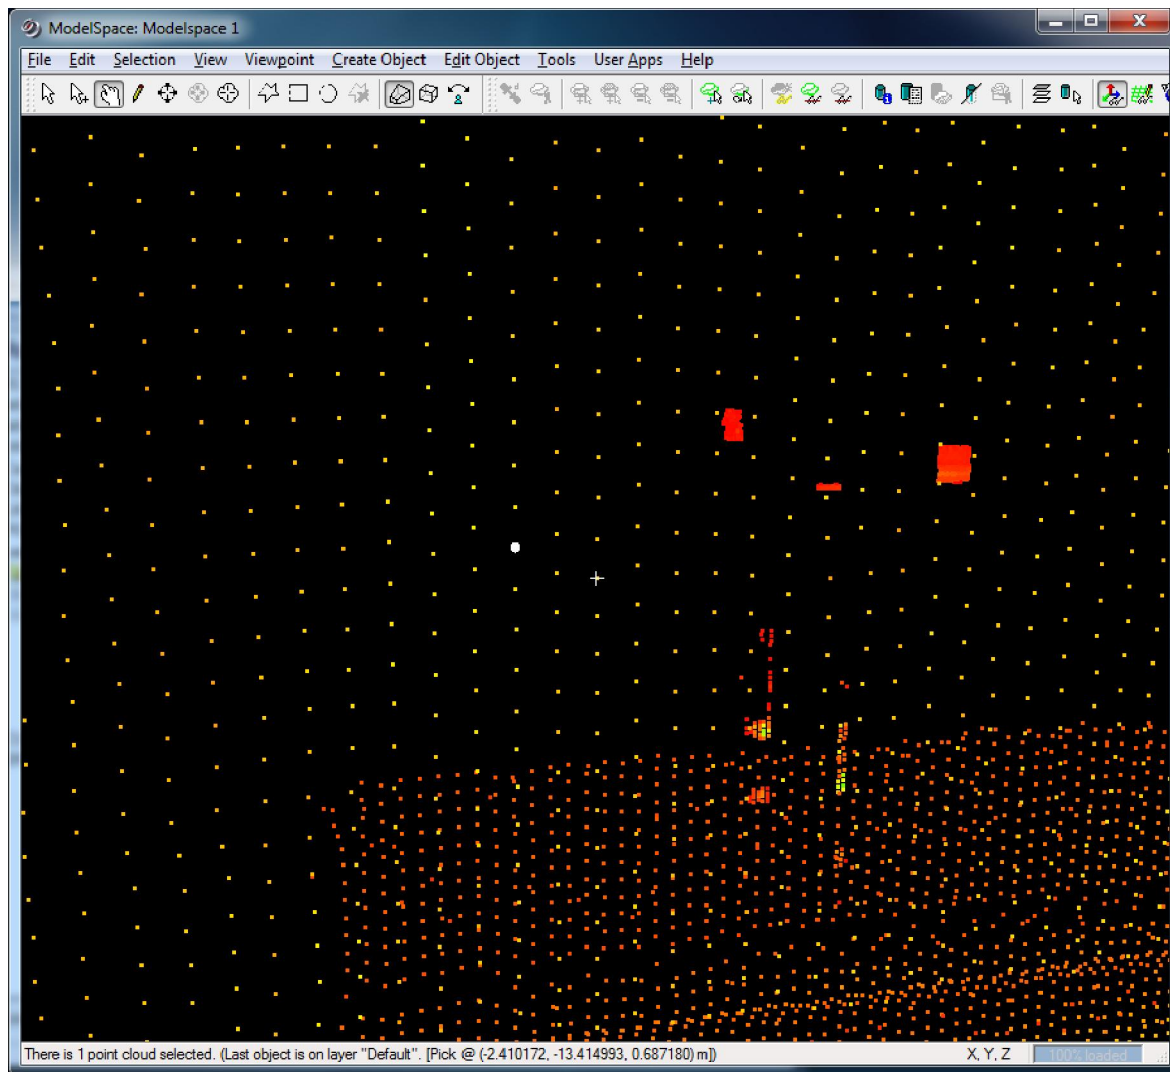


Figure 3: In most processing software you can zoom into the scan points as close as you like, and they remain points with a very long number of digits after the decimal place. This has nothing to do with the uncertainty, accuracy or the physical process of creating that very precise number, which generally involves a laser beam with a particular physical size, length, and directional uncertainty. The spot of the laser is far larger than the dots shown, and often, the scan density is so high that the ellipsoids of uncertainty for each "point" overlaps completely the other "point" ellipsoids. The value of such scan density is questionable indeed.

In addition, real world problems involve comparing measurements taken at different times, environmental conditions, and likely with completely different instruments. The holding tank might be sitting in a factory quite distant from the window, and a practical person would not ship it to the job site if it was not going to fit through the window. A practical person would order a different tank. This means that when we talk about the size of the tank, a meter must mean the same thing for the instrument measuring the tank and the other instrument measuring the

window. This might seem like an arcane point, but it is important enough for there to be a term for this – both instruments measurements uncertainty must be “traceable” to the same standard.

Knowing the uncertainty of a measurement – sometimes called the error rate – is the distinguishing characteristic between simple observations and measurements with a characterized uncertainty. In a court of law, a collection of data without a known error rate would not be considered evidence, because it fails the Daubert test of possessing a known error rate. Such data could only be considered “demonstrative”, not “evidentiary”. Leica Geosystems data sheets provide all of the information needed to compute the uncertainty of the data produced by the scanner, because each instrument is observed over the working envelope in a factory calibration. It is the process of calibration, described in this document, that allows Leica to stand behind the P and C series units as survey grade laser scanners.

Key Points:

1. A Measurement needs bounded uncertainty for many practical uses.
2. Enormous detail does not replace a bounded uncertainty
3. Bounded uncertainty requires traceability to standards
4. Data without bounded uncertainty is not considered evidence in legal proceedings

Resolution, Precision and Calibrations

Data sheets for laser scanners sometimes talk about resolution (range and angular) and precision, while these terms are for the most part missing from the Leica data sheets. What is resolution? If you purchase a ruler from the store, it would have tick marks dividing the length. The smallest spacing of the tick marks is the resolution of the ruler. What is precision? Suppose you replace this ruler with a laser range finder. If you measure, or rather call it observe, the same object over and over again, each observation would not be exactly the same, but they would all be very similar. The similarity can be characterized many ways – for example, by a standard deviation of the group of values (observations), or by an absolute bound on the values – and such descriptions of how tightly these repeated observations are clustered describe the precision of a measurement.




Figure 4: Precision and Resolution of a measurement. The hash marks on a rule provide a good example of resolution, in this case, in both metric and imperial units.

Precision and resolution do not tell the entire story, however, about uncertainty. Uncertainty is generally considered to be composed of both random error and systematic error. If you take many measurements with the steel ruler of the same object, you might expect to get some random error approximately the size of the spacing of the hash marks of the ruler. Averaging these multiple measurements could increase the effective resolution of the ruler, if the individual measurements were done properly (super resolution). However, if the hash marks in the rules are spaced too close together, every measurement of an object will be very similar, perhaps to the resolution of the ruler, but they will all be wrong by ratio of the actual distance between the hash marks and the expected value of the hash marks. This is an example of what is known as a systematic error, which for precise measurement devices is characterized by the Parts Per Million

(PPM) error and perhaps an offset. No amount of averaging can remove systematic error, a point which is often overlooked when seeing the vast amount of data a laser scanner has produced. Because the data is all collected by the same device, the information can be completely self-consistent and be consistently wrong. Having more data is very useful for reducing the impact of random error, but of less use with systematic error.

The full story with this steel ruler has another important element. Steel rulers can be calibrated. You can have this ruler checked against a NIST traceable artifact and have a conversion factor for the measurements that incorporates the systematic error that is literally built in. Unfortunately rulers do change length with temperature. Fortunately, they are generally short enough that the change in length, at approximately 11 part per million per degree C, is well within the measurement resolution for a reasonable range of temperatures. Measuring longer distances is a little more hazardous, and if angles are involved, the story is quite involved.

Precision and resolution, while interesting characteristics of an instrument, are of also limited useful value because while related to accuracy, they do not tell you anything about the accuracy of the numbers produced. Certainly if an observation has a fine resolution and is very repeatable, then one could take a set of measurements of objects with known characteristics, with respect to some standard, and construct a table or formula which converts the measurements produced by the devices into the traceable values and characterize how well the formula fits. The converted values now are measurements with a known uncertainty, or error rate.

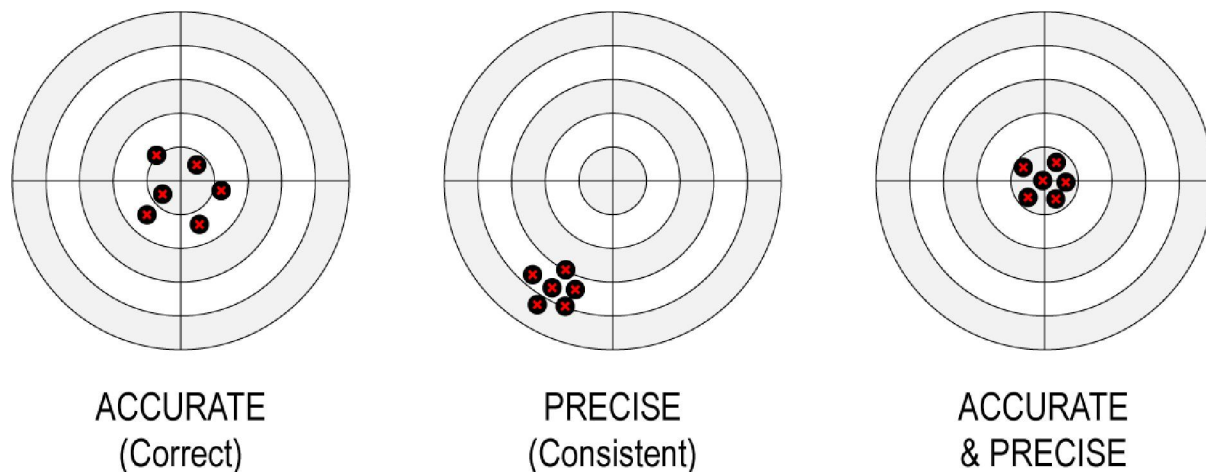


Figure 5: An accurate measurement delivers measurements that are close to what is referred to as ground truth, that is, the value with respect to standard units. A precise instrument returns consistent results and such consistency means that with a proper calibration, the returned values can be both accurate and precise, as shown on the right.

The process of taking a set of observations of known objects and building a map between the observations and the traceable values is referred to as a "Calibration". A calibrated instrument produces measurements with a known error rate provided it is used in conditions that are significantly the same as the calibration. The extra step of calibration and also of constructing an instrument so that it is stable enough for the calibration to be meaningful adds considerable expense to a device, but there is no particular choice in the matter if you are interested in taking measurements with a bounded error rate.



Figure 6: The traceable thermometer on the left is far more expensive than the Acurite® thermometer on the right. The name of the thermometer aside, the uncertainty of the readings on the right are not known explicitly given, because they are likely not known. For the purpose of determining if you need to wear a jacket or not, the uncertainty is not very important.

The precision and resolution of an instrument (and stability) are very important for calibration, but calibration is an impractical burden to impose upon customers using a laser scanner and hence such numbers are for the most part impractical. Because figures on precision and resolution are usually much smaller than the accuracy, they are also rather misleading. For this reason, resolution will not typically be found on Leica data sheets, and precision manifests as range noise.

Key Points:

1. Resolution and Precision indicate potential accuracy, but are not accuracy.
2. The uncertainty (accuracy) in a measurement consists of both a random element (precision) and a systematic element.
3. Averaging does not reduce a systematic error.
4. Self-consistent data can contain systematic error.
5. Calibration is the process by which an instrument uncertainty is controlled and characterized over typical usage conditions, including temperature.

Angular Accuracy

Most of the examples provided in this section talk only about range accuracy, as range accuracy is a very relatable concept. The angular accuracy of a laser scanner, however, is also needed to characterize the uncertainty of the measurements provided. Usually there is little random noise in an angular measurement, so systematic error dominates the specification. Why do you need the angular accuracy? Suppose you needed to determine the space between two buildings. The laser scanner would be placed somewhere near these two buildings and the scene would be scanned. In the office, one would either select points, accepting the random error due to range noise in the uncertainty, or perhaps solve for particular geometric features of the buildings such as walls or corners using many points. Even using many points, systematic range and angular error would remain.

Let's compute the uncertainty in the case with extracted corners, using the systematic error bounds given in the P-series scanner data sheet. Given each corner is, say, 35 meters away from

the scanner, then the lateral angular uncertainty for data produced by a Leica P-series scanner on each corner is 8 arc seconds * (4.85e-6 radians/arc second) * 35 meters = 1.4 millimeters, plus or minus. The range uncertainty is 35 meters * 10 ppm + 1.2 mm = 1.5 millimeters, about the same. This represents a plus or minus figure, so when comparing two measurements, one could have an error off to one side, the other measurement, on the complete opposite side. If we form a box with depth is 3 mm and whose sides are 2.8 mm vertically and horizontally, then we have a worst case error of the diagonal length through the box, given by $\sqrt{2.8^2 + 2.8^2 + 3^2} = 5$ mm, between two measurements. That is about 1/4 of an inch, not plus or minus, as I doubled sides.




Figure 7: Measuring the space between two buildings. The diagonal worst case error is swiftly computed from the specifications and the measurement situation. Without the angular accuracy specification there is no way to compute the error bound.

How could you possibly get an angle wrong? Imagine now you decided to measure the distance between the buildings again. This time, however, you rotated the scanner on the tripod. One might expect to get exactly the same angle between the two building corners, provided the buildings did not move. By taking this additional measurement, and perhaps many more such measurements with different rotations of the scanner, you are performing what is known as a "Circle Test", from the STN ISO 8322 standard, because unfortunately whether or not you get the same angular distance between the corners depends strongly on how the optical encoders are constructed and, yes, calibrated.

Angular measurements are made in principle no differently than reading marks on a steel ruler, except that the marks are in a circle. There could be flaws in the fabrication of the disk, the printing of the marks, or the interpolation of the spaces between the marks. These are the more obvious kinds of mistakes that can be made. An example of a less obvious error is how well centered the disk is on the shaft. If the disk is mounted slightly to one side, the spacing between the marks on the disk appears wider and later more narrow than they are meant to be. This means the angle measured between two points for a single read head will depend on the relative orientation of the instrument, as shown in Figure 8. A poorly centered optical encoder would fail

the ISO circle test, because the exact same angular distance measured across different parts of the disk appear to be two different angles θ_1 and θ_2 .

How to correct for this and other similar errors? The other sources of error are as you might imagine legion. Turns out, balls in ball bearings are not all exactly the same size. And this matters. How the disk is attached to the hub, usually two different materials with different thermal expansion coefficients, is a detail that matters, greatly.

While careful design and assembly is essential, at the levels of accuracy Leica delivers, accuracies measured in arc seconds, additional angular measurements are needed. For the P and C series scanner, each angle measured is actually the composite of four different read heads angle measurements, which are instantiated as small cameras. These redundant readings can automatically detect and remove centration error and many bearing effects. So each point produced by a Leica instrument has not only a precision calibrated range finder but also 8 small cameras observing the M-codes on the encoder disks. This is why Circle tests with Leica scanners produce errors far smaller than what is listed on the specification sheet.

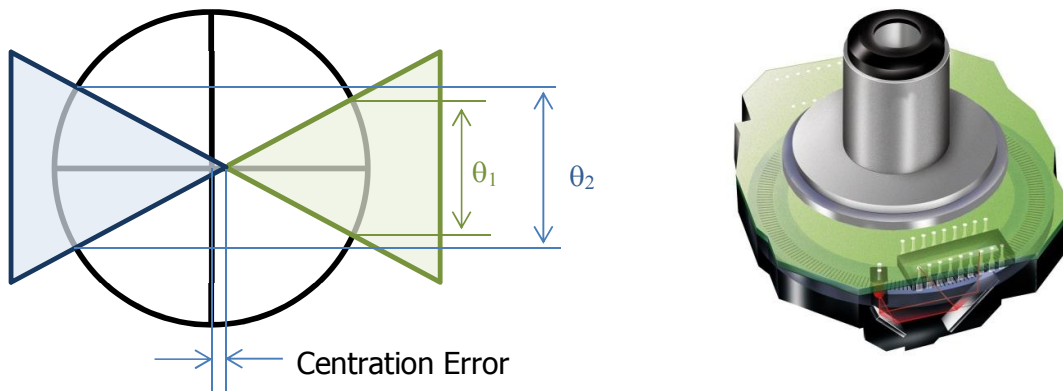


Figure 8: Optical encoder disk centration error causes significant error. Leica manufactures each of the optical encoders and mounts them using a special fixture with a microscope and gentle striker. The unit is rotated slowly before the glue dries and adjusted by hand with the impact striker to minimize centration error.

The units in which Leica describes the angular accuracy of the P series and C series instruments is arc seconds, a unit so small it is perhaps difficult to grasp. Take the angular accuracy of the tilt measurement – that is, 1.5 arc seconds. That is to say, the tilt sensor in the P40 can locate the direction of gravity to 1.5 arc seconds, which is $1.5 * 4.85e-6$ radians = $7.275e-6$ radians. What does that possibly mean? You can think of an angle as the size of something at a distance. For example, a US dime is 17.91 mm in diameter. There are 1609.34 meters in a mile. For the dime to subtend 1.5 arc seconds, it would have to be placed about 1.5 miles away.




Figure 9: A dime placed 1.5 miles away subtends approximately 1.5 arc seconds.

Why does Leica use such a small unit to specific angular accuracy? Because arc seconds is the appropriate unit for instruments that have more than a few meters worth of useful range. Angular accuracy specifications are multiplied by range – range acts as a lever magnifying even very small systematic errors. Surveyors are interested in making measurements with uncertainties on the order of some fraction of an inch, or in metric, of a few millimeters. Having a specification in the arc second range, which is 4.85 micro-radians, means that when it is multiplied by typical ranges as in the example above, the error rate is controlled to the millimeter level.

Key Points

1. Both range and angular uncertainties are needed to compute the uncertainty (accuracy) of data from a laser scanner.
2. Angular error rates are multiplied by range, so must be exceptionally small.
3. Redundant measurements are used in Leica P and C series instruments to address errors that arise from manufacturing tolerances.
4. In order to meet the strict angular accuracy specifications, the angular encoder and tilt sensors used in Leica P and C series scanners are designed and built by Leica.

3. Leica Scanner Manufacturing and Calibration

This section describes some of the process steps taken to assemble a P and C series laser scanner on the production floor in Widnau, Switzerland. Similar steps are taken at service centers, provided they have the capital equipment and training necessary to perform these steps properly.



Figure 10: A set of P40 Laser Scanners awaiting calibration on the Manufacturing Floor in Widnau, Switzerland.

The processes might at first glance seem over-elaborate and lengthy, and their cost inspires many extraordinarily heated debates with Leica, so you can know that processes and parts that are part of the creation of each P and C series scanner are there because they absolutely must be. Even the colors are carefully chosen. We know this because at Leica we check during the design process and on every single instrument we produce. A very tight specification is a fine goal but the most important aspect of delivering a quality measurement is knowing the uncertainty, even if the measurement instrument manufacture does not guarantee at the arc seconds or PPM level, they can, at least, guarantee something at some level.

Assembly of a P and C Series Laser Scanner

When a Leica laser scanner of any sort is assembled, care has to be taken about the order in which the parts are attached, the formulation and condition of the adhesives used, and the torque of each screw. The P and C series scanners have each of these specified and technicians use specially and regularly calibrated torque wrenches during assembly. These steps ensure that the each instrument manufactured by Leica is consistent within expected tolerances. Unfortunately the specifications that the scanners are well beyond human/hand adjustment and assembly ability to reach, so special processes and tools have been developed to ensure these quality metrics are achieved.



Figure 11: An array of calibrated torque wrenches on the P and C series manufacturing floor at an assembly station, being checked against a calibrated torque reference.

A calibration quality is clearly limited by the precision of the instrument which is calibrated. If that cluster of point moves from day to day, the error bound must be grown to cover all possible movements of the cluster, not reduced to the size of the cluster itself. A great deal of effort goes into designing the interface between the side cover in the frame, for example, and into designing the process of tightening the bolts that attach the side cover to the frame, and ensuring the technicians follow these processes to the letter. Even so, when a series of bolts are tightened, the side cover (and every other part in the scanner) will have some residual internal stress between it and the frame. Since you are never going to move the bolt (an adhesive blend is used on the bolts), this seems like an unimportant point, but not for Leica.

In the world of arc seconds, which, because of the distances involved in laser scanning matters, this residual stress is a detail which matters. Normal handling processes will vibrate and thermally shock the machine, and over time, these stresses will be worked out of the instrument. There is neutral assembly state where the machine will end up in the normal course of events. You cannot so much as attach a side cover and start there. In order for the calibration to be meaningful, each scanner after assembly must be moved to that neutral assembly state.



Figure 12: Environmental chamber for thermally cycling P and C series laser scanners. This station is called ESS, or Environmental Stress Screening.

For this reason, each scanner after assembly is environmentally stressed, visiting the specified temperature ranges and beyond, over a period of several days in environmental chambers. In addition, the scanner is operated inside an environmental chamber before calibration over the operating range to verify operation. Why so much time is spent on this process? So much time is spent because it is necessary and we know this because we have checked over a variety of instruments.

Angular Calibration

Angular calibration, that is, developing a map between the angles measured at the axis and the direction of the laser beam, is required for each laser scanner produced, and for each laser scanner, the behavior over the full range of operating temperatures must be observed and corrected for. One would imagine that with self-compensating quadruple read head angular encoders, on each axis, with an assembled accuracy of approximately an arc second, that additional steps would not be necessary. Already you might suspect the answer is No. There are additional details that matter.

The angles are a great starting place, to be sure. The direction of the laser beam, in reality, depends on a series of optical and mechanical components of which the measured angles are just

a part. The axes themselves, even relative to each other are not exactly intersecting or perpendicular. The laser beam, formed by a source and a collection of transmit optics, is not directly in line or aligned with the elevation axis. Nor does the elevation mirror deflect this laser beam precisely perpendicular. Nor is the transmit window to perfectly flat as to not slightly deflect the laser beam. With angles in particular, very small movements of any of these parts is sufficient to cause very small angular deflections of the laser beam. So if we are concerned with "arc seconds", with is on the order of 5 micro-radians, deflections of mechanical parts on the order of "microns" become interesting and important to observe.



Figure 13: A collection of P40 Scanners waiting for their turn on the ATEC Calibration Rig. To the human eye, they are perfectly identical, but the human eye is not designed to discern microns. To the Antipodal Telescope Environmental Chamber (ATEC), they are all snowflakes.

When considering microns, recall that human hair ranges between approximately 50 and 150 microns in diameter. Unavoidably, microns, or fractions of a human hair, matter, for surveying grade laser scanners. This poses yet another challenge for angular calibration –since the unit is made up of different materials such as steel, aluminum, and engineered plastics, and these assemblies have different thermal expansion coefficients, the exact positions and relative orientations of the individual components will move as the instrument changes temperature. The shift is on the order of microns and hence not visible to the eye, but it does have an impact on the direction of the laser beam in the arc second regime. Consequently we must ensure the movement is repeatable (within a known tolerance), know how the instrument changes over temperature, and provide enough internal sensing so the instrument is aware and can correctly compensate for this motion. A proper (and practical for the customer) calibration of a survey

grade laser scanner observes the scanner over the span of useful environmental conditions, particularly temperatures. Yet again, environmental chambers are needed, this time, with windows.




Figure 14: A Leica Scanstation P40 being loaded into the ATEC Rig.

The antipodal telescope environmental chamber, or ATEC, is an instantiation of a patented (Patent US7643135 "Telescope Based Calibration of a three dimensional optical scanner") process for executing the full temperature angular calibration of P and C series laser scanners. Of course Leica have developed many different angular calibration processes for their products over the years – the ATEC simply represents the current technological pinnacle of this process. What distinguishes the ATEC from previous efforts using laser trackers or high accuracy surveying total stations is the use of auto-collimating telescopes, pointed at each other (that is, mounted in anti-podes) to locate the laser beam. The use of telescope might remind surveyors of total station calibration stations. This is not accidental.

A collimating telescope is focused very far away, ideally at infinity. When a laser beam is directed into such a telescope, the arrangement of lenses acts like an optical lever, converting very small angular displacements of the laser beam into distance displacement of the beam at the focal point. Lateral displacements of the beam a similarly attenuated. Mechanical levers act in a similar way, as understood since the time of Archimedes, with the ratio of the lever lengths from the fulcrum acting as a multiplier of the displacements. Similarly, the focal length and the focal points becomes a multiplier. This means a telescope focused very far away, a collimator, can detect the minutest change in the angular direction of a laser beam. All that is needed is an

image sensor array to digitize this change. The ATEC telescopes have a sub-arc second ability to resolve the direction of the laser beam.




Figure 15: Data from the ATEC telescopes after a laser scan. The center of each of the laser spots is used to determine the direction of the laser beam.

The laser beam direction found by the telescopes is relative to the telescopes themselves, even though they come attached to particular scanner angle measurements. Knowing the relative position of the telescopes with respect to the scanner is needed to solve for the calibration. This is the point at which the large pieces of granite encircling the environmental chamber and holding the various telescopes come into play. The telescope mounts are very stable and are observed in two faces by the scanner. This would be a sufficient set of observations to solve for both the scanner and telescope relative positions and kinematics if the scanner had the kinematics of a surveying transept, but it possesses a more complex behavior because of the mounting of the laser. Hence the telescopes are mounted so that they can observe each other, that is, become self-referencing, when the scanner is not between them, as in Figure 16.




Figure 16: Antipodal placement of the ATEC telescopes allows the telescopes to observe their relative placement, when the scanner is removed. This is a patented twist on the idea of auto-

collimation techniques that critically allows for arc second level observation and calibration of laser scanners.

The angular calibration process in the ATEC still takes the better part of an entire day, however, because the instrument must be observed at different temperatures, and in order to emulate customer use, the instrument has to sit for some considerable time in the deep cold and for similarly in the blazing heat for a complete picture of the scanner behavior to be made. Each machine that is manufactured by Leica has this behavior recorded and stored internally for later use in correcting the data.

These complexities are managed by the scanner in the field, so a customer does not need to know anything about these processes or manage them, except to bring the unit in for service at the periodically recommended intervals. Using this internally stored data and the temperature sensors distributed throughout the interior of the instruments, the scanner automatically adjusts produced data continuously as it is being produced.

Tilt Calibration

The tilt sensor mounted in the P and C series scanners is the same unit used in high accuracy Leica total stations. The principle of operation is shown in Figure 17. At the heart of the tilt compensator is a single piece glass bowl containing special oil forming an artificial horizon. Tilt sensing is accomplished by bouncing a light pattern off of this special oil surface. Double reflection of a particularly designed target at the surface of an oil level, imaged on a small line camera, allows continuous and simultaneous detection of tilt in transversal and longitudinal direction




Figure 17: Leica tilt sensors are also based on using a small line camera, much like the angular encoders. Much like the angular encoders, their sensitivity is very high, less than 1 arc second.

The bowl containing the artificial horizon is made from a single piece of engineered material, so as it changes temperature, as long as the mounting is stress free, the bowl itself remains dimensionally stable. The mounting of course will move with temperature and is measured in the ATEC, additionally, the ATEC contains a tilt table and reference tilt sensor. This allows for the complete calibration with respect to this reference sensor of the tilt sensor while the unit is inside the ATEC.

A separate calibration stand exists for the tilt sensor, shown in Figure 18. The tilt sensor room temperature behavior is sometimes identified on this tilt calibration stand. The scanner is placed in the stand and the laser beam pointed at a specially modified reference unit, shown to the left of the scanner in the figure. The reference unit is mounted with a camera able to view the laser beam; this allows the software to tie the two coordinate systems (scanner, and reference) together. The tilt table can then perform a series of motions and compare the two sensors outputs and in a straight forward way determine the adjustment parameters of the Scan Station C10. In addition, if any tilt sensor is out of specifications, it will be rejected on this stand. Temperature

behavior of the tilt sensor is determined in the environmental chamber and again, units with out of specified behavior are identified and removed.



Figure 18: Tilt calibration station for the Leica Scanstation C10. The tilt sensor as it is mounted in the scanner is also observed over temperature.

Camera Calibration

The third calibration fixture is for calibrating the internal camera. The internal camera can take pictures over the entire field of view of the scanner, and these images are assembled into mosaics which are then applied to the point clouds. Hence the objective of the camera calibration is to ensure a good match between measured LIDAR points and pixels on the camera. For this reason, the camera calibration is performed after LIDAR calibration. Because the camera mounting is stable enough, and because the pixels are essentially huge, subtending on the order of 150 micro-radians, a room temperature calibration is sufficient for the camera to keep the pixels aligned to the scans within expectations.

A special set of targets (see Figure 19) has been created to calibrate the camera. These targets are white disks on a black background, and are automatically located by the scanner and by the camera. The camera will take photographs of these targets from many different positions, both from the front and from the back. The calibration control software, again in MATLAB and C++, identifies the target centers both in the scans and in the images, and solves for the matching correspondences, as shown in Figure 9. Once a large set of corresponding matches is assembled, a bundle adjustment again is run to solve for the camera parameters. Parameters

include the mounting of the camera inside the scanner as well as the distortion to the image caused by the lenses.




Figure 19: A view of the camera calibration stand. Two sets of targets are placed in front of the scanner and are both imaged and scanned.




Figure 20: Images from a Leica C10 scanner as viewed in MATLAB. MATLAB has drawn on the image guess locations of the target centers using an approximate camera model. These guess locations are used to solve the target correspondence problem, and are generally off by a gross amount such as 15 or 20 pixels. This is close enough to prepare the matches.



Figure 21: Tone map of the Atrium at the San Ramon, California office of Leica Geosystems. This image is a spherical projection of a panorama assembled from 260 HDR images taken with a Scanstation P40.

4. Scanner Testing and Verification

Though born in essentially laboratory conditions, the Leica P and C series scanners is not designed nor intended for use in such benign circumstances. Out in the real world, where our customers need to collect data, the weather can be extremely cold, or hot, or dirty. The transport situation might not be very smooth, and the support infrastructure at the job site might be less than one was grown accustomed to. After all, bringing a surveyor to a job site is often a first step in creating the very infrastructure we take for granted. The instruments are designed to operate in a variety of environments and collect data at full accuracy, and to be handled and transported following the Leica field handling standards. The instrument is an accurate, precise arc second quality device and should be treated with reasonable care, but that caution cannot pose undue burdens on the customer's use of the equipment. As much as I would like to, it is not to be.



Figure 22: Infrastructure at the job sit might be lacking, and often is.

The P series instruments are qualified over -20 C to +50 C, and are IP54. What is IP54 mean? IP stands for "Ingress Protection", meaning the level at which the instrument must not let the outside environment in. We do design the unit to meet these kinds of environmental, handling, vibration, shock, and ingress protection standards from the outset. As you can imagine, however, the real world is extremely creative and such standards can only verified through testing: testing, redesign, and repeated testing.





Figure 23: A mosaic of testing events for the C and P series scanners. There is something such as a thing as a regulation shower head, among other things.

For the author, the most dreaded tests involve shocks. By shocks, we mean dropping the instrument onto unforgiving surfaces. Sometimes the scanner is in the case, sometimes, the scanner is not in the case. There are also specially designed tables to shake the instrument in many directions with variable amounts of force. Scanners must retain full accuracy through a variety of events – just placing the scanner in the tripod roughly can deliver a considerable shock, let alone placing the unit in the care of baggage handlers at an airport, even if it is in a case. The trunk of car in the outback of Australia can get surprisingly hot. The scanner and its calibration must handle these reasonable events and deliver accurate measurements at the job site.




Figure 24: Talcum powder particles are very small and surprisingly invasive, particularly if they have a full day to work their way into the interior of a laser scanner.

In previous versions of this document, I have detailed what is referred to as a mini-QA, or quality assurance. Leica has also constructed a not so mini QA, used for the verification of laser scanners in general. The scale of this testing field is very large, spanning hundreds of meters, and because it is the large size you cannot depend on buildings remaining the same place from day to day or even over the course of one day. Even large, stable concrete buildings shown in the figure essentially breathe (move) with the rising of the sun and its setting, for example, this motion is observable with Leica equipment.



Figure 25: The mounting station for outdoor scanner verification, with a TS30.

In the outdoor quality verification test site, a surveying total station is placed next to the laser scanner and collects, in two faces, the coordinates of the targets observed by the scanner at as close to the same time as possible. Clearly the targets are special targets, which dimensionally stable surfaces and mounted with both targets that can be precisely located by the scanner and by the total station.





Figure 26: Image mosaic of a variety of jointly observable target sets, with weather hoods.



In each corner of a target plate is a precisely mounted prism target for the total station, which then can locate the plate geometry, as well as different kinds of scanner targets, which the scanner can scan and then locate, in both faces. As with the indoor mini-QA testing field, a large number of these targets are mounted to a variety of surfaces and buildings throughout the field of view of the laser scanner and total station.

The calibrated TS30 provides a reference from which we can check the data produced by the scanner. Just as a point of reference, the TS30 has an angular accuracy of $\frac{1}{2}$ arc second, which is better by more than a factor of 10 over the scanner, just barely. Only Leica, at this time, makes total stations with that level of accuracy.




Figure 27: Derived objects, such as building corners, are located within the scan data separately for each face of the scan and compared against the location of the corner found by the TS30. The plates are built and measured in advance of mounting so the relationship between the total station prisms and the corner is known in advance.

And for the mathematically minded readers, yes, we do have targets placed at high elevations taken in two faces. Such targets are needed to verify accuracy.

5. Summary

Leica Geosystems is in the business of selling measuring tools and takes great care to ensure that every product that is shipped meets or exceeds our specifications. These specifications include error rates on both the range and angular components of the data produced; both are needed in order to characterize the uncertainty or accuracy of the point cloud. The specifications are tight but necessary for devices with ranges greater than a few meters. A large number of details matter when producing an instrument at characterized by arc seconds or error. Details matter with survey grade laser scanners. An operator of a laser scanner does not need to know or measure these details, because Leica stands behind our specification sheet with confidence borne of the experience that comes with producing high quality surveying products for nearly 100 years.