Characterizing Accuracy of Trimble Pathfinder Mapping Receivers

Trimble Navigation Limited
Surveying and Mapping Systems
Commercial Systems Group

This paper was presented at the second annual Trimble Surveying and Mapping User Conference and Exposition on October 2nd and 3rd, in San Jose, California, 1996. The performance results presented here are derived from data collected during the spring of 1996 at and near Trimble World Headquarters in Sunnyvale, California. These results represent nearly optimum performance of current Trimble Mapping receivers and processing software, since the data collection sites were chosen to have low multipath and a clear view of the sky. Typically, these sites were on the roof tops of buildings. Your performance results will vary depending on the site conditions where you collect data and the atmospheric conditions as well.

CONTENTS

Introduction

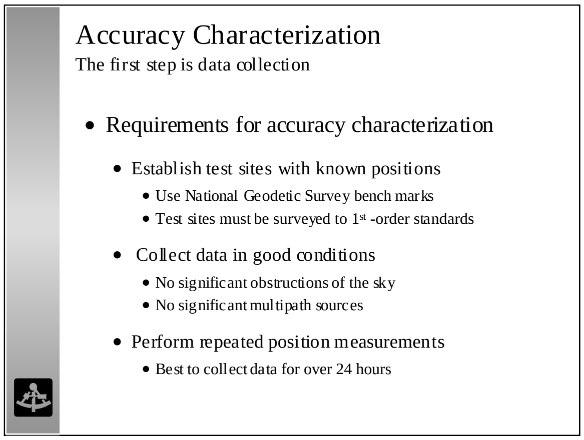

The first step in any accuracy characterization experiment is to collect some GPS data. When we talk about the accuracy of positions, we are making a comparison between our experimental results and some reference position whose coordinates we know to be true. In order to make this comparison between experimental results and true position, we must establish test sites with known positions. National Geodetic Survey bench marks[1], whose coordinates are published, can make good test sites. These bench marks usually have coordinates that have been established by 1st-order surveys. In any case, you should know the true coordinates of your test site with better accuracy than the specified accuracy of the equipment that you are planning to characterize. Additionally, you should consider the GPS observing conditions of your test sites. If your goal is to assess the optimum performance of the equipment, you should pick test sites with good GPS observing conditions. Good observing conditions means no obstructions of the sky and very little multipath. On the other hand, you may want to characterize performance at a site that does not have good observing conditions. Just keep in mind that performance varies depending on the GPS observing conditions. The performance results presented in this paper were derived from GPS data collected in good site conditions, with a stationary antenna. Once you have established a test site, with known coordinates, then collect repeated position measurements at that site. It is best to collect data for periods exceeding 24 hours. If you exceed 24 hours, then any variability in the position measurements due to differing satellite constellations will be captured, since all possible constellations, during that period, will have been sampled.

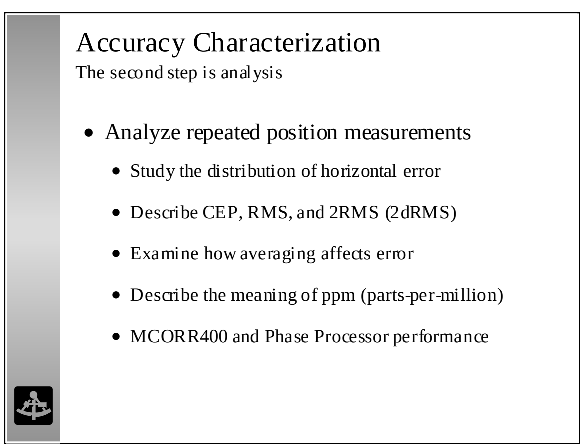

The second step in our characterization experiment is the analysis of the repeated position measurements. To analyze our repeated position measurements we will study the distribution of the horizontal error. This distribution will be examined with respect to the common reference terms for accuracy such as the Circular Error Probable (CEP), the Root Mean Square error (RMS), and twice the root mean square error, sometimes called the 2dRMS. Next, we will examine how averaging affects error. Then we will discuss the meaning of ppm or parts-per-million and describe how to interpret the accuracy specifications on product data sheets. Finally, the results presented will characterize the performance of the MCORR400

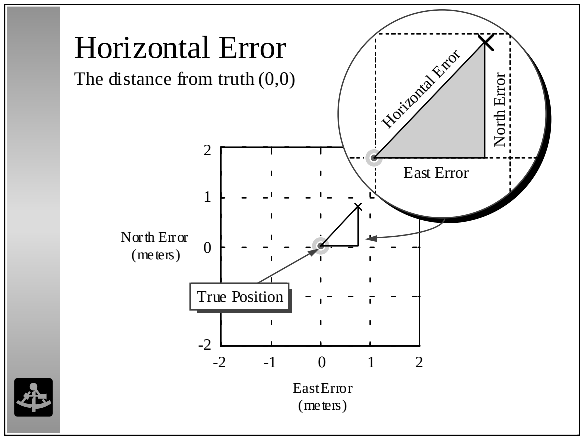

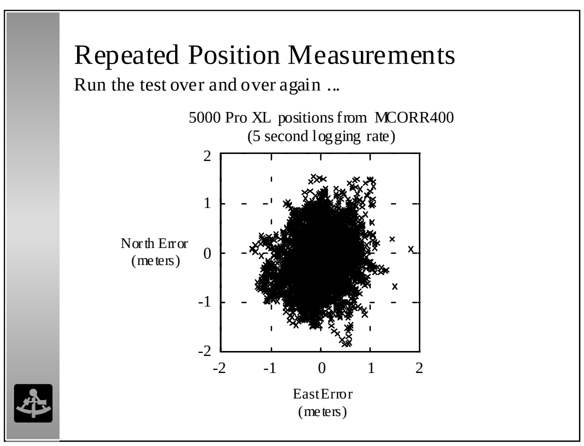

Base receiver:
Pro XL with firmware version 3.08
Rover receiver:
Pro XL with firmware version 3.08

and Phase Processor software with all of the current Trimble Pathfinder mapping receivers - GeoExplorer, Pro XL, Pro XR

Just what exactly is horizontal error? Horizontal error is the horizontal distance of a GPS position from the known true position of the test site. This graph shows a plot of north error, on the vertical axis, versus east error, on the horizontal axis. Our test site is at the center of the graph at the point (0,0). For any GPS position that we measure, shown by the 'x' on the graph, we can define a north error, an east error, and the horizontal error. The horizontal error is simply related to the north and east errors by the Pythagorean theorem:

$$horizontal\ error = \sqrt{(east\ error)^2 + (north\ error)^2}$$

This calculation is only valid when we know the true position of our test site. Knowing the true position, we can subtract the coordinates of any position measurement from the true coordinates and then compute the east, north and horizontal errors. In this presentation, we are going to focus on the horizontal errors. Vertical errors for GPS code derived positions follow the usual rule-of-thumb and are about 2 times the horizontal error.

MCORR400 Discussion

Once we have decided how to compute the horizontal error of a particular GPS position measurement, we then repeat our measurements many, many, times. This graph of the north versus east errors of 5,000 GPS positions is derived from approximately 7 hours of positions collected with a Pro XL receiver at a 5 second logging rate. The positions have been differentially corrected by MCORR400 with base station data also from a Pro XL receiver. The base station was less than 1 kilometer away from the rover. Later in this talk, we will discuss the dependence of error on the separation distance between the rover and the base. These 5,000 positions are just a subset of the data used in our characterization experiments, where we analyzed continuous data collections of 4 to 6 days in length. All of the positions in this plot have a PDOP of 4 or less and were derived from constellations of 5 or more satellites. A satellite elevation mask of 15 degrees was used for all data collections. What this graph shows is that the horizontal accuracy of individual positions varies significantly over this 7 hour period. Some positions have very small errors and others have errors as large as 1.5 meters, or more. With this kind of plot it is very difficult to say very much about the accuracy. So many of the positions are overlapping each other it is difficult to see how many are near the origin - the true position of our test site. A better way to view these data is with a histogram of the horizontal error.

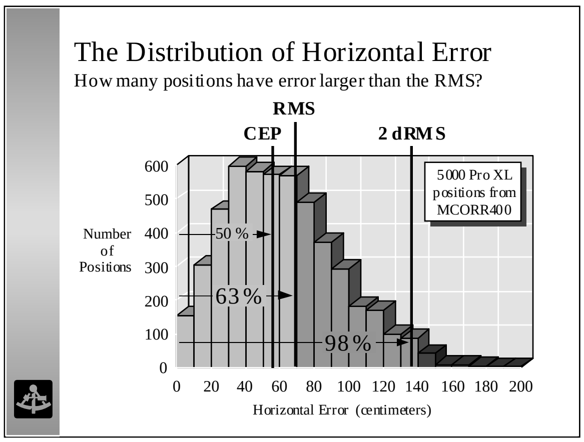

This graph shows a histogram of the horizontal errors for the same 5,000 Pro XL differentially corrected positions. Along the horizontal axis we have plotted the horizontal error in 10-centimeter-wide bins up to 200 centimeters, or 2 meters. The vertical axis shows the number of positions that fall into each bin of horizontal error. Most positions have a horizontal error near 50 centimeters, but it is clear that there are some with smaller and larger horizontal errors as well. Also shown with this histogram are 3 lines marking where the commonly used reference terms for accuracy, the CEP, RMS, and 2dRMS, fall within the distribution. For each of these markers, an arrow with a percentage value is shown. These values show the percentage of the 5,000 positions that have horizontal error equal to or less than the reference values. The CEP reference value bounds 50% of the positions while the RMS bounds approximately 63%. Twice the RMS, the 2dRMS, bounds approximately 98% of the positions. At every site where you make GPS measurements, there is a distribution of horizontal error similar to the one shown in this graph. If you make many repeated measurements, you can expect that approximately 37% of the positions will have a horizontal error greater than the RMS error. But, only a few percent of the measurements will have a horizontal error greater than twice the RMS. The percentage for the RMS depends on the shape of the distribution. In a number of simulated random data sets, we observed the RMS to range from 63 to 68%. Similarly, we observed the 2dRMS fell in the range 95-98%. The CEP had a relatively constant percentage. See the appendix for the RMS calculation.

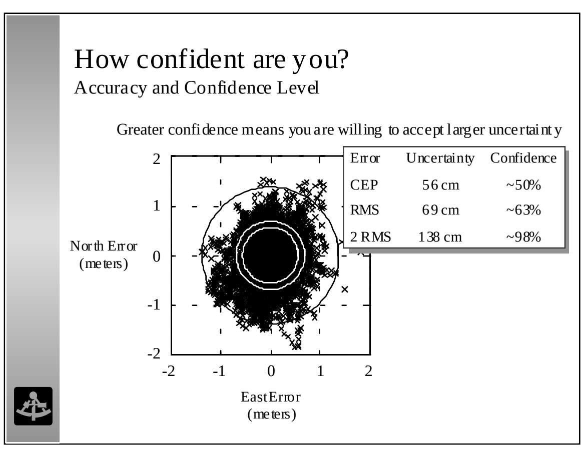

How confident are you that your position measurement has the error that you think it does? Confidence level refers directly to the same percentage values that we just discussed in regard to the reference values of the distribution of horizontal error. That is, the CEP represents 50% of the distribution so using the CEP as an indication of horizontal error for a particular position measurement means that we have 50% confidence that our position measurement is accurate to plus or minus 56 centimeters. Similarly, using the RMS means that we have 63-68% confidence that our position measurement is accurate to plus or minus 69 centimeters. And, if we use the 2dRMS to represent our uncertainty, we can say, with 95-98% confidence, that our position measurement is accurate to plus or minus 138 centimeters. To have greater confidence in your accuracy statements, you must be willing to accept a larger uncertainty.

We have examined the horizontal error associated with instantaneous differentially corrected position measurements and have discussed the meaning of the CEP, RMS, and 2dRMS in terms of the percentage of the distribution that they represent, and found that we can interpret those percentages as confidence levels. So, how does averaging affect horizontal error? Now, we will look at the effect of averaging GPS code positions.

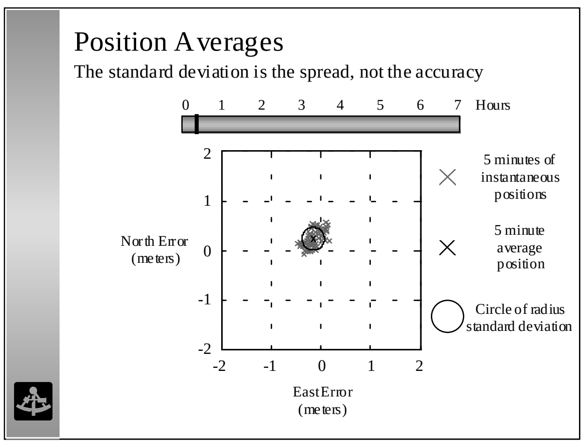

In the following sequence of graphs we are going to look at the same 7 hour Pro XL data. Each graph in the sequence is a 5 minute subset of the 7 hour data set. The time line above the graph shows the entire 7 hour period and the black bar is the 5 minute subset that is plotted on this graph. The light gray ‘x’s are the 5 minutes of instantaneous (5-second logging rate) positions, the dark gray ‘x’ is the average position for the 5 minute occupation, and the black circle is drawn around the average position with a radius equal to the standard deviation. As we progress through the sequence, the black bar on the time line will indicate which 5 minute period is currently plotted on the graph. Previous 5 minute periods will be shown with lighter gray ‘x’s. See the appendix for the standard deviation computation.

In the original presentation this sequence was composed of 25 graphs. In this printed version there are only 7.

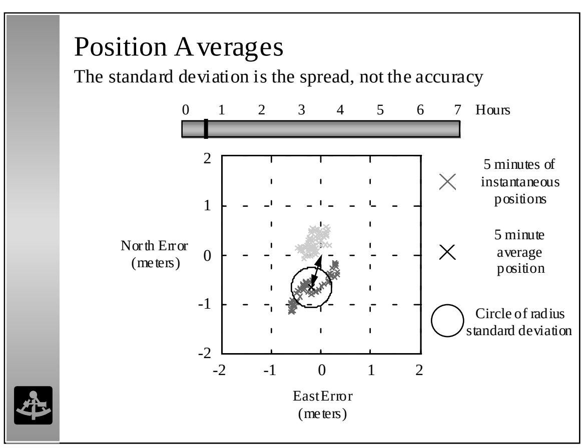

The second 5 minute occupation in the sequence. As we examine the 5 minute occupations, it is clear that positions from different time periods don’t always overlap. It is important to notice that the circle of radius equal to the standard deviation does not represent the accuracy of the average position, but does represent the spread of the positions during that period. For example in this 5 minute occupation, the average position is farther from the origin (the true position of the test site) than the standard deviation.

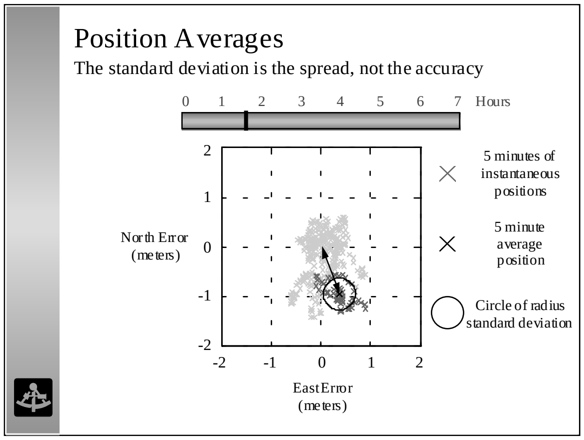

Be careful to note here that the standard deviation is much smaller during this 5 minute occupation than the average position's accuracy (distance from the origin).

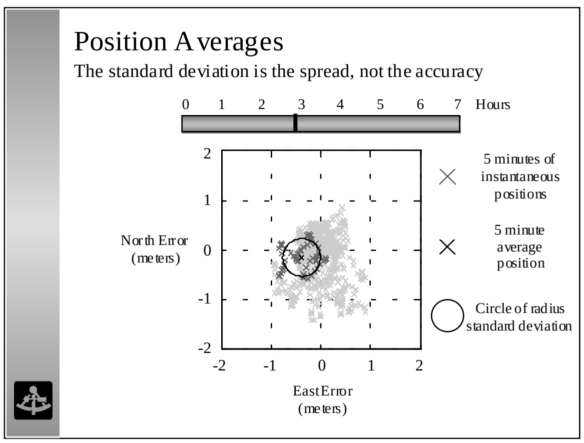

Some 5-minute periods have large standard deviations and some have smaller standard deviations.

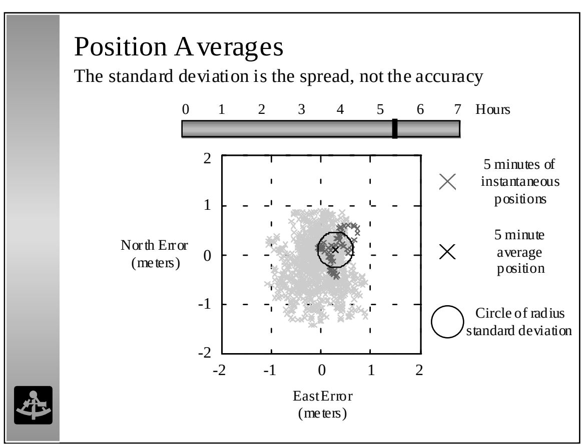

Here is a 5-minute occupation that has a standard deviation about equal to the accuracy of the 5-minute average position.

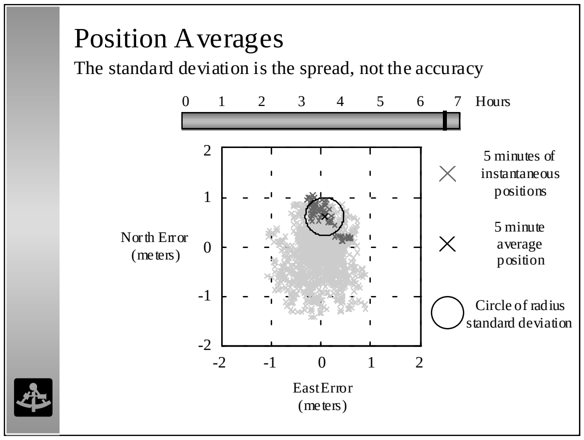

The last 5 minute occupation in the sequence.

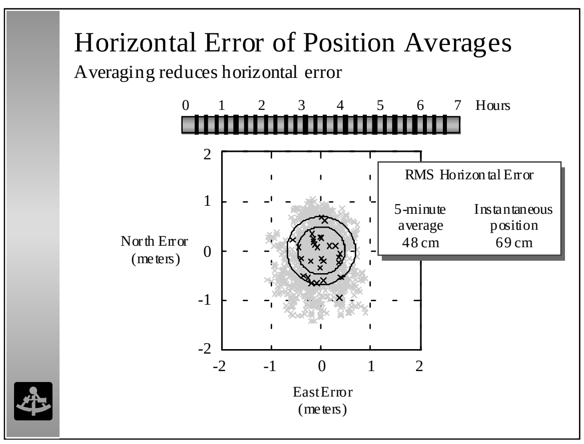

Finally, this graph shows all the 5-minute occupations together. There are two distributions shown on this plot. One distribution shows the instantaneous positions for all the 5-minute occupations that we studied (light gray ‘x’s). This distribution has an RMS horizontal error of 69 centimeters. The second distribution shows the 5-minute average positions (dark gray ‘x’s). The 5-minute averages have an RMS horizontal error of 48 centimeters. The two black circles are drawn with radii of 48 and 69 centimeters to illustrate the relative size of the RMS errors. So, averaging for 5 minutes reduces the horizontal error and increases accuracy. What about other occupation times?

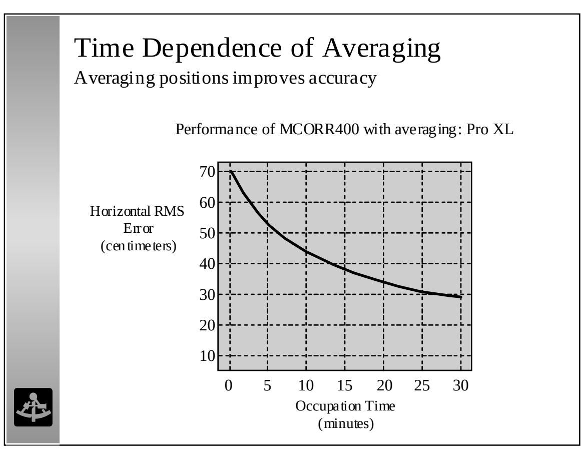

The time dependence of averaging indicates how accuracy improves with occupation time. The vertical axis is the RMS horizontal error in centimeters, and the horizontal axis shows increasing occupation times. On the left side of the plot at zero occupation time the graph shows the instantaneous position accuracy for the Pro XL, which is ~70 centimeters. Moving to the right on the graph, towards longer occupation times, results in lower RMS horizontal error. Thus, averaging positions over longer occupation times improves accuracy.

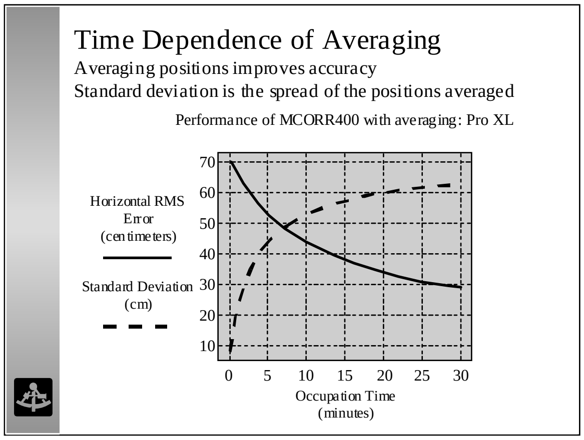

While averaging improves accuracy, the standard deviation of the positions averaged does not indicate the accuracy of the average position. This graph shows the dependence of both the RMS horizontal error and the standard deviation for increasing occupation times. When we were viewing the individual 5-minute occupations earlier in this presentation, we noticed that the standard deviations were often not very similar to the accuracy of the average position. This graph illustrates that the standard deviation of position is not simply related to the accuracy of the average position at any occupation time (except where the two curves cross, but this point varies with different distances between the base and rover receivers). In particular, for short occupations the standard deviation can be very small, while the error of the average position is much larger. The standard deviation simply measures the spread of positions during the occupation; from the spread, we can not easily assess the accuracy of the average position. However, if the standard deviation or spread of positions is more than twice the specified instantaneous position accuracy of the particular receiver for the base-rover separation distance being used, we would suspect that this occupation has degraded accuracy. Degraded accuracy could result from bad site conditions or high PDOP, for example.

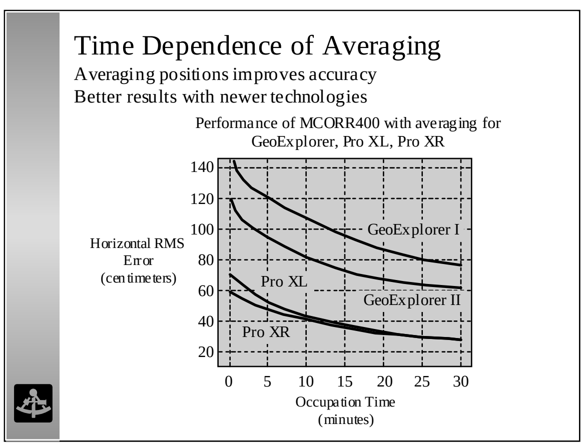

Now, we can examine the effect of averaging for all of the Trimble Pathfinder mapping receivers, the GeoExplorer I, the GeoExplorer II (both using internal antenna's), the Pro XL (using Compact Dome antenna), and the Pro XR (using integrated GPS-Beacon antenna). For all receivers, averaging over longer occupation times improves accuracy. Also, as Trimble has developed newer technologies, better and better accuracy has been obtained.

Base receiver:
- Pro XL with firmware version 3.08

Rover receivers:
- GeoExplorer I with firmware version 2.09
- GeoExplorer II with firmware version 2.09
- Pro XL with firmware version 3.08
- Pro XR with firmware version 1.00

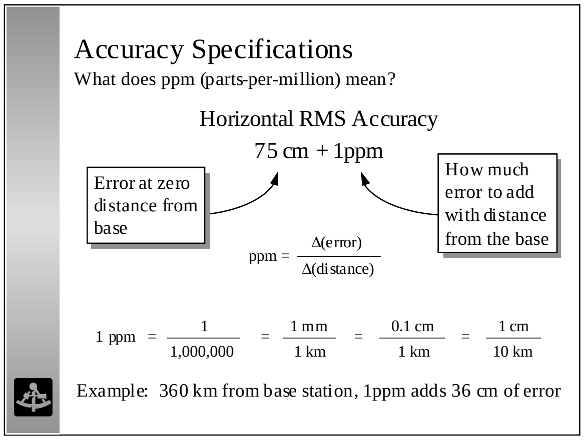

So far, we have been discussing the time dependence of accuracy and the effect of averaging positions. Accuracy is also dependent on the distance between the base and rover receivers. It is the distance dependence that is more commonly discussed in the literature and in product specifications. The distance dependence of accuracy is expressed by a fractional error called ppm, or parts-per-million. For example, the Trimble Pro XR specification[2] for horizontal RMS accuracy is stated as 75 cm plus 1 ppm. There are two parts to this accuracy specification. The left-hand side (75 cm) tells you how large the error is at zero distance from the base station. The right-hand side (1 ppm) tells you how much error to add with distance from the base station. In this example, the specification says 1 ppm, or 1 part-per-million which tells you how much error to add to the zero distance error (75 cm) per unit of distance from your base station. The 1 ppm value refers literally to the fraction 1 over 1,000,000, or one part in one million parts. This fraction can be written many different ways by noting that there are 1 million millimeters in one kilometer, and then performing some units transformations. So, 1 ppm also means 1 mm/km or 0.1 cm/km or 1 cm/10 km. As an example, if we are 360 km from our base station, 1 ppm adds 36 cm of error to the zero distance error of 75 cm for a total of 111 cm RMS horizontal error. If the ppm value is 2, 3, or 5, then the fractional errors just change to 2 mm/km, 3 mm/km, or 5 mm/km, respectively.

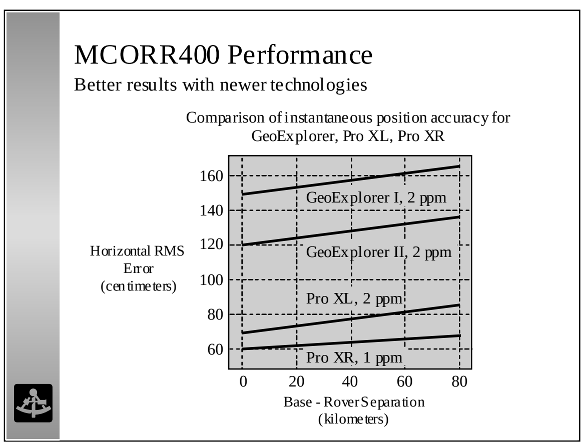

If we examine the distance dependence of error for all of the Trimble Pathfinder mapping receivers, we can see that newer technologies have led to better accuracy. This graph shows the horizontal RMS error on the vertical axis and the distance between the base and rover receiver on the horizontal axis. For the GeoExplorer I, II and the Pro XL, error increases with distance at 2 ppm. For the Pro XR, the error increases at 1 ppm.

Base receiver:
    Pro XL with firmware version 3.08
Rover receivers:
    GeoExplorer I with firmware version 2.09
    GeoExplorer II with firmware version 2.09
    Pro XL with firmware version 3.08
    Pro XR with firmware version 1.00

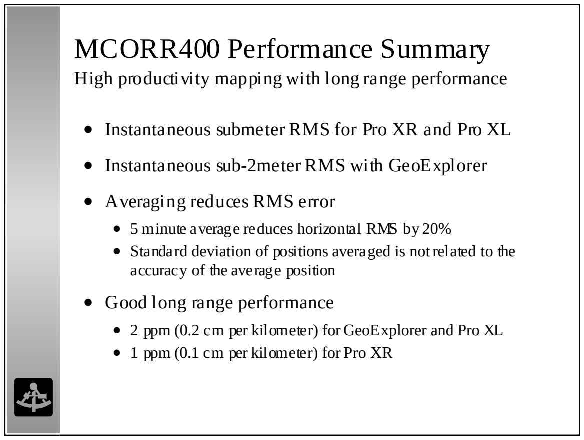

MCORR400 Conclusions

Here, we summarize the main points that we have discussed about the performance of Trimble Pathfinder mapping receivers when differentially corrected with MCORR400. The rest of this paper focuses on the Phase Processor software which is now a standard component of the Trimble Pathfinder™ Office software.

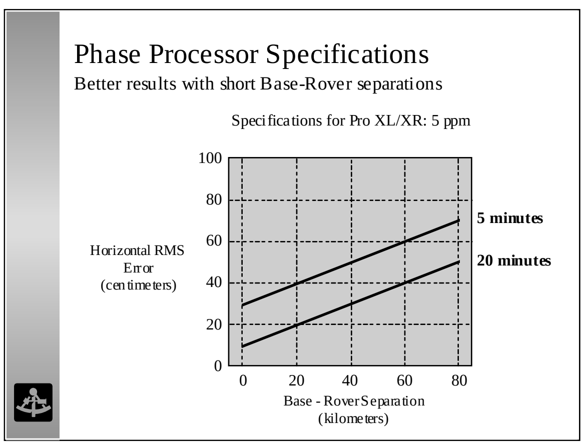

Phase Processor Discussion

The Phase Processor software can be used with all Trimble Pathfinder mapping receivers that support the collection of carrier phase data. On this graph of horizontal RMS error versus base-rover separation, the two lines show the specified accuracy for Pro XL/XR receivers for occupation times of 5 and 20 minutes. Note that with the Phase Processor software, the specified degradation of accuracy with distance from the base station is 5 ppm. However, at zero-distance the error is very small. Thus, the Phase Processor software is best used when mapping features within the vicinity of a nearby base station, or for mapping the relative positions of features using a local temporary base station. The specification says 30 cm + 5 ppm for a 5 minute occupation and 10 cm + 5 ppm for a 20 minute occupation. Again, in our discussion of the Phase Processor performance, we will focus on the horizontal accuracy. Unlike GPS positions corrected with MCORR400, the vertical error for Phase Processor solutions is about equal to the horizontal error.

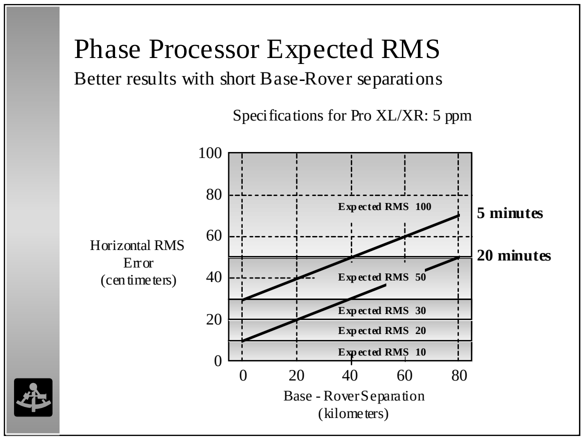

The new Phase Processor software version 2 produces an estimate of the horizontal RMS error for each solution. This graph shows the two specification lines for the 5 and 20 minute occupation times, and also the 5 Expected RMS values that the Phase Processor software will produce. The Expected RMS can take on the values 10, 20, 30, 50, and 100 centimeters. To interpret the Expected RMS correctly, we have to remember our discussion of the distribution of horizontal error. In that discussion we found that the RMS does not represent the error of a particular solution, but that it represents a point in the distribution of horizontal error that bounds about 63-68% of all solutions. The statement that we can make about any one solution with an Expected RMS of X is that we have 63-68% confidence that our solution is accurate to plus or minus X cm of truth. If we want to have a higher confidence, we can double the Expected RMS value to 2X and then we have 95-98% confidence that our solution is accurate to plus or minus 2X of the truth. We must acknowledge that to have greater confidence in our results means that we must accept larger uncertainty. In addition to these 5 values, the Expected RMS can display "Unreliable." This can occur when very poor quality data have been collected, or when the occupation time is too short (less than two minutes), or the baseline length is very long, or when the proper carrier phase data collection methods have not been followed.

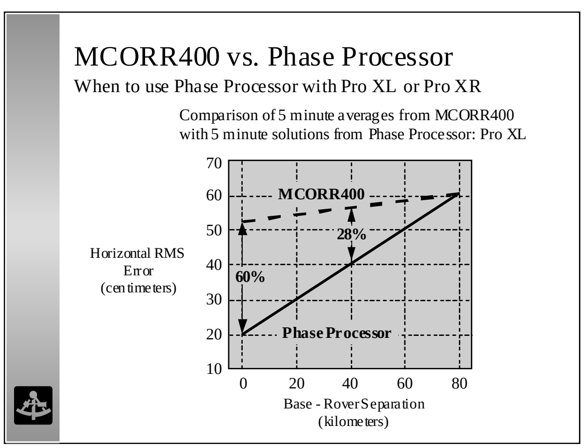

Base receiver:

Pro XL with firmware version 3.08

Rover receiver:

Pro XL with firmware version 3.08

When does the Phase Processor software provide a significant advantage over using MCORR400 for processing? This graph shows the horizontal RMS error, on the vertical axis, versus the distance between the base and rover receivers, on the horizontal axis. The two lines on the graph compare the difference between 5-minute averages of differentially corrected positions computed with MCORR400 (dashed line) and 5-minute occupations processed with the Phase Processor software (solid line). The data are from a Pro XL receiver. The two lines have different slopes because of the different parts-per-million (ppm) values associated with the two kinds of processing - 2 ppm for MCORR400 (1 ppm with Pro XR), and 5 ppm for the Phase Processor software. The space between the two lines indicates the improvement in accuracy when using the Phase Processor software. For very short distances from the base receiver, the Phase Processor provides 60% smaller horizontal RMS error than MCORR400. As the distance from the base increases, the improvement in accuracy is reduced; at 40 km, the Phase Processor software provides 28% smaller horizontal RMS error. For 5 minute occupations, there is not a significant advantage to using the Phase Processor software if you are more than 60 km from your base receiver.

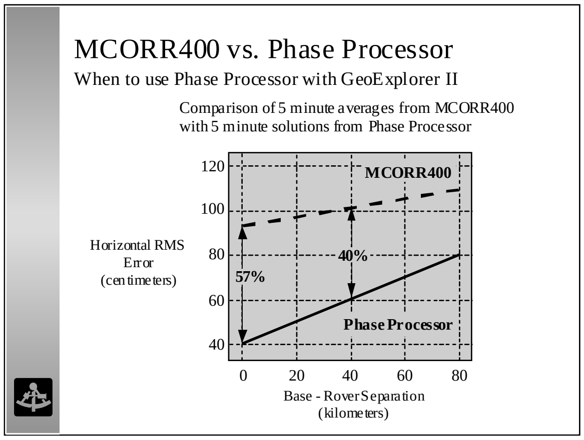

Base receiver:
Pro XL with firmware version 3.08
Rover receiver:
GeoExplorer II with firmware version 2.09

We can perform the same analysis with the GeoExplorer. This graph compares 5-minute averages of differentially corrected positions from MCORR400 (dashed line) with 5-minute solutions from the Phase Processor software (solid line) for data from a GeoExplorer II (using it's internal antenna). Again, the slopes of the two lines are different because of the different parts-per-million (ppm) values associated with the two different processing techniques - 2 ppm for MCORR400, and 5 ppm for the Phase Processor software. The Phase Processor software provides a significant improvement in accuracy for the shortest distances from the base receiver, about 57% lower horizontal RMS than MCORR400. As the distance from the base receiver increases, the improvement in accuracy is smaller; at 40 km, the Phase Processor software provides 40% lower horizontal RMS for the GeoExplorer II. The improvement in accuracy appears to be significant at longer distances from the base receiver than with the Pro XL/XR because of the relatively less accurate MCORR400 performance of the GeoExplorer II. While there is a significant accuracy improvement at longer distances, ionospheric noise at these distances can cause these solutions to be less reliable. Trimble recommends the use of the Phase Processor software at distances of 50 km or less.

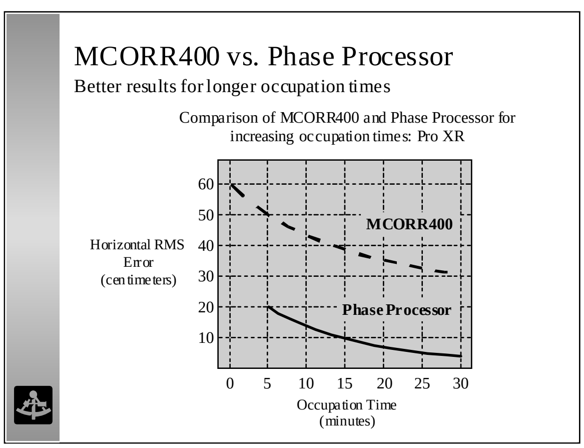

Base receiver:
Pro XL with firmware version 3.08
Rover receiver:
Pro XR with firmware version 1.00

As we discussed earlier in this presentation, better accuracy can be obtained by averaging MCORR400 differentially corrected positions over longer occupation times. Longer occupation times also improve the accuracy of Phase Processor solutions. This graph shows a plot of horizontal RMS error versus occupation time and compares averaged differentially corrected positions from MCORR400 (dashed line) with Phase Processor solutions (solid line). Longer occupation times improve the accuracy of Phase Processor solutions in a very similar way to averaging of MCORR400 corrected positions. The results shown in this plot are from a Pro XR rover receiver using a Pro XL base station at a very short distance (less than 1 km) from the rover.

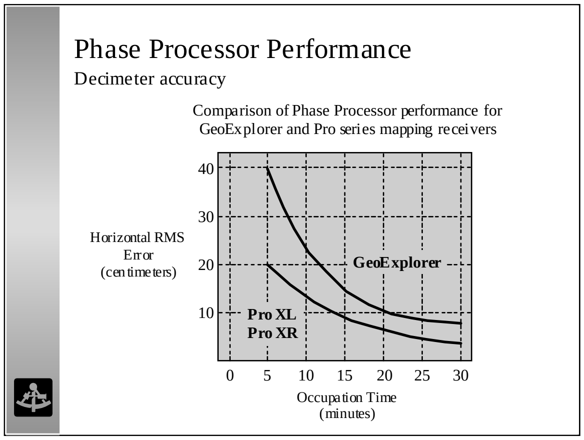

Base receiver:
- Pro XL with firmware version 3.08

Rover receivers:
- GeoExplorer I & II with firmware version 2.09
- Pro XL with firmware version 3.08
- Pro XR with firmware version 1.00

Now we can compare the Phase Processor performance of the GeoExplorer and the Pro XL/XR receivers. At all occupation times the Pro XL/XR receivers provide twice the accuracy of the GeoExplorer. For all receivers, longer occupation times improve accuracy. Decimeter accuracy is achieved after ~15 minutes for the Pro XL/XR and after ~22 minutes with the GeoExplorer.

Phase Processor Performance Summary

High accuracy mapping with exceptional short range performance

- Pro XR and Pro XL
  - Decimeter RMS with 15-minute occupation
  - Max. 60% reduction in horizontal RMS error over MCORR400
- GeoExplorer
  - Sub-50 centimeter RMS with 5-minute occupation
  - Max. 57% reduction in horizontal RMS error over MCORR400
- Longer occupation times improve accuracy
- Good short range performance
  - 5 ppm (0.5 cm per kilometer) limits range
  - Ionospheric noise dominates error at distances >100 km

Phase Processor Conclusions

Here, we summarize the performance of the Phase Processor software when used with Trimble's Pathfinder mapping receivers. The Phase Processor software is now a standard component of the Pathfinder Office™ software and provides exceptional short range accuracy. This short range accuracy is especially useful when absolute position mapping is not a requirement, but the relative positions of mapped features are important. In this case, setting up a local base station within your field area and mapping features with carrier phase data allows the relative positions of the local features to be very accurately determined.

For more information about GPS and Trimble products, visit the Trimble Navigation Limited WEB page at www.trimble.com

Notes

1. To search for National Geodetic Survey bench mark data sheets for your field area, try using the NGS Products WEB page at:

http://www.ngs.noaa.gov/page2.html

2. Trimble Surveying and Mapping Products: GPS Pathfinder Pro XR. Submeter GPS mapping/GIS data capture. TID10724 (9/96)

Appendix

Computing standard deviations

From any GPS position data set that has been converted into a local north and east coordinate system you can compute the average of the east and north components of position. Here, we denote the average in the east as $\bar{e}$ and the average in the north as $\bar{n}$; N is the total number of GPS positions. Utilizing these average position components, we can then compute the standard deviation of the east and north components. The standard deviation in the east component (sd_e) and the standard deviation in the north component (sd_n) can be combined to compute the horizontal standard deviation (sd_h).

$$\bar{e} = \frac{1}{N}\sum_{i=1}^{N} east_i$$

$$\bar{n} = \frac{1}{N}\sum_{i=1}^{N} north_i$$

$$sd_e = \sqrt{\frac{1}{N-1}\left(\sum_{i=1}^{N}\left(\bar{e} - east_i\right)^2\right)}$$

$$sd_n = \sqrt{\frac{1}{N-1}\left(\sum_{i=1}^{N}\left(\bar{n} - north_i\right)^2\right)}$$

$$sd_h = \sqrt{\left(sd_e{}^2 + sd_n{}^2\right)}$$

Computing true errors

If we know the true position of a site that we have made GPS measurements on, we can denote these truth coordinates as e_T and n_T. By subtracting each of our measured east and north components from the true coordinates, we can find the true error of our GPS data. These east (Ee_i) and north (En_i) components of error can be combined to find the horizontal error (Eh_i) for each GPS position. Utilizing the horizontal error, we can compute the horizontal RMS error for a set of GPS positions, where N is the number of positions in the set.

$$Ee_i = east_i - e_T$$

$$En_i = north_i - n_T$$

$$Eh_i = \sqrt{\left(Ee_i{}^2 + En_i{}^2\right)}$$

$$RMS = \sqrt{\frac{1}{N}\left(\sum_{i=1}^{N}\left(Eh_i\right)^2\right)}$$