Field Test Results for a Self-Calibrating Pseudolite Array

Edward A. LeMaster, *Stanford University*
Stephen M. Rock, *Stanford University*

BIOGRAPHY

Edward A. LeMaster is a Ph.D. candidate in Aeronautics and Astronautics at Stanford University. He received his B.S. in Aeronautical and Astronautical Engineering from the University of Washington in 1995 and his M.S. from Stanford in 1996. Since 1996 Edward has been working on applying GPS to robotic systems. Previously, he developed an integrated GPS/computer-vision navigation system for the Stanford HUMMINGBIRD autonomous helicopter. His current research focuses on applying pseudolite technologies to Mars exploration, specifically as an aid to rover navigation and cooperative task execution.

Stephen M. Rock is an Associate Professor of Aeronautics and Astronautics at Stanford University. He received his S.B. and S.M. in Mechanical Engineering from MIT in 1972 and Ph.D. in Applied Mechanics from Stanford University in 1978. Dr. Rock joined the Stanford faculty in 1988 where he teaches courses in dynamics and control and pursues research in developing and applying advanced control techniques for vehicle and robot applications. Prior to joining the Stanford faculty, Dr. Rock led the advanced controls group of Systems Control Technology.

ABSTRACT

Pseudolites can extend the availability of GPS-type positioning systems to a wide range of applications not possible with satellite-only GPS. One such application is Mars exploration, where the centimeter-level accuracy and high repeatability of CDGPS would make it attractive for rover positioning during autonomous exploration, sample collection, and habitat construction if it were available. Pseudolites distributed on the surface would allow multiple rovers and/or astronauts to share a common navigational reference. This would help enable cooperation for complicated science tasks, reducing the need for instructions from Earth and increasing the likelihood of mission success.

Conventional GPS pseudolite arrays require that the devices be pre-calibrated through a survey of their locations, typically to sub-centimeter accuracy. This is a problematic task for robots on the surface of another planet. By using the GPS signals that the pseudolites broadcast, however, it is possible to have the array self-survey its own relative locations, creating a Self-Calibrating Pseudolite Array (SCPA). This requires the use of GPS transceivers instead of standard pseudolites. Surveying can be done either at carrier- or code-phase levels. An overview of SCPA capabilities, system requirements, and self-calibration algorithms is presented in [1].

The Aerospace Robotics Laboratory at Stanford has developed a fully operational prototype SCPA. The array is able to determine the range between any two transceivers with either code- or carrier-phase accuracy, and uses this inter-transceiver ranging to determine the array geometry. This paper presents results from field tests conducted at Stanford University demonstrating the accuracy of inter-transceiver ranging and its viability and utility for array localization, and shows how transceiver motion may be utilized to refine the array estimate by accurately determining carrier-phase integers and line biases. It also summarizes the overall system requirements and architecture, and describes the hardware and software used in the prototype system.

INTRODUCTION

The benefits of using pseudolites to augment the Global Positioning System are well documented. First, they provide a method to rapidly initialize carrier-phase integers through rapid geometry change. This has been demonstrated for ground applications such as autofarming [2], and forms the central basis for the Integrity Beacon

Landing System [3]. Second, pseudolites can provide additional ranging sources for use in areas of limited sky coverage such as urban canyons [4] or for specialized applications such as open pit mining [5]. This use as an additional ranging source extends to space applications, where it has been proposed to aid spacecraft formation flying [6][7]. Pseudolites can also be used to increase integrity and reliability for applications with life-safety implications, such as with the FAA's proposed LAAS system [8].

In some applications it is desirable to have GPS-type positioning without the availability of GPS satellites. Pseudolite-only systems provide this capability. The use of pseudolite-only systems for indoor use was pioneered by Zimmerman [9] and has since been used successfully in a number of different applications [10][11]. In all of these applications, the locations of the pseudolites must be accurately known (typically to sub-centimeter level) in order to achieve centimeter-level positioning. Kee et al inverted the problem and showed in simulation that a receiver placed at precisely known locations can enable one to survey in the pseudolite locations [12]. This method does not work, however, when no device positions are known with certainty.

The current research at Stanford is aimed at developing a pseudolite-based navigation system for use in Mars exploration [1][13].[1] The concept calls for placing several pseudolites on the surface to allow local-area navigation and assist cooperative tasks by multiple rovers and/or astronauts, as is illustrated in Figure 1. One of the major difficulties with this proposal is that, as was just described, the pseudolite locations must be known with high accuracy for such a system to work. Such precise surveying is extremely difficult for an autonomous system on the surface of another planet.
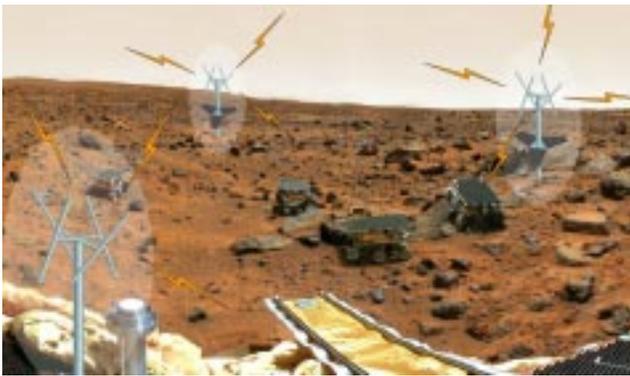
Figure 1: Mars SCPA

[1] This is a complementary system to the Mars Network proposed by JPL, which would give intermittent global positioning (~ 10 m accuracy) via orbiting satellites [14].

The solution to this problem is a new type of pseudolite positioning system called a Self-Calibrating Pseudolite Array (SCPA), in which each device is a full transceiver capable of both receiving and transmitting GPS signals. These transceivers exchange signals, allowing them to determine the ranges between the devices and therefore the corresponding array geometry. The positioning accuracy of a completely static array is limited by code-level noise. The array can provide carrier-level accuracy, however, if transceiver motion or other methods such as multiple-frequency pseudolites [15] are used to resolve the carrier-cycle (integer) ambiguities.

ARRAY CALIBRATION

SCPA self-calibration is a four-step process. First, the transceivers exchange signals to determine their relative ranges in a process called inter-transceiver ranging. Second, these ranges are combined to determine the array geometry. Third, an initial guess of the carrier-cycle ambiguities is made using the code-level solution. Finally, motion of a mobile transceiver (such as one carried by a rover) is used to refine this estimate of the integers and line biases and also to survey in the locations of the static transceivers to centimeter-level accuracy.

Inter-Transceiver Ranging

The primary measurement observable used for both positioning and array calibration is inter-transceiver bidirectional ranging. This is a double-difference scheme between a single pair of GPS transceivers in which each receiver takes a single difference between its own transmitted signal and the signal it receives from the other transceiver. These two single differences are then combined to determine both the range between the two devices and their relative clock offsets, as shown in Equation 1.

$$\begin{Bmatrix} \Delta \tau^{ij} \\ R_{ij} \end{Bmatrix} = \frac{1}{2}\begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}\begin{Bmatrix} \Delta \phi_i - \Delta b_i \\ \Delta \phi_j - \Delta b_j \end{Bmatrix} \quad (1)$$

with

Δb_i Line biases in Rec i

R_{ij} Range between device antennas

$\Delta \phi_i$ Single difference at Rec i

$\Delta \tau^{ij}$ Clock bias between TR i and TR j

Inter-transceiver ranging can be done at either the code or carrier level. Code ranging has the advantage that it is unambiguous, and can therefore provide a rough initial estimate of the carrier-phase integers. Moreover, if the

array is very sparse (2 static transceivers) or if high accuracy is not critical, code-ranging is sufficient to allow coarse (approx. 1-2 meter) positioning of a third mobile transceiver. One disadvantage with code ranging is that the line biases must be accurately calibrated beforehand, and even then noise and temperature-related drift can create sizable errors. An additional source of code ranging error is multipath, which can be extreme for SCPAs because of the two-dimensional geometry; signals reflect off the ground between the two transceivers, and also from buildings or other vertical surfaces surrounding the array. Carrier measurements are much less affected by these factors, but do have a cycle ambiguity. This can be resolved either by initializing the system with the transceivers a known distance apart, or by using an adaptive method to refine an initial guess.

Array Geometry Reconstruction

The methods currently used for converting these ranges into the array geometry include triangulation as well as non-linear iterative least squares when more than two ranges to a particular transceiver are available. Despite the simple methods used this is not always an unambiguous process. The solution space is highly non-linear, and any algorithm used must take into account singularities, multiple solutions, unobservable positions, and warping due to ranging errors. Figure 2 shows some of these effects for the case of triangulating the position of one transceiver with respect to two others, when both the ranging measurements are off by 2 meters. The gray area shows locations where a solution is impossible, while the contours show the magnitude of the positioning errors over the entire workspace. Other ranging errors give similar elliptical or hyperbolic solution spaces.
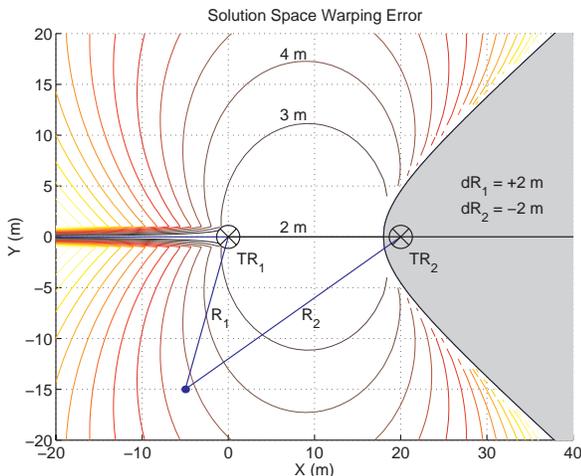

Figure 2: Triangulation Errors

Using more than two ranges to determine positions via a non-triangulation technique adds a great deal of robustness but also adds the difficulty of switching between solution methods depending upon the number of

range measurements available, especially when dropping down to only two. This switching causes discontinuities in the solutions space, which can pose a problem for tasks involving precise relative navigation.

Array Refinement

The techniques described above are not sufficient for centimeter-level accuracy because the unknown bias terms in the ranging measurements warp the solution space. It is possible to determine these biases, however, by looking at the ranging measurements over time during some change of array geometry, and then solving for the biases that are consistent with those measurements.

The current system accomplishes this by having the fourth of the four transceivers mobile, and collecting carrier-range data at widely separated locations. This data is then processed in a batch computation to determine the actual positions of the static transceivers and the mobile transceiver at each of the sample points, as well as the carrier-cycle ambiguities and other range biases. The algorithm employs non-linear iterative least squares, and works very well under two conditions. First, the array must maintain at least a rough position fix on the mobile transceiver throughout the course of its motion. If enough ranging measurements are lost simultaneously, the biases will in general be different upon reacquisition. Since these biases must be resolved to centimeter-level, this requires restarting of the array refinement estimator. Second, there must be enough relative motion between the transceivers to give observability of the positions of the static transceivers. This generally requires motion outside the bounds of the static array such as by looping around the transceivers, as is demonstrated in Figure 14.

EXPERIMENTAL SYSTEM

An SCPA is a complex distributed system, relying on a large number of separate components. The most important are the transceivers, which in this system are self-differencing transceivers or 'differlites' [16]. The pseudolite signal is split, with a component fed directly into the receiver via a dedicated front end. This gives the advantage that the receiver is virtually guaranteed to be able to receive its corresponding transmitted signal regardless of interference from other pseudolites, at the expense of a small amount of added complexity. Other components include the antenna system, the wireless communications system for data collection, the central ground station processing computer, and the array management software.

The current system includes four operational transceivers. This is the minimum number needed for both unambiguous dynamic positioning and for the array

refinement algorithm, necessitating nearly constant tracking of all pseudolites on all receiver channels in order to achieve continuous localization. Performance may be improved by adding redundant static transceivers to the array.

GPS Pseudolites

The pseudolite portion of each transceiver is an IntegriNautics IN200C signal generator. These pseudolites are programmable for different PRN numbers and output power levels, several different pulsing schemes and data messages, and optional frequency offsets from L1. The operating mode used for this research is an RTCM pulsing sequence with a 9% duty cycle. This pulsing helps reduce the effects of the near-far problem − wherein receivers have trouble tracking signals from sources at different ranges because of the widely varying power levels – by taking advantage of the non-linear saturation characteristics of the receivers. The data message is a single repeated frame with an incrementing TOW value in each subframe, allowing unambiguous timing for a maximum of 15 seconds before aliasing occurs. The total combined output power of the entire array is less than 1μW.[2] This limits the effective range of the array to approximately 20 meters.

GPS Receivers

The receiver portion of each transceiver is a Mitel Orion receiver. Because the source code is available, these receivers are fully programmable to accept the pseudolite data structure and to output the tracking data needed for operation of the SCPA. The modified versions used in this research also have two separate RF front ends.[3] Each pseudolite output signal can therefore be split and fed directly into a dedicated front end on its corresponding receiver. Although the separate transmission path is not necessary for tracking, it does give an added measure of signal separation and tends to speed acquisition times. Each receiver is equipped with an RS-232 serial link for data collection, and gives an output data rate of 5Hz.
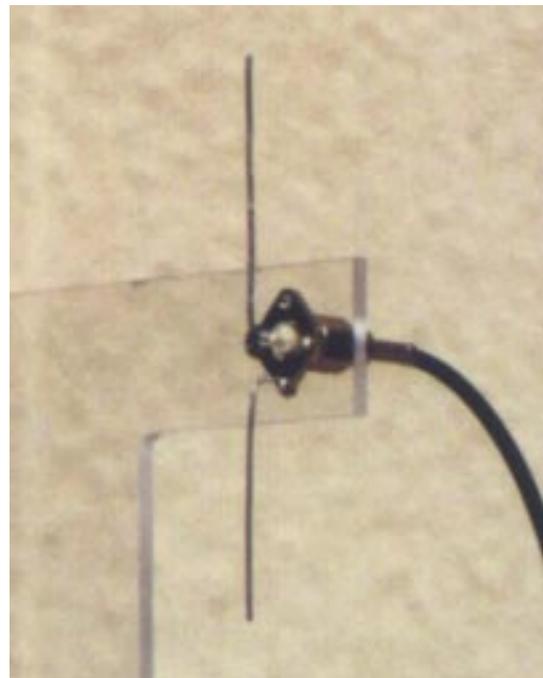
Antennas

Each transceiver has an antenna pair, one dedicated to transmitting the pseudolite signal and one for receiving. Although an earlier prototype used commercial patch antennas, the current system uses custom dipole antennas. This is because the antenna pattern desirable for an SCPA – a 360° omnidirectional pattern at low elevations − is not readily available in commercial devices. A more complicated antenna such as a Lindenblad would give a circularly polarized waveform and therefore improve ground-reflection multipath rejection. Simple dipoles are preferable, however, because of the ease of construction and because they do not experience phase windup when the mobile transceiver moves in trajectories around the static transceivers. Table 1 summarizes some of the design tradeoffs with this particular antenna design.

Table 1: Dipole Antenna Tradeoffs

Advantages	Disadvantages
Omnidirectional pattern	Poor out-of-band signal rejection
No phase windup	
Low Cost	No multipath rejection

Figure 3 shows the construction of these antennas. The base is a female bulkhead SMA connector, to which are soldered pieces of bus-wire. The overall length of the antenna is slightly less than the theoretical 9.5 cm for a half-wave dipole because of the impedance properties of the connector. These antennas are individually tuned by connecting them to a network analyzer and trimming the radiating elements until maximum transmission is close to L1. Figure 4 shows the transmission characteristics of a typical antenna, together with a commercial patch antenna and a bare SMA bulkhead connector for comparison (the plot shows the magnitude of the non-radiated signal, so a smaller value indicates greater transmission). The antenna is a fairly good radiator at L1, with approximately 98% of the input energy radiated out. Compared with the commercial antenna, however, it does not offer much out-of-band signal rejection.


Figure 3: Dipole Antenna

---

[2] These limits are set by the FCC, which allows users with an experimental license to intentionally broadcast on L1 with a maximum continuous power of 1μW.

[3] Receiver hardware and initial tracking loop modifications by Dr. Eric Olsen, Stanford University.

In the transceivers these antenna pairs are arranged vertically, with a spacing of 26 cm and with each antenna in the null of the other's beam pattern, as shown in Figure 5. Although near-field effects do allow some direct transmission between the antennas, this arrangement does reduce the level of cross-talk interference. This is important because the transmitted signal may be 60 to 80 dB stronger than the received signal as a result of near-far effects.
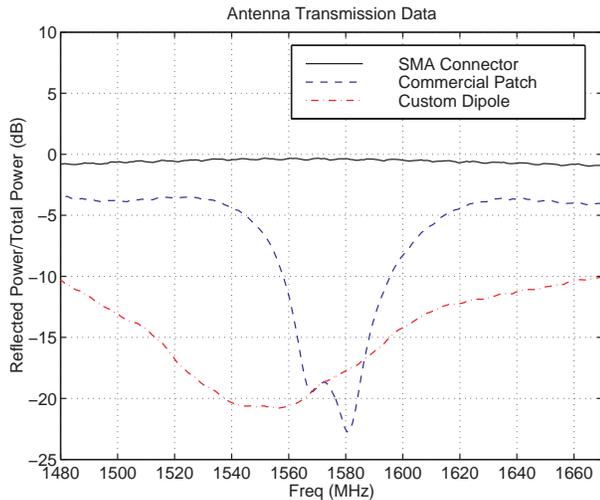
Figure 4: Antenna Transmission Characteristics

Communications

Because of the distributed nature of the SCPA, a wireless communications system is necessary to transmit the raw receiver tracking information to a central location for processing. At low data rates (potentially up to 500 bps) this could be done by piggybacking on the pseudolite data message itself. The current system uses a commercial wireless product, however, in order to speed development and to achieve higher data rates.

The product is a RangeLan2 network, manufactured by Proxim. Each transceiver is equipped with a Serial Adapter which takes the raw RS-232 output from the receiver and converts it to TCP/IP packets, broadcast at up to 1.6 Mbps. The ground-station computer then collects these packets using an Ethernet access point. The Ethernet protocol guarantees that no data is lost, but it does also add variable latencies into the system. These latencies are handled by the array management software.

Transceiver Totes

Special 'Transceiver Totes' (Figure 6) carry the components of each transceiver and make the entire system portable. Each tote holds a pseudolite, a receiver, a wireless unit, a 4.4 A-hr NiCd battery pack which provides roughly 4 hours of continuous operation, and the corresponding cables and wiring harnesses. All of the

components are secured inside the case for ease of operations and safety.

Figure 5: Static Transceiver

Figure 6: Transceiver Tote

Ground Station

The ground-station computer used for central data processing is a 133MHz Pentium laptop running a Windows NT operating system. This processor speed is adequate for full operation of the SCPA due to the relatively low data rates. However, NT is a poor approximation for a real-time operating system and offers limited time management capabilities. This forces the use of extensive buffering and time stamping to ensure that no receiver data is lost or otherwise mishandled.

Array Management Software

In order to handle the data management and processing duties for the array the SCPA uses a custom application called the GPSMixer. The name was inspired by mixing consoles used in lighting and sound applications, and it performs a similar function: bringing in raw data from the widely distributed components of the array and combining it to form useful output. It also allows the user to move through the array hierarchy, controlling the array at levels ranging from a single unified entity down to the level of individual receiver tracking channels. In addition, the software provides the capability to save and replay data for future analysis.

The GPSMixer has four primary levels. The Receiver Interface is at the lowest, and presents receiver tracking information such as SNR values, oscillator frequency offset, and code bit and frame sync. The operator can change relevant tracking parameters at will. The next higher level is the Pair Manager, which combines the raw data into code- and carrier-phase ranging solutions between all the possible transceiver pairs. This level also keeps track of biases and carrier-cycle ambiguities, and checks for cycle slips and other error sources. It also manages data latency and dropouts. The third level is the Array Manager, which converts the range data into an array geometry. The Array Manager includes separate estimators for code and carrier positioning, because both measurement types are not always available simultaneously. The highest level is the Bias Estimator, which monitors the array geometry over time and runs the refinement algorithms to determine carrier-cycle integers and line biases, and improve the estimate of the locations of the static transceivers.

FIELD TEST LOCATIONS

Tests of the array were conducted at two different locations on the Stanford campus. The first was the Science and Engineering Quadrangle (SEQ), a square open-air plaza roughly 80 meters across. (Figure 7) The tile-paved center section of the quad is surrounded by a grassy area 10 meters wide, which is in turn enclosed by covered pedestrian walkways and then the surrounding buildings. Trees are scattered throughout the plaza. The multipath environment in the SEQ is fairly high, and very difficult to characterize accurately because of the many smaller obstacles. Testing was performed in the grassy strips at the perimeter of the quad. When possible these tests avoided having intervening sections of pavement between the transceivers in order to minimize ground reflections.

Figure 7: Science and Engineering Quadrangle

The second location at Stanford was the Oval, located between Palm Drive and the Main Quad (Figure 8). This is a flat, open, grassy area roughly 200 meters long and 100 meters wide. The Oval is about as benign a multipath environment as exists in an urban environment, and test results from here show the clear benefits.
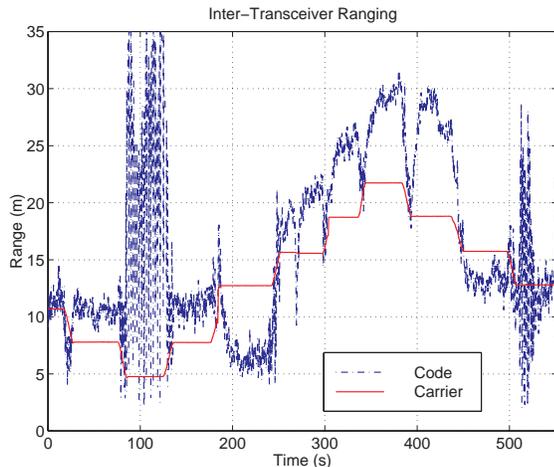

Figure 8: The Oval

EXPERIMENTAL RESULTS

Inter-Transceiver Ranging

The core algorithm for the array self-calibration relies upon direct inter-transceiver ranging, both at the code and carrier levels. Figure 9 shows ranging data taken in the SEQ, a high-multipath environment. One of the transceivers in the pair is fixed, and the other moves

relative to the first in roughly 3 meter increments. Carrier cycle ambiguities are eliminated by starting the devices at a known distance from each other. The carrier measurements are quite stable and except for an undetected cycle slip of roughly 2 meters (at T = 190 seconds), track the motion of the transceiver extremely well. The code measurements clearly show the effects of multipath, however, exhibiting considerable divergence from the carrier during the motion and in one instance indicating the complete opposite direction of motion.

Figure 9: Ranging in a High-Multipath Environment

Figure 10 shows a similar test on the Oval, a low-multipath environment. The code visibly tracks the carrier, and there are no indications of any divergence such as was present during the tests in the SEQ. In environments such as this, averaged code measurements can be used quite effectively to provide initial estimates for carrier-cycle initialization. Further testing must be done to assess the multipath level in environments more similar to that of the Martian surface.
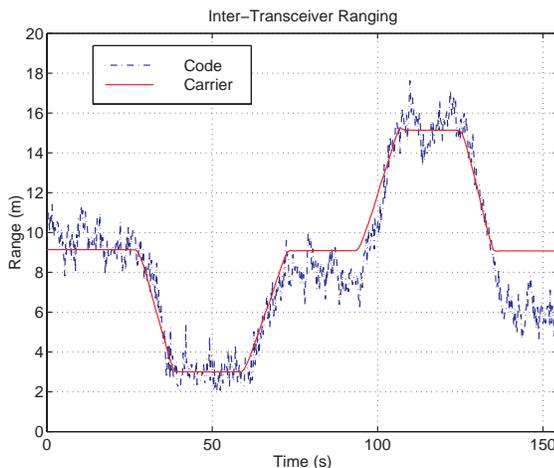
Figure 10: Ranging in a Low-Multipath Environment

Overall, the ranging accuracy in these tests is roughly what would be expected from the GPS signal structure.

Table 2 shows typical observed RMS ranging accuracies. The data used for these values comes from a total of 12 minutes of raw data collected using 6 different transceiver pairings (at known locations to eliminate the bias terms). Throughout these tests, near-far causes the received signal power levels to vary by up to 35dB.

Table 2: Ranging Accuracies

Code	Carrier
1.23 m	0.76 cm

Static Positioning

The geometry of an SCPA and the positions of the transceivers may be deduced once the ranges between the devices are known. Figures 11 and 12 show positioning data for static transceiver TR_3 from the triangular array depicted in Figure 13, computed using simple triangulation. The ellipses show the 2-σ positioning error bounds, the shapes of which are driven by the array geometry (X-DOP = 1.36, Y-DOP = 0.82). The sizes of the ellipses exhibit good correspondence with the raw ranging accuracy presented above.
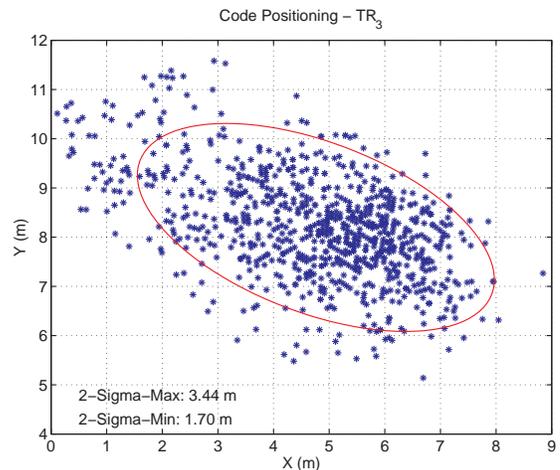

Figure 11: Transceiver 3 Static Code Positioning
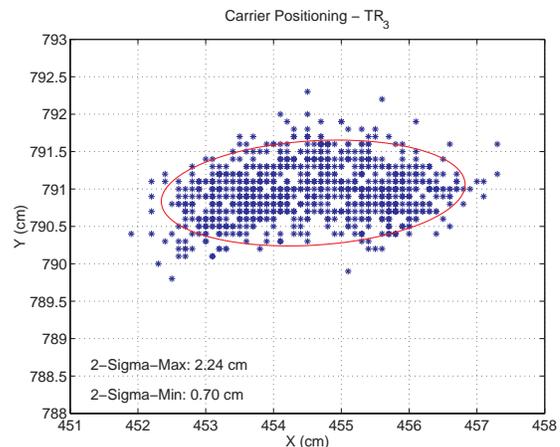
Figure 12: Transceiver 3 Static Carrier Positioning

Dynamic Positioning

During array calibration, the state of primary interest is the position of the mobile transceiver as it moves through the array. Several tests were therefore conducted in order to determine the accuracy of this dynamic positioning. Figure 13 shows a typical configuration for these tests, which includes 3 static transceivers and a fourth mobile transceiver. The static transceiver positions are pre-surveyed, as is the starting location of the mobile transceiver. Because the current system does not have an accurate dynamic truth reference, several reference points are pre-surveyed and marked prior to each test. The survey method is via tape measure, with an overall accuracy of roughly 20-30 cm. Transceiver motion is imparted by simply hand-carrying the device, and so the actual ground tracks are not precise.
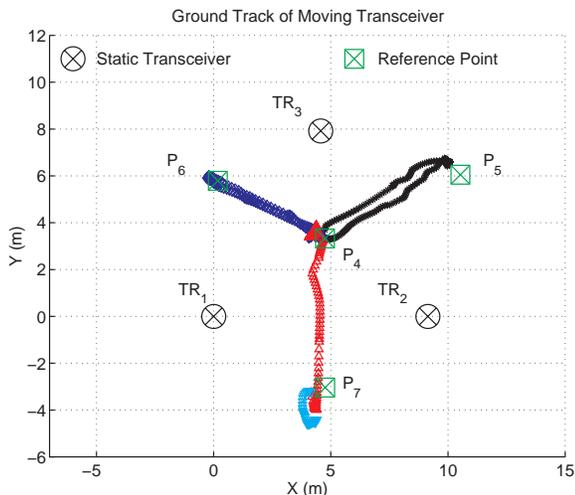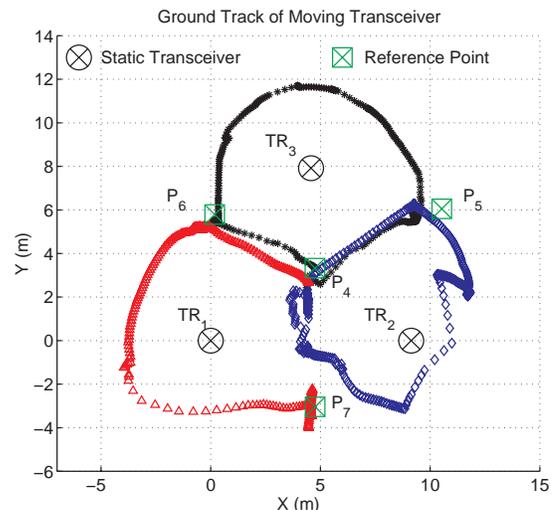

Figure 13: Trajectory #1

Figure 13 shows the calculated trajectory of the mobile transceiver using carrier-phase measurements only. In this trajectory, the mobile transceiver starts at point P_4. It then moves to P_5 and back, then to P_6 and back, and finally to P_7. The small cyan loop near P_7 is a result of the tester rotating the mobile transceiver in order to provide it with better antenna reception towards the other devices. Since the antenna is offset from the center of the antenna tripod, this causes it to trace out an arc. The paths do not terminate exactly at the reference points because the pre-surveying is of limited accuracy, resulting in slightly inaccurate estimates of the carrier-cycle ambiguities and a small warping of the solution space.

Figure 14 shows another trajectory, this one designed to give the observability necessary to refine the estimate of the array positions. The mobile transceiver starts at the center (P_4). It then traces the top (black) path clockwise around TR_3 followed by the right (blue) path counterclockwise around TR_2, and finally moves

counterclockwise on the left (red) path around TR_1 to point P_7, where the receiver lost lock.

Figure 14: Trajectory #2

This data clearly shows some of the difficulties encountered in converting the ranges to an array geometry. The apparent backtracking on the blue loop near P_4 is actually rapid switching between different solution and range combinations as the signals fade in and out of lock. The mobile transceiver then lost its ranging signal from TR_3 completely, causing the algorithm to rely wholly on triangulation. The jumps on the outside of the blue loop are a result of the transceiver passing over a region of no possible solutions, such as is illustrated in Figure 2.[4] By the time it starts the red loop, however, the transceiver has regained lock and reinitialized the corresponding biases. More intelligent algorithms are being developed to better handle these signal dropouts and solution-space handoffs.

Array Calibration

The locations of the static transceivers in the array and the carrier-cycle ambiguities can be accurately determined by using a batch refinement algorithm on data collected during the transceiver motion. Although the data from Figure 15 is not sufficiently consistent to demonstrate this algorithm because of the loss of lock in the middle and the subsequent reinitialization,[5] its effectiveness can be determined by using a custom simulator which mimics the effects of an actual transceiver moving through the array. The simulator does this by generating raw receiver tracking loop output such as would be seen during an actual field test, and then injecting appropriate noise levels and additional biases. The simulator output is then

---

[4] The widely separated points are artifacts of a smoothing filter on the output.

[5] The algorithm successfully converges, but gives position errors of roughly one meter.

fed directly into the same GPSMixer processing system used for the hardware field tests.
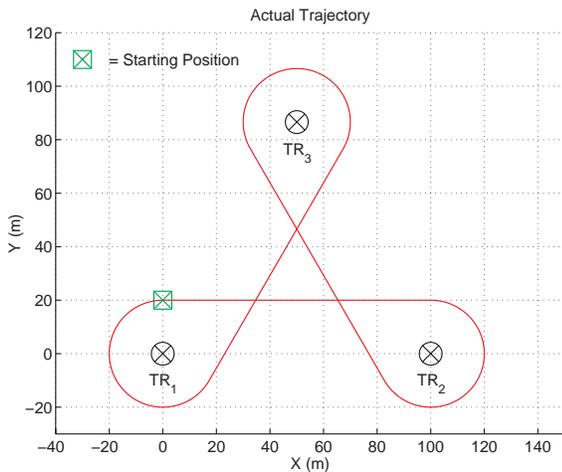
Figure 15: Simulation Trajectory

Figure 15 shows the actual trajectory traveled through the array. Note that the array is 100 meters across, about the size expected for use on the Martian surface. Figure 16 shows the path computed by the array management algorithm, but before the batch array refinement step is completed. Code noise was roughly 2 meters and carrier noise 1 cm, while integer-valued ranging biases varied from 1 to 5 meters. The discontinuities in the position estimate are caused by switching between various solution subspaces as they become available.
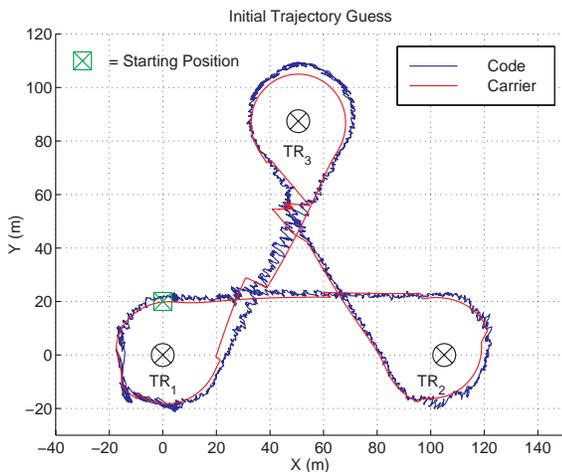
Figure 16: Pre-Refinement Trajectory Estimate

Using Figure 16 as an initial guess for the trajectory and the static transceiver locations along with an initial guess of no range biases, the batch ILS algorithm correctly converged upon the actual trajectory, positions, and biases in three iterations. The RMS residuals were 0.578 mm, and the RMS positioning and bias errors were 1.732 mm. A further breakdown of the solution values appears in Tables 3 and 4.

Table 3: Static Transceiver Locations

Initial Positions		Final Positions	
X (m)	Y (m)	X (m)	Y (m)
0	0	0	0
105.003	0	100.003	0
50.560	87.436	50.002	86.602

Table 4: Solution Bias Values

Final Ranging Biases (m)			
TR #	1	2	3
2	-5.002		
3	-1.000	-3.002	
4	1.000	-1.002	2.999

CONCLUSIONS

GPS pseudolites constitute a useful and viable method to achieve CDGPS-type positioning in locations without GPS satellite coverage. One of the primary challenges with using pseudolite arrays – that of surveying in the locations of the pseudolites − can be overcome by using a special variant called a Self-Calibrating Pseudolite Array. In an SCPA each device is a full transceiver, and they exchange signals to compute their relative ranges and hence the array geometry.

Field tests of the prototype system have successfully achieved both code- and carrier-phase inter-transceiver ranging, with corresponding accuracies of less than 1.5 meters and 1 cm, respectively. Positioning of both the static and mobile devices in the array and the corresponding determination of array geometry has been experimentally demonstrated to corresponding levels of accuracy. Simulations have also shown how transceiver motion can be used to further refine the array position and bias errors, provided that reasonable carrier lock is maintained during the maneuver.

Several modifications can be made to the current prototype system to improve performance. Adding extra static transceivers would greatly increase system redundancy and reduce dropouts. The effective range of the array can also be improved, both by increasing pseudolite signal strength and by rewriting the receiver tracking loops to include gain scheduling. The algorithms used for transceiver position estimation can also be refined to reduce switching between solution subspaces and to more smoothly combine the different range measurements.

Future field tests will be conducted to further study these issues, and will demonstrate the use of this system together with other Mars navigation systems upon the NASA Ames Research Center K9 rover. These latter tests are scheduled to begin in late October, 2000.

REFERENCES

[1] LeMaster, E., Rock, S., "Self-Calibration of Pseudolite Arrays Using Self-Differencing Transceivers", Proceedings of the Institute of Navigation GPS-99 Conference, Nashville, TN, Sept. 1999, pp. 1549-1558.

[2] Elkaim, G., et al., "System Identification and Robust Control of Farm Vehicles Using CDGPS", Proceedings of the Institute of Navigation GPS-97 Conference, Vol. 2, Kansas City, MO, Sept. 1997, pp. 1415-1424.

[3] Pervan, Boris, *Navigation Integrity for Aircraft Precision Landing Using the Global Positioning System*, Ph.D. Thesis, Stanford University, March 1996.

[4] Altmayer, Christian, "Experiences Using Pseudolites to Augment GNSS in Urban Environments", Proceedings of the Institute of Navigation GPS-98 Conference, Nashville, TN, Sept. 1998, pp. 981-991.

[5] Stone, J., and Powell, J.D., "Precise Positioning with GPS near Obstructions by Augmentation with Pseudolites", Proceedings of IEEE PLANS 1998, palm Springs, CA, pp. 562-569.

[6] Bauer, Frank H., et al., "Enabling Spacecraft Formation Flying through Spaceborne GPS and Enhanced Autonomy Technologies", Proceedings of the Institute of Navigation GPS-99 Conference, Nashville, TN, Sept. 1999, pp. 369-384.

[7] Corazzini, T., and How, J., "Onboard GPS Signal Augmentation for Relative Navigation", Proceedings of the Institute of Navigation GPS-99 Conference, Nashville, TN, Sept. 1999, pp. 1559-1568.

[8] Levy, Larry, et al., "GPS Risk Assessment for Civil Aviation", Proceedings of the Institute of Navigation GPS-99 Conference, Nashville, TN, Sept. 1999, pp. 2121-2130.

[9] Zimmerman, Kurt, *Experiments in the Use of the Global Positioning System for Space Vehicle Rendezvous*, Ph.D. Thesis, Stanford University, December 1996.

[10] Olsen, Eric, et al., "3D Formation Flight Using Differential Carrier-phase GPS Sensors", Proceedings of the Institute of Navigation GPS-98 Conference, Nashville, TN, Sept. 1998, pp. 1947-1956.

[11] Teague, Harris, *Flexible Structure Estimation and Control Using the Global Positioning System*, Ph.D. Thesis, Stanford University, May 1997.

[12] Kee, Changdon, et al., "Precise Calibration of Pseudolite Positions in Indoor Navigation System", Proceedings of the Institute of Navigation GPS-99 Conference, Nashville, TN, Sept. 1999, pp. 1499-1507.

[13] LeMaster, E., Rock, S., "Mars Exploration Using Self-Calibrating Pseudolite Arrays", Proceedings of the Institute of Navigation GPS-98 Conference, Nashville, TN, Sept. 1998, pp. 1967-1974.

[14] Ely, Todd A., et al., "Mars Network Constellation Design Drivers and Strategies", AAS/AIAA Astrodynamics Specialist Conference, Girwood, Alaska, Aug. 1999. (AAS 99-301)

[15] Stone, J.M., Powell, J.D., "Carrier Phase Integer Ambiguity Resolution Using Dual Frequency Pseudolites", Proceedings of the Institute of Navigation GPS-98 Conference, Nashville, TN, Sept. 17, 1998, pp. 961-968.

[16] Stone, J., et al., "GPS Pseudolite Transceivers and their Applications", Proceedings of the 1999 Institute of Navigation National Technical Meeting, San Diego, CA, Jan. 1999.