An Improved Solution Algorithm for Self-Calibrating Pseudolite Arrays

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, and will complete his Ph.D. in early 2002. Since 1997 Edward has been working on applying GPS pseudolites to robotic systems, specifically as an aid to Mars rover navigation and cooperative task execution. Previously he developed an integrated GPS/computer-vision navigation system for the Stanford HUMMINGBIRD autonomous helicopter.

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

Tasks envisioned for future generation Mars rovers – sample collection, area survey, resource mining, habitat construction, etc. – will require greatly enhanced navigational capabilities over those possessed by the Mars Sojourner rover. Many of these tasks will require cooperative efforts by multiple rovers and other agents, adding further requirements both for accuracy and commonality between users. Although a Mars-based global satellite navigation system analogous to GPS could provide the necessary capability, the cost of such a system with sufficient coverage and accuracy would likely be prohibitively high, at least within the foreseeable future. Precise positioning can be accomplished in the absence of a satellite navigation system, however, by using GPS pseudolites or similar ground-based transmitters to mimic the effects within a local area. Stanford University has developed such a system to aid NASA in its future Mars exploration goals.

The greatest difficulty in using pseudolites on the Martian surface is determining the precise locations of the transmitters themselves, a prerequisite for successful navigation to occur. The Self-Calibrating Pseudolite Array (SCPA) developed at Stanford solves this problem by utilizing GPS transceivers for improved geometric observability and limited motion of a rover-mounted transceiver to resolve the carrier-phase cycle ambiguities.

This paper presents recent advances in the development of the operational SCPA prototype, expanding upon the preliminary methods and results presented in [1]. It begins by summarizing the overall system requirements and architecture and then discusses the array self-calibration methods themselves, focusing on a new quadratic algorithm for precisely determining the locations of the pseudolite transceivers. This solution method offers a factor of 2-3 improvement in terms of solution convergence over the previous linear algorithm, allowing for the calibration of smaller-baseline arrays in high-multipath environments. When coupled with a stochastic multiple-solution approach, the combined algorithm yields successful self-calibration under most possible array configurations. Finally, the paper reviews the results of field trials performed at NASA Ames Research Center with the K9 Mars rover research platform which validate both the navigation and self-calibration capabilities of the system. By carrying an onboard GPS transceiver, K9 was successfully able to calibrate the system using no a priori position information, and localized the pseudolite beacons to under 5 cm RMS.

INTRODUCTION

Mars surface exploration presents many challenges to designers of robotic systems. Long communication time delays (up to 40 minutes round trip) and limited bandwidth dictate high levels of rover autonomy. The rovers, however, are also operating in a very uncertain and potentially hostile environment. In order to achieve successful autonomy, rovers must be able to sense and make sense of the environment around them. This sensing requirement becomes even more stringent when multiple rovers are attempting to cooperate in a common area to do joint tasks. Such cooperative tasks can include surveying, resource mining and utilization, and habitat construction.

A GPS-analogue system would be of great benefit for future Mars operations, and in fact NASA is considering using an orbiting satellite constellation to aid in its Mars exploration program. The JPL proposal, called the Mars Network, would incorporate approximately a half-dozen satellites to provide both global communications and approximately 10 meter positioning accuracy at fix intervals of several hours [2]. Although the proposed system would be very valuable, it does not provide the continuous high-accuracy positioning required for many tasks, especially cooperative operations between several rovers. A complementary system to provide this capability is therefore highly desirable.

GPS pseudolites provide one viable alternative to satellite-based navigation. The use of pseudolite-only systems for indoor use was pioneered by Zimmerman [3] and has since been used successfully in a number of different applications [4][5]. The current research at Stanford is aimed at developing a pseudolite-based navigation system for use in Mars exploration. 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 the pseudolite locations must be accurately known (typically to centimeter level) in order to achieve centimeter-level positioning. Such precise surveying is extremely difficult for an autonomous system on the surface of another planet, and is of course the impetus for the creation of the system itself.

The Self-Calibrating Pseudolite Array created at Stanford overcomes this difficulty by solving for the locations of the stationary broadcast elements simultaneously with the determination of the carrier-phase integers and the location of the mobile units. This is accomplished through the use of GPS transceivers instead of separate transmitter and receiver units in order to increase the geometric observability of the system. These transceivers are able to directly determine the range between any pair of devices through a process called bidirectional ranging. Special algorithms have been developed which take in range data between the transceivers over the course of a calibration trajectory performed by a transceiver-bearing rover, and are then able to back out the system states from a batch calculation.

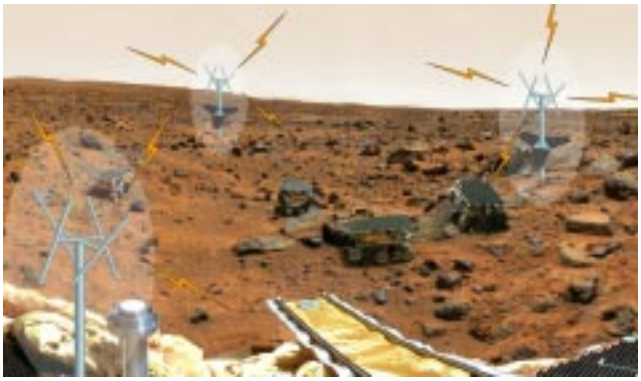

Figure 1: Mars SCPA

An experimental hardware prototype constructed at Stanford and operated at NASA Ames Research Center has validated the basics of transceiver operations, and has successfully shown full array self-calibration in the presence of moderate errors in the initial location estimates. This previous work is presented in [1].

The current paper describes new algorithms that have been developed to improve self-calibration performance in the presence of very large errors in the initial location estimates, and also shows simulation results which validate the improved self-calibration capability. Additionally, these new algorithms have been successfully applied to the prototype system.

SYSTEM OVERVIEW

An SCPA is a distributed system consisting of several GPS transceivers together with a common base computer for data processing. The transceivers exchange bidirectional ranging signals between themselves, and triangulation methods then enable relative positioning of the devices. This section provides a brief overview of both the navigation and self-calibration functions of the array.

Basic Navigation

Navigation using an SCPA follows the same principle as satellite-based differential GPS, and can be accomplished at both the code or carrier levels. Details of conventional GPS navigation can be found in [6]. In order to achieve precise navigation without using atomic clocks, a double-difference ranging solution has been developed between GPS transceivers with both receiving and transmitting

elements in a common device. The resulting bidirectional ranging solution involves exchanging ranging signals (corrupted by clock biases) between device pairs. It then cancels out the clock biases associated with the transmitter oscillators through the differencing process, as is presented in detail in [7].

Determination of the array geometry and the location of the rover are accomplished by combining the range measurements between transceiver pairs, either using triangulation or standard non-linear optimization techniques. Code-level positioning is available instantaneously, allowing a rough (2-4 meter) navigation capability to all users within the array. Although uncalibrated line and system biases can further degrade the accuracy, code-based ranging is sufficient for many tasks such as general navigation between points and collision avoidance.

If more precise navigation is required – such as for more complex or repetitive tasks like cooperative manipulation or construction – carrier-phase positioning may be performed. Raw carrier-phase ranging accuracy using the SCPA has been demonstrated to better than 0.8 cm RMS [1]. Achieving such accurate positioning is only possible after an additional calibration step is used to resolve the associated integer ambiguities.

Array Self-Calibration

Array self-calibration to determine these carrier-phase integers follows a multiple-step process (Figure 2). Following array deployment, initial coarse calibration is obtained by using code-level bidirectional ranging between the transceivers to triangulate their relative positions. The self-calibration process itself then utilizes the relative motion of a transceiver-bearing rover to alter the array geometry over time. During this motion the unknown carrier-phase integers remain constant.

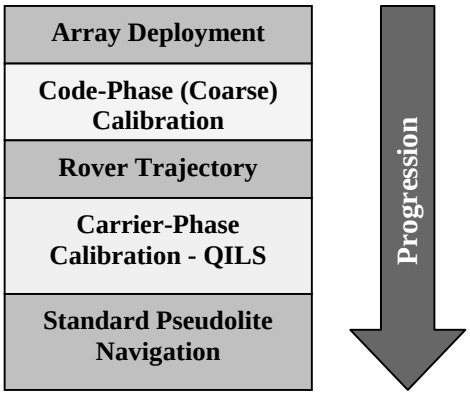

Figure 2: Self-Calibration Process

A batch process collects carrier-range data during the course of this maneuver, and is subsequently able to determine both the integers and the actual positions of the static transceivers to centimeter-level accuracy via a non-linear iterative optimization process. The next section of this paper begins an examination of the actual algorithms used for this carrier-phase self-calibration process.

At least three range measurements from the rover to the static transceivers must be available for self-calibration, and rover motion must be considerable – but not unreasonably so – for successful convergence. For example, a complete circumnavigation of the array by the rover is sufficient. Note that the rover does not have to drive a tightly defined trajectory in order to calibrate the array, since the algorithm backs out the actual rover trajectory as part of its solution. This calibration process can also be used to remove unknown line biases from the code-range solution.

LINEAR CALIBRATION ALGORITHM

The positions of receivers using conventional satellite-based GPS – both with conventional and differential techniques – is generally determined using linear algorithms. This is possible because the only major unknown system states are associated with the receiver itself, and because the great distance to the satellites makes the system largely unaffected by errors in the initial system estimate. In contrast, determination of the locations of the transceivers in an SCPA is a much more difficult task. None of the system states are known beforehand to high accuracy, and the near-field geometry makes the solution very sensitive to small changes in device location. As a consequence, completely linear algorithms are of limited effectiveness in the SCPA calibration process.

In situations where the initial array estimates are reasonably close to the true system states, pseudo-linear methods may still be useful. The most straightforward extension of standard GPS solution techniques to the challenges presented by the SCPA is iterative least squares (ILS). This is also the technique used in the previous work on SCPAs, and therefore makes a useful baseline with which to compare the latest generation of self-calibration algorithms.

During ILS the array geometry is linearized around the initial estimate of the system state (generated by code-phase ranging), and the predicted range values are compared to the actual bidirectional range measurements. Differences from this comparison are then applied to the linear system gradient, which moves the state estimate to a new location. The process is then repeated until the solution converges.

Although the baseline ILS algorithm is able to successfully self-calibrate the array when the initial estimation errors are small, it provides inadequate performance when those estimation errors become larger. Such conditions could potentially occur with an SCPA deployed on the Martian surface. This shortcoming of the traditional ILS techniques is demonstrated in the following section.

Monte-Carlo Simulation

Because of the non-linear nature of the SCPA there is no known closed-form solution for evaluating the effectiveness of the ILS solution algorithm. Extensive Monte-Carlo simulations are therefore the most effective method for determining self-calibration effectiveness under the wide range of possible array configurations.

Figure 3 presents the standard configuration used for these simulations. Three stationary transceivers – the minimum number – are arrayed in a triangle. One is at the origin and the second along the x-axis. The final stationary transceiver is nominally located at the point (0.5,1.0), although it may be moved in order to change the array shape.

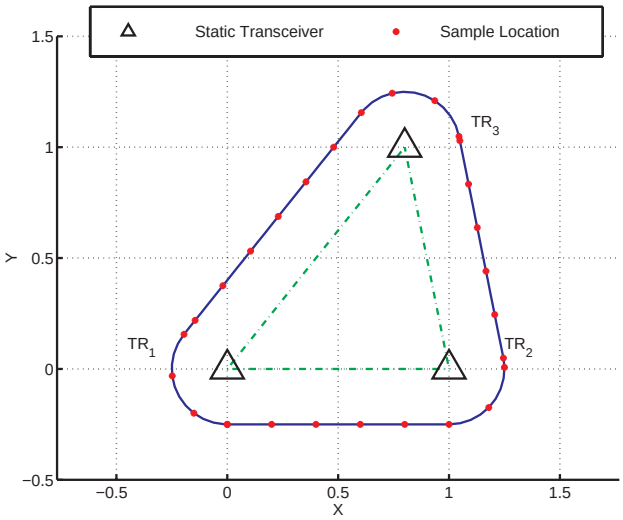

Figure 3: Standard Array Configuration

A fourth transceiver is located on a mobile platform and moves in a looping trajectory such as the one shown. The radius of the curves – and hence the proximity of the rover to the static transceivers – may be varied as well. Range measurements are taken at periodic sample points. The range measurements between the transceivers contain biases of common magnitude but random direction. Using the same magnitude range bias for each transceiver pair is physically unrealistic because the equivalent code-based errors would be mostly random, but it does allow for easy evaluation of the algorithms with respect to the bias size. The biases also affect the ranges between the stationary transceivers themselves, and therefore warp the initial estimate of the array configuration.

Figure 4 shows results from simulating the baseline ILS self-calibration algorithm. Transceiver #3 is located at its nominal position to form a near-equilateral triangle, while the bias magnitude and the trajectory curve radius is varied. As long as the mobile transceiver maintains a reasonable separation from the array itself, the ILS algorithm converges to the correct array configuration 100% of the time for biases less than or equal to 10% of the array size, and with approximately 80% success for biases equal to 20% of the array size. With greater biases, convergence drops off rapidly.

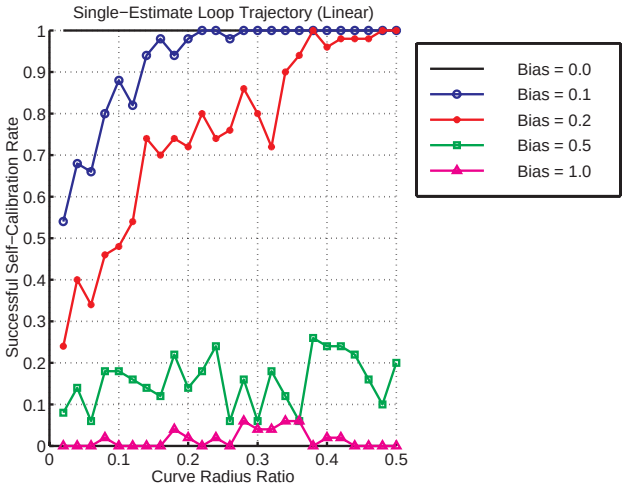

Figure 4: ILS Success vs. Curve Radius

Such convergence properties may be acceptable with larger arrays in low-multipath environments: for an array 100 meters across the initial multipath errors or other biases could be as great as 10 meters and not adversely affect the self-calibration. Smaller arrays or more severe multipath environments could lead to unsuccessful self-calibration, however, so it is important to have more effective algorithms for array self-calibration.

QUADRATIC CALIBRATION ALGORITHM

The need for more effective self-calibration algorithms for use in situations with large initial error estimates has led to a re-examination of the standard linear ILS techniques. The most important factor in the failure of the self-calibration attempts is that the linearization about the estimated states ceases to become an accurate representation of the true array geometry when those estimates are poor. It is therefore highly desirable to have a solution algorithm that retains some of these non-linear characteristics. The search for such an algorithm has led to the development of a solution technique called Quadratic Iterated Least Squares (QILS). Although this is

still not a closed-form solution and does not guarantee successful convergence, the non-linear terms included in the iteration process yield a significant increase in self-calibration effectiveness, especially in the presence of large biases.

This section presents the derivation of the QILS algorithm. Some of the notation used throughout the discussion appears below.

$\mathbf{p}^{(i)}$	≡	Position of stationary transceiver *i*
$\mathbf{p}_{m,s}$	≡	Position of mobile transceiver at sample point *s*
$r_s^{(i)}$	≡	Range between transceiver *i* and the mobile transceiver at point *s*
$b_m^{(i)}$	≡	Bias between transceiver *i* and the mobile transceiver

Derivation of the Quadratic Iterative Least Squares algorithm follows the same progression as for standard iterative least squares, with exception that 2nd-order perturbations around the estimated array configuration are included instead of merely the linear gradient. The relation for the perturbation in a single range measurement may be written as

$$\delta r_s^{(i)} \cong \nabla f\left(b_m^{(i)}, \hat{\mathbf{p}}^{(i)}, \hat{\mathbf{p}}_{m,s}\right) + \nabla^2 f\left(b_m^{(i)}, \hat{\mathbf{p}}^{(i)}, \hat{\mathbf{p}}_{m,s}\right) \quad (1)$$

where the linear term is

$$\nabla f\left(b_m^{(i)}, \hat{\mathbf{p}}^{(i)}, \hat{\mathbf{p}}_{m,s}\right) = \begin{bmatrix} 1 & \dfrac{\partial \hat{r}_s^{(i)}}{\partial \mathbf{p}^{(i)}} & \dfrac{\partial \hat{r}_s^{(i)}}{\partial \hat{\mathbf{p}}_{m,s}} \end{bmatrix} \begin{bmatrix} \delta b_m^{(i)} \\ \delta \hat{\mathbf{p}}^{(i)} \\ \delta \hat{\mathbf{p}}_{m,s} \end{bmatrix} \quad (2)$$

and the quadratic term is

$$\nabla^2 f\left(b_m^{(i)}, \hat{\mathbf{p}}^{(i)}, \hat{\mathbf{p}}_{m,s}\right)$$

$$= \frac{1}{2}\begin{bmatrix} \left[\delta \hat{\mathbf{p}}^{(i)}\right]^T & \left[\delta \hat{\mathbf{p}}_{m,s}\right]^T \end{bmatrix} \begin{bmatrix} \dfrac{\partial^2 \hat{r}_s^{(i)}}{\left[\partial \hat{\mathbf{p}}^{(i)}\right]^2} & \dfrac{\partial^2 \hat{r}_s^{(i)}}{\partial \hat{\mathbf{p}}^{(i)} \partial \hat{\mathbf{p}}_{m,s}} \\ \text{Sym.} & \dfrac{\partial^2 \hat{r}_s^{(i)}}{\left[\partial \hat{\mathbf{p}}_{m,s}\right]^2} \end{bmatrix} \begin{bmatrix} \delta \hat{\mathbf{p}}^{(i)} \\ \delta \hat{\mathbf{p}}_{m,s} \end{bmatrix}$$

$$= \frac{1}{2}\begin{bmatrix} \left[\delta \hat{\mathbf{p}}^{(i)}\right]^T & \left[\delta \hat{\mathbf{p}}_{m,s}\right]^T \end{bmatrix} \begin{bmatrix} \boldsymbol{\Psi}_s^{(i)} & -\boldsymbol{\Psi}_s^{(i)} \\ \text{Sym.} & \boldsymbol{\Psi}_s^{(i)} \end{bmatrix} \begin{bmatrix} \delta \hat{\mathbf{p}}^{(i)} \\ \delta \hat{\mathbf{p}}_{m,s} \end{bmatrix} \quad (3)$$

Computation of the second order gradient terms is greatly simplified by noting that for each range measurement, the Hessian matrix is composed of identical sub-blocks. In 2-dimensions, each of these sub-blocks is given by

$$\boldsymbol{\Psi}_s^{(i)} = \frac{1}{\left[\hat{r}_s^{(i)}\right]^3} \begin{bmatrix} \left(\Delta y_s^{(i)}\right)^2 & -\Delta x_s^{(i)} \Delta y_s^{(i)} \\ \text{Sym.} & \left(\Delta x_s^{(i)}\right)^2 \end{bmatrix} \quad (4)$$

A similar expression applies for 3-dimensional geometries.

The range equations for a single sample point s along the trajectory are easily combinable into a single equation involving the vectors of range measurements $\delta \mathbf{r}_{s,k}$ and the system states $\delta \mathbf{z}_{s,k}$ at the iteration step k.

$$\delta \mathbf{r}_{s,k} = \mathbf{G}_{s,k} \cdot \delta \mathbf{z}_{s,k} + \mathbf{Z}\left(\delta \mathbf{z}_{s,k}\right) \cdot \mathbf{H}_{s,k} \cdot \delta \mathbf{z}_{s,k} \quad (5)$$

$$\mathbf{Z}\left(\delta \mathbf{z}_{s,k}\right) \equiv \begin{bmatrix} \left[\delta \mathbf{p}^{(1)}\right]^T & \left[\delta \mathbf{p}_{m,s}\right]^T & \cdots & \mathbf{0}_{1,D} & \mathbf{0}_{1,D} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \mathbf{0}_{1,D} & \mathbf{0}_{1,D} & \cdots & \left[\delta \mathbf{p}^{(N)}\right]^T & \left[\delta \mathbf{p}_{m,s}\right]^T \end{bmatrix} \quad (6)$$

$$\mathbf{H}_{s,k} \equiv \begin{bmatrix} & \boldsymbol{\Psi}_s^{(1)} & \cdots & \mathbf{0}_D & -\boldsymbol{\Psi}_s^{(1)} \\ & -\boldsymbol{\Psi}_s^{(1)} & \cdots & \mathbf{0}_D & \boldsymbol{\Psi}_s^{(1)} \\ \mathbf{0}_{D(N+1),N} & \vdots & \ddots & \vdots & \vdots \\ & \mathbf{0}_D & \cdots & \boldsymbol{\Psi}_s^{(N)} & -\boldsymbol{\Psi}_s^{(N)} \\ & \mathbf{0}_D & \cdots & -\boldsymbol{\Psi}_s^{(N)} & \boldsymbol{\Psi}_s^{(N)} \end{bmatrix} \quad (7)$$

In these relations N is the number of stationary transceivers composing the array, and D is the dimension (2-D or 3-D) of the array.

The equations for each of the sample points may then be combined into a single global equation describing the entire array, including all of the points along the trajectory. The outward form of the resultant relation is the same as Equation 5, although it naturally contains many more terms.

$$\delta \mathbf{r}_k = \mathbf{G}_k^* \cdot \delta \mathbf{z}_k + \mathbf{Z}\left(\delta \mathbf{z}_k\right) \cdot \mathbf{H}_k^* \cdot \delta \mathbf{z}_k \quad (8)$$

At this point it is necessary to impose coordinate constraints on the global set of equations, because an SCPA only provides relative positioning within the array and does not locate the array with respect to inertial space. These constraints may be applied through Lagrange multipliers, or by simply specifying a coordinate system. For this discussion it is simplest to specify that stationary transceiver #1 is located at the origin and stationary transceiver #2 is along the x-axis.

This eliminates several columns in $\mathbf{G}_k^*$ and $\mathbf{H}_k^*$, causing both to become full rank.

Solving Equation 8 is still difficult because of the presence of the system states in the block diagonal matrix $\mathbf{Z}(\delta\mathbf{z}_k)$. Because it involves perturbations around a nominal state estimate, however, it is possible to solve it in an iterative manner using a 2-step cascaded solution method. Although the quadratic term is not represented exactly, it still exerts a strong influence on the local gradient and greatly improves the performance of the algorithm in situations with large initial estimate errors.

The first stage of the solution process is to solve the linearized system of equations as in the basic ILS algorithm. Taking the left pseudo-inverse of the first half of Equation 8 gives

$$\delta\mathbf{z}'_k = \left(\mathbf{G}_k^{*\,\mathrm{T}} \cdot \mathbf{G}_k^*\right)^{-1} \mathbf{G}_k^* \cdot \delta\mathbf{r}_k \qquad (9)$$

The resulting state perturbation estimate $\delta\mathbf{z}'_k$ is used to construct the matrix $\mathbf{Z}'(\delta\mathbf{z}'_k)$, which is then substituted into Equation 8. This results in a new linear equation of the form

$$\begin{aligned}\delta\mathbf{r}_k &= \left(\mathbf{G}_k^* + \mathbf{Z}'(\delta\mathbf{z}'_k)\cdot\mathbf{H}_k^*\right)\cdot\delta\mathbf{z}_k \\ &= \mathbf{J}_k \cdot \delta\mathbf{z}_k\end{aligned} \qquad (10)$$

which is then solvable by taking the left pseudo inverse of the resultant linear matrix.

The QILS algorithm progression is summarized in Figure 5.

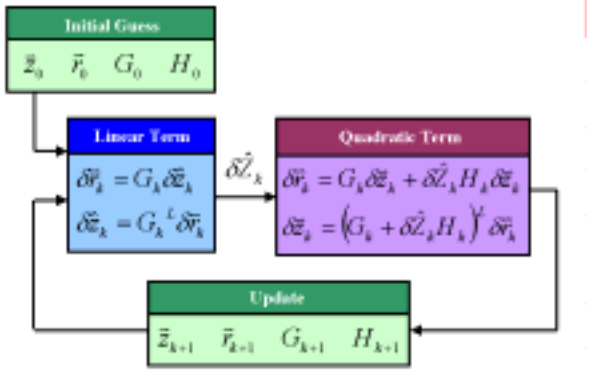

Figure 5: QILS Algorithm

Simulation Results

Figure 6 presents results from the same set of simulations as were applied to the linear ILS algorithm, but now applied to the QILS algorithm. For smaller biases (up to 20% of the array size) there is only limited improvement, because in this situation the linearization is reasonably accurate. With larger biases, however, the QILS algorithm shows marked improvement, raising the success rate for a 50% bias from under 20% to approximately 70%. Even biases as large as the array itself are now removable, although at a low success rate.

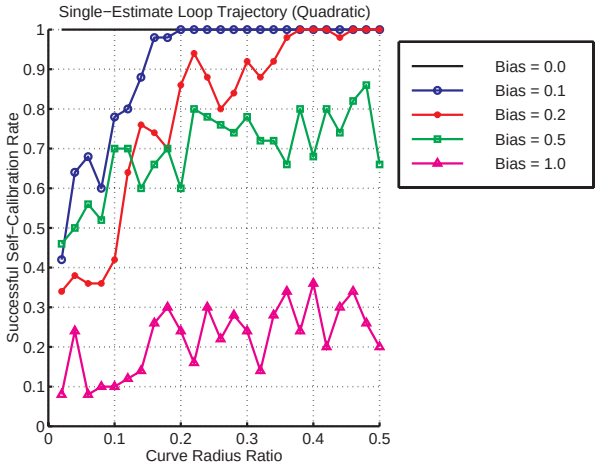

Figure 6: QILS Success vs. Curve Radius

Figure 7 shows the ability of the QILS algorithm to self-calibrate arrays of different shape than the nominal configuration used before. The curve radius and the bias size are now fixed at 0.25 units and 0.2 units, respectively, and the location of transceiver #3 now varies with respect to the rest of the array. The rover trajectory shifts to pass outside of transceiver #3. It is apparent that the self-calibration process has roughly equivalent success (approximately 80-90%) if transceiver #3 moves within a disk-shaped region nearly as large as the array itself.

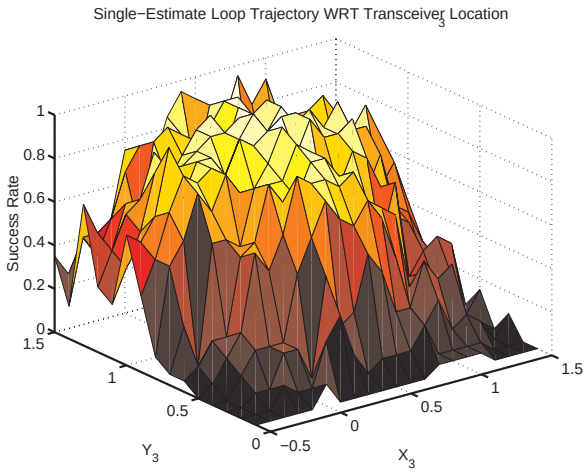

Figure 7: QILS Success vs. TR_3 Location

The improvement that the QILS algorithm provides in situations with large biases is evident when one examines the relative magnitudes of the linear and quadratic components of the solution. Figure 8 presents the ratio of the matrix norms of $\mathbf{G}_k^*$ and $\mathbf{Z}'(\delta\mathbf{z}'_k)\cdot\mathbf{H}_k^*$ from a set of Monte-Carlo simulations similar to those yielding Figure 6. It is apparent that when the biases become equal to 30% of the array size the quadratic term – which is completely ignored in the standard linear formulation – becomes as large as the linear term, and for larger biases it completely dominates the solution. (For very large biases the maximum influence of the quadratic term is actually under-represented in this plot because the cascaded solution method forces solution of the linear problem first. If the linear method completely diverges then the quadratic term can no longer be calculated and the ratio of the matrix norms is undefined.)

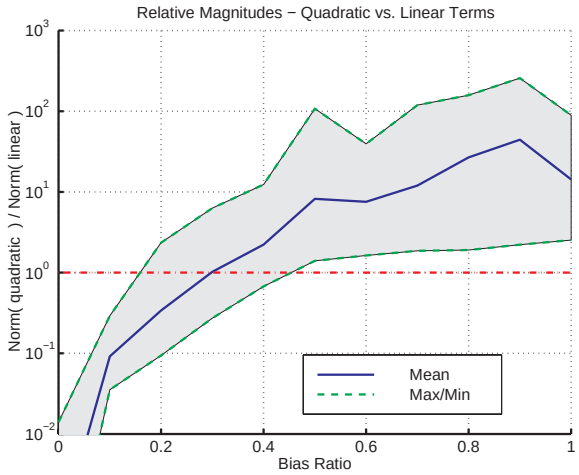

Figure 8: Linear/Quadratic ILS Comparison

MULTI-ESTIMATE ALGORITHM

Significant improvement to the QILS success rate may be made by careful examination of the modes in which it fails. Except near singularities the general failure mode is convergence to false local minima, as opposed to a general failure to converge at all. In addition, correct convergence can sometimes be highly sensitive to both the actual transceiver locations and the initial position estimate; one test showed that moving one of the stationary transceivers a distance equivalent to 1% of the array size reduced the QILS algorithm success from nearly 100% to only 50%.

It is possible to turn these apparent difficulties into very beneficial attributes through a special multiple-estimate solution algorithm. The initial positions are computed using code-based ranging as before, and then the QILS algorithm is run multiple times on the batch ranging data as usual. Before each run, however, the estimated locations of the stationary transceivers are shifted by random distances of roughly 10-20% of the array size. This change in the 'seed' values for the QILS algorithm causes some of the solutions to fall into local minima, while others will converge to the correct global minimum. A comparison of the RMS residuals ($\delta r^T \delta r$) is usually able to quickly eliminate the erroneous solutions. More sophisticated heuristics for the discriminant could further improve algorithm performance.

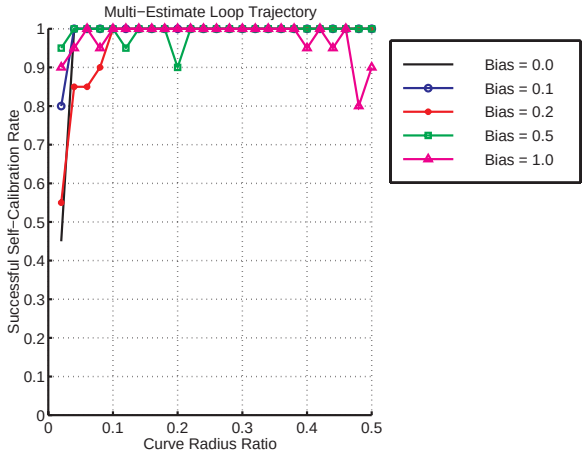

Figure 9: Multi-Estimate Success vs. Curve Radius

Figure 9 shows the QILS success rate with respect to loop radius for the same Monte-Carlo simulation performed earlier. Twenty different starting points for the algorithm were employed. In this case successful convergence to the correct solution is achieved nearly 100% of the time, except when the rover passes very near the singularities associated with the positions of the stationary transceivers.

Figure 10 presents the results of similar simulations as the location of transceiver #3 is changed (biases are ±0.2). Again, the successful convergence rate is 100% over a very wide range of array geometries.

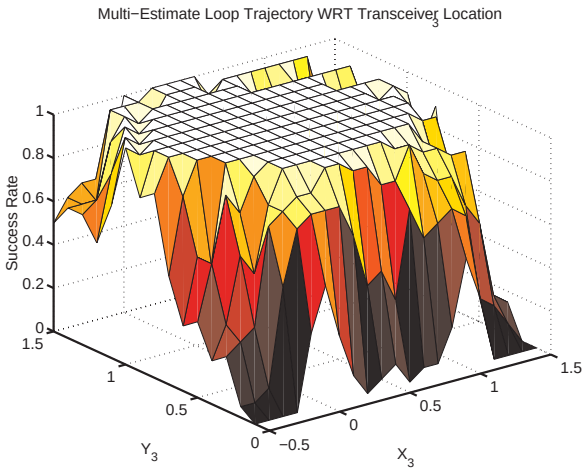

Figure 10: Multi-Estimate Success vs. TR_3 Location

EXPERIMENTAL SYSTEM

The current prototype SCPA includes four operational transceivers: three in stationary locations and one mounted on the rover. This is the minimum number of static transceivers needed for both unambiguous dynamic positioning of the rover and for the array self-calibration algorithm. System performance and robustness may be improved by adding redundant transceivers to the array. Other system components include the ground station computer and the rover testbed.

Each of the main components is described briefly below. A more comprehensive description of the experimental system appears in [1].

GPS Transceivers

Each transceiver consists of a single GPS receiver and a separate pseudolite signal generator. The receiver monitors the pseudolite output signal to form a self-differencing transceiver, as is described in [8]. The receiver is a slightly modified Mitel Orion receiver with custom tracking loops for the non-standard pseudolite data message. The pseudolite is an IntegriNautics IN200C signal generator utilizing a 3% duty cycle RTCM pulsing scheme to help combat the near/far problem associated with near-field operations. The total combined broadcast power of the current experimental system is less than 1μW. (This limit is set by the FCC, which allows users with an experimental license to intentionally broadcast on L1 with a maximum continuous power of 1μW.) The low signal power limits the range of operation of the prototype system to about 30-50 meters. Higher power levels will enable operation over baselines of kilometers, provided that line-of-sight is maintained.

Figure 11 shows the three stationary transceivers from the prototype system together with the K9 rover. The custom-built dipole broadcast and receive antennas are located on the transparent plastic plates on top of the tripods. Using dipoles instead of commercial GPS patch antennas allows 360° operation around the transceivers because of the omnidirectional pattern and the lack of circular polarization, although this comes at the penalty of losing some multipath rejection. The tote-buckets beneath the tripods hold the transceiver components themselves. In addition to the receiver and pseudolite each bucket contains a 1.6 Mbps Proxim RangeLan2 wireless link for data collection, a 4.4 A-hr NiCd battery pack which gives roughly 4 hours of continuous operation, and RF power amplifiers to improve signal acquisition and tracking.

Figure 11: Experimental System

Ground Station

The ground-station computer – a 133MHz Pentium laptop running the Windows NT operating system – runs a custom software program that collects the raw data from the transceiver wireless units, combines common-epoch measurements into ranges between transceiver pairs, and computes the corresponding array geometry. This program also allows remote control and diagnostics of the receivers.

K9 Rover

The K9 rover used for the experiments at NASA Ames is also shown in Figure 11. K9 is a variant of the FIDO rover under development at JPL for future Mars missions. It features a rocker-bogie suspension system, 360° variable steering, and an onboard dead-reckoning system. Top speed of operation is roughly 10 cm/sec. The large sensor mast holds a stereo camera pair used for terrain mapping. A scanning laser rangefinder is mounted on the front of the rover for obstacle detection. The GPS antennas used for the onboard transceiver are located on a small plastic plate at the back corner of the rover.

Test Site

Testing at Ames is done in a large empty lot near the inlet of the 80’ by 120’ subsonic wind tunnel. This lot is surrounded by a chain-link fence, yielding a moderately high multipath environment. Figure 12 shows the experimental system in operation, including all three static transceivers (placed in a triangle approximately 20 meters apart) and the K9 rover.

Figure 12: NASA Ames Test Site

EXPERIMENTAL RESULTS

Experimental results from a successful self-calibration of the prototype SCPA at NASA Ames are presented in [9]. Because these tests demonstrate the effectiveness of the QILS algorithm for situations with large initial estimation errors, these results are reviewed below. This section also provides a description of the actual experimental procedure followed during these tests.

Code-Phase Calibration

The testing process for the SCPA follows the same steps as the general self-calibration progression described earlier. The stationary transceivers are arrayed in the test area in an triangular configuration 20 meters to a side. The rover starts outside of the array close to one edge of the triangle. The locations of these transceivers are pre-surveyed to provide a truth metric; knowledge of these positions, however, is not used at any time during the self-calibration process. Once the array is in place, averaged code-range measurements between the transceivers are used to generate an initial estimate of the transceiver locations. Figure 13 shows the transceiver locations as determined by this coarse calibration step. The actual locations are at the corners of the large dotted triangle, while the true rover starting position is at the small circle underneath the triangle.

Table 1 shows the corresponding position errors. The errors for the stationary transceivers are 2.76 meters RMS, an acceptable result for code-phase positioning. The positioning error for K9 is greater than 20 meters, however, most likely due to strong multipath from the surrounding fence. With such a small array, errors of this magnitude would greatly cripple the navigational effectiveness unless they were removed during the self-calibration process.

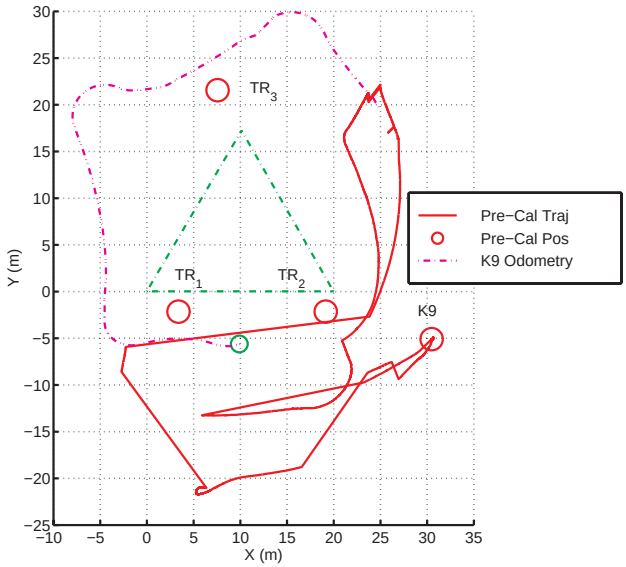

Figure 13: Pre-Calibration Positions and Trajectory

Table 1: Initial (Code-Phase) Position Errors

Transceiver	X (m)	Y (m)
1	3.38	-2.13
2	-0.87	-2.13
3	-2.52	4.27
K9	20.42	0.60

Calibration Trajectory

The K9 rover then circumnavigates the array to provide the geometry change needed for self-calibration. The overall trajectory is approximately 100 meters in length, and takes 20 minutes to complete. During the trajectory carrier-range data is collected between the transceiver on K9 and each of the stationary transceivers, the carrier phase integers having been estimated beforehand from the results of the code-phase calibration. Figure 13 presents the rover trajectory as determined from these carrier-phase range measurements. Rather than a smooth loop around the array, the large errors in the integer estimates have produced an almost unrecognizable hash of segments and jumps. For comparison the path computed by the wheel encoders onboard K9 is presented as the dashed line. Although the encoder data does not show K9 returning to the starting point like it does in the true trajectory, the character of the loop is readily apparent.

Carrier-Phase Calibration

The self-calibration algorithm mentioned earlier is now applied to the range data collected during the preceding trajectory. Even with such large initial errors in the state estimate, the algorithm successfully converges to the correct array geometry and rover path. Figure 14 shows the results, and the corresponding errors in the locations of the stationary transceivers are displayed in Table 2.

The calculated trajectory now matches the true trajectory to within centimeters, and the RMS position errors for the stationary transceivers have been reduced to 4.2 cm RMS. The error associated with K9 is slightly higher because a hardware failure in one of the receivers during testing caused a loss of clock synchronization, creating a slight drift in the measured ranges. Overall, the self-calibration process reduced the positioning errors in this experiment by three orders of magnitude.

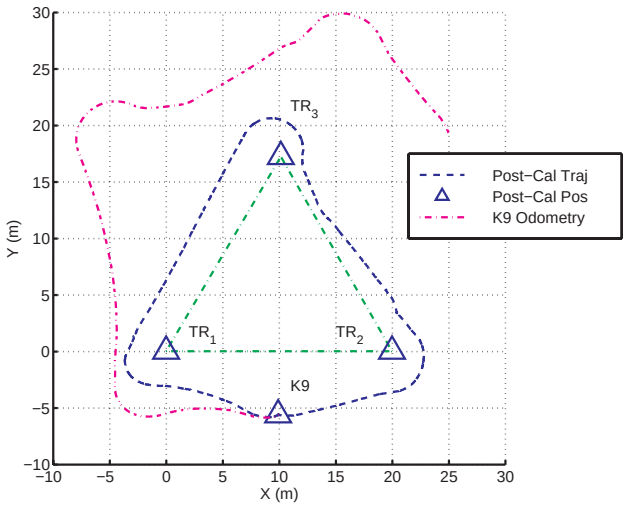

Figure 14: Post-Calibration Positions and Trajectory

Table 2: Final (Carrier-Phase) Position Errors

Transceiver	X (m)	Y (m)
1	-0.01	0.03
2	-0.05	0.03
3	0.06	-0.05
K9	-0.16	0.09

CONCLUSIONS

The solution methods described in this paper greatly improve the performance of Self-Calibrating Pseudolite Arrays, yielding successful surveying of the array elements even in the presence of large relative code-phase errors such as occur with small arrays or in high-multipath environments. The QILS algorithm, by including the more dominant second-order effects, increases the calibration success rate by a factor of 2-3 times for intermediate-sized biases (i.e. up to 50% of the array size). The improvement to approximately 20% raw calibration success for even larger biases (up to 100% of the array size) allows the incorporation of the stochastic multiple-estimate solution method, yielding almost universal self-calibration success for most conceivable array configurations. These improvements make it more practical to implement an SCPA as a critical component in an integrated Mars-based navigation system.

The field tests conducted at NASA Ames using the K9 rover and the Stanford SCPA testbed demonstrate both successful transceiver operations under realistic near/far conditions and successful carrier-phase self-calibration of the array to better than 5 cm RMS. Rover position during the trajectory is computed to similar accuracy. Much of this error was in fact due to a minor hardware failure during the testing, and it is therefore expected that self-calibration accuracy in a properly synchronized system would be near that of the raw carrier-phase tracking loops themselves.

Although the Self-Calibrating Pseudolite Array described in this paper is capable of providing extremely accurate and repeatable navigation without any additional augmentation, an SCPA on the Martian surface would ideally be used in conjunction with a complementary set of sensors in order to provide additional information beyond the scope of the raw GPS-based position data. Computer vision or scanning lasers, for example, would be required for obstacle detection and avoidance, and would also be useful for fine servoing control. Additionally, blending the SCPA navigation data with an inertial navigation or dead reckoning system would provide an additional level of robustness in case of GPS cycle slips or signal loss due to intervening terrain or other obstacles. Because of its capability for centimeter-level, drift-free positioning for multiple agents, an SCPA would be a critical enabling technology for such an integrated sensing and navigation system.

ACKNOWLEDGEMENTS

This research has been conducted under NASA grants NCC2-1154 and NAG2-1464 as part of a joint effort between the Stanford University Aerospace Robotics Laboratory and the NASA Ames Research Center. We would especially like to thank Maria Bualat, Mike Fair, and Anne Wright at ARC for their efforts in making these joint field tests possible.

REFERENCES

[1] LeMaster, E., Rock, S., ”Field Test Results for a Self-Calibrating Pseudolite Array”, Proceedings of the Institute of Navigation GPS-2000 Conference, Salt Lake City, UT, Sept. 2000, pp. 1046-1055.

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

[3] Zimmerman, Kurt, *Experiments in the Use of the Global Positioning System for Space Vehicle*

Rendezvous, Ph.D. Thesis, Stanford University, December 1996.

[4] 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.

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

[6] Parkinson, B. et al., ed., Global Positioning System: Theory and Applications, Vols. I & II, American Institute of Aeronautics and Astronautics, 1996.

[7] LeMaster, E.A., Rock, S.M., "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.

[8] 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.

[9] LeMaster, E.A., Rock, S.M., "Field Demonstration of a Mars Navigation System Utilizing GPS Pseudolite Transceivers", Proceedings of the 2002 IEEE Position, Location, and Navigation Symposium, Palm Springs, CA, April 2002.