

Slip-compensated path following for planetary exploration rovers

DANIEL M. HELMICK^{1,*}, STERGIOS I. ROUMELIOTIS², YANG CHENG¹,
DANIEL S. CLOUSE¹, MAX BAJRACHARYA¹ and LARRY H. MATTHIES¹

¹*Jet Propulsion Laboratory, 4800 Oak Grove Drive, Pasadena, CA 91109, USA*

²*Department of Computer Science and Engineering, University of Minnesota, 200 Union Street SE,
Minneapolis, MN 55455, USA*

Received 3 March 2006; accepted 14 April 2006

Abstract—A system that enables continuous slip compensation for a Mars rover has been designed, implemented and field-tested. This system is composed of several components that allow the rover to accurately and continuously follow a designated path, compensate for slippage and reach intended goals in high-slip environments. These components include visual odometry, vehicle kinematics, a Kalman filter pose estimator and a slip-compensated path follower. Visual odometry tracks distinctive scene features in stereo imagery to estimate rover motion between successively acquired stereo image pairs. The kinematics for a rocker–bogie suspension system estimates vehicle motion by measuring wheel rates, and rocker, bogie and steering angles. The Kalman filter processes measurements from an inertial measurement unit and visual odometry. The filter estimate is then compared to the kinematic estimate to determine whether slippage has occurred, taking into account estimate uncertainties. If slippage is detected, the slip vector is calculated by differencing the current Kalman filter estimate from the kinematic estimate. This slip vector is then used to determine the necessary wheel velocities and steering angles to compensate for slip and follow the desired path.

Keywords: Rover navigation; visual odometry; slip compensation; Kalman filter; rover kinematics.

1. INTRODUCTION

This paper describes the design, implementation and experimental results of an integrated system for Mars rover navigation in high-slip environments (Fig. 1). The presented algorithms enable the rover to accurately follow a designated path, compensate for slippage and reach intended goals independent of the terrain over which it traverses (within the mechanical constraints of the mobility system). The proposed system is comprised of several key components that were developed

*To whom correspondence should be addressed. E-mail: dmh@jpl.nasa.gov



Figure 1. Rocky 8 on a sandy slope.

and refined for this task and are described in detail below. These components include visual odometry, full vehicle kinematics, a Kalman filter pose estimator and a slip-compensated path-following algorithm. Figure 2 provides a high-level functional block diagram of the system. Visual odometry is an algorithm that relies on stereo imagery to estimate rover motion independent of mechanical terrain properties and is described in Section 2. The full vehicle kinematics, described in Section 3, uses position sensor inputs from the joints and wheels of the rocker-bogie mobility system (see Fig. 4) to estimate rover motion. The Kalman filter combines inertial measurements from the onboard inertial measurement unit (IMU) with motion estimates from the visual odometry to estimate rover motion at high sampling rates and is detailed in Section 4. Since both the IMU measurements and the visual odometry estimates are independent of the vehicle's interaction with the ground, the motion estimates from the Kalman filter can be compared to those computed based on the vehicle kinematics (which are highly dependent upon the vehicle's interaction with its environment) to determine if any statistically significant slippage has occurred. If there is no slippage, the vehicle kinematic motion estimates can contribute to the Kalman filter pose estimates. If, however, slippage is detected, then the kinematic estimate and the Kalman filter estimate are differenced, resulting in a rover slip vector. This slip vector is then used by the path-following algorithm to calculate the rover velocity commands required for following

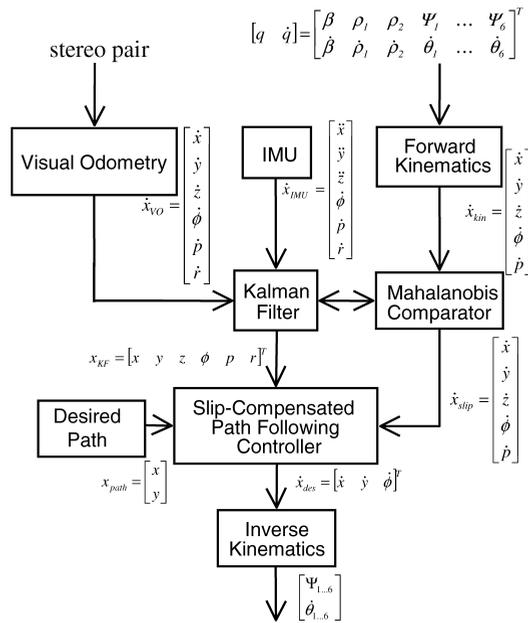


Figure 2. Block diagram of the slip-compensated path-following system.

a path while compensating for slip. This algorithm is described in greater detail in Section 6.

The individual components of the system as well as a simplified integrated system have been tested onboard a rover. Several independent tests were performed using Rocky 8 (see Fig. 1), a Mars rover research platform. In the first test, visual odometry was tested onboard the rover in the JPL Mars Yard over two 25-m traverses. Under nominal conditions, wheel odometry accuracy is not better than 10% of the distance traveled and, in higher-slip environments, it can be significantly worse. Results from our first tests showed that at least 2.5% accuracy can be achieved by using visual odometry, regardless of the mechanical soil characteristics. The second test was a field test that used the slip compensation system described above, minus the Kalman filter. This test was a traverse of over 50 m on sandy slopes. The third set of tests was performed on a tiltable platform measuring 5×5 m. In this last set of tests, the continuous rover-slip compensation algorithm (see Section 6) was employed instead of stopping the rover to estimate slippage, as was the case in previous tests. Results from these three sets of experiments are provided in Section 7. This paper extends the work described in Refs [1, 2]. Related work includes rover trajectory generation [3], rover navigation [4–6] and path following [7]. To the best of our knowledge this is the first time that an integrated robotic system for slip detection and compensation is being presented and tested in realistic conditions.

2. VISUAL ODOMETRY ALGORITHM

Long-distance mobile robot navigation on a remote planetary body requires an accurate method for position estimation in an unknown or poorly known environment. Visual odometry, or image-based ego-motion estimation, was originally developed by Matthies [8]. Following this work, a number of variations and modifications have been suggested for improving its robustness and accuracy [9, 10]. The key idea of this method is to determine the change in position and attitude by solving for the maximum likelihood estimate of the transformation between a selection of three-dimensional (3-D) features extracted from consecutive stereo images. The basic steps of this algorithm are described below. A more detailed description of the ensuing derivations can be found in Ref. [2].

2.1. Feature selection

The first step in the visual odometry algorithm is to select features that can be easily matched between stereo pairs and tracked across a single image step. To achieve this, the Forstner interest operator [12] is applied to the left image of the first stereo pair. The pixels with higher interest values are better features. In order to ensure that the detected visual points are evenly distributed across the image scene, a minimum distance between any two features is enforced. Furthermore, the image scene is divided into grids, with the grid size significantly larger than the minimum distance between features, so as to reduce the volume of data that needs to be sorted. The one feature with the highest interest value in each grid is selected as a viable candidate. Finally, a fixed number of these candidates with the highest interest values which also meet the minimum distance constraint is selected.

2.2. Feature gap analysis and covariance computation

The 3-D positions of the selected features are determined by stereo matching. A template around each feature in the left image is correlated to a location in the right image. Knowing the location of the features in the left and right images, a ray corresponding to the feature can be projected out of each camera. Under perfect conditions, the rays of the same feature from the left and right images should intersect in space. However, due to image noise and matching error, they do not always intersect. The gap (the shortest distance between the two rays) indicates the goodness of the stereo matching. Features with large gaps are eliminated from further processing. Additionally, the error model is a function of the gap. This effect is incorporated in the covariance matrix computation described below.

Assuming the stereo cameras are located at $C_1(X_1, Y_1, Z_1)$ and $C_2(X_2, Y_2, Z_2)$ (Fig. 3), r_1 and r_2 are two unit rays from the same feature in both images. As mentioned before, due to noise, r_1 and r_2 do not always intersect in space. The stereo point is selected to be the midway between the closest points of the two rays.

Assuming the closest points between the two rays are P_1 and P_2 we have:

$$P_1 = C_1 + r_1 m_1,$$

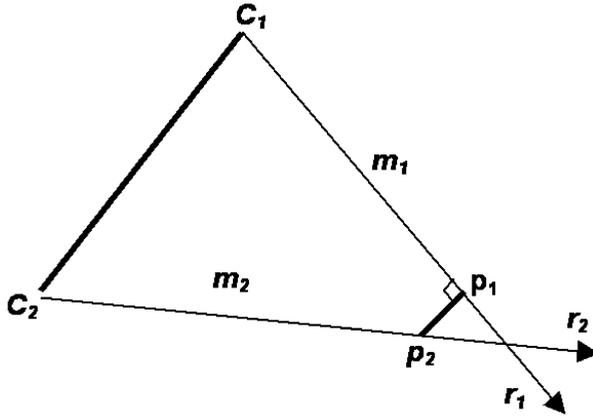


Figure 3. Feature gap.

$$P_2 = C_2 + r_1 m_2, \quad (1)$$

where m_1 and m_2 are the lengths of $P_1 C_1$ and $P_2 C_2$, respectively. Since P_1 and P_2 represent the same point in space, we have:

$$\begin{aligned} 0 &= (P_2 - P_1) \cdot r_1 = (C_2 - C_1 + r_2 m_2 - r_1 m_1) \cdot r_1, \\ 0 &= (P_2 - P_1) \cdot r_2 = (C_2 - C_1 + r_2 m_2 - r_1 m_1) \cdot r_2, \end{aligned} \quad (2)$$

Solving for m_1 and m_2 yields:

$$\begin{aligned} m_1 &= \frac{1}{\delta} ((r_1 \cdot r_2) r_2 - r_1) \cdot (C_2 - C_1), \\ m_2 &= \frac{1}{\delta} (r_2 - (r_1 \cdot r_2) r_1) \cdot (C_2 - C_1), \end{aligned} \quad (3)$$

where $\delta = (r_1 \cdot r_2)^2 - 1$. The coordinates of point P are given by:

$$P = (P_1 + P_2)/2, \quad (4)$$

where P_1 and P_2 are computed by substituting (3) into (1). The covariance for the position of point P is:

$$\Sigma_P = J \begin{bmatrix} \Sigma_l & 0 \\ 0 & \Sigma_r \end{bmatrix} J^T, \quad (5)$$

where J is the Jacobian matrix, i.e., the matrix of partial derivatives of P with respect to the 2-D feature locations in the left and right images, and Σ_l and Σ_r are 2×2 matrices whose elements are the curvatures of the biquadratic polynomial along the vertical, horizontal and diagonal directions, which can be obtained directly from subpixel interpolation. The quality of a 3-D feature is a function of its relative location, the gap between the two stereo rays and the sharpness of the correlation peak. This covariance computation fully reflects these three factors. For a more detailed derivation, see Ref. [2].

2.3. Feature tracking

After the rover moves some distance, a second pair of stereo images is acquired. The features selected from the previous image are then projected into the second pair using the knowledge of the approximated motion provided by the onboard wheel odometry (forward kinematics). The features are first matched in the new left image by searching an area around the projected feature locations. Stereo matching is then performed on these tracked features on the second pair to determine their new 3-D positions. As the 3-D positions of those tracked features are already known from the previous step, the stereo matching search range can be greatly reduced. Features whose initial and final 3-D positions differ by too large an amount are filtered out.

2.4. Motion estimation

Given two sets of corresponding 3-D features, the transformation between them is determined using a motion estimation algorithm that takes into account the matching covariance of each feature. This algorithm is decomposed into two sequential steps. Coarse motion is first computed with Schonemann motion estimation and then a more accurate motion estimate is determined by maximum likelihood (ML) motion estimation.

Schonemann motion estimation [13] uses singular value decomposition (SVD) with an orthogonal constraint to estimate the rotation matrix, R , and translation vector, T , that transforms the feature positions in I_1 to those found in I_2 . The Schonemann method is simple and fast; however, it is highly unstable when large errors are involved. In order to overcome this problem, a least-median-of-squares approach [14] is adopted. In this method, a subset of features is randomly selected. Then each feature from the previous frame is projected to the current frame, and the distance error between that projection and the position of the corresponding feature in I_2 is computed. The total count of features under a given error tolerance is calculated. This procedure is repeated multiple times. The motion with the largest number of agreeable features is chosen as the best motion.

The best motion estimate determined by employing the above procedure is refined using ML motion estimation. The ML estimation algorithm takes into account the 3-D feature position differences and the associated error models in order to estimate motion. Let $P_j^{(1)}$ and $P_j^{(2)}$ be the observed feature positions before and after a robot motion. Then, we have:

$$P_j^{(2)} = RP_j^{(1)} + T + e_j, \quad (6)$$

where $R = R(\Theta)$ and T are the rotation and translation of the robot, and e_j is the combined errors in the observed positions of the j th features. In this estimation process, the three-axis rotations (Θ) and translation (T) are directly determined by minimizing the summation in the exponents $\sum_j e_j^T W_j e_j$ where $e_j = P_j^{(2)} - RP_j^{(1)} - T$ and W_j is the inverse covariance matrix of e_j . The minimization of this nonlinear problem is implemented by iterative linearization [8]. Two nice properties of

ML estimation make this algorithm significantly more accurate compared to the Schonemann method. First, it estimates the three-axis rotations (Θ) directly so that it eliminates the error caused by rotation matrix estimation (which occurs with least-squares estimation). Second, it incorporates error models for the feature positions in the estimation process (see (5)), which greatly improves the achieved accuracy.

3. KINEMATIC ALGORITHMS

The forward and inverse kinematics were determined for the rover (Rocky 8) shown in Fig. 2. The forward kinematics of the vehicle are used for estimating rover motion given the wheel rates, and rocker, bogie and steering angles. The inverse kinematics of the vehicle are employed for calculating the necessary wheel velocities and steering angles to create a desired rover motion.

These algorithms are specific to the rocker–bogie configuration with six steerable wheels (Fig. 4). However, the techniques used to derive these algorithms could be applied to any vehicle configuration, though there may be a fewer number of observable degrees of freedom (d.o.f.) for different configurations. Additionally, these forward kinematic algorithms can be used directly for rovers with a subset of functionality (e.g., a rocker–bogie rover with only four of six wheels capable of steering, such as the Mars Exploration Rovers (MERs), or a rover with only four wheels with or without a rocker, etc.) simply by restricting the relevant parameters to be constant.

The motivation for developing the full kinematics of this class of vehicles, rather than adopting the more common planar simplification, is 2-fold. First, it allows for the observation of 5 d.o.f., whereas the planar approximation limits this to 3 d.o.f. Second, as terrain becomes rougher, the errors due to the planar assumption grow.



Figure 4. Rocky 8.

These errors can become significantly large (up to 30% of the distance traveled) and affect the slip calculations and, consequently, the slip compensation controller.

The formulation of the forward and inverse kinematics closely follows that of Refs [15, 16], with significant extensions being made for six-wheel steering. A detailed description of the kinematic derivations can be found in Refs [2, 15–17].

3.1. Rocker–bogie configuration

The rocker–bogie configuration is a suspension system that is commonly used for planetary rovers and their prototypes. The configuration analyzed in this work consists of 15 d.o.f.: six steerable/drivable wheels (12 d.o.f.), a rocker and two bogies (all of these d.o.f. are sensed using either encoders or potentiometers). It is beyond the scope of this paper to describe the benefits of such a mobility system; the interested reader is referred to Ref. [18] for a discussion on this topic. In what follows, we show that under two assumptions the rocker–bogie system allows for the observation of 5 of the 6 d.o.f. of the rover. These assumptions are: (i) the wheel/terrain contact point is in a constant location relative to the wheel axle, and (ii) slip between the wheel and the terrain only occurs about the steering axis (e.g., no side or rolling slip). These slip assumptions, however, are only made for the kinematics algorithm and not for the slip compensation system as a whole. The second assumption is what makes the difference between the Kalman filter motion estimate and the kinematic motion estimate a measurement of vehicle slippage.

3.2. Denavit–Hartenberg (D–H) table formulation

The D–H convention was adopted for defining the frames of each of the 15 d.o.f. [19]. From the frame definitions (Fig. 5), a unique set of D–H parameters can be derived that completely describes the kinematics of the rover. From these parameters, wheel Jacobians are computed as described in Ref. [2].

3.3. Forward kinematics

Once the wheel Jacobians are known, rover motion estimation can be performed using the least-squares formulation:

$$\begin{bmatrix} v \\ \dot{\eta} \end{bmatrix} = (A^T A)^{-1} A^T \cdot J_{\text{comp}} \dot{q}_{\text{comp}}, \quad (7)$$

where $v = [\dot{x} \ \dot{y} \ \dot{z} \ \dot{\phi} \ \dot{p} \ \dot{r}]^T$ is the vector of rover linear and rotational velocities, $\dot{\eta}$ is a vector of unobservable wheel turning rates (the sum of the rotational slip rate and the steering rate), A is the matrix of the unsensed elements of the Jacobians, J_{comp} is the block-diagonal composite matrix of the wheel Jacobians and \dot{q}_{comp} is the composite vector of the measured joint kinematic rates. For details of the derivation and simplified solution to this least-squares problem, see Refs [15, 17]. Note that it is not necessary to actually perform the inversion

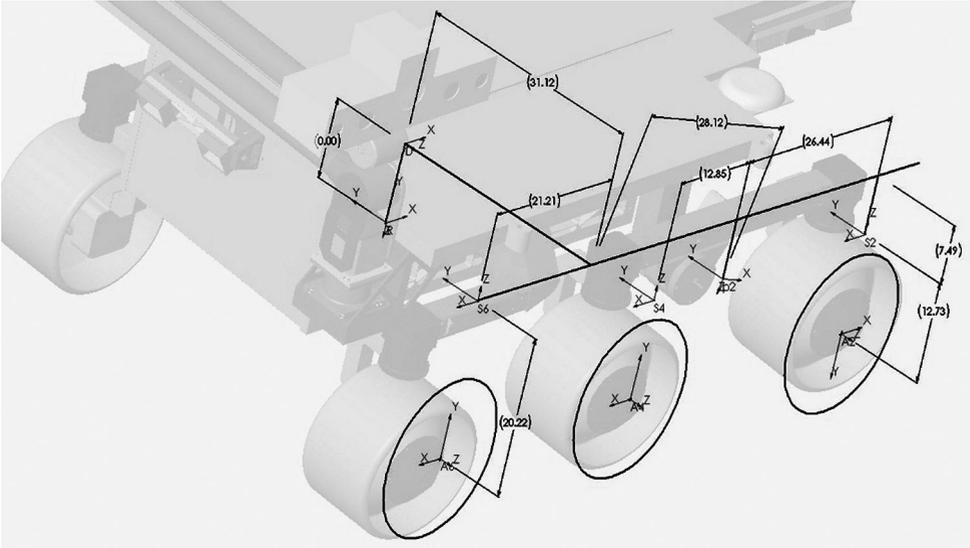


Figure 5. Coordinate frame definition for the right side of the rover (all dimensions in cm).

of $A^T A$. The matrix equations can be greatly simplified algebraically to make these computations significantly more efficient.

3.4. Inverse kinematics

As depicted in Fig. 2, the inverse kinematics algorithm receives as input the commanded rover motion, and the current kinematic angles and angle rates, and produces six steering angles and six wheel rates. An interesting feature of the six steerable wheels is the fact that this configuration creates a holonomic rover, under the assumption of instantaneous steering. Consequently, all three controllable d.o.f. of the rover, $[\dot{x}, \dot{y}, \dot{\phi}]$, are independent, which enables the isolation of several different control loops as described in Section 6.2.

The first step of the inverse kinematics algorithm is to calculate an instantaneous center of rotation, $[x_O \ y_O]$, in the rover frame, by employing the following equations

$$\begin{aligned} x_O &= \dot{y}_{\text{cmd}} / \dot{\phi}_{\text{cmd}}, \\ y_O &= \dot{x}_{\text{cmd}} / \dot{\phi}_{\text{cmd}}, \end{aligned} \quad (8)$$

where \dot{x}_{cmd} , \dot{y}_{cmd} and $\dot{\phi}_{\text{cmd}}$ are the commanded vehicle velocities from the slip-compensated path follower (see Section 6). As the z -axis of each wheel is the steerable axis, the instantaneous center of rotation needs to be projected from the rover frame $\{R\}$ to the x - y plane of the steering contact coordinate frame $\{M_i\}$, $i = 1 \dots 6$, by employing the first two rows of the transformation matrix $T_R^{M_i}$

between these frames:

$$\begin{bmatrix} x_O \\ y_O \end{bmatrix}_{M_i} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix} \cdot T_R^{M_i} \cdot \begin{bmatrix} x_O \\ y_O \\ 0 \\ 1 \end{bmatrix}_R. \quad (9)$$

The steering angle for each wheel is then calculated using:

$$\psi_i = \arctan(x_{O,M_i}/y_{O,M_i}), \quad i = 1 \dots 6, \quad (10)$$

and each wheel rate is determined as:

$$\dot{\theta}_i = [J_{ai}^T \Delta(J_{ui}) J_{ai}]^{-1} J_{ai}^T \Delta(J_{ui}) E \cdot v_{\text{cmd}}, \quad i = 1 \dots 6, \quad (11)$$

where:

$$\Delta(J_{ui}) = J_{ui} (J_{ui}^T J_{ui})^{-1} J_{ui}^T - I \quad (12)$$

$$E = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix}^T \quad (13)$$

$$v_{\text{cmd}} = [\dot{x}_{\text{cmd}} \quad \dot{y}_{\text{cmd}} \quad \dot{\phi}_{\text{cmd}}]^T, \quad (14)$$

and J_{ai} and J_{ui} are the actuated and unactuated Jacobians defined in Ref. [17]. Note that (11) is the actuated inverse solution from Muir and Neumann [17]. As before, these matrix inversions can be algebraically simplified so that each wheel rate calculation is relatively simple and computationally efficient.

4. KALMAN FILTER

In this section, we present our approach for estimating the position and orientation of the rover using inertial measurements from the IMU and relative pose (position and orientation measurements) from visual odometry and vehicle odometry (forward kinematics). Since our formulation is based on sensor modeling, we use the indirect form of the extended Kalman filter (EKF) that estimates the errors in the estimated states instead of the states themselves. The interested reader is referred to Refs [21, 23, 24] for a detailed description of the advantages of the indirect *versus* the direct Kalman filter. Within this framework, the IMU measurements are integrated in order to propagate the state estimate [20, 22], while the visual odometry and vehicle kinematics (only when no slip has occurred) are employed for updating the state estimate and providing periodic corrections. The equations of the EKF designed for this particular nonlinear system are described in detail in Ref. [22].

4.1. System propagation model

The state vector of interest in this estimation problem is:

$$x^T = [q^T \quad b_g^T \quad v^T \quad b_a^T \quad p^T], \quad (15)$$

where q is the 4×1 quaternion of rotation that represents the attitude of the vehicle, v and p are the 3×1 vectors of linear velocity and position of the rover, and b_g and b_a are the 3×1 vectors of biases in the gyroscope and accelerometer signals. The corresponding error state vector is:

$$\Delta x^T = [\delta\theta^T \quad \Delta b_g^T \quad \Delta v^T \quad \Delta b_a^T \quad \Delta p^T], \quad (16)$$

where $\Delta\zeta = \zeta - \hat{\zeta}$ is the algebraic difference (error) between the real value of a state $\zeta \in \{b_g, v, b_a, p\}$ and its estimate $\hat{\zeta}$, and $\delta\theta$ is the 3×1 error vector of the tilt angles determined based on the small angle approximation:

$$q = \delta q \otimes \hat{q}, \quad \delta q^T = [(1/2)\delta\theta^T \quad 1], \quad (17)$$

The continuous-time equation for the error-state propagation is:

$$\delta\dot{x}(t) = F_c(t)\delta x(t) + G_c(t)w(t), \quad (18)$$

where $w(t)$ is the white zero-mean continuous time noise process affecting the IMU signals:

$$F_c(t) = \begin{bmatrix} -[\hat{\omega}] & -I_{3 \times 3} & 0_{3 \times 3} & 0_{3 \times 3} & 0_{3 \times 3} \\ 0_{3 \times 3} & 0_{3 \times 3} & 0_{3 \times 3} & 0_{3 \times 3} & 0_{3 \times 3} \\ -C^T(\hat{q})[\hat{a}] & 0_{3 \times 3} & 0_{3 \times 3} & -C^T(\hat{q}) & 0_{3 \times 3} \\ 0_{3 \times 3} & 0_{3 \times 3} & 0_{3 \times 3} & 0_{3 \times 3} & 0_{3 \times 3} \\ 0_{3 \times 3} & 0_{3 \times 3} & I_{3 \times 3} & 0_{3 \times 3} & 0_{3 \times 3} \end{bmatrix},$$

$$G_c(t) = \begin{bmatrix} -I_{3 \times 3} & 0_{3 \times 3} & 0_{3 \times 3} & 0_{3 \times 3} \\ 0_{3 \times 3} & I_{3 \times 3} & 0_{3 \times 3} & 0_{3 \times 3} \\ 0_{3 \times 3} & 0_{3 \times 3} & -C^T(\hat{q}) & 0_{3 \times 3} \\ 0_{3 \times 3} & 0_{3 \times 3} & 0_{3 \times 3} & I_{3 \times 3} \\ 0_{3 \times 3} & 0_{3 \times 3} & 0_{3 \times 3} & 0_{3 \times 3} \end{bmatrix}.$$

$[\hat{\omega}]$ ($[\hat{a}]$) is the skew-symmetric matrix of the vector $\hat{\omega} = \omega_m - \hat{b}_g$ ($\hat{a} = a_m - \hat{b}_a$), ω_m (a_m) is the measured rotational velocity (linear acceleration) vector and $C^T(\hat{q})$ is the rotational matrix that projects vectors from the sensor to the global coordinate frame. By discretizing (18) we obtain:

$$\delta x_{k+1} = F_k(t_{k+1}, t_k)\delta x_k + w_k, \quad (19)$$

where $F_k(t_{k+1}, t_k) = e^{\int_{t_k}^{t_{k+1}} F_c(\tau) d\tau}$, and w_k is the equivalent discrete-time white zero-mean noise process with covariance $Q_k = \int_{t_k}^{t_{k+1}} F_k(t_{k+1}, \tau) G_c(\tau) Q_c G_c^T(\tau) F_k^T(t_{k+1}, \tau) d\tau$, Q_c is a block diagonal matrix with matrix elements $\sigma_i^2 I_{3 \times 3}$, $i = 1 \dots 4$, and σ_i^2 is the power of the gyroscope, gyroscope-bias, accelerometer and

accelerometer-bias continuous-time noise components, respectively. The interested reader is referred to Ref. [22] for the details of the derivation of (18) and (19).

4.2. Relative pose measurement model

The motion measurement provided by the visual odometry and/or the forward kinematics of the vehicle corresponds to a relative pose measurement, i.e., the algebraic difference in position z_p and attitude quaternion z_q . In order to process measurements that relate state estimates at different time steps (e.g., t_k and t_{k+m}), the state vector needs to be augmented with a copy of the vehicle state estimate at t_k . This process enables the filter to explicitly account for the correlations of the vehicle state estimates between the time instants that, for example, two sets of images were recorded [25]. In what follows, we assume that at time t_k the vehicle is at position ${}^G p(t_k) = p_1$ with quaternion attitude ${}^1_G q(t_k) = q_1$, and after m steps it has moved to position ${}^G p(t_{k+m}) = p_2$ with attitude ${}^2_G q(t_{k+m}) = q_2$. Frames $\{G\}$, $\{1\}$ and $\{2\}$ are the inertial frames of reference attached to the vehicle at times t_0 , t_k and t_{k+m} respectively.

The errors in the relative position and attitude (pose) measurements are given by:

$$\begin{aligned} \Delta \tilde{z}_{k+m} &= \begin{bmatrix} \Delta z_p \\ \Delta \tilde{z}_q \end{bmatrix} = \mathcal{X} \begin{bmatrix} \Delta z_p \\ \Delta z_q \end{bmatrix} \\ &= \Gamma \begin{bmatrix} J_{x_1} & J_{x_2} \end{bmatrix} \begin{bmatrix} \Delta x_1 \\ \Delta x_2 \end{bmatrix} + \mathcal{X} n_r = H \begin{bmatrix} \Delta x_1 \\ \Delta x_2 \end{bmatrix} + \tilde{n}_r, \end{aligned} \quad (20)$$

where Γ is a 6×6 block diagonal matrix with both 3×3 matrix elements equal to the rotational matrix ${}^1_G C(q)$, and J_{x_1} , J_{x_2} are the measurement Jacobians with respect to the state vector at the two time instants t_k and t_{k+m} , correspondingly. In this last expression, \mathcal{X} is the 6×7 block diagonal matrix with matrix elements the 3×3 identity matrix and the 3×4 Jacobian matrix that projects the algebraic quaternion difference to the tilt angles' error. All these quantities are described in detail in Ref. [22]. Both the original measurement noise n_r and the projected \tilde{n}_r are assumed to be a zero-mean white noise Gaussian processes with covariances:

$$R_r = E[n_r n_r^T] = \begin{bmatrix} R_q & R_{pq} \\ R_{pq} & R_q \end{bmatrix}, \quad \tilde{R}_r = E[\tilde{n}_r \tilde{n}_r^T] = \mathcal{X} R_r \mathcal{X}^T.$$

As is evident from (20), the relative pose measurement error is expressed in terms of the current $\Delta x_2 = \Delta x(t_{k+m})$ and the previous $\Delta x_1 = \Delta x(t_k)$ (error) state of the system. Note that t_k and t_{k+m} are the time instants when, for example, the two images (encoder readings) processed by the visual (vehicle) odometry algorithm were recorded and thus the relative pose (motion estimate) measurement provided by it corresponds to the time interval $[t_k \ t_{k+m}]$.

4.3. Augmented-state propagation

If $\Delta x_{k/k}$ is the state estimate at time t_k (when the first image or encoder measurement was recorded) we augment the state vector with a second copy of this estimate:

$$\Delta \check{x} = \begin{bmatrix} \Delta x_{k/k}^T & \Delta x_{k/k}^T \end{bmatrix}^T.$$

Since initially, at time t_k , the two copies of the state estimate are identical, the covariance matrix for the augmented system would be:

$$\check{P}_{k/k} = \begin{bmatrix} P_{kk} & P_{kk} \\ P_{kk} & P_{kk} \end{bmatrix},$$

where P_{kk} is the covariance matrix for the (error) state of the vehicle at the time t_k . In order to conserve the estimate of the state at t_k , necessary for evaluating the relative pose measurement error at t_{k+m} , only the second copy of the state estimate is propagated during this interval, while the first remains stationary. The propagation equation for the augmented system is:

$$\begin{bmatrix} \Delta x_s \\ \Delta x \end{bmatrix}_{k+1/k} = \begin{bmatrix} I & 0 \\ 0 & F_{k+1} \end{bmatrix} \begin{bmatrix} \Delta x_s \\ \Delta x \end{bmatrix}_{k/k} + \begin{bmatrix} 0 \\ I \end{bmatrix} w_k,$$

or:

$$\Delta \check{x}_{k+1/k} = \check{F}_{k+1} \Delta \check{x}_{k/k} + \check{G}_{k+1} w_k,$$

where Δx_s is the stationary copy of the error state of the vehicle and w_k is the system noise due to errors in the IMU measurements [22]. The covariance of the augmented system is propagated and after m steps it is:

$$\check{P}_{k+m/k} = \begin{bmatrix} P_{kk} & P_{kk} \mathcal{F}^T \\ \mathcal{F} P_{kk} & P_{k+m/k} \end{bmatrix}, \quad (21)$$

where $\mathcal{F} = \prod_{i=1}^m F_{k+i}$ and $P_{k+m/k}$ is the propagated covariance of the evolving state at time t_{k+m} .

At this point, we should note that by employing the error models for the propagation and update (see Section 4.2), the state vector can be updated every time a relative pose measurement from the visual odometry or the kinematics algorithm becomes available. The equations for the augmented-state Kalman filter are described in detail in Refs [22, 25].

5. MAHALANOBIS COMPARATOR (SLIP ESTIMATION)

In this section, we present our approach to rover slippage detection. Based on the kinematic equations of the rover and assuming no wheel slippage, the wheel and rocker-bogie joint measurements are processed to produce a relative position and orientation measurement over a certain time (sampling) interval. Before updating the state estimate of the EKF, this kinematics-based displacement measurement

needs to be validated by comparing it to the equivalent one estimated by the EKF. If significant wheel slippage has occurred, the residual (i.e., difference between odometry-based and EKF-based displacement) will be significantly larger compared to the case where the rover moves on solid ground without any of the wheels slipping. A statistical measure for assessing the validity of these measurements is the Mahalanobis distance:

$$m_d = r_{k,k+m}^T S^{-1} r_{k,k+m}, \quad (22)$$

where $r_{k,k+m}$ is the residual and S is the corresponding residual covariance matrix. In the case of this vehicle's odometry measurement, the Mahalanobis distance follows a χ^2 distribution with 5 d.o.f. A sufficient test for validating vehicle odometry measurements $z_{k,k+m}$ is to require that these match the expected (estimated by the EKF) measurements $\hat{z}_{k,k+m}$ of the same quantities with a certain level of confidence. By requiring the fit between the expected and actual measurements to be valid with probability, e.g., $P = 95\%$, odometric measurements are processed by the EKF only when $m_d \leq t$, with $t = 11.07$. If this inequality does not hold, these measurements are discarded and wheel slippage is detected. In this case, the residual is provided to the slip compensation algorithm for appropriately modifying the rover steering and driving commands.

The incentive to include the kinematics measurement in the Kalman filter estimator, when it passes the Mahalanobis distance criterion (as shown in Fig. 2), is mainly for off-nominal conditions, when visual odometry does not perform at 2.5% accuracy (e.g., under rare conditions, when it does not converge to a solution). Under nominal conditions for visual odometry, and when the vehicle is slipping, is when the vehicle kinematics are most useful since they allow for the estimation of slippage.

6. SLIP-COMPENSATED PATH FOLLOWER

At the center of Fig. 2 is the slip-compensated path-following controller. This is comprised of two algorithms that are used in close conjunction to achieve the system's end-goal of enabling the traversal of a desired path through high-slip environments. At the highest level, the algorithms receive a 3×1 slip vector $[\dot{x} \ \dot{y} \ \dot{\phi}]_{\text{slip}}$, a 3×1 rover pose vector $[x \ y \ \phi]_{\text{pose}}$ and a $2 \times n$ desired path vector (where n is the number of waypoints) $[x \ y]_{\text{path}}$, and then outputs a 3×1 commanded rover velocity vector $[\dot{x} \ \dot{y} \ \dot{\phi}]_{\text{cmd}}$.

6.1. Carrot heading algorithm

The carrot heading algorithm calculates the desired heading, ϕ_{carrot} , of the vehicle given the desired path and the current rover pose. This algorithm was employed for its robustness to path error [26, 27]. The desired path consists of a set of linear segments between waypoints. The waypoints, however, can be spaced at

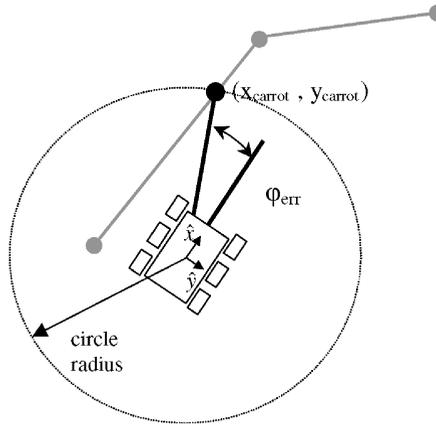


Figure 6. Carrot heading calculation.

any distance apart, thus allowing for paths of arbitrary complexity. The algorithm determines the desired heading by calculating the intersection of a circle centered on the rover frame origin with the desired path and computing the direction of that point (Fig. 6). The intersection point, $(x_{\text{carrot}}, y_{\text{carrot}})$, that is furthest along the path is always selected. The heading error is then computed as:

$$\phi_{\text{err}} = \phi_{\text{carrot}} - \phi_{\text{pose}}. \quad (23)$$

A large radius results in a smooth motion of the rover, but tends to filter out small features of the path. A small radius reduces the total path following error. However, it requires large heading changes of the rover for small path errors, which is extremely inefficient. A circle radius is selected that balances the desire to closely follow the path and the magnitude of the heading changes. Under nominal conditions, the rover path error will always be smaller than the circle radius. If this is not the case, then the radius is grown until an intersection occurs.

6.2. Slip-compensated path-following algorithm

When the Mahalanobis comparator determines that slippage has actually occurred, the calculation of rover slip is made by comparing the output from the Kalman filter and the output from the forward kinematics. If statistically significant slippage is not detected, then the slip vector consists of zeros and the compensation algorithm described below converges to a heading controller.

The slip compensation algorithm consists essentially of two separate control loops. The first control loop, the heading controller, is described by the equation:

$$\dot{\phi}_{\text{cmd}} = (K_1 \cdot \phi_{\text{err}} + K_2 \cdot \dot{\phi}_{\text{slip}}) / T_s. \quad (24)$$

This loop determines the commanded yaw rate of the vehicle as a combined function of the heading error, ϕ_{err} (as calculated by the carrot heading algorithm), and the yaw slip rate, $\dot{\phi}_{\text{slip}}$. It attempts to achieve the optimal heading, determined by the carrot

algorithm, even when slipping in the yaw direction. The second loop is described by the equation:

$$\dot{y}_{\text{cmd}} = K_3 \cdot \dot{y}_{\text{slip}} / T_s. \quad (25)$$

This loop calculates the rate of the rover along the y -axis based entirely on the slip in the y direction during the previous sample period. A \dot{y}_{cmd} command results in a crabbing maneuver, where all six wheels have a steering angle offset in the same direction. \dot{x}_{slip} is implicitly compensated for by the rover driving for a longer period of time along the path towards the goal. In (24) and (25), K_1 , K_2 and K_3 are tuned controller gains, and T_s is the controller sample period. \dot{x}_{cmd} is then determined to be the maximum value allowed that keeps the rover within its operational constraints (i.e., the maximum speed of the drive motors). These rover commands, $[\dot{x} \ \dot{y} \ \dot{\phi}]_{\text{cmd}}$, are then passed to the inverse kinematics.

6.3. Continuous motion

With minor changes made to the algorithm it was possible to allow continuous motion of the rover while compensating for slippage. Images are taken, visual odometry is run, the slip vector is estimated and the compensation command is then computed, all while the rover is moving.

There are many benefits for selecting a continuous motion profile: (i) more efficient motion (both in energy and time), (ii) less slippage due to smaller changes in momentum (this becomes particularly significant as the rover increases in mass and the terrain becomes steeper) and (iii) it enables a higher rate of slip compensation, thus achieving greater accuracy when following a path. It also removes the requirement of many navigation algorithms to use arcs as the fundamental motion of the rover and permits more complex spline paths to be the basis of path planning. Another advantage is that it supports more optimal spacing between navigation stops (allowing for other requirements, such as IMU bias zeroing, to determine the appropriate spacing between stops).

7. EXPERIMENTAL RESULTS

Three sets of experiments have been performed using Rocky 8 (Fig. 4), a Mars rover research platform developed at JPL. This rover has a very similar mobility system to Sojourner, MERs and the current design of the 2009 Mars Science Laboratory (MSL) rover. Rocky 8 utilizes a rocker–bogie suspension (see Section 3.1) which allows for excellent mobility over very rough terrain [18]. Specifications of the Rocky 8 rover are given in Table 1.

With the processor onboard the rover, visual odometry runs at approximately 1 Hz. This rate dictates the sample period of the slip-compensated path follower, T_s . Several parameters of the slip-compensated path follower also had to be tuned in the field. The three gains of the control loops K_1 , K_2 and K_3 in (24) and (25)

Table 1.
Rocky 8 rover specifications

Number of driveable wheels	6
Number of steerable wheels	6
Wheel diameter (m)	0.20
Wheel width (m)	0.10
Mass (kg)	60
Wheelbase (m)	0.75
Track width (m)	0.62
Maximum linear velocity (m/s)	0.09
Maximum rotational velocity (rad/s)	0.24
Camera resolution (pixels)	640 × 480
Camera field of view (horizontal deg × vertical deg)	79.5 × 64.0
Camera baseline (m)	0.084
Camera angle from horizontal (deg)	45.0
Processor type and speed (GHz)	Pentium III 1.2

were tuned by first increasing K_1 until the heading error was maintained at an acceptably low value without inducing oscillations in the heading of the rover. Then, K_2 was increased until the rover acceptably compensated for yaw slippage under the conditions experienced in the field. Finally, K_3 was increased until the slippage in the y direction was effectively compensated for without inducing crabbing oscillations. Additionally, a radius of the carrot heading algorithm was determined that balanced the need to closely follow the path with the need to minimize heading changes. An acceptable value for Rocky 8 was found to be 0.50 m.

The first set of experiments was performed in the JPL Marsyard, a 20×20 -m space designed as an analogue (in rock size/distribution and soil characteristics) to the Viking Lander sites. The experiment consisted of two consecutive 25-m runs with visual odometry running onboard. The second set of experiments was performed in the Mojave Desert. The terrain of this area contains slopes up to 25° consisting of loose granular sand (Fig. 1). This set of experiments was a test of a simplified integrated slip-compensated path-following system. It was simplified in the sense that the Kalman filter and Mahalanobis comparator had not yet been implemented; however, a slip estimate was calculated and compensated for every time the visual odometry provided a new estimate, which was approximately every 20–30 cm. Another simplification, due to limitations of the vehicle, was to assume the rocker and bogie angles were zero. The third set of experiments was performed on JPL's tiltable platform, a 5×5 -m 'sandbox' that could be tilted in discrete increments from 0 to 30° . The terrain on the tiltable platform can be characterized as cohesionless beach sand. In the experiments presented here it was set at 10° . This set of experiments consisted of multiple runs at varying approach angles (ranging from straight up the 10° slope to horizontally across the slope) and used the updated continuous motion algorithm described in Section 6.2.

In the first two sets of experiments, ground truth data were collected with a Leica Total Station (LTS), which is a laser-based position measurement system. The LTS was used to measure the absolute position of four prisms mounted on the vehicle (Fig. 4) every time the rover was stationary (approximately every 20–30 cm). This system gives an accuracy of ± 2 mm in position and $\pm 0.2^\circ$ in attitude. In the field test experiments the waypoints for the rover were also designated using the LTS and a single prism. In the third set of experiments, the LTS could only be used for ground truth at the beginning and end of the runs because the rover did not stop during each test.

7.1. Visual odometry results

Visual odometry results are shown from the Marsyard (Figs 7 and 8) and the Mojave Desert experiments (Fig. 9). The errors at the end of both Marsyard runs are less than 2.5% of the distance traveled. As shown in Fig. 9, the error (0.37 m) at the end of the field test run, at the Mojave Desert, is less than 1.5% of the distance traveled (29 m).

7.2. Slip-compensated path follower results

Results of the slip-compensated path-following algorithm are shown in Figs 10 and 11. The entire section of the path shown in Fig. 10 was on a slope of between 10 and 15°. Figure 10 is an expansion of the box shown in Fig. 9 (note that the axes of Fig. 10 have been rotated in order to minimize the vertical size of the plot). Figure 11 is an expansion of the box shown in Fig. 10. These two figures show three important pieces of information that the slip-compensated path-following algorithm uses to calculate the rover commands: visual odometry pose (as described in Section 2), kinematics pose (as described in Section 3) and the desired path. From this information the slip-compensated path follower calculates the rover commands as described in Section 6. Carrot heading, which is calculated in an intermediate step, is also shown. In Fig. 10, the rover was able to accurately and efficiently follow the desired path, despite significant slippage. As can be seen in Fig. 11, there is a noticeable bias between the visual odometry pose and the kinematics pose in the y direction. This is due to the downhill slippage of the rover; this bias is being compensated for by the slip compensation algorithm, as is evident from the fact that the rover accurately follows the desired path.

7.3. Continuous motion results

Results from the continuous motion slip compensation experiments are shown in Fig. 12 (see Section 7.2 for a description of the elements of this plot). The top graph in Fig. 12 shows a 4.5-m run using continuous slip compensation on a 10° slope (the slope is going down towards the positive y -axis). Again, notice the constant offset between the kinematic and visual odometry estimates, indicating slippage. The

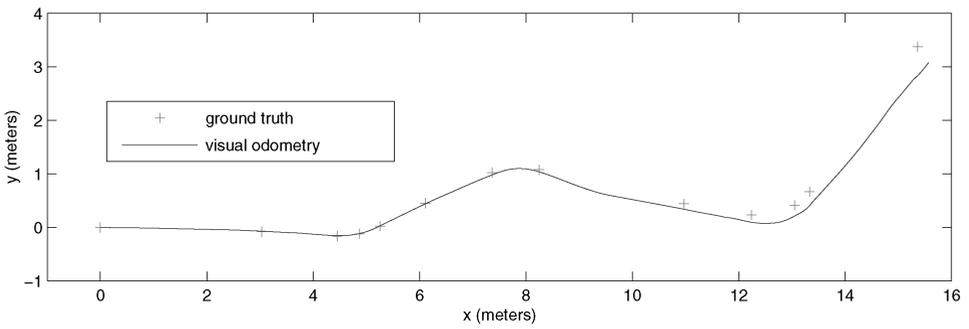


Figure 7. Marsyard run 1 visual odometry results.

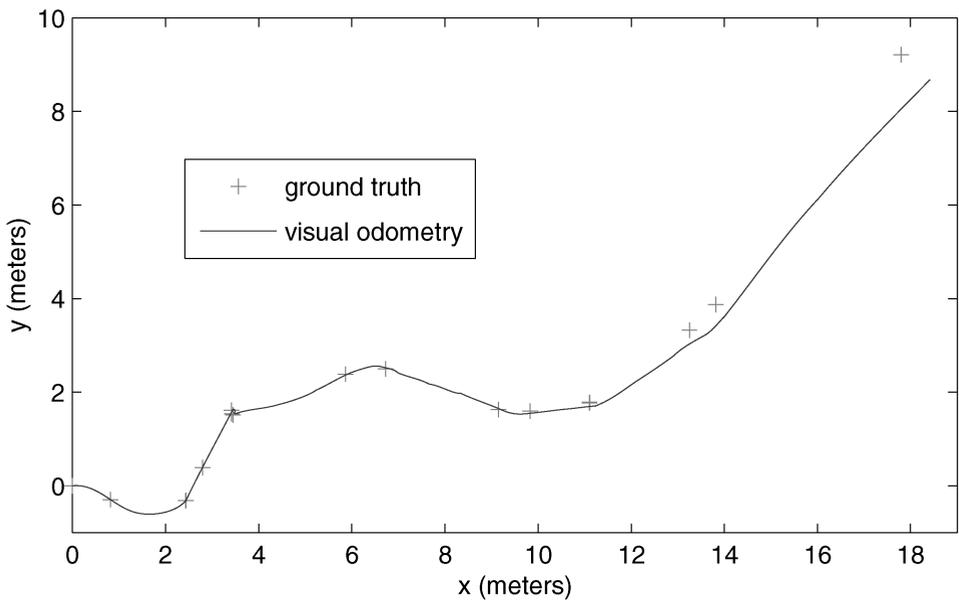


Figure 8. Marsyard run 2 visual odometry results.

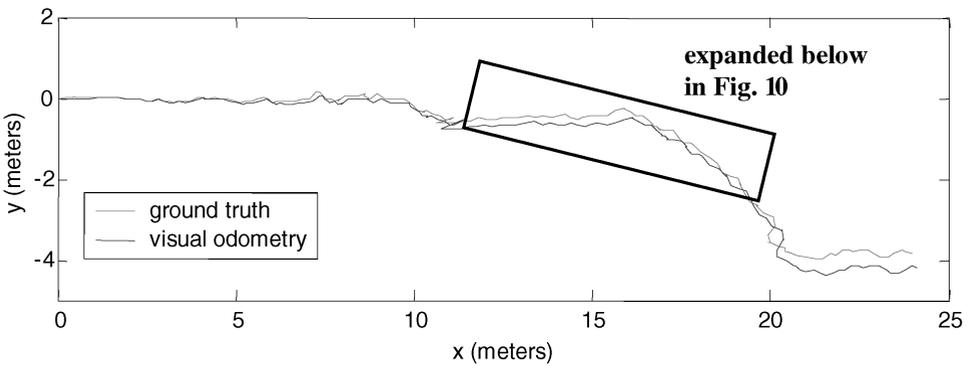


Figure 9. Field test visual odometry results.

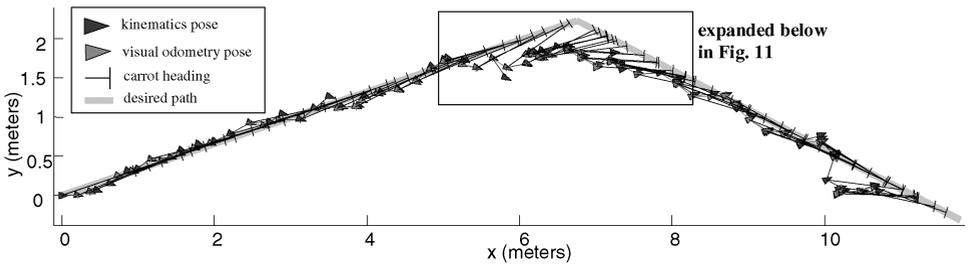


Figure 10. Field test slip compensation results.

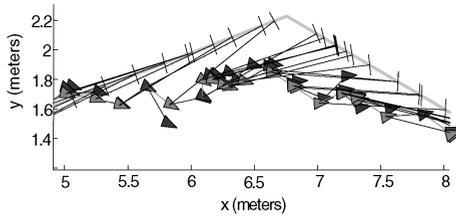


Figure 11. Expanded slip compensation results.

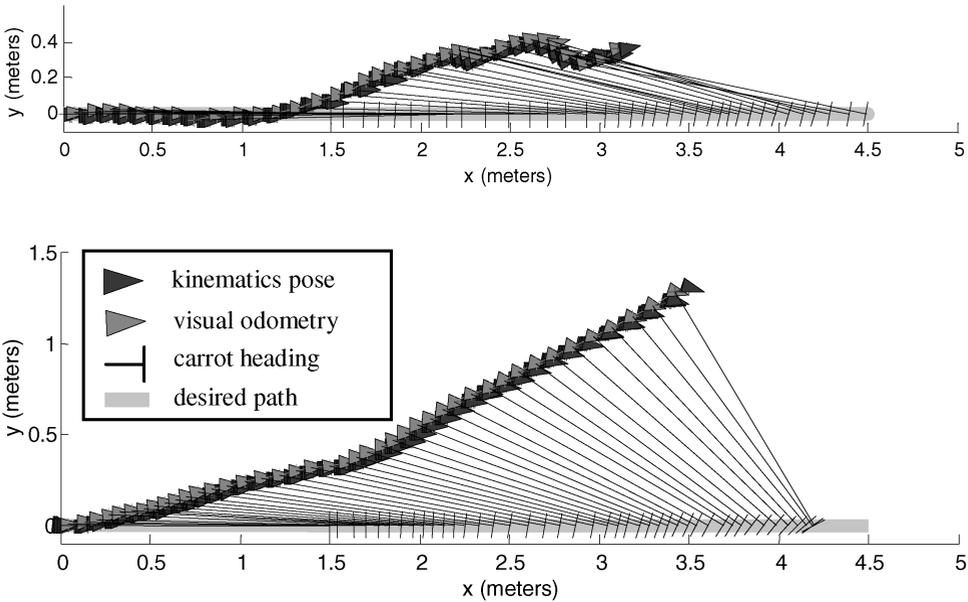


Figure 12. Tilttable platform slip compensation results (top: slip compensation enabled; bottom: slip compensation disabled).

bottom graph in Fig. 12 shows an identical run without using slip compensation. As can be seen, the significant slippage shown in the bottom graph was compensated for in the run shown in the top graph. The mean square path offset error is 0.05

and 0.54 m² for the runs with and without slip compensation, respectively. This is a significant improvement in performance.

8. CONCLUSIONS

In this paper, we have described the design, implementation and testing of a system that enables a rover to accurately follow a designated path, compensate for slippage and reach intended goals, independent of terrain geometry and soil characteristics along the path (within the mechanical constraints of the mobility system). Individual components have been simulated and tested; additionally, an integrated system has been tested onboard a rover in a desert field test. The results from the individual and integrated tests are encouraging. Visual odometry is able to consistently estimate rover motion to within 2.5% of distance traveled. Given this knowledge, the slip-compensated path-following algorithm is able to accurately estimate and effectively compensate for slip, and thus accurately follow a desired path and reach the intended goal while traversing through a high-slip environment. The algorithm has also been extended to allow for continuous motion slip compensation. Future work includes integration of this slip-compensated path-following algorithm with a path-planning/obstacle avoidance system such as Morphin or GESTALT [28, 29].

Acknowledgements

The research described in this publication was carried out at the Jet Propulsion Laboratory (JPL), California Institute of Technology under contract from the National Aeronautics and Space Administration (NASA), with funding from the Mars Technology Program, NASA Science Mission Directorate. The work of S. I. R. was supported by the JPL (grant no. 1248696) and the National Science Foundation (ITR-0324864 and MRI-0420836).

REFERENCES

1. D. Helmick, S. I. Roumeliotis, Y. Cheng, D. Clouse, M. Bajracharya and L. Matthies, Slip compensation for a Mars rover, in: *Proc. IEEE Int. Conf. on Intelligent Robots and Systems*, Edmonton, pp. 1419–1426 (2005).
2. D. M. Helmick, Y. Chang, S. I. Roumeliotis, D. Clouse and L. Matthies, Path following using visual odometry for a Mars rover in high-slip environments, in: *Proc. IEEE Aerospace Conf.*, Big Sky, MT (2004).
3. A. Kelly and B. Nagy, Reactive nonholonomic trajectory generation via parametric optimal control, *Int. J. Robotics Res.* **22**, 583–601 (2003).
4. S. Hayati, R. Volpe, P. Backes, J. Balam, R. Welch, R. Ivlev, G. Tharp, S. Peters, T. Ohm, R. Petras and S. Laubach, The Rocky 7 rover: a Mars sciencecraft prototype, in: *Proc. IEEE Int. Conf. on Robotics and Automation*, Albuquerque, NM (1997).
5. R. Volpe, J. Balam, T. Ohm and R. Ivlev, Rocky 7: a next generation Mars rover prototype, *Advanced Robotics* **11**, 341–358 (1997).

6. P. Schenker, P. Pirijanjan, J. Balaram, K. S. Ali, A. Trebi-Ollennu, T. L. Huntsberger, H. Aghazarian, B. A. Kennedy, E. T. Baumgartner, K. Iagnemma, A. Rzepniewski, S. Dubowsky, P. C. Leger and Apostopolous, Reconfigurable robots for all terrain exploration, *Proc. SPIE* **4196**, 454–468 (2000).
7. R. Hogg, A. Rankin, S. Roumeliotis, M. McHenry, D. Helmick, C. Bergh and L. Matthies, Algorithms and sensors for small robot path following, in: *Proc. IEEE Int. Conf. on Robotics and Automation*, Washington, DC (2002).
8. L. Matthies, Dynamic stereo vision, PhD thesis. Carnegie Mellon University (1989).
9. C. F. Olson, L. H. Matthies, M. Shoppers and M. Maimone, Robust stereo ego-motion for long distance navigation, in: *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, Vol. 2, (2000).
10. C. F. Olson, L. H. Matthies, M. Shoppers and M. Maimone, Stereo ego-motion improvements for robust rover navigation, in: *Proc. IEEE Int. Conf. on Robotics and Automation* (2001).
11. Y. Cheng, M. Maimone and L. Matthies, Visual odometry on the Mars Exploration Rovers, in: *Proc. IEEE Conf. on Systems, Man and Cybernetics*, The Big Island, HI (2005).
12. R. Deriche and G. Giraudon, A computational approach for corner and vertex detection, *Int. J. Comput. Vis.* **10**, 101–124 (1993).
13. P. H. Schonemann, A generalized solution of the orthogonal Procrustes problem, *Psychometrika* **31**, 1–10 (1966).
14. P. J. Rousseeuw, Least median-of-squares regression, *J. Am. Stat. Ass.* **79**, 871–880 (1984).
15. M. Tarokh, G. McDermott, S. Hayati and J. Hung, Kinematic modeling of a high mobility Mars rover, in: *Proc. IEEE Int. Conf. on Robotics and Automation* (1999).
16. M. Tarokh, G. McDermott and J. Hung, Kinematics and control of Rocky 7 Mars rover, *Preliminary Report*. Department of Mathematics & Computer Sciences, San Diego State University (1998).
17. P. F. Muir and C. P. Neumann, Kinematic modeling of wheeled mobile robots, *J. Robotics Syst.* **4**, 281–340 (1987).
18. D. Bickler, A new family of JPL planetary surface vehicles, in: *Proc. Missions, Technologies, and Design of Planetary Mobile Vehicles*, Toulouse, pp. 301–306 (1992).
19. J. J. Craig, *Introduction to Robotics*, 2nd edn. Addison-Wesley, Reading, MA (1989).
20. B. Friedland, Analysis strapdown navigation using quaternions, *IEEE Trans. Aerospace Electron. Syst.* **14**, 764–768 (1978).
21. E. J. Lefferts and F. L. Markley, Dynamics modeling for attitude determination, *AIAA Paper 76-1910* (1976).
22. S. I. Roumeliotis, A Kalman filter for processing 3-D relative pose measurements, *Technical Report*, Robotics Laboratory, California Institute of Technology (2001); http://robotics.caltech.edu/~stergios/tech_reports/relative_3d_kf.pdf.
23. S. I. Roumeliotis, G. S. Sukhatme and G. A. Bekey, Circumventing dynamic modeling: evaluation of the error-state Kalman filter applied to mobile robot localization, in: *Proc. IEEE Int. Conf. on Robotics and Automation*, Detroit, MI (1999).
24. S. I. Roumeliotis, Robust mobile robot localization: from single-robot uncertainties to multi-robot interdependencies, PhD Thesis, Electrical Engineering Department, University of Southern California (2000).
25. S. I. Roumeliotis and J. W. Burdick, Stochastic cloning: a generalized framework for processing relative state measurements, in: *Proc. IEEE Int. Conf. on Robotics and Automation*, Washington, DC, pp. 1788–1795 (2002).
26. A. Kelly, A feedforward control approach to the local navigation problem for autonomous vehicles, *Technical Report CMU-RI-TR-94-17*. Robotics Institute, Carnegie Mellon University (1994).

27. S. Singh, D. Feng, P. Keller, G. Shaffer, W. Shi, D. H. Shin, J. West and B. X. Wu, FastNav: a system for fast navigation, *Technical Report CMU-RI-TR-91-20*. Robotics Institute, Carnegie Mellon University (1991).
28. S. B. Goldberg, M. W. Maimone and L. H. Matthies, Stereo vision and rover navigation software for planetary exploration, in: *IEEE Aerospace Conf. Proc.*, Big Sky, MT, Vol. 5 (2002).
29. S. Singh, K. Schwehr, R. Simmons, T. Smith, A. Stentz, V. Verma and A. Yahja, Recent progress in local and global traversability for planetary rovers, in: *Proc. IEEE Int. Conf. on Robotics and Automation* (2000).

ABOUT THE AUTHORS



Daniel Helmick received his BS degree in Mechanical Engineering from Virginia Polytechnic Institute and State University and his MS degree in Mechanical Engineering with a specialization in controls from Georgia Institute of Technology, in 1996 and 1999, respectively. Since June 1999 he has been working at the Jet Propulsion Laboratory on robotics research projects involving vision/sensor-based control of robots, state estimation, and navigation and mobility algorithms. He has worked on robotic vehicles covering a wide range of functionality, including: Mars research rovers for rough terrain mobility; small, tracked robots for urban mobility; a cryobot for ice penetration; and reconfigurable wheeled robots for Mars exploration. His research interests include: sensor-based control of robots, sensor fusion and state estimation, and rover navigation and mobility.



Stergios Roumeliotis received his Diploma in Electrical Engineering from the National Technical University of Athens (NTUA), Greece in 1995, and his MS and PhD degrees in Electrical Engineering from the University of Southern California (USC), Los Angeles, CA, in 1996 and 2000, respectively. Between 2000 and 2002, he was a Postdoctoral Fellow at the Division of Engineering and Applied Science at the California Institute of Technology, Pasadena, CA. Currently, he is an Assistant Professor in the Department of Computer Science and Engineering and a Faculty Affiliate with the Digital Technology Center at the University of Minnesota. His research has focused on inertial navigation, sensing and estimation for distributed autonomous systems, sensor networks, and fault detection and identification. He has authored or coauthored more than 30 journal and conference papers in the above areas. He was the recipient of the Myronis fellowship (USC, 1998–2000). Finally, he has received grants from the Jet Propulsion Laboratory and the National Science Foundation for work pertinent to autonomous vehicle state estimation.



Yang Cheng is a Senior Research Staff Member in the Machine Vision group of the Jet Propulsion Laboratory (JPL). He earned his PhD in Remote Sensing from the Geography Department of the University of South Carolina and was then a Staff Member at Oak Ridge National Laboratory. Since starting at JPL in 1999, he has worked on several space robotic research projects including the Landmark Based Small Body Navigation System, Vision and Navigation Subsystem of the FIDO rover, and Passive Imaging Based Spacecraft Safe landing. His research interests include robotic navigational autonomy, computer vision, cartography, map projection, geographic information systems, parallel computing, etc.



Daniel Clouse received his BA degree in Computer Science from the University of California, Berkeley in 1982, his MS degree in Computer Science from the University of California, San Diego (UCSD) in 1992, and his PhD degree in Cognitive Science and Computer Science from UCSD in 1998. He is a Senior Member of the Technical Staff at the Jet Propulsion Laboratory, California Institute of Technology in the Machine Vision group of the Mobility Section. His interests include computer vision, machine learning, language translation and word sense disambiguation.



Larry Matthies obtained a PhD degree in Computer Science from Carnegie Mellon University in 1989 and then moved to the Jet Propulsion Laboratory (JPL), where he is currently a Principal Member of Technical Staff and Supervisor of the Machine Vision Group. His research has focused on terrain sensing and obstacle avoidance algorithms for autonomous navigation of robotic vehicles. At JPL, he pioneered the development of real-time algorithms for stereo vision-based obstacle detection and he contributed to the development of the structured light sensor used by the Sojourner Mars rover. He has also developed algorithms for visual motion estimation from image sequences, three-dimensional scene reconstruction from image sequences, real-time terrain classification using multispectral imagers, and environmental mapping using sonar and stereo vision sensors. His group currently has research projects on computer vision for robotic vehicles sponsored by NASA, DARPA and the US Army; these projects include work on navigation of Mars rovers, asteroid and comet landers, and Earth-based robotic vehicles for urban and cross-country missions. He is a member of the Editorial Board of *Autonomous Robots* and an adjunct member of the Computer Science Department at the University of Southern California. He has also been an Invited Speaker at the Frontiers of Engineering Symposium organized by the National Academy of Engineering.