Autonomous Docking of a Smart Wheelchair for the Automated Transport and Retrieval System (ATRS)

Chao Gao, Thomas Miller, and John R. Spletzer
Computer Science and Engineering
Lehigh University
Bethlehem, Pennsylvania 18015
e-mail: chg205@cse.lehigh.edu, spletzer@cse.lehigh.edu

Ira Hoffman and Thomas Panzarella
Freedom Sciences, LLC
4601 South Broad Street
Philadelphia, Pennsylvania 19112
e-mail: ihoffman@freedomsciences.com, tpanzarella@freedomsciences.com

Received 27 August 2007; accepted 30 January 2008

The Automated Transport and Retrieval System (ATRS) represents a technology-based alternative to van conversions for automobile drivers in wheelchairs. Rather than requiring dramatic, permanent, and expensive modifications to the host vehicle, ATRS employs robotics and automation technologies and can be integrated noninvasively into a standard minivan or sport utility vehicle. At the core of ATRS is a "smart" wheelchair system that autonomously navigates between the driver's position and a powered lift at the rear of the vehicle, eliminating the need for an attendant. From an automation perspective, autonomously docking the wheelchair onto the lift platform presented the most significant technical challenge due to limited clearance between the chair wheels and the lift platform rails. To solve the docking task, we employed a light detection and ranging (LIDAR)–based approach for wheelchair localization coupled with a hybrid motion controller design. Extensive testing from the localization subsystem to the complete ATRS was conducted under representative usage conditions. This included 3 days of public demonstrations indoors at the World Congress on Disabilities, where potential end users were able to evaluate the system. In this environment, ATRS performed more than 300 consecutive cycles without failure. During 2 days of outdoor reliability testing, 97.5% docking reliability was observed. The system is scheduled to enter the commercial market in 2008.

Journal of Field Robotics 25(4–5), 203–222 (2008) © 2008 Wiley Periodicals, Inc.
Published online in Wiley InterScience (www.interscience.wiley.com). • DOI: 10.1002/rob.20236

1. INTRODUCTION AND MOTIVATION

According to the U.S. Department of Transportation, more than six million people with disabilities have difficulties in obtaining the transportation they need (National Council on Disabilities, 2005). This is a major contributor to the unemployment rate of the disabled population nationally, estimated at more than 65% by the U.S. Census Bureau (Stern & Brault, 2005). These statistics are reinforced by research from the Pennsylvania State Board of Vocational Rehabilitation, which found that transportation was critical for Americans with disabilities to participate fully in basic activities such as employment, education, worship, job training, commerce, recreation, and other activities of community life that most people take for granted (Pennsylvania Rehabilitation Council, 2006).

A van conversion is the de facto personal transportation solution for an individual in a power wheelchair. Van conversions start with a standard van produced by a major automotive manufacturer. The van is subsequently modified by another company specializing in mobility equipment. The modifications are permanent and include extensive changes to the chassis, frame, and interior. Typical modifications include removing and lowering the vehicle floor and relocating or replacing major subsystems such as the gas tank, fuel system, and heating/cooling systems of the vehicle (Wood, 2004). While enabling independent mobility, van conversions represent a costly and unsafe transportation solution for wheelchair users.

Safety is compromised by the modifications, but more significantly by the paradigm itself. Van conversions place the operator *in-wheelchair* behind the steering wheel of the vehicle. Entry/exit to the vehicle is also accomplished in-wheelchair via a ramp or powered lift device. Such a design has significant safety shortcomings. Wheelchairs do not possess levels of crash protection similar to those afforded by traditional motor vehicle seat systems, and the provisions used for securing them are often inadequate. It should not be surprising then when research by the U.S. National Highway Traffic Safety Administration (NHTSA) showed that 35% of all wheelchair/automobile-related deaths were the result of inadequate chair securement. Another 19% were associated with vehicle lift malfunctions (NHTSA, 1997).

To eliminate these shortcomings, we have developed a technology-based alternative to van conversions for wheelchair users: the Automated Transport and Retrieval System (ATRS). ATRS employs robotics and automation technologies and can be integrated into a standard minivan or sport utility vehicle (SUV). At the core of ATRS is a "smart" wheelchair system that autonomously navigates between the driver's position and a powered lift at the rear of the vehicle. A primary benefit of this paradigm is that the operator and chair are separated during vehicle operations as well as entry/exit. This eliminates the potential for injuries or deaths caused by both improper securement (as the operator is seated in a crash-tested seat system) as well as lift malfunctions. Furthermore, by eliminating the drastic and permanent vehicle modifications associated with van conversions, ATRS will cost significantly less. ATRS represents the rare case when an assistive technology provides additional functionality at a lower cost than current alternatives.

This paper reviews the estimation and control techniques used for autonomously docking the smart wheelchair onto a power lift platform. This eliminates the need for an attendant to shuttle the wheelchair back-and-forth from the lift platform to the driver's seat—enabling personal independence through technology.

2. RELATED WORK

Extensive work has been done in order to increase the safety levels of power wheelchairs while minimizing the level of human intervention. In these systems, the human operator is responsible for high-level decisions while the low-level control of the wheelchair is autonomous.

The Tin Man system (Miller & Slack, 1995), developed by the KISS Institute, automates some of the navigation and steering operations for indoor environments. The Wheelesley project (Yanco, 1998), based on a Tin Man wheelchair, is designed for both indoor and outdoor environments. The chair is controlled through a graphical user interface that has successfully been integrated with an eye tracking system and with single-switch scanning as input methods. The TAO Project (Gomi & Griffith, 1998) provides basic collision avoidance, navigation in an indoor corridor, and passage through narrow doorways. The system also provides landmark-based navigation that requires a topological map of the environment. The NavChair assistive wheelchair navigation system (Simpson & Levine, 1999) uses

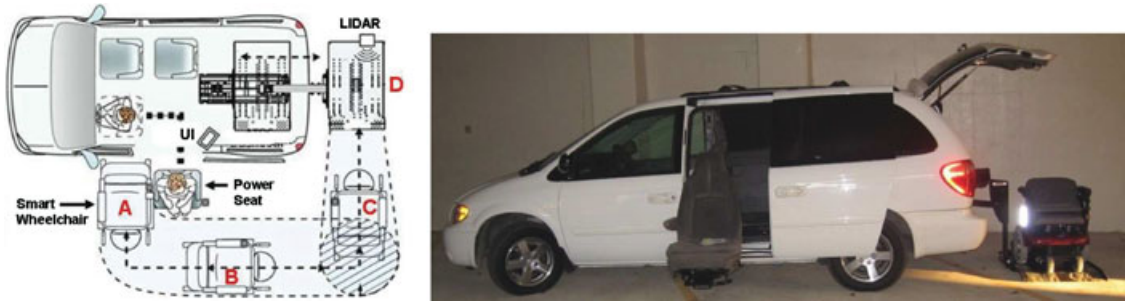

Figure 1. (Left) ATRS concept diagram. (Right) Beta ATRS showing the power seat, smart wheelchair, and lift platform subsystems.

feedback from ultrasonic sensors and offers obstacle avoidance, door passage, and wall-following modes. More recently, the SmartChair (Parikh et al., 2003; Parikh, Grassi, Kumar, & Okamoto, 2004) uses a virtual interface displayed by an onboard projection system to implement a shared control framework that enables the user to interact with the wheelchair while it is performing an autonomous task.

A common theme in the above research is that robotics technology has been applied to assist or augment the skills of the chair operator. In contrast, the ATRS wheelchair is in fact capable of autonomous vehicle navigation in outdoor environments. This can be realized because the operator is never seated in the chair during autonomous operations, and the chair always operates in the vicinity of the automobile. The former constraint mitigates operator safety issues, while the latter provides significant, invariant landmarks/features in an otherwise unstructured environment. What also makes the ATRS wheelchair attractive is that it is commercially viable, providing a safer alternative to van conversions at a significantly lower cost.

We should also point out that for users of manual wheelchairs, there are alternatives to van conversions for personal automobility. These range from "muscling" the collapsed wheelchair into the rear seat for those with sufficient flexibility and upper-body strength, to more dramatic commercial systems such as the ChairTopper™, which integrates a lift into a container that is mounted to the roof of an automobile (Braun Corporation, 2006). However, these solutions are inappropriate for powerchairs, which typically weigh 125–150 kg.

3. SYSTEM OVERVIEW

The ATRS can be decomposed into five primary components: a "smart" power wheelchair system, a light detection and ranging (LIDAR) system for localization, a powered lift platform, a traversing power seat, and a touch-screen user interface (UI) computer. These are illustrated in the concept drawing and the ATRS beta system shown in Figure 1. From a robotics perspective, the smart wheelchair and localization systems are the heart of ATRS. Combined, these two subsystems allow the operator to be separated from the chair and eliminate the need for an attendant.

In describing the ATRS operational procedures, we refer to Figure 1 (left). When the operator returns to his/her automobile, a keyless entry is used to both unlock the vehicle and deploy the traversing driver's seat. The operator then positions the wheelchair and performs a seat-to-seat transfer (pose A). After this, the wheelchair is deployed to the rear of the vehicle (pose B). In our proof-of-concept system, this side traversal was completely autonomous (Sermeno-Villalta & Spletzer, 2006). In the current system, referred to colloquially as "ATRS-Lite," the wheelchair is remotely controlled by the vehicle operator via a joystick located at the UI. Once the chair enters the handoff site at the rear of the vehicle (pose C), it is automatically tracked by the localization system used for docking. The wheelchair then switches to "docking" mode (either automatically or via a UI input from the operator for the ATRS-Lite system), which enables the vanside computer to transmit real-time control inputs to the chair over a dedicated radio frequency (RF) link for reliable docking (locking in place) onto the lift platform (pose D). With

the chair docked, the operator actuates the lift via the UI, stowing the platform and chair in the vehicle cargo area. The process is repeated in reverse when disembarking from the automobile. We should emphasize that when not operating autonomously, the ATRS wheelchair is placed in "manual mode" and operates no differently from any other powered wheelchair.

The primary focus of this paper is the development of a reliable, autonomous means for docking (and undocking) the ATRS wheelchair onto (and off of) the vehicle's lift platform. Autonomous docking was dictated by the narrow clearances available on the lift platform; teleoperation proved an unreliable proposition even for trained operators. Our initial efforts in this area investigated a vision-based control approach (Sermeno-Villalta & Spletzer, 2006). Pathological failure modes with the passive vision system led us to our current configuration, which integrates a SICK LMS291 LIDAR for estimating the chair pose. We should also point out that whereas the paper focuses on wheelchair docking, the same algorithms are used in reverse during the undocking phase (although the undocking problem is significantly less difficult).

4. WHEELCHAIR CONTROL

The overall architecture for wheelchair control is presented in Figure 2. Recall that there are two primary modes for controlling the ATRS wheelchair without an operator: remote control and autonomous operations. In both cases, control inputs to the chair are sent via a RF link from the vanside computer. From an automation perspective, two aspects to the control problem must be considered: *motor* control and *motion* control. The motion controller generates higher level velocity commands vanside based on the current chair pose as estimated via the localization system presented subsequently. These are in turn transmitted to the powerchair, which regulates the wheel velocities to achieve the objective linear and angular velocities for docking.

4.1. Motor Control

The wheelchair employs a differential drive system and as such can be accurately modeled through the corresponding kinematic model

$$\begin{bmatrix} w_R \\ w_L \end{bmatrix} = \begin{bmatrix} 1 & b \\ 1 & -b \end{bmatrix} \begin{bmatrix} v \\ \omega \end{bmatrix}, \tag{1}$$

where (w_R, w_L) are the right and left wheel velocities in meters/second and b is the differential drive wheelbase. Referring to Figure 2, the motion controller transmits objective linear and angular velocities, which are in turn mapped to wheel velocities via Eq. (1). These are then regulated by the chair via a proportional–integral–derivative (PID) controller

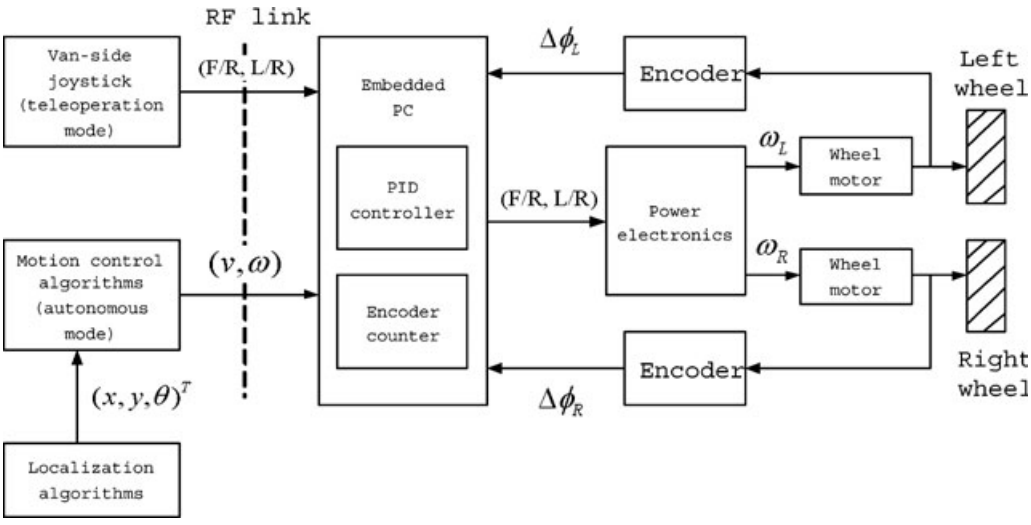

Figure 2. Control architecture for passengerless wheelchair operation.

implemented in software on the chair's embedded personal computer. Feedback to the PID is provided via high-resolution quadrature encoders that measure right and left wheel travel ($\Delta\phi_L$, $\Delta\phi_R$) at 100 Hz.

4.2. Motion Control

The design of the higher level motion controller was influenced by real-world constraints associated with system use. These included actuation latency, docking clearances, and the constrained ground area adjacent to the vehicle for navigation. Furthermore, the chair entry onto the lift platform has to be with sufficient velocity to ensure that a plough installed on its base will strike with sufficient momentum to positively engage the docking mechanism that automatically secures the chair to the lift.

As such, our motion planner employed a hybrid control design consisting of two primary controller modes: course correction and path following. In this paradigm, gross alignment errors were first corrected (when necessary) in the course-correction phase before proceeding to path following for docking. Admittedly, this is a very conservative approach. However, it served to minimize the impact of actuation latency through the elimination of gross pose errors and ensured that the wheelchair remained in close proximity to the vehicle.

To ground the notation, coordinate frames used in this paper are defined as shown in Figure 3. We now describe each mode in greater detail.

4.2.1. Path-Following Phase

Although the motion problem might be classified as point-to-point, there is one caveat. The velocity of the wheelchair at its objective pose must be significantly greater than zero. This is a function of the docking procedure, which requires that a plough mounted to the chair bottom strike the dock with sufficient momentum to actuate the locking mechanism. As a result, we choose instead to treat it as a particular case of path following and employ a traditional proportional–derivative (PD) controller derived using input/output feedback linearization techniques (Deluca, Oriolo, & Samson, 1998):

$$\omega = -k_v \tan\theta - \frac{k_p y}{v\cos\theta}, \tag{2}$$

where ω and v are the desired linear and angular velocities transmitted to the chair, $v(t)$ is assumed piecewise constant, k_v, k_p are positive controller gains, and y, θ are with respect to the world frame W. Fixing the value for v ensured that the singularity at $v = 0$ is avoided. However, whereas Eq. (2) implicitly defines a trajectory, actuator constraints must also be accommodated. For safety considerations, we specify maximum linear and angular velocities $[v_{\max}, \omega_{\max}]^T$ for the chair. Typical values were 0.4 m/s and 0.9 rad/s in practice. (It should be noted that these are significantly less than what can be achieved by the actual hardware.) To accommodate these limits, we borrow from Oriolo, Luca, and Vendittelli (2002) and constrain the actual controller inputs to

$$\begin{aligned} \omega_{\text{act}} &= \mathcal{S}(\omega)\arg\min\{|\omega|, \omega_{\max}\}, \\ v_{\text{act}} &= \frac{\omega_{\text{act}}}{\omega} v_{\max}, \end{aligned} \tag{3}$$

where $\mathcal{S}()$ corresponds to the sign function. These constraints ensure that whereas the wheelchair will no longer follow the same trajectory specified by Eq. (2), it will follow the same path while protecting against actuator saturation.

Last, one further refinement was made to the path-following mode. By exploiting the two degrees of mobility offered by the differential drive system, Eq. (2) was immediately preceded by an orientation correction. The intent was to find an initial orientation θ^* such that the magnitude of ω_0 was minimized—and ideally zero. Squaring both sides of Eq. (2), setting $w = \omega^2$, differentiating w with respect to θ, and setting the result equal to zero, we obtain two possible solutions:

$$\theta^* = \left\{ -\arcsin\left(\frac{k_p y}{k_v v}\right), -\arcsin\left(\frac{k_v v}{k_p y}\right) \right\}. \tag{4}$$

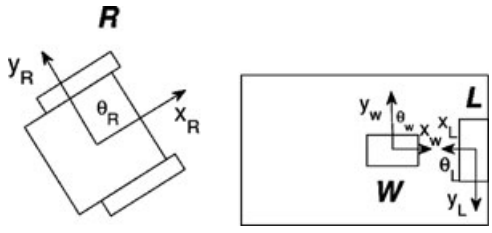

Figure 3. World W, LIDAR L, and wheelchair R coordinate frames.

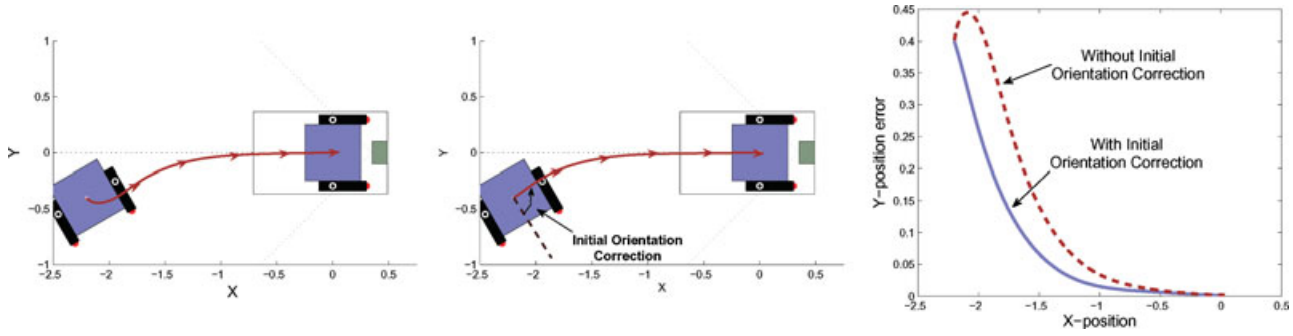

Figure 4. Wheelchair path without (left) and with (center) orientation correction. The latter not only improves the settling distance but also ensures that the chair remains in closer proximity to the vehicle (right).

The first also corresponds to directly setting Eq. (2) to zero. So, for the case where $|(k_p y)/(k_v v)| \leq 1$ there is an initial orientation for our path follower that requires zero initial angular velocity. Fortunately, our configuration parameters allow such an orientation to be readily achieved. Thus, all initial orientation error can be removed prior to initializing the path-follower controller. This reduces the settling distance of the wheelchair, as well as constraining the wheelchair path closer to the vehicle. The effect is illustrated in a simulation trial for a representative gain set (Figure 4).

An additional advantage of the orientation correction phase is that the path-follower controller operates away from the singularity at $\theta = \pm\pi/2$. Because typical gain sets are $\{k_p, k_v\} \approx \{1, 2\}$ and $v \approx 0.4$ m/s, $|\theta^*| \leq 30$ deg. Furthermore, with appropriate gains Eq. (2) ensures that the value for θ decreases from θ^* across the trajectory. This assumes that y-position errors are $\leq$0.4 m. This was ensured (when necessary) through a course-correction phase described below.

4.2.2. Course-Correction Phase

To enhance wheelchair docking reliability, a course-correction mode is also incorporated to address gross y-position errors. This controller phase is activated after initial localization in autonomous model only if it is determined that the path-following mode would be at risk for failing to dock the chair at the hand-off location provided by the operator (e.g., for large y-position errors). In this event, we again exploit the chair's two degrees of mobility to align the chair along the x axis in our world frame. This is accomplished through the following pair of control inputs that are processed serially:

$$d\theta_R = -\theta_{\max}\mathcal{S}(y_0) - \theta_0, \tag{5}$$

$$dx_R = \frac{\text{abs}(y_0)}{\sin\theta_{\max}}, \tag{6}$$

where $(d\theta_R, dx_R)$ are with respect to the wheelchair frame, y_0, θ_0 denote the initial y position and orientation of the wheelchair, respectively, and $\theta_{\max}$ corresponds to a maximum allowable orientation angle for the wheelchair that ensures that both features being tracked by the LIDAR system would remain visible. (Typical values during development were $\theta_{\max} = 60$–75 deg.)

Figure 5 illustrates the effect of these control inputs on a representative docking trial. Figure 5(a) shows the wheelchair's initial pose. Figures 5(b) and 5(c) illustrate the orientation and y-position correction modeled by Eqs. (5) and (6), respectively. At this point, nearly all y-position error has been eliminated from the wheelchair pose. Finally, Figure 5(d) shows the wheelchair reorienting to the optimal orientation defined by Eq. (4), which immediately precedes the path-following control phase from Eq. (2). The net result is a dramatic reduction in the settling distance. This, in conjunction with the orientation correction phase prior to path following, also serves to eliminate vehicle operations around the singularity condition in Eq. (2), i.e., when $\theta = \pm 90$ deg.

4.2.3. Motion Controller Simulation Results

To demonstrate the efficacy of the hybrid control approach, simulations were first conducted to assess docking performance. Docking trials were conducted

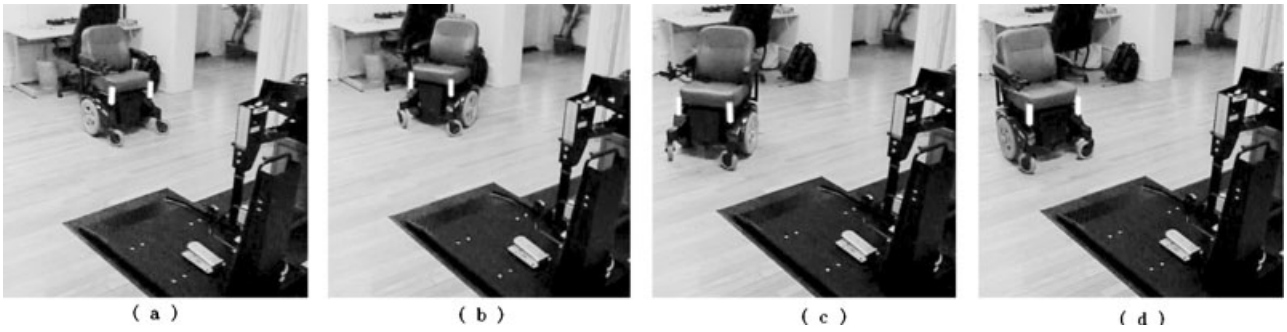

Figure 5. Laboratory docking trial test illustrating the course correction mode (a–c) and orientation correction phase of the path-follower mode (d). This hybrid approach allows for reliable docking across the operational envelope of initial chair poses.

over a 1.2-m^2 handoff area in front of the lift platform using uniform sampling with a discrete position resolution of 1 cm. For each trial, docking was considered successful if immediately prior to reaching the platform ramp, y-position errors were $\leq$5 cm and orientation errors were $\leq$10 deg. These criteria were based on empirical observations that would lead to a successful wheelchair dock.

Simulations were conducted for representative gain sets using purely a path-following controller, path following preceded by orientation correction, and finally with a course-correction mode integrated. Results from a representative set of trials are illustrated in Figure 6. In this figure, the darker dots correspond to positions where docking failed at at least one initial orientation. Dramatic improvements in controller performance can be seen through the integration of the orientation-correction and course-correction modes.

5. WHEELCHAIR LOCALIZATION

5.1. Vision-based Approach

The localization approach for ATRS has evolved significantly over the past two and a half years, and we would admit that as of this writing it is still under development. Our initial implementation demonstrated on the proof-of-concept system was vision based (Sermeno-Villalta & Spletzer, 2006). A high-resolution, black-and-white Point Grey Dragonfly camera using a wide-field-of-view (90-deg) lens was mounted to the inside of the vehicle liftgate. With the liftgate opened, the camera had a "bird's-eye" view of the wheelchair and lift platform as shown in Figure 7. Using normalized intensity distribution as a similarity metric (Fusiello, Trucco, Tommasini, & Roberto, 1999), binary fiducials were tracked on the wheelchair armrests. Assuming a ground plane constraint, this allowed the pose of the wheelchair to be directly estimated. The vision-based solution was demonstrated to industry representatives in June 2005 and subsequently evaluated on the system through October 2005. During this time, the vehicle was driven more than 5,000 km, and the camera did not require recalibration.

There were several advantages to the vision-based approach. First, it was highly accurate. Based on analysis of reprojection residuals from camera calibration, the estimated positioning accuracy of the localization system was subcentimeter. This was more than sufficient for the docking task. The camera system also had a very compact form factor and was relatively inexpensive. However, its performance could suffer in degraded weather conditions (e.g., snow or rain). Furthermore, we identified pathological failure modes even in benign weather conditions. For example, with the vehicle parked underneath a tree on a sunny day, the fiducials being tracked would transition across regions of high brightness (from the sunlight) to deep shade (from the tree leaves). The limited dynamic range of the camera charge-coupled device could not compensate sufficiently for these differences in illumination, resulting in a lost fiducial track and localization failure. The shortcomings of the vision system led us to migrate to the LIDAR-based approach currently used in the ATRS beta model. Similar techniques have seen

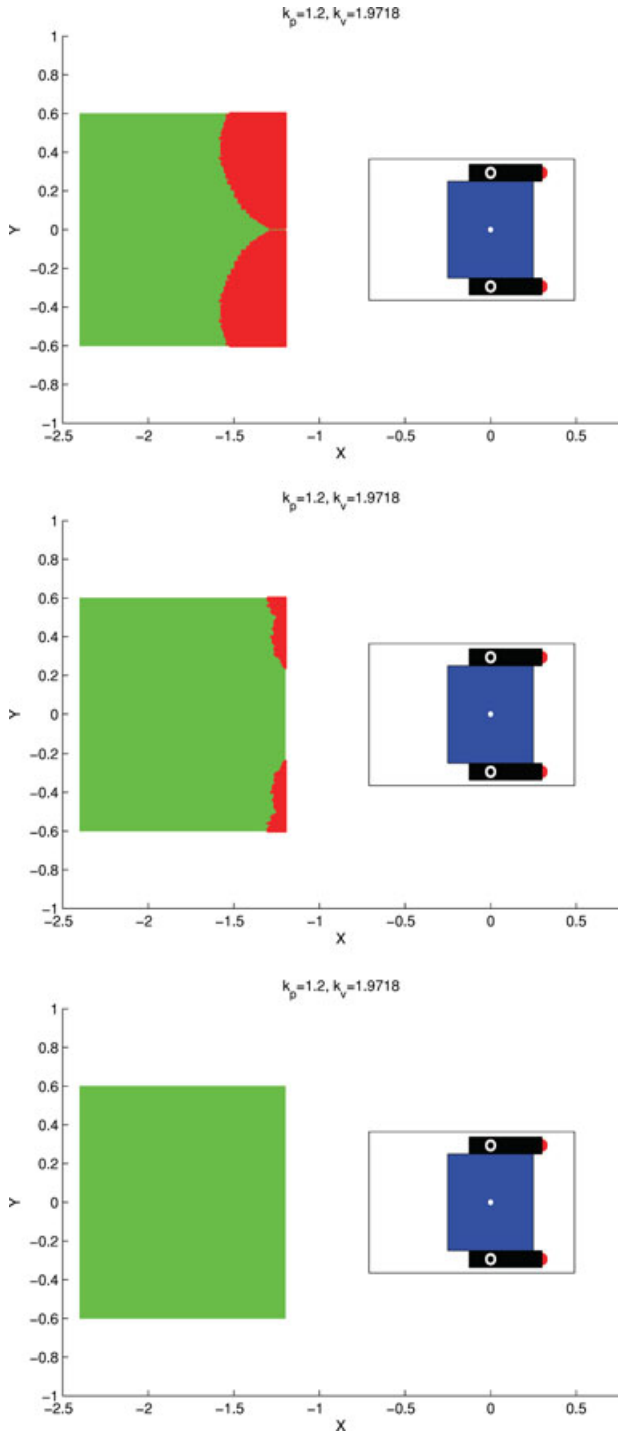

Figure 6. Controller performance for a sample gain set with the PD control law (top), orientation correction integrated (center), and course correction added (bottom). The latter eliminated residual poses associated with docking failure. Simulation resolution was 1 cm^2.

widespread used in the robotics field (Dissanayake, Newman, Clark, Durrant-Whyte, & Csorba, 2001; Howard, Parker, & Sukhatme, 2006; Leonard & Durrant-Whyte, 1991; Lingemann, Surmann, & Hertzberg, 2005; Lionis & Kyriakopoulos, 2002), and a LIDAR/beacon approach was deemed a low-risk solution for the docking task.

5.2. LIDAR-based Approach

In the beta ATRS, the primary sensor used for estimating the wheelchair pose with respect to the life platform is a SICK LMS291 LIDAR. Figure 8 (top) illustrates a typical integration of the LIDAR into the vehicle lift platform. The LMS291 measures the line-of-sight range to objects in the environment over a 90-deg field of view with a discretization of 0.5 deg. Each of these measurements can be written as a tuple of the form $z_m = [r, \alpha, \gamma]_m^T$, $m = 0 \ldots 180$, where r_m and γ_m denote the measured range to and reflectivity of the mth feature at a bearing of $\alpha_m = m/2 - 45$ deg with respect to the LIDAR sensor frame L. To simplify the feature segmentation process, two cylindrical retroreflector fiducials were permanently affixed to the front of the wheelchair, as shown in Figure 8 (center). When imaged by the LIDAR, a significant portion of the incident beam is reflected directly back to the detector, saturating the photodiode. This allows a simple threshold on reflectivity $\gamma_{\min}$ to be used as the primary filter for segmenting the target features. An additional level of filtering is based on a range constraint $r_{\max}$. As the wheelchair is presented in the immediate vicinity of the lift platform, targets at excessive ranges (e.g., >4 m away) can immediately be disqualified from potential features of interest. From these two filters and assuming a ground plane constraint, we construct a valid feature set:

$$F = \begin{bmatrix} r_n \cos \alpha_n \\ r_n \sin \alpha_n \end{bmatrix}, \text{ s.t. } r_n < r_{\max}, \ \gamma_n > \gamma_{\min}. \quad (7)$$

During all testing, $\gamma_{\min} = 250$, $r_{\max} = 400$ cm (in hindsight, $\gamma_{\min}$ could have been set to 255 based on the test results for an even stronger validation gate). This is illustrated in a representative LIDAR scan in Figure 8 (bottom).

After this filtering phase, valid features were clustered in Euclidean space as "candidate fiducials" based on the known fiducial size ($d = 5.2$ cm). Every hypothetical fiducial pair was then evaluated against the known fiducial baseline ($|t_l - t_r| = 44$ cm). Any

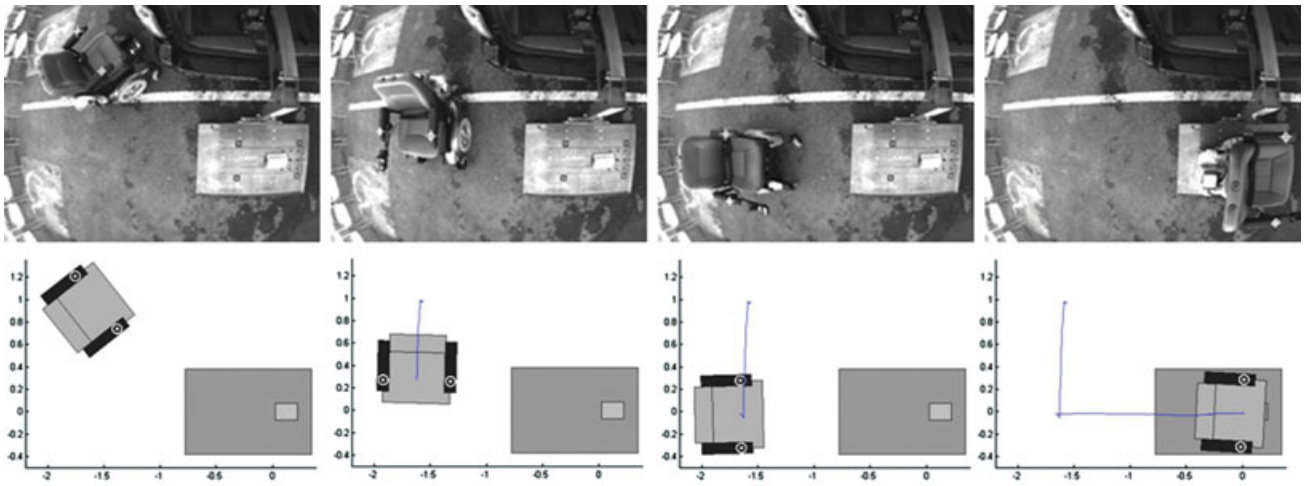

Figure 7. Sample docking trial with the proof-of-concept system, showing the camera perspective and the estimated chair position over time.

pair within preset tolerance of the known baseline was considered a valid chair pose. If and only if one valid chair pose was obtained, the wheelchair segmentation was considered to be successful. Otherwise, the user would be notified via the UI to take corrective action (i.e., reposition the chair) and the segmentation process repeated. This approach has proven to be extremely reliable in real-world conditions. However, to further enhance robustness to outliers in extreme conditions (e.g., heavy rain), the recovered fiducial positions are not directly used. Instead, median filters (typically seven elements) are applied to the recovered range to each fiducial cluster. These filtered ranges are then used to estimate the final fiducial positions. With the position of both retroreflectors known, estimating the chair position and orientation (assuming a ground plane constraint) was straightforward.

Whereas we previously investigated using an extended Kalman filter (EKF) to fuse LIDAR and wheelchair odometry measurements (Gao, Hoffman, & Spletzer, 2007), the current localization implementation uses a simpler approach whereby the wheelchair position and orientation are estimated directly at each timestep from the recovered fiducial positions. This simpler implementation was motivated from significant testing of the localization subsystem, as well as system-level testing outdoors. First, we determined that the positioning accuracy of the LIDAR system was subcentimeter without temporal filtering. Second, significant wheel slippage from gravel, sand, etc., was observed when testing the wheelchair outdoors. Such uncertainty in the odometry measurements results in an EKF implementation that relies primarily on the LIDAR measurement updates anyway and places very little weight on the odometry measurements. Finally, the use of median filters on the range measurements to reject outliers introduced from rain, etc., introduces its own temporal filtering and raises questions as to how the measurements should be correctly fused in the EKF. Although the EKF implementation may be revisited, the test results discussed in Section 8 relied on the direct localization approach.

6. FAILURE RECOVERY

Regardless of the controller design, there inevitably remains the potential for unmodeled disturbances to compromise controller performance and cause docking failure. Without corrective action, this could leave the vehicle operator stranded with the wheelchair trapped in a "partially docked" position on the lift platform. To minimize the potential for such a failure, we leverage the robustness of the localization system. Immediately prior to the wheelchair reaching the ramp of the lift platform, the vanside computer makes a *go/no-go* decision based on the estimated chair pose. If docking is not assured with a very high confidence based on this pose estimate, the

Figure 8. (Top) Tracker^{TM} lift platform with integrated LMS291. The LIDAR housing provides environmental protection and integrates a visor to improve system SNR. (Center) Powerseat and smartchair with retroreflective targets clearly shown. (Bottom) Range and reflectivity data from a single LIDAR scan. The reflectivity measurements of the fiducial targets (twin peaks) facilitate feature segmentation.

chair is stopped and the operator notified via the UI to reposition the chair via the remote-control joystick so that autonomous docking can be reattempted.

However, should the LIDAR system fail, autonomous docking would not be feasible. To address this case, a camera system capable of streaming video to the UI is colocated with the LIDAR on the lift platform to support teleoperation. As mentioned previously, a teleoperation mode was not sufficiently reliable for day-to-day operations. However, operators were able to successfully dock the chair most of the time. As such, it provides a fallback docking modality in the event of a primary sensor failure.

7. LOCALIZATION SUBSYSTEM TESTING

Extensive testing of the localization subsystem was conducted to characterize its performance in both benign and degraded conditions. Most of this testing was accomplished using the modified turntable assembly shown in Figure 9. This test fixture design offered two significant benefits. First, the simulated wheelchair position estimates could be projected to a single point (the center of rotation) so that "ground truth" localization accuracy could be well characterized. Second, the rotation allowed us to simultaneously simulate wheelchair motion in a compact footprint. The fiducials were mounted on an arm and spaced 44 cm apart (the same distance as on the actual wheelchair used in development). The standard turntable angular velocity (33-1/3 rpm) corresponded to an instantaneous fiducial velocity

Figure 9. Test fixture used throughout localization subsystem testing. The turntable modeled wheelchair motion while allowing the simulated wheelchair position estimates to be projected to a single point.

of 0.77 m/s. The actual chair geometry and maximum wheelchair linear and angular velocities used during development (0.4 m/s and 0.9 rad/s, respectively) translated to a worst-case fiducial velocity of 0.68 m/s. As such, the turntable rotation was only slightly more aggressive than the "worst case" motion for the wheelchair.

During all testing with the turntable, it was assumed that the orientation of the wheelchair would be constrained to [−75 deg, 75 deg] to protect against self-occlusion of the retroreflector targets. Test data were collected and evaluated only in this range. The SICK itself was configured with a 500-kbps connection, which has achieved valid scan rates of approximately 70 Hz during laboratory testing.

7.1. Baseline Performance Testing

All baseline testing was conducted at midday (from 11 am to 2 pm) under clear skies to ensure high solar load conditions in an attempt to maximize photodetector noise. Ambient light levels were measured using a light meter and typically exceeded 100 klx.

The turntable used for testing was initially placed at a nominal position of $(x,\ y) = (1.0\text{ m},\ 0.0\text{ m})$ with respect to the LIDAR coordinate frame. The turntable was then actuated, and the fiducials tracked by the SICK for a minimum of 5,000 scans. The x distance to the turntable was then increased by 0.5 m, and the process was repeated. This procedure was iterated for $x = 1.0$–4.0 m and for $y = \{0.0, 0.5\}$ m. The entire test procedure for all ranges was then repeated with the turntable and fiducials facing the sun.

During these tests, the number of hits per fiducial was also recorded to determine whether we could approach the theoretical tracking limits based on fiducial dimensions, angular discretization, and beam divergence. Results from specific subtests are as follows.

7.1.1. Reflectivity Stability

Recall that the motivation for using retroreflector fiducials was to simplify the feature segmentation task by using the reflectivity measurements from the LMS291. As such, the goal of this subtest was to determine whether a single fixed threshold could be used for filtering measurements based on the reflectivity value.

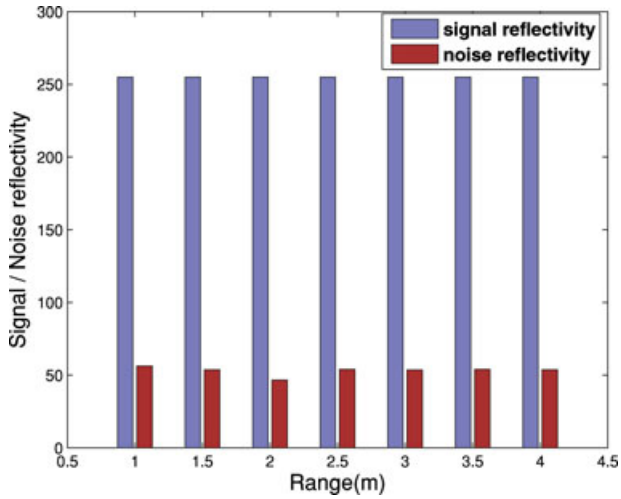

Figure 10. Mean reflectivity values for the fiducials and background at all baseline ranges. The fiducial reflectivity was maximum (255) for all scans across all test configurations.

The results showed that under all test configurations listed in Section 7.1, the reflectivity measurement value from the retroreflectors saturated the SICK's 8-bit buffer ($\gamma = 255$). In fact, we continued to increase the distance to the fiducials beyond the baseline conditions and determined that this would remain the case at ranges in excess of 14 m! This is far in excess of the planned tracking distance of <4 m, providing a significant safety margin for thresholding. Reflectivity "noise" levels (the reflectivity of the background) also seemed unaffected by the sunlight or sun orientation and remained at approximately 50 throughout all trials. Summary subtest results are presented in Figure 10. Whereas there is always potential for direct or reflected sunlight to dazzle the SICK, the integrated visor and occlusion by the wheelchair body has prevented this from happening to date. From this, we concluded that a single fixed threshold could be reliably used as one binary filter for fiducial segmentation.

7.1.2. Localization Accuracy

Scan data were evaluated to characterize the accuracy of the localization approach. A sample plot for nominal x–y coordinates of (2.0, 0) is shown in Figure 11 (left), and summary data for all positions are in Figure 11 (right). Assuming an unbiased estimator,

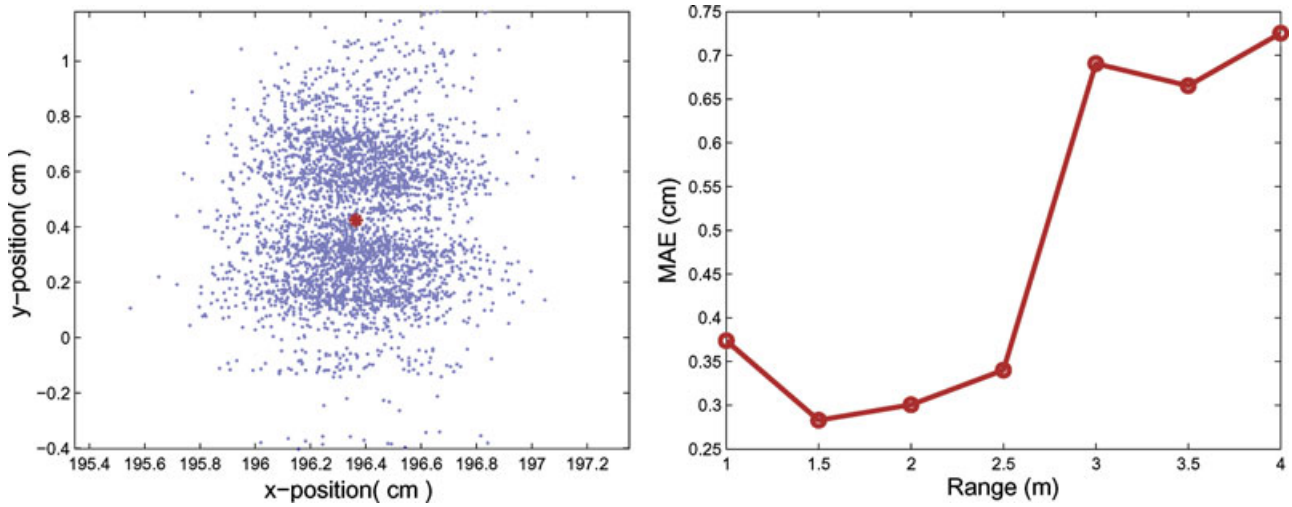

Figure 11. (Left) Distribution of wheelchair position estimates for a sample trial. (Right) Variance of x–y position and mean square error.

that is, the mean position of the distribution is the actual center of rotation, the mean absolute error at all ranges was <7.5 mm; at ranges less than 2.5 m, it was <4 mm. This is an order of magnitude smaller than the docking tolerance (≈4–5 cm), which highlights the efficacy of the localization approach.

7.1.3. Fiducial Segmentation Reliability

In conjunction with the localization testing, data were also collected to determine the number of LIDAR hits per fiducial as a function of target range. Although this could be estimated from the angular resolution and beam divergence specifications from SICK AG (2003), it was not known what fraction of the spot size would still result in saturation of the reflectivity measurement and as a consequence a fiducial detection. As such, empirical data were used to quantify this. To this end, more than 360,000 LIDAR scans were evaluated. A histogram showing summary data from one set of trials can be found in Figure 12. With the exception of 4 m, a minimum of two hits were recorded at every range for every scan. At the 4-m range, 99.46% of all scans had two or more hits. As the final LIDAR scan rate (75 Hz) will be far in excess of the control updates rate to the chair (15 Hz), these could easily be filtered so that a minimum of two features can be required to define a viable fiducial cluster for wheelchair localization.

7.2. Damaged Fiducial Tests

The purpose of these subtests was to simulate the effects of wear and tear on the retroreflector fiducials and estimate its impact on localization performance. During these subtests, the left retroreflector fiducial was covered on the rear side by a 2-cm black stripe. The front sides of both fiducials were left unchanged. This test setup is illustrated in Figure 13. As a result, when the front side of the turntable arm assembly was visible to the LIDAR, we collected data for the undamaged baseline configuration. The rear side corresponded to our "damaged" fiducial case. This allowed a simultaneous comparison without any other changes in the test configuration.

More than 10,000 scans were made to ensure a minimum set size of 5,000 samples for both the damaged and undamaged retroreflector configurations. These trials were conducted for $x = \{1.0, 2.0, 3.0, 4.0\}$ m, $y = 0.0$ m. A sample plot for nominal x–y coordinates of (1.0 m, 0.0 m) is shown in Figure 14 (top). Whereas the variance of the position estimates from the damaged fiducials was greater than that of the baseline system, mean absolute errors were still <7 mm for all trials. This is illustrated in the summary data in Figure 14 (center). We also observed a typical 3-mm y-position bias between the mean position of the two distribution sets at each range. This can also be seen in Figure 14 (bottom). Again, this also was not unexpected, as we deliberately covered only

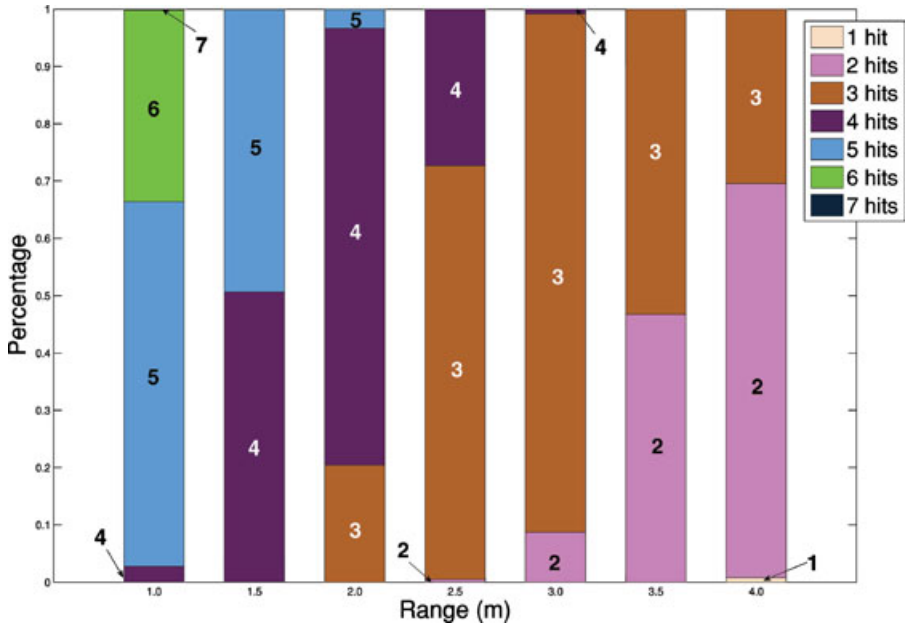

Figure 12. Histogram of scan hits per fiducial. Results indicate that a minimum of two hits can be required to validate a cluster, which aids in outlier rejection.

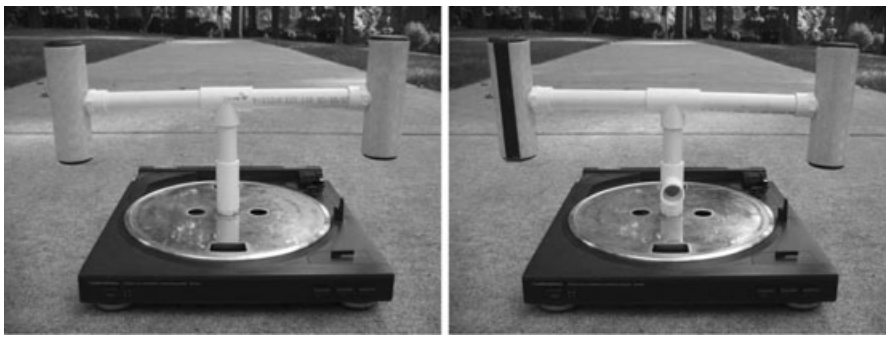

Figure 13. (Left) Undamaged fiducials on the front side. (Right) Simulated damaged fiducials on the arm's rear side. This allowed a simultaneous comparison of damaged fiducial position accuracy versus an undamaged baseline system.

one fiducial to investigate the bias effect as such errors cannot be removed through the EKF. However, even with this bias the y-position errors were still less than 1 cm, which again is far below our docking tolerances.

7.3. Headlight Interference Tests

The primary objectives of this test were as follows:

1. Determine whether secondary light sources (e.g., headlights from neighboring automobiles) could saturate the LMS291 reflectivity measurement or dazzle the SICK.
2. Determine whether the current filtering approach was sufficient to segment out the fiducials in the presence of headlight interference.

Two different subtests were conducted. The first was to determine whether headlights could in fact interfere with the LIDAR system. These were conducted with the SICK staring at an automobile with halogen headlights at ranges of 2.5, 5.0, 7.5, and 10.0 m. Data were then recorded without and

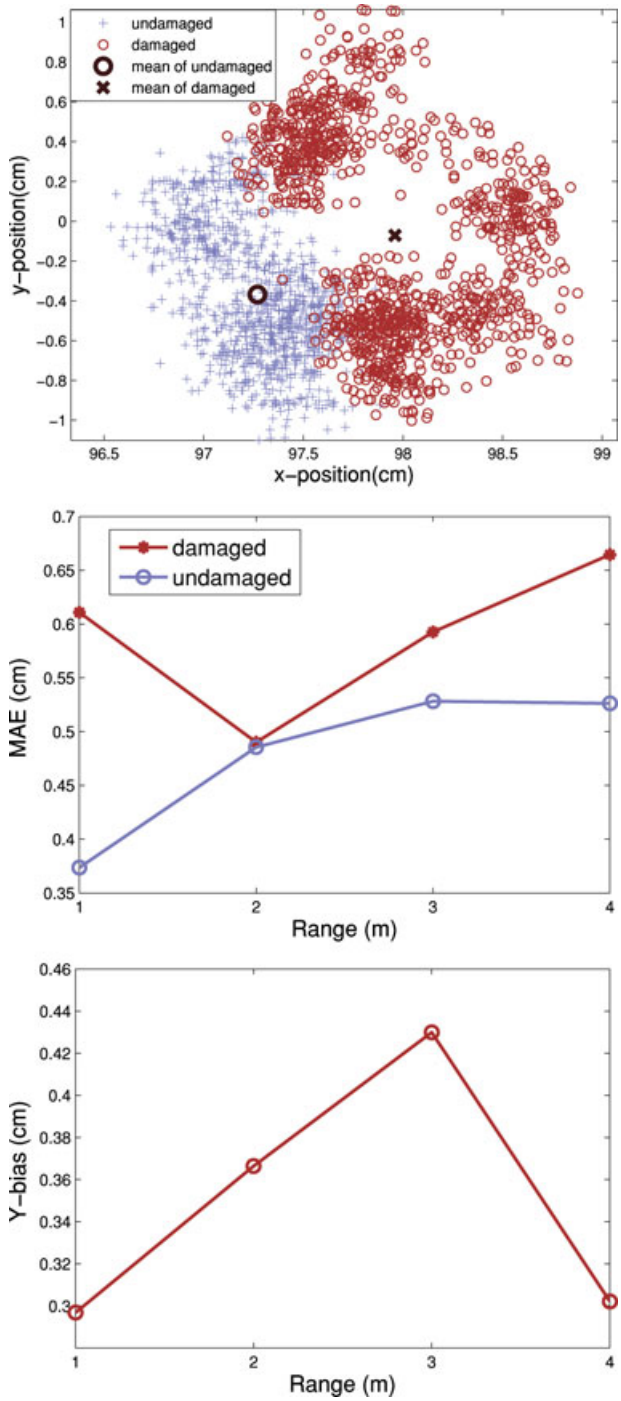

Figure 14. (Top) Position estimate distributions for damaged and undamaged fiducials at 1 m. (Center) Mean absolute error comparison for the damaged and undamaged fiducials. In both cases, this was still <7 mm for all trials. (Bottom) *y*-position bias from one-sided damage was 3–4 mm at all ranges.

with the headlights on to ensure a proper baseline. After this, an additional trial was conducted with the turntable assembly placed between the car and the SICK LIDAR to determine whether the fiducials could still be successfully segmented.

From these tests, we determined that under appropriate conditions the car headlights would both saturate the reflectivity measurement ($\gamma = 255$) and dazzle the SICK photodiode. This occurred at all ranges tested, and both phenomena could occur simultaneously. A sample trial from our fiducial tracking experiment is shown in Figure 15. The left-hand panel shows the raw reflectivity measurements. The headlight effects are labeled as noise. Note also the immediate drops in reflectivity after the noise spikes. These correspond to the SICK being dazzled. The right-hand panel shows the resulting fiducial segmentation. In this case, the segmentation is correct and the headlight effects are ignored based solely on range thresholds. It must be emphasized that in practice this will not always be the case. It would be quite common for vehicles to be within the 4.5-m-range threshold in parking lots. However, during testing it was also observed that always only a single bearing angle would have a $\gamma = 255$ reflectivity value required by the thresholding filter. As a result, requiring a minimum valid cluster size of two elements as outlined in Section 7.1.3. will also prevent this false positive from being identified.

7.4. Rain Testing

The objectives of these subtests were as follows:

1. Estimate the impact of rainfall on fiducial segmentation performance.
2. Determine the impact of rainfall on wheelchair localization accuracy.

To protect the turntable assembly electronics, these tests were conducted using two static fiducials. As the LMS291 integrates rain protection in the near field (ranges <2 m), rain testing focused on standoff distances of 3–4 m from the LIDAR. Trials were conducted under five different conditions: no rain (baseline) and light, medium, heavy, and very heavy rain (i.e., peak rainfall during a thunderstorm). The positions of the fiducials and the LIDAR were fixed throughout the testing process to ensure fair comparison of the localization accuracy.

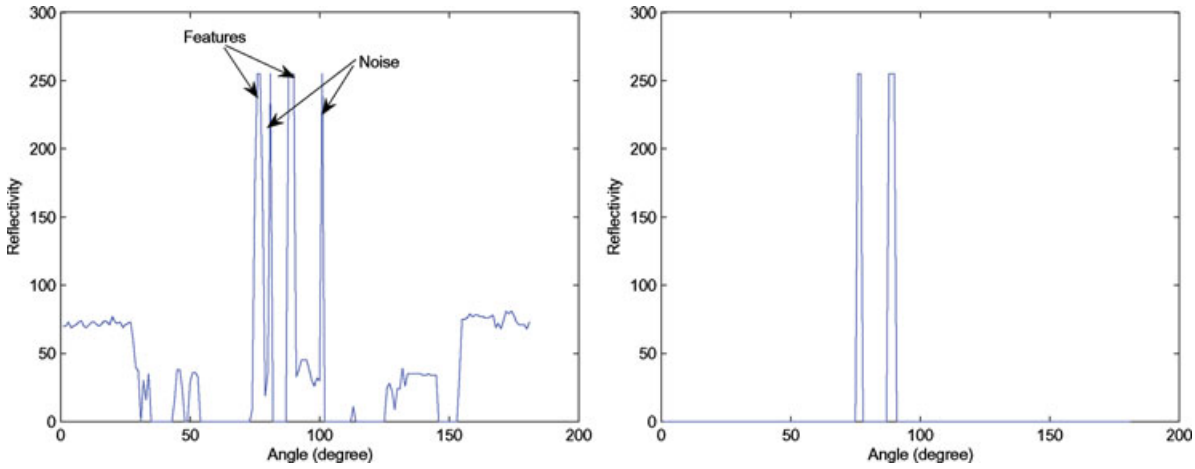

Figure 15. Raw (left) and filtered (right) reflectivity data from headlight interference testing. In this case, the range filter allowed the fiducials to be reliably segmented.

Figure 16 summarizes the mean absolute error for the four rain scenarios at 3-m distance. We see that light to medium rain had little to no effect on the feature segmentation and localization accuracy; mean absolute error was less than 5 mm. However, heavy rain increased the "random" error component to >2 cm.

A second impact of rainfall was that outliers in range measurements were observed. When this occurred, a significantly shorter than actual range was obtained. Bearing angle and reflectivity measurements remained consistent throughout the trials. As these outliers were not observed during testing under benign conditions, we inferred that raindrops caused sufficiently large returns so as to queue the "start of pulse" for the range measurement and the pulse was of sufficient length to register the associated fidicuial reflectivity. Figure 17 illustrates this effect during a sample rain trial, where the positions of both the left and right fiducial have been estimated from the raw range and reflectivity measurements.

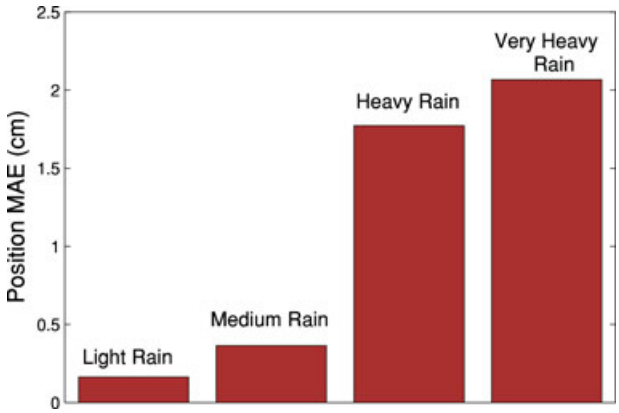

Figure 16. Mean absolute error for rain trials. This "random" error component increased dramatically with the rainfall rate.

Several significant outliers can be observed in the top fiducial position estimate. Without additional filtering, these outliers can lead to two different, seemingly valid—but incorrect—wheelchair poses. The first case stems from when the range error induces a failed fiducial segmentation (case A in Figure 17). The net result is a false cluster being introduced into the scan: one from the left and two from the right fiducials. In this case, by chance the two clusters from the right fiducial better matched the known fiducial baseline and are designated the optimal cluster pair. If this pair was accepted, the resulting localization result would be shown in Figure 18 (left).

The second, perhaps the more anticipated case, was where the fiducial segmentation was correct but the range error led to a large error in pose (case B in Figure 17). Such an error has the potential to pass through several filters currently used in the chair localization process, including using geometric constraints of known fiducial baseline. An example of this case is illustrated in Figure 18 (right).

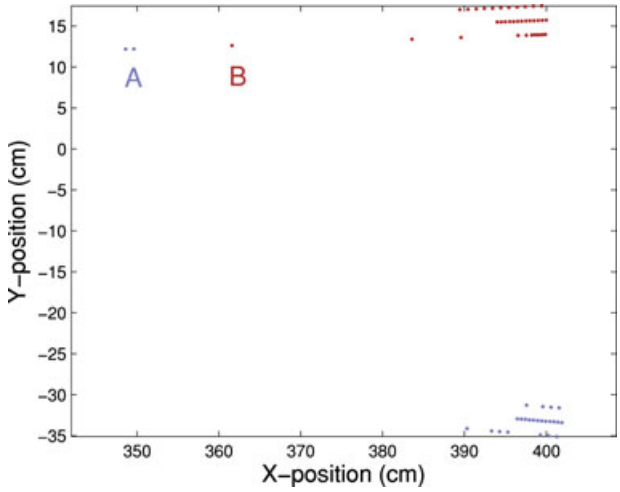

Figure 17. Fiducial position estimates during rainfall testing.

In this case, a bad range estimate to the right fiducial leads to dramatic position and orientation error. Fortunately, integrating either of our two validation gates—accepting only fiducial clusters with two or more hits (Section 7.1.3) and accepting only scans that have a single valid fiducial pair (Section 5.2)—would enable both of these cases to be rejected as outliers.

To further improve fiducial segmentation robustness and reject outlier range measurements, median filters have been added that track the raw range and bearing measurements. Filtered measurements are then used as validation gates for outlier rejection.

8. SYSTEM-LEVEL TESTING

For all system-level testing, ATRS was integrated into a Chrysler minivan. The power lift was a Freedom Lift Tracker^{TM} system. The wheelchair was Invacare Pronto M91 w/MK5 Electronics. This lift platform/power chair combination allowed for approximately ± 6 cm of clearance between the chair wheels and the lift rails for the docking task.

8.1. Public Demonstrations

There have been numerous demonstrations of ATRS both indoors and outdoors to potential end users, members of the media, industry personnel, and state and local government representatives. The largest of these included 3 days of continuous demonstrations at the World Congress Exposition on Disabilities (WCD 2006) in November 2006. This was done not only to verify docking robustness, but also to provide potential end users (conference participants) with an opportunity to use the system and provide feedback on such issues as human factors engineering and human–robot interaction (HRI). More than 300 cycles of docking and undocking were conducted during this time without a single failure. A sample trial is illustrated in Figure 19.

8.2. Outdoor Testing

ATRS has been informally tested outdoors during the past year of development. In this section, we summarize our most recent results, which consisted of 3 days of continuous system testing. Two days were dedicated to reliability testing, and the third sought to characterize the system performance under atypical conditions.

Reliability Testing In its current configuration, the intended operational environment for ATRS is paved, locally flat parking areas. This includes operations on moderate slopes so long as the surface between the lift platform and chair is approximately planar.

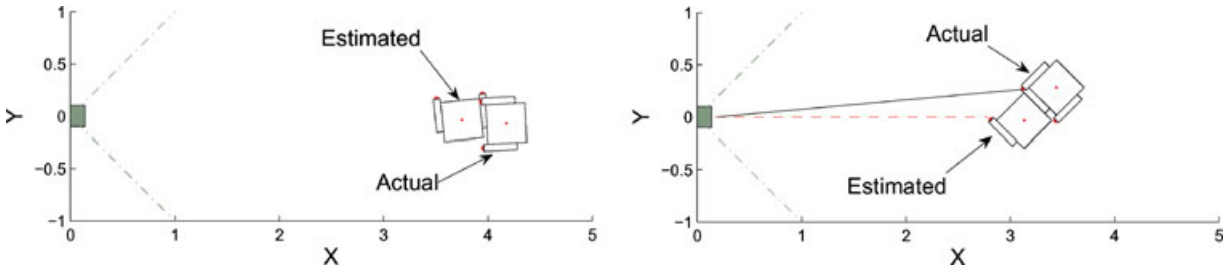

Figure 18. Chair pose outliers resulting from failed fiducial segmentation (right) and poor fiducial position estimate (left). Both were caused by range measurement outliers but are filtered by current validation gates.

Figure 19. ATRS demonstration at WCD 2006. The operator performs a seat-to-seat transfer (left), and then remotely controls the wheelchair to the vicinity of the lift platform (center). At this point, the LIDAR tracks the chair while sending real-time control inputs over a dedicated RF link. Autonomous operations conclude with the chair successfully docked (right).

Although this may appear restrictive, these requirements are significantly less constraining than the Americans with Disabilities Act (ADA) standards for accessibility. These specify a maximum slope for handicapped parking spaces not exceeding 1:50 (2%) in all directions (28 CFR Part 36, 1994). Whereas the beta ATRS design can accommodate significantly greater grades as well as departures from the ground plane assumption of nearly twice this magnitude, the 2 days of reliability testing were conducted in parking spaces conforming to ADA requirements.

Over the 2-day test period, 121 docking runs were made. A run was considered successful if the plough mounted to the wheelchair base successfully engaged the dock on the lift platform, thereby locking the wheelchair into place for safe transport. This could be objectively determined from visual inspection. (There is also an electronic indicator from the dock itself.) Of the 121 runs, the wheelchair successfully docked without further input from the operator on 118 trials for a 97.5% success rate. Ideally, we would like to see 100% success under such conditions. However, in each of the three failed runs the wheelchair could be recovered by the user through teleoperation via the UI's joystick. A subsequent attempt to dock was successful in each case, so we could consider the system to have achieved 100% docking reliability during this subtest.

Regardless, further examination of the failed runs is warranted. Figure 20 summarizes the test results, showing the initial and final positions of the wheelchair as estimated by the localization system. Successful runs are shown as + signs, and failed runs as x's. In one of the three failed trials, the culprit appears to have been our "heartbeat thread." This serves as a watchdog for communication failure and e-stops the chair if constant communication is not maintained between the wheelchair and van-side computers. In this case, the only thing required for resuming docking was for the user to again press the "dock" button at the UI. We expect this problem can be ameliorated through software changes alone.

In the remaining two failed runs, the wheelchair drove onto the dock with too much y-position error for the docking mechanism to engage the wheelchair plough. Examining the log files from these runs indicates that the PID motor controllers were dramatically failing to achieve the setpoints dictated by the linear and angular velocities specified by the motion planner. This can be attributed to a single fixed gain set being used for the motor controllers. Per the suggestion of the reviewers, we plan to integrate adaptive gains for the PID controllers to resolve this issue. It should also be noted that the go/no-go behavior discussed in Section 6 was not integrated during this testing. Although this would not have improved the docking performance, it would have simplified the teleoperation recovery task for the two failed trials.

Operational Testing in Nonconforming Environments The third day of testing was dedicated to testing ATRS in more challenging conditions than ADA-conforming parking spaces. The primary goal of these tests was to better characterize the operational envelope of the system. Specific subtests included docking with the lift platform elevated (from coming to rest on a brick approximately 5 cm thick), as well as docking trials on broken asphalt and sand/gravel-covered surfaces. The test conditions were sometimes

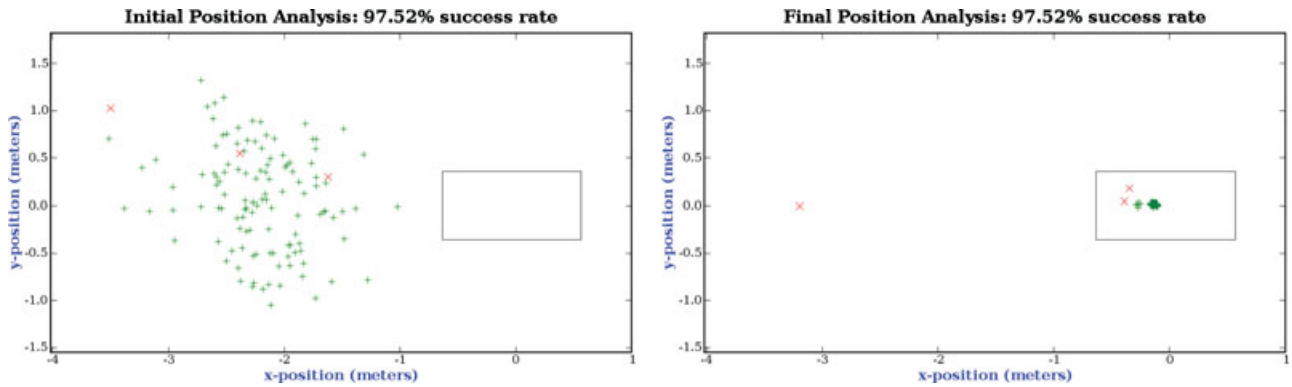

Figure 20. Summary results for 2 days of docking reliability testing showing the initial (left) and final (right) distributions of chair positions as estimated by the localization system. The three failed trials are marked with x.

further aggravated by the deliberate placement of sticks and stones in the wheelchairs's path, as well as the introduction of vehicle exhaust and sand/dust streams in front of the LIDAR window. Images from several trials can be found in Figure 21.

In general, docking performance under these conditions was good. If the initial pose of the wheelchair was reasonable, the chair could still dock reliably on these surfaces. Localization system performance was unaffected by the sand and dust streams, and the chair was able to reliably dock on the raised platform. However, it was also easy to deliberately construct a condition in which the chair would fail to dock. Docking was less robust to gross initial pose errors than during reliability testing. Furthermore, we observed significant wheel slippage from the sand and gravel. In one specific case, the wheelchair had to be manually recovered via the UI when the left wheel entered a "depression" in the ground and could not gain sufficient traction to exit.

Whereas these results are less objective than the reliability testing, they indicate that an automated wheelchair has sufficient mobility for docking on a wide range of surfaces and system docking reliability can be made quite good through the ability of the operator to recover the chair via teleoperation. However, we should emphasize that the scope of testing outdoors did not encompass the full range of environmental conditions anticipated during ATRS operations. Additional outdoor trials are planned over the coming months, which will help to more fully characterize the ATRS operational envelope.

8.3. Video Supplements

Several videos from ATRS outdoor testing, the WCD demonstration, the proof of concept demonstration, etc., can be viewed online at the project webpage http://vader.cse.lehigh.edu/projects/atrs/.

9. DISCUSSION

The primary objective of this work was to develop a reliable, robust means for autonomously docking a wheelchair onto a lift platform to eliminate the need for an attendant. Ultimately, this was accomplished through LIDAR-based localization and a hybrid control design. Under normal conditions, the localization paradigm achieved subcentimeter positional accuracy. It has also proven to be robust in heavy rain, as well as to significant damage to the retroreflector targets.

In parking areas conforming to the accessibility standards of the ADA, the overall system has proven to be reliable in outdoor testing to date. However, operations in environments that significantly deviate from these standards can cause issues. On surfaces with gravel or sand, significant wheel slippage has been observed. We are considering the integration of a gyroscope on the production system to detect these occurrences. On surfaces that deviate significantly from the ground-plane assumption, the LIDAR scan may miss the retroreflectors completely. The operational envelope in such areas is a function of the fiducial size. Unfortunately, growing the cylinder height significantly beyond the present configuration begins

Figure 21. (Clockwise from upper left): Docking trials with vehicle exhaust for laser interference and the lift platform elevated, on broken asphalt, and on a sand and gravel surface. The wheelchair successfully docked in each of these trials.

to compromise the aesthetics of the wheelchair. This is a significant design consideration, as ATRS will eventually become a commercial product. To accommodate a larger operational envelope, we are currently investigating two enhancements. The first is dynamically regulating the LIDAR tilt angle to accommodate deviations from the ground-plane assumption. The second is a structured lighting solution with a significant vertical field of view (e.g., 60 deg). Although our previous work with passive vision systems identified limitations in this application (Sermeno-Villalta & Spletzer, 2006), we expect that this active vision system paradigm will address the signal-to-noise issue associated with segmenting the wheelchair features.

We are also investigating reviewer suggestions for improving PID performance. System-level testing to date has examined road surfaces ranging from a "like new" parking lot to broken asphalt, gravel, and sand. Whereas docking performance has been acceptable, we have observed conditions in which the motor controllers fail to achieve velocity setpoints, sometimes by a significant margin. This was also identified as a source of docking failures during reliability testing. We are currently investigating adaptive parameter tuning and/or user input through the computer UI to improve motor controller performance. Additional reviewer comments regarding the potential for LIDAR dazzle at sunrise/sunset are also being investigated.

ATRS has the potential to be one of the first widespread applications of mobile robotics outdoors that operate in the vicinity of and in service to people. This is made possible by having a limited and well-defined requirement for autonomous operations. It also relies on a HRI model that leverages both human decision making and the accurate automation capabilities of ATRS. In this "supervised autonomy" approach, the operator provides a safety margin for ATRS operation in instances when the system is uncertain how to proceed and for docking recovery. Although not the focus of this paper, HRI aspects are key to system robustness and a primary enabler of ATRS commercialization. We expect that a similar model will be followed for the foreseeable future to facilitate human-robot integration.

ACKNOWLEDGMENTS

Special thanks to the entire ATRS Team, especially Mike Martin and Tom Panzarella, Sr. (Cook Technologies), and Humberto Sermeno-Villalta (University of San Salvador) for his early work on the system. Also, thanks are due the reviewers for providing valuable suggestions to improve both the paper and the system itself. Our research results are based on work supported by the National Science Foundation Partnerships for Innovation (PFI) Program under Grant No. 0650115. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation. This project was also funded in part by a grant from the Commonwealth of Pennsylvania, Department of Community and Economic Development.

REFERENCES

28 CFR Part 36 (1994). ADA Standards for Accessibile Design. U.S. Department of Justice.

Braun Corporation (2006). Products: ChairTopper. http://www.braunmobility.com.

Deluca, A., Oriolio, G., & Samson, C. (1998). Feedback control of a nonholonomic car-like robot. In Robot motion planning and control (pp. 171–253). Springer-Verlag.

Dissanayake, M. W. M. G., Newman, P., Clark, S., Durrant-Whyte, H. F., & Csorba, M. (2001). A solution to the simultaneous localization and map building (SLAM) problem. IEEE Transactions on Robotics and Automation, 17(3), 229–241.

Fusiello, A., Trucco, E., Tommasini, T., & Roberto, V. (1999). Improving feature tracking with robust statistics. Pattern Analysis and Applications, 2, 312–320.

Gao, C., I. Hoffman, T. P., & Spletzer, J. (2007). Automated transport and retrieval system (ATRS): A technology solution to automobility for wheelchair users. In Proceedings of the 2007 International Conference on Field and Service Robotics (FSR '07), Chamonix, France.

Gomi, T., & Griffith, A. (1998). Developing intelligent wheelchairs for the handicapped. In Lecture Notes in Computer Science, 1458, 150–178.

Howard, A., Parker, L. E., & Sukhatme, G. S. (2006). Experiments with large heterogeneous mobile robot team: Exploration, mapping, deployment and detection. International Journal of Robotics Research, 25(5), 431–447.

Leonard, J., & Durrant-Whyte, H. F. (1991). Mobile robot localization by tracking geometric beacons. IEEE Transactions on Robotics and Automation, 7(3), 376–382.

Lingemann, K., H. Surmann, A. N., & Hertzberg, J. (2005). High-speed laser localization for mobile robots. Journal of Robotics and Automation Systems, 51(4), 275–296.

Lionis, G. S., & Kyriakopoulos, K. J. (2002). A laser scanner based mobile robot SLAM algorithm with improved convergence properties (volume 1, pp. 582–587). In Proceedings of the 2002 IEEE International Conference on Intelligent Robots and Systems.

Miller, D., & Slack, M. (1995). Design and testing of a low-cost robotic wheelchair prototype. Autonomous Robots, 2(1).

National Council on Disabilities (2005). The current state of transportation for people with disabilities in the United States. Washington, DC: National Council on Disabilities. Available at http://www.ncd.gov/newsroom/publications/2005/current_state.htm.

NHTSA (1997). Research note: Wheelchair user injuries and deaths associated with motor vehicle related incidents. Washington, DC: National Highway Traffic Safety Association. Available at http://www.nrd.nhtsa.dot.gov/Pubs/97.850.pdf.

Oriolo, G., Luca, A., & Vendittelli, M. (2002). WMR control via dynamic feedback linearization: Design, implementation and experimental validation. IEEE Transactions on Control Systems Technology, 10(6), 835–852.

Parikh, S., Grassi, V., Kumar, V., & Okamoto, J. (2004). Incorporating user inputs in motion planning for a smart wheelchair (volume 2, pp. 2043–2048). In Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), New Orleans, LA.

Parikh, S., Rao, R., Jung, S., Kumar, V., Ostrowski, J., & Taylor, C. (2003). Human robot interaction and usability studies for a smart wheelchair (pp. 3206–3211). In Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), Las Vegas, NV.

Pennsylvania Rehabilitation Council (2006). PARC annual report: Position papers. http://www.parac.org.

Sermeno-Villalta, H., & Spletzer, J. (2006). Vision-based control of a smart wheelchair for the automated transport and retrieval system. In Proceedings of the 2006 IEEE International Conference on Robotics and Automation (ICRA '06), Orlando, FL.

SICK AG (2003). Technical manual, LMS2xx laser measurement systems. Reute, Germany: SICK AG.

Simpson, R., & Levine, S. (1999). Automatic adaptation in the NavChair assistive wheelchair navigation system. IEEE Transactions on Rehabilitation Engineering, 7(4), 443–451.

Stern, S., & Brault, M. (2005). Disability data from the American Community Survey: A brief examination of the effects of question redesign in 2003. Washington, DC: U.S. Census Bureau.

Wood, T. (2004). When your plan is a van. Quest—Journal of the Muscular Dystrophy Association, 11(2). Available at http://www.mda.org/publications/Quest/q112vanplan.html.

Yanco, H. (1998). Wheelesley, a robotic wheelchair system: Indoor navigation and user interface. Lecture Notes in Artificial Intelligence: Assistive Technology and Artificial Intelligence, 1458, 256–268.