

An Automated Asset Locating System (AALS) with Applications to Inventory Management

Thomas H. Miller, David A. Stolfo, and John R. Spletzer

Abstract In this work, we present a proof-of-concept Automated Asset Locating System (AALS) for enhancing inventory management. AALS integrates LIDAR and RFID sensor measurements into a Rao-Blackwellized particle filter for simultaneously localizing its pose with the positions of assets in the environment. We present significant experimental results where the proof-of-concept system successfully traveled a total distance of 1.4 km autonomously, while detecting and mapping all 143 available assets in real-time, and with a mean position error of < 80 cm.

1 INTRODUCTION

Radio Frequency Identification (RFID) systems use radio frequency to identify, locate and track features of interest. The technology sees widespread use in commercial applications to include baggage handling, passport readers, and toll collection to name but a few [1]. There are several RFID variants: passive, semi-passive, and active. In this work, we limit our discussion to the former. A passive RFID system is composed of three primary components: a reader (RF transmitter/receiver), a passive tag, and a host computer. The tag is composed of an antenna coil and an integrated circuit that contains both modulation circuitry and non-volatile memory. The tag is energized by the RF carrier signal transmitted by the reader. Using this scavenged energy, the information stored on the tag – to include a unique identifier for that tag instance – can be transmitted back to the reader [2]. The strength of RFID is that it explicitly solves the data association problem. As each tag is associated with a unique identifier, false correspondences across tag detections are eliminated.

In this work, we investigate the potential for applying RFID and robotics technologies to inventory management tasks. Manual intervention in material tracking systems is labor intensive, costly, and error-prone [3]. Furthermore, low-frequency “scheduled scanning” approaches cannot ensure that inventory remains up-to-date. The ability to automate the material tracking task can dramatically enhance as-

T. Miller, D. Stolfo, and J. Spletzer
Lehigh University, Bethlehem, PA, USA, e-mail: thm204,das611,josa@lehigh.edu

set visibility. To this end, we demonstrate an Automated Asset Locating System (AALS) that integrates LIDAR and RFID sensing on a mobile robot base for enhanced inventory management. The RFID system’s role is dual purpose. First, the tags serve to identify assets to be tracked. Second, they are integrated into the environment as correspondence-free landmarks. In this role, they effectively introduce dramatic, artificial asymmetries into the environment. This enables reliable robot localization indoors even in largely symmetric environments, and where the scale of the environment was large compared to the range of the robot’s sensors – conditions which could be problematic for traditional SLAM and localization approaches. The RFID tag’s extremely compact size ($\approx 10\text{-}30\text{ cm}^2$ stickers) and low cost ($\$0.1\text{-}1.0$) allows them to be discretely integrated into the environment. The net result is an automated system capable of reliably locating assets in the environment.

2 RELATED WORK

Several researchers have investigated the convergence of robotics and RFID technologies. Most related to our work is that of Hähnel *et al* [4], where a Pioneer 2 Robot equipped with a Sick LMS200 and an RFID reader was manually steered through the environment. Using a map generated *a priori*, the authors employed Monte-Carlo localization (MCL) to estimate the position of RFID tags in the environment. Formal results on tag localization accuracy were not provided. However, they demonstrated that using these same tags as landmarks, robot localization could be achieved using only RFID measurements (although not to the same level of accuracy as when the LIDAR system was used). Schneegans *et al* [5] built on this to demonstrate a system for robot localization using a more sophisticated sensor model, and whereby an RFID snapshot was associated with a database of learned features. They compared their approach with those from [4], and found comparable accuracy in the end position estimate of the robot, but a significantly faster filter convergence rate. This work was also done off-line.

There is also significant work that has emphasized using RFID to assist in localization tasks. Kulyukin *et al* incorporated RFID into a robotic assistant for the visually impaired [6]. Tsukiyama demonstrated a limited implementation where RFID tags served as topological landmarks enabling the robot to correctly follow a path [7]. Mapping the position of assets was not considered. Chae and Han used a topological approach with RFID and a vision sensor [8]. Experimental results were again off-line. Miah and Gueaieb examined using tag received power (TRP) to estimate the distance from the robot to the tag [9]. However, their implementation was limited to simulations. Milella *et al* developed an RFID-assisted mobile robot system for mapping and surveillance using fuzzy inference methods [10]. In terms of asset tracking task, Ehrenberg *et al* investigated the use of a LibBot to locate books in a library environment [11]. They localized densely packed, short range tags by again employing a probabilistic RFID antenna model. The actual implementation was rather limited however, with experiments only over a single library shelf.

Our work differs from these efforts in several ways. First, we employ a Rao-Blackwellized particle filter for the simultaneous localization of the robot pose and

mapping of asset positions in the environment. Second and more significantly, unlike these efforts we provide significant experimental results with AALS operating on-line. In our experiments, AALS is completely responsible for its own navigation as it self-localizes and maps the positions of assets in the environment. These results show that AALS is capable of reliably detecting and mapping the position of assets in the environment in real-time.

3 THE DEVELOPMENT PLATFORM

The AALS proof-of-concept system was built upon an iRobot Create robotics development platform. The Create is an excellent low-cost research platform, combining a robust mobile chassis with a higher level motor control interface through RS-232 communication, odometry feedback, limited sensing, and 5V DC power output. The other primary components of AALS are:

Computing With the exception of motor control which ran on the Create's embedded computer, all computing was done on a Lenovo X200 laptop with a 2.4 GHz Core 2 Duo processor and 2 GB memory.

LIDAR. The primary exteroceptive sensor for AALS was a Hokuyo URG-04LX LIDAR. The URG-04LX provides a 240° field of view with an angular resolution of 0.36°. It offers an advertised range of up to 5.6 meters, although in this application we found a more accurate estimate to be ≤ 4.5 meters.

RFID. The RFID transceiver used in this work was a Skyetek M9 operating at 862-955 MHz. We deliberately chose an UHF module to maximize range. The reader was multiplexed to a pair of antennae oriented to maximize detection coverage to the front and sides of the robot. To date, all development has been done using the Alien Technology ALN-9534 Gen 2 tag. In an evaluation of available Gen 2 tags, this model provided acceptable detection ranges (up to 4.0 meters) while exhibiting fairly good omnidirectional performance in a compact footprint. Images of AALS, showing the integration of on-board computing, the URG-04LX, the M9, multiplexer, and antennae are at Figure 1.



Fig. 1: Top and side profiles of AALS showing the integration of Hokuyo URG-04LX LIDAR, RFID reader, and on-board computing.

4 ROBOT LOCALIZATION & ASSET TRACKING

For robot localization and asset tracking, we employed a Rao-Blackwellized Particle Filter (RBPF). Such approaches were first introduced to the robotics community by Doucet *et al* [12], who observed that the simultaneous localization and mapping (SLAM) problem could be factored into two sub-problems

$$p(x_{1..t}, l_{1..m} | z_{1..t}, u_{1..t-1}) = p(x_{1..t} | z_{1..t}, u_{1..t-1}) \prod_{i=1}^m p(l_i | x_{1..t}, z_{1..t}) \quad (1)$$

where $x_{1..t}$ denotes the robot pose over time, $l_{1..m}$ the m landmark positions, z the sensor measurements, and u the control inputs. The left term on the right side of (1) corresponds to the robot localization problem, and the right term to estimating the position of m conditionally independent landmarks in the map. This partitioning enabled the robot localization problem to be solved using a traditional particle filtering approach, while allowing the mapping problem to be estimated through analytical methods. The significance of this factorization was that it mitigated the otherwise exponential increase of particle samples with increases in state space dimension (*i.e.*, the number of landmarks). This result was leveraged by Montemerlo *et al* in developing FastSLAM [13], where mapping was accomplished by associating m Extended Kalman Filters (EKF) with each particle to independently track the m landmarks $l_{1..m}$. We employed a similar approach, using Monte-Carlo Localization (MCL) to estimate the robot pose, and Kalman Filters for asset tracking.

4.1 Sensor Model Development

The LIDAR Sensor Model. AALS relies heavily upon the Hokuyo URG-04LX for localization. The URG-04LX is extremely compact and lightweight compared to the ubiquitous Sick LMS2xx series LIDARs, which made it well suited for our proof-of-concept. However, they are also myopic, demonstrating an effective range of ≤ 4.5 meters in our experiments. Such limited range can be a challenge for MCL approaches, which solves a data association problem relating the robot pose $[x(t), y(t), \theta(t)]^T$ vs. time through asymmetries in the environment. Our development site consisted of two building wings connected by a corridor ≈ 40 meters in length with little asymmetry. To mitigate the potential for the filter converging too quickly (and likely incorrectly), the conditional density functions which model the uncertainty in LIDAR measurements (and are used to weight the individual particles) were dramatically smoothed. As a result, even a relatively improbable measurement was unlikely to penalize a particle dramatically. We found that such a PDF would ensure that the robot’s pose would eventually converge to the correct position/orientation even without input from the RFID sensors and regardless of the initial robot pose.

The RFID Sensor Model. Two different models were used for the RFID sensor depending upon its given role. The primary purpose of landmark tags (with positions known *a priori*) was to provide a low-cost mechanism for enhancing localization ro-

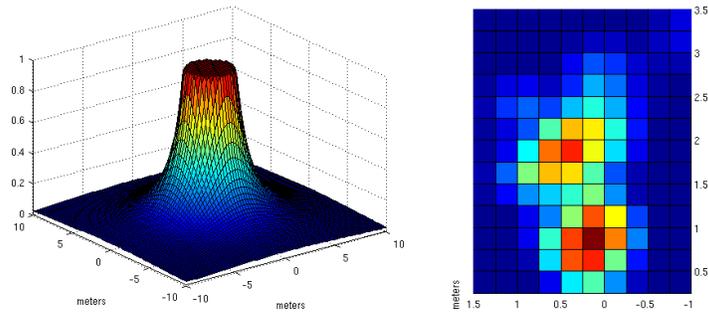


Fig. 2: (Left) Weighting function for landmark tags used in the robot localization process. (Right) PDF for asset tag detection generated empirically.

business, as there was no potential for data association errors. We considered their ability to assist in pose estimates and improve filter convergence as demonstrated in [4] of secondary importance. Therefore, we assumed no relative orientation information was available and a symmetric scaling function S was used to reflect the likelihood of landmark detection by the robot. To model this, we defined a critical radius r^* around each landmark where detection was expected based upon empirical results. With r^* so defined, the weight function used was

$$S(i) = \left[(d(i) < r^*) + (d(i) > r^*) \frac{r^*}{d(i)^2} \right] \quad (2)$$

where $d(i) = \|(x,y)^T - (x_i,y_i)^T\|$ was the Euclidean distance from the robot to the i^{th} landmark. When used in conjunction with the MCL process, particles within the critical radius of a detected landmark are unaffected, while the weights of those outside are scaled inversely proportional to the squared distance to the landmark. This is illustrated at Figure 2 (left). The motivation for the quadratic model is the Friis Transmission Equation, which shows that the power ratio between receiving and transmitting antennae are inversely proportional to their distance squared [14]. The placement of only several landmark tags in the environment dramatically accelerated particle filter convergence during our experiments.

For asset detection, we assumed that the estimated robot pose was approximately correct. As such, the sensor model was directional to reflect the relative robot/asset tag orientation. We initially generated a discrete PDF model empirically by collecting detection data as a function of tag position, orientation, and height as in [5]. The resulting two-dimensional PDF estimate in the antenna frame is shown at Figure 2 (right). The PDF is highly non-Gaussian, and does not lend itself to a Kalman filter implementation. However, in reality this model – as well as those typically used in related work – is ad-hoc. Antenna performance is strongly environment specific. Signal is strongly tied to reflections from the floor, walls, ceiling, obstacles, signal absorption, the amount of metal in the environment, tag line-of-sight, the object to

which the tag is affixed, *etc.*. In fact, in preliminary testing we compared a voting approach based upon our discrete PDF model with a pure Kalman filter using an overly conservative approximation of this PDF. The latter demonstrated equal or better performance, and as such we ultimately employed such an approach.

4.2 Robot Localization & Asset Position Estimation

For the most part, robot localization was accomplished using a traditional MCL approach [15]. The time update phase corresponded to the transformation of the particles' poses using a unicycle model for robot motion. Measurement updates using the LIDAR were also straightforward. However, an additional measurement update stage was integrated for whenever a landmark tag was detected. In this event, samples were re-weighted based upon $w_{k+1}(i) = S(j)w_k(i)$ where $w_k(i)$ denotes the current weight of the i^{th} particle at time-step k , and $S(j)$ the scaling function defined by (2). After re-weighting, the particle set was re-sampled. The net effect was that particles far away from landmark j were quickly killed off.

With the ability to reliably localize the robot, we turn to the case of mapping assets. To this end, each particle p_i , $i = 1 \dots n$, in our RBPF maintains a Kalman filter that propagates an estimate for the position and positional covariance $\{\mathbf{x}(i, j), \Sigma(i, j)\}$, $j = 1 \dots m$, for each of the m assets detected. Note that RFID asset detections are *not* used to refine the robot pose estimate, so the asset position estimates remain uncorrelated. As a result, only n of the mn total Kalman filters need be updated for a given asset detection.

We model each RFID asset detection as a direct estimate of the asset's position, *i.e.*, $z = {}^W T_A \mathbf{x}_A$ where \mathbf{x}_A is the tag position estimate in the antenna frame, and ${}^W T_A$ maps points from the antenna frame to world frame. The associated measurement covariance is then $\Sigma_R = R(\theta_R + \theta_A)\Sigma_A R(\theta_R + \theta_A)^T$ where Σ_A denotes the estimated uncertainty in the antenna frame, and R is a 2-D rotation matrix associated with the robot θ_R and antenna θ_A orientations in the world and robot frames, respectively. The measurement update is then textbook Kalman Filter, and since the asset position is assumed static there is no process update.

5 EXPERIMENTAL RESULTS

5.1 Component Level Testing

As part of the proof-of-concept, we performed component level testing to determine the robustness of tag detection as a function of tag density. Of concern was the potential for message collisions if multiple irradiated tags attempted to transmit at the same time. To this end, we examined both linear arrays of 5-15 tags, and grid arrays of 12 tags (3×4) with inter-tag spacings ranging from 0-45 cm. This also included different heights above the ground plane. A representative linear array configuration with 10 cm spacing is shown at Figure 3 (left). For each test geometry, AALS was driven multiple times past the tag array at standoff distances consistent with an

expected detection based upon the sensor model derived in Section 4.1. A tag was considered detected if it was successfully identified at least one time while AALS traversed the array. Summary statistics are shown at Figure 3 (center).

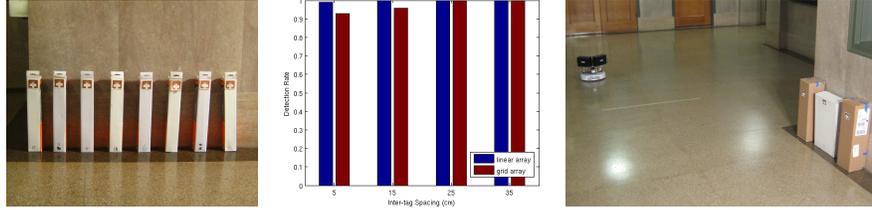


Fig. 3: (Left-Center) Sample RFID linear array used during component level testing. All configurations demonstrated at least a 93% success rate. (Right) Sample asset configuration during system level testing.

There were 908 true positives, 17 false negatives, and 0 false positives. Sixteen of the 17 false negatives were with grid arrays with inter-tag spacing of 5 cm (14) and 15 cm (2). These corresponded to tag densities of 100 and 30 tags/m², and detection rates were 93% and 96%, respectively. These results indicate that the anti-collision protocols employed by the system worked very well for the range of geometries tested even under very high tag densities.

5.2 System Level Testing

To demonstrate the system level proof-of-concept, we conducted a series of experiments using the fourth floor of Packard Laboratory at Lehigh University as the development site. This constituted a region $\approx 48 \times 14$ meters. Our map M representation was an occupancy grid with a cell resolution of 10 cm, and was provided to AALS *a priori*. The map was constructed from digital blue prints. While nominally correct, there were significant inconsistencies between this map and the actual floorplan. Only the most serious of these were corrected. One final alteration to M included the introduction of 4 landmark tags with positions also known *a priori* by the robot. These were spaced approximately every 15 meters in our corridor set. Finally, 10-15 assets (*i.e.*, cardboard boxes and plastic bins with tags affixed) were placed in random locations throughout the environment. A representative configuration is at Figure 3 (right).

For global path planning, AALS was provided a route network graph $G(V, E)$ that delineated in continuous space the intended paths for navigation. Waypoints in the route network corresponded to vertices $v_i \in V$ of G , and the edge set $E \subseteq G$ corresponded to path segments where each $e_{ij} \in E$ connected a pair of waypoints (v_i, v_j) . The desired path for a given mission was then specified via a waypoint sequence (v_i, v_j, \dots, v_n) . For motion planning, AALS relied upon 2 modes: obstacle avoidance, and path following. Prior to particle filter convergence or in the event that the specified route segment was blocked, AALS would operate in obstacle avoidance

mode. For path-following, a PD controller was used where the normal distance to the current route segment was employed as an error metric. The typical mission for AALS entailed a complete circuit of the test area. This corresponded to a mission length of ≈ 125 meters.

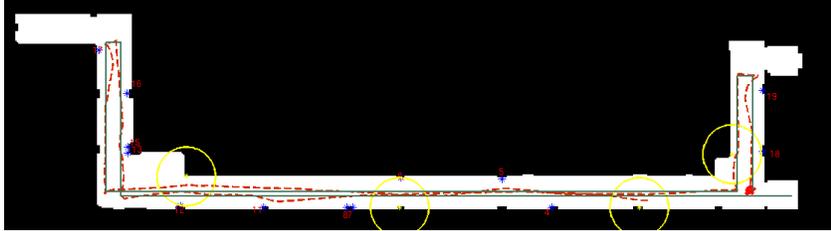


Fig. 4: Mission results showing the actual (blue “*”) and estimated (red “numbers”) asset locations. The mean position error in this trial was 54 cm.

After preliminary testing to characterize the system, a total of 12 missions were conducted. During these trials, the starting point was varied, as were the position and orientation of assets. This ensured that asset detection and mapping was possible with tag orientations parallel and orthogonal to the robot path. The geometry changes were also done to ensure that the sensor model for the Kalman filter was not deliberately biased. For each mission, AALS drove at a nominal linear velocity of 0.3 m/s. At the initiation of each trial, 10,000 particles were used to instantiate the prior for the robot pose. This number was reduced dynamically to as few as several hundred particles using the second-order statistics to infer convergence of the particle set. To further support real-time computation, LIDAR range measurements were sub-sampled to an angular resolution of 1.08° . The target update rate for AALS was 2 Hz. At the conclusion of a given mission, the estimate for the position of assets was determined from

$$\begin{bmatrix} x \\ y \end{bmatrix}_i = \sum_{j=1}^n w(j) \begin{bmatrix} x \\ y \end{bmatrix}_{ij}, \quad i = 1 \dots m \quad (3)$$

where $[x, y]_i^T$ denotes the position of the i^{th} landmark, $[x, y]_{ij}^T$ the i^{th} landmark position as estimated by the Kalman filter of the j^{th} particle, and w_j is the corresponding sample weight at mission completion. Results from a representative mission are at Figure 4. This shows the route network (green lines), the path as estimated by the robot (red dashed lines), the position of landmark tags (yellow circles), and the actual (blue “*”) and estimated (red “numbers”) positions of assets.

Of the 12 missions, 11 were completed successfully. The one failure occurred when an asset was deliberately placed across the path. The motion planner incorrectly determined the path was not traversable, and aborted the mission. The motion planner was subsequently modified, and this same configuration was successfully re-tested. The 11 completed missions constitute a total distance traveled of 1.4 km.

During this time, all 143 assets that were placed in the environment were detected. The estimated asset positions were then compared with hand-measured ground-truth values. Statistics for the different configurations are shown at Table 1. Border and interior configurations discriminate as to whether the asset was located on the map border or in the interior. Parallel/normal to path refers to the antenna orientation with respect to the robot’s primary direction of travel.

Asset Configuration	Number Samples	Number Detected	MPE (cm)	σ (cm)
All	143	143	79.2	49.5
Border	130	130	79.4	49.8
Interior	13	13	77.1	48.3
Parallel to path	91	91	86.8	55.5
Normal to path	52	52	65.8	33.3

Table 1: Mean Position Error (MPE) for detected assets as a function of geometry.

From these, we see that the average position error was <80 cm. There was little difference between assets that were located within the interior or along the border of the map (we should note that no optimizations were done to asset location estimates that were outside the boundary of the map, which would have improved results). We do note a fairly significant difference between tag orientations that were normal vs. parallel to the robot’s direction of travel. This appears to be attributed to the normal antennae being detected at longer standoff distances, and the associated Kalman filters seeing a larger number of measurement updates. However, further analysis is needed to support this hypothesis. We should note that in a warehouse or similar environment where such a system would be used, tag orientation would typically be parallel to the direction of travel and as such these errors are more representative.

For portions of three trials, we also estimated robot position using a Sick LMS291-S14 to track a retro-reflector affixed to the robot. Using this as ground truth, the mean absolute position error of the robot localization system was 53.3 cm ($\sigma_x = 49.3$ cm, $\sigma_y = 19.1$ cm). The bias was not surprising due to the strong symmetry and limited configuration space in the x and y directions, respectively. Taking these findings into consideration, a more accurate estimate of tag localization performance would be a MAE of ≈ 60 cm.

6 DISCUSSION

In this work, we demonstrated a proof-of-concept Automated Asset Locating System (AALS) that integrates LIDAR and RFID sensing on a mobile robot base. The RFID system’s role was dual purpose in this application – identifying both asset and landmarks tags in close proximity to the robot platform. These measurements enabled the position of asset tags in the environment to be estimated with a mean error of <80 cm. Furthermore, they were able to augment the limited range of the Hokuyo URG-04LX by not only accelerating the filter’s convergence rate, but also

ensuring against divergence in areas of low feature asymmetry. A natural question regarding this approach is the use of MCL vs. SLAM. This decision was made so that landmark tags with known “absolute” positions in the map could readily be integrated to protect against localization failures (*e.g.*, incorrect loop closures). We are currently investigating a hybrid approach which integrates both aspects, and working with members of the NSF Center for Engineering Logistics and Distribution to evaluate AALS in a larger scale, representative environment.

Acknowledgments. Special thanks to Mr. Sean Kelly (Lehigh University) for his efforts. This work was funded in part as a National Science Foundation Center for Engineering Logistics and Distribution (CELDi) center designated project. 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 NSF.

References

1. Federal Trade Commission, “RFID: Applications and Implications for Consumers,” Workshop Report, March 2005.
2. Pete Sorrells, “Passive rfid basics,” Tech. Rep. AN680, Microchip Technology, Inc., 2002.
3. “RFID Solutions & Applications,” <http://www.gaorfidassettracking.com>, 2009.
4. D. Hähnel, W. Burgard, D. Fox, K. Fishkin, and M. Philipose, “Mapping and localization with RFID technology,” in *ICRA*, 2004.
5. Sebastian Schneegans, Philipp Vorst, and Andreas Zell, “Using RFID snapshots for mobile robot self-localization,” in *Proceedings of the 3rd European Conference on Mobile Robots (ECMR 2007)*, Freiburg, Germany, September 19-21 2007, pp. 241–246.
6. V. Kulyukin, c. Gharpure, J. Nicholson, and S. Pavithran, “RFID in Robot-Assisted Indoor Navigation for the Visually Impaired,” in *IROS*, 2004.
7. H. Chae and K. Han, “World map based on rfid tags for indoor mobile robots,” in *Proceeding of SPIE - The International Society for Optical Engineering*, 2005, vol. 6006.
8. H. Chae and K. Han, “Combination of rfid and vision for mobile robot localization,” in *Proceeding of the 2005 International Conference on Intelligent Sensors, Sensor Networks and Information Processing Conference*, 2005, pp. 75–80.
9. Md. Suruz Miah and Wail Gueaieb, “Mobile robot navigation using custom-made RFID tag system,” in *Proceedings of the 5th Scientific Research Outlook*, Fes, Morocco, October 26–30 2008.
10. A. Milella G. Cicirelli and A. Distante, “Rfid-assisted mobile robot system for mapping and surveillance of indoor environments,” *Industrial Robot: An International Journal*, vol. 35, no. 2, pp. 143–152, 2008.
11. Christian Floerkemeier Isaac Ehrenberg and Sanjay Sarm, “Inventory management with an rfid-equipped mobile robot,” in *Proceedings of the 3rd Annual IEEE Conference on Automation Science and Engineering*, Scottsdale, AZ, USA, Sept 2007, pp. 1020–1026.
12. Arnaud Doucet, N. de Freitas, K. Murphy, and S. Russell, “Rao-blackwellised particle filtering for dynamic bayesian networks,” in *In Proceedings of the Sixteenth Conference on Uncertainty in Artificial Intelligence*, 2000, pp. 176–183.
13. M. Montemerlo, S. Thrun, D. Koller, and B. Wegbreit, “Fastslam: A factored solution to the simultaneous localization and mapping problem,” in *In Proceedings of the AAAI National Conference on Artificial Intelligence*. 2002, pp. 593–598, AAAI.
14. C. Balanis, *Antenna Theory: Analysis and Design*, Wiley, 2005.
15. S. Thrun, D. Fox, W. Burgard, and F. Dellaert, “Robust monte carlo localization for mobile robots,” *Artificial Intelligence*, vol. 128, no. 1-2, pp. 99–141, 2000.