Planning with Uncertainty in Position Using High-Resolution Maps

Juan Pablo Gonzalez and Anthony (Tony) Stentz

The Robotics Institute
Carnegie Mellon University
5000 Forbes Ave, Pittsburgh PA 15213, USA
{jgonzale, axs}@ri.cmu.edu

Abstract— **We present a novel approach to mobile robot navigation that enables navigation in outdoor environments without GPS. The approach uses a path planner that calculates optimal paths while considering uncertainty in position and that uses landmarks to localize the vehicle as part of the planning process. The landmarks are simple, possibly aliased, features that have been previously identified in a high-resolution map. These landmarks are combined with an estimate of the position of the vehicle to create unique and robust features. This approach reduces or eliminates the need for GPS and enables the use of prior maps with imperfect map registration.**

I. INTRODUCTION

Navigating autonomously is probably the most important problem facing outdoor mobile robots. This task can be extremely difficult if no prior information is available and would be trivial if perfect prior information existed. In practice prior maps are usually available, but their quality and resolution varies significantly.

When accurate, high-resolution prior maps are available and the position of the robot is precisely known, many existing approaches can reliably perform the navigation task for an autonomous robot. However, if the position of the robot is not precisely known, most existing approaches would fail or would have to discard the prior map and perform the much harder task of navigating without prior information.

Most outdoor robotic platforms have two ways of determining their position: a dead-reckoning system and a Global Position System (GPS). The dead reckoning system provides a locally accurate and locally consistent estimate that drifts slowly, and the GPS provides globally accurate estimate that does not drift but is not necessarily locally consistent. A Kalman filter usually combines these two estimates to provide an estimate that has the best of both position estimates.

While for many scenarios this combination suffices, there are many others in which GPS is not available or its reliability is compromised by different types of interference such as mountains, buildings, foliage or jamming. In these cases, the only position estimate available is that of the dead-reckoning system, which drifts with time and does not provide a position estimate accurate enough for most navigation approaches.

This paper presents a new approach to mobile robot navigation that addresses some of the issues mentioned above. We propose a path planner for autonomous ground vehicles that calculates optimal[1] paths while considering uncertainty in position, and uses landmarks to localize the vehicle as part of the planning process. The planner uses simple, possibly aliased, features that have been previously identified in a high-resolution map, and combines them with an estimate of the position of the vehicle to create unique and robust features. This approach reduces or eliminates the need for GPS and enables the use of prior maps with imperfect map registration.

II. RELATED WORK

While many existing approaches provide robust localization, few combine localization with the path planning process. To the best of our knowledge, the approach presented here is the first to address the combined challenges of planning with uncertainty and using landmarks to reduce uncertainty while optimizing a continuous objective function in an outdoor setting.

There is however, significant work on some of the parts of the problem: In the field of classical path planning, Latombe [2][3] has an extensive review on the state of the art as of 1991. Since then, important contributions by Lazanas and Latombe [4], Bouilly [5][6], Haït [7], Fraichard [8] and others have not only expanded the theoretical approaches to planning with uncertainty, but have also addressed some of

[1] optimal given the approximations made and the solution space selected

its practical limitations. There is, however, little work aimed at creating paths that are optimal with respect to more general objective functions. Although the planner proposed by Bouilly [5] calculates an optimal path with respect to uncertainty or path length, the approach is not applicable to finding optimal paths with respect to other important criteria such as mobility cost, risk, or energy expended. Gonzalez and Stentz [13] proposed a planner that considers uncertainty in position while optimizing a non-binary objective function. Their approach, however, does not use landmarks, does not deal with aliasing, and is limited to results in simulation. While some of the approaches mentioned above use landmarks as part of the planning process, none addresses the possibility of aliasing in the landmarks.

In the field of Partially Observable Markov Decision Processes (POMDPs), the problem of planning with uncertainty has been frequently addressed. However, most algorithms become computationally intractable when dealing with worlds that have a large number of states. Only Roy and Thrun [9] have solved the problem of finding optimal paths for large, continuous-cost worlds in the presence of uncertainty. The planner they propose includes some of the elements of the planner proposed here, but it is based on an approximate solution to a POMDP. Their approach requires pre-processing of all the states in the search space, in order to later allow for very fast planning. However, the total planning time (including the pre-processing stage) can take from several minutes to a few hours [10]. The planner optimizes uncertainty rather than expected cost and does not deal with aliasing of features.

The research presented here extends the work of Gonzalez and Stentz [13] by adding landmarks and providing tools to deal with ambiguity and aliasing. We also present some experimental results that show the feasibility of the approach.

III. PROBLEM STATEMENT

The problem we are trying to solve is navigating autonomously between two points without GPS and using a high-resolution prior map.

We assume an accurate, high-resolution map that allows the identification of landmarks and the approximate estimation of terrain types by automatic or manual methods. The high-resolution map is translated into a cost grid in which the value of each cell corresponds to the cost of traveling from the center of the cell to its nearest edge. Non-traversable areas are assigned infinite cost and considered obstacles.

Additionally we use the following simplifying assumptions:

- Low initial uncertainty (smaller than the sensor range of the vehicle)
- Low uncertainty rate (less than 10% of distance traveled)

- Availability of basic landmarks that can be reliably detected.

These assumptions are easily met in a mobile robot that has wheel encoders, a fiber-optic gyro for dead-reckoning, and a good initial estimate of the position of the vehicle.

The resulting path should minimize the expected value of the objective function along the path, while ensuring that the uncertainty in the position of the robot does not compromise its safety or the reachability of the goal.

IV. PROPOSED APPROACH

Without simplifying assumptions, the solution of the problem described above would require solving a Partially Observable Markov Decision Process (POMDP), since we are combining planning, uncertainty, and sensing.

However, POMDPs are intractable for most large problems, and although efficient approximations exist, they are not as efficient as deterministic search, especially A*.

We present an alternative approach that takes advantage of the low drift rate in the inertial navigation system of many outdoor mobile robots. The planner uses a Gaussian distribution to model position uncertainty and uses deterministic search to efficiently find resolution-optimal paths that minimize expected cost while considering uncertainty in position.

A. Perception

The vehicle is assumed to have a range sensor with a maximum detection range R and 360° field of view. The vehicle is assumed to be able to reliably detect the landmarks in the map and to be able to detect obstacles not present in the prior map. Vandapel *et al* [14] have shown that many natural and man-made structures can be reliably detected in laser data.

The results presented in this paper use electric poles as landmarks since they are widely available in our test location and can be reliably detected at distances of up to 10 meters. With little modification the approach could be modified to detect tree trunks and other similar features.

B. Motion Model and Uncertainty Propagation

The first-order motion model for a point-sized robot moving in two dimensions is:

$$\dot{x}(t) = v(t)\cos\theta(t)$$
$$\dot{y}(t) = v(t)\sin\theta(t) \qquad (1)$$
$$\dot{\theta}(t) = \omega(t)$$

where the state of the robot is represented by x(t), y(t) and θ(t) (x-position, y-position and heading respectively), and the inputs to the model are represented by v(t) and ω(t) (longitudinal speed and rate of change for the heading respectively). Equation (1) can also be expressed as:

$$\dot{\mathbf{q}}(t) = f(\mathbf{q}(t), \mathbf{u}(t)) \qquad (2)$$

where $\mathbf{q}(t) = (x(t), y(t), \theta(t))$ and $\mathbf{u}(t) = (v(t), \omega(t))$.

A typical sensor configuration for a mobile robot is to have an odometry sensor and an onboard gyro. We can model the errors in the odometry and the gyro as errors in the inputs where $w_v(t)$ is the error in $v(t)$ (error due to the longitudinal speed control), and $w_\omega(t)$ is the error in $\omega(t)$ (error due to the gyro random walk).

Incorporating these error terms into (1) yields:

$$\dot{x}(t) = (v(t) + w_v(t))\cos\theta(t)$$
$$\dot{y}(t) = (v(t) + w_v(t))\sin\theta(t) \qquad (3)$$
$$\dot{\theta}(t) = \omega(t) + w_\omega(t)$$

or, in discrete-time:

$$x^{k+1} = x^k + (v^k + w_v^{\,k})\cos\theta^k \Delta t$$
$$y^{k+1} = y^k + (v^k + w_v^{\,k})\sin\theta^k \Delta t \qquad (4)$$
$$\theta^{k+1} = \theta^k + (\omega^k + w_\omega^{\,k})\Delta t$$

Using the extended Kalman filter (EKF) analysis for this system, which assumes that the random errors are zero-mean Gaussian distributions, we can model the error propagation as follows:

$$\mathbf{P}^{k+1} = \mathbf{F}^k \cdot \mathbf{P}^k \cdot (\mathbf{F}^k)^T + \mathbf{L}^k \cdot \mathbf{Q}^k \cdot (\mathbf{L}^k)^T \qquad (5)$$

where

$$\mathbf{P}^k = E(\hat{\mathbf{q}}^k(\hat{\mathbf{q}}^k)^T) \quad \mathbf{Q}^k = \frac{1}{\Delta t}\begin{pmatrix} \sigma_{vv} & 0 \\ 0 & \sigma_{\omega\omega} \end{pmatrix} \qquad (6)$$

$$\mathbf{F}_{ij} = \left.\frac{\partial f(q_i(t), u_j(t))}{\partial q_i}\right|_{t=k\Delta t}$$
$$\mathbf{F} = \begin{pmatrix} 1 & 0 & -v^k \sin(\theta^k)\cdot\Delta t \\ 0 & 1 & v^k \cos(\theta^k)\cdot\Delta t \\ 0 & 0 & 1 \end{pmatrix} \qquad (7)$$

$$\mathbf{L}_{ij} = \left.\frac{\partial f(q_i(t), u_j(t))}{\partial u_j}\right|_{t=k\Delta t}$$
$$\mathbf{L} = \begin{pmatrix} \cos(\theta^k)\cdot\Delta t & 0 \\ \sin(\theta^k)\cdot\Delta t & 0 \\ 0 & \Delta t \end{pmatrix} \qquad (8)$$

C. Prior Map

A prior map is necessary to provide estimates of the cost to traverse different areas and to provide landmarks for navigation.

1) Cost Map

The cost map is the representation of the environment that the planner uses. It is represented as a grid in which the cost

Fig 1. Cost map: lighter regions represent lower cost, and darker regions represent higher cost. Green areas are manually labeled buildings.

of each cell corresponds to the cost of traveling from the center of the cell to its nearest edge. Non-traversable areas are assigned infinite cost and considered obstacles.

The resulting path calculated by the planner minimizes the expected value of the cost along the path, while ensuring that the uncertainty in the position of the robot does not compromise its safety or the reachability of the goal.

The procedure to create a cost map from a prior map depends on the type of prior map used. If elevation maps are available, cost is usually calculated from the slope of the terrain. When only aerial maps are available, machine learning techniques such as those in [15] can be used. Fig 1 shows the cost map for the test area used in the experimental results presented here. The cost map was created by training a Bayes classifier and adding manual annotations to the resulting map. The table below shows the cost assigned to the different terrain types.

TABLE I

COST VALUES FOR DIFFERENT TERRAIN TYPES.

Terrain Type	Cost
Paved Road*	5
Paved Road 2	10
Dirt Road	15
Grass	30
Trees	40
Water	250
Buildings*	255

* Items manually labeled.

2) Landmarks

Landmarks are features in the prior map that can be detected with the on-board sensors. They can come from a separate database of features or can be extracted directly from the prior map if the resolution of the map is high enough. In our approach, we use aerial maps with a resolution of 0.3 meters per cell. At this resolution, many features are clearly visible and can be identified using manual labeling. Automatic extraction is also possible for some types of features, but the state-of-the-art for automatic feature detection does not yet allow for reliable extraction of most features. Fig 2 shows a small section of our test area with some electric poles labeled as landmarks.

3) Map Registration

A prior map that is not correctly registered to the position of the vehicle is of little use for most planning approaches. The error in map registration usually comes from two main sources: error in the estimation of the position of the vehicle, and error in the estimation of the position of the map. The approach presented here uses the prior map as the reference for all planning and execution. Since the planner considers uncertainty in position, the error in map registration can be modeled as being part of the error in the position of the robot. This makes use of the information in the map in a way that considers the total uncertainty in the position of the robot.

D. Unique Detection Regions and Aliasing

The planning approach proposed requires the presence of reliable landmarks. Reliable landmarks are features in the map that can be reliably detected in both aerial images and the onboard perception system. In order to improve the reliability in the detection of the landmarks we chose very simple yet easy to detect vertical features such as electric poles and trees. Detection of electric poles and trees can be reliably achieved with existing approaches such as the one presented in [14] .

The challenge with simple features such as electric poles and trees is that they cannot be uniquely identified in a reliable manner. It is easy to find an electric pole in a high-resolution aerial image and in a range image from the on-board perception of the vehicle, but it is very hard to uniquely identify which electric pole or tree we are looking at.

However, if we know that the position of the robot is within a certain error distribution, the number of features that are visible within a given detection range are significantly fewer. And if we choose features and positions well, we can often make sure that there is only one feature within the detection range of the robot, in which case the feature detected becomes a unique feature.

The central idea is to identify those areas in which a given feature can be uniquely identified. We call these

Fig 2.Detail of test area showing features of interest

regions *unique detection* regions. Assuming flat terrain, 360^{o} field of view, and a detection range R, the detection region for a point feature i such as an electric pole is a circle of radius R. If the robot is located within this region we can guarantee that only feature i can be detected.

If there is no overlap between detection regions, each circle would be a *unique detection* region. However, if there are other features within a *2R* radius of the feature, the other features would reduce the unique detection region of the original feature. These overlapping regions would be the unique detection regions for groups of two or more features. However, multiple features are harder to detect than individual ones because of occlusions and visibility constraints. For this reason, the current approach only uses *unique detection* regions generated by single features.

Fig 3 shows the same area as in Fig 2, with the unique detection regions highlighted for a detection range R=10
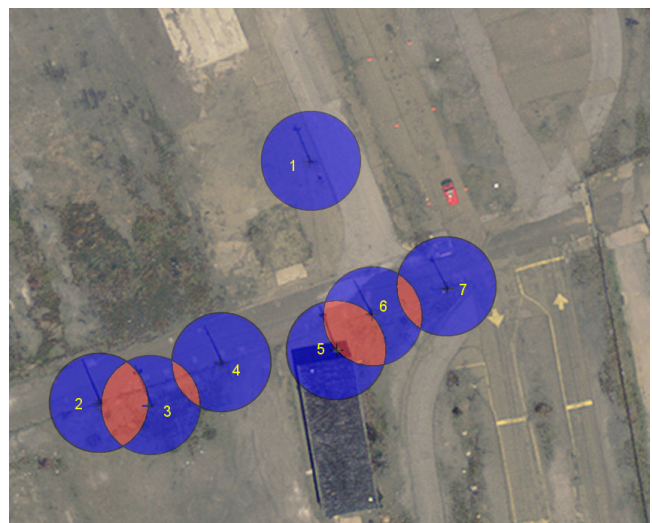

Fig 3.Unique detection regions

meters. The dark(blue) shading shown in feature number 1 indicates *unique detection* regions. In this case, since there are no other features in a *2R* radius the whole circular detection region is a *unique detection* region. The light shading(red) shown in parts of the detection regions of all the other features indicates a part of the detection region where more than one feature can be detected, therefore excluding that area from the *unique detection* region.

E. State Space Representation

Gonzalez and Stentz [13] showed that an isometric Gaussian is an appropriate upper bound in the uncertainty propagation for planning horizons up to a few kilometers. Based on their results we model the probability density function (pdf) of the error as a Gaussian distribution, centered at the most likely location of the robot at step *k*:

$$\mathbf{q} = (x, y)$$
$$\mathbf{q} : N(\mathbf{q}_c{}^k, \sigma^k) \tag{9}$$

where $\mathbf{q}_c^k$ is the most likely location of the robot at step *k*, and $\sigma^k = \sigma_x^k = \sigma_y^k$ is the standard deviation of the distribution at step *k*.

Let us define:

$$\varepsilon^k = 2 \cdot \sigma^k \tag{10}$$

We can then model the boundary of the uncertainty region as a disk centered at $\mathbf{q}_c^k$ with a radius ε^k. This model is a conservative estimate of the true error propagation model, and, depending on the type of error that is dominant in the system, it can provide an accurate approximation of the true model.

F. State Propagation

In order to use a deterministic planner to plan we need to define the transition cost between adjacent cells. In our 3-D configuration space, we are interested in calculating the cost of moving between the configuration $\mathbf{r}^k$ (at path step k) and an adjacent configuration $\mathbf{r}^{k+1}$ (at path step *k+1*). This is equivalent to calculating the expected cost of going from a most likely workspace location $\mathbf{q}_c{}^k$, with uncertainty ε^k to an adjacent most likely workspace location $\mathbf{q}_c{}^{k+1}$ with uncertainty ε^{k+1}:

$$C(\mathbf{r}^k, \mathbf{r}^{k+1}) = E\left[C\left((\mathbf{q}_c{}^k, \varepsilon^k), (\mathbf{q}_c{}^{k+1}, \varepsilon^{k+1}) \right) \right] \tag{11}$$

Equivalently,

$$E\left[C\left((\mathbf{q}_c{}^k, \varepsilon^k), (\mathbf{q}_c{}^{k+1}, \varepsilon^{k+1}) \right) \right] = \tag{12}$$
$$\sum_{\forall i} \sum_{\forall j} C_o\left(\mathbf{q}_i{}^k, \mathbf{q}_j{}^{k+1} \right) \cdot p\left(\mathbf{q}_i{}^k, \mathbf{q}_j{}^{k+1} \mid \mathbf{q}_c{}^k, \varepsilon_k, \mathbf{q}_c{}^{k+1}, \varepsilon^{k+1} \right)$$

where $\mathbf{q}_i{}^k$ is each of the *i* possible states at path step *k*, $\mathbf{q}_j{}^{k+1}$ is each of the *j* possible states at path step *k+1*, and
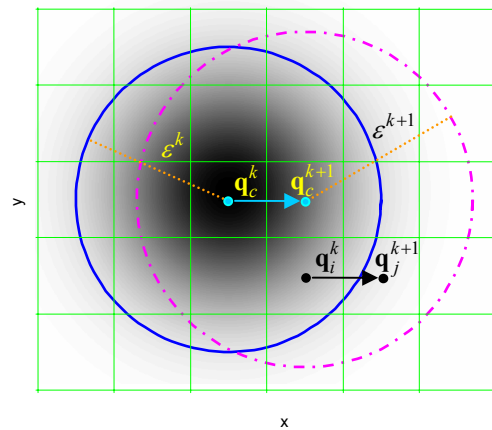

Fig 4. $p\left(\mathbf{q}_i^k \mid \mathbf{q}_c^k, \varepsilon_k \right)$ and the state transitions from $\mathbf{r}^k = (\mathbf{q}_c{}^k, \varepsilon^k)$ to $\mathbf{r}^{k+1} = (\mathbf{q}_c{}^{k+1}, \varepsilon^{k+1})$

$C_o(\mathbf{q}_i{}^k, \mathbf{q}_j{}^{k+1})$ is the deterministic cost of traveling from $\mathbf{q}_i{}^k$ to $\mathbf{q}_j{}^{k+1}$ (see Fig 4)[2].

Since we are assuming a low uncertainty rate ($\alpha_u < 0.1$), we can make additional simplifications that transform (12) into:

$$E\left[C\left((\mathbf{q}_c{}^k, \varepsilon^k), (\mathbf{q}_c{}^{k+1}, \varepsilon^{k+1}) \right) \right] = \tag{13}$$
$$a \cdot C_E(\mathbf{q}_c{}^k, \varepsilon_k) + b C_E(\mathbf{q}_c{}^{k+1}, \varepsilon^k)$$

where *a* and *b* are constants determined by the relative position of $\mathbf{q}_c{}^k$ and $\mathbf{q}_c{}^{k+1}$, and

$$C_E(\mathbf{q}_c{}^k, \varepsilon^k) = \sum_{\forall i} C_o\left(\mathbf{q}_i{}^k \right) \cdot p\left(\mathbf{q}_i{}^k \mid \mathbf{q}_c{}^k, \varepsilon^k \right) \tag{14}$$

is the expected cost of traversing cell $\mathbf{q}_c^k$ if the uncertainty at this location is ε^k. Therefore,

$$C(\mathbf{r}^k, \mathbf{r}^{k+1}) = a \cdot C_E(\mathbf{q}_c{}^k, \varepsilon_k) + b C_E(\mathbf{q}_c{}^{k+1}, \varepsilon^k). \tag{15}$$

The planner used for planning with uncertainty is a modified version of A* in 3-D in which the successors of each state are calculated only in a 2-D plane, and state dominance is used to prune unnecessary states [13].

1) Outside of Unique Detection Regions

Since the dominant term in the error propagation for our planning horizon is linear with distance traveled, equation (5) can be simplified to the following model to propagate uncertainty:

$$\varepsilon(\mathbf{q}_c{}^k) = \varepsilon(\mathbf{q}_c{}^{k-1}) + \alpha_u d(\mathbf{q}_c{}^{k-1}, \mathbf{q}_c{}^k) \tag{16}$$

where α_u is the uncertainty accrued per unit of distance traveled, $\mathbf{q}_c^{k-1}$ is the previous position along the path, and $d(\mathbf{q}_c^{k-1}, \mathbf{q}_c^k)$ is the distance between the two adjacent path

---

[2] As mentioned previously, we assume a discretized grid of states corresponding to a known map.

locations $\mathbf{q}_c^{k-1}$ and $\mathbf{q}_c^k$. The uncertainty rate α_u is typically between 0.01 and 0.1 (1% to 10%) of distance traveled.

By modeling the error propagation in this manner, we are assuming that the dominant term is the uncertainty in the initial angle. Even though we are not explicitly modeling θ as a state variable, the effects of uncertainty in this variable are accounted for in the uncertainty propagation model for $\mathbf{q}=(x,y)$.

State propagation outside *unique detection* regions uses (15) for cost propagation, and (16) to update the uncertainty at each step. No sensing of landmarks takes place when planning outside of *unique detection* regions, which allows the planner to use deterministic search in the state expansion.

2) Inside Unique Detection Regions

If all the possible locations for a configuration $\mathbf{r}^k$ are inside a *unique detection* region, we can guarantee that the feature that created the region can be detected and that no other features will be visible within the field of view of the robot.

For practical purposes we make the simplifying assumption that a circle with radius $\varepsilon^k = 2 \cdot \sigma^k$ completely contains all possible locations on (x,y) of a given state $\mathbf{r}^k$. Therefore, if a circle of radius ε^k centered at q_c^k is completely contained within a u*nique detection* region $i,$ we can guarantee that feature i will be detected. This approximation allows us to model the detection of landmarks in a deterministic way, therefore allowing the use of deterministic search for this part of the state propagation as well. This assumption is only valid if we can *reliably* detect landmark i.

G. Discussion

The approach proposed here finds the path that has the lowest expected cost and guarantees the reachability of the goal within the given error bounds. The solution is resolution-optimal as long as the landmarks can be reliably detected.

However, in some scenarios, the best approach would be to have a policy instead of a path. A policy would consider the detection of features as a non-deterministic event and would produce a set of actions to be performed depending on the outcome of the detection of features. Because of this added flexibility a policy could have a lower expected cost than the path found by our approach. But finding an optimal policy would require solving a POMDP, which would be intractable to solve or would take significantly longer to plan even with an approximate solution as in [9]

V. RESULTS

A. Simulation Results

Fig 5 shows a sample cost map with some landmarks. Shades of gray indicate different costs in the cost map: areas with lighter color have lower cost, and areas with darker color have higher cost. Solid green areas are obstacles. The start location is a small square on the left, and the goal is a small circle on the right. As a reference, this figure also shows the shortest path that guarantees reachability of the goal for this cost map. The uncertainty rate is 10% of distance traveled. The yellow circles indicate the $\varepsilon = 2\sigma$ contours of the error distribution at each step along the path. The expected cost for this path is 8586.9, and the uncertainty at the goal is $\varepsilon_f = 2\sigma_f = 3.8$ m.

The following figures show the paths found by our approach under different constraints for uncertainty at the goal. They also illustrate the advantages of minimizing the expected cost of the path instead of minimizing the path length or the uncertainty of the path.

Fig 6 shows the lowest expected cost path with an uncertainty rate $k_u=10\%$ if the maximum uncertainty allowed at the goal is 12 m. The path found has an expected cost of 1485 (82% lower than the shortest path), and the uncertainty at the goal is 11 m. Because the uncertainty allowed at the goal is large, the planner has enough freedom to look for a low cost path, even if a low cost path is longer and has higher uncertainty. Only one of the localization regions can provide an improvement in the total cost, therefore the planner only includes that landmark in the final path. The
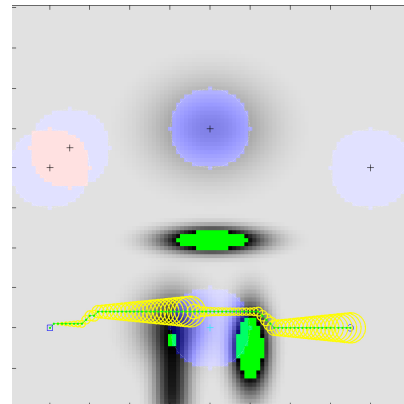

Fig 5. Planning with uncertainty rate $k_u=10\%$ and using landmarks for localization (shortest path).
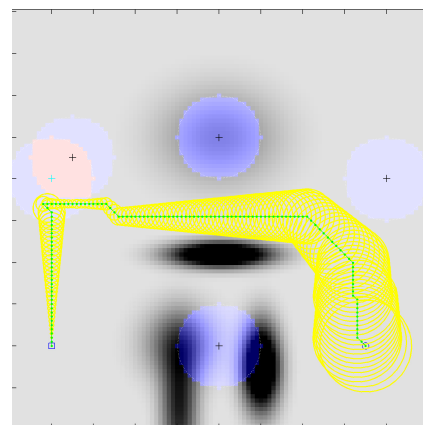

Fig 6. Planning with uncertainty rate $k_u=10\%$ and maximum uncertainty at the goal of 12 m.
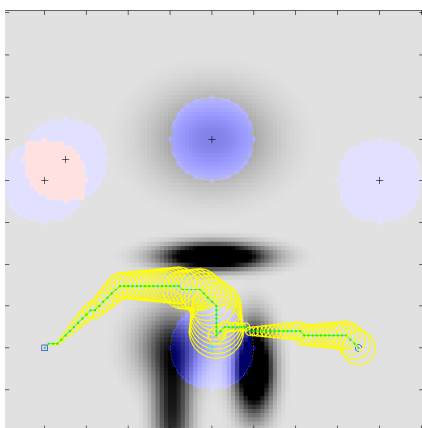
Fig 7. Planning with uncertainty rate k_u=10% and maximum uncertainty at the goal of 3.8 m

planner also avoids the aliased region between the two landmarks on the left, and localizes only with the leftmost landmark.

If the maximum uncertainty allowed at the goal is small, the planner trades off lower cost solutions in order to satisfy the uncertainty constraint. Fig 7 shows the lowest expected cost path when the maximum uncertainty allowed at the goal is reduced to 3.8 m. Even with a maximum uncertainty at the goal equal to that of the shortest path there can be significant advantages in minimizing the expected cost instead of the uncertainty or the path length. The expected cost is now 4710 (still 47% lower than the shortest path), and the final uncertainty is ε=3.8 m. Although the last segment of the path is the shortest path for that segment, the first segment is able to look for a less expensive path than the shortest path and the resulting path is significantly less costly than the shortest path.
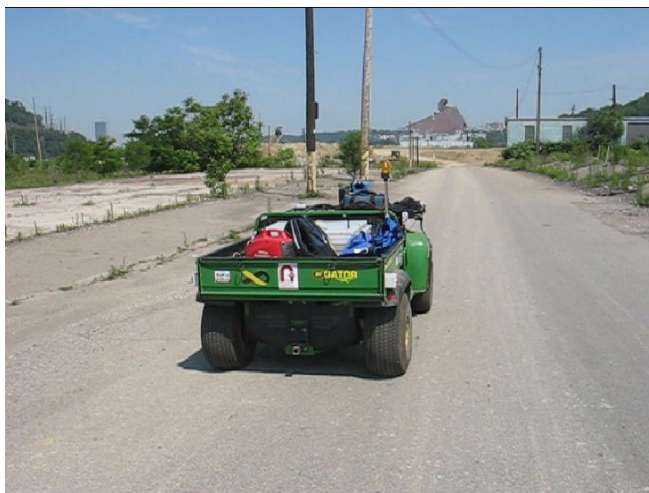

Fig 8. E-gator autonomous vehicle used for testing and electric poles used for localization at test site. The vehicle equipped with wheel encoders and a KVH E- core 1000 fiber-optic gyro for dead reckoning, and a tilting SICK ladar and onboard computing for navigation and obstacle detection

B. Field Tests

In order to validate the results experimentally the following field test was carried out on the e-gator autonomous vehicle shown in Fig 8: a path was planned between a location S and a location G, assuming initial uncertainty σ=2.5m, uncertainty rate of 5% of distance traveled and maximum uncertainty of 10 m, using electric poles for localization (Fig 9). Notice how the path follows a road in order to minimize the expected cost along the path (instead of just minimizing the length of the path). The path also visits detection regions as needed to maintain a low cost

Fig 9. Path planned assuming initial uncertainty σ=2.5m, uncertainty rate of 5% of distance traveled and maximum uncertainty of 10 m. The expected cost of the path is 3232, and the final uncertainty is σ=2.7m. Top: aerial image and unique detection regions. Bottom: cost map used.
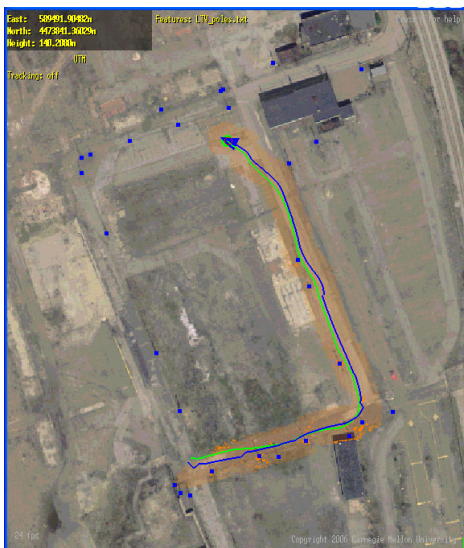
Fig 10. Path planned and executed without GPS. Blue dots show the location of landmarks. The blue line is the position estimate of the Kalman filter on the robot and the green line is the position reported by a WAAS differential GPS with accuracy of approximately 2 meters (for reference only).

path and avoids narrow areas that could not be safely avoided if the position of the robot is not accurately known.

Also notice that the final uncertainty of the path is relatively high ($\varepsilon = 5.4$ m). This is because the maximum uncertainty allowed was set to 10 meters, and the planner only tries to reduce the uncertainty if the reduction in uncertainty will reduce the expected cost of the path. In the last segment the path is going through a large paved area and there is no increase in cost because of the higher uncertainty. For this reason, the planner does not try to detect any features in last 50 meters of the path.

Fig 10 shows the path executed by the robot. The blue line is the position estimate of the robot according to the onboard Kalman filter that combines the dead reckoning sensors and the landmark detections (no GPS). The green line shows the position estimate according to the GPS (shown as a reference only). Notice how the blue line stays very close to the GPS estimate and jumps in a few places after detecting a landmark.

VI. CONCLUSIONS AND FUTURE WORK

We have introduced a novel approach for mobile robot navigation that allows robust and efficient navigation without GPS. The approach uses landmarks in the environment that have been manually identified in a high-resolution prior map to reduce the uncertainty in the robot's position as part of the planning process. The resulting path minimizes the expected cost along the route considering the uncertainty in the position of the robot. We have also shown

experimental results of the system, showing navigation capabilities similar to those of a robot equipped with GPS.

The current version of the algorithm uses light poles as its landmarks, and assumes that the landmarks will always be detected in both planning and execution. Future work includes using other common features for localization such as trees, buildings and roads. This will require a version of the algorithm that allows for more complex representations of the error propagation model. Relaxing the assumption that landmarks will always be detected in planning and execution is another area of future work that we will explore.

REFERENCES

[1] B. Bonet and H. Geffner, "Planning with incomplete information as heuristic search in belief space," in Proceedings of the 6th International Conference on Artificial Intelligence in Planning Systems (AIPS), pp. 52-61, AAAI Press, 2000

[2] J.C. Latombe, A. Lazanas, and S. Shekhar, "Robot Motion Planning with Uncertainty in Control and Sensing," Artificial Intelligence J., 52(1), 1991, pp. 1-47.

[3] J.C. Latombe, *Robot Motion Planning*. Kluwer Academic Publishers. 1990

[4] A. Lazanas, and J.C. Latombe, "Landmark-based robot navigation". In Proc. 10th National Conf. on Artificial Intelligence (AAAI-92), 816--822. Cambridge, MA: AAAI Press/The MIT Press. http://citeseer.nj.nec.com/lazanas92landmarkbased.html

[5] B. Bouilly. "Planification de Strategies de Deplacement Robuste pour Robot Mobile". PhD thesis, Insitut National Polytechnique, Tolouse, France, 1997

[6] B. Bouilly, T. Siméon, and R. Alami. "A numerical technique for planning motion strategies of a mobile robot in presence of uncertainty". In Proc. of the IEEE Int. Conf. on Robotics and Automation, volume 2, pages 1327--1332, Nagoya (JP), May 1995. http://citeseer.nj.nec.com/bouilly95numerical.html

[7] A. Haït, T. Simeon, and M. Taïx, "Robust motion planning for rough terrain navigation".Published in IEEE Int. Conf. on Intelligent Robots and Systems Kyongju, Korea, 1999. http://citeseer.ist.psu.edu/319787.html

[8] Th. Fraichard and R. Mermond "Integrating Uncertainty And Landmarks In Path Planning For Car-Like Robots" Proc. IFAC Symp. on Intelligent Autonomous Vehicles March 25-27, 1998. http://citeseer.nj.nec.com/fraichard98integrating.html

[9] N. Roy, and S. Thrun, "Coastal navigation with mobile robots". In Advances in Neural Processing Systems 12, volume 12, pages 1043—1049, 1999.

[10] N. Roy, Department of Aeronautics and Astronautics, MIT. Private Conversation. September 1, 2004

[11] K. Goldberg, M. Mason, and A. Requicha. "Geometric uncertainty in motion planning". Summary report and bibliography. IRIS TR. USC Los Angeles, 1992.

[12] A. Kelly, "Linearized Error Propagation in Odometry", The International Journal of Robotics Research. February 2004, vol. 23, no. 2, pp.179-218(40).

[13] J.P. Gonzalez and A. Stentz, "Planning with Uncertainty in Position: An Optimal and Efficient Planner," Proceedings of the IEEE International Conference on Intelligent Robots and Systems (IROS '05), August, 2005.

[14] N. Vandapel, D. Huber, A. Kapuria, and M. Hebert, "Natural Terrain Classification using 3-D Ladar Data," IEEE International Conference on Robotics and Automation, April, 2004.

[15] B. Sofman, E. Lin, J. Bagnell, N. Vandapel, and A. Stentz, "Improving Robot Navigation Through Self-Supervised Online Learning," Proceedings of Robotics: Science and Systems, August, 2006.