Motion Planning Under Uncertainty: Application to an Unmanned Helicopter

Joshua D. Davis* and Suman Chakravorty†
Texas A&M University, College Station, Texas 77843

A methodology is presented in this work for intelligent motion planning in an uncertain environment using a nonlocal sensor, such as a radar. This methodology is applied to an unmanned helicopter navigating a cluttered urban environment. It is shown that the problem of motion planning in an uncertain environment, under certain assumptions, can be posed as the adaptive optimal control of an uncertain Markov decision process, characterized by a known, control dependent system, and an unknown, control independent environment. The strategy for motion planning then reduces to computing the control policy based on the current estimate of the environment, also known as the "certainty equivalence principle" in the adaptive control literature. The methodology allows the inclusion of a nonlocal sensor into the problem formulation, which significantly accelerates the convergence of the estimation and planning algorithms. Further, the motion planning and estimation problems possess special structure which can be exploited to significantly reduce the computational burden of the associated algorithms. The methodology is applied to the problem of motion planning for an unmanned helicopter through a partially known model of the Texas A&M campus and the testing is done on a flight simulator.

I. Introduction

IN THIS paper, a methodology for the motion planning of an autonomous agent in an uncertain environment is proposed and applied to an unmanned helicopter navigating a cluttered urban environment. The optimal path for the unmanned helicopter is planned using a priori knowledge about the environment and the sensor data received as the helicopter navigates the obstacles in the environment. The motion planner involves a high-level planner which plans against the uncertainty in the environment and issues its commands as a series of waypoints for a lower-level controller to track. A commercially available lower-level controller from Rotomotion LLC, called the automated flight control system (AFCS), was interfaced with the high-level motion planner so that the motion planning algorithm could be implemented on a 6 degrees of freedom flight simulator.

The high-level controller plans the motion of the agent in an uncertain environment using a radar sensor. However, the methodology is applicable for any nonlocal sensor. The state space of any motion planning problem can be expressed as the ordered pair $(s, q(s))$ where s represents the system state and $q(s)$ represents the state of the environment at the state s. For example, in the case of an unmanned helicopter exploring an urban environment, s corresponds to the (x, y, z) coordinates of the helicopter and $q(s)$ corresponds to the presence or absence of an obstacle at the point (x, y, z). The goal of the motion planning strategy is to use all available information about the environment, until the current time instant, to plan the "best possible" path. It is known that the planning problem can be modeled as a Markov decision process (MDP), characterized by a known, control dependent exploration system and unknown, uncontrollable environment [1,2]. Our formulation allows the integration of a radar or a similar nonlocal sensor into the planning methodology. The planning and estimation problems, as formulated in this paper, have special structure which can be exploited to significantly reduce the dimensionality of the associated algorithms.

There has been substantial research in the adaptive control of controlled Markov chains, or Markov decision processes, in the past two decades. In indirect adaptive control, the transition probabilities of the underlying Markov chain are estimated and the control is applied based on the most recent estimate of the transition probabilities [3–5]. This is known as the so-called "certainty equivalence principle." The "direct" approach to stochastic adaptive control falls under the category of "reinforcement learning" methodologies wherein the optimal control is calculated directly without resorting to estimating the transition probabilities of the underlying Markov chain [6–8]. Underlying all these methods is Bellman's "principle of optimality" or dynamic programming, which is a methodology for sequential decision making under uncertainty [9,10]. In this work, it is shown that the motion planning problem can be reduced to the adaptive optimal control of a Markov decision process and thus the above methodologies can be applied to the same. The indirect approach to adaptive control is adopted because mapping the environment is of interest too. As additional states are added to a dynamic programming problem there is a geometric growth in computation, which is considerably more attractive than a direct enumeration method for control determination, which would give an exponential growth in computation [11]. The reason that direct enumeration requires such a large amount of computation is that it determines all possible control sequences and compares them to determine an optimal path. Even though dynamic programming is more efficient computationally than direct enumeration, the greatest challenge associated with dynamic programming is still "the curse of dimensionality." This means that as the number of dimensions or states increase, the computational requirements rapidly increase and can cause the solution to become computationally unfeasible.

Motion planning considers how the state space (configuration space) is represented. The dimension of the configuration space depends upon the degree of freedom (DOF) of the robot being used in the motion planning [12]. Another consideration is if the system is modeled as continuous or discrete. If the system is modeled as discrete, the grid size (number of states) will greatly effect the computation time.

Various approaches have been developed for collision-free motion planning of unmanned systems in known environments [13]. In the past decade, there has been an increasing interest in the case when the environment in which the system is operating is partially or

completely unknown. The uncertainty in the environment is treated as a deterministic worst case [14,15] or on a probabilistic average case basis [16]. In "probabilistic robotics," there has been substantial research in the localization of a mobile robot while simultaneously mapping the environment [17–21]. In recent years, there has been a blending of previously different areas of study: planning, control theory, and artificial intelligence for motion planning. In the past, planning has focused more on planning the trajectory of a vehicle, while control theory has focused on the response of differential equations to control inputs [22]. The area of artificial intelligence in the past tended to focus on problem solving in a discrete state space [23,24]. The development of algorithms for autonomous systems to navigate through obstacle fields has caused the differences between planning algorithms and control theory to become less distinguished.

Many of the motion planning techniques that have been implemented are essentially graph search methods, such as breadth first search, depth first search, and Dijkstra's algorithm. Breadth first search considers all possible paths that are equal distance from the starting point. Although it is not the most computationally efficient it guarantees the optimal path is determined [25]. Depth first search considers the cost from a starting point to the goal state along a feasible path and may later consider alternate paths [26]. Dijkstra's algorithm [27] is a single source shortest path algorithm and was developed as a special form of dynamic programming. Dijkstra's algorithm includes several heuristics to draw the vehicle toward the goal state and help eliminate unnecessary computations by removing unnecessary state analysis.

One of the most well-known graph search methods for motion planning is A* and was created as a method for determining an optimal path or trajectory by including heuristics. It determines an admissible, "optimistic," solution. A* is effective for a priori planning in a static environment but requires complete recalculation if the environment changes [28]. A* and Dijkstra's algorithm are essentially the same except for the function which is used to sort the vector recording the cost of feasible paths [23]. A development in A* research led to a class of algorithms called focused dynamic A* (D*) [29], which includes heuristics as well as incremental search techniques so that a complete recalculation of the costs is not necessary for a dynamically changing environment. D* has been implemented for various motion planning problems, such as indoor robots, outdoor unmanned ground vehicles in the Defense Advanced Research Projects Agency (DARPA) Unmanned Ground Vehicle Program [30], urban robots, and even the Mars rover [31]. A newer version of D* called D* Lite has been developed that is at least as efficient and often more efficient than D*. D* Lite is also more intuitive to understand than D* and has been rigorously analyzed mathematically [32].

Logic based methods also exist as a possible solution to motion planning and can be implemented similar to graph search methods, but take a somewhat different approach to motion planning [33,34]. Logic methods often focus on partial plans and subgoals. There are various implementations to logic based planning, for example, planning graphs can be analyzed by layer construction, or the planning problem can be tackled using a Boolean approach.

Probabilistic road maps (PRM) are a recent development in motion planning and are an efficient method for determining an optimal path. They are essentially a sample-based approach to navigation of an environment [35]. The road map is a graph of randomly generated collision-free paths, which are connected by some simple fast planning method [36]. The advantage of using a road map is the efficiency with which the nodes (waypoints) are connected in the configuration space. If part of the road map is not used then those computations are wasted. However there are probabilistic road map variants, which minimize unnecessary computations. The samples can be chosen randomly for the state space (configuration space) to obtain meaningful information in developing a model of the environment. In the study of road maps important information includes the denseness of the samples chosen as well as the method of choosing the samples. Road maps require preprocessing before the vehicle begins navigation and essentially act as a network of multiple pairs of initial-goal points. In the

implementation of the algorithm the road map construction phase is used to generate the nodes and connect them, while the query phase is used to evaluate which path is optimal. The "probabilistic" part of the name comes from the fact that the method performs sampling in a probabilistic fashion. However, probabilistic road maps do not perform control in a probabilistic fashion.

The original contributions of the current work are as follows. We propose a hierarchical motion planner wherein, the problem of "intelligent motion planning" for the high-level planner is reduced to the adaptive optimal control of an uncertain Markov decision process, characterized by a known, control dependent system and an unknown, control independent environment; and a low-level controller is then used to track the commands from the high-level planner. It is noted here that the methods of implementing the higher level planner are different for different methods, for instance, see [37]. However, the underlying philosophy of these methods is the same: the higher level planner plans at the global environmental level based on some coarse uncertain knowledge of the environment, while the local planner implements the global plans making sure that it avoids the obstacles and constraints arising due to the incomplete knowledge of the environment. The methodology allows for the inclusion of nonlocal sensors, which significantly reduce the computational burden of the estimation and planning algorithm in an uncertain environment. The motion planning methodology is applied to the problem of an unmanned aerial vehicle (UAV) helicopter navigating the Texas A&M campus and is tested on a 6 degrees of freedom flight simulator of the helicopter.

The rest of the paper is organized as follows. Section II details the structure of the motion planner. Section III contains the formulation of the high-level planner as a Markov decision problem. In Sec. IV, we present the results of implementing the planning methodology on a UAV helicopter navigating the Texas A&M campus by testing the algorithms on a 6 DOF flight simulator.

II. Controller Structure

The motion planning algorithm is implemented in an hierarchical fashion wherein a high-level motion planner determines the optimal waypoints for the low-level controller to track during its traverse to the goal state. The planner takes into consideration the shortest distance to the goal state, local and nonlocal environment sensing, as well as the current state of the helicopter when it makes decisions. The high-level planner minimizes the distance traveled by choosing waypoints between the initial and goal states while avoiding obstacles during the flight. The high-level controller determines the sequence of waypoints and issues it to the low-level controller. The low-level controller receives the waypoint command and issues the necessary commands to the flight surfaces. Once the waypoint is achieved the high-level planner receives new sensor information about the environment and determines the next waypoint. This process continues until the goal state is achieved. Figure 1 outlines the basic motion planning architecture.

The low-level flight controller used in the simulations was the automated flight control system, which was developed by Rotomotion, LLC, and is capable of executing waypoint commands given to the helicopter. The AFCS has a series of proportional–integral–derivative (PID) control loops that stabilize the attitude, position, and velocity of a remote control helicopter. The control loops use an observer (state estimator), which is implemented using a Kalman filter. As the helicopter flies, the Kalman filter also gives state feedback to the high-level motion planner so that the high-level planner knows that a waypoint has been achieved. The motion planner is thus in the form of a composite feedback control wherein the high-level planner plans against the uncertainty in the environment on a longer time/length scale while the lower-level controller accounts for the uncertainties at a shorter time/length scaling, and also for the uncertainties in the system dynamical model. In the following, we shall use the terms planner and controller interchangeably, as the two terms are synonymous in the context of this paper.
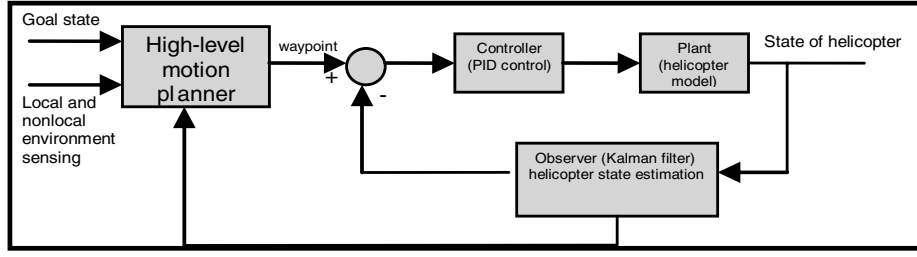
Fig. 1 Control architecture for the high-level motion planner and low-level controller.

III. Planning Under Uncertainty Using Nonlocal Sensing

First, we recount some results that will be required for the motion planning [1,2]. The development in the rest of this section is exclusively for the higher level planner.

A. Preliminaries

Let the state of the exploration system be denoted by s, $s \in S$. Denote the state of the environment at the system state s by $q(s)$. For example, in the case of robotic exploration of unknown terrain, s corresponds to the (x, y) coordinates of the robot and $q(s)$ corresponds to the height of the terrain $z(x, y)$ at the point (x, y). In the case of a UAV navigating enemy territory while avoiding radar detection, the s variable corresponds to the position (x, y, z) of the UAV while $q(s)$ corresponds to the binary valued variable indicating the presence or lack of radar coverage at the point (x, y, z).

We assume that the state set S is discrete, that is, the planning domain has been gridded into a finite number of cells. From here on assume that the system state is sensed perfectly and only the environment is sensed imperfectly (in the context of the previous section, it means that the particular grid that the vehicle is in is known with certainty). Let the number of system states be N and let the number of possible environment states, at any system state s, be D and denote this set by Q. Denote the local state of the exploration system by the ordered pair $(s, q(s))$. Let the set of control actions be denoted by U and let the total number of control actions possible be denoted by M. The control of the high-level planner is discrete and corresponds to the adjacent grid cell that the UAV needs to move to, given its current location cell. Denote any particular control action by u. The following Markovian assumption is made about the system. Let $F^t = \{(s_0, q_0(s_0)), u_0, \ldots, (s_{t-1}, q_{t-1}(s_{t-1})), u_{t-1}\}$ represent the history of the process until time t.

A 3.1: The current system state s_t is dependent only on the system state and control input at the previous time instant, that is,

$$p(s_t/F^t) = p(s_t/s_{t-1}, u_{t-1}) \tag{1}$$

A 3.2: The environment process is "incoherent," that is, the environment process is spatially uncorrelated and temporally stationary. In other words, if $\{q_t(s), s \in S\}$ denotes the environment process, $q_t(s)$ is a stationary process for all $s \in S$. Moreover, $q_t(s)$ is independent of $q_\tau(s')$ whenever $s \neq s'$, for all t, τ. Note that $q_t(s)$ is a random variable and the above assumption is used in a probabilistic sense.

The second assumption above essentially means that the environmental random process (note that the environment is modeled as a random process here) is temporally stationary and spatially uncorrelated/independent. Deterministic environments (like an unstructured terrain) automatically satisfy the above assumptions.

Proposition 1: Under assumptions A 3.1 and A 3.2, the following holds:

$$p((s_t, q_t(s_t))/F^t) = p(s_t/(s_{t-1}, u_{t-1}))p(q_t(s_t)) \tag{2}$$

The above result is a direct consequence of assumptions A3.1 and A3.2, please see Chakravorty and Junkins [1,2] for more details. The transition probabilities $p(s_t/(s_{t-1}, u_{t-1}))$ quantify the control

uncertainties inherent in the system and are assumed to be known beforehand. The environmental uncertainty $p(q(s))$ is unknown and successive estimates are made of this uncertainty as the planning proceeds to completion. Motion planning may be framed as an infinite horizon discounted stochastic optimization problem, that is, given the initial state $(s_0, q_0(s_0))$, the optimal control policy $\mu^*(s_0, q_0(s_0)) = \{u_1, u_2, \ldots\}$ is defined by (see Bertsekas and Tsitsiklis [6] for more details)

$$\mu^*(s_0, q_0(s_0)) = \arg\min_\mu E_\mu$$

$$\times \left(\sum_{t=1}^{\infty} \beta^t c((s_t, q_t(s_t)), (s_{t-1}, q_{t-1}(s_{t-1})), u_{t-1})/(s_0, q_0(s_0)) \right) \tag{3}$$

where $c((s_t, q_t(s_t)), (s_{t-1}, q_{t-1}(s_{t-1})), u_{t-1})$ is a positive predefined cost that the system incurs in making the transition from state $(s_{t-1}, q_{t-1}(s_{t-1}))$ to $(s_t, q_t(s_t))$ under the control action u_{t-1}, $E_\mu(\cdot)$ denotes the expectation operator with respect to the policy μ, and $\beta < 1$ is a given discount factor.

The following environment sensing model is adopted:

At every instant t, the system (UAV) at state s_t can observe the environment state $q_t(s)$, (i.e., the current environment state at state s), if $s_t \in F(s) \subseteq S$, where $F(s)$ is assumed to be known beforehand. The set $F(s)$ constitutes a "footprint" of the sensor system.

Associated with every observation-vantage point pair, $(q(s), s')$, $s' \in F(s)$, [i.e., we are observing the environment at s, $q(s)$, from the state s'], there exists a known measurement error model, $p(\hat{q}(s)/q(s), s')$, $\hat{q}(s)$, $q(s) \in Q$, and s, $s' \in S$, that is, the probability that $\hat{q}(s)$ is observed when the environment is actually at the state $q(s)$, for an observation made from system state s'. Such a model may be deduced from sensor calibrations.

B. Estimation

Consider the following relationship:

$$\pi(\hat{q}(s)) = \frac{1}{\pi(F(s))} \sum_{s' \in F(s), q(s)} p(\hat{q}(s)/q(s), s')p(q(s))\pi(s') \tag{4}$$

where

1) $\pi(\hat{q}(s))$ denotes the probability of observing the noise corrupted environment state $\hat{q}(s)$ during the course of the exploration, that is, the fraction of the time (or periodicity) that the environment at state s is observed to be at $\hat{q}(s)$ during the course of the exploration.

2) $p(q(s))$ denotes the true probability that the environment state is $q(s)$ at the state s.

3) $\pi(s')$ denotes the fraction of the time that the system is at state s' and $\pi(F(s))$ represents the fraction of time that the system spends in the footprint set $F(s)$.

Note that the above equation can be derived heuristically by considering $\pi(\hat{q}(s))$ as a probability and then conditioning it on random variables $q(s)$ and s. The above equation states that the frequency of observing a particular value of the environment during the course of the exploration process is related to the actual probability of the environment taking that value, the noise model, and the frequency of visiting the states of the system. Because the noise model is known, and the values of $\pi(\hat{q}(s))$ and $\pi(s)$ can be estimated during the course of the exploration using the Monte Carlo

method, it is possible to obtain the true environment probabilities using Eq. (4). Mathematically:

$$\pi_t(\hat{q}_n(s)) := \frac{1}{t}\sum_{n=1}^{t} 1(\hat{q}_n(s) = \hat{q}(s)) \tag{5}$$

$$\pi_t(s) := \frac{1}{t}\sum_{n=1}^{t} 1(s_n = s) \tag{6}$$

where $1(A)$ denotes the indicator function of the event A. Then, the true probabilities of the environment process $p(q(s))$ can be obtained recursively as

$$P_t(s) := \arg\min_{P\in\bar{V}} \|\Pi_t(s) - \Gamma_t(s)P\|^2 \tag{7}$$

where

$$P_t(s) = [p_t(q_1(s)),\ldots, p_t(q_D(s))]' \tag{8}$$

$$\Pi_t(s) = [\pi_t(\hat{q}_1(s)),\ldots, \pi_t(\hat{q}_D(s))]' \tag{9}$$

$$\Gamma_t(s) = \lambda(s)[\gamma_t^{ij}(s)] \tag{10}$$

$$\gamma_t^{ij}(s) = \sum_{s'\in F(s)} p(q_i(s)/q_j(s), s')\pi_t(s') \tag{11}$$

$\|\cdot\|$ denotes the Euclidean norm in $\Re^D$, and $\bar{V}$ represents the space of all probability vectors in $\Re^D$. Note that we call the algorithm recursive because the quantities $\Pi_t(s)$ and $\Gamma_t(s)$ are obtained in a recursive fashion. Thus, keeping account of the probabilities, $\pi(s)$ and $\pi(\hat{q}(s))$, the true probabilities of the environment process can be recovered asymptotically. The term $\lambda(s)$ is equal to $1/\pi(F(s))$ and is included to make $\Gamma_t(s)$ stochastic, that is, such that the sum of the elements of every column of the matrix is unity.

C. Control

Consider the stochastic optimal control problem posed in Eq. (3). Using the Bellman principle of optimality [6], it can be shown that the optimal policy is stationary, that is, the optimal control is independent of time, and that the optimal control at the state $(s, q(s))$, $\mu^*(s, q(s))$ is given by the following equation:

$$u^*(s, q(s)) = \arg\min_u \sum_{(r,\bar{q}(r))} p(r/s, u)p(\bar{q}(r))$$
$$\times [c((r, \bar{q}(r), (s, q(s)), u) + \beta J^*(r, \bar{q}(r))] \tag{12}$$

where $J^*(r, \bar{q}(r))$ is the optimal cost to go from the state $(r, \bar{q}(r))$. Moreover, J^* satisfies the following fixed point equation:

$$J^*(s, q(s)) = \min_u \sum_{(r,\bar{q}(r))} p(r/s, u)p(\bar{q}(r))$$
$$\times [c((r, \bar{q}(r), (s, q(s)), u) + \beta J^*(r, \bar{q}(r))] \tag{13}$$

The problem of path planning is one of adaptive control of an uncertain Markov decision process (because the probabilities of the environment process are not known). However, the system state and the local environment are assumed to be sensed perfectly and thus the problem is not a partially observed Markov decision process [11]. In such a scenario, the strategy of adaptive control is to use the policy that is optimal with respect to the current estimate of the system, because it corresponds to the current knowledge of the system that is being controlled and is referred to as the certainty equivalence principle in adaptive control [3].

Let $T = \{t_1, t_2, \ldots, t_k, \ldots\}$ denote the set of all times at which the control policy is updated during the path planning. Let the updated control policy at time instant t_k be denoted by $\mu_k(s, q(s))$. Let $p_t(q(s))$ denote the estimated environmental uncertainty at time t, obtained from the estimation Eq. (7). Then, the control update at time $t_k \in T$, $\mu_k(\cdot)$ using the principle of optimality and the certainty equivalence principle, is given by

$$\mu_k(s, q(s)) = \arg\min_u \sum_{(r,\bar{q}(r))} p(r/s, u)p_{t_k}(\bar{q}(r))$$
$$\times [c((r, \bar{q}(r)), (s, q(s)), u) + \beta J_k(r, \bar{q}(r))] \tag{14}$$

where

$$J_k(s, q(s)) = \min_u \sum_{(r,\bar{q}(r))} p(r/s, u)p_{t_k}(\bar{q}(r))$$
$$\times [c((r, \bar{q}(r)), (s, q(s)), u) + \beta J_k(r, \bar{q}(r))]$$

Next, a few of the salient properties of the control problem posed are listed. The results are given without proof; please see [1,2] for more details about the proofs. The following result establishes the convergence of the cost-to-go functions to the optimal cost-to-go function when the estimates of the environmental process converge.

Proposition 2: Under assumptions A3.1 and A3.2, the cost-to-go functions $J_t(s, q(s)) \to J^*(s, q(S))$ as $t \to \infty$, if $p_t(q(s)) \to p(q(s))$, for all $s \in S$, $q(s) \in Q$.

The optimality equations can be simplified based on the special structure of the path planning problem due to the incoherent environment assumption, which allows us to reduce the dimensionality of the dynamic programming problem. Consider the optimality equation:

$$J(s, q(s)) = \min_u \sum_{(r,\bar{q}(r))} p(r/s, u)p(\bar{q}(r))$$
$$\times [c((r, \bar{q}(r), (s, q(s)), u) + \beta J(r, \bar{q}(r))] \tag{15}$$

Let

$$\bar{c}(s, u, q(s)) = \sum_{(r,q'(r))} p(r/s, u)p(\bar{q}(r))c((r, \bar{q}(r)), (s, q(s)), u) \tag{16}$$

Noting that

$$p((r, \bar{q}(r))/(s, q(s)), u) = p(r/s, u)p(\bar{q}(r)) \tag{17}$$

it follows that

$$J(s, q(s)) = \min_u \left[\bar{c}(s, u, q(s)) + \beta\sum_r p(r/s, u)\bar{J}(r)\right] \tag{18}$$

where

$$\bar{J}(s) = \sum_{q(s)} p(q(s))J(s, q(s)) \tag{19}$$

Noting that

$$u^*(s, q(s)) = \arg\min_u \left[\bar{c}(s, u, q(s)) + \beta\sum_r p(r/s, u)\bar{J}(r)\right] \tag{20}$$

it can be concluded that an average value of the cost-to-go function $J(s, q(s))$, at the system state s, namely, $\bar{J}(s) = \sum_{q(s)} p(q(s))J(s, q(s))$, is required to be able to evaluate the optimal control at any state $(s, q(s))$. This allows us to carry an "average" feedback control. However, there still remains the problem of estimating the average cost-to-go vector $\bar{J}(s)$. To answer this question, note that

$$\bar{J}(s) = \sum_{q(s)} p(q(s))\min_u \left[\bar{c}(s, u, q(s)) + \beta\sum_r p(r/s, u)\bar{J}(r)\right] \tag{21}$$

Hence, $\bar{J}$ is the fixed point of the average dynamic programming operator $\bar{T} : \Re^N \to \Re^N$, defined by the following equation:

$$\bar{T}\,\bar{J}(s) = \sum_{q(s)} p(q(s))\min_u$$

$$\times \left[\bar{c}(s,u,q(s)) + \beta\sum_r p(r/s,u)\bar{J}(r) \right], \quad \forall\; s \qquad (22)$$

The following proposition states that the "average DP operator" $\bar{T}$ is a contraction mapping that maps $R^n \to R^n$ and thus, the optimal average cost-to-go vector, which is a unique fixed point, can be obtained using successive approximations [37]. This allows us to significantly reduce the computational burden of the planning algorithm. Hence, we have the following result.

Proposition 3: The average DP operator $\bar{T}$ is a contraction mapping in $\Re^N$ under the maximum norm.

IV. Testing on an Unmanned Helicopter

The methodology presented thus far in this paper was used for motion planning of an unmanned helicopter in an uncertain environment. The motion planning algorithms were implemented in 6 DOF simulations of an unmanned helicopter navigating through an urban environment model of the Texas A&M campus. The results show that the planning methodology proposed may be suitable for autonomous navigation in a cluttered urban environment.

A. Developing an Urban Environment

An environment model of the inner part of the Texas A&M campus was developed in MATLAB, using maps as well as aerial photos of the campus from Google Earth. Once a three-dimensional campus model was developed, it was overlaid with the aerial images for visualization purposes.

The system state s represents the (x, y) grid point coordinate of the vehicle, and the environment state $q(s)$ represents the presence of an obstacle or otherwise at that particular grid point. A 750×550 m area on the campus of Texas A&M was discretized into a 75×55 grid of points (i rows, j columns) so that the number of possible system states N was equal to 4125. The environment state at any grid point has two possible values, obstacle or no obstacle. At any given location the helicopter had four possible high-level control actions: Go (forward/back/right/left) to the adjacent grid point (see Fig. 2).

B. Nonlocal State Estimation Using a Radar

A nonlocal sensor for motion planning would ideally have an unlimited range, a 360 deg field of view, and no uncertainty. However, a perfect level of accuracy is not possible in practical applications, and there are multiple sensors available that can be used in motion planning. Researchers have used millimeter wave radar, LIDAR (light detection and ranging), infrared sensors, and ultrasonic sensors. LIDAR sensors as well as millimeter wave radars have been implemented on full size helicopters to warn pilots of obstacles. A millimeter wave radar (24.7 GHz) called the Eaton VORAD (vehicle on-board radar) was chosen as the nonlocal sensor to be modeled in this research. The VORAD has an operating range of 1–107 m (3–350 ft), with an uncertainty of 5%, and a field of view of 12 deg, with an uncertainty of 0.2 deg. The radar antenna, onboard processor, and batteries require a helicopter with a lift capability of 6–7 lbs. The radar software was developed to track up to seven obstacles every
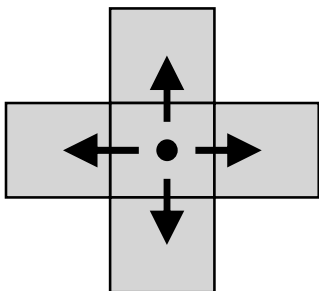

Fig. 2   Possible high-level control actions.

67 ms (15 Hz), by reporting azimuth to each obstacle, range to each obstacle, and gain of the return from each obstacle. The simulation treats all obstacles as having the same radar reflectivity characteristics. Two changes were made to the sensing model to accommodate the simulation: first, the field of view was changed to 90 deg and second, radar tracking was set to nine obstacles.

C. Interfacing with a Low-Level Flight Controller

A low-level flight controller known as the AFCS, developed by Rotomotion, LLC, for unmanned helicopters, was used to track the high-level flight control commands. The AFCS uses two extended Kalman filters for state estimation: the first is a seven-state Kalman filter, four states for quaternion and three states for gyro bias, and the second is a six-state Kalman filter, three states for the North–East–Down (NED) position and three states for orthogonal velocity components (UVW). The flight controller implements nine PID loops—three for position (x, y, z), three for velocity $(\dot{x}, \dot{y}, \dot{z})$, and three for attitude (ψ, θ, ϕ). The AFCS primary position sensor uses the wide area augmentation system global positioning system and gives position measurements to within 3 m accuracy 95% of the time.

D. Motion Planning in the Flight Simulator

1. Navigation and Exploration

The high-level motion planner calculates an initial cost-to-go map $\bar{J}(s)$, which is an array of i rows and j columns, based upon the a priori environment data available. Before the high-level motion planner implements a control, it senses and updates the environment model (in terms of the environmental probabilities) at eight adjacent grid points. If the helicopter is commanded in a direction that has not been sensed by the radar sensor, the helicopter changes its direction to the commanded direction and performs new measurements of the environment states. This is because the radar only has a 90 deg field of vision, and thus, cannot sense all directions at once. Then the motion planner recalculates the control and the command is issued to the low-level flight controller for execution.

The operation of the high-level planner involves four different tuning parameters. The cost associated with hitting an obstacle is $\zeta = 100$, while the transition cost to visit a state with no obstacle is $\eta = 1$. Various tuning values were used for ζ and η. The only importance of the numbers for ζ and η is their magnitude relative to each other. An infinite horizon discount factor of $\beta = 0.99$ was used to ensure that states in the immediate future have a greater effect on the control than states farther out on the horizon. For lower β values the vehicle is more likely to be trapped once the vehicle enters a boxed in area. Also a tuning parameter of $\alpha = 6$ was included in the simulations to represent a weighting value for distance to the goal. For large α values, the vehicle is drawn to the final destination but hits obstacles, whereas the value of $\alpha = 6$ draws the vehicle to the final destination without hitting obstacles. It is necessary to tune the parameters because the path planning problem is posed as an infinite horizon problem. The structure of the control is given by

$$u^*(s, q(s)) = \arg\min_u \left[\bar{c}(s,u,q(s)) + \beta\sum_r p(r/s,u)\bar{J}(r) \right]$$

where the cost to go is

$$\bar{J}(s) = \sum_{q(s)} p(q(s))\min_u \left[\bar{c}(s,u,q(s)) + \beta\sum_r p(r/s,u)\bar{J}(r) + \alpha d(r) \right]$$

The $\bar{c}(s, u, q(s))$ term represents the average transition cost from a state s to an adjacent state r, given that the control is u. The term $\alpha d(r)$ represents the tuning parameter α multiplied by the distance of the future state r to the goal state.

The results shown in Figs. 3–15 are for an a priori environment model accuracy of 90%. This number quantifies the reliability of the a priori environment model. Various levels of a priori environment model accuracy were used in the initial testing of the algorithm. If the a priori environment model accuracy was treated as 50%, then the a priori model did not have a meaningful contribution to the motion
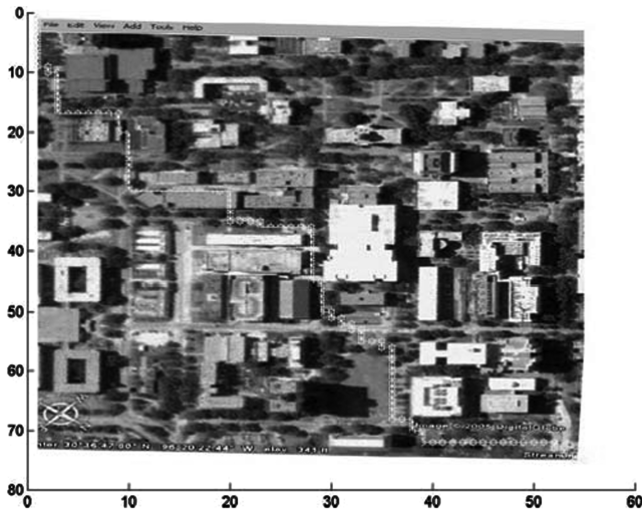
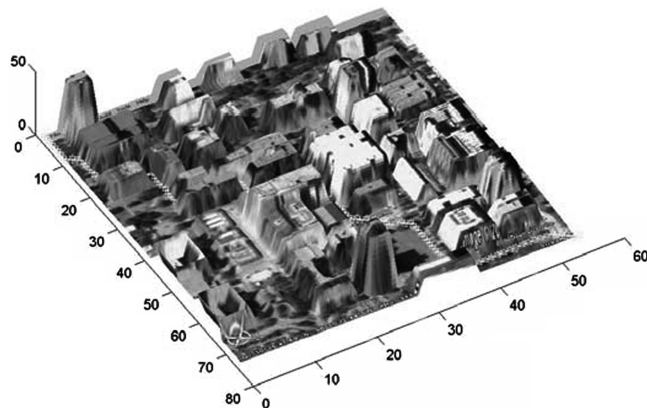**Fig. 3 Top view of navigation through campus.**

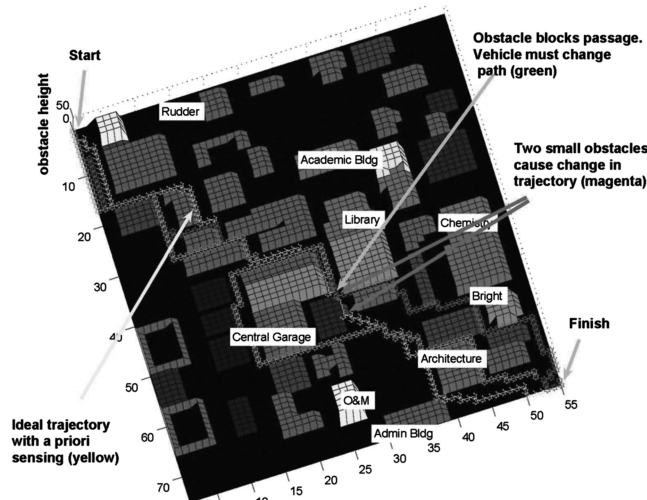**Fig. 4 Angled view of navigation through campus.**

Fig. 5 Obstacles and blocked passage.

planning (because a 50% reliability is as good as flipping a fair coin). As the a priori accuracy of the model increased from 50 to 90%, the initial motion plan became more reliable and less likely to hit obstacles.

The sensor model was generated by simulating sensor data using a Monte Carlo approach. Each sensor measurement was the true value corrupted by a Gaussian noise term based upon the distance to the
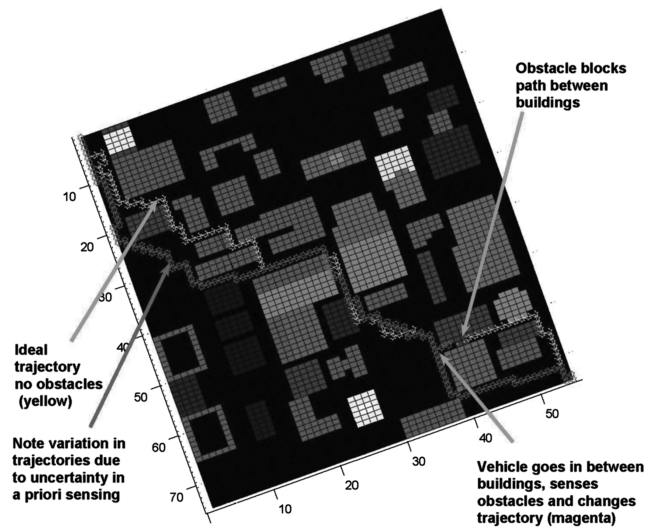

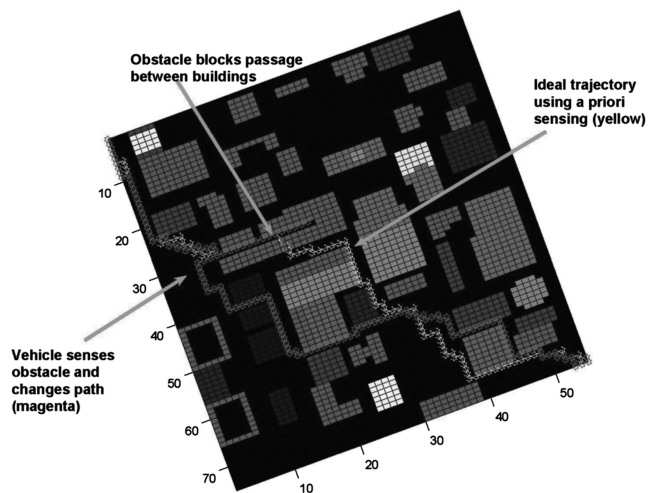**Fig. 6 Blocked passage by Langford.**

Fig. 7 Blocked passage by Heldenfelds.

obstacle. The simulations using the Eaton VORAD model had a 95% accuracy for the nonlocal state sensing.

The simulations were also performed with various different uncertainty models so that the effects of using sensors, with different accuracy levels, on the performance of the algorithm could be quantified. For example, in Figs. 5–7 the sensor model has an accuracy of 96% for the local sensing, an accuracy of 93% at a distance of 10 m, and an accuracy of 91% at a distance of 14 m. The accuracy of the radar for a given range was found by linearly interpolating between the accuracy of the sensor at 1 m (99%) and 100 m (57%). For a sensor uncertainty of 50%, the vehicle performed as if no sensor data were available. For sensor uncertainties between 70 and 80%, the performance increased significantly. For uncertainties in the range of 90–99%, the performance remained approximately the same. The motion planning software was tested multiple times before executing the waypoints in the flight simulator. Figures 3 and 4 have been included to show a path that the high-level flight controller generated to navigate through campus while using a priori map data as well as radar sensor information. The plots have been overlaid with aerial photos to display the simulated vehicle trajectory through the A&M campus.

2. Avoiding Obstacles

The simulations in Figs. 3 and 4 used a static environment model. However, it is realistic to consider the situation where the

environment model is dynamic. If a newly sensed obstacle such as a vehicle or fallen building is encountered, the high-level flight controller should be able to adapt to the change. When the helicopter radar senses a new obstacle, it creates a large transition cost so that the helicopter avoids flying into the newly sensed obstacle. For instance, in Figs. 5–7, additional obstacles were added to the map after the helicopter started moving, so that the helicopter encountered a dynamic obstacle field and blocked passages at various points on the A&M campus. At the algorithmic level, this alters the cost structure of the DP algorithms and thus, the helicopter has to replan when such an event occurs, that is, reevaluate the cost-to-go map of the entire planning domain.

3. Replanning at a Blocked Passage

Occasionally, the vehicle may encounter an obstacle that entirely blocks a passage. The vehicle may not have time to recalculate the entire cost map J, that is, the new optimal cost to go for the entire environment. Practically, updating the local region's cost map (a 10×10 grid) is sufficient for the vehicle to replan its trajectory to get out of such an impasse. This is a heuristic modification of the theoretical framework which makes it practically implementable. The size of the local region to update is a heuristic choice and is addressed experimentally. In this case, updating the local 10×10 grid was sufficient. It is conceivable that if the vehicle is unable to find a path out of a local impasse due to a blocked passage, then the global cost-to-go map may need to be recalculated. Figures 5–7 display the motion planner's performance with respect to blocked passages at various locations on campus.

4. Motion Planning in the Flight Simulator

Multiple tests were run while the high-level motion planner and low-level flight controller were connected together. The high-level controller computation time was relatively minor compared to the time necessary to execute the waypoint commands. Running each MATLAB simulation took approximately 15–20 s, while a simulation of actually flying the helicopter to 180 waypoints, which were 10 m apart, took approximately 12–15 min. If 180 waypoints were executed during the initial cost-to-go calculation then approximately $180^*N^*M \approx 3,000,000$ major computations were required, which requires a computation time of approximately (5–10 s).

During the MATLAB simulations the vehicle made a minimum of one local and eight nonlocal sensor measurements and up to four local and 32 nonlocal sensor measurements per waypoint. This resulted in approximately 1800 optimization runs for the various environment probabilities and 720 waypoint calculations if 180 waypoints were executed. This requires a computation time of approximately 5–10 s. Figures 8–15 demonstrate the performance of the motion planner and the flight simulator while connected together. Figures 8 and 9 were generated during simulator testing 2 using the campus map along with a fictitious blocked passage by the library and a blocked passage by the pavilion. The obstacles were introduced after the helicopter started its flight. The vehicle does not always choose the same path. In this testing the vehicle did not take the anticipated route of going beside the library because it accounted for sensor uncertainties and planned a route by Dunn.

Figures 10 and 11 were generated during simulator testing 3 using the campus map with an a priori unknown blocked passage by the library. The helicopter sensed the blocked passage by the library and moved back and forth several times until it recalculated a local cost map. It took several local cost map recalculations before the helicopter was able to determine a path between the library and the library annex.

In a few simulations it was observed that a local cost map recalculation may not be sufficient to update the cost map and guide the helicopter out of a blocked passage. By only performing a local cost map update, it is possible to get trapped in a local minimum. The local cost map recalculations may need to be enlarged until the helicopter is able to find its way out of an area. If a local cost map is unable to update the cost map sufficiently for the helicopter to find its
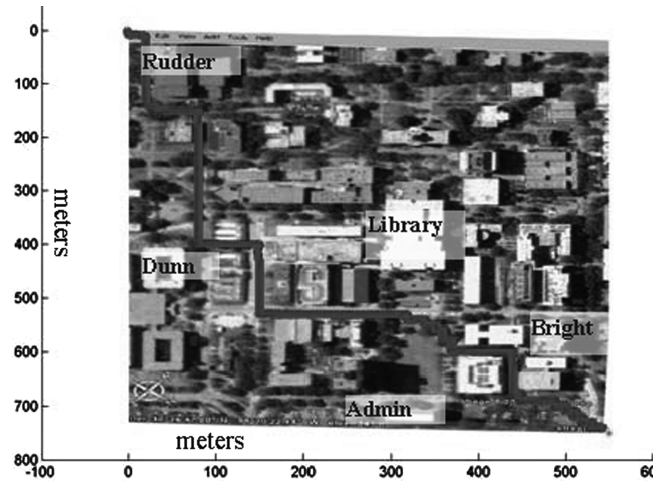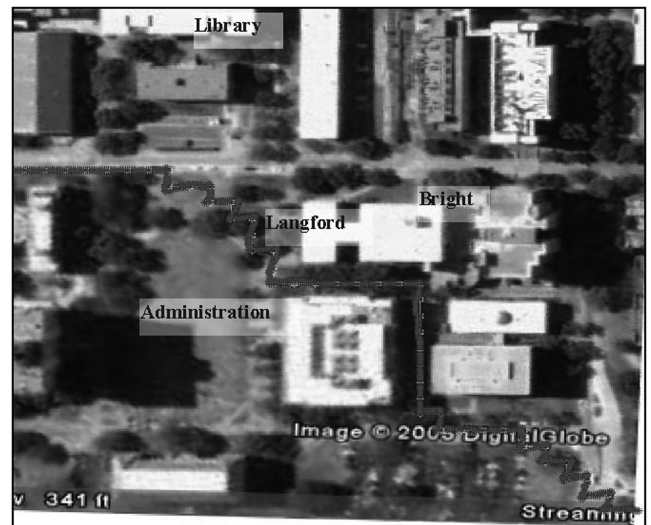

Fig. 8 Simulator testing 2—large view.

Fig. 9 Simulator testing 2—detailed view.

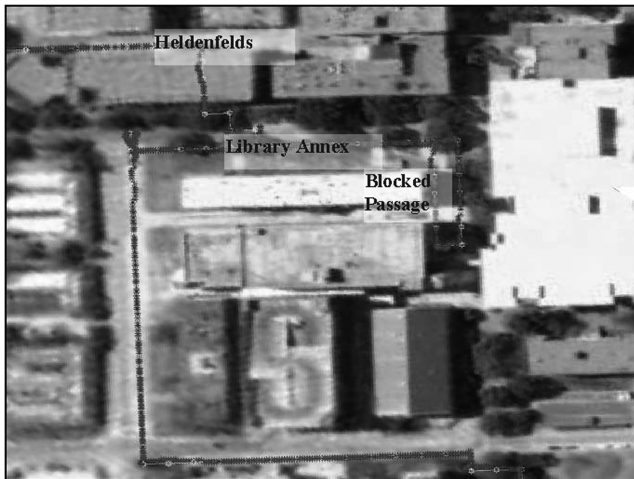
Fig. 10 Simulator testing 3—large view.
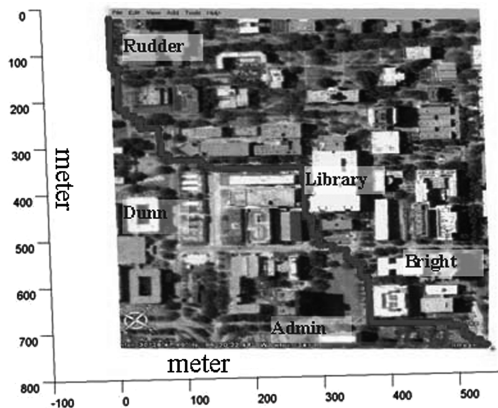
Fig. 11   Simulator testing 3—detailed view.

Fig. 14   Simulator testing 5—large view.

Fig. 12   Simulator testing 4—large view.

Fig. 15   Simulator testing 5—detailed view.

Fig. 13   Simulator testing 4—detailed view.

way out of an area, then it may be necessary to perform an entire recalculation of the cost map. Because dynamic programming by definition is a global optimization, if the entire cost map is recalculated the helicopter will be able to determine the optimal path to the final destination.

Figure 10 includes a blocked passage by the library in which the helicopter recalculated the local cost map to exit.

Figures 12 and 13 were generated during simulator testing 4 using the obstacle map of the campus with a blocked passage by
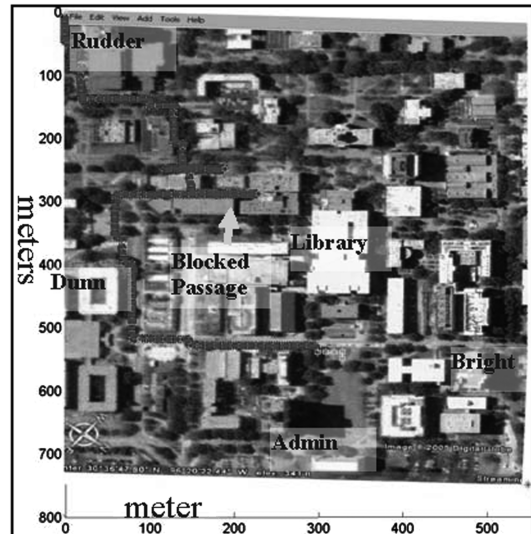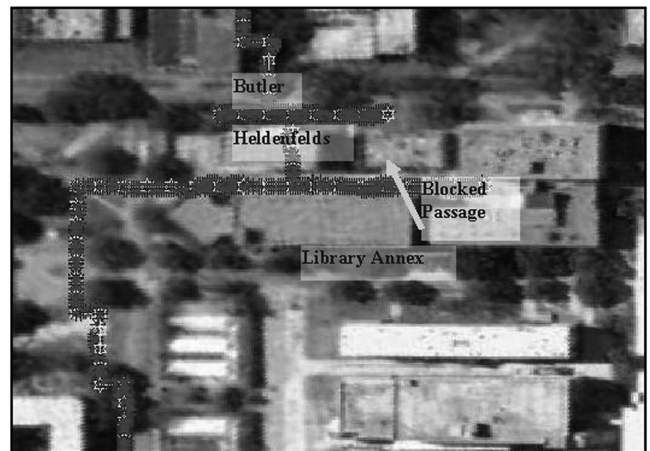
Heldenfelds and Langford. The blocked passages and two obstacles near the library were introduced after the helicopter started its flight. The helicopter encountered only one of the obstacles and the result is represented in Fig. 13.

Figures 14 and 15 were generated during the simulator testing 5 using the obstacle map of the campus with a blocked passage by Heldenfelds and Langford. The helicopter encountered only the blocked passage by Heldenfelds. After several recalculations of the local cost map the helicopter was able to leave the boxed-in area.

V.   Conclusions

In this work, a methodology was presented for intelligent exploration of a partially known environment using a nonlocal sensor. It was shown that a motion planning problem, under certain assumptions, can be reduced to the adaptive control of an uncertain Markov decision process, consisting of a known control dependent system state and an unknown control independent environment. The feasibility of the planning methodology was illustrated by testing on an unmanned helicopter navigating through an urban environment in a 6 DOF flight simulation.

The frequency of the cost-to-go update is of considerable interest. At one end of the spectrum, the entire cost-to-go map could be updated at every time instant, which might be computationally infeasible, while at the other end of the spectrum, only an initial cost-to-go map could be calculated before the navigation begins.

However, both these extremes are possibly not "optimal" and the best solution might be somewhere midway. It was surmised that the cost to go may need to be changed when the environments, which are sensed after the helicopter begins its flight, start looking "significantly different" from the estimates developed according to a priori data. However, these are qualitative statements and need to be quantified in terms of algorithms.

Another area of interest is to consider implementing heuristics in the dynamic programming algorithm that would allow for an accelerated computation of the cost-to-go map by initially excluding the evaluation of states that are far from the desired path of the helicopter. Also, the controller architecture used in this research only uses the high-level motion planner to detect and avoid obstacles, while the low-level controller is assumed to track the desired trajectories closely enough to ensure obstacle avoidance. A low-level controller needs to be designed that guarantees the local avoidance of obstacles while tracking the high-level motion plans, so that integration of the high-level motion planner and low-level flight controller can result in a truly intelligent autonomous system. Further, the amalgamation of the methodology presented in this paper with existing motion planning methods such as probabilistic road maps or D* might lead to more robust, near real-time implementable motion planning algorithms.

References

[1] Chakravorty, S., and Junkins, J. L., "Intelligent Exploration of Unknown Environments with Vision Like Sensors," *Proceedings of the 2005 IEEE/ASME International Conference of Advanced Intelligent Mechatronics*, IEEE, Piscataway, NJ, 2005, pp. 1204–1209.

[2] Chakravorty, S., and Junkins, J. L., "A Methodology for Intelligent Path Planning," *Proceedings of the 2005 IEEE International Symposium on Intelligent Control*, IEEE, Piscataway, NJ, 2005, pp. 592–597.

[3] Kumar, P. R., and Varaiya, P., *Stochastic Systems: Estimation, Identification and Adaptive Control*, Prentice–Hall, Englewood Cliffs, NJ, 1986.

[4] Borkar, V., and Varaiya, P., "Adaptive Control of Markov Chains I: Finite Parameter Set," *IEEE Transactions on Automatic Control*, Vol. 24, No. 6, 1979, pp. 953–958.

[5] Mandl, P., "Estimation and Control in Markov Chains," *Advances in Applied Probability*, Vol. 6, No. 1, March 1974, pp. 40–60.

[6] Bertsekas, D. P., and Tsitsiklis, J. N., *Neuro-Dynamic Programming*, Athena Scientific, Belmont, MA, 1996.

[7] Sutton, R. S., and Barto, A. G., *Reinforcement Learning: An Introduction*, MIT Press, Cambridge, MA, 1998.

[8] Sutton, R. S., Barto, A. G., and Williams, R. J., "Reinforcement Learning is Direct Adaptive Optimal Control," *IEEE Control Systems Magazine*, Vol. 12, No. 2, 1992, pp. 19–22.

[9] Bellman, R. E., *Dynamic Programming*, Princeton Univ. Press, Princeton, NJ, 1957.

[10] Bellman, R. E., and Dreyfus, S. E., *Applied Dynamic Programming*, Princeton Univ. Press, Princeton, NJ, 1962.

[11] Kirk, D. E., *Optimal Control Theory: An Introduction*, Prentice–Hall, Englewood Cliffs, NJ, 1970.

[12] Lozano-Pérez, T., "Spatial Planning: A Configuration Space Approach," *IEEE Transactions on Computers*, Vol. C-32, No. 2, 1983, pp. 108–120.

[13] Latombe, J. C., *Robot Motion Planning*, Kluwer, Boston, MA, 1991.

[14] Latombe, J. C., Lazanas, A., and Shekhar, S., "Robot Motion Planning with Uncertainty in Control and Sensing," *Artificial Intelligence*, Vol. 52, No. 1, Nov. 1991, pp. 1–47.

[15] Mason, M. T., "Automatic Planning of Fine Motions: Correctness and Completeness," *Proceedings of the IEEE Conference on Robotics and Automation*, IEEE, Piscataway, NJ, 1989, pp. 484–489.

[16] Lavalle, S. M., "Robot Motion Planning: A Game-Theoretic Foundation," *Algorithmica*, Vol. 26, Nos. 3–4, March 2000, pp. 430–465.

[17] Thrun, S., "A Probabilistic On-Line Mapping Algorithm for Teams of Mobile Robots," *The International Journal of Robotic Research*, Vol. 20, No. 5, May 2001, pp. 335–363.

[18] Thrun, S., "Probabilistic Algorithms in Robotics," *AI Magazine*, Vol. 21, No. 4, 2000, pp. 93–109.

[19] Burgard, W., Fox, D., Jans, H., Matenar, C., and Thrun, S., "Sonar-Based Mapping of Large-Scale Mobile Robot Environments Using EM," *Proceedings of the International Conference on Machine Learning*, Morgan Kaufmann, San Francisco, CA, 1999, pp. 23–29.

[20] Castellanos, J. A., Montiel, J. M., Neira, J., and Tardos, J. D., "The SP Map: A Probabilistic Framework for Simultaneous Localization and Mapping," *IEEE Transactions on Robotics and Automation*, Vol. 15, No. 5, Oct. 1999, pp. 948–953.

[21] Dissanayake, G., Durant-White, H., and Bailey, T., "A Computationally Efficient Solution to the Simultaneous Localization and Mapping Problem," *ICRA'2000 Workshop W4: Mobile Robot Navigation and Mapping*, IEEE, Piscataway, NJ, April 2000.

[22] LaValle, S. M., *Planning Algorithms*, Cambridge Univ. Press, New York, 2006.

[23] Korf, R. E., "Artificial Intelligence Search Algorithms," *Algorithms and Theory of Computation Handbook*, CRC Press, Boca Raton, FL, 1999.

[24] Nilsson, N. J., *Principles of Artificial Intelligence*, Tioga Publishing Co., Wellsboro, PA, 1980.

[25] Cormen, T. H., Leiserson, C. E., Rivest, R. L., and Stein, C., *Introduction to Algorithms*, 2nd ed., MIT Press, Cambridge, MA, 2001.

[26] Goodrich, M. T., and Tammasia, R., *Algorithm Design: Foundations, Analysis and Internet Examples*, Wiley, New York, 2002.

[27] Dijkstra, E. W., "A Note on Two Problems in Connection with Graphs," *Numerische Mathematik*, Vol. 1, No. 1, Nov. 1959, pp. 269–271.

[28] Pearl, J., *Heuristics: Intelligent Search Strategies for Computer Problem Solving*, Addison–Wesley, Boston, MA, 1985.

[29] Stentz, A., "The Focused D* Algorithm for Real-Time Replanning," *Proceedings of the International Joint Conference on Artificial Intelligence*, Morgan Kaufmann, San Francisco, CA, 1995, pp. 1652–1659.

[30] Stentz, A., and Herbert, M., "A Complete Navigation System for Goal Acquisition in Unknown Environments," *Autonomous Robots*, Vol. 2, No. 2, 1995, pp. 127–145.

[31] Hebert, M., McLachlan, R., and Chang, P., "Experiments with Driving Modes for Urban Robots," *Proceedings of SPIE: The International Society for Optical Engineering*, Vol. 3838, 1999, pp. 70–78.

[32] Koenig, S., and Likhachev, M., "Fast Replanning for Navigation in Unknown Terrain," *IEEE Transactions on Robotics and Automation*, Vol. 21, No. 3, 2005, pp. 354–363.

[33] Ghallab, M., Nau, D., and Traverso, P., *Automated Planning: Theory and Practice*, Morgan Kaufman, San Francisco, CA, 2004.

[34] Russell, S., and Norvig, P., *Artificial Intelligence: A Modern Approach*, 2nd ed., Prentice–Hall, Englewood Cliffs, NJ, 2003.

[35] Kavraki, L. E., Svestka, P., Latombe, J. C., and Overmars, M. H., "Probabilistic Roadmaps for Path Planning in High-Dimensional Configuration Spaces," *IEEE Transactions on Robotics and Automation*, Vol. 12, No. 4, 1996, pp. 566–580.

[36] Song, G., Shawna, T., and Amato, N., "A General Framework for PRM Motion Planning," *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, Piscataway, NJ, 2003, pp. 4445–4450.

[37] Khalil, H. K., *Nonlinear Systems*, 2nd ed., Prentice–Hall, Upper Saddle River, NJ, 1996, p. 64.