

Optimal Path-planning Under Finite Memory Obstacle Dynamics Based on Probabilistic Finite State Automata Models[†]

Ishanu Chattopadhyay[‡]
 ixc128@psu.edu

Asok Ray[‡]
 axr2@psu.edu

Abstract—The ν^* -planning algorithm is generalized to handle finite memory obstacle dynamics. A sufficiently long observation sequence of obstacle dynamics is algorithmically compressed via Symbolic Dynamic Filtering to obtain a probabilistic finite state model which is subsequently integrated with the navigation automaton to generate an overall model reflecting both navigation constraints and obstacle dynamics. A ν^* -based solution then yields a deterministic plan that maximizes the difference of the probabilities of reaching the goal and of hitting an obstacle. The approach is validated by simulated solution of dynamic mazes.

Index Terms—Language Measure; Probabilistic Finite State Machines; Robotics; Path Planning; Supervisory Control

1. INTRODUCTION & MOTIVATION

Recently, a novel path planning algorithm ν^* was reported that models the navigation problem in the framework of Probabilistic Finite State Machines and computes robust optimal plans via optimization of the PFSA from a strictly control-theoretic viewpoint. In this paper, we present a significant improvement; ν^* is generalized to handle finite memory obstacle dynamics. A sufficiently long observation sequence of obstacle movement in terms of an evolving obstacle map is symbolically compressed to obtain a finite state probabilistic model of obstacle dynamics which is then used to compute optimal plans in dynamic cluttered environments. It is important to note that the problem considered in this paper is different from the ones that modify plans on-the-fly (e.g. D^*), to incorporate newly learnt information about obstacle locations; in this paper we propose an approach for optimally incorporating apriori knowledge of the expected dynamical evolution of obstacles. The key advantages are:

- 1) **Pre-processing is cheap:** The cellular decomposition required by ν^* is simple and computationally cheap. The cells are mapped to PFSA states which are defined to have identical connectivity via symbolic inter-state transitions.
- 2) **Fundamentally different from search:** ν^* optimizes the resultant PFSA via a iterative sequence of combinatorial operations which elementwise maximizes the language measure vector [1][2].
- 3) **Computational efficiency:** The time complexity of each iteration step can be shown to be linear in problem size implying significant numerical advantage over search-based methods for high-dimensional problems.
- 4) **Global monotonicity:** The solution iterations are globally monotonic. The final waypoint sequence is generated essentially by following the measure gradient which is maximized at the goal. The measure gradient is reminiscent of potential field methods [3]. However, ν^* automatically generates the measure gradient; no potential function is necessary. Furthermore, the potential function based planners often get trapped in local minimum which can be shown to be a mathematical impossibility for ν^* .

The paper is organized in five sections including the present one. Section 2 briefly explains the language-theoretic models considered in this paper, reviews the language-measure-theoretic optimal control

of probabilistic finite state machines and presents the necessary details of the reported ν^* algorithm. Section 3 presents the approach for symbolic compression of observed obstacle dynamics. Section 4 formulates the integration of the obstacle automaton and the navigation automaton to generate the overall model to be subsequently optimized via ν^* planning. Simulation examples are included for illustration. The paper is summarized and concluded in Section 5 with recommendations for future work.

2. LANGUAGE MEASURE-THEORETIC OPTIMIZATION

This section summarizes the signed real measure of regular languages; the details are reported in [1]. Let $G_i \equiv \langle Q, \Sigma, \delta, q_i, Q_m \rangle$ be a trim (i.e., accessible and co-accessible) finite-state automaton model that represents the discrete-event dynamics of a physical plant, where $Q = \{q_k : k \in \mathcal{I}_Q\}$ is the set of states and $\mathcal{I}_Q \equiv \{1, 2, \dots, n\}$ is the index set of states; the automaton starts with the initial state q_i ; the alphabet of events is $\Sigma = \{\sigma_k : k \in \mathcal{I}_\Sigma\}$, having $\Sigma \cap \mathcal{I}_Q = \emptyset$ and $\mathcal{I}_\Sigma \equiv \{1, 2, \dots, \ell\}$ is the index set of events; $\delta : Q \times \Sigma \rightarrow Q$ is the (possibly partial) function of state transitions; and $Q_m \equiv \{q_{m_1}, q_{m_2}, \dots, q_{m_l}\} \subseteq Q$ is the set of marked (i.e., accepted) states with $q_{m_k} = q_j$ for some $j \in \mathcal{I}_Q$. Let Σ^* be the Kleene closure of Σ , i.e., the set of all finite-length strings made of the events belonging to Σ as well as the empty string ϵ that is viewed as the identity of the monoid Σ^* under the operation of string concatenation, i.e., $\epsilon s = s = s \epsilon$. The state transition map δ is recursively extended to its reflexive and transitive closure $\delta : Q \times \Sigma^* \rightarrow Q$ by defining $\forall q_j \in Q, \delta(q_j, \epsilon) = q_j$ and $\forall q_j \in Q, \sigma \in \Sigma, s \in \Sigma^*, \delta(q_j, \sigma s) = \delta(\delta(q_j, \sigma), s)$

Definition 2.1: The language $L(q_i)$ generated by a DFSA G initialized at the state $q_i \in Q$ is defined as: $L(q_i) = \{s \in \Sigma^* \mid \delta^*(q_i, s) \in Q\}$ The language $L_m(q_i)$ marked by the DFSA G initialized at the state $q_i \in Q$ is defined as: $L_m(q_i) = \{s \in \Sigma^* \mid \delta^*(q_i, s) \in Q_m\}$

Definition 2.2: For every $q_j \in Q$, let $L(q_i, q_j)$ denote the set of all strings that, starting from the state q_i , terminate at the state q_j , i.e., $L_{i,j} = \{s \in \Sigma^* \mid \delta^*(q_i, s) = q_j \in Q\}$

The formal language measure is first defined for terminating plants [4] with sub-stochastic event generation probabilities i.e. the event generation probabilities at each state summing to strictly less than unity.

Definition 2.3: The event generation probabilities are specified by the function $\tilde{\pi} : \Sigma^* \times Q \rightarrow [0, 1]$ such that $\forall q_j \in Q, \forall \sigma_k \in \Sigma, \forall s \in \Sigma^*,$

- (1) $\tilde{\pi}(\sigma_k, q_j) \triangleq \tilde{\pi}_{jk} \in [0, 1]; \sum_k \tilde{\pi}_{jk} = 1 - \theta$, with $\theta \in (0, 1)$;
- (2) $\tilde{\pi}(\sigma, q_j) = 0$ if $\delta(q_j, \sigma)$ is undefined; $\tilde{\pi}(\epsilon, q_j) = 1$;
- (3) $\tilde{\pi}(\sigma_k s, q_j) = \tilde{\pi}(\sigma_k, q_j) \tilde{\pi}(s, \delta(q_j, \sigma_k))$.

The $n \times \ell$ event cost matrix is defined as: $\tilde{\Pi}_{ij} = \tilde{\pi}(q_i, \sigma_j)$

Definition 2.4: The state transition probability $\pi : Q \times Q \rightarrow [0, 1]$, of the DFSA G_i is defined as follows: $\forall q_i, q_j \in Q, \pi_{ij} = \sum_{\sigma \in \Sigma \text{ s.t. } \delta(q_i, \sigma) = q_j} \tilde{\pi}(\sigma, q_i)$ The $n \times n$ state transition probability matrix is defined as $\Pi_{jk} = \pi(q_j, q_k)$

The set Q_m of marked states is partitioned into Q_m^+ and Q_m^- , i.e., $Q_m = Q_m^+ \cup Q_m^-$ and $Q_m^+ \cap Q_m^- = \emptyset$, where Q_m^+ contains all good marked states that we desire to reach, and Q_m^- contains all bad marked states that we want to avoid, although it may not always be possible to completely avoid the bad states while attempting to reach the good states. To characterize this, each marked state is assigned a

[†]This work has been supported in part by the U.S. Army Research Laboratory and the U.S. Army Research Office under Grant No. W911NF-07-1-0376 and by the Office of Naval Research under Grant No. N00014-08-1-380. Any opinions, findings and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the sponsoring agencies.

[‡]The Pennsylvania State University, University Park, PA 16802, USA.

real value based on the designer's perception of its impact on the system performance.

Definition 2.5: The characteristic function $\chi : Q \rightarrow [-1, 1]$ that assigns a signed real weight to state-based sublanguages $L(q_i, q)$

is defined as: $\forall q \in Q, \chi(q) \in \begin{cases} [-1, 0), & q \in Q_m^- \\ \{0\}, & q \notin Q_m \\ (0, 1], & q \in Q_m^+ \end{cases}$ The state

weighting vector, denoted by $\chi = [\chi_1 \chi_2 \cdots \chi_n]^T$, where $\chi_j \equiv \chi(q_j) \forall j \in I_Q$, is called the χ -vector. The j -th element χ_j of χ -vector is the weight assigned to the corresponding terminal state q_j .

In general, the marked language $L_m(q_i)$ consists of both good and bad event strings that, starting from the initial state q_i , lead to Q_m^+ and Q_m^- respectively. Any event string belonging to the language $L^0 = L(q_i) - L_m(q_i)$ leads to one of the non-marked states belonging to $Q - Q_m$ and L^0 does not contain any one of the good or bad strings. Based on the equivalence classes defined in the Myhill-Nerode Theorem, the regular languages $L(q_i)$ and $L_m(q_i)$ can be expressed as: $L(q_i) = \bigcup_{q_k \in Q} L_{i,k}$ and $L_m(q_i) = \bigcup_{q_k \in Q_m} L_{i,k} = L_m^+ \cup L_m^-$ where the sublanguage $L_{i,k} \subseteq G_i$ having the initial state q_i is uniquely labelled by the terminal state $q_k, k \in I_Q$ and $L_{i,j} \cap L_{i,k} = \emptyset \forall j \neq k$; and $L_m^+ \equiv \bigcup_{q_k \in Q_m^+} L_{i,k}$ and $L_m^- \equiv \bigcup_{q_k \in Q_m^-} L_{i,k}$ are good and bad sublanguages of $L_m(q_i)$, respectively. Then, $L^0 = \bigcup_{q_k \notin Q_m} L_{i,k}$ and $L(q_i) = L^0 \cup L_m^+ \cup L_m^-$.

A signed real measure $\mu^i : 2^{L(q_i)} \rightarrow \mathbb{R} \equiv (-\infty, +\infty)$ is constructed on the σ -algebra $2^{L(q_i)}$ for any $i \in I_Q$; interested readers are referred to [1] for the details of measure-theoretic definitions and results. With the choice of this σ -algebra, every singleton set made of an event string $s \in L(q_i)$ is a measurable set. By Hahn Decomposition Theorem [5], each of these measurable sets qualifies itself to have a numerical value based on the above state-based decomposition of $L(q_i)$ into L^0 (null), L^+ (positive), and L^- (negative) sublanguages.

Definition 2.6: Let $\omega \in L(q_i, q_j) \subseteq 2^{L(q_i)}$. The signed real measure μ^i of every singleton string set $\{\omega\}$ is defined as: $\mu^i(\{\omega\}) \equiv \tilde{\pi}(\omega, q_i)\chi(q_j)$. The signed real measure of a sublanguage $L_{i,j} \subseteq L(q_i)$ is defined as: $\mu_{i,j} \equiv \mu^i(L(q_i, q_j)) = \left(\sum_{\omega \in L(q_i, q_j)} \tilde{\pi}[\omega, q_i] \right) \chi_j$

Therefore, the signed real measure of the language of a DFSA G_i initialized at $q_i \in Q$, is defined as $\mu_i \equiv \mu^i(L(q_i)) = \sum_{j \in I_Q} \mu^i(L_{i,j})$. It is shown in [1] that the language measure can be expressed as $\mu_i = \sum_{j \in I_Q} \pi_{ij} \mu_j + \chi_i$. The language measure vector, denoted as $\boldsymbol{\mu} = [\mu_1 \mu_2 \cdots \mu_n]^T$, is called the $\boldsymbol{\mu}$ -vector. In vector form, we have $\boldsymbol{\mu} = \boldsymbol{\Pi} \boldsymbol{\mu} + \boldsymbol{\chi}$ whose solution is given by $\boldsymbol{\mu} = (\mathbf{I} - \boldsymbol{\Pi})^{-1} \boldsymbol{\chi}$. The inverse exists for terminating plant models [4] because $\boldsymbol{\Pi}$ is a contraction operator [1] due to the strict inequality $\sum_j \Pi_{ij} < 1$. The residual $\theta_i = 1 - \sum_j \Pi_{ij}$ is referred to as the termination probability for state $q_i \in Q$. We extend the analysis to non-terminating plants with stochastic transition probability matrices (i.e. with $\theta_i = 0, \forall q_i \in Q$) by normalizing the language measure [1] with respect to the uniform termination probability of a limiting terminating model as described next.

Let $\tilde{\boldsymbol{\Pi}}$ and $\boldsymbol{\Pi}$ be the stochastic event generation and transition probability matrices for a non-terminating plant $G_i = \langle Q, \Sigma, \delta, q_i, Q_m \rangle$. We consider the terminating plant $G_i(\theta)$ with the same DFSA structure $\langle Q, \Sigma, \delta, q_i, Q_m \rangle$ such that the event generation probability matrix is given by $(1 - \theta)\tilde{\boldsymbol{\Pi}}$ with $\theta \in (0, 1)$ implying that the state transition probability matrix is $(1 - \theta)\boldsymbol{\Pi}$.

Definition 2.7: (Renormalized Measure:) The renormalized measure $\nu_\theta^i : 2^{L(q_i(\theta))} \rightarrow [-1, 1]$ for the θ -parametrized terminating plant $G_i(\theta)$ is defined as: $\forall \omega \in L(q_i(\theta)), \nu_\theta^i(\{\omega\}) = \theta \mu^i(\{\omega\})$. The corresponding matrix form is given by $\boldsymbol{\nu}_\theta = \theta \boldsymbol{\mu} = \theta [\mathbf{I} - (1 - \theta)\boldsymbol{\Pi}]^{-1} \boldsymbol{\chi}$ with $\theta \in (0, 1)$. We note that the vector representation allows for the following notational simplification $\nu_\theta^i(L(q_i(\theta))) = \boldsymbol{\nu}_\theta^i$. The renormalized measure for the non-terminating plant G_i is defined to be $\lim_{\theta \rightarrow 0^+} \boldsymbol{\nu}_\theta^i$.

A. Event-driven Supervision of PFSA

Plant models considered in this paper are deterministic finite state automata (plant) with well-defined event occurrence probabilities. In other words, the occurrence of events is probabilistic, but the state at which the plant ends up, given a particular event has occurred, is deterministic. Since no emphasis is placed on the initial state and marked states are completely determined by χ , the models can be completely specified by a sextuple as: $G = (Q, \Sigma, \delta, \tilde{\boldsymbol{\Pi}}, \chi, \mathcal{C})$

Definition 2.8: (Control Philosophy) If $q_i \xrightarrow{\sigma} q_k$, and the event σ is disabled at state q_i , then the supervisory action is to prevent the plant from making a transition to the state q_k , by forcing it to stay at the original state q_i . Thus disabling any transition σ at a given state q results in deletion of the original transition and appearance of the self-loop $\delta(q, \sigma) = q$ with the occurrence probability of σ from the state q remaining unchanged in the supervised and unsupervised plants. For a given plant, transitions that can be disabled in the sense of Definition 2.8 are defined to be controllable transitions. The set of controllable transitions in a plant is denoted \mathcal{C} . Note controllability is state-based.

B. The Optimal Supervision Problem: Formulation & Solution

A supervisor disables a subset of the set \mathcal{C} of controllable transitions and hence there is a bijection between the set of all possible supervision policies and the power set $2^{\mathcal{C}}$. That is, there exists $2^{|\mathcal{C}|}$ possible supervisors and each supervisor is uniquely identifiable with a subset of \mathcal{C} and the language measure ν allows a quantitative comparison of different policies.

Definition 2.9: For an unsupervised plant $G = (Q, \Sigma, \delta, \tilde{\boldsymbol{\Pi}}, \chi, \mathcal{C})$, let G^\dagger and G^\ddagger be the supervised plants with sets of disabled transitions, $\mathcal{D}^\dagger \subseteq \mathcal{C}$ and $\mathcal{D}^\ddagger \subseteq \mathcal{C}$, respectively, whose measures are $\boldsymbol{\nu}^\dagger$ and $\boldsymbol{\nu}^\ddagger$. Then, the supervisor that disables \mathcal{D}^\dagger is defined to be superior to the supervisor that disables \mathcal{D}^\ddagger if $\boldsymbol{\nu}^\dagger \succeq_{(\text{Elementwise})} \boldsymbol{\nu}^\ddagger$ and strictly superior if $\boldsymbol{\nu}^\dagger >_{(\text{Elementwise})} \boldsymbol{\nu}^\ddagger$.

Definition 2.10: (Optimal Supervision Problem) Given a (non-terminating) plant $G = (Q, \Sigma, \delta, \tilde{\boldsymbol{\Pi}}, \chi, \mathcal{C})$, the problem is to compute a supervisor that disables a subset $\mathcal{D}^* \subseteq \mathcal{C}$, such that $\boldsymbol{\nu}^* \succeq_{(\text{Elementwise})} \boldsymbol{\nu}^\dagger \forall \mathcal{D}^\dagger \subseteq \mathcal{C}$ where $\boldsymbol{\nu}^*$ and $\boldsymbol{\nu}^\dagger$ are the measure vectors of the supervised plants G^* and G^\dagger under \mathcal{D}^* and \mathcal{D}^\dagger , respectively.

Definition 2.11: We note that algorithms reported in [2] compute a lower bound for the critical termination probability for each iteration of such that the disabling/enabling decisions for the terminating plant coincide with the given non-terminating model. We define $\theta_{\min} = \min_k \theta_\star^{[k]}$ where $\theta_\star^{[k]}$ is the termination probability computed in the k^{th} iteration.

Definition 2.12: If G and G^* are the unsupervised and supervised PFSA respectively then we denote the renormalized measure of the terminating plant $G^*(\theta_{\min})$ as $\boldsymbol{\nu}_\#^i : 2^{L(q_i)} \rightarrow [-1, 1]$ (See Definition 2.7). Hence, in vector notation we have: $\boldsymbol{\nu}_\# = \theta_{\min} [\mathbf{I} - (1 - \theta_{\min})\boldsymbol{\Pi}^\#]^{-1} \boldsymbol{\chi}$ where $\boldsymbol{\Pi}^\#$ is the transition probability matrix of the supervised plant G^* , we note that $\boldsymbol{\nu}_\# = \boldsymbol{\nu}^{[K]}$ where K is the total number of iterations required for convergence.

C. Formulating A PFSA Model of Autonomous Navigation

Purely for expositional simplicity, we consider a 2D workspace for the mobile agents. The workspace is first discretized into a finite grid. The underlying theory does not require the grid to be regular; however for the sake of clarity we shall present the formulation under the assumption of a regular grid. The obstacles are represented as blocked-off grid locations in the discretized workspace. We specify a particular location as the fixed goal and consider the problem of finding optimal and feasible paths from arbitrary initial grid locations in the workspace. Figure 1(a) illustrates the basic problem setup. We further assume that at any given time instant the robot occupies

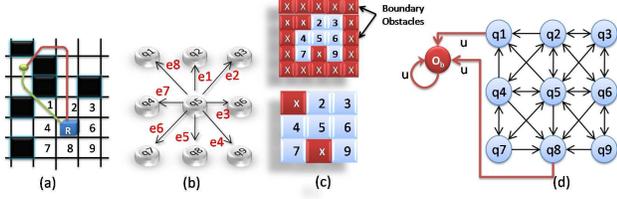


Fig. 1. (a) shows the vehicle (marked "R") with the obstacle positions shown as black squares. The green dot identifies the goal (b) shows the finite state representation of the possible one-step moves from the current position. (d) shows uncontrollable transitions "u" from blocked states

one particular location (*i.e.* a particular square in Figure 1(a)). As shown in Figure 1, the robot has eight possible moves from any interior location. The possible moves are modeled as controllable transitions between grid locations since the robot can "choose" to execute a particular move from the available set. We note that the number of possible moves (8 in this case) depends on the chosen fidelity of discretization of the robot motion and also on the intrinsic vehicle dynamics. The complexity results presented in this paper only assumes that the number of available moves is significantly smaller compared to the number of grid squares, *i.e.*, the discretized position states. Specification of inter-grid transitions in this manner allows us to generate a finite state automaton (FSA) description of the navigation problem. Each square in the discretized workspace is modeled as a FSA state with the controllable transitions defining the corresponding state transition map. The formal description of the model is as follows:

Let $\mathbb{G}_{\text{NAV}} = (Q, \Sigma, \delta, \bar{\Pi}, \chi)$ be a Probabilistic Finite State Automaton (PFSA). In the absence of dynamic uncertainties and state estimation errors, the alphabet contains one uncontrollable event *i.e.* $\Sigma = \Sigma_C \cup \{u\}$ such that Σ_C is the set of controllable events corresponding to the possible moves of the robot. The uncontrollable event u is defined from each of the blocked states and leads to q_{\ominus} which is a deadlock state. All other transitions (*i.e.* moves) are removed from the blocked states. Thus, if a robot moves into a blocked state, it uncontrollably transitions to the deadlock state q_{\ominus} which is physically interpreted to be a collision. We further assume that the robot fails to recover from collisions which is reflected by making q_{\ominus} a deadlock state. We note that q_{\ominus} does not correspond to any physical grid location. The set of blocked grid locations along with the obstacle state q_{\ominus} is denoted as $Q_{\text{OBSTACLE}} \subseteq Q$. Figure 1 illustrates the navigation automaton for a nine state discretized workspace with two blocked squares. Next we augment the navigation FSA by specifying event generation probabilities defined by the map $\bar{\pi} : Q \times \Sigma \rightarrow [0, 1]$ and the characteristic state-weight vector specified as $\chi : Q \rightarrow [-1, 1]$. The characteristic state-weight vector [2] assigns scalar weights to the PFSA states to capture the desirability of ending up in each state.

Definition 2.13: The characteristic weights are specified for the navigation automaton as follows: $\chi(q_i) = \begin{cases} -1 & \text{if } q_i \equiv q_{\ominus} \\ 1 & \text{if } q_i \text{ is the goal} \\ 0 & \text{otherwise} \end{cases}$

In the absence of dynamic constraints and state estimation uncertainties, the robot can "choose" the particular controllable transition to execute at any grid location. Hence we assume that the probability of generation of controllable events is uniform over the set of moves defined at any particular state.

Definition 2.14: Since there is no uncontrollable events defined at any of the unblocked states and no controllable events defined at any of the blocked states, we have the following consistent specification of event generation probabilities: $\forall q_i \in Q, \sigma_j \in \Sigma, \bar{\pi}(q_i, \sigma_j) = \begin{cases} \frac{1}{\text{No. of controllable events at } q_i} & \text{if } \sigma_j \in \Sigma_C \\ 1 & \text{otherwise} \end{cases}$

The boundaries are handled by "surrounding" the workspace with

blocked position states shown as "boundary obstacles" in the upper part of Figure 1(c).

Definition 2.15: The navigation model is defined to have identical connectivity as far as controllable transitions are concerned implying that every controllable transition or move (*i.e.* every element of Σ_C) is defined from each of the unblocked states.

D. Problem Solution as PFSA Optimization

The above-described probabilistic finite state automaton (PFSA) based navigation model allows us to compute optimally feasible path plans via the language-measure-theoretic optimization algorithm [2] described in Section 2. We refer to the language-measure-theoretic algorithm as ν^* in the sequel. For the unsupervised model, the robot is free to execute any one of the defined controllable events from any given grid location (See Figure 1(b)). The optimization algorithm selectively disables controllable transitions to ensure that the formal measure vector of the navigation automaton is element-wise maximized. Physically, this implies that the supervised robot is constrained to choose among only the enabled moves at each state such that the probability of collision is minimized with the probability of reaching the goal simultaneously maximized. *Although ν^* is based on optimization of probabilistic finite state machines, it is shown that an optimal and feasible path plan can be obtained that is executable in a purely deterministic sense.*

Let \mathbb{G}_{NAV} be the unsupervised navigation automaton and $\mathbb{G}_{\text{NAV}}^*$ be the optimally supervised PFSA obtained by ν^* . We note that $\nu_{\#}^i$ is the renormalized measure of the terminating plant $\mathbb{G}_{\text{NAV}}^*(\theta_{\min})$ with substochastic event generation probability matrix $\bar{\Pi}^{\theta_{\min}} = (1 - \theta_{\min})\bar{\Pi}$. Denoting the event generating function (See Definition 2.3) for $\mathbb{G}_{\text{NAV}}^*$ and $\mathbb{G}_{\text{NAV}}^*(\theta_{\min})$ as $\bar{\pi} : Q \times \Sigma \rightarrow Q$ and $\bar{\pi}^{\theta_{\min}} : Q \times \Sigma \rightarrow Q$ respectively, we have $\bar{\pi}^{\theta_{\min}}(q_i, \epsilon) = 1$ and $\forall q_i \in Q, \sigma_j \in \Sigma, \bar{\pi}^{\theta_{\min}}(q_i, \sigma_j) = (1 - \theta_{\min})\bar{\pi}(q_i, \sigma_j)$

Notation 2.1: For notational simplicity, we use $\nu_{\#}^i(L(q_i)) = \nu_{\#}(q_i) = \nu_{\#|i}$ where $\nu_{\#} = \theta_{\min}[I - (1 - \theta_{\min})\bar{\Pi}^{\theta_{\min}}]^{-1}\chi$

Definition 2.16: (ν^* -path): A ν^* -path $\rho(q_i, q_j)$ from state $q_i \in Q$ to state $q_j \in Q$ is defined to be an ordered set of PFSA states $\rho = \{q_{r_1}, \dots, q_{r_M}\}$ with $q_{r_s} \in Q, \forall s \in \{1, \dots, M\}, M \leq \text{CARD}(Q)$ such that

$$q_{r_1} = q_i, q_{r_M} = q_j, \forall i, j \in \{1, \dots, M\}, q_{r_i} \neq q_{r_j} \quad (1a)$$

$$\forall s \in \{1, \dots, M\}, \forall t \leq s, \nu_{\#}(q_{r_s}) \leq \nu_{\#}(q_{r_t}) \quad (1b)$$

We reproduce without proof the following key results pertaining to ν^* -planning as reported in [6].

- 1) There exists an enabled sequence of transitions from state $q_i \in Q \setminus Q_{\text{OBSTACLE}}$ to $q_j \in Q \setminus \{q_{\ominus}\}$ in $\mathbb{G}_{\text{NAV}}^*$ if and only if there exists a ν^* -path $\rho(q_i, q_j)$ in $\mathbb{G}_{\text{NAV}}^*$.
- 2) For the optimally supervised navigation automaton $\mathbb{G}_{\text{NAV}}^*$, we have $\forall q_i \in Q \setminus Q_{\text{OBSTACLE}}, L(q_i) \subseteq \Sigma_C^*$
- 3) (**Obstacle Avoidance:**) There exists no ν^* -path from any unblocked state to any blocked state in the optimally supervised navigation automaton $\mathbb{G}_{\text{NAV}}^*$.
- 4) (**Existence of ν^* -paths:**) There exists a ν^* -path $\rho(q_i, q_{\text{GOAL}})$ from any state $q_i \in Q$ to the goal $q_{\text{GOAL}} \in Q$ if and only if $\nu_{\#}(q_i) > 0$.
- 5) (**Absence of Local Maxima:**) If there exists a ν^* -path from $q_i \in Q$ to $q_j \in Q$ and a ν^* -path from q_i to q_{GOAL} then there exists a ν^* -path from q_j to q_{GOAL} , *i.e.*, $\forall q_i, q_j \in Q \left(\exists \rho_1(q_i, q_{\text{GOAL}}) \wedge \exists \rho_2(q_i, q_j) \Rightarrow \exists \rho(q_j, q_{\text{GOAL}}) \right)$

3. OBSERVATION-DRIVEN IDENTIFICATION OF OBSTACLE DYNAMICS

The ν^* algorithm assumes a known obstacle map to compute the optimal plan. In this paper, we investigate scenarios where the obstacle map varies with time and solve the problem under the assumption that there exists a finite state symbolic representation

of the obstacle dynamics. This requires that the number of possible configurations for the obstacle locations be finitely bounded. Also, if each configuration is mapped to an unique abstract symbol, then the sequence obtained from a discrete observation sequence of the dynamic obstacles, must have bounded memory.

A. Data-driven PFSA Construction Algorithms

The Data-driven PFSA construction approach is built upon the concepts of Symbolic Dynamics [7], Probabilistic Finite State Automata [8], [9] and Pattern Recognition [10] as means to capture the information on fast scale dynamical behavior in terms of symbol sequences. Partitioning the space of process dynamics yields an alphabet which is subsequently used to obtain symbol sequences from time-series data. Among several papered algorithms [9], the following are relevant for this paper:

- 1) **Causal-State Splitting Reconstruction: CSSR** [11] constructs PFSA which belong to the class of sofic shifts [7] and has an a priori unknown structure; it yields optimal pattern discovery in the sense of mutual information [12].
- 2) **Symbolic Dynamic Filtering: SDF** [13] constructs PFSA which belong to the more restricted class of shifts of finite type [7] and has an a priori known structure that can be freely chosen.

Both algorithms stated above exhibits remarkable insensitivity to spurious noise and exogenous disturbances due to:

- Inherent "coarse graining" [14] of the time series data in the process of space partitioning [15].
- Very small variance in the estimated parameters of probability distribution, which is a consequence of repeated recurrences of paths in the graph of the finite state machine with a relatively small number of states and a very large number of sample points [1].

It is important to mention here that conventional *HMM*-based approaches could be sensitive to variations in initial conditions especially if the resulting model is initially off-phase with the current behavior. Also, in Dynamic Bayesian Networks [16], the state machines are hand-coded and the probabilities are estimated with Expectation-Maximization algorithms that are computationally much slower than either *SDF* or *CSSR* [9].

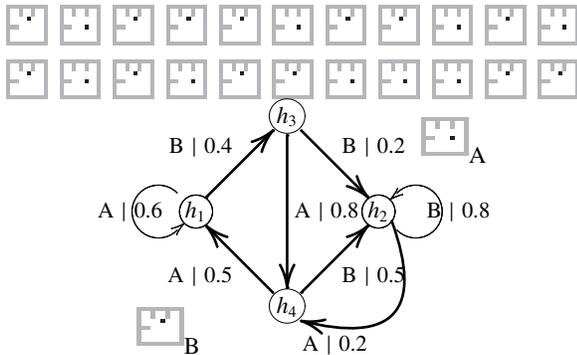


Fig. 2. Symbolic Obstacle dynamics model generation: The left plat illustrates the first few observed obstacle maps with the black spot denoting the moving obstacle and gray denoting static obstacles. White denotes open space. The plate on the right illustrates the constructed \mathbb{H}_{Obs} which is shown to have 4 states and a two event alphabet $\{A, B\}$, where the maps corresponding to A and B are also shown. Note $\mathcal{M}(h_1) = \mathcal{M}(h_4) = \{B\}$ and $\mathcal{M}(h_2) = \mathcal{M}(h_3) = \{A\}$

B. Symbolic Identification of Obstacle Dynamics

Let $\{\mathcal{M}_i\}_{i=1, \dots, N}$ denote the discrete observed sequence of obstacle maps with \mathcal{M}_i being the map observed at the i^{th} observation instant. The observed sequence is symbolized by using Algorithm 1. Note that Algorithm 1 generates both the symbolic sequence $\omega_{\mathcal{M}}$

```

input : Model  $\{\mathcal{M}_i\}_{i=1, \dots, N}$ 
output: Symbolic Sequence  $\omega_{\mathcal{M}}$ , Alphabet  $\Sigma_{\mathcal{M}}$ 
1 begin
2   Let  $\Sigma = \sigma_0$ ;
3   Assign  $\mathcal{M}_1 \mapsto \sigma_0$ ;
4   Set  $k = 1$ ;
5   for  $i = 1 : N$  do
6     for  $r = 1 : k$  do
7       if  $d(\mathcal{M}_i, \mathcal{M}_r) < T$  then
8         /*  $d$ : appropriate metric,  $T$ : threshold */
9         Assign  $\mathcal{M}_i \mapsto \sigma$  where  $\mathcal{M}_r \mapsto \sigma$ ;
10        else
11           $\Sigma_{\mathcal{M}} = \Sigma_{\mathcal{M}} \cup \{\sigma_{k+1}\}$ ; /* Augment alphabet */
12           $k = k + 1$ ;
13          Assign  $\mathcal{M}_i \mapsto \sigma_k$ ;
14        end if
15      end for
16    end for

```

Algorithm 1: Symbolic Observation Sequence Generation

and the appropriate alphabet $\Sigma_{\mathcal{M}}$. The size of the alphabet is dependent on the value of the classification threshold T (See Line 7 of Algorithm 1). Thus, if the exact configuration space is infinite, we obtain a finite state approximation by using a non-zero threshold over an appropriate metric d . The exact metric is unimportant; if the obstacle maps are represented as matrices, then d can be chosen to be the infinity norm of the matrix difference. The generated symbolic sequence can now be compressed to a PFSA. Note the assumption that $N < \infty$ guarantees that one can always construct a finite state model from $\omega_{\mathcal{M}}$; the trick, however, is to construct one that has a small number of states and yet captures the relevant obstacle dynamics. Here we assume that the sequence length N is large enough for convergence of the PFSA construction algorithms. In the sequel, we will assume that, inspite of the suboptimal output, SDF is used for this PFSA construction. This choice simplifies the approach to a considerable degree, as will be elucidated in the sequel.

Definition 3.1: The probabilistic obstacle automaton $\mathbb{H}_{\text{Obs}} = (\mathcal{Q}_{\mathcal{M}}, \Sigma_{\mathcal{M}}, \delta_{\mathcal{M}}, \bar{\Pi}_{\mathcal{M}})$ is defined to be the PFSA obtained via Symbolic Dynamic Filtering (SDF) on the symbolic observation sequence $\omega_{\mathcal{M}}$. As before, $\mathcal{Q}_{\mathcal{M}}$ is the set of states, $\Sigma_{\mathcal{M}}$ is the event alphabet, $\delta_{\mathcal{M}} : \mathcal{Q}_{\mathcal{M}} \times \Sigma_{\mathcal{M}} \rightarrow \mathcal{Q}_{\mathcal{M}}$ is the transition map and $\bar{\Pi}_{\mathcal{M}}$ is the event generation probability matrix.

Lemma 3.1: Given the comparison metric $d(\cdot, \cdot)$ and the classification threshold T , each state $q_i \in \mathcal{Q}_{\mathcal{M}}$ of \mathbb{H}_{Obs} can be uniquely associated with a set of observed obstacle maps $\mathcal{M}(q_i)$ satisfying

$$\forall \mathcal{M}_r, \mathcal{M}_s \in \mathcal{M}(q_i), d(\mathcal{M}_r, \mathcal{M}_s) < T \quad (2)$$

Proof: SDF output has a special structure which is generally denoted as D-Markov PFSA models [13] or equivalently as Probabilistic Suffix Automata [9]. PFSA states are right invariant equivalence classes [17] of symbolic sequences; for D-Markov models we have

$$\forall q_i \in \mathcal{Q}_{\mathcal{M}}, \exists s, \tau(q_i) \in \Sigma_{\mathcal{M}}^*, \text{ s.t. } q_i \equiv \{\omega \in \Sigma_{\mathcal{M}}^* : \omega = s\tau(q_i)\} \quad (3)$$

where $\tau(q_i)$ is a unique string depending on q_i . It follows that each state q_i in $\mathcal{Q}_{\mathcal{M}}$ can be uniquely associated with a symbol from $\Sigma_{\mathcal{M}}$, namely the terminating symbol for $\tau(q_i)$. The result then follows from noting that each symbol in $\Sigma_{\mathcal{M}}$ corresponds to a set of observed obstacle maps that satisfy Eq.(2) (See Algorithm 1). \square

Each observed obstacle map \mathcal{M}_i is a specification of the obstacle locations at the corresponding observation instant and therefore is a specification of the characteristic weights on the states of the navigation automaton (See Definition 2.13) at that instant. Let for each \mathcal{M}_i , the corresponding characteristic weight vector be denoted as $\chi^{\mathcal{M}_i}$. Once \mathbb{H}_{Obs} is constructed, we are required to assign the state

characteristic weights, which, in light of Lemma 3.1, can be done easily as follows:

Definition 3.2: The characteristic weights on the states of \mathbb{H}_{Obs} is defined as follows:

$$\forall h_i \in \mathcal{Q}_{\mathcal{M}}, \chi(h_i) = \frac{1}{\text{CARD}(\mathcal{M}(h_i))} \sum_{\mathcal{M}_r \in \mathcal{M}(h_i)} \chi^{\mathcal{M}_r} \quad (4)$$

where $\mathcal{M}(h_i)$ is as defined in Eq. (2).

Thus, each $h_i \in \mathcal{Q}_{\mathcal{M}}$ can be associated with a single symbol (Lemma 3.1), and the characteristic of h_i quantifies an average effect of all obstacle maps that are identified to this symbol by Algorithm 1. The procedure is illustrated in Figure 2, where we end up with a 4 state \mathbb{H}_{Obs} and $\mathcal{M}(h_1) = \mathcal{M}(h_4) = \{B\}$ and $\mathcal{M}(h_2) = \mathcal{M}(h_3) = \{A\}$, where A, B are the two distinct obstacle maps observed. Next we integrate the constructed \mathbb{H}_{Obs} and the navigation automaton \mathbb{G}_{Nav} to obtain the overall description of estimated obstacle dynamics and navigation constraints and transitions in the PFSA framework.

4. INTEGRATION OF NAVIGATION AUTOMATON & OBSTACLE AUTOMATON

We note that the event alphabet Σ defined by \mathbb{G}_{Nav} consists of controllable transition (except the collision transition u), while $\Sigma_{\mathcal{M}}$ in \mathbb{H}_{Obs} consists of uncontrollable transitions that capture the obstacle dynamics. Also, the generation probabilities (specified by $\tilde{\Pi}_{\mathcal{M}}$) of the events in \mathbb{H}_{Obs} must be preserved in the integrated model, implying that a simple product construction would not suffice. We employ the a input-output automata [18] based approach to define the syntactics of the integration. Namely, we define the states of the integrated model to be of the type (q, h) , where the first coordinate specifies the current state in the underlying navigation model. The second coordinate relates to the state of the obstacle dynamics. However, we need two distinct states $(q, h), (q, h^\circ)$ for the same underlying navigation state q and the same state h of the obstacle dynamics to capture the required segregation of navigation events and obstacle-dynamics events. Thus for the state set \mathcal{Q}^\otimes of the integrated model, we have $\text{CARD}(\mathcal{Q}^\otimes) = 2\text{CARD}(\mathcal{Q})\text{CARD}(\mathcal{Q}_{\mathcal{M}})$. It is important to note that the obstacle locations in the modified model are specified by the individual obstacle maps which evolve according to \mathbb{H}_{Obs} . Hence, the navigation automaton $\mathbb{G}_{\text{Nav}} = (\mathcal{Q}, \Sigma, \delta, \tilde{\Pi}, \chi, \mathcal{C})$ no longer needs to have the obstacle state q_\ominus . Hence, in the sequel, \mathbb{G}_{Nav} is assumed to represent purely the navigation constraints and the uncontrollable collision event u is eliminated from Σ . Also, note that this implies every event in Σ is now controllable. The formal construction is presented next:

Definition 4.1: The integrated model denoted as $\mathbb{G}_{\text{Nav}} \otimes \mathbb{H}_{\text{Obs}} \triangleq (\mathcal{Q}^\otimes, \Sigma^\otimes, \delta^\otimes, \tilde{\Pi}^\otimes, \chi^\otimes, \mathcal{C}^\otimes)$ is constructed from the navigation automaton \mathbb{G}_{Nav} and the obstacle automaton \mathbb{H}_{Obs} by the following relations. Note that $\mathcal{Q}_{\mathcal{M}}^\circ$ is an isomorphic copy of $\mathcal{Q}_{\mathcal{M}}$ and h_j° is the image of h_j in the isomorphic copy.

- (1) $\mathcal{Q}^\otimes = (\mathcal{Q} \times \mathcal{Q}_{\mathcal{M}}) \cup (\mathcal{Q} \times \mathcal{Q}_{\mathcal{M}}^\circ)$
- (2) $\Sigma^\otimes = \Sigma \cup \Sigma_{\mathcal{M}}$
- (3) $\forall q_i \in \mathcal{Q}, h_j \in \mathcal{Q}_{\mathcal{M}}, \delta^\otimes((q_i, h_j), \sigma) = \begin{cases} (\delta(q_i, \sigma), h_j^\circ), & \text{if } \sigma \in \Sigma \\ \text{Undefnd.} & \text{, otherwise} \end{cases}$
- (4) $\forall q_i \in \mathcal{Q}, h_j^\circ \in \mathcal{Q}_{\mathcal{M}}^\circ,$
 $\delta^\otimes((q_i, h_j^\circ), \sigma) = \begin{cases} (q_i, \delta_{\mathcal{M}}(h_j, \sigma)) & \text{, if } \sigma \in \Sigma_{\mathcal{M}} \\ \text{Undefnd.} & \text{, otherwise} \end{cases}$
- (5) $\forall q_i \in \mathcal{Q}, h_j \in \mathcal{Q}_{\mathcal{M}}, \tilde{\Pi}^\otimes((q_i, h_j), \sigma) = \begin{cases} \tilde{\Pi}(q_i, \sigma) & \text{, if } \sigma \in \Sigma \\ 0 & \text{, otherwise} \end{cases}$
- (6) $\forall q_i \in \mathcal{Q}, h_j^\circ \in \mathcal{Q}_{\mathcal{M}}^\circ, \tilde{\Pi}^\otimes((q_i, h_j^\circ), \sigma) = \begin{cases} \tilde{\Pi}_{\mathcal{M}}(h_j, \sigma) & \text{, if } \sigma \in \Sigma_{\mathcal{M}} \\ 0 & \text{, otherwise} \end{cases}$
- (7) $\forall q_i \in \mathcal{Q}, h_j \in \mathcal{Q}_{\mathcal{M}}, \chi^\otimes((q_i, h_j)) = 0$
- (8) $\forall q_i \in \mathcal{Q}, h_j^\circ \in \mathcal{Q}_{\mathcal{M}}^\circ, \chi^\otimes((q_i, h_j^\circ)) = \xi(\chi_{\mathcal{M}}(h_j), \chi(q_i))$

where $\xi: [0, 1]^2 \rightarrow [0, 1]$ is defined as $\xi(a, b) = \begin{cases} a & \text{, if } b = 0 \\ b & \text{, otherwise} \end{cases}$

- (9) $\sigma \in \Sigma$ are controllable; $\sigma \in \Sigma_{\mathcal{M}}$ are uncontrollable

It follows that $\chi^\otimes(q, h^\circ)$ is -1 if the location denoted by q is blocked in the obstacle map currently active, and $\chi^\otimes(q, h^\circ)$ is 1 , i.e., (q, h°) is a goal if q is the goal in the underlying navigation automaton and it is not currently blocked. It is interesting to note that in the trivial case of $\text{CARD}(\mathcal{Q}_{\mathcal{M}}) = 1$, i.e., where the obstacles are static, we obtain twice the number of states as compared to \mathbb{G}_{Nav} . However, it can be easily shown, that the integrated model obtained in this special case is in fact a non-minimal realization [17] of \mathbb{G}_{Nav} .

A. v^* -Optimization of the Integrated PFSA

The integrated PFSA model $\mathbb{G}_{\text{Nav}} \otimes \mathbb{H}_{\text{Obs}}$ can be optimized to obtain optimal plans in a straightforward manner by using the v^* planning algorithm. The semantics of the plan plan execution is however, somewhat different. Algorithm 2 summarizes the overall approach.

We note that the computed optimal measure vector $v_\#$ induces

```

input : Observation sequence  $\{\mathcal{M}_i\}_{i=1, \dots, N}$ 
output: Optimal plan and execution
1 begin
2   Compute  $\mathbb{H}_{\text{Obs}}, \mathbb{G}_{\text{Nav}}, \mathbb{G}_{\text{Nav}} \otimes \mathbb{H}_{\text{Obs}}$ ;
3    $v^*$ -Optimize  $\mathbb{G}_{\text{Nav}} \otimes \mathbb{H}_{\text{Obs}}$  and obtain  $v_\#$ ;
4   Set initial state  $(q, h)$ ;
5   while GOAL_NOT_REACHED do
6     Compute  $W = \{\text{set of reachable states from } (q, h)\}$ ;
7     Compute  $(q_{\text{next}}, h^\circ) = \underset{w \in W}{\text{argmax}} v_\#(w \in W)$ ;
8     Move to  $(q_{\text{next}}, h^\circ)$ ; /* Controllable Move */
9     Set  $q = q_{\text{next}}$ ;
10    Estimate current  $\mathbb{H}_{\text{Obs}}$  state  $h_{\text{next}}$ ;
11    Move to  $(q, h_{\text{next}})$ ; /* Uncontrollable Move */
12    set  $h = h_{\text{next}}$ ;
13
14 end

```

Algorithm 2: Plan computation & Execution

a gradient on the state space \mathcal{Q}^\otimes , which the robot can follow to reach the goal. Thus the computed plan is not probabilistic; given the current state, the plan deterministically specifies the next move. In the case of static obstacles [6], it was shown that the planned path never collides. However, in the generalized situation, this is not guaranteed deterministically. Due to the symbolic compression of the obstacle dynamics, there is always a small probability that the obstacle appears where it is not generally expected, resulting in a small probability of collision. The key point to be noted in Algorithm 2, is the alternating sequence of controllable and uncontrollable moves. The robot can decide which controllable move to execute; but not how the obstacle map will evolve. However, incorporation of the symbolic dynamic model of obstacle evolution in the planning phase implies that the robot can decide upon the best controllable move at any point. Also, there is the implicit assumption that the time scales of the obstacle dynamics is slower compared to the move execution time scale for the robot. This is critical, since, if the obstacle map switches faster than the robot can move, then the state estimates will be erroneous leading to incorrect execution. We have the following result:

Proposition 4.1: The proposed planning and execution approach solves the following optimization problem: Maximize $\Psi = (\wp_{\text{GOAL}} - \wp_{\text{OBSTACLE}})$ under the navigation constraints and observed history of obstacle dynamics, where \wp_{GOAL} is the probability of reaching the goal and \wp_{OBS} is the probability of hitting an obstacle.

Proof: We recall that the language-measure-theoretic optimization of PFSA accomplishes the maximization of $\wp^T \chi$ where \wp is the stationary probability vector on the automaton states [2]. Since $\chi(\text{GOAL}) = 1$ and $\chi(\text{OBSTACLE}) = -1$ and all other states have zero characteristic, it follows that Ψ gets maximized in the optimization. \square

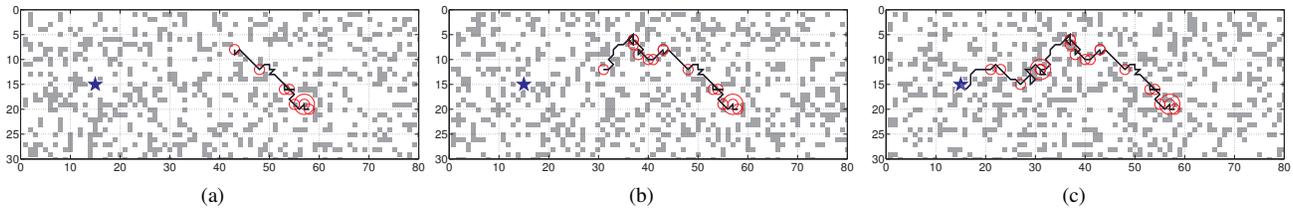


Fig. 3. Simulated execution on a 30×80 dynamic maze. The plates (a) - (c) are snapshots at different times. Note that the obstacle map is different in the three cases. The red circles indicate locations where the simulated robot decided to wait for one or more ticks as the obstacles continued to evolve

B. Simulation Examples on Dynamic Mazes

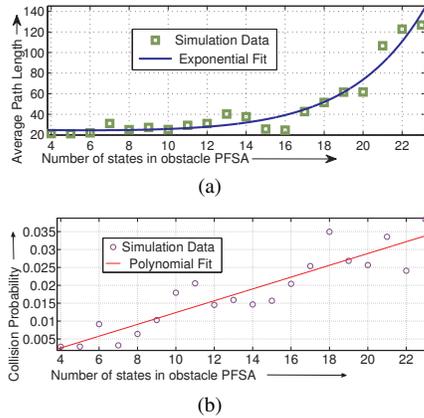


Fig. 4. Effect of increasing complexity of obstacle dynamics: (a) shows the exponential increase in average path lengths to reach goal, (b) shows linear increase of collision probability

The proposed algorithm is validated on simulated models. Figure 3 illustrates one such simulation run on a 30×80 grid. Plates (a) - (c) capture snapshots of the plan execution at different times. Note that the obstacle map (gray for blocked locations; white for open spaces) is different in the three plates. Also note, that the executed path is significantly rough. However, the optimization result of Proposition 4.1 guarantees that the executed path indeed is optimal. The obstacle automaton in this case has 10 events and 100 states. The red circles in Figure 3 indicates where the simulated robot waited for one or more ticks, *i.e.*, decided not to execute any controllable move. Note that while it is futile to “wait” at any location for a static obstacle map; it is often the optimal course of action when the obstacles are dynamic.

We simulate the effect of increasing complexity of obstacle dynamics by considering \mathbb{H}_{Obs} with larger and larger number of states. The results are shown in Figure 4(a)-(b). Note that while the collision probability increases only linearly with increased complexity of the obstacle dynamics, the average path length required to reach the goal increases exponentially.

5. SUMMARY & FUTURE RESEARCH

The ν^* planning algorithm is augmented to handle dynamic obstacles. The obstacle dynamics is first encoded symbolically from a finite but sufficiently long observation sequence as a probabilistic finite state machine. This symbolic dynamic model of the obstacles is then integrated with the navigation automaton and is subsequently optimized by the language-measure theoretic approach. It is shown that the proposed approach maximizes the difference in the probability of reaching the goal and the probability of hitting an obstacle and is optimal in that sense. Future work will extend the language-measure

theoretic planning algorithm to address the following problems:

- 1) **Multi-robot coordinated planning:** Future work will address multi-robot scenarios, with each robot treating the remaining group as moving obstacles.
- 2) **Hierarchical implementation to handle very large workspaces:** Large workspaces can be solved more efficiently if planning is done when needed rather than solving the whole problem at once.
- 3) **Handling partially observable dynamic events:** Physical errors and onboard sensor failures may need to be modeled as unobservable transitions and will be addressed in future publications.

REFERENCES

- [1] I. Chattopadhyay and A. Ray, “Renormalized measure of regular languages,” *Int. J. Control*, vol. 79, no. 9, pp. 1107–1117, 2006.
- [2] I. Chattopadhyay and A. Ray, “Language-measure-theoretic optimal control of probabilistic finite-state systems,” *Int. J. Control*, August, 2007.
- [3] J. Barraquand, B. Langlois, and J.-C. Latombe, *Robot motion planning with many degrees of freedom and dynamic constraints*. Cambridge, MA, USA: MIT Press, 1990.
- [4] V. Garg, “An algebraic approach to modeling probabilistic discrete event systems,” *Proceedings of 1992 IEEE Conference on Decision and Control*, pp. 2348–2353, Tucson, AZ, December 1992.
- [5] W. Rudin, *Real and Complex Analysis, 3rd ed.* McGraw Hill, New York, 1988.
- [6] I. Chattopadhyay, G. Mallapragada, and A. Ray, “ ν^* : a robot path planning algorithm based on renormalized measure of probabilistic regular languages,” *International Journal of Control*, in press.
- [7] D. Lind and M. Marcus, *An Introduction to Symbolic Dynamics and Coding*. Cambridge University Press, United Kingdom, 1995.
- [8] J. E. Hopcroft, R. Motwani, and J. D. Ullman, *Introduction to Automata Theory, Languages, and Computation, 2nd ed.* Addison-Wesley, 2001.
- [9] K. Murphy, “Passively learning finite automata,” 1996.
- [10] R. Duda, P. Hart, and D. Stork, *Pattern Classification*. John Wiley & Sons Inc., 2001.
- [11] C. R. Shalizi, K. L. Shalizi, and J. P. Crutchfield, “An algorithm for pattern discovery in time series,” *Technical Report, Santa Fe Institute*, October 2002.
- [12] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. John Wiley, New York, 1991.
- [13] A. Ray, “Symbolic dynamic analysis of complex systems for anomaly detection,” *Signal Processing*, vol. 84, no. 7, pp. 1115–1130, 2004.
- [14] C. Beck and F. Schlogl, *Thermodynamics of chaotic systems: an introduction*. Cambridge University Press, United Kingdom, 1993.
- [15] V. Rajagopalan and A. Ray, “Symbolic time series analysis via wavelet-based partitioning,” *Signal Processing*, vol. 86, no. 11, pp. 3309–3320, 2006.
- [16] E. Denis, N.; Jones, “Spatio-temporal pattern detection using dynamic bayesian networks,” *In Proceedings of 42nd IEEE Conference on Decision and Control*, vol. 5, pp. 4533–4538, 9-12 Dec. 2003.
- [17] I. Chattopadhyay and A. Ray, “Structural transformations of probabilistic finite state machines,” *International Journal of Control*, vol. 81, pp. 820–835, May 2008.
- [18] N. Lynch and M. Tuttle, “An introduction to input/output automata,” Technical Memo MIT/LCS/TM-373, Massachusetts Institute of Technology, November 1988.