

Decision-Theoretic Coordination and Control for Active Multi-Camera Surveillance in Uncertain, Partially Observable Environments

Prabhu Natarajan, Trong Nghia Hoang, Kian Hsiang Low, and Mohan Kankanhalli
Department of Computer Science, National University of Singapore, Republic of Singapore
{prabhu, nghiaht, lowkh, mohan}@comp.nus.edu.sg

Abstract—A central problem of surveillance is to monitor multiple targets moving in a large-scale, obstacle-ridden environment with occlusions. This paper presents a novel principled Partially Observable Markov Decision Process-based approach to coordinating and controlling a network of active cameras for tracking and observing multiple mobile targets at high resolution in such surveillance environments. Our proposed approach is capable of (a) maintaining a belief over the targets’ states (i.e., locations, directions, and velocities) to track them, even when they may not be observed directly by the cameras at all times, (b) coordinating the cameras’ actions to simultaneously improve the belief over the targets’ states and maximize the expected number of targets observed with a guaranteed resolution, and (c) exploiting the inherent structure of our surveillance problem to improve its scalability (i.e., linear time) in the number of targets to be observed. Quantitative comparisons with state-of-the-art multi-camera coordination and control techniques show that our approach can achieve higher surveillance quality in real time. The practical feasibility of our approach is also demonstrated using real AXIS 214 PTZ cameras.

I. INTRODUCTION

Monitoring, tracking, and observing multiple mobile targets in a large-scale, obstacle-ridden environment (e.g., airport terminals, railway stations, bus depots, shopping malls, etc.) is a central problem of surveillance. It is often necessary to acquire high-resolution videos/images of these targets. Traditionally, such high-quality surveillance is achieved by placing a large number of static cameras to completely cover the large environment. This is impractical in terms of equipment, installation, and maintenance costs. Recent works ([6], [9], [14]) have employed a heterogeneous network of wide-view static camera(s) to detect and track the targets within the environment at low resolution and some active pan/tilt/zoom (PTZ) cameras to be directed and focused on these targets to observe them at high resolution. Such surveillance systems face two serious practical limitations: (a) The obstacles in the environment (e.g., physical structures like walls, pillars, and barriers) are likely to occlude the fields of view (fov) of the static cameras and hence they cannot detect or track the targets that reside within these *occluded regions*. Since the surveillance system is not informed of these targets, the active cameras may not be directed to observe them; and (b) when the targets move further away from the low-resolution static cameras, their measured locations become less accurate regardless of the calibration method. The vision algorithms to

detect and recognize the targets also grow less reliable.

More importantly, the above limitations raise a practical implication affecting real-world multi-camera surveillance in general: The exact locations of the targets may not be observed directly by (or fully observable to) the cameras at all times. Such an environment is said to be *partially observable* [5]. Instead of introducing additional static cameras to resolve this issue of partial observability, we propose an alternative that maintains a probabilistic belief over the targets’ possible locations, directions, and velocities in the environment. Our proposed alternative offers a practical advantage over ([6], [9], [14]) by eliminating the dependence on wide-fov static cameras to track the targets’ locations and enabling the active PTZ cameras to perform dual roles of tracking the targets as well as observing them with high resolution. Hence, we will focus on using (though not limited ourselves to) only active PTZ cameras in this paper.

This paper presents a novel principled decision-theoretic approach to coordinating and controlling a network of active cameras for tracking and observing multiple mobile targets at high resolution in uncertain, partially observable surveillance environments. Our proposed approach stems from framing the surveillance problem formally using a rich class of decision making under uncertainty models called the *Partially Observable Markov Decision Process* (POMDP) (Section III). Specifically, it overcomes the above limitations by (a) modeling a belief over the targets’ states (i.e., locations, directions, and velocities) and updating the belief in a Bayesian paradigm (Section III-D) based on probabilistic models of the targets’ motion (Section III-B) and the active cameras’ observations (Section III-C); (b) coordinating the active cameras’ actions to simultaneously improve the belief over the targets’ states and maximize the expected number of targets observed with a guaranteed pre-defined resolution (Sections III-E and IV); and (c) exploiting the inherent structure of our surveillance problem to improve its scalability (i.e., linear time) in the number of targets to be observed (Section IV). Our POMDP-based approach is empirically evaluated in simulations over various realistic surveillance environments (Section V-A) and with real AXIS 214 PTZ cameras to demonstrate its practical feasibility (Section V-B).

II. RELATED WORKS

As mentioned earlier, existing multi-camera coordination and control techniques have to operate in a fully observable surveillance environment where the locations, directions, and velocities of all the targets can be directly observed/estimated by either using additional low-resolution static cameras and sensors ([2], [6], [9], [14], [15]) or configuring one or more active cameras to zoom out to their wide view ([4], [11], [12], [13]). They use these targets' information to predict their trajectories in order to schedule, coordinate, and control the network of active cameras to focus on and observe these targets at high resolution.

The major drawbacks of these techniques are: (a) They cannot be deployed in real-world surveillance environments with occlusions. In this case, they cannot observe the targets that reside in the occluded regions, hence limiting the active cameras' full surveillance capability. In contrast, our approach does not assume that all targets can be fully observed at every time instance, and hence models a belief of the targets' states to keep track of them when they are not observed by any of the cameras; (b) since the resolution of the wide-view static cameras is low, they often produce inaccurate locations of the targets. This in turn induces errors in targets' directions and velocities which consequently affect the prediction capability of existing surveillance systems. On the other hand, our approach uses only active cameras to observe the targets at high resolution, thus allowing location errors to be kept minimal; and (c) many existing techniques have serious issues of scalability in the number of targets to be observed. Our approach extends our previous work [9] to achieve scalability in partially observable surveillance environments.

III. SYSTEM OVERVIEW AND PROBLEM FORMULATION

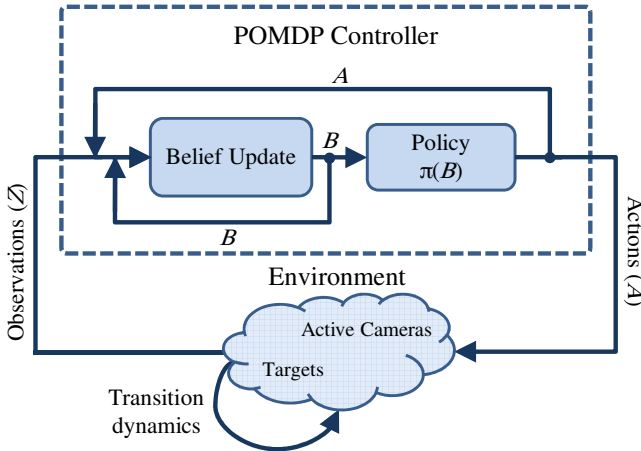


Fig. 1. POMDP controller for coordinating active cameras to perform high-quality surveillance in a partially observable environment.

A POMDP controller models the interaction between the active cameras and the partially observable surveillance environment. In particular, it is responsible for coordinating the cameras' actions to achieve a high-level surveillance goal which can be defined formally using a real-valued objective

function and, in the context of this paper, is to maximize the number of targets observed with a guaranteed resolution (Section III-E). By calibrating the active cameras, they can determine the locations of the targets observed in their fov, which are communicated to the controller. The targets are assumed to be non-evasive and hence their motion cannot be controlled or influenced by the cameras. The targets' correspondences across multiple cameras are resolved by distinct features like color and texture.

Formally, a POMDP controller is defined as a tuple $(S, \mathcal{A}, \mathcal{Z}, T_f, O_f, R)$ consisting of

- a set S of joint states of active cameras and targets in the surveillance environment (Section III-A);
- a set \mathcal{A} of joint actions of active cameras (Section III-A);
- a set \mathcal{Z} of joint observations of the targets taken by the cameras (Section III-A);
- a transition function $T_f : S \times \mathcal{A} \times S \rightarrow [0, 1]$ denoting the probability $P(S'|S, A)$ of going from the current joint state $S \in S$ to the next joint state $S' \in S$ using the joint action $A \in \mathcal{A}$ (Section III-B);
- an observation function $O_f : S \rightarrow [0, 1]$ denoting the probability $P(Z|S)$ of observing the joint observation $Z \in \mathcal{Z}$ given the joint state $S \in S$ (Section III-C); and
- a real-valued objective/reward function $R : S \rightarrow \mathbb{R}$ representing a high-level surveillance goal (Section III-E).

At any given time, the exact state of the environment is not fully observable to the POMDP controller. Instead, it maintains a belief B over the set S of all possible states (Section III-D), that is, $B(S)$ is the probability that the environment is in the state $S \in S$ such that $\sum_{S \in S} B(S) = 1$. As shown in Fig. 1, at every time step, the POMDP controller issues an action $A \in \mathcal{A}$ and makes an observation $Z \in \mathcal{Z}$ from the environment. Based on the action A and observation Z , the prior belief B is updated by Bayes' rule to the posterior belief B' as follows:

$$B'(S') = \eta P(Z|S') \sum_{S \in S} P(S'|S, A) B(S) \quad (1)$$

where $\eta \triangleq 1/P(Z|B, A)$ is a normalizing constant. A policy π for the POMDP controller is defined as a mapping from each belief B to an action A (Section IV). Solving a POMDP involves choosing the optimal policy π^* that maximizes the expected reward for any given belief B :

$$\pi^*(B) = \arg \max_{A \in \mathcal{A}} \sum_{Z' \in \mathcal{Z}} R(B') P(Z'|B, A) .$$

When the number of targets and active cameras increases, the state space and hence the belief space of the POMDP grow exponentially (Section III-A). Therefore, computing the optimal policy incurs exponential time. Fortunately, by exploiting the structure of our surveillance problem (Sections III-B and III-C), the optimal policy for a given belief B can be computed efficiently (Section IV).

A. States, Actions, and Observations

A joint state $S \in S$ of the POMDP controller is defined as a pair of joint states $T \in \mathcal{T}^m$ of m targets and $C \in \mathcal{C}^n$ of

n active cameras where \mathcal{T} and \mathcal{C} denote sets of all possible states of each target and active camera, respectively. That is, $S \triangleq (T, C)$ and $\mathcal{S} = \mathcal{T}^m \times \mathcal{C}^n$. Let $T \triangleq (t_1, t_2, \dots, t_m) \in \mathcal{T}^m$ and $C \triangleq (c_1, c_2, \dots, c_n) \in \mathcal{C}^n$ where $t_k \in \mathcal{T}$ and $c_i \in \mathcal{C}$ denote the corresponding states of target k and camera i . Let $t_k \triangleq (t_{l_k}, t_{d_k}, t_{v_k}) \in \mathcal{T}_l \times \mathcal{T}_d \times \mathcal{T}_v$ where t_{l_k} , t_{d_k} , and t_{v_k} denote target k 's location, direction, and velocity, respectively. That is, $\mathcal{T} = \mathcal{T}_l \times \mathcal{T}_d \times \mathcal{T}_v$.

The state space \mathcal{C} of an active camera is a finite set of discrete pan/tilt/zoom positions. Let $fov(c_i) \subset \mathcal{T}_l$ be a subset of target locations lying within the fov of camera i in its state c_i . The joint fov of all cameras in joint state C is defined as $fov(C) = \bigcup_{i=1}^n fov(c_i)$. The depth of fov of each active camera is limited such that (a) imageries of the targets detected within its fov satisfy a pre-defined resolution, and (b) the observed locations of the targets detected within its fov are of minimal location error. This is done by adjusting the zoom parameter of each camera based on its position.

The joint actions of the POMDP controller are PTZ commands that move the corresponding cameras to their specified states. Let a joint action of the n cameras be denoted by $A \triangleq (a_1, a_2, \dots, a_n) \in \mathcal{A}$ where a_i denotes the PTZ command of camera i .

Let $\mathcal{Z} \triangleq \mathcal{T}_l \cup \{\phi\}$ denote a set of all possible observations of a target comprising the set \mathcal{T}_l of all possible locations of the target in the environment and a null observation ϕ when the target is not observed by any of the cameras. Let an observation of target k be denoted by $z_k \in \mathcal{Z}$ and a joint observation of the m targets be denoted by $Z \triangleq (z_1, z_2, \dots, z_m) \in \mathcal{Z}^m$. That is, $\mathcal{Z} = \mathcal{Z}^m$.

B. Transition Model T_f

By exploiting the following structural assumptions in the state transition dynamics of the surveillance environment:

- camera i 's next state c'_i is conditionally independent of the other $n - 1$ cameras' states and actions and the m targets' states given its current state c_i and action a_i for $i = 1, \dots, n$ and
- target k 's next state t'_k is conditionally independent of the n cameras' states and actions (i.e., target's motion is not affected by the cameras' states and actions) and the other $m - 1$ targets' states (i.e., every target moves independently) given its current state t_k for $k = 1, \dots, m$,

the transition model T_f can be factored into transition models of individual targets and active cameras, hence significantly reducing the time incurred to compute the optimal policy π^* for a given belief B (Section IV). Furthermore, since the modern active cameras are able to move to their specified positions accurately [1], it is practical to assume the transition model of each individual camera to be deterministic and consequently represented by a function τ that moves camera i from its current state c_i to its next state $\tau(c_i, a_i)$ by the action a_i . Then, the transition model of the POMDP controller can be simplified to

$$P(S'|S, A) = \prod_{k=1}^m P(t'_k|t_k) \prod_{i=1}^n \delta_{\tau(c_i, a_i)}(c'_i) \quad (2)$$

where $\delta_x(x')$ is a Kronecker delta function of value 1 if $x' = x$, and 0 otherwise. Details on the derivation of (2) are reported in [9]. The state transition of target k from t_k to t'_k includes stochastic transitions of its location from t_{l_k} to t'_{l_k} , its direction from t_{d_k} to t'_{d_k} , and its velocity from t_{v_k} to t'_{v_k} . So, the transition probability of target k can be factored into transition probabilities of its location, direction, and velocity:

$$P(t'_k|t_k) = P(t'_{l_k}|t_{l_k}, t'_{d_k}, t'_{v_k}) P(t'_{d_k}|t_{d_k}) P(t'_{v_k}|t_{v_k}).$$

The transition probabilities $P(t'_{d_k}|t_{d_k})$ and $P(t'_{v_k}|t_{v_k})$ of the target's direction and velocity are, respectively, modeled as Gaussian distributions $\mathcal{N}(\mu_d, \sigma_d)$ and $\mathcal{N}(\mu_v, \sigma_v)$ with the means μ_d and μ_v being the current direction and velocity of the target, and σ_d and σ_v being the variance parameters which are learned from a dataset of the targets' trajectories in the environment. The transition probability $P(t'_{l_k}|t_{l_k}, t'_{d_k}, t'_{v_k})$ of the target's next location is constructed using the general velocity-direction motion model, as described in [9].

C. Observation Model O_f

Similar to the factorization of the transition model T_f , the observation model O_f can also be factored into observation models of individual targets using the following structural assumption: The observed location $z_k \in \mathcal{Z}$ of target k is conditionally independent of the observed and true states of the other $m - 1$ targets and its true direction $t_{d_k} \in \mathcal{T}_d$ and velocity $t_{v_k} \in \mathcal{T}_v$ given its true location $t_{l_k} \in \mathcal{T}_l$ and the joint state $C \in \mathcal{C}^n$ of the n active cameras for $k = 1, \dots, m$. As a result, the time incurred to compute the optimal policy π^* for a given belief B can be significantly reduced (Section IV). Then, the observation model of the POMDP controller can be simplified to

$$P(Z|S) = \prod_{k=1}^m P(z_k|t_{l_k}, C). \quad (3)$$

The derivation of (3) is reported in Appendix A. The observation probability $P(z_k|t_{l_k}, C)$ of target k depends on whether the target lies within the joint fov of the active cameras. When the target lies within the cameras' joint fov corresponding to their joint state C (i.e., $z_k \neq \phi$), the observation model of target k becomes deterministic:

$$P(z_k|t_{l_k}, C) = \begin{cases} 1 & \text{if } z_k = t_{l_k} \wedge t_{l_k} \in fov(C), \\ 0 & \text{otherwise.} \end{cases}$$

On the other hand, when target k does not lie within the joint fov of the active cameras corresponding to their joint state C , the observation probability of target k is uniformly distributed over the locations not covered by the joint fov (i.e., $\overline{fov(C)}$):

$$P(z_k = \phi|t_{l_k}, C) = \begin{cases} \frac{1}{|\overline{fov(C)}|} & \text{if } t_{l_k} \notin fov(C), \\ 0 & \text{otherwise.} \end{cases}$$

D. Bayesian Belief Update

By making use of independence assumptions similar to that in the transition model (Section III-B), a belief B can be factored into beliefs of individual targets and cameras:

$$\begin{aligned} B(S) &= P((T, C)) = P(T)P(C) \\ &= \prod_{k=1}^m P(t_k) \prod_{i=1}^n P(c_i) = \prod_{k=1}^m b_k(t_k) \prod_{i=1}^n \delta_{\hat{c}_i}(c_i) \end{aligned} \quad (4)$$

where b_k denotes a belief over the set \mathcal{T} of all possible states of target k (i.e., $b_k(t_k)$ is the probability that target k is in state t_k) and \hat{c}_i is the current state of camera i that, unlike a target's state, is fully observable to the POMDP controller since its position can be directly read from its port. Hence, the probability $P(c_i)$ of a state c_i of camera i can be represented by a Kronecker delta $\delta_{\hat{c}_i}(c_i)$ and the last equality in (4) follows.

The POMDP controller issues a joint action A to move each camera i from current state \hat{c}_i to next state \hat{c}'_i , receives an observation z_k of each target k , and then updates the prior belief B to the posterior belief B' using Bayes' rule (1). Similar to the factorization of the prior belief B above, the posterior belief B' can also be factored into posterior beliefs of individual targets and cameras:

$$B'(S') = \prod_{k=1}^m b'_k(t'_k) \prod_{i=1}^n \delta_{\hat{c}'_i}(c'_i) \quad (5)$$

where the posterior belief b'_k of target k is defined as

$$b'_k(t'_k) \triangleq \eta_k P(z_k | t'_k, C') \sum_{t_k \in \mathcal{T}} P(t'_k | t_k) b_k(t_k), \quad (6)$$

$C' \triangleq (c'_1, \dots, c'_n)$, and $\eta_k \triangleq 1/P(z_k | b_k, C')$ is a normalizing constant. The derivation of (5) is reported in Appendix B.

E. Objective/Reward Function R

The goal of the surveillance system is to maximize the number of targets observed with a guaranteed resolution. This can be achieved by defining a reward function that measures the total number of targets lying within the joint fov of the active cameras corresponding to their joint state C :

$$R(S) = R((T, C)) \triangleq \sum_{k=1}^m \tilde{R}(t_k, C) \quad (7)$$

where

$$\tilde{R}(t_k, C) \triangleq \begin{cases} 1 & \text{if } t_{l_k} \in \text{fov}(C), \\ 0 & \text{otherwise.} \end{cases}$$

Since the exact locations of the targets may not be fully observable to the cameras at all times, the POMDP controller has to track the joint belief B of the targets and consider the *expected* reward with respect to this belief instead:

$$R(B) \triangleq \sum_{S \in \mathcal{S}} R(S) B(S) = \sum_{k=1}^m \tilde{R}(b_k, \hat{C}) \quad (8)$$

where $\hat{C} \triangleq (\hat{c}_1, \dots, \hat{c}_n)$ and

$$\tilde{R}(b_k, C) \triangleq \sum_{t_k \in \mathcal{T}} \tilde{R}(t_k, C) b_k(t_k). \quad (9)$$

The derivation of (8) is reported in Appendix C.

IV. POLICY COMPUTATION

Recall that a policy π for the POMDP controller is a mapping from each belief B to a joint action $A \in \mathcal{A}$ of the n cameras. At every time step, the POMDP controller determines the optimal policy π^* for the belief B such that the expected number of observed targets in the next time step is maximized. Since the observations of the m targets taken by the cameras in the next time step are not known to the POMDP controller, it has to consider the *expected* reward with respect to these future observations. Then, the optimal policy π^* for a given belief B becomes

$$\pi^*(B) = \arg \max_{A \in \mathcal{A}} V(B, A) \quad (10)$$

where

$$V(B, A) = \sum_{Z \in \mathcal{Z}} R(B') P(Z | B, A). \quad (11)$$

Computing the policy π^* (10) for a given belief B incurs $\mathcal{O}(|\mathcal{A}| |\mathcal{Z}|^m |\mathcal{T}|)$ time which is exponential in the number m of targets. Fortunately, by exploiting simplified transition and observation models due to conditional independence assumptions (i.e., (2) and (3)), this computational cost can be significantly reduced. In particular, it is derived in Appendix D that the value function $V(B, A)$ of m targets can be simplified to comprise a sum of value function $\tilde{V}(b_k, C')$ of individual target k for $k = 1, \dots, m$:

$$V(B, A) = \sum_{k=1}^m \tilde{V}(b_k, (\tau(\hat{c}_1, a_1), \dots, \tau(\hat{c}_n, a_n))) \quad (12)$$

where

$$\tilde{V}(b_k, C') \triangleq \sum_{z_k \in \text{fov}(C')} \sum_{t'_k \in \mathcal{T}} \tilde{R}(t'_k, C') \hat{b}'_k(t'_k) \quad (13)$$

and \hat{b}'_k is the unnormalized belief of target k (i.e., $\hat{b}'_k(t'_k) = b'_k(t'_k)/\eta_k$). Using (12) and (13), we obtain the following result:

Theorem 1: If (2) and (3) hold, then computing policy π^ (10) for a given belief B incurs $\mathcal{O}(|\mathcal{A}| |\mathcal{Z}| |\mathcal{T}| m)$ time.*

Computing the value function $\tilde{V}(b_k, C')$ (13) for a single target k incurs $\mathcal{O}(|\mathcal{Z}| |\mathcal{T}|)$ time. For m targets, the value function $V(B, A)$ (12) therefore incurs $\mathcal{O}(|\mathcal{Z}| |\mathcal{T}| m)$ time. Finally, computing the optimal policy π^* (10) for a given belief B incurs $\mathcal{O}(|\mathcal{A}| |\mathcal{Z}| |\mathcal{T}| m)$ time which is linear in number m of targets.




Fig. 2. Setups of simulated surveillance environments: (a) Hall ($|\mathcal{T}_l| = 20 \times 10$ target locations), (b) corridor ($|\mathcal{T}_l| = 40 \times 5$ target locations), and (c) junction ($|\mathcal{T}_l| = 168$ target locations) with its corresponding real-world map shown in (d).

V. EXPERIMENTS AND DISCUSSION

This section evaluates the performance of our proposed POMDP controller in simulations over various realistic surveillance environments using Player/Stage simulator [3] and with real AXIS 214 PTZ cameras to show its practical feasibility in real-world surveillance. Our POMDP-based approach (denoted by P in Fig. 3) that uses only active PTZ cameras is compared against the following state-of-the-art multi-camera coordination and control techniques under partially observable surveillance environments:

- *MDP with only PTZ cameras (MP)*: This approach uses a Markov Decision Process (MDP) controller [9] to coordinate and control the active cameras. There is no static camera to directly observe the targets' locations. Hence, they are observed only from the active cameras' fov;
- *MDP with static and PTZ cameras (MSP)*: This approach uses the MDP controller [9] to coordinate and control the active cameras that are supported by wide-view static cameras. A Gaussian noise is added to the location of each target observed by the static camera such that the Gaussian variance increases with greater distance of the target from the static camera;
- *Systematic Approach (Sys)*: The active cameras are panned systematically to each of its states in a round robin fashion for every time step; and
- *Static Approach (Stat)*: The active cameras are fixed at a particular state such that they observe the maximum area of the environment.

The performance metric used to evaluate the above approaches is given by

$$PercentObs = \frac{100}{\tau M_{tot}} \sum_{i=1}^{\tau} M_{obs}^i$$

where τ (i.e., set to 100 in simulations) is the total number of time steps taken in our experiments, M_{obs}^i is the total number of targets observed by the active cameras at a given time step i , and M_{tot} is the total number of targets present in the environment. That is, the *PercentObs* metric averages the percentage of targets being observed by the active cameras over the entire duration of τ time steps.

A. Simulated Experiments

Fig. 2 shows three different setups of simulated surveillance environments: (a) hall, (b) corridor, and (c) junction. The




Fig. 3. Graphs of *PercentObs* vs. number m of targets with $n = 4$ active cameras for the (a) corridor, (b) hall and (c) junction setups.

junction setup simulates a surveillance environment within our university campus which consists of obstacles (black shades in Fig. 2c) like buildings and walls. In order to introduce more occlusions in the environment, we have added a virtual pillar in the center of the junction setup (Fig. 2c). The active cameras are simulated in Player/Stage simulator by configuring the states of the cameras across various pan angles, as discussed in Section III-A. There are $n = 4$ active cameras with $|\mathcal{C}| = 3$ states each. The targets' trajectories are generated in the simulator based on the velocity-direction motion model (Section III-B) which resembles real human motion in a surveillance environment. Every target can move in one of the 8 possible discretized directions $\mathcal{T}_d = \{0^\circ, 45^\circ, \dots, -90^\circ, -45^\circ\}$ with an assumed velocity of 1.5 cells per time step. The transition model, observation model, and reward function for a single target are computed and stored offline for the above setups.

Fig. 3 shows the comparison of performance of different approaches for up to $m = 20$ targets for all three setups. It can be observed that our POMDP controller outperforms the other evaluated approaches in all three setups. The detailed observations from the experiments are as follows:

Our POMDP controller performs better than the *MP* approach because (a) when the targets leave the fov of any of the cameras and enter an occluded region, the active cameras in the *MP* approach have no idea where the targets will be moving to in the next few time steps, and (b) when the targets enter the fov of any of the active cameras from an occluded region, the directions of the targets are wrongly interpreted by the MDP controller. This is a serious limitation of the *MP* approach, i.e., there is no way of knowing the direction of the targets when they are in an occluded region. In contrast, the Bayesian belief update process in our POMDP controller helps to trace the locations and directions of the targets, even

TABLE I
PERFORMANCE FOR REAL CAMERA EXPERIMENTS.

No. of targets m	1	2	3	4	5
<i>PercentObs</i>	98.2	96.6	93.3	91.5	87

when they are not observed in any of the cameras. Hence, the cameras are controlled based on the belief of the targets.

Our POMDP controller performs better than the *MSP* approach because when the static cameras observe the targets that are far away, they obtain noisy locations of the targets. This in turn induces the errors in the direction and velocity of the targets. Hence, when the noisy targets' information is used in the MDP controller, it predicts the expected locations of the targets poorly, which consequently affects the performance of the *MSP* approach. In contrast, for our POMDP-based approach, the targets' locations are observed by high-resolution active cameras whose calibration error is bounded by limiting the depth of its fov (Section III-A). Since the observations (i.e., locations of the targets) for POMDP are more accurate than in the *MSP* approach, the predictions of the targets' locations and directions through the Bayesian belief update process are also more accurate.

Our POMDP controller performs much better than the *Sys* and *Stat* baseline approaches because, for our approach, the active cameras are controlled based on the targets' predicted motion and observations taken by the active cameras. But, for the *Sys* approach, the cameras are panned without accounting for the targets' information such as locations and direction while, for the *Stat* approach, every camera is fixed in one of the states.

To summarize, our POMDP-based approach performs better than the *MP* approach due to its ability to keep track of the targets' locations and directions through its Bayesian belief update process. It outperforms the *MSP* approach because the observations (i.e., target's location) taken by the active cameras in our POMDP controller are more accurate as compared to the noisy observations taken by the static cameras in the *MSP* approach. Lastly, the *Sys* and *Stat* approaches suffer from performance degradation because the cameras are controlled independently of the targets' information.

B. Real Camera Experiments

The feasibility of our POMDP controller is tested using real AXIS 214 PTZ cameras to monitor Lego robots (targets) over the environment of size $|\mathcal{T}_i| = 10 \times 8$ grid cells. We have $n = 3$ PTZ cameras, each of which has $|\mathcal{C}| = 3$ states. These cameras are calibrated in each of its state and the depth of the fov of these cameras are determined empirically for each of them. The Lego robots are programmed to move based on the velocity-direction motion model. Table I shows the performance of our approach in real camera experiments. Due to space limitation, we showcase our detailed results of real camera experiments in a demo video¹.

The limitations of our POMDP-based approach are as follows: (a) it scales well only in number of targets and needs improvement in scalability in the number of cameras; and (b) it

works well only if the underlying computer vision algorithms for target detection and recognition perform accurately. For our future work, we would like to extend this work by scaling to a large number of cameras and also accounting for the uncertainties arising from the underlying vision algorithms. We would also like to deploy active cameras along with a team of robots ([7], [8]) for indoor surveillance.

VI. CONCLUSION

This paper describes a novel POMDP-based approach to coordinating and controlling a network of active cameras for maximizing the number of targets observed with a guaranteed resolution in an uncertain, partially observable surveillance environment. Specifically, our approach helps to eliminate the dependency on wide-view static cameras for tracking the targets' locations and simultaneously performs the tracking and observation of the targets at high resolution. We have exploited the conditional independence property in the targets' motion and observation for our surveillance problem in order to reduce the exponential policy computation time to linear time in the number of targets. The experimental evaluation shows that our proposed POMDP controller performs better than the state-of-the-art approaches and is feasible and practical to operate in real-world environments.

REFERENCES

- [1] AXIS 232D+ Network Dome Camera datasheet (<http://www.axis.com>).
- [2] C. Costello and I.-J. Wang. Surveillance camera coordination through distributed scheduling. In *Proc. CDC*, pages 1485–1490, 2005.
- [3] B. P. Gerkey, R. T. Vaughan, and A. Howard. The Player/Stage project: Tools for multi-robot and distributed sensor systems. In *Proc. ICAR*, pages 317–323, 2003.
- [4] C. Huang and L. Fu. Multitarget visual tracking based effective surveillance with cooperation of multiple active cameras. *IEEE Trans. Syst., Man, Cybern. B*, 41(1):234–247, 2011.
- [5] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra. Planning and acting in partially observable stochastic domains. *AIJ*, 101:99–134, 1998.
- [6] N. Krahnstoeber, T. Yu, S.-N. Lim, K. Patwardhan, and P. Tu. Collaborative Real-Time Control of Active Cameras in Large Scale Surveillance Systems. In *Proc. M2SFA2*, 2008.
- [7] K. H. Low, W. K. Leow, and M. H. Ang, Jr. Task allocation via self-organizing swarm coalitions in distributed mobile sensor network. In *Proc. AAAI*, pages 28–33, 2004.
- [8] K. H. Low, W. K. Leow, and M. H. Ang, Jr. Autonomic mobile sensor network with self-coordinated task allocation and execution. *IEEE Trans. Syst., Man, Cybern. C*, 36(3):315–327, 2006.
- [9] P. Natarajan, T. N. Hoang, K. H. Low, and M. Kankanhalli. Decision-theoretic approach to maximizing observation of multiple targets in multi-camera surveillance. In *Proc. AAMAS*, pages 155–162, 2012.
- [10] P. Natarajan, T. N. Hoang, K. H. Low, and M. Kankanhalli. Decision-theoretic coordination and control for active multi-camera surveillance in uncertain, partially observable environments. arXiv:1209.4275, 2012.
- [11] E. Sommerlade and I. Reid. Probabilistic surveillance with multiple active cameras. In *Proc. ICRA*, pages 440–445, 2010.
- [12] B. Song, C. Soto, A. K. Roy-Chowdhury, and J. A. Farrell. Decentralized camera network control using game theory. In *ICDSC*, pages 1–8, 2008.
- [13] C. Soto, B. Song, and A. K. Roy-Chowdhury. Distributed multi-target tracking in a self-configuring camera network. In *Proc. CVPR*, pages 1486–1493, 2009.
- [14] W. Starzyk and F. Z. Qureshi. Multi-tasking smart cameras for intelligent video surveillance systems. In *Proc. AVSS*, pages 154–159, 2011.
- [15] C. Ward and M. Naish. Scheduling active camera resources for multiple moving targets. In *Proc. CCECE*, pages 528–532, 2009.

¹<http://www.comp.nus.edu.sg/~lowkh/camera.html>

A. Observation model factorization

$$\begin{aligned}
P(Z|S) &= P(Z|T, C) \\
&= P(z_1, z_2, \dots, z_m | t_1, t_2, \dots, t_m, C) \\
&= \prod_{k=1}^m P(z_k | t_k, C) \\
&= \prod_{k=1}^m P(z_k | t_{l_k}, C).
\end{aligned}$$

The last two equalities are due to the conditional independence assumption in the observation model (Section III-C).

B. Posterior belief decomposition

$$\begin{aligned}
B'(S') &= \eta P(Z|S') \sum_{S \in \mathcal{S}} P(S'|S, A) B(S) \\
&= \eta \prod_{k=1}^m P(z_k | t'_{l_k}, C') \sum_{T \in \mathcal{T}^m} \sum_{C \in \mathcal{C}^n} \prod_{k=1}^m P(t'_k | t_k) \\
&\quad \prod_{i=1}^n \delta_{\tau(c_i, a_i)}(c'_i) \prod_{k=1}^m b_k(t_k) \prod_{i=1}^n \delta_{\hat{c}_i}(c_i) \\
&= \eta \prod_{k=1}^m P(z_k | t'_{l_k}, C') \left(\sum_{T \in \mathcal{T}^m} \prod_{k=1}^m P(t'_k | t_k) \prod_{k=1}^m b_k(t_k) \right) \\
&\quad \left(\sum_{C \in \mathcal{C}^n} \prod_{i=1}^n \delta_{\tau(c_i, a_i)}(c'_i) \prod_{i=1}^n \delta_{\hat{c}_i}(c_i) \right) \\
&= \eta \prod_{k=1}^m P(z_k | t'_{l_k}, C') \prod_{k=1}^m \sum_{t_k \in \mathcal{T}} P(t'_k | t_k) b_k(t_k) \\
&\quad \prod_{i=1}^n \delta_{\tau(\hat{c}_i, a_i)}(c'_i) \\
&= \sum_{C' \in \mathcal{C}^n} \prod_{k=1}^m \eta_k \prod_{i=1}^n \delta_{\hat{c}'_i}(c'_i) \\
&\quad \prod_{k=1}^m P(z_k | t'_{l_k}, C') \sum_{t_k \in \mathcal{T}} P(t'_k | t_k) b_k(t_k) \prod_{i=1}^n \delta_{\hat{c}'_i}(c'_i) \\
&= \prod_{k=1}^m \eta_k P(z_k | t'_{l_k}, C') \sum_{t_k \in \mathcal{T}} P(t'_k | t_k) b_k(t_k) \prod_{i=1}^n \delta_{\hat{c}'_i}(c'_i) \\
&= \prod_{k=1}^m b'_k(t'_k) \prod_{i=1}^n \delta_{\hat{c}'_i}(c'_i).
\end{aligned}$$

The first equality is due to (1). The second equality follows from (2), (3), and (4). The fifth equality follows from $\eta = \sum_{C' \in \mathcal{C}^n} \prod_{k=1}^m \eta_k \prod_{i=1}^n \delta_{\hat{c}'_i}(c'_i)$ (Section E). The last equality is due to (6).

C. Reward function decomposition

$$\begin{aligned}
R(B) &= \sum_{S \in \mathcal{S}} R(S) B(S) \\
&= \sum_{(T, C) \in \mathcal{S}} R((T, C)) B((T, C)) \\
&= \sum_{C \in \mathcal{C}^n} \sum_{T \in \mathcal{T}^m} \sum_{k=1}^m \tilde{R}(t_k, C) \prod_{k=1}^m b_k(t_k) \prod_{i=1}^n \delta_{\hat{c}_i}(c_i) \\
&= \sum_{k=1}^m \sum_{t_k \in \mathcal{T}} \tilde{R}(t_k, \hat{C}) b_k(t_k) \sum_{T_{-k} \in \mathcal{T}^{m-1}} \prod_{j \neq k} b_j(t_j) \\
&= \sum_{k=1}^m \sum_{t_k \in \mathcal{T}} \tilde{R}(t_k, \hat{C}) b_k(t_k) \\
&= \sum_{k=1}^m \left(\sum_{t_k \in \mathcal{T}} \tilde{R}(t_k, \hat{C}) b_k(t_k) \right) \\
&= \sum_{k=1}^m \tilde{R}(b_k, \hat{C})
\end{aligned}$$

where $T_{-k} = (t_1, \dots, t_{k-1}, t_{k+1}, \dots, t_m)$. The third equality is due to (4) and (7). The fifth equality follows from our independence assumption similar to that in (4) and the law of total probability:

$$\sum_{T_{-k} \in \mathcal{T}^{m-1}} \prod_{j \neq k} b_j(t_j) = \sum_{T_{-k} \in \mathcal{T}^{m-1}} P(T_{-k}) = 1.$$

D. Value function decomposition

$$\begin{aligned}
V(B, A) &= \sum_{Z \in \mathcal{Z}} R(B') P(Z|B, A) \\
&= \sum_{Z \in \mathcal{Z}} \sum_{k=1}^m \tilde{R}(b'_k, \hat{C}') \prod_{j=1}^m P(z_j | b_j, \hat{C}') \\
&= \sum_{k=1}^m \sum_{Z \in \mathcal{Z}} \tilde{R}(b'_k, \hat{C}') \prod_{j=1}^m P(z_j | b_j, \hat{C}') \\
&= \sum_{k=1}^m \sum_{z_k \in \mathcal{Z}} \tilde{R}(b'_k, \hat{C}') P(z_k | b_k, \hat{C}') \sum_{Z_{-k} \in \mathcal{Z}^{m-1}} \prod_{j \neq k} P(z_j | b_j, \hat{C}') \\
&= \sum_{k=1}^m \sum_{z_k \in \text{fov}(\hat{C}')} \tilde{R}(b'_k, \hat{C}') P(z_k | b_k, \hat{C}') \\
&= \sum_{k=1}^m \sum_{z_k \in \text{fov}(\hat{C}')} \sum_{t'_k \in \mathcal{T}} \tilde{R}(t'_k, \hat{C}') b'_k(t'_k) P(z_k | b_k, \hat{C}') \\
&= \sum_{k=1}^m \sum_{z_k \in \text{fov}(\hat{C}')} \sum_{t'_k \in \mathcal{T}} \tilde{R}(t'_k, \hat{C}') \hat{b}'_k(t'_k) \\
&= \sum_{k=1}^m \tilde{V}(b_k, \hat{C}') = \sum_{k=1}^m \tilde{V}(b_k, (\tau(\hat{c}_1, a_1), \dots, \tau(\hat{c}_n, a_n)))
\end{aligned}$$

where $\widehat{C}' \triangleq (c'_1, \dots, c'_n)$ and $Z_{-k} = (z_1, \dots, z_{k-1}, z_{k+1}, \dots, z_m)$. The first equality is due to (11). The second equality is obtained using (8) and $\eta^{-1} = \sum_{C' \in \mathcal{C}^n} \prod_{k=1}^m \eta_k^{-1} \prod_{i=1}^n \delta_{c'_i}(c'_i)$ (Section E). The fifth equality follows from $P(Z_{-k}|B_{-k}, A) = \prod_{j \neq k} P(z_j|b_j, \widehat{C}')$

where $B_{-k}(S) = \prod_{j \neq k} b_j(t_j) \prod_{i=1}^n \delta_{c_i}(c_i)$ and then the law of total probability: $\sum_{Z_{-k} \in \mathcal{Z}^{m-1}} \prod_{j \neq k} P(z_j|b_j, \widehat{C}') = 1$. Also, note

that when $z_k \notin \text{fov}(\widehat{C}')$, $\widetilde{R}(b'_k, \widehat{C}') = 0$. The sixth equality is due to (9). Since the normalizing constant of $b'_k(t'_k)$ is $1/P(z_k|b_k, \widehat{C}')$, the seventh equality follows.

$$E. \text{ Derivation of } \eta = \sum_{C' \in \mathcal{C}^n} \prod_{k=1}^m \eta_k \prod_{i=1}^n \delta_{c'_i}(c'_i)$$

$$\begin{aligned} \eta^{-1} &= P(Z|B, A) \\ &= \sum_{S' \in \mathcal{S}} P(Z|S')P(S'|B, A) \\ &= \sum_{S' \in \mathcal{S}} P(Z|S') \sum_{S \in \mathcal{S}} P(S'|S, A)P(S|B) \\ &= \sum_{S' \in \mathcal{S}} P(Z|S') \sum_{S \in \mathcal{S}} P(S'|S, A)B(S) \\ &= \sum_{C' \in \mathcal{C}^n} \sum_{T' \in \mathcal{T}^m} \prod_{k=1}^m P(z_k|t'_k, C') \\ &\quad \sum_{C \in \mathcal{C}^n} \sum_{T \in \mathcal{T}^m} \left(\prod_{k=1}^m P(t'_k|t_k) \prod_{i=1}^n \delta_{\tau(c_i, a_i)}(c'_i) \right) \left(\prod_{k=1}^m b_k(t_k) \prod_{i=1}^n \delta_{c_i}(c_i) \right) \\ &= \sum_{C' \in \mathcal{C}^n} \sum_{T' \in \mathcal{T}^m} \prod_{k=1}^m P(z_k|t'_k, C') \left(\sum_{T \in \mathcal{T}^m} \prod_{k=1}^m P(t'_k|t_k) \prod_{k=1}^m b_k(t_k) \right) \\ &\quad \left(\sum_{C \in \mathcal{C}^n} \prod_{i=1}^n \delta_{\tau(c_i, a_i)}(c'_i) \prod_{i=1}^n \delta_{c_i}(c_i) \right) \\ &= \sum_{C' \in \mathcal{C}^n} \sum_{T' \in \mathcal{T}^m} \prod_{k=1}^m P(z_k|t'_k, C') \prod_{k=1}^m \sum_{t_k \in \mathcal{T}} P(t'_k|t_k) b_k(t_k) \\ &\quad \prod_{i=1}^n \delta_{\tau(\hat{c}_i, a_i)}(c'_i) \\ &= \sum_{C' \in \mathcal{C}^n} \sum_{T' \in \mathcal{T}^m} \prod_{k=1}^m P(z_k|t'_k, C') \prod_{k=1}^m \sum_{t_k \in \mathcal{T}} P(t'_k|t_k) b_k(t_k) \\ &\quad \prod_{i=1}^n \delta_{c'_i}(c'_i) \\ &= \sum_{C' \in \mathcal{C}^n} \prod_{k=1}^m \sum_{t'_k \in \mathcal{T}} P(z_k|t'_k, C') \sum_{t_k \in \mathcal{T}} P(t'_k|t_k) b_k(t_k) \prod_{i=1}^n \delta_{c'_i}(c'_i) \\ &= \sum_{C' \in \mathcal{C}^n} \prod_{k=1}^m \eta_k^{-1} \prod_{i=1}^n \delta_{c'_i}(c'_i) \end{aligned}$$

where $\eta_k^{-1} = \sum_{t'_k \in \mathcal{T}} P(z_k|t'_k, C') \sum_{t_k \in \mathcal{T}} P(t'_k|t_k) b_k(t_k) = P(z_k|b_k, C')$. The fourth equality follows from (2), (3), and (4). It follows that $\eta = \sum_{C' \in \mathcal{C}^n} \prod_{k=1}^m \eta_k \prod_{i=1}^n \delta_{c'_i}(c'_i)$.