

CameraNets: Coverage and Data Management Protocols for Smart Camera Networks^{*}

Vikram P. Munishwar
Computer Science
State University of New York, Binghamton
vmunish1@cs.binghamton.edu

Nael B. Abu-Ghazaleh
Computer Science
Carnegie Mellon University, Qatar, and
State University of New York, Binghamton
naelag@cmu.edu

Abstract

The dissertation aims to address a number of important systems-related problems associated with smart camera networks, consisting of autonomous and cooperative camera nodes, to enable capturing most useful physical phenomena and managing the captured data efficiently. Specifically, a smart camera network should autonomously configure itself in order to maximize the total useful information it captures. Next, the captured information should be managed efficiently, either by delivering it to the end users or by archiving it within the network for later retrieval. In this work, we focus on a subset of the systems-related challenges that should be addressed to achieve these goals.

1 Introduction

Recent technological developments in processing and imaging have created an opportunity for the development of smart camera nodes that can operate autonomously and collaboratively to meet an application's requirements. Cameras are increasingly being used at places where there is a need for visual monitoring, but a lack of enough human resources. Smart camera networks have a wide range of applications in areas such as security monitoring and surveillance, habitat monitoring, traffic management, health care and telemedicine. The self-configuring nature of camera nodes coupled with wireless connectivity enable them to easily become a part of the existing fixed infrastructure of surveillance cameras, or achieve ad hoc surveillance in situations where infrastructure is unavailable or expensive, or quick deployment is desired.

Existing camera-based surveillance systems use fixed communication infrastructure, where high-end CCTV cameras are connected with a central controlling location using wires [4]. Fixed and costly infrastructure of such systems and the need for a human-in-the-loop limits their usability for many applications that need low-cost, short-term surveillance or use mobile cameras.

The goal of this work is to identify and address a subset of the systems-related challenges associated with advancing the state of the traditional surveillance systems into autonomously configurable, scalable, easily deployable, and

cost and energy efficient smart camera networks. Specifically, a smart camera network should autonomously configure itself in order to maximize the total useful information it captures. This leads to camera coverage specific challenges. Next, the captured information should be managed efficiently, either by delivering it to the end user or by archiving it within the network for query-based later retrieval. This leads to networking, data storage and retrieval related challenges. These are important systems-related challenges that are a part of the most of the smart camera network deployments, and this dissertation attempts to advance the current state of camera networks by addressing a subset of these challenges.

As a first step to address these challenges, we construct a design space for camera network applications in order to understand requirements of different applications and the solution alternatives that are necessary to address them. The approach that we used to construct the design space is briefly discussed in Section 2.

Among the coverage related problems, we focus on three types of coverage goals. The first goal focuses on continuously tracking maximum number of targets of interest. However, since not all applications are interested tracking targets continuously, and may demand covering each target for at least a certain duration of time, the second goal focuses on developing an appropriate schedule for cameras. The last coverage-specific goal focuses on ensuring that not a single target passing through a given region remains undetected. This type of coverage is known as a *barrier* coverage, and it is useful for monitoring country borders or sensitive areas. Coverage problems and their solution approaches are discussed in Section 3.

The data management related challenges can be split into networking-related challenges and storage/indexing related challenges. For both types of challenges, we attempt to exploit the inter-linking between them and the camera coverage policies. From the networking related challenges, we focus on developing application-driven network management approach to jointly adapt the networking and coverage protocols to ensure application-specific data accuracy for the end-users. Furthermore, from the coordination perspective, we focus on providing efficient and reliable communication of control data among cameras to meet their coverage-specific needs. Section 4 discusses networking related approaches. From the data management specific challenges, we focus

^{*}This research work is funded by Qatar National Research Fund (QNRF) under the National Priorities Research Program (NPRP) Grant No.:08-562-1-095

on achieving efficient, distributed garbage collection of the stored visual information. Section 5 presents the distributed garbage collection problem and its solution approach. Finally, Section 6 presents concluding remarks and directions for future research.

2 Design Space for Smart Camera Networks

Designing a smart camera network based on the users' requirements is a complex task due to the large variety of parameters (in terms of resources, protocols, architectures, etc.) that need to be tuned appropriately. We begin by identifying the common tasks that need to be performed by majority of the camera networks. The required accuracy of these tasks is mainly influenced by the application's requirements. Thus, we extract general application requirements into a problem space. Specifically, we classify the problem space based on four dimensions: *where and how to deploy the network?*, *what to monitor?*, *how to monitor?*, *how to handle captured data?*. Furthermore, we also consider common requirements related to cost and power budget, and privacy and security concerns, which potentially affect the solution space for all of the above problems. Each requirement can be addressed by a spectrum of alternatives, which we denote as a parameter-space. For instance, the deployment size could range from small-scale to large-scale, the architecture could range from flat to centralized, etc. The solution-space essentially focuses on providing insights about appropriate tuning of parameters and selection of algorithms to address the requirements for different scenarios.

3 Coverage

In this section, we present solution approaches for the three types of coverage problems: tracking-based, time-based, and barrier coverage.

3.1 Tracking-based Coverage

The goal of tracking-based coverage is to ensure that maximum useful targets are being tracked at any given time. Such type of coverage is important for tracking suspicious people in homeland security application, or tracking mischievous kids in a kindergarten monitoring application. Existing approaches propose heuristics for providing maximum coverage for static targets [2, 3], or focus on enabling multi-object tracking, without consideration for utility maximization [6, 14].

Coverage for static targets: We first focus on the basic instance of the problem, where we assume that *static* targets are present in a cameras deployment region having no obstacles. Furthermore, cameras are assumed to take discrete pans. We first formulate the problem as an Integer Linear Programming Problem. Since the optimization problem is NP-hard, we present a polynomial-time solvable heuristic, centralized force-directed approach (CFA), that provides very close to optimal solution. Essentially, the key idea in CFA is to iteratively select most beneficial camera-pan pairs, till each camera is assigned a pan. However, as the scale of the network grows, the centralized solutions cease to be useful. To address these problems, we propose to use hierarchical mechanism for coverage maximization in large-scale networks. The intuition is to spatially decompose cameras

within the network into small-sized partitions by exploiting the separation among cameras due to their geographical placement or presence of obstacles. The details of the proposed approaches are given in our paper [11].

We present simulation-based performance evaluation of the proposed coverage policies (including existing approaches: Centralized Greedy Approach (CGA) and Distributed Greedy Approach (DGA) [2]). Each camera is set with a pre-defined maximum and minimum lengths of the FoV, 100 meters, and 0 meters, and FoV angle to be 45°. The camera can choose from 8 discrete pan values as its orientation. For the hierarchical approach, the maximum allowed cluster size is set to be 10. Each experiment was run for 500 seconds. Coverage gain is represented using *percent coverage*, which is a ratio of the total targets covered to the total potentially-coverable targets. Overhead of the policy is measured in terms of worst-case end-to-end response time. Cameras and targets are deployed randomly on a terrain of dimensions 1000x1000 meters.

As it can be observed from Figure 1 and Figure 2, the CFA not only outperforms the other heuristics, but also gives extremely close to optimal results. Moreover, the hierarchical policy also tracks the optimal very closely even when the terrain is densely populated with targets. As noted from the overhead graph in Figure 3, the overhead of centralized policies increases drastically as the number of cameras increases, while the increase in overhead for the distributed and semi-centralized approaches is minimal. Also, the Hierarchical approach outperforms DGA, even for networks with larger scale, since the DGA is inherently iterative in its approach.

Extension for mobile targets: In order to provide solution for mobile targets, we ask two basic questions: *How to compute camera configurations?* and *When to perform the former?* The first question can be addressed by using a centralized (by asking a central node to compute optimal solution), collaborative (by enabling cameras to collaborate with their neighbors), or local (essentially greedy) approach.

To answer the second question, we propose two approaches that work by frequently reporting the current state, which includes the current targets' locations and camera configuration parameters, to a central node and obtaining optimal camera configurations. Specifically, in the first approach, cameras periodically (after fixed interval) report the state information, while in the second approach, the central node decides when to compute the optimal solution, based on the measure of acceptable loss of coverage for the underlying application.

Further extensions: We plan to include the notion of penalty, as a cost required to configure cameras, and evaluate the proposed coverage policies on a small testbed of cameras, to verify their correctness. In addition, we plan to extend the formulation to consider all PTZ parameters as well as presence of obstacles.

3.2 Time-based coverage

For some application scenarios, it may not be required to track targets continuously; it is sufficient to cover targets intermittently. For instance, it is sufficient to track each animal for a small duration in an animal tracking application [7]. The goal for such applications is to cover possibly all targets

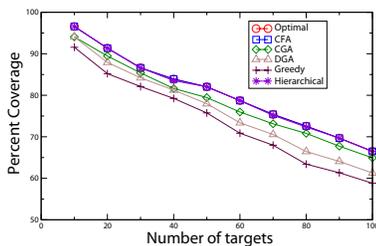


Figure 1. Study of varying number of targets. Number of cameras: 50.

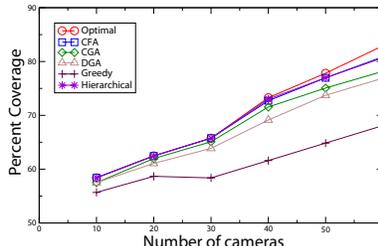


Figure 2. Study of varying number of cameras. Number of targets: 100.

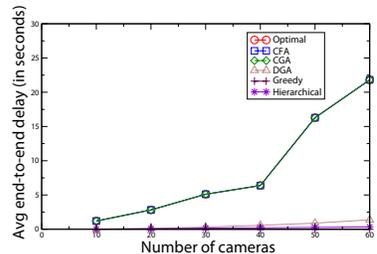


Figure 3. Study of overhead for varying number of cameras. Number of targets: 100.

in the cameras deployment region, at least for a certain duration, such that the cameras collectively maximize the monitoring duration of maximum number of targets.

We plan to formulate the camera scheduling problem as an optimization problem. Further, we plan to decompose the optimization problem in to multiple sub-problems, which can be solved individually based on certain parameters obtained from the master problem. Existing approaches also formulate the problem as an optimization problem, and propose heuristics to solve it [8, 9, 15], however we plan to work on improving the overall coverage gain by proposing a distributed solution that converges to optimal.

3.3 Barrier coverage

The goal of a barrier coverage problem is to ensure cameras placement and configuration such that there is *no undetected breach* possible through the given area. The existing works mainly focus on providing barrier coverage by selecting a subset of pre-deployed cameras (or directional sensors) [1, 10], and do not focus on the cameras' placement.

We plan to use a GIS-based data structure to obtain a representation of an irregular terrain that can be used to discretize the space, and carry out visibility analysis from discrete points in order to select smallest number of discrete points where cameras can be placed.

4 Networking-related Challenges

In this section, we discuss about application-driven network management for data delivery and camera coordination related problems.

4.1 Application-driven network management

Due to the potentially large amount of data generated within the network, it is critical to manage the network to satisfy application-specific needs in terms of perceived quality and jitter. Specifically, data captured by each camera can be of varying interest to the application, and thus its delivery can be prioritized by adapting the video quality (e.g. by adjusting data rate) and resource allocation. We consider a class of applications where video streams from multiple cameras need to be delivered together to the end user(s). For instance, such video streams can be combined or stitched together to create a panoramic view or stored at external storage location(s). Video stream aggregation near the cameras is vital for such applications in order to reduce network burden by removing spatial redundancies. Note that the temporal redundancies will already be removed by using appropriate codec (e.g.

MPEG-4). Further, each camera may or may not capture interesting information for the application. This can be exploited to assign different resolutions to different streams. We propose to attain these goals by coupling the coverage problem with the networking problem. Specifically, we plan to create an optimization framework for application-specific resource allocation that takes coverage specific information (such as FoV configurations of cameras) to allocate networking resources such as paths for data delivery and data rate. Next, we plan to propose a distributed mechanism to achieve application specific resource allocation.

4.2 Reliable group coordination

The need to improve the camera coordination performance was motivated from the delay observed in the overall response time for obtaining optimal camera configurations, as shown in Figure 3. Here, the end-to-end delay is defined as the difference of time when the first camera sent its request to the base-station, and when the last camera received the optimal configuration from the base-station. Since, it involves control data communication over wireless channel, we have used TCP as a reliable transport-layer protocol. As it can be observed, the delay for centralized approaches is very high, especially for larger number of cameras. This is because the overall response time is severely affected even if the data for one camera is delayed from the base station, since the base station backs off multiplicatively on packet losses, based on the TCP's congestion control mechanism.

Although, this problem was detected for the centralized coverage maximization approach, it is prevalent in majority of the centralized approaches [2, 5, 11, 16]. The approach used in such scenarios involves communicating data from all the nodes to a central place, computing optimal solution centrally, and disseminating the optimal configurations back to each node.

We identify two important design goals for coordination in centralized approaches: (1)*Reliability*: it is important to ensure that all the data is reliably communicated to and from the central node; (2)*Timeliness*: smaller the response time, smaller will be the loss of overall utility. Timeliness is especially important where the state may change over a period of time. We plan to propose a hierarchical aggregation based technique to address these problem. Moreover, to reduce the response time further, we plan to identify and use the only cameras that can contribute the most towards obtaining

a close-to-optimal solution.

5 Efficient Garbage Collection

When the data is stored within the network, the storage space at camera nodes may get exhausted quickly due to the large size of visual data. Thus, we focus on the problem: *How to achieve efficient storage reclamation in smart camera networks?*

Due to the limited number of allowed write operations and block-erasable nature of flash memories, it is a challenging problem to ensure energy efficiency and wear-leveling while supporting efficient garbage collection of video data. Existing techniques focus on providing efficient garbage collection on a single node [12, 13]. However, coverage overlaps may lead to spatial redundancies or correlations across cameras, which can be exploited for performing garbage collection. In this work, we plan to propose a light-weight distributed garbage collection policy by exploiting meta-information from the underlying coverage policy.

6 Conclusion and Future Work

Traditional multi-camera systems often require a human-in-the-loop to control cameras and analyze the video streams, which is a costly and unscalable approach. Instead, a camera network should be able to autonomously adapt its coverage and data management protocols to meet the underlying application's needs. This dissertation focuses on advancing the current state of the art by identifying and addressing important systems-related challenges associated with meeting this goal. Specifically, from camera coverage perspective, we focus on enabling cameras to autonomously configure their PTZ parameters to maximize the total useful information they capture. From the data management perspective, we focus on addressing the networking and storage related challenges, by exploiting their relationship with the underlying coverage policy.

Beyond the focus of this dissertation, other important systems-related problems need to be addressed to reach the overall goal. For instance, this work primarily focuses on addressing challenges with static cameras deployment, which can be used as a stepping stone to focus on scenarios where a subset of cameras are mobile (mounted on robots or carried by people), and address problems related to a variety of application scenarios, including ubiquitous computing, interactive remote-spaces monitoring, etc.

7 References

- [1] J. Adriaens, S. Megerian, and M. Potkonjak. Optimal worst-case coverage of directional field-of-view sensor networks. *Sensor and Ad Hoc Communications and Networks, 2006. SECON'06. 2006 3rd Annual IEEE Communications Society on*, 1, 2006.
- [2] J. Ai and A. A. Abouzeid. Coverage by directional sensors in randomly deployed wireless sensor networks. *Journal of Combinatorial Optimization*, 2006.
- [3] Y. Cai, W. Lou, M. Li, and X. Li. Target-Oriented Scheduling in Directional Sensor Networks. *IEEE Infocom*, 2007.
- [4] R. Collins, A. Lipton, T. Kanade, H. Fujiyoshi, D. Duggins, Y. Tsin, D. Tolliver, N. Enomoto, O. Hasegawa, P. Burt, et al. *A System for Video Surveillance and Monitoring*. Carnegie Mellon University, the Robotics Institute, 2000.
- [5] J. Elson, L. Girod, and D. Estrin. Fine-grained network time synchronization using reference broadcasts. *ACM SIGOPS Operating Systems Review*, 2002.
- [6] A. Ercan, A. El Gamal, and L. Guibas. Object tracking in the presence of occlusions via a camera network. *Proceedings of the 6th international conference on Information processing in sensor networks*, 2007.
- [7] R. Kays, B. Kranstauber, P. Jansen, C. Carbone, M. Rowcliffe, T. Fountain, and S. Tilak. Camera traps as sensor networks for monitoring animal communities. In *IEEE 34th Conference on Local Computer Networks, 2009. LCN 2009*, pages 811–818, 2009.
- [8] N. Krahnstoeber, T. Yu, S. Lim, K. Patwardhan, and P. Tu. Collaborative Real-Time Control of Active Cameras in Large Scale Surveillance Systems. In *Workshop on Multi-camera and Multi-modal Sensor Fusion Algorithms and Applications*, 2008.
- [9] S. Lim, L. Davis, and A. Mittal. Task scheduling in large camera networks. *Computer Vision—ACCV 2007*, pages 397–407, 2007.
- [10] S. Megerian, F. Koushanfar, M. Potkonjak, and M. Srivastava. Worst and best-case coverage in sensor networks. *IEEE Transactions on Mobile Computing*, pages 84–92, 2005.
- [11] V. Munishwar and N. Abu-Ghazaleh. Scalable Target Coverage in Smart Camera Networks. In *ACM/IEEE International Conference on Distributed Smart Cameras*, 2010.
- [12] V. Munishwar, S. Tilak, and N. B. Abu-Ghazaleh. File-system support for adaptive resolution storage in sensor networks. In *Proceedings of Adhoc-Now*, 2006.
- [13] S. Nath. Energy efficient sensor data logging with amnesic flash storage. In *Proceedings of the 2009 International Conference on Information Processing in Sensor Networks—Volume 00*, pages 157–168. IEEE Computer Society, 2009.
- [14] F. Porikli and A. Divakaran. Multi-camera calibration, object tracking and query generation. In *Multimedia and Expo, 2003. ICME'03. Proceedings. 2003 International Conference on*, volume 1, 2003.
- [15] F. Z. Qureshi and D. Terzopoulos. Surveillance camera scheduling: a virtual vision approach. In *VSSN '05: Proceedings of the third ACM international workshop on Video surveillance & sensor networks*, pages 131–140, New York, NY, USA, 2005. ACM.
- [16] Y. Wu, S. Fahmy, and N. Shroff. Optimal sleep/wake scheduling for time-synchronized sensor networks with QoS guarantees. *IEEE/ACM Transactions on Networking (TON)*, 17(5):1508–1521, 2009.