

# Snapshot: A Self-Calibration Protocol for Camera Sensor Networks

Xiaotao Liu, Purushottam Kulkarni, Prashant Shenoy and Deepak Ganesan

Department of Computer Science

University of Massachusetts, Amherst, MA 01003

Email: {xiaotaol, purukulk, shenoy, dganesan}@cs.umass.edu

**Abstract**—A camera sensor network is a wireless network of cameras designed for ad-hoc deployment. The camera sensors in such a network need to be properly calibrated by determining their location, orientation, and range. This paper presents *Snapshot*, an automated calibration protocol that is explicitly designed and optimized for camera sensor networks. *Snapshot* uses the inherent imaging abilities of the cameras themselves for calibration and can determine the location and orientation of a camera sensor using only four reference points. Our techniques draw upon principles from computer vision, optics, and geometry and are designed to work with low-fidelity, low-power camera sensors that are typical in sensor networks. An experimental evaluation of our prototype implementation shows that *Snapshot* yields an error of 1-2.5 degrees when determining the camera orientation and 5-10cm when determining the camera location. We show that this is a tolerable error in practice since a *Snapshot*-calibrated sensor network can track moving objects to within 11cm of their actual locations. Finally, our measurements indicate that *Snapshot* can calibrate a camera sensor within 20 seconds, enabling it to calibrate a sensor network containing tens of cameras within minutes.

## I. INTRODUCTION

### A. Motivation

Recent advances in embedded systems technologies have made the design of camera sensor networks a reality. A camera sensor network is an ad-hoc wireless network of low-power imaging sensors that are connected to networked embedded controllers. Today, available camera sensors range from tiny, low-power cameras such as Cyclops to “cell-phone-class” cameras and from inexpensive webcams to high-resolution pan-tilt-zoom cameras.

Typical applications of camera sensor networks include active monitoring of remote environments and surveillance tasks such as object detection, recognition, and tracking. These applications involve acquisition of video from multiple camera sensors and real-time processing of this data for recognition, tracking, and camera control. Video acquisition and processing involves interaction and coordination between multiple cameras, for instance, to hand-off tracking responsibilities for a moving object from one camera to another. Precise calibration of camera sensors is a necessary pre-requisite for such coordination. Calibration of a camera sensor network involves determining the location, orientation, and range of each camera sensor in three dimensional space as well as the overlap and spatial relationships between nearby cameras.

Camera calibration is well studied in the computer vision community [8], [18], [21], [23], [24]. Many of these techniques

are based on the classical Tsai method—they require a set of reference points whose true locations are known in the physical world and use the projection of these points on the camera image plane to determine camera parameters. Despite the wealth of research on calibration in the vision community, adapting these techniques to sensor networks requires us to pay careful attention to the differences in hardware characteristics and capabilities of sensor networks.

First, sensor networks employ low-power, low-fidelity cameras such as the CMUcam [16] or Cyclops [11] that have coarse-grain imaging capabilities; at best, a mix of low-end and a few high-end cameras can be assumed in such environments. Further, the cameras may be connected to nodes such as the Intel Motes or Intel Stargates that have two or three orders of magnitude less computational resources than PC-class workstations. Consequently, calibration techniques for camera sensor networks need to work well with low-resolution cameras and should be feasible on low-end computational platforms. Vision-based techniques that employ computationally complex calibration algorithms are often infeasible on sensor platforms. Instead, we must draw upon techniques that are computationally-efficient and require only a few reference points for calibration.

Second, vision-based calibration techniques typically assume that the location of all reference points is accurately known. In contrast, an automated procedure to calibrate cameras in a sensor network will depend on a positioning system (e.g., GPS or ultrasound) to determine the coordinates of reference points. All positioning systems introduce varying amounts of error in the coordinates of reference points, and consequently, the calibration technique must determine the impact of such errors on the computed camera location and orientation. The impact of using imprecise reference point locations on calibration error has not been addressed in vision-based calibration techniques [8], [18], [21], [23], [24].

Finally, a camera sensor network will comprise tens or hundreds of cameras and any calibration technique must scale to these large environments. Further, camera sensor networks are designed for ad-hoc deployment, for instance, in environments with disasters such as fires or floods. Since quick deployment is crucial in such environments, it is essential to keep the time required for calibrating the system to a minimum. Scalability and calibration latency have typically not been issues of concern in vision-based methods.

Automated localization techniques are a well-studied prob-

lem in the sensor community and a slew of techniques have been proposed. Localization techniques employ beacons (e.g., IR [1], ultrasound [2], RF [3]) and use sophisticated triangulation techniques to determine the location of a node. Most of these techniques have been designed for general-purpose sensor networks, rather than camera sensor networks in particular. Nevertheless, they can be employed during calibration, since determining the node location is one of the tasks performed during calibration. However, localization techniques are by themselves not sufficient for calibration. Cameras are *directional* sensors and camera calibration also involves determining other parameters such as the orientation of the camera (where a camera is pointing) as well as its range (what it can see). In addition, calibration is also used to determine overlap between neighboring cameras. Consequently, calibration is a harder problem than pure localization.

The design of an automated calibration technique that is cost-effective and yet scalable, efficient, and quickly deployable is the subject matter of this paper.

### B. Research Contributions

In this paper, we propose *Snapshot* a novel wireless protocol for calibrating camera sensor networks. We draw upon calibration techniques from the vision community and develop a variant that is particularly suited to the constraints and needs of sensor networks. *Snapshot* requires only four reference points to calibrate each camera sensor and allows these points to be randomly chosen. Both properties are crucial for sensor networks, since fewer reference points and fewer restrictions enable faster calibration and reduce the computational overhead for subsequent processing. Further, unlike sensor localization techniques that depend on wireless beacons, *Snapshot* does not require any specialized positioning equipment on the sensor nodes. Instead, it leverages the inherent picture-taking abilities of the cameras and the onboard processing on the sensor nodes to calibrate each node. *Snapshot* uses a positioning system to calculate locations of reference points, which in turn are used to estimate the camera parameters. Since positioning technologies introduce error in the reference point determination, we conduct a detailed error analysis to quantify how the error in reference points impacts calibration error.

Our techniques can be instantiated into a simple, quick and easy-to-use wireless calibration protocol—a wireless calibration device is used to define reference points for each camera sensor, which then uses principles from geometry, optics and elementary machine vision to calibrate itself. When more than four reference points are available, a sensor can use median filter and maximum likelihood estimation techniques to improve the accuracy of its estimates.

We have implemented *Snapshot* on a testbed of CMUcam sensors connected to wireless Stargate nodes. We have conducted a detailed experimental evaluation of *Snapshot* using our prototype implementation. Our experiments yield the following key results:

- 1) *Feasibility*: By comparing the calibration accuracies of low and high-resolution cameras, we show that it is

feasible to calibrate low-resolution cameras such as CMUcams without a significant loss in accuracy.

- 2) *Accuracy*: Our error analysis of *Snapshot* shows that the calibrated parameters are more sensitive to random errors in reference point locations than correlated errors. We experimentally show that *Snapshot* can localize a camera to within few centimeters of its actual location and determine its orientation with a median error of 1.3–2.5 degrees. More importantly, our experiments indicate that this level of accuracy is sufficient for tasks such as object tracking. We show that a system calibrated with *Snapshot* can localize an external object to within 11 centimeters of its actual location, which is adequate for most tracking scenarios.
- 3) *Efficiency*: We show that the *Snapshot* algorithm can be implemented on Stargate nodes and have running times in the order of a few seconds.
- 4) *Scalability*: We show that *Snapshot* can calibrate a camera sensor in about 20 seconds on current hardware. Since only a few reference points need to be specified—a process that takes a few seconds per sensor—*Snapshot* can scale to networks containing tens of camera sensors.

The rest of this paper is structured as follows. Section II presents some background and the problem formulation. Sections III, IV and V present the design of *Snapshot* its instantiation into a protocol and its use in an application. We present the error analysis of *Snapshot*, our prototype implementation and our experimentation evaluation in Sections VI, VII and VIII. Section IX describes related work and Section X presents our conclusions.

## II. PROBLEM FORMULATION

A camera sensor network is defined to be a wireless network of camera sensors, each connected to an embedded controller. A typical realization of a camera sensor node consists of a low-power camera such as the CMUcam [16] or Cyclops [11] connected to an embedded sensor platform such as the Intel Mote or the Intel Stargate. The sensor platform consists of a programmable microprocessor, memory, and a wireless interface for communication. Not all cameras in the system are homogeneous; in general, a small number of higher resolution cameras may be deployed to assist the low-fidelity cameras in performing their tasks.

Consider an ad-hoc deployment of  $N$  heterogeneous camera sensor nodes in an environment. An ad-hoc deployment implies that cameras are quickly placed without *a priori* planning. Given such an ad-hoc deployment, the location, orientation and the range of each camera sensor needs to be automatically determined. The goal of our work is to design a wireless protocol to automatically derive these parameters for each camera node. Specifically, the calibration protocol needs to determine the  $(x, y, z)$  coordinates of each camera, which is defined as the coordinates of the center of the camera lens. The protocol also needs to determine the camera orientation along the three axes, namely the *pan*  $\alpha$ , *tilt*  $\beta$  and *roll*  $\gamma$  of the camera respect to the left handed coordinate system. Finally, the protocol needs to determine the field of view of

each camera (i.e., what it can see) and the degree of overlap with neighboring cameras (i.e., the common regions visible to both cameras).

Our work assumes that the focal length  $f$  of camera lens is known to the calibration protocol. This is a reasonable assumption since lens parameters are typically published in the camera specifications by the manufacturer or they can be estimated offline for each camera prior to deployment [18]. Further, sensor nodes are assumed to lack specialized positioning devices such as GPS receivers, which suffer from 5-15m locationing error. Instead, our goal is to devise a protocol that exploits the inherent imaging abilities of each camera and the onboard processing on each sensor node to determine the above calibration parameters.

### III. SNAPSHOT DESIGN

Snapshot draws inspiration from a class of vision-based techniques called extrinsic camera calibration (extrinsic calibration determines external camera parameters such as its location and orientation, as opposed to intrinsic or internal parameters such as focal length and distortion of the lens). Our technique is similar in spirit to [8], [23], which use four reference points to determine extrinsic camera parameters; however, the technique used in *Snapshot* has been adapted to the specific needs of sensor networks. The basic *Snapshot* protocol involves taking pictures of a small randomly-placed calibration device. To calibrate each camera, at least four pictures of the device are necessary, and no three positions of the device must lie along a straight line. Each position of the device serves as a reference point; the coordinates of each reference point are assumed to be known and can be automatically determined by equipping the calibration device with a locationing sensor (e.g., ultra-sound Cricket receiver). Next, we describe how *Snapshot* uses the pictures and coordinates of the calibration device to estimate camera parameters. We also discuss how the estimates can be refined when additional reference points are available.

#### A. Camera Location Estimation

We begin with the intuition behind approach. Without loss of generality, we assume all coordinate systems are left handed, and the z-axis of the camera coordinate system is co-linear with the camera's optical axis. Consider a camera sensor  $C$  whose coordinates need to be determined. Suppose that four reference points  $R_1, R_2, R_3$  and  $R_4$  are given along with their coordinates for determining the camera location. No assumption is made about the placement of these points in the three dimensional space, except that these points be in visual range of the camera and that no three of them lie along a straight line. Consider the first two reference points  $R_1$  and  $R_2$  as shown in Figure 1. Suppose that point objects placed at  $R_1$  and  $R_2$  project an image of  $P_1$  and  $P_2$ , respectively, in the camera's image plane as shown in Figure 1. Further, let  $\theta_1$  be the angle incident by the the reference points on the camera. Since  $\theta_1$  is also the angle incident by  $P_1$  and  $P_2$  on the camera lens, we assume that it can be computed using elementary optics (as discussed later). Given  $\theta_1, R_1$  and  $R_2$ ,




Fig. 1. Projection of reference points on the image plane through the lens.




Fig. 2. Geometric representation of possible camera locations.

the problem of finding the camera location reduces to finding a point in space where  $R_1$  and  $R_2$  impose an angle of  $\theta_1$ . With only two reference points, there are infinitely many points where  $R_1$  and  $R_2$  impose an angle of  $\theta_1$ . To see why, consider Figure 2(a) that depicts the problem in two dimensions. Given  $R_1$  and  $R_2$ , the set of possible camera locations lies on the arc  $R_1CR_2$  of a circle such that  $R_1R_2$  is a chord of the circle and  $\theta_1$  is the angle incident by this chord on the circle. From elementary geometry, it is known that a chord of a circle inscribes a constant angle on any point on the corresponding arc. Since we have chosen the circle such that chord  $R_1R_2$  inscribes an angle of  $\theta_1$  on it, the camera can lie on any point on the arc  $R_1CR_2$ . This intuition can be generalized to three dimensions by rotating the arc  $R_1CR_2$  in space with the chord  $R_1R_2$  as the axis (see Figure 2(b)). Doing so yields a three dimensional surface of possible camera locations. The nature of the surface depends on the value of  $\theta_1$ : the surface is shaped like a football when  $\theta_1 > 90^\circ$ , is a sphere when  $\theta_1 = 90^\circ$ , and a double crown when  $\theta_1 < 90^\circ$ . The camera can lie on any point of this surface.

Next, consider the third reference point  $R_3$ . Considering points  $R_1$  and  $R_3$ , we obtain another surface that consists of all possible locations such that  $R_1R_3$  impose a known angle  $\theta_2$  on all points of this surface. Since the camera must lie on both surfaces, it follows that the set of possible locations is given by the intersection of these two surfaces. The intersection of two surfaces is a closed curve and the set of possible camera locations is reduced to any point on this curve.

Finally, if we consider the pair of reference points  $R_2$  and  $R_3$ , we obtain a third surface of all possible camera locations. The intersection of the first surface and the third yields a second curve of possible camera locations. The camera lies on the intersection of these two curves, and the curves can intersect in multiple points. The number of possible camera locations can be reduced further to at most 4 by introducing the fourth reference point  $R_4$ . Although 4 reference points give us up to 4 possible camera locations, we observe that, in reality, only one of these locations can generate the same

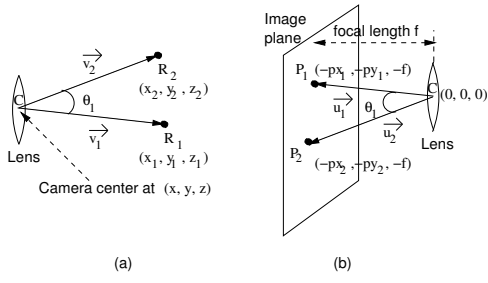


Fig. 3. Vector representation of reference points and their projections.




Fig. 4. Relationship between object location and its projection.

projections as  $R_1$ ,  $R_2$ ,  $R_3$ , and  $R_4$  on the image plane. Using elementary optics, it is easy to eliminate the false solutions and determine the true and unique location of the camera.

With this intuition, we now present the details of our technique. Consider a camera  $C$  placed at coordinates  $(x, y, z)$ , and four reference points  $R_1, \dots, R_4$  with coordinates  $(x_1, y_1, z_1) \dots (x_4, y_4, z_4)$ . The line joining the camera  $C$  with each reference point defines a vector. For instance, as shown in Figure 3(a), the line joining  $C$  and  $R_1$  defines a vector  $\overrightarrow{CR_1}$ , denoted by  $\vec{v}_1$ . The components of  $v_1$  are given by

$$\vec{v}_1 = \overrightarrow{CR_1} = \{x_1 - x, y_1 - y, z_1 - z\}$$

Similarly, the vector  $\vec{v}_i$  joining points  $C$  and  $R_i$  is given as

$$\vec{v}_i = \overrightarrow{CR_i} = \{x_i - x, y_i - y, z_i - z\} \quad 1 \leq i \leq 4$$

As shown in Figure 3(a), let  $\theta_1$  denote the angle between vectors  $\vec{v}_1$  and  $\vec{v}_2$ . The dot product of vectors  $\vec{v}_1$  and  $\vec{v}_2$  is given as

$$\vec{v}_1 \cdot \vec{v}_2 = |\vec{v}_1| |\vec{v}_2| \cos \theta_1 \quad (1)$$

By definition of the dot product,

$$\vec{v}_1 \cdot \vec{v}_2 = (x_1 - x)(x_2 - x) + (y_1 - y)(y_2 - y) + (z_1 - z)(z_2 - z) \quad (2)$$

The magnitude of vector  $\vec{v}_1$  is given as

$$|\vec{v}_1| = \sqrt{(x_1 - x)^2 + (y_1 - y)^2 + (z_1 - z)^2}$$

The magnitude of  $\vec{v}_2$  is defined similarly. Substituting these values into Equation 2, we get

$$\cos(\theta_1) = \frac{\vec{v}_1 \cdot \vec{v}_2}{|\vec{v}_1| \cdot |\vec{v}_2|} \quad (3)$$

Let  $\theta_2$ , through  $\theta_6$  denote the angles between vectors  $\vec{v}_1$  and  $\vec{v}_3$ ,  $\vec{v}_1$  and  $\vec{v}_4$ ,  $\vec{v}_2$  and  $\vec{v}_3$ ,  $\vec{v}_2$  and  $\vec{v}_4$  and  $\vec{v}_3$  and  $\vec{v}_4$  respectively. Similar expressions can be derived for  $\theta_2, \theta_3, \dots, \theta_6$ .

The angles  $\theta_1$  through  $\theta_6$  can be computed using elementary optics and vision, as discussed next. Given these angles and the coordinates of the four reference points, the above expressions yield six quadratic equations with three unknowns:  $x, y$  and  $z$ .

A non-linear solver can be used to numerically solve for these unknowns.

**Estimating  $\theta_1$  through  $\theta_6$ :** We now present a technique to compute the angle between any two vectors  $\vec{v}_i$  and  $\vec{v}_j$ . Consider any two reference points  $R_1$  and  $R_2$  as shown in Figure 3 (a). Figure 3 (b) shows the projection of these points through the camera lens onto the image plane. The image plane in a digital camera consists of a CMOS sensor that takes a picture of the camera view. Let  $P_1$  and  $P_2$  denote the projections of the reference points on the image plane as shown in the Figure 3(b), and let  $f$  denote the focal length of the lens. For simplicity, we define all points with respect to the camera's coordinate system: the center of the lens is assumed to be the origin in this coordinate system. Since the image plane is at a distance  $f$  from the lens, all points on the image plane are at a distance  $f$  from the origin. By taking a picture of the reference points, the coordinates of  $P_1$  and  $P_2$  can be determined. These are simply the pixel coordinates where the reference points project their image on the CMOS sensor; these pixels can be located in the image using a simple vision-based object recognition technique.<sup>1</sup> Let the resulting coordinates of  $P_1$  and  $P_2$  be  $(-px_1, -py_1, -f)$  and  $(-px_2, -py_2, -f)$  respectively. We define vectors  $\vec{u}_1$  and  $\vec{u}_2$  as lines joining the camera (i.e., the origin  $C$ ) to the points  $P_1$  and  $P_2$ . Then, the angle  $\theta_1$  between the two vectors  $\vec{u}_1$  and  $\vec{u}_2$  can be determined by taking the dot product of them.

$$\cos(\theta_1) = \frac{\vec{u}_1 \cdot \vec{u}_2}{|\vec{u}_1| |\vec{u}_2|}$$

The inverse cosine transform yields  $\theta_1$ , which is also the angle incident by the original reference points on the camera.

Using the above technique to estimate  $\theta_1$ – $\theta_6$ , we can then solve our six quadratic equations using a non-linear optimization algorithm [5] to estimate the camera location.

## B. Camera Orientation Estimation

We now describe the technique employed by *Snapshot* to determine the camera's orientation along the three axes. We assume that the camera location has already been estimated using the technique in the previous section. Given the camera location  $(x, y, z)$ , our technique uses three reference points to determine the pan, tilt, and roll of the camera. Intuitively, given the camera location, we need to align the camera in space so that the three reference points project an image at the same location as the pictures takes by the camera. Put another way, consider a ray of light emanating from each reference point. The camera needs to be aligned so that each ray of light pierces the image plane at the same pixel where the image of that reference point is located. One reference point is sufficient to determine the pan and tilt of the camera using this technique and three reference point are sufficient to uniquely determine all three parameters: pan, tilt and roll. Our technique uses the actual coordinates of three reference points and the pixel coordinates of their corresponding images to determine the

<sup>1</sup>In *Snapshot* the calibration device contains a colored LED and the vision-based recognizer must locate this LED in the corresponding image.

unknown rotation matrix  $R$  that represents the pan, tilt and roll of the camera.

Assume that the camera is positioned at coordinates  $(x, y, z)$  and that the camera has a pan of  $\alpha$  degrees, a tilt of  $\beta$  degrees, a roll of  $\gamma$  degrees. The composite 3x3 matrix corresponding to matrices representing the pan, tilt and roll rotations of the camera is denoted by  $\mathbf{R}$ .

If an object is located at  $(x_i, y_i, z_i)$  in the world coordinates, the object's location in the camera coordinates  $(x'_i, y'_i, z'_i)$  can be computed via Equation 4.

$$\begin{bmatrix} x'_i \\ y'_i \\ z'_i \end{bmatrix} = \mathbf{R} \times \begin{bmatrix} x_i - x \\ y_i - y \\ z_i - z \end{bmatrix} \quad (4)$$

Intuitively, we can construct and solve a set of linear equations (see Equation 5) where  $(x_1, y_1, z_1)$ ,  $(x_2, y_2, z_2)$ , and  $(x_3, y_3, z_3)$  are the world coordinates of 3 reference points, and  $(x'_1, y'_1, z'_1)$ ,  $(x'_2, y'_2, z'_2)$ , and  $(x'_3, y'_3, z'_3)$  are the corresponding camera coordinates to estimate  $\mathbf{R}$ , and then estimate  $\alpha$ ,  $\beta$ , and  $\gamma$  from  $\mathbf{R}$ . The three sets of linear equations in Equation 5 have unique solution for  $\mathbf{R}^T$  (since the right-hand-side matrix is non-singular).

$$\begin{bmatrix} x_1 - x & y_1 - y & z_1 - z \\ x_2 - x & y_2 - y & z_2 - z \\ x_3 - x & y_3 - y & z_3 - z \end{bmatrix} \times \mathbf{R}^T = \begin{bmatrix} x'_1 & y'_1 & z'_1 \\ x'_2 & y'_2 & z'_2 \\ x'_3 & y'_3 & z'_3 \end{bmatrix} \quad (5)$$

As shown in Figure 4, an object's location in the camera coordinates and the projection of the object on the image plane have the following relation:

$$\begin{bmatrix} x'_i \\ y'_i \\ z'_i \end{bmatrix} = \frac{D_i}{D_p} \times \begin{bmatrix} px_i \\ py_i \\ f \end{bmatrix} \quad (6)$$

where:

$$D_i = \sqrt{(x_i - x)^2 + (y_i - y)^2 + (z_i - z)^2} \text{ and} \\ D_p = \sqrt{px_i^2 + py_i^2 + f^2}$$

$D_i$  and  $D_p$  represent the magnitude of the object to camera center vector and the projection on image plane to camera center vector respectively.

Therefore, we can compute the location of an object in the camera coordinate system using Equation 6. The actual location of each reference point and its location in the camera coordinates can then be used in Equation 5 to determine the rotation matrix  $R$ . Given  $R$ , we can obtain pan  $\alpha$ , tilt  $\beta$ , and roll  $\gamma$  as follows:

$$\begin{aligned} \alpha &= \begin{cases} \arctan\left(\frac{r_{31}}{r_{33}}\right) + 180^\circ & \text{if } \frac{r_{31}}{\cos(\beta)} < 0 \text{ and } \frac{r_{33}}{\cos(\beta)} < 0 \\ \arctan\left(\frac{r_{31}}{r_{33}}\right) - 180^\circ & \text{if } \frac{r_{31}}{\cos(\beta)} >= 0 \text{ and } \frac{r_{33}}{\cos(\beta)} < 0 \\ \arctan\left(\frac{r_{31}}{r_{33}}\right) & \text{otherwise} \end{cases} \\ \beta &= \arcsin(r_{32}) \\ \gamma &= \begin{cases} \arctan\left(\frac{r_{12}}{r_{22}}\right) + 180^\circ & \text{if } \frac{r_{12}}{\cos(\beta)} < 0 \text{ and } \frac{r_{22}}{\cos(\beta)} < 0 \\ \arctan\left(\frac{r_{12}}{r_{22}}\right) - 180^\circ & \text{if } \frac{r_{12}}{\cos(\beta)} >= 0 \text{ and } \frac{r_{22}}{\cos(\beta)} < 0 \\ \arctan\left(\frac{r_{12}}{r_{22}}\right) & \text{otherwise} \end{cases} \end{aligned} \quad (7)$$

**Eliminating False Solutions:** Recall from Section III-A that our six quadratic equations yields up to four possible solutions for the camera location. Only one of these solutions is the true camera location. To eliminate false solutions, we

compute the pan, tilt and roll for each computed location using three reference points. The fourth reference point is then used to eliminate false solutions as follows: for each computed location and orientation, we project the fourth reference point onto the camera's image plane. The projected coordinates are then matched to the actual pixel coordinates of the reference point in the image. The projected coordinates will match the pixel coordinates only for the true camera location. Thus, the three false solutions can be eliminated by picking the solution with the smallest re-projection error. The chosen solution is always guaranteed to be the correct camera location.

### C. Determining Visual Range and Overlap

Once the location and orientation of each camera have been estimated, the visual range of cameras and the overlap between neighboring cameras can be determined. Overlap between cameras is an indication of the redundancy in sensor coverage and can also be used to localize and track moving objects.

The visual range of a camera can be approximated as a polyhedron. The apex of the polyhedron is the location of the camera's lens center and height of the pyramid is the maximum viewable distance of the camera. An object in the volume of the polyhedron is in the visual range of the camera. The viewable range of a camera is assumed to be finite to avoid distant objects appearing as point objects in images, which are not useful for detection and recognition tasks. After determining the camera location and orientation using *Snapshot*, the polyhedron visual range of each camera can be determined and computational geometry techniques for polyhedron intersection can be used to determine the overlap between cameras.

### D. Iterative Refinement of Estimates

While *Snapshot* requires only four reference points to calibrate a camera sensor, the estimates of the camera location and orientation can be improved if additional reference points are available. Suppose that  $n$  reference points,  $n \geq 4$ , are available for a particular sensor node. Then  $\binom{n}{4}$  unique subsets of four reference points can be constructed from these  $n$  points. For each subset of four points, we can compute the location and orientation of the camera using the techniques outlined in the previous sections. This yields  $\binom{n}{4}$  different estimates of the camera location and orientation. These estimates can be used to improve the final solution using the median filter method.

This method calculates the median of each estimated parameter, namely location coordinates  $x, y, z$ , pan  $\alpha$ , tilt  $\beta$ , and roll  $\gamma$ . These median values are then chosen as the final estimates of each parameter. The median filter method can yield a final solution that is different from all  $\binom{n}{4}$  initial solutions (since the median of each parameter is computed independently, the final solution need not correspond to any of the initial solutions). The median filter method is simple and cost-effective, and performs well when  $n$  is large.

## IV. A SELF-CALIBRATION PROTOCOL

In this section, we describe how the estimation techniques presented in the previous section can be instantiated into a

simple wireless protocol for automatically calibrating each camera sensor. Our protocol assumes that each sensor node has a wireless interface that enables wireless communication to and from the camera. The calibration process involves the use of a wireless calibration device which is a piece of hardware that performs the following tasks. First, the device is used to define the reference points during calibration—the location of the device defines a reference point, whose coordinates are automatically determined by equipping the device with a positioning sensor (e.g., ultrasound-based Cricket). Second, the device also serves as a point object for pictures taken by the camera sensors. To ensure that the device can be automatically detected in an image by vision processing algorithms, we equip the device with a bright LED sensor (which then serves as the point object in an image). Third, the device serves as a “wireless remote” for taking pictures during the calibration phase. The device is equipped with a switch that triggers a broadcast packet on the wireless channel. The packet contains the coordinates of the device at that instant and includes a image capture command that triggers a snapshot at all camera sensors in its wireless range.

Given such a device, the protocol works as follows. A human assists the calibration process by walking around with the calibration device. The protocol involves holding the device at random locations and initiating the trigger. The trigger broadcast a packet to all cameras in the range with a command to take a picture (if the sensor node is asleep, the trigger first wakes up a node using a wakeup-on-wireless algorithm). The broadcast packet also includes the coordinates of the current position of the device. Each camera then processes the picture to determine if the LED of the calibration device is visible to it. If so, the pixel coordinates of the device and the transmitted coordinates of the reference point are recorded. Otherwise the camera simply waits for the next trigger. When at least four reference points become available, the sensor node then processes this data to determine the location, orientation and range of the camera. These parameters are then broadcast so that neighboring cameras can subsequently use them for determining the amount of overlap between cameras. Once a camera calibrates itself, a visual cue is provided by turning on an LED on the camera so that the human assistant can move on to other sensors.

## V. AN OBJECT TRACKING APPLICATION

In general, the accuracy desired from the calibration phase depends on the application that will subsequently use this calibrated sensor network. To determine how calibration errors impact application accuracy, we consider a simple object localization and tracking example. This scenario assumes that the calibrated sensor network is used to detect external objects and track them as they move through the environment. Tracking is performed by continuously computing the coordinates of the moving object. A camera sensor network can employ triangulation techniques to determine the location of an object—if an object is simultaneously visible from at least two cameras, and if the locations and orientations of these cameras are known, then the location of the object can be calculated by




Fig. 5. Object localization using two cameras.

taking pictures of the object and using its pixel coordinates to compute its actual location.

To see how this is done, consider Figure 5 that depicts an object  $O$  that is simultaneously visible in cameras  $C_1$  and  $C_2$ . Since both cameras are looking at the same object, the lines connecting the center of the cameras to the object, should intersect at the object  $O$ . Since the locations of each camera is known, a triangle  $C_1OC_2$  can be constructed as shown in the figure. Let  $D_1$  and  $D_2$  denote the distance between the object and the two cameras, respectively, and let  $D_{12}$  denote the distance between the two cameras. Note that  $D_{12}$  can be computed as the Euclidean distance between the coordinates  $C_1$  and  $C_2$ , while  $D_1$  and  $D_2$  are unknown quantities. Let  $\theta_1$ ,  $\theta_2$  and  $\phi$  denote the internal angles of the triangle as shown in the figure. Then the Sine theorem for a triangle from elementary trigonometry states that

$$\frac{D_1}{\sin(\theta_1)} = \frac{D_2}{\sin(\theta_2)} = \frac{D_{12}}{\sin(\phi)} \quad (8)$$

The angles  $\theta_1$  and  $\theta_2$  can be computed by taking pictures of the object and using its pixel coordinates as follows. Suppose that the object projects an image at pixel coordinates  $(-px_1, -py_1)$  at camera  $C_1$ . Let  $f_1$  denote the focal length of camera  $C_1$ . Then projection vector  $\vec{v}_1 = (px_1, py_1, f)$  is the vector joining the pixel coordinates to the center of the lens and this vector lies along the direction of the object from the camera center. If  $\vec{v}$  is the vector along the direction of line connecting the two cameras, the angle  $\theta_1$  can be calculated using the vector dot product:

$$\vec{v} \cdot \vec{v}_1 = |\vec{v}| |\vec{v}_1| \cos(\theta_1) \quad (9)$$

The angle  $\theta_2$  can be computed similarly and the angle  $\phi$  is next determined as  $(180 - \theta_1 - \theta_2)$ .

Given  $\theta_1$ ,  $\theta_2$  and  $\phi$  and the distance between two cameras  $D_{12}$ , the values of  $D_1$  and  $D_2$  can be computed using the Sine theorem as stated above.

Given the distance of the object from the cameras (as given by  $D_1$  and  $D_2$ ) and the direction along which the object lies (as defined by the projection vectors  $\vec{v}_1$  and  $\vec{v}_2$ ), the object location can be easily computed. Note that the orientation matrices of the cameras must also be accounted for when determining the world coordinates of the object using each camera. In practice, due to calibration errors, the object location as estimated by the two cameras are not identical. We calculate the mid-point of the two estimates as the location of the object.

Thus, two overlapping cameras can coordinate with one another to triangulate the location of an external object. We will use this object localization application in our experiments




Fig. 6. *Snapshot* hardware components.

to quantify the impact of calibration errors on the application tracking error.

## VI. SNAPSHOT ERROR ANALYSIS

As described in Section IV, locations of reference points are estimated using a positioning system (ultrasound based Cricket locationing system) which are further used for calibration. The estimated locations of reference points have uncertainties due to errors in ultrasound based range estimates. The average location error using Cricket (measured in terms of Euclidean distance) is empirically estimated to be 3-5 cm. The error in reference point locations impacts the calculated calibration parameters and we study the sensitivity of calibrated parameters to these errors. Consider four reference points with true locations  $(x_1, y_1, z_1)$ ,  $(x_2, y_2, z_2)$ ,  $(x_3, y_3, z_3)$  and  $(x_4, y_4, z_4)$  which estimate the location of the camera as  $(x_c, y_c, z_c)$  and orientation angles as  $\alpha$ ,  $\beta$  and  $\gamma$ . Further, we assume that the error in each dimension of the reference point location is specified by a normal distribution  $\mathcal{N}(0, \sigma^2)$ , with zero mean and variance  $\sigma^2$ . Given  $n$  reference points, an error component is added to each reference point  $(x_i, y_i, z_i)$  as follows,

$$x'_i = x_i + e_1; \quad y'_i = y_i + e_2; \quad z'_i = z_i + e_3;$$

where,  $e_i$  is randomly sampled from the normal distribution  $\mathcal{N}$ . The  $\binom{n}{4}$  updated reference point subsets are then used to compute the camera location  $(x'_c, y'_c, z'_c)$  and orientation parameters  $\alpha'$ ,  $\beta'$ ,  $\gamma'$ . The relative error in calibration as result of the error in reference point locations is measured as,

$$loc_{err} = \sqrt{(x'_c - x_c)^2 + (y'_c - y_c)^2 + (z'_c - z_c)^2} \quad (10)$$

$$pan_{err} = \|\alpha' - \alpha\| \quad (11)$$

$$tilt_{err} = \|\beta' - \beta\| \quad (12)$$

$$roll_{err} = \|\gamma' - \gamma\| \quad (13)$$

where,  $loc_{err}$  is the relative location error, measured as the Euclidean distance between the estimated camera locations and  $pan_{err}$ ,  $tilt_{err}$  and  $roll_{err}$  are the relative orientation errors of pan, tilt and roll angles respectively.

The sensitivity of the calibration parameters is estimated by measuring the relative location and orientation errors for different (increasing) variances of the error distribution. We test sensitivity for *random error*—errors in each dimension of every reference point are randomly sampled and *correlated error*—errors for each dimension are sampled randomly but are same for all reference points. We present the experimental results of the sensitivity analysis in Section VIII.

## VII. SNAPSHOT IMPLEMENTATION

This section describes our prototype implementation.

### A. Hardware Components

The *Snapshot* wireless calibration device is a Mote-based Cricket ultrasound receivers equipped with a LED that turns itself on during calibration (see see Figure 6(a)). We assume that the environment is equipped with Cricket reference beacons, which are used by a Cricket receiver to compute its location coordinates during calibration [13].

We use two types of camera sensors in our experiments: the CMUcam vision sensor [16] and a Sony webcam. The CMUcam comprises of a OV6620 Omnivision CMOS camera and a SX52 micro-controller and has a resolution of 176x255. In contrast, the Sony webcam has a higher resolution of 640x480. We use the high resolution webcam to quantify the loss in accuracy when calibrating low-resolution cameras such as the CMUcam. Although beyond the scope of the current paper, our ongoing work focuses on calibrating a second low-resolution camera sensor, namely the Cyclops [11]. All camera sensors are connected to Intel Stargates (see Figure 6(b)), which is a PDA-class sensor platform equipped with a 400MHz XScale processor and running the Linux operating system. Each Stargate also has a Intel Mote connected to it for wireless communication with our Mote-based calibration device.

Finally, we use a digital compass, Sparton 3003D, to quantify the orientation error during calibration. The compass has resolution of 0.1 degrees and accuracy of 0.3 degrees.

### B. Software Architecture

Our Mote-based calibration device runs TinyOS with the Cricket toolkit. The *Snapshot* software on the Mote is simple: each human-initiated trigger causes the Mote to determine its coordinates using Cricket, which are then embedded in an “image-capture” trigger packet that is broadcast to all nodes, using the wireless radio.

**Camera Calibration Tasks:** Every time a trigger packet is received from the calibration device, the Stargate sends a set of commands over the serial cable to capture an image from the CMUcam. The image is processed using a vision-based recognition algorithm; our current prototype uses background subtraction and a connected components algorithm [15] to detect the presence of the calibration device LED. If the device is found, the pixel coordinates of the LED and the Cricket coordinates of the Mote are stored as a new reference point. Otherwise the image is ignored.

Once four reference points become available, the Stargate estimates location, orientation and range of the camera. A non-linear solver based on the interior-reflective Newton method [5], [6] is used to estimate the camera location. We use the methods discussed in Section III to eliminate false solutions, and iteratively refine the location estimate.

**Object Localization and Tracking:** Finally, we implement our object localization and tracking (described in Section V) application on the Stargates. If an object is simultaneously viewed by two cameras, the cameras exchange their parameters, location and orientation, and the objects projection coordinates on its image plane. This information is used by each camera to localize the object and estimate its location.

Continuous localization can be used at each node to track an object of interest.

## VIII. EXPERIMENTAL EVALUATION

In this section, we evaluate the efficacy of Snapshot, quantify the impact of using Cricket, and evaluate the impact of Snapshot on our object tracking application.

### A. Experimental Setup

The setup to evaluate the accuracy and sensitivity to system parameters of *Snapshot* consisted of placing the two types of cameras, CMUcam and the Sony MotionEye webcam, at several locations. To simplify accurate location measurements we marked a grid to place the position sensor objects. Each camera took several pictures to estimate the parameters. The difference between the estimated parameter value and the actual value is reported as the measurement error. The Cricket sensors on the objects received beacons from a set of pre-calibrated Cricket sensor nodes placed on the ceiling of a room. The digital compass was attached to the two cameras in order to measure the exact orientation angles.

### B. Camera Location Estimation Accuracy

To evaluate *Snapshot's* performance with camera location estimation, we place tens of reference points in the space, and take pictures of these reference points at different locations and orientations. We measure the location of these reference points by hand (referred as without Cricket) which can be considered as the object's real location and by Cricket [13] (referred as with Cricket) where we observed a 2–5cm error.

For each picture, we take all the combinations of any four reference points in view (not any 3 points in the same line), and estimate camera's location accordingly. We consider the distance between the estimated camera's location and the real camera's location as the location estimation error.

As shown in Figure 7(a), our results show: (i) the median errors using webcam without Cricket and with Cricket are 4.93cm and 9.05cm, respectively; (ii) the lower quartile and higher quartile errors without Cricket are 3.14cm and 7.13cm; (iii) the lower quartile and higher quartile errors with Cricket are 6.33cm and 12.79cm; (iv) the median filter (referred as M.F.) improves the median error to 3.16cm and 7.68cm without Cricket and with Cricket, respectively.

Figure 7(b) shows: (i) median errors using CMUcam without Cricket and with Cricket are 6.98cm and 12.01cm, respectively; (ii) the lower quartile and higher quartile errors without Cricket are 5.03cm and 10.38cm; (iii) the lower quartile and higher quartile errors with Cricket are 8.76cm and 15.97cm; (iv) the median filter improves the median error to 5.21cm and 10.58cm without Cricket and with Cricket, respectively.

1) *Effect of Iteration on Estimation Error*: As our protocol proceeds, the number of available reference points increases. As a result, the number of combinations of any four reference points also increases, and we have more location estimations available for the median filter. Consequently, we can eliminate tails and outliers better. In this section, we study the effect of the iterations of our protocol's runs on camera location

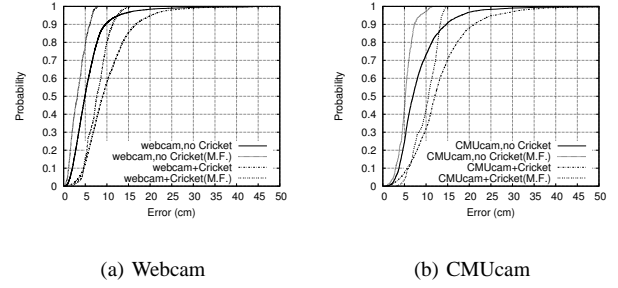


Fig. 7. Empirical CDF of error in estimation of camera location.




Fig. 8. Effect of number of reference points on location estimation error.


estimation error by plotting the median versus the number of available reference points.

Figure 8 shows: (i) the median errors using webcam drop from 4.93cm to 2.13cm and from 9.05cm to 6.25cm as the number of reference points varies from 4 to 16 for without and with Cricket, respectively; (ii) the median errors using CMUcam drop from 6.98cm to 2.07cm and from 12.01cm to 9.59cm as the number of reference points varies from 4 to 16 for without and with Cricket, respectively. The difference in the location estimation errors (with and without Cricket) are due to the position error estimates in Cricket and also due to errors in values of camera intrinsic parameters.

### C. Camera Orientation Estimation Error

Next, we evaluate *Snapshot's* accuracy with estimation of camera orientation parameters. We used the two cameras, the CMUcam and the Sony MotionEye webcam, to capture images of reference points at different locations and different orientations of the camera. We used estimated location of the camera based on exact locations on reference points and Cricket-reported locations of reference points to estimate the orientation parameters of the camera. The orientation of the camera was computed using the estimated camera location. We compared the estimated orientation angles with the measured angles to calculate error. Figure 9(a) shows the CDF of the error estimates of the pan, tilt and roll orientations respectively using the CMUcam camera. Figure 9(b) show the CDF of the error of the three orientations using Cricket for location estimation. The cumulative error plots follow the same trends for each of the orientation angles. The median roll orientation error using Cricket and without Cricket for camera location estimations is 1.2 degrees. In both cases, the 95th percentile error is less than 5 degrees for the pan and tilt orientation and less than 3 degrees for the roll orientation. The slight





(a) CDF without Cricket

(b) CDF with Cricket

Fig. 9. Empirical CDF of error in estimating orientations with the CMUcam.

discrepancies in the error measurement of the two cases is due to the use of the digital compass to measure the orientation of the camera.


Thus, we conclude the Cricket’s positioning errors do not add significant errors in estimation of camera orientation parameters. In our experiments, we find that a median location estimation error of 11cm does not affect the orientation estimation significantly.

#### D. Sensitivity Analysis

As described in Section VI, we evaluate the sensitivity of calibrated parameters to uncertainty in reference point locations. We varied the standard deviation of the error distribution in each dimension from 1cm to 8cm and numerically computed its impact on the calibration parameters. As shown in Figure 10(a), the estimated locations are less sensitive to the correlated error, but are highly sensitive to the random error. Further, the results in Figure 10(b) shows that: (i) orientation estimation is insensitive to the correlated error, the mean error is always very close to zero; and (ii) the orientation estimation is very sensitive to the random error, the mean error increases by a factor of four as the standard deviation increases from 1cm to 8cm. The calibrated parameters are less sensitive to correlated errors as all reference points have the same error magnitudes and the camera location shifts in the direction of the error without affecting the estimated orientation. With random errors in each dimension of the reference points, all reference points shift to different directions by different offsets, and as a result, calibration errors are larger. However, the error in a real Cricket system is neither correlated nor random, it is somewhere between these two cases, and has intermediate sensitivity. The previous experimental results verify this hypothesis.

#### E. Object Localization

In this section, we study the performance of object localization using *Snapshot*. We use *Snapshot* to estimate camera locations and their orientations, and then in turn use the calibrated parameters to triangulate an object via the technique described in Section V. Similar to Section VIII-B, we use the empirical CDF of object’s location estimation error to measure the performance. Our results (see Figure 11) show that: (i) the median localization error using webcams is 4.94cm and 5.45cm without and with Cricket, respectively; (ii) the median localization error using CMUcams is 11.10cm and 11.73cm



(a) Location

(b) Orientation (CMUcam)

Fig. 10. Sensitivity of estimation to uncertainty in reference point location.




Fig. 11. Empirical CDF of error in estimation of object’s location.

without and with Cricket, respectively; (iii) localization without Cricket outperforms localization using Cricket for all cameras; and (iv) localization using webcams outperforms that with the CMUcams due to its higher fidelity.

#### F. Runtime Scalability

Task	Duration(ms)
Snap Image	$178 \pm 2$
Recognize Object Location	$52 \pm 0.1$
Location Estimation	$18365 \pm 18$

Fig. 12. Runtime of different calibration tasks.

Using our prototype implementation of we measure the runtime of the *Snapshot* protocol. Figure 12 reports runtime of different tasks of the *Snapshot* calibration protocol executing on the Intel Stargate platform with the camera attached to a USB connector (the transfer of an image on the serial cable with the CMUcam requires additional time). As seen from the table, the location estimation task which uses a non-linear solver, has the highest execution time. The time to calibrate an individual camera is,  $4 \times (178 \text{ ms} + 52 \text{ ms})$  – time to snap four images and recognize the location of object in each and 18365 ms for the location and orientation estimation, which is total time of 19.285 seconds. Thus, with a time of approximately 20 seconds to calibrate a single camera, *Snapshot* can easily calibrate tens of cameras on the scale of a few minutes.

## IX. RELATED WORK

Camera calibration using a set of known reference points is well studied in the computer vision community. Methods developed in [8], [18], [23], [24] are examples of techniques that estimate both the intrinsic and extrinsic parameters of a camera using a set of known reference points. These efforts present techniques to estimate the complete set of twelve parameters and also for a partial set (extrinsic parameters) of

camera parameters. *Snapshot*, designed to estimate only the extrinsic parameters, draws inspiration from these techniques and extends them to suit resource-constrained sensor network environments and works with uncertainty in reference point locations. A recent effort [21] also estimates only the extrinsic parameters with four reference points, with the requirement that three out of the four are co-linear. *Snapshot* is a more general technique and does not impose such a requirement. Further, unlike [21], we demonstrate the feasibility of our approach through a detailed experimental evaluation.

Several studies have focused on the design and implementation of camera sensor networks. SensEye [12] is a multi-tier camera sensor network that exploits heterogeneous cameras sensors and computing platforms to provide benefits over single-tier camera sensor networks. Panoptes [19] is an example of video sensor node that implements power-efficient video delivery mechanisms and techniques to handle long periods of disconnections. Panoptes nodes can be incorporated in the SensEye architecture and can also be used in design of single-tier networks [9]. [20] presents an architecture to quickly compose sensor networks incorporating multi-modal sensors (along with video sensors) with optimized application-specific algorithms. *Snapshot* can be used to automatically calibrate camera sensors used in the networks described above. Several other efforts [10], [14] have also studied the problem of video surveillance but without considering resource constraints of camera sensor networks.

Localization is well studied in the sensor networks community [7], [17], [22]. All these techniques assume a sensor node capable of position estimation. For example, a temperature sensor can use its RF wireless communication link to send and receive beacons for location estimation. *Snapshot* does not require any position estimation capability on the nodes and directly uses the imaging capability of the cameras for localization and calibration.

Several positioning and self-localization systems have been proposed in the literature. Active Badge [1] is a locationing system based in IR signals, where badges emit IR signals are used for location estimation. A similar successor system based on ultrasound signals is the Active Bat [2] system. Several other systems use RF signal strength measurements, like RADAR [3], for triangulation based localization. While most of these techniques are used indoors, GPS [4] is used for outdoor localization. While any of these methods can be used by the *Snapshot* calibration device instead of the Cricket, each has its own advantages and disadvantages. Based on the environment and desired error characteristics a suitable positioning system can be chosen.

## X. CONCLUSIONS

In this paper, we presented *Snapshot*, an automated calibration protocol that is explicitly designed and optimized for sensor networks. Our techniques draw upon principles from vision, optics and geometry and are designed to work with low-fidelity, low-power camera sensors that are typical in sensor networks. Our experiments showed that *Snapshot* yields an error of 1-2.5 degrees when determining the camera

orientation and 5-10cm when determining the camera location. We argued that this is a tolerable error in practice since a *Snapshot*-calibrated sensor network can track moving objects to within 11cm of their actual locations. Finally, our measurements showed that *Snapshot* can calibrate a camera sensor within 20 seconds, enabling it to calibrate a sensor network containing tens of cameras within minutes.

## ACKNOWLEDGMENT

This work was supported in part by National Science Foundation grants EEC-0313747, CNS-0219520, CNS-052072 and EIA-0098060.

## REFERENCES

- [1] Andy Harter and Andy Hopper. A Distributed Location System for the Active Office. *IEEE Network*, 8(1), January 1994.
- [2] Andy Ward and Alan Jones and Andy Hopper. A New Location Technique for the Active Office. *IEEE Personal Communications*, 4(5):42–47, October 1997.
- [3] P. Bahl and V. N. Padmanabhan. RADAR: An in-building RF-based user location and tracking system. In *INFOCOM*, pages 775–784, 2000.
- [4] R. Bajaj, S. L. Ranaweera, and D. P. Agrawal. Gps: Location-tracking technology. *Computer*, 35(4):92–94, March 2002.
- [5] T. F. Coleman and Y. Li. On the convergence of reflective newton methods for large-scale nonlinear minimization subject to bounds. *Mathematical Programming*, 67(2):189–224, 1994.
- [6] T. F. Coleman and Y. Li. An interior, trust region approach for nonlinear minimization subject to bounds. *SIAM Journal on Optimization*, 6(2):418–445, 1996.
- [7] T. He, C. Huang, B. Blum, J. Stankovic, and T. Abdelzaher. Range-Free Localization Schemes in Large Scale Sensor Networks. In *MOBICOM*, 2003.
- [8] B. K. P. Horn. *Robot Vision*. The MIT Press, first edition, 1986.
- [9] L. Jiao and Y. Wu and G. Wu and E. Y. Chang and Y. F. Wang. The Anatomy of a Multi-camera Security Surveillance System. *ACM Multimedia System Journal Special Issue*, October 2004.
- [10] M. Chu and J. E. Reich and F. Zhao. Distributed Attention for Large Video Sensor Networks. In *IDS*, 2004.
- [11] M. Rahimi and D. Estrin and R. Baer and H. Uyeno and J. Warrior. Cyclops, Image Sensing and Interpretation in Wireless Networks. In *SensSys*, 2004.
- [12] P. Kulkarni and D. Ganesan and P. Shenoy. SensEye: A Multi-tier Camera Sensor Network. In *ACM Multimedia*, 2005.
- [13] N. B. Priyantha, A. Chakraborty, and H. Balakrishnan. The cricket location-support system. In *MobiCom*, 2000.
- [14] R. Collins and A. Lipton and T. Kanade. A System for Video Surveillance and Monitoring. In *ANS 8th International Topical Meeting on Robotics and Remote Systems*, 1999.
- [15] A. Rosenfeld and J. L. Pfaltz. Sequential Operations in Digital Picture Processing. *J. ACM*, 13(4):471–494, 1966.
- [16] A. Rowe, C. Rosenberg, and I. Nourbakhsh. A Low Cost Embedded Color Vision System. In *IROS*, 2002.
- [17] A. Savvides, C.-C. Han, and M. B. Strivastava. Dynamic fine-grained localization in ad-hoc networks of sensors. In *MOBICOM*, 2001.
- [18] R. Y. Tsai. A Versatile Camera Calibration Technique for High-Accuracy 3D Machine Vision Metrology Using Off-the-Shelf TV Cameras and Lenses. *IEEE T-RA*, 3(4):323–344, August 1987.
- [19] W. Feng and B. Code and E. Kaiser and M. Shea and W. Feng and L. Bavoil. Panoptes: A Scalable Architecture for Video Sensor Networking Applications. In *ACM Multimedia*, 2003.
- [20] W. Feng and N. Bulusu and W. Feng. Dissecting the Video Sensing Landscape. In *ACM NOSSDAV*, 2005.
- [21] F. Y. Wang. A Simple and Analytical Procedure for Calibrating Extrinsic Camera Parameters. *IEEE T-RA*, 20(1):121–124, February 2004.
- [22] K. Whitehouse and D. Culler. Calibration as Parameter Estimation in Sensor Networks. In *WSNA*, 2002.
- [23] J.-C. Yuan. A general photogrammetric method for determining object position and orientation. *IEEE T-RA*, 5(2):129–142, 1989.
- [24] Z. Y. Zhang. A Flexible New Technique for Camera Calibration. *IEEE TPAMI*, 22(11):1330–1334, November 2000.