Proceedings of IMECE 2001: ASME International Mechanical Engineering Congress, New York, New York, Nov. 11-16, 2001

DISPATCHING OF COORDINATED PROXIMITY SENSORS FOR OBJECT SURVEILLANCE

Michael D. Naish **Elizabeth A. Croft**[†] **Beno Benhabib**

Computer Integrated Manufacturing Laboratory
University of Toronto
Toronto, Ontario M5S 3G8
Canada
[naish, beno]@mie.utoronto.ca

[†]Industrial Automation Laboratory
University of British Columbia
Vancouver, British Columbia V6T 1Z4
Canada
ecroft@mech.ubc.ca

ABSTRACT

This paper presents a method of selecting and positioning groups of sensors in a coordinated manner for the surveillance of a maneuvering object. The object trajectory is discretized into a number of demand instants (data acquisition times) to which groups of sensors are assigned, respectively. Heuristic rules are used to evaluate the suitability of each sensor for servicing (observing) a demand instant, determine the composition of the sensor group, and, in the case of dynamic sensors, specify the position of each sensor with respect to the object. This approach aims to improve the quality of the surveillance data in three ways: (1) the assigned sensors are maneuvered into "optimal" sensing positions, (2) the uncertainty of the measured data is mitigated through sensor fusion, and (3) the poses of the unassigned sensors are adjusted to ensure that sensing-system can react to object maneuvers. Simulations with proximity sensors demonstrate the advantages of dispatching dynamic sensors over similar static-sensor systems.

INTRODUCTION

This paper considers the principles of dispatching applied to the surveillance of a maneuvering object. Just as the system response for a taxi company can be improved by effective dispatching, the quality of the sensor data acquired by a set of sensors can be improved through appropriate selection and positioning; however, sensors introduce an additional consideration. Dispatching multiple sensors (as opposed to only one) to observe a moving target at a particular location provides an opportunity to significantly improve the quality and robustness of the data. Specifically, sensor fusion may be used to combine information from these multiple, coordinated sensors into a single representation [1].

Strategies and techniques used for scheduling and dispatching service vehicles have been investigated by a number of researchers [2–7]. Determining which vehicle will be sent to service a particular demand (fixed pickup/drop-off location) is the role of the dispatcher. Typically, each vehicle is assessed for assignment to a demand based on a set of evaluation criteria. One can note that, in the two exemplary approaches reviewed below as well as others, the assignment of a vehicle to a particular demand also includes its (desired) position. Once assigned, the vehicle must move to the (fixed) position of the demand.

A rolling-horizon algorithm was developed by Psaraftis [6] to assign cargo to ships destined to various ports. Cargoes are tentatively assigned to eligible ships; permanent assignments are made for those cargoes that fall at the beginning of the rolling horizon (more immediate events are known with greater certainty). The rolling horizon is then shifted to the next time step. For the dynamic vehicle allocation (DVA) problem, Powell [7] discusses and compares a number of optimization approaches (e.g., deterministic and stochastic networks, Markov decision theory).

However, optimization techniques used in service vehicle dispatching are not, in general, particularly well suited to real-time surveillance applications. The alternate approach is to utilize heuristics. Though not a rigorous treatment of the problem (i.e., stability and optimality cannot be guaranteed), heuristics can provide a tractable and timely solution.

On-line dispatching for surveillance involves both selecting an appropriate subset of sensors to be used in a data fusion process as well as maneuvering all the sensors in response to the motion of the object. Namely, the sensors provide information of sufficient quality for the task at hand while ensuring adequate response to future object maneuvers (keeping all sensors "in the game"). This is a reactive procedure; therefore, no absolute condition of optimality can be practically imposed. For static (fixed-position) sensors, dispatching considers which sensors should be included in the sensor fusion process during data-acquisition. Dynamic (mobile) sensor dispatching must, in addition, address the motion of each sensor.

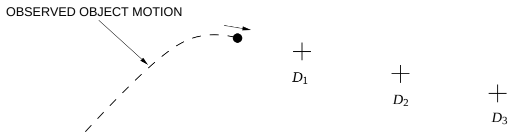

Figure 1. Future demand points.

The system proposed in [8] uses a "generate-and-test" method to determine the best sensor position to optimize both feature visibility and measurement reliability. Zhang (1995) aimed to minimize the magnitude of the sensor measurement co-variances in optimally placing multiple sensors to be used for sensor fusion. The system proposed in [10] proposes to handle moving objects by discretizing time and computing new viewing configurations for each time interval while attempting to minimize changes in sensor position from one time interval to the next. A different approach is taken by Abrams *et al.* [11], which attempts to find viewpoints that satisfy the task constraints over the entire task interval; if none exist, the interval is divided until satisfactory viewpoints are found. Matsuyama *et al.* [12] determine camera layouts (in 2-D) through optimization (off-line); camera actions and temporal camera switching and coordination are determined by heuristics (on-line) to account for deviations. Finally, an agent-based approach is taken in [13]. The sensors are fixed in space and the workspace is divided into a number of sectors, each with a managing agent. The managing agents attempt to recruit sensors (each with a corresponding agent) to scan the workspace for targets and, once found, provide synchronized measurements that may be fused to estimate the real-time object location.

The objective of this paper is to utilize the general dispatching approaches previously proposed for service vehicles to better survey a moving object using multiple sensors. This approach is targeted towards object surveillance for subsequent robotic interception.

PROBLEM FORMULATION

Surveillance

Surveillance of a maneuvering object involves providing estimates of particular object parameters at predetermined times or positions along the object trajectory. Herein, the parameter of interest is the Cartesian position of a single point on the target. The times at which information is desired are fixed. These predetermined times are referred to as demand instants, t_j. Often, the interval between each demand instant is a constant value, Δ. Thus, lacking *a priori* knowledge of the object trajectory, surveillance refers herein to the prediction of the position of the target at each demand instant, Fig. 1. The position of the target for a particular demand instant is referred to as a demand point, D_j. As the prediction accuracy improves, the estimation of the demand point would change; however, the demand instant remains constant.

In this paper, the sensors that comprise the sensing-system may be static or dynamic; the speed and maneuverability of dynamic sensors are considered to be inferior to the object. As a result, it would be necessary to distribute the sensors throughout the workspace to ensure a degree of data consistency over the entire object trajectory.

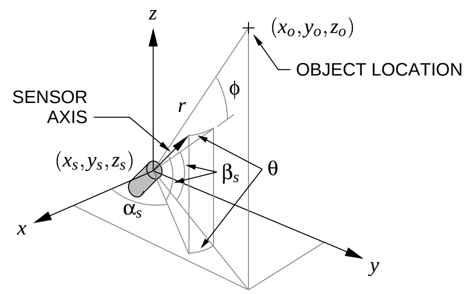

Figure 2. Measurement variables for a proximity sensor.

Data Acquisition

Herein, a three-dimensional proximity measurement is defined as the range, bearing, and elevation to the object from a known sensor pose. For the sensor shown in Fig. 2, the range, r, is the linear distance between the object and the sensor frame. The bearing, θ, is the radial difference between the orientation of the sensor axis with respect to the x-z plane, α_s, and the object position. The elevation, ϕ, is the radial difference between the orientation of the sensor axis with respect to the x-y plane, β_s, and the object position.

There exist a number of different non-contact sensors which may be used to measure the proximity of an object. These generally utilize laser-triangulation, phase- or amplitude-modulation based electro-optical transducers, or ultrasonic transducers [14]. A pair of calibrated CCD cameras can also be used in conjunction with stereo image-processing techniques to estimate the distance between an object and a centre point between the cameras [15].

The task of fusing multisensor parametric data to yield an improved estimate of the state of an entity may be addressed by a number of different techniques. Specific methodologies suitable for parameter estimation include the least squares estimator (LS) and its variations: WLS, BWLS, MLE, MSE [16], geometric fusion [17], and the Kalman Filter (KF) [18, 19].

Dispatching

Within the context of optimal dispatching, sensor fusion does not need to combine data from all of the sensors in the system. Instead, a subset of sensors ($k \leq n$, where k is the subset size and n is the total number of sensors) is selected, or assigned, to survey the object at a particular demand instant, Fig. 3. Herein, this group of sensors is referred to as a fusion subset. A sensor may belong to multiple fusion subsets.

The general dispatching problem addressed in this paper may be stated as: *Given a set of sensors and a set of time-varying*

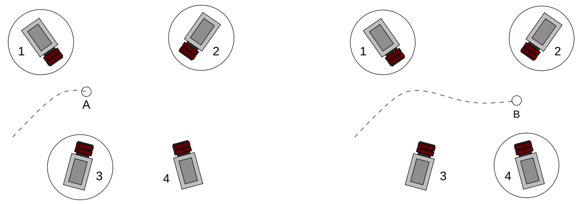

(a) Fusion subset: 1-2-3. (b) Fusion subset: 1-2-4.

Figure 3. Coordination strategy.

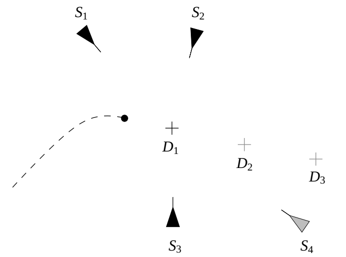

Figure 4. Assignment and positioning of S_1, S_2, and S_3 to D_1. S_4 is unassigned.

demand points (based on the predicted motion of a maneuvering object at the corresponding set of demand instants), determine the subsets of sensors (and their corresponding poses) that will (optimally) sense the object at each demand instant and, furthermore, ensure that the remaining sensors are adequately distributed throughout the workspace for possible future use.

Dispatching may be accomplished using two complementary strategies. First, a "coordination strategy" specifies which sensors will be used for surveillance at each and every demand instant. This selection should be based on a logical search through the sensor set, using an objective function, to evaluate the fitness of each sensor. Next, a "positioning strategy" specifies the pose of each of the sensors, whether it has been assigned to the first demand instant at hand, or to any one of the future demand instants.

A HEURISTIC DISPATCHING APPROACH

There exist two primary differences between a global optimization method, carried out off-line, and the on-line, heuristic dispatching solution to the multisensor surveillance problem proposed in this paper. First, the heuristic approach does not attempt to optimize fusion subsets in a combinatoric manner, but rather considers each sensor individually. Secondly, instead of considering all of the demand instants concurrently (i.e., the complete object trajectory), each is considered sequentially, limited to a finite segment (the rolling horizon) of the object trajectory. This segment is defined by a limited number of demand instants, $t_1 \ldots t_m$. The object location is *predicted*[1] at each of these demand instants. Note that, the absolute clock time of each demand instant is fixed; however, the estimated location of the object at each—the demand point—would vary over time as the predictions of the object motion become more accurate.

In the proposed methodology, an "assignment and positioning" search method considers only the first demand point (to which a subset of size k sensors will always be assigned), while a "preassignment and prepositioning" search method considers the remaining demand instants on the rolling horizon, as discussed below.

[1]Herein, it is assumed that, the prediction of the demand points is performed by an external prediction subsystem.

Assignment and Positioning

Given the first demand point on the rolling horizon, D_j, the assignment and positioning method selects an optimal subset of sensors (of size k) from the set of all sensors, n, to service this demand instant (i.e., coordination). The search method also specifies a desired pose for each sensor at the time of data acquisition (i.e., positioning), Fig. 4. Assignment occurs once during each search interval (the time between the previous and current demand instants). (It is triggered by an object entering the workspace or the completion of the previous search interval.) Namely, the selection of the sensor set cannot be altered until the first demand instant on the rolling horizon has been serviced; however, the pose of each sensor may be altered in real-time. Pose adjustment is handled by a replanning method described later in this paper.

The general approach to the assignment and positioning of sensors for a demand point can be summarized as follows:

1. Predict the object's pose, D_j, with respect to the world coordinate frame, at the demand instant, t_j.
2. For every sensor, $S_i,\ i = 1 \ldots n$:
   (a) Determine its best achievable pose with respect to D_j, and
   (b) Assess the corresponding (single sensor) visibility of D_j, v_s, from the best achievable pose.
3. Rank all S_i according to the achievable visibility, v_s, from highest to lowest.
4. Assign the top k ranked sensors to t_j. (The desired pose of each assigned sensor is the best achievable pose determined in Step a.)

Preassignment and Prepositioning

Once an assignment has been made for the first demand instant of the rolling horizon, a preassignment and prepositioning search method selects sensors for pseudo-assignment to subsequent demand instants. The objective here is to position the unassigned sensors in anticipation of future service requirements. All sensors are considered for pseudo-assignment; however, only sensors that have not been assigned by the Assignment & Positioning Module (sub-system) may be pseudo-assigned to a future demand instant, Fig. 5. Note that, previously assigned sensors are

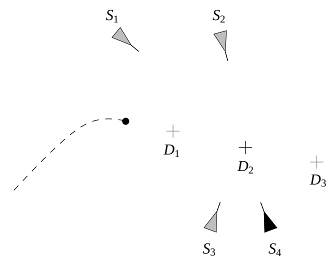

Figure 5. Preassignment and prepositioning of only S_4 to D_2.

evaluated from their (assigned) desired poses and may, in effect, be assigned to multiple future demand instants, though their future position will not be determined until the next search interval. This approach aims to service each demand with the sensors that can provide the highest quality data, rather than by those that are least utilized. Additional demand instants are considered until either all sensors have been pseudo-assigned, the maximum number of future demand points, m (equal to the rolling horizon size), has been reached, or the search interval has expired.

The general preassignment and prepositioning procedure can be summarized as follows:

1. Let $p = 1$
2. Predict the object's pose, D_{j+p}, with respect to the world coordinate frame, at the demand instant, t_{j+p}.
3. For each sensor, $S_i,\ i = 1 \ldots n$:
   (a) Determine its best achievable pose with respect to D_{j+p}, and
   (b) Assess the (single sensor) visibility of D_{j+p}, v_s, from the best achievable pose.
4. Rank S_i according to best visibility, v_s, from highest to lowest.
5. For the top k ranked sensors, $S_i,\ i = 1 \ldots k$, determine whether any has been assigned to an earlier demand point. Those that have not been are assigned to t_{j+p}.
6. Determine whether there remain any unassigned sensors and whether $p < m$ (i.e., have all points on the rolling horizon been considered). If both are true, let $p = p + 1$ and return to Step 2.

The above approach is nearly identical to that used for assignment and positioning. The searches, however, are separated here to emphasize that different criteria may be used for assignment and preassignment, depending on the requirements of the application and the capabilities of the system. For example, a faster algorithm may be used for preassignment, allowing more elaborate measures to be taken for the evaluation of sensors during the assignment search.

In conclusion to the above discussion of assignment & positioning and preassignment & prepositioning, one must note that the success of sensor dispatching would be dependent on the initial pose of each sensor within the workspace. This is especially true for sensors with limited dynamic capabilities with respect to the object. In general, the slower the sensors are, the more widely distributed through the workspace the sensors should be. Thus, if any part of the object trajectory is known *a priori*, an optimal initial configuration of the sensors must be determined. Techniques used for the determination of the initial sensing-system configuration are beyond the scope of this paper, though the effect of proper selection is discussed through exemplary simulations.

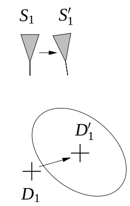

Figure 6. Replanning for D_1. D_1' is the updated prediction of the demand point.

Replanning

Replanning is initiated upon the assignment of sensors to a demand point. The desired pose of each assigned or pseudo-assigned sensor is continuously reassessed during the search interval. The goal of replanning is to compensate for errors introduced by the uncertainties associated with the object's predicted demand-point locations. If the newly predicted position lies outside of the desired confidence interval for the demand point at the time of assignment, the demand point is replaced with the new prediction and the desired pose of the sensor is determined anew, Fig. 6.

Visibility Measure for Proximity Sensing

The algorithms used for sensor dispatching presented herein utilize a visibility metric can be used to evaluate the quality of information that a sensor, or a group of sensors, can provide about a demand point. The visibility measure for a single proximity sensor is considered in this paper to be:

$$v_s = \begin{cases} \dfrac{1}{||R||} & \text{if the demand point is unoccluded,} \\ 0 & \text{otherwise,} \end{cases} \tag{1}$$

where $||R||$ is the Euclidean norm of the covariance matrix associated with the sensor measurement. For the proximity sensor in Fig. 2, the variance in r, θ, and ϕ may be expressed as:

$$\sigma_r^2 = \begin{cases} a + b_1(r - r^*)^2 & \text{if } r < r^*,\ r \in \mathcal{R}^+, \\ a + b_2(r) & \text{otherwise,} \end{cases} \tag{2}$$

$$\sigma_\theta^2 = \begin{cases} c + d\,\theta^2 & \text{if } |\theta| < \theta_{\max},\ \theta \in \mathcal{R}, \\ \infty & \text{otherwise,} \end{cases} \tag{3}$$

$$\sigma_\phi^2 = \begin{cases} e + f\,\phi^2 & \text{if } |\phi| < \phi_{\max},\ \phi \in \mathcal{R}, \\ \infty & \text{otherwise,} \end{cases} \tag{4}$$

where a, b_1, b_2, c, d, e, and f are characteristic constants. r^* is the range between the sensor and the object at which the variance is minimal; here, the variance is equal to the constant error a. If the range is small, the variance increases proportional to b_1; if the range is large, the variance increases proportional to b_2. Similarly, for the variance in bearing and elevation, c and e are the constant measurement errors, while d and f represent the increase in variance incurred by deviations of the object position from the sensor axis. θ_{max} and ϕ_{max} limit the field of view of the sensor. Assuming that σ_r^2, σ_θ^2, and σ_ϕ^2 are uncorrelated, the covariance matrix R may be expressed in Cartesian coordinates as follows:

$$R = \begin{bmatrix} \sigma_x^2 & \sigma_{xy} & \sigma_{xz} \\ \sigma_{xy} & \sigma_y^2 & \sigma_{yz} \\ \sigma_{xz} & \sigma_{yz} & \sigma_z^2 \end{bmatrix}, \tag{5}$$

where σ_x^2, σ_y^2, σ_z^2, σ_{xy}, σ_{xz}, and σ_{yz} are functions of σ_r^2, σ_θ^2, σ_ϕ^2, r, θ, ϕ, α_s, and β_s.

Visibility for a subset of k sensors, whose measurements are combined using sensor fusion, is defined as:

$$v_f = \frac{1}{||P||}, \tag{6}$$

where P represents the fused covariance matrix,

$$P = \left[\sum_{i=1}^{k} R'^{-1}_i\right]^{-1}, \tag{6a}$$

and

$$R'_i = \begin{cases} R_i & \text{if demand point is unoccluded,} \\ \emptyset & \text{otherwise.} \end{cases} \tag{6b}$$

Namely, the covariance matrix of the ith sensor, R_i, (and hence its visibility) is considered only if the sensor has an unoccluded view of the demand point.

SIMULATED DISPATCHING EXAMPLES

The examples presented in the following sections serve to illustrate different performance aspects of the proposed approach to dispatching. In particular, surveillance-system performance is evaluated under changing object trajectories and sensor dynamics. For ease of illustration, a 2-D workspace is assumed. For all examples, the object moves at approximately 0.2 m/sec. Demand points are predicted using a Kalman-filter based predictor [20], on the basis of lower-resolution images acquired from an overhead camera, Fig. 7. These observations are corrupted by Gaussian noise with σ = 0.02 m. In each simulation, fusion subsets of size 3 were used with a rolling horizon size of 3 demand instants separated by 0.6 seconds. The sensor parameters were set as follows: $a = 2.5e^{-5}$ m, $b_1 = 1.25e^{-3}$, $b_2 = 6.25e^{-5}$, $r^* = 0.05$ m, $c = 8e^{-5}$ rad, $d = 50$, and $\theta_{max} = \frac{\pi}{4}$ rad.

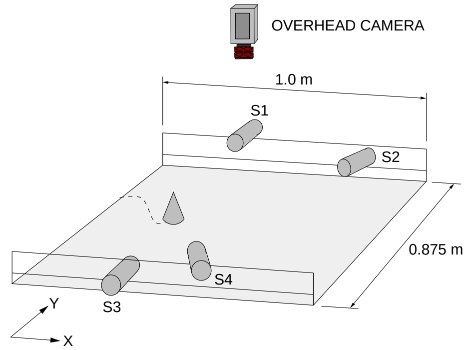

Figure 7. Overview of example workspace configuration.

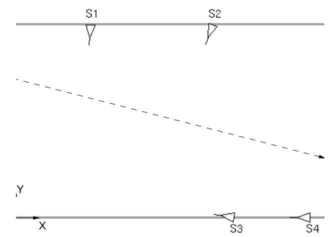

Figure 8. Optimal initial sensing-system configuration for sensors with $\dot{x}$ = 0.750 m/sec and $\dot{\alpha} = \frac{\pi}{3}$ rad/sec.

Dispatching

The following example illustrates how a surveillance system is reconfigured in real-time using the dispatching approach. Fig. 9 shows snapshots of a sample run for a noise-corrupted straight-line object trajectory. Here, the maximum translational velocity of each sensor is $\dot{x} = 0.1$ m/sec; the maximum rotational velocity is $\dot{\alpha} = \frac{\pi}{3}$ rad/sec. The system starts from an initial configuration optimized for these sensors and this object trajectory [21], Fig. 8. (In each figure, the predicted demand point locations are indicated by a cross (+). The actual demand point locations, if the system were capable of perfect prediction, are indicated by circles ($\circ$). Assigned sensors appear as black and preassigned sensors as grey.)

Sensitivity to Dynamic Characteristics of Sensors

The above example (Fig. 9) demonstrates how reconfiguration performs for a surveillance system with modest dynamic capabilities. Here, the effect of changing the dynamic characteristics of the surveillance system on the performance of dispatching is investigated. Fig. 10 presents the overall visibilities for four different surveillance systems surveying an object that is following a straight-line trajectory. The performance of these systems range from very fast (at least an order of magnitude faster than the object) to a static system in which the sensors have no dynamic capability at all. (Note that, while the static sensors cannot move, dispatching is still utilized to determine the appropriate subsets of sensors (the coordination strategy) in real-time.) For each system, its own optimal initial surveillance-system configuration was determined.

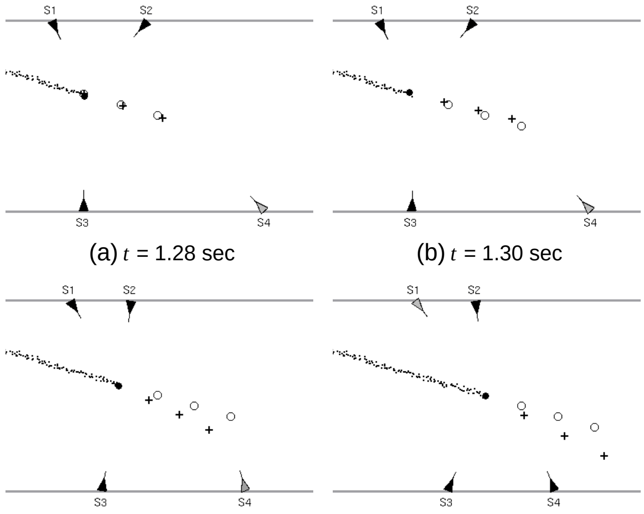

(a) t = 1.28 sec (b) t = 1.30 sec (c) t = 1.90 sec (d) t = 2.50 sec

Figure 9. Straight-line trajectory.

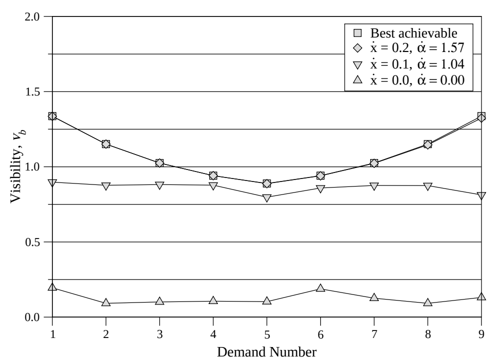

Figure 10. Observed visibilities of a straight-line trajectory for different surveillance systems.

From Fig. 10, as expected, it can be seen that dynamic surveillance systems outperform the static surveillance system. The variation in visibility for the fast surveillance systems is an artifact of the workspace constraints placed on the sensors. By constraining the sensors to rails, they cannot maintain a constant range from the object. Thus, as the object approaches the centre of the workspace, the visibility drops due to increased range. The upper curve in Fig. 10 represents the best achievable visibility under these conditions; here, for each demand point, the assigned sensors match the x-position of the object and align their axes directly with the object ($\theta = 0$).

A comparison between the "best" system and the other systems indicates that increasing the speed of the surveillance system beyond an upper limit is not particularly valuable (e.g., the achieved visibilities for $\dot{x} = 2.5$ m/sec, $\dot{\alpha} = 2\pi$ rad/sec (best achievable) and $\dot{x} = 0.2$ m/sec, $\dot{\alpha} = \frac{\pi}{2}$ rad/sec are practically the same). However, it is clear that providing a surveillance system with even limited dynamic capabilities (e.g., $\dot{x} = 0.1$ m/sec and $\dot{\alpha} = \frac{\pi}{3}$ rad/sec) may significantly improve its effectiveness.

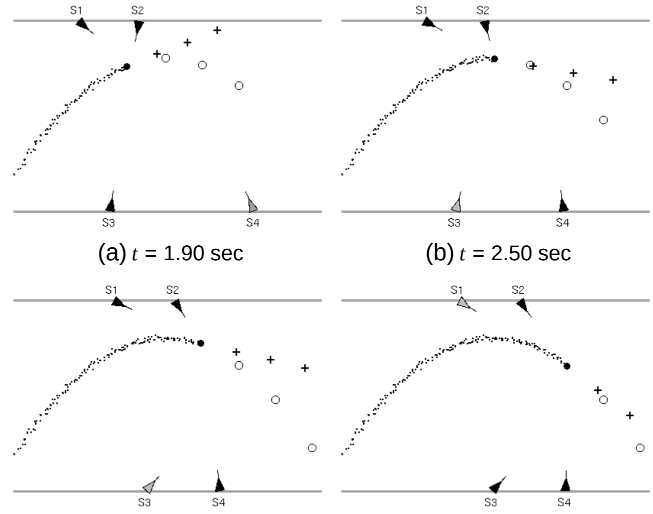

(a) t = 1.90 sec (b) t = 2.50 sec (c) t = 3.10 sec (d) t = 3.70 sec

Figure 11. Parabolic trajectory.

Robustness to Trajectory Variation

This section presents an example that demonstrates how dispatching can be used to adapt a surveillance system, expecting a particular object trajectory (i.e., for which the initial configuration is optimized), to a different object trajectory. The initial configuration of the surveillance system is identical to that used for Example 1. Fig. 11 illustrates how the surveillance system adapts to a parabolic object trajectory.

The proposed dispatching methodology adjusts the fusion subsets (coordination strategy) and the individual sensor poses (positioning strategy) in response to an *a priori* unknown object trajectory. Fig. 12 presents the overall visibilities for this trajectory using four different surveillance systems, with dynamic capabilities ranging from static to very fast.

The performances of the four surveillance systems considered for the parabolic trajectory (Fig. 11) are shown in Fig. 12(a). As expected, the faster dynamic system ($\dot{x} = 0.2$ m/sec, $\dot{\alpha} = \frac{\pi}{2}$ rad/sec) performs almost perfectly, the slower dynamic system ($\dot{x} = 0.1$ m/sec, $\dot{\alpha} = \frac{\pi}{3}$ rad/sec) fairs somewhat worse, and all dynamic systems outperform the static system. The drops in visibility for the first demand are a result of the difference between the expected and actual trajectories. All sensors were in their initial poses under the assumption that the object would be following a straight-line trajectory, entering from the upper-left corner of the workspace. When the object entered from the lower-left, the sensors did not have an opportunity to react (sensing of the first demand point is almost instantaneous and determined through the initial configuration, not dispatching). As a result, the sensors were not in an optimal pose, nor was the initial fusion subset appropriate.

Fig. 12(b) illustrates how the performance of each system would be altered, if provided with prediction information prior to surveillance of the first demand point. In this case, observations of the object (using the overhead camera) began outside of the workspace. These observations allowed for the *a priori* ini-

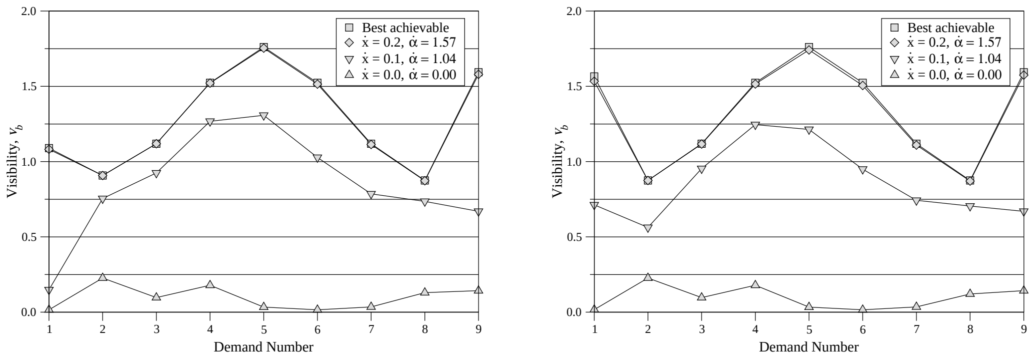

(a) Parabolic (b) Parabolic with initial reaction

Figure 12. Observed visibilities for different surveillance systems expecting a straight-line trajectory.

tialization of the Prediction Module, which in turn provided the dispatching and positioning modules with better estimates of the first demand point and allowed them sufficient time to optimally position the sensors for best visibility. For all of the dynamic systems, the performance is improved.

Several observations may be made from this example: First, these simulations confirm that the surveillance system can still provide valuable information, even when the actual object trajectory deviates significantly from the expected object trajectory. Second, adaptation of the surveillance system to the trajectory is a function of the dynamic capabilities of the sensors. In other words, while the surveillance system's performance may degrade as the actual trajectory deviates from expectation, the degradation is more marked for slower sensors. (It is important to note that even static systems still provide surveillance information, just at a lower visibility). This would indicate that the slower the system, the more important a reasonable initial guess of the object trajectory becomes.

Dispatching vs. Non-Dispatching Systems

Fig. 13 compares a number of systems using the dispatching methodology presented in this paper with a system that does not use dispatching at all. The dispatching systems select subsets of three sensors from the four available for use in a sensor fusion process; the non-dispatching system simply fuses the measurements from all four sensors. In other words, the non-dispatching approach may be stated as: "all the sensors, all the time".

In the comparative runs considered in Fig. 13(a), each system started with an initial configuration that was optimized for a straight-line object trajectory and observed the same trajectory "with noise". In Fig. 13(b), each surveillance system was configured under the assumption that the object would follow a straight-line trajectory, while the actual observed trajectory was parabolic.

From these Figures, it is clear that dispatching systems outperform non-dispatching systems, provided that they have at least limited dynamic capabilities. This is best exemplified by the dynamic system having only rotational capability ($\dot{x} = 0$ m/sec, $\dot{\alpha} = \frac{\pi}{3}$ rad/sec) that still significantly outperforms the non-dispatching system. One should not conclude from this that $\dot{\alpha}$ capability is more important than $\dot{x}$; in fact, if the sensors can move as fast as the object along the rails, rotational ability will not improve the object visibility at all. As such, further study of the relationship between $\dot{x}$ and $\dot{\alpha}$ is warranted.

However, for static systems, the advantage of the dispatching approach is not apparent. From a performance perspective, the user of a static surveillance system may be better off to simply use all of the sensors at once. The use of dispatching for a static system may only make sense if the costs associated with processing the data from all of sensors at once compromises the real-time performance of the system (e.g., high-resolution image processing). In this case, dispatching provides an effective mechanism to select an appropriate subset for processing.

CONCLUSIONS

A method for maximizing the effectiveness of moving-object surveillance using multiple sensors is presented in this paper. The overall goal of the method is to position sensors in response to changing demands. This is shown to be possible using a two-part dispatching strategy, comprising coordination and positioning strategies. The motion of each sensor is evaluated based on the quality of information that each can provide for specified object locations. From this, a group of sensors (for use in a sensor fusion context) may then be assigned to a particular sensing demand. In addition, the sensors that are not required for the most imminent demand, are assigned to future predicted demands. This ensures that as many sensors as possible are maneuvered in anticipation of upcoming requirements rather than remaining idle or moving randomly. The use of this dispatching methodology, combined with dynamic sensors and sensor fusion, has been shown to provide a considerable benefit over (fixed) static-sensor surveillance systems. The increased accuracy and robustness of a dynamic system makes it suitable for tasks that require high-quality, real-time sensor information.

ACKNOWLEDGMENT

The financial support of the Natural Sciences and Engineering Research Council of Canada is gratefully acknowledged.

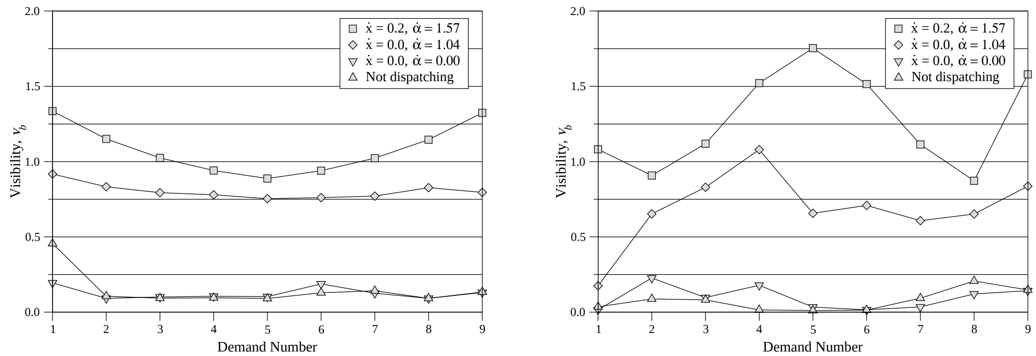

(a) Straight-line (b) Parabolic

Figure 13. Observed visibilities for dispatching vs. non-dispatching surveillance systems.

REFERENCES

[1] Luo, R. C. and Kay M. G., 1992, "Multisensor integration and fusion for intelligent machines and systems," in *Data Fusion in Robotics and Machine Intelligence* (Abidi, M. A. and Gonzalez, R. C., eds.), Academic Press, San Diego, CA, pp. 7–136.

[2] Viswanadham, N. and Narahari, Y., 1992, *Performance Modeling of Automated Manufacturing Systems*, Prentice Hall, Englewood Cliffs, NJ, pp. 90–94.

[3] Shrivastava, M., Chande, P. K., and Monga, A. S., 1997 "Taxi dispatch: A fuzzy rule approach," in *Proc. 1997 IEEE Conf. on Intelligent Transportation Systems*, pp. 978–982.

[4] Potvin, J.-Y., and Shen, Y., 1991, "A neural network approach to the vehicle dispatching problem," in *Proc. IEEE Int. Joint Conf. on Neural Networks*, **2**, pp. 1230–1235.

[5] Benyahia, I. and Potvin, J.-Y., 1998, "Decision support for vehicle dispatching using genetic programming," *IEEE Trans. SMC–Part A*, **28**, pp. 306–314.

[6] Psaraftis, H. N., 1988, "Dynamic vehicle routing problems," in *Vehicle Routing: Methods and Studies* (Golden, B. L. and Assad, A. A., eds.), North-Holland, Amsterdam, pp. 223–248.

[7] Powell, W. B., 1988, "A comparative review of alternative algorithms for the dynamic vehicle allocation problem," in *Vehicle Routing: Methods and Studies* (Golden, B. L. and Assad, A. A., eds.), North-Holland, Amsterdam, pp. 249–291.

[8] Trucco, E., Umasuthan, M., Wallace, A. M., and Roberto, V., 1997, "Model-based planning of optimal sensor placements for inspection," *IEEE Trans. Robotics and Automation*, **13**, pp. 182–194.

[9] Zhang, H., 1995, "Two-dimensional optimal sensor placement," *IEEE Trans. SMC*, **25**, pp. 781–792.

[10] Niepold, R., Sakane, S., and Shirai, Y., 1987, "Vision sensor set-up planning for a hand-eye system using environmental models," in *Proc. of the Society of Instrument and Control Eng. of Japan*, (Hiroshima, Japan), pp. 1037–1040.

[11] Abrams, S., Allen, P. K., and Tarabanis, K., 1999, "Computing camera viewpoints in an active robot work cell," *Int. J. of Robotics Research*, **18**, pp. 267–285.

[12] Matsuyama, T., Wada, T., and Tokai, S., 1999, "Active image capturing and dynamic scene visualization by cooperative distributed vision," in *Advanced Multimedia Content Processing* (Nishio, S. and Kishino, F., eds.), Springer-Verlag, Berlin, pp. 252–288.

[13] Horling, B., Vincent R., Mailler R., Shen J., Becker R., Rawlins, K., and Lesser, V., 2000, "SPT: distributed sensor network for real time tracking," Tech. Rep. 00-49, University of Massachusetts, Amherst, MA 01003.

[14] Saad, R. E., Bonen, A., Smith, K. C., and Benhabib, B., 1999, "Proximity sensing for robotics," in *The Measurement, Instrumentation, and Sensors Handbook* (J. G. Webster, ed.), CRC Press, Boca Raton, Florida, pp. 8–1—8–16.

[15] Ruocco, S. R., 1987, *Robot Sensors and Transducers*, Open University Press, Milton Keynes, England.

[16] Hall, D. L., 1992, *Mathematical Techniques in Multisensor Data Fusion*, Artech House, Boston, MA.

[17] Nakamura, Y. and Xu, Y., 1995, "Geometrical fusion method for multisensor robotic systems," in *Multisensor Integration and Fusion for Intelligent Machines and Systems* (Luo, R. C. and Kay, M. G., eds.), Ablex Publishing Corporation, Norwood, NJ, pp. 241–259.

[18] Bar-Shalom, Y. and Li, X.-R., 1993, *Estimation and Tracking: Principles, Techniques, and Software*, Artech House, Boston, MA.

[19] Broida, T. J., 1992, "Kinematic and statistical models for data fusion using Kalman filtering," in *Data Fusion in Robotics and Machine Intelligence* (Abidi, M. A. and Gonzalez, R. C., eds.), Academic Press, San Diego, CA, pp. 267–309.

[20] He, D., Hujic, D., Mills, J. K., and Benhabib, B., 1996, "Moving-object recognition using premarking and active vision," in *Proc. 1996 IEEE Int. Conf. on Robotics and Automation*, **3**, pp. 1980–1985.

[21] Naish, M. D., Croft, E. A., and Benhabib, B., 2001, "Simulation-Based Sensing-System Configuration for Dynamic Dispatching," in *Proc. 2001 IEEE Int Conf. on Systems, Man, and Cybernetics*.