Multidim Syst Sign Process (2009) 20:135–164
DOI 10.1007/s11045-008-0071-2

Wireless image sensor networks: event acquisition in attack-prone and uncertain environments

Alexandra Czarlinska · Deepa Kundur

Received: 30 April 2008 / Revised: 30 April 2008 / Accepted: 16 September 2008 /
Published online: 9 October 2008

Abstract Wireless Image Sensor Networks (WISNs) consisting of untethered camera nodes and sensors may be deployed in a variety of unattended and possibly hostile environments to obtain surveillance data. In such settings, the WISN nodes must perform reliable event acquisition to limit the energy, computation and delay drains associated with forwarding large volumes of image data wirelessly to a sink node. In this work we investigate the event acquisition properties of WISNs that employ various techniques at the camera nodes to distinguish between event and non-event frames in uncertain environments that may include attacks. These techniques include lightweight image processing, decisions from n sensors with/without cluster head fault and attack detection, and a combination approach relying on both lightweight image processing and sensor decisions. We analyze the relative merits and limitations of each approach in terms of the resulting probability of event detection and false alarm in the face of occasional errors, attacks and stealthy attacks.

Keywords Image sensor networks · Lightweight event acquisition · Sensor network security

1 Introduction

Wireless Image Sensor Networks (WISNs) are envisioned for a variety of innovative applications such as distributed surveillance, intelligent infrastructure monitoring and scientific data collection (Akyildiz et al. 2007; Feng et al. 2001). To realize this vision, WISN research must overcome challenges associated with the increased processing, transmission and bandwidth costs required for image data compared with conventional sensor data (Basharat et al. 2005; Soro and Heinzelman 2005; He and Wu 2006). A variety of interesting approaches have been

This work is supported in part by NSF grants ECCS-0735114 and EEC-0649142.

A. Czarlinska (✉) · D. Kundur
Department of Electrical & Computer Engineering, Texas A&M University, College Station, TX, USA
e-mail: czlinska@ece.tamu.edu

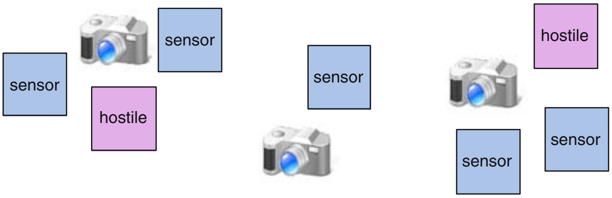

Fig. 1 A general heterogeneous visual sensor network comprised of untethered camera nodes and supporting sensors in the possible presence of a distributed attacker

proposed to address these issues including exploiting spatial and temporal data correlations (Chow et al. 2007; Ma and Liu 2005) as well as utilizing collaboration between camera nodes and sensors (Veeraraghavan et al. 2005). The latter approach has shown particular promise for detecting events of interest occurring in the environment (He et al. 2006; Rahimi et al. 2005).

In this work we examine the event acquisition properties of WISNs pre-deployed randomly or deterministically in unattended outdoor environments for the purpose of collecting relevant surveillance data regarding an event of interest. As shown in Fig. 1, we consider a heterogeneous WISN comprised of sensors and untethered camera nodes (battery operated nodes with wireless data transmission to the sink; He et al. 2006). If the sensor and camera node deployment is random, each camera may or may not find itself within a close range of one or more supporting sensors. Such ad hoc arrangements may arise for example in rapid deployment scenarios where sensors are dropped aerially into an area of possible danger. Deterministic deployments on the other hand may allow for the arrangement of n sensors within a desired radius around each camera to assist the camera in triggering and/or in identifying events of interest. To offer such decision support, a variety of sensors such as magnetometer or motion sensors might be chosen as dictated by the application. In this work we only require that the sensor be capable of providing a binary "yes/no" output decision regarding the presence of events given its sensory input. We also consider a general deployment where each camera node may or may not find itself in the vicinity of n supporting sensors.

Whether with or without the assistance of supporting sensors, the camera nodes must operate reliably under significant resource constraints and may thus require the use of lightweight image processing (LIP) algorithms to perform event acquisition (Rahimi et al. 2005; Veeraraghavan et al. 2005). Under the collaboration paradigm, camera nodes may receive supporting decisions about the presence or absence of an event from the distributed sensors (Veeraraghavan et al. 2005; Basharat et al. 2005). Due to their unattended deployment in potentially hostile regions however, the sensors may experience errors due to an attack that is perpetrated by a hostile entity (for example a distributed hostile network as depicted in Fig. 1; Czarlinska et al. 2007; Raymond and Midkiff 2008; Buttyan and Hubaux 2002). In this work we wish to investigate the detection and false alarm characteristics of such heterogeneous WISNs in uncertain environments where outdoor conditions may present challenges for the camera LIP algorithms while occasional faults and deliberate attacks may present a challenge for the supporting sensors. In particular we wish to study the relative merits and limitations of the following cases:

1. Lightweight Image Processing (LIP) Approach: preliminary results from real-world testbeds of low-power low-computation wireless camera nodes suggest that LIP algorithms may achieve a probability of detection and false alarm ($P_D - P_{FA}$) that may be acceptable for certain applications (Rahimi et al. 2005; Veeraraghavan et al. 2005). Though we do

not set out to improve any particular LIP algorithm, in this work we wish to understand the underlying analytical properties of simple threshold based LIP algorithms and examine their $P_D - P_{FA}$ performance for a variety of real world surveillance sequences. The overarching goal is thus to understand the limits and suitability of LIP algorithms and to provide a framework for enhancing their performance with sensors if required for a given application.

2. Sensor Decisions Approach: in unattended outdoor settings, sensors may be prone to occasional errors due to faults or may experience stealthy attacks (designed to avoid detection; Czarlinska et al. 2007; Raymond and Midkiff 2008; Buttyan and Hubaux 2002). Although quality testing may provide an estimate for the probability of a fault, an estimate for the probability of an attack may not be generally available a priori. Without verification mechanisms, the reliability of the sensor decisions may thus not be adequate for some applications despite node redundancy. In this work we study and compare the $P_D - P_{FA}$ performance of different fault/attack verification mechanisms at the cluster head for the case of occasional errors, attacks and stealthy attacks. We also comment on the level of node redundancy (cluster size) required to achieve a desired event acquisition performance under each scenario.
3. Combined Decisions Approach: camera nodes may rely on a combination of decisions from the sensors and the LIP algorithm for event acquisition (Rahimi et al. 2005; Veeraraghavan et al. 2005; He et al. 2006). We wish to study the characteristics of such combined decisions for various levels of node redundancy to exploit the desirable qualities of both approaches and avoid the degradation in performance that each method may experience in certain settings.

The remainder of this paper is organized as follows. Section 2 provides important motivation for the WISN event acquisition problem and overviews recent advances salient to our work. Section 3 describes the details of the event acquisition problem in uncertain environments. In Sect. 4 we focus on lightweight image processing (LIP) algorithms and their performance. In Sect. 5 we focus on the reliability of sensor decisions and study the performance of fault/attack detectors. In Sect. 6 we examine the combined decisions approach with/without fault and attack detectors. Finally in Sect. 7 we summarize our findings and conclusions.

2 Background and recent advances

For many applications, the viability of Wireless Image Sensor Networks (WISNs) depends on the resolution of significant design issues centered around network reliability (Czarlinska and Kundur 2008; Eltoweissy et al. 2006) longevity (Yu et al. 2007; Maniezzo et al. 2002) and security (Eltoweissy et al. 2005; Chan et al. 2003). The specific issues include everything from energy-efficient capture of images as well as their processing and routing (Chow et al. 2007; Veeraraghavan et al. 2005; Chow et al. 2006; Rodriguez 2003), to economical network design relying on node heterogeneity with sleep/wake-up cycles (Bandyopadhyay and Coyle 2003; He et al. 2006), to the network's robustness to attack and compromise of privacy (Olariu et al. 2007). WISNs thus present a very wide range of timely challenges. In this section we wish to briefly outline some recent advances most salient to the focus of our work.

In He et al. (2006) describe VigilNet, a prototype implementation of a heterogeneous image sensor network for energy efficient surveillance missions. In the experimental setup, 70 Mica2 motes are deployed to detect and track the passing of a vehicle while triggering cameras. The authors demonstrate how a multi-tier sleep/wake-up system consisting of motes and mote leaders called sentries can extend the lifetime of the network. Importantly, a sentry decides whether an event of interest is occurring in the environment by counting the number of "yes" votes it receives from the motes which are utilizing magnetometer sensors. To accurately capture events, the probabilities of false positives and false negatives are balanced by carefully selecting the detection threshold (i.e. the number of "yes" votes required to declare that an event has occurred, referred to by the authors as the Degree of Aggregation DoA). Importantly, the DoA is selected *experimentally* and as the authors suggest, a framework with adjustable sensitivity for this selection could largely improve the system's detection performance. In comparison, in our work we study a framework that enables the cluster head (i.e. sentry) to make optimal decisions based on sensor inputs with a flexible level of sensitivity.

Although heterogeneous multi-tier systems greatly improve the longevity of the *overall* image network, practical designs must account for the substantial power consumption of individual image-capture devices and of their image processing algorithms which handle the acquired frames. Recent advances in CMOS imaging technology have produced a new breed of low-power camera devices. Unfortunately these devices are generally intended for higher-power hosts and are thus not suitable for sensor networks. To address this issue (Rahimi et al. 2005), present a seminal camera device named Cyclops. Cyclops provides an electronic interface between a low-power low-computation camera module based on CMOS imagers and a lightweight camera host such as a mote. While providing a critical bridge and enabling use in visual sensor networks, Cyclops still suffers from extreme constraints in its computational power and processing delay, necessitating judicious use of its resources. For instance, Cyclops's complex programmable logic device (CPLD) can perform simple operations on the frames at *capture time*, such as background subtraction and frame differentiation. Performing such simple operations at capture time instead of post-processing the frames greatly reduces the energy consumption and delay of the device. In our work, we study how lightweight image processing (LIP) compatible with these ideas can improve event detection performed by the cluster head and its associated sensors. Importantly cyclops indeed possesses an asynchronous trigger input that can be connected to sensors (such as a passive IR detector, microphone or magnetometer) to trigger the camera and improve the overall system's performance.

3 The WISN event acquisition problem in uncertain environments

Upon deployment, the goal of a typical image network is to *capture relevant* visual surveillance pertaining to an event of interest and to forward this surveillance to a sink where further analysis might be performed. To capture relevant surveillance, the network should generally exhibit a high probability of event detection P_D (true positive) and a low probability of false alarm P_{FA} (false positive). These probabilities not only affect the relevance of the collected materials to the surveillance task, but also have an impact on the network's energy consumption and thus on its longevity. Specifically, the erroneous identification of "non-event" frames as "significant" and their subsequent processing and transmission through a (wireless) medium needlessly drains the nodes' battery resources and burdens the sink with non-content. On the other hand, the omission of "event" frames may significantly compromise the quality of the surveillance mission. The $P_D - P_{FA}$ characteristics of the image network should also ideally exhibit the highly desirable property of being adjustable based on the requirements of the application (such as its surveillance or energy requirements). The detection performance should also ideally exhibit some optimality in the sense of being "the best" achievable performance given the practical challenges of WISNs.

Event acquisition challenges experiences by nodes in a WISN generally stem from more than one source. The first such source originates from the hardware and energy limitations of the camera nodes themselves (Rahimi et al. 2005). Specifically, the camera nodes may not have the *capability* of applying advanced image processing to the captured frames. The processing delay (per frame) as well as the energy and memory utilization generally render such processing infeasible even when it is available at the camera nodes. However lightweight image processing (LIP) is often feasible on such devices and provides very basic in situ analysis of the frames' content. Unfortunately the practical $P_D - P_{FA}$ performance of LIP varies widely depending on the specific environmental conditions (for example, lighting, size and speed of the moving object(s) and movement of "background" objects such as trees). The performance of LIP for any given arbitrary image sequence is consequently not truly predictable or controllable, and thus not inherently "adjustable" to meet application requirements. Nevertheless, real-world experiments with Cyclops based on a LIP algorithm have demonstrated an average $P_D \approx 78\%$ and a $P_{FA} \approx 22\%$. While this level of $P_D - P_{FA}$ performance may potentially be acceptable for certain applications, it is important to investigate whether this performance can be improved through the use of collaborating sensors.

Exploiting the information collected by sensors might generally improve the $P_D - P_{FA}$ performance of heterogeneous WISNs. The use of sensors however introduces two new sources of error that must be considered. The first such source comes from occasional sensor faults or errors that occur with some small but non-zero probability at each sensor. Aside from quality testing prior to deployment (which might be selective or altogether absent due to the large number of sensors), the general approach is to employ sensor *redundancy* to reduce the chance of false reporting (Raymond and Midkiff 2008). Nevertheless it is not always clear what *level of redundancy* (number of sensors) is required to achieve a given $P_D - P_{FA}$ performance, especially if the camera nodes are already performing a basic level of detection via a LIP algorithm. The issue of redundancy becomes even more salient when we consider the second possible source of sensor error, that is, error due to a persistent and distributed attack. In particular, WISNs are intended for deployment in unattended and possibly hostile regions. In such scenarios, an opponent can clearly gain physical access to the sensors with the possibility of destroying them, capturing them for reprograming purposes, or interfering with their readings via actuator devices (Czarlinska and Kundur 2008; Czarlinska et al. 2007). Despite the possibility of tampering or error, the use of sensors to achieve reliable and energy efficient image networks is highly enticing if these issues can be resolved (Akyildiz et al. 2007; Rahimi et al. 2005).

In this work we wish to study the WISN event acquisition problem in uncertain environments where the sensors are prone to either occasional faults or persistent attacks and where the camera devices perform very basic event detection using lightweight image processing (LIP) algorithms with unpredictable performance due to varying conditions. Specifically we wish to understand how the role of sensor redundancy changes if uncertainty in the environment shifts from mere faults to hostile attacks. For instance we wish to understand if LIP algorithms alone are sufficient in certain cases (such as the case of a severe sensor attack) or if their performance should be augmented with that of sensors.

3.1 System model

In this section we wish to detail the specific system setup analyzed in this work. As shown in Fig. 2, a camera node has the capability of performing lightweight image processing (detailed in Sect. 4) to perform event acquisition. The camera node may however also rely on input

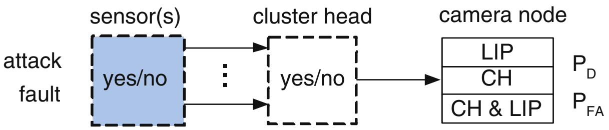

Fig. 2 Each camera node may employ lightweight image processing (LIP) to determine if an event of interest has occurred in a collected frame or it may rely on decisions from the sensor(s)/cluster head (CH) or a combination of both (CH & LIP)

Table 1 Methods for marking a frame as an "Event" at camera node

Method	Action for marking frame as an event
LIP	Mark based on lightweight image processing (LIP)
CH	Mark based on sensor(s) with/without Cluster Head (CH) detection
CH & LIP	Mark if *either* LIP *or* CH detects an event

from one or more (error or attack-prone) sensors regarding the presence or absence of an event of interest (shown in Fig. 2 as binary "yes/no" decisions). Information from the sensor(s) may be directly fed into the camera or it may first pass through a cluster head (CH) where some form of attack/error detection is performed. In that case, the camera node receives a decision about the presence or absence of an event from the cluster head instead of directly from the sensor(s).

Based on this setup, a camera node may receive information about the presence or absence of an event from more than one source (i.e. from the sensor(s)/cluster head and from its own frame processing). Since it is not known a priori which source will be more reliable under a given setting, the camera node faces several possible methods of utilizing the received information. Specifically, the camera node must decide which source to trust when the sources are in disagreement (i.e. one source reports an event of interest while the other reports no event). As shown in Table 1, one possible approach to resolve a disagreement is for the camera to trust its lightweight image processing (LIP) as the more "reliable" element which is not prone to attack. Indeed such a strategy might be fruitful in favorable lighting and background-motion conditions. Another obvious approach is for the camera node to trust the sensor/cluster head decision (CH in Table 1) and treat the LIP as the more volatile element. Finally the third approach listed in Table 1 instructs the camera to mark a frame as an "event" if *either* of the two sources reports an event. This approach *may* prevent the missed detection of certain events but could produce many false reports if at least one of the sources experiences significant errors. Thus under arbitrary environmental conditions, it is not clear which of the techniques will produce a better *overall* $P_D - P_{FA}$ performance (with P_D as close to 1 as possible and P_{FA} as close to 0 as possible) and how sensor redundancy will affect this performance.

3.2 Sensor error and attack models

As described in Sect. 3.1, in this work we assume that the sensors provide binary "yes/no" decisions about the presence or absence of an event as a result of their sensed observations. The abstraction is detailed in Fig. 3a where each sensor utilizes an adjustable threshold T_h (within technology limits of the sensor) to decide whether an event of interest has occurred

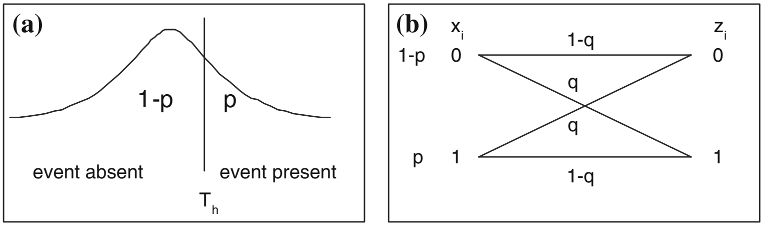

Fig. 3 **a** Binary sensor model with sensing threshold T_h and resulting probability of witnessing an event $0 \le p \le 1$. **b** Basic bit error model due to fault or (unstealthy) attack where $0 \le q \le 1$ is typically small for faults but may be arbitrarily large for attacks

to produce a resulting decision bit of value 1/0. Decisions based on the use of a threshold result in a given probability p of witnessing an event (bit of value 1). The output decision of a sensor i is thus described by a Bernoulli random variable which we denote by X_i as shown in Eq. 1.

$$X_i = \begin{cases} 1 & \text{w.p. } p \\ 0 & \text{w.p. } 1-p \end{cases} \tag{1}$$

$$Y_i = \begin{cases} 1 & \text{w.p. } q \\ 0 & \text{w.p. } 1-q \end{cases} \tag{2}$$

$$Z_i = X_i \oplus Y_i \tag{3}$$

$$Z_i = \begin{cases} 1 & \text{w.p. } r \\ 0 & \text{w.p. } 1-r \quad \text{where} \quad r = q + p - 2pq \end{cases} \tag{4}$$

The notion of an occasional sensor fault is captured through the familiar bit error model depicted in Fig. 3b. In the model, a bit may be reported erroneously with a probability q. The error at sensor i is thus modeled by another Bernoulli random variable which we denote by Y_i as shown in Eq. 2. The decision bit produced by a sensor i in the presence of possible errors is thus given by the Bernoulli random variable Z_i as shown in Fig. 3b and given by Eqs. 3 and 4.

Importantly we note that the probability q of an error due to an occasional sensor fault might be *small* and may possibly be estimated or upper-bounded from experimental setups. If however the sensor error is caused by a hostile attack in an unattended environment, we may no longer conclude anything specific regarding the probability q a priori. Indeed if no *verification* mechanisms are in place, the value of q at a given sensor may take on *any* value in the permissible range of $0 \le q \le 1$ (a value of $q = 1$ at a given sensor indicates that the sensor always gives a decision opposite from the real sensed event such as for instance in the case when a node is captured). This possibility leads back to the question of the level of redundancy (number of sensors) that should be utilized to support a camera node. The issue is further complicated if the attacker utilizes a *distributed* attack approach through sensor-actuator nodes (Akyildiz and Kasimoglu 2004). Sensor-actuator nodes deployed in an environment interfere with the readings taken by other sensors in their vicinity without the need to physically "capture" the sensor and break its physical/cryptographic mechanisms for the purposes of reprograming (Czarlinska et al. 2007). Depending on the effectiveness of such a hostile deployment, many sensor nodes might be affected and produce erroneous readings with some non-negligible probability. A form of verification against occasional errors as well as wide-spread errors (due to attack) is highly desirable.

In Sect. 3.1 we described a system model that may include a cluster head (CH) which collects individual sensor decisions as in many proposed systems and practical implementations (Akyildiz et al. 2007; He et al. 2006; Rahimi et al. 2005). In such systems the cluster head requires that a specified number of sensors c reports an event before deciding that an event has most likely occurred and passing this information to the camera node. The requisite number of sensors c [also known as the "weight" (Czarlinska et al. 2007) or the degree of aggregation DoA (He et al. 2006)] may be determined experimentally, approximated based on expectations or obtained theoretically (Sect. 5). For instance to perform an approximation, if there are n sensors and each sensor has a probability p of witnessing an event (based on its threshold T_h), then the average expected number of sensors that report an event is $c \approx np \pm \epsilon$ where ϵ may be determined experimentally. The cluster head thus receives decisions from n sensors which for the case of no errors is denoted in vector form by $\mathbf{X} = [x_1, x_2, \ldots, x_n]$ and under the case of possible errors by $\mathbf{Z} = [z_1, z_2, \ldots, z_n]$. If we denote by w the weight (number of 1's) contained in the vector received by the cluster head, then the cluster head computes $w(\mathbf{Z})$ (or $w(\mathbf{X})$ if there are no errors) and compares it to the expected weight c. Based on this simple form of error/attack filtering, the cluster head is effectively deciding between two hypotheses regarding the received sensor data as shown in Eq. 5. The H_0 or null hypothesis is that the received data $\mathbf{Z}$ comes from the $Bern(p)$ distribution and is thus error free (in this case $\mathbf{Z}$ is really $\mathbf{X}$). The H_1 or alternative hypothesis is that the received data contains errors due to attack or fault.

$$\begin{aligned} H_0 &: \text{normal operation, } \mathbf{Z} \sim Bern(p) \\ H_1 &: \text{attack (or fault), } \mathbf{Z} \sim Bern(r) \end{aligned} \tag{5}$$

The preceding consideration of sensor error due to attack effectively treats the attack as a regular fault with the difference that the fault may possibly be unrestricted in value due to the attacker's choice of action. However an attacker that wishes to effectively misguide the event acquisition process of a WISN may possibly take the cluster head detector into account. Specifically, the attacker may wish to determine the optimal probability of attack q^* that causes erroneous decisions while minimizing the chance of being detected by the cluster head. This problem which we refer to as a *stealthy* attack is captured by Eq. 6. The attacker wishes to choose an optimal value of attack parameter q^* such that the weight of the attacked data (which depends on probabilities p and q) generally matches (in terms of the probability of occurrence) the weight of the unaltered data (which depends on the probability p alone). In general the attacker need not know the probability p (since it depends on the sensor threshold T_h) and the optimization might be performed through game theoretic optimization where the sensors with unknown parameter p are treated as an opponent (the attacker is treated as the other player with unknown parameter q). In this work we wish to consider the effect of occasional errors, unconstrained attacks and stealthy attacks on the WISN event acquisition process and determine what level of sensor redundancy is required to support the lightweight image processing available at the camera nodes.

$$q^* = \max \quad {}_{0 \le q \le 1} \quad Pr\{w(\mathbf{X}_p) = w(\mathbf{Z}_{p,q})\} \tag{6}$$

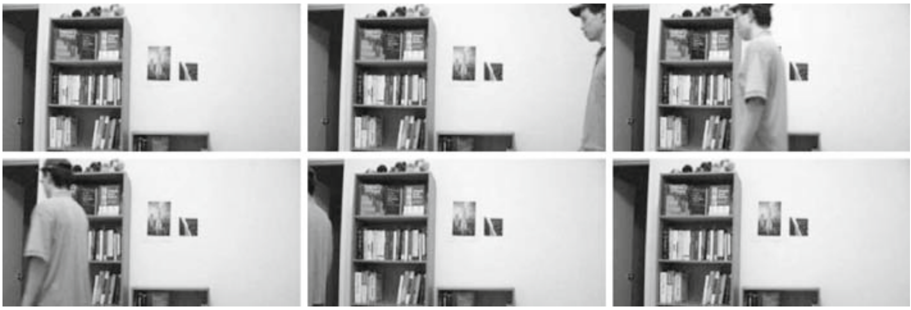

Fig. 4 Seq. 1 with frames (**a**–**f**) from top left to bottom right: indoor test conditions with constant lighting and no background changes

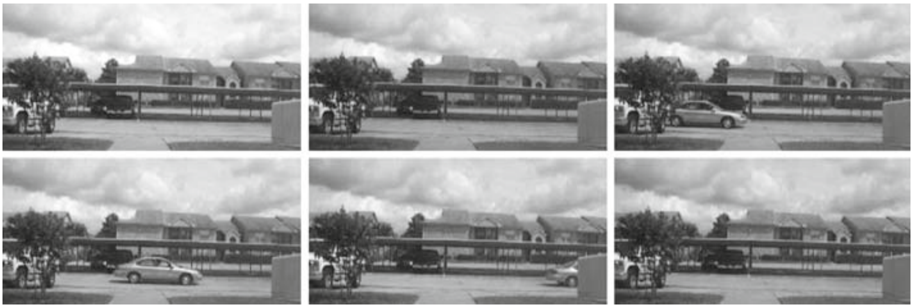

Fig. 5 Seq. 2 with frames (**a**–**f**) from top left to bottom right: outdoor variable lighting due to clouds. Ex: The light intensity changes by 70% between frames (**a**) and (**b**). Additional background movement due to shrub

4 Lightweight image processing in WISNs for event acquisition

4.1 Lightweight image processing and sequence characteristics

In the spirit of lightweight image processing (LIP) (Rahimi et al. 2005; Rosin 2002; Wu and Chen 2007), we consider a relatively simple and general event acquisition algorithm (i.e. the approach is not tailored to the detection of any *specific* type of object). Examination of the proposed algorithm is intended to provide more insights into the properties of simple visual algorithms and serve as an illustration of their performance in the context of energy and computation-limited camera nodes. To assess this generic algorithm for WISNs we consider its properties in analytic form and obtain the algorithm's $P_D - P_{FA}$ performance for surveillance sequences under varying conditions (Sect. 4.3). The real-world image sequences used in our testing are shown in Figs. 4–7. The characteristics of these sequences are important in understanding the suitability of the proposed algorithm for event acquisition in WISNs. We thus describe the test sequences prior to outlining the visual event acquisition algorithm and its properties.

The sequence of Fig. 4 is an idealized indoor test where the lighting and background conditions do not change appreciably over time. The only significant change comes from the event of interest in the form of a test subject entering the camera's field of view. The dominant source of noise in this case is internal camera noise and flicker. The sequence of Fig. 5 shows outdoor parking-lot surveillance on a windy day, where the event of interest is the passing of an unidentified car. The event acquisition task in this sequence is complicated

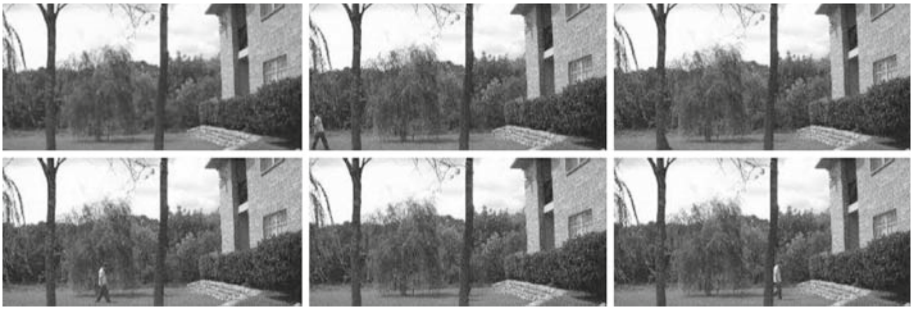

Fig. 6 Seq. 3 with frames (**a**–**f**) from top left to bottom right: changing outdoor light and background (swaying trees). The subject temporarily disappears behind a tree in frames (**c**) and (**e**)

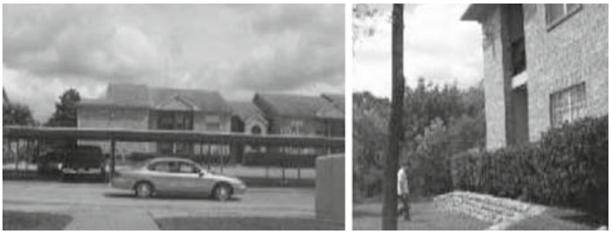

Fig. 7 **a** Seq. 4a showing Seq. 2 modified to remove the shrub. **b** Seq. 4b showing Seq. 3 modified to remove the swaying trees

by the presence of a nearby shrub which experiences significant swaying of its branches over time. Furthermore the background lighting changes visibly with cloud movement (between frames 5a, b for example). The sequence of Fig. 6 also experiences changes due to swaying trees and variable light conditions. The event of interest is the appearance and movement of a test subject which temporarily disappears behind a tree in frames 6(c) and 6(e). Finally Fig. 7 shows a truncation of the sequences of Figs. 5 and 6 where the camera's field of view now excludes the shrub and trees.

Statistical analysis of the image sequences in Figs. 5 and 6 [such as Levine's Test and the t-test (Ott and Longnecker 2001)] reveal that the mean and standard deviation are not reliable indicators of an event of interest occurring even after various filtering mechanisms are employed. This can be seen intuitively from the fact that the subjects of interest (person walking and car driving-by) do not occupy a much larger percent of a frame's pixels than the other randomly moving objects (shrub and trees). Hence the mean and variance of the frames *do* change based on the appearance of the subject, but these differences are not statistically distinguishable. In essence, the pixels corresponding to the person and car are getting dwarfed by the presence of many shrub and tree pixels which are also changing over time. Truncating the frames as shown in Fig. 7 to exclude the vegetation does indeed improve the statistical difference between an event and non-event frame. However for the general WISN deployment case (with cameras facing in various directions), we do not wish to select an event acquisition technique which relies on the truncated assumption. Based on the observed statistical similarity of event and non-event frames, we wish to determine an event detector suitable for WISNs. In addition to its generality (detection not tailored to a specific type of object) and good detection performance, the chosen event detector should be implementable in the simple WISN devices. In addition to their hardware and general processing limitations, WISNs process a large volume of surveillance frames which must in turn be transmitted wirelessly to the sink if they contain an event of interest. Analysis of frames at the small block or pixel level may consequently not always be a suitable or possible approach for event acquisition.

Instead we seek a simple form for the detector where a single frame statistic is compared to a threshold in order to determine the presence or absence of an event. However as discussed, event and non-event frames from real-world surveillance sequences have similar statistics. Furthermore it can be shown that a difference image $D = B - A$ computed from two consecutive frames A and B is not perfectly Gaussian (as often assumed) but rather contains significant outliers for both event and non-event frames. An optimal non-parametric (robust) detector is thus more appropriate for this case of statistical similarity and presence of outliers. However we show that a simple "chi-squared" detector (relying on a comparison of a frame statistic to a threshold) is *equivalent in form* to the robust detector and can thus be used in WISNs (Sect. 4.3). Furthermore, through the use of composite hypothesis testing, we show that the chi-squared detector can be made uniformly most powerful (UMP) through proper threshold selection. The UMP property signifies that the detector achieves a probability of detection P_D *higher or equal to* the detection of all other detectors given the worst-case scenario probability of false alarm P_{FA}. In other words, no detector performs better given the same probability of false alarm.

4.2 Lightweight image processing detector

The simple algorithm we selected is based on difference images, similar to the techniques found in the image change detection literature (Radke et al. 2005). We describe this detector, which we refer to as the "chi-squared" detector, in relation to the detector proposed by Aach and Kaup (1993, 1995), where we use entire difference frames instead of blocks (noting that the technique is also applicable to blocks of any size). We now overview the basics of the technique. In essence, a difference image $D = B - A$ between two consecutive frames A and B reveals all the pixels that have changed between these frames containing both relevant and irrelevant changes (such as the tree swaying). The Mean Squared Error (MSE) of the difference image is computed as the relevant statistic, and it is compared to a theoretically-obtained robust threshold T. We now present the specific details of this detector.

In Aach and Kaup (1993, 1995) [and in (Radke et al. 2005)], the difference image D is computed and divided into smaller blocks. Importantly, each pixel of the *difference* image is modeled as a Gaussian random variable with 0 mean and variance σ_i^2, where $i = 0$ corresponds to a non-event frame and $i = 1$ corresponds to an event frame. In order to conserve computational energy, in this work we use the entire difference image instead of the block-based solution. The resulting detector hypothesis test can be summarized as:

$$\mathcal{H}_0 : \text{no event}, \ D_k \sim \mathcal{N}(0, \sigma_0^2) \ \forall k \tag{7}$$

$$\mathcal{H}_1 : \text{event}, \ D_k \sim \mathcal{N}(0, \sigma_1^2) \ \forall k \tag{8}$$

with $\sigma_0^2 < \sigma_1^2$ and where D_k is the kth difference pixel in $D = B - A$. Since the entire difference image is utilized in the detection, instead of considering individual pixels we may consider a new random variable defined as:

$$X = \sum_{k=1}^{n} D_k^2 = \sigma_j^2 \sum_{k=1}^{n} \frac{D_k^2}{\sigma_j^2} = \sigma_j^2 Y, \quad \text{for } J \in \{0, 1\} \tag{9}$$

where Y has distribution chi-squared with n degrees of freedom and where n is the total number of pixels in the difference frame. The new detection hypothesis test is thus given by:

$$\mathcal{H}_0 : X \sim \frac{1}{\sigma_0^2} f_{\chi^2,n}\left(\frac{x}{\sigma_0^2}\right) \tag{10}$$

$$\mathcal{H}_1 : X \sim \frac{1}{\sigma_1^2} f_{\chi^2,n}\left(\frac{x}{\sigma_1^2}\right) \tag{11}$$

where $f_{\chi^2,n}(x)$ is the probability density function (pdf) of the chi-squared distribution with n degrees of freedom.[1] Significantly, the hypothesis test to distinguish between an event and non-event is given by the comparison of a single statistic (x) to a threshold T as shown in Eqs. 12 and 13, where σ_0^2 is the variance of a null frame, $F^{-1}_{\chi^2,n}$ is the inverse chi-squared cumulative distribution function (cdf) and α is the desired probability of false alarm.[2]

$$x \underset{H_0}{\overset{H_1}{\gtrless}} T \tag{12}$$

$$T = \sigma_0^2 F^{-1}_{\chi^2,n}(1-\alpha) \tag{13}$$

4.3 Lightweight detector properties

In this Section we wish to analyze some of the properties of the simple chi-squared detector of Eqs. 12 and 13. We begin by showing that the simple chi-squared detector can be made uniformly most powerful (UMP) (Van Trees 2001). To achieve this we show that if there exists a real positive number γ, such that $\sigma_0^2 < \gamma$ and $\sigma_1^2 > \gamma$, where the actual σ_0^2, σ_1^2 are *unknown*, then there exists a UMP detector where a realization x from Eq. 9 is compared to a threshold T, such that the probability of false alarm $P_{FA} = \alpha$ is given by:

$$\alpha = \sup_{\sigma_0^2 < \gamma} \int_T^\infty \frac{1}{\sigma_0^2} f_{\chi^2,n}\left(\frac{x}{\sigma_0^2}\right) dx \tag{14}$$

Proposition 1 *Suppose there exists a $\gamma > 0$, such that $\sigma_0^2 < \gamma$ and $\sigma_1^2 > \gamma$ in Eqs. 10 and 11. Then there exists a UMP test of the form*

$$x \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma f^{-1}_{\chi^2,n}(1-\alpha) \tag{15}$$

for false alarm rate not exceeding α. Proposition 1 is a composite hypothesis test in which the parameters for the null and alternate hypotheses are unknown, but the regions for these parameters are divided by a threshold γ. The proposition says that if the parameter space is divided as thus, then a test that compares the actual x in Eq. 9 to a threshold, achieves optimal detection when the worst case false alarm is considered (the use of sup *in Eq. 14).*

[1] We note that σ_i appears in the detector as σ_0^2 while it appears as $\frac{1}{\sigma_0^2}$ and $\frac{1}{\sigma_1^2}$ in the distributions of the two hypotheses.

[2] The threshold T is obtained directly by writing $\alpha = Pr\{\text{announce}\mathcal{H}_1|\mathcal{H}_0\}$ which results in an integration of the null hypothesis pdf $\frac{1}{\sigma_0^2} f_{\chi^2,n}(\frac{x}{\sigma_0^2})$ over the interval from T to ∞.

Proof If we can show that the likelihood ratio is monotonically increasing in x for $\sigma_1^2 > \sigma_0^2$, then the UMP test of the form in Eq. 15 follows from a theorem on composite hypothesis testing (Van Trees 2001). It can easily be shown that the log-likelihood ratio is given by $\frac{1}{2}\left(\frac{1}{\sigma_0^2} - \frac{1}{\sigma_1^2}\right) x + \frac{n}{2} \ln \left(\frac{\sigma_0^2}{\sigma_1^2}\right)$. Since $\sigma_1^2 > \sigma_0^2$, this ratio is strictly increasing in x. To show that T is as given on the left side of Eq. 15, we note that the probability of false alarm is given by $1 - f_{\chi^2,n}(T/\sigma_0^2)$ by applying an integration change of variable in Eq. 14. To get the sup in Eq. 14, it suffices to set $\sigma_0^2 = \gamma$.

Having established that the simple detector of Eqs. 12 and 13 can be made uniformly most powerful, we next show that the form of the detector is equivalent to that of a robust (non-parametric) detector. This is an important property given that the statistical similarity of event and non-event frames along with difference-frame distributions that are not quite Gaussian render $\mathcal{H}_1$ and $\mathcal{H}_0$ almost indistinguishable when the entire frame is used. Thus we would like to maximize the event detection assuming that $\sigma_1^2 \approx \sigma_0^2$ rather than assuming that the statistics are significantly different. This can be re-phrased as

$$\max \frac{\partial \beta}{\partial \sigma_1^2} |_{\sigma_1^2 = \sigma_0^2} \tag{16}$$

where $\beta = Pr\{\text{declare } H_1 \mid H_1 \text{ occurs}\}$ is the probability of detection.

Proposition 2 *The test*

$$x \underset{H_0}{\overset{H_1}{\gtrless}} T \tag{17}$$

maximizes Eq. 16 for a false alarm rate not exceeding α*, i.e.* T *is chosen so that*

$$\alpha > \int_T^\infty \frac{1}{\sigma_0^2} f_{\chi^2,n}\left(\frac{x}{\sigma_0^2}\right) dx. \tag{18}$$

Proof By the proof of the Neyman-Pearson lemma (Van Trees 2001), the optimal test can be shown to be of the form

$$\frac{\frac{\partial \frac{1}{\sigma_1^2} f_{\chi^2,n}\left(\frac{x}{\sigma_1^2}\right)}{\partial \sigma_1^2}\Big|_{\sigma_1^2 = \sigma_0^2}}{\frac{1}{\sigma_0^2} f_{\chi^2,n}\left(\frac{x}{\sigma_0^2}\right)} \underset{H_0}{\overset{H_1}{\gtrless}} \tilde{T}, \tag{19}$$

which is equivalent to

$$\frac{x - n\sigma_0^2}{2\sigma_0^4} \underset{H_0}{\overset{H_1}{\gtrless}} \tilde{T}. \tag{20}$$

Letting $T = 2\sigma_0^4 \tilde{T} + n\sigma_0^2$ proves the proposition.

In summary, given the actual statistics of the difference image, a non-parametric (robust) detector is appropriate to perform event detection. However the simple chi-squared detector is equivalent in form to the robust detector and can be made uniformly most powerful through threshold selection. The simple image difference test may thus be used at the camera nodes with acceptable performance within its class of algorithm complexity.

Table 2 Event detection based on lightweight image processing (LIP) organized in order of decreasing detection performance P_D

Image sequence	Description	P_D	P_{FA}
Seq. (1)	Indoor walking	1.0	0.13
Seq. (4a)	Outdoor car (no trees)	0.98	0.17
Seq. (2)	Outdoor car (with trees)	0.87	0.26
Seq. (4b)	Outdoor walking (no trees)	0.50	0.03
Seq. (3)	Outdoor walking (with trees)	0.05	0.23

4.4 Lightweight detector performance

Given these desirable properties, we would like to examine how the visual event acquisition algorithm performs on the real-world surveillance sequences described in Sect. 4.1. Table 2 shows the performance results obtained for these sequences arranged in order of *decreasing* performance. We make several key observations regarding these results. The first observation is that the *median* performance result (corresponding to Seq. 2) is quite similar to the results obtained in Rahimi et al. (2005) despite differences in the form of the exact LIP algorithm that is utilized [in (Rahimi et al. 2005) the average reported $P_D = 0.78$ compared with our $P_D = 0.87$ and the average reported $P_{FA} = 0.22$ compared with our $P_{FA} = 0.26$]. This result is encouraging in that Seq. 2 corresponds to an unknown object moving in difficult outdoor conditions with significant lighting changes and the presence of extraneous motion. Thus despite their simple nature, LIP algorithms for event acquisition do hold some promise. The second observation from Table 2 is that the *actual* $P_D - P_{FA}$ performance varies greatly depending on the specific image sequence. It is thus very difficult to guarantee a given level of performance in the camera node for an *arbitrary* sequence.

If we classify the image sequences into broad categories based on their characteristics, a coarse level of performance prediction may be possible. For instance, sequences with minimal levels of extraneous motion achieve a better overall performance than sequences afflicted with such motion. Sequences where an object occupies a larger portion of the overall frame (such as a car rather than a person) also show improved $P_D - P_{FA}$ performance. Though intuitive, these observations do not provide much assistance for the general WISN case where camera nodes may encounter conditions that vary appreciably over time. We thus seek a collaborative approach between camera nodes equipped with LIP algorithms and sensors to help capture the value of visual detection while addressing the large variability in its performance.

5 Reliability of sensor decisions in uncertain environments

In Sect. 4 we examined the properties and performance of a generic lightweight image processing algorithm (LIP) and determined that although promising, the performance exhibited considerable variability depending on conditions. To tap into the promise of LIP algorithms and address this variability, we wish to investigate the role of sensors to improve the event acquisition performance. As discussed in Sect. 3.1, sensor decisions regarding the presence or absence of an event can be made available to the camera nodes directly. This simple augmentation may result in a performance improvement (Sect. 6.1) with the caveat that sensors may themselves be prone to fault or attack in unattended environments. An alternative approach is to employ an intermediate mechanism where an entity (such as a cluster head) receives readings from one or more sensors as shown in Fig. 2 and performs attack/fault detection. In this Section we wish to investigate the form and performance of such a cluster head (CH) detector in the face of both occasional errors (faults) and in the face of deliberately stealthy attacks (Sect. 3.2). We proceed by obtaining the form of the optimal Neyman-Pearson (NP) detector in Sect. 5.1, analyzing the effects of a stealthy attack in Sect. 5.2 and finally obtaining the performance of the CH detector in Sect. 5.3.

5.1 Cluster head (CH) detector

In detection problems we are generally faced with the task of deciding between two or more hypotheses based on received data. The Neyman-Pearson (NP) is an optimal detector appropriate for cases where a priori probabilities of the hypotheses are not available, and for cases where the probability of detection P_D and the probability of false alarm P_{FA} may not be of equal significance to the application (otherwise a Bayesian detector may be appropriate). According to the NP approach, we obtain a detector by maximizing P_D for a desired false alarm rate $P_{FA} = \alpha$. The resulting optimal detector is a *likelihood ratio* detector $\Lambda(\mathbf{z})$ given by Eq. 21, where $\mathbf{z}$ is the received data vector and where the comparison threshold $\mathcal{T}$ is chosen according to Eq. 22.

$$\Lambda(\mathbf{z}) = \frac{p(\mathbf{z}; \mathcal{H}_1)}{p(\mathbf{z}; \mathcal{H}_0)} \underset{H_0}{\overset{H_1}{\gtrless}} \mathcal{T} \tag{21}$$

$$P_{FA} = \int_{\mathbf{z}:\Lambda(\mathbf{z})>\mathcal{T}} p(\mathbf{z}; \mathcal{H}_0)\, d\,\mathbf{z} \leq \alpha \tag{22}$$

For the case of n binary sensors of Sect. 3.2, the data vector $\mathbf{z}$ consists of Bernoulli random variables from a distribution which is $Bern(p)$ (hypothesis $\mathcal{H}_0$) or a distribution Bern(r) (hypothesis $\mathcal{H}_1$). By applying Eq. 21 it can easily be shown that the NP detector for this case is given by Eq. 23, where $w(\mathbf{z})$ is the weight (the number of 1 s) in the data vector $\mathbf{z}$.

$$\Lambda(\mathbf{z}) = \frac{r^{w(\mathbf{z})}(1-r)^{n-w(\mathbf{z})}}{p^{w(\mathbf{z})}(1-p)^{n-w(\mathbf{z})}} \underset{H_0}{\overset{H_1}{\gtrless}} \mathcal{T} \tag{23}$$

$$\sum_{\mathbf{z}:\Lambda(\mathbf{z})>\mathcal{T}} p^{w(\mathbf{z})}(1-p)^{n-w(\mathbf{z})} \leq \alpha \tag{24}$$

The threshold $\mathcal{T}$ is chosen to satisfy a desired α based on Eq. 24, where the summation is over all the possible data vectors $\mathbf{z}$ such that $\Lambda(\mathbf{z})$ exceeds $\mathcal{T}$. However this is equivalent to summing over all possible weights w for $w \in [0, n]$ as shown in Eq. 25. The notation $\mathbf{I}_P$ denotes the indicator function which is equal to 1 if the proposition P is true and is equal to 0 otherwise. Finally, the probability of detection β resulting from the use of the $\Lambda(\mathbf{z})$ detector is given by Eq. 26.

$$\sum_{w=0}^{n} \binom{n}{w} p^{w(\mathbf{z})}(1-p)^{n-w(\mathbf{z})}\, \mathbf{I}_{\,\Lambda(\mathbf{z})>\mathcal{T}} \ \leq \alpha \tag{25}$$

$$\beta = \sum_{w=0}^{n} \binom{n}{w} r^{w(\mathbf{z})}(1-r)^{n-w(\mathbf{z})}\, \mathbf{I}_{\,\Lambda(\mathbf{z})>\mathcal{T}} \tag{26}$$

To distinguish between normal operation and an attack/fault, the cluster head must therefore employ the detection statistic (likelihood ratio Λ) of Eq. 23 and compare it to the threshold T that is set based on Eq. 25 with the resulting probability of detection given by Eq. 26. Importantly, the detection statistic $\Lambda(p, r)$ depends on the value of p and r. Hence in the case of an attack or fault with unpredictable probability, the detection statistic depends on the *unknown* underlying parameter q. The detection statistic's dependence on q also translates into difficulties in determining the probability of false alarm and detection based on the dependence of Eqs. 25 and 26 on Λ. Thus although we have determined an optimal attack/fault detector for the cluster head, it is not implementable in its current form unless the parameter q is known. Fortunately we can re-arrange the likelihood ratio Λ as shown in Eq. 27 where $\frac{(1-r)^n}{(1-p)^n}$ is equal to some positive constant $k > 0$ for all values of $0 \le r \le 1$, $0 \le p < 1$ and n.

$$\Lambda(z) = \frac{r^{w(\mathbf{z})}(1-r)^{n-w(\mathbf{z})}}{p^{w(\mathbf{z})}(1-p)^{n-w(\mathbf{z})}} = \frac{(\frac{r}{1-r})^{w(\mathbf{z})}}{(\frac{p}{1-p})^{w(\mathbf{z})}} \frac{(1-r)^n}{(1-p)^n} = k\left(\frac{r}{p}\right)^{w(\mathbf{z})}\left(\frac{1-p}{1-r}\right)^{w(\mathbf{z})} \tag{27}$$

Let us for the moment assume that $r > p$ in Eq. 27. Then it can easily be seen that $\Lambda(w) = k\left(\frac{r}{p}\right)^w\left(\frac{1-p}{1-r}\right)^w$ is monotonically increasing in w where we have written Λ in terms of w to simplify the notation and emphasize the role of the "aggregate" statistic (the weight w) in lieu of the original data vector $\mathbf{z}$. Based on the monotonicity of Λ, we may now invoke the Karlin-Rubin theorem (Van Trees 2001) to obtain an alternative form for the CH detector with the same $P_D - P_{FA}$ performance. The alternative form for the CH detector is given by Eqs. 28–30 where p' is a probability of mixing between the two hypotheses if w is precisely equal to T (the mixing probability p' is set based on Eq. 29 for a desired α). As shown in Eq. 31, based on the assumption that $r > p$, these equations are valid for the interval where $p < \frac{1}{2}$. When $p \in [\frac{1}{2}, 1]$, the hypotheses in Eq. 28 are switched.

$$w \underset{H_0}{\overset{H_1}{\gtrless}} \mathcal{T} \quad \text{if} \quad w = \mathcal{T} \quad \text{then declare} \quad H_1 \quad \text{w.p} \quad p' \tag{28}$$

$$\alpha = \sum_{w>\mathcal{T}}^{n} \binom{n}{w}\left(\frac{p}{1-p}\right)^w (1-p)^n + p'\binom{n}{\mathcal{T}}\left(\frac{p}{1-p}\right)^{\mathcal{T}}(1-p)^n \tag{29}$$

$$\beta = \sum_{w>\mathcal{T}}^{n} \binom{n}{w}\left(\frac{r}{1-r}\right)^w (1-r)^n + p'\binom{n}{\mathcal{T}}\left(\frac{r}{1-r}\right)^{\mathcal{T}}(1-r)^n \tag{30}$$

$$r > p \Rightarrow \quad p + q - 2pq > p \quad \Rightarrow \quad p < \frac{1}{2} \tag{31}$$

We make a few key observations regarding this result. The detection statistic w and the comparison threshold T_h in Eq. 28 no longer require the cluster head to know the value of q and thus the detector is implementable at the cluster head. We note however that in order to determine the resulting probability of detection P_D, the value of r (and thus q) is still required in Eq. 30. Thus for the case of an attack, analysis of the optimal attack q is beneficial in determining the detector's performance and we address this issue in Sect. 5.2.

We also note that the detection statistic is now a simple weight and thus the cluster head must merely count the number of 1's that it has received from the n sensors and compare this count to a threshold. The optimal NP detector at the cluster head is thus identical in form to the detectors commonly used in practice as discussed in Sect. 3.2 (where the weight was set based on experimental trials or based on an expected average count of $c \pm \epsilon$). The optimal NP detector however makes use of a threshold T that is set based on a desired probability of false alarm α and based on the probability p of an event. Setting the threshold based on Eq. 29 thus provides a greater level of control and flexibility to meet the P_{FA} requirements of the application. Furthermore, this detector is guaranteed to provide the best probability of detection P_D for a chosen $P_{FA} = \alpha$ (a property of Neyman-Pearson detectors). In Sect. 5.3 we examine the actual performance of this cluster head (CH) detector and compare it to the performance of a detector based on the expected average $c \pm \epsilon$.

5.2 Attack analysis results

As discussed in Sect. 3.2, sensors deployed in unattended environments may experience errors due to occasional faults or distributed (stealthy or unstealthy) attacks (Buttyan and Hubaux 2002; Czarlinska and Kundur 2008). In the case of occasional faults, the probability of error q for sensor i may be quite small while in the case of a general attack, this probability may be arbitrarily large and unknown a priori. Instead of proceeding with an arbitrary attack strategy however, an intelligent attacker may take into account the presence of an attack/fault detection mechanism such as the NP-based detector or the $c \pm \epsilon$ detector at the cluster head. In this case the attacker wishes to select a probability of attack q that minimizes the probability of attack detection. This requirement translates into a stealth condition where the attacker wishes to select q such that $Pr\{w(\mathbf{X}_p) = w(\mathbf{Z}_{p,q})\}$ is maximized. Through combinatorial analysis (Czarlinska et al. 2007), this probability can be expressed as shown in Eqs. 32 and 33 where we have generalized the stealth condition to allow a deviation of $\epsilon_r \in \mathcal{Z}$ away from perfect stealth and where a, b and c are binomial coefficients defined in Eq. 33.

$$Pr\left\{\left|w(\mathbf{Z}_{p,q}) - w(\mathbf{X}_p)\right| < \epsilon_r\right\} = \sum_{m=\lceil \frac{l-\epsilon_r}{2} \rceil}^{\lfloor \frac{l+\epsilon_r}{2} \rfloor} \sum_{k=1}^{n} \sum_{l=1}^{n} a\,(k,m)\, b\,(k,l,m)\, c(k) \cdot p^k (1-p)^{n-k} \cdot q^l (1-q)^{n-l} \tag{32}$$

$$a(k,m) = \begin{cases} \binom{k}{m} & \text{if } k \geq m \\ 0 & \text{o.w} \end{cases} \tag{33}$$

$$b(k,l,m) = \begin{cases} \binom{n-k}{l-m} & \text{if } n-k \geq l-m \\ 0 & \text{o.w} \end{cases} \tag{34}$$

$$c(k) = \begin{cases} \binom{n}{k} & \text{if } n \geq k \\ 0 & \text{o.w} \end{cases} \tag{35}$$

The stealth condition of Eq. 32 is unfortunately cumbersome to inspect. Plotting Eq. 32 for different values of cluster size n and probability of an event p nevertheless yields a unique value of probability q that maximizes the equation (i.e. it is the global peak of Eq. 32). The results of such plotting are summarized in Table 3 which shows the optimal value of attack probability q^* for each pair (n, p) (the stealth condition is symmetric in p and thus the effect of p is the same as the effect of $1 - p$ and we only consider $p \in [0, 0.5]$). We make a few key observations regarding this result. For a given value of cluster size n (i.e. a row in

Table 3 Optimal q^* value for cluster size n and probability of event p

n	$p = 0.01$	$p = 0.05$	$p = 0.1$	$p = 0.2$	$p = 0.3$	$p = 0.4$	$p = 0.5$
1	0	0	0	0	0	0	0
2	0.9990	0.9990	0.9990	0.9990	0.9990	0.9990	0.9990
5	0.4050	0.4050	0.4050	0.4550	0.5050	0.5050	0.5050
10	0.2050	0.2050	0.2050	0.2550	0.2550	0.9990	0.9990
20	0.1010	0.1050	0.1100	0.1220	0.1350	0.1470	0.1520
30	0.0670	0.0700	0.0740	0.0820	0.0910	0.0990	0.1030
40	0.0500	0.0530	0.0550	0.0620	0.0680	0.0750	0.0780
50	0.0400	0.0420	0.0440	0.0490	0.0550	0.0600	0.0620
100	0.0200	0.0210	0.0220	0.0250	0.0280	0.0300	0.0310

Table 3), q^* is almost constant to within one significant digit irrespective of the value of probability p. Thus if the attacker knows the cluster size n, he can determine the optimal value of attack without having to know the probability p. This is significant since the value of p depends in part on a sensor's threshold selection T_h and may not always be available to the attacker.

Importantly, the results of Table 3 can also be obtained analytically from Eq. 32 by applying game theoretic analysis where the attacker is treated as one "player" who does not know the value of p and by treating the sensor network as the other "player" who does not know the value of attack q (Czarlinska et al. 2007). Such analysis also reveals that the optimal value of p for $p \in [0, 0.5]$ is p^* *small* (the optimal value of p for $p \in [0.5, 1]$ is p^* *large*). This suggests that to improve the attack detection, the sensors should be calibrated to have a small (or large) value of p through threshold T_h selection if such selection is possible (depending on the underlying technology of the sensor). Indeed if we examine more significant digits in the results of Table 3, the optimal value of q^* does indeed appear to decrease with decreasing p with possible ramifications for attack detection as will be shown in Sect. 5.3.

Examination of Table 3 also yields important insights regarding the relationship between the cluster size n and the optimal attack parameter q^*. We observe that as n increases, q^* decreases for all values of p. This result can be understood in the context of typical sets if we consider the n sensor decisions as a string of length n. The typical set is usually a small set but with probability of occurrence close to 1. When n is small, the typical set is small but relatively large compared to the set of all possible strings of length n. When n increases, the set of all possible strings of length n grows to be very large and the size of the typical set is *relatively* much smaller. Thus it becomes more difficult for the attacker to attack the "string" and still remain in the typical set. This implies that the chance of attack on the sensors decreases but this decrease may also carry ramifications for the *detection* of such an attack as we will explore in Sect. 5.3.

Finally we make an observation regarding the optimal strategy for the $n = 1$ case. When there is only one sensor and there is *no attack detection mechanism*, intuitively the attacker should always attack. However if an attack detection mechanism is present and the attacker wishes to be *stealthy*, the optimal value of attack is $q^* = 0$ (with only 1 bit, the attacker has a very low probability of fooling the detector and should avoid an attack altogether). The situation changes dramatically for $n > 1$ due to the underlying combinatorics in Eq. 32 which dictate that the probability of attack should be high but decreasing with increasing n. The implications of these results for attack detection are examined in Sect. 5.3.

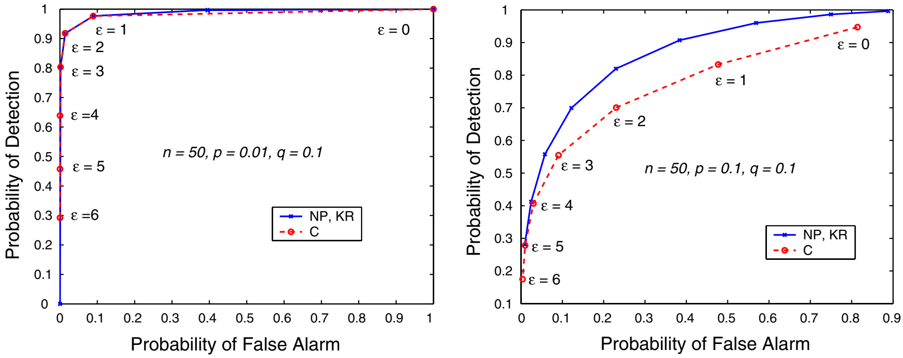

Fig. 8 **a** $n = 50$, $p = 0.01$ and $q = 0.1$. **b** $n = 50$, $p = 0.1$ and $q = 0.1$

5.3 Cluster head detector performance

As discussed in Sect. 5.1, a detector for attack/fault identification at the cluster head may be based on the Neyman-Pearson design or on an average expected count c where $c \approx np \pm \epsilon$ for a cluster size of n sensors with probability of event p. In this section we wish to compare the performance of these two approaches and also obtain some general insights into the $P_D - P_{FA}$ performance curve for different values of p and n as well as for various attack probabilities q.

In Figs. 8 and 9, the $P_D - P_{FA}$ performance of the NP and $c \pm \epsilon$ detectors are depicted for $n = 50$ sensors, various (p, q) pairs and various values of ϵ "slack" in the c detector. As can be seen from these figures, the performance of the average expected c detector follows the general trend of the NP detector although it typically does not achieve the same overall performance. Nevertheless, by choosing a different value of ϵ, it is possible to achieve a desired trade-off between the probability of detection P_D (vertical axis) and the probability of false alarm P_{FA} (horizontal axis). The average expected c detector may thus be useful for certain applications, in particular ones where the probability of an event p is expected to be small (based on the phenomenon of interest and the selection of the sensor threshold T_h) as in Fig. 8a. As a side note, we observe that the performance of the NP detector based on Eqs. 23, 25 and 26 which we denote by "NP" is the same as the performance of the NP detector based on the Karlin-Rubin simplification of Eqs. 28–30 which we denote by "KR" in Figs. 8 and 9. Thus we are justified in utilizing the simplified form of the NP detector to obtain the same performance.

Finally based on Figs. 8 and 9, we observe that the NP detector performs better for smaller values of probability p for $p \in [0, 0.5]$ (by symmetry for $p \in [0.5, 1]$ it performs better for values of p closer to 1). This result is consistent with the results of Sect. 5.2 where based on analysis of the stealth condition of Eq. 32 we noted that it was best to calibrate the sensors to a small value of p or to a large value of p. Indeed based on Figs. 8 and 9 and the results of Table 3, the worst $P_D - P_{FA}$ performance is obtained for values of p closest to $p = 0.5$. This can be understood from Eq. 31 where $r = p + q - 2pq$. When $p = 0.5$, $r = 0.5$ and thus the detector is not able to distinguish between the H_0 and H_1 hypotheses. At an intuitive level, when $p = 0.5$, the probability of obtaining a decision of value 1 is the same as the probability of obtaining a sensor decision of value 0. This situation corresponds to the largest level of uncertainty that the cluster head can experience and makes it easier for an attacker to fool the detector.

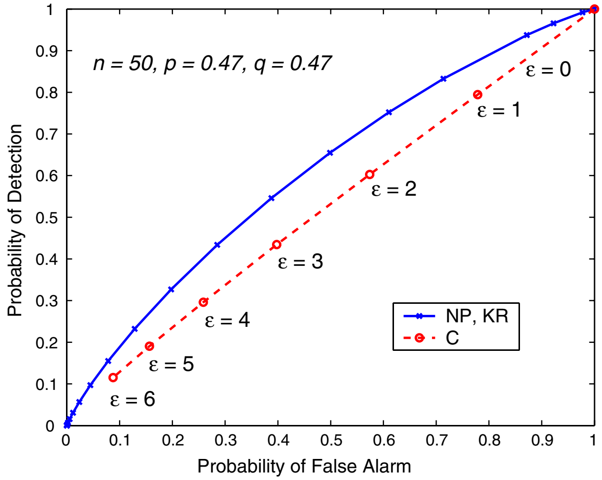

Fig. 9 $n = 50$, $p = 0.47$ and $q = 0.47$

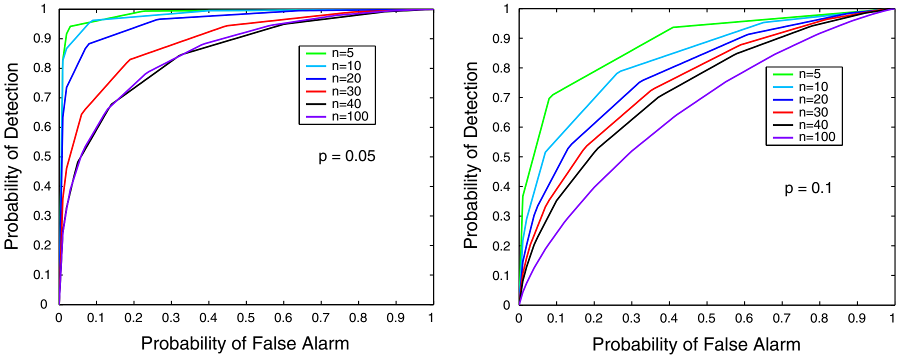

Fig. 10 The $P_D - P_{FA}$ performance curves for various values of n and the optimal value of q^* corresponding to that n for **a** $p = 0.05$ and **b** $p = 0.1$

Having compared the performance of the NP and c detectors for various values of p, we now wish to investigate the performance of the NP detector for an *optimal* attack parameter q^*. Fig. 10a depicts the $P_D - P_{FA}$ performance of the NP detector for $p = 0.05$ and for various values of cluster size n. Crucially, each performance curve is obtained assuming the value of attack parameter q that is *optimal* for the given (p, n) pair as obtained in Table 3. Figure 10b is obtained similarly but for a value of $p = 0.1$. Based on these two figures, we make the important observation that the $P_D - P_{FA}$ performance *decreases* as the number of sensors n is *increased*. This somewhat surprising result is an outcome of the stealth condition of the attacker. That is, *if* the attacker is *not* stealthy, increasing the number of sensors will increase the detection performance. If however the attacker is *stealthy*, then he selects an optimal value of q^* based on Eq. 32 or Table 3. As discussed in Sect. 5.2, this optimal value of q decreases with increasing cluster size n. Thus the attack becomes more rare (which is a desirable property) but by the same token becomes more difficult to detect when it does occur.

Based on the performance of the NP detector in the face of an optimal attack q^*, we are brought back to the question of the level of sensor redundancy that should be employed to detect the attack and to suitably complement the event acquisition already provided through lightweight image processing.

6 Performance of combined approach for event acquisition

As discussed in Sect. 4, lightweight image processing (LIP) may offer a suitable $P_D - P_{FA}$ performance for event acquisition in WISNs. This performance however exhibits significant variability depending on the conditions experienced during image capture. To exploit the potential of LIP algorithms while reducing their variability, we wish to augment the event acquisition process with sensors. As explored in Sect. 5 however, sensors are prone to occasional faults or deliberate and stealthy attacks. We may thus consider a variety of methods at the camera nodes for exploiting the sensor decisions and the LIP algorithm to improve reliability as discussed in Sect. 3.1. In this section we wish to explore the implications of utilizing a cluster head fault/attack detector mechanism before making the sensor decisions available to the camera nodes. Based on this input, the camera nodes may trust the sensor decisions, the LIP decision or rely on a combination of both (Sect. 6.2). For perspective, we wish to compare this scenario with the case where sensor decisions are made available directly to the camera nodes *without* a cluster head fault/attack detection mechanism (Sect. 6.1).

6.1 Direct sensor decisions approach

In this Section we focus on the $P_D - P_{FA}$ performance of a camera node augmented with a single sensor decision that is made directly available to the camera node as depicted in Fig. 11. Thus to perform event acquisition under this scenario, a camera node has access to two sources of information regarding the possible occurrence of an event. As shown in Fig. 11, if there is disagreement between the two sources regarding the presence of an event, the camera node may trust the lightweight image processing (LIP) over the potentially faulty or attack-prone sensor. Alternatively, the camera node may trust the sensor decision (SN) in lieu of the variability-prone LIP algorithm or may take the "safe" strategy of declaring an event if *either* of the sources reports an event.

$$P_D = \frac{\text{No.(event | event frame)}}{\text{No.(total frames)}} \tag{36}$$

$$P_{FA} = \frac{\text{No.(event | non-event frame)}}{\text{No.(total frames)}} \tag{37}$$

Figures 12 and 13 show the simulation results obtained for this scenario where we use the notation LIP to denote lightweight image processing, SN to denote the sensor decision and "SN & LIP" to denote use of the combination. Figure 12a depicts the results obtained for the image sequence of Fig. 6 where an unidentified individual walks through an environment with substantial background movement and is periodically obscured due to the presence of trees. Figure 12b shows the results for a truncated version of this sequence corresponding to Fig. 7b where the camera's field of view largely excludes the trees. Figure 13a shows the results for

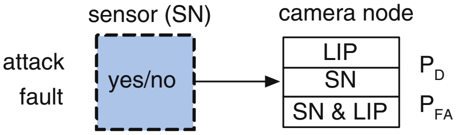

Fig. 11 For event acquisition, each camera node may utilize lightweight image processing (LIP), a sensor decision (SN), or rely on both to determine the presence or absence of an event

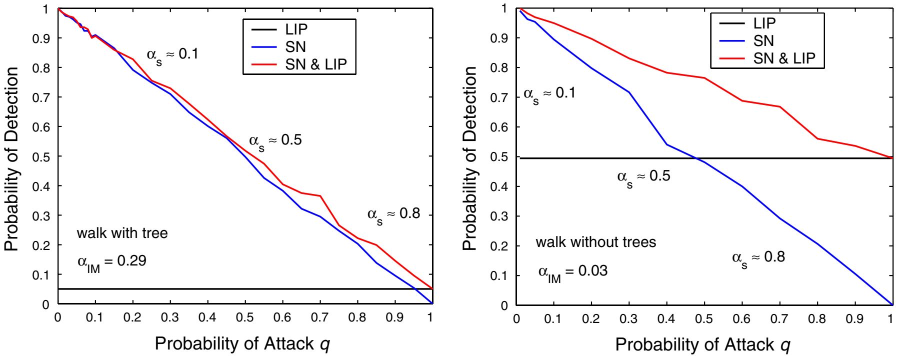

Fig. 12 Probability of detection P_D versus probability of sensor error q for the LIP, SN and SN & LIP approaches for image sequence **a** walking with trees from Fig. 6 **b** walking without trees from Fig. 7b

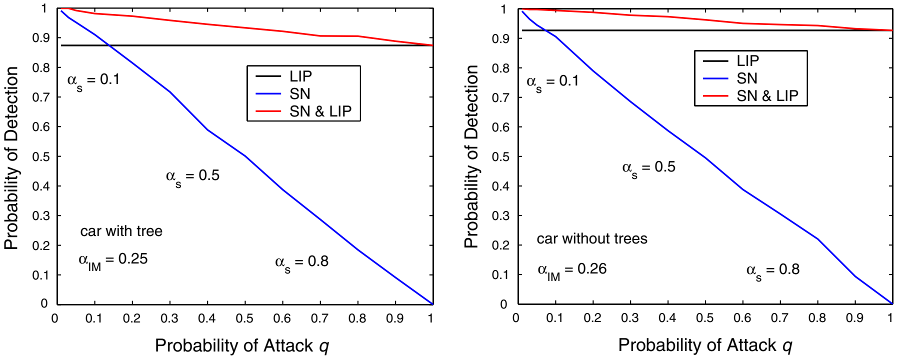

Fig. 13 Probability of detection P_D versus probability of sensor error q for the LIP, SN and SN & LIP approaches for image sequence **a** car with trees from Fig. 5 **b** car without trees from Fig. 7a

the image sequence of Fig. 5 where the passing of an unknown vehicle is captured in the presence of significant background variability. Finally Fig. 13b shows the results obtained for a truncated version of the passing car and corresponds to the image sequence depicted in Fig. 7a. Importantly, the horizontal axes in Figs. 12 and 13 correspond to the probability q of sensor error (due to attack or fault) and the vertical axes correspond to the probability of detection P_D. The resulting probability of false alarm for the sensor α_s is also shown in the figures for each P_D segment along with the probability of false alarm for the LIP algorithm α_{IM}. These probabilities are obtained experimentally based on Eqs. 36 and 37 where No. denotes the number of frames where a certain type of decision was made.

Based on Figs. 12 and 13 we confirm that the $P_D - P_{FA}$ performance of the lightweight image processing exhibits great variability from sequence to sequence. As expected, this performance tends to improve for image sequences with better characteristics (such as less background variability). The level of improvement itself however experiences variability as can be seen by comparing Fig. 12a, b with Fig. 13a, b. We also note the inherent result that the probability of detection P_D and the probability of false alarm P_{FA} for the image processing algorithm remain constant over the entire range of q. This is fully expected since the visual algorithm at the camera node is independent of the sensor readings. The lack of predictability and control in the LIP algorithm for an arbitrary sequence is a visible disadvantage. However the constancy of the performance over the range of q is a clear advantage as can be seen by comparing the SN and SN & LIP curves in Figs. 12 and 13.

Specifically, the $P_D - P_{FA}$ performance of the single sensor is excellent for a *small* probability of error q. However if this probability of error is caused by an attack, it may become arbitrarily large and dramatically decrease the event acquisition performance (the performance decreases linearly with increasing q). Indeed in this setup the sensor does not perform a fault/attack detection and therefore an attacker *need not be stealthy* in the attack, but rather choose *any* implementable probability of attack q. The combined SN & LIP approach however inherits the best of both approaches; the good performance of the sensor given a small q and the invariance of the LIP approach over the range of q. Thus we observe that the $P_D - P_{FA}$ performance of the combined approach is always *better than or equal to* the *best* performance from among the other two methods. Specifically, for a fixed probability of false alarm P_{FA}, the probability of detection P_D of the combined approach is described by Eq. 38. It is important to note that the probability of false alarm of the sensor α_s varies depending on q. Thus for a direct comparison of P_D with the LIP, we have to locate the point on the SN & LIP curve where $\alpha_s \approx \alpha_{IM}$ in order to obtain the result of Eq. 38 (as can be seen from the figures however, the SN & LIP curves lie above the LIP curve for all values of q).

$$P_{D_{\text{combined}}} \geq \max(P_{D_{\text{LIP}}}, P_{D_{\text{SN}}}) \tag{38}$$

Since combining the LIP with a *single* sensor decision *without* cluster head checking improves the performance in uncertain environments, it is important to determine if including cluster head checking and increasing the number of sensors results in a justifiably improved level of performance. Indeed the performance achieved with the "direct sensor approach" may be sufficient for certain applications.

6.2 Cluster head aided event acquisition

In this Section we wish to investigate the role of cluster head checking based on the CH detector of Sect. 5 and sensor redundancy n in improving the $P_D - P_{FA}$ performance in uncertain environments.

We begin by comparing the relative performance of the LIP algorithm with the performance of the CH detector (based on decisions from n sensors). Figures 14–17 show the event acquisition performance of multiple sensors *with* cluster head detection. Figure 14 corresponds to the image sequence of an individual walking with the presence of trees while Fig. 15 corresponds to the image sequence of the individual without the background trees. Figure 16 corresponds to the car sequences with the presence of trees while Fig. 17 corresponds to the car sequence without background trees. As before the horizontal axis represents the probability of attack q while the vertical axis corresponds to the probability of detection P_D. The probability of detection P_D of the LIP algorithm for each sequence is also shown in the figures for comparison and all probabilities are determined experimentally from Eqs. 36 and 37. As before, the notation IM denotes the image processing (LIP) based performance. The number of sensors n reporting their decisions to the cluster head is varied from $n = 1$ to $n = 40$.

In assessing the relative performance of the LIP algorithm and the CH detector with n sensors in Figs. 14–17 we maintain the *same* probability of false alarm α. That is, we set $\alpha_s = \alpha_{IM} \doteq \alpha$. This is made possible through the adjustability of the CH detector from Sect. 5 through Eq. 29. In contrast, in Sect. 6.1 a single sensor was used instead of a CH detector. Thus the probability of false alarm α_s was not adjustable but rather varied with the attack parameter q.

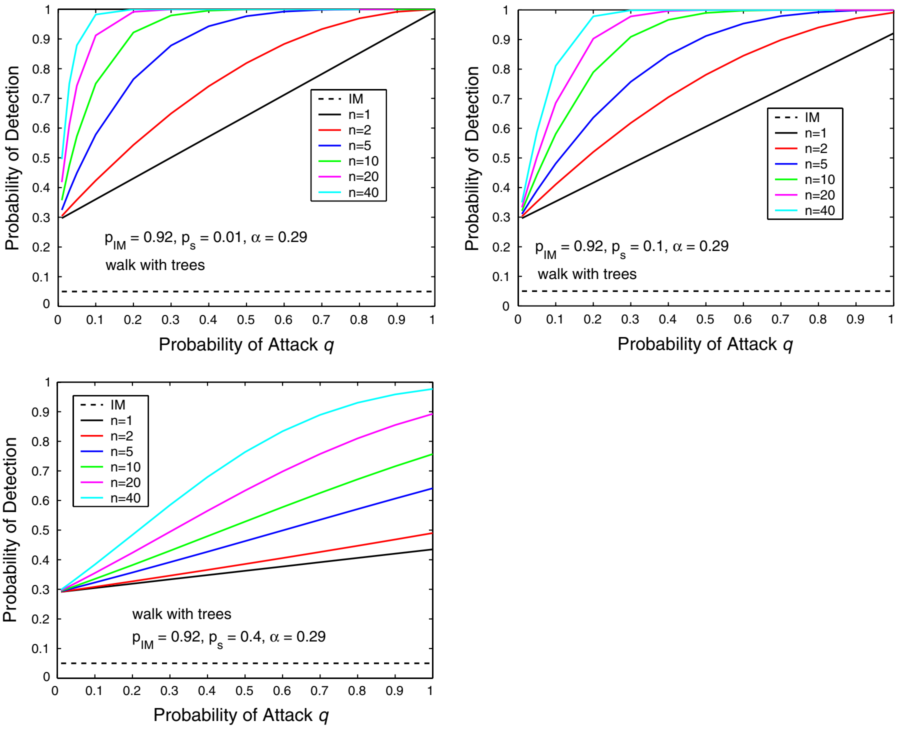

Fig. 14 P_D versus q for walking with trees from Fig. 6. **a** $p_s = 0.01$. **b** $p_s = 0.1$. **c** $p_s = 0.4$

Furthermore in Figs. 14–17 we use the notation p_s to denote the probability p of an event as witnessed by a sensor (the subscript s is used to emphasize that this probability corresponds to the conditions experienced by the sensor). We use the notation p_{IM} to denote the probability that a *frame* in a given image sequence contains an event. In Sect. 6.1, p_s was implicitly set to p_{IM}, however in this section we relax this constraint to investigate the role of the probability of an event at a sensor p_s as well as the role of cluster size n.

We make some key observations based on the results of Figs. 14–17. The first key point is that use of the CH detector fundamentally changes the relationship between the probability of detection P_D and the probability of sensor error q. Specifically, use of the CH detector eliminates the linear decrease in sensor performance with increasing q. Indeed the sensors achieve a better detection performance for higher values of q which also results in a relatively good P_D over a much wider range of q. This result is an inherent outcome of the properties of detectors which perform better when there is a significant difference between the hypotheses (in this case the values of probabilities p_s and q). Use of the detector will thus increase the detection performance for the case of unstealthy attacks (i.e. attacks with a large probability q relative to the cluster size n). This is in contrast with the results of Sect. 6.1 where an increase in q degraded the detection performance. Unfortunately use of the detector alone (without LIP) for the case of very small q (such as due to stealthy attacks or occasional errors) may not be sufficient, and as evidenced in the plots of Figs. 14–17, may necessitate the use of a higher value of n.

The second key observation is that choosing a threshold T_h that results in a smaller probability of event p_s results in a better sensor performance (in accordance with the results of Sect. 5.2). Indeed for a small value of p_s, it may be possible to obtain the desired $P_D - P_{FA}$

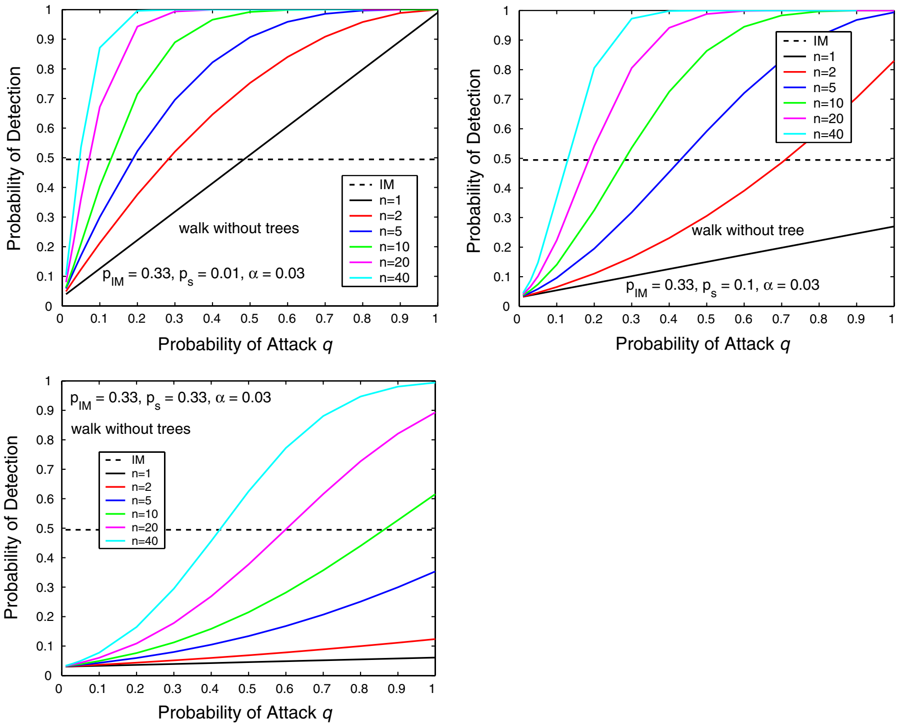

Fig. 15 P_D versus q for walking without trees from Fig. 7b. **a** $p_s = 0.01$. **b** $p_s = 0.1$. **c** $p_s = 0.33$

performance with a smaller number of sensors n. As can be seen by comparing parts (a), (b), and (c) of Figs. 14–17, choosing $p_s = 0.01$ (part a) offers a better performance than setting $p_s = 0.1$ (part b). Thus in general we do not wish to set $p_s = p_{IM}$ (part c) since p_{IM} may be arbitrarily large depending on the image sequence. Finally we observe that for certain values of p_s and n, the lightweight image processing algorithm (LIP) still achieves a better detection performance for certain values of q.

Based on these results we now wish to investigate the performance of a camera decision that is based on combining the CH detector with the LIP algorithm. Figure 18a shows the P_D versus q performance for the image sequence of an individual walking in the presence of trees (from Fig. 6) while Fig. 18b depicts this performance for the truncated walking sequence (from Fig. 7b). Figure 19a shows the performance for the image sequence of a vehicle in the presence of trees (from Fig. 5) while Fig. 19b depicts this performance for the truncated car sequence (from Fig. 7a).

In the figures, the dashed lines represent the performance based on the CH detector alone while the solid lines represent the performance of the combined decisions. As can be seen from these figures, the combined decisions achieve a better (or equal) detection performance P_D (for the same probability of α) than the CH decisions or the LIP decisions alone. This is consistent with the results obtained in Sect. 6.1 (where a single sensor decision was utilized without CH detection). However by utilizing the CH detector, we avoid the degradation of the detection performance for large q while by utilizing the LIP algorithm we avoid degradation of detection for small q. Thus based on the selection of a suitably small p_s and/or the selection of a suitably large cluster size n, we are able to adjust the $P_D - P_{FA}$ performance to suit the application requirements over the entire range of error q due to error or attack.

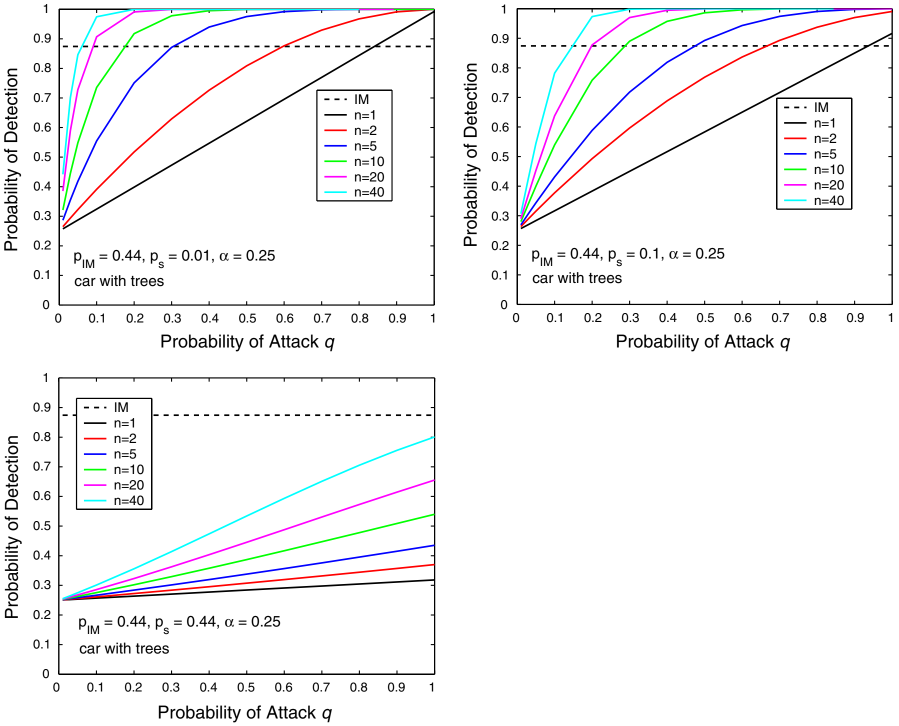

Fig. 16 P_D versus q for car with trees from Fig. 5. **a** $p_s = 0.01$. **b** $p_s = 0.1$. **c** $p_s = 0.44$

7 Summary and conclusions

Wireless Image Sensor Networks (WISNs) consisting of untethered camera nodes and sensors may be deployed in a variety of unattended and possibly hostile environments to obtain surveillance data. In such settings, the WISN nodes must perform reliable event acquisition to limit the energy, computation and delay drains associated with forwarding large volumes of image data wirelessly to a sink node.

In this work we investigated the event acquisition properties of WISNs that employ various techniques at the camera nodes to distinguish between event and non-event frames in uncertain environments that may include attacks. These techniques include lightweight image processing, decisions from *n* sensors with/without cluster head fault and attack detection, and a combination approach relying on both image processing and sensor decisions. In closing, we summarize the resulting properties and observations for event acquisition in WISNs:

1. Lightweight Image Processing (LIP) Approach in Uncertain Environments:

   a. LIP algorithms are generally compatible with low-power, low-complexity camera nodes (Rahimi et al. 2005). Indeed analysis demonstrates that through proper threshold selection (Sect. 4.3), simple LIP algorithms (such as based on the comparison of a single frame statistic to a threshold) can be made robust and achieve the best detection performance for a given worst-case probability of false alarm.

   b. The typical probabilities of detection and false alarm obtained in our experiments were consistent with the average probabilities reported in the literature with our $P_D = 0.87$ and our $P_{FA} = 0.26$.

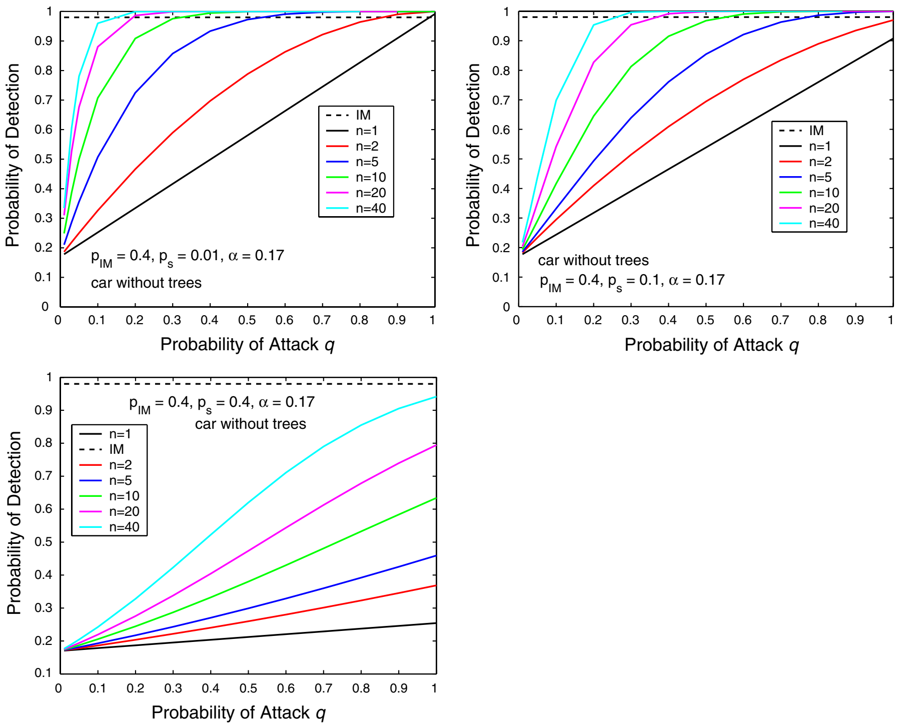

Fig. 17 P_D versus q for car without trees from Fig. 7a. **a** $p_s = 0.01$. **b** $p_s = 0.1$. **c** $p_s = 0.44$

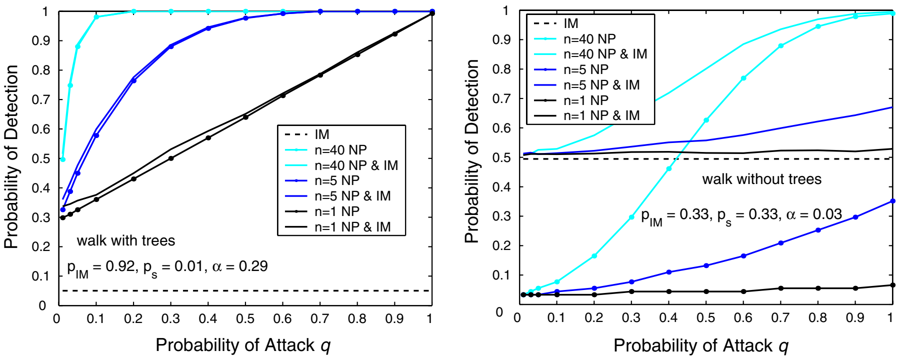

Fig. 18 **a** P_D versus q for walking sequence *with* tree from Fig. 6 for $p_s = 0.01$. **b** P_D versus q for walking sequence *without* tree from Fig. 7b for $p_s = 0.33$

c. While achieving an average $P_D - P_{FA}$ performance that may be suitable for some applications, the LIP algorithm exhibited large *variability* in its performance from sequence to sequence depending on environmental conditions. LIP algorithms alone thus may not offer the level of performance control and flexibility required in many applications.

2. Sensor Decisions Approach in Uncertain Environments:

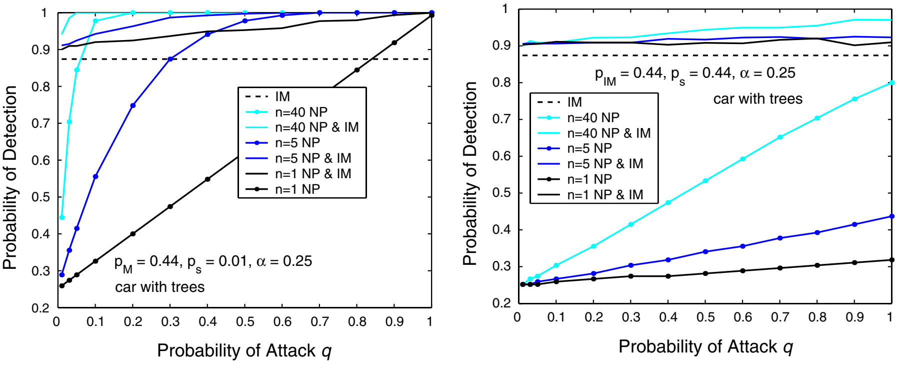

Fig. 19 P_D versus q for car sequence with tree from Fig. 5 for (**a**) $p_s = 0.01$ (**b**) $p_s = 0.44$

a. Sensors deployed in unattended outdoor environments may be prone to occasional faults or deliberate attacks that may be carried out in a stealthy manner (i.e. such as to avoid detection). Although quality testing may provide an estimate for the probability of a fault, an estimate for the probability of an attack may not be generally available a priori. Without verification mechanisms, the reliability of the sensor decisions may not be adequate for some applications.

b. An optimal Neyman-Pearson (NP) fault/attack detector based on the comparison of a single statistic to a threshold can be implemented at the cluster head to verify the sensor decisions. A detector where the comparison threshold is based on the average expected weight (i.e. the count or degree of aggregation) can also be implemented. This detector follows the performance of the NP detector closely, especially for small values of the probability of an event p.

c. For the case of stealthy attacks, use of a cluster head (CH) detector forces the attacker to select a smaller probability of attack q, especially for a larger cluster size n. This has the dual effect of rendering attacks more rare but also harder to detect despite the ability to predict the optimal attack parameter as given in the analysis of Sect. 5.2.

3. Combined LIP and Sensor Approach:

a. For the case of *no* cluster head attack/fault verification, combining a single sensor decision with a decision based on a LIP algorithm does provide an improved event acquisition performance. Specifically, for a fixed probability of false alarm, the combined decision achieves a probability of detection P_D higher than or equal to the P_D of LIP and SN over the entire range of the probability of sensor error q. Although better than LIP or SN alone, the detection performance does decrease with increasing q, which without cluster head detection, may be arbitrarily large depending on the attacker.

b. Combining decisions from n sensors with cluster head verification and LIP decisions provides the best overall performance over the entire range of attack probability q. Specifically, by utilizing the CH detector we avoid the degradation of the detection performance for large q while by utilizing the LIP algorithm we avoid degradation of detection for small q.

c. Choosing a sensor threshold T_h that results in a smaller probability of event p_s (depending on the underlying sensor technology) results in a better sensor performance and allows the use of a smaller cluster size n to achieve the desired $P_D - P_{FA}$ performance.

Importantly we note that in this work we have focused on the sensor-camera collaboration paradigm under different error and attack scenarios. Paradigms that additionally exploit spatial correlations among the cameras may further improve the overall event detection performance of the network. Although such paradigms may require additional setup information to determine the visual correlation among the camera nodes, a study of the trade-offs of such systems is of great interest for future work.

References

Aach, T., & Kaup, A. (1993). Statistical model-based change detection in moving video. *Signal Processing, 31*, 165–180.

Aach, T., & Kaup, A. (1995). Bayesian algorithms for adaptive change detection in image sequences using markov random fields. *Signal Processing: Image Communication, 1*(2), 147–160.

Akyildiz, I., & Kasimoglu, I. (2004). Wireless sensor and actor networks: Research challenges. *Ad Hoc Networks Journal, 2*(4), 3351–3677.

Akyildiz, I., Melodia, T., & Chowdhury, K. (2007). A survey on wireless multimedia sensor networks. *Computer Networks, 51*(4), 921–960.

Bandyopadhyay, S., & Coyle, E. (2003). An energy efficient hierarchical clustering algorithm for wireless sensor networks. In *INFOCOM 2003*, pp. 1713–1723.

Basharat, A., Catbas, N., & Shah, M. (2005). A framework for intelligent sensor network with video camera for structural health monitoring of bridges. In *Proceedings Third IEEE International Conference on Pervasive Computing and Communications Workshops, PerCom 2005 Workshops*, pp. 385–389.

Buttyan, L., & Hubaux, J. P. (2002). Report on a working session on security in wireless ad hoc networks. *Mobile Computing and Communications Review, 6*(4), 1–17.

Chan, H., Perrig, A., & Song, D. (2003). Random key predistribution schemes for sensor networks. In *IEEE Security and Privacy*, pp. 197–213.

Chow, K. Y., Lui, K. S., & Lam, E. (2006). Balancing image quality and energy consumption in visual sensor networks. In *IEEE International Symposium on Wireless Pervasive Computing* (p. 5). Phuket, Thailand.

Chow, K. Y., Lui, K. S., & Lam, E. (2007). Efficient on-demand image transmission in visual sensor networks. *Eurasip Journal on Advances in Signal Processing, V*, 11 pp.

Czarlinska, A., & Kundur, D. (2008). Reliable scalar-visual event-detection in wireless visual sensor networks. In *IEEE CCNC '08: Consumer Communications & Networking Conference* (pp. 660–664). Las Vegas, NV.

Czarlinska, A., Luh, W., & Kundur, D. (2007). Attacks on sensing in hostile wireless sensor-actuator environments. In *IEEE Globecom '07: Global Telecommunications Conference* (pp. 1001–1005). Washington, DC.

Eltoweissy, M., Moharrum, M., & Mukkamala, R. (2006). Dynamic key management in sensor networks. *IEEE Communications Magazine, 44*(4), 122–130.

Eltoweissy, M., Wadaa, A., Olariu, S., & Wilson, L. (2005). Scalable cryptographic key management in wireless sensor networks. *Journal of Ad Hoc Networks: Special Issue on Data Communications and Topology Control in Ad Hoc Networks, 7*, 796–802.

Feng, W. C., Walpole, J., Feng, W. C., & Pu, C. (2001). Moving towards massively scalable video-based sensor networks. In *Workshop on New Visions for Large-Scale Networks: Research and Applications* (p. 385). Washington, DC.

He, T., Krishnamurthy, S., Luo, L., Yan, T., Gu, L., Stoleru, R., Zhou, G., Cao, Q., Vicaire, P., Stankovic, J. A., & Abdelzaher, T. F. (2006). Vigilnet: An integrated sensor network system for energy-efficient surveillance. *ACM Transactions on Sensor Networks, 2*(1), 1–38.

He, Z., & Wu, D. (2006). Resource allocation and performance analysis of wireless video sensors. *IEEE Transactions on Circuits and Systems for Video Technology, 16*(5), 590–599.

Ma, H., & Liu, Y. (2005). Correlation based video processing in video sensor networks. In *IEEE International Conference on Wireless Networks, Communications and Mobile Computing* (p. 987). Maui, Hawaii.

Maniezzo, D., Yao, K., & Mazzini, G. (2002). Energetic trade-off between computing and communication resource in multimedia surveillance sensor network. In *4th IEEE Conference on Mobile and Wireless Communications Networks (MWCN)* (p. 373). Stockholm, Sweden.

Olariu, S., Eltoweissy, M., & Younis, M. (2007). ANSWER: Autonomous networked sensor system. *Journal of Parallel and Distributed Computing, 67*(1), 111–124.

Ott, R., & Longnecker, M. (2001). *An introduction to statistical methods & data analysis*. Duxbury Press.

Radke, R., Al-Kofahi, S. A. O., & Roysam, B. (2005). Image change detection algorithms: A systematic survey. *IEEE Transactions on Image Processing, 14*(3), 294–307.

Rahimi, M., Baer, R., Iroezi, O. I., Garcia, J. C., Warrior, J., Estrin, D., & Srivastava, M. (2005). Cyclops: In situ image sensing and interpretation in wireless sensor networks. In *ACM SenSys '05*, pp. 192–204.

Raymond, D., & Midkiff, S. (2008). Denial-of-service in wireless sensor networks: Attacks and defenses. *IEEE Pervasive Computing, 7*(1), 74–81.

Rodriguez, V. (2003). Resource management for scalably encoded information: The case of image transmission over wireless networks. In *IEEE Proceedings 2003 International Conference on Multimedia and Expo (ICME 2003)* (pp. I-813–816). Baltimore, MD.

Rosin, P. L. (2002). Thresholding for change detection. *Computer Vision and Image Understanding, 86*(2), 79–95.

Soro, S., & Heinzelman, W. (2005). On the coverage problem in video-based wireless sensor networks. *IEEE Broadband Networks, 2*, 932–939.

Van Trees, H. L. (2001). *Detection, estimation, and modulation theory part I.* John Wiley & Sons, Inc.

Veeraraghavan, K., Peng, D., & Sharif, H. (2005). Energy efficient multi-resolution visual surveillance on wireless sensor networks. In *IEEE International Conference on Electro Information Technology* (6 pp.). Lincoln, NE.

Wu, M., & Chen, C. (2007). Collaborative image coding and transmission over wireless sensor networks. *EURASIP Journal on Advances in Signal Processing, 2007*, 1–9.

Yu, C., Soro, S., Sharma, G., & Heinzelman, W. (2007). Lifetime-distortion trade-off in image sensor networks. *IEEE ICIP, V*, 129–132.

Author Biographies

Alexandra Czarlinska is a Ph.D. candidate in the Wireless Communications Group (WCL) in the Department of Electrical & Computer Engineering at Texas A&M University. She received her B.A.Sc. degree in Engineering Science (Electrical Option) in 2002 from the University of Toronto Canada where she was the recipient of the National Scholarship Award. Her current research focuses on the identification and prevention of new security attacks in mobile wireless sensor and actuator networks and on applications of game theory.

Deepa Kundur received the B.A.Sc., M.A.Sc., and Ph.D. degrees all in Electrical and Computer Engineering in 1993, 1995, and 1999, respectively, from the University of Toronto, Canada. In January 2003, she joined the Department of Electrical and Computer Engineering at Texas A&M University, where she is currently an Associate Professor and leads the SeMANTIC (Sensor Media Algorithms and Networking for Trusted Intelligent Computing) Research Group of the Wireless Communications Laboratory. Her research interests include security and privacy for scalar and broadband sensor networks, multimedia security, digital rights management, and steganalysis for computer forensics. She has given tutorials in the area of information security at ICME-2003 and Globecom-2003, and was a Guest Editor of the June 2004 Proceedings of the IEEE Special Issue on Enabling Security Technologies for Digital Rights Management. She currently serves as the Vice Chair for the Security Interest Group of the IEEE Multimedia Communications Technical Committee and is an Associate Editor for the IEEE Transactions on Multimedia, IEEE Communication Letters, and EURASIP Journal on Information Security.