Attack vs. Failure Detection in Event-Driven Wireless Visual Sensor Networks

Alexandra Czarlinska
Electrical & Computer Engineering
Texas A&M University
czlinska@ece.tamu.edu

Deepa Kundur
Electrical & Computer Engineering
Texas A&M University
deepa@ece.tamu.edu

ABSTRACT

In *event-driven* wireless Visual Sensor Networks (wVSNs), video nodes have access to additional data from scalar-sensors such as temperature or motion. The scalar-data may be used locally by the nodes instead of (or in conjunction with) vision technologies to control the potentially energy-costly transmission and storage of video frames and must thus be reliable. In this work we focus on the detection of *occasional* errors in such scalar-data sensors under both the scenario of *harsh* environmental conditions, and the scenario of *hostile* conditions involving an attacker. In the hostile case, the attack statistics may not be known to the cluster-head performing the error detection. We hence propose the use of a *count* detector in conjunction with Nash equilibrium analysis for the hostile case. We compare the detection performance of the count detector in hostile conditions to the performance of the optimal Neyman-Pearson (NP) detector which may be used under harsh conditions (scenario where the error statistics may be estimated). Through analysis and simulations we conclude that in this severe regime of attack with missing statistics, the count detector performs reasonably well compared with the optimal NP detector with significance for reliable event-driven wVSN.

Categories and Subject Descriptors

H.4 [**Information Systems Applications**]: Miscellaneous; C.4 [**Performance of Systems**]: Reliability, availability, and serviceability

General Terms

Reliability, Security

Keywords

Wireless Visual Sensor Networks, Event-Driven, Error Detection, Attack

MM&Sec'07, September 20–21, 2007, Dallas, Texas, USA.

1. INTRODUCTION

Ongoing improvements in miniaturized video technology, proof-of-concept testbeds and the lure of innovative applications have all continued to fuel interest in Visual Sensor Networks (VSNs) both within wired and wireless settings (wVSNs) [8]. Indeed the riveting technical challenges of such networks and the appeal of their possible applications have led some researchers to name wVSNs as the next evolutionary step in sensor and ad-hoc networks [6]. Possible applications of wVSNs include intelligent infrastructure monitoring, scientific and wildlife data collection, health and safety monitoring in child-care and retirement centers, and a variety of homeland-security related applications [1], [7].

To realize the potential of such wVSN applications, a variety of technical challenges must be addressed. Some of these challenges are particularly onerous in the case of a wireless setup involving multimedia, and a variety of approaches have recently been proposed in hopes of realizing the wVSN vision [6]. Any such effective approach must necessarily address the issues of limited energy, bandwidth and storage at the nodes, as well as the need for sophisticated image processing (context-based fusion, filtering, and/or aggregation of various video feeds) which may be taxing for certain large-scale wireless setups [8]. Furthermore, various trade-offs between computation and communication in multimedia networks and between power consumption in video encoding and wireless transmission must be understood [10], [7]. Ideally, a wVSN framework would also aim to minimize the time delay caused by the transmission of images and by their processing (for real-time applications, especially involving security monitoring) [1].

One possible approach is to develop wVSNs which are fully or partially *event-driven* [1]. According to this model, the wireless video nodes are equipped with additional scalar-valued sensors (i.e. sensors that collect temperature/pressure readings or detect motion). Such additional sensors could either be integrated directly into the nodes to form one unit as depicted in Figure 1a, or be deployed as separate devices that form their own network as in Figure 1b. Such augmentation of video data with scalar readings may be used in a variety of data fusion approaches [13]. In this scenario however the scalar readings are used by the nodes to determine whether recorded video frames should be transmitted wirelessly to a base station (or cluster head), stored for a period of time in memory, or be discarded if no event of interest occurred (they could also be used to control the nodes' sleep-mode). The motivation for such an approach is based on the observation that scalar-data might be easier to pro-

cess for the nodes in terms of computational time, battery energy and algorithm complexity.

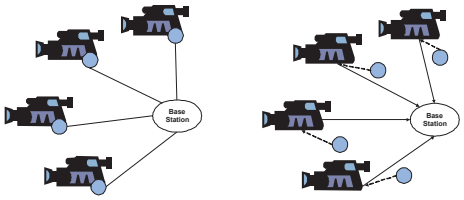

Figure 1: Even-driven wVSN setup (a) Nodes equipped with cameras and sensor(s) (b) Sensors form a separate support network for the cameras.

A variety of complimentary approaches to *even-driven* wVSNs have also been proposed, offering interesting trade-offs. One school of thought is to focus on energy-efficient compression and/or power trade-offs between encoding and transmission at the nodes [7], [10], [11]. The advantage of these approaches is that *all* the video/image feeds are transmitted to the base station (though possibly compressed in a lossy way) and are hence available for later audits (such as following a security incident when previous, seemingly innocent footage must be reviewed). Another complimentary proposal to facilitate the wireless transmission of video data is the use of free-space optical transmission, which consumes less energy than traditional radio transmission and which offers very high bandwidth [8].

These approaches alone however may not be fully optimized for low security, extended-duration applications where events of interest may occur sparsely in time (for example, wildlife monitoring where animals may be out of camera range for long periods of time) [2]. In particular, these approaches do not exploit the possibility of eliminating irrelevant or non-interesting frames at the source (i.e. at the nodes). Such selection of frames may be possible using context-based information processing where nodes employ basic vision technologies to discern motion in the video or to pick out relevant objects [13]. One challenge of this approach is to tailor the vision technology to the low-energy, low-computation, low-storage paradigm of wireless nodes. In comparison, the event-driven wVSN approach attempts to achieve the objective of selective recording/transmission based on relevant scalar data. In practice, the event-driven approach may ultimately be used in conjunction with other techniques. In this work however we focus on the reliability of the event-driven approach depicted in Figure 1.

A key element in unlocking the potential of the event-driven approach is to understand, and ultimately control, the probability of erroneous event-reporting by the scalar-sensors (especially if the event-driven approach is used without vision-based information processing). Such erroneous readings in the scalar-sensors might result in the dismissal of a relevant image, or conversely, in the transmission of an irrelevant image, hence wasting limited energy and bandwidth. Scalar-sensors deployed in harsh or hostile environments may be prone to *occasional* Byzantine-like errors. Unlike permanent defects, such occasional errors are much harder to detect during network operation. For instance, *harsh* environments of interest might include infrastructure in seismically active areas, or wildlife monitoring in remote rainforests [1], [2]. In *hostile* environments on the other hand, faulty sensor reporting may be caused by a malicious opponent engaged in false-packet injection, actuation of sensed-data or other sensor network attacks [4]. Some critical applications such as battle-field surveillance may experience both harsh and hostile conditions.

Importantly, occasional errors caused by harsh or hostile elements may occur with vastly different frequency than wireless transmission errors [4]. Whereas coded wireless transmission and decoding might produce a bit error with probability between 10^{-6} to 10^{-9}, the erroneous reporting rate of a particular scalar-sensor is not as easy to ascertain. Depending on the conditions of the harsh or hostile environment, we might expect a probability of sensor error that changes largely over time and over a wide range of values [5]. Furthermore, whereas we may be able to estimate the probability of sensor error due to harsh conditions, such estimation may not be possible in the case of an attacker who changes the attack statistics to avoid detection [5]. Hence whether deployed in normal, harsh or hostile conditions, it is imperative to prevent and detect errors in scalar-reporting in order to achieve a reliable event-driven wVSN operation.

1.1 Focus and Contribution of Paper

In this paper we examine the probability of detecting occasional errors in the scalar-sensors that support event-driven wireless Visual Sensor Networks (wVSNs). Our specific focus is on comparing the detection performance achieved under a hostile scenario (i.e. attack) versus the detection performance achieved under a harsh-environment scenario. The key detection-related difference between these two cases is that in a hostile scenario, we assume that information about the probability distribution function (PDF) of the attacker is not available to the detector (e.g. the attacker may keep changing his attack statistics to avoid detection). In this case we propose the use of a *count* (or *type*) detector at the cluster-head (or base station). We employ the game theoretic concept of a Nash equilibrium to determine the best strategy of the attacker and it's resulting impact on detection when the *count* detector is employed by the cluster-head. We perform simulation experiments to determine the resulting probability of detecting the attack and compare it to the harsh-environment case where we assume the error PDF information is available (such as through estimation).

2. EXISTING WORK

Given the breadth of work relevant to wVSNs and given our focus on sensor error detection in the harsh/hostile setting, we delimit our overview to relevant wVSN frameworks and implementations, as well as to pertinent detection concepts.

Feng *et al.* present a vision for massively scalable video-based sensor networks (VBSN) [6]. The authors identify and overview key technical issues that must be addressed in order to accommodate video feeds from possibly thousands of nodes. In their vision, wireless battery-run video nodes pass their data feeds directly to a higher-powered "base station" (or cluster-head) where the data aggregation and filtering first occur. Alternatively, the nodes are equipped with context-based vision technologies (i.e. using objects in the video data) to track objects or to filter information at the nodes. Basharat *et al.* propose and implement an event-driven framework for using wireless sensor nodes with wired video cameras to monitor the structural health of bridges [1]. The approach is aimed at eliminating vast amounts of

unnecessary data intelligently, at prolonging the system's life and rendering the system robust to errors and flexible to deploy. Scalar-data nodes collect temperature and vibration readings on the bridge at different sampling rates. In the system's Passive Mode, only a few select nodes do the sensing. Elevated readings however trigger the remaining scalar nodes to enter Active Mode to increase the spatial-coverage and detail of the sensing. If abnormal readings are detected while in Active Mode, the wired cameras begin their operation, otherwise the nodes return to Passive Mode.

Liu and Sayeed examine the use of a general type of detector called the *type* detector, with a special focus on wireless sensor networks [9]. In their setup, each sensor node collects a sequence of observations of length n. Based on the *Method of Types*, this approach does not require the nodes to use PMF (or PDF) information to select a hypothesis (though this information is known at the Base Station). For a discrete alphabet of observations (for example binary-valued observations), the nodes act as simple *counters*, recording the relative frequencies with which each symbol (1 or 0 for binary alphabets) occurred in the sequence. After determining the count of each symbol, a node transmits this information to the Base Station which then performs a joint "type" detection to select a final hypothesis about the data. It is shown that for the case of binary *i.i.d* observations, the method approaches the performance of optimal centralized detection, while reducing the communication of the nodes (as well as possibly the storage and computation required by the nodes).

3. PRELIMINARIES

3.1 Notation

Throughout this paper we use the notation P_D to denote the probability of detection and P_{FA} to denote the probability of false alarm. The notation $P_D = \beta$ and $P_{FA} = \alpha$ where $\alpha \in [0, 1]$ and $\beta \in [0, 1]$ is also used to denote specific values of P_D and P_{FA}. The notation P_M is used to denote the probability of miss, where $P_M = 1 - P_D$. The probability mass function (PMF) of a random variable X is denoted by $p_X(x)$ or $p(x)$ for short. For a binary-valued vector $\mathbf{x}$ of length m, we use the notation $w(\mathbf{x})$ or w for short, where $w\{0, \ldots, m\}$, to denote the *weight* of the vector (the number of 1s contained in the vector). The notation $\mathbb{I}_P$ denotes the *indicator* function which is equal to 1 if the proposition P is true and is equal to 0 otherwise.

3.2 Neyman-Pearson and Count Detectors

In detection problems we are generally faced with the task of deciding between two or more hypothesis 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 P_D and P_{FA} may not be of equal significance (otherwise a Bayesian detector may be appropriate).

We consider the case where the data vector $\mathbf{z}$ consists of n *i.i.d* Bernoulli random variables coming either from a PMF $Bern(p)$ (hypothesis $\mathcal{H}_0$) or a PMF $Bern(r)$ (hypothesis $\mathcal{H}_1$). For this case it can be shown that the NP detector is given by Eq. 1, where $w(\mathbf{z})$ is the weight (the number of 1s) in the data vector $\mathbf{z}$.

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

The threshold $\mathcal{T}$ is chosen to satisfy a desired α based on Eq. 2, 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. 3, where $\Lambda(w) = r^w(1-r)^{n-w} \div p^w(1-p)^{n-w}$. Finally, the probability of detection β resulting from the use of the $\Lambda(\mathbf{z})$ detector is given by Eq. 4.

$$\sum_{\mathbf{z}:\Lambda(\mathbf{z})>T} p^{w(\mathbf{z})}(1-p)^{n-w(\mathbf{z})} \leq \alpha \tag{2}$$

$$\sum_{w=0}^{n} \binom{n}{w} p^w (1-p)^{n-w} \; \mathbb{I}_{\,\Lambda(w)>T} \; \leq \alpha \tag{3}$$

$$\beta = \sum_{w=0}^{n} \binom{n}{w} p^w (1-p)^{n-w} \; \mathbb{I}_{\,\Lambda(w)>T} \tag{4}$$

In cases where the PMF under different hypotheses is not available at the *sensors* (only at the cluster-head), a *count* or *type* detector may be appropriate as described in Section 2 pertaining to reference [9]. For cases where PMF information is missing altogether, such as in the case of a changing attacker, game theoretic analysis may yield insights into best-response PMFs.

3.3 Game Theory Nash Equilibria

Game Theory (GT) is generally concerned with "decision-making" (optimization of one's choice) in the case of multiple players whose choices affect the outcomes achievable by other players. In the basic "strategic" form, the game theoretic concept of a Nash equilibrium allows a player to determine which action she should select to achieve her best possible outcome *given* that she knows all the players' *preferences* for all the outcomes (but *not* which actions they *actually* choose in that instance of play). The concept of Nash equilibrium proves useful in determining the best actions of an attacker (worst-case scenario for the wVSN) who must select a distribution $Bern(q)$ with which to carry out the attack. Given such information, the attack detection performance may be better understood.

4. PROBLEM FORMULATION

In this paper we consider an *event-driven* wireless Visual Sensor network (wVSN) and focus on the detection of *occasional* scalar-sensor errors caused by either harsh or hostile conditions. Specifically, we consider n wireless battery-run cameras, *each* of which is supported by a scalar-data sensor as depicted in Figure 1 (the type of scalar-data sensor is application-dependent). We adopt the configuration shown in Figure 1a (a study of the benefits and challenges of 1b is beyond the current scope, as is the inclusion of vision-processing).

The model of event-detection at the scalar-sensors is largely sensor and application-dependent (i.e. the detection model of underground optical pressure sensors would differ from a model of temperature-collecting sensors). For purposes of generality and tractability we therefore choose to model the scalar-sensor event-detection via a sensor's ultimate output decision, 'event present' (bit 1) or 'event absent' (bit 0). Specifically, the detection of an event of interest in the environment at node i is modeled as a random variable X_i according to the Bernoulli distribution $Bern(p)$ of Eq. 5.

Hence a realization $x_i = 1$ denotes a scalar-sensor i having detected an event of interest.

$$X_i = \begin{cases} 1 & \text{w.p. } p \\ 0 & \text{w.p. } 1-p \end{cases} \quad (5)$$

Each scalar-sensor transmits its binary data to the cluster-head (or base station), reporting whether an event-of-interest occurred at the location of the node. This binary data may be appended to packets (such as in the header) that are already being sent to the cluster-head for other purposes. Due to delays it is important for these packets to be synchronized (or ordered in time) through appropriate techniques [12]. Hence we assume that after some time the cluster-head receives a data vector $\mathbf{z}$ of length n which corresponds to the n sensors' decisions for a specific observation period. This data vector is used by the cluster-head for fault and attack detection. Feedback from the cluster-head regarding the status of the scalar-data (valid or faulty/attacked) is used by a node to control whether a given frame in its short-term storage should be transmitted to the cluster-head or discarded. For scope reasons, in this work we do not elaborate on the local frame control and delay aspects but focus on the detection task at the cluster-head.

An error at a scalar-sensor i due to a fault (harsh conditions) or an attack (hostile conditions) may be modeled in a variety of ways. For tractability and generality we model the error as a random variable Y_i distributed according to a Bernoulli distribution $Bern(q)$. The effect of the error is modeled as changing the reading of a scalar-sensor, from declaring a 1 (event occurred) to declaring a 0, or vice-versa. This bit-flipping affects both the local frame control and the binary information z_i sent by each node to the cluster-head. The overall effect of the error on each bit *sent* to the cluster-head is given by Eq. 6, where X_i is the true event in nature, Y_i is the sensor error and Z_i is the resulting bit reported to the cluster-head. It can be shown that the distribution of the random variable Z_i is given by $Bern(r)$ where r is given by Eq. 7.

$$Z_i = X_i \oplus Y_i \quad (6)$$

$$r = (1-p)q + (1-q)p = q + p - 2pq \quad (7)$$

Hence at any given time, the cluster-head receives a data-vector $\mathbf{z}$ of length n that it uses to ascertain whether the network is functioning correctly. The cluster-head must decide between two alternative hypothesis:

$$\begin{aligned} H_0 &: \text{ normal operation, } PMF \sim Bern(p) \\ H_1 &: \text{ erroneous operation, } PMF \sim Bern(r) \end{aligned}$$

In the case of a harsh environment, prior sensor quality testing may yield a reasonable estimate for the probability of sensor error q. *If* additionally a similar probability estimate can be obtained for the occurrence of an event of interest (which realistically may not be available), then the probability p may also be obtained. In such a case, the optimal NP binary-data detector described in Section 3.2 may be used to distinguish between the hypotheses H_0 and H_1.

In the case of hostile conditions however, it may not be possible to obtain an estimate for q given that it is controlled by an attacker who is free to change his attack statistics over time. Indeed the reason why an attacker may employ a random attack in lieu of a deterministic one is precisely to thwart attempts by the wVSN to estimate its attack patterns. Assuming that an average count np is available (the average number of sensors that report a bit 1 based on the weak law of large numbers for n large), we propose the use of a modified *count* detector described in Section 2 and reference [9]. The simple modification is based on the observation that in practice the sequence of observations collected by a single node over *time* is most likely not truly *i.i.d* (i.e. p may not be *i.i.d* over time). Furthermore, we wish to detect an attack or malfunction early (otherwise video frames may be mishandled by the nodes) and hence do not wish to collect a relatively long sequence of observations. Instead we assume that in a *given* time interval, p is *i.i.d spatially* within a cluster of nodes and q (which is unknown) is also *i.i.d* spatially within this cluster. Hence we propose the use of the detector $D(\mathbf{z})$ given in Eq. 8, where $c = np$ is the average number of 1s the cluster-head expects to receive from the n sensors (for n large, by the weak law of large numbers), $w(\mathbf{z})$ is the *actual* count (number of 1s) received in the data vector $\mathbf{z}$, and ϵ is a variance-related "slack-factor" allowing the cluster-head to relax or tighten the detection constraint.

$$D(\mathbf{z}) = |w(\mathbf{z}) - c| \underset{H_0}{\overset{H_1}{\gtrless}} \epsilon \quad (8)$$

5. ANALYSIS

In the rest of this paper we wish to understand 1 - how an attacker should choose his attack parameter q if he knew (but did not control) p, such that he maximizes the probability of evading the detection that is performed by the cluster-head using Eq. 8, and 2 - how the detection performance in the hostile case (using Eq. 8) compares to the harsh conditions case (using Eqs. 1, 3 and 4).

5.1 Performance in Hostile Environment

We consider the case where the attacker's goal is to cause occasional errors in the scalar-sensors while minimizing the chance of getting detected by the cluster-head (detection would prevent the attacker from continuing to misguide the network). An alternative goal for the attacker may be to sacrifice some of the detection-evasion if it allows for more nodes to be attacked. This interesting alternative scenario is the study of ongoing work and is not considered in this paper. Based on the $D(\mathbf{z})$ detector and the outlined goal, the attacker wants to maximize the condition given by Eq. 9, where we have used basic combinatorics as in [5] and where $w(\mathbf{Z}) = w(\mathbf{X} \oplus \mathbf{Y})$ is based on Eq. 6. The binomial coefficients a, b and c are defined in Eq. 10.

$$\begin{aligned} & Pr\{|w(\mathbf{X} \oplus \mathbf{Y}) - c)| < \epsilon\} \\ = & \sum_{m=\lceil \frac{l-\epsilon}{2} \rceil}^{\lfloor \frac{l+\epsilon}{2} \rfloor} \sum_{k=1}^{n} \sum_{l=1}^{n} a(k,m)\, b(k,l,m)\, c(k) \\ & \cdot p^k (1-p)^{n-k} q^l (1-q)^{n-l} \end{aligned} \quad (9)$$

$$\begin{aligned} a(k,m) &= \begin{cases} \binom{k}{m} & \text{if } k \geq m \\ 0 & \text{o.w} \end{cases} \\ b(k,l,m) &= \begin{cases} \binom{n-k}{l-m} & \text{if } n-k \geq l-m \\ 0 & \text{o.w} \end{cases} \\ c(k) &= \begin{cases} \binom{n}{k} & \text{if } n \geq k \\ 0 & \text{o.w} \end{cases} \end{aligned} \quad (10)$$

It can be shown that the condition in Eq. 9 is concave in p with peak at $p = \frac{1}{2}$ and semi-concave in q in the asymptotic

case of large n. Recalling that $p \in [0,1]$ and $q \in [0,1]$, the Nash equilibria of Eq. 9 in the asymptotic n case can be shown to exist for $(p,q) = (\delta_p, \delta_q)$ and $(p,q) = (1-\delta_p, \delta_q)$ where δ_p and δ_q are small numbers approaching 0. From this result we conclude that in order to minimize the probability of being detected, the attacker should choose a small value of q (i.e. δ_q). We also notice that the detector should perform best for either small p (i.e. δ_p) or for large p (i.e. $1-\delta_p$). That is, the count detector $D(\mathbf{z})$ should perform best when the event of interest is either rare (small probability of occurrence) or very common. This result agrees with the intuition that it is easier for an attacker to fool the cluster-head if the latter is expecting a bit of value 1 or 0 with equal probability (i.e. $p = \frac{1}{2}$). Figure 2a shows a plot of Eq. 9 over the range of p and q with $\epsilon = 0$ and $n = 50$ nodes.

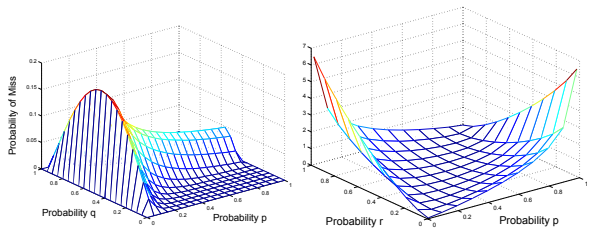

Figure 2: (a) P_M for various p and q, $\epsilon = 0$, $n = 50$. (b) $D(p \parallel r)$ over the range of p and q.

5.2 Performance in Harsh Environment

Generally we expect that the detection performance in the harsh conditions case (using the NP detector of Eqs. 1, 3 and 4) will be superior to the hostile conditions case where information about probability q is missing (using the count detector of Eq. 8). We wish to understand however if and how the difference in performance between the two cases varies with the values of p and r where r is given by Eq. 7.

In the case of *i.i.d* random variables used for a hypothesis test between two alternative PMFs p and r (such as in our case), we may use the Kullback Leibler (KL) distance $D(p \parallel r)$ viz. Stein's Lemma [3]. According to Stein's Lemma, the larger the KL distance between the distributions, the better the detection performance. It must be noted that the KL distance is not a true measure and is an asymptotic distance (for a large number of *i.i.d* samples n). The KL distance between two Bernoulli PMFs p and r is given by Eq. 11.

$$D(p \parallel r) = p \log_2 \frac{p}{r} + (1-p) \log_2 \frac{1-p}{1-r} \qquad (11)$$

Table 1 shows $D(p \parallel r)$ for various values of p and q, arranged in order of an expected (relative) *decrease* in performance. For instance, we *expect* that detection in the $(p,q) = (0.1, 0.99)$ case will be better than detection for the $(p,q) = (0.1, 0.1)$ case since the KL distance in the former case is larger. Figure 2 b) shows a plot of $D(p \parallel r)$ over the entire range of p and r.

6. RESULTS AND DISCUSSION

We wish to determine and compare the detection performance of the NP detector of Eqs. 1, 3 and 4 and of the count detector of Eq. 8. Based on the analysis of Section 5.1 for the hostile environment, we expect the count detector to perform worst in cases where the attacker chooses a small value of q, and/or when the event of interest occurs with probability p close to $\frac{1}{2}$. On the other hand, in the case of a harsh environment as examined in Section 5.2, the NP detector need not perform poorly for small values of q as long as the KL distance between p and r is relatively large.

Table 1: KL-Distance

Value of p	Value of q	$D(p \parallel r)$
0.1	0.99	2.44
0.1	0.1	0.0361
0.6	0.3	0.0105
0.47	0.47	0.0023
0.5	0.5	0

Figures 3 and 4 show the probability of detection P_D versus the probability of false alarm P_{FA} for the NP detector corresponding to the values of p and r shown in Table 1. As expected, the performance of the NP detector is best for cases where $D(p \parallel r)$ is largest and the relative performance follows that of Table 1. For instance, in Figure 3a where $D(p \parallel r)$ is largest, we see excellent performance where a high P_D is achieved for all P_{FA}. In contrast, in Figure 4a where the $D(p \parallel r) = 0$, we see the worst possible (P_D, P_{FA}) performance.

Superimposed on these plots, is the P_D versus P_{FA} for the corresponding count detector. As expected, its performance is lower than that of the optimal NP detector. Interestingly, altering the value of ϵ, which may be thought of as a slack factor, has the effect of moving the (P_D, P_{FA}) performance of the count detector along the $P_D - P_{FA}$ curve of the NP detector. For instance, setting $\epsilon = 0$ increases P_D in the count detector close to its maximum value of 1, but it also produces an undesirable P_{FA} close to its maximal value of 1. For certain cases, selecting a value of ϵ equal to 2 or 3 produces a more desirable (P_D, P_{FA}) pair.

Finally, we examine the effect of cluster-size on detection performance. It is known that for any optimal NP (or Bayesian) detector, the performance generally increases with an increase in the number of samples. This implies that a larger number of nodes n reporting their scalar-data readings to a cluster-head should perform better for the *harsh* conditions case. We obtain a similar result for the *hostile* conditions case. Figure 4b shows a plot of P_M (obtained from Eq. 9) for $\epsilon = 0$, $p = 0.1$ and various q and n. The plot shows that the probability of attack success (corresponding to P_M) decreases as the size of the cluster n increases.

The implication of these results for reliable event-driven wVSNs is that detection of occasional errors due to hostile attacks is possible even though the attack parameter may not be known. Harsh conditions generally yield better detection performance than hostile conditions. If the resulting reliability of the scalar-data is suitable for a given wVSN application, feedback from the cluster-head regarding the scalar-data may be used to perform local filtering and selection of video frames in lieu of more energy-consuming image-based processing. For general wVSN applications with differing security requirements, a combination of techniques involving efficient image-processing, compression and scalar-data assistance will likely yield the best overall trade-offs. A study of such trade-offs under hostile and harsh conditions is the subject of ongoing study.

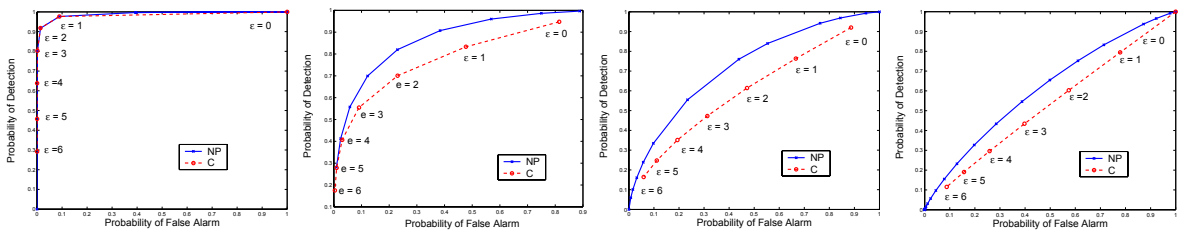

Figure 3: NP **vs.** C **detectors for** $n = 50$**, (a)** $p = 0.99$, $q = 0.1$**. (b)** $p = 0.1$, $q = 0.1$**. (c)** $p = 0.6$, $q = 0.3$**. (d)** $p = 0.47$, $q = 0.47$**. Vertical Axis:** P_D**, Horizontal Axis:** P_{FA}**. The** C **detector is varied over a range of** ϵ**.**

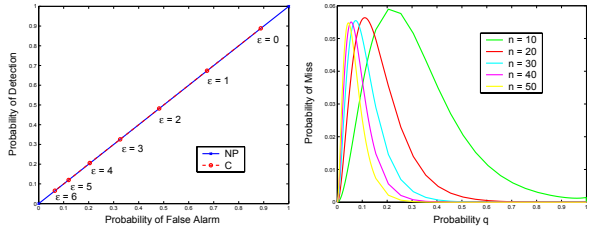

Figure 4: (a) NP **vs.** C **for** $n = 50$**,** $p = 0.50$**,** $q = 0.50$ **and various** ϵ**. (b)** P_M **for** $p = 0.1$**, various** q **and** n**.**

7. CONCLUSIONS

In this work we examine event-driven wireless Visual Sensor Networks (wVSNs) with a focus on the detection of occasional errors in the scalar-data sensors under harsh vs. hostile conditions. The significant difference between the harsh and hostile detection case is that in the latter, PMF information required for detection may not be known or estimated. To address this unknown, we propose a modified count detector for use at the cluster-head of a wVSN. We compare the detection performance of the count detector against the optimal Neyman-Pearson (NP) detector which may be used in the harsh conditions case. We determine that though it does not perform as optimally as the NP, for many cases the count detector performs reasonably well, with its P_D-P_{FA} adjusted through selection of the ϵ parameter.

ACKNOWLEDGMENTS The authors would like to thank the anonymous reviewers for their insightful suggestions.

8. REFERENCES

[1] A. Basharat, N. Catbas, and M. Shah. 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*, page 385, March 2005.

[2] D. Biello. *Camera Captures Image of Mysterious Creature in Borneo*. December 7 2005.

[3] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. John Wiley & Sons Inc., 1991.

[4] A. Czarlinska and D. Kundur. Distributed actuation attacks in wireless sensor networks: Implications and countermeasures. *Second IEEE Workshop on Dependability and Security in Sensor Networks and Systems*, pages 3–12, April 2006.

[5] A. Czarlinska, W. Luh, and D. Kundur. Attacks on sensing in hostile wireless sensor-actuator environments. In *IEEE Globecom*, Washington, D.C., November 26–30.

[6] W.-C. Feng, J. Walpole, W.-C. Feng, and C. Pu. Moving towards massively scalable video-based sensor networks. In *Workshop on New Visions for Large-Scale Networks: Research and Applications*, page 385, Washington, DC, March 2001.

[7] Z. He and D. Wu. Resource allocation and performance analysis of wireless video sensors. *IEEE Transactions on Circuits and Systems for Video Technology*, 16(5):590–9, May 2006.

[8] D. Kundur, W. Luh, and U. N. Okorafor. Emerging security paradigms for distributed multimedia sensor networks. *Proceedings of the IEEE Special Issue on Distributed Multimedia*, to appear 2007.

[9] K. Liu and A. Sayeed. Asymptotically optimal decentralized type-based detection in wireless sensor networks. In *IEEE International Conference on Acoustics, Speech, and Signal Processing*, page 873, Montreal, May 2004.

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

[11] S. Pradhan and K. Ramchandran. Distributed source coding: Symmetric rates and applications to sensor networks. In *Proc. DCC'00*, Snowbird, UT, March 2000.

[12] F. Sivrikaya and B. Yener. Time synchronization in sensor networks: A survey. *IEEE Transactions on Circuits and Systems for Video Technology*, 18(4):45–50, July/August 2004.

[13] N. Su, H. Park, E. Bostrom, J. Burke, M. Srivastava, and D. Estrin. Augmenting film and video footage with sensor data. In *2nd IEEE Annual Conference on Pervasive Computing and Communications*, pages 3–12, Orlando, Florida, March 2004.