

A framework for QoI-inspired analysis for sensor network deployment planning*

Sadaf Zahedi[†]
Univ. of California, Los Angeles
Electrical Engineering Department
Los Angeles, CA 90095-1594, USA
szahedi@ee.ucla.edu

Chatschik Bisdikian
IBM Corporation
Thomas J. Watson Research Center
Hawthorne, NY 10532, USA
bisdik@us.ibm.com

ABSTRACT

The quality of information (QoI) that sensor networks provide to the applications they support is an important design goal for their deployment and use. In this paper, we introduce a layered framework for QoI-centered evaluation of sensor network deployment. The layered framework allows decomposing the deployment evaluation in three steps: input pre-processing, core analysis, and result post-processing. The layering allows the creation of a rich, modular toolkit for QoI-centered analysis that can accommodate both existing and new system modeling and analysis techniques. We demonstrate the utility of the framework by comparing the QoI performance of finite-sized sensor networks with general deployment topology. We also derive some new analysis results for the class of applications considered herein.

1. INTRODUCTION

With advances in computing and communication technologies, low(er) cost, intelligent networked sensor systems find their way in a multitude of application environments in areas as diverse as the military intelligence gathering, habitant monitoring, forest monitoring, utility grid monitoring, environmental control, machinery control, and so on.

For the past several years, considerable amount of research work has been conducted regarding the internal operation

*Research was sponsored in part by the U.S. Army Research Laboratory and the U.K. Ministry of Defence and was accomplished under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the author(s) and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

[†]Ms. Zahedi performed this work during her summer 2007 internship at IBM Research.

of a sensor network including ad hoc deployment and operation of sensor networks, energy-aware architectures and protocols, coverage and localization, efficient query dissemination, etc. While the value and necessity of these studies is unquestionable, since sensor networks are deployed for a purpose –to serve an application–, we find it appropriate that additional studies are also needed that bridge the behavior of the sensor networks with the application(s) they support.

We have elected to use the concept of *quality of (sensor) information* (QoI) as a means to capture an application’s information needs from the sensors. Research on QoI, also referred to as IQ (for *information quality*, of course) as well, has its roots in the study of the condition of structured information stored in database systems with respect to its consistency, completeness, currency, etc. [5]. Within the context of sensor networks, a comprehensive definition of QoI is still an inconclusive endeavor, but for the purpose of this paper, we have adopted the following definition:¹

DEFINITION 1. *QoI is the collective effect of the available knowledge regarding sensor-derived information that determines the degree of accuracy and confidence by which those aspects of the real world (that are of interest to the user of the information) can be represented by this information.*

It follows from the above definition that QoI is a multi-dimensional concept that is applicable (at least some aspects of it) to a user, i.e., an application. The multi-dimensionality of QoI pertains to attributes that “codify” and possibly quantify the knowledge about the information. For example, knowledge about the *timeliness* of the available information, its *reliability*, its (arithmetic) *precision*, and so on, may be used to describe how accurately we can represent the real world using the sensor-derived information [1, 2, 3, 5].

Which information attributes are more important and/or what is the range of acceptable values for them, will depend on application(s) and their needs as specified by an application planner. Given the broad application space for sensor networks, we have elected to use for our research work a general class of applications that are part of many decision making operations. This class of applications relates to event detection that find applications in surveillance and intelligence gathering operations like detecting presence of

¹We have actually based our wording of the definition on a paraphrase of ITU’s Rec. E.800 definition of quality of service provided by a telecommunications services provider.

enemy weaponry, hostile activities (e.g., gunfire, explosions), monitoring remote territories, and so on.

Upon detection of an event of interest, an operative can take an action to mitigate the effects of the event. The effectiveness and the severity of the action, (or the inaction, if the event is not detected) will depend on the QoI that is provided to the operative. Thus, for the class of event detection applications, we elect as QoI attributes of importance the *detection probability* P_d of correctly detecting the occurrence of the event and the *false alarm rate* P_f , i.e., the probability of declaring an event occurrence when no event really occurred.

With the above as the motivating background, in this paper, we introduce a QoI analysis framework instantiated through a toolkit system. The toolkit (and the analysis framework it represents) serves as a computational aid for a sensor systems designer to evaluate the performance of his/her design based on deployment and QoI constraints provided by the application planner. Considerable amount of work has been performed analyzing detection systems for different system models such as signal-to-noise ratio (SNR), channel fading, spatial correlation, and so on [13, 4, 11, 6, 7]. In [11] a homogenous system is considered that results in identical SNR levels at the various sensors nodes, while in [6] non-identical SNR levels are considered as a result of the spatial distribution of the nodes and channel fading. In [7] a decision fusion algorithm is also studied and an approximation is derived for the system-level decision threshold that provides performance guarantees for the system under fixed level, persistent events, as most studies assume. Detection systems have also studied with relation to energy requirements as well. For example, in [10] and [9] hybrid (neither centralized nor distributed) energy-driven detection schemes are proposed based on a binary observation model for the sensors that provides the flexibility to trade-off detection accuracy and energy consumption. We do not considered energy efficiency in this paper, which is left for future study.

We have recognized that past research in the area comprises of specific performance analysis methodologies that are applicable to the specific system models studies. Contrary, in this paper, and motivated by our previous work [1, 2], we take a general, system-level approach to detection evaluation that considers QoI analysis techniques for a class of centralized, distributed, as well as hybrid detection architectures in a unified way. Furthermore, we also consider general events whose (signal) signatures with arbitrary support and amplitude, with special emphasis in *transient* and *decaying* events. Studying general classes of events, brings a new dimension to research in the area. This is because samples taken at different times by the various sensors, either due to their geographical distribution or the use of different sampling policies, may result in a sensor measuring different signal amplitude levels over time which may significantly impact the performance of the detection systems.

The organization of the paper is as follows: In section 2, we introduce the reference model for the system under consideration and the general toolkit framework. In section 3, we present the core analysis approach based on hypothesis testing. In section 4, we include example uses of the framework




Figure 1: Functional architecture of the reference detection system.

while, at the same time, deriving some new results for the analysis of networks with finite number of sensor and transient events. We conclude in section 5 with some concluding remarks.

2. THE REFERENCE DETECTION MODEL AND QOI ANALYSIS FRAMEWORK

We start this section by introducing a reference architecture of our sensor-enabled, detection system and then we introduce the QoI analysis toolkit that is built around the reference model.

2.1 The reference detection system

Figure 1 shows the reference architecture of our sensing system. It comprises three functional subsystems: (a) the *sensor subsystem* or *sampler*; (b) the *fusion subsystem*; and (c) the *detection subsystem*. The sensor subsystem comprises M sensors that sample the physical world and provide their samples to the fusion subsystem. The fusion subsystem, comprising a collection of L fusion centers, operates on the samples it receives (which could be corrupted by noise) to produce a “summary” description of the samples. The summary, in turn, is used by the detection subsystem to decide whether an event of interest has occurred or not. These three subsystems may be collocated or separated as could be the case of a networked sensing system.

According to the reference model, the sensor subsystem takes samples $\{s_1, s_2, \dots\}$ of the event signature (if the event occurred) as it is experienced locally by the sensors (the index k in the figure identifies “things” related to the k -th sensor). The sampled signature is distorted by noise and the combination of the two (or just the noise component, when the event did not occurred) produce recordable observations $\{r_1, r_2, \dots\}$. These observations are then processed by the fusion subsystem to generate the sample summaries.

If the event occurred, the sequence $\{s_1, s_2, \dots\}$ can be thought of as samples of “projections” of the original event signature $s^*(t)$ at the locations of the sensors. These projections accommodate the impact of the signal propagation, or, in other words, they are reflective of the physical geography of the system and the medium properties. Since, decision




Figure 2: The QoI analysis framework and toolkit usage architecture

making is really made based on the recorded observations, if the above projections were known, one could have proceeded with the detection analysis unbeknownst to the physical geometry of the system. In other words, a core fusion and detection analysis engine can be developed that is independent of “external factors” by assuming knowledge of the signal projections at the sensor locations. Anchored on this core analysis engine, a system-level analysis framework can be developed that accommodates the remaining system parameters, like the deployment and observation topology, the application domain (and hence the collection of event signatures $s^*(t)$ to be encountered), the signal propagation models, the noise models, and so on.

Based on the above observation and the reference system model in figure 1, next we propose a toolkit architecture and a framework for a QoI analysis system.

2.2 The QoI analysis toolkit architecture

Figure 2 shows the component and “usage” architecture of the proposed toolkit system. It comprises three major functional blocks responsible for: (a) *input pre-processing*; (b) *detection* (QoI) *analysis* (the core analysis engine); and (c) *output post-processing*. The input pre-processing deals with all those aspects of the toolkit that generally relate to the constraints imposed on the signal(s) to be detected. These constraints include the deployment and observation topologies (which determine where the sensors are located and where the events occur), the signal propagation and attenuation models (which determine how the original signal projects itself at the sensor locations), the sampling policies (which determine which sensors contribute which samples to the detection process), the noise models (which determine the distortion process of the signal), and so on. All of the above contribute to the creation of the observation sequence $\{r_1, r_2, \dots\}$ that feeds into the core analysis engine which then calculates the QoI attributes for a very specific (well defined) set of system parameters. The figure also shows the special case of additive noise, where, typically, noise samples $\{n_1, n_2, \dots\}$ are simply added to the signal samples $\{s_1, s_2, \dots\}$; this case will be studied further in the next sec-

tion. Given the application requirements, post-processing of the QoI analysis results may be necessary, for example, to calculate averages over an observation region, or calculate optimal position of sensors, and so on. During post-processing, the services of the core analysis engine may be requested again for the QoI analysis of the system for a different set of system parameters.

In addition to the toolkit itself, the figure also shows the relationship between the application planner, system designer and the toolkit. The planner provides the problem definition and constraints, e.g., provide the deployment and observation regions, any cost constraints, the application domain, the desired QoI level, and so on. The designer provides system level model libraries, e.g., propagation models, noise models, communication and interference models, multi-path models, and so on. When a deployment plan with associated QoI performance levels has been produced, the planner may take a decision as to whether the expected performance of the system will be satisfactory enough, or may request additional analysis and trade-off studies for “what if” situations.

The layered analysis framework suggested by the toolkit allows us to place (research) emphasis on different aspects of the system, while still be able to relate these aspects back to the overall system design objectives. For example, in the next section, we focus on the core analysis engine independently of the specifics of the sensor deployment. Special cases of the latter then are considered during the derivation of numerical results in section 4.

We close this section with an example of input pre-processing; this will be discussed further later on in section 4. The example relates to the physical topology (geometry) of the sensor network. Let the physical topology of the system be represented by the distance vector² $\mathbf{d} = [d_1, \dots, d_M]^T$, where d_k is the distance of the path k that a signal takes from the event location to sensor k . Over path k , let $a_k(t)$ represent the cumulative environmental impact, e.g., the propagation *attenuation*, on the signal, and let v be the propagation speed. Assuming that an event occurs at time $t = 0$ and possesses the (transient) event signature $s^*(t)$, the signal signature seen by sensor k (the k -th event signature projection), $1 \leq k \leq M$, would be (excluding any noise components):

$$s_k(t) = a_k(t)s^*(t - \tau_k)u(t - \tau_k), \quad (1)$$

where $u(t)$ is the unit step function. The *time shift* τ_k is due to the propagation delay to sensor k and equals $\tau_k = d_k/v$. As discussed earlier, if $s_k(t)$ were known, a QoI analysis methodology could be employed independently of the original signal $s^*(t)$.

3. THE CORE QOI ANALYSIS ENGINE: HYPOTHESIS TESTING

Hypothesis testing is one of the primary tools used for the performance analysis of detection systems [12, 8]. We will base our core QoI analysis engine on hypothesis testing as well. In this section, we will quickly review the key results from hypothesis testing and we will (re)interpret them

²Bold letters represent the (column) vector version of a corresponding collection of parameters. Unless otherwise stated, T represents the matrix transposition operation.

within the context of our toolkit for sensor networks. In the course of doing so, as a result of our focus on general signals, some new results for detection in sensor networks will also be derived.

In hypothesis testing, a number of hypotheses is made, e.g., no event occurred (the *null* event), an event occurred, an event of type E occurred, and so on. Then, based on observations made, sampled data in the context of sensors, one of these hypotheses is declared to hold true. The selection of a hypothesis is done to satisfy certain performance objectives. In Bayesian hypothesis testing, which we will also adopt in this paper, the objective is to minimize the average cost of making a decision.

For the core QoI analysis engine in our toolkit, we start by considering binary hypotheses, i.e., event occurrence, (hypothesis H_1) vs. the null event (hypothesis H_0). The general QoI-influenced hypothesis testing formulation for a single sensor system were presented in [1, 2] and is based on the following traditional formulation [12, 8]:

$$\begin{aligned} \text{hypothesis } H_1 : r_i &= s_i + n_i, & i = 1, \dots, N, \\ \text{hypothesis } H_0 : r_i &= n_i, & i = 1, \dots, N; \end{aligned} \quad (2)$$

where, under hypothesis H_1 , s_i represents the value of the signal at the i -th sampling instance, while, under both hypotheses, n_i represents an additive noise component that is added to the i -th sample, and r_i represents the i -th measurement that is contributed to the fusion subsystem. A decision is made based on the *likelihood ratio test* (LRT):

$$\Lambda(\mathbf{r}_N) = \frac{f_{\mathbf{R}_N|H_1}(\mathbf{r}_N)}{f_{\mathbf{R}_N|H_0}(\mathbf{r}_N)} \underset{\text{select } H_0}{\overset{\text{select } H_1}{\geq}} \eta, \quad (3)$$

where $f_{\mathbf{R}_N|H_i}(\cdot)$ represents the probability density function for the N observations conditioned on hypothesis H_i , $i \in \{0, 1\}$ ³; for notational brevity, in the sequel, we will skip the size index N unless necessary. The threshold η is calculated from the a priori probabilities for the two hypotheses, P_0 and P_1 . If cost of one unit is incurred (only) when a wrong decision is made $\eta = P_0/P_1$, and the Bayesian test minimizes the risk of making a wrong decision. When the noise is described by a zero mean, additive and stationary Gaussian process with covariance matrix $\mathbf{C} = E\{\mathbf{n}^T \mathbf{n}\}$, where $\mathbf{n} = [n_1, \dots, n_N]$, then the test in (3) reduces to the following:

$$l \triangleq \mathbf{r}^T \mathbf{C}^{-1} \mathbf{s} \underset{H_0}{\overset{H_1}{\geq}} \eta^* \triangleq \ln(\eta) + \frac{1}{2} \mathbf{s}^T \mathbf{C}^{-1} \mathbf{s}. \quad (4)$$

The parameter l is referred to as the *sufficient statistic* and represents the summary operation performed by the fusion subsystem on the measurements \mathbf{r}_i . Note that the execution of the comparison between l and η^* , for deciding in favor of the one or the other hypothesis, is the responsibility of the detection subsystem in figure 1. Finally, in (4)

$$\psi^2 \triangleq \mathbf{s}^T \mathbf{C}^{-1} \mathbf{s}, \quad (5)$$

is reflective of the *signal-to-noise ratio* (SNR) for this setup.

In the special case of a single sensor system in additive, white Gaussian (AWG) noise process with zero mean and variance

³ \mathbf{R}_N represents a vector random variable for the N observations and \mathbf{r}_N a sample realization of N observations.

σ^2 (a $\mathcal{N}(0, \sigma)$ random variable), the QoI performance metrics, i.e., the probability of detection P_d and false alarm rate P_f , can be derived from (4), and are given by [2, 8]:

$$P_d = \Pr(l \geq \eta^* | H_1) = 1 - \Phi\left(\frac{\ln(\eta)}{\psi} - \frac{\psi}{2}\right), \quad \text{and} \quad (6a)$$

$$P_f = \Pr(l \geq \eta^* | H_0) = 1 - \Phi\left(\frac{\ln(\eta)}{\psi} + \frac{\psi}{2}\right), \quad (6b)$$

respectively, where $\Phi(\cdot)$ is the cumulative distribution function of a $\mathcal{N}(0, 1)$ random variable. The square of the parameter ψ is reflective of the signal-to-noise ratio (SNR) and in particular:

$$\psi^2 = \frac{1}{\sigma^2} \sum_{i=1}^N s_i^2. \quad (7)$$

Within the context of a sensor network, the covariance matrix \mathbf{C} of the noise process in (4) may capture correlations across both the spatial (i.e., across the sensors) and temporal dimensions. We are currently investigating the implications of this fact. In this paper, we consider the special case where the covariance matrix is diagonal with variance σ_k^2 for the measurements related to sensor k , $1 \leq k \leq M$.

Due to the additive nature of the terms in (4), it could be possible that (under certain conditions) the ‘‘contributions’’ from the various sensors to (4) are separable. In this case it will be easier to describe the system-wide QoI performance using the performance formulations for single sensor systems in (6) and (7). We will study the implications of this possibility in the next two subsections where we instantiate the above discussion for two different fusion architectures.

In a sensor network, a fusion architecture captures how sensors, and their samples, relate to the fusion and ultimately detection subsystems, see figure 1. In one case, we may have all sensors feeding their samples to a single fusion subsystem; the $M \rightarrow 1$ or centralized case. In another, more general case, we may have multiple fusion and detection subsystems in the system, with each sensor associated with only one of them; the $M \rightarrow L$ case, where $1 \leq L \leq M$. Just as we consider one or more fusion centers, we can also consider the case of having multiple decision subsystems tied, eventually, to a single system-level decision subsystem. Each fusion subsystem will be associated with only one of the decision subsystems. The possibilities of such architectures are numerous and choosing the right one will depend on some form of cost vs. performance trade-off analysis as well; this cost could include the cost of deployment, cost of nodes of various types, cost of energy paid, communication cost and so on. Next we consider two extreme cases: (a) $L = 1$ and (b) $L = M$.

We start by introducing first some notations. Let \mathcal{S}_e^M represent the set of all sensors. Let also \mathcal{R}^N and \mathcal{S}_p^N represent the set of all the sensor observations and the set of all the corresponding signal projections, respectively, involved in a single instance of the detection process; in a sense, the sets \mathcal{R}^N and \mathcal{S}_p^N collect all the r_i 's and s_i 's in (2). The cardinality of the sets is $|\mathcal{S}_e^M| = M$ and $|\mathcal{R}^N| = |\mathcal{S}_p^N| = N$. Let also \mathcal{S}_p^k represent the set of signal (projection) samples $s_1^k, \dots, s_{N_k}^k$ due to sensor k , where $N_k = |\mathcal{S}_p^k| \geq 0$. Hence,

$\mathcal{S}_p = \cup_{k=1}^M \mathcal{S}_p^k$ and $N = \sum_{k=1}^M N_k$. Note that the N_k 's need not be equal to accommodate, for example, cases where the sampling rates of the various sensors are not equal. Finally, if \mathbf{I}_N represent the $N \times N$ identity matrix, then the block diagonal covariance matrix considered hereafter will satisfy:

$$\mathbf{C} = \text{diag}(\sigma_1^2 \mathbf{I}_{N_1}, \sigma_2^2 \mathbf{I}_{N_2}, \dots, \sigma_M^2 \mathbf{I}_{N_M}). \quad (8)$$

3.1 L=1: Centralized detection

When $L = 1$, all N measurements from all M sensors are fed to a single fusion and detection subsystem. Hence, in this case, we can apply LRT to the entire set of observations \mathcal{R}^N and signal projections \mathcal{S}_p^N to derive the ‘‘system-wide’’ versions of the QoI performance metrics. It follows from (5) and (8) that the system-wide SNR is given by:

$$\psi_{sys}^2 = \sum_{k=1}^M \sum_{i=1}^{N_k} \frac{(s_i^k)^2}{\sigma_k^2} = \sum_{k=1}^M \psi_k^2, \quad (9)$$

where ψ_k^2 is the SNR due to the signal projection at the location of the k -th sensor. It can be easily shown that the test in (4) results in QoI performance metrics as in (6) with the per sensor ψ in (7) replaced by the system-wide ψ_{sys} in (9).

DEFINITION 2. *The equivalent sensor of a multi-sensor system is a single-sensor sensing system that achieves the same QoI level as the multi-sensor system using the same observation set.*

It follows from (9) that the centralized detection system considered here possesses an equivalent detector.

3.2 L=M: Distributed detection

When $L = M$, detection occurs in two steps. Firstly, a detection decision is made by each sensor (or more accurately, by each detection subsystem that is associated with one-and-only-one sensor) based on its own measurements. Secondly, the detection decision is made at the system level based on the local detection decisions.⁴ It follows from section 3.1 that the QoI performance at the sensor level is given by (6) again, where we use ψ_k instead for ψ .

Next, to obtain the system level detection probabilities, we need to use a system-wide *detection policy*. A detection policy describes how the sensor-level detection decisions combine to generate the system-wide decision. Such a policy will typically be based on some form of a counting strategy, e.g., decide that the event has occurred if at least, say, Q out of the M sensors indicate that it has. Note that while one may tend to use a majority rule, where $Q > M/2$, we do not believe that this is necessarily the best policy for the general case. The use of the majority rule tacitly assumes that each sensor detects the event with equal probability, which may not always be true for any sensor deployment geography. We elaborate on this further in section 4.1

In accordance to our toolkit’s usage architecture, see figure 2, the detection policy is passed to the core analysis engine by the system designer. Any fine tuning of the detection policy to increase performance, e.g., comparing different decision thresholds, is made via output post-processing.

⁴These two steps can be generalized to a hierarchy of decision steps and to the case where $1 < L \leq M$.

Assuming a ‘‘counting’’ detection policy, let \mathcal{S}_q^M represent the collection of all subsets of sensors that contain q sensors; the cardinality of \mathcal{S}_q^M is $M!/(q!(M-q)!)$. To declare that the event has occurred, when it has really occurred, i.e., under hypothesis H_1 , there should be at least one set of sensors $\mathbf{x}_q \in \mathcal{S}_q^M$ with $q \geq Q$ for which, *all* the sensors in \mathbf{x}_q indicate that the event has occurred.

Each sensor makes a local decision based on its own measurements. Thus, given these measurements, each sensor makes a detection decision independently of the other sensors. This conditionally independent decision making process should not be confused with the fact that measurements made by the sensors may relate to each other. Let P_d^k and P_f^k represent the probability of detection and false alarm for sensor k , respectively. Then, the system-wide probability of detection $P_d(Q; M)$ for a given threshold Q is given by:

$$P_d(Q; M) = Pr(q \geq Q | H_1) = \sum_{q=Q}^M \left\{ \sum_{\mathbf{x}_q \in \mathcal{S}_q^M} \left[\left(\prod_{\substack{\text{sensor} \\ k \in \mathbf{x}_q}} P_d^k \right) \left(\prod_{\substack{\text{sensor} \\ k \notin \mathbf{x}_q}} (1 - P_d^k) \right) \right] \right\}. \quad (10)$$

A similar expression for $P_f(Q, M)$ can be obtained by substituting P_d^k with P_f^k in the above.

While highly unlikely, when all the P_z^k 's, $z \in \{d, f\}$, are equal to P_z^* independently of k , e.g., when the event (if it occurs) is equidistant from all the sensors, then (10) –and likewise for the $P_f(Q; M)$ – reduces to the tail distribution of a binomially distributed random variable with parameters M and P_z^* and thus:

$$P_z(Q; M) = \sum_{q=Q}^M \left[\frac{M!}{q!(M-q)!} (P_z^*)^q (1 - P_z^*)^{(M-q)} \right]. \quad (11)$$

A system with the above binomial case was considered in [11]. Note that for $Q \leq M-1$, (11) can be conveniently evaluated recursively:

$$P_z(Q; M) = (1 - P_z^*)P_z(Q; M-1) + P_z^*P_z(Q-1; M-1). \quad (12)$$

4. NUMERICAL STUDIES

In this section, we supplement the core QoI analysis of the previous section with specific deployment models to derive performance results for the overall resultant systems. Specifically, we consider two cases. In the first case, we compare the QoI performance of a totally distributed system with its corresponding centralized system as discussed in sections 3.1 and 3.2. In the second case, we study the behavior of a sensor system by deriving upper and lower bounds for the probability of detection P_d . The use of the framework allows us a systematic step-wise development of the performance results reusing analysis approaches as necessary.

For both cases, we use a four-sensor system described in table 1. We assume two families of decaying electromagnetic signals, the (relatively) ‘‘slow’’ decaying $s^*(t) = 1 - t^n$ and ‘‘fast’’ decaying $s^*(t) = (1 - t)^n$ that last for one time unit. The family of event signatures chosen, which may represent envelopes of more realistic signatures, is a convenient one. For an appropriate selection of the *decay parameter* n ($n \geq 0$), the signatures can span the range from a unit pulse

Topology	$M = 4$ sensors $\mathbf{d} = [d_1, d_2, d_3, d_4]^T = [1, 2, 4, 8]^T$
Sampling policy	$[N_1, N_2, N_3, N_4] = [20, 20, 20, 20]$ N_k samples in 1 time unit
Attenuation model	$a_k(t) = a_k = 1/(1 + d_k^2)$
Propagation model	$\tau_k = d_k/\nu$ ($\nu = 3 \times 10^8$ m/s)
Noise variance	$\sigma_1^2 = \sigma_2^2 = \sigma_3^2 = \sigma_4^2 = \sigma^2$
event signature	$s^*(t) : 1 - t^n$ and $(1 - t)^n$

Table 1: Parameters for the reference system




Figure 3: Centralized vs. distributed architectures detection ($M = 4$, $\eta = 1$ and $1 \leq Q \leq M$)

(or, even a fixed-amplitude constant signal) to a unit delta function. Per the toolkit figure 2, the parameters populating table 1 are provided by the application planner and/or the system designer.

4.1 Centralized vs. distributed architectures

In this section, we study the QoI merits of the centralized ($L = 1$) and fully distributed ($L = M$) architectures. Furthermore, we will study their performance as our predisposition of whether the event has occurred (i.e., the priori probability for hypothesis H_1) changes. We consider $\eta = P_0/P_1$ (see discussion about η following (3)). We first consider $\eta = 1$, i.e., $P_0 = P_1$ which will serve as our reference point and then $\eta = 2$ (> 1) and $\eta = 0.5$ (< 1) as well. The signal signature will be $s^*(t) = 1 - t^2$, and each sensor contributes 20 samples. We consider a “lucky” system that happens to start sampling just after the signal signature arrives at its sensors. We will discuss this idealistic case further in section 4.2.

When $\eta = 1$, figure 3.a and 3.b show the P_d and P_f as a function of σ (the standard deviation of the noise) for the centralized architecture and for the distributed architecture for various decision thresholds Q . As expected, in the case of the distributed architecture, both probabilities decrease with increasing Q as the detection policy decides in favor of H_0 more liberally; this trend holds true for all cases of η studied. When comparing the QoI performance metrics of the two architectures, the metrics for the architecture lie between that of the distributed ones for $Q = 2$ and 3, which




Figure 4: Centralized vs. distributed architectures detection ($M = 4$, $\eta = 2$ and $1 \leq Q \leq M$)

is above the “middle point” ($Q = 2$).

Naturally, one will expect that the higher the P_d , and the lower the P_f , the better the QoI will be. To achieve a more direct comparison of the two architectures, we use the average of two metrics represented by the *probability of error* $P_e = P_0 P_f + P_1 (1 - P_d) = P_1 (\eta P_f + (1 - P_d))$, which is shown in figure 3.c. The centralized architecture achieves the best (i.e., smallest) P_e . That is not surprising as this architecture makes best use of the sampled observations. However, an interesting byproduct of this analysis is that setting Q at the middle point ($Q = 2$) attains the best (i.e., smallest) P_e for the distributed architecture. On the other hand, when $Q = 1$ or $Q = M$, due to their poor P_d and P_f performance, these two extreme cases achieve the worst (i.e., highest) P_e .

The case where $\eta = 2$ is shown in figure 4. First of all, with $\eta > 1$ (i.e., $P_0 > P_1$), the condition for deciding in favor of H_1 becomes harder to achieve when compared with the $\eta = 1$ case, see (4). Therefore, P_d and P_f decrease in magnitude when compared with these probabilities when $\eta = 1$ (for the same σ) for both architectures. Comparing the two architectures, the P_d and P_f for the centralized one lie between $Q = 1$ and 2 of the distributed one, which is below the “middle point.” Comparing the P_e performance of the two architectures, we again see the centralized one achieving the best performance. With respect to the distributed architecture, the best distributed detection policy is attained when Q is set to 2 for smaller σ and 1 for larger σ . The reduction of the optimal Q threshold in the case where $\eta > 1$ when compared with that for $\eta = 1$ follows from the fact that with decreasing P_1 it becomes harder for any sensor to declare in favor of H_1 . Therefore, having even a small(er) number of sensors declaring in favor of H_1 is reason enough to decide in favor of H_1 system-wide.

Finally, the case where $\eta = 0.5$ is shown in figure 5. Arguing as before, the behavior of the QoI metrics in this case relative to those of the $\eta = 1$ case is in reverse order to the behavior experienced when $\eta = 2$. With respect to P_e , while the centralized architecture still performs the best, the best




Figure 5: Centralized vs. distributed architectures detection ($M = 4$, $\eta = 0.5$ and $1 \leq Q \leq M$)

distributed detection policy is attained when Q is set to 3 for smaller σ and 4 for larger σ .

Based on the behavior of the two systems when compared on the basis of P_e , we conjecture that in general the best detection policy for the distributed architectures is achieved at a threshold Q that decreases away from $M/2$ (and possibly toward 1) as η increases above 1. On the other hand, the threshold Q increases away from $M/2$ (and possibly toward M) as η decreases below 1.

4.2 The lucky and unlucky sensor

In the previous subsection, we made reference to the lucky sensor k , that starts sampling the signal just after the signal projection first arrives at the sensor location (and then takes additional $N_k - 1$ samples as well). For the decaying signal considered, this lucky sensor picks the strongest possible signal (or, better, event signature projection) measurements, and hence achieves the highest possibility of detecting the event, hence, $P_{d;l}^k \geq P_{d;r}^k$, where the indexes l and r stand for lucky and real (k -th sensor), respectively. Similarly, we may have an unlucky sensor that samples just before the projection arrives at the sensor and, hence, when the signal samples are taken, they are the weakest possible, which results in $P_{d;u}^k \leq P_{d;r}^k$.

In general, under hypothesis H_1 , a sensor will take its first signal y_k time units after the signal projection arrives at the k -sensor, see figure 1, where $0 < y_k < T_k$; where T_k is the sampling period of the k -th sensor. The i -th sample contributed to the decision process by the k -th sensor is

$$\hat{s}_i^k = a_k * s^*(t_i - \tau_k + y_k) + n_i, \quad 1 \leq i \leq N_k. \quad (13)$$

Since, y_k is in general unknown (could be modeled as a random variable), the convenience of using the lucky and unlucky sensors instead to bound the system performance for the case of decaying signals becomes apparent. We will study these bounds next; this analysis can be generalized for non-decaying signals as well.




Figure 6: Performance region for $M = 4$ sensors, $s^*(t) = (1 - t)^4$ and $\eta = 1$




Figure 7: Performance comparison for fast-decaying signals $(1 - t)^n$ and slow decaying signals $1 - t^n$

We will consider the class of fast decaying signals $s^*(t) = (1 - t)^n$, applied to a centralized system with parameters as in table 1, we also assume that $P_0 = P_1$, i.e., $\eta = 1$. Along with the upper and lower bounds derived from the lucky and unlucky sensors, we also consider two additional approximations: (a) an approximation derived when $y_k = T_k/2$ for all sensors (we may drop the index k in this case); and (b) an average signal approximation given by

$$\hat{s}_i^k = \frac{a_k}{T_k} \int_0^{T_k} s^*(t_i - \tau_k + y) dy. \quad (14)$$

Figure 6 shows the P_d for the four different approximations as a function of σ . As expected, the lucky system (where all the sensors are lucky) achieves the highest probability of detection and the unlucky system the lowest for all values of σ . Interestingly, the case where $y = T/2$ and the average signal approximation in (14) result in very close results. In general, it is expected that the higher the sampling rate becomes, the upper and lower bounds will become tighter, and so will any approximation that restricts itself between these two bounds. We also notice that as σ increases, the P_d decreases toward 0.5 for all cases (and, hence, for the real system too). Looking at (6a), this should be expected since as σ increases (or ψ decreases) and $\eta = 1$, the probability of detection tends toward 0.5. Applying this observation to the equivalent sensor of our centralized system results in the behavior seen in figure 6.

Finally, figure 7 shows a comparison of the performance behavior for fast and slow decaying signals as a function of n , for the same system as in figure 6. Specifically, we show the integral between the upper and lower performance bounds (noted as performance area in the figure) and the maximum difference (MAX_{diff}) between the bounds for each class of signals. It can be easily shown that, as n increases, the fast decaying signals tend to the unit delta function and the successive samples taken by a sensor differ significantly from one sample to next. Thus, the performance differences between the lucky and unlucky sensors for fast decaying signals increase. The reverse behavior is exhibited for the slow decaying signals that tend to the unit pulse with n and hence the differences between successive samples diminish. Note that when $n = 0$, we have a constant amplitude (persistent) signal and, hence, the upper and lower bounds due to the lucky and unlucky sensors coincide in this case.

5. CONCLUDING REMARKS

In this paper, we have used QoI as the means to capture an application's information needs from an underlying sensor networks. We then introduce a QoI-based framework for the evaluation of sensor network deployments. The deployment evaluation aids application planners to anticipate the performance capabilities of the network, thus, allowing them to possibly develop contingency plans in case the expected performance is below desired QoI levels.

The framework is reflective of how application planners and system designers will use it, and input parameters to it. It considers a class of event detection applications and comprises: (a) an input preprocessing step, where the physical constraints of the system deployment are considered; (b) a core analysis step, where a common performance analysis approach is used based on hypothesis testing; and (c) a result post-processing step, where specific solutions obtained through core analysis are combined and processed further to derive the desired end-results. We have applied the framework on analyzing a rather non-homogenous system comprising a finite-size sensor network with transient (non-constant) signals, arbitrary sensor deployment, and different noise levels at each sensor. Note that persistent events, typically considered in research work on sensor networks, can be studied as limiting cases of transient events, but this is not necessarily true the other way around. We have also analyzed the performance of "lucky" vs. "unlucky" systems with respect to sampling times thus developing performance bounds for various sampling possibilities.

We have compared the QoI performance of centralized vs. distributed detection architectures. For distributed schemes, we have also studied the selection of an optimal threshold for counting-based, system-wide detection policies. In the latter case, we have demonstrated how a priori knowledge about event occurrence influences the selection of the best detection policy threshold. This analysis is also applicable for the case where local fusion is performed but the detection decision-making is still centralized, e.g., when the sum on the RHS of (4) is calculated in a distributed fashion, but the inequality comparison is performed at a central location.

The proposed framework facilitates the decoupling of the three steps, and the mix-and-match of different analysis, and

modeling approaches (both existing and new ones), e.g., use of a Neyman-Pearson instead of Bayesian detection testing for the core analysis engine. The decoupling permits us to focus on separate aspects of the problem at different times, thus simplifying the analysis, and modeling approaches engaged at each step. This various approaches can then become components of a toolkit that can be used for the analysis of a broad collection of sensing systems.

6. REFERENCES

- [1] C. Bisdikian. On sensor sampling and quality of information: A starting point. In *Proc. IEEE PerSeNS Workshop*, White Plains, NY, USA, March 19, 2007.
- [2] C. Bisdikian. Quality of information trade-offs in the detection of transient phenomena. In *SPIE Unattended Ground, Sea, and Air Sensor Technologies And Applications IX Conference*, Orlando, FL, USA, April 9-13 2007.
- [3] E. Blasch and S. Plano. Dfig level 5 (user refinement) issues supporting situational assessment reasoning. In *8th International Conference on Information Fusion*, 25-28 July 2005.
- [4] R. Blum, S. Kassamand, and H. Poor. Distributed detection with multiple sensors: part ii-advanced topics. In *Proceedings of the IEEE*, volume 85, 1997.
- [5] S. Ehikioya. A characterization of information quality using fuzzy logic. In *18th Int'l Conf. of the North American Fuzzy Information Processing Society, (NAFIPS 1999)*, June 10-12 1999.
- [6] S. K. Jayaweera and K. Al-Tarazi. Large system decision fusion performance in inhomogeneous sensor networks. *Proc. of the IEEE Int'l Conf. on Video and Signal Based Surveillance*, page 72, 2006.
- [7] N. Katzenka, E. Levina, and G. Michailidis. Local vote decision fusion for target detection in wireless sensor networks. In *Joint Research Conference on Statistics in Quality Industry and Technology*, Knoxville, TN, USA, June 7-9, 2006.
- [8] S. M. Kay. *Fundamentals of Statistical Signal Processing: Detection theory*. Prentice Hall, 1998.
- [9] A. E. L. Yu. *Detection, Energy, and Robustness in Wireless Sensor Networks*. John Wiley & Sons, 2006.
- [10] L. Yu, L. Yuang, G. Qu, and A. Ephremides. Energy-derived detection scheme with with guaranteed accuracy. In *in Proc. IPSN'06*, Nashville, TN, USA, April 9-21 2006.
- [11] R. Niu and P. Varshney. Distributed detection and fusion in a large wireless sensor network of random size. In *EURASIP Journal on Wireless Communications and Networking*, volume 5, 2005.
- [12] H. L. van Trees. *Detection, Estimation, and Modulation Theory, part I*. John Wiley & Sons, 1968.
- [13] R. Viswanathan and P. Varshney. Distributed detection with multiple sensors: part i-fundamentals. In *Proceedings of the IEEE*, volume 85, 1997.

6.1 Acknowledgments

The authors would like to thank Prof. M. B. Srivastava for his support of this research work.