Protecting with Sensor Networks: Attention and Response

Jeffrey V. Nickerson
Stevens Institute of Technology
jnickerson@stevens.edu

Stephan Olariu
Old Dominion University
olariu@cs.odu.edu

Abstract

Sensor networks are expected to play an important role in hybrid protection infrastructures when combined with robots and human decision makers. In order to be effective, detection of intruders needs to be attended to, and a timely response needs to be made. Intruders, on the other hand, if they can't avoid detection may make use of strategies to overload attention or dilute the response. Such strategies are related to Denial of Service attacks in electronic security. This paper develops the conceptual framework for modeling the game between intruders and defenders.

1. Introduction

The use of sensors and actuators in surveillance applications goes back more than a century [14]. These sensors were generally bulky devices wired to a central control unit whose role was to collect, process, and act upon the data gathered by individual sensors. Wireless sensor network research, as we know it today, had its origins in the DARPA-sponsored *SmartDust* program [16]. The vision of SmartDust was to make machines with self-contained sensing, computing, transmitting, and powering capabilities so small and inexpensive that they could be released into the environment in massive numbers. These small devices have come to be called *motes* and serve as nodes in a sensor network.

Since building massively-deployed sensor networks is prohibitively expensive under current technology, in the past few years we have witnessed the deployment of *small-scale* sensor networks in support of a growing array of applications. These prototypes provide solid evidence of the usefulness of sensor networks and suggest that the future will be populated by pervasive sensor networks, that will redefine the way we live and work [1, 6, 7, 11, 12]. It is, thus, expected that in the near future, in addition to the existing implementations, a myriad of other applications, including battlefield command and control, disaster management and emergency response, will involve sensor networks as a key mission-critical component.

Current sensor networks, however, are for the most part modeled after conventional computing networks under centralized control and involve a small number of sensors usually deployed as the lowest layer in a multi-layer hybrid network. In critical applications the underlying sensor network is usually augmented by a second layer consisting of (mobile) robots monitored by human beings or, perhaps, by a combination of humans and robots. For, in spite of phenomenal advances in robotics, human intelligence is more capable of interpreting situations than machine intelligence. Besides, human experts are better at evaluating the broad context of the event at hand and are held responsible for their actions – thus, they ultimately will want to participate in crucial decisions.

We can imagine a sensor network of motes, small *motion detectors*, *metal detectors*, *pressure detectors, vibration detectors*, and the like, deployed around a valuable asset, say an electrical transformer. These motes may have the mission of detecting human intruders. While sensors do quite well in detecting motion and other characteristics, including weight and the presence of metallic objects, they may not do so well in differentiating animals from humans; or, they may not do so well in differentiating the friendly mechanic coming to maintain the device from the unfriendly saboteur.

Even if sensors were able to accurately classify intruders, we would probably want a human to decide what to do in response, as such a decision might depend on the available resources, the time to deploy them, the value of the defended object, and the political environment of the surrounding community.

In situations related to security in which sensor networks are used, it makes sense to discuss human-sensor network interaction the same way we discuss human-robot interaction. Indeed, if we expand our transformer example to include a deployable robotic guard, we are then clearly concerned with human-robot interaction.

In a previous paper, we discussed the use of perimeter defenses [16]. We pointed out that an axial model might be sometimes better than a perimeter model. We developed heuristics based on security as a Quality of

Proceedings of the 40th Annual Hawaii International Conference on System Sciences (HICSS'07)
0-7695-2755-8/07 $20.00 © 2007 IEEE
1530-1605/07 $20.00 © 2007 IEEE

Service (QoS) parameter, as other authors have done [9]. However, we did not consider the nature of human attention, or the nature of a response. Such considerations are generally absent from the sensor network literature, and they shouldn't be.

We consider how human attention can be modeled in relation to a sensor network, as reflected through the allocation of function between humans and sensors. First, we will conceptually discuss security in relation to sensor networks and humans. Then, we will develop a physical scheme for detection. Following this, we will discuss the human role in handling the false positives that will be part of any detection system.

2. Conceptualizing security

2.1 The nature of the problem

In this section we develop a vocabulary for sensor security. As with any human-machine symbiosis, much of our concerns will be over the allocation of function [3]. For, it is clear we would like the sensor network to do as much as possible with as little human supervision and direct intervention as possible.

So, we might imagine the ideal network would be capable of automatically detecting an intruder with no errors. Since there would be no errors, the sensor network could also automatically block or restrain the intruder until humans could intervene.

The problem, of course, is that there will be false positives: classifiers are seldom perfect. And the inadvertent restraint of bystanders is not popular.

Even short of automatic restraint, automatic perception is difficult. Part of the difficulty of designing a fully automated intruder detection system is that the sensor network will have to perform most (if not all) of the aggregation and fusion of the raw sensory data collected. While this task is, in theory, well understood, the sensors need external supervision at least initially. Indeed, a sensor network has considerable learning capabilities, often times referred to as "wisdom of the crowd" [14]. However, in order for the sensor network to reach its full capabilities it needs to be trained (supervised learning). At the moment this supervision comes in the form of remote experts interacting (e.g., by satellite) with the sensor network. Such is the case in most NASA missions where the sensory data collected, say, on Mars, is fused and interpreted in the mission control room here on Earth.

In addition, certain tasks, such as recognition and response strategy, can be performed better by humans, if they can focus their attention. Notice that the allocation of function here will have a different slant than that in, say, emergency response. For, in a security situation two entities are opposing each other: the *defender* and the *intruder,* an instantiation of an *adversary*. Since both the intruder and the defender are capable of intentional behavior, the conflict can be modeled as a game. The intruder seeks to penetrate a sensor network. We will simplify the intruder's goal to be one of reaching the center of the defended territory, where there is something of value. The defender tries primarily to prevent such an intrusion, and, secondarily, to capture any intruder so as to prevent recurrent attacks.

A moderately sophisticated adversary may engage in a *Denial-of-Service* (DoS) attack, flooding the network with bogus intrusion events in order to trigger a *massive* number of alarms that human responders cannot attend to. This is likely to have one of two possible effects: (1) either the humans will come to think that the sensor network does not function properly, and will perhaps disconnect it, or (2) they may decide that the sensor network does function properly but is *over-reacting,* in which case they will tend to ignore subsequent alarms. In addition, the DoS attack can serve the purpose of allowing a *potential* adversary to test the network in order to unearth the allocation strategies of the network, to aid in planning a future intrusion.

2.2 Perimeter and access security

Many have discussed the distinction between perimeter and access security, e.g., [10]. While perimeter security seeks to isolate, access security needs to let people pass. So, for example, museums use access security. Pure perimeter security is relatively rare; even when we fence off an area, someone needs to get inside, and so we add a door, which leads us to practice access security.

Access security leads to tradeoffs between false positives and false negatives. We might not want to let any intruder get in, and therefore may bias toward searching everyone. When this doesn't work, we may bias the opposite way, searching nobody.

Intrusion is particularly difficult to defend against in access control situations. As intrusion is by its nature rare, we sometimes over-react to detection events, even though by prior probabilities, we should be less alarmed. One study, in analyzing the prior probabilities, identified the false positives as the limiting factor in computer intrusion detection systems [2]. (While computer intrusion is distinct from physical intrusion, there are interesting analogies, as sensor systems have both a physical and a network characteristic.) From our perspective, the false positive is an issue for two related reasons.

First, human intervention is usually called for when an intrusion is detected. Therefore, attention will be distracted by too many false positives. If there are a large number of false positives over a sustained period of time,

then the alarms tend to be ignored. We witness such a phenomenon in every day life when car alarms go off and excite no one into action.

Second, an intruder can take advantage of false positives. For example, as already pointed out, the intruder can do something which will flood the network. The defender will initially be distracted chasing down false positives, and eventually will cease to respond to any alarm, allowing the intruder through.

We note that there is a hard physical reality to multiple attacks; given any kind of defender, human or automated, a large enough attack will overwhelm the defender. Therefore, a protection network will have a certain capacity, as in QoS – more resource may provide defense against larger teams of intruders.

The problem is also one of human attention. As the number of attacks increase, attention is split, and eventually fatigued. In addition, unaided human cognition will tend to over-estimate the likelihood of an intrusion given an alarm; this tendency can be exploited by intruders, who may attempt a DoS attack on the sensor network.

Ideally, a DoS attack should be recognized as just that and both attention and response resources held back until an intruder is identified with more certainty. Sometimes delay in responding might be the better course of action.

Later, we will look at how this decision to hold back might be made. In any case, human response to an alarm does not scale – humans are prone to fatigue and unless a large team of responders is available, cannot attend to or confront concerted attacks. This state of affairs makes it imperative to enlist the help of a hybrid system composed by a reliable sensor network apt to filter out the vast majority of false positives and that cannot be easily fooled into a false negative (i.e., failing to report an intrusion when one occurs). The topic of building adaptive sensor networks in support of reliable intrusion detection will be further discussed in Section 3.

3. Building adaptive sensor configurations

In this section, we focus on the construction of the hardware of sensor configurations, in preparation for a discussion in Section 4 of the human component of the sensor system.

3.1 A layered response architecture

One of the key advantages of a sensor network is it functional versatility. Indeed, while the sensors once deployed do not move, functionality can migrate freely in the resulting network. In particular, individual sensors can be activated or deactivated, placed in high alert or in stand-by mode. By selectively activating groups of sensors several defensive configurations can be obtained as shown below.

Consider a sensor deployment around a central asset and refer to Figure 1. Each sensor has a sensing range (often denoted by a disk centered at the sensor) and a transmission range (not shown). The sensors self-organize into a wireless sensor network as detailed in [16, 20, 22-24]. In this work we assume a virtual infrastructure grafted on top of the set of sensors, compatible with the ideas discussed in [16, 17]. At deployment time the sensors are inactive (in the sequel, blue sensors are inactive, red ones are active).

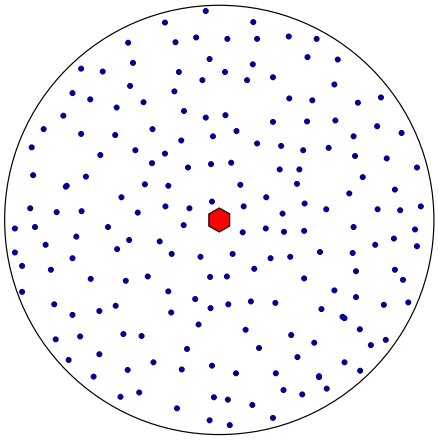

Figure 1. The original deployment

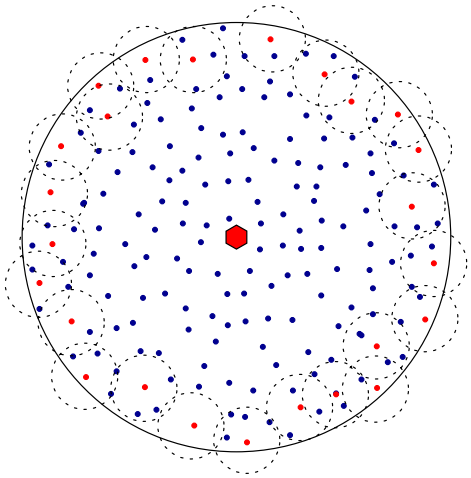

Figure 2. Randomized activation of outermost corona

To begin, following a randomized scheme where the sensors become active with a pre-determined probability p, the sensors close to the outer boundary of the

deployment area (outermost corona, in the terminology of [16]) become active as illustrated in Figure 2.

The intention is to set up a reliable *early-warning system* (EWS, for short) that can alert the defender of incipient intrusion events. The activation probability p of individual sensors can be determined accordingly. We refer the reader to Subsection 3.2 for the technical details related to computing p. There is an obvious tradeoff here: the higher the activation probability, the more sensors will be awake per time unit, the better and more reliable the EWS, the shorter the time to certainty, but also the higher the number of potential false positives. In practice, the activation probability will be set as a function of other components of the EWS (clearly, no EWS should rely on a single technology).

Now assume that the EWS is alerting the defender to the possibility of intrusion: in other words, one or more of the sensors in the EWS have detected a possible intruder that has penetrated the guarded area. Notice that there is no immediate response from the part of the defender. However, at this moment, a further set of sensors are activated. For example an axial configuration may be set up as discussed in [16]. Such an axial configuration is featured in Figure 3, where there are eight axes each at an angle of 45^o from its neighbors. It is worth noting that the axes are set up by a simple broadcast message from the sink. Namely, all the sensors in a small wedge about the desired axes are activated. The directions of the axes are also communicated by the sink. The sensors in the respective wedges will activate themselves. Of course, this is predicated on the sensors being aware of the angular distance from the sink. We refer the reader to [16] where this issue is discussed further.

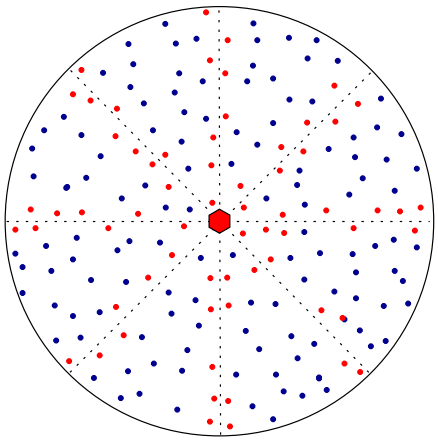

Figure 3. Setting up an axial system

As pointed out in [16] the axial defensive system is of assistance in determining whether or not an intrusion has taken place and to pinpoint its location. We refer the interested reader to [16] for an analytical derivation of the detection probability afforded by an axial defensive system.

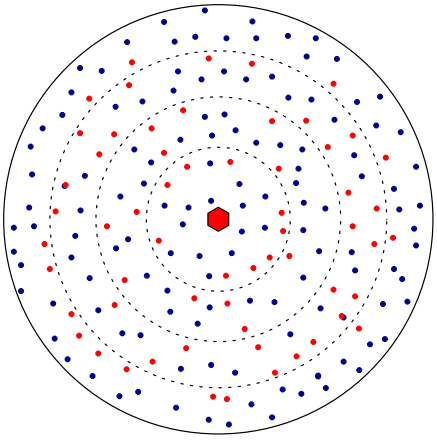

Figure 4. Setting up circular perimeters

In some cases the defender may choose to forgo the setup of an axial defensive system preferring a circular perimeter system instead as illustrated in Figure 4. The details are as follows. Imagine the deployment area partitioned into concentric disks of decreasing radii centered at the sink. These disks partition the deployment areas into *coronas* which will play an important role in our discussions of intervention and restraint in later sections of the paper. The task of setting up the coronas and the wedges (determined by the axis system above) is referred to as *training*. We refer the reader to [17, 23] for efficient training protocols.

With training in place, the sensors around the boundaries of the coronas can be activated by a simple broadcast message form the sink. Importantly, in order to save energy individual sensors that qualify for being activated do so with a certain probability that depends on the perceived level of danger.

3.2 Computing the activation probability

The main goal of this subsection is to provide a closed form for the probability with which a sensor needs to be activated in order to obtain a circular defensive perimeter. For this purpose, assume that the sensors were deployed uniformly at random in the coverage area with density ρ.

Consider a corona of width d bounded by the disks of radii R-d and R, and refer to Figure 5. We begin by determining a virtual circle Γ, shown in dotted lines in

Figure 5, that corresponds to the expected distance of a sensor to the sink.

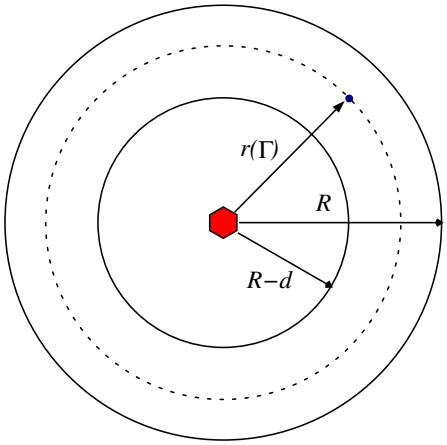

Figure 5. Illustrating r(Γ)

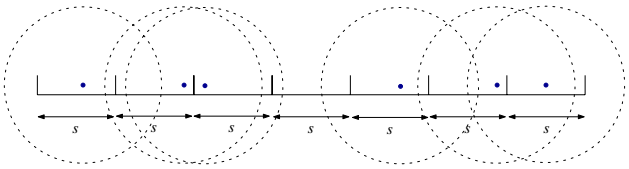

Figure 6. Filling the detection gaps

It is easy to see that the expected radius of Γ is

$$r(\Gamma) = \frac{2}{d(2R-d)} \int_{R-d}^{R} x^2 dx = \frac{2}{3}\left[2R - d - \frac{R(R-d)}{2R-d}\right]. \quad (1)$$

Next, notice that the number N of sensors deployed in the corona is given by

$$N = \rho\pi d\left[2R - d\right].$$

In an expected sense, we may think of these N sensors as being deployed on the circle Γ itself. Now, assuming that the sensing radius of a sensor is s, the minimum number of active sensors needed to cover the circumference of Γ is $\frac{2\pi r(\Gamma)}{2s} = \frac{\pi r(\Gamma)}{s}$.

It follows that the probability p with which a sensor in the given corona is activated is

$$p = \frac{\pi r(\Gamma)}{sN} = \frac{r(\Gamma)}{s\rho d(2R-d)}.$$

Now, replacing the value of $r(\Gamma)$ obtained in (1) we obtain the desired closed form for p. Specifically, we can write

$$p = \frac{2}{3\rho d}\left[1 - \frac{R(R-d)}{(2R-d)^2}\right]. \quad (2)$$

3.3 Making detection watertight

The activation probability derived in (2) is in some sense very optimistic: it assumed that the sensing areas are nicely lined up "shoulder to shoulder" making it impossible for the intruder to sneak by undetected. In other words it provides a probabilistic detection that is good on the average. The randomness of the process makes things different in practice. To understand the problem refer to Figure 6 where the circumference of Γ has been partitioned into segments of size s (the detection radius). Assume, further, that only the sensors featured in Figure 6 have been activated. Visual inspection shows that there is a gap in the coverage through which an intruder gets in undetected.

Thus, if detection with high probability is desired then a larger number of stations have to be active and, consequently, the activation probability has to increase. The insight as to how to proceed is suggested by Figure 6. Perceiving the segments of size s as *bins* and the activated sensors as *balls*, we are in the presence of a classic *balls-and-bins* problem. This particular instance of the problem asks for the least number of balls that have to be thrown into the bins in such a way that no bins are unoccupied. The reader should have no difficulty to confirm if all the bins are occupied then the detection is guaranteed. It turns out that this problem is also known as the *Coupon Collector's Problem* [22 (Theorem 8.2)], after a frivolous application in which each cereal box contains one coupon. Given that there are k different coupons and that the coupons have been placed uniformly at random, one per cereal box, how many cereal boxes does a housewife need to buy before she has all the k coupon varieties. It turns out that the answer is $k \ln k$ where ln stands for the natural logarithm.

In our case the number of coupons is

$$k = \frac{2\pi r(\Gamma)}{s}.$$

And, consequently, the number of active sensors should be

$$\frac{2\pi r(\Gamma)}{s}\ln\frac{2\pi r(\Gamma)}{s}.$$

With the new activation probability (that, at the risk of some overload, we also denote by p) can be written as

$$p=\frac{\frac{2r(\Gamma)}{s}\ln\frac{2\pi r(\Gamma)}{s}}{\rho d\left[2R-d\right]}.$$

4. Focusing human attention

4.1 The geometry of attention

Once a sensor network has detected a possible intrusion event, the next issue becomes the differentiation of real intruders from false positives.

The typical and natural way is to send a human being out to investigate. This has many advantages, as a person can look at the potential intruder, and sometimes recognize instinctively if the person is a threat or not. If the person is a threat, then the defender can attempt a capture. However, in an environment with many false positives, this technique does not make sense.

It is clear that the longer one waits, the more information one gets, and the more certain one can be about whether or not an alarm is associated with an intruder. For example, approaches using temporal Bayesian nets will feed previous states to the current state, which in turn should adjust the probabilities of an entity being an intruder [29]. We can say that *P(I|A)* (the conditional probability of an intruder existing given an alarm) for a real intruder will become higher over time, and the false positives *P(B|A)* (the conditional probability of a bystander having triggered an alarm) will go down over time. In the language of signal detection theory, the receiver operating characteristics for classification will get better the more time we allow.

However, there is a tradeoff. For, there are other times that are important to consider in the network. One is the *time to intervention*. This is the amount of time it would take to both make the decision to intercept a possible intruder and actually accomplish the intervention. From the intruder's perspective, there is the intruder's *time to the target*, in this set of scenarios the time to reach the center.

It is useful to imagine a *time to certainty*. If the probability of an intruder is above a certain point, then we will consider that entity an intruder at that instant and, at least, move to physically intercept. This concept is discussed in more depth in the next subsection.

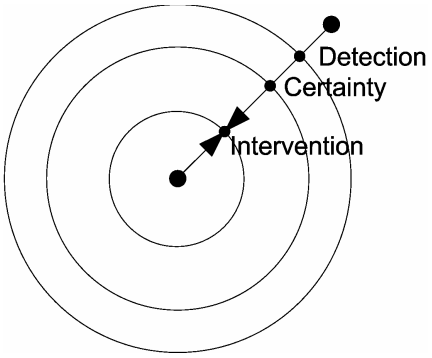

Figure 7. Detection is the first step; the intruder will continue to move while the defender achieves certainty, and moves to intervene.

This is shown in Figure 7. Integrating the time measures together, we can say that a defender wishes that the following always holds:

time to detection + time to intervene < intruder's time to target

In other words, the intruder needs to be seen before it is too late. However, this may not be sufficient, for in situations where there are many false positives, we probably can't deploy intervention resources to attend to every alarm; what we want is the following:

time to certainty + time to intervene < intruder's time to target

This is a harder constraint, as *time to certainty > time to detect.* Detection has to happen early enough so that certainty of detection can follow while still allowing time for intervention.

The previous discussion in section 3 on the density of sensors assumes a greater importance as we contemplate this diagram. If we assume the defender is centrally located, and moves at the same speed as the intruder, then we can see that between the time that the defender is sure there is an intruder and the time of intervention, the intruder can halve the distance to a target.

In order to provide a greater spatial buffer, the detection corona needs to be pushed back; as this pushes back, the number of sensors needed in the network increases quadratically in the newly added corona width. This is a QoS tradeoff – we can buy more time to attend and respond by extending the sensor network; however the quadratic increase in sensors means this will eventually become infeasible.

4.2 Time to certainty

While the word *certainty* connotes an absolute, it is clear in many decision making situations we rarely achieve full certainty. Instead, we pick a probability threshold and decide to act when we exceed the threshold. *Time to certainty* is the time to cross this threshold.

There are three ways we can imagine getting to this threshold.

Using passive sensors, we can wait until enough samples have been taken to form a trajectory signature, which we know from experience has a high probability of predicting an intruder.

Alternatively, we can take advantage of the human ability to recognize. Let us assume that, in addition to motion sensors, we have control of visual sensors which can see any designated part of the guarded circle. Then we might achieve certainty by viewing the potential intruder through the camera. Such visual monitoring is already a part of most corporate security, and in many cases the monitoring is remote.

With this method, the time to certainty might involve the time to view the image and form an impression. If there are multiple intruders, then a queue will be established. Each intruder will be looked at, a decision made, and the next one evaluated.

This brings to mind issues discussed in the research on teleoperation of robots [5, 21]. There will be switching costs associated with moving attention from one physical context to another. Attention will sometimes be spread across a range of events, but sometimes attention will be totally focused on one situation. This will tend to create a queue of information unprocessed by humans. The time the information is ignored is equivalent to what Olsen and Goodrich called *neglect time* [18]. In the case contemplated here, the neglect of the adversary extends the time to certainty, and thus the time to initiate a response.

There is a third path – to decide to intervene even in uncertain cases, so that certainty takes place at the time of intervention. This probably works in the case of uncommon and isolated intrusion detection, but may call for too much resource in conditions where many false positives are present

There are, then, multiple paths to certainty. One is to wait for the network to form an impression; the longer the intruder moves over the sensors, the more the network learns. A second path is to look at the intruder, using human pattern recognition capabilities. A third path might intervene, also focusing human attention.

Why use the later methods in preference to the automated method? The human may perform better in some situations. A human will be able to tell the difference between, say, an animal and a human quickly. In addition, we imagine a common problem in sensor network security may be the labeling of a friend as an intruder; humans will recognize members of their team.

4.3 States of the system

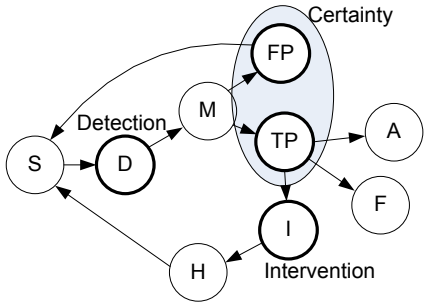

Figure 8. The system, starting at S, moves through a set of states.

We can now make some observations about the overall flow of the human-sensor network system. Figure 8 shows the states through which a protective sensor system will move. After the detection state (D), the system will move into some form of monitoring state (M), which might involve human attention. The monitoring state will eventually transfer to a state of certainty – the alarm was either a false positive (FP) or a true positive (TP). If there is really an intruder, the intruder may flee (F). If the defender responds well, then the system may reach a state of intervention (I). However, this state is likely to take a lot of time; resources in the system won't be freed up for a while; once the intruder is safely held (H) then the system returns to the start state. If the defender cannot intercept the intruder then the intruder may succeed in taking the system to a state of attack (A).

The overall capacity of the system can be measured in the number of simultaneous intruders the system can intercept for a given rate of false positives.

The rate of false positives drives human attention. This is because each detection event may create a task for a person.

Human attention is not fully understood; there appear to be multiple types of attention, and, while it is clear there are capacity limitations, the hardness of these limitations is still a topic of investigation [19, 27]. For the time being, we assume that there is some limit per person to the number of possible intruders that could be monitored, say, using video cameras.

As attacks are by their nature unusual, it is fair to expect that the conditional probability of an intruder, given a detection event is low. Then, the corona distance from the place at which detection occurs to the place at which certainty occurs, presuming humans decide, will be

related to the number of simultaneous detections that occur. If, for example, 12 detections occur at the periphery of the network, and a human can process 2 at a time in a one-minute period, in a situation with 3 humans, the time to certainty will be 2 minutes, as a second set of 6 false positives will queue while the first set are being considered. This 2 minutes can be converted to a distance in figure 8 by making assumptions about how fast an intruder can move.

4.4 Simultaneous intrusion

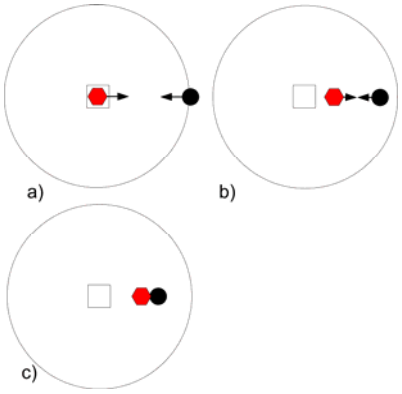

Figure 9. A defender in black intercepts an intruder.

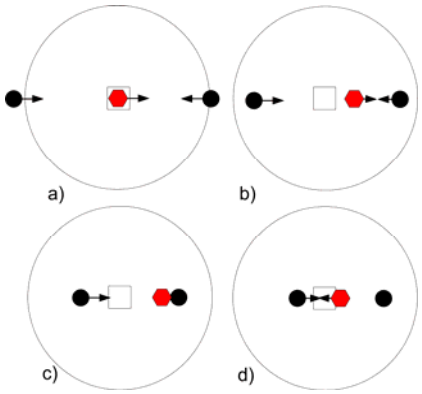

Figure 10. A defender intercepts a possible intruder, while another possible intruder attacks from the other side. The defender needs to double back.

While we have considered multiple false positives, here we consider the possibility of a coordinated attack. If an intrusion is detected, a natural response would be to intercept the intruder by going to meet them, as in figure 9. However, this may increase vulnerability to a second attack, as shown in figure 10.

The geometry of the situation works in the intruder's favor if the intruder outnumbers the defender and can coordinate an attack. In addition, if the defender does not have a way of remotely clearing false positives, then an intruder may watch for a response to detection on the opposite side, and then move in. Even without superior force, an intruder can use false positives as a method to dilute attention and response.

From this analysis we are led to consider how the network itself might play a role in response.

5. In-network response

Having seen the difficulties of maintaining enough capacity to intervene as a result of detecting intruders, we explore an alternative.

Let us imagine that we have a sensor network with some kind of actuator capabilities that are capable of restricting the movement of an intruder. Such an architecture might be designed so that it (1) is sufficiently forgiving that, say, children chasing a ball that was tossed into the system are not harmed, and (2) ensures that any *bona-fide* intruder is restricted in some form or another before human response is dispatched to the scene.

The former case can be easily handled by some form of an *invisible fence* (that should also keep out stray animals) involving, perhaps, ultra-sound barriers that can be activated as a result of the sensor network reporting an incipient intrusion. We note that there is a difference in scale between radio communications and mechanical movement, and, consequently, the system can be designed to react adaptively to the perceived velocity of the intruder.

The latter case is best handled in an application-dependent fashion that takes into account the equipment at hand, the value of the asset, and the local bylaws in the community. Indeed, it is clear that the response to confirmed intrusion into the physical space of an airport should be handled differently from intrusion into a museum, for example.

In either case, the response infrastructure may be pre-deployed as shown in Figure 11. Barriers can be set up that restrict an intruder into the *sector* determined by adjacent corona and adjacent wedge boundaries where the defender can then confront the intruder. We note that this method is especially useful in the case of concerted attacks, where a number of attackers penetrate the system at the same time. The idea is that the attackers coming from different directions are naturally segregated making it easy for the human response team to address them individually.

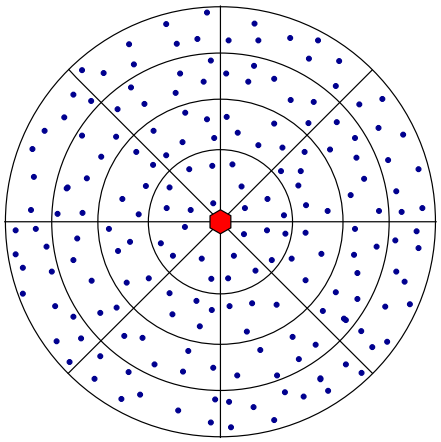

Figure 11. Illustrating the response system

Airport security police on occasion practice a similar tactic; when it seems as though someone has gone through a metal detector with a dangerous object, and the person evades initial search, then parts of the airport are quarantined.

6. Concluding remarks

This paper continues the series titled *Protecting with sensor networks* where the authors set out to evaluate various possibilities for using a wireless sensor network as the basic layer of a hybrid defensive system.

We outlined novel a way of thinking about sensors as providing early warning about a possible intruder; either the network itself or a human monitor may want to wait to gain certainty about an intruder before intervening. The time to certainty, combined with the time to intervene should be faster than the time for the intruder to reach a target. These times can be seen in relationship to the physical nature of the sensor network, as more time to respond can be bought by increasing the extent of the network. An important determining factor in the network design is the anticipated frequency of false positives, as these will in general call for human attention.

The concepts discussed in the paper may be useful to those involved in the design of sensor networks, as it suggests that the physical network and the attention of its operators are intertwined; the physical network is important, and so is the geometry of attention.

Acknowledgements

This research was supported, in part, by the Office of Naval Research, Swampworks, grant #N00014-05-1-00632 and by the National Science Foundation under grant #ITR-032630.

References

[1] I. F. Akyildiz, W. Su, Y. Sankarasubramanian, and E. Cayirci, Wireless sensor networks: A survey, *Computer Networks*, 38(4), 2002, 393-422.

[2] S. Axelsson, The base-rate fallacy and its implications for the difficulty of intrusion detection, *Proc 6th ACM Conference on Computer and Communications Security*, 1999.

[3] C. E. Billings, *Aviation automation: the search for a human centered approach*, Mahwah, N.J.: Lawrence Erlbaum Associates Publishers, 1997.

[4] S. Chakrabarti and A. Strauss, Carnival booth: an algorithm for defeating the computer-assisted passenger screening system, *First Monday*, 7(10), 2002.

[5] J. W. Crandall, M. A. Goodrich, D. R. O. Jr., and C. W. Nielsen, Validating Human-Robot Interaction Schemes in Multi-Tasking Environments, *IEEE Transactions on Systems, Man, and Cybernetics Part A*, 35(4), 2005, 438-449.

[6] D. Culler, D. Estrin and M. Srivastava, Overview of sensor networks, *IEEE Computer*, 37(8), 2004, 41-49.

[7] D. Estrin, D. Culler, K. Pister and G. Sukhatme, Instrumenting the physical world with pervasive networks, *Pervasive Computing*, 1(1), 2002, 59-69.

[8] D. Estrin, R. Govindan, J. Heidemann and S. Kumar, Next century challenges: Scalable coordination in sensor networks, *MOBICOM*, Seattle, WA, August 1999.

[9] R. Iyer and L. Kleinrock, QoS control for sensor networks, *Proc. ICC*, 2003.

[10] S. H. Jacobson, J. E. Kobza, and A. S. Easterling, A detection theoretic approach to modeling aviation security problems using the knapsack problem, *IEE Transactions*, vol. 33, 2001 747-759.

[11] J. M. Kahn, R. H. Katz, and K. S. J. Pister, Next century challenges: Mobile support for Smart Dust, *MOBICOM*, Seattle, WA, August 1999, 271-278.

[12] K. Martinez, J. K. Hart and R. Ong, Environmental sensor networks, *IEEE Computer*, 37(8), 2004, 50-56.

[13] Mitchell, M., An Introduction to Genetic Algorithms, The MIT Press, 1999

[14] National Research Council, *Embedded Everywhere*, National Academy Press, 2001.

[15] S. Olariu, A. Wadaa, L. Wilson and M. Eltoweissy, Wireless sensor networks: leveraging the virtual infrastructure, *IEEE Network*, 18(4), 2004, 51-56.

[16] S. Olariu, J. V. Nickerson, Protecting with Sensor Networks: Perimeters and Axes, *MILCOM* 2005.

[17] A.Wadaa, S. Olariu, L. Wilson, M. Eltoweissy and K. Jones, Training a wireless sensor network, *Mobile Networks and Applications*, 10, 2005, 151-167.

[18] D. R. Olsen and M. A. Goodrich, Metrics for Evaluating Human-Robot Interaction, *Permis*, 2003.

[19] H. E. Pashler, *The psychology of attention*, Cambridge, Mass.: MIT Press, 1998.

[20] P. Saffo, Sensors, the next wave of innovation, *Communications of the ACM*, 40(2), 1997, 93-97.

[21] J. Scholtz, Theory and Evaluation of Human Robot Interactions *Hawaii International Conference on System Sciences*, 2003.

[22] R. Sedgewick and P. Flajolet, *An introduction to algorithm analysis*, Addision-Wesley, 1996.

[23] K. Sohrabi, J. Gao, V. Ailawadhi and G. Pottie, Protocols for self-organization of a wireless sensor network, *IEEE Personal Communications*, 7(5), 2000, 16-27.

[24] K. Sohrabi, W. Merrill, J. Elson, L. Girod, F. Newberg and W. Kaiser, Methods for scalable self-assembly of ad hoc wireless sensor networks, *IEEE Transactions on Mobile Computing*, 3(4), 2004, 317-331.

[25] B. Warneke, M. Last, B. Leibowitz and K. Pister, SmartDust: communicating with a cubic-millimeter computer, *IEEE Computer*, 34(1), 2001, 44-55.

[26] A. D. Wood and J. A. Stankovic, Denial of Service in sensor networks, *IEEE Computer*, 35(10), 2002.

[27] C. D. Wickens and J. G. Hollands, *Engineering psychology and human performance*, 3rd ed. Upper Saddle River, NJ: Prentice Hall, 2000.

[28] J. Yang, C. K. Mohan, K. G. Mehrotra, and P. K. Varshney, A Tool for Belief Updating over Time in Bayesian Networks, *International Conference on Tools with Artificial Intelligence*, 2002.

Proceedings of the 40th Annual Hawaii International Conference on System Sciences (HICSS'07)
0-7695-2755-8/07 $20.00 © 2007 IEEE