SMART PERIMETER SECURITY

Fiber SenSys

SMART PERIMETER SECURITY

WHITE PAPER

BY

Duwayne Anderson

Duwayne Anderson is Chief Technologist at Fiber SenSys and has over 20 years experience in the area of fiber optics. Prior to joining Fiber SenSys Mr. Anderson was Principal Engineer at Phoseon where he designed high-power solid state lighting systems. Mr. Anderson's career also spanned 15 years as a Principal Engineer at Tektronix where he was responsible for the design of fiber-optic test equipment for telecommunications companies. His other industry experience includes several years each at Honeywell and Goodyear Aerospace. Mr. Anderson has 28 patents and is the principal author of a text book on trouble shooting fiber-optic networks.

INTRODUCTION

Perimeter protection represents one of the most important yet difficult post-911 security challenges facing the nation and global economy today. High-profile, high-value assets are often large, with extended circumferences offering a tempting target for those trying to inflict maximum economic and physiological harm. Especially valuable assets often have relatively long perimeters, and these can present serious challenges for designers. Airports, rail stations, large petrochemical plants, government buildings and nuclear power plants are all at risk.

The difficulty in developing optimum perimeter security systems belies the simplicity of listing their three major components:

1. A physical barrier to prevent or delay intrusion
2. Sensors to detect and warn of attempted intrusions
3. Sensors to assess and track attempted or real intrusions

Clearly, sensors are critical to highly secure perimeters, and the two principal challenges for perimeter sensors are:

1. Very high probability of detection for real threats (PD)
2. Very low nuisance alarm rate from non-threats (NAR)

These challenges are generic, and apply to any sensing technology. Their importance derives chiefly from both security and economic concerns. High PD obviously enhances security directly, while systems prone to high NAR inhibit security because they are soon ignored or switched off by frustrated personnel. High NAR also carries an economic penalty because of the excessive costs associated with investigating and clearing nuisance alarms.

In this paper we shall examine the factors influencing PD and NAR and how to control them for maximum benefit. We'll begin with a short summary of fiber sensors, generally considered the best technological solution to perimeter sensing. Then we'll review the definitions of PD and NAR and describe the reasons for using multi-parameter analysis for maximum performance. Next we'll discuss the practical challenges when using multi-parameter systems, culminating in

discussion about smart software that enables sensor systems that learn, thus minimizing training demands on those responsible for perimeter security systems. We'll close with a brief description of commercially available multi-parameter fiber-optic sensor systems and summarize test data illustrative of best-in-class performance.

ADVANTAGES OF OPTICAL FIBER

For many reasons optical fiber is the obvious choice for sensors on long perimeters. Fiber is inexpensive (costing just pennies per foot) and unlike metallic cable and wire sensors it requires virtually no maintenance, making it the least expensive solution for long-range sensors. Fiber is easily available in lengths exceeding 50 km, and in standard cable configurations that are extraordinarily robust and deployable in the most extreme environments. Unlike metallic sensors, fiber is all-dielectric so it's impervious to EMI, making it the only viable perimeter sensor in explosive environments and locations with extreme electromagnetic interference like those found in power distribution facilities and around high-voltage lines. Unlike metallic cable and wire systems fiber has virtually no loss over most perimeters, resulting in excellent linear performance. And unlike metallic cable and wire sensors fiber is unaffected by the damaging electromagnetic fields inherent in lightning strikes. Once installed fiber maintains its sensitivity without loss due to corrosion or other problems that might arise with metallic sensors after prolonged exposure to the elements.

NUISANCE ALARM RATE AND THE PROBABILITY OF DETECTION

When evaluating perimeter security systems the most obvious performance measurand is the probability of detecting an intruder, or PD. This metric, though conceptually simple, is surprisingly complex and nuanced. One measures the PD by instigating, under controlled conditions, a certain number of intrusion attempts (typically at least 10) and measuring the

number of times the perimeter security system alarms on the intrusion. It's essential that the conditions be carefully controlled because the PD is often a strong function of specific circumstances. Stealthy intrusions made with padded ladders typically have lower PD than robust intrusions made in hast. Depending on the way the sensor operates the PD may also be a factor of environmental conditions. For example, wind may affect the PD if the sensor system adjusts parameters in the presence of wind in order to avoid nuisance alarms. Thus PD must always carry qualifiers that describe the type of perimeter (fence, buried, type of fence, etc.), the weight of the intruder, tools and techniques used in the intrusion simulation, and environmental conditions at the time of the simulations.

One of the biggest problems for long perimeters is the nuisance alarm rate, or NAR. The NAR is expressed as the number of nuisance alarms in a given amount of time, normalized to a particular length of perimeter. For example, the NAR for a particular installation might be 1 nuisance alarm per kilometer of perimeter, per month. Mathematically we write

$$NAR = \frac{Number_{of\ nuisance\ alarms}}{Length_{Perimeter} \cdot Time}$$

Notice that NAR is a rate, and not a probability. The NAR is, however, a function of the probability that something ***capable*** of causing a false alarm will actually ***do it***. We call this the probability of nuisance, or PN. The mathematical expression is

$$NAR = PN \cdot \frac{Number_{of\ possible\ sources\ of\ nuisance\ alarms}}{Length_{Perimeter} \cdot Time}$$

There are a couple of ways to determine the number of possible sources of nuisance alarms. One way is to walk along the perimeter, counting all the sources. Another, more practical way, is to determine the density of the possible sources of nuisance alarms and then multiply that density by the length of the perimeter.

Let's pause a moment and look at a hypothetical example. Suppose we have a 20-km perimeter that averages 100 rabbits per km during a 24-hour period, giving us an average of 20·100 = 2,000 rabbits every day. Next, suppose that a single rabbit carries with it a 2% chance of causing a nuisance alarm. The NAR is

$$NAR = \frac{0.02 \cdot alarm}{rabbits} \cdot \frac{2{,}000 \cdot rabbits}{20km \cdot 24hr} = \frac{14 \cdot alarms}{km \cdot week}$$

This is a very high number, especially so for our hypothetical perimeter which, at 20 km, can expect a total of 280 nuisance alarms each week. The purpose of this example is to illustrate how a relatively modest PN (2%) can result in a very unacceptable NAR. That's because, in practice, there are many different sources of nuisance alarms along a typical perimeter, and so the NAR is a very complicated entity that depends on the type of perimeter, its location, animal and plant life, weather, etc. In practice the NAR is determined heuristically and has considerable uncertainty due to many conditions that are practically uncontrolled but still significant, including the type of perimeter installation (type of fence, or type of burial/soil) as well as local weather, plants growing adjacent to the fence, wild or domestic animals that might brush up against the perimeter, traffic, etc.

From the example above it's clear that the reason for expressing NAR with inverse units of time·distance is that the probability of encountering something that might cause a nuisance alarm (whether it's a squirrel or dust devil) is statistically proportional to the elapsed time and the length of the perimeter. If you wait one month you're more likely to encounter a nuisance alarm than if you wait 10 minutes. Similarly, if you have a perimeter that's 50 km long you'll be more likely to encounter a nuisance alarm during the next 10 minutes than if you have a perimeter that's only 100 meters long.

This simple fact means that sensors designed for long distances must have extraordinarily low PN because the NAR is proportional to PN multiplied by the length of the perimeter. Thus, even very low PN can result in unacceptably high NAR when multiplied by lengths of tens of km or more.

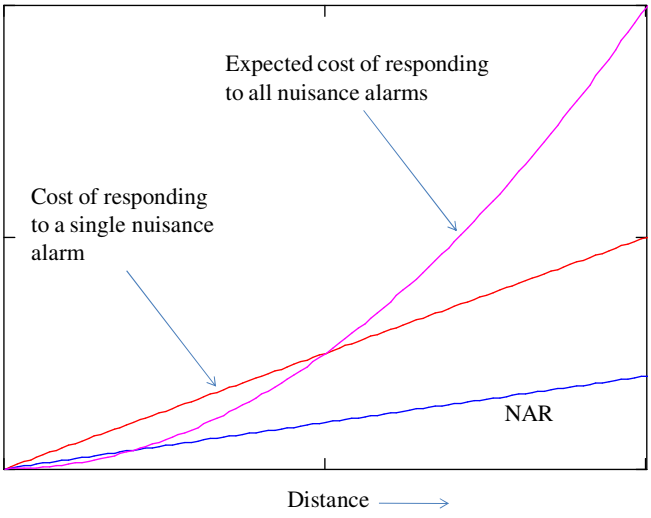

For long perimeters NAR is more than a nuisance, it's an economic factor that can play heavily into total operating costs. For a perimeter security system to be functional it must provide actionable data that is, in fact, acted upon. This means that when the system sounds an alarm it must be

investigated. This investigation can range from directing a camera to look at the suspected area to dispatching armed personnel for an on-site investigation. In all cases the alarm costs money, but the cost of investigating all nuisance alarms increases much faster than the length of the perimeter. This is because while both the NAR and the cost of responding to a single alarm increase linearly with perimeter length, the mean cost of responding to all nuisance alarms is the product of these two linear functions and increases as a second-order polynomial. Thus the cost per km (attributed to personnel investigating nuisance alarms) is greater for long perimeters than for shorter ones, and even systems with moderate NAR can become prohibitively expensive when deployed along moderate to long perimeters. For such perimeters economic considerations demand that the NAR be extraordinarily low.

DESIGNING FOR LOW NUISANCE ALARM RATE

When a perimeter security system experiences excess NAR the problem tends to be systemic. Although minor improvements can sometimes be made by tuning the system for lower PD, usually the problem must be solved by system design considerations. Given a physical sensor with well-designed hardware and good signal-to-noise ratio (SNR) a very effective way to improve the NAR is to design the software components of the system with multi-parametric algorithms that employ the power of statistical coincidence. Let's consider a simple example.

Suppose we're designing a fiber-optic interferometer that attaches to a fence around a 10 km perimeter. This sensor is designed to detect intruders that try to cut through or climb over the fence. It works because small vibrations in the fence cause tiny stress fluctuations in the fiber that the sensor detects and converts into a proportional voltage that's digitized and analyzed in a digital signal processor.

Initially we try a very simple algorithm that measures the change in light and, when the change gets too big, sounds an alarm. We test our system and find we have a 99.5% probability of detecting a person attempting to breach the perimeter. But when we test the NAR, by letting the system run continuously for several days, we find that our sensor generates about 10 nuisance alarms each day. This is a very high figure so we investigate and find the system alarming on all

sorts of non-lethal objects like wind vibrating the fence, birds alighting on it, and small animals browsing along the perimeter.

To reduce the NAR we begin by making adjustments to our algorithm. Instead of alarming on a single instance of large signal we also consider the length of time the signal is above a pre-set level, requiring the signal to be both large and prolonged before we generate an alarm. A bird alighting on the fence won't produce a signal as prolonged as that created by a person trying to climb the fence.

We test our system again and find the NAR greatly reduced. Not only does the new algorithm improve the NAR for birds, it also reduces it for wind. But the changes don't eliminate the NAR completely, so we investigate further and find that human intruders create a distinct spectral signature when they try to climb over or cut through the fence. When we look at the spectral distribution from events caused by human intruders we see that most of the signal energy lies well above 300 Hz, while the spectral components from nuisance events caused by wind and small animals tend to cluster below this level. So we adjust our algorithms again, this time examining the following three parameters:

- Signal level
- Signal duration
- Signal spectral content

We might continue with this process, adding more and more parameters to our algorithms. Not surprisingly we find that the more parameters we analyze the more accurately we can discriminate between what we want to detect and the extraneous signals that are of no interest. Although the scenario we've just described is fictional the lesson is real, namely that (all other things being equal) sensors employing multiple parameters generally have better NAR than systems using less sophisticated forms of analysis.

Mathematically there's a simple way to understand why this happens. Suppose we have "n" parameters. We label the first one p_1, the second p_2 and so on. We might use a shorthand notation where we label an arbitrary parameter p_i where i is a counting index ranging from 1 to n. For each parameter there is a probability of detecting the intruder using just that parameter, and a probability of getting a nuisance alarm using just that parameter. We'll call these PD_i and PN_i where, as before, the "i" is a counting index corresponding to the appropriate parameter. That is, PD_1 and PN_1 are the probability of detection and probability of nuisance alarm (respectively) when analyzing data using only parameter p_1, and so-forth.

Now the reason multi-parameter analysis works so well is that the PD for all well-chosen parameters tends to be relatively high, perhaps in excess of 0.995. Meanwhile the PN for the

various parameters tends to be in the range of about 0.1. When we use all the parameters, and logically conjoin them with AND statements, the composite probability of detection equals all the probabilities multiplied together:

$$PD_{Composite} = \prod_{i=1}^{i=n} PD_i$$

In the hypothetical example we just considered the probability of detection was 0.995 for each parameter, so the composite probability (conjoining all the parameters with logical "and") is $0.995^3 = 0.985$. Similarly the composite PN is:

$$PN_{composite} = \prod_{i=1}^{i=n} PN_i$$

In our hypothetical example the composite PN is $0.1^3 = 0.001$ (0.1%). Notice that the composite PN is 100 times smaller than the individual PN_i while the composite PD is almost identical to the individual PD_i. This is exactly the solution we are looking for: high PD, and low PN.

This hypothetical example shows that multi-parameter analysis works well because high probabilities (near 1) tend to remain un-changed when raised to a power, while small numbers get much smaller really fast. The graph to the right plots the function $f^n(x)$ for several values of n (the vertical axis uses a logarithmic scale). As you can see, as we increase n (add more parameters) the NAR (near the lower left of the graph) rapidly diminishes while the PD (near the upper right) stays mostly unchanged. This is exactly what we need; high PD and low PN (and low NAR). Intuitively we can see that multi-parameter algorithms work so well because circumstances become increasingly unlikely, as we add more and more parameters, that a nuisance signal will look just like the intrusion signals that generate alarms.

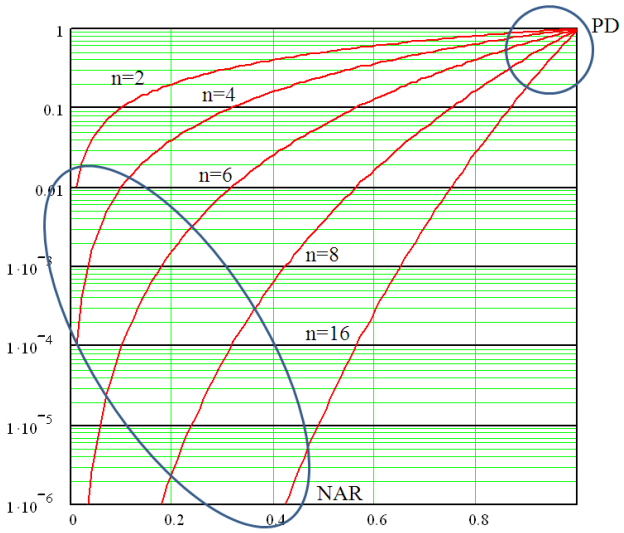

TUNING

Now that we've found a way to get the NAR low, and the PD high, we might be tempted to think the problem is solved, but multi-parameter analysis is only part of the solution. While multi-parameter analysis is a powerful analytic tool it can be difficult to use, and in some cases nearly impossible for a human operator to optimize. And if not used correctly, multi-parameter analysis, like many other tools, can lead to problems as severe as those it was designed to fix.

Let's return to our previous example. When we began with that example we had one parameter. Let's suppose we divided that parameter into 10 settings, in which case we would have had 10 ways to tune the sensor. Then we added the second parameter and in doing so we increased our options. Let's suppose the second parameter also had ten settings, in which case the system (with two parameters) would have 10·10=100 possible tuning configurations. Now suppose that when we added the third parameter it also had 10 settings. With the third parameter, along with the other two, we would have 10·10·10=1,000 possible system configurations.

Now we see a pattern developing and recognize that the total number of system configurations is equal to the product of the number of settings for each of the parameters:

$$Number\ of\ system\ configurations = \prod_i (Number\ of\ settings\ for\ ith\ parameter)$$

It's a conundrum. On one hand we need many parameters and settings to achieve a low NAR and good PD; many parameters provide optimum differentiation between real intrusions and nuisances. But as we add more and more parameters the resulting system has so many possible configurations we risk making it practically impossible to optimally tune. And if we can't optimally tune the system we won't get the best PD/NAR, even though we have lots of parameters. And since optimum PD/NAR was the reason for adding the parameters in the first place we've come full circle in our search for a perimeter security system with low NAR and good PD.

Actually, the situation isn't quite this dire because in real systems we aren't just setting parameters randomly. We usually have an idea of what makes a real intrusion signal unique, and we can do experiments and see how various parameters affect the system's ability to discriminate among real and nuisance alarms. We can formulate general rules and teach these rules to the professionals that tune the systems. So through effective study and training we can still get most (if not all) of the benefit from having many tuning parameters.

LEVERAGING SMART SOFTWARE

Can we find a better solution? Is there an easier way to take full advantage of highly tunable, multi-parameter systems without placing such onerous training requirements on the people that install and maintain the perimeter security systems? The answer is yes, and the solution is a sensor system using multi-parameter analysis where the *system* chooses the parameters based on non-mathematical, non-technical inputs from the user. This sophisticated system is characterized by many adjectives, including "smart," "intelligent," "adaptive," and "optimized."

A smart security perimeter system is one that can be taught, and learns from experience. The training process is similar to teaching a child to read. The user puts the system into its learning mode and then simulates various intrusion attempts. The system stores the signatures from these intrusions and then simulates how it would have responded had it been configured with hundreds or even thousands of possible parametric settings. Using high-speed computing and efficient learning algorithms the system quickly finds the optimum parameters that produce the lowest NAR and the best PD.

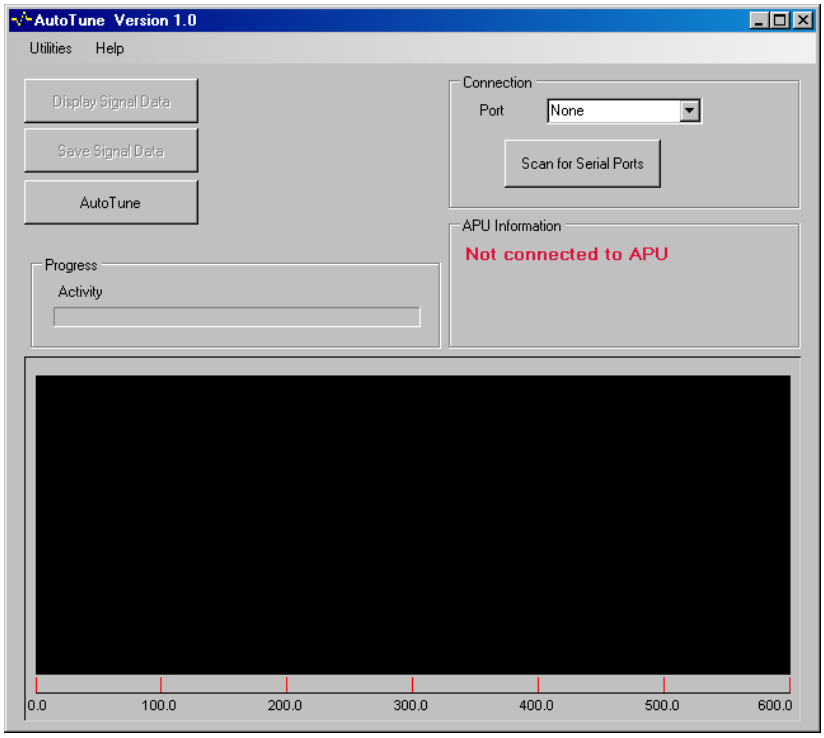

Fiber SenSys has developed software that provides this level of optimum tuning with their fiber-optic perimeter sensors. These sensors provide up to 9 user-selectable parameters that, when optimally selected, provide near-optimum PD and NAR. Software under the brand name AutoTune™ allows the user to teach the system through a process of simulating various sources of nuisance alarms and real intrusions. Just as we've described above, AutoTune™ uses powerful search algorithms and simulations to rapidly test and evaluate how the system would perform when tuned using hundreds of different possible parametric configurations, and then it picks the best configuration, the configuration that optimizes both PD and NAR.

Since complexity is the problem being solved simplicity is a key requirement for smart tuning software. The user interface needs to be clean and easy to understand. Visibility is another requirement. For example, the Fiber SenSys solution displays the optimum tuning parameters, as determined by the software, and gives the user the option of saving them, trying again, or making modifications.

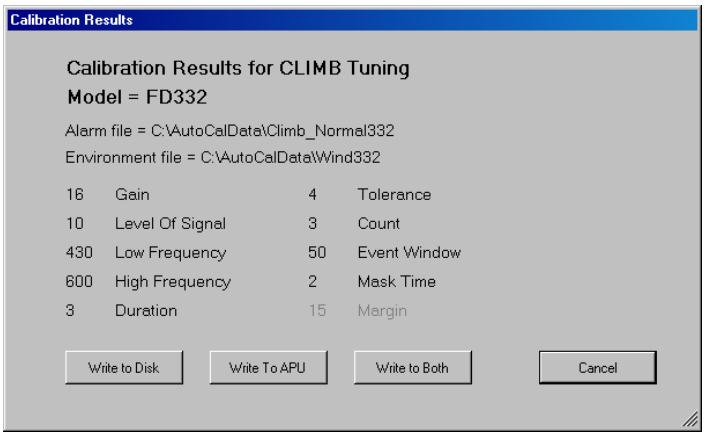

One of the major challenges in designing such smart perimeter security solutions is testing and verifying performance. As with other smart computing structures the optimum test is often a competition between the software and a grand champion. The designers of Deep Blue tested their chess software against Garry Kasparov and the designers of AutoTune™ tested their software against expert technical installers.

In this procedure two identical fiber-optic perimeter security systems are installed in parallel and in close proximity on a representative perimeter (a fence, for example) and labeled system A and system B. Then AutoTune™ is used to select the optimum parameters for system A while the expert tunes system B, after which the experimenters conduct a series of controlled intrusions while recording all alarms generated from the two systems. When these experiments are completed they leave the systems unmolested for several days to test their susceptibility to nuisance alarms. Next the experimenters repeat the process, but this time they use AutoTune™ to select the parameters for system B, and have the expert tune system A. In this way the experimenters can recognize and remove any biases due to slight differences in the installations of the two systems.

In practice the scientists and engineers that developed the AutoTune™ software conducted many such tests at various sites and using multiple experts. Such careful testing is required because of the site-specific and time-specific nature of nuisance alarms and the subjective manner in which different individuals simulate intentional intrusions. The results of these tests have been

impressive. In all cases AutoTune™ either matched or beat the performance (gauged in terms of PD and NAR) of the expert installer/tuner.

As perimeter security systems become more adaptable, capable, and intelligent smart software like AutoTune™ will increasingly play an important role. These systems offer the hope that smart perimeter security systems will relieve personnel from the onerous task of managing large amounts of data so they can focus on making accurate and timely high-level decisions while, at the same time improving performance, making perimeter security systems more accurate, reliable, and less expensive to maintain.