Discrimination Performance Requirements for Ballistic Missile Defense

Stephen D. Weiner and Sol M. Rocklin

A missile defense system must be able to deal with an attack containing decoys in addition to warheads. If the defense system does not have enough interceptors to shoot at all the incoming objects, it must be able to discriminate between decoys and warheads. This discrimination process is not perfect and results in two types of errors: leakage (not shooting at warheads) and false alarms (shooting at decoys). This article describes a methodology for analyzing the consequences of these discrimination errors and determining how well discrimination must perform in a variety of defense scenarios. The analysis focuses on game-theoretic solutions in which the defense can achieve its overall objective of surviving the attack regardless of the tactics used by the offense.

Since the early 1960s, Lincoln Laboratory has worked on many programs involving almost all aspects of ballistic missile defense (BMD). Initial BMD research focused on the problem of ballistic missile interception as expressed by the question "Can a bullet hit another bullet?" Researchers soon realized that an even more serious problem was "Can we find and identify the bullet that needs to be hit?" The process of identifying the target to be intercepted is generally called *discrimination*. Discrimination is a remote-sensing operation wherein sensor measurements of target observables are interpreted to identify the target as threatening or nonthreatening.

Traditionally, we have thought of threatening objects as reentry vehicles (RV) carrying lethal warheads and nonthreatening objects as decoys that are deliberately flown to confuse the defense. Decoys degrade the defense in two related ways. If the defense mistakenly thinks a decoy is an RV and shoots at it, the defense might exhaust its supply of interceptors prematurely. On the other hand, if the defense mistakenly thinks an RV is a decoy and doesn't shoot at it, the RV will penetrate the defense and the consequences will be more immediate and more serious. Mistaking a decoy for an RV is called a *false alarm*, while mistaking an RV for a decoy is called *leakage*. In addition to deliberate decoys, an attack typically contains incidental debris and deployment hardware that must also be discriminated from RVs. As we show below, for a given level of discrimination performance the defense can change decision thresholds to trade off leakage and false alarms to balance the consequences of these two types of errors. Furthermore, we quantify how improving discrimination performance improves overall BMD system performance.

The work reported here was started in 1976 to answer questions like "How good does discrimination have to be?" and "What options are available if discrimination is not good enough?" Over the past two decades the focus of the national BMD research program has changed significantly but the need for discrimination and the need to assess discrimination-performance requirements continue to be integral to all candidate missile defense systems.

Two broad classes of BMD systems have been analyzed. The first class is defense of high-value targets such as cities or defense radars for which low leakage must be achieved. This defense often consists of two or more layers that each thin out the incoming attack. The second class is defense of multiple military targets such as missile silos. Higher leakage is acceptable for these targets, but the cost of the defense must be kept low. The analyses presented in this article are drawn from both of these BMD classes. As we analyze the various types of defense systems for these two classes, we emphasize the features they have in common as well as their important differences.

Modeling of BMD Systems

Figure 1 shows the three major functions a BMD system must perform; these functions are *search and detection*, *discrimination*, and *interception*. The search-and-detection stage scans the region of space through which attacking objects must pass and detects anything that might be an RV. Objects—including RVs and decoys—that are not detected at this stage travel unimpeded to their destination, and the defense is completely unaware of their existence (until it is too late). Detected objects then pass through the discrimination stage in which the objects are identified as RVs or decoys. In the interception stage, all objects that are classified as RVs will be fired at if a sufficient number of interceptors are available. Interceptors will not be fired at RVs that are misclassified by the discrimination stage as decoys. RVs that are not fired at, or are fired at but not killed, will continue to their destination. As illustrated in Figure 1, each of the functions of a BMD system might allow RVs to leak through the system; the total leakage may well be higher than the defense can bear.

To reduce the possibility of leakage, the defense can employ a parallel redundancy or series redundancy, as shown in Figure 2. In parallel redundancy, two or more elements perform the same function and RVs will leak only if all the parallel elements fail. In series redundancy, the sequence of functions shown in Figure 1 is repeated, frequently in another stage of the missile trajectory, and the defense gets another chance to detect, identify, and intercept those RVs which leaked through the initial defense layer. In a two-layer system, for example, if the layers operate independently and each layer has 10% leakage, then the leakage through two layers will be only 1%. Several of the examples in later sections show how we can produce more effective discrimination by combining discriminants from different sensors either within the same layer or in subsequent layers.

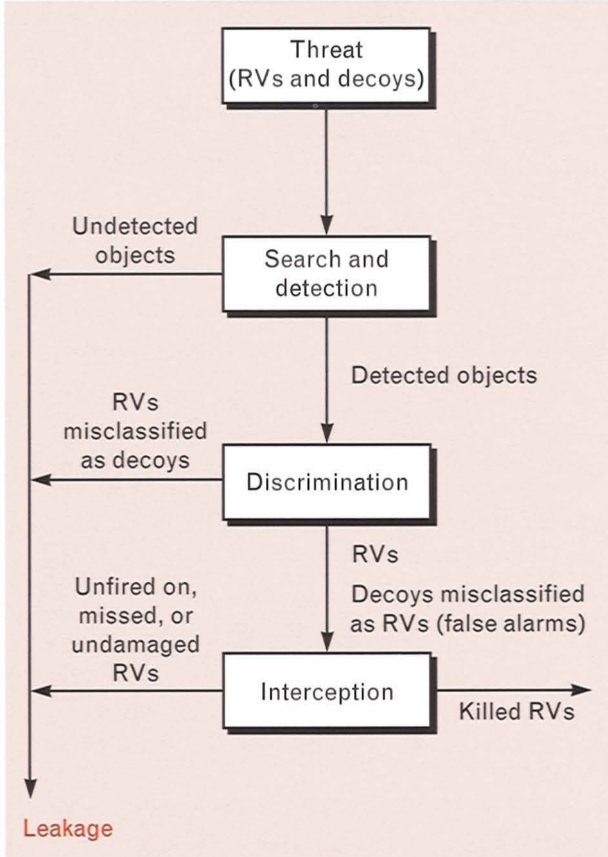

FIGURE 1. Functional flow in a single-layer ballistic missile defense (BMD) system. The major functions are shown along with the paths taken by RVs and decoys through these functions.

Search and Detection

The search-and-detection process involves distinguishing between a measurement containing a target signal plus noise and a measurement containing only noise. This process is well understood because the statistical characteristics of both the signal and the noise can be modeled in detail (see the sidebar entitled "Detection of Signals in Noise"). The detection process is a balance between leakage (failure to detect) and false-alarm rate (noise exceeding detection threshold). By increasing the sensor signal-to-noise level, the defense can reduce the detection leakage while maintaining acceptable false-alarm rates. This trade-off represents a cost trade because increasing the signal-to-noise ratio (SNR) requires building a better sensor or operating the detection system at shorter ranges and deploying more sensors to maintain the needed coverage.

In modeling the search-and-detection process, many parameters affect the performance as well as the cost needed to achieve this performance. These parameters include search time, range, field of view, type of sensor used, target signal expected, and required leakage level. Similar trade-offs occur when we consider the discrimination function.

Discrimination

The discrimination function is similar to the search-and-detect function in that it involves deciding between two possibilities—namely, is the detected object an RV or non-RV. For several reasons, however, the problem is more difficult than distinguishing between "signal plus noise" and "noise only." First, the physical connection between the actual targets and the sensor observables used for discrimination might not be fully understood. Second, many different types of RVs and non-RVs must be distinguished. Finally, for the case of deliberate decoys, the attacker is trying to match the observables of the RVs and the decoys. This matching can be done by making the decoys look like RVs (simulation), by making the RVs look like decoys (antisimulation), or by deliberately varying or obscuring the observables of both RVs and decoys (masking or confusion).

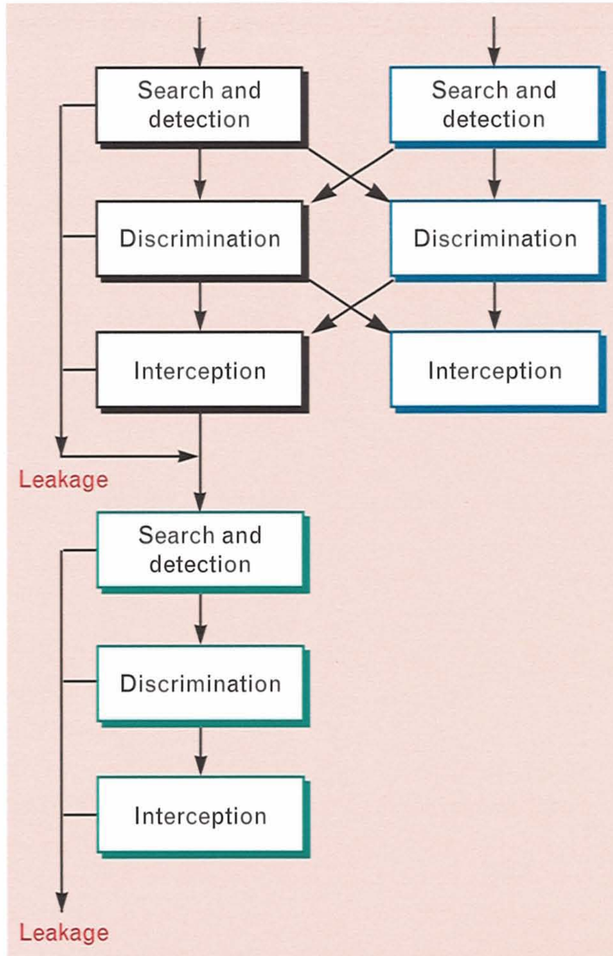

FIGURE 2. Functional flow in a parallel or multiple-layer BMD system. The blue boxes illustrate parallel redundancy within a layer while the green boxes illustrate an additional defense layer in series with the first layer.

Figure 3 illustrates the possibilities for error that can occur in the discrimination process. The RVs and the decoys usually differ in mass, but the sensors cannot measure this difference directly. Instead they measure parameters that can be combined to yield an estimate of the mass or an estimate of the likelihood that an object is an RV. Because the connection between estimated mass and true mass (for example, through kinematic variables) is not unique, even a perfect sensor with no measurement errors produces a spread in the values of estimated mass, as illustrated in Figure 3(a). For the case of good decoys, as shown in Figure 3(b), this spread can result in an overlap between RVs and decoys, which causes decision errors even with a perfect sensor. For actual sensors, the measurement errors further broaden the distributions and increase the overlap between RVs and decoys. For crude sensors, the measurement errors can dominate the inherent target spread and result in overlap even for poor decoys, as shown in Figure 3(c), as well as for good decoys, as shown in Figure 3(d). The overlap can be reduced by using more accurate sensors, as shown in Figures 3(e) and 3(f), but the measurement errors can never be less than the inherent target spreads shown in Figures 3(a) and 3(b). Thus there is a point of diminishing returns in improving sensor quality.

For the purpose of assessing discrimination performance requirements, we can construct a simplified mathematical model of the discrimination process. Figure 4 illustrates the essence of this model. Figure

DETECTION OF SIGNALS IN NOISE

A CLASSIC PROBLEM for sensors is the detection of a signal in noise. The sensor must distinguish between the two possibilities: signal plus noise and noise alone. This problem is completely analogous to the discrimination problem of distinguishing RVs from decoys. The sensor sets a decision threshold and is subject to leakage (failure to detect) and false alarms (detections of noise). By varying the threshold, the sensor can trade off the two types of errors. By increasing the SNR the sensor can simultaneously reduce both types of errors. Curves of P_L versus P_{FA} for various values of SNR are called *operating curves*. They are analogous to the curves shown in Figure 4 with the k factor playing the role of SNR.

For different types of sensors and targets, there are different statistical distributions for both signals and noise. Passive sensors (radio receivers and optical sensors) measure only amplitude, while active sensors (radars and ladars) measure both amplitude and phase. Signals may be constant or fluctuating, while noise may have Gaussian or other statistical distributions, depending on the source of the noise.

We consider the simple case of a radar detecting a constant signal in a background of Gaussian noise. Figure A illustrates how detection and false-alarm probabilities are calculated. S is the constant signal taken as a vector along the x axis. N is a random-noise vector whose x and y coordinates represent the in-phase and out-of-phase noise components, which are each Gaussian. The density functions represent the probability distributions for noise alone and signal plus noise. A threshold on received power is a circle centered at the origin. The false-alarm probability P_{FA} is the integral of the noise distribution outside the threshold circle; it is a function only of the threshold-to-noise ratio. The leakage probability P_L is the integral of the signal-plus-noise distribution inside the threshold circle; it depends on both the signal-to-threshold ratio and the SNR. The resulting operating curves can be found in any radar systems book; a sample is given in Figure B. We see that for fixed SNR, we can trade off leakage probability and false-alarm probability by varying the detection threshold. To reduce both types of error simultaneously, we must increase the SNR. This is exactly the role played by the k factor in the article.

Many similarities exist between classical detection analysis and the k factor approach to discrimination analysis, but there are also important differences. In detection

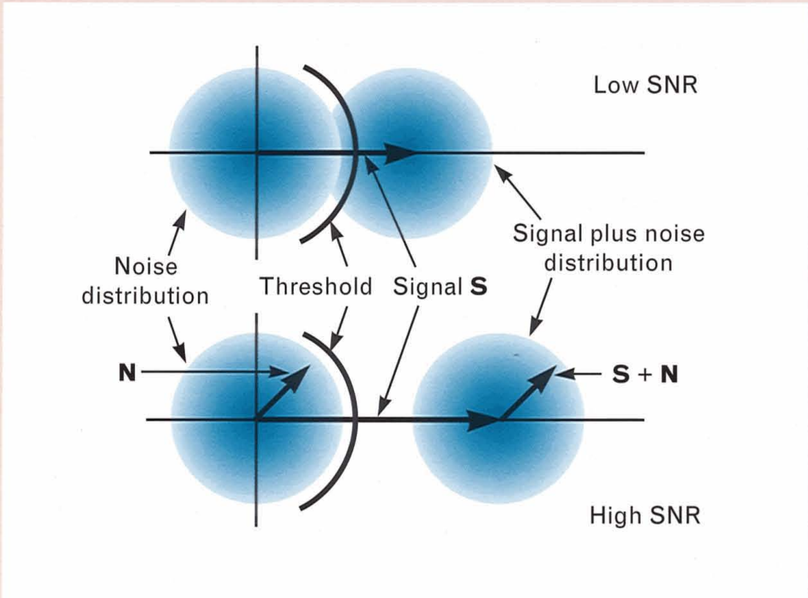

FIGURE A. Signal-vector and noise-vector probability distributions. These two-dimensional vectors have amplitude and phase; the threshold is set on amplitude alone.

we often require and achieve false-alarm probabilities of 10^{-8} or even lower, while we would never need or expect this performance in rejecting decoys. For detection in noise, the sensor must examine millions of resolution elements, which would cause it to become overwhelmed if P_{FA} were too high. Because the properties of Gaussian noise (and many other types of noise) are well understood, we can be confident of reducing P_{FA} by raising the detection threshold. For decoys and RVs the connection between physical characteristics and sensor observables is not understood well enough nor is repeatable enough for us even to hope to achieve values of P_{FA} as small as 10^{-8}. Fortunately, we don't expect to face an attack with millions of decoys so we don't need such low values of P_{FA}.

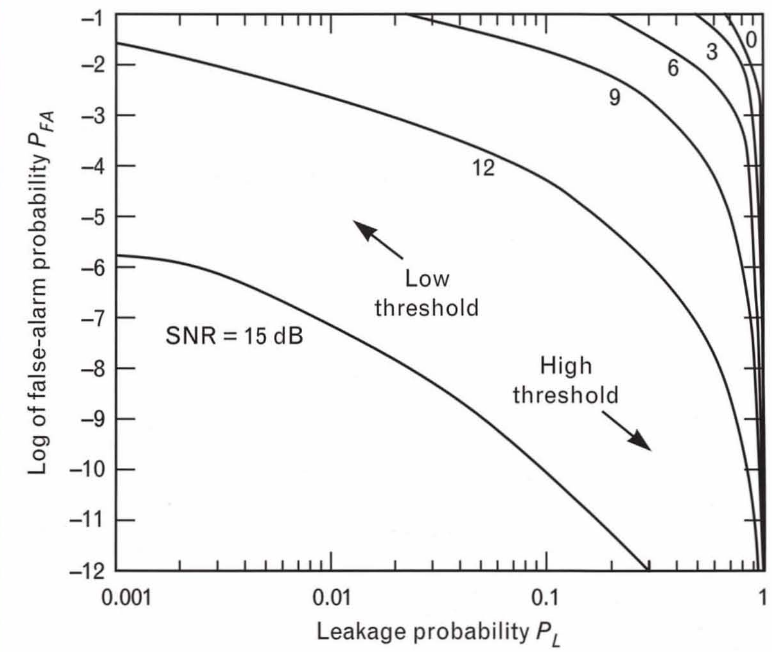

FIGURE B. Operating curves for detection of a constant signal in Gaussian noise. This figure shows how leakage and false alarms can be traded off as a function of signal-to-noise ratio.

4(a) shows typical observable distributions for two populations of targets, represented by probability density functions $f_R(x)$ and $f_D(x)$, for RVs and decoys, respectively. The variable x represents the measured (or computed) value of the observable (or combination of observables); throughout this article we colloquially say that the larger the observable x is, the more RV-like the object is. The process of discrimination amounts to setting a threshold τ and classifying every object whose observable is larger than τ as an RV. As illustrated in the figure, some decoys exceed τ and they are mistakenly classified as RVs (i.e., they are false alarms). More serious, however, are the RVs whose observable lies below τ, for these objects are misclassified as decoys (i.e., they are discrimination leakers) and ignored by the defense. If we lower the threshold τ by moving it to the left, then we can decrease the number of leaking RVs, but only at the price of creating more false alarms. These false alarms will ultimately cause us to waste scarce interceptors.

Thus there is an inherent trade-off between false alarms and leakage. This trade-off can be illustrated by plotting the probability of leakage P_L (the area under $f_R(x)$ to the left of the threshold t) versus the probability of false alarm P_{FA} (the area under $f_D(x)$ to the right of τ). For different values of τ we can generate a P_L-P_{FA} operating curve, as illustrated in Figure 4(b). This figure shows that anytime we want to reduce the rate of discrimination leakage we pay a price in increased false-alarm rate; the two types of errors are inextricably linked through the P_L-P_{FA} operating curve.

In general, decoy density curves can be multimodal if the attack contains a mixture of decoy types. Under these more general circumstances the simple discrimination model illustrated in Figure 4(a) with a single moving threshold might appear to be an oversimplifi-

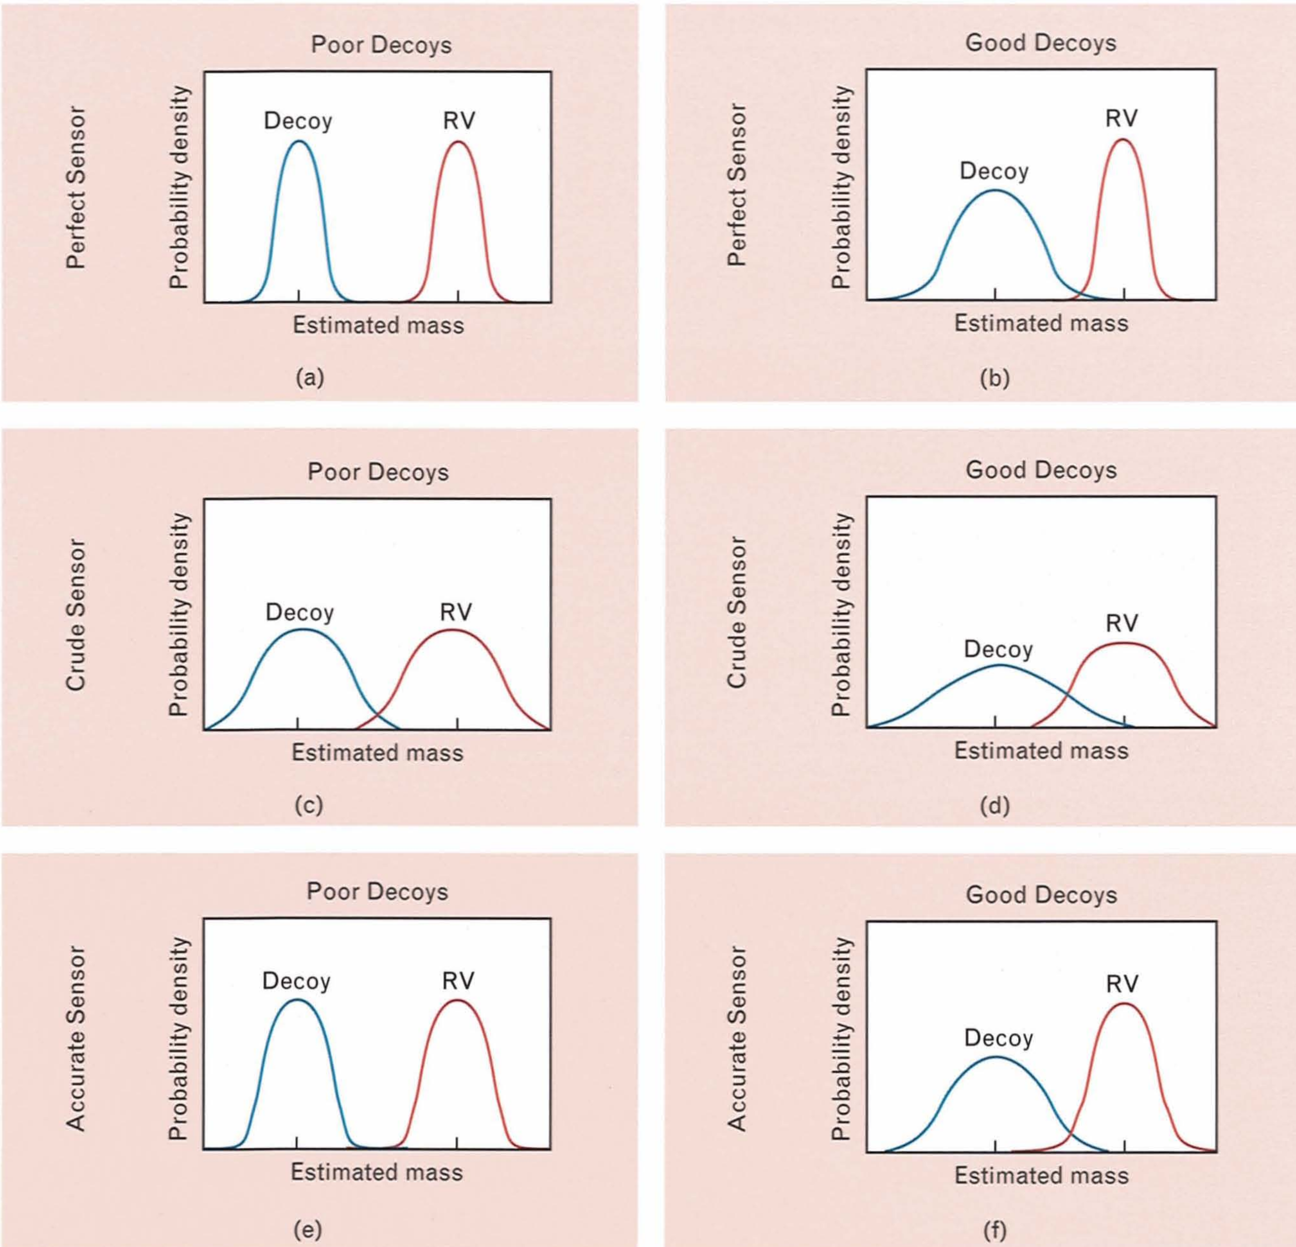

FIGURE 3. Discriminant probability densities for different attack situations. The spread in the probability curves arises from measurement errors and inherent target variability. (a) With a perfect sensor, poor decoys are well separated from RVs while (b) good decoys are closer to RVs and may have a larger inherent spread in their observables. Sensors are not perfect, however, and they typically introduce measurement errors that significantly broaden the probability curves. (c) With a crude sensor, even poor decoys overlap with RVs while (d) good decoys show substantial overlap with RVs. With accurate sensors, as in (e) and (f), the discrimination performance improves up to the limit of a perfect sensor.

cation. The appendix entitled "An Optimal One-Dimensional Representation for a Sensor" shows that the observable space (the *x* axis in Figures 3 and 4) can be transformed into a space characterized by the relative probability or likelihood that the target is an RV or a decoy. In this transformed space, the resulting density functions are as shown in Figure 4(a) with a single optimum decision threshold. Thus, throughout this article, for binary discrimination (in which targets are grouped into two classes—RVs and decoys) we use the single discrimination-threshold model illustrated in Figure 4.

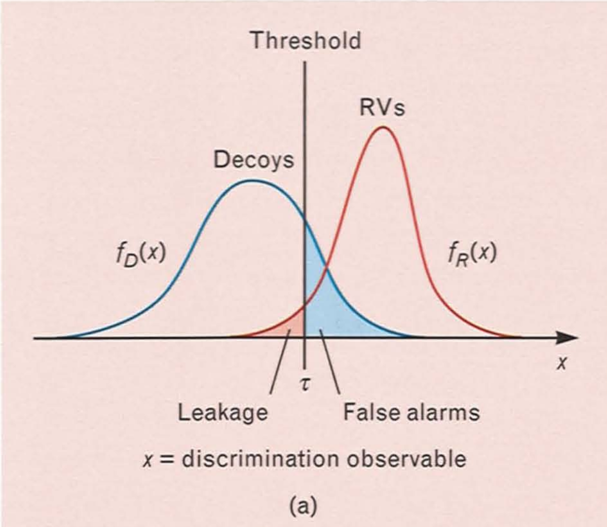

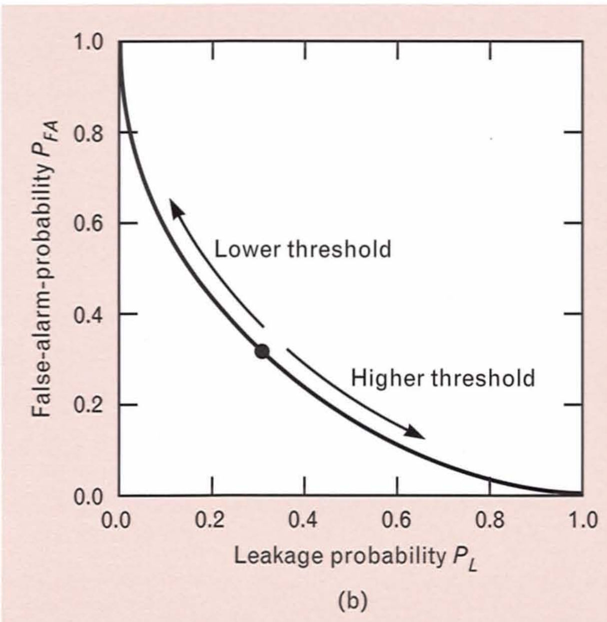

FIGURE 4. The discrimination decision process and the two types of measurement errors. (a) The RV and decoy probability densities with a decision threshold results in *leakage* (i.e., an RV is classified as a decoy) and *false alarms* (i.e., a decoy is classified as an RV). (b) By shifting the decision threshold, we can obtain an operating curve that gives a trade-off between leakage and false-alarm probabilities.

For unimodal single discrimination-threshold models, such as that of Figure 4, discrimination performance can be approximately characterized by a fixed parameter known as the *Bhattacharyya distance.* This number is determined by dividing the distance between the two peaks of the density functions by an appropriate average of the two standard deviations. To investigate many complex missile defense models, our ability to characterize discrimination performance by using a single parameter greatly aids in the analysis, with little loss in accuracy.

Figure 5 shows a simplified model we can use for many classes of discrimination sensors. This model assumes that the observables for both RVs and decoys have Gaussian distributions with equal standard deviations. As illustrated in Figure 5(a), the distance between the peaks of the RV and decoy densities, divided by the (common) standard deviation, is known as the *discrimination factor,* or *k factor.* (Note that the *k* factor is the special case of the Bhattacharyya distance when the standard deviations are equal.) For such Gaussian distributions, the *k* factor completely characterizes the discrimination performance because it unambiguously defines a P_L-P_{FA} operating curve. Figure 5(b) shows some representative operating curves for different *k* factors. The higher the value of *k*, the closer the operating curves hug the P_L-P_{FA} axes. For a *k* factor of zero the operating curve is a straight line; it represents no discrimination ability whatsoever (i.e., random guessing). If the P_L and P_{FA} values are plotted on a probability scale, as shown in Figure 5(c), the resulting operating curves become parallel straight lines.

Interception

We analyze BMD systems by using two general classes of interceptors—perfect and imperfect. A perfect interceptor is defined as one that hits and kills its target with a probability of one. Thus with perfect interceptors any object classified as an RV will be killed, provided a sufficient number of interceptors are available in the inventory. Perfect interceptors are a useful theoretical tool to investigate whether serious performance weaknesses exist in the discrimination function.

In contrast, imperfect interceptors have a nonzero probability of failure to kill. Such a failure can be caused by an inaccurate interception (i.e., the interceptor misses the target) or by insufficient destructive capability. To ameliorate this shortcoming the defense can choose to fire several interceptors at once at each object, or shoot one interceptor, perform a kill assessment, and then shoot at the object again if necessary. The first strategy is preferred if time is short and the

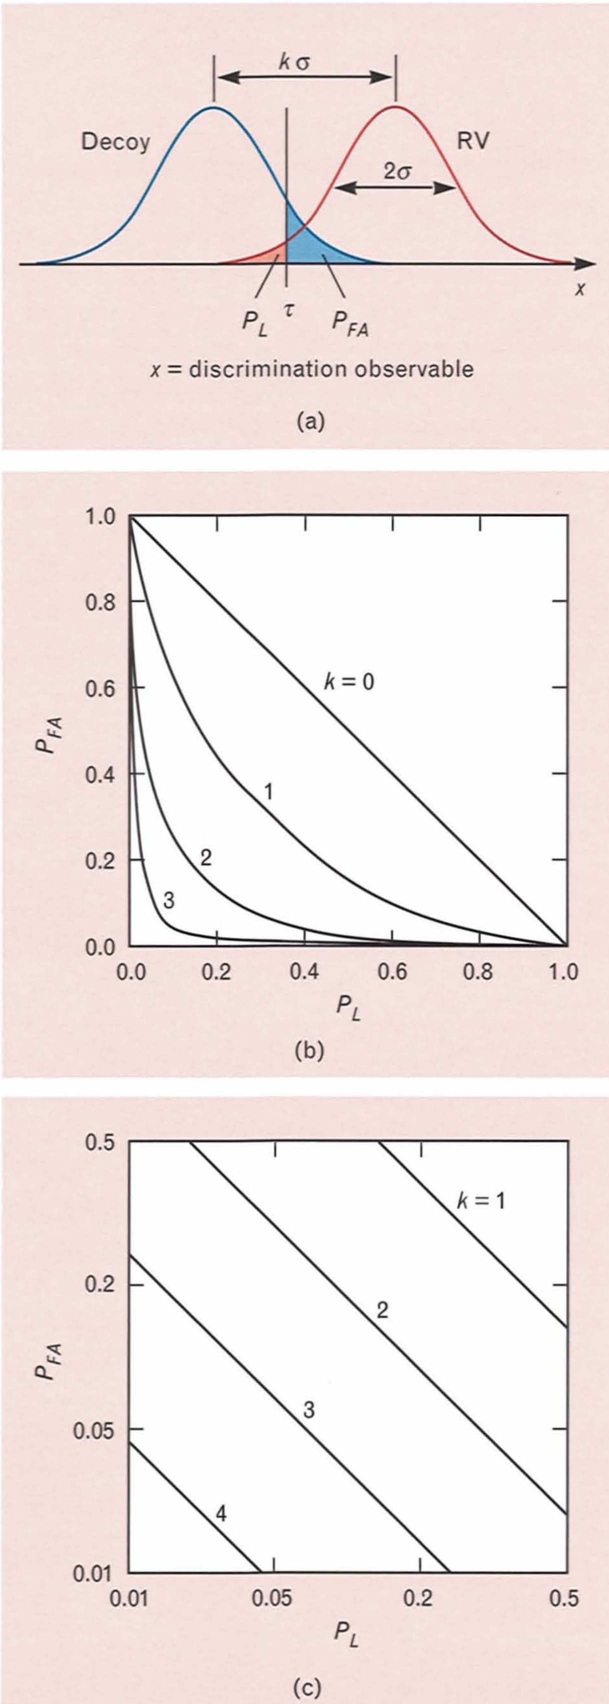

inventory is large. The second strategy is preferred if the inventory is modest and sufficient time exists to perform a kill assessment.

Threat

To carry out discrimination performance analysis, we must also include a model of the threat. A BMD system must be able to handle a variety of incoming ballistic missiles carrying a variety of payloads. These include both theater missiles and strategic missiles with single or multiple RVs, together with deliberate decoys and a spectrum of incidental deployment hardware. Each missile has a given maximum payload; if the offense wants to use additional decoys, then some of the RVs must be removed or smaller RVs must be substituted for the nominal RVs. In any case, the offense must give up something to make room for the decoys.

The offense can design heavy decoys that are likely to be a good match to the RVs (i.e., the *k* factor has a low value), but the offense can deploy only a few of these decoys for each RV replaced. On the other hand, the offense can deploy a much larger number of lightweight decoys but they are likely to be a poorer match to the RVs (i.e., the *k* factor has a high value). Such decoys are called *traffic decoys* and are discussed later in the article. Thus the decoy size, or, equivalently, the *RV/decoy exchange ratio*, is one of the offense variables that must be included in our model of the threat.

The other offense variable is the *offloading*, or the number of RVs to be removed and replaced with decoys. Figure 6 illustrates some of the trade-offs the offense must consider when selecting a decoy design and deciding what mix of RVs and decoys to use on each missile. Each point in the plane in the figure de-

FIGURE 5. A simplified model of the discrimination process. (a) RV and decoy probability distributions are taken to be equal-variance Gaussian curves with means separated by k times their standard deviation. The parameter k is known as the *discrimination factor*, or *k factor*. The trade-off between leakage and false alarm is shown for different values of k on (b) a linear scale and (c) a probability scale. We can trade off leakage and false alarms at fixed values of k by changing the decision threshold τ, and we can simultaneously reduce leakage and false alarms by increasing the value of k.

termines an RV/decoy exchange ratio and the number of RVs offloaded for decoys. If no RVs are offloaded, the defense does not need discrimination because there are no decoys. If all the RVs are offloaded, again the defense does not need discrimination (or anything else) because there are no RVs. For some intermediate value of offloading (which depends on many parameters such as the k factor and the number of interceptors available to the defense) the requirements for discrimination will be most stressing; this is the situation that the offense should select and the defense must design for. In the next section, we indicate how the offense and defense each select the most effective component options in their multidimensional parameter spaces.

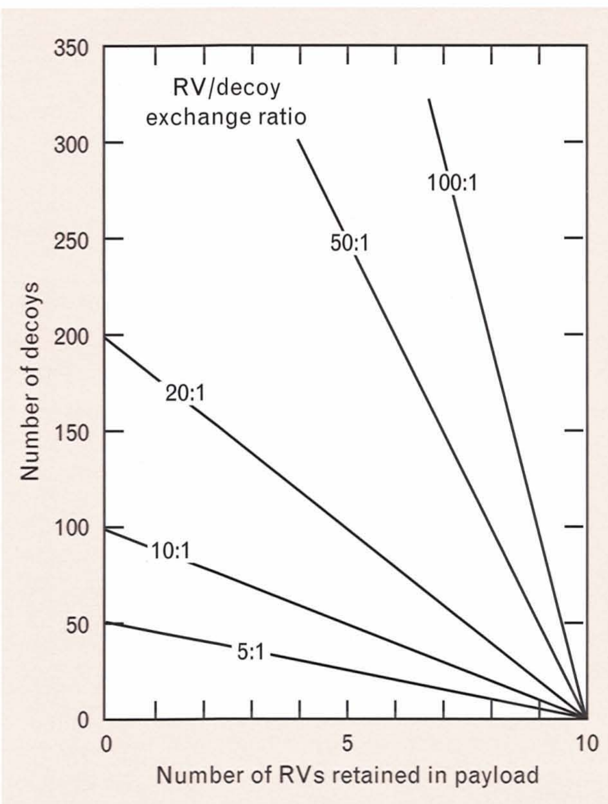

FIGURE 6. An illustration of the payload options available to the offense. In this example, the total offense payload is a maximum of ten RVs, and each RV can be replaced by a specific number of decoys (this number is called the *RV/decoy exchange ratio*). The spectrum of possible payloads is determined in RV-decoy space by a point on a line for the appropriate exchange ratio. The defense must be prepared to deal with any feasible payload combination.

The Fundamental Optimization Problem: Offense and Defense Strategies for Single-Layer Defense

Before we present results for a variety of offense and defense scenarios, let us review all the decision variables controlled by the offense and defense, and whether and when the actual value of these variables will be known by the opponent. Table 1 summarizes these decision variables and classifies them into either long-term, or *strategic*, decisions, such as what type and quantity of components are deployed, or short-term, or *tactical*, decisions relating to how these components are used.

The offense and defense have conflicting goals and differing strategies to implement these goals, which places the problem in a game-theoretic context rather than that of a straightforward optimization. Focusing on what the offense and defense know and when they know it gives the problem a hierarchical decision-theory interpretation. If the defense knew exactly what the offense attack would consist of in terms of number and type of RVs and decoys, it would know exactly what combinations of k factor and number of interceptors would meet the overall BMD-system leakage requirements. This solution is easy to obtain and is shown in the next section as a stepping stone toward the more general results. If the offense knew the number of defense interceptors and the quality of defense discrimination, it would select the mix of RVs and decoys that would maximize the leakage. To be certain of achieving its objective, the defense must be sized to handle all possible offense attacks.

Referring to the variables listed in Table 1, let M be the number of RVs and N be the number of decoys in the offense threat. Let I be the number of interceptors and k be the discrimination k factor comprising a simple defense deployment. Let τ be the discrimination threshold. The overall system leakage L can be written as a function of these five variables, or

$$L(M, N, I, k, \tau).$$

The leakage L is the probability that one or more RVs will penetrate the defense either through discrimination leakage or interceptor exhaustion (or detection

Table 1. Offense and Defense Variables

	*Offense Variable**	*Defense Variable***
Long Term (before attack)	Maximum number of RVs RV/decoy exchange ratio	Number of interceptors *k* factor for each decoy type
Short Term (during attack)	Numbers of RVs and decoys used (subject to payload constraints)	Decision threshold

* Offense knows number of interceptors
Offense may know *k* factor
Offense doesn't know threshold

** Defense knows maximum number of RVs
Defense may know exchange ratio and total number of objects

failure or interceptor failure, if appropriate). A number of scenarios are possible, depending on whether the defense observes all the attacking objects before making discrimination decisions or sees the attacking objects one at a time. The calculation of L for these cases is presented in the appendix entitled "Leakage for Different Discrimination Scenarios."

For each scenario the offense chooses M and N to maximize L while the defense chooses I, k, and τ to minimize L. The net leakage result and the optimum strategies chosen by the offense and defense depend on the order in which these minimizations and maximizations are done.

We assume that the sequence of events is as follows: (1) the defense deploys with given values of I and k; (2) the offense selects optimum values of M and N subject to a payload constraint; and (3) the defense sets an optimum threshold τ. Assuming that the defense has chosen a deployment, this sequence can be represented symbolically by

$$\hat{L}(I,k) = \max_{M,N}\left\{\min_{\tau} L(M,N,I,k,\tau)\right\}.$$

The offense optimization is done by systematically searching all possible values of M and N subject to the constraint

$$M + \frac{N}{R} = M_0 = \text{total payload} = \text{constant},$$

where R is the RV/decoy exchange ratio. The defense optimization can be done by setting a threshold τ such that the expected number of interceptors used is

$$M(1 - P_L) + NP_{FA} = I\,, \qquad (1)$$

or by always shooting at the I most threatening targets. The differences between these two discrimination strategies are discussed in the appendix entitled "Leakage for Different Discrimination Scenarios."

Offense Offloading

As discussed above, the offense attempts to optimize its attack by a judicious mixture of RVs and decoys. Figure 7 shows an example of how we determine the most stressing attack. Consider a simplified example of a threat with a maximum of ten RVs, each of which can be replaced by ten decoys, for a maximum of one hundred decoys. This corresponds to an RV/decoy exchange ratio R of 10. The defense consists of eight, ten, or twelve perfect interceptors and a sensor that has a k factor of 3. We consider several different attacks, each represented by a point on the abcissa of the figure. The defense looks at all the threat objects in the attack and uses all of its eight, ten, or twelve interceptors. The expected number of leaking RVs is calculated for each attack. Leakage curves are plotted for each level of interceptor inventory, and the most damaging attack is highlighted with a small circle. Notice that the most damaging attack for each curve has a higher fraction of decoys as the interceptor inventory increases, but the overall leakage decreases as interceptor inventory increases.

To illustrate some defensive issues in the next examples, we consider offensive threats that are fixed but not necessarily the most stressing. In subsequent examples we include worst-case scenarios. Each example forms an increasingly larger subsystem of the redundant systems and the multiple-layer systems illustrated in Figure 2.

Single Defense Layer Using Single Discriminant

A single-layer configuration using a single discriminant forms the most basic defense architecture. It is the configuration illustrated in Figure 1, and it serves as a good introduction to the issues involved in determining discrimination performance. We are primarily interested in the trade-off between sensor quality, as expressed by the k factor, and the interceptor inventory I, to achieve a fixed system leakage objective L. For simplicity we consider perfect interceptors and a fixed normalized threat consisting of $M = 1$ RV (i.e., the number of RVs is normalized to unity) and $N = 10$ decoys.

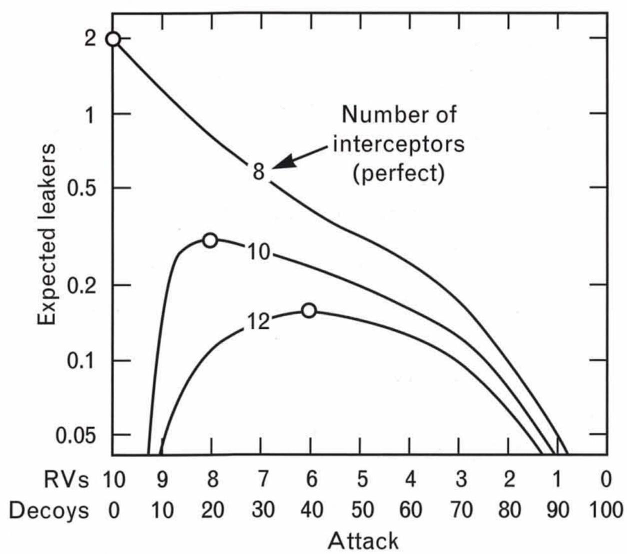

FIGURE 7. Possible leakage for different attack scenarios. The offense has a payload equivalent to ten RVs, and any or all of the RVs can be replaced with ten decoys each. For each feasible attack, we calculate the expected number of RVs leaking through the defense system for a defense with discrimination k factor equal to 3 and for eight, ten, or twelve interceptors. The optimum attack and the resulting leakage for each defense case is shown by a circle. For the case of eight interceptors, the best attack does not use any decoys at all.

Figure 8 shows the operating curves that illustrate the trade-offs involved in a single-layer configuration. Figure 8(a) is an example of the P_L-P_{FA} operating curves for this single-layer system. Superimposed on these operating curves are straight lines of constant I (i.e., number of interceptors fired) as obtained from Equation 1. By cross-plotting with I as an abcissa we see that each P_L-P_{FA} operating curve corresponds to a unique leakage-inventory (P_L-I) performance curve, as shown in Figure 8(b). The two sets of performance curves are entirely equivalent for a specified threat. Most analyses of the discrimination performance of a particular defense deployment use the P_L-I performance curves. By specifying a fixed level of leakage, such as $P_L = 0.1$, we can generate an interceptor versus sensor-quality (I-k) trade-off curve, as shown in Figure 8(c). These curves are particularly useful to quantify how an increase in interceptor inventory can be used to offset a decrement in sensor quality.

For the case of a single RV and perfect interceptors, the probability of overall system leakage L is the same as the probability of discrimination leakage P_L. For some of the subsequent examples, the defense must factor in the additional effects of detection leakage, interceptor leakage, and interceptor exhaustion in converting the requirements on L into requirements on P_L.

Single Defense Layer: Multiple Discriminants with a Traffic-Limited Sensor

As discussed in the threat subsection above, the offense has the option of deploying large numbers of lightweight decoys, or traffic decoys, that will not be a close match to the RVs. If the defense sensor can make accurate measurements on these traffic decoys, it can easily discriminate them. If the defense sensor must handle large numbers of targets, however, it may only be able to make preliminary rough measurements on each object, which reduces the discrimination performance. In this subsection we address the issues related to traffic limitation on discrimination sensors.

We assume a discrimination sensor with a limited

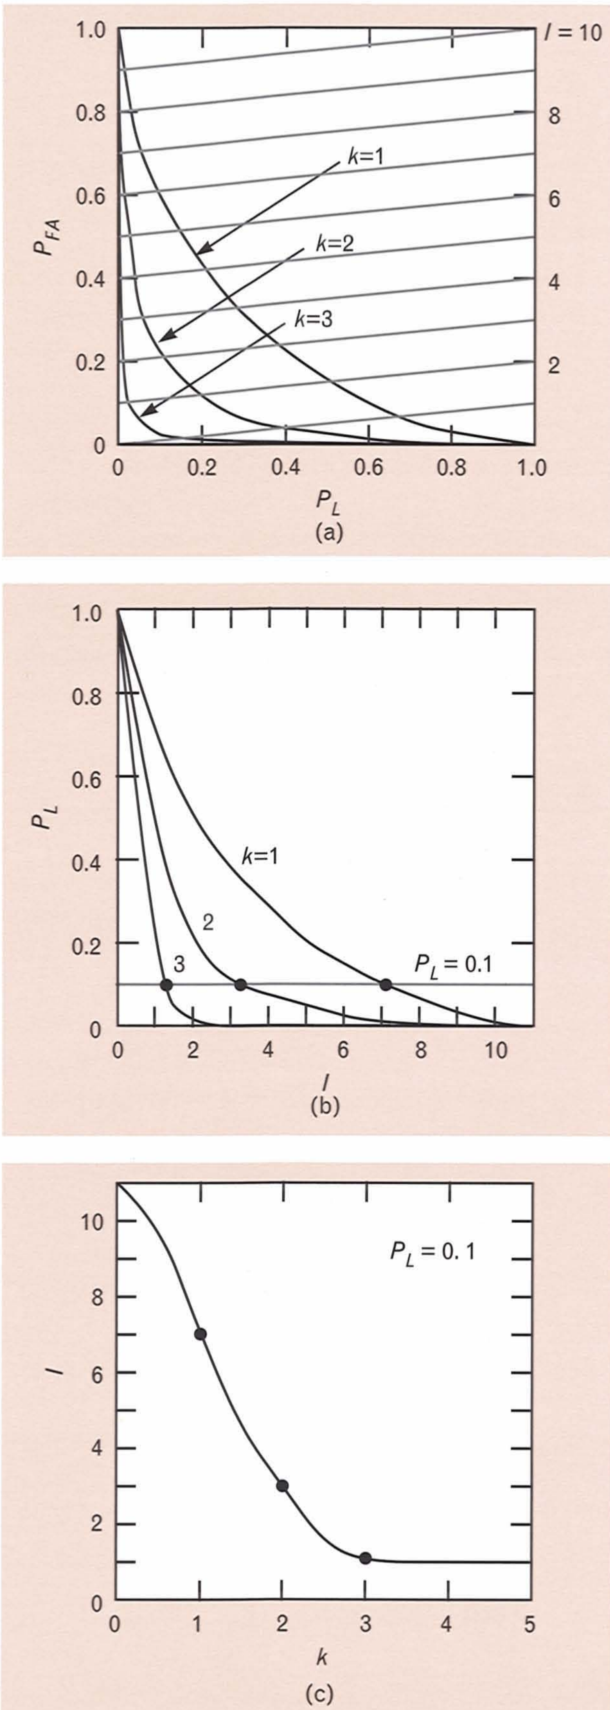

traffic capacity T_0 such that it can measure up to T_0 objects and discriminate them with a given k factor. If faced with a threat of M RVs and N decoys, with the total number of threat objects $T = M + N > T_0$, then the best strategy is to make measurements on T_0 objects, and select objects to be intercepted on the basis of the number of interceptors available. The first targets shot at should be measured objects that appear RV-like. If other interceptors are available, they could be fired at random at unmeasured objects. Finally, if still more interceptors are available, they could be fired at measured objects that appear decoy-like. A precision sensor utilized in this manner is said to have a fixed traffic capacity T_0.

Figure 9 shows curves of P_L versus interceptor inventory I for two sensors with k factor equal to 3; one sensor has no traffic limit, which is identical to the lowest curve in Figure 8(b), while the other sensor has a traffic limit T_0 equal to 5. For leakage values above $1 - T_0/T$, the curves are similar. For smaller values of P_L the traffic-limited curve falls off with a constant slope corresponding to random shooting at objects not measured. Finally, when there are enough interceptors to shoot at all unmeasured and RV-like targets, the leakage reduces asymptotically to zero.

The performance of the traffic-limited sensor in Figure 9 is relatively poor. In this case, a better solution for the defense would be to measure all incoming objects first with a relatively crude sensor or discriminant that cannot be saturated. We call this measurement function *bulk filtering*, or *bulk discrimination*. The general defense strategy for using multiple discriminants with a traffic-limited sensor is to precede a precision sensor with a bulk filter. For precision sensors with a fixed traffic capacity, the bulk filtering can be performed in an optimal way.

FIGURE 8. Generation of defense performance curves for an attack of one RV and ten decoys. A set of discrimination operating curves together with curves of expected numbers of interceptors used is shown in (a) as a function of leakage probability and false-alarm probability. The leakage probability and number of interceptors are cross plotted in (b). Finally, (c) shows how the number of interceptors and the k factor may be traded off for a fixed leakage probability of 0.1. The dots in parts b and c represent corresponding leakage and k-factor values.

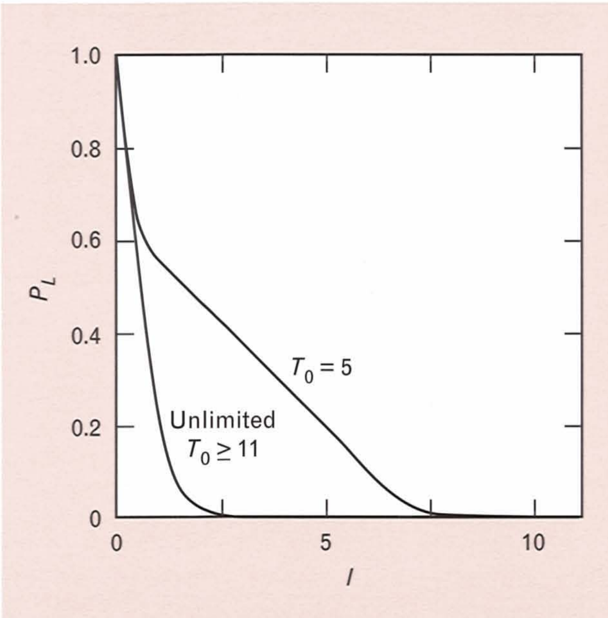

FIGURE 9. Effect of traffic limits on discrimination performance. The unlimited traffic curve is a repeat of the $k = 3$ curve in Figure 8(b). For a traffic-limited sensor that can handle only five of the eleven targets, the performance is clearly poorer. The straight portion of the curve arises from firing randomly at objects that the sensor cannot measure.

Figure 10 illustrates the process of sequential discrimination. Figure 10(a) shows a bulk filter (sensor S_1) in series with a precision sensor S_2. T is the incoming traffic and T_0 is the maximum traffic capacity of sensor S_2. Figure 10(b) shows the procedure to be used in measurement space, where x_1 is the bulk discriminant and x_2 is the precision discriminant. The distributions of the observables in sensors S_1 and S_2 are indicated along the x_1 and x_2 axes. Note that the k factor for S_2 is considerably larger than the k factor for S_1. Probability contours for the joint distributions are also shown. The notation R_1, R_2, D_1, and D_2 denote objects that are classified as RVs and decoys, respectively, by the appropriate sensor. For the bulk filter S_1, two traffic thresholds τ_a and τ_b are set. Objects that fall outside these thresholds are classified specifically as RVs or decoys. Objects that fall between the thresholds are classified as *uncertain*, and they are selected for further measurement by the precision sensor. Because precision sensor S_2 has a fixed traffic capacity of T_0, the thresholds τ_a and τ_b are chosen so that the ex-

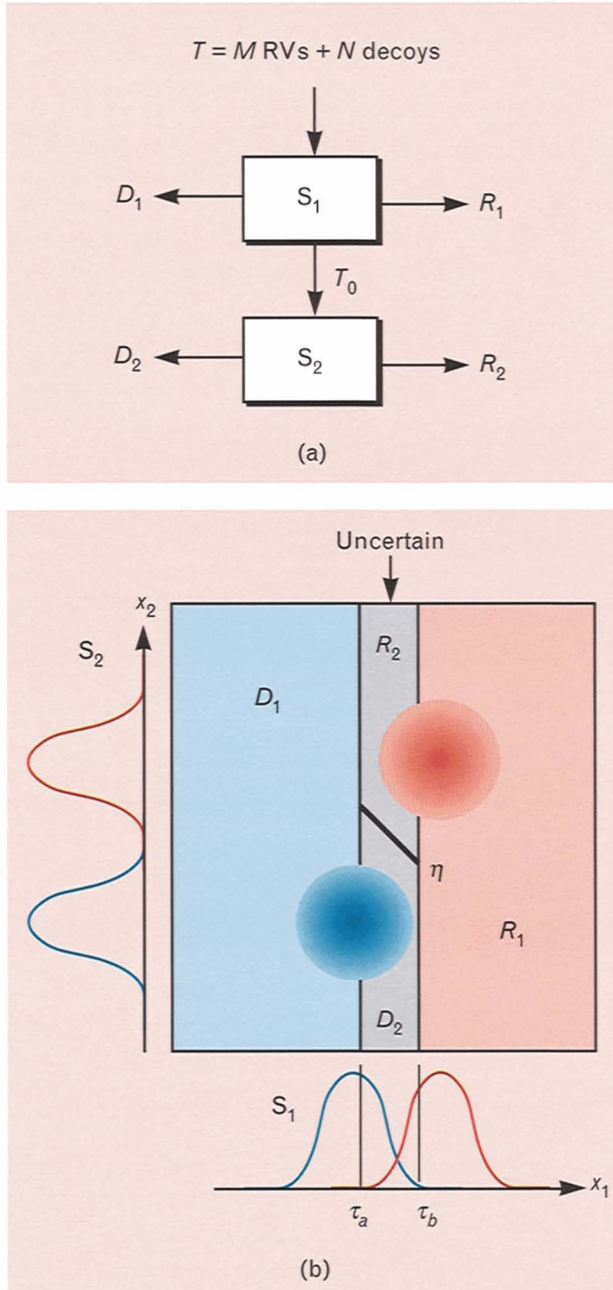

FIGURE 10. Discrimination by a bulk-filter sensor S_1 followed by a traffic-limited precision sensor S_2. (a) The bulk filter reduces a total of T objects to T_0 uncertain objects that are subsequently measured by the precision sensor. (b) The measurement space of the two sensors shows the corresponding RV and decoy density functions and the two traffic thresholds τ_a and τ_b. The bulk filter operates on measurement x_1, while the precision discriminant uses x_1 and x_2 on those uncertain objects in the region between the thresholds. Regions R_1, R_2, D_1, and D_2 denote objects that are classified as RVs or decoys by the appropriate sensor.

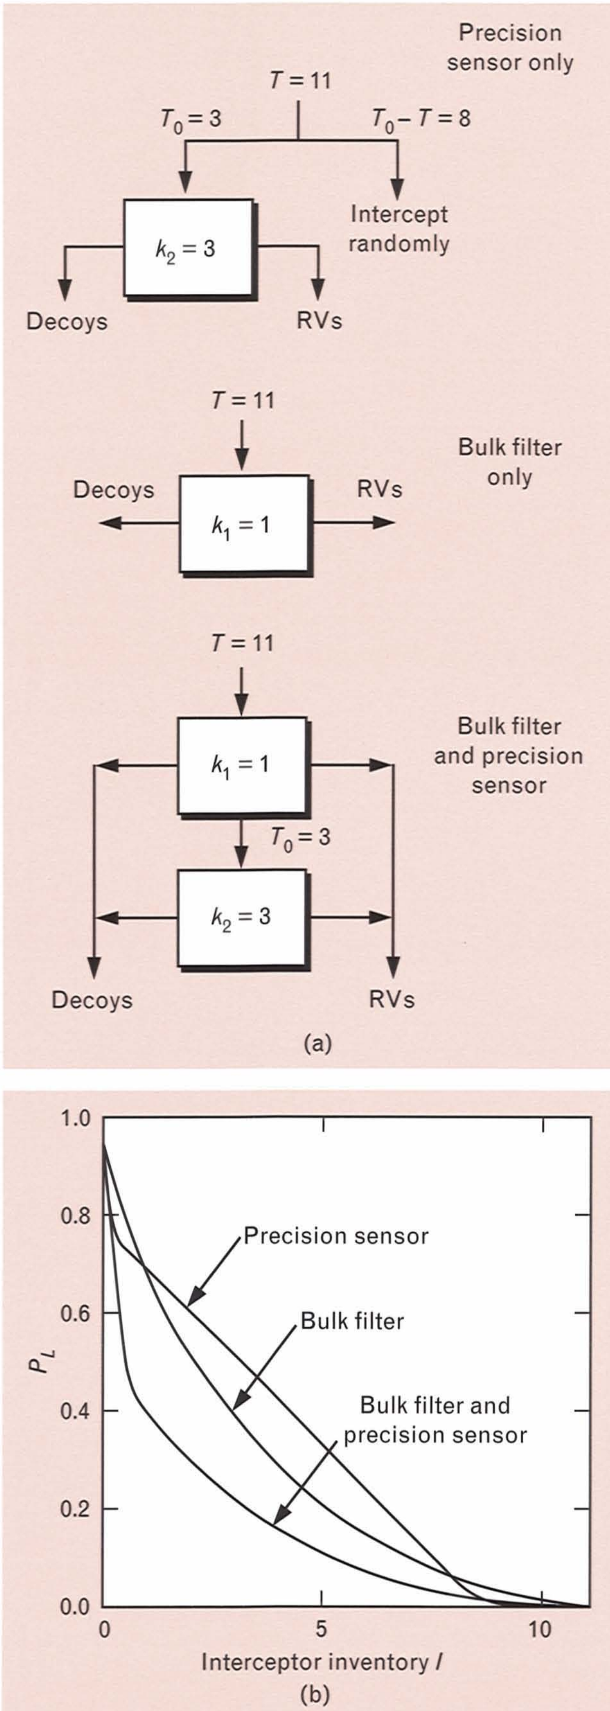

pected number of objects classified as uncertain is equal to T_0.

We want to stress an important point here. Many pairs of traffic thresholds that contain T_0 objects between them can be chosen along the x_1 axis. Only one pair is optimal, however, and they can be determined by the optimal discrimination rules referred to in the appendix entitled "Sequential Discrimination with Limited Resources." This appendix illustrates how the attacking objects can be identified into more than two classes (i.e., RVs, decoys, uncertain), subject to several defense constraints (such as number of interceptors and sensor traffic capacity). Let τ_a and τ_b be the optimal traffic thresholds in the bulk sensor. Next, all objects classified as uncertain—i.e., all objects in the center area in Figure 10(b)—are measured by the precision sensor S_2 and classified as RVs and decoys by using a threshold η in the two-dimensional x_1-x_2 measurement space. The total number of objects classified as RVs—namely, the area $R_1 + R_2$—must be equal to I, the total number of interceptors.

Figure 11(a) shows different combinations of a bulk filter and a precision sensor in series, for specific k factors and traffic capacity T_0. Figure 11(b) shows the performance curves for the combination and also for the bulk filter alone and the precision sensor alone. Of particular importance is the region for small values of the inventory I in which the combined sensor performance is far superior to that of either the bulk filter or the precision sensor alone. For large values of I, the defense is shooting mostly at decoys, and precise discrimination is not that necessary.

Figure 12 illustrates different trade-offs that can occur for a bulk sensor and a precision sensor in series with an attack of one RV and ten decoys. Figure 12(a) shows the trade-offs in k factor necessary to achieve a

FIGURE 11. Discrimination performance for a bulk filter and precision sensor in series. The bulk filter has no traffic limit and a discrimination k factor equal to 1, while the precision sensor has a discrimination k factor equal to 3 but can handle only three of the eleven objects in the attack. (a) The three flow diagrams show the object flow through each sensor alone and through the combination of the two sensors in series. (b) The resulting performance curves show leakage as a function of interceptor inventory for the three sensor cases.

prescribed level of leakage, namely, $L = P_L = 0.2$. The fixed traffic capacity of the precision sensor is the independent parameter. When the precision sensor is not traffic limited at all (i.e., $T_0 = T = 11$), the circular arc

$$\bar{k} = \sqrt{k_1^2 + k_2^2}$$

defines the required combined k factor, where k_1 is the k factor for the bulk filter and k_2 is the k factor for the precision sensor. When the precision sensor is traffic limited, a minimum quality of bulk filtering must be available to achieve the required leakage. Figure 12(b) shows how the minimum required k factor for the bulk filter varies as we change the required level of leakage. Figure 12(c) considers a pair of sensors with fixed k factors ($k_1 = 1$, $k_2 = 3$) and shows how increasing the interceptor inventory I is a possible (but expensive) way to compensate for a limited traffic capacity in the precision sensor.

In designing a defense system, we must be able to select combinations of sensors and interceptors that

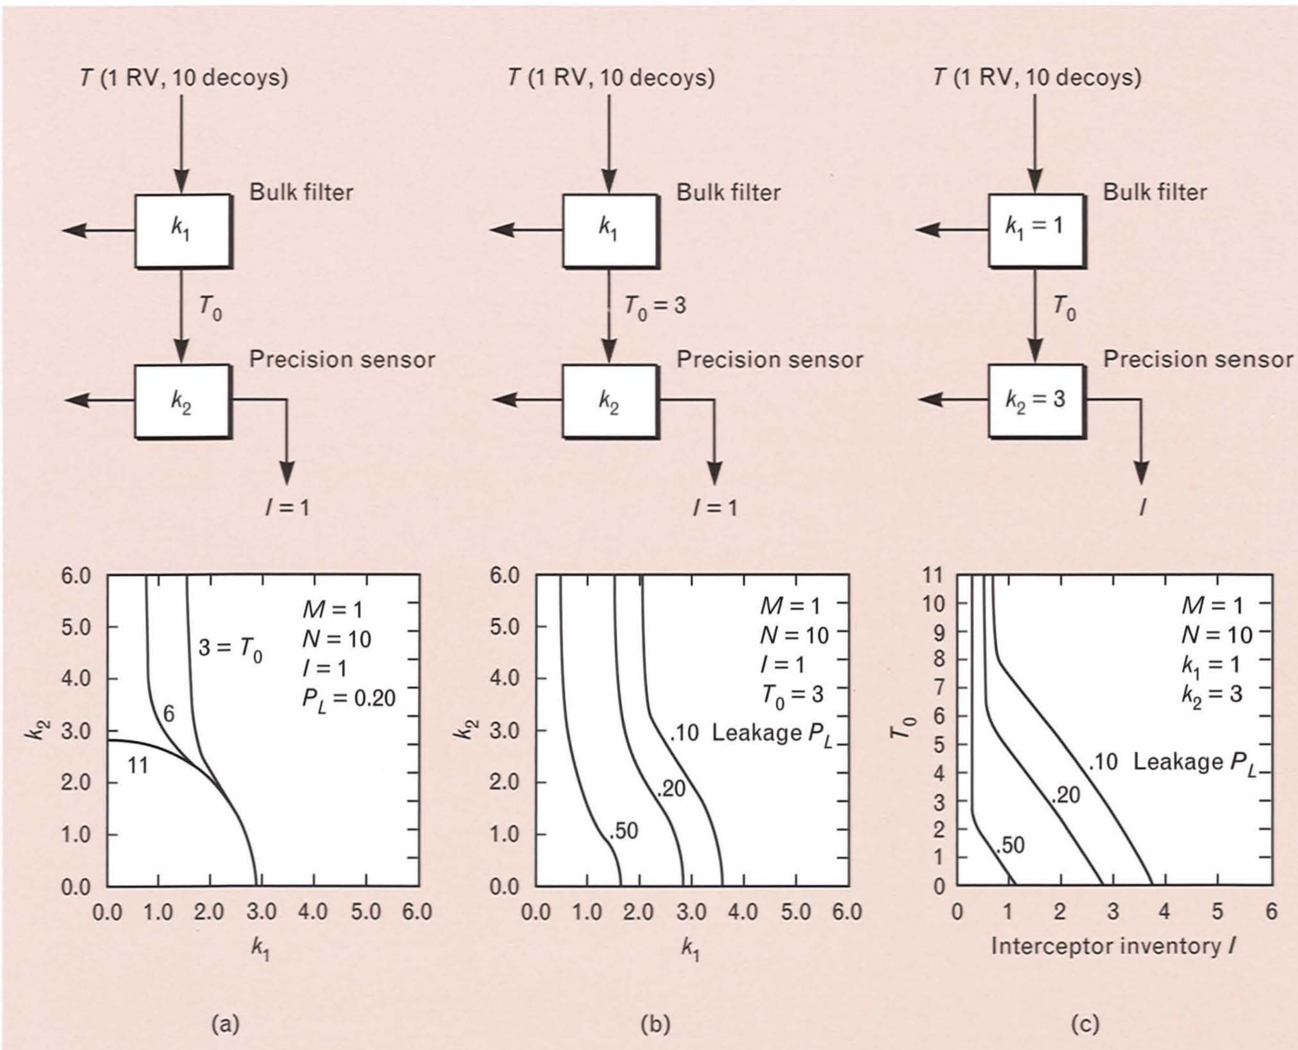

FIGURE 12. Trade-offs for a bulk filter and a precision sensor in series. All results are for an attack of one RV and ten decoys. (a) The k-factor requirements for the two sensors to achieve a leakage of 20% or less are shown as a function of the traffic capacity of the precision sensor. (b) The k-factor requirements with a traffic capacity of 3 are shown as a function of the leakage requirement. (c) The trade-off between precision-sensor traffic capacity and the number of interceptors is shown as a function of the leakage requirement. In each case, the associated block diagram shows the target flow through the sensors.

can handle the entire spectrum of attacks. Bulk filters must be available to handle numerous crude decoys, precision sensors must be available to handle better decoys, and enough interceptors must be available to handle the RVs and the very-high-quality decoys that appear RV-like even to a precision sensor. If the defense is lacking in any one of these areas, the offense is likely to try to exploit this weakness.

Multiple-Layer Defenses

A major reason the defense operates with multiple layers is leakage reduction. If a single layer destroys 90% of the attacking RVs, two independent layers operating sequentially can destroy 99% of the RVs. By definition, each layer in a multiple-layer defense has its own complement of interceptors and its own interception phases. In the case of a two-layer defense, there are two distinct interception phases (as indicated by the two layers in series in Figure 2). This defense configuration is a significantly different architecture from the configuration of two sensors in series within a single layer. In this section we also drop the assumption of perfect interceptors and assume that each interceptor has its own associated probability of failing to kill its target; that is, a single interceptor has an associated leakage l. The second layer can cover leakage from the first layer because of either discrimination failure or interceptor failure.

Figure 13 shows one mode of operation of a two-layer defense. We have not shown the search-and-detection functions in this figure because we want to highlight the passage of discrimination information between the first (or *midcourse*) and second (or *terminal*) layers. Figure 14 illustrates the options for operating the system in Figure 13. In all three cases x_1 is the measurement observable for the first layer and x_2 is the measurement observable for the second layer. Probability contours for the RV and decoy measurements are indicated. The mode of operation known as *handover RVs*, which is illustrated in Figures 13 and 14(a), leads to the following scenario. Incoming objects are measured by sensor S_1 and a discrimination threshold $\tau(x_1)$ is set based only on measurement x_1 and a midcourse inventory constraint I_1. Objects classified as RVs (shown in green and blue to the right of the threshold) are fired on. Those RVs which survive because of interceptor failure continue to be tracked and their discrimination measurement x_1 is retained and handed over to the second layer. Objects classified as decoys do not have their measurements retained.

At the second layer, all these objects (correctly identified decoys and misidentified RVs) are reac-

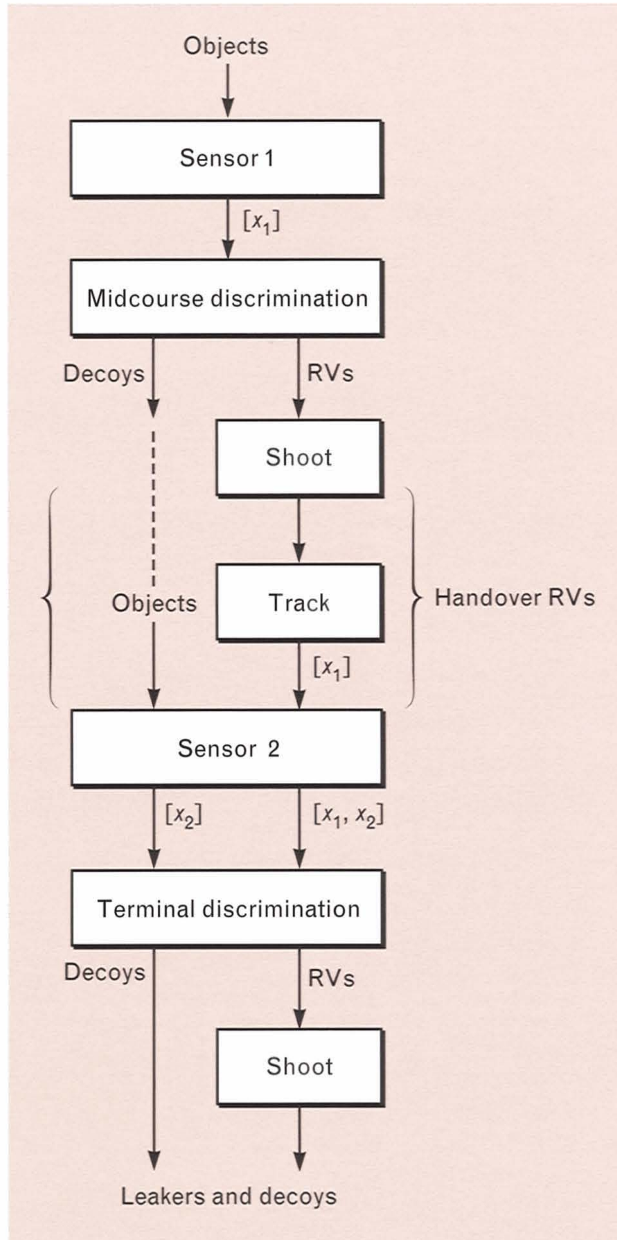

FIGURE 13. Target flow through a two-layer (midcourse and terminal) BMD system. In this case, the objects identified as RVs in the first layer are kept in track while those identified as decoys are dropped (other cases will be considered in subsequent figures). Note that interceptors are fired in both layers.

quired and then measured by sensor S_2. A new discrimination threshold for the second layer is set; this threshold takes the discontinuous form shown in Figure 14(a). Threshold $\tau(x_2)$ is applied only to the subcollection of objects classified as decoys in the first layer; it depends only on measurement x_2 because the x_1 measurement was not retained for these objects. Threshold $\tau(x_1, x_2)$ is applied to the subcollection classified as RVs in the first layer; it depends on both of the measurements x_1 and x_2. The thresholds $\tau(x_2)$ and $\tau(x_1, x_2)$ are set according to a terminal inventory constraint I_2. These two thresholds minimize overall leakage according to the optimal discrimination rules set forth in the appendixes entitled "An Optimal One-Dimensional Representation for a Sensor" and "Sequential Discrimination with Limited Resources."

Figures 14(b) and 14(c) illustrate the form of the optimal thresholds for the cases in which no information and all information, respectively, is handed over to the second layer.

We now discuss some results for these three handover strategies, which we denote as *RVs only*, *none*, and *all*. The following parameters are used for the figures below: the inventory constraint I_1 can vary, but I_2 is related to I_1 by the equation $I_2 = I_1/10$; single-interceptor failure probabilities in each layer are described by the relation $l_1 = l_2 = 0.05$; and the offensive RV/decoy exchange ratio R is 10. For any given defensive deployment (k_1, k_2, I_1, I_2), we assume the offense will optimize its offloading fraction of RVs and decoys to stress the defense maximally.

Figure 15 shows the sensor discrimination trade-

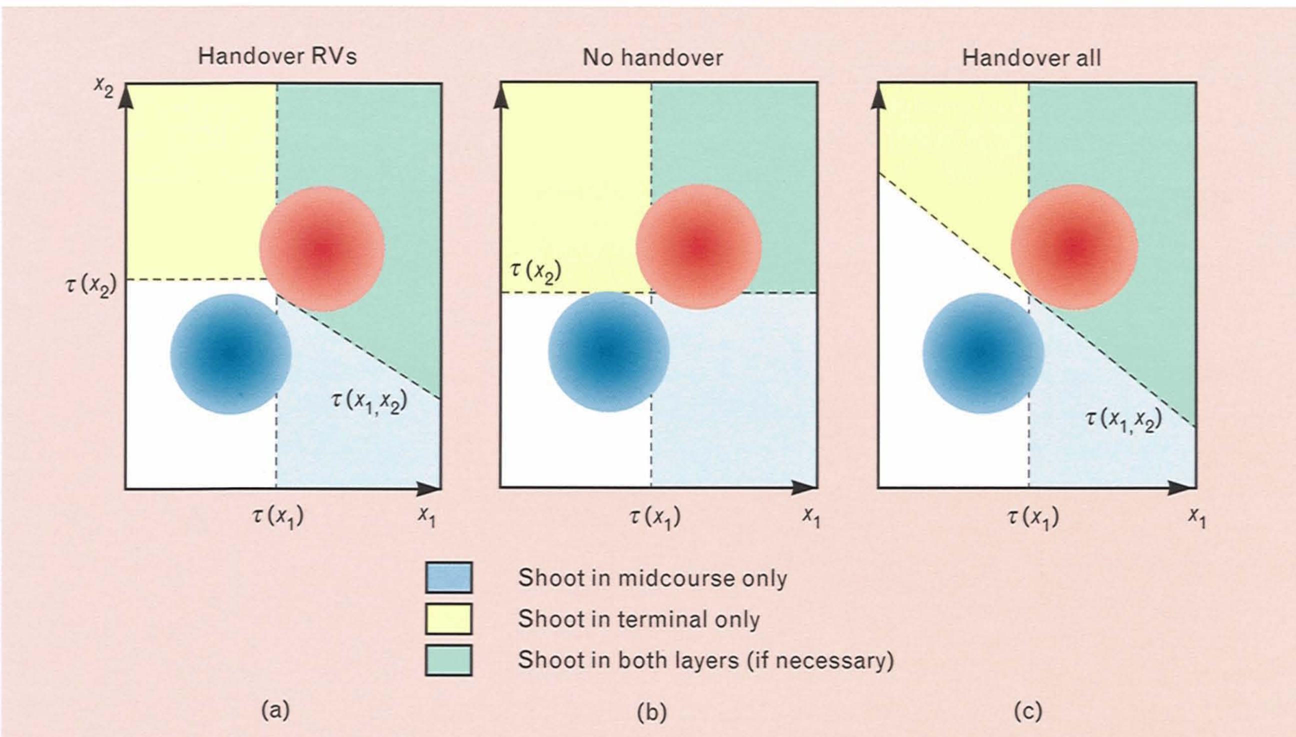

FIGURE 14. Discrimination in a two-layer BMD system. The two-dimensional measurement space of the midcourse (x_1) and terminal (x_2) sensors is shown together with decision thresholds for three different handover strategies. (a) The handover-RVs case corresponds to the flow diagram in Figure 13, while (b) in the no-handover case no communication occurs between layers, and (c) in the handover-all case all information obtained by the midcourse layer is passed to the terminal layer. Because the decision to fire midcourse interceptors is based on midcourse data only, the x_1 threshold is vertical. Depending on the handover option, the terminal threshold is a function of x_2 only for some objects and a function of x_1 and x_2 for others. Objects in the green region (upper right) are shot at in both layers, objects in the blue region (lower right) are shot at only in the midcourse layer, while objects in the yellow region (upper left) are shot at only in the terminal layer. The leakage arises from RVs that are in the white region (lower left) or are shot at and missed.

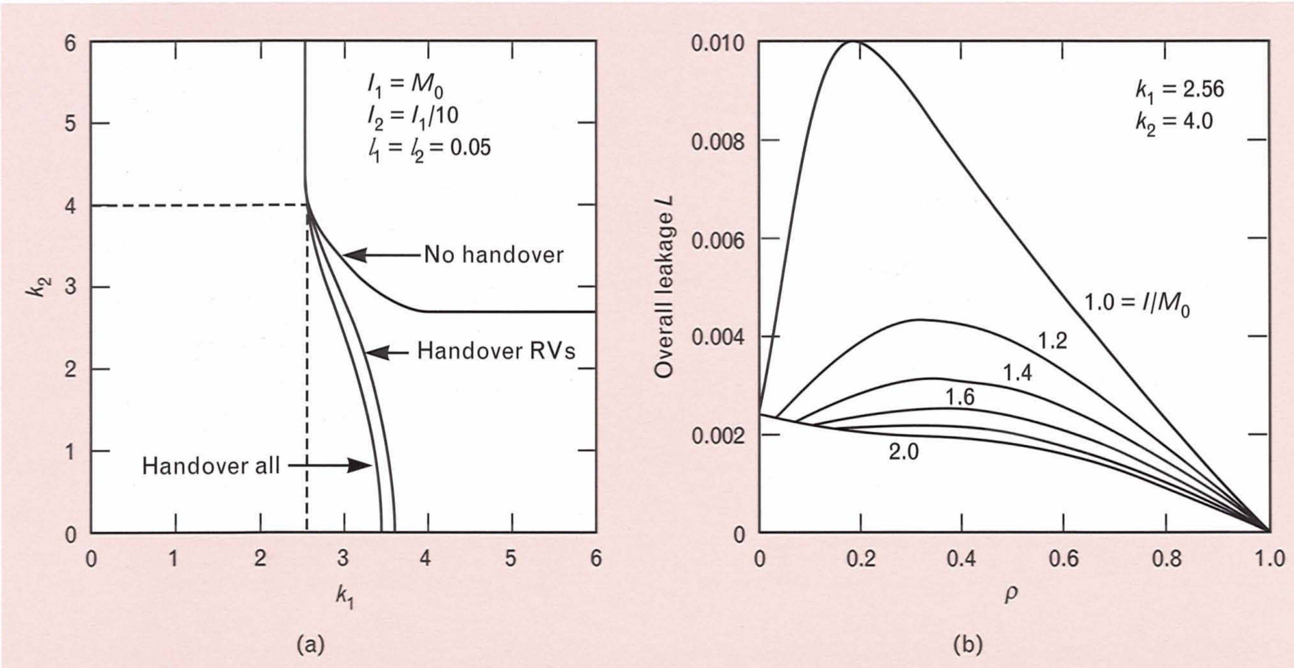

FIGURE 15. Defense and offense trade-offs for a two-layer BMD system. (a) The k-factor requirements necessary to achieve a 1% overall leakage are shown for the case of one midcourse interceptor and 0.1 terminal interceptors for each RV in the baseline (all RV) offense payload. The three handover strategies are those illustrated in Figure 14. (b) The point ($k_1 = 2.56$, $k_2 = 4.0$) examined shows overall leakage as a function of offense offload fraction ρ for several different interceptor inventories for an RV/decoy exchange ratio of 10. The optimum offload is the one that maximizes the leakage. For sufficiently high interceptor inventorys, the best attack does not use decoys. This curve is analogous to the curve shown in Figure 7 for a single-layer system.

off curves required to achieve an overall leakage L of 0.01, or 1%, through both layers. For this case we have assumed perfect detection. An interceptor inventory I_1 for this example is set equal to the offensive payload M_0; that is, the midcourse interceptor inventory would equal the number of RVs in the attack if there were no offloading. We see that if information from the first layer is handed over to the second layer, then the discrimination requirements on the S_2 sensor can be made arbitrarily small if sensor S_1 is good enough. If no discrimination information is handed over to the second layer, then minimum values for both k_1 and k_2 are required to achieve an overall leakage of 1%. Values of k_1 and k_2 above or to the right of the curves in Figure 15(a) will result in an overall leakage of less than 1% for the appropriate handover configuration.

By comparing the curves in Figure 15(a) we see that the defense can achieve its objective in a variety of ways, including different combinations of k_1 and k_2 and different handover strategies. The difference in k-factor requirements between the handover case and the no-handover case is particularly important when the terminal layer has poor discimination performance and the information handed over would be most welcome. In this case, there is relatively little difference between the curves for handing over measurements on all targets and handing over only measurements on the objects classified as RVs; this fact is particularly important because the number of objects classified as RVs is typically a small fraction of the total number of objects in the attack, and the communications requirements could be reduced.

Figure 15(b) shows the result of the offense optimization that was used to calculate a single point on one of the k_1-k_2 trade-off curves. The total offense payload is M_0 RVs, any or all of which can be replaced by ten decoys per RV. The offense offload fraction ρ is defined as the fraction of RVs replaced by decoys. The resulting system leakage L is a function of ρ and the

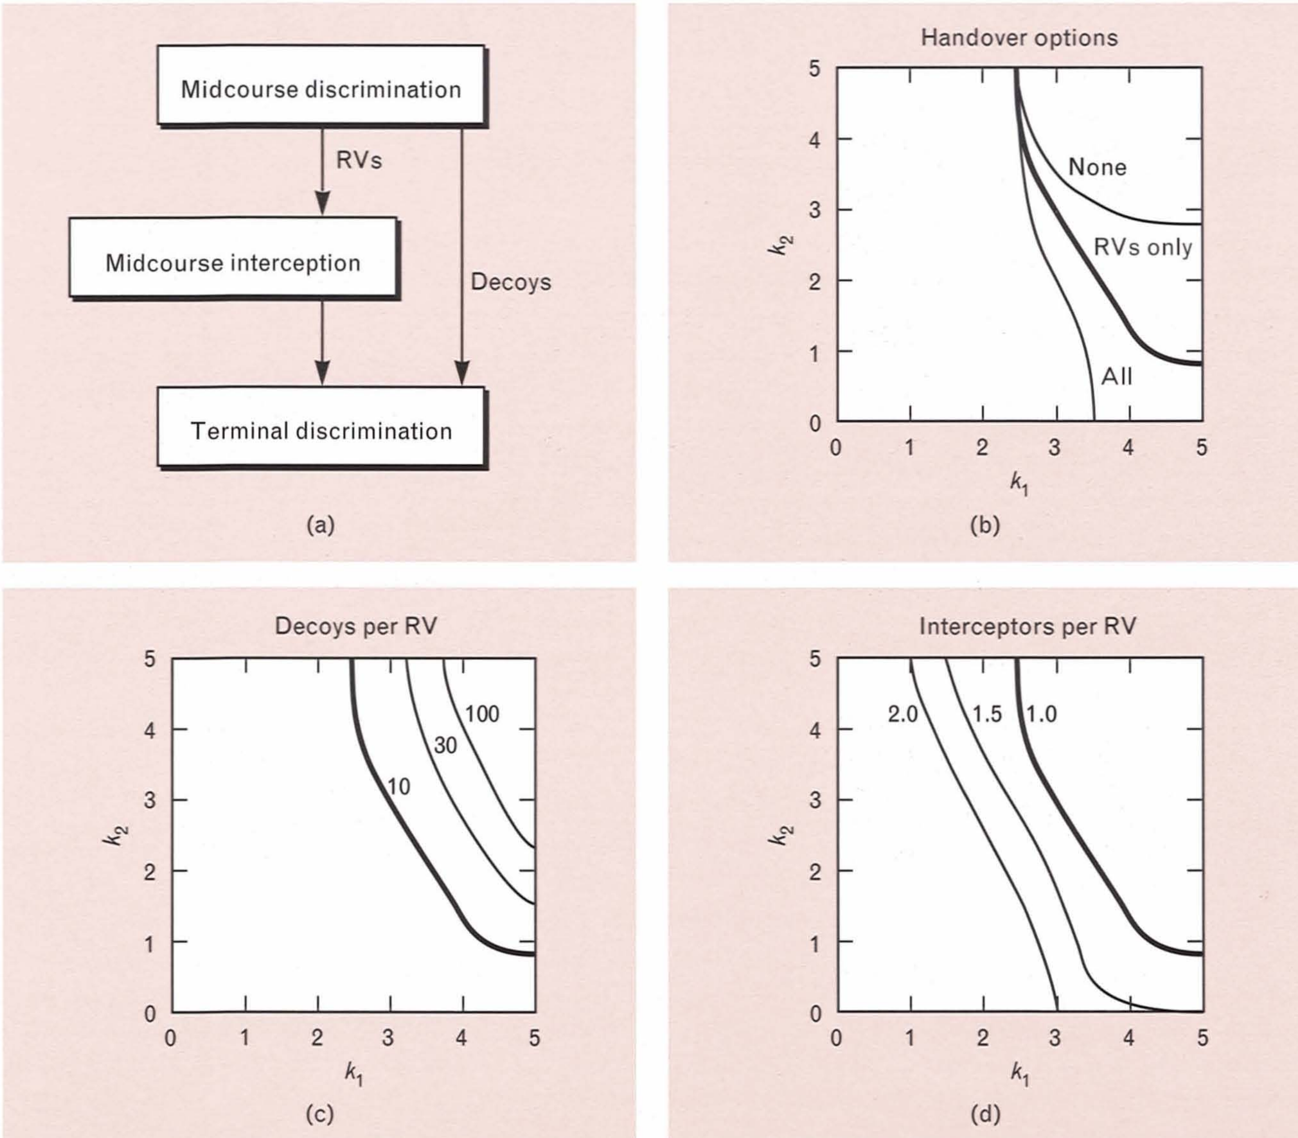

FIGURE 16. Discrimination requirements for two-layer BMD systems. (a) The flow diagram illustrates the flow of objects through the two layers. The requirements on the k factor in the two layers are traded off in various ways in the three sets of curves shown here. (b) This figure compares the different handover options; (c) this figure shows the increase in required k factor with RV/decoy exchange ratio; (d) this figure shows the decrease in required k factor with the number of interceptors. In all cases the detection leakage is 1% per layer, the interceptor leakage is 5% per layer, and the overall leakage requirement is 1%. The heavy curves in each of the three graphs have parameters in common; all other curves are obtained by varying one parameter at a time.

defense parameters I_1, I_2, k_1, k_2, and the discrimination thresholds. Figure 15(b) shows how the system leakage L depends on ρ and the relative defense inventory I_1/M_0 for the point $(k_1, k_2) = (2.56, 4.0)$ from the handover-RVs curve in Figure 15(a). The upper offense operating curve in Figure 15(b) shows that the optimal offloading fraction is 0.18; this offloading fraction gives an overall leakage of 1% when the interceptor inventory $I_1/M_0 = 1$. Offloading curves for other interceptor inventories (which were used in other k_1-k_2 trade-off curves) are also shown. These trade-off curves are similar to those shown in Figure 7 for a single-layer system with perfect interceptors. As the interceptor inventory increases the fraction of decoys in the optimum attack also increases. As the inventory I_1/M_0 increases above the value 1.6, however, the optimal offloading fraction ρ jumps to zero offload. In this case the defense discrimination is sufficient to drive the offense away from using decoys.

Figure 16 shows examples of three different trade-offs in k_1-k_2 space, again to achieve an overall leakage of 1%. For these examples we include a detection leakage of 1% per layer. Thus 1% of all targets are not detected in midcourse and pass through to the terminal layer, where an independent search and detection is carried out; 1% of these targets are also not detected. This leakage, combined with additional leakage in the intercept stage, forces us to reduce the discrimination leakage to achieve the overall system leakage requirement. By using the same type of analysis shown in Figure 15, we present requirements for k_1 and k_2 for different parameter trade-offs. For all three examples, the overall leakage requirement is met for k factors above the curves, while the leakage exceeds 1% for k factors below the curves. The heavy curve in each graph is the parameter value common to the other graphs.

Figure 16(b) shows the effect of different handover strategies. By comparing this figure with Figure 15(a), we see that the handover-RVs curve now requires a minimum k_2 to achieve 1% overall leakage when detection leakage is nonzero. Figure 16(c) shows that discrimination needs to be better (and is likely to be better) for lighter decoys. Figure 16(d) shows that discrimination can be poorer if the defense has more interceptors; note that this figure is the two-layer version of Figure 8(c). Note that for all cases shown in Figure 16, both the offense and defense use their optimum strategies and tactics, subject to the appropriate constraints.

Summary

We have described a number of analyses of discrimination performance requirements for a variety of BMD systems. All these analyses represent a compromise between fidelity and flexibility. They involve multidimensional optimization of both offense and defense tactics that require extensive computation to examine fully the relevant offense and defense system parameters.

The general conclusions of these analyses are (1) for each defense system, there is a threat (including number of RVs, type of decoy, and mix of RVs and decoys) that most stresses discrimination performance; (2) beyond a certain level of discrimination performance the offense does not benefit by using decoys (the defense does not need better discrimination than this); and (3) the defense can compensate for poorer discrimination by deploying more interceptors (at low levels of discrimination this option becomes expensive).

Acknowledgments

Many people in the Analysis and Systems Group have contributed to this effort over the past fifteen to twenty years. The earliest work in the field was done by A.V. Mrstik, K.P. Dunn, and J. Rheinstein. Major participants in specific analyses include J.W. Tolleson, C.B. Chang, R.B. Holmes, G.N. Sherman, S.H. Levine, H.Z. Ollin, P.W. Lert, and R.F. Mozzicato.

APPENDIX 1: AN OPTIMAL ONE-DIMENSIONAL REPRESENTATION FOR A SENSOR

ANY BINARY DECISION problem involving sensors with measurement uncertainties that are described by probability densities or, equivalently, any discrimination problem between two classes of objects, can be recast in the form shown in Figure 4, namely, as two unimodal distributions along a one-dimensional feature axis with a decision region determined by a single threshold. In keeping with the notation of this article we refer to the two classes of objects as RVs and decoys, but the underlying results are more general.

Let the functions $g_R(z)$ and $g_D(z)$ be the probability distributions of the RV and decoy populations, respectively, where z represents any point in a (possibly) n-dimensional measurement space Ω. Suppose the measurement space is partitioned into two regions Ω_R and Ω_D so that if measurement $z \in \Omega_R$ we classify the object as an RV; likewise, if measurement $z \in \Omega_D$ we classify the object as a decoy. Suppose we seek optimal discrimination regions Ω_R^* and Ω_D^* that minimize the expected leakage

$$P_L = \int_{\Omega_D} g_R(z)\, dz\,,$$

subject to an interceptor inventory constraint

$$I = M \int_{\Omega_R} g_R(z)\, dz + N \int_{\Omega_R} g_D(z)\, dz = I_0\,,$$

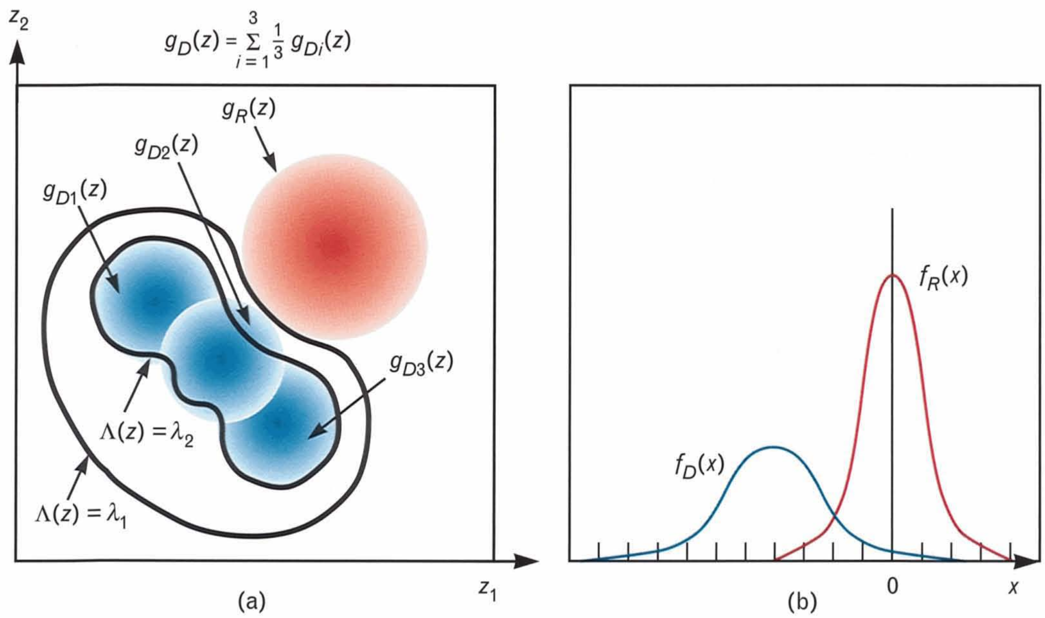

FIGURE A. RV and decoy probability distributions in a multidimensional measurement space. (a) The two-dimensional space (z_1-z_2) shows the probability density contours for an RV and a mixture of three types of decoys. Contours of constant likelihood ratio $\Lambda(z)$ connect points in the space that are equally likely to be RVs. (b) By using these likelihood ratios, we can transform the multidimensional space into a one-dimensional space in which RV and decoy probability density distributions are functions of the variable x.

where M and N are the RV and decoy populations, respectively. Then the optimal discrimination region Ω_R^* is determined by

$$\Omega_R^* = \{z : \Lambda(z) \geq \lambda\},$$

where the quantity

$$\Lambda(z) = \frac{g_R(z)}{g_D(z)}$$

is the likelihood ratio and λ is a constant that depends on the prescribed inventory I_0. The set of points z satisfying $\Lambda(z) = \lambda$ determines a multidimensional threshold surface in the measurement space.

Figure A shows an example of the probability contours of a multimodal decoy population consisting of three subtypes. A unimodal RV population is also shown on this two-dimensional measurement space. Two threshold contours λ_1 and λ_2 are superimposed to show how the threshold can vary in shape as the parameter λ changes. As λ varies, the probability of leakage and the probability of false alarm, given by

$$P_L = \int_{\Omega_D(\lambda)} g_R(z)\, dz$$

and

$$P_{FA} = \int_{\Omega_R(\lambda)} g_D(z)\, dz\,,$$

respectively, trace out a P_L-P_{FA} operating curve similar to that shown in the article in Figure 4(b).

Because all optimal thresholds are surfaces of constant likelihood ratio in the measurement space, we can show that the RV and decoy distributions can be mapped onto unimodal one-dimensional distributions as illustrated in part *b* of Figure A and in the article in Figure 4(a). These unimodal distributions, which can be constructed in a simple manner from the P_L-P_{FA} operating curve, have the important property that their likelihood ratio is monotonic along the x axis. In particular, by moving a single threshold along the x axis, we can generate the same P_L-P_{FA} operating curve. Thus the two representations shown in Figure A are equivalent.

APPENDIX 2: LEAKAGE FOR DIFFERENT DISCRIMINATION SCENARIOS

AS WE DISCUSS in the body of the article, the key calculation done by both the offense and the defense is the determination of the overall system leakage L, where $L(M, N, I, k, \tau)$ is a function of M (the number of RVs), N (the number of decoys), I (the number of interceptors), k (the discrimination k factor), and τ (the discrimination threshold). In this appendix we assume that the defense is operating its discrimination optimally by using some prescribed P_L-P_{FA} operating curve (i.e., a given k factor). As we shall see, the term "optimally" depends on how the defense views the attack and makes its discrimination decision (i.e., chooses a threshold τ). The functional dependence of the leakage on the remaining three variables—M, N, and I—also reflects the defense scenario.

Case 1: Defense Views One Object at a Time

In this case the objects arrive one at a time in random order and the defense must decide whether to intercept each object as it arrives. The defense knows M and N but not the order of arrival of objects. The pertinent equation for $L(M, N, I)$ is

$$L(M,N,I) = \begin{cases} 1 & I < M \\ 1 - Q & I \geq M \end{cases},$$

where

$$Q = 1 - \frac{(1-P_L)^M}{\binom{M+N}{N}} \times \left[\sum_{i=0}^{I-M} \binom{M-1+i}{i} + \sum_{i=I-M+1}^{N} \binom{M-1+i}{i} \sum_{j=0}^{I-M} \binom{i}{j} P_{FA}^j (1-P_{FA})^{i-j} \right],$$

and where P_L and P_{FA} are related by the k factor. The index i represents the number of decoys that arrive before the last RV in the attack. In the first summation, if the number of decoys that precede the last RV is such that $i \le I - M$, then we don't care whether they are classified as RVs or not; the interceptor supply will not be exhausted. In the second summation, where $i \ge I - M + 1$, we must include the probability that at most $I - M$ of these decoys are classified as RVs.

In the equation for Q above, the defense chooses the unique point on the P_L-P_{FA} operating curve that maximizes the function Q. In general, these will not be the same values that satisfy the inventory equation given in case 3.

Case 2: Defense Views Entire Attack

If the defense observes the entire threat before committing its interceptors, it can achieve better performance. It can measure the observables of all the incoming RVs and decoys first, and then shoot at the I most threatening objects. The leakage probability L is

$$L(M,N,I) = 1 - \int_{-\infty}^{\infty} \left[\int_{x}^{\infty} h_R(y)\, dy \right] h_D(x)\, dx\,,$$

where

$$h_R(y) = M\left[1 - F_R(y)\right]^{M-1} f_R(y)$$

is the probability density of the smallest RV, and

$$h_D(x) = \frac{N!}{(N-I+M-1)!\,(I-M)!} \times \left[F_D(x)\right]^{N-I+M-1} \left[1 - F_D(x)\right]^{I-M} f_D(x)$$

is the probability density of the $N-(I-M)$th largest decoy (i.e., there are $I - M$ decoys larger than this one).

The double integral above is the probability that the smallest RV is larger than the $N-(I-M)$th largest decoy. That is, the double integral determines the probability that all M RVs will be included in the I largest objects. The functions F and f, with the appropriate subscripts, are the cumulative distribution and density functions of the RV and decoy observables. F for RV observables is defined as

$$F_R(x) = \int_{-\infty}^{x} f_R(z)\, dz\,.$$

Notice that in this case there is no need for a threshold because all objects are observed before a decision is reached.

Case 3: Simplified Model—Expected Leakage

In this case the threshold τ is chosen so that the expected number of objects classified as RVs is equal to

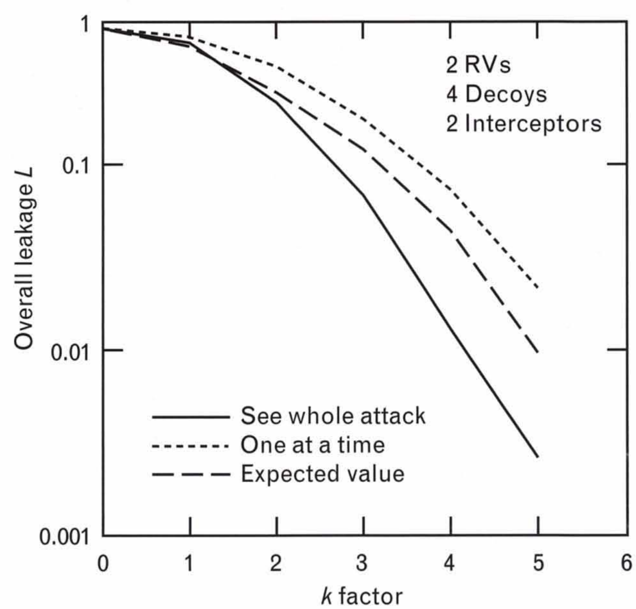

FIGURE A. Leakage for different defense scenarios. The leakage is least if the defense can wait to see the entire attack before commiting any interceptors. The leakage is highest if the defense must make a defensive decision on each object one at a time. The expected-value result is intermediate to these two leakage scenarios and is the easiest to calculate.

the interceptor inventory I, or

$$I = M\int_{\tau}^{\infty} f_R(x)\,dx + N\int_{\tau}^{\infty} f_D(x)\,dx$$
$$= M(1 - P_L) + NP_{FA}\,.$$

The expected probability of leakage is

$$L(M, N, I) = 1 - (1 - P_L)^M\,.$$

This approach results in a greatly simplified analysis, with little loss in accuracy.

Numerical Results

To illustrate how results for the three cases differ, we consider a simple attack containing two RVs and four decoys against a defense with two interceptors (i.e., $M = 2$, $N = 4$, and $I = 2$). Figure A shows the overall leakage L as a function of k factor for the three cases. We can see that the defense always does better if it sees the whole attack rather than if it treats the objects one at a time. Furthermore, in the low-leakage region, the simple expected-value calculation gives results intermediate to the other cases and should be a reasonable approximation in many situations.

APPENDIX 3: SEQUENTIAL DISCRIMINATION WITH LIMITED RESOURCES

CONSIDER THE RV and decoy distributions shown in Figure A in the appendix entitled "An Optimal One-Dimentional Representation for a Sensor" in their original measurement space. Suppose that after measuring all the objects we wish to classify them as RVs or decoys or uncertain, and that a final disposition of objects classified as uncertain awaits further measurements or data processing. As we see below, such situations arise naturally in discrimination systems with traffic-limited sensors or in multiple-layer systems with data handover.

We have developed an optimal discrimination rule that includes the category of objects designated as uncertain; we present this rule here informally. Let us denote RVs as objects of class 1, decoys as objects of class 2, and uncertain as objects of class 0. By using the notation from the appendix entitled "An Optimal One-Dimentional Representation for a Sensor," we seek a partition of our measurement space into optimal discrimination regions Ω_0^*, Ω_1^*, and Ω_2^*, so that if a measurement x lies in a region of Ω_i^* then we classify that object as class i. These regions are optimal if they minimize the expected leakage L, where

$$L = 1 - \sum_j \left[\int_{\Omega_j} r_j P_j f_j(x)\,dx \right]. \tag{1}$$

This relation is subject to either an interceptor inventory constraint

$$T\sum_j \left[\int_{\Omega_j} c_j M(x)\,dx \right] = \text{constant}, \tag{2}$$

or to a traffic constraint

$$(T - T_0)\sum_j \left[\int_{\Omega_j} c_j M(x)\,dx \right] = \text{constant}, \tag{3}$$

where

$$M(x) = \sum P_j f_j(x)$$

is the overall population probability distribution over the measurement space, P_i is the fraction of objects of type i, T is the total number of objects in the threat, and T_0 is the traffic capacity of the precision sensor.

The weighting constants r_j in Equation 1 represent the contribution to expected leakage caused by classifying an RV (class 1) into category j. In some cases, such as when an RV is identified as a decoy, the leakage is immediate and complete. Even if an RV is correctly classified as an RV, however, there may be a leakage contribution due to interceptor unreliability. This leakage would be accounted for by the weighting constants r_j. For RVs initially classified as uncertain (class 0), the expected leakage results from the probability of subsequent misclassification and the probability of interceptor failure, even if the subsequent classification is correct. The weighting constants c_j in Equations 2 and 3 represent those objects in class j which contribute to the inventory constraint. For an interceptor constraint, only objects classed as RVs contribute. For a sensor traffic constraint, only objects classified as uncertain contribute.

The optimal discrimination rules are based on the values of a likelihood function similar to that discussed in the appendix entitled "Leakage for Different Discrimination Scenarios." The function

$$\Lambda_i(x) = \frac{P_i f_i(x)}{M(x)}$$

represents the likelihood function for class i. The optimal discrimination rules follow from considering two inequality conditions. Let x be the measurement observable of any object. Suppose

$$r_i \Lambda_i(x) < \lambda T c_i \quad \text{for all } i\,. \tag{4}$$

If this inequality condition is satisfied, then the object is classified as uncertain. If this inequality condition is not satisfied, then the second inequality condition

$$r_i \Lambda_i(x) - r_j \Lambda_j(x) \geq \lambda T(c_i - c_j) \quad \text{for all } j \neq i \tag{5}$$

is considered. If this second inequality condition is satisfied, then the object is classified as class i. If the interceptor inventory constraint in Equation 2 is applicable, then the optimal discrimination rules in Equations 4 and 5 reduce to that in the appendix entitled "An Optimal One-Dimensional Representation for a Sensor." If imperfect interceptors are considered, then the discrimination rules for the multiple-layer defense result (see Figure 14).

On the other hand, for traffic-limited sensors Equation 3 is applicable and the optimal discrimination rules are applied as follows. The constant λ is a threshold that is adjusted to satisfy the traffic constraint. Because the inequalities given in Equation 4 must be satisfied for both the RV and decoy likelihood functions, the proper choice of λ will yield a pair of thresholds that delineate the region Ω_0^* as uncertain. In the case of the bulk filter and the precision-discriminant sensor in the limited-traffic-capacity example, the proper choice of λ yields the two optimal traffic thresholds shown in Figure 10(b).

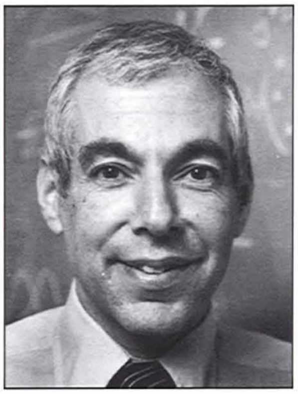

STEPHEN D. WEINER
is a senior staff member in the System Testing and Analysis group. He received S.B. and Ph.D. degrees in physics from MIT in 1961 and 1965, respectively. Since then he has worked at Lincoln Laboratory in many areas of research in ballistic missile defense. These areas include system and sensor design, sensor tracking and discrimination measurements, and interceptor guidance. He has also written several book chapters on different aspects of the history and technology of ballistic missile defense. He has participated in a number of national studies, and for a time he held the world record for the viewgraph with the greatest number of overlays.

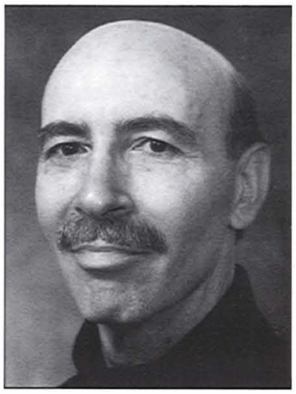

SOL M. ROCKLIN
is a staff member in the System Testing and Analysis group. He received a B.S. degree from Auburn University, an M.S. degree from Princeton University in aeronautics and astronautics, and a Ph.D. in Engineering Science in 1973 from the University of California at Berkeley. He held a postdoctoral fellowship in the Department of Biophysics at the University of Chicago, and he was a National Institutes of Health postdoctoral fellow in theoretical biology at Berkeley. His postdoctoral research included a dynamical model of the early immune system response and a model of chaotic dynamics in multispecies population genetics. In 1978 Sol joined Bell Laboratories, where he worked in the area of telephone traffic theory, and where he developed an inventory-type connect/disconnect policy for the trunking system. Since 1980 he has worked at Lincoln Laboratory in the area of sensor performance. His interests include dynamical systems, statistics, optimization, and mathematical biology.