

The Bang-Soak Theory of Missile Attack and Terminal Defense

By Alan Washburn
Operations Research Department
Naval Postgraduate School
washburn@nps.navy.mil

Executive Summary

This paper describes a spreadsheet-level model for analyzing attacks by a small, mixed collection of ICBMs, perhaps including decoys, on a set of targets individually defended by terminal ABMs. The central questions are how a fixed supply of ABMs should be divided up among the targets, and the resulting effectiveness of the optimized defense. All ABMs are assumed perfect in the sense that each ABM eliminates the reentry vehicle at which it is aimed. Since the ABM assignments are apparent to the attacker, he can “soak them up” by presenting the appropriate number of his least effective reentry vehicles to the subject target. The target is then vulnerable to any remaining “bangs” among the attackers.

The Bang-Soak Theory of Missile Attack and Terminal Defense

By Alan Washburn
Operations Research Department
Naval Postgraduate School
washburn@nps.navy.mil

Abstract

Analyses of the terminal defense of a set of targets by anti-ballistic missiles are generally conducted under the assumption that all attacking reentry vehicles are identical. This paper generalizes to the case where the attacking arsenal is mixed, the main motivation being that a mixed attacking arsenal can contain decoys that are harmless to targets, but which can still “soak up” defenders. All defenders are assumed to be perfect, and the main focus of the paper is on cases where the attack is comparatively weak. A kind of Prim-Read defense results. A simple method for deriving an upper bound on damage is described and illustrated.

Introduction

As nuclear weapons and the means to deliver them by ballistic missiles proliferate (Feickert, 2004), it becomes increasingly likely that a credible threat to use them might eventually be made or even carried out. Concern about this has led to various proposals for defense against small ICBM attacks, with anti-ballistic missiles (ABMs) being the centerpiece. For example, on December 17, 2002, the President of the United States gave ABM development for such contingencies as a justification for withdrawing from the 1972 ABM treaty (Bush, 2002). See Missile Defense Agency (2004) for an overview of the current United States ballistic missile defense program.

This paper introduces a method for assessing the effectiveness of terminal ABMs for defending against an attack by ICBMs carrying a *mixture* of reentry vehicles (RVs) of different types. All of these RVs are assumed to be indistinguishable to the defense when assigning ABMs. The primary reason for considering a mixture of RVs is to allow for decoys, which can be thought of as RVs with no effectiveness against targets. There is also, of course, the possibility that the RVs are not decoys, but simply differ in effectiveness.

We assume throughout that ABMs are perfect in the sense that each ABM will reliably eliminate one RV. The simplicity of this assumption will make it easier to explore the real focus of our inquiry, which is the mixed nature of the attack. If the RVs are not mixed, then there are existing methods where the assumption of ABM perfection is not necessary (Burr, et al., 1985).

ICBMs are vulnerable in their boost, midcourse (exoatmospheric) and terminal phases. Boost phase intercepts are in a sense ideal, since the ICBM is destroyed before it can deploy its payload, but difficult to carry out. Terminal defenses benefit from the effect of the atmosphere, which strips away some of the decoys that are effective outside of the atmosphere. Midcourse (exoatmospheric) ABMs offer the flexibility of being able to defend many targets, rather than just one. Effective defensive systems can be constructed out of the joint use of *both* midcourse and terminal ABMs, although such systems are comparatively difficult to design and study. Miercort and Soland (1971) study the attack of such a system under the assumption that all interceptors (midcourse and terminal) are perfect, taking the defensive assignments as given. Our ambition is to solve a two-sided problem with a mixed attack, so we make the simplifying assumption

that all ABMs are terminal, except possibly for a midcourse ABM system whose reach includes all targets, and which is also perfect.

To some extent, the mixed context justifies our assumption that ABMs are perfect, which is operationally equivalent to assuming that ABMs are sufficiently threatening that the attacker chooses to neutralize them with decoys or small RVs, rather than subjecting one of his more powerful RVs to possible intercept.

Even though the atmosphere is an ally in stripping away light decoys, it is still possible to design functional terminal decoys. Such decoys are heavy enough to displace other RVs, so the attacker has a payload tradeoff to make between firepower and dilution of terminal ABM defenses. This tradeoff is central to our problem.

General Analysis

Terminal ABMs must wait until incoming RVs are close to their targets before attempting an intercept. Operationally, “close” means that ABMs must be assigned to individual targets, with no possibility that an ABM assigned to one target can defend another. The attacker must similarly partition his forces, with the outcome at each target depending only on the forces assigned to that target.

The sequence in which the assignments are made, and the knowledge that each side has when making assignments, are crucial. Each side would, of course, prefer to make its assignments knowing the assignments that have already been made by the other. The scenario most favorable to the defense is when the ABM force can examine the whole attack, all at once, before committing any ABMs. Given the difficulty of determining the attacker’s ultimate intention when the first ICBM is launched, this assumption is rarely made. Our analysis here is based on the opposite assumption, where

the physical deployment of ABMs is observable and reveals to the attacker the extent to which targets are being defended. The intermediate case where both sides must commit forces without knowing the other's intention is a kind of Blotto game (Washburn, 2003) that generally goes by the name of "pre-allocation" (Eckler and Burr, 1972; Bracken, et al, 1987). In a preallocated defense, a secret but fixed number of interceptors is irrevocably committed to each target. If the number of attackers exceeds that number, then further interceptors are not committed, even if the command and control system would permit it. This seems unlikely when discussing a defense against small attacks. Besides, Blotto analyses have been performed only for homogeneous attacking forces, and we have no useful generalization to offer.

When the attacker has the last move, the best that the defense can hope for is that his ABMs will eliminate the most powerful attacking weapons, one for one. If the defense also incorporates the idea that no attacking RV will be deliberately allowed to penetrate while the means of intercepting it are still available, as we assume to be the case, then the best that can be hoped for is that the ABMs will intercept the *least* effective RVs. The reason for this is that the attacker can always ensure that result by attacking each target at its known defense level, using the least effective RVs, before going on to attack the now undefended targets with more powerful RVs. The defense may be even less effective, since heavily defended targets may simply be ignored by the attacker while undefended or lightly defended targets are overwhelmed. In that case some ABMs might eliminate no attackers at all. We call situations where some ABMs are wasted in this manner "defense dominant", since they characterize the attack of a strong defense by a weak attacker. It is defense dominant situations that are our primary concern.

In the defense dominant case, it is natural for the defense to plan his ABM allocations in such a manner that the attacker sees no weak spots, in the sense that there are no targets where there is a large return per attacker. Such defenses are called Prim-Read (Read, 1958; Karr, 1981; Burr, et al, 1985) defenses, for which there are well known optimal allocation methods when all attackers are identical. Our intention is to generalize this work so that the attackers can differ in terms of the amount of target killing power.

The killing power of a nuclear arsenal has a number of measures. One is a simple count of the numbers of warheads. Nuclear arms limitation treaties deal mainly with this measure, since it is the easiest to verify. Other measures depend on the yield (Y , in megatons) and accuracy C (“circular error probable” in nautical miles) of the warheads. $Y^{2/3}$ is the warhead’s Equivalent MegaTons (EMT), essentially a measure of the amount of two-dimensional area that can be covered lethally. EMT measures a warhead’s capability to kill soft countervalue targets such as urban areas. $Y^{2/3}/C^2$ is a warhead’s Counter Military Potential (CMP), a measure of its ability to kill hardened point targets. Both CMP and EMT are cumulative measures; that is, the effectiveness of an arsenal is just the total effectiveness of the RVs composing it. During the Cold War, the arsenals of the United states and the Soviet Union differed strongly in these measures, with the United States (Soviet Union) having more CMP (EMT) (*Current News*, 1976).

The most natural measure of power for a small attack is probably EMT, but we do not wish to be specific, so we refer in the sequel to a general cumulative measure called “bangs” that could be either EMT or CMP. However, a bang must be one thing or the other; that is, the target set cannot consist of a mixture of hard and soft targets. If the

target set is mixed in that manner, then both the total EMT and CMP of the attacker's arsenal are important, and the method described below will not apply.

The attacker's arsenal is to be expended in attacking a collection of targets indexed by i , each of which has an associated value function $V_i(x)$ that represents "expected value killed if the target is undefended and attacked by x bangs". If target i is defended by d perfect ABMs, then the first d attacking RVs will be eliminated, and only the bangs associated with RVs in excess of d will affect the target. The first d attackers "soak up" the ABMs, so it is useful to think of each RV as possessing a single "soak", in addition to its bangs. Obviously, if the vehicles differ bangwise, then the attacker will use small ones (especially decoys) to soak up the defense. Since the defenses are assumed known when the attack is being optimized, our object is to calculate the minimum (over defenses) of the maximum (over attacks) of the total target value killed, or, more concisely, the min max value.

The numbers of ABMs and soaks are required to be integer-valued, but the number of bangs possessed by an RV can be any nonnegative number. Every RV has exactly one soak, but the number of bangs depends on the physical system of units in use — an RV with only .01 bangs might actually be very powerful. The total supply of bangs is assumed to be infinitely divisible, so that a given quantity of bangs can be distributed arbitrarily over the targets. This assumption makes most sense when there is a diversity of RVs, but is uniformly employed below without further inquiry into the question of diversity. In essence, once the set of targets at which the defenses have been soaked up is determined (this set automatically includes all undefended targets), the attacker is free to allocate the total bangs of all remaining RVs to that set of targets.

We first consider the case where only two types of RV are available to the attacker, one being decoys. The mix of RVs is subject to a linear payload constraint, but is otherwise determinable after the defense has been observed by the attacker. The constraint is $b + \lambda s \leq B$, where b and s are the number of bangs and soaks, respectively. Parameter B is the total ICBM payload measured in bangs, and each decoy displaces λ bangs, where λ is some positive constant. We refer to this as the “free” case, since the attacker is assumed free to determine s (a nonnegative integer) and b (a nonnegative real number) after observing the defenses. We will then consider “constrained” problems where the attacking arsenal is a given, arbitrary mixture of vehicle types. The allocation of the attacking arsenal can be made after observing the defenses even in the constrained case, but the arsenal itself is considered fixed and known to the defense.

The following two examples of the free case illustrate the attack dominant and defense dominant cases.

Example 1: Suppose there are a small target and a large target, with the two value functions being $V_1(x) = (1 - \exp(-x))$ and $V_2(x) = 2(1 - \exp(-x))$. There are four ABMs, and the attacker is constrained by $b + 0.5s \leq 3$. It is easy to show that the optimal defense is to use all four ABMs to defend the second target. Even knowing this distribution, the attacker should still soak up all four ABMs and use his remaining bang to achieve 1.285, splitting the bang (0.15, 0.85) between the two targets. If the attacker does not soak up all four ABMs, he can achieve at most 1 by attacking the first target.

If the attacker actually had six identical RVs with 0.5 bang each, then the best “split” of the bang remaining after soaking up the four ABMs would be to use the whole

bang on target 2, this being better than the (0.5, 0.5) split achieved using one RV on each target. However, as mentioned earlier, we assume that bangs can be split arbitrarily among the targets, even in small examples such as this one.

Example 2: This is the same example except that $B = 2$; i.e., $b + 0.5s \leq 2$.

Against the same defense, the attacker would ignore the second target, achieving $1 - \exp(-2) = 0.865$ on the first. The best defense is to assign 1 ABM to the first target and 3 to the second. The best attack is then on the second target, with the payoff being $2(1 - \exp(-0.5)) = 0.787$. This is the min max total value killed. Unlike example 1, this example is defense dominant because at least one defended target is not attacked when both sides allocate optimally.

General analysis of the free case

Let T be the set of targets, let D be the total number of ABMs, and let the min max value of the payoff to the attacker from attacking T be $M(B,D)$. Thus $D = 4$ and $M(3,4) = 1.285$ in example 1, or $M(2,4) = 0.787$ in example 2. Our object is to find a general upper bound for $M(B,D)$.

The problem of optimally allocating bangs when $D = 0$ is a separable, concave maximization problem. It closely resembles the search theory problem of allocating search effort to a set of places where a target might be, so a variety of efficient solution techniques is known (see Washburn(1981), or more generally Ibaraki and Katoh (1988)). If S is any subset of the targets, let the maximum total value achievable with B bangs when the targets are undefended be $V(S,B)$. For convenience, we define $V(S,B) \equiv 0$ for $B < 0$.

If a target is defended by d ABMs, then the attacker should either soak up all d of them before applying bangs, or else ignore the target entirely. Therefore, once the set of targets to be attacked has been selected and the price of admission paid in soaks, the attacker faces the concave optimization problem described above. Formally, let $I(S)$ be the number of interceptors assigned to target subset S . The function $I()$ is to be chosen by the defense, subject only to the constraint that no more than D interceptors can be used in defending the whole target set T . Once the defense is committed, the attacker chooses a subset S to attack. Thus

$$M(B,D) = \min_{I(T) \leq D} \max_{S \subset T} V(S, B - \lambda I(S)) \quad (1)$$

Since the attacker is free to attack the entire target set, it follows from (1) that $M(B,D) \geq V(T, B - \lambda D)$, with equality holding in the attack dominant case. Theorem 1 gives circumstances where equality must hold.

Theorem 1: Suppose that $V_i(0)=0$ and $V_i(x)$ is an increasing value function for $x \geq 0$, with $\lim_{x \rightarrow \infty} V_i(x) \equiv V_i^* > 0$, for $i=1, \dots, n$. Then there exists a B^* such that the attack is dominant for $B > B^*$.

Proof: Let $V_{TOT} \equiv \sum_{i=1}^n V_i^*$, and let $i^* \equiv \arg \max_i V_i^*$. Since $V(T,B)$ increases with B to V_{TOT} , there exists some B^* such that $V(T, B^* - \lambda D) > V_{TOT} - V_{i^*}$. For $B > B^*$, the attacker can gain more than $V_{TOT} - V_{i^*}$ by engaging the entire target set T , whereas he will gain at most that amount by engaging any set of targets that does not include i^* . Therefore target i^* must be attacked, and an optimal defense is to concentrate all D ABMs on i^* . All D ABMs will be soaked up, so the attack is dominant. QED

In examples 1 and 2, with parameters other than B as specified, B^* is at most 3 and at least 2.

Except in attack dominant cases, B^* will be greater than B . $M(B,D)$ could still be determined using (1), but the number of evaluations of $V()$ could easily be very large. If there are N targets, then the number of possible defenses is the number of combinations of $N+D - 1$ things taken D at a time. For each of these defenses, $2^N - 1$ attacks must be considered unless B is small enough to make the attack of some target subsets infeasible. If $N = 10$ and $D = 10$, the total number of evaluations would be $(92,378)(1,023)$, approximately 10^8 . Even though any given evaluation of $V()$ is not difficult, this combinatorial explosion will make an exact evaluation of $M(B,D)$ difficult in even moderately large problems.

Our interest is mostly in the defense dominant case, which is characterized by the attacker ignoring some targets and looking for cheap victories at others. We will find an upper bound on $M(B,D)$ by having the defense control the return per bang that is available to the attacker. If the target is defended by d ABMs, and if the attacker chooses to soak up those ABMs and commit an additional x bangs, then the return per bang (converting the soaks into bangs at the rate λ) is $V_i(x)/(x+\lambda d_i)$. If the return per bang is to be bounded above by $r > 0$, then it is necessary that $V_i(x)/(x+\lambda d_i) \leq r$ for all x , or, equivalently, that $\lambda d_i \geq F_i(r)$, where

$$F_i(r) \equiv \max_{x \geq 0} \{V_i(x)/r - x\}. \quad (2)$$

The function $F_i(r)$ exists and decreases from ∞ to 0 on the interval $(0, \infty)$, as long as $V_i(x)$ is as specified in Theorem 1. We can now define the smallest return per bang that is permitted by an inventory of D ABMs when the cost of a soak is $\lambda > 0$:

$$r(\lambda, D) \equiv \min \left\{ r : \sum_{i=1}^n \lceil F_i(r) / \lambda \rceil \leq D \right\}, \quad (3)$$

where $\lceil z \rceil$ is the smallest integer that is not smaller than z ; i.e., z rounded up. If the set over which the minimum is taken is empty, we take $r(\lambda, D) = \infty$. Since $F_i(r)$ approaches infinity monotonically as r approaches 0, a finite minimum will be found somewhere in $(0, \infty)$ as long as $\lambda > 0$. The individual terms of the sum in (3) are the defense levels d_i that are required to assure $r(\lambda, D)$. A notable feature of this defense is that it does not require knowing B . Indeed, this feature is sometimes taken to be definitive of the Prim-Read defense

Figure 1 shows a plot of the sum in (3) versus r for Examples 1 and 2. When $\lambda = 0.5$, the smallest value of r for which the sum is smaller than the given number of ABMs is $r(0.5, 4) = 0.520$.




Figure 1: Showing the total number of ABMs required (D) versus the return per bang (r) in examples 1 and 2.

Once $r(\lambda, D)$ is obtained, the upper bound is obtained by multiplying by B :

$$M(B,D) \leq Br(\lambda,D) \tag{4}$$

This bound is 1.560 in example 1, or 1.040 in example 2, compared to the true values of 1.285 and 0.787.

Our contention is that the upper bound is only slightly larger than the true min max value in defense dominant problems where there are many targets of different types. We have given little evidence of this so far, since 1.040 is not a particularly good approximation to 0.787 in Example 2. But consider what happens when Example 2 is scaled up so that there are two targets of each type, two “small” and two “large”, with B and D each likewise doubled. If the defense chooses to defend both small targets with 1 ABM each, and both large targets with 3 ABMs each, then the attacker should attack one target of each type, leaving him with 2 bangs after soaking up four ABMs. The small target should get 0.65 bangs and the large target 1.35 bangs, for a total score of $0.48 + 1.48 = 1.96$. Furthermore, *any* defense must leave some pair of small and large targets with at most 4 ABMs between them, so the attacker can do at least this well regardless of the defense (if two ABMs are assigned to each target, for example, the attacker could attack both large targets and achieve 2.53). Thus $M(B,D) = 1.96$. The upper bound in this scaled-up problem is 2.08, only 6% larger than $M(B,D)$.

Intuitively, the attacker has so many choices of what set of targets to attack in large problems that he is almost bound to be able to partition his force so that each attack is carried out at a level that achieves nearly $r(\lambda,D)$ per attacker. Further evidence of this phenomenon is shown in Figure 2, which shows the best possible return per attacker in a problem where there are 10 identical “small” targets, 10 ABMs, and $\lambda=0.5$. The ABMs should of course be distributed one per target. The best possible return per attacker is

then $r(0.5, 10) = 0.424$, achieved when the attacker uses 0.858 bangs per attacked target after soaking up the single ABM, a total of 1.358 bangs of payload per attacked target. With only a small total payload, say $B = 1$, the attacker has no choice but to attack one target by soaking up its ABM and then assigning 0.5 bangs, which is inefficient. On the other hand, the return per attacker is maximal when B is nearly a multiple of 1.358, or when B is large enough to allow some flexibility in the number of targets attacked. Once B is large enough to make attacking all targets optimal, the defense is no longer dominant and the concavity of the value functions causes the return per attacker to begin decreasing, as in Figure 2.

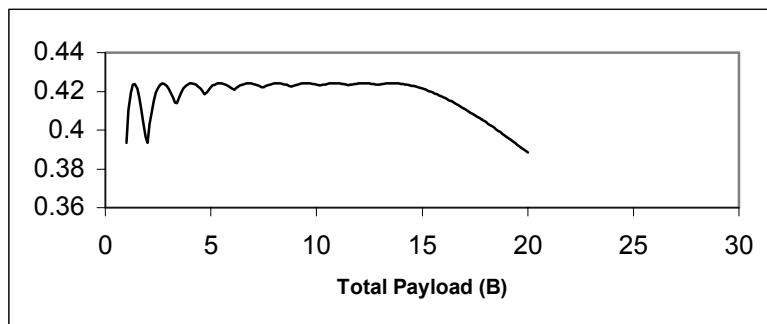


Figure 2: Return per attacker in a problem with 10 identical targets. The defense is dominant for $B \leq 13$.

The actual optimization problem faced by the attacker for a given defense is separable, but the presence of a defense complicates the solution because of the “buy-in” cost of soaking up the interceptors. A variety of methods are available (see Ibaraki and Katoh, 1988) for solving such problems, but, since our object is merely to show that the results of this optimization will be close to the upper bound when the defense is dominant, we content ourselves with (4).

Fixed Arsenal analysis

So far we have been assuming that the attacker is free, after observing the defense, to convert bangs into soaks at the rate λ . In this section, we assume instead that the attacker possesses a fixed arsenal of RVs, each with specified numbers of bangs in the interval $[\beta_{\min}, \beta_{\max}]$. For example, the attacker might have 10 RVs with a total of 8 bangs, the bang vector being $\beta=(0, 0, 0, 0.4, 0.8, 0.8, 1, 1, 2, 2)$. Here $\beta_{\min} = 0$ and $\beta_{\max} = 2$. The first three RVs are decoys, with the rest listed in order of increasing bang measure. The problem is otherwise as described in the introductory section, with a fixed number D of perfect ABMs and a value function for each target. The bang vector is assumed known to the defense. Let the min max value killed by the attacker be $M(\beta, D)$.

The attacker still has the problem of overcoming defenses, and should use the least powerful RVs, beginning with the decoys, to do so. A given RV can be used either for its bangs or its soaks, but not both, so there is a tradeoff between the two measures. The efficient frontier can be shown as a concave relationship between bangs and soaks. Figure 3 shows such a diagram for the example above, together with one of many possible tangents. The point of tangency is one where the first 8 RVs are used as soaks, with the last two contributing 4 bangs. As before, we assume that the bangs can be distributed arbitrarily over the targets, even to more than 2 targets if the attacker desires.




Figure 3: The tradeoff between bangs and soaks in an example with 10 RVs.

Consider a relaxation of the attacker’s problem where the efficient frontier is replaced by one of its tangents. Since the tangent lies entirely above the efficient frontier, any solution of the relaxed problem is an upper bound on the total value that the attacker can kill. Let the expression of the tangent curve be $b + \lambda s \leq B(\lambda)$, where $B(\lambda)$ is the vertical intercept and $-\lambda$ is the slope. Since this is the same form as the “free” problem considered earlier, it follows that $M(\beta, D) \leq B(\lambda)r(\lambda, D)$. Since λ is an arbitrary number in the interval $[\beta_{\min}, \beta_{\max}]$, it follows that

$$M(\beta, D) \leq \min_{\beta_{\min} \leq \lambda \leq \beta_{\max}} B(\lambda)r(\lambda, D) \quad (5)$$

Formula 5 is the desired upper bound on the min max value in the constrained case.

Example 3: This is the same as example 1 except that there are three small and three large targets, $D = 12$, and the bang vector β is as given above and shown in Figure 3. The minimizing value for λ is 1, which corresponds to a defense where large (small) targets are defended by 3(1) ABMs. Since $r(1, D) = 0.348$ and $B(1) = 12$, $M(\beta, D) \leq 4.176$ according to (5). Against that defense, the best attack is against 1 small

and 2 large targets, which requires 7 soaks and leaves the 3 largest RVs. Distributing the 5 associated bangs optimally over the three targets produces a total value killed of $M(\beta, D) = 4.101$, slightly smaller than the bound.

Given the lack of variety in the targets, this comparison between value and bound could be made worse by postulating a number of defenders that is not a multiple of 3. Suppose $D = 14$ instead of 12. The computation of $r(\lambda, D)$ does not change, since it is based on the target that is easiest to attack, so neither does the bound. There are 2 ABMs left over whose use is not guided by theory. If these two are allocated to large targets, so that two of the large targets have 4 ABMs each, then the attacker should ignore both of those targets, attacking the other large target and the three small ones. This will achieve 3.938 in total, while the bound is still 4.176. If the leftover ABMs are used in any other way, the attacker can still achieve 4.101, so the two leftovers should indeed be allocated to large targets. If $D=13$, then neither the bound nor $M(\beta, D)$ changes; that is, the thirteenth ABM is useless against the postulated attack.

An extended example of the free type

We consider here an extended example, assuming a defense dominant scenario with the United States being the defender. The threatened nuclear attack is to be entirely against urban areas, with the goal being to cause as many casualties as possible. We assume that each equivalent megaton will cover 58 km^2 lethally (SPARTA (2001)), so the attacker should seek to attack urban areas with a high population density.

Table 1 shows the population and area of the five largest Continuously Built up Urban Areas (CBAU) in the United States according to Demographia (2003). We assume that the population is uniformly distributed over the stated area; that assumption

is significantly wrong, but still sufficient for our purposes here. We also assume that all CBUAs are within reach of the attackers. If an undefended CBUA has population P and area A in km^2 , then the value function when x EMT are allocated is $V_{LIN}(x) = P \min(1, 58x/A)$. This (truncated) linear function incorporates the implicit assumptions that edge effects and the difficulty of covering a region with circles can be ignored, and that RVs are perfectly accurate. For comparison we also analyze the exponential value function $V_{EXP}(x) = P(1 - \exp(-58x/A))$. The exponential case assumes a poorly coordinated attack where the area covered by x EMT is randomly distributed within the CBUA. The two cases should bracket the practical case where RVs are aimed in the presence of targeting errors. In general $0 \leq V_{EXP}(x) \leq V_{LIN}(x)$.

Assume first that there are 100 ABMs ($D=100$), and that the linear case applies. Using (2) and (3), we find that $r(0.5, D) = 110,959$ and $r(0.25, D) = 120,317$, with corresponding ABM distributions as shown in Table 1. If there were no ABMs, the largest number of casualties per EMT would be 158,278, obtained by attacking Los Angeles. Thus, the presence of ABMs has the effect of reducing the number of casualties per EMT available to the attacker. As λ decreases, the ABMs tend to be concentrated on the most vulnerable target, and the overall effect of the ABM defense decreases because it becomes increasingly easy to soak them up with decoys. The total number of casualties, of course, depends on the unspecified magnitude of the attack.

If λ is held constant at 0.5 while D is varied in the linear case, the dotted curve in Figure 4 results. By the time $D=3000$, the first 61 CBUAs are defended, rather than only the first 5. Those 61 include Washington, DC and finally Atlanta. It takes about 1000

ABMs to reduce the casualty rate by a factor of 2, compared to no defense. Further increases in the number of ABMs have increasingly small effects.

| CBUA | Population | Area (km ²) | ABMs |
|------------------------|------------|-------------------------|---------|
| Los Angeles | 11,789,000 | 4,320 | 64 (94) |
| San Francisco-San Jose | 5,479,000 | 2,582 | 10 (5) |
| Miami | 1,915,000 | 914 | 3 (1) |
| New York | 17,800,000 | 8,684 | 22 (0) |
| New Orleans | 1,009,000 | 513 | 1(0) |

Table 1: The optimal distribution of 100 ABMs in the linear case over Continuously Built Up Areas in the United States when each decoy displaces 0.5 EMT (0.25 EMT).

The exponential case is shown by a solid line in Figure 4. The linear and exponential value functions have the same initial slope, so both cases have the same maximum casualties per EMT when there are no ABMs. As the number of ABMs grows, so does the difference in casualty levels between the two cases. For any given level of interceptors, the attack is more spread out over CBUAs in the exponential case because of decreasing returns to the attacker at targets with high population density. The defense is spread out similarly; when D=100, fifteen cities are defended, as opposed to 5 in the linear case, and only 45 ABMs defend Los Angeles, as opposed to 64 in the linear case.

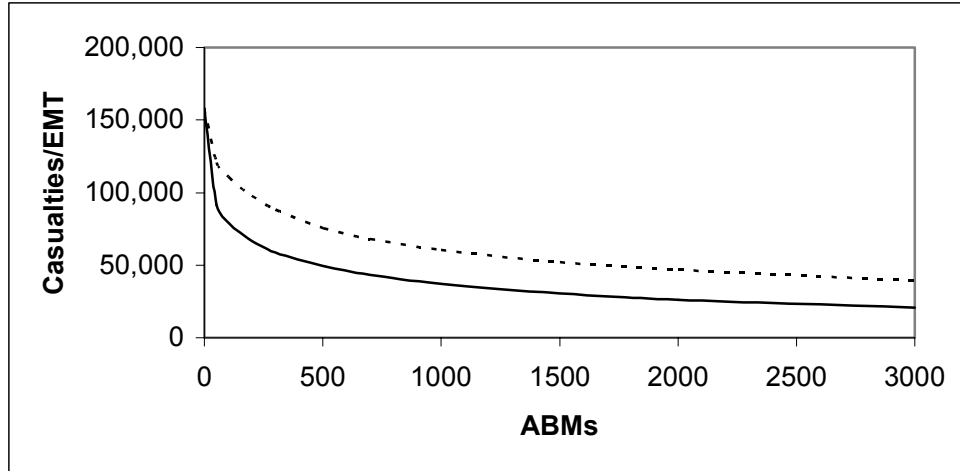


Figure 4: Showing the effect of increasing numbers of ABMs when $\lambda=0.5$.

Solid curve is the exponential case, dotted curve is the linear case.

The introduction of ABMs, even thousands of them, succeeds in lowering the attacker's return by only a factor of 2 or 3 over what it would be without them. This low effectiveness is not due to imperfections in the ABMs, since they have all been assumed to function perfectly. It is due, instead, to a profusion of attractive targets, together with the rule that (terminal) ABMs can only protect the target to which they are assigned. A layered defense that includes flexible, midcourse interceptors would be more effective, assuming that the midcourse interceptors could not themselves be easily decoyed.

In the free case, the ultimate (bound on the) damage caused by the attacker is $Br(\lambda, D)$, which is of course proportional to B , the payload that attacks the terminal defenses. The effect of any midcourse interceptors is essentially subtractive on B . Suppose, for example, that there are I midcourse interceptors, that effective midcourse decoys cost μ each in terms of payload, with μ presumably less than λ because exoatmospheric ABMs can be light, and that the initial payload is A . Then, assuming that

the interceptors get soaked up by decoys, $B = A - \mu I$. Damage to the targets would be additionally reduced, of course, if midcourse interceptors could discriminate RVs from decoys, or even if they could attack RVs at random.

The reader who wishes to vary the assumptions made above, or to investigate the computations lying behind Figures 1-4, may wish to download *BangSoak.xls*, a workbook available at <http://diana.or.nps.navy.mil/~washburn/>. That workbook was used in making most of the computations reported above.

Summary

We have introduced a method for quickly approximating the effects of perfect terminal ABM systems against a mixed attack, under the assumption that the attacker is in control of *which* RVs are intercepted by the ABMs. The method is simple enough to be implemented in a spreadsheet. The effect of terminal defenses is highly dependent on the number of decoys present, or, in the free case, on the payload displaced by a decoy.

References

Missile Defense Agency. 2004 “Ballistic Missile Defense System Overview”, downloaded from <http://www.acq.osd.mil/mda/mdalink/pdf/bmdsbook.pdf>, October, 2004.

Bracken, J.; Brooks, P.; and Falk, J. 1987. “Robust Preallocated Preferential Defense”, *Naval Research Logistics* **34**, pp. 1-22.

Burr, S.; Falk, J.; and Karr, A. 1985. “Integer Prim-Read Solutions to a Class of Target Defense Problems”, *Operations Research* **33**, #4, pp. 726-745.

Bush, G. 2002, press briefing downloaded from <http://www.whitehouse.gov/news/releases/2002/12/20021217.html>, October, 2004.

Current News. 1976. Issue #142, 23 September.

Demographia, 2003. downloaded from <http://www.demographia.com/db-intlua-cover.htm>, October, 2003.

Feickert, A. 2004. “Missile Survey: Ballistic and Cruise Missiles of Foreign Countries”, Congressional Research Office, download from <http://www.fas.org/spp/starwars/crs/RL30427.pdf>, October, 2004.

- Ibaraki, T., and Katoh, N. 1988. *Resource Allocation Problems*, MIT Press.
- Karr, A.F. 1981. "Nationwide Defense Against Nuclear Weapons: Properties of Prim-Read Deployments". P-1395. *Institute for Defense Analyses*, Alexandria, VA.
- Miercort, F., and Soland, R. 1971, "Optimal Allocation of Missiles Against Area and Point Defenses", *Operations Research* **19**, #3, pp. 605-617.
- Eckler, A. R. and Burr, A. 1972. *Mathematical Models of Target Coverage and Missile Allocation*, Military Operations Research Society, chapter 5.
- Read, W. T., Jr. 1958. "Tactics and Deployment for Anti-missile Defense", Bell Telephone Laboratories, Whippany, N.J.
- SPARTA, Inc. and DFI International, 2001. "A Study for The Defense Threat Reduction Agency Advanced Systems and Concepts Office", p. 26.
- Washburn, A. 1981. "Note on a Constrained Maximization of a Sum", *Operations Research* **29**, #2, 411-414.
- Washburn, A. 2003. *Two-Person Zero-Sum Games*, INFORMS Topics in OR series, chapter 6.