

INFORMING AND CALIBRATING A MULTIREOLUTION EXPLORATORY ANALYSIS MODEL WITH HIGH RESOLUTION SIMULATION: THE INTERDICTION PROBLEM AS A CASE HISTORY

Paul K. Davis
James H. Bigelow
Jimmie McEver

RAND
1700 Main Street
Santa Monica, CA 90407-2138, U.S.A.

ABSTRACT

Exploratory analysis uses a low-resolution model for broad survey work. High-resolution simulation can sometimes be used to inform development and calibration of such a model. This paper is a case history of such an effort. The problem at issue was characterizing the effectiveness, in interdicting an invading army, of long-range precision fires. After observing puzzling results from high-resolution simulation, we developed a multiresolution personal-computer model called PEM to explain the phenomena analytically. We then studied the simulation data in depth to assess, adjust, and calibrate PEM, while at the same time discovering and accounting for various shortcomings or subtleties of the high-resolution simulation and data. The resulting PEM model clarified results and allowed us to explore a wide range of additional circumstances. It credibly predicted changes in effectiveness over two orders of magnitude, depending on situational factors involving C4ISR, maneuver patterns, missile and weapon characteristics, and type of terrain. The insights gained appear valid and a simplified version of PEM could be used for scaling adjustments in comprehensive theater-level models.

1 INTRODUCTION

1.1 Background

An important operational challenge for the U.S. military is being able to halt an invading armored force *early*. Being able to do so is plausible, for some scenarios of considerable interest, if U.S. forces acquire the appropriate weapon systems, organization, and doctrine. Doing so would be part of the larger effort to “transform U.S. military forces” as called for in the Quadrennial Defense Review (Cohen 1997, see also Davis et al. 1998).

Meeting the challenge will not be easy and it is by no means clear as yet how successful ongoing developments and efforts to change organization and doctrine will be. Analysis can contribute to understanding what is needed and how the effectiveness of interdiction efforts would vary with scenario. This paper, based on a much more lengthy report (Davis, Bigelow, and McEver 2000), and building upon a preliminary presentation (Bigelow, Davis, and Bigelow 2000), is a case history of how one can use a family-of-models approach that includes multiresolution exploratory-analysis models and a pre-existing high-resolution simulation.

1.2 The Analytical Problem

The idealized military problem here is one in which an invading armored force with many armored fighting vehicles (AFVs) pours across a border along one or a few major roads and moves rapidly toward an objective that may be some 100s of km from the border. An important issue is whether a defender could halt that invading army quickly by interdicting it with long-range precision fires in the form of aircraft with precision weapons such as JDAM or JSOW, and long-range missiles such as the Army ATACMS firing the “brilliant” munition BAT. The answer, of course, should depend upon the number of forces immediately present, deployment rates, weapon effectiveness, the size and speed of the advance, and so on. Clarifying those and other dependences is a natural task for exploratory analysis.

An exploratory analysis in 1997 highlighted some of the key issues for desert circumstances (Davis and Carrillo 1997), some of which were examined in considerable detail the next year (Ochmanek et al. 1998). It was evident, however, that mixed-terrain cases would be different—although it was widely believed that the differences would not be dramatic because roadways are often open. More detailed modeling was needed to understand the issues.

In fact, RAND has long used a detailed, entity-level force-on-force suite of models in defense studies. A number of these studies, led by Randall Steeb and John Matsumura, have considered long range precision fires, as well as upgraded ground forces, in scenarios dealing with defense against an armored invasion. In work done for the 1996 and 1998 Summer Studies, and for the Army, RAND found (Matsumura et al. 1997, Matsumura et al. 1999) very different effectiveness for long range fires of the ATACMS/BAT combination. In simulations for the DSB '96 summer study, ATACMS/BAT killed about three Red vehicles per missile, which was already lower than many individuals expected of this weapon. So it was a most unwelcome surprise when, in simulations for the DSB '98 summer study, the same weapon killed a factor of 5-10 fewer Red vehicles per missile.

There were some possible explanations. In DSB '96 the terrain was entirely open, while the DSB '98 terrain had a good deal of tree cover. In DSB '96 almost all the Red vehicles were armored fighting vehicles (AFVs), while in DSB '98 fewer than 20 percent of the Red vehicles were AFVs. And in DSB '96 the Red vehicles were in dense formations (50–100 meter spacing), while in DSB '98 vehicles were much more dispersed.

These arguments might seem to rationalize results, but counter arguments suggested caution. For example, in DSB '98 the missiles were aimed only at clearings, and only when the human-in-the-loop targeter (an Army officer trained in ATACMs doctrine) projected AFV arrivals based on simulated C4ISR information. One could therefore argue that tree cover should reduce the number of missiles launched, but not the effectiveness per missile. Also, the BAT submunition preferentially homes in on AFVs, so the presence of trucks should make little difference. And the large footprint of the ATACMS/BAT (a radius of at least four kilometers) should negate the large separations between vehicles. It seemed clear that we should not be satisfied with glib rationales, but should instead study the issues more carefully.

Doing so was not straightforward. Although the entity-level simulation is a rich description of phenomenology and has clear physics algorithms and a rooting in "hard" weapon data, the analytical implications for operational-level effectiveness are often difficult to understand because of the simulation's bottom-up character and huge number of variables. Further, using the simulation is manpower-intensive; it is not practical to consider a wide range of scenarios (although many variations can be made in, say, weapon Pks, detection probabilities, and entity characteristics).

To try to understand better the higher-level issues while connecting them with the microscopics, we built a personal computer model called the PGM Effectiveness Model (PEM) (Davis, Bigelow and McEver, 2000). This was an intermediate-level model in that, in *some* respects,

it had low resolution in comparison with RAND's entity-level simulation suite, but higher resolution than models used for theater-level analysis. PEM was focused on a small part of the overall problem: effectiveness of long-range fires in interdicting a particular group of AFVs amidst terrain. PEM is not small enough to be written on the back of an envelope, but it is nonetheless quite small and simple. The conceptual core is based on simple physics. We implemented it in Analytica™, a very flexible visual modeling tool.

One can think of PEM as a stochastic, physically and mathematically-based, scaling model (not a mere statistical fit) that adjusts the effect of long range precision fires for the influence of a variety of factors. These factors include the time of last update, which operates through the error in the missile arrival time; the footprint of the weapon; the openness of the terrain; and the density of the Red formation. We have used the model to investigate interactions among the factors. We have also developed an even simpler deterministic version of the model that could be used as a subroutine to incorporate these factors in other models, such as EXHALT (McEver, Davis, and Bigelow 2000), RAND's JICM, or even DoD's emerging JWARS model.

1.3 Relating the Work to Generic Issues

Our work is an example of multi-resolution modeling (MRM) (Davis and Bigelow 1998), which is the practice of building mutually consistent models or families of models. PEM itself has multiple levels of resolution, which are related cleanly through hierarchical design. RAND's high-resolution simulation suite was developed years ago and is by no means integrated with PEM. However, we could do family-of-models work by investing the time necessary to understand relationships and accomplish some calibrations. Doing so would illustrate common difficulties in working with legacy models.

Multiresolution modeling is important for many reasons, the most fundamental of which is perhaps the need of humans to reason at different levels of detail. Such reasoning requires variables that can be manipulated (i.e., inputted) at those different levels to discuss cause and cause-effect relationships. This implies the need for models at different. Even excellent high-resolution models do not meet this need. Other reasons for MRM and its cousin multiresolution, multiperspective modeling (MRMPM) (Davis 2000) involve tradeoffs between agility and phenomenological richness, the need to connect to different levels of empirical data, costs, time, and the treatment of uncertainty. For example, we need low-resolution models for exploratory analysis, but we need high-resolution models to understand underlying phenomena, to provide links to physical entities and

concrete low-level options, and, sometimes, to calibrate the lower-resolution models.

With this background, let us now proceed as follows. Section 2 describes the models we used for our analysis; Section 3 describes our analysis of high-resolution data; and Section 4 draws some lessons learned. Davis developed the PEM model; Bigelow did the extensive data analysis reported here; and McEver developed a simplified “repro model” version of PEM.

2 THE MODELS

2.1 RAND’s Force-on-Force Modeling Suite

The high resolution models that produced the provocative results motivating our study provide RAND with a valuable capability for high fidelity analysis of force-on-force encounters. In this suite, the RAND version of JANUS (a model originally developed by the Lawrence Livermore Laboratories) serves as the primary force-on-force combat effectiveness simulation and provides the overall battlefield context, modeling as many as 1500 individual systems on a side. The Seamless Model Interface (SEMINT) integrates JANUS with a host of other programs into one coordinated system, even though the participating models may be written in different programming languages, and run on different hardware under different operating systems. In effect, SEMINT gives RAND the ability to augment a JANUS simulation with specialized high fidelity models, without modifying the basic JANUS algorithms. The result is distributed and sometimes interactive simulation (DIS) for analysis, although the models in this case are all located in the same laboratory. The system was developed prior to the High Level Architecture (HLA) that is becoming the DoD standard geographically distributed work.

As currently configured, JANUS conducts the ground battle, calling on the RAND Target Acquisition Model (RTAM) to provide more accurate calculation of detection probabilities of special low observable vehicles. The Model to Assess Damage to Armor by Munitions (MADAM), developed originally by the Institute for Defense Analyses, simulates the effects of smart munitions, including such aspects as chaining logic, multiple hits, and unreliable submunitions, while the Acoustic Sensor Program (ASP) provides a detailed simulation of acoustic phenomenology for such systems as air-delivered acoustic sensors and wide-area munitions. Should the conflict involve helicopter or fixed wing operations, the flight planners BLUE MAX II (fixed wing) and CHAMP (helicopter) determine flight paths for the missions, flown against the actual JANUS threat, and RAND’s Jamming and Radar Simulation (RJARS) conducts the defense against the aircraft, including detection, tracking, jamming and SAM operations. The

Cartographic Analysis and Geographic Information System (CAGIS), developed originally at the Johns Hopkins University Applied Physics Laboratory, provides consistent geographic information to all the simulations, while SEMINT passes messages among the models, and maintains a Global Virtual Time to keep the models in synchronization.

For our purposes, the Model to Assess Damage to Armor by Munitions (MADAM) is key. RAND has upgraded MADAM so that it models the technologies associated with the following munitions:

- Seek And Destroy ARMor (SADARM)
- Sensor-Fused Weapons (SFW-Skeet)
- Damocles
- Low-Cost Anti-Armor Submunition (LOCAAS)
- Terminally-Guided Weapon/Projectile (TGW/TGP)
- Precision Guided Mortar Munition (PGMM) (Infra-Red (IR) & Millimeter Wave (MMW))
- Brilliant Anti-Tank (BAT)
- Wide Area Munitions (WAM)

The model simulates target seeking logic, false alarm rates, hulks, submunition reacquisition, shots, hits and kills, as well as bus, munition, and submunition reliability. For example, to estimate how many vehicles are killed by a BAT, MADAM simulates the separation of the bus from the launch vehicle, the separation of submunitions from the bus, several stages of acoustic seeking and deployment by the submunitions as they descend, an IR detection stage and a final shot/hit/kill event for each submunition. The outcome at each stage is determined, in part, by a random draw.

2.2 High Resolution Study of Long Range Fires

RAND has used this suite of models as follows to study long range precision fires (Matsumura et al. 1997, 1999). JANUS simulates the movement of each vehicle in a Red force across a terrain. The analyst scripts this movement by specifying the initial position and nominal velocity of each vehicle, as well as the path the vehicle will follow. For a road march, the path is a road. If Red is attacking a Blue position, the path will include off-road maneuver.

Periodically, say every five minutes of simulated time, a snapshot of Red vehicle positions is provided to a man-in-the-loop who decides the aim points and impact times of the long range precision weapons (sometimes it is possible to automate this function). Each snapshot provides incomplete information on the positions of Red vehicles. If a vehicle is obscured by foliage, there is a probability P_{TREE} of seeing it. If the vehicle is in the open, the probability P_{OPEN} of seeing it is larger. These probabilities can be adjusted to represent different qualities of C4ISR.

Based on the vehicles he sees, the man-in-the-loop selects aim points and impact times for the long range weapons. He will aim only at open areas, because we have assumed that a vehicle obscured by foliage or hidden in a town or city is not vulnerable. In addition, the time between the snapshot and the impact time of any salvo based on that snapshot must be at least as long as a specified latency period, which includes the time to collect the information in the snapshot, plus a decision time, plus the flight time of the weapon. In DSB '98, the man-in-the-loop would identify a group of vehicles to shoot at, estimate how long it would take to arrive at a nearby clearing, and lead the target as a duck hunter would lead his flying prey.

Finally, MADAM simulates the effect of the weapon on the Red vehicles near its aim point at its time of impact.

2.3 PGM Effectiveness Model (PEM) Concepts

Figure 1 illustrates the concepts on which PEM is built. PEM assumes that a column of Red vehicles is traveling along a road through a clearing of width W . Rather than being uniformly spaced, the Red vehicles are grouped into packets, perhaps representing platoons. Each packet has N AFVs separated from one another by a distance S . Successive packets (not shown) are separated by a distance P , which is larger than S . This column of vehicles moves through the clearing at a velocity V .

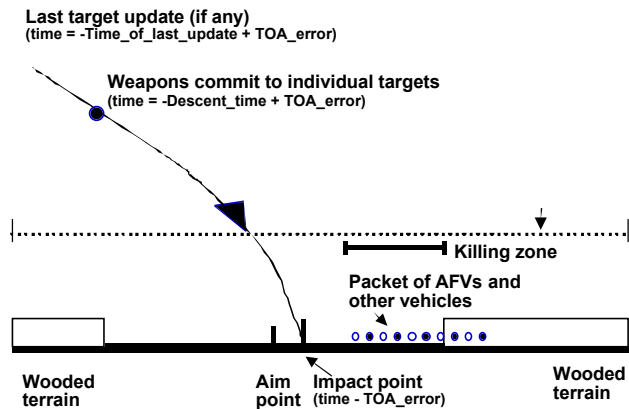


Figure 1: PEM Concepts

Blue attacks the column by firing a salvo of one or more missiles at the clearing, timed to arrive when a selected packet is expected to be in the center of the clearing. There is a random error in impact point and, more significant, a random error in the arrival time (TOA_error), the standard deviation of which is proportional to the time since the missile last received information about the projected position of the target packet ($Time_of_last_update$). If TOA_error is too large, the target packet may have passed completely through the

clearing, or beyond the weapon's footprint F , whichever is larger; or (if the missile arrives early) the target packet may not have entered the clearing or the footprint. A smaller error will find the target packet not centered in the clearing, and part of it may be hidden in the trees (or urban clutter) on either end of the clearing.

Depending on the various parameters, parts of the packets just forward and rearward of the target packet may be in the killing zone. The presence or absence of these neighboring packets can change in kills per salvo by a factor of 2 or 3. The size of this effect depends on the separation between packets and the standard deviation of TOA_error .

Except for a few details, this determines how many Red AFVs are in the killing zone of the weapon at its time of impact. PEM assumes the number of AFVs actually killed will be a specified fraction of the vehicles in the killing zone, up to a specified maximum. Actually, this was initially a hypothesis to be tested by comparing with the high-resolution data. As we shall see, truth is more complex.

3 CALIBRATING PEM TO JANUS/MADAM

A major part of our effort was data analysis—treating simulation data very much like experimental data (including recognition that the experimental conditions were sometimes not what they at first appeared to be, that the experiment [simulation model] was imperfect, that some of the data was flawed, and that not enough information was retained to determine all the causes and effects).

In what follows, we discuss four aspects of calibrating PEM to the high resolution models. First, we determine the lengths of clearings to use in PEM. In the high resolution model, this corresponds to selecting candidate aim points for the long range precision weapon. Next, we determine the Red order of march, i.e., the PEM parameters of vehicles per packet, vehicle and packet separations, and vehicle velocity. Third, we estimate the parameters that determine the missile arrival time error. In the high resolution model, these two aspects of calibration correspond to identifying a lucrative group of vehicles to target, and estimating when the group will arrive in a clearing. Fourth, we estimate weapon effectiveness.

3.1 Lengths of Clearings

A clearing is a basic PEM concept entirely characterized by its width. But no such concept exists in JANUS/MADAM: users of the model viewing displays of simulated behavior "see" clearings that are consequences of the microscopic terrain data bases, but they are visual abstractions, not something built in. One of our most

interesting discoveries (something weapon engineers have undoubtedly come across over the years, but something to which we were previously insensitive) was that the definition of a clearing depends on point of view. The man-in-the-loop sees a snapshot of Red vehicle positions from the point of view of a long range reconnaissance device, perhaps a UAV or J-STARS orbiting a hundred kilometers or more from the target area. Its field of view must be wide enough to take in a large portion of the Red formation. Thus, the man-in-the-loop may miss small clearings, but may also fail to see that what appear to be unbroken open areas are in fact cluttered with small stands of trees or villages.

By contrast, when the weapon arrives over the aim point, it sees the local terrain at much higher resolution. What the man-in-the-loop thought was a clearing may be chopped up into very short stretches of open road (middle of Figure 2). What the man-in-the-loop thought was an unbroken stretch of trees may contain a long, open corridor (lower left of Figure 2). Such corridors along roads are common, though for different reasons, in both the real world and in the model.




Figure 2: A Clearing at High Resolution

Of course, whether vehicles in an open area are really vulnerable depends also on weapon characteristics, including their tolerance for “clutter” in the target area.

The ambiguity in how to define a clearing in the JANUS/MADAM simulations means that we cannot calibrate PEM simply by measuring open spaces in the high resolution representation of the terrain. PEM needs the distribution of open-area widths as seen by the arriving weapon—i.e., the distribution of small open-area widths (the small white areas in Figure 2). Since measuring these areas precisely is tedious, the distribution of widths considerable, and the precise distribution dependent on the particular physical area of battle, we calibrated PEM only approximately using a triangular distribution. The mode of that distribution for open-area width is a variable parameter in exploratory analysis.

3.2 Red Order of March

In PEM, the Red order of march is specified by four parameters: the number of AFVs in a packet, the distance between successive vehicles within a packet, the separation between packets, and the speed of vehicles across clearings. In any particular PEM run, all these parameters are constant, though we investigate their influence by changing them stochastically or parametrically from one run to another. Thus PEM assumes the Red formation is highly regular, described by only a few parameters. In the high resolution cases, the Red formations are rather irregular, so we must approximate.

In the DSB ‘98 study, the Red columns deviated from this simple description in several ways. First, there were three types of packets. One consisted of AFVs, the other two of scout vehicles and (mostly) trucks, respectively. In total, only 104 of the 543 Red vehicles in the simulation were AFVs. Second, not all AFV packets had the same numbers of vehicles separated by the same distances. AFV packets had from three to ten vehicles (average 6.7), and were separated by 150 to 600 (average 350) meters. And successive packets of AFVs were separated by from 1000 to 3800 meters. The speed of AFVs, however, was nearly constant at 76 kilometers per hour. This was a deliberately stressful case for long-range fires, much more so than in most previous RAND work.

The three kinds of packets in the DSB ‘98 summer study moved at different speeds. Reconnaissance, AFV, and truck packets began the simulation in overlapped positions, but as the simulation progressed the reconnaissance packets pulled ahead and the truck packets fell behind. All packets of the same kind moved forward in lockstep at a nearly constant speed.

In the DSB ‘96 summer study, the Red order of march was quite different. First, virtually all vehicles (458 out of 504) were AFVs. Second, they did not move in column formation for the whole simulation. Rather, they moved in columns for roughly the first 30 minutes of the simulation, and then redeployed for an attack on a Blue position. As the Red force redeployed into attack formation, vehicles were still densely packed. However, since they were no longer following the roads, it became harder to predict where they were going, and hence to lead them with long range fires. On the other hand, once the Red vehicles left the roads, they slowed. About two-thirds of the salvos were aimed at Red vehicles in column formation and the remainder at vehicles in attack formation.

While they were in column formation, the AFVs were packed much more densely than in the DSB ‘98 cases. Separations of 50 to 100 meters between AFVs were typical (even these separations, of course, are higher than in most historical battles). The groups of AFVs we might identify as packets for PEM often contained 50 vehicles or more.

3.3 Error in Weapon Arrival Time (TOA_error)

In PEM, the error in the weapon arrival time is the difference between the time the weapon arrives and the time the target packet is centered in the clearing. A negative TOA_error indicates that the weapon has arrived early, and a positive error that it has arrived late. We have assumed that the error is random with nearly zero mean and a standard deviation that is proportional to the time of last update.

The time of last update is the last time the shooter has the opportunity to adjust the aim point or impact time of the weapon. In the DSB '98 study, the man-in-the-loop had to specify the aim point and impact time of each ATACMS based on information that would be, for some cases at least 11 minutes old, and for other cases as much as 20 minutes old, at the time of impact.

The standard deviation of the TOA_error, as mentioned earlier, is proportional to the time of last update; and the constant of proportionality is the fractional error in the shooter's estimate of the speed of the AFVs along the road or track. The fractional speed error must have been considerable in the DSB '98 cases, since most ATACMS missile salvos found few AFVs in their footprints when they arrived over their aim points. Some of this was due to the fact that the shooter misguessed the route a group of Red vehicles would take, i.e., thought they would turn right instead of left. But most missiles fell on the routes that the Red columns actually did follow. Note, that the error in movement rate would not be remedied by a more accurate radar measuring instantaneous velocity, because the issue is estimating future speed along a curved and sometimes complex road.

We could not estimate the fractional-speed or time-of-arrival error directly, because the man-in-the-loop did not keep records on which group of vehicles he was targeting with each shot. But we could do it indirectly. Imagine that a camera floats above the aim point of a particular ATACMS salvo, and counts the number of vehicles in the ATACMS footprint (a circle with a nominal ATACMS kill radius). If we plot that number as a function of time, we obtain a figure such as Figure 3. The upper curve is the count of all vehicles; the lower curve is of AFVs only.

Two things stand out. First, the number of AFVs in the footprint is highly variable and even small deviations from the optimal impact time will reduce the number of targets substantially. Second, there are frequently many vehicles in the footprint but no AFVs. Thus, if the shooter can't distinguish AFVs from other vehicles, he will surely waste a substantial number of shots.

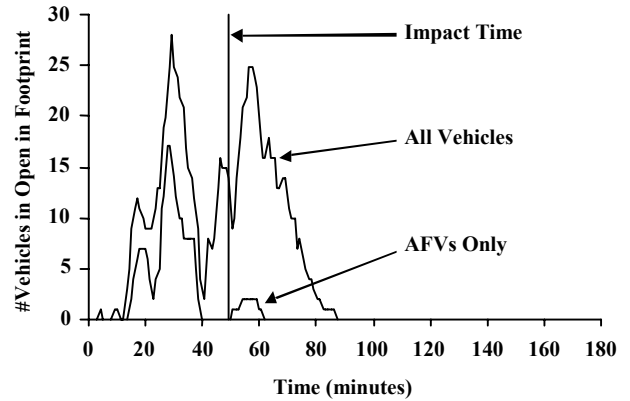


Figure 3: Vehicles in a Typical DSB '98 Footprint

Following up on this last point, we can define two windows of opportunity (WOP). The first is the interval of time from the entry of the first vehicle of any kind into the footprint, to the exit of the last vehicle from the footprint. The second is the time interval similarly defined for AFVs rather than for all vehicles.

Table 1 evaluates the timing of salvos in three DSB '98 cases, with times of last update of 20, 15, and 11 minutes, respectively. In all three cases, virtually all shots fell within the window of opportunities for all vehicles, suggesting that the shooter could assess when some part of the whole column would be in the footprint, even if he could not do so for a specific packet.

Table 1: Measures of Success in Timing of Shots

	Cases		
	X	Y	Z
Delay	20	15	11
	Percent of shots		
In WOP for all vehicles	100	94	100
In WOP for AFVs only	65	56	71
100*WOP(AFV)/WOP (All)	38	39	37
	Avg. AFVs in footprint		
Actual shots	1.85	1.44	3.42
In WOP for all vehicles	1.49	1.61	1.40
In WOP for AFVs only	3.88	4.03	3.73

The shooter was less successful at estimating when the AFVs in the column would pass through the footprint, though he did better than chance: In each case, the window of opportunity for AFVs was about three-eighths of the window of opportunity for all vehicles, but the shooter was able to put more than half of his shots into the AFV window. The reason for this was probably not due to the shooter's ability to distinguish one vehicle type from another, although in principle he had the information to do so. Rather, the shooter had a "template" of the Red column that placed scout vehicles in front, followed by combat vehicles. The template placed support vehicles at the rear.

The shooter targeted the portions of the column that his template suggested would contain the combat vehicles.

Next we compare the number of AFVs in the footprints in the actual shots with the average numbers in the two windows of opportunity. In cases X and Z, the actual number is between the average numbers in the two windows, and it would be in case Y as well, if we eliminated the two shots that missed the larger window altogether. It is plausible to argue, therefore, that the shooter was unable to time his shots well enough to hit individual packets. He could have done as well simply to establish a window somewhere between the all-vehicle and AFV windows we have defined, and picked an impact time at random within his window.

Finally, then, we have an indirect way to calibrate PEM's TOA_error distribution to the high resolution results. From the JANUS/MADAM results we may take the variation over time in the number of AFVs in a footprint (as in Fig. 3), and turn it into a frequency distribution—i.e., we can determine the fraction of time in some window of opportunity that there are zero AFVs in the footprint, or one AFV, or any other number of AFVs. In PEM there is a probability distribution of AFVs in the footprint, and the variation in this distribution is affected by the standard deviation of TOA_error. We need only set the standard deviation so that the PEM probability distribution looks similar to the JANUS/MADAM frequency distribution.

If we place a camera above a typical aim point in the DSB '96 cases, we see a very different picture. In particular, the number of AFVs in the footprint is large (Figure 4) compared with the numbers we saw in the DSB '98 chart. As we shall see shortly, when the number of AFVs in the footprint is large enough, kills per ATACMS/BAT missile reaches a maximum, after which a further increase in the number of AFVs has no effect. For this footprint there is a window of 25 or 30 minutes within which a missile impact should achieve that maximum number of kills. Thus, precise timing is unnecessary against Red formations as dense as those in the DSB '96 study.

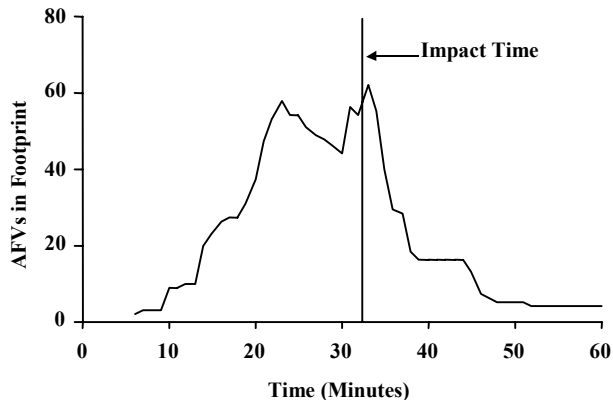


Figure 4: AFVs in a Typical DSB '96 Footprint

3.4 AFVs Killed per Salvo

In PEM, we calculate the number of AFVs killed per salvo as a function of the number of AFVs in the open in the salvo's footprint at the time of impact. The function we used initially was simple. The number of AFVs killed is proportional to the number in the footprint, up to a maximum (reflecting the finite number of submunitions in a BAT). As a result of the analysis below, we ended up adding a stochastic factor.

When we compare the initial, simple representation with actual salvos simulated in the DSB '98 and DSB '96 cases, inevitably we see errors. Part of the error is due to the fact that MADAM performs Monte Carlo trials to estimate kills for a given shot, and PEM's original relationship was deterministic. Another part is due to the fact that several types of vehicles are vulnerable to ATACMS/BAT, but they are not all equally vulnerable. In the high resolution simulations there are also such factors as background noise and dead vehicles that can interfere with the performance of the BAT submunition. Indeed, it is because of such terrain-and-case-dependent factors that the numbers we present here should be considered illustrative and unclassified. It is also why we consider PEM rather more of a scaling model than as a self-contained complete model. It cannot fully substitute for higher-resolution work.

One of the more interesting sources of error lies in the fact that BAT has differing effects against vehicles laid out in different patterns. ATACMS/BAT is particularly effective against a linear pattern of vehicles. Depending on the search algorithm assumed, it may be much less successful against a pair of crossed lines of vehicles. The crossed pattern confused the baseline search algorithm assumed for the BAT submunitions, and most of them fell harmlessly between the two lines. We emphasize, however, that alternative search algorithms exist and can also be used. The point here is that such details can matter a good deal. Such factors can only be understood with high-resolution work, not something as simple as PEM.

Whatever the reason, the number of AFVs killed per salvo is highly variable, even when one controls for the number of AFVs in the footprint. Figure 5 shows the variation across all the hundreds of simulated shots from the DSB '98 study that had exactly five AFVs in the footprint.

Figure 6 plots the number of AFVs killed per salvo versus the number of AFVs in the footprint. We have included data from both the DSB '98 and DSB '96 studies. The points from the DSB '98 study have a maximum of 14 AFVs per footprint. Each point represents the average kills from all shots with the same number of AFVs in their footprints. Points from the DSB '96 study include salvos with very large numbers of AFVs in their footprints, and each point represents a single salvo of two ATACMS/BAT missiles. The solid line represents a plausible relation to use in PEM for

calculating AFVs killed from AFVs in the killing zone. We fit it by eyeball, not by statistical methods, and ignored some points we had reason to believe were artifacts.




Figure 5: Variation in Kills for Salvos with Five AFVs in the Footprint




Figure 6: Kills per Salvo vs. AFVs in Footprint

Up to about eight AFVs in the footprint, the average number of AFVs killed per salvo is approximately 40 percent of the number of AFVs in the footprint. This portion of the data is dominated by results from the DSB '98 study. When there are more than eight AFVs in the footprint, the trend line begins to flatten out. We knew it had to do so, since a salvo of two ATACMS dispenses a finite number of BAT submunitions, each of which is capable of a maximum of one kill. This portion of the data consists entirely of results from the DSB '96 study.

In the current version of PEM, the relationship between AFVs in the footprint and AFVs killed is treated deterministically. This discussion, however, suggests that it would be better to introduce a stochastic component if PEM is used in a context where statistical variation would be of interest. Otherwise, adding another stochastic

component (as distinct from using distributions to reflect uncertainty in parameter values) is unnecessary.

4 LESSONS LEARNED

Many lessons could be learned from this work viewed as a case study. Some of the lessons were military; others were methodological.

4.1 Lessons About Long Range Precision Fires

In a sense, the reason each TACMS/BAT killed more AFVs in the DSB '96 study than in the DSB '98 study is simply that there turned out to be more AFVs in the weapons' footprints. However, building and exercising PEM has taught us *why*: what factors influenced the numbers of AFVs in the footprint, and how they interacted. The substantive lessons were reported in (Davis, Bigelow, and McEver 1999) and (Defense Science Board 1998).

Most of these results are reported elsewhere (Davis, Bigelow & McEver 2000). Briefly, the effectiveness of such fires can vary by two orders of magnitude, depending on the time of last update (which operates through the error in the weapon's time of arrival), the footprint of the weapon, the openness of the terrain, and the density of the Red formation. Moreover, these factors interact. Some examples: (1) if one shoots at small clearings, it becomes unimportant to use a weapon with a large footprint; (2) against low density Red formations it is vital that the weapon arrive at just the proper time; (3) against high density formations the TOA_error isn't important; and (4) if the weapon can loiter, TOA_error will also be less important.

The interaction of these factors affects one's choice of weapons. ATACMS/BAT has a large footprint but also a large TOA_error, while aircraft-delivered sensor fused weapons (SFW) have a small footprint but a low TOA_error. The ATACMS/BAT has an advantage over the SFW against high density Red formations in open terrain (large clearing), but loses its advantage if either Red's density is low or clearings are small.

In retrospect, none of these conclusions is counterintuitive, but before-the-fact intuition had been poor and PEM allowed us to crystallize a better intuition and develop quantitative relationships. Thus, we can estimate just how small the clearing must be, or how low the density of the Red vehicles, before the relative effectiveness of ATACMS/BAT versus SFW falls to any specified threshold. This would permit cost-benefit judgements about when to switch from one weapon to the other, or what weapon mix to use.

4.2 Lessons About Model Families

The exercise of building and calibrating PEM has provided lessons for the practice of building multi-resolution model

families that involve a high-resolution legacy model that cannot readily be integrated with the family's higher level models. First, in any model at any level of resolution, variation of an output quantity from one case to another will be explained in part by variations in input parameters, and in part as the result of random events. Because a high resolution model will have more input parameters than a companion low resolution model, there is more scope to explain output variations by variations in input parameters. To achieve the same degree of variation in the low resolution model, it will often be necessary to introduce a random process. Parameters that were available to explain this part of the variation in the high resolution model, but which are missing in the low resolution model, are called hidden variables.

An example of this is the determination of the number of AFVs killed per salvo. In PEM, this is a function only of the number of AFVs in the footprint, and any variation from this function must be represented by a random process. In the high resolution model MADAM, it is a complicated function of the positions and types of all the vehicles—AFVs or not—in the neighborhood of the aim point, and the background noise generated by vehicles not very near the aim point. By changing the types and positions of vehicles for MADAM, we can produce variations in the number of AFVs killed per salvo, without changing the number of AFVs in PEM's footprint.

A second lesson is that concepts that seem well defined at one level of resolution may be ambiguous at another. The notion of a clearing is well defined in PEM. It has a definite width, and there is nothing more to know about it. But in the high resolution model, clearings must be identified from the description of the terrain and the roads. Is a stretch of road a clearing if trees line the road closely but the road itself is clear? Is it a clearing if it is interrupted by a few very short stands of trees? Must the treeless area be sufficiently large to be identified by a reconnaissance platform from miles away, or only large enough so trees don't interfere with the terminal search algorithm of a submunition?

Third, a high resolution model and its companion low resolution model may differ not only in their level of detail, but in their scope as well. Indeed, limiting the scope of a model as well as its detail is a way of keeping it small and agile. We built PEM to investigate long range precision fires against moving armored columns, and thus omitted all other vehicle types. The JANUS/MADAM suite of models includes the other vehicle types (e.g., trucks), and as we have seen, their presence affected the selection of impact times by the man-in-the-loop. He had to try to distinguish AFVs from other types of vehicles, and he did so with imperfect success.

ACKNOWLEDGMENTS

The authors benefited from collaboration and discussions with colleagues Randy Steeb, John Matsumura, and Tom Herbert, who lead RAND's high-resolution simulation work. The research reported here was sponsored by the Office of the Secretary of Defense.

REFERENCES

- Bigelow, J.H., P.K. Davis, and J. McEver. 2000. Calibrating a policy-level exploratory analysis model with high-resolution simulation data for the interdiction problem. *Proceedings of the SPIE*. Vol. 4026.
- Cohen, W.. 2000. *Report from the quadrennial defense review*. Department of Defense, Washington, D.C.
- Davis, P.K. 2000. Exploratory analysis enabled by multiresolution, multiperspective modeling. In *Proceedings of the 2000 Winter Simulation Conference*.
- Davis, P. K. and J. H. Bigelow. 1998. *Experiments in multiresolution modeling (MRM)*. MR-1004-DARPA, The RAND Corporation, Santa Monica, CA.
- Davis, P. K., J. H. Bigelow, and J. McEver. 2000. *Effects of Terrain, Maneuver Tactics, and C4ISR on the Effectiveness of Long Range Precision Fires: A Stochastic Multi-Resolution Model (PEM) Calibrated to a High Resolution Simulation*. MR-1138-OSD, The RAND Corporation, Santa Monica, CA.
- Davis, P.K., J.H. Bigelow, and J.McEver. 1999. *Analytical methods for studies and experiments on transforming the force*. CB-278-OSD, The RAND Corporation, Santa Monica, CA.
- Davis, P.K. and M. Carrillo. 1997. *Exploratory analysis of the halt problem: A briefing on methods and initial insights*. DB-232, The RAND Corporation, Santa Monica, CA.
- Davis, P.K., D. Gompert, R. Hillestad, and S. Johnson. 1998. *Transforming the force: Suggestions for DoD strategy*. RAND Issue Paper, Santa Monica, CA.
- Davis, P. K. and R. Hillestad. 2000. *Exploratory analysis for strategy problems with massive uncertainty*. Draft in review. The RAND Corporation, January 23.
- Defense Science Board. 1998. *1998 Summer study task force on operations superiority in the 21st century*. Office of the Under Secretary of Defense for Acquisition, Department of Defense, Washington, D.C.
- Matsumura, J., R. Steeb, T. Herbert, M. Lees, S. Eisenhard, and J. Gordon. 1997. *Analytic support to the defense science board. Tactics and technology for 21st century military superiority*. DB-198-A, The RAND Corporation, Santa Monica, CA.

- Matsumura, J., R. Steeb, E. Isensee, T. Herbert, S. Eisenhard, and J. Gordon. 1999. *Joint operations superiority in the 21st century*. Analytic Support to the 1998 Defense Science Board. DB-260-A/OSD, The RAND Corporation, Santa Monica, CA.
- McEver, J., P.K. Davis, and J.A. Bigelow. 2000. *EXHALT: An interdiction model for the halt phase of armored invasions*. MR-1137, The RAND Corporation, Santa Monica, CA.
- Ochmanek, D., Harshberger, E., T. Thaler, and G. Kent. 1998. *To find , and not to yield: How advances in information and firepower can transform theater warfare*. MR-958, The RAND Corporation, Santa Monica, CA.

AUTHOR BIOGRAPHIES

PAUL K. DAVIS is a senior scientist and Research Leader at RAND, and a Professor of Public Policy in the RAND Graduate School. He holds a B.S. from the U. of Michigan and a Ph.D. in Chemical Physics from the Massachusetts Institute of Technology. Dr. Davis is a recipient of the Wanner Memorial Award of the Military Operations Research Society. He is a member of the Naval Studies Board under the National Research Council and has served on a number of studies for the Council, the Defense Science Board, and the National Institute for Standards and Technology. His addresses are <pdavis@rand.org> and <www.rand.org/personal/pdavis>.

JAMES H. BIGELOW is a senior mathematician at RAND. He holds a B.S. from the University of California at Berkeley and a Ph.D. in Operations Research from Stanford University. His research has dealt with medical, environmental, and transportation problems, as well as military problems such as logistics for the Army and Air Force, total force structure costs for all services, and total-Army design. His address is <James_Bigelow@rand.org>.

JIMMIE McEVER is an associate physical scientist at RAND. He holds a B.S. and Ph.D. in physics from the Georgia Institute of Technology and a Masters in Public Policy from Harvard's Kennedy School. Dr. McEver's current research involves applications of information technology in the Army, information security, science and technology issues in organizational settings, and alternative Air Force strategies for regional defense problems. His address is <Jimmie_McEver@rand.org>.