



A Comparison of Tactical Ballistic Missile Defense Engagement Coordination Schemes

Simon Moskowitz, Rebecca J. Gassler, and Bart L. Paulhamus

Dealing with large ballistic missile raids will require an unprecedented level of synergy between dissimilar air defense units because any individual unit will likely be overwhelmed. Shooter selection schemes currently in use include unbidden first launch and unique defended asset assignment. These methods only attempt deconfliction to avoid a “free fire” situation. Past analysis has shown that a scheme exists that yields much better results than these unit-centric schemes. That scheme fits into the distributed engagement decision paradigm of preferred shooter selection. This article uses the APL Coordinated Engagement Simulation (ACES) to contrast the performance of the aforementioned schemes in a realistic simulation environment. The distributions of raid outcomes for each scheme under two configurations, ideal and achievable, are compared. Operating conditions are examined through selecting real-world options and modeling underlying probability distributions, where applicable.

INTRODUCTION

When tactical ballistic missiles (TBMs) are used in the future, they will likely come in a massive raid and be more advanced than in the past. Dealing with such a raid will require an unprecedented level of synergy between dissimilar air defense units because any individual unit will likely be overwhelmed. Continued unit-centric and service-independent doctrine and warfare impede achievement of this goal. Shooter selection schemes currently in use include unbidden first launch and unique defended asset assignment. These methods only attempt deconfliction to avoid a “free fire” situation. In addition to minimizing the number of unintentional overengagements, which these schemes strive for, the more important aspiration of minimizing leakers

(i.e., TBMs that penetrate defenses) must be addressed. Minimizing leakers means minimizing free riders (i.e., unengaged threats) and maximizing depth of fire (i.e., number of salvos separated by kill assessment periods), as well as selecting the engagements that have the greatest chance of success. For area tactical ballistic missile defense (TBMD) systems, reengagement after kill assessment is seldom possible.

Past analysis has shown that a scheme exists that yields much better results than the unit-centric schemes. That scheme fits into the distributed engagement decision (DED) paradigm of preferred shooter selection. In this scheme, individual units estimate their own effectiveness and schedule their own shots,

the information is distributed, and then all units use a common algorithm to arrive at a synchronized engagement assignment solution for the force. Units defer their own shots if they are not preferred.

This article contrasts results using the aforementioned schemes (free fire, first launch, unique defended asset assignment, and DED) in a realistic simulation environment. The distributions of raid outcomes (e.g., leakers, free riders, overengagements) for each scheme under two configurations, ideal and achievable, are compared. Operating conditions are examined through selecting real-world options and modeling underlying probability distributions, where applicable. With the creation of the APL Coordinated Engagement Simulation (ACES), generation of the air picture is modeled at an adequate level of fidelity, including multi-object threats with aspect-dependent radar cross section (RCS), “just-in-time” radar search setup, launch event counting (i.e., estimating how many TBMs were launched), clustering (i.e., deciding which objects originated from which TBM), primary object track (POT) selection (i.e., which objects should be chosen to represent each TBM), and correlation. Communications over a network similar to Link 16 account for latency via bandwidth, time, and content-constrained transmission. Probability of kill depends on other actions and outcomes during the raid, such as whether the lethal object is being tracked by a radar.

ACES ENGAGEMENT COORDINATION SCHEMES

The current version of ACES, version 0.9, implements four different engagement coordination schemes: free fire, simple sectorized with unique defended asset assignment, first launch, and DED. This section briefly explains the first three schemes and then presents a detailed discussion of the last scheme. Information on engagement coordination scheme taxonomy and future schemes to be implemented in ACES is discussed in the article by Shafer et al., this issue.

Free Fire

The free fire engagement scheme is the most primitive of the four coordination schemes. Actually, this scheme involves *no* coordination. All of the defending units act independently and are free to fire at whatever they want. This scheme is often used to illustrate the need for coordination because it usually performs poorly. Units often end up firing at the first threats to appear. This results in the units expending all their radar resources, thereby overengaging (i.e., launching more defensive missiles than the force firing policy calls for—one salvo per threat in this article) a few threats while later threats become free riders (i.e., threats that are never shot at; they receive a “free ride”). In general,

the more the coverage of each unit overlaps, the worse free fire will perform.

Simple Sectorized

Simple sectorized engagement schemes are unbidden, static schemes that do not rely on any communications. In the most familiar variant of the simple sectorized category (the one discussed in this article), all of the defended assets are divided among the units in a force so that each asset is assigned to only one unit. The units only engage threats that are aimed at one of their own defended assets. Much preplanning is involved so that the defended assets are assigned to the different units in an intelligent manner.

One advantage of this scheme is that it is very simple to implement because it does not use any communications. Another advantage is that it usually produces a very low number of overengagements because the engagement decision is well defined. The only time an overengagement occurs is when a threat’s predicted impact area overlaps two assets defended by two different units or when two units disagree on the target’s predicted impact area.

However, the scheme’s main disadvantage occurs when one unit’s assets receive most of the threats in a raid. This usually causes the unit to be overwhelmed and unable to engage all the threats, resulting in free riders. At the same time, simple sectorized does not permit a neighboring unit to help the overwhelmed unit, even if the neighboring unit’s assets are not being attacked and the neighboring unit can engage threats aimed toward the overwhelmed unit’s assets. This leads to the force underutilizing its resources while letting threats attack some defended assets without any engagement.

First Launch

The third engagement coordination scheme implemented by ACES is first launch. The rules of this scheme allow the units in a force to act independently while planning engagements. From a force command and control (C^2) perspective, units are not aware of inorganic (i.e., other units’) threat–weapon pairing until they are notified of (defending) missile launch. However, once a unit launches against a threat, it sends out an engagement message stating that it has launched against that track. Upon receiving the message, all other units that were planning an engagement against the same track defer. If, for some reason, there is an interceptor failure, the firing unit will send out another engagement message with the status “engagement terminated.” This will allow other units to then schedule an engagement against that track number again. The scheme is unbidden and dynamic, and relies on a simple communications network. In a low-density (spatial and temporal) raid, it may be possible to implement

a variant of this scheme without any communications network; units would track other friendly units' missile launches and predict which threat they were heading toward.

One main advantage of the first launch scheme is that it attempts to minimize the number of free riders by allowing every unit that can shoot against a threat to do so unless another unit has already engaged it. When compared against simple sectored, this scheme does a good job at letting units help an overwhelmed unit. Another big advantage is that it is easy to implement; it does not require changes to the communications network and only impinges minimally on unit-level autonomy.

In first launch, the units' unique capabilities versus certain threats do not affect the choice of shooter. All Navy area TBMD units have nearly the same predicted intercept point (on the target trajectory), regardless of their location, because of the small endo-atmospheric intercept window. The unit farthest away from the predicted intercept point has the longest interceptor time of flight (TOF). Therefore, the farthest unit is the first one to launch, if it is not already too busy. Thus, often a unit that is uniquely capable against a certain threat will have launched against threats that other units could have engaged, and then that unit may not have any (radar) resources left to engage the ones only it could have taken, resulting in free riders. The farthest engagements may also be the engagement opportunities with the least chance of success, increasing the number of leakers in the scenario.

Another disadvantage is the scheme's reliance on the communications network. If the network is slow, the engagement messages may arrive too late for the other units to cancel their own engagements, resulting in overengagements. Overengagement can also occur, no matter how fast the communications network is, when units are positioned on contours of constant TOF, causing their launch times to be identical. Dual tracks will inevitably also lead to an overengagement because the units think they are engaging different threats. Also, incorrectly correlating tracks can result in one unit confusing the engagement message against a particular threat as an engagement message against a different threat, resulting in free riders. Of course, these last two disadvantages apply to all dynamic schemes.

Distributed Engagement Decision

ACES currently implements a scheme from the DED category. It is the only bidirectional scheme that has been implemented in ACES, version 0.9. DED schemes are also dynamic and rely heavily on a communications network. Because this is a very complex scheme that may be new to many readers, it will be described in more detail than the three previous schemes.

In the DED scheme, all units in a force share information about their possible engagements and then the force

makes intelligent decisions about which units should engage which threats. From a force C^2 perspective, early awareness of threat–weapon pairings provides a basis for command reaction, as needed. The fundamental idea behind the scheme is as follows: if all the units use the same engagement information with the same engagement selection algorithm, all the units will arrive at the same engagement decisions.

Like first launch, the DED scheme has the advantage over simple sectored of allowing neighboring units to help each other. However, the main advantage of this scheme is its bidirectional property. Unlike first launch, where the coordination takes place after the engagement has been made, DED makes decisions *before* the actual engagements. This allows the scheme to make intelligent decisions toward balancing and optimizing all of the engagements for the force.

The main disadvantage of DED is its dependence on a communications network. The communications network is even more important to the DED scheme than the first launch scheme because much more information is being shared. Another disadvantage is its susceptibility to errors caused by miscorrelation.

The next four sections give more details about the DED scheme. First, a database is described for organizing all of the possible engagement information for the entire force. Then the selection algorithm for deciding which of the engagements should be carried out is explained. As part of this algorithm, the shot elimination tests and criteria are defined. Both the database and the selection algorithm exist on every unit.

Data Management

Because of the complexity of the DED scheme, data management is very important. Information must be shared, stored, and processed uniformly across all units. Engagement information for the entire force is shared via engagement messages and stored in a force engagement schedule (FES) on each unit. The FES contains every possible engagement within the force. The entries on the FES are defined by a threat, unit, weapon, and shot number. The threat is the TBM track that is expected to impact a defended asset, the unit is the unit that has the shot against the threat, the weapon is the unit's weapon for the shot, and the shot number is the number of the shot in a salvo. Currently in ACES all units have only one weapon, Standard Missile-2 Block IVA, and all engagements are dual salvos. Therefore, for each unit, threat pairing can only result in two entries on its FES (one for each shot in a salvo).

Each unit uses the FES to decide which one should shoot at which threats. To make that decision, the entries on the FES are used to build a temporary hierarchical database with three levels: group (the highest level), threat, and possible shot. Within each group are threats that are spatially very close together. Within

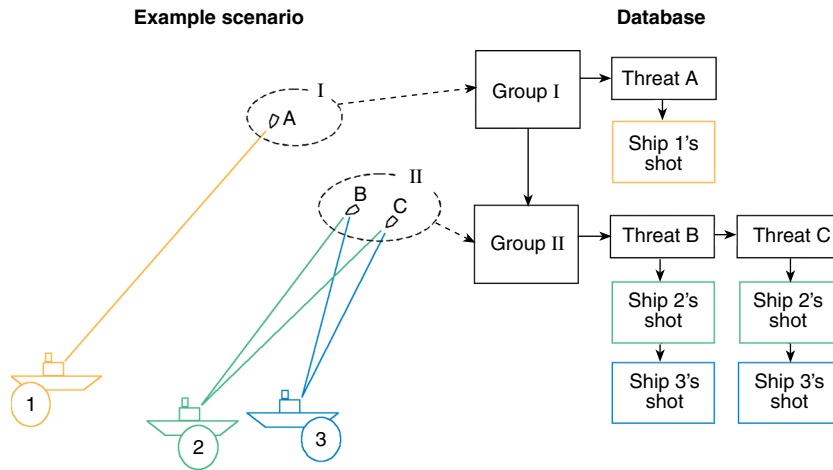


Figure 1. Example illustrating DED database construction.

each threat level are the possible shots against that threat. Figure 1 shows an example of the database from a simple scenario, which contains three threats (A, B, and C) and three units (1, 2, and 3). A line connecting a unit to a threat denotes that the unit is engageable against that threat. In this example, unit 1 is only engageable against threat A. The dashed ovals represent threat groups. In the figure, threats B and C have been assembled into group II. The database format shows how this example scenario would be organized within the DED database.

Engagement Selection Algorithm

The engagement selection algorithm has two main parts. The first part builds the database and the second selects a preferred shot for each threat in the database.

Building the database. The following five steps build the database:

1. Build the threat level.
2. Add the possible shots to each threat.
3. Combine threats into groups.
4. Freeze groups.
5. Order the groups.

Step 1 creates the threat level of the database by making every unique threat on the FES a threat in the database. Step 2 adds the possible shots to each threat in the database by walking through the FES and adding each shot to its corresponding threat. Once the threat and shot levels of the database are created, step 3 creates the group level by collecting threats that have a geometric separation distance less than the group distance threshold into the same group. When objects are closely spaced, they may become cross-correlated (e.g., when there are really two threats and two units agree that there are two threats, but one unit labels them A and B and the other unit labels them B and A). If one

unit is declared to have all the preferred shots for every threat within a group, an overengagement and a free rider can be avoided. If all of the threats within a group do not have at least one possible shot from the same unit, the group is divided. A threat can only belong to one group, and every threat is in a group, even if there is only one threat in a group.

After the entire database is built, the information is organized before being processed by the preferred shot selection algorithm. Step 4 freezes certain groups in the database. A group is frozen when there is not enough time to act on a change made to the preferred shot of one

of the threats in the group (i.e., it is within minimum coordination time of planned time to launch). If a group is frozen, none of its threats' preferred shots can be changed. Therefore, the group is skipped in the selection algorithm.

The final step in the process orders the groups in the database according to the following criteria: frozen, size, uniquely capable, maximum threat priority, and maximum threat number. The database is sorted because the preferred shot selection algorithm is run on the database so that higher priority is given to the more important threats.

Running the preferred shot selection algorithm. After the database is built and organized, the preferred shot selection algorithm is run. The goal of the algorithm is to remove the less desirable possible shots from each threat until only one possible shot remains for each threat. That shot is then called the "preferred shot" for the threat. The unit with the preferred shot against a threat is the preferred shooter for that threat. In Aegis, the second shot of the dual salvo is not scheduled until the first shot is in flight. Thus, in ACES, each (Aegis area TBMD) unit only has one possible scheduled shot against each threat, so finding the preferred shooter for a threat is the same as finding the preferred shot. In addition, because the same unit must engage every threat in a group, finding the preferred shooter for a group is the same as finding the preferred shot for every threat within that group. Therefore, another way of stating the goal of the elimination tests is to find a preferred shooter for each group.

Elimination Tests

Elimination tests are used to find a preferred shot for every threat in the scenario by sequentially removing possible shots from each threat. This is accomplished by sieving each group through a series of eight criteria. After each elimination test, if only one possible shot

remains for each threat in the group, sieving on that group is complete. Many possible criteria could be used in the elimination tests, and the order in which the tests are performed is just as important as the tests themselves. The eight tests and the order in which they are performed in ACES is just one possibility. After ACES is able to model more diverse systems, such as the Navy Theater Wide system, and more scenarios have been examined, other DED variants with different tests and/or different orderings of tests will be investigated. This section explains, in order, the eight elimination tests used in this particular DED scheme.

1. Notification time test eliminates the possible shots that do not allow the force enough time to coordinate. This is done by removing all shots that have a latest time-to-launch that is less than the current time plus the minimum coordination time. The minimum coordination time is an estimate of the least amount of time necessary to guarantee that—after a decision is made on one unit—there is enough time for the same decision to be made on all other units.

2. Maximum concurrent engagement test removes possible shots that cannot be carried out because the shot's unit has already reached its hard engagement capacity (i.e., maximum number of engagements that a unit can carry out at one time). To determine if a shot causes its shooter to exceed its hard engagement capacity, the shot's number of simultaneous discriminations is calculated. Discrimination in this context is expending additional radar resources to gather metrics to determine which TBM object to intercept. The number of simultaneous discriminations for a particular shot is found by counting all of the unit's discriminations (current engagements and predicted engagements of higher priority) that overlap the discrimination time needed for that shot. The test eliminates a shot from a threat if the shot's number of simultaneous discriminations is greater than the unit's hard engagement capacity.

3. Inventory sufficiency test eliminates possible shots from a threat's possible shot list if the shot's shooter does not have enough inventory to carry out the engagement. To support one of its shots being selected as a preferred shot, a unit must have enough planned inventory to engage all the threats within the group. A unit's planned inventory is the amount of inventory left after accounting for all of its other preferred shots of higher priority.

4. Engagement loading test tries to balance the number of simultaneous engagements for a unit across all the units in the force to facilitate engagement of a future threat that may be engageable by only a specific unit. This is done by comparing the number of simultaneous discriminations calculated in the maximum concurrent engagement test to the unit's soft engagement capacity. The latter is the maximum number of engagements at one time that can be handled by a unit without

stress. It must be large enough so that other criteria can be used to make preferred shot decisions. The actual test removes a shot if the shooting unit's number of simultaneous discriminations is greater than the soft engagement capacity of the unit.

For any threat in the group, if all the shots on the threat's possible shot list are removed by the engagement loading test, it means that all of the units that can engage the threat are already under stress. For this reason, all the possible shot lists for every threat in the group are restored to their pretest status. After all, engaging a threat under stress is better than not engaging it.

5. Reengagement test tries to maximize the number of reengagement opportunities. It eliminates possible shots that do not allow for another intercept opportunity. A shot allows a reengagement against a threat if the shot's latest time-to-launch is greater than or equal to the predicted time of intercept, plus the kill assessment time, plus the minimum coordination time. The test works by checking if any of the shots within the group allow for a reengagement. If a unit has no shots that allow for a reengagement, all of that unit's shots are removed from the threat's possible shot lists.

For any threat in the group, if all the shots on the threat's possible shot list are removed by the reengagement test, it means that none of the units has a reengagement opportunity. Therefore, all the possible shot lists for every threat in the group are restored to their pretest status.

6. Maximum inventory test attempts to balance the amount of available inventory for each unit across the force. The test removes a unit's possible shot from a threat if another unit with significantly more inventory has a possible shot (i.e., their inventory difference is greater than the inventory test tolerance). It uses the planned inventory associated with each shot from the inventory sufficiency test.

7. Time-of-flight test breaks a tie in case a group still has more than one possible preferred shooter. For each group, a unit's shot TOFs against each threat in that group are added together. The result is the total TOF necessary if that unit is selected as the preferred shooter. The unit with the least amount of total TOF is selected as the preferred shooter against the group, and for each threat in the group, all shots on the threat's possible shot list not belonging to the preferred shooter are removed.

8. Unit number test is executed in the rare case that two units have shots with exactly the same amount of total TOF against all the threats in a group. This test selects the unit with the lowest unit number as the preferred shooter. Then all shots on each threat's possible shot list not performed by the preferred shooter are removed. Because this is the last of the elimination tests, it must leave only one shot on each threat's possible shot list. Since ACES only has one weapon type

and does not divide salvos, examining the unit number is enough to ensure this condition. However, when ACES expands to using multiple weapons and dividing salvos among units, this test will also need to be expanded.

Elimination Criteria Constants

Table 1 lists the values of the elimination criteria constants used for this analysis.

Constant name	Usage	Value used in this analysis
Minimum coordination time	Notification time test	4.375 s (ideal); 12.25 s (achievable)
Soft engagement capacity	Engagement loading test	3
Hard engagement capacity	Maximum concurrent engagement test	5
Inventory test tolerance	Maximum inventory test	3
Group distance threshold	Grouping	1000 m

Asset no.	Threat nos.
6	20
7	1, 3-5, 17
8	None
12	13, 14, 16
13	12
15	11
20	7, 8, 19
22	None
24	2, 6, 9, 10, 15, 18
25	None

ACES ANALYSIS

Scenario

This scenario uses selected TBM launch and impact points from day D+5 from the Northeast Asia scenario used in the Capstone and Navy TBMD Cost and Operational Effectiveness assessments. Threat trajectories were generated using the Defense Intelligence Agency/Missile and Space Intelligence Center’s six-degree-of-freedom model, DICE. Four units defend against 20 threats launched from five launch areas. Fifteen of the threats are Scud Cs and five are No Dongs, which have reentry vehicles (RVs) that separate from the spent booster shortly after burnout. There are 29 entries on the defended asset list for this scenario, but only 10 of them can be defended from all launch areas by at least one of the ships. Seven of these 10 are near the aimpoints and actual impact points in the raid.

Table 2 shows the defended assets used, as well as the number of threats that impact at each one. However, the defended asset sizes were increased to encompass the impact points because the aimpoints were at facilities not actually contained within the assets. For example, a fuel storage tank on a dock was targeted and the assigned defended asset was the naval base on the other side of the dock. Impact time for the 20 threats was randomly distributed over 10 s for each Monte Carlo run. The four unit locations were based on those used for the aforementioned Cost and Operational Effectiveness assessments.

In this scenario, not every unit is in position to engage every threat; this is realistic and challenging. Figure 2 shows which units are physically engageable against each threat. In fact, none of the 20 threats is engageable by all four units. Nine threats are engageable by three units, seven are engageable by two units, and four are engageable by only one unit. There is no chance for reengagement to occur in this scenario because of the size of the area TBMD engagement volume and speed of the threats.

The radar setup is based on planned pairs of launch areas and defended assets (see Modeling Fidelity and Assumptions). The yellow shading in the boxes in Fig. 2 signifies that—for the given unit—the indicated threat was from a planned launch area/defended asset pair. In planning, it was

thought that these threats would be physically engageable, but as shown in the figure, this is not always true because of differences in modeling.

Communications Network

Navy Network 021A was used¹ as the basis for developing a reasonable Link 16 network (Fig. 3) for this scenario. It is assumed that the Cooperative Engagement Capability (CEC) of this era is still limited to anti-air warfare, although future analysis with ACES will include a space track picture based on CEC-like sensor fusion. The dashed lines in Fig. 3 indicate network participation groups (NPGs) relayed by an E-2C Hawkeye or E-3A Airborne Warning and Control System. Although not shown on the maps, this network was sized for 20 of the C² units and 36 of the fighters present in the theater. On the surveillance NPG, time slots were allocated as follows: 18 for the Theater High-Altitude Area Defense battery, 30 for the Patriot Battalion Information Coordination Central, and 8 for each of the other C² units (including Aegis area TBMD ships); this allows all units at least one opportunity to transmit every 1.5 s. On the Engagement Coordination NPG, each of the potential shooters was given 14 time slots, yielding a transmission opportunity every 0.875 s.

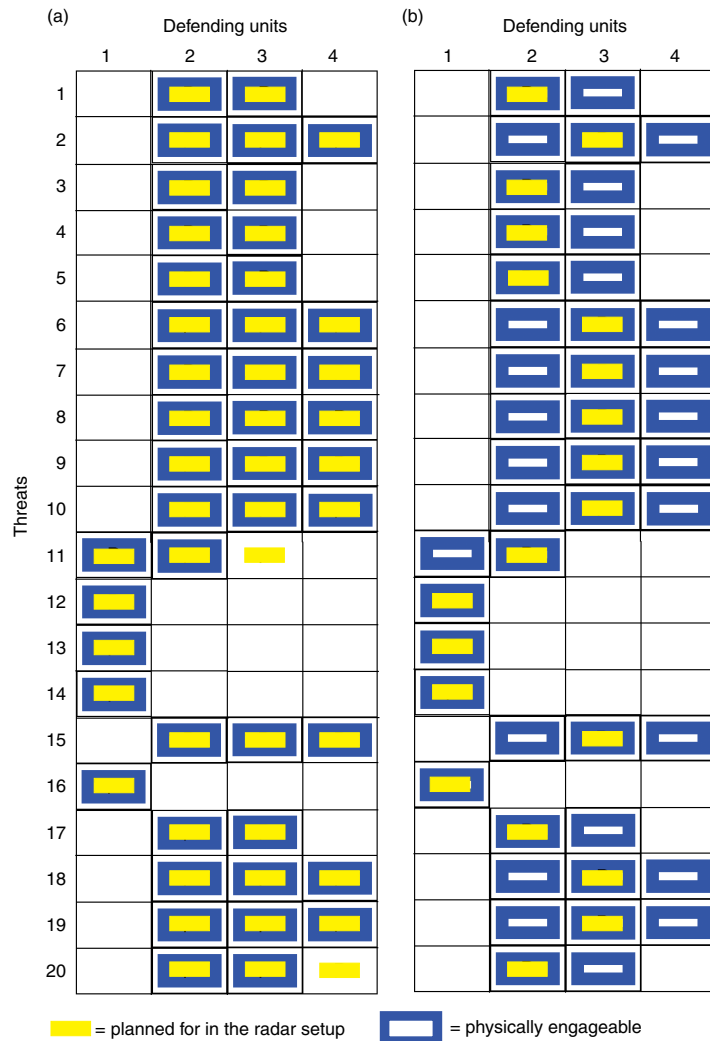


Figure 2. Planned pairs for (a) free fire, first launch, and DED and (b) simple sectored.

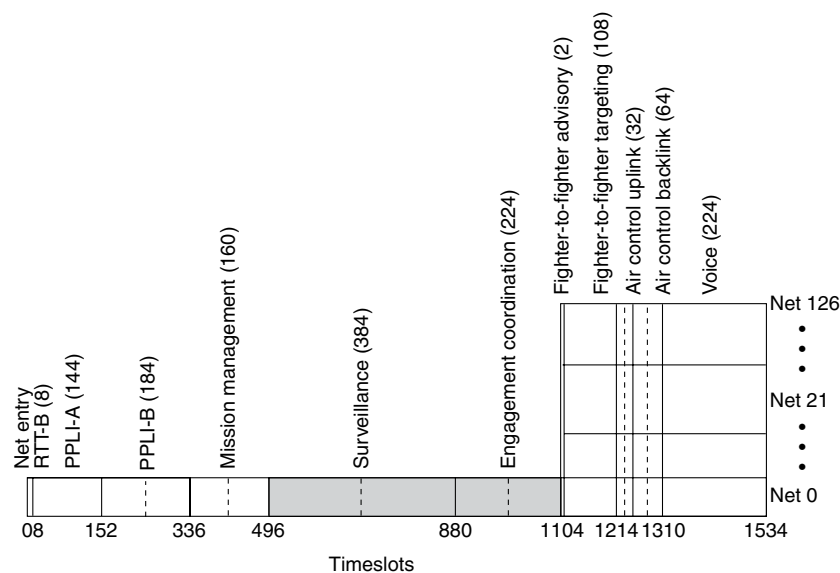


Figure 3. Modified Link 16 network (RTT = round-trip timing, PPLI = precise participant location and ID).

Modeling Fidelity and Assumptions

ACES version 0.9 was used as the modeling environment to conduct the study documented in this article. ACES is being created to provide a modeling and simulation framework for the development and analysis of advanced distributed weapons coordination concepts. As such, it helps analyze the performance and interactions of multiple weapon and sensor systems when they are deployed as an integrated theater air and missile defense force facing the full spectrum of theater air and missile threats. ACES functionality accurately represents unit-level engagement decision, control, and execution processes in a relatively complex environment (i.e., multi-sensor, multi-weapon, and multi-target). Weapon and sensor control functions, functional flow, and information exchange are being modeled with sufficient fidelity to facilitate the evaluation of alternative concepts for automated engagement coordination and the assessment of force- and unit-level effectiveness and efficiency in performing multiple simultaneous theater air and missile defense missions.

ACES version 0.9 is limited to Aegis area TBMD; other TBMD unit types and warfare areas will be implemented in the future. Engagement coordination performance was compared under two sets of conditions:

- Ideal: an abstract set without problems caused by limited search setup and resources, track inaccuracy, correlation errors, lack of communication bandwidth, or imperfect interceptor performance
- Achievable: a more realistic set including these disturbances to the air picture and nonunitary probability of kill

Unlike in anti-air warfare, TBMD hemispherical search coverage is not possible with current radar resources because of the need to detect targets at much greater distances. For the ideal set, this reality was ignored and a hemisphere of maximum waveform was used. For Aegis Baseline 6 Phase 3, the solution implemented was construction of search sectors with the goal of detecting TBMs only early enough to complete a minimum reaction time engagement for intercept at

the planned altitude. The APL-developed Array Radar Guaranteed Useful Search tool was used to generate such types of search sectors (waveform, slant range limits, and slant range rate limits as functions of beam position, along with the search frame time) for use in the achievable set. Each ship was assigned an Aegis resource planning and assessment (RP&A) extended air defense mission (goal of 95% probability of track initiation at least 20 s before the special auto-TBMD reaction time for the planned launch). The RCS values used in planning represent the constant mean value for the lethal object used in a Swerling IV distribution (i.e., what RP&A intends to use). However, for detection and tracking, ACES uses a roll-averaged, frequency-averaged (over AN/SPY-1B[V]/D band), vertical-vertical polarized, aspect-dependent RCS for each TBM piece as the mean in a Swerling IV distribution because this approach is believed to be more accurate.

In ACES, once the target is in track, the accuracies of the measurement are based on the signal-to-noise ratio return. A generic approach is used to manipulate these raw accuracies to approximate the effects of filtering. The track measurements are created by randomly drawing from a normal distribution using the calculated standard deviations of error and the target ground truth values as the means. Local bias errors are then added to the measurements. These biases represent residual errors in sensor registration and navigation. At the beginning of the simulation run, position and orientation bias errors are randomly drawn for each sensor, and these biases are held constant throughout. When reporting on a track, the standard deviation of the residual network bias is added to the covariance of the track state. The radar measurement and track accuracies used for the achievable set were commensurate with those from the AN/SPY-1B(V)/D radar present on Aegis area TBMD ships. Residual bias standard deviations representative of CEC operations were used because operational numbers for Link 16 were not available. The Link 16 requirement is 8.7 mrad and 150 m in range,² and the Aegis ownship sensor installation and navigation misalignment requirement is significantly better than the Link 16 requirement, but not as good as that achieved by CEC. It is known that Aegis ships achieve much better sensor registration than the Link 16 requirement, and it was assumed that sensor registration should provide some benefit over ownship misalignment. For the ideal set, all measurements and tracks were assumed to have no position or velocity errors, and correlation of remote tracks to local tracks was perfect. More details of the ACES detection, track, and correlation modeling methodology are available in the article by Bates et al., this issue.

For the nonideal set of runs, two messages were allowed per time slot. For the ideal set, this limit was removed. ACES exchanges data in the J3.6 space

track message and the J3.0 reference point message (to communicate impact point) on the Surveillance NPG. ACES also exchanges data in an enhanced engagement status message (including J10.2I and J10.2C2 engagement status words); this is necessary for the DED scheme to function. Further details on how communications networks operate in ACES are available in the article by McDonald et al., this issue.

ACES currently performs ground truth clustering; an RV is selected as the POT and clustered with its spent booster if they are both in track. Future versions of ACES will perform clustering and linking on the local track picture based on logic consistent with Aegis Baseline 6 Phase 3.

ACES threat assessment, engageability, and scheduling are performed on local tracks in a manner consistent with Aegis Baseline 6 Phase 3, although scheduling required changes previously mentioned (i.e., the addition of the “deferred” engagement status and use of the preferred shooter indicator) to accommodate the DED scheme. Future versions of ACES will also allow these functions to be performed on remote tracks. Threat assessment projects the track state of the POT and forms a $3\text{-}\sigma$ impact point error ellipse. The POT is given the priority of the highest-priority defended asset, which has any area within the ellipse. Engageability projects the POT (again ignoring drag) through the appropriate irregular hexagon in downrange vs. altitude space as a function of cross range and specific energy.

The number of simultaneous target discriminations allowed to be scheduled by a platform is five; this is the number chosen for past analysis and construction of the scenario, and represented the best estimate of capability at the time. The threshold and objective requirements for the number of simultaneous discriminations are contained in the Naval TBMD Operational Requirements Document. The priorities for choosing which engagements to schedule, in order, are as follows: (1) if the unit is the preferred shooter on this track, (2) higher threat priority, and (3) earlier latest time to launch. These functions, as well as force engagement selection (i.e., deciding on preferred shots), are performed periodically every second.

ACES engagement outcome has several probabilistic tests: launcher failures, missile prelaunch reliability, missile in-flight reliability, uplink/downlink reliability, handover/designation, and lethality. These tests do not depend on geometry (but *do* depend on whether the RV is in track). Geometric dependence is planned for future versions of ACES.

Results

Because this scenario is saturated (i.e., number of simultaneous discriminations \times number of shooters \times depth of fire per shooter = number of threats), each overengagement will lead to a free rider, except in cases

where not all discriminations overlap. Each unit can only perform five simultaneous discriminations. In certain raid timings, units may be able to engage six threats because one intercept occurs before the next discrimination must begin. In the free rider, overengagement, and leaker figures shown throughout the text, the orange horizontal lines indicate the actual results of the 100 Monte Carlo runs; the light yellow portions represent a 95% confidence interval. In the ideal case, the leaker and free rider charts are the same because the effective probability of kill of the missile equals 1.

Ideal

Free fire. In free fire, there are almost always four or five free riders, as seen in Fig. 4. This can be explained by the way the units determine which threats to engage. Because the units in the ideal case have detected each threat with plenty of time to engage it, they have a fairly complete picture of the raid. When a unit schedules intercepts, it prioritizes the threats by threat priority and then by the earliest latest time to launch. By looking at the threat levels and which unit can engage which threat, it can be determined that only two out of six threats to asset 24 (see Table 2) will be engaged. This means that each run is guaranteed at least four free riders. It is also true that two units will each engage four out of five threats to asset 7. Depending on the latest times to launch for each unit, this could lead to one free rider and will guarantee three or four overengagements. Figure 5 shows that in cases where a unit can engage six threats, there is almost always another overengagement.

Simple sectored. Because of the way the assets were assigned in the simple sectored case and the way the raid actually occurs, several units have more threats against their assets than they can possibly defend. This situation leads to a higher number of free riders than with free fire. Also, one unit has none of its defended assets threatened and therefore does not contribute to the defense. Table 3 shows the number of threats that each unit must defend against.

Unit 1 will be able to engage all of its threats. In most cases, unit 2 will miss two of its threats and unit 3 will miss four. This accounts for the large percentage of cases with six free riders. For a smaller percentage of the time there are only five free riders because unit 2 or 3 was able to engage six threats, as shown in the table.

As seen in Fig. 5, there are never any overengagements in the simple sectored case. Because this engagement scheme has static defended asset assignments and the units have perfect knowledge of the track states in the ideal runs, they will never confuse the impact point of a threat. This means that a unit will never believe that a threat is targeting its defended asset if it is not. Therefore, the units will only engage threats to their defended assets.

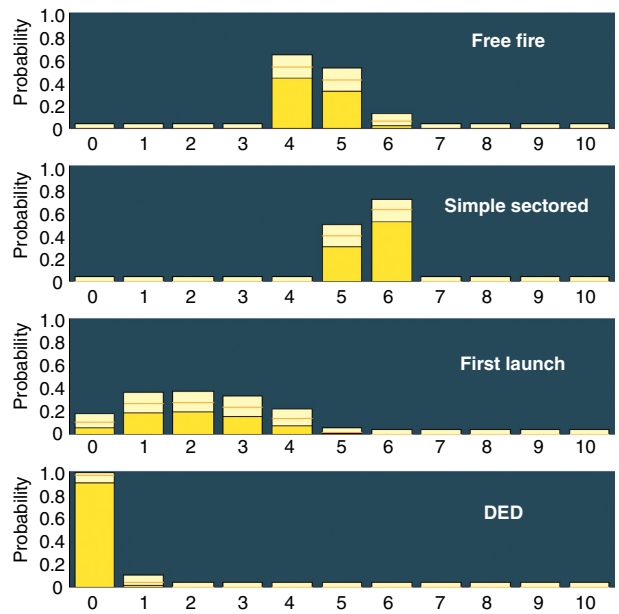


Figure 4. Free riders (ideal).

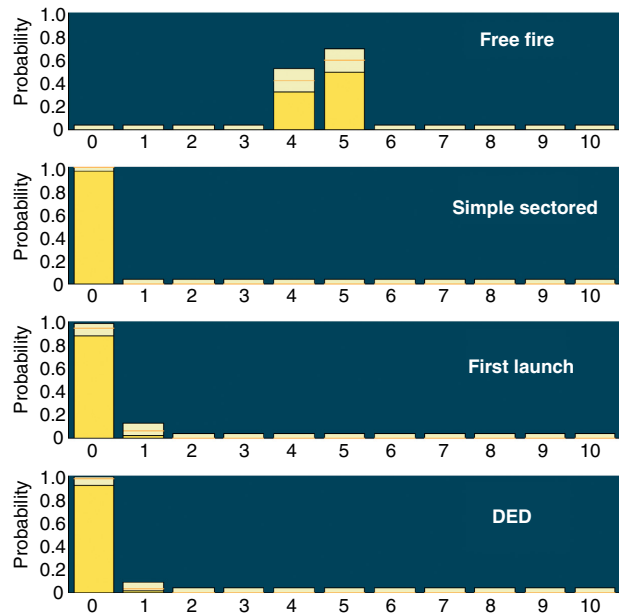


Figure 5. Unintentional overengagements (ideal).

Unit	No. of threats to assigned assets
1	4
2	7
3	9
4	0

First launch. As shown in Fig. 4, first launch performs better than both free fire and simple sectored. Like free fire, this scheme allows all units to potentially engage all threats that are physically engageable and in local track. Unlike free fire, an engagement status message is sent once a unit has launched to prevent other units from engaging the same threat. This dramatically cuts down the number of overengagements between free fire and first launch. The only time an overengagement can occur in the ideal case is when two units must launch at nearly the same time. This does not allow time for an engagement status message to be received before launcher assignment/missile commit. By having fewer overengagements, the units are available to fire at what may have become free riders in the free fire case. This does not mean that there are never free riders. In this scenario, units 1 and 4 may each be left with less than five threats that they can engage because units 2 and 3 fired first. This leaves several threats that units 1 and 4 cannot engage, while units 2 and 3 are occupied.

Distributed engagement decision. Figures 4 and 5 show that under the ideal configuration, the DED scheme outperforms all the others. Overengagements seldom occurred, leaving each unit free to carry out its maximum number of simultaneous engagements. The three times overengagements happened were when a dual track resulted from a unit that never locally tracked the spent booster of one of the No Dongs; its RV track could not inherit the remote track number from the booster as did the other units. Occasionally, free riders can still occur because of the sequential nature of the current DED algorithm. This scheme had many of the advantages of first launch with the added advantage that decisions were made well before any of the engagements took place. In the ideal situation, all the targets are detected at least 40 s before the first planned launch time. Pairing of units to targets before launch solved the communications latency problem present in first launch.

Achievable

As expected, most of the results under the achievable configuration are worse (a greater number of free riders and overengagements) than those obtained under the ideal configuration. Most of the discrepancies can be attributed to the inferior air picture created in the achievable configuration. The achievable air picture had fewer threats detected, and those that were detected were seen later. Multiple tracks (multiple track numbers assigned to the same object) and mis correlations (multiple objects assigned to the same track number) further degraded the information available.

Free fire. The results of the achievable case are almost the same as the ideal case. However, some of the achievable runs had only three free riders. The reason

for this is likely that fewer threats were detected by the units, which then changed the order in which the intercepts were scheduled. This would mean that if a higher-priority threat was not detected, a lower-priority threat that was never engaged in the ideal case is now engaged. Figure 6 shows for each planned pair (as seen in Fig. 2a) in all runs the amount of time between detection and an ideal planned time to launch. Those detected later than planned may not yield adequate time to plan and execute an intercept. There are 708 threats (out of a possible 4500) that are not detected at all by the planned unit. The RP&A methodology failed to meet its goal (95% probability of track initialization at least 20 s before the special auto-TBMD reaction time for the planned launch) because the RCS values used by RP&A are too high in this case. The goal is met less often with the nominal mission (at the special auto-TBMD reaction time); thus, the extended mission was chosen for this article.

Simple sectored. Figures 7 and 8 show that the simple sectored results from the achievable configuration were nearly the same as the results obtained from the ideal configuration (Figs. 4 and 5). The defended assets were assigned so that the track errors did not create any confusion as to which unit's defended asset was being targeted by a threat. Therefore, no overengagements occurred. For the reason previously explained in the ideal simple sectored case, there will typically be five or six free riders. In other instances, because of statistical variation, the number of free riders decreased to four because units 2 and 3 each had six shots.

First launch. As expected, the results in Figs. 7 and 8 show that first launch performed worse under the achievable conditions than under the ideal conditions. Unlike free fire or simple sectored, the overengagement results are considerably different than the ideal

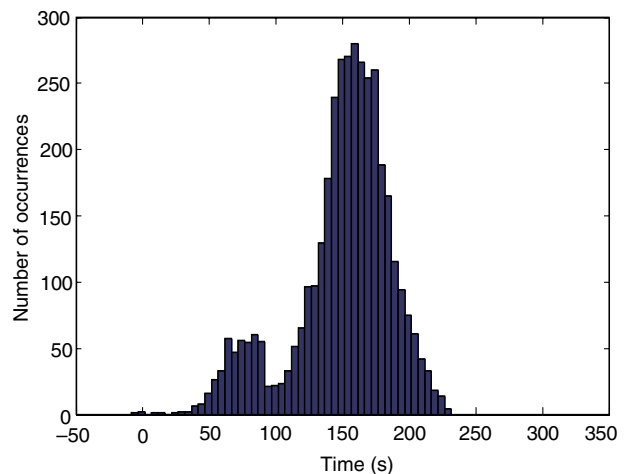


Figure 6. Time from detection to planned launch (achievable). A total of $84.1\% \pm 0.5\%$ ($1-\sigma$) of all objects were detected in time; 708 objects were not detected at all.

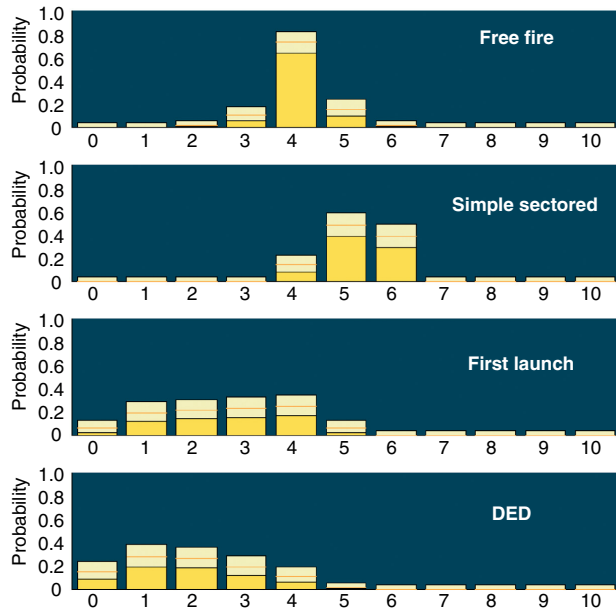


Figure 7. Free riders (achievable).

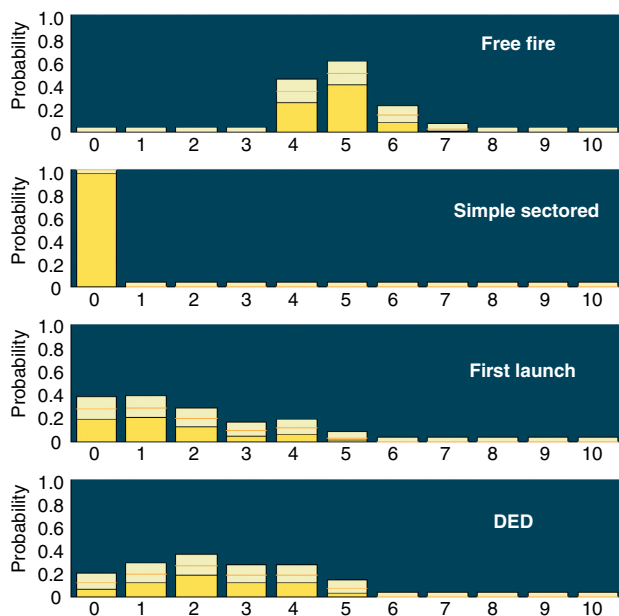


Figure 8. Overengagements (achievable).

configuration results. This can be explained by the fact that first launch relies on a common air picture and a communications network, whereas free fire and simple sectored do not. In the achievable runs, errors in the track picture led to multiple track numbers for the same object. Because the force treats that object as more than one different threat, overengagements usually result. This is a product of the bad air picture and not the engagement coordination scheme. In the first launch results, 96% of overengagements are caused by multiple track numbers for the same object. The remainder of the

overengagements can be attributed to engagement status messages arriving too late (nearly simultaneous launches). The number of free riders did not increase significantly because the units that overengaged threats had not always reached their maximum capacity in the ideal case.

Distributed engagement decision. Similar to first launch, the results in Figs. 7 and 8 show that the DED outcomes became considerably worse under the achievable conditions. Just like first launch, the degradation can be contributed to the unclear air picture. In the DED results, 88% of the overengagements are caused by multiple track numbers. Figure 9 shows the variability over Monte Carlo runs of the number of unique remote track numbers at the earliest planned launch time for DED (other schemes are similar). The orange portion of the bars indicates boosters with remote track numbers; usually at least one booster was reported (because someone was not tracking the corresponding RV), but not always. The goal was to report one track for each of the 20 events; this never happened. Often there were fewer unique remote track numbers than the number of objects (25), but the average was 28 (or 1.3 unique remote track numbers per object with a remote track number). These extra remote track numbers are the result of a failure to correlate either initially or because recorelation is not allowed after decorrelation. Miscorrelations (more than one object with the same remote track number) are more rare, as seen in Fig. 10.

The air picture also varies significantly with time. For Monte Carlo run 24, Fig. 11 shows the number of objects in track by at least one unit (cyan), the number of miscorrelations (yellow), and the number of unique remote track numbers (magenta). In this run, several decorrelations occur after the first launch at 407 s and before the first intercept at 469 s, which can clearly cause problems. After successful intercepts occur, the

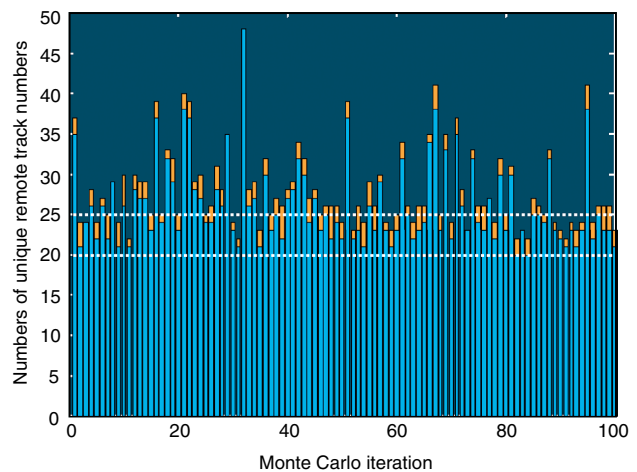


Figure 9. Multiple remote tracks (snapshot, DED, achievable).

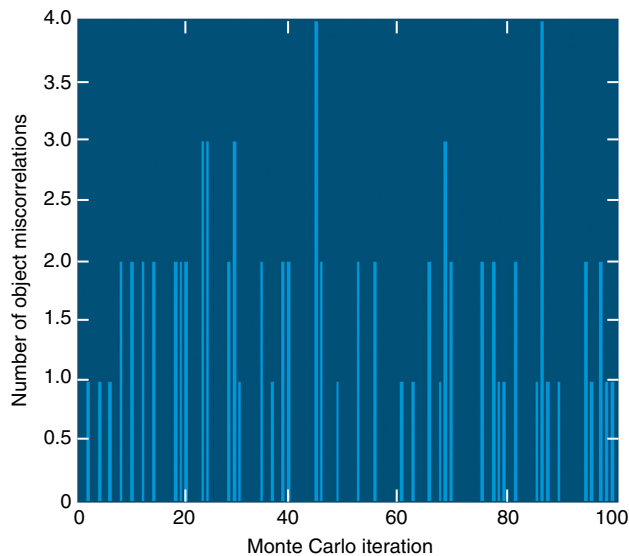


Figure 10. Number of miscorrelations (snapshot, DED, achievable).

local tracks on these objects degrade before they are dropped, leading to further decorrelations.

As previously mentioned, not all of the planned detections always occurred. This reduces the solution space (i.e., eligible number of engageable threat/unit pairs) for the algorithm. Launches can occur before detection of all objects, as seen in Fig. 12. It is impossible to guarantee a correct solution if decisions must be made before complete information availability.

More launches occur in DED than in first launch because units 1 and 4 nearly always had five targets to engage. More launches cause more missile failures (duds and failure of uplink or downlink). Each missile failure causes an engagement status message with the status set to “engagement terminated” to be sent. Once this occurs, the algorithm decides again which unit has the best shot on that particular track. Currently in the model, this is being counted as an overengagement because two units engaged the same track. In the future, these overengagements will be counted separately from the unintentional overengagements. In DED, the other units are more likely than in first launch to be available to shoot because of the load-leveling aspects of the algorithm.

Finally, for all schemes, the number of leakers will usually exceed the number of free riders in the real world because missiles are not infallible. Figure 13 shows the number of leakers for the achievable case for all four schemes. Only in 1 run out of the 400 did no leakers make it through; it was for DED.

CONCLUSIONS

The results reveal two important truths. Under ideal conditions, a DED scheme performs significantly

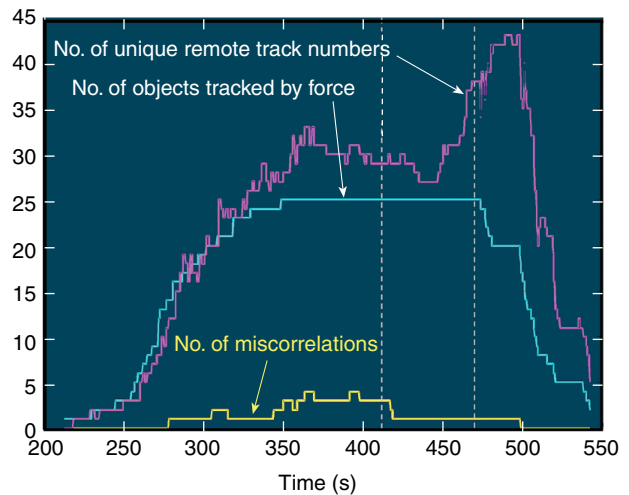


Figure 11. Run 24 air picture (DED, achievable). The dashed lines represent the first scheduled launch time and the first scheduled intercept time.

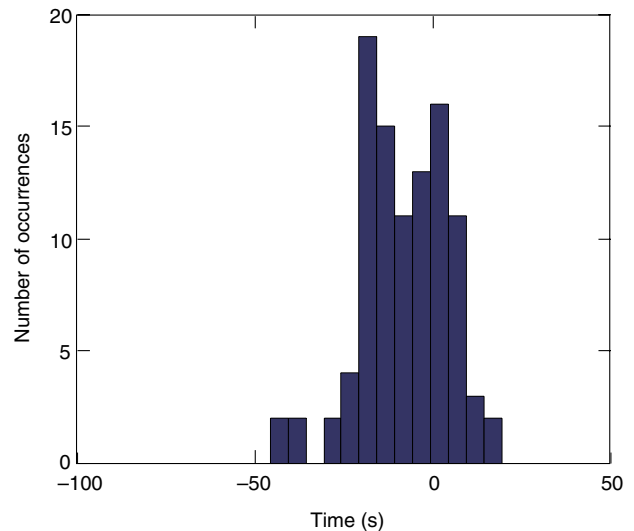


Figure 12. Time from latest detection to earliest planned launch (achievable).

better (and nearly perfectly in this scenario) than other schemes. However, realistic shortcomings in the information available to the units blur the differences between DED and first launch.

The results of both setup configurations (i.e., ideal and achievable) can be used to show a general comparison of the four engagement coordination schemes implemented in ACES. Using free riders as the metric to compare performance, simple sectored proved to be the least effective in this particular scenario, closely followed by free fire. The two dynamic schemes were significantly better, with the DED scheme modeled outperforming first launch.

The achievable air picture formed in this scenario leaves much to be desired. None of the schemes was

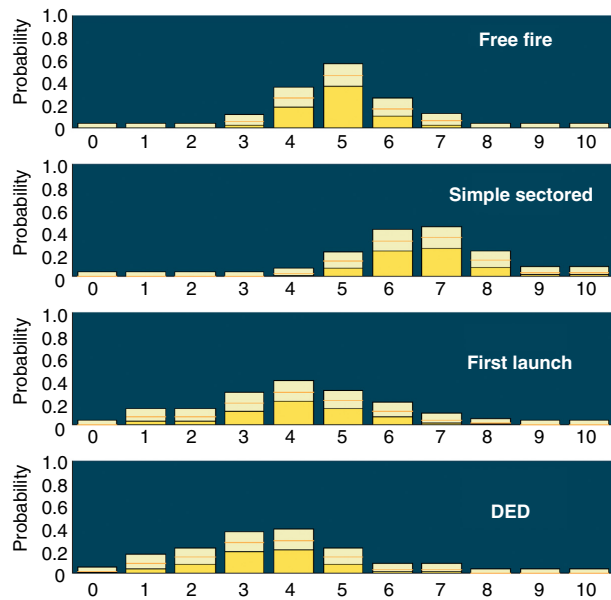


Figure 13. Leakers (achievable).

able to overcome numerous instances of multiple tracks (averaging 1.3 unique remote tracks per object with a remote track number), several miscorrelations, and too many targets going undetected by units that planned to see them. The effects of communication latency and changing planned launch times seemed to have less of an effect than air picture quality.

FUTURE WORK

There are five main areas for further research.

1. Investigate excursions in the achievable conditions (primarily radar setup and correlation/decorrelation rules) to determine what can be changed in the real world to improve performance.
2. Analyze other scenarios (i.e., different theaters, raids, and unit locations) to ensure that the conclusions here extend to them as well.
3. Examine other engagement coordination schemes that exist, including those being considered for use.
4. Include more defensive systems (e.g., Navy Theater Wide, Patriot, and Theater High-Altitude Area Defense) that will be present in future conflicts.
5. Allow other warfare areas (i.e., Anti-Air Warfare and Overland Cruise Missile Defense) to compete in real time for defensive resources.

REFERENCES

- ¹JNL 198 Quick Reference Guide, NTSCI (1 Nov 1998).
²Link-16 Operational Specification OS-516.2, Change 2, 3, Published under the Direction of CNO (N62) by the Navy Center for Tactical Systems Interoperability (Mar 2000).

ACKNOWLEDGMENTS: The authors would like to especially thank the Office of Naval Research and the Navy Theater Wide Program Office for funding the research discussed in this article. The authors also want to thank the other members of the ACES development team (S. R. Allen, C. W. Bates, M. J. Burke, J. F. Engler Sr., J. M. Henly, B. L. Holub, E. M. McDonald, and K. E. Shafer) for their continuing contributions to the modeling and analysis efforts discussed here. For further information, please contact R. A. Phillippi, APL Program Manager for Joint Missile Defense Projects.

THE AUTHORS



SIMON MOSKOWITZ earned a B.S. in aerospace engineering from the University of Virginia in 1993 and an M.S. in systems engineering from the University of Arizona in 1994. He is a Senior Professional Staff member of the Air Defense Systems Engineering Group. He joined APL in 1994. Theater Air and Missile Defense radar search and engagement coordination algorithm development and modeling comprise his current primary focus. His expertise includes TBMD threat characterization, concept of operations/scenario development, effectiveness analysis, visualization, requirements development, and specification review. His e-mail address is simon.moskowitz@jhupl.edu.



REBECCA J. GASSLER graduated from Virginia Tech with a B.S. in aerospace engineering. She is currently pursuing an M.S. in information systems and technology from The Johns Hopkins University. She joined the Air Defense Systems Engineering Group in ADSD in June 2000 and has dedicated most of her time to working on ACES. Ms. Gassler's professional interests concern modeling, simulation, and analysis of current and future theater- and campaign-level operations. Her e-mail address is rebecca.gassler@jhuapl.edu.



BART L. PAULHAMUS is a member of APL's Associate Professional Staff in the Air Defense Systems Department. He received B.S. and M.S. degrees in computer science and engineering from Penn State University in 1998 and 2000, respectively. He joined APL in 2000 and is currently working on developing modeling and simulation software for the Air Defense Systems Engineering Group. His e-mail address is bart.paulhamus@jhuapl.edu.