DRDC External Publication

Produced for: 36th International Symposium on Military Operational Research (ISMOR)
Royal Holloway, University of London
London, UK
23 – 26 July 2019

Leveraging physics-based simulations for Operational Analysis:
Task Group Air Defence Case Study

Peter J. Young[1]
Centre for Operational Research and Analysis
Defence Research and Development Canada

Abstract

Physics-based Monte Carlo simulations are extensively used for combat system engagement analysis. Applications include combat effectiveness assessment, tactics and doctrine development, data construction for implementation into systems, training, planning for trials and post-trial reconstruction/analysis. The development and use of such simulations are pursued with engineering disciplines and form mainstream capabilities owned and employed directly by military organisations. The military requirement for consideration of a wider operational context, coupled with scalable simulation designs and high throughput computing capabilities, makes the employment of these tools attractive for Operational Analysis (OA) purposes. Their usage, however, can still be compounded by the large parameter spaces typically present in OA studies, particularly for optimisation problems. Traditional OA techniques involving study design, metric definition, problem abstraction, and employment of stochastic techniques to account for uncertainties can play a role to complement and leverage off the physics-based simulations. This theme is discussed in context of an ongoing study investigating air defence for a naval Task Group (TG). Physics-based simulations are employed to generate single defender – single threat engagement zones. The engagement zones are then used by a simplified multi-defender – multi-threat model to investigate TG operations.

Acknowledgements: To the officers and staff of CFMWC AWB, M&S and ORT who have contributed to the TG AD OA study.

[1] Operational Research Team, Canadian Forces Maritime Warfare Centre, CFB Halifax, Nova Scotia, Canada. Email: peter.young2@forces.gc.ca

Introduction

In late 2017 the Canadian Forces Maritime Warfare Centre (CFMWC) received an Operational Analysis (OA) tasking to study naval Task Group (TG) Air Defence (AD). The CFMWC is the Canadian Forces centre of excellence for the development and delivery of maritime tactics and operational manoeuvre doctrine in support of Canada`s maritime forces. To achieve this mission the CFMWC is structured into the Above Water, Under Water and Joint Technology and Innovation Battlespaces, AWB, UWB and JITB. The CFMWC is also supported by an Operational Research Team (ORT) staffed by three defence scientists from the Centre for Operational Research and Analysis (CORA), Defence Research and Development Canada (DRDC). The CFMWC has specific expertise in Anti-Air Warfare (AAW, in the AWB) and Modelling and Simulation (M&S, in the JITB supporting the other battlespaces). A study team was therefore formed drawing on subject matter experts (SMEs) from these areas with ORT support.

The TG AD study is now ongoing at the CFMWC. Physics-based simulations acquired to support AWB requirements are being employed to underpin the OA study. Simulation results are being directly used within a simplified modelling approach permitting exploration of a large parameter space and aggregation of results into metrics appropriate for the study. The following sections provide an overview of the study and how physics-based simulations are being leveraged to support it. Information gained from this case study is used to explore the theme of how physics-based simulations complement traditional OA techniques involving study design, metric definition, problem abstraction and employment of stochastic techniques for undertaking studies.

Modelling and Simulation support for the Royal Canadian Navy

The CFMWC is the primary provider of M&S capability supporting the Royal Canadian Navy (RCN). The CFMWC's M&S capability consists of a suite of models and simulations and a computing infrastructure permitting their usage to support a range of RCN requirements. This section provides an overview of this capability.

M&S capability

Figure 1 shows a modelling hierarchy to support naval defence analysis at tactical, operational and strategic levels. Typical to other defence areas, the fidelity of representation decreases from high, at the tactical level, to low, at the strategic level, while the degree of abstraction increases.

The CFMWC employs a suite of computer-based models and simulations for the areas in the rounded rectangles of Figure 1. The models are used to support the following activities:

- Concept Development and Experimentation (CD&E). Focused at the operational level with human operators controlling entities through networked simulations.
- Training. Across all levels. One focus at the operational level is for maritime warfare training.
- Tactics development. Engagement simulations at tactical and operational levels used to develop Maritime Tactical Instructions (MARTIs).
- Doctrine development. Engagement simulations at the tactical level used to derive input data for command management systems onboard RCN ships.

- Operational Test and Evaluation (OT&E). Engagement simulations at the tactical level used to support planning for at-sea trials and to perform post-trial reconstruction and analysis.
- Requirements development and verification. Engagement simulations at tactical and operational levels used to support development and verification of future requirements.

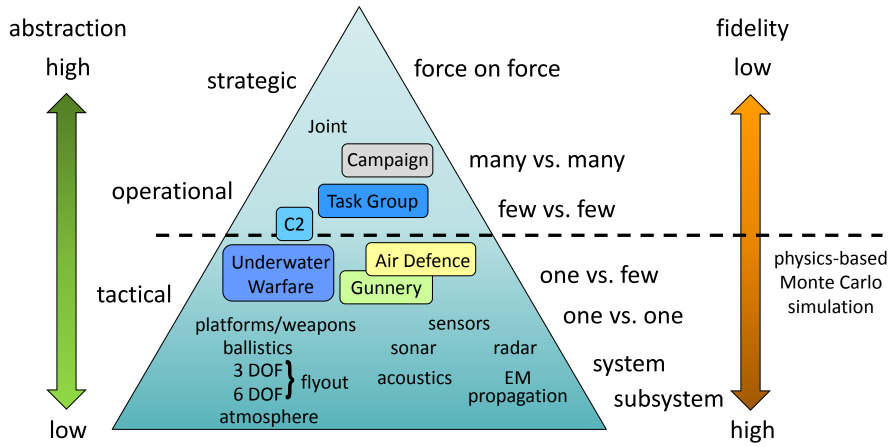

Figure 1. *Modelling hierarchy to support naval defence analysis*

Development and usage of the above models and simulations has had a significant OA influence, e.g. general naval operations analysis [1], CD&E [2], and OT&E [3]. They are now part of main-stream CFMWC capabilities used and maintained by dedicated naval and engineering staff.

The tactical models in Figure 1 are physics-based and typically Monte Carlo simulations. They are further supported by embedded or integrated system and subsystem models of high fidelity. The naval AD model is a commercial simulation with integrated modules developed by DRDC and Allied partners employing radar and missile theory [4, 5] with environmental conditions specified through the Advanced Propagation Model [6] and an atmosphere model based on [7]. The AD model has also been used as a test-bed supporting research into Command and Control (C2) for AAW [8].

Computing infrastructure

A schematic of the CFMWC computing infrastructure is shown in Figure 2. This infrastructure is supported by dedicated staff who perform configuration and management roles for networks, models and databases, and provide expertise for the running of the models and subsequent analysis of results. The infrastructure with hosted models and simulations is used to perform studies and conduct training and experimentation, the latter two involving human operators controlling entities via computer stations in syndicate rooms. The computing infrastructure is also used to store and perform analysis of data collected from at-sea trials.

Industrialised M&S

The CFMWC's simulation capability, in combination with its computing infrastructure, provide, what may be termed, an industrialised M&S capability for the RCN. A key aspect of the computing infrastructure is its ability to support High Throughput Computing (HTC), as depicted in Figure 3. Batch jobs for a simulation with input parameters covering a range of values can be submitted to the computing cluster, this comprising in excess of 1000 processors. Following run completion, output data from the simulations are parsed and stored in databases. Queries can then be used to extract specific results from the databases. This approach is particularly useful for Monte Carlo simulations, for which each combination of input parameters require multiple runs to obtain mean results with variances.

HTC enables Monte Carlo simulation as a viable means for exploring large parameter spaces using high fidelity models. Parameter space dimensionality, run-time speeds and data management can, however, still provide limitations for this. The TG AD study adopted an analysis approach to capitalise on the capabilities of physics-based Monte Carlo simulations while balancing this with a simplified modelling layer to permit investigation of a large parameter space for the study problem.

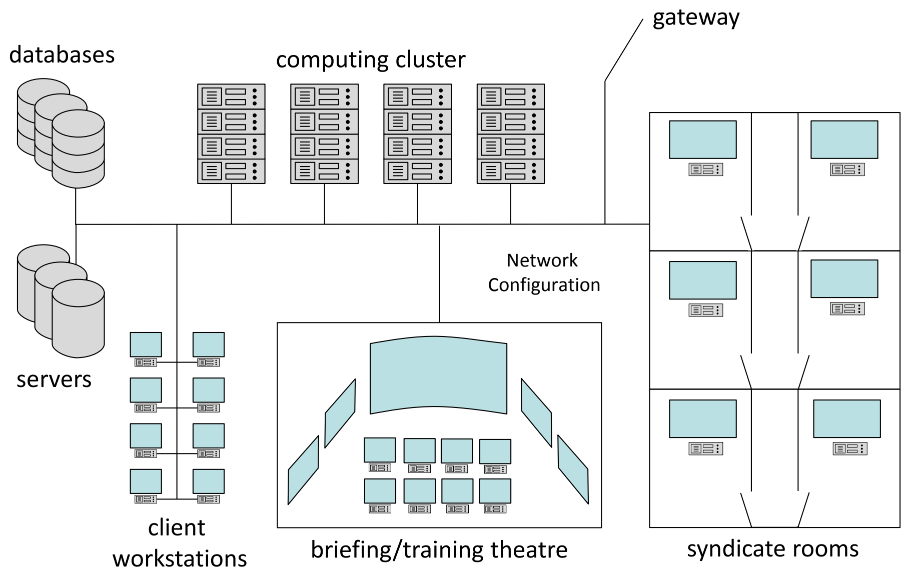

Figure 2. *Schematic of the CFMWC computing infrastructure*

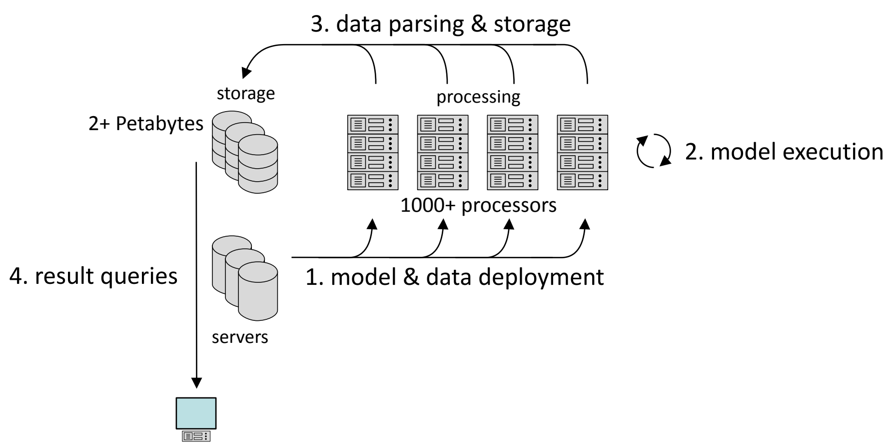

Figure 3. *High Throughput Computing (HTC) capability for M&S*

Case Study: Task Group Air Defence

Initial CFMWC discussions identified the key characteristics of the systems that needed to be studied and the models/simulations that could be used for this. Central to this would be the employment of CFMWC's primary naval AD model. This Monte Carlo simulation has capability to represent a naval TG under attack by multiple air threats. The simulation with supporting models provides medium to high fidelity for representing search and tracking radars, C2 for controlling coordination between AD ships, and Surface-to-Air Missiles (SAM) using 3 degree-of-freedom (DOF) and 6 DOF fly-out algorithms. The simulation has been extensively used for the study of a single AD ship defending itself or a High Value Asset (HVU) against a single threat, with additional application including multiple defenders and multiple threats. Outputs from these studies have supported development of tactics, doctrine for current AD capabilities, and requirements for future capability.

Further discussions for the study focused on the context in which the TG was operating, the range in which input parameters should be varied, and the metrics that should be used to measure TG AD effectiveness. Through these discussions a study approach was adopted to balance analysis using the high fidelity simulations with simplified modelling to permit coverage of the identified parameter space and aggregation/roll-up of results into suitable metrics. The following sub-sections provide details of the study design and how a simplified modelling approach was adopted to leverage off of the physics-based simulations.

Study design considerations

Scenario aspects of the problem include the following:

- TG configuration. Number of AD ships and their stationing about a HVU. The objective for the AD ships is to protect the HVU from an air attack.
- AD capabilities. Types of search and tracking radars. Types and loadouts of short range (SR) and long range (LR) SAM systems on each AD ship. AD coordination and firing policies across the TG.
- Air threat. Types and numbers of threats in the air raid. The main threat axis relative to the TG and individual spacing and timing of threats about this axis. All threats are assumed to target the HVU.
- Environmental conditions including ducting height affecting radar propagation and atmospheric conditions affecting missile fly-out.

System/subsystem aspects include:

- Radar. Performance characteristics of search and tracking radars.
- SAM. Fly-out, guidance and warhead characteristics for each SAM type.
- Air threat. Signature and flight characteristics for each threat type. Vulnerability of each threat type to each type of SAM warhead.

The key Measure of Effectiveness (MoE) identified for assessing TG effectiveness was the overall SAM expenditure for the TG to counter a threat raid by achieving a desired Probability of Kill (P_K) against each threat in the raid. A study objective was therefore to find TG configurations which would reduce this MoE whilst achieving the desired P_K for each threat. Adoption of firing policies with assessments between shots, e.g. Shoot Look Shoot (SLS), instead of sequential shots with no assessments between shots, e.g. Shoot Shoot (SS), would result in smaller SAM expenditures whilst achieving the same P_K. Employment of SR SAMs over LR SAMs may also be preferred as AD ship loadouts for the latter, more capable and expensive missile type are smaller. Preferred firing policies will therefore require an increased battlespace, i.e. sufficient time from the initial detection of a threat to when it enters a keep-out range about the HVU, to enable multiple shots using preferred weapons, with assessments in between shots. Coordinated AD across multiple ships optimally stationed within a TG can help achieve this increased battlespace and thereby improve TG AD effectiveness.

The relative bearing of the main threat axis to the TG and the TG configuration, as specified by its ship stationing, define the initial conditions for an air raid on the TG. In general, the actual threat axis will not be known in advance and therefore has to be considered within a range of values. These factors, in combination with the numbers of threats and AD ships within the TG, define a scenario parameter space as shown in Figure 4. In this example, a set of 20 threat axes are placed in a symmetrical fan with even spacing and horizontally centred about the vertical axis through the HVU, each axis directed towards the HVU. One AD ship is placed in a 20 km by 20 km region in close proximity to the HVU and horizontally centred about the vertical axis, and two additional AD ships are placed in similarly sized regions up-threat and symmetrically opposed relative to the vertical axis. The resulting scenario parameter space contains a symmetry which permits the region for the first ship to be reduced to the right half, as shown in Figure 4. The resulting number of TG configurations obtained when employing grid structures with 1 km spacing for each of the ship regions is approximately 45 million. This yields, with the 20 threat axes, approximately 9 x 10^8 cases to consider. The parameter space would be further increased when considering alternative firing policies given a TG configuration and main threat axis. Employment of a Monte Carlo simulation would require multiple runs per case, e.g. 50 to 100, further increasing the overall number of simulation runs required to exhaustively cover the complete parameter space.

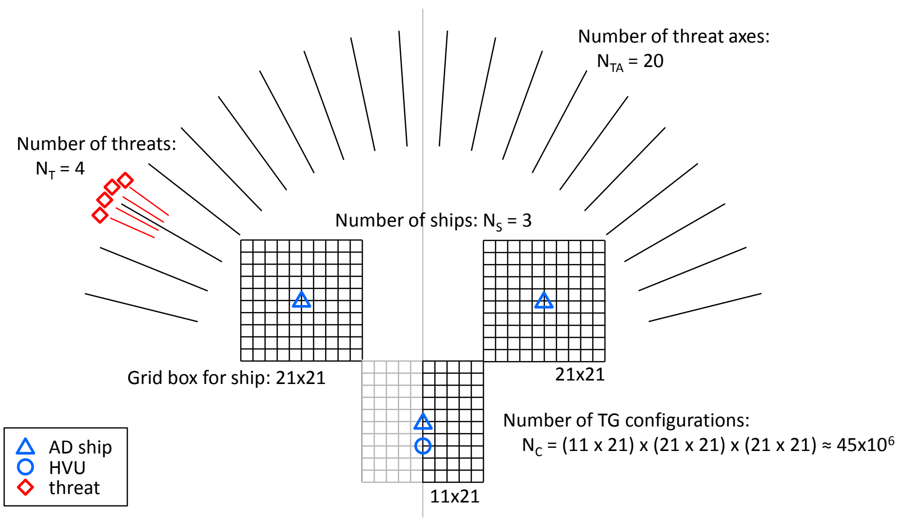

Figure 4. *Scenario factors defining the parameter space for the TG AD study*

Analysis approach

The physics-based AD model is particularly suitable for capturing technical system/subsystem aspects of surface-to-air engagements including radar detection and tracking of threats, threat evaluation and weapon assignment, and SAM launch, fly-out and intercept. The scenario size, based on number of threats and defending ships, and the requirement for multiple runs of the Monte Carlo simulation would restrict the amount of the parameter space associated with varying threat axis and TG ship stationing that could be explored. An alternative analysis approach was therefore adopted in which the high fidelity simulation would be used to perform single defender – single threat engagement assessments and a simplified model, drawing on these engagement assessments, would be used to analyse ship stationing and firing policies for a range of threat axes. Outputs from the ship stationing analysis would then be used to define scenarios for more focused study with the high fidelity simulation. At this time only hard kill involving employment of SAMs is being considered, although the AD model permits soft kill to be also included.

The following analysis approach was therefore adopted:

1. Use the high fidelity AD model to construct single defender – single threat Missile Engagement Zones (MEZs) for each SAM/threat combination.
2. Employ the MEZs in a lookup manner to determine the optimal TG firing policy given a specific threat presentation and TG configuration.
3. Conduct ship stationing analysis by varying TG ship positions to determine optimal locations for given threat axis assumptions.
4. Cross-validate ship stationing results with and perform focused runs using the high fidelity model.

To date, the study has progressed steps 1 to 3 for selected threats, methodology details of which are given below. For step 4, cross-model validation between the SSM and AD model for single defender – single threat engagements and SME validation of SSM results for complete scenarios have been performed. Further validation is planned as larger scenarios are implemented in the AD model.

Metrics

Standard metrics, as described in [9], are adopted for evaluating TG configurations. The following metrics are used to evaluate each threat engagement by the TG:

- Pareto Efficient Boundary (PEB) showing threat P_K as a function of cost for alternative firing policies. The PEB is used to select a firing policy option which achieves a desired P_K whilst minimising cost. A cost function is used to derive cost from SR and LR SAM expenditure. This function permits different cost values for SR and LR SAMs to be represented, e.g. based on associated loadouts and/or the value of each missile type to the TG Commander.
- Achieved threat P_K for the selected firing option, noting this may vary from the desired threat P_K. In some cases the desired P_K may not be achieved while in other cases the achieved P_K will be for the firing policy option with P_K that exceeds but is closest to the desired P_K.
- Mean SR and LR SAM expenditure and associated cost for individual ships and the TG.

For a specific threat axis with *N* threats the following metrics are used to assess TG performance against the air raid:

- Mean P_K for a single threat.
- Mean number of SR and LR SAMs expended and associated cost.
- Mean number of threats killed.
- Probability of Air Raid Annihilation (PARA), probability all threats in the raid are killed.
- Probability distribution for *i* threats killed, *i* varying from 0 to *N*.
- Probability distribution for cost spent derived from SR and LR SAM expenditure.
- Distribution showing number of threats killed as a function of expended cost.

For a set of main threat axes:

- Mean SR and LR SAM expenditure and associated cost as a function of bearing angle for the main threat axis.

Aggregation using a threat axis probability distribution is applied to the results for the range of specific threat axes considered. This yields associated metrics for the TG configuration given the threat axis assumptions.

For a given set of TG configurations with varying ship positions and a selected threat axis assumption distribution, the optimal TG configuration is obtained by finding that which minimises cost whilst achieving the desired P_K for each threat. If no configuration achieves this, selection is made based on maximising threat P_K's, noting this places no constrains on cost.

MEZ construction using physics-based simulations

The naval AD model was employed to construct a set of MEZs, one for each combination of SAM system and threat type. The primary characteristics for a MEZ, illustrated in Figure 5, are the cross-range of the approaching threat and the threat location, in the down-range direction, for events associated with SAM firing initiated, SAM launch, SAM intercept and subsequent engagement assessment. The MEZ is axisymmetric about the vertical axis through zero cross-range.

Each MEZ was constructed by running the AD model multiple times for each potential launch point with non-negative cross-range for a grid structure oriented about the defender. AD model outputs were then used to obtain the launch point, mean single-shot P_K and associated fly-out time for SAM intercept with the threat. Symmetry considerations were used to provide results for negative cross-ranges. The results are captured in the form of look-up tables and provide a range of firing options, with associated timelines and single-shot P_K's, for engagement of a threat with a given cross-range.

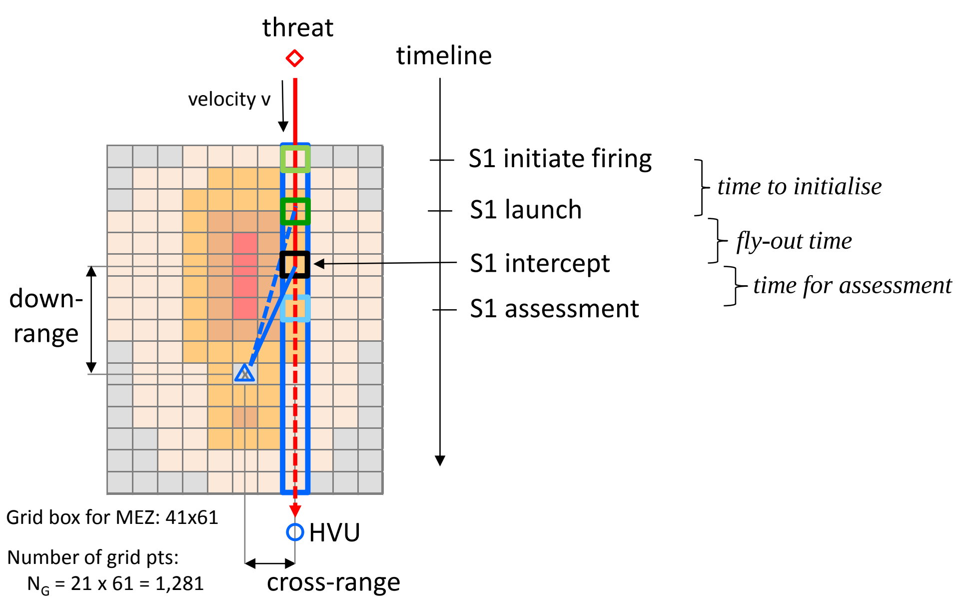

Figure 5. *Missile Engagement Zone (MEZ) characteristics*

Each MEZ captures key aspects of a complete single shooter – single target engagement as represented in the naval AD model. These aspects include:

- Detection by a search radar, which must be achieved in order for a launch at the desired point in the MEZ to occur.
- Tracking by a fire control radar, which must be maintained for mid-course guidance or until the SAM acquires the target, if it is active.
- Launch and fly-out of the SAM with associated fly-out time.
- Intercept of the SAM with the target with a kill assessment based on the SAM's warhead capabilities and endgame conditions.

Multiple runs of the AD model permit a single-shot P_K to be estimated for each grid cell in the MEZ from the ratio of successful engagements to the total number engagements for the cell.

Various MEZ dependencies include:

- Target type and flight profile (height & velocity).
- Search and tracking radar performance.
- SAM guidance and flight capabilities.
- SAM warhead effectiveness against the target).
- Environment conditions (duct height, atmosphere).

Ship stationing analysis

Steps 2 and 3 of the study approach involve a ship stationing analysis conducted through simplified modelling implemented in a Ship Stationing Model (SSM). SSM inputs include:

- SAM – threat MEZs produced from the naval AD model. The SSM incudes functionality to visual these MEZs, an illustrative example of which is given in Figure 6.
- Number of AD ships and threats.
- Range of TG configurations.
- Range of threat axes and threat axis assumptions.
- Firing doctrine constraints.
- AD coordination.

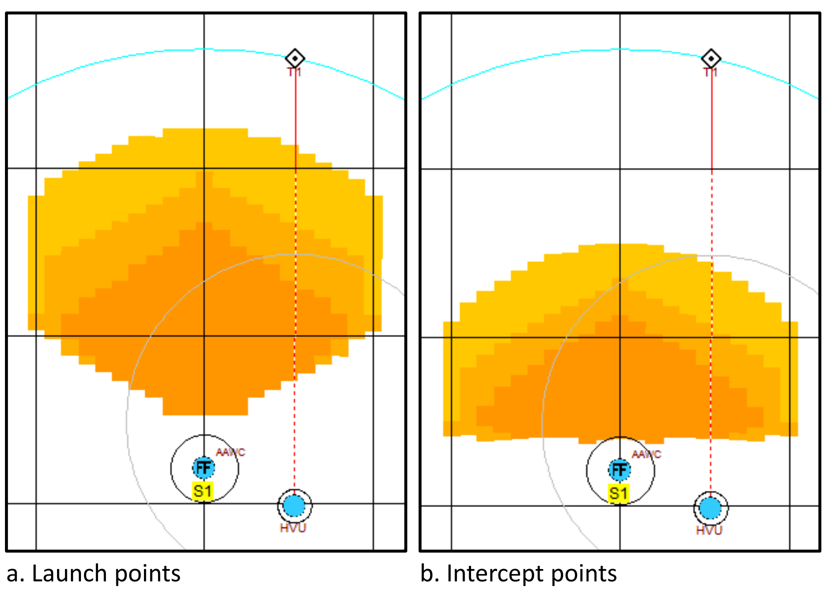

Figure 6. *Illustrative MEZ in the Ship Stationing Model:*
a. MEZ showing launch points; b. MEZ showing intercept points

The SSM is used to conduct ship stationing analysis through the following:

- For each threat axis/TG configuration, determine the optimal firing policy to achieve a desired P_K against each threat that minimises cost associated with missile expenditure.
- For each threat axis assumption, determine the TG configuration that minimises mean cost associated with mean missile expenditure.

This functionality is further described in the following two subsections.

Determination of an optimal firing policy

MEZs from the naval AD model are used in the SSM to construct a tree of firing policy options for each threat in the air raid. Given a threat's cross-range to each of the AD defenders, the possible engagement cells from the defenders' MEZs are meshed together to form a sorted set in time. This set is then used to recursively construct a tree of firing policy options. The tree starts with selection of the first engagement cell for which a firing policy is created using initiate, launch, intercept, and assessment times from the associated MEZ. Engagement cells with initiate times following the assessment time are then used to consider follow-on shots to the first Shoot Look, each of these forming sub-branches for further firing policies. This process is repeated for subsequent engagement cells, each being considered for the first shot. Tracked with each firing policy are the shots made, missile expenditures by the engaging MEZs, and the associated P_K's. The tree construction procedure is illustrated in Figure 7 for a threat ingressing through LR and SR MEZs for two ships. A firing policy comprising a shot fired in engagement cell 11 followed by an assessment and second shot fired in cell 24 is highlighted in the set of meshed engagement cells. This firing policy is also shown in the tree of permissible policies on the right with $S_{11}LS_{24}$ indicating Shoot in cell 11, Look, Shoot in cell 24, The associated P_K and cost are computed from the engagement cells and stored with the firing policy. The procedure accommodates multiple shots in a salvo to a maximum salvo size for each of the MEZs and applies launcher availability constraints for determining permissible launches.

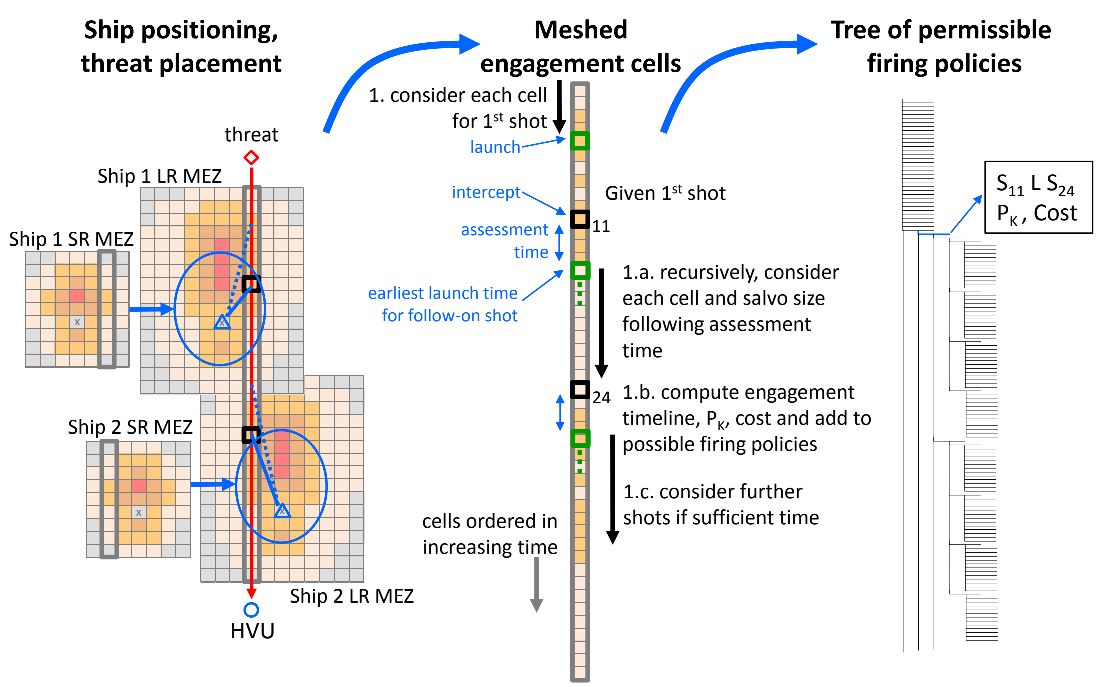

Figure 7. *Determination of a firing policy tree for a target ingressing through multiple MEZs*

The resulting set of permissible firing policies can be used to form a plot of mean cost, representing overall SAM expenditure, and P_K for each firing policy. Such a plot for an illustrative example is shown in Figure 8. A feature of these plots is the presence of a PEB, the solid curve to the upper left in Figure 8. All firing policy points to the lower right of this curve are dominated by a firing policy on the curve, i.e. there is at least one point on the PEB which has higher P_K and lower mean cost to each of the dominated points. Valid solutions for selecting a firing policy which maximises P_K and minimises mean cost must therefore lie on the PEB. The PEB is therefore used to select a firing policy which achieves the desired P_K but minimises mean cost, this being illustrated in Figure 8 for a main threat axis at 0^0 bearing to the TG. In this figure a desired P_K of 0.95 is used to find the closest firing policy on the PEB, this comprising $s_2Ls_2Ls_1s_1$: ship 2 shoots a SR SAM, look, ship 2 shoots a second SR SAM, look, ship 1 shoots two SR SAMs. The mean cost associated with this firing policy is 0.941, this being computed using a SR SAM cost of 0.5 per missile and single shot PKs of 0.55 for the first SAM fired and 0.52 for the remaining shots.

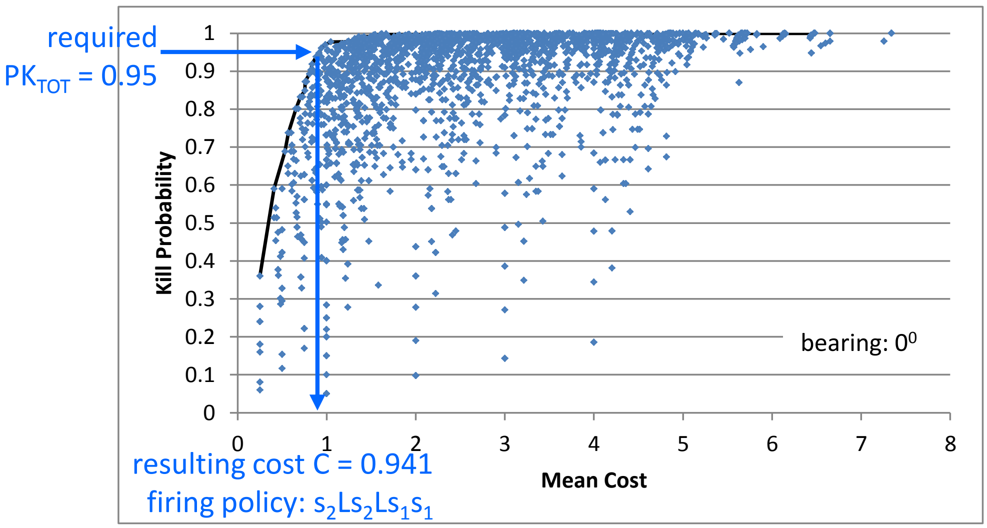

Figure 8. *Pareto Efficient Boundary (PEB) for firing policy tree*

By default, a brute force construction of the firing policy tree involves creation of a large number of leaf nodes, some 232876 nodes for one engagement example using the illustrative MEZ from Figure 6 and a second SR SAM MEZ. The Pareto dominance condition can be used to perform leaf pruning during tree construction to reduce this number to 138027 nodes. Further pruning, based on terminating sub-branching once a parent node is found to be dominated, reduces the final firing policy tree to 2792 nodes, the mean cost - P_K graph for which is shown in Figure 8.

The above approach provides a means to select a firing policy which achieves a desired PK against a single threat whilst minimising cost. Launch times associated with this selected firing policy are recorded and act as constraints for consideration of further threat engagements. The procedure considers the threats in the raid sequentially, but varies the threat prosecution order to find that which minimises overall mean cost for countering the air raid. A final PARA is calculated from the achieved P_K's for the threats.

Assessing TG configurations in context of threat axis assumptions

The previous subsection detailed how an optimal firing policy, with associated mean cost and PARA, can be obtained for a given TG configuration and main threat axis. As indicated earlier, there is uncertainty as to what the actual threat axis for an air raid would be. Results for each TG configuration are therefore obtained for a range of threat axis values, e.g. the threat axis varying from -90^0 to 90^0.

Illustrative results are presented in Figures 9 and 10 for three TG configurations defending against an air raid of four threats. A desired P_K of 0.95 against each threat was used for the calculations. Figure 9 shows the three configurations and engagement results for a single threat axis. Figure 10 shows cost results, reflecting SR and LR SAM expenditures, as functions of threat axis bearing angle for the three TG configurations from Figure 9. The blue sections of these graphs indicate the desired P_K has been achieved by the selected firing policy option, while the red sections indicate this P_K has not been achieved. Figure 10.a indicates TG configuration 1 has best results when the threat axis has bearing angle in the vicinity of 0^0, while TG configuration 3 in Figure 10.c does best when the angle approaches +/- 90^0.

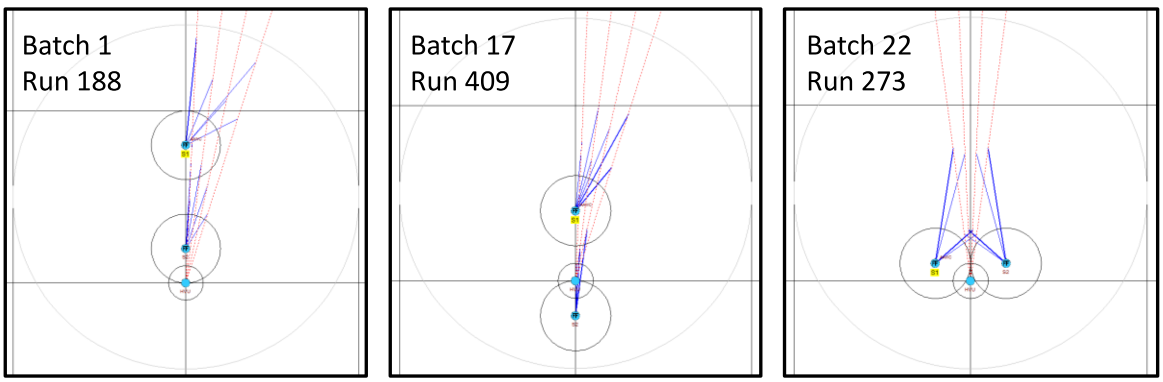

a. TG configuration 1 b. TG configuration 2 c. TG configuration 3

Figure 9. *TG configurations with engagement results*
a. TG conf. 1, 10^0 threat axis; b. TG conf. 2, 10^0 threat axis; c. TG conf. 3, 0^0 threat axis

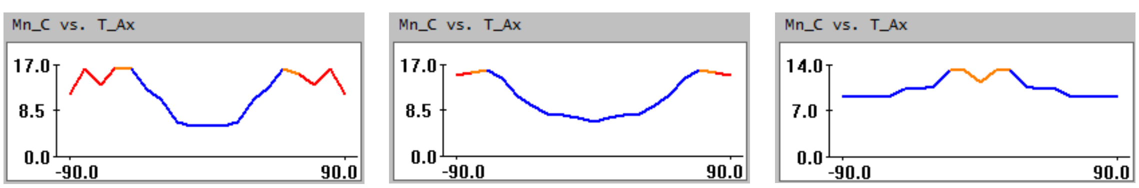

a. TG conf. 1 b. TG conf. 2 c. TG conf. 3

Figure 10. *TG configuration cost results as functions of threat axis bearing angle:*
a. TG conf. 1; b. TG conf. 2; c. TG conf. 3

The results of Figure 10 give an indication of how TG configurations compare to each other for varying threat axis bearing angle, but they do not take into account any a priori information that the TG Commander may have on the threat. This aspect has been introduced through the definition of a Threat Axis Distribution (TAD), reflecting the TG Commander's assumptions on where the threat will come from. A TAD defines a probability distribution across a range of threat axis bearing angles. Figure 11 shows plots for three example TADs where the horizontal axis is the threat axis bearing angle and the vertical axis is the probability. TAD 1 "narrow" reflects a narrow distribution centred about 0^0 bearing angle (reflecting a high certainty as to the direction of the threat), TAD 2 "wide" is a wider distribution (less certainty as to where the threat will come from); and TAD 3 "flat" is a flat distribution (all bearing angles having equal probability).

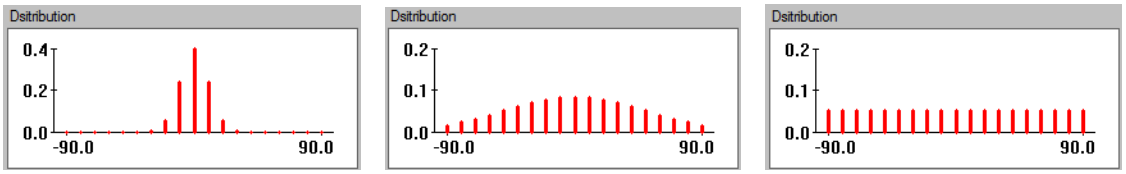

a. TAD 1: “narrow” b. TAD 2: “wide” c. TAD 3: “flat”

Figure 11. *Threat Axis Distributions (TAD) reflecting a commander’s threat assessment a. narrow TAD; b. wide TAD; c. flat TAD*

A TAD can be used to aggregate the results of Figure 10 to produce a single mean cost, PARA and other metrics for each TG configuration. Illustrative results using this approach for the three TG configurations in Figure 9 and the three TADs in Figure 11 are shown in Table 1. Shown for each TG configuration, indicated by batch and run number, and TAD are the mean PARA (P_ARA), mean PK (Mn_PK) against a single target, mean cost (Mn_Cost) for countering the air raid, and mean number of targets killed (Mn_TK) for a single raid. The mean cost information can be used to determine, for a set of runs, the preferred TG configuration for each threat axis assumption. The mean cost information from Table 1 is reorganised in Table 2 and shows that TG configuration 1 (batch 1 run 188) is preferred for TAD 1, TG configuration 2 (batch 17 run 409) for TAD 2, and TG configuration 3 (batch 22 run 273) for TAD 3. Tables 1 and 2 are produced directly by the SSM to present results for a set of loaded runs.

Table 1. *Numerical results for the TG configurations*

RS Batch Mn_C aggregation over Threat Axis Distribution

BATCH_RUN	TAD	P_ARA	Mn_PK	Mn_C	Mn_TK
1_188	1	0.845	0.959	5.735	3.835
1_188	2	0.813	0.949	10.472	3.797
1_188	3	0.780	0.939	11.762	3.757
17_409	1	0.840	0.957	6.893	3.830
17_409	2	0.831	0.955	9.778	3.818
17_409	3	0.817	0.951	11.320	3.802
22_273	1	0.813	0.950	12.512	3.798
22_273	2	0.829	0.954	11.174	3.817
22_273	3	0.832	0.955	10.577	3.820

Table 2. *Preferred TG configuration for each TAD*

RS Batch Mn_C aggregation over Threat Axis Distribution

BATCH_RUN	TAD 1	TAD 2	TAD 3
1_188	5.735	10.472	11.762
17_409	6.893	9.778	11.320
22_273	12.512	11.174	10.577

A systematic exploration of TG configurations can be used to produce heat maps for each TAD, such as those shown in Figure 12 where over 40000 TG configurations were considered. The heat maps permit identification of optimal ship locations, shown in the red grid cells, for given assumptions on the main threat axis. In context of the examples here, Figure 12.a shows TG configuration 1 is preferred across all configurations considered for the “narrow” TAD 1, this comprising one AD ship close to the HVU and the second positioned in the direction of the centre bearing angle of the TAD. Figure 12.b shows the preferred TG configuration for the “wide” TAD 2 is to move the first ship behind the HVU and the second ship closer to the HVU. For the “flat” TAD 3, the preferred TG configuration in Figure 12.c is to have the two ships close but on opposite sides of the HVU.

The heat map for each TAD also gives an indication in the yellow/orange grid cells of permissible locations which will allow the desired P_K to be achieved for varying mean costs. Increasing shades of yellow to orange indicate positions with reduced cost to achieve the desired P_K. The blue regions in these heat maps indicate a firing policy solution cannot be found which meets the desired P_K, the shading now indicating the achieved P_K, with darker shading being higher. The results show, as intuitively expected, the region of permissible positions reduces as the uncertainty in the threat axis increases.

The optimal ship positions identified through the heat maps are those that permit engagements with assessments between shots and lower mean cost, this being weighted across the uncertainties in the threat axis. For the TG with two defenders, these positions generally involve the two ships being placed in a line towards the central bearing angle of the TAD with one ship pushed up threat. This buys increased battlespace permitting better opportunity for SLS firing policies using SR SAMs. Increasing uncertainty in direction of the threat axis results in the ships being pulled closer to the HVU. The battlespace is now reduced and resultant costs are increased.

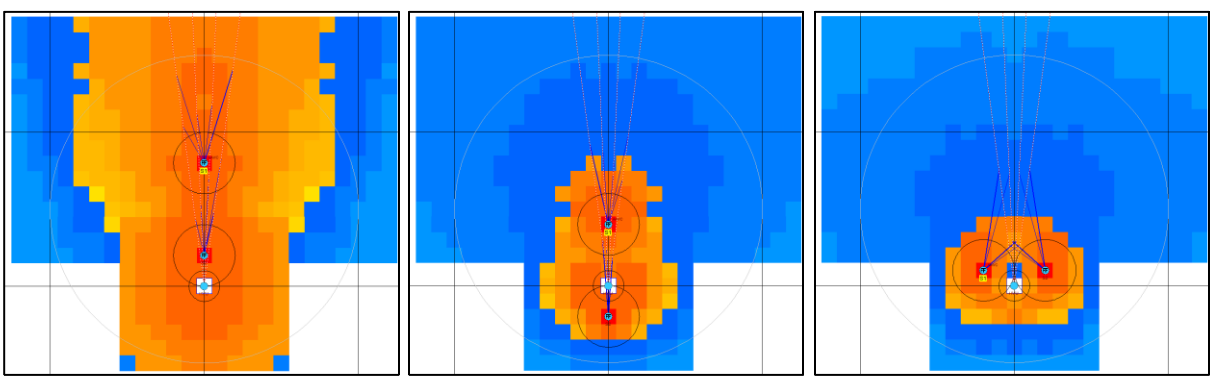

a. TAD 1 with TG conf. 1 b. TAD 2 with TG conf. 2 c. TAD 3 with TG conf. 3

Figure 12. *Heat maps showing optimal TG configurations: a. TAD 1; b. TAD 2; c. TAD 3*

A further feature of the heat maps is their ability to support what-if analysis. This is illustrated in Figure 13 for the three TADs. In this example the location of one ship’s position has been fixed in the green grid cell. This situation could arise for a TG when other mission objective’s require one ship to be at a specific location, e.g. to prosecute a submarine threat. In Figure 13.a for TAD 1, the updated heat map for the second ship identifies an optimal location in red, as well as a set of permissible locations in yellow/orange so that the desired P_K can be met. Figures 13.b and c show updated heat maps for TADs 2 and 3 for which the desired P_K cannot be met and the second ship’s location has been selected to maximise P_K.

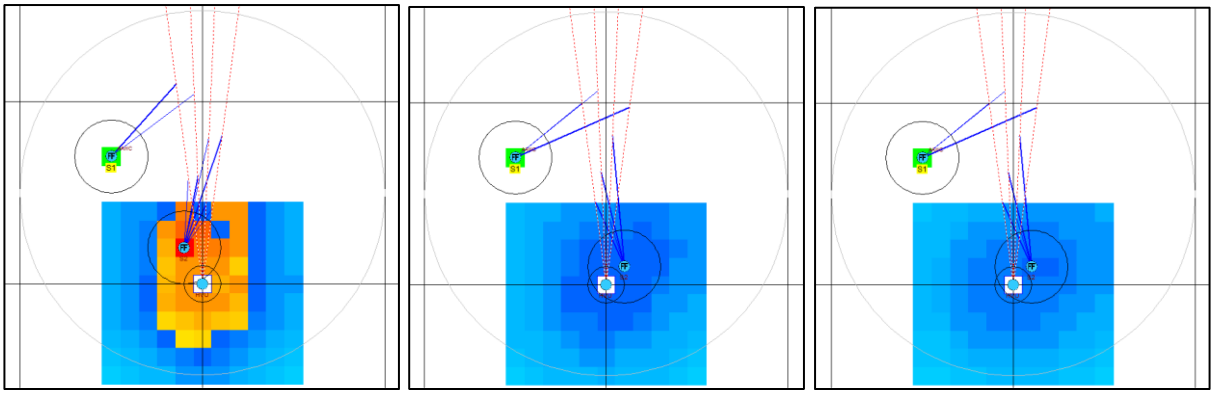

a. Optimal TG conf. for TAD 1 b. Optimal TG conf. for TAD 2 c. Optimal TG conf. TAD 3

Figure 13. *Revised heat maps with one ship's position fixed: a. TAD 1; b. TAD 2; c. TAD 3*

Discussion

The theme raised in the introduction to this paper was how physics-based simulations can complement traditional OA techniques for undertaking studies. The TG AD case study provides an example of where a physics-based simulation was leveraged to underpin an OA study. For the study, the high fidelity AD model has the underlying analysis capability to fully address study objectives. It usage, however, is compounded by a large parameter space that, even with the simulation horsepower afforded by HTC, could only be partially explored. The TG AD study therefore adopted a two-phased approach: first, to develop and use a simplified model drawing on engagement outputs from the AD model to explore the full parameter space; and second, to perform cross-model validation and focused runs on full scenarios using the AD model.

The TG AD study is demonstrating the value of complementing high fidelity models with simpler models for exploring a large parameter space. The SSM incorporates an engagement meta-model developed directly from the AD model and used in a fast running, lookup fashion. On top of this meta-model, the SSM provides two additional areas of functionality: an engagement planner that determines an optimal[2] firing policy for a TG to counter an air raid, and methods to aggregate results across a range of uncertainties through which TG configurations can be identified that provide the best defence for given threat axis assumptions. This approach permits larger scenarios to be more easily considered and the analysis to be extended to larger parameter spaces. It does, however, raise specific study validation requirements, which are now briefly discussed.

Verification and validation of the SSM engagement meta-model has been performed through direct comparison of results with outputs from the AD model for single defender – single threat scenarios. SSM functionality was developed to visualise AD model MEZs, engagement timelines and SAM launcher scheduling. SME review throughout this process has helped in the definition and production of MEZs by the AD model, and in the identification of functionality requirements for the SSM to use them.

The SSM engagement planner is new functionality for which validation, to date, has been primarily through SME review. Data inputs for the planner permit a varying degree of AD

[2] More specifically, a "pseudo-optimal" firing policy obtained by sequential processing of targets with identification of the preferred target prosecution order.

coordination and firing policy constraints to be specified. Planned, cross-model validation with the AD model for multi-defender – multi-threat scenarios is expected to identify differences, which will be subject to SME review and possibly lead to further study refinement. The AD model's engagement planner provides mechanisms for specifying coordinated AD across a group of ships, but its usage has been more focused on smaller scenarios. Future work will therefore support development of analysis capability in this area through using the AD model`s built-in engagement planner and external C2 models, such as [8].

SSM development has been based on representing key aspects of AD engagements in a simple manner so as to permit rapid analysis of larger scenarios. Current functionality reflects implicit and simplifying assumptions, e.g. all threats target the HVU, lack of soft kill, or no interference in SAM guidance given presence of multiple targets. Such assumptions may be relaxed through future work, either by expanding definition of MEZs generated by the AD model, or increased functionality in the SSM with calibrations from the AD model. In both cases the SSM can act as an aggregating tool permitting extension of analysis to TG operations with conditions of uncertainty.

A further aspect that became very evident through the TG AD study was the analysis potential provided by HTC. Some illustrative numbers help to put this in perspective. For production of a MEZ the AD model, with radar modelling using the APM, takes 5 minutes to simulate a single engagement. Using a 40 x 60 grid with 50 Monte Carlo runs per grid cell, this results in 120,000 runs taking 10,000 hours of processing. The run batch can be completed in 10 hours when executing this over a computing infrastructure of 1,000 processors. Typical turn-around times for the AD model producing MEZs for the TG AD study were 1 to 2 days, with a further 2 to 3 days for data parsing and storage.

The CFMWC's HTC capability was also employed by the SSM for ship stationing analysis. The SSM could typically analyse a TG configuration across a range of 20 threat axis bearings in one minute. Batches of up to 100,000 TG configurations were submitted to the HTC, these taking 1.5 to 2 hours to complete. Subsequent transfer of output files could take up to 6 hours. Given this latter overhead, the SSM was modified to log minimal data for a TG configuration and write this information directly to a storage database. This information was used to produce heat maps from which TG configurations of interest, with full logging of output data, would be obtained by rerunning the SSM. HTC therefore permits very large parameter spaces to be analysed in a systematic fashion, although data management becomes more challenging. Increases in problem parameter spaces can also quickly outstrip a HTC capacity. The 45×10^6 TG configurations in Figure 4 would take over 500 hours on a 1,000 computer cluster. Adding a fourth ship with 21x21 positions would then require 11 years of processing. Care, therefore, has to be taken in study design to focus modelling and computing capabilities on specific areas of the problem parameter space so as to gain best benefit.

The approach taken by the TG AD study follows practices being adopted by the M&S community, e.g. Design of Experiments (DoE) and data farming as developed in [10]. One aspect introduced in the TG AD study is the usage of a simplified modelling layer over high fidelity modelling performed by a physics-based simulation. Data from the simulation, in this context, is "farmed" so as to feed into the second modelling layer, which, in turn, is further farmed to facilitate exploration of a large parameter space. The resulting modelling/data "ecosystem" is proving to be useful for analysis of TG AD operations. Further work to perform cross-model validation and focused runs using the high fidelity simulation will help to establish the robustness of this approach.

Concluding remarks

A two-layered approach comprising simplified modelling using the SSM underpinned by detailed modelling using the physics-based AD model is being applied to a study on TG AD operations. The approach is permitting exploration of a large parameter space for the problem that is not possible using the AD model alone. The embedding of MEZs from the AD model into the SSM has provided a firm basis and credibility for the resulting ship stationing analysis. The production of MEZs using the AD model has also gained endorsement with CFMWC M&S now applying this approach to support development of plans for an upcoming at-sea trial.

Physics-based simulations are established as an M&S capability to supporting engineering requirements for military clients. Coupled with HTC, they offer increased capability for undertaking of OA studies. There is still a role for simplified models to complement the physics-based simulations for studies. They can help provide a structured approach to undertaking a study, map results directly into study metrics and permit larger parameter spaces to be explored. A final benefit is potential exploitation of algorithms or outputs from the simplified models to help improve current modelling capability. Regarding this last point, two aspects from the SSM to be further explored are usage of TG configuration heat maps to support tactical planning and versatility of the engagement planner for determination of TG AD firing policies.

References

[1] Wagner, D.H., W.C. Mylander & T.J. Sanders, Naval Operations Analysis, 3rd ed., Naval Institute Press, Annapolis, MD, USA, 1999.

[2] Alberts, D.S., R.E. Haynes, Code of Best Practice for Experimentation, CCRP Publication Series, DoD, USA, 2005.

[3] Giadrosich, D.L., Operations Research Analysis in Test and Evaluation, AAIA, Washington, DC, USA, 1995.

[4] Skolnik, M.J., Radar Handbook, 2nd ed., McGraw-Hill, Inc., New York, NY, USA, 1990.

[5] Strickland, J., Missile Flight Simulation, 2nd ed., Lulu Press, Inc. Raleigh, NC, USA 2015.

[6] Barrios, A.E., Considerations in the development of the Advanced Propagation Model (APM) for U.S. Navy applications, Proceedings of the International Conference in Radar (RADAR 2003), Adelaide, Australia, 2003.

[7] Manual of the ICAO Standard Atmosphere. ISO Doc. 7488/3, 1993.

[8] Benaskeur, A.R., F. Kabanza & E. Beaudry, CORALS: A Real-Time Planner for Anti-Air Defence Operations, ACM Transactions on Intelligent Systems and Technology, Vol. 1, No. 2, Article No. 13, Nov. 2010.

[9] Nguyen, B., Comparison of metrics for missile defence between perfect coordination and no coordination, DRDC-RDDC-2015-R228, 2015.

[10] Data Farming in support of NATO, Final Report of Task Group MSG-088, STO/NATO, 2014.