Operational Planning Decision Support using Multi-Dimensional Data Farming

Authors

MSG-186: LCDR Bernt Åkesson, MRS Maude Amyot-Bourgeois, DR Sreerupa Das, DR Gunar Ernis, DR Andrew Gill, LTCOL Esa Lappi, DR Bao Nguyen, MR Chris Rolfs, LTCOL Stephan Seichter, MS Lynne Serré, COL Vadym Slyusar, DR Alessio Vaghi, MR Peter Volbach, MR Alexander Zimmermann

ABSTRACT

Multi-Dimensional Data Farming (MDDF) uses Machine Learning (ML) to automate data farming to allow improved and faster decisions in highly complex multi-scale, multi-domain, and multi-level hybrid war campaigns. This has significant utility when used in support of Operational planning, allowing multiple Courses of Action (CoA) to be rapidly developed and evaluated prior to the execution of any operation. Using MDDF enables decision-makers to explore the problem space and identify multiple optimal solutions significantly faster than current techniques.

MSG-186 has applied MDDF in a sand-box environment to an illustrative combined strategic campaign and tactical hybrid warfare operation resource allocation problem considering the balance between local and global optimal solutions. We have tested the technical feasibility of implementing MDDF within the Federated Mission Network operational environment at Coalition Warrior Interoperability Exercise (CWIX).

Through MDDF, we aim to show it is possible to combine ML techniques exploring operations at multiple scales (Multi-Domain Operations and targeted-fidelity modelling) and optimize the strategic/operational level goal, by selecting the correct resource allocation scheme at the tactical level. This paper describes an ML-based assistant able to conduct MDDF experiments and optimization tasks on an automated basis, which was examined in detail during CWIX in 2024.

1.0 INTRODUCTION

Data Farming (DF) is the analytical process of exploring a simulation of a real dynamical system in operation to infer properties of that system which may inform decisions about its efficient and effective management. Figure 1 summarizes the DF process into four main steps: create the scenario to investigate, define values and ranges of the scenario parameters, run the scenario, and analyze and visualize the data.

Previous research conducted by MSG-088, MSG-124 and MSG-155 has codified the DF concept [1], turned the concept into actionable DF decision support [2] and developed data farming services (DFS) as a web-based environment enabling analysts to conduct DF experiments on their system of interest [3 – 6]. Traditional DF is an interactive man-in-the-loop process, primarily to support military decision-making throughout the development, analysis and refinement of Courses of Action (COA). However, the DF process may be time consuming, even with the software tools or dedicated services developed so far. Military decision-makers want and need to make decisions fast. MSG-186 is extending this research: 1) to support a multi-model, multi-operation approach to simulation-based decision-making in highly complex multi-scale, multi-domain, and multi-level hybrid war campaigns and 2) to automate some of the steps in the DF process by leveraging innovative artificial intelligence (AI) techniques. This Multi-Dimensional Data Farming (MDDF) approach has been outlined at the Symposia MSG-197 [7] and MSG-207 [8].

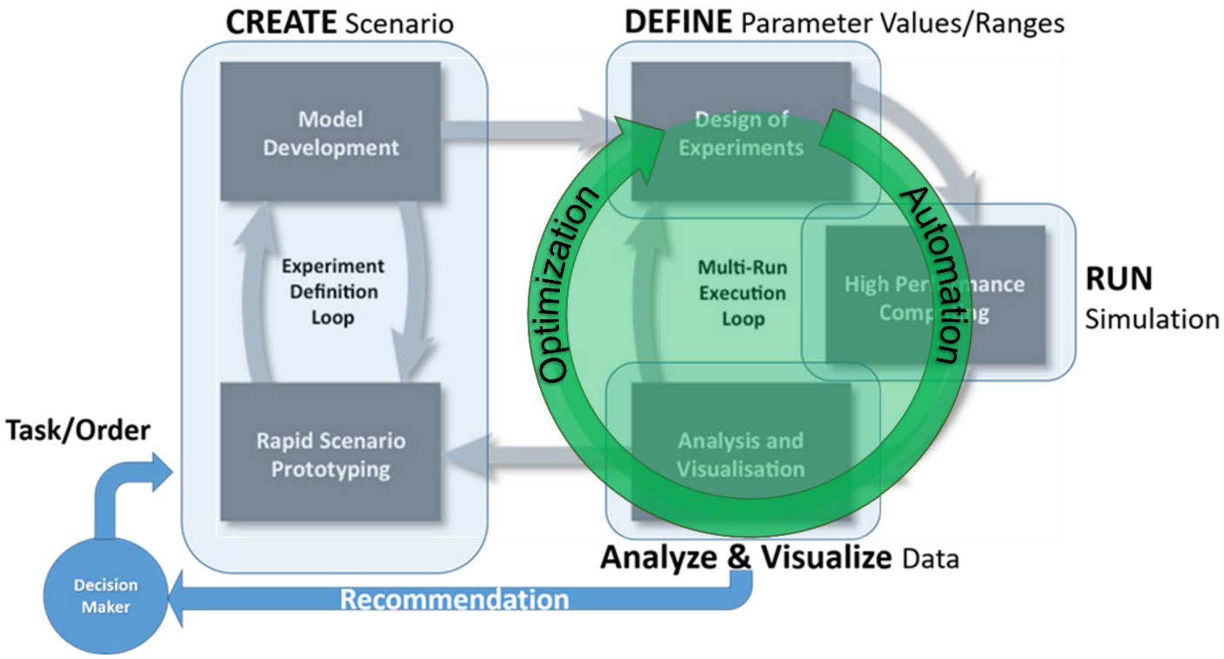

Figure 1: Data Farming Decision Process distinguished in four main steps interconnected by two iterative loops: Experiment Definition Loop and Multi-Run Execution Loop. MDDF allows to automate and optimize the latter.

The previous DF research task groups have actively participated in the Coalition Warrior Interoperability eXercise (CWIX). MSG-124 demonstrated the applicability of DF to military decision making at CWIX in 2016 and 2017. MSG-155 tested the DFS architecture during CWIX cycles 2018 – 2020 to ensure interoperability in a fully Federated Mission Networking (FMN) compliant integration. This testing at CWIX has continued with new tools developed by MSG-186. In this paper we present results from implementing MDDF within the FMN environment at CWIX 2024. In Section 2 we introduce the multi-scale resource allocation problem, consisting of a multi-faction, multi-domain campaign in which a tactical-level operation has been embedded. In Section 3 we present results from CWIX 2024 on the use of MDDF to this problem. In Section 4 we discuss the impact on MDDF.

2.0 HYBRID WAR SCENARIO & RESOURCE ALLOCATION PROBLEM

NATO defines unconventional warfare as military activities conducted through or with underground, auxiliary or guerrilla forces to enable a resistance movement or insurgency to coerce, disrupt, or overthrow a government or occupying power [9]. In the same standard, NATO also defines a hybrid threat as one that combines conventional, irregular, and asymmetric activities in time and space, and methods include disinformation, cyber-attacks, economic pressure, deployment of irregular armed groups and use of regular forces. The migrant crisis on the Belarus-Poland border in 2021 has been regarded by European Union (EU) and United States leaders as a hybrid attack [10]. As NATO has an extensive land border to non-NATO and non-EU states, an unconventional warfare campaign embedding a Border Operation scenario is an interesting case study for a MDDF experiment.

In this paper, we therefore consider a strategic Unconventional Campaign where two armed allied forces, Blue and Green, fight against a Red one in both the physical and cyber domains. Furthermore, Red launches a hybrid operation on Green, by additionally triggering a Border Operation, in support of two strategic objectives: to force Green to divert resources away from the campaign and to lower Green public opinion (both thus reducing the chances of campaign success). Our case study aims to illustrate how tactical operations may influence campaign outcomes through MDDF.

In [7] we formulated a multi-faction multi-domain model (called ACE) describing the Unconventional Campaign. The ACE model relies on extensions of classical Lanchester and epidemic based equations to describe attritions in physical and cyber domains. In [8] we developed the Border Operation in MANA [11], focusing on a small segment of a fictitious border zone with terrain features including roads, fields, evergreen forests, and frozen bodies of water. Migrant agents cross the border at various times and if not intercepted, move deep into Green territory towards a safe destination. Green controls its border with a combination of fixed and moving border patrol agents, sensors (small UAVs patrolling the border and mini-UAVs organic to some of the border patrols) and fences. Sensor detections communicate the position to border patrols, which move to intercept and detain the detected migrants. While apprehending migrant agents, they are unable to intercept others, but once completed the patrol returns to its earlier position or patrol route. The simulation ends when all non-intercepted migrant agents reach the safe Green area.

The decision factors for the Border Operation relate to the selected border control assets and the measure of effectiveness is the proportion of intercepted migrants. We link the inputs and outputs between the campaign and the Border Operation, as follows. Choosing a setting for each of the decision factors allows a total cost for that COA to be determined. When the Border Operation begins, the Green force is partitioned into a force continuing to prosecute the campaign and this COA allocated to the Border Operation. When the Border Operation concludes, those allocated resources return to the campaign. The final linkage between the Border Operation and Campaign is a function penalizing the initial supply rate of Green. This function encodes a proxy for Green public opinion, which decreases as the percentage of non-intercepted migrants increases indicating that the supply rate of Green to the Campaign diminishes if the Border Operation is less successful. The resource allocation problem for this combined campaign and tactical operation is to choose the COA which maximizes the probability of campaign success (Red being annihilated before Blue (or Green) according to the dynamics of the ACE Campaign).

3.0 CWIX 2024

Our goals in attending CWIX 2024 were to demonstrate the applicability and usability of the enhanced DFS-System (iDFS) in an FMN compliant environment and to address the analytical questions of if and how good the automation and optimization approach performs in such an operational environment.

3.1 OPERATIONAL PLANNING

In CWIX 2024, we extended the Border Operation scenario, by testing the ability to react to new information and adapt existing plans according to new simulation runs. This involved three geographically distributed groups: a Finnish operational planner, a German analysis cell and a Canadian modelling cell. The initial planning was created using the DF process, and the resource allocation was made under the assumption of migrants walking on foot. Then new information arrives: Red is going to provide fat bikes to the migrants, to increase their speed and ability to avoid interception by the border patrols. The new information changes the initial Border Operation parameters, implying more effort in the Multi-Run Execution loop to get insights to the adapted scenario.

The planning for CWIX 2024 was: 1) The Finnish operational planner receives information that migrants have fat bikes to rapidly cross the border, requiring to update the resource allocation solution. 2) The new information is passed via secure channel to the German analysis cell. 3) The required model modifications are passed to the Canadian modelling cell. 4) The updated model is distributed to the German analysis cell. 5) The

Design, Analyze and Visualize Experiments (DAVE) agent executes the optimization simulation loop. 6) The German analysis cell shares the updated optimal resource allocation solution with the Finnish operational planner. To compare traditional DF and the new approach based on automation and optimization, we first conduct the Multi-Run Execution loop with a static design of experiment (DoE), and then we apply DAVE.

3.2 DAVE: AUTOMATION & OPTIMIZATION APPROACH

DAVE consists of a stand-alone application for the automated execution of simulations as well as intelligent optimization and analysis of multi-dimensional scenarios. It allows decision-makers to simulate a conflict scenario in multi-stage steps and to carry out analyzes and make decisions based on AI-based recommendations. It will autonomously explore the multi-dimensional parameter landscape, optimizing the probability of campaign success while considering resource constraints and the potential negative impacts of poorly managed Border Operations. As a standalone tool, DAVE allows for the pre-setting of the initial campaign and the definition of an optimization goal. For all this DAVE is based on the DFS platform introduced in the prior MSG-155 [3 – 6]. It automates time consuming manual DFS processes, including DoE generation and the execution of simulation runs. By leveraging Bayesian Optimization (BO) techniques, it accelerates the development of meta-models within the Multi-Run Execution loop of the DF process, enabling faster identification of the best, worst, and most promising COAs in complex scenarios. This allows decision-makers to gain actionable insights efficiently.

The concept of DAVE is depicted in Figure 2. Based on a decision question, DoEs are generated, simulated, and analyzed. In contrast to the existing DFS instance, all steps are automated, and an AI-based optimization of the design generation is carried out iteratively considering the initial question until a specified target value/goal is achieved. All optimization progress is tracked and displayed via a dashboard, which gives information about current meta-model metrics and simulation results.

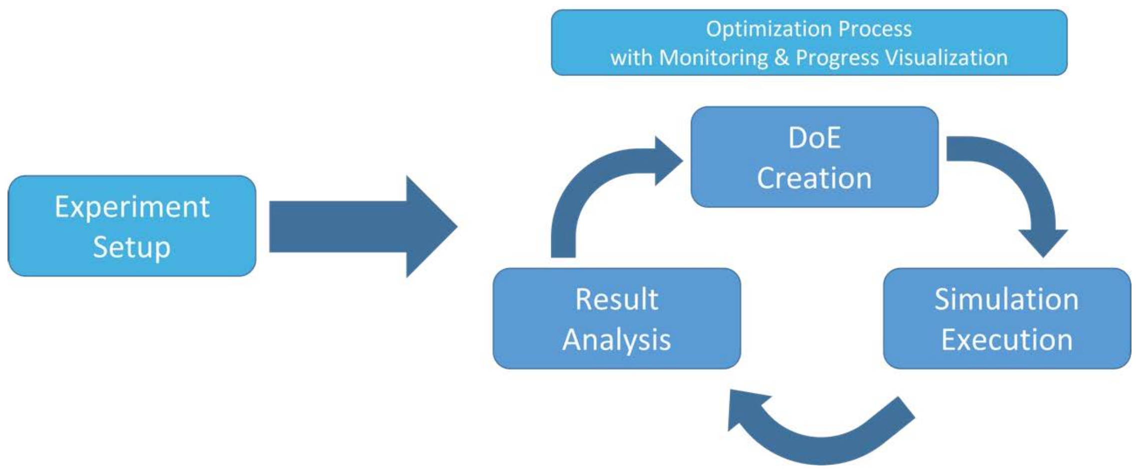

Figure 2: DAVE concept from experiment setup to optimization process.

For Operational Planning the “Experiment Setup” consists of initializing the Border Operation: Configuration of start time and conditions, as a reaction from "Red". After that the “Optimization Process” starts with configurable hyperparameters to identify optimal decision points. These include the optimization goal, an initial design for the influencing factors of the Border Operation, the meta-model for optimization, and an upper limit for the number of iterations. After starting the optimization, DAVE automatically generates a design in the background, performs simulations, and updates the optimization progress after each iteration.

During the process and after reaching a termination criteria the optimization is completed, and the results are analyzed. The optimization goal is considered and the "best", "most promising", and "worst" result can be selected and visualized.

3.2.1 AI-DRIVEN OPTIMIZATION APPROACH

BO is a framework for finding the optima of black-box functions, such as response functions from DF experiments, which are challenging to evaluate due to their computational costs [12, 13]. The framework consists of a surrogate model, used to approximate the (black-box) objective function relying on the current data set; an acquisition function, used to select the next point where the objective function should be evaluated; and the actual process of evaluating the black-box function. The acquisition function balances between the utilisation of already successful regions (exploitation) and the investigation of less examined areas (exploration) in order to learn more about the objective function.

DAVE exploits the BO framework for sequential decision-making. Its goal is to find a global optimum within a fixed parameter value range by means of iterative experiments to optimally improve the surrogate model with each new data point. To achieve this, BO comprises methods from DoE to create an initial experiment plan and Machine Learning (ML) to train a surrogate model. Experiment points or function evaluations are treated as data points and are used to generate the posterior distribution of the objective function. The most common models for defining the prior/posterior pairs are Gaussian Processes (GPs) [14], which can be regarded as generalisations of Gaussian distributions. Being an ML model, GPs can be adjusted to the problem at hand by choosing appropriate hyperparameters, which also involve designing so-called kernel functions. Kernel functions in GPs are a way to measure how similar two points are. They help the model to understand the relationship between different data points and use the similarity to make smooth predictions across data. The following describes all sequential steps in the BO applied for DAVE:

1) Create an initial experiment plan using a static DoE method and run simulations based on it.
2) Fit an initial GP to the simulated data.
3) Use the GP to select the next sampling point analytically by optimizing the acquisition function.
4) Evaluate the new sampling point by running the simulation given its parameters.
5) Update the GP with the simulation result of the new sampling point.
6) Repeat the acquisition, simulation and update steps (3 – 5) until the optimal value is reached or a pre-defined stopping criterion is met.

Since a good outcome from the model and optimization heavily depends on the data generated initially, the choice of an appropriate DoE prevents difficulties from the outset. Within DF a proven choice are Nearly Orthogonal Latin Hypercube (NOLH) designs, which typically strike a good balance between the required accuracy for estimating meta-model parameters and the number of experiments needed [15, 16]. Definitive Screening Designs (DSDs) have been developed to significantly reduce experiments while maintaining desirable orthogonality [17, 18]. These designs serve as a basis for applying BO methods.

3.3 IMPLEMENTATION & TESTING

The system components of DAVE are distributed among three different architectural layers: the data layer, the backend, and the frontend (see Figure 3). Each layer comprises containerized services that are required to control the overall optimization process. The data layer is central to underlying data structures that are needed for the optimization. The backend is responsible for carrying out the optimization and associated automated simulation runs of DFS experiments. The frontend provides a dashboard for controlling the

DAVE agent and for visualizing the process. DFS instead is composed of several backend services and repositories to conduct ACE and MANA simulations. An application programming interface conforming to constraints of the representational state transfer architecture (REST-API) allows DAVE to communicate with corresponding DFS services that are required during optimization.

3.3.1 ARCHITECTURE AND INTEROPERABILITY

Inside the Joint Force Training Centre (JFTC) during CWIX there were low and high confidential domains connected in a simulated way, via the SDoT security gateway [19]. Figure 3 illustrates domains and network connections between DFS and DAVE within CWIX infrastructure.

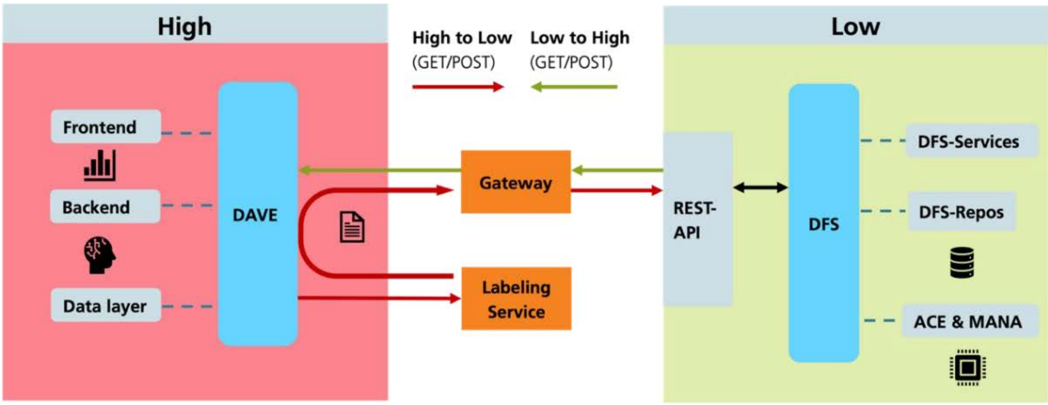

Figure 3: Illustration of DAVE and DFS interaction within CWIX infrastructure.

The DFS platform was placed in the low confidential domain, while DAVE was in the high confidential domain. The reason is that individual simulation runs are not classified as critical as the analysis results derived from all simulations. This holds especially with the derived optimal resource allocation results shared afterwards with the operational planner. To satisfy a secured and controlled interaction between DAVE and DFS the gateway was configured to control incoming GET and POST requests. Any request sent from the high to the low confidential domain must first be labelled by the labelling service. Once the label has been generated for the request, it is resent to the gateway. The gateway checks whether incoming requests from the high to the low confidential domain are labelled. Only if the label is set correctly, the request is forwarded. This process is mostly applied for sending optimization results. On the contrary, requests from low to high are directly forwarded by the gateway and need no labelling.

As described in Section 3.1, we simulated an operational environment at CWIX 2024. The interoperability challenge was to link four different operational domains as shown in 4. The Field HQ and the DF analysis support cell were planned to be located outside of CWIX training centre, connected via FMN compliant channels. Therefore, all information needed for the scenario update, model adaption and analysis results should be transferred via secure channels.

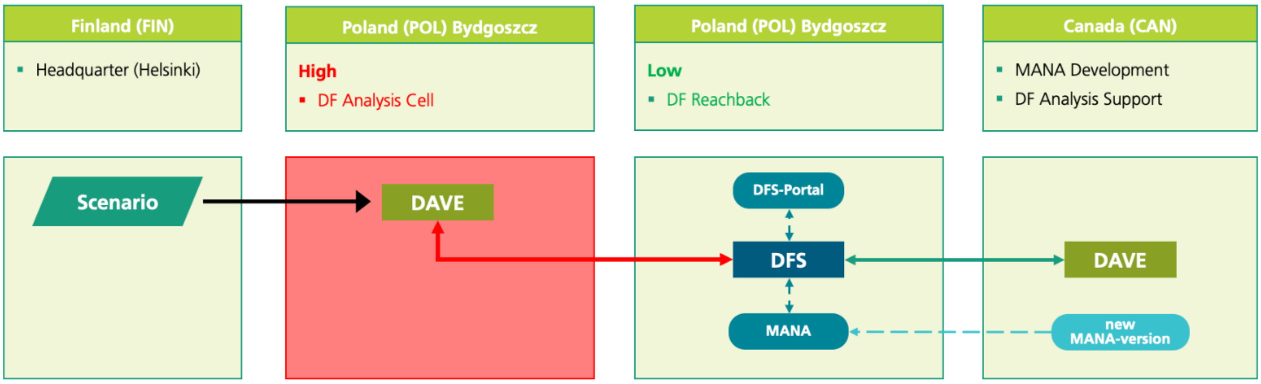

Figure 4: CWIX setup of high and low domain communication to test interoperability.

3.3.2 TEST CASES AND CONDITIONS

Overall, we performed 26 different test-cases during CWIX 2024. They are separated into technical and functional tests over three major phases as depicted in Figure 5.

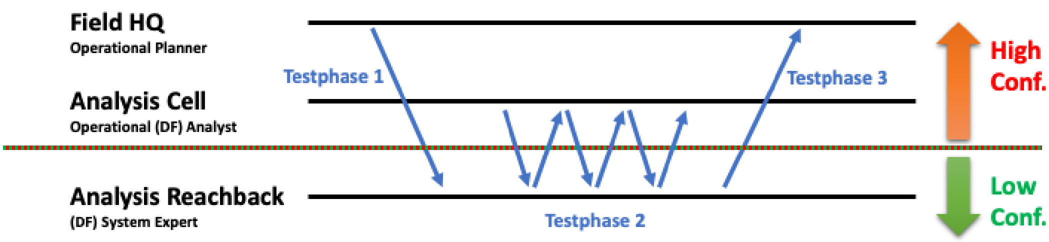

Figure 5: CWIX Test cases diagram.

Test phase 1: The Field HQ identifies a specific event (decisive point) during a long-term campaign, which may influence overall success. As operational planner, it delivers information (operational picture and scenario question) to the downstream analysis cell for obtaining insight into the operational question and its influence at the campaign level. The analysis cell will create all necessary elements to setup, conduct the analysis and deliver the information to the analysis reach-back downstream via the security gateway to set up all technical preparations.

Test phase 2: An initial DoE and MANA simulation results are required to start the optimization based on the MDDF goal. Given the MDDF goal, operational information and an initially generated DoE the optimization process using DAVE is started. The optimization process runs automatically and iteratively generates new experiment results used for retraining the ML model, until the optimization goal or stopping criteria is reached. During the optimization process, information (e.g., current status, intermediate results) is sent to DAVE via the security gateway. Especially during this test phase, we iterated the optimization process with different conditions to obtain more insights into the configuration and its impact on the optimization as described in Section 3.2. The results obtained with different configurations are summarized in Section 3.4.

Test phase 3: The System Expert in the analysis reach-back will make all relevant information from the simulation and optimization process available to the upstream analysis cell via the security gateway. The analysis cell will elaborate this information and create a thorough analysis that is delivered back to the operational planner via FMN compliant channels. Based on this analysis the operational planner can take a decision on the scenario question from Test phase 1.

All test phases followed the protocol of initially executing the tests locally on a single machine with all systems available; executing the tests in a distributed environment without security gateway; and executing the tests in a distributed environment with security gateway, as described in the initial setup. In this way we were able to identify and resolve technical questions early and not influence the analytical questions in later iterations.

3.4 RESULTS

The case study in CWIX was focused on the simulation process, to illustrate how DF can be applied to support operations as additional information becomes available by updating existing plans. The process worked: information was successfully transmitted and data adjustment to the existing scenario occurred sufficiently rapid. We performed the simulation manually before testing the automated process. The automated process worked: the plan could be improved, and the automated process was both faster and provided a better resource allocation solution than the one determined from simulations manually.

Regarding our goals in attending CWIX 2024, we demonstrated the combined use of DAVE and DFS under the regulations of an FMN compliant environment. The technical tests using the SDoT Security Gateway showed consistent successful results. The only limitation was the connection to external partners outside of JFTC. We couldn’t establish a direct technical link to the Canadian partner to send and receive scenario data (bypassed by transferring via air-gap). In the following, we present the analytical results of CWIX 2024being important for the future development and usage of MDDF.

3.4.1 ANALYTICAL RESULTS

We compared traditional DF with the automated approach based on BO. The main difference is that traditional DF tries to cover the input space with as many experiments as possible to have good estimates of the target functions at any given point (e.g., 200 Latin Hypercube (LH) simulations) while the automated approach based on BO starts out with a relatively small number of experiments, which are supplemented during the optimization process. We used different sized LH experimental designs, which were sequentially enriched with further simulation runs recommended by the BO algorithm, and as metrics, we chose the mean number of intercepted migrants (MIM) based on the initial (resp. optimized) experiment plan, see Table 1. We observe that MIM improves with more simulations coming from the BO recommendations, however the simulation run time also increases. The purpose of the last two experiments was to study whether it is more useful to start with a larger number of initial simulations (e.g., 50 LH simulations) and a smaller number of sequential simulations or to start with a smaller number of initial simulations (e.g., 10 randomly distributed simulations) and a larger number of recommended simulations. We find that with the same total number of simulations (in this case, 100), the combined approach based on a standard experiment plan first and then adding recommended simulations provides significantly better results than starting out with a smaller number of initial simulations. In the random design experiment, the initial value of MIM (i.e. 3.50) is larger than the other starting design (i.e. 1.82) because good starting points were simply picked by chance.

Table 1: CWIX simulation and optimization results. Mean intercepted migrants (MIM) is used as an evaluation metrics.

Initial DoE	Optimization runs	Total simulation time [minutes]	MIM (initial)	MIM (optimized)	Best run
Traditional DF LH (200)	-	141	1.78	-	13/32 (40.6 %)
LH (20)	30	37	1.84	6.53	12/24 (50.0 %)
LH (50)	30	61	1.98	3.77	9/11 (81.8 %)
LH (50)	50	85	1.82	6.76	15/28 (53.6 %)
Random (10)	90	84	3.50	5.77	13/33 (39.4 %)

This has several implications: if the total run time is crucial, e.g., because further decisions based on the simulation results must be taken, then the time is best employed when the algorithm can decide which simulations to run next. Regarding the coverage of the input space with simulations, we observe that the state after the optimization is comparable to the traditional DF approach, with the difference being that the factor combinations are not generated as a grid but rather in a way where more interesting areas of the target function are covered with more simulations.

Furthermore, the meta-model can be used for more than recommending further simulations. Once well trained, the meta-model can serve as a shortcut to estimate the objective function without requiring further simulations. This can save even more time which can be employed instead for human decision making. Note that the meta-model provides not only a point estimate but also a measure of uncertainty of the prediction. This allows to judge the goodness of model predictions with respect to true values.

3.4.2 Discussion of Final Results and Limitations

During CWIX 2024 we successfully demonstrated that:

1) DAVE can be deployed as a standalone tool in an FMN compliant environment.
2) DAVE can efficiently automate and optimize the Multi-Run Execution Loop of DF.

The first result is very important, as it confirms that DAVE can be deployed in an FMN compliant environment. All information and required data could be exchanged without delay or loss. Especially the high data exchange during simulation execution raised no issues in the distributed setup over high and low confidentiality networks. The second result is promising. We demonstrated that DAVE can enhance traditional DF with a (static) design even without fine-tuning the optimization method.

The final meta-model, created during the optimization process, can even be applied for answering questions not directly related to the optimization goal. Theoretically, the final model should reproduce the investigated black-box function very precisely around its optima and more approximately elsewhere, but this aspect was not examined in this work. Hence, the meta-model could be used to predict simulation run results. Such predictions are produced much faster than actual simulations. Since the GP provides uncertainty estimations of its predictions, additional simulations are needed, when the uncertainty is large. Alternatively, we can rely entirely on the meta-model predictions saving a lot of computing time for simulations. Note that the fine-tuning of the GP model and acquisition function may lead to even better results, as here we have only taken standard configurations.

4.0 IMPACT ON MDDF AND FUTURE STEPS

With the results achieved in CWIX 2024, we have taken a major step towards MDDF. We are now able to perform optimization-based analysis in an automated way in DF experiments. In the future, this will allow to address analytical questions in even higher dimensions in a timelier manner than before.

MDDF operates at multiple scales and levels. As described in Section 2 we aim to analyze the linkages of the strategic level with the campaign modelled by ACE and the tactical level with the Border Operation simulated in MANA. Currently we are considering only one tactical operation influencing the campaign level. This is one interconnection between factors affecting resource allocation, which should be optimized with respect to the success of the overall campaign. In the future it will be possible to investigate several tactical operations simultaneously, as we have now automated the necessary optimization-based DF experiments. This will enable the investigation of more complex scenarios for the overall campaign and determine the best, most likely and worst solutions with respect to campaign success.

In addition, the meta-models created during DF experiments can be reused for different "what-if" questions during the evaluation of the higher-level campaign without repeating the local DF experiments or adding new ones. Another advantage results from the specific type of meta-model created during the optimization as mentioned in Section 3.2. This specific meta-model may be applied to predict the effect of new factor combinations. Depending on the prediction uncertainty one can then choose to apply the meta-model or to refine it through new simulation runs.

Several aspects of the optimization process can be improved. First, the continuous fine-tuning of the meta-models (especially, for the hyper-parameters and the acquisition function) could enhance model efficiency and accuracy. Then, employing alternative DoE methods for the initial fitting of the GP model (such as DSDs) could yield more robust initial approximations. Moreover, a detailed adjustment of meta-models (in particular of kernel selection, cross-validation techniques, and setting hyper-parameters) should improve the adaptability and precision of our predictive models.

Finally, the MSG-186 Research Task Group aims to extend the test cases simulated at CWIX 2024 to a complete Multi-Layer scenario composed by a strategic campaign embedding the Border Operation. This can be achieved by extending technical and functional parts of the present scenario and it will prove the FMN compliance of MDDF.

ACKNOWLEDGEMENT

This paper is dedicated in memory of MSG-186 co-chair Stephan Seichter and his inestimable contributions to the development of Data Farming.

5.0 REFERENCES

[1] Research Task Group MSG-088 (2014) Data Farming in Support of NATO – Final Report of Task Group MSG-088. STO Technical Report STO-TR-MSG-088. Neuilly-Sur-Seine, France: NATO Science and Technology Organization, 31st March 2014. ISBN: 978-92-837-0205-4. DOI: 10.14339/STO-TR-MSG-088.

[2] Research Task Group MSG-124 (2018) Developing Actionable Data Farming Decision Support for NATO - Final Report of MSG-124. STO Technical Report STO-TR-MSG-124. Neuilly-sur-Seine, France: NATO Science and Technology Organization, 30th July 2018. ISBN: 978-92-837-2151-2. DOI: 10.14339/STO-TR-MSG-124.

[3] Research Task Group MSG-155 (2017) Data Farming Services (DFS) for Analysis and Simulation-based Decision Support, NATO Science and Technology Organization.

[4] Huber D, Horne G, Hodicky J, Kallfass D, De Reus N (2019) Data Farming Services: Micro-Services for Facilitating Data Farming in NATO, in: N. Mustafee, K.-H. G. Bae, S. Lazarova-Molnar, M. Rabe, C. Szabo P, Haas Y-J S (Eds.), Proceedings of the 2019 Winter Simulation Conference. Available on https://www.informs-sim.org/wsc19papers/239.pdf.

[5] Seichter S, Åkesson B, Richter M, Muts N (2020) Data Farming Services (DFS) for Analysis and Simulation-based Decision Support, CA2X2 Forum 2020 Papers Collection. Rome, Italy: The NATO Modelling and Simulation Centre of Excellence, 2021, pp. 120–128. Available on https://www.mscoe.org/document/nato-ca2x2-forum-2020-papers-collection/download.

[6] Seichter S, Horne G (2020) Data Farming Services (DFS) for Analysis and Simulation and Simulation-Based Decision Support Based Decision Support, Symposium MSG-177 on Towards On-Demand Personalized Training and Decision Support. Virtual, 2020: ISBN 978-92-837-2321-9. DOI: 10.14339/STO-MP-MSG-177.

[7] Research Task Group MSG-186 (2022) Modelling a Multi-Faction Conflict in Multi-Domain Operations, Symposium MSG-197 on Emerging and Disruptive Modelling and Simulation Technologies to Transform Future Defence Capabilities. Bath, United Kingdom, 20-21st October 2022. ISBN: 978-92-837-2428-5. DOI: 10.14339/STO-MP-MSG-197.

[8] Research Task Group MSG-186 (2023) Multi-Dimensional Data Farming, Symposium MSG-207 on Simulation: Going Beyond the Limitations of the Real World. Monterey, United States of America, 19-20th October 2023: ISBN 978-92-837-2491-9. DOI: 10.14339/STO-MP-MSG-207.

[9] NATO Standardization Office (2020), AAP-6, NATO Glossary of Terms and Definitions, pp. 64, 133. Available on https://www.coemed.org/files/stanags/05_AAP/AAP-06_2020_EF_(1).pdf.

[10] Margesson R, Mix DE, Welt C (2023) Migrant Crisis on the Belarus-Poland Border. Available on https://crsreports.congress.gov/product/pdf/IF/IF11983.

[11] McIntosh GC, Galligan DP, Anderson MA, Lauren MK (2007) MANA (Map Aware Non-Uniform Automata) Version 4 User Manual, DTA Technical Note 2007/3, Defence Technology Agency, New Zealand.

[12] Garnett R (2023) Bayesian Optimization, Cambridge University Press. Available on https://bayesoptbook.com.

[13] Frazier PI (2018) A Tutorial on Bayesian Optimization, arXiv: 1807.02811.

[14] Rasmussen CE, Williams CKI (2006) Gaussian Processes for Machine Learning, MIT Press. Available on http://gaussianprocess.org/gpml.

[15] Kleijnen JPC, Sanchez SM, Lucas TW, Cioppa TM (2005) A User's Guide to the Brave New World of Designing Simulation Experiments, INFORMS Journal on Computing 17(3), pp. 263–289. DOI: 10.1287/ijoc.1050.0136.

[16] Cioppa TM, Lucas TW (2012) Efficient Nearly Orthogonal and Space-Filling Latin Hypercubes, Technometrics 49, pp. 45-55. DOI: 10.1198/004017006000000453.

[17] Jones B, Nachtsheim CJ (2011) A Class of Three-Level Designs for Definitive Screening in the Presence of Second-Order Effects, Journal of Quality Technology, 43(1), pp. 1–15. DOI: 10.1080/00224065.2011.11917841.

[18] Jones B, Nachtsheim CJ (2012) Definitive Screening Designs with Added Two-Level Categorical Factors, Journal of Quality Technology, 45(2), pp. 121–129. DOI: 10.1080/00224065.2013.11917921.

[19] NATO Communications and Information Agency, SDoT Security Gateway (2022). Available on https://www.ia.nato.int/niapc/Product/SDoT-Security-Gateway_711.

Appendix 1: DETAILED LIST OF AUTHORS

Esa Lappi
National Defence University, Finland
Department of Military Technology
PO Box 7 FI-00861 Helsinki
FINLAND
esa.lappi@mil.fi

Bao Nguyen
Defence R&D Canada and University of Ottawa
101 Colonel By Dr, Ottawa, ON
CANADA
p.bao.u.nguyen@gmail.com

Stephan Seichter
Bundeswehr Office for Defence Planning
D-82024 Taufkirchen
GERMANY
PlgABwIV24ITWissUstgNTSuTBw@bundeswehr.org

Sreerupa Das
Lockheed Martin Corporation
6801 Rockledge Drive, Bethesda MD 20817
USA
sreerupa.das@lmco.com

Bernt Åkesson
Finnish Defence Research Agency
PO Box 10, FI-11311 Riihimäki
FINLAND
bernt.akesson@mil.fi

Maude Amyot-Bourgeois
Defence R&D Canada
60 Moodie Drive, Ottawa, ON
CANADA
maude.amyot-bourgeois@forces.gc.ca

Lynne Serré
Defence R&D Canada
60 Moodie Drive, Ottawa, ON
CANADA
lynne.serre@forces.gc.ca

Gunar Ernis
Fraunhofer IAIS
Schloss Birlinghoven, D-53757 Sankt Augustin
GERMANY
gunar.ernis@iais.fraunhofer.de

Peter Volbach
Fraunhofer IAIS
Schloss Birlinghoven, D-53757 Sankt Augustin
GERMANY
peter.volbach@iais.fraunhofer.de

Alexander Zimmermann
Fraunhofer IAIS
Schloss Birlinghoven, D-53757 Sankt Augustin
GERMANY
alexander.zimmermann@iais.fraunhofer.de

Vadym Slyusar
Central Research Institute of Armaments and Military
Equipment of Armed Forces of Ukraine
Air Force Avenue 28 B, Kyiv, 03049
UKRAINE
swadim@ukr.net

Chris Rolfs
Cervus Defence and Security Ltd.
Porto Science Park, Bybrook Road, Porton
Down, SP4 0BF Salisbury
UNITED KINGDOM
chris@cervus.ai

Alessio Vaghi
armasuisse Science and Technology
CH-3602 Thun
SWITZERLAND
alessio.vaghi@armasuisse.ch

Andrew Gill
Defence Science and Technology Group
PO Box 1500, Edinburgh, SA
AUSTRALIA
andrew.gill@defence.gov.au