POLARIS

TECHNICAL REVIEW
INNOVATION JOURNAL

THE SYNTHETIC ECOSYSTEM

From Digital Twins to Distributed Multi-Domain LVC Simulations

54-2026

LEONARDO

contents

editorial

In the Aerospace, Defence and Security sector, complexity is no longer an exception. It is the normal operating condition. Increasingly interconnected multi-domain systems, compressed development cycles, rapid technological evolution, and the need to contain costs and risks make it progressively less sustainable to rely exclusively on physical experimentation. In this context, Modelling & Simulation (M&S) has gradually moved beyond the role of a mere support tool, to establish itself not only as a Science but also among the primary enablers of innovation and operational capability.

Today, M&S can no longer be considered a stand-alone computational technique. Rather, it is a cross-cutting Science that applies its methodology to the entire system life cycle, from the conceptual phase, where virtual prototyping accelerates design iterations, through validation and certification, up to training and to operational decision support. Whether simulating the behaviour of a next-generation aircraft, analysing the resilience of a space system, or assessing complex electronic warfare scenarios, simulation provides a controlled, repeatable and safe environment in which to explore solutions that would be costly or risky in the real world.

This paradigm is fully aligned with the approach adopted by the Italian Armed Forces, which refers to M&S as a key instrument for capability development. Not only training, therefore, but also conceptual experimentation, planning support, assessment of new capabilities and reduction of decision-making risk. In a Multi-Domain Operations context, the ability to integrate and orchestrate land, air, naval, space and cyber assets within federated simulation environments represents a decisive advantage.

Technological evolution has played a determining role in this transformation. High-Performance Computing has made high-fidelity simulations accessible within timeframes compatible with operational and industrial needs, while immersive technologies - from virtual reality to augmented reality - have made simulation results more understandable and usable, fostering collaboration among multidisciplinary teams. Within this landscape, concepts such as the Digital Twin have emerged, enabling the creation of dynamic virtual replicas of real assets, alongside the integration of Artificial Intelligence and Machine Learning, which are increasingly central to generating complex scenarios, optimising models and transforming large volumes of data into predictive information.

Within this framework, the synergy between Leonardo's and the Italian Army's M&S Test Centres is of particular significance, as it represents a concrete example of effective collaboration between industry and the Armed Forces. The integration of their respective knowledge, skills, infrastructures and methodological approaches enables mutual capability growth, aimed not only at improving individual capabilities but, above all, at addressing complex problems that require a systemic vision. The ability to jointly experiment with models, tools and scenarios fosters continuous alignment between operational needs and technological solutions, reducing capability maturation times and increasing the overall system effectiveness.

A further step forward is represented by the Modelling & Simulation as a Service (MSaaS) paradigm. The ability to access models, tools and data through shared and scalable infrastructures directly addresses the growing need - also felt by the Italian Armed Forces - to operate in an interoperable and collaborative manner, both at the joint level and within NATO and European contexts. MSaaS enables the sharing of simulation capabilities among different entities, reducing duplication and promoting the reuse of validated assets.

Close and continuous collaboration with stakeholders and the Armed Forces is fundamental to this journey. The co-creative approach by M&S Test Centres enables aligning model development with real operational needs. Simulation as an engineering tool is therefore turned into an actual force multiplier in the field. In this direction also lies Leonardo's commitment to making available its permanent M&S infrastructures, capable of supporting the sharing of expertise, as well as the adoption of common best practices and the coherent implementation of the MSaaS paradigm.

Ultimately, Modelling & Simulation can be seen as a common language that enables translating complexity into understanding, uncertainty into awareness, and risk into opportunity. For the Italian Armed Forces and the defence industry, investing in M&S does not simply mean adopting new technologies, but building a collaborative ecosystem capable of supporting the development of the capabilities required to address the operational domain of the future.

The present issue of the Polaris Innovation Journal explores such a context and some of the key technological drivers described. From adoption of MSaaS, to use of Digital Twins and High-Performance Computing capabilities - such as the Leonardo davinci-1 system - through to the evolution of Advanced Training towards the Integrated Live-Virtual-Constructive paradigm, in which real pilots interact with synthetic entities and augmented reality content within a single, coherent and immersive scenario.

The issue also examines the growing role of Artificial Intelligence, through the implementation of Adaptive Training ecosystems based on Machine Learning techniques, capable of personalising training paths by dynamically adapting them to individual performance. The picture is completed by themes related to Cyber Resilience with projects such as ARTIC, and the RIACE system, focused on the management of complex multi-domain scenarios and distributed mission simulations within secure networks. This further confirms how M&S is becoming a central element also in the protection of critical infrastructures and digital capabilities.

In the operational domain of the future, those who can simulate better will not be merely able to anticipate such a change: they will shape it.

Ten. Col. Tito Luise
Italian Army
Technical Director
Innovation Hub - M&S Test Centre
EDA – IT MoD/CGE Captech Simulation

Agatino Mursia
Leonardo Corporate
Head of Digital Engineering/Simulation & Training

EDF FEDERATES Project

Agatino Mursia[1], Davide Bruzzi[2], Andrea Rivetta[2]

[1]Leonardo - Corporate/Strategy and Innovation, [2]Leonardo – Electronics Division

This paper looks at how the EDF FEDERATES (FEDerated Ecosystem of EuRopean Simulation Assets for Training and DEcision Support) Project will realize the first European MSaaS (M&S as a Service) Ecosystem prototype. Leonardo is involved in it since the ideation of the proposal, with Leonardo Corporate and most of its Divisions (Electronics, Aeronautics, Helicopters and Cyber), being the Company a member of the "Core" team composed by the most important European companies in the Aerospace and Defence sectors. The project is coordinated by Rheinmetall (Germany), while Leonardo acts as the National lead for the Italian group of companies participating in the project (MBDA Italy, CETENA, TXT, STAM and ITRES). The project is now entering its last year and has already successfully passed the first two Demonstrations in June and November 2025. The final Demonstration, aiming to show the FEDERATES prototype in its full capabilities and functionalities, is foreseen by the second half of 2026.

INTRODUCTION

he European Commission (EC) has taken proactive steps to ensure that the EU builds upon its strategic autonomy through the presence of strong, full-spectrum and well-trained military forces. Therefore, enhanced EU military training and decision-making capabilities across all Member States (MS) and Norway are crucial to achieve and maintain forces' readiness in the EU. Distributed simulation and training are key enablers that can provide high responsiveness and flexibility to meet current and future training needs. However, Modelling & Simulation (M&S) applications and services often use different technologies and standards. Furthermore, the national Training Centres often operate in isolation, making the pooling of resources difficult. Differences in policies and procedures for networking assets in different security domains or classification levels are significant challenges when creating distributed federated simulation environments. As a result, simulation assets are not as interoperable, configurable or scalable as it is required by a multi-functional, multidomain, multi-national federation of simulation systems.

This makes preparing distributed simulation environments complex, challenging, expensive and time-consuming. According to current approaches, lead times for exercise preparation go lasting many months or even years, which is unacceptable for missions that require rapid training availability. Furthermore, current simulation and training approaches need more support for incorporating all domains, including land, air, maritime, space, cyber, human and cognitive aspects and even medical operations in multi-domain scenarios.

The paper covers:

- FEDERATES Organization and Approach;
- FEDERATES Innovation;
- FEDERATES Validation Campaign through Demonstrations;
- Leonardo Contribution and Lesson Learned;
- Conclusions and Recommendations.

FEDERATES Organization and Approach

The FEDERATES project proposes the study, design, prototyping, testing and validation of a unique and unified EU capability to develop a European Modelling and Simulation as a Service (MSaaS) solution that supports Distributed Synthetic Training (DST) and decision-making, by using assets from MS (Member State) and Norway available in a common Ecosystem [1]. A close cooperation among major defence industries, research organisations and innovative small and medium enterprises from 13 MS and Norway will ensure the availability of M&S applications and services across EU in the FEDERATES Ecosystem Prototype (Figure 1). This will allow MS and Norway to operate national simulation environments and to connect with other MS' assets and services, to compose multi-national simulation events, even with non-European systems using compatible International and NATO standards.

At operational level (Figure 2), the FEDERATES Ecosystem vision is the interconnection of national M&S clouds and national training centres, which allows the composition of EU-wide simulation systems that can support concurrent multi-domain training and decision support activities, also in collaboration with non-EU entities like NATO.
From the functional perspective, the FEDERATES Ecosystem consists of three major elements:

- National Portals to get access;
- Common Foundation Services;
- Content consisting of M&S Services and External Systems.

The FEDERATES Foundation includes common services to discover, compose, orchestrate and execute simulation environments and is based on common standards and processes, including business models.

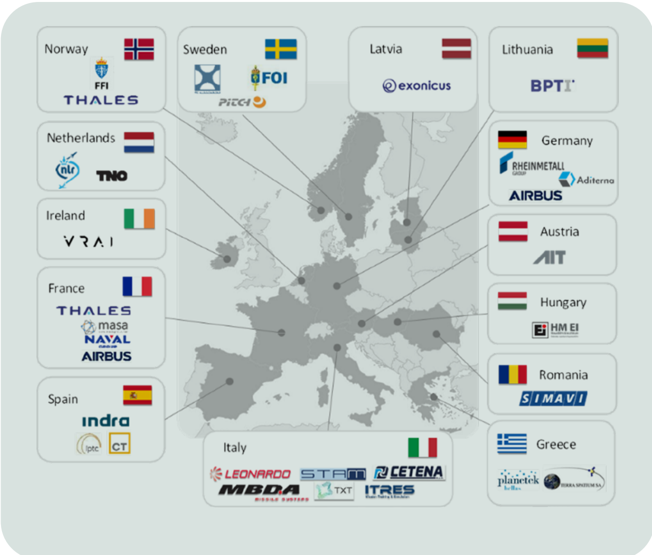

1–FEDERATES Project Members

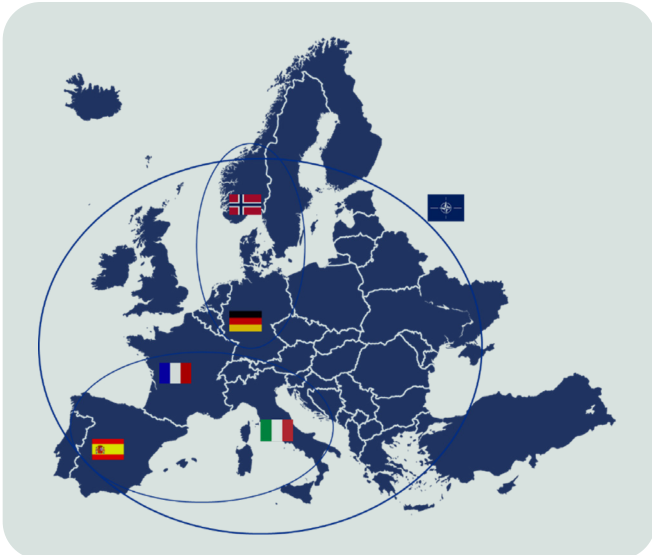

2–FEDERATES Operational View

FEDERATES Innovation

FEDERATES brings innovation to three different aspects such as infrastructure, applications & services, and security:

- In terms of infrastructure, the FEDERATES Consortium will overcome current limitations to setup distributed training, by favouring the integration between legacy systems and new systems, and the interoperability among them. Thus, such a federation at European level will allow MS and Norway to access and use MSaaS from other FEDERATES members using assets from different domains at different levels of classification.

- In the application domain, FEDERATES will examine “cutting edge” technologies e.g. Artificial Intelligence, Extended Reality (Virtual, Augmented and Mixed Reality) and Cloud Computing as the future functional building blocks that enhance the training and simulation capabilities. Thanks to these technologies, new and much more realistic contents can be created in the shared FEDERATES Ecosystem Prototype.
- Finally, the project will apply a “secure-by-design” approach from the requirements definition to the final verification phase, to set up a secure solution, ready for the future certification process. Dedicated analysis on how to realise cyber-resilient MSaaS capabilities will be performed, starting from risks analysis and mitigation, threats identification and impact assessment. Moreover, as this solution will operate simultaneously in different domains at different levels of classification, the information security requirements to control the communication flow among these domains and the integrated external systems will be defined and a dedicated cross-domain solution will be implemented.

The design of the FEDERATES Foundation abstraction layer and Services based on it, is an important key innovation, as it allows its end-user to access the M&S related capabilities (e.g. preparation, execution and control, after-action review) within the MSaaS paradigm. The FEDERATES Foundation is the key element that allows for seamless integration of different National MSaaS implementations. It will allow for easier and more efficient way of setting up, configure and execute multi-national, multi-domain exercises by using native MSaaS Resources and External Systems (Figure 3).

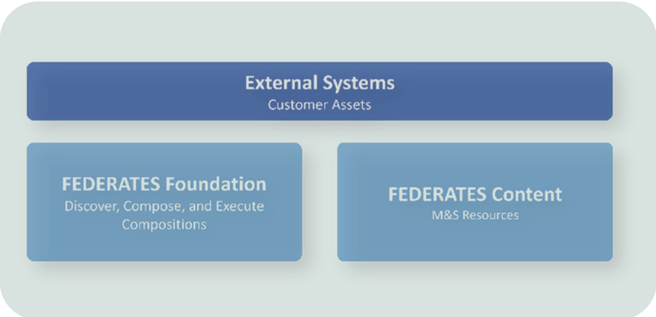

3–FEDERATES Basic Concept View

As shown in Figure 4 and Figure 5, FEDERATES Customer/User according to his MSaaS role, acts in the Ecosystem to perform a specific function.

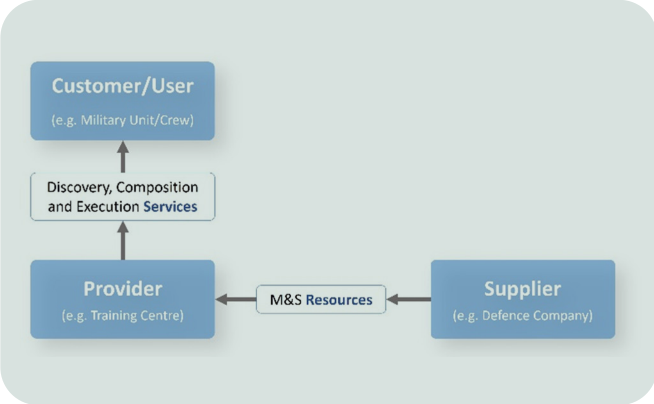

4–FEDERATES MSaaS Roles

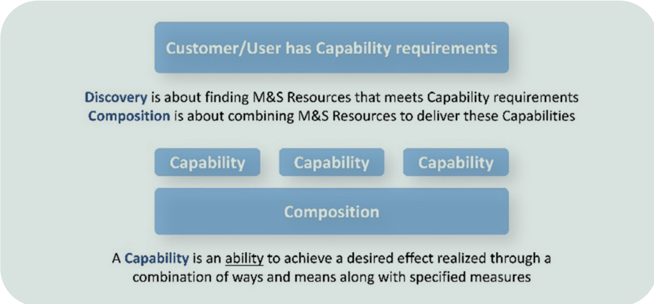

5–FEDERATES Discover and Compose

Leonardo has participated since the very beginning to the NATO long standing activities on MSaaS performed by the NMSG (NATO M&S Group), which are at the heart of the FEDERATES approach and architecture, as shown in Figure 6, and are well described in a set of NATO documents [2].

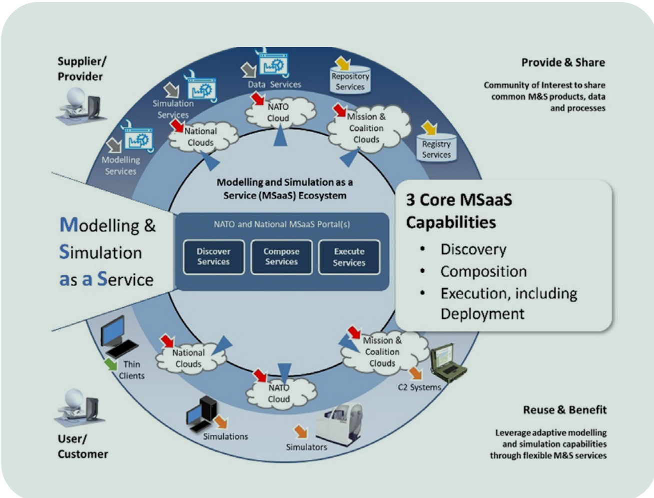

6–NMSG Allied Framework for MSaaS

Here is the description of the different components shown in the picture:

- Clouds all around the three Core MSaaS capabilities (The Foundation in FEDERATES) represent the **virtualisation** (*Red Arrows*);
- **Clouds** host products and services from various organizations (*Grey Arrows*);
- **Suppliers and Providers** provide and share resources based on the use of registry and repository services (*Yellow Arrows*);
- **Customers/Users** reuse and benefit from MSaaS services coupled with their systems (*Orange Arrows*);
- All roles access these services through web browsers (*thin clients as Green Arrows*).

The comprehensive Reference architecture of FEDERATES is then declined in a Solution architecture and then in the Ecosystem Design, as described in Figure 7. Along the life of the project, an agile approach has been used. This has been realized through a series of three iterations, each of one being based on a cycle of Design, Development and Testing, and concluded by a Demonstration event.

The three planned Demonstration events are also aimed at involving the FEDERATES MoD Stakeholders, to obtain feedback useful to realign the project after each iteration. With this approach, the final FEDERATES Ecosystem Prototype is aiming to get to a TRL 6 that will allow for a final Demo where all the requested functionalities will be tested and tuned at the end of the three iterations cycles.
During the design phases, an Enterprise Architect tool has been used to depict all the relationship among the different M&S resources used to build the Solution architecture from the general Reference one.

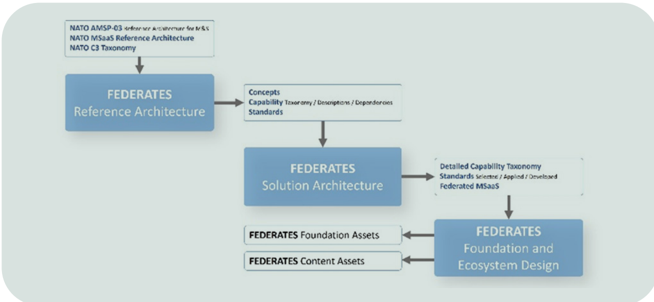

7 – FEDERATES RA and SA Architectures

FEDERATES Validation Campaign through Demonstrations

The last Working Package (WP8 led by Leonardo) of the FEDERATES project has the scope of detailing, finalizing, preparing, executing and documenting the validation of the implemented FEDERATES Ecosystem Prototype, through a Demonstration Campaign. This will be reached through:

- Use of the Verification and Validation Plan throughout all the Validation Phase;
- Use Cases refinement and finalization considering available resources and relevant requirements;
- Demonstration Events preparation, execution and documentation.

The campaign is based on a distributed set-up to involve as many as possible locations, to validate the MSaaS, Cloud-based and Federated approach on which FEDERATES is based.
The WP8 has identified four different Use Cases that exploit the FEDERATES potentialities, trying to use most of the M&S resources made available by the Project's partners. Each of the Use Case is steered by a lead company that coordinates sub-teams with their respective participating partners.

The four Use Cases are:

- LVC Use Case (Live, Virtual, Constructive with Decision Support and After Action Review) led by Rheinmetall Germany;
- Land Use Case (Collaborative Training with allies at Brigade/Battalion Level) led by Airbus Defence France;
- Naval ASW Use Case (Anti-Submarine Warfare) led by Thales France;
- MDO Use Case (Multi-Domain Operations) led by Leonardo Italy.

The way the Use Cases have been identified is functional to demonstrate the flexibility of the FEDERATES Ecosystems to accommodate any potential combinations of the M&S Components and Resources available within the Ecosystem.
The pictorial in Figure 9 shows how this approach works. Every Demo event will show a combination of M&S Components, with a specific Use Case that addresses a particular set of requirements.

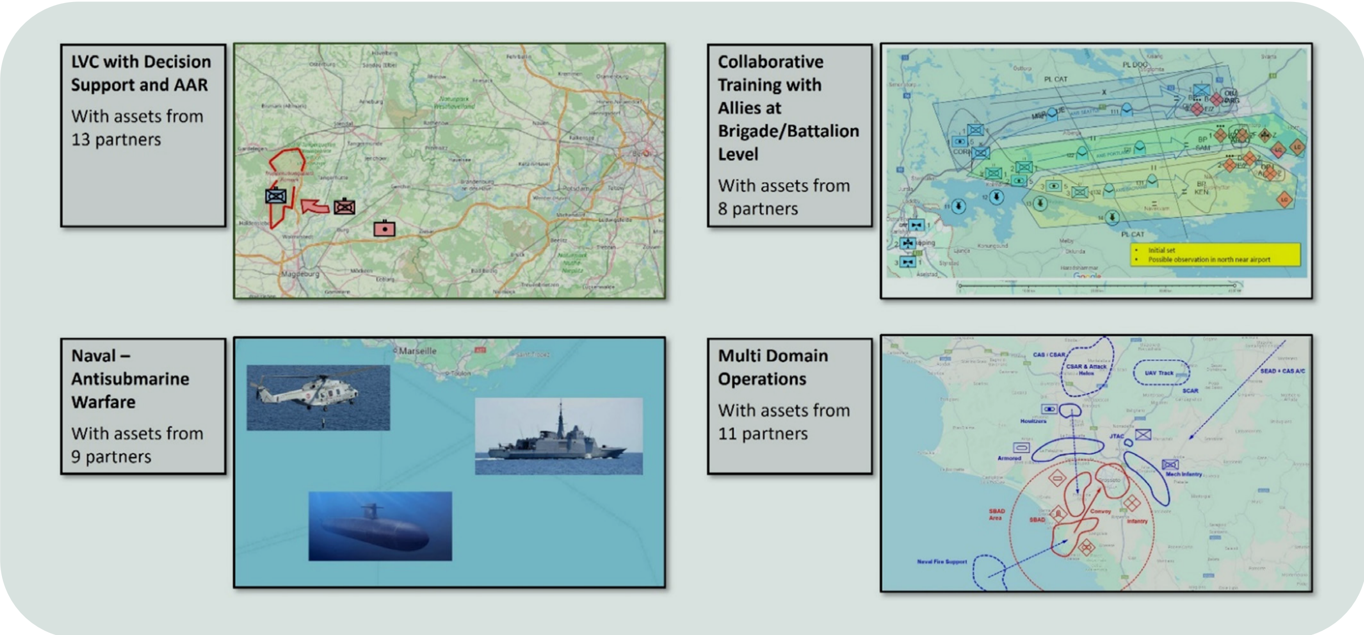

8 – FEDERATES Use Cases

LEONARDO CONTRIBUTION AND LESSONS LEARNED

Leonardo participates to FEDERATES with a large presence of most of its Divisions, including Leonardo Corporate (Digital Engineering – Simulation & Training). The coordination has been provided by Leonardo Corporate and Leonardo Electronics, while the other participants are Aeronautics, Helicopters, Cyber and Security Solutions Divisions.

The participation of Leonardo is very extensive and covers a wide range of topics inside the project. Starting from its role as the Technical Coordinator of the project (assumed during the latest year, in a rotation format with Thales and Rheinmetall), to definition of the CONOPS of FEDERATES, the Solution Architecture, the Foundation design and finally the leading of the Validation Phase through the Demonstration events.

Figure 10 shows an example of the complexity of the HLA (High Level Architecture) that has been used to implement the MDO Use Case led by Leonardo during Demo 1.

Looking at the scheme, it must be highlighted the role of the LED OCEAN product as the MSaaS framework through which all the M&S components (not only the Leonardo's ones but also those by its Partners) are seen in the FEDERATES Ecosystem. RIACE, also from LED, represents the common Synthetic Environment that produces and receives all the entities (including the ones generated by the others) simulated in the overall scenario.

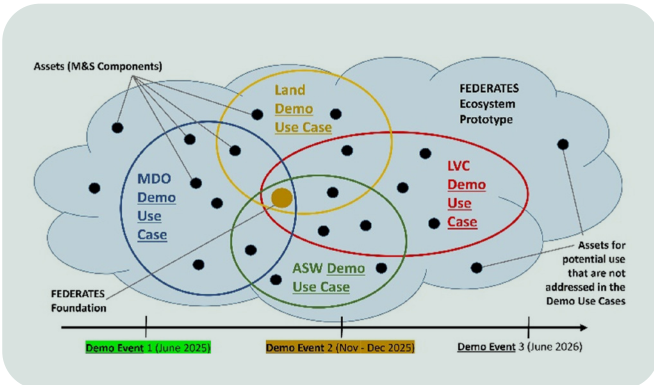

9 – Demo Event vs Demo Use Cases

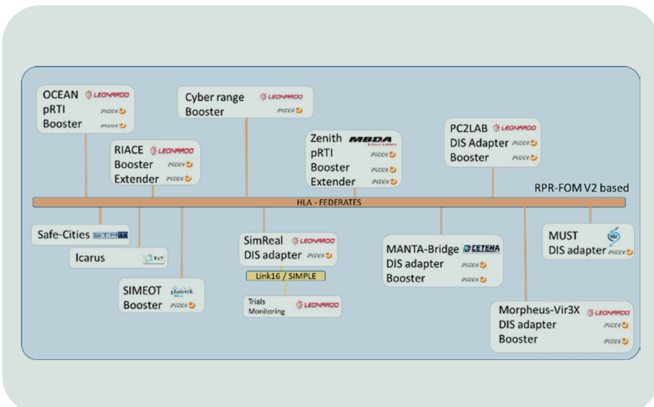

10 – MDO Use Case Demo Setup

PC2LAB from LAD simulates the aerial domain of the scenario, while the Cyber Range from LCD takes care of simulating the cyber effects that could alter the Common Operating Picture. In addition, all the components from the other partners add complexity, by taking care of the naval domain (CETENA MANTA-Bridge), space domain (SIMEOT), drones' behaviour (ICARUS and MUST) and civil crowd behaviour (Safe-Cities).
The MDO Use Case network is based on the Leonardo internal SHORE network connecting all the LDO M&S systems needed for running the demo as shown in Figure 11. The SHORE network is then connected through dedicated VPNs with other partners (NLR in the Netherland, Planetek in Greece and HMEI in Hungary) to complete the overall MDO network.
Moving into the Operational aspect of the MDO Use Case, the overall storyboard has been divided into four different vignettes:

- **TMB** (Tactical Ballistic Missile) that involves Cyber, Land, Maritime and Space domains;
- **OCA/DCA** (Offensive/Defensive Counter Air) that involves Air and Land domains;
- **Blue Water** that involves Air and Maritime domain;
- **Ground** that involves Air, Land and Maritime domains.

Each of them, as described in Figure 12, advances through a specific CoA (Course of Action). The CoAs are exercised in parallel during the Demo and interoperate to complete the MDO Use Case scenario that has a series of specific purposes:

- Demonstrate and validate the effectiveness in repelling a small-sized (battalion) invasion combining **A**ir, **C**yber, **L**and, **M**aritime and **S**pace engagements;
- Coordinate multiple C2s in the execution of several aerial, maritime and land warfare missions;
- Engage surface and sub-surface opponents by coordinating own surface/subsurface and airborne ASuW/ASW assets;
- React of hostile cyber-attacks attempting to compromise the Recognized Air Picture.

The storyboard of the action that initiates the scenario is as follows:

- SARDINIA (RED Party) invades the region surrounding the air base of Grosseto, to stop air policing operations over CORSICA.
- NORTHERN FEDERATION (BLUE Party) Army retrocedes on a line of defense located in the Apennine Mountains and requires the support of allied European Nations.

After the execution of each Demonstration event (Demo 1 and Demo 2 have been hold so far), a survey workshop (after Demo1) and a collection of comments and feedback have been conducted, to get feedback from the Demonstration audience. In addition, also specific post-demo meetings inside the Project's teams have been organized, to go through pros and cons of each event.
All of this had already produced very useful information for fixing and/or improving the way the project is conducting its activities and will be also included in the foreseen Lesson Learned document that is part of the overall FEDERATES deliverables.

11–Leonardo SHORE network for MDO Use Case Demo

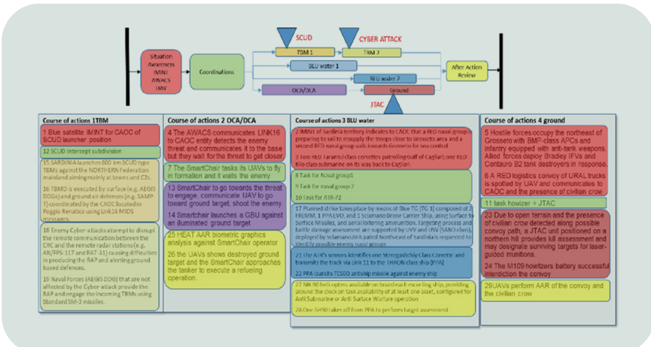

12–MDO Use Case Course of Actions

CONCLUSIONS AND RECOMMENDATIONS

The proposed FEDERATES Ecosystem will provide the technical foundation, initial content and the processes required to allow MS and Norway to share, compose and deliver M&S services and assets by using a federated secure networked environment of different security domains. This will enable the participating MS and partners to access a more comprehensive set of shared and pooled M&S assets (models, data and knowledge).

Furthermore, this will increase interoperability among training centres, simulation systems and forces' training standards, to overcome limitations in versatility, applicability and interoperability for multi-national, multi-domain missions and operations. Finally, from the strategic point of view, this means realising a common approach for more efficient European Forces preparation using state-of-the-art simulation techniques.

From the Italian and Leonardo's point of view it is recommended that:

- Leonardo and all the partners should continue focusing on sharing M&S capabilities and building a MSaaS mindset, also involving the customer communities, to promote knowledge sharing and awareness of this innovative approach.
- Leonardo's capabilities, already compliant with the MSaaS approach, should be properly exploited: OCEAN that has already demonstrated many times its flexibility and maturity on several occasions as the MSaaS orchestration platform, several state-of-the-art tools coming from all the Divisions (***RIACE*, Vir3x, PC2LAB, Cyber Range, Helicopter Sim**) and the ***SHORE*** network operating as the powerful enabler of such approach.
- A follow-up project of FEDERATES should continue to invest for creating a European permanent MSaaS infrastructure. This will allow for creating new opportunities for the EU M&S industry, specifically SMEs, to provide their available services and components to a broader community of M&S users in the EU. With MSaaS and FEDERATES, new business models using subscription and pay-per-use could complement the more traditional perpetual licence models.

As final remark, I'd like to thank all the colleagues from the Divisions for their efforts in the project. The competence and knowledge they have put in the FEDERATES activities so far is a valuable company asset that has to be properly valorised.

Co-funded by the European Union

The FEDERATES project (101121330) is co-funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the granting authority can be held responsible for them.

Agatino Mursia: agatino.mursia@leonardo.com

REFERENCES

[1] EDF FEDERATES Project: FEDERATES-Factsheet_EDF22.pdf

[2] NATO. (2024) NATO webpage for the NMSG on MSaaS technical reference architecture. [Online]. Available: https://www.sto.nato.int/document/modelling-and-simulation-as-a-service-msaas-technical-reference-architecture/

International Cooperation in Uncertain Times: How Modelling & Simulation as a Service helps countries come together

Luke Hamnett[1], Agatino Mursia[2], Gerrard (Ged) Cunliffe[2]
[1]Leonardo – Electronics Division, [2]Leonardo – Corporate

This paper looks at how Leonardo in the UK and Italy is working together across and between national boundaries to bring more certainty to warfighting. Through novel service-based approaches for modelling and simulation, help can be given to warfighters to work through more assumptions and plans more frequently, with better efficiency from the convenience of their own facilities. Leonardo is becoming a better industrial partner by not only developing highly capable technology and tools, but through the strength of its people and processes. This paper explores how Leonardo is improving modelling and simulation across three main areas: engineering, operational analysis, and training. All three are beneficial to the customer community.

INTRODUCTION

The world is becoming increasingly more uncertain. With international politics potentially changing on a weekly basis and technological advances occurring at an exponential rate, the warfighter will likely need to adapt to their changing circumstances more rapidly than has been in the past decades. The conflict in Ukraine has seen new capabilities fielded that have never been fielded before in a major conflict, such as the extensive use of comparatively low-cost uncrewed vehicles. A change that may just alter the way all future conflicts unfold.

With shifts in multiple landscapes, now is the time to further strengthen already robust relationships and prepare for more possibilities that perhaps may not have been considered before. One aspect to leverage is the immense power of digital transformation and its enabling technologies to quickly explore scenarios and assumptions with increasing variation. Modelling and Simulation (M&S) is evolving accordingly and can be a cornerstone of evaluating solutions addressing these rapid changes. Applying a Modelling and Simulation as a Service approach (MSaaS) gives even greater flexibility for both industry and warfighters alike.

Leonardo has a pedigree in M&S with various Simulation Labs across the world and extensive experience into providing state-of-the-art solutions in this field. This article explores how Leonardo in the UK and Italy is leveraging strengths of its people, novel approaches and solutions to bring about positive change for engineering efficiency, operational analysis and training readiness.

The paper will cover:

- Why today?
- Leonardo Digital Continuum and Modelling and Simulation;
- What MSaaS is and its Benefits;
- People and Culture;
- Conclusions and Recommendations.

WHY TODAY?

With the uncertain circumstances the world finds itself in, nations want and need to be ready at potentially a moment's notice as enduring peace no longer appears to be a given.

Nations need their capabilities available, not off on a training exercise or in a field experimenting with what can be done better. Where in either case, those military assets could become damaged, have excessive life added or are just not in the right place in preparation for a dire eventuality.
Of course, it is recognised that warfighters must ensure they are ready to operate in the real-world so will still need to do real-world exercises and should not be replaced. It must be vastly better supported and augmented with other methods. Extensive military exercises have a value to demonstrate to potential adversaries what nations and alliances such as North Atlantic Treaty Organisation (NATO) are capable of [1], it may also signal tactics and approaches adversaries can learn from or counter. Publicly demonstrating such approaches may also be a strategy to give misinformation to adversaries. However, smaller exercises can still cause problems at military ranges that adversaries are watching closely. In an Information Age, hiding what you are capable of may be just as valuable as showing or misdirecting your capabilities. The digital world can provide a safe harbour, working through much more than could be logistically organised in the field at a rapid pace digital options can offer.
The nature of alliances and cooperations are changing. With the Global Combat Air Programme (GCAP), the three nations of the UK, Italy and Japan are separated by almost 10,000km. It is infeasible to expect GCAP to be developed through past methods with the timescales it has. It is embracing more digitally centric approaches, with modelling and simulation likely to be a major aspect to its development. In the NATO context, the Distributed Synthetic Training (DST) initiative has become important as a High Visibility Project (HVP) signed by more than eighteen NATO Countries. DST is aiming to build a flexible and ready-to-use distributed simulation capability, allowing nations more efficient and frequent training. Leonardo is contributing from Italy and the UK, to support the development of such an M&S capability having the right skills and competencies in all the domains in which it will be used.
The technological nature of nation-level warfare is changing. Thousands of uncrewed air vehicles, known as drones, have been used in the ongoing Ukraine Conflict with Ukraine reportedly producing over 1 million drones in 2024 alone [2] with targets for over 4 million. Over 1,000 drones are alleged to have been used in the 2025 Israel-Iran conflict in just 12 days [3]. Drones of all shapes and sizes will now likely be a normalised warfare component. However, many drones used are a form of rotary wing aircraft and almost all drones in these conflicts are propeller driven. Rotary and propeller-based air vehicles are lower cost than jet-based ones but at the sacrifice, in general, of velocity.

Most drones produced and used by Ukraine are known as First Person View (FPV) drones that are guided by operators [4]. In all likelihood, this is a 1:1 relationship of one operator to one FPV drone. Missiles by comparison in many cases are autonomous once launched and require limited to no input from an operator, allowing that person to move onto other actions. How would combining more autonomy with rotary wing aircraft further change the battlespace? Leonardo is developing just such a capability with the 3000kg Proteus autonomous rotary aircraft and already has the 200kg AWHERO autonomous rotary wing craft. Exploring the usage of such capabilities in an MSaaS way and how it can be integrated in synthetic training will be necessary, as capabilities must mature and evolve quickly.

1-Renders of Proteus (Left) and AWHERO (Right)

Throughout the previous sections, three key themes emerge:
- ***Operational Analysis Improvements*** - Rapidly explore how quickly changing capabilities will impact the battlespace;
- ***Training Improvements*** - Rapidly prepare warfighters to use those capabilities effectively without giving insight to adversaries;
- ***Engineering Efficiency Improvements*** - Increase the pace of changes to capabilities or fielding new ones as the operational environment changes. It can no longer take decades to go from concept to being in the hands of the warfighter.

Service-based approaches allow people to flex what exactly it is they want. For example, online streaming is a service of delivering video content from almost any location in the world. A user can increase the number of videos they can watch at once, modify the quality of those videos and can stop using that service whenever they want with no disposal considerations needed.
Modelling and simulation allows people to have representations at their fingertips of the real world, the information world, people and systems or equipment. While a model or simulation is not going to be an exact replica from the real world, they are becoming much closer through validation and assurance that they are good enough to represent the real world.

LEONARDO DIGITAL CONTINUUM AND MODELLING AND SIMULATION

Leonardo has started an internal process to improve the way our company can answer this ever-changing environment we have described in previous sections. The following Figure 2 is the Leonardo Multi-Domain approach, showing, on the left, key features that must be taken into consideration and, on the right, enablers to make this happen in the most efficient way.

As seen from Figure 2, M&S and the Training application are among these key features. The architecture in Figure 3 below outlines the comprehensive capabilities for a single, globally operable Synthetic Environment, operational within Leonardo in support of its Space, Rotary Wing (RW), Fixed Wing (FW), Electronics, and Cyber domains. It is designed to support a wide range of use cases, from early customer engagement to advanced multi-domain operations and battlefield simulation.

Core principles for global operation includes:

- *Scalability & Resilience*: Designed to handle vast geographic scales, high user concurrency, and continuous operation with minimal downtime.
- *Data Consistency & Fidelity*: Ensuring accurate and up-to-date representation of the global environment and domain-specific assets across all operational locations.
- *Interoperability & Standards*: Adherence to international and industry standards for seamless integration and collaboration.
- *Security & Compliance*: Robust security measures and adherence to global regulatory frameworks (e.g., ITAR, export controls, data privacy) for sensitive information.
- *Low Latency & High Bandwidth*: Optimized network infrastructure to support real-time interaction and data exchange across continents.

Key aspects of such a synthetic environment architecture include:

- *Environment Generation & Representation*: These elements focus on creating and defining the synthetic world, with a strong emphasis on global coverage and Leonardo's core domains.
- *Simulation Core*: These capabilities form the fundamental physics engine driving the synthetic environment's behaviour across all domains making use of high-fidelity models of Leonardo assets and operating globally.
- *Rendering & Visualization*: These elements focus on highly realistic and domain-relevant visual representation that performs consistently worldwide.
- *User Interaction & Control*: These elements enable users to interact with and control elements within the synthetic environment from any location.
- *Data Management, Analysis & Replay*: These elements are crucial for evaluating performance, debriefing, and improving systems and training, with global data integrity and accessibility.
- *System Infrastructure & Management*: These are the underlying capabilities that support the entire synthetic environment system within an enterprise context, with explicit focus on worldwide operation, multi-domain use cases, and partner/customer integration.
- *High Performance Computing (HPC) Capability*: Leveraging Leonardo Da-Vinci1 HPC to provide powerful and flexible computing and virtualization capabilities.

This comprehensive architecture provides a robust framework for a global synthetic environment, addressing the multi-domain focus of Leonardo and enabling seamless, secure collaboration with partners and customers across a diverse set of critical use cases. Leonardo already has key capabilities in all six areas as well as working with trusted partners and suppliers around the world. Leonardo in the UK and Italy have a focus on creating a 'Digital Highway' between the two nations to further facilitate a multi-domain architecture as shown in Figure 4.

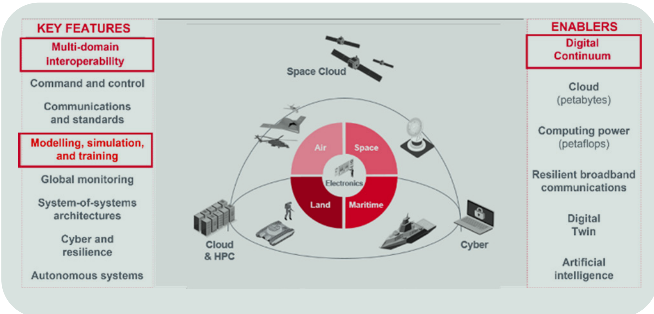

2 - Leonardo's perspective on Multi-Domain key elements

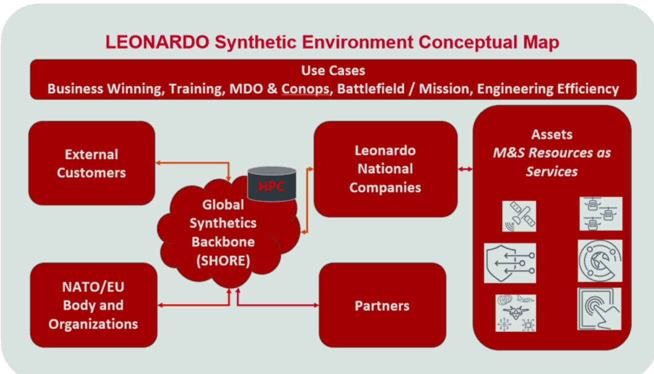

3 - Leonardo's perspective on a global and multi-domain synthetic environment

WHAT IS MODELLING AND SIMULATION AS A SERVICE?

4 – Leonardo's view on a Digital Highway for the UK and Italy

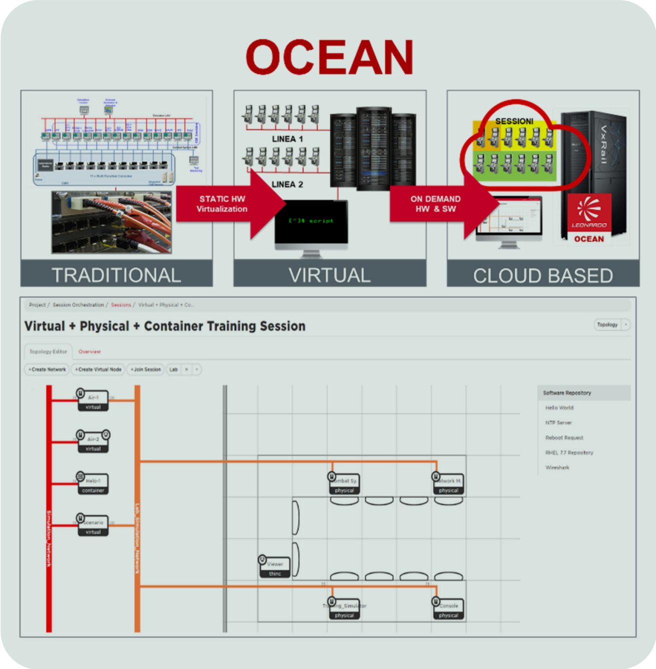

5 – OCEAN MSaaS

MSaaS is an approach that provides simulation capabilities through cloud-based and service-centric architectures that allows greater flexibility and scalability. MSaaS, developed in the NATO Modelling and Simulation Group (NMSG) [5], effectively enables access M&S capabilities on demand, when and where it is beneficial for users. Like online video streaming, people are getting access to powerful tools instead. MSaaS allows activities to be quickly set-up, elements composed in varying combinations and supports interoperability inherently due to cloud-based approaches.
Leonardo has developed a set of capabilities and tools that supports MSaaS specifically. These are:

- ***OCEAN*** – An easy-to-use 'drag & drop' orchestration platform capability for experimentation and training that pulls together hardware, software and facilities from across diverse locations in a scalable way.
- ***SHORE*** – A network of infrastructure to flexibly connect capabilities for Leonardo.

In addition to these capabilities, Leonardo is keenly aware of substantial effort and investment continuously pours into commercial capabilities. Leonardo has prioritised key areas of synthetics engines for military applications, orchestration and infrastructure but continues to leverage excellent capabilities of others. While commercial capabilities may have primarily been developed for the gaming industry, commercial high-fidelity graphics and physics simulations are becoming more regularly used in defence. Customers are going to demand higher quality visuals for training environments to make them as close to reality as possible, as well as being familiar to what they may use in their personal lives.
Leonardo must continue to take benefit from such capabilities but not become over reliant on them longer term. In recent years, commercial tools are starting to be taken off the market [6] which could cause challenges if those capabilities are no longer readily available to the defence industry or its customers. Leonardo has a long history of integrating commercial tools with specific internal capabilities.

WHAT ARE BENEFITS OF A SERVICE-BASED APPROACH TO MODELLING & SIMULATION?

There are several benefits to MSaaS but broadly fall into categories of:

- Financial
- Operational
- Engineering

Looking at financial benefits, by providing M&S via service-based approaches it becomes substantially more cost effective to create something once and use many times.

For example, keeping with the video content analogy, providing video content via discs comes with additional expense even though the content itself is used many times. Costs consist of the physical disc copying, physical storage and logistics. You also cannot keep that content updated to latest formats or improve its quality. Whereas if you store the video on a server and connect to a network with access controls, as many people with connections can access that one copy. You can swap it with an improved version when available and people can still access that new version. The same applies to M&S where users can take the latest and best models or simulations when needed. They do not need to wait for the model to arrive in the mail on a disc, making planning savings. Similarly, the need for keeping expensive hardware available at the user's end becomes less necessary as cloud-based hardware reduces user investment and overheads, making it more financially attractive for customers. By being cloud-based, the supplier of M&S can veer and haul how much of that hardware they need, when they need it, to get the most value for money. If their user base increases, bringing online more hardware will be possible so more people can access MSaaS. When user numbers go down perhaps after a substantial training exercise, then cloud hardware can be released and costs reduced. This means the supplier is not paying for hardware it is not using either. This is particularly beneficial when both M&S suppliers and consumers are based in different nations as network connectivity is all that is needed, with appropriate security and export control.

Moving onto operational benefits, MSaaS gives customers great adaptability. Normally, if a customer wanted to set up a synthetic-centric training exercise, the hardware with the right models and simulations needs to be set-up far in advance along with rehearsals to make sure it achieves its purpose. Leading to training programme being broadly fixed unless forethought considers spare models, simulations and scenarios.

If a customer has decided to work with a service-based approach, then it opens the ability to change training on the fly. Are attendees getting through the planned scenarios faster than expected? Access more challenging scenarios from servers with more challenging assets being simulated. Has another unit become available from another country at short notice? Ramp up the hardware and M&S capability available to conduct inter-unit training. Is there a new threat that's become available since the training plan was put together while deployed to another nation? Reach back to the model supplier and have it updated and uploaded. Having such flexibility before, during and after an exercise while working internationally would be substantially more difficult to orchestrate through other methods.

Lastly but by no means least, MSaaS has engineering benefits. One of the main benefits is engineers across multiple countries can share their developments, forming a collaborative MSaaS. By doing it this way, it adds further efficiencies by leveraging the M&S strengths from each country involved, not duplicating efforts and maximising the quality possible from available effort. By it being a service, the usage of models and simulations can be better tracked to refresh ones that are being less utilised with new features. Simultaneously, monitoring where issues may be occurring with content so it can be fixed in a timelier manner adding further efficiency. If the content were deployed to fixed hardware, engineers would not know or be able to fix common issues until reports are sent or hardware with deployed content is returned for review. To have more immediate fixes today, engineers are directly deployed to customer activities, which adds further expense. A pilot project is on the way in Italy to use OCEAN, running on the Da-Vinci1, and MSaaS approach to allow for an easy and efficient solution to handle Virtual Business Networks among Engineering Labs.

PEOPLE AND CULTURE

People and culture are a key aspect. Without capable people that are enthusiastic about working across existing boundaries, MSaaS will not work. Leonardo has committed to supporting this with a UK-Italy Synthetics Steering Group where key people come together to champion MSaaS ways of working and opportunities that will help warfighters. The UK-Italy Synthetics Steering Group are specifically focussed on generating recommendations that will benefit both nations such as: investments, common architectures and technology roadmap alignments. Similarly, several Communities of Interest and Communities of Practice are in place to help more informal culture growth and knowledge sharing to further build people capability. These groups have people from across different lines of business, divisions and nations to support working across boundaries. Though fostering a more outcome-based approach than an organisational-based approach allows our people to focus much closer on our customers.

This can be seen through how well supported NATO Working Groups are by members of Leonardo, particularly in the M&S domain. By focusing more on desired outcomes, we can have move service-based thinking and co-operation.
Passionate people involved with the various groups are helping others take on a more 'digital-first' mindset. What is a digital-first mindset? This is where people embed the use of digital technologies, data, models and simulations up front as a key element of development. This is opposed to where perhaps in the past, it was considered when needed. By flipping to being more digital-first, it helps us shift to being as proactive as possible with reuse opportunities, rather than being reactionary.
A culture of reuse and knowledge sharing is extended to the use of tools and capabilities. As discussed in previous sections, Leonardo Italy has great capabilities such as OCEAN. These are looking to be shared with Leonardo UK through technology transfer to give both nations a common set of capabilities to work together more fluidly, as well as having a greater pool of people able to leverage the power of those tools. Why is this important? Common approaches and understanding how to work together with the same tools gives a shared development language for both people and projects. Those involved know what they are getting and so makes it more efficient for our engineers to give customers services that they would want.

CONCLUSIONS AND RECOMMENDATIONS

As the world becomes more uncertain, countries cooperating and collaborating is necessary. MSaaS is one way to help with bringing them closer together. It can be achieved through its inherent aspects such as a boundary crossing cloud-based approach and flexibility to ramp up or down as customer needs require, for experimentation of operational analysis or training exercises. It is recommended that:

- The UK and Italian customers create small scall experiments to pilot MSaaS approaches further, sharing outcomes with NATO.
- Leonardo must continue focusing on sharing it capabilities and continue with building a critical number of people across the UK and Italy with a 'digital-first' mindset. Ideally involving the customer communities in knowledge sharing and building activities.
- Leonardo has several capabilities that assist the MSaaS approach. These are the ***OCEAN*** orchestration platform for training and experimentation, the ***SHORE*** network for flexible connectivity, the ***RiACE*** simulation engine designed for defence and the ***SynAPSIS*** simulation integration platform for digital twin creation. These should be experimented with for inclusion into MSaaS approaches further.

To summarise conclusions:

- Leonardo is utilising MSaaS approaches as a way of building stronger, collaborative relationships for and across the UK and Italy.
- By utilising more modelling and simulation, customer communities across boundaries can work through uncertainty much quicker and more flexibly with a service-based approach.
- It takes the enthusiasm and drive of people to push for a more 'digital-first' mindset that enables MSaaS approaches. Tools and process alone are not enough.
- MSaaS approaches are critical to improvements in engineering efficiency, operational analysis outcomes and training readiness to navigate a world heading into an increasingly uncertain future.

Luke Hamnett: uksystemsmailbox@leonardo.com

REFERENCES

[1] Reuters. (2024) Reuters article on NATO exercise.
[Online]. Available: https://www.reuters.com/world/europe/nato-kick-off-biggest-drills-decades-with-some-90000-troops-2024-01-18/
[2] D. Axe. (2025) Forbes article on Ukraine drone production.
[Online]. Available: https://www.forbes.com/sites/davidaxe/2025/03/12/45-million-drones-is-a-lot-of-drones-its-ukraines-new-production-target-for-2025/
[3] J Epstein. (2025) Business Insider article on drone usage in Iran/Israel conflict.
[Online]. Available: https://www.businessinsider.com/huge-iranian-attacks-helped-israel-upgrade-top-ballistic-missile-shield-2025-6
[4] Reuters. (2024) Reuters webpage on drone usage in Ukraine.
[Online]. Available: https://www.reuters.com/graphics/UKRAINE-CRISIS/DRONES/dwpkeyjwkpm/
[5] NATO. (2024) NATO webpage for the NMSG on MSaaS technical reference architecture.
[Online]. Available: https://www.sto.nato.int/document/modelling-and-simulation-as-a-service-msaas-technical-reference-architecture/
[6] CAE. (2021) CAE webpage on Presagis products discontinuation.
[Online]. Available: https://www.cae.com/defense-security/what-we-do/training-systems/presagis-modeling-and-simulation-tools/

The Future of Military Flight Training: Moving from Legacy Live Exercises to Integrated LVC

Francesco Ollosu, Giuseppe Pietroniro

Leonardo – Aeronautics Division

Live exercises are the most effective way for building fundamental airmanship, but they require several assets that make the military flight training expensive. Thus, the issue is how to make pilots ready for more complex multi-domain battlespace, enhancing effectiveness of their training while reducing costs. The Live-Virtual-Constructive concept links real aircraft in flight, flight simulators on the ground, and additional virtual entities, to obtain a more challenging and complete air picture, with no need for airborne participant. The Leonardo solution bases on a training datalink that shares data between real assets and simulators, while live pilots go interacting with other entities playing a role in the scenario and operate simulated sensors and weapons. The Live-Virtual-Constructive (LVC) has now evolved into the Integrated Live-Virtual-Constructive (I-LVC) that merges live, virtual, and constructive nodes, into a single battlespace. From the operator perspective, there is no distinction between real and synthetic elements. This is obtained thanks to three main technological pillars: augmented reality (AR), modular embedded simulation for onboard and ground systems, and smart computer-generated forces operating onboard and on ground within a flexible synthetic environment. Pilots can practice engagements against threats controlled by the embedded simulation by means of datalink connection. Thanks to AR, the all-weather training takes place on virtual arenas that make it independent from military ranges.

INTRODUCTION

Traditionally, military flight training has been relying on live exercises as its primary vehicle for teaching aircrew all the principles of the aviation, from the basics to the advanced tactical manoeuvres. Exercises immerse pilots in real-world flight dynamics, weapon employment and formation flying. They still remain the most effective way for building up fundamental airmanship. Live training, however, requires multiple assets: aircraft, large reserved airspace, ground and/or naval range support, and large numbers of safety and instructional personnel. Therefore, it is an expensive task in defence. Planning a single high-end exercise occupies several assets, people and equipment, for a large amount of time. Due to technological changes in technologies, as seen in the recent international combat theatres, the question is not anymore whether live training is indispensable, but how to amplify its benefits, reduce its costs and prepare pilots for an increasingly complex multi-domain battlespace.

EMERGING THREATS DEMAND NEW TRAINING MODELS

As shown in the recent international events, the battlespace has evolved beyond purely kinetic exchanges. Adversaries deploy advanced surface-to-air missiles and weaponize drones. Moreover, they pursue attacks in the electromagnetic and cyber domains too. Cyber-attacks can hijack datalinks, manipulate GPS inputs or degrade airborne networks, while electronic jamming can deactivate radars or disturb voice communications [1]. In this newest warfare environment, pilots must be able to operate with unstable situational awareness, to aggregate multiple sensor inputs and to coordinate with other blue forces, even in case of deteriorated network. Legacy live exercises, while continuing to be invaluable for mastering flight feelings and manoeuvres, can face big difficulties in managing such a multi-domain complexity by operating traditional static/analog range equipment.

LIVE-VIRTUAL-CONSTRUCTIVE: THE LEGACY PARADIGM

Beginning in mid-2000, the Live-Virtual-Constructive (LVC) became a concept well established in the military community. Leonardo operated to link together real aircraft (Live), networked simulators (Virtual) and computer-generated forces (Constructive). By linking real aircraft in flight with high-fidelity flight simulators on the ground and virtual entities controlled by synthetic models, exercises carried out with the LVC paradigm enable a larger, more challenging air picture, without requiring every participant to be airborne. The Leonardo solution, by means of a training datalink, shares data between real assets and simulators, while live pilots can interact with the entities taking part in the scenario, by leveraging on simulated sensors and weapons. The LVC model brings clear benefits in terms of cost savings and scenario complexity, but it has also revealed gaps: the lack of visualization of constructive entities in Within Visual Range (WVR) reduces fidelity of the exercise, the detail of interactions especially around electronic warfare, weapon effects, updating rate of constructive entities. Moreover, the transfer speed of the Datalink was sufficient for mid-2000 phase IV training but it can't be totally useful for training scenarios of today and of the next generation.

INTEGRATED LVC: THE EMERGING TECHNOLOGIES

The Integrated Live-Virtual-Constructive (I-LVC) paradigm represents the next logical evolution of the simulation-enabled training. Instead of networking discrete live, virtual, and constructive nodes, the I-LVC is engineered to merge them into a single, evolved battlespace in which distinctions between real and synthetic elements disappear from the operator perspective. Three key technological pillars are at the base of the I-LVC: augmented reality (AR), modular embedded simulation for onboard and ground systems, and smart computer-generated forces operating seamlessly both onboard and on ground, within a flexible synthetic environment.

Augmented reality brings constructive elements into the pilot's out-the-window view. High-resolution 3D models of constructive entities (aircraft, ground vehicles and buildings, ships, etc.), surface-to-air missile countermeasures and even synthetic clouds, appear into the real world. Of course, this is carried out taking always in consideration the safety of flight. Pilots can practice engagements against threats that are controlled by the embedded simulation onboard the aircraft and/or ground synthetic environment, by means of datalink connection. Thanks to AR, the training can take place on virtual arenas, thus reducing dependence on military ranges, even simulating weather characteristics suitable to the specific training.

Embedded simulation, as Leonardo Embedded Training System (L-ETS), extends synthetics directly into the aircraft's avionics, by simulating sensors and weapons that are actually not installed on the real platform. This approach enables pilots to practice multi-sensor fusion, EW detection and avoidance and safe weapons employment procedures within a complex virtual training arena. Scenario fidelity rises because the pilot perceives the synthetic sensor feedback as the real ones, maintaining the real pilot's workload and task flows.

Smart computer-generated forces (CGFs) employ artificial intelligence and machine-learning algorithms to move beyond scripted patterns. These CGFs monitor the evolving tactical scenario, adapt their tactics and coordinate with other AI agents that represent entities of different domains. The CGFs learn from repeated engagements, increasing their unpredictability and ensuring a wide set of responses against trainee's manoeuvres. This dynamic behaviour mirrors real-world opponents and requests the pilots to leverage on their cognitive ability, thus it can be further tracked and analysed on ground by means of dedicated AI-based algorithms.

In the "background of the previous elements" there is a flexible synthetic environment hosted on a scalable, infrastructure, seamlessly on ground or on-board. This environment models terrain, weather, electronic emissions, and network traffic across domains and manages the orchestration of constructives. Additional scenario "players" can connect from dispersal locations, while sharing a single, time-synchronized battlespace, leveraging on high performances Datalink, eventually empowered with radio bridges to extend the ranges.

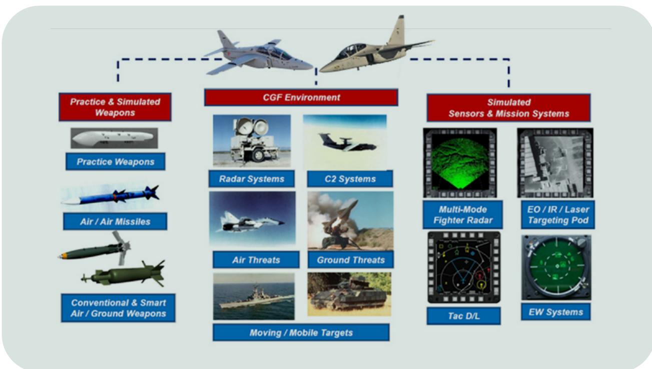

1-Leonardo Embedded Training System (L-ETS)

REDUCING COSTS, INCREASING EFFECTIVENESS

By shifting much of the high-intensity threat replication and multi-domain complexity into synthetic domains, the I-LVC reduces the cost of current and new generation training, while keeping high the level of training effectiveness. In parallel, leveraging on synthetic capabilities, the cancellation of sorties due to availability of the physical platform or to non-optimal weather conditions, is minimized.
Pilots can fly more sorties per training day, encountering a broader set of threats and receive immediate, objective feedback through sophisticated Mission Debrief tools that merge live flight data with synthetic event logs, further analysable with training instructors.

EXPANDING TO MULTI-DOMAIN TRAINING

One of the I-LVC's most profound impacts is its multi-domain capability breaking legacy training silos. High performances datalink and synthetic environment can host protocol simulation [2] in other to seamlessly integrate different real platform within the virtual arena or can even simulate those entities with high fidelity, smart CGFs.

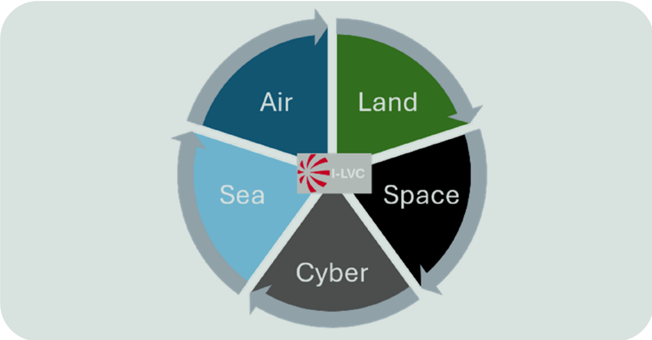

2-I-LVC, a connection between different domains

CHALLENGES TO ADOPTION

The promise of the I-LVC brings technical and organizational challenges. Latency and synchronization between live cockpits and synthetic feeds must be tightly controlled, to preserve realism and safety. Security shall be increased to avoid external menaces both at software and at radio/frequency transmissions.
Software certification shall be further analysed in order to manage AI features within traditional software development. Additionally, the I-LVC should be inserted within Air Forces syllabus, thus requiring a set of changes in legacy instructions in order to map competencies to I-LVC capabilities.
Institutionally, Air Forces need to invest in simulation network infrastructure, develop open architectural standards and certification methodology for the new technologies and introducing validation standards for smart CGF entities.

CONCLUSIONS

In a period of strong changes in the international theatre and in threats employment, the Integrated Live-Virtual-Constructive paradigm offers a sustainable path to preserve and increase benefits of the live flight training, by leveraging on new technologies: augmented reality, embedded simulation, smart computer-generated forces, innovative synthetic environments and powerful training Datalinks. The I-LVC reduces costs, increases training fidelity, prepares trainees and pilots to multi-domain challenges, while decreasing the carbon footprint [2] of the legacy flight training.

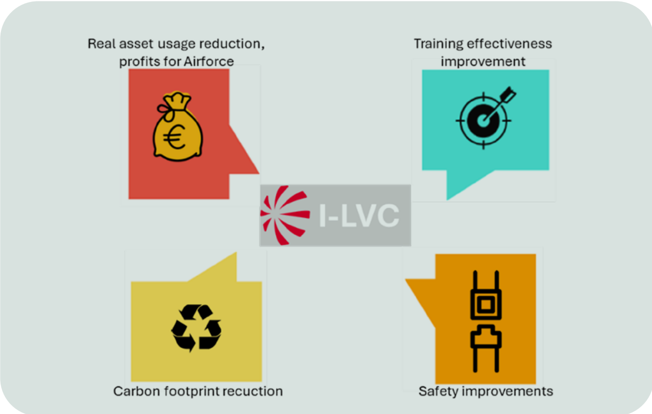

3-I-LVC benefits

Francesco Ollosu: francesco.ollosu@leonardo.com
Giuseppe Pietroniro: giuseppe.pietroniro@leonardo.com

REFERENCES

[1] NATO – AMSP-03 - Multi-Domain Operations and Digital Transformation: Enabling Converged Effects in the Modern Battlespace
[2] Council of the European Union – ART Analysis and Research Team-Greening the Armies

Effective LVC Multi-Domain Mission Simulation with RIACE

Luca Cistriani

Leonardo - Electronics Division

The RIACE Distributed Mission Training (DMT) System is employed by the Italian Air Force to simulate complex multi-domain missions in NATO large scale distributed simulation exercises. The success of the training activities is mainly linked to a powerful simulation system deployed by using a modular, robust and resilient HW & SW architecture that allows secure distributed training on multi-national secure networks up to NATO SECRET. RIACE supports several operational levels, from tactical (platform level) to higher Command and Control (C2) level, either as a combination of HICON/LOCON levels represented as role-players or effectively feeding real C2 systems. RIACE also allows effective Live-Virtual-Constructive (LVC) simulations linking real assets to complex simulated scenarios. The system can be used effectively beyond training, to support CONOPS development or REHEARSAL of critical missions. The paper presents a panoramic of the more recent developments of the RIACE System and illustrates a number of hypothetical multi-domain LVC scenarios supported by RIACE.

INTRODUCTION

The RIACE Synthetic Environment, as part of the RIACE Distributed Mission Training (DMT) system has been delivered by Leonardo Electronics Division to the Italian Air Force to provide an effective distributed mission simulation solution supporting a heterogeneous training audience. RIACE supports EF2000 and Tornado crew training in the respective Full Mission Simulators (FMS), Control Reporting Centre (CRC) operators training and overall coalition mission accomplishment training at NATO level ([1]). So far, RIACE has "participated" in the following major distributed simulation and Live-Virtual-Constructive (LVC) exercises:

- July 2018: Spartan Alliance 18-8 ([1], [2])
- February 2019: Spartan Warrior 19-2 ([2])
- September 2020: Spartan Warrior 20-9 ([11])
- September 2024: Spartan Warrior 24-2 ([6])
- September 2025: Centurion Warrior 25 ([7])

Spartan Warrior and Spartan Alliance are Distributed Simulation Exercises organized by the USAFE/USAFRICOM Warfare Center (UAWC) in Ramstein, Germany.

RIACE has been used extensively not only during the exercises themselves, but also during the planning and preparation phases. Simulated missions have included Combined Aerial Operations (COMAOs), Close Air Support (CAS), Offensive and Defensive Counter Air (OCA/DCA), Air-to-Air Refuelling (AAR) and Ground Based Air Defence (GBAD). RIACE has also provided a test-bed for technical man-in-the loop trials and mission rehearsals performed in the Tornado simulators in Ghedi and in the EF2000 simulators in Gioia del Colle and Grosseto, prior to the execution of the exercises themselves. RIACE supports training at several operational levels, from tactical (platform level like in EF2000 and Tornado FMS) to Command and Control (e.g. Control Reporting Centre level), either as a combination of HICON/LOCON levels represented as role-players or feeding real C2 systems in an effective LVC simulation ([4]). As a practical example, an actual mobile Control Reporting Centre (CRC) re-deployed in Gioia del Colle has been used during Spartan Alliance 18-8, while real CRCs located in Licola and Bari have been used during Spartan Warrior 24-2.

All these systems have taken part in the simulated scenario, thanks to a suite of simulated sensors (e.g. radars), simulated radio communications (Voice Over DIS - VODIS) and a JREAP-C interface to the simulated Link16 environment. The system can be used for effective training at many different levels ranging from Advanced Flight Training (Phase 3 of Common NATO Flight Training – [5]) to Operational Readiness Training (Phase 6). Beyond training, RIACE can effectively support CONOPS development or REHEARSAL of critical missions.

THE RIACE DISTRIBUTED MISSION TRAINING (DMT) SYSTEM

The RIACE Distributed Mission Training (DMT) system, with its modular configurations adaptable to different needs, is a turnkey solution that allows the user to start training immediately at tactical and C2 level through the distributed simulation of complex multi-domain military operations.

The RIACE DMT System is composed by:

- One or more racks 42U 19” hosting the computing system and the network devices engineered to support different system configurations including large Control rooms. The systems generally hosted inside a DMT rack (depending on configuration) are:
    - Domain Server: controls user accesses implementing the IT Security Infrastructure;
    - Firewall: protects the internal network;
    - One or more Synthetic Environment Servers (as many as the simultaneous simulation sessions foreseen) provide the actual RIACE Synthetic Environment simulation;
    - Data Storages (as many as the levels of classification supported): contain the RIACE Synthetic Environment libraries and databases, store exercises and doctrines, etc.
    - One or more HMI Servers (as many as the monitors of the Operator Consoles);
    - Additional components required for networked exercises and linking of external components as:
        1. Gateway/Bridge (including the DIS Audio Gateway and other specific interfaces);
        2. Network Control Station;
- One or more (as many as needed) Operator Consoles including:
    - One or two (arranged vertically) 55” 4K monitors;
    - Keyboard & Mouse;
    - Headset (for the simulation of radio communications);
    - Optional Hand On Throttle And Stick (HOTAS) manual control of constructive entities.

The system is designed to be deployed in classified areas (up to NATO SECRET), being its computing system (e.g. the modular 42U 19” racks) hosted in a conditioned Computing Room separated from the Operators Room where the HMI consoles are. The Operators Room can be the Instructor Operators (IOS) Room of a Full Mission Simulator (like the EF2000 and Tornado FMS) or a Control Room for distributed mission simulation (like the Control Room available at the Italian Air Force Air and Space Operations Preparation Centre (ASOPC) in Poggio Renatico shown in Figure 20 (from the exercise Spartan Warrior 24-2).

1-RIACE DMT Consoles in the ASOPC at Poggio Renatico during exercise Spartan Warrior 24-2 (photo from [6])

THE RIACE TYPHOON MISSION TRAINER (TMT) SYSTEM

The RIACE Mission Trainer is an operational desktop trainer solution customised to simulate specific RIACE High Fidelity assets with higher level of interaction with the user.
Currently, the RIACE Typhoon Mission Trainer (TMT) is the main application already delivered to the Italian Air Force. It is a portable Part-Task Trainer of the Typhoon Fighter Aircraft (EF2000 P1Eb), specifically designed for mission-oriented training.
The simulator is designed to complement the higher level EF2000 Full Mission Simulators, providing also the capability to simulate High Fidelity manned opponents (e.g. MiG-29, SU-27) for the main simulators force.

The TMT simulation includes the following A/C subsystems:

- Aircraft flight dynamics and systems model;
- Weapons and Store Management System model;
- Radar Simulation model;
- IRST Simulation model;
- DASS Simulation model;
- IFF Simulation model;
- MIDS Simulation model.

The RIACE TMT Hardware includes:

- A 42U 19” rack hosting the computing system;
- A wheeled, adjustable cockpit structure with chair satisfying the requirements of compactness and transportability, accommodating a 95 Percentile Male/Female in accordance with MIL-STD-1472;
- Out-of-the-Window (OTW) display system providing an immersive 3D rendered view of the outside simulated world including the aircraft Head Up Display (HUD) Symbology;
- Cockpit display system including the aircraft front panel and a selection of operational aircraft controls;
- Specific EF2000 Hands-On-Throttle and Stick (HOTAS) controls.

The HMI of the RIACE Mission Trainer System can be customised to represent several types of platforms (e.g. aircraft, helicopters, tanks, warships, etc.) including their different operator stations.
The current architecture of the RIACE TMT is specifically designed to allow Quick Reaction Alert (QRA), OCA/DCA - specifically Beyond Visual Range (BVR) combat - CAS, Strike and AAR training in a compact environment, additional devices like Head Mounted Devices (HMD) or different types of consoles can be provided to support different training tasks (e.g. Within Visual Range - WVR –combat training).

2 – RIACE TMT flight simulator during exercise Spartan Warrior 24-2 (photo from [6])

THE SPARTAN WARRIOR 24-2 EXERCISE

The exercise Spartan Warrior 24-2, executed in the framework of Mission Training through Distributed Simulation (MTDS), has seen a broad participation of NATO countries including Canada, Denmark, Estonia, France, Greece, Italy, Latvia, Netherlands, UK and USA ([6]). With this exercise, the UAWC has promoted synthetic training in air warfare operations through many federated simulation systems, real weapons systems (e.g. Ballistic Missile Defence Systems) and Tactical Command and Control (TAC C2), being either present at the UAWC itself or federated through the Combined Federated Battle Laboratories Network (CFBL Net).

3 – The Mobile Control Reporting Centre in Bari during exercise Spartan Warrior 24-2 (photo from [6])

The Italian Air Force has participated with the following federated LVC assets powered by RIACE DMT and TMT systems:

- Modelling & Simulation Control Room in the ASOPC in Poggio Renatico (Ferrara);
- 2 TMT simulators at the 4th Wing in Grosseto;
- Mobile Control Reporting Centre in Bari;
- Control Reporting Centre at the 22nd Group DAMI (“*Difesa Aerea e Missilistica Integrata*”) in Licola (Naples).

The simulated assets included virtual EF2000 simulated by the TMT simulators and piloted by Combat Ready pilots of the 4th Wing, and constructive assets simulated by RIACE in the ASOPC Control Room:

- Several EF2000 flights for DCA operations controlled by artificial intelligence or by the Blue Force Cell of the ASOPC;
- A tanker KC767 for AAR operations controlled by artificial intelligence;
- RAT-31DL Mobile radars to provide the Recognized Air Picture (RAP) for the CRCs;
- SAMP-T Batteries, controlled by Italian Air Force DAMI specialists of the 2nd wing to counter Air Breathing Targets (e.g. hostile strike, anti-radar and bomber aircraft), Cruise and Ballistic missiles.

During the exercise, the Integrated Air Defence System (IADS) composed by the CRCs, SAMP-T batteries and EF2000 CAP patrols has defended the NATO territory from the hostile intrusions, coordinating with the nearby components (NATO AWACS and US fighters) under control of the Air and Space Operations Centre (ASOC) located at the UAWC.

THE CENTURION WARRIOR 25 EXERCISE

The exercise Centurion Warrior 25 ([7]) has been a proof of concept for the fully integrated Live, Virtual and Constructive (LVC) environment that will be represented by the Operational Training Infrastructure (OTI) of the Italian Air Force.

Live fighters of the $4+^{th}$ gen. (EF2000) and the 5^{th} gen. (F35) have operated in conjunction with their Virtual counterparts (TMT and FENIX simulators) in complex Combined Aerial Operations (COMAO) scenarios in the airspace of the “Poligono Interforze del Salto di Quirra” (PISQ) in Sardinia, protecting a flight of Constructive Tornado IDS controlled by the ASOPC Control Room in Poggio Renatico. All airborne assets (either Live, Virtual or Constructive) have operated under control of the Virtual Control Reporting Centres (V-CRC) in Poggio Renatico and Licola using Link-16 and VODIS Radios.

The LVC environment was based on two main enablers:

- The simulated Link-16 network in the HLA/DIS environment was connected to a “real” (radio-frequency) Link-16 network through a prototype of Multi-Datalink Processor (M-DLP) specifically modified and produced by the Leonardo Electronics site of Taranto. This allowed the Live aircraft to receive a simulated RAP from ground and the Virtual and Constructive assets to receive kinematic, sensors and weapons data in the HLA/DIS environment from the Live aircraft thanks to an LVC gateway.
- The VODIS radios were connected to the “real” radios using a Radio-Bridge realized by the Italian Company SITTI in cooperation with Leonardo. This allowed Live aircraft to talk with the Virtual Pilots and C2 operators.

The Blue Forces have faced a Constructive hostile force, controlled by the ASOPC Control Room in Poggio Renatico, represented by a combination of the $4+^{th}$ gen. fighters (simulated by RIACE Constructive entities) and two-digits Surface to Air Missile (SAM) systems. The Constructive SAM systems were actually the “Digital Twin” of the real systems located in the PISQ Firing Range, constituted by emitters reproducing actual electro-magnetic threats combined with a full-scale mock-up of the actual SAM system able to provide effective training for the 5^{th} gen. fighters.

4 – Use of the RIACE Distributed Mission Training (DMT) console during exercise Centurion Warrior 25 (photo from [7])

THE BVR COMBAT DEMONSTRATION SCENARIO

As most of the above-mentioned LVC training activities are executed in a CLASSIFIED environment, it is not possible to share details of the missions performed. However, in the past years the RIACE team has implemented a number of demonstration scenarios for several test, development and demonstration purposes.

The first scenario is a Beyond Visual Range (BVR) combat scenario implemented for the purpose of testing that demonstrates the use and customization of internal doctrines that control the Computer Generated Forces (CGF) behaviours as well as the simulated “mechanization” of combat operations.

The mechanization of combat operations involves the combined use of:

- CGF sensors (e.g. Radar, Radar Warning Receiver, Missile Approach Warner, Datalink);
- CGF weapons as Semi-Active Radar Homing (SARH) missiles, Inertial-Midcourse Active Radar Homing (INS/ARH) missiles, Infra-Red (IR) missiles and guns;
- CGF counter-measures as jammers, towed decoys, chaffs and flares.

The simulated scenario implements a Blue DCA mission with a 4-ship, two sections flight of fighters on a racetrack Combat Air Patrol (CAP) attempting to defend a Fighter Area Of Responsibility (FAOR) from hostile airborne intruders. Figure 5 shows the geometry of the engagement area.

The dissimilar aircraft combat scenario includes fighters of the 4+th gen. (Italian Air Force EF2000) on the Blue Side (Friendly) armed with a combination of INS/ARH and IR missiles and of the 4th gen. fighters (MiG-29) on the Red Side (Hostile) armed with a combination of SARH and IR missiles.

The scenario gives a clear advantage to the Blue Side that is partially balanced by preventing the EF2000 CGFs from employing their ECM suite (switched OFF).

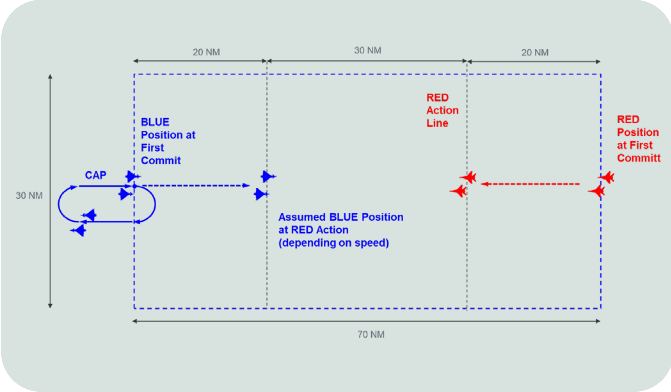

5–BVR Scenario: geometry of engagement area

The Italian Air Force EF2000s on the Blue Side are complemented by friendly EF2000s of other NATO nations (UK and Germany) in adjacent FAORs and a NATO AWACS acting as an Airborne Early Warning (AEW) and Flying C2 for the Link-16 network.

The purpose of the test/demonstration environment is to check the Identification, Engagement and BVR tactics implemented in the RIACE Doctrine Editor, as well as to verify the overall "credibility" of the simulation evolution and outcome (e.g. winner and loser, ammunition expenditure, performance aspects, etc.). The Red MiG-29s can be instructed to follow different attack game plans with their artificial intelligence models either switched OFF (to simply provide targets for the Blue game plan) or ON (to attempt an effective contrast of the Blue game plan).

The several runs of the scenario executed with different settings and Red game-plans, have demonstrated a very good BVR combat simulation with respect to the simulated performance of aircraft, systems and weapons, as well as the implemented doctrines.

Figure 8 and Figure 9 show some snapshots taken from the RIACE HMI during the different runs of the simulation.

6–BVR Scenario: details of CGF opponents

7–BVR Scenario: overall scenario with Blue & Red forces

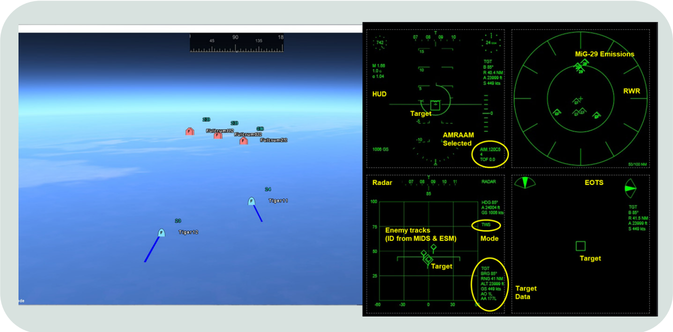

8–BVR Scenario: initial engagement and CGF sensors view on a Friendly fighter

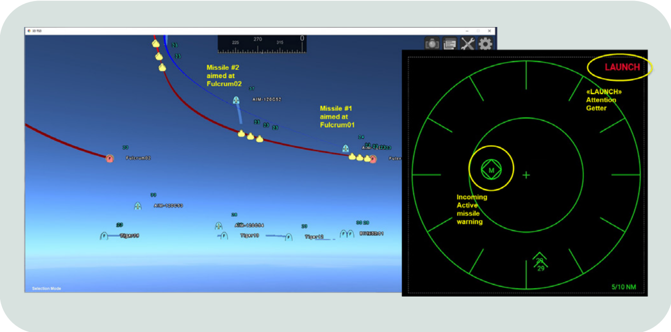

9–BVR Scenario: Red CGFs attempting to evade INS/ARH missiles gone active (see the RWR view) and dispensing chaffs

THE EDF-FEDERATES MULTI-DOMAIN DEMONSTRATION SCENARIO

A complex multi-domain scenario has been prepared for the European Defence Funded (EDF) Project FEDERATES (FEDerated Ecosystem of euRopean simulation Assets for Training and dEcision Support).
The FEDERATES project ([8]) proposes a European Modelling and Simulation as a Service (MSaaS) solution for Distributed Synthetic Training and decision-making, by using existing MS assets, leading to new training opportunities, pooling of simulation resources, access to training, reduced set-up times, reduced costs and faster development of future solutions, as well as to a new marketplace for services.
For this project, RIACE provides a complex, immersive Synthetic Environment with several "vignettes" running simultaneously in a multi-domain warfare environment for the "Multi Domain Operations Use Case".
The goal and scope of the Use Case is to demonstrate and validate the effectiveness of a multi-national EU/US force in:

- repelling a small-sized (battalion) invasion combining land, maritime, air, cyber and space engagements;
- coordinate multiple Air C2s in the execution of several aerial warfare missions (air-to-air and air-to-surface);
- engage surface and sub-surface opponents coordinating own surface/subsurface and airborne ASuW/ASW assets.

The scenario includes the use of hostile cyber-attacks attempting to compromise the Recognized Air Picture (RAP) and massive use of RIACE auto-generated civilian traffic ([10]) in the air, ground and surface domain.
The geo-political scenario is derived from the Spartan Alliance 18-8 exercise ([2]) and assumes that a hostile nation (Sardinia) invades with battalion level force the area of Grosseto. The Blue Side (European and US Forces) isolate the landing force and counter-attacks in the ground warfare domain. At the same time, a naval block is issued in front of Grosseto (Anti-Surface and Anti-Submarine Warfare) and air superiority missions (DCA/OCA) are flown by assets of the 5th Gen. over the Tyrrenian Sea.
Simultaneously, Ballistic Missile Defence (BMD) operations are carried on to defend the mainland from Tactical Ballistic Missiles (600 km category) launched from Sardinia.
The four Multi-domain vignettes inside the overall geo-political scenario thus are:

- OCA/DCA vignette;
- Ground Warfare vignette;
- Blue Water ASuW/ASW vignette;
- Tactical Ballistic Missile Defence (TBMD) vignette.

10 – 5th Gen. RIACE automated opponents for DCA/OCA

In the first vignette (OCA/DCA) the Blue assets of the 5th gen. perform Beyond Visual Range (BVR) engagements against Constructive SU-57 simulated by RIACE. The automated opponents include the S-70 Collaborative Combat Aircraft (CCA) drones and are controlled by a customizable expert system that provides a capable opponent proficient in BVR tactics as we have seen in the example of the previous chapter. The added value of this simulation with respect to the BVR combat demonstration seen before is the increased complexity given by the "stealth" characteristics of the SU-57 platform (e.g. low observability and weapons in internal bays) and the INS/ARH weapons carried.

In the Ground vignette a complex warfare scenario (at battalion level) is simulated between the town of Grosseto and the Tyrrenian coast: infantry supported by Armored Personnel Carriers (APCs) and Infantry Fighting Vehicles (IFVs) are facing on the north-east side of the town, while Main Battle Tanks (MBTs) fight for the control of the airport on the western side.
A Joint Terminal Attack Controller (JTAC) directs ground, air and surface fires (joint fires) to support the ongoing fight.

The multi-domain scenario includes several air, ground, and surface assets including EF2000, EA/18G, AH-64D, HH-60G, M109 Howitzers, Ariete and M1 Abrams MBTs, Centauro B2, Bradely M2, Lince, Humvee and several Unmanned Air Vehicles (UAVs) as the Leonardo Xplorer Medium Altitude Long Endurance (MALE) UAV.
A convoy attempting to resupply the Red fighters in the town from the coast is contrasted by artillery fire and a SEAD mission aimed to suppress the Red Surface Based Area Defence (SBAD) is flown by Blue airborne assets.

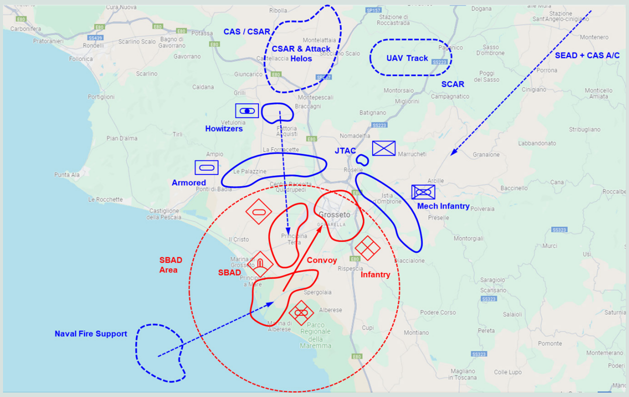

11 – Ground Warfare multi-domain vignette

12 – Leonardo Xplorer UAV overflying Grosseto

In the last vignette, Al-Hussein (SCUD type) Tactical Ballistic Missiles are launched against the Blue ASOC (assumed to be located in Poggio Renatico).
The Ballistic Missiles are detected and tracked by the multi-domain chain of sensors including the Italian Air Force remote early warning radars located according to Wikipedia ([9]) in the midst of RIACE auto-generated civilian airliners traffic ([10]). In the Tyrrenian Sea, a MMI Orizzonte class destroyer, a FREMM class frigate and a United States Navy (USN) Arleigh-Burke AEGIS class destroyer (armed with SM-3 missiles) contribute to tracking and engagement of the TBM. The defence of the Terminal area is assigned to a SAMP-T battery located in Poggio Renatico. The engagement of the TBM is coordinated using Link16 among the different multi-domain systems. Figure 14 shows the Al Hussein 600 km class missile along its ballistic trajectory to the intended target in Poggio Renatico; in the RIACE 2D Tactical Situation Display (2D TSD) in the bottom left of the figure the launch area and impact are marked as uncertainty ellipses as estimated by the (unclassified) RIACE early warning radar simulation model.

In the Blue Water vignette, several multi-domain Anti-Surface Warfare (ASuW) and Anti-Submarine Warfare (ASW) operations are simulated, including the launch of Teseo anti-ship missiles from Italian Navy FREMM Frigates upon target designation from an Italian Air Force ATR-72MP using simulated Link11 communications.
The Figure 13 shows an example of ASW operations with an Italian Navy (MMI) SH-90 ASW helicopter locating a Kilo class diesel submarine (SSK) by dispensing sonobuoys and providing the target data to the C2 ship using Link16.

13 – ASW operations: MMI SH-90 dispensing sonobuoys around a Kilo SSK contact

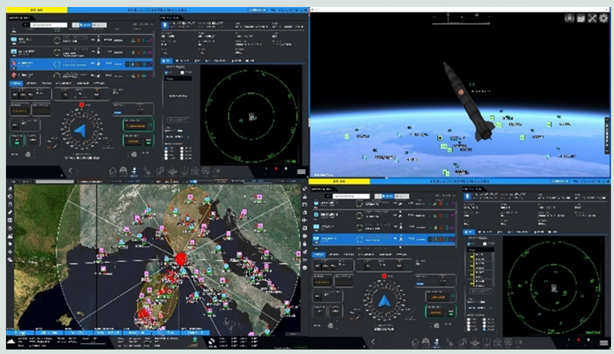

14 – TBMD vignette: Al Hussein Missile near apogee

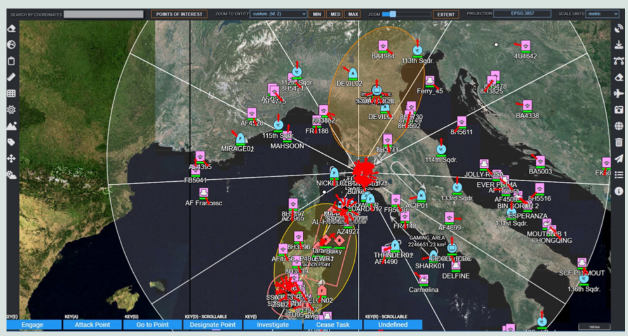

15 – TBMD vignette: Enlarged view of the 2D TSD showing the Radar TBM launch/impact point estimation functionality

CONCLUSIONS

The paper presented a panoramic of the more recent developments of the RIACE System and illustrates a number of hypothetical multi-domain LVC scenarios supported by RIACE.

Luca Cistriani: luca.cistriani@leonardo.com

REFERENCES

[1] L. Cistriani, "*RIACE – In missione nel Synthetic Environment*", published on Rivista Aeronautica, Italy, June-July 2018.
[2] S. Bonfiglio, L. Cistriani, P. Caldarelli, G. Di Mauro, "The RIACE Synthetic Environment in the Spartan Alliance and Spartan Warrior Exercises", in Polaris Innovation Journal, Nr. 38, 2019.
[3] S. Rigoni, "*Tecniche di Distributed Synthetic Training per lo sviluppo di ambienti simulativi complessi*", Master Thesis Work in Software Engineering, Università Telematica "e-Campus", A.A. 2024/2025.
[4] L. Cistriani, "Agent-Oriented Distributed Simulation of C3I Systems with RIACE", presented at NATO International Conference on Military Communication and Information Systems (ICMCIS), Udine, Italy, 2022.
[5] Hawker J., "NATO Flight Training Europe - NFTE", presented at International Fighter Conference 2025, Rome, Italy.
[6] https://www.difesa.it/smd/news-italia/conclusa-spartan-warrior-24-02/57962.html
[7] https://www.aeronautica.difesa.it/news/esercitazioni-conclusa-la-centurion-warrior-2025/
[8] https://defence-industry-space.ec.europa.eu/document/download/2323e186-afd1-43e8-860a-4c910563a680_en?filename=FEDERATES%20-%20Factsheet_EDF22.pdf
[9] https://en.wikipedia.org/wiki/NATO_Integrated_Air_Defense_System
[10] A. Amorosi, S. Bonfiglio, P. Caldarelli, L. Cistriani, "Simulation Model of the Civilian Traffic in the Italian Upper Airspace for Synthetic Environment Applications", Proceedings of the 10th MIMOS Conference, Rome, Italy, 2012.
[11] https://www.af.mil/News/Article-Display/Article/2367116/spartan-warrior-20-9-nato-strengthens-air-power-through-distributed-training/

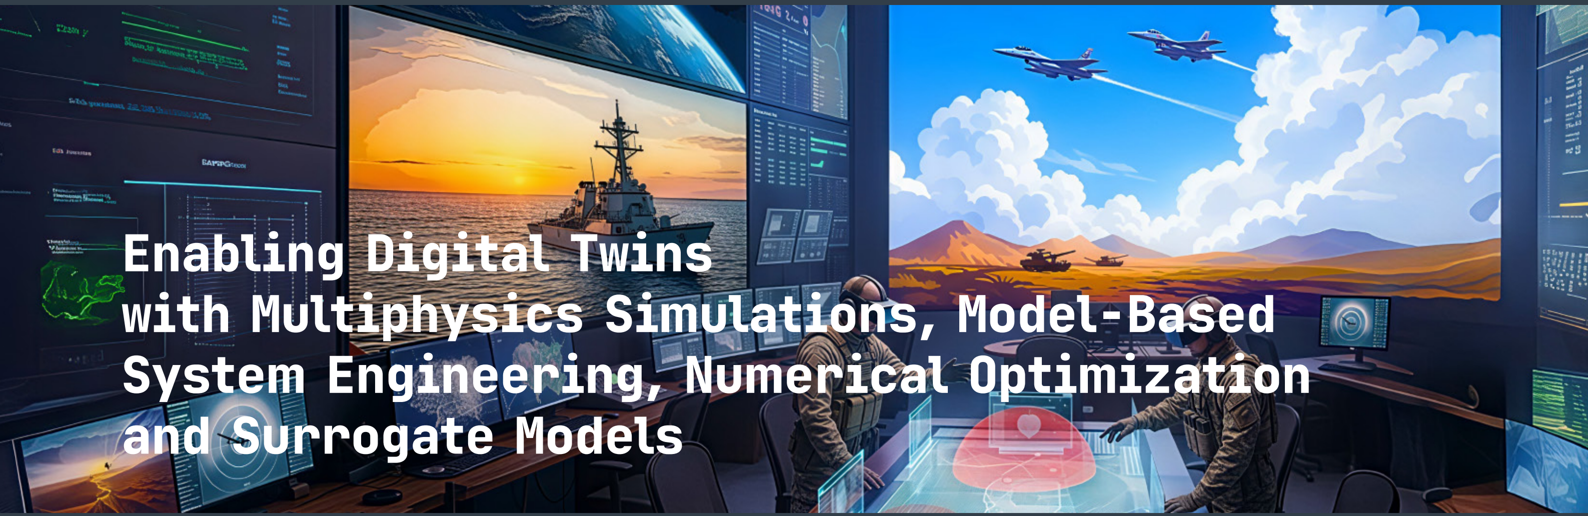

Enabling Digital Twins with Multiphysics Simulations, Model-Based System Engineering, Numerical Optimization and Surrogate Models

Joel Guerrero, Stefano Centomo, Alberto Dagna, Damiano Paniccia, Viola Rossano, Andrea Lopez, Tony Di Fabbio, Eric Segalerba, Gian Marco Zampa, Alberto Artoni, Daniele Fontana, Erika Pellegrino

Leonardo – Corporate/Strategy and Innovation

This technical white paper presents a comprehensive overview of the Digital Twins and Advanced Simulations research group's approach to implementing multi-physics Digital Twins. We adopt a broad definition of a Digital Twin as "a digital representation of a planned or existing physical asset, process, or requirement. It is the embodiment of the physical world within the digital realm.". Our discussion explores the digital product lifecycle and highlights the key pillars and enabling technologies that accelerate the deployment of multi-physics Digital Twins. Finally, we showcase the depth and breadth of the technologies and methodologies used by our research group through selected applications that demonstrate advanced capabilities in modelling and simulation, numerical optimization and workflow automation and orchestration.

INTRODUCTION

Businesses across all industries are embracing digitalization to drive innovation, accelerate scientific discovery, and create new products and services with greater speed, efficiency, and quality. Digital Twins (DT) and Modelling and Simulation (M&S) technologies are at the forefront of this transformative digital revolution and evolution. Digital Twins are ubiquitous in the Aerospace and Defence (A&D) sector, Leonardo's core business, and are part of Leonardo's integral vision and roadmap towards the future, "**Be Tomorrow – Leonardo 2030**". While definitions of Digital Twins vary across sectors (such as energy, materials, industrials, utilities, healthcare, finance, consumer goods, information technology, communications, real estate, etc); the concept broadly refers to a digital model of an intended or actual real-world physical product, system, or process that serves as a digital counterpart of it for purposes such as modelling, simulation, integration, testing, monitoring, and maintenance [1]. Digital Twins reduce reliance on physical prototypes, shorten development cycles, and enhance safety, quality, and operational flexibility. They enable informed decision-making across disciplines, significantly mitigating risks and minimizing material waste.

Ultimately, the scope and design of a Digital Twin depend on its intended purpose.

Within the context of this article, we adopt the following general definition of a Digital Twin:

A Digital Twin is a digital representation of a planned or existing physical asset, process, or requirement. It is the embodiment of the physical world within the digital realm.

This definition underscores the importance of uncertainty quantification, which integrates the stochastic nature of real-world phenomena into the deterministic models and governing equations used in the digital domain.

The tenet of the Digital Twin for the Digital Twins and Advanced Simulations research group (DTAS) is illustrated in Figure 1. This figure represents the digital product life-cycle, that is, the complete journey of a digital product from conception to retirement. It encompasses all stages, from initial idea generation and development to digital engineering (which leverages computational tools, digital models, and connected data to support design, testing, monitoring, and maintenance), to its phase-out.

Within this lifecycle, we envision a seamless connection between **digital engineering** and **digital manufacturing** (which is the use of computational tools to optimize manufacturing services, supply chains, products, and processes), creating an integrated approach that bridges design and production. The digital lifecycle is inherently iterative, continuous feedback loops across all stages ensure that both digital and physical assets adapt to evolving market demands and technological trends. Through iterative digital design, we enable product optimization and informed decision-making using advanced methodologies such as: multi-disciplinary design optimization (MDO), uncertainty quantification (UQ), robust design, and certification by simulation (CbS) methodologies. These approaches collectively enhance efficiency, reduce risk, and accelerate innovation throughout the product lifecycle.

The principal technologies and methodologies employed by the DTAS group are illustrated in Figure 2. Central to these is modelling and simulation (M&S), which encompasses multi-physics simulations conducted at varying levels of fidelity: low, medium, or high. This foundational capability enables the integration and coupling of diverse disciplines, including aerodynamics, heat transfer, electromagnetism, mechanics, structural dynamics, finite rate chemistry, acoustics, and more. Multi-physics simulations often produce vast amounts of complex data, and this is where simulation process and data management (SPDM) prove essential.

SPDM refers to the systematic framework for organizing, storing, and controlling simulation workflows and their associated data, ensuring consistency, traceability, and efficient reuse across projects.

Another key technology is model-based systems engineering (MBSE), which is closely linked to SPDM. Within MBSE, functional models are developed using data derived from multi-physics simulations or other authoritative sources of information. Beyond model creation, MBSE provides a structured framework for process standardization and change management, ensuring consistency, traceability, and alignment across complex engineering projects. To support these engineering tasks, we employ high-performance computing (HPC) resources, both CPU and GPU, as well as advanced high-performance databases. MDO enables the integration of diverse disciplines, while reduced-order models (ROM) and surrogate models facilitate efficient optimization studies. In addition, artificial intelligence and machine learning (AI/ML) techniques are leveraged to accelerate data consumption and analysis. All processes depicted in Figure 2 are interconnected, enabling bidirectional exchange of information. The supporting technologies are inherently intensive in data, software, and hardware requirements, while each pillar remains firmly grounded in physics-based principles. Effective use of these technologies demands a deep understanding of the underlying physics, rigorous validation and verification, and careful uncertainty quantification (whether aleatoric or epistemic) of the models and assumptions employed. The ultimate objective is certification by simulation (CbS): reducing reliance on physical tests and prototypes prior to production or entry into service through the deployment of digital twins.

In summary, our foundational pillars are modelling and simulation, MBSE and SPDM. These are complemented by MDO, AI/ML, numerical optimization and workflow automation, which collectively enable iterative design and accelerate data consumption. All of these technologies are further empowered by HPC, with the Davinci supercomputer serving as our most important physical asset. Effective use of these technologies demands a deep understanding of the underlying physics, rigorous validation and verification, and careful uncertainty quantification (whether aleatoric or epistemic) of the models and assumptions employed. The ultimate objective is certification by simulation (CbS): reducing reliance on physical tests and prototypes prior to production or entry into service through the deployment of digital twins. In summary, our foundational pillars are modelling and simulation, MBSE and SPDM.

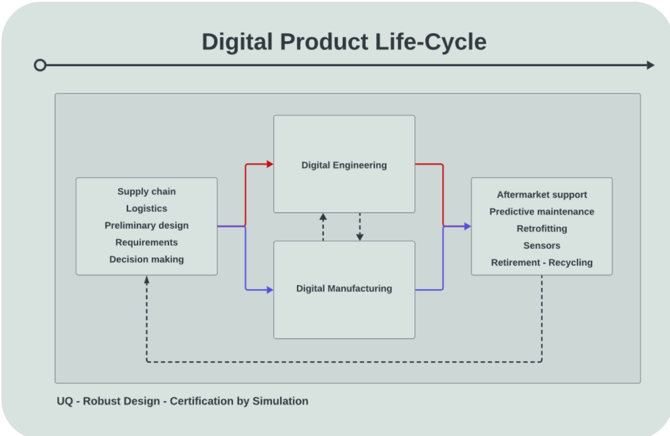

1 – Digital product life-cycle from conception to retirement

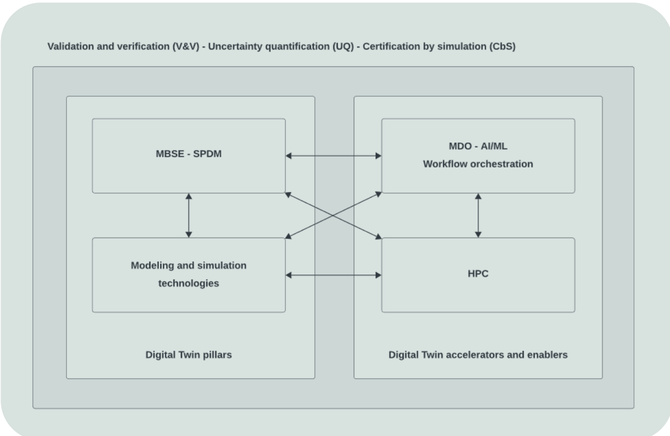

2 – Digital Twin pillars and accelerators and bidirectional exchange of information among processes

These are complemented by MDO, AI/ML, numerical optimization and workflow automation, which collectively enable iterative design and accelerate data consumption. All of these technologies are further empowered by HPC, with the Davinci supercomputer serving as our most important physical asset.
In the following sections, we highlight the breadth and depth of the different technologies, methodologies, and applications explored by the DTAS research group. The selected applications focus on:

- Model-based system engineering (MBSE) and digital threads;
- Leveraging Reduced-Order Models (ROM) and machine learning for system performance prediction;
- Multiphysics simulations (MPS) and the orchestration and automation of complex modelling and simulation workflows;
- Optimization of physical models (turbulence modelling) through data-driven approaches;
- Inverse design and adjoint optimization.

We strive to implement and deploy multi-physics Digital Twins to support the entire lifecycle of Leonardo's products and services.

Application 1. MBSE for Digital Thread and Digital Twin: Integrating Heterogeneous Models for Simulation-Based Validation

In complex aerospace systems engineering, it is essential to identify tools, technologies, and modelling languages that are robust, maintainable, and valid over the long term. Concepts such as LOTAR (Long-Term Archival and Retrieval) [2] play a fundamental role in preserving engineering data, models, and design knowledge for decades, ensuring that digital assets remain accessible, interpretable, and trustworthy throughout the lifecycle. This long-term perspective sets the foundation for developing interoperable and traceable system models within a Digital Twin and Digital Thread framework.
This research stream is focus on advancing methodologies and technologies that enable the structured definition, integration, and verification of complex systems through digital models. While systems modelling language v1 (SysML v1) [3] has been widely adopted in the past, modern systems increasingly require the integration of heterogeneous models across multiple domains, including physics-based, control, and data-driven components. Our activities focus on ensuring interoperability, traceability, and digital validation across these domains. SysML v2 [4] is crucial for our approach, the next-generation modelling language that improves expressiveness, precision, and tool interoperability. SysML v2 provides a rigorous semantic foundation, a standardized API, and an intuitive syntax, allowing system architects and engineers to define structures and behaviours that are directly executable and integrable with simulation tools.
To enable multi-domain simulation, we combine Functional Mock-up Interface (FMI) and System Structure and Parameterization (SSP) [5]. FMI allows the exchange and co-simulation of Functional Mock-up Units (FMUs) between heterogeneous tools, while SSP defines the structure, hierarchy, and parameterization of complete systems. Together, they ensure interoperability, traceability, and reuse of simulation assets across lifecycle phases. This approach allows each modeler to work with their preferred tool while integrating seamlessly within a unified digital environment.
Standards such as AP242 [6] for CAD and ReqIF [7] for requirements ensure cross-domain interoperability, complementing the long-term accessibility perspective provided by LOTAR. In the broader Digital Thread, which spans the entire lifecycle from design through manufacturing, MBSE plays a central role in capturing all structural and functional information of the system under development. Open Services for Lifecycle Collaboration (OSLC) [8] is currently being studied by our group to enable future connection and interoperability between tools, providing a framework for secure, non-duplicative linking of resources while preserving a consistent digital thread across all lifecycle phases.
The use case presented hereafter, demonstrates this methodology through a modular Simulink aircraft model integrating: model-based functional components (flight dynamics, propulsion, flight control), data-driven models (neural-network-based aerodynamic predictions), and Integrated requirements monitoring implemented directly in the simulation.
This case study illustrates how heterogeneous models from different domains can be integrated and executed in a coherent digital environment, enabling early verification of system requirements, optimization of system parameters, and traceability across the lifecycle. To support this long-term and interoperable approach, Figure 3 presents the overarching Digital Thread vision developed within our research unit. The diagram synthesizes the methodological framework adopted, showing how Digital Requirements, Digital System Models, Digital Verification, and Digital Manufacturing are organized into a coherent, standards-based lifecycle.

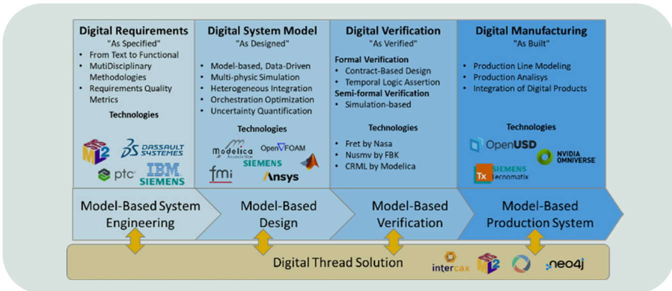

3–High-level overview of the integrated lifecycle vision adopted by the MBSE group

Each layer reflects the technologies, modelling formalisms, and interoperability mechanisms that our team considers foundational for enabling robust Digital Twin solutions. This conceptual view also highlights the central role of MBSE as the structural backbone connecting heterogeneous models, ensuring semantic consistency and traceability across all lifecycle phases. The methodology demonstrated in our use case is directly derived from this framework and represents a practical instantiation of our research vision in a simulation-based environment.

This methodology provides clear benefits, including:

- Reduced gap between design and verification through simulation;
- Enhanced interoperability among tools and disciplines;
- Reusable, modular models supporting multi-domain collaboration;
- Improved efficiency in meeting system requirements and constraints;
- Enabling the Digital Twin and Digital Thread as practical, executable frameworks spanning from design to manufacturing, with MBSE at the core of capturing all system information.

Use case discussion

To apply and test the approach we have just described, and thereby bridge the gap between requirements and their verification in a digital environment, we employed a simple use case. The use case consists of an aircraft model, which accounts for the following design specifications and requirements:

- Point-mass 3-DOF flight mechanics [9];
- Data-Driven model for the prediction of airfoil Aerodynamics [10];
- Atmosphere model based on ISA [11];
- Turbo Jet model for propulsion [12];
- Flight controls for throttle and pitch commands.

Each of these, alongside the definition of their interfaces, has been first defined in SysML v2, allowing to define a common ground between multiple potential modeler and different simulation tools. Then, a model for each has been generated, either in Simulink [13], AmeSIM [14] or OpenModelica [15]. The FMI standard has been used to create interoperable black boxes of these models, which were finally integrated in a single simulation environment (Simulink in the present case). The resulting model is shown in Figure 4.

This model simulates the simplified behaviour of a generic aircraft during flight operations. As such, it is possible to implement the requirements that guide the design of a certain aircraft configuration in the model itself, allowing to verify in a digital manner if such aircraft can respect these requirements. Since it implements a propulsion system model, one such behaviour is the consumption of fuel during flight. Therefore, it is possible to verify that the on-board fuel is sufficient to complete certain flight scenarios, such as:

"The aircraft shall have enough fuel to complete a flight mission lasting "x" seconds".

Such requirement has been implemented in the form of a state-machine, which monitors along the simulated mission whether the aircraft has fuel remaining in the tanks. The model has then been simulated with a starting fuel on-board of 200 kg, which, as it can be seen in Figure 5, it is not enough to complete a generic mission lasting 300 seconds. In fact, the top plot shows that the fuel remaining becomes negative at around 100 seconds, with the Boolean output of the requirement (in the bottom plot) switching from True (1) to False (0). This prompted the need to load more fuel: thus, 350 kg has been set as the starting value, which allowed to complete the mission. However, it is noticeable that around 35 kg of fuel are still remaining, indicating there is room for optimization. Therefore, a last mission has been simulated with 315 kg of fuel, which resulted in just a couple of kilograms of fuel left.

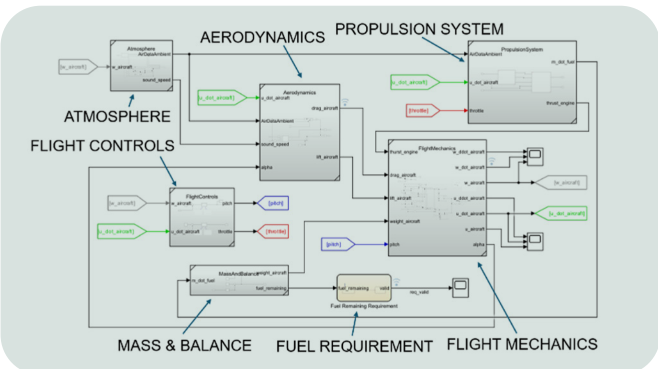

4 – Simulink model used as use case for requirement implementation

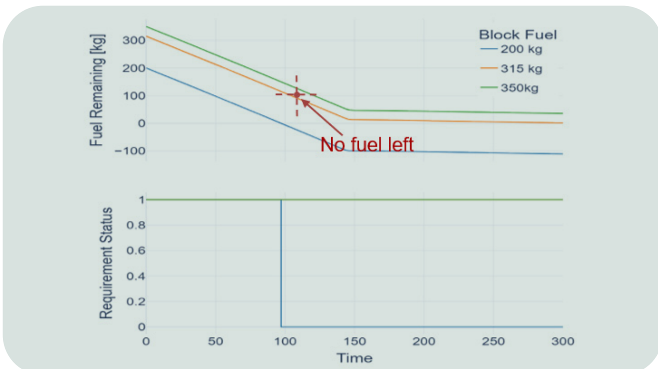

5 – Plots for the fuel remaining (top) and requirement verification (bottom) of the three simulated scenarios

Summary

To summarize, the use case demonstrated, in a simplified fashion, the possibility and advantage of being able to implement the requirements that guided the design directly in a simulation model, as it can be immediately confirmed, without extensive post-processing and analysis, whether or not they are satisfied by the designed architecture and configuration. This approach could be further expanded, through the inclusion of multi-fidelity and multi-domain models, gradually increasing the capability of digital simulations of verifying digitally the product requirements. Additionally, with continuous additions and structuring of generated entities, a true digital thread technological solution could be achieved, enabling the development of new systems in a purely digital environment, eliminating the need for extensive physical testing and prototyping.

Application 2. A Supervised Machine-Learning Approach for Turboshaft Engine Dynamic Modelling Under Real Flight Conditions

Real-time control of helicopter turboshaft engines is challenging due to the strong engine–drivetrain coupling, which requires robust yet efficient models [16]. High-fidelity simulations are accurate but computationally demanding, motivating the use of reduced-order and transfer-function models [17], typically based on static maps and linear corrections for environmental effects [18]. Although reliable, these physics-based approaches require extensive calibration and flight testing, and their accuracy may degrade as operating conditions evolve. The nonlinear and fast-varying nature of rotorcraft engines further reinforces the need for accurate dynamic models for design, diagnostics, monitoring and emerging predictive-maintenance applications.
The rise of smart technologies has introduced AI and ML as effective tools for capturing nonlinear turbomachinery behaviour [19], enabling applications in flow control [20], deterioration modelling [21], and fault prediction [22], despite ongoing concerns about interpretability and safety [23]. Approaches such as dynamic mode decomposition (DMD), Koopman Analysis, and in particular SINDy [24] offer a compromise between data-driven modelling and physical interpretability, and have recently shown strong performance in turboshaft engine modelling.
Building on this progress, this work investigates ML methods for modelling the AW189 (shown in Figure 6) prototype engine using real flight data. Multi-Input Single-Output models based on Feed-Forward Neural Networks (FFNN) and Long Short-Term Memory (LSTM) networks are trained to predict engine torque (TRQ) from engine and environmental variables, while SINDy is used to derive a compact single-input single-output (SISO) model linking fuel flow (WF) to torque. This combined strategy yields accurate nonlinear dynamic prediction with minimal complexity, demonstrating improved adaptability and real-time potential over traditional transfer-function approaches, and highlighting its suitability for future health-monitoring and predictive-maintenance frameworks.

6 – Leonardo's AW189 twin-engine helicopter

Methodology and framework

The methodology combines Neural Network (NN) models and Sparse Identification of Nonlinear Dynamics (SINDy) to model helicopter engine behaviour. Two NN architectures were implemented using PyTorch: a Feed-Forward Neural Network (FFNN) for capturing nonlinear relationships between input and output variables at each time step, and a Long Short-Term Memory (LSTM) network to account for temporal dependencies in sequential data. Both were trained to predict engine torque (TRQ) from multiple input variables using Mean Squared Error as the loss function.

Their architectures are sketched in Figure 7 and Figure 8, respectively. In parallel, the SINDy algorithm was applied to derive a low-dimensional, interpretable model linking fuel flow (WF) and torque (TRQ). By identifying a minimal set of governing equations through sparse regression, SINDy balances simplicity and accuracy, offering a physics-consistent view of engine dynamics complementary to the purely data-driven NN models.

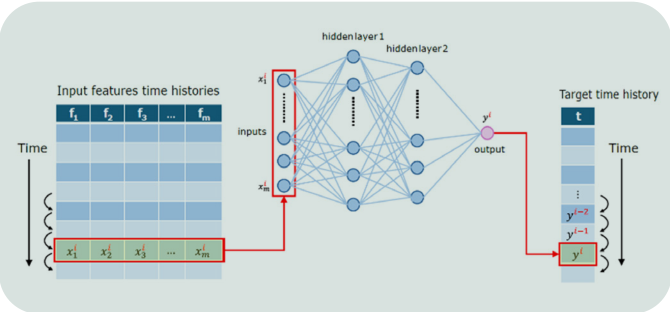

7 – A simple illustration of the FFNN architecture for the torque time-history prediction

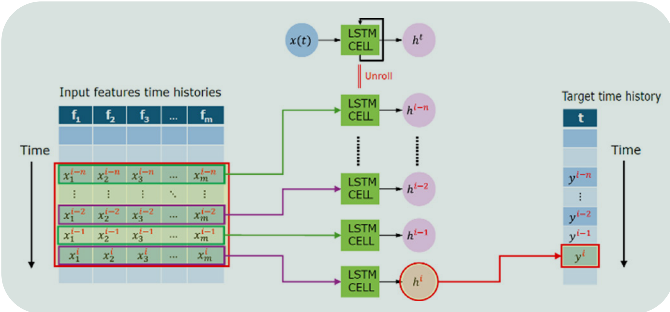

8 – A simple illustrration of the LSTM architecture for the torque time-history prediction

MISO Neural Network Results

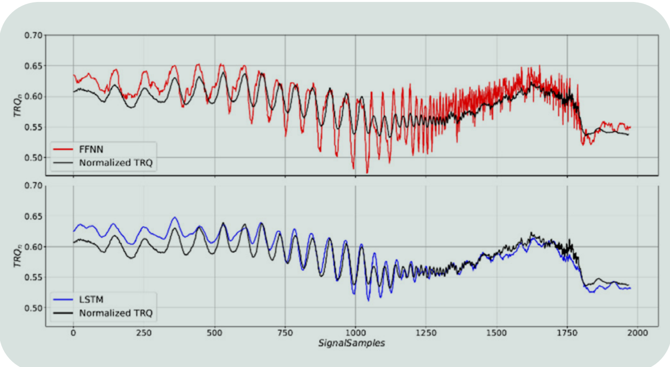

9 – Comparison of the normalized TRQ predictions of the FFNN and LSTM models

Figure 9 shows the results for a manoeuvre from a given flight (ID1 in this case), comparing the normalized TRQ predictions of the FFNN and LSTM models. Although their rMAE values appear similar ($rMAE_{FNN}$=3.75% and $rMAE_{LSTM}$=2.34%), the FFNN clearly fails to reproduce the engine dynamics. Its output is noticeably noisy and unable to capture rapid variations, particularly during manoeuvres that deliberately excite the engine response, such as collective sweeps. This lack of smoothness arises from the FFNN architecture, which processes each time step independently and disregards temporal dependencies. In contrast, the LSTM predictions follow the true signal more closely thanks to their recurrent structure, which effectively models sequential data and captures long-term temporal patterns.

SINDy Results

Figure 10 shows the same flight manoeuvre previously analysed for the MISO neural networks. The simulation results obtained with the SINDy-based model are highly accurate, reproducing the TRQ dynamics very well, with only minor under- or overestimations around some oscillation peaks. The overall error, expressed through the rMAE for the test flight, remains below 3%. These results highlight the capability of the SINDy framework to extract a simple and interpretable dynamic model even without detailed prior knowledge of the system physics. In this case, achieving this level of accuracy relies on extending the model to second order: the inclusion of the fuel-flow time derivative as an additional input, together with the extra initial condition on the time-derivative of TRQ, appears to be the key factor enabling the improved performance.

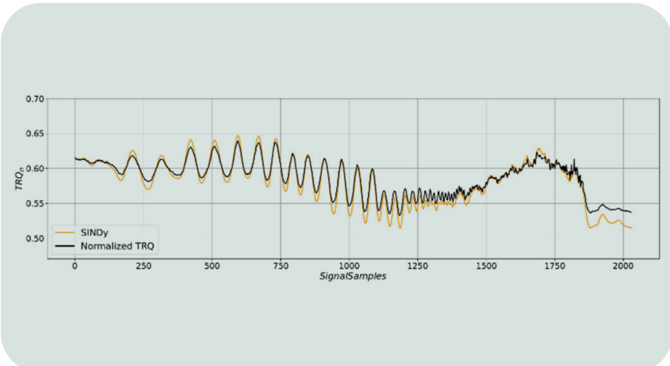

10 – TRQ predictions using the MISO model

Summary

This application explored three data-driven methods to model helicopter turboshaft engine dynamics using real flight data from Leonardo's AW189. Two Neural Network architectures, a Feed-Forward (FFNN) and a Long Short-Term Memory (LSTM), were trained to predict engine torque from engine and environmental variables. While the FFNN struggled with unstable outputs, the LSTM delivered smoother and more realistic torque predictions, proving effective for capturing temporal dynamics. However, its performance degraded over time, highlighting the need for periodic retraining to account for engine aging, an insight directly relevant to long-term health monitoring and predictive-maintenance planning.
To address the "black-box" nature of neural networks, the Sparse Identification of Nonlinear Dynamics (SINDy) approach was tested as a transparent, physics-inspired alternative. The resulting second-order SISO SINDy model, although still preliminary and focused on a single engine component, successfully captured torque behaviour with high accuracy and robustness, even under complex manoeuvres. Its interpretability and compactness make it particularly suitable for on-board monitoring, anomaly detection, and early-warning maintenance frameworks.
Overall, these findings suggest a promising hybrid future: combining the adaptability of neural networks with the interpretability of sparse, physics-informed models to create next-generation digital twins for helicopter engines, digital twins explicitly designed not only for improved control and simulation, but also as core enablers of proactive diagnostics and predictive maintenance.

Application 3. Enabling Digital Twins with Multiphysics Simulations and Data-Driven Reduced Order Models

The design of modern aircraft has become a highly complex, multidisciplinary task. Aerodynamic, structural, electromagnetic, thermal and propulsion requirements, among others, often conflict, and improvements in one area can impede performance in another. Iterative testing and coordination among specialists further complicate the process, making it essential to adopt integrated, collaborative approaches from the earliest design stages.
To meet these challenges, an automated, multi-user workflow that spans across low-, mid-, and high-fidelity Multiphysics simulation approaches is presented in this application case. By enabling easy data sharing, intuitive workflow management, and outputs ready for real-time simulation and digital-twin use, the framework accelerates design processes and supports more informed decision-making. By orchestrating both commercial and in-house multiphysical solvers and centralizing data management, the system ensures consistency, traceability, and rapid uptake by existing teams.
Together with the involvement of the Preliminary Design Department and the Aeronautics Division's Innovation Management group, this study extends the methodologies introduced in reference [25] and addresses the conceptual and preliminary design phases of next-generation aircraft, where performance trade-offs are most critical. Hereafter, we employ multidisciplinary optimization (MDO) techniques to explore vast solution spaces, identify Pareto-optimal configurations, and accelerate decision-making.
Although tailored to fighter aircraft, the methodology can be extended to any aircraft category, from unmanned aerial vehicles to commercial airliners. The approach addresses three common roadblocks in complex aerospace system design:

- Physics conflicts: Reconciling antithetical requirements, such as lift-optimized geometries versus radar cross section (stealth) through concurrent optimization rather than sequential iteration;
- Collaborative complexity: Allowing multiple experts to contribute data and models in a shared environment, supporting teamwork and coordinated decision-making;
- Discipline integration and usability: the workflow is designed to scale across multiple domains including aerodynamics, radar signature prediction, structural analysis, and thermal management, while automating repetitive tasks and interfacing seamlessly with both commercial software and in-house tools. By consolidating these capabilities, the approach supports shorter development cycles and enables design teams to deliver innovative, high-performance aircraft more efficiently. The automated workflow consists of a sequence of integrated computational steps, beginning with parametric geometry generation, followed by aerodynamic analysis, detailed structural evaluation with mass estimation, and flight-mechanics performance assessment as shown in Figure 11.

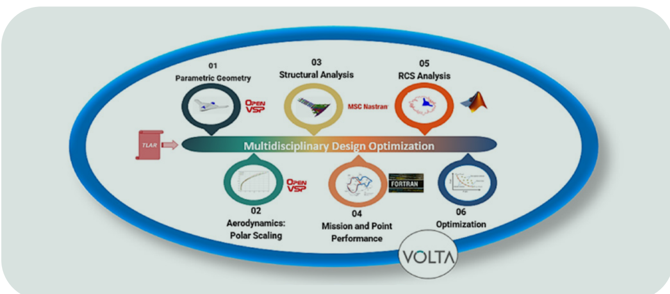

11-The proposed workflow developed using X47-B test case

Methodology and framework

A central aspect of this study is the integration of multi-fidelity computational tools from different disciplines. This approach balances accuracy and computational cost, enabling fast and reliable evaluations of aerodynamic performance, structural behaviour, and radar cross-section (RCS) within a single, unified framework.

The X-47B [26] was selected as the representative use case due to the accessibility of relevant data in open literature and its status as a relatively open case, facilitating collaboration and data exchange with external software vendors.

The methodology leverages platforms such as VOLTA [27] and ModeFrontier [27] to orchestrate and automate interconnected workflows as shown in Figure 12. In Figure 13, a comparative analysis between the baseline configuration (in blue), the discipline-specific optima, and the compromise solution (in red) obtained through multidisciplinary approach is shown. Clear geometric trends emerge across the single-discipline optimizations. The performance-oriented optimum shows noticeable refinements of the leading and trailing edges, with an effective increase in wing aspect ratio aimed at reducing induced drag and improving aerodynamic efficiency.

These changes, while beneficial aerodynamically, tend to enlarge specular reflection regions, resulting in higher radar cross-section (RCS). Conversely, the RCS-oriented optimum exhibits more faceted shaping and increased sweep, characteristic of stealth-driven design. These modifications contribute to reduced RCS but degrade aerodynamic efficiency.

The trade-off solution lies consistently between the two extremes, demonstrating how the MDO framework negotiates antagonistic requirements. The resulting configuration preserves aerodynamic characteristics closer to the performance optimum while retaining significant stealth improvements relative to the baseline. This highlights the effectiveness of the proposed methodology in identifying balanced solutions that would be difficult to obtain through sequential or discipline-isolated optimization.

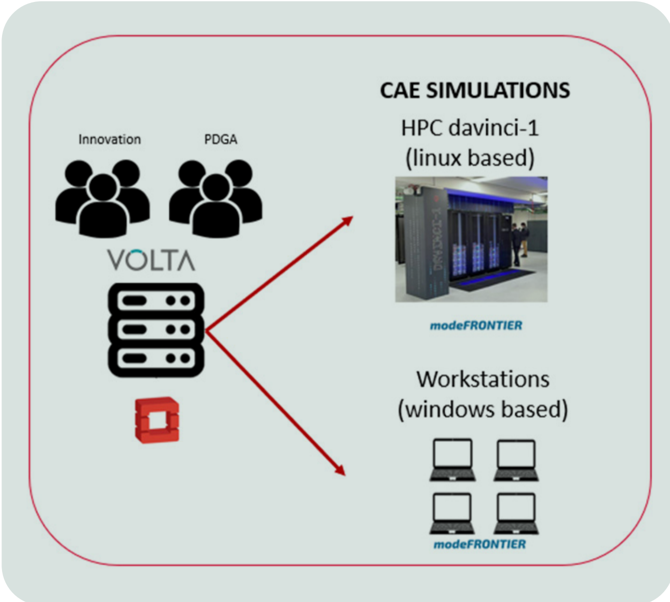

12-Computing Architectures

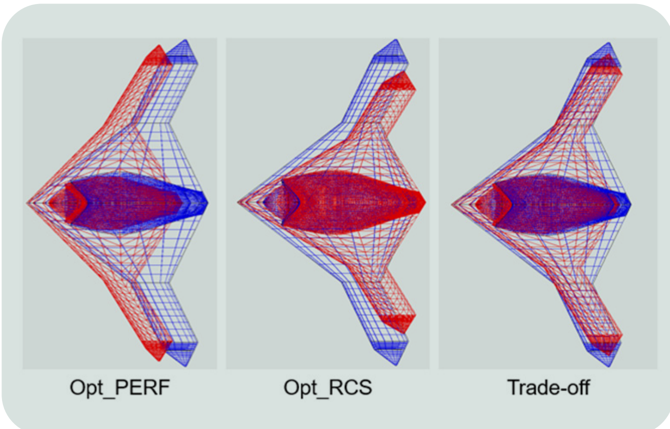

13-Comparison with the baseline configuration: performance optimum (left), RCS optimum (centre), and multidisciplinary trade-off solution (right)

Summary

This study presents a fully integrated, multi-fidelity, multi-domain optimization framework designed to manage the competing objectives of next-generation aircraft design. By unifying parametric modelling, aerodynamic analysis, structural assessment, and RCS evaluation within a cloud-enabled automated environment, the framework enables rapid exploration of extensive design spaces and supports collaborative decision-making across engineering teams. The results underscore the value of MDO methodologies in generating highly efficient aircraft configurations, sometimes non-intuitive, particularly useful during conceptual and preliminary design stages, where design flexibility is greatest and early choices carry long-term implications. Future work will extend the workflow to include thermal management and propulsion integration and will incorporate surrogate models to further reduce computational cost and enable real-time design-space interrogation.

Application 4. Data-driven RANS Turbulence Model for External Aerodynamics Applications

The development of data-driven turbulence models for Reynolds Averaged Navier-Stokes (RANS) simulations represents a major step forward in external aerodynamics, which is crucial in aircraft design, influencing performance, efficiency, and safety. While wind tunnel and flight testing offer high accuracy, their cost and time demands make them impractical for iterative design. Consequently, numerical simulations, particularly RANS models, are widely used due to their computational efficiency. However, this comes at the cost of reduced accuracy in capturing complex flow phenomena such as flow separation, shock–boundary layer interaction, and vortex dynamics. Improving turbulence modelling without compromising efficiency remains a major challenge.

Several studies have introduced data-driven corrections to traditional models, using evolutionary optimization for delta-wing flows and complex geometries and for turbomachinery applications [25]-[26]. Field inversion and Machine Learning (FIML) methods have also shown predictive capabilities, though with high computational cost and limited interpretability [27].
This application case introduces a data-driven enhancement to traditional RANS models, leveraging high-fidelity and experimental datasets to improve accuracy without sacrificing efficiency. The result is a new generation of turbulence models that combine the robustness of RANS with the adaptability of modern data-driven techniques, offering a practical solution for aerodynamic design challenges.

Methodology and framework

The proposed framework is based on the RANS equations, which remain the standard for industrial aerodynamic simulation. However, conventional turbulence closures such as Spalart-Allmaras (SA) or k-omega SST [28] rely on simplified assumptions that fail in complex regimes, particularly under transonic conditions with shock–boundary layer interaction and vortex breakdown. To address these limitations, this work introduces a physics-informed, data-driven correction to the baseline model.
The approach integrates high-fidelity reference data, obtained from large eddy simulations (LES), direct numerical simulations (DNS), and validated experiments, into the turbulence modelling process. Rather than adopting black-box surrogate models, which lack interpretability and physical consistency, we employ Gene Expression Programming (GEP), an evolutionary algorithm that generates explicit correction terms suitable for direct implementation in CFD solvers. During optimization (Figure 14), candidate corrections are inserted into the closure, tested through RANS simulations, and evaluated against reference data on pressure distribution, flow separation, and vortex behaviour. The best candidates evolve over successive generations toward a robust correction. Unlike single-case optimization, a multi-case training strategy is adopted to reduce case dependency and improve generalization.
The framework is fully open source, combining Julia and Python for algorithm development and OpenFOAM framework [29] for CFD simulations. Initial training employs canonical two-dimensional benchmarks from NASA's Turbulence Modeling Resource [30], including the wall-mounted hump and backward-facing step, which provide well-documented experimental data.
Validation then extends to three-dimensional configurations from the High-Lift Prediction Workshop (HLPW5) [31] and ONERA wind tunnel datasets [32], representative of realistic external aerodynamics. Figure 15 shows the qualitative results of a typical configuration exhibiting side of body separation. Figure 16 depicts the quantitative results of the drag polar, where the error bars of the numerical results were plotted using aleatoric and epistemic uncertainty quantification (UQ) models.

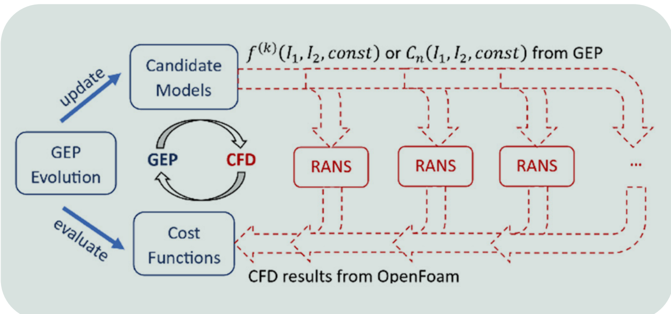

14 - Data-driven optimization scheme

Summary

The data-driven RANS model achieved clear improvements over conventional closures. Predictions of pressure distribution, skin friction, and separation points showed much closer agreement with experimental data. In particular, the model captured separation and reattachment more accurately, a longstanding weakness of traditional RANS approaches. These improvements were achieved without notable increases in computational cost, making the approach viable for industrial design cycles. Overall, the results indicate that data-driven corrections can generalize effectively when trained on sufficiently diverse datasets. The integration of data-driven techniques into RANS turbulence modelling marks a transformative step for aerodynamic simulation. By combining the efficiency of RANS with insights from high-fidelity data, this approach offers a practical solution to improving predictive accuracy without compromising industrial feasibility. This reduces reliance on costly physical testing and supports faster, more reliable design cycles. As aerodynamic analysis moves toward digital twins and automated optimization, data-driven turbulence modelling provides a bridge between traditional CFD and emerging design needs. Future work will extend the approach to additional training cases and explore more complex correction structures. The long-term objective is to develop a robust, interpretable, and efficient turbulence model suitable for realistic external-aerodynamics applications, enhancing the reliability of industrial CFD tools.

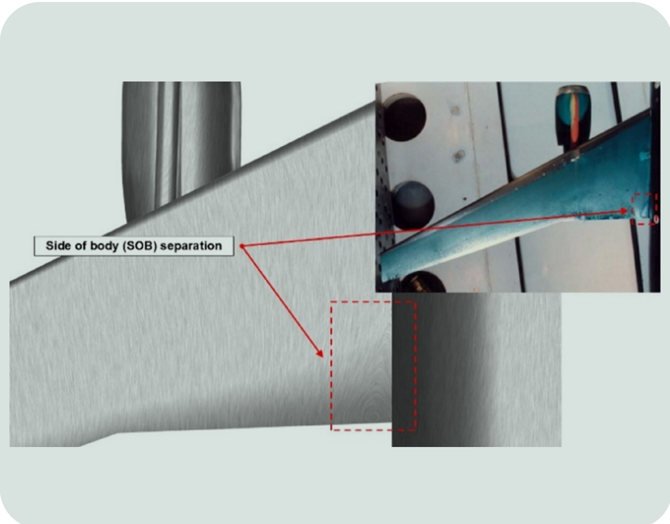

15 - Test case from HLPW5 [31, 32]. Qualitative comparison of experimental results (inset) against numerical results

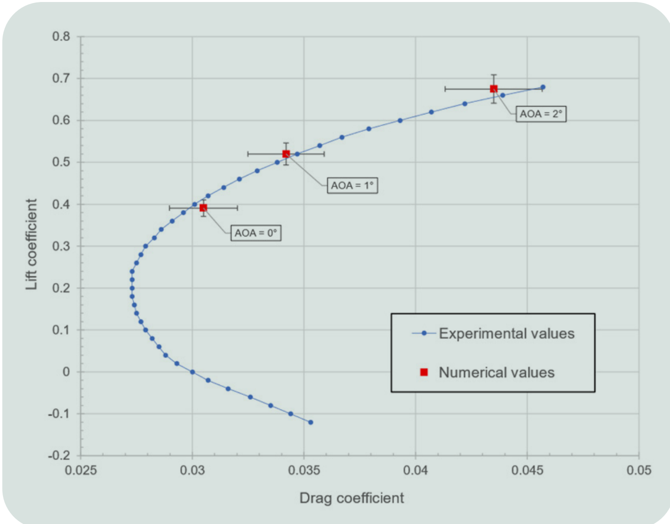

16 – Polar plot of the test case from HLPW5 [31, 32]. Quantitative comparison of experimental results against numerical results. The error bars of the numerical results were plotted using aleatoric and epistemic uncertainty quantification models

Application 5. An Adjoint-Based Optimization Method for Inverse Design and Wake Recognition

This application case presents a novel inverse design approach using adjoint-based optimization to reconstruct airfoil geometries based on wake velocity data. Unlike traditional methods that rely on surface pressure information, this technique utilizes velocity measurements taken downstream of the airfoil. Hereafter we introduce an inverse design methodology based on adjoint optimization, leveraging flow information collected far downstream from the controlled surface.
The adjoint-based inverse design approach, originally proposed by Jameson [33], demonstrated the ability to reshape airfoil geometries to achieve a prescribed pressure distribution. Building upon this foundation, the current study aims to reconstruct the geometry of an airfoil that generates a specific wake signature. To achieve this goal, wake velocity profiles are sampled at a fixed distance from the trailing edge using an array of probes. These measurements are then compared to a target wake profile generated by a reference airfoil. The difference between the measured and target velocities defines the objective function, while control points on the airfoil surface serve as the design variables.

By computing the gradient of the objective function with respect to these variables, the method guides a gradient-based optimizer, which iteratively refines the airfoil shape to minimize the wake mismatch. This approach enables the recovery of aerodynamic profiles based solely on wake characteristics, offering a powerful tool for flow-driven design in scenarios where direct surface data is unavailable or impractical.

Adjoint-equation method overview

For a detailed derivation of the adjoint equation problem, readers are referred to Giannakoglou et al. [34]. To achieve the inverse design, wake velocity profiles are sampled at a fixed distance from the trailing edge. These measurements are compared to a target wake profile generated by a reference airfoil. The discrepancy between the measured and target velocities defines the Objective Function (J), defined as the variance of the velocity in the present work, while selected control points on the airfoil surface serve as the design variables. To efficiently minimize this objective function, the adjoint-equation method is utilized to compute the gradient of J with respect to the design variables (the control points on the airfoil surface). By solving the corresponding adjoint governing equations, it is possible to obtain the necessary adjoint quantities, which are then used to calculate the desired gradients. This gradient is supplied to a gradient-based optimizer, which iteratively refines the airfoil shape to minimize the wake mismatch.

Discrete formulation and DAFoam

This work builds upon the discrete formulation of the adjoint problem, implemented within the DAFoam framework [35] and [36], which is freely available and widely adopted across both academia and industry. Among the two main approaches to adjoint-based optimization, discrete and continuous, the discrete method offers notable advantages. Most significantly, it provides higher accuracy in the adjoint field and ensures independence from the convergence level of the primal solution. Its primary drawback lies in the increased computational cost associated with gradient evaluation. However, in the present study, this limitation is mitigated by the relatively simple nature of the simulations, which involve feasible mesh sizes and can be executed efficiently on local computers. DAFoam offers a unified framework capable of solving both primal and adjoint equations. When coupled with OpenMDAO [37] and mesh morphing tools, it enables a fully integrated aerodynamics shape optimization (ASO) workflow.

Use case description

To evaluate the proposed methodology, a test case involving the identification of a 2D airfoil based on wake velocity data is conducted, as illustrated in Figure 17. Initially, a primal simulation is performed for a NACA 4-digit asymmetrical airfoil, and the resulting wake velocity field is recorded. A second simulation is then set up using a different geometry, specifically the NACA0012 airfoil, and the previously recorded wake data is used to define the objective function. The optimization process aims to minimize the discrepancy between the wake velocity profiles of the two cases. Convergence is achieved when the morphed NACA0012 airfoil closely matches the original target geometry. The accuracy of the reconstruction is assessed using the mean squared error metric. Given that the airfoil spans the domain x=(0,1), the y-coordinates of both profiles are resampled uniformly, allowing for a straightforward computation of the Mean Squared Error (MSE) and a quantitative evaluation of the optimization performance.

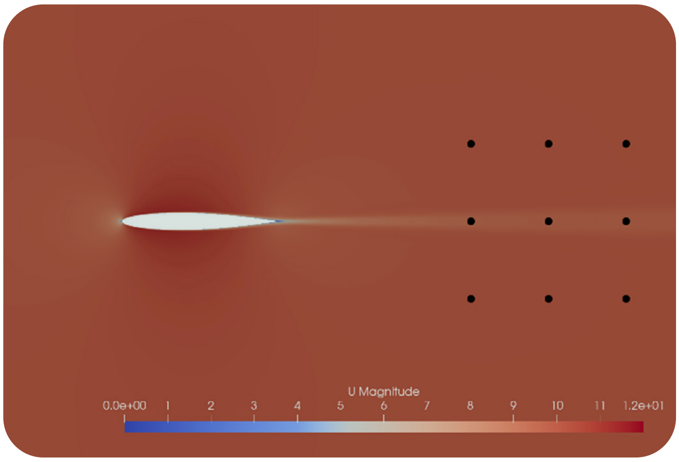

17–Velocity field around the airfoil and measurement of the wake velocity data (black dots in the wake behind the airfoil)

Results and Discussion

Hereafter, the adjoint-based inverse design methodology is demonstrated through the reconstruction of the NACA 6412 airfoil. In Figure 18 the plot shows the superimposed profiles of the target airfoil (red) and the optimized airfoil (blue) for the final iteration. In all the cases the process aimed to match the target airfoil's wake velocity signature calculated both with equal and different fluid conditions (freestream velocity and angle of incidence) as the starting one.

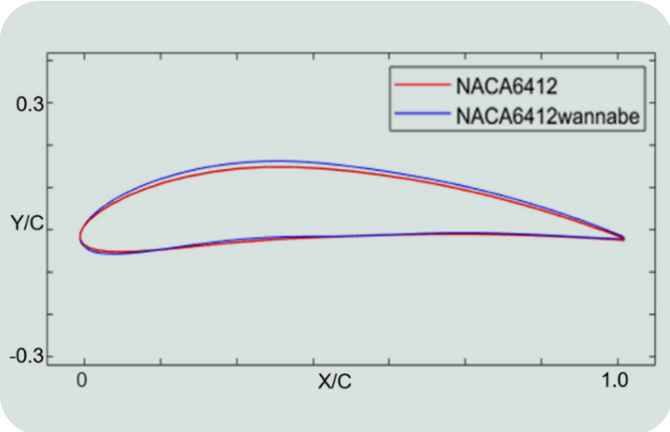

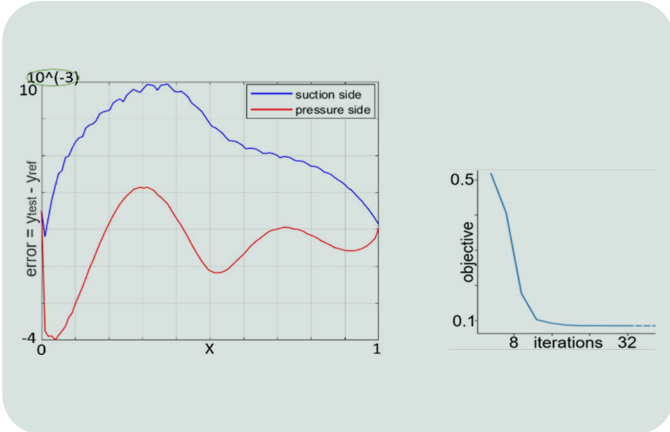

18 – Comparison between refence and optimized airfoil shape

19 – Error distribution along the airfoil chord (left), Variance function Trend (right)

Figure 19 provides a quantitative assessment of the final reconstruction accuracy. The suction side error (blue line) is generally larger and suggests that while the overall shape is reconstructed well, the fine features in the high-curvature areas are the most challenging to resolve perfectly from far-field wake data.

Summary

The proposed implementation demonstrates robust performance and promising initial results, successfully enabling the reconstruction of aerodynamic shapes based solely on wake flow information. The method proves effective in capturing key geometric features through indirect measurements, highlighting its potential as a powerful tool for inverse design. While the current tests validate its core functionality, further investigations are necessary to assess its limitations, scalability, and sensitivity to noise or measurement uncertainty. Despite these open questions, the approach represents a significant innovation in the field of aerodynamic optimization, offering broad applicability across both academic research and industrial design. Its versatility could be further enhanced by integrating complementary technologies such as image recognition, data assimilation techniques, or deep learning models, particularly for real-world applications involving complex geometries or limited sensor data. These extensions may pave the way for more intelligent, data-driven design frameworks capable of operating in diverse engineering environments.

CONCLUSIONS AND OUTLOOKS

Through five illustrative applications, we showcase the technologies, methodologies, and capabilities developed by the DTAS research group to advance multi-physics and multi-fidelity Digital Twins.

Our mission and vision are ambitious: to shape the evolution of this field by harnessing Leonardo's HPC resources and ensuring comprehensive coverage of the digital product lifecycle across all Leonardo production units.

Joel Guerrero: joelenrique.guerrerorivas@leonardo.com
Stefano Centomo – Application 1: stefano.centomo@leonardo.com
Damiano Paniccia – Application 2: damiano.paniccia@leonardo.com
Viola Rossano – Application 3: viola.rossano@leonardo.com
Tony Di Fabbio – Application 4: tony.difabbio@leonardo.com
Eric Segalerba – Application 5: eric.segalerba01@leonardo.com

REFERENCES

[1] S. Haag, R. Anderl. Digital twin – Proof of concept. Manufacturing Letters. 15: 64–66. January 2018.
[2] Aerospace Industries Association, LOTAR Best Practices, 2021. Available: https://www.aia-aerospace.org
[3] Object Management Group (OMG), Systems Modeling Language v1.0 Specification, 2022. Available: https://www.omg.org/sysml/sysmlv1/
[4] Object Management Group (OMG), Systems Modeling Language v2.0 Specification, 2022. Available: https://www.omg.org/spec/SysML/2.0
[5] Modelica Association, Functional Mock-up Interface (FMI) Standard Version 3.0, 2023; SSP Standardization Group, System Structure and Parameterization (SSP) 2.0 Specification, 2024. Available: https://fmi-standard.org, https://ssp-standard.org
[6] ISO 10303-242, Managed Model-Based 3D Engineering (AP242), 2014. Available: https://www.iso.org/standard/66654.html
[7] Object Management Group (OMG), Requirements Interchange Format (ReqIF) Specification, 2011. Available: https://www.omg.org/spec/ReqIF
[8] OASIS OSLC Core Technical Committee, OSLC Core 3.0 Specification, OASIS, 2017. Available: https://open-services.net
[9] Stevens and F. Lewis, Aircraft Control and Simulation. 2nd ed. Ho John Wiley & Sons, 2003
[10] P. Sharpe, NeuralFoil: An airfoil aerodynamics analysis tool using physics-informed machine learning, 2023. Available: https://github.com/peterdsharpe/NeuralFoil
[11] U.S. Standard Atmosphere. Washington, D.C.: U.S. Government Printing Office, 1976
[12] P.M. Sforza. Theory of Aerospace Propulsion. Elsevier Science, 2016
[13] Simulink, The Mathworks Inc. Available: https://it.mathworks.com/products/simulink.html
[14] Simcenter Amesim, Siemens. Available: https://plm.sw.siemens.com/it-IT/simcenter/systems-simulation/amesim/
[15] OpenModelica, Modelica Association. Available: https://openmodelica.org/
[16] J. Mihaloew, M. Ballin, and G. Ruttledge. Rotorcraft Flight-Propulsion Control Integration: An Eclectic Design Concept. Technical Report 2815. NASA, 1988.
[17] I. Fuksman and S. Sirica. "Real-Time Execution of a High Fidelity Aero- Thermodynamic Turbofan Engine Simulation". In: Journal of Engineering for Gas Turbines and Power 134.5 (Mar. 2012), p. 054501. issn: 0742-4795. doi: 10. 1115 / 1. 4004856.
[18] A. Volponi. "Gas Turbine Parameter Corrections". In: International Gas Turbine and Aeroengine Congress and Exhibition. 1998.
[19] T. Castiglione et al. "Linear Model of a Turboshaft Aero-Engine Including Components Degradation for Control-Oriented Applications". In: Energies 16.6 (Mar. 2023), p. 2634. issn: 1996-1073. doi: 10. 3390 / en16062634.
[20] H Badihi, A Shahriari, and A. Naghsh. "Artificial neural network application to fuel flow function for demanded jet engine performance". In: 2009 IEEE Aerospace conference. IEEE. 2009, pp. 1–7.
[21] S Sina Tayarani-Bathaie, ZN Sadough Vanini, and Khashayar Khorasani. "Dynamic neural network-based fault diagnosis of gas turbine engines". In: Neurocomputing 125 (2014), pp. 153–165.
[22] A Vatani, K Khorasani, and Nader Meskin. "Health monitoring and degradation prognostics in gas turbine engines using dynamic neural networks". In: Turbo Expo: Power for Land, Sea, and Air. Vol. 56758. American Society of Mechanical Engineers. 2015, V006T05A030.
[23] S. Brunton et al. "Data-driven aerospace engineering: reframing the industry with machine learning". In: Aiaa Journal 59.8 (2021), pp. 2820–2847.
[24] S. Brunton, J. Proctor, and J. Kutz. "Discovering governing equations from data by sparse identification of nonlinear dynamical systems". In: Proceedings of the national academy of sciences 113.15 (2016), pp. 3932–3937.
[25] M. Reissmann, Y. Fang, A. S. H. Ooi, and R. D. Sandberg, "Symbolic turbulence model development for complex-geometry flows exploiting language model-based transfer learning," Phys. Fluids, vol. 37, no. 9, p. 095180, Sep. 2025.
[26] T. Di Fabbio, Y. Fang, E. Tangermann, R. D. Sandberg, and M. Klein, "Strategies for Enhancing One-Equation Turbulence Model Predictions Using Gene-Expression Programming," Fluids, vol. 9, no. 8, p. 191, 2024.

[27] A. Ferrero, A. Iollo, and F. Larocca, “Field Inversion for Data-Augmented RANS Modelling in Turbomachinery Flows,” Comput. Fluids, vol. 201, p. 104474, 2020.
[28] D. Wilcox. Turbulence Modeling for CFD. DCW Industries, Third Edition, 2006.
[29] OpenFOAM website. [Online]. Available: https://www.openfoam.org/
[30] NASA’s Turbulence Modeling Resource website. [Online]. Available: https://tmbwg.github.io/turbmodels/
[31] 5th AIAA CFD High Lift Prediction Workshop (HLPW-5) website. [Online]. Available: https://hiliftpw-committee.github.io/HLPW/index-workshop5.html
[32] S. Mouton, G. Charpentier, and A. Lorenski, “Test Summary of the Full-Span High-Lift Common Research Model at the ONERA F1 Pressurized Low-Speed Wind Tunnel,” AIAA Paper No. 2023-0823, AIAA SCITECH Forum, Jan. 2023.
[33] A. Jameson. Aerodynamic Design via Control Theory. Journal of Scientific Computing 3(3), 1988.
[34] E. M. Papoutsis-Kiachagias, K. C. Giannakoglou. Continuous Adjoint Methods for Turbulent Flows, Applied to Shape and Topoly Optimization: Industrial Applications, Archives of Computational Methods in Engineering 23:255-299, 2016
[35] P. He, C. Mader, J. Martins, K. Maki. DAFoam: An open-source adjoint framework for multidisciplinary design optimization with OpenFOAM. AIAA Journal, 58:1304-1319, 2020.M. Wegmuller, J.P. von der Weid, P. Oberson and N. Gisin, “High resolution fiber distributed measurements with coherent OFDR”, in Proc. ECOC’00, 2000, paper 11.3.4, p. 109.
[36] P. He, C. Mader, J. Martins, K. Maki. An aerodynamic design optimization framework using a discrete adjoint approach with OpenFOAM. Computer & Fluids, 168:285-303, 2018.
[37] J. S. Gray, J. T. Hwang, J. R. R. A. Martins, K. T. Moore, and B. A. Naylor, “OpenMDAO: An Open-Source Framework for Multidisciplinary Design, Analysis, and Optimization,” Structural and Multidisciplinary Optimization, 2019.

ARTIC Project – A Cyber Attacks & Effects Emulation and Simulation Platform

Stefano Solari, Samuele Bertelli, Alessandro Oneto, Gianluca Ceccoli

Leonardo – Cyber & Security Solutions Division

The project Affordable, Reusable and Truly Interoperable Cyber Ranges (ARTIC), funded by the SERICS Foundation, adopts cloud and orchestration technologies for delivering services, designed to realize Cyber Digital Twins, i.e. cyberspace emulation & simulation environments that can integrate and extend into the cyber-physical domain. The project leverages the adoption of infrastructure-as-a-service and lightweight virtualization solutions to maximize the efficiency of designing, developing, executing, highly complex cyber scenarios, including the definition and the orchestration of cyber-attacks, and all the relevant data collection and analysis. The ARTIC solutions have been also enhanced by Leonardo CSSD Cyber Fast Prototyping Lab with a special component, to assure the interoperability of ARTIC with all simulation federation technologies (DSS, HLA), with cyber-physical systems emulation frameworks. Further Cyber Fast Prototyping Lab research is ongoing towards full automation of the deployment, configuration, cyber effect and events management process, through the use of Scenario Definition Language (SDL) frameworks.

INTRODUCTION

For many years, simulation technologies have been enabling the digital representation of real-world objects, systems, or processes. The most recent manifestation of simulation technologies integrated with advanced digital technologies, has produced the so-called "Digital Twins", i.e. technological solutions that allow for rapid simulation of different scenarios and prediction of potential issues. They help to optimize the performance of the asset (or process) while ensuring safety and keeping costs.

Digital Twins result from the wide and deep digitalization that is affecting every area of business and industry. The same concept has since made its way into various other sectors, such as the energy sector, construction/real estate, smart cities, healthcare, and also military platforms.

One of the most interesting applications of the Digital Twins concept is cybersecurity, which has been a constant focus of attention in recent years, due to the exponential increase in threats.

IT infrastructures are more complex than they once were, as they are distributed and hybrid by definition. The components of an IT ecosystem are multiple. They go ranging from cloud resources to network equipment, from traditional (monolithic) applications to cloud-native, including operating systems, middleware, and interconnection nodes. In many cases, a simple damage affecting one component is enough to impact the behaviour of the entire ecosystem in any sector that relies on multiple technological domains.

The real advantage, which underlies the relationship between Digital Twins and cybersecurity, is that all elements of both the IT and the various cyber physical ecosystems (typical of military C2 platforms and all tactical assets) produce data and information, making it possible to create digital models that, in response to certain actions, replicate the behaviour of systems and/or the system of systems as a whole.

This isn't simple or immediate but given respectively the national security or mission assurance implications for the military sector, and the multi-million-dollar costs for civilian market caused by data breaches and downtime, it's certainly worth it.
Having a virtual replica makes it possible to study the system's behaviour in the face of real cyberattacks, thus preventing them by modifying the reaction and interactions between components of the ecosystem. In practical terms, this means being able to model and execute a potentially infinite number of cyber threats, leveraging all the latest trends, techniques, and potential vulnerabilities, without impacting the production infrastructure in any way. The ARTIC project has focused on the design and construction of a framework derived from previous approaches, used by the so called "cyber range platforms", that satisfies properties of:

- "affordability", since the framework manages efficiently both large and small implementations;
- "extensibility", because it gives the possibility to dynamically modify its components;
- "interoperability", being it able to communicate in an easy way with other physical or logical components.

"IAAC" AS THE INITIAL APPROACH

The utilization of Infrastructure as a Code (IaaC) platforms with language-based orchestrators on top, is one of the most diffused approaches for the simulation of cyber scenarios that has been adopted in both academical and industrial communities. The Leonardo Cyber & Security Solutions Division has developed a technological stack for its "cyber range platform" based on the TOSCA language as its standard. Such approach adopts the Infrastructure as a Code paradigm to design, compose, and reuse full "cyber theatres of cyber operations," i.e., virtualized IT infrastructures for cyber procedures training and testing. In this approach, modelling is expressed with a TOSCA Simple Profile in YAML, treated as a declarative domain-specific language that represents applications and infrastructures as topologies of nodes and relationships with properties, requirements/capabilities, and lifecycle operations.

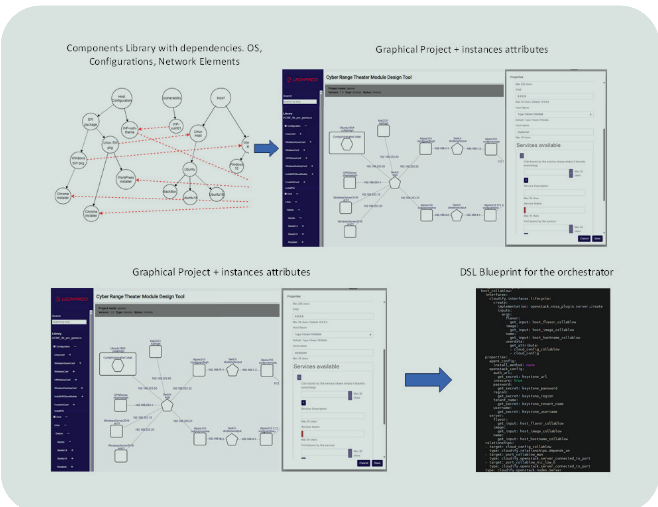

1 – IaaC based Cyber Scenario visual modelling

A TOSCA dialect interpreted by a TOSCA-compliant orchestrator enables automated deployment of complete environments on a virtualization platform via OpenStack technology as the IaaC layer, allowing orchestrators to instantiate normative node types portably, without any imperative deployment logic.
To avoid manual authoring of blueprints, a graphical interface has been provided on top of TOSCA. A curated library of components encodes hierarchy and dependency knowledge for operating systems, software installers, configuration bundles, virtual networks, firewalls, routers, switches, and their related elements. Components expose capabilities and requirements so that only valid connections are permitted. Designers compose topologies by dragging components onto a canvas, while each node exposes instance attributes - such as IP addressing, DNS, and service ports - together with validation rules that preserve consistency as scenarios are customized.
When a design is complete, the tool compiles the diagram into a valid TOSCA YAML blueprint for Cloudify. The orchestrator executes the blueprint to provision hosts, configures networking, installs software, and applies configurations automatically on the virtualization platform through OpenStack. As inputs are fully parameterized, the same blueprint is reusable across teams and exercises, enabling rapid creation of scenario variants by changing only IP plans, exposed ports, or credentials. This workflow - visual modelling with a dependency-aware library, automatic blueprint generation, and orchestration through Cloudify - turns the scenario design into a code that is versioned, testable, and repeatable, thus improving the governance and accelerating the preparation for training and testing activities.

THE "ARTIC" APPROACH

Starting from the IaaC background experience and competence of the Leonardo Cyber Fast Prototyping Lab team, the ARTIC framework enables the creation and management of cyber ranges that are affordable, extensible, and interoperable. Its architecture adopts a microservice-oriented and cloud-native paradigm, leveraging containerization and pods to ensure modularity, scalability, and dynamic adaptability. Each node in ARTIC is implemented as a pod consisting of multiple containers with distinct roles, named: Workload, Init, Sidecar, Management. This design allows flexible orchestration and runtime updates without any service disruption.

At the core of ARTIC also lies a Message Queue that serves as the communication backbone for publish/subscribe and RPC-based interactions among components. This decoupled messaging system supports orchestration, lifecycle management, and scoring, also ensuring resilience and scalability across distributed environments. The Orchestrator component manages the network creation and removal, pod lifecycle operations, and recovery mechanisms that restore dynamically added containers after failures, by replaying persisted RPC streams.

The dynamic extensibility is a key feature of ARTIC: both its pods (nodes) and sidecar containers can be added or removed at runtime by sending specific RPC commands to the message broker. This capability allows operators to modify the composition of a scenario without restarting the system, enabling real-time adaptation during exercises. For example, new monitoring agents or scoring modules can be injected into a running pod, or additional nodes can be instantiated to simulate network expansion or adversarial presence. Similarly, sidecars such as probes or agents can be created on demand to inspect or alter the Workload state, supporting advanced introspection and automation.

The Workload container represents the primary execution unit and can host lightweight services, full virtual machines (via QEMU/KVM for near-native performance and strong isolation), or complex elements such as simulated Internet environments.

The Init containers perform the pre-configuration tasks, including the network setup and cloud-init provisioning, before the main Workload starts.

The Sidecar containers extend Workload functionality with specialized roles: noVNC provides browser-based GUI access to VMs, DHCP handles dynamic IP assignment, Green manages VM lifecycle through QEMU Machine Protocol (QMP), and Sender Commands executes remote commands inside the VM using QEMU Guest Agent. The Management container enables dynamic creation or removal of containers within a pod using Podman, supporting real-time scenario evolution. ARTIC employs Object Storage (MinIO) to maintain per-node configuration files and cloud-init templates, enabling automated VM initialization and parameterized deployments. Image cloning with differential overlays reduces the memory footprint and facilitates the reuse across multiple nodes.

The Scenario definition relies on Makefiles and YAML templates, allowing declarative specification of network topology, IP addressing, and Workload parameters. This approach ensures consistency and accelerates scenario deployment. The framework integrates seamlessly with an external module called Automatic Attack Execution (AAE) that introduces red-team capabilities by deploying attack agents within pods. The Scoring Engine collects performance metrics and then it streams them to Logstash for visualization, while a Topology Modelling function automates the scenario generation from graph-based models, including network creation, object storage population, and Ansible configuration. Advanced features enhance realism and usability. Yellow nodes simulate user activity, automating interactions with browsers, office suites, and email clients. A simulated Internet emulates Internet-like connectivity and services within a single host, enriching realism of the training. Additionally, ARTIC supports the process evaluation by tracking incident response workflows against BPMN models using conformance checking, providing real-time dashboards and scoring.

By combining containerized design, declarative orchestration, and modular integrations, ARTIC delivers a versioned, testable, and reusable framework for cyber-attack realistic experiences. Its dynamic architecture supports rapid scenario customization, ability to perform federation across platforms, and automated provisioning of complex environments, accelerating preparation for training and testing activities.

THE ARCHITECTURE

Architectural pillars of the ARTIC approach are the provisioning of a reusable set of components, libraries, and standards for operations, the leveraging and integration of the benefits of modern microservice-oriented and loosely coupled architectures (e.g., cloud-native paradigms), and the leveraging of containerization by using pods and containers.

The architecture is basically composed by a Backend, a Message Queue and the Cyber Scenario Components. The Message Queue is the core element for communication within the architecture. This means that two elements must exchange a message using the message queue to communicate with each other. Each scenario component runs in a Pod (a collection of multiple, tightly coupled containers); it works as a "decorator" for an actual component (aka Workload). Workload is the only element accessible to the Blue Team.
Virtual Machine is a type of Workload that runs the VM images saved in the Object Storage using QEMU. Physical entity is a type of Workload that creates the connection with an external physical entity.

A Registry provides a catalogue of ready-to-use and configurable "init" containers, probes, and agents. Examples of available probes include those that sense the internet connectivity status, the emulated service status, and the integrity or presence of files. Examples of available agents are Ansible scripts executors, console access agents through web interfaces, user (yellow team) activities (GUI or script automation), attackers (red team) activity generators. The orchestrator element performs two different operations: network management and recovery. The network management is responsible for dynamically creating and deleting networks. The recovery element is responsible to perform the recovery operation. A Time Series DB is responsible for persistently storing the messages transmitted with the correct order. A dashboard element is responsible for showing the result of the score metrics. An Object Storage is responsible for storing the configuration files of every pod.

THE AUTOMATED CYBER ATTACK & DEFENCE EXECUTION PLATFORM

Using the Leonardo Cyber Range Automated Attack/Defence Execution (AA/DE) component whose internal architecture is shown here, cyber threats can be rendered against assets that can be internally virtualized, or even physically connected. Assets can range from ICT applications to 5G software-defined radio networks, radars systems, data links, or against large-scale reproduction of ICS/OT critical infrastructures. Such component has been interconnected as a part of the ARTIC framework. The Attack & Defence Automation Execution subsystem performs three main roles. It is a library of software tools available to perform attack and defence tasks, plus the required support to deploy them inside a training session. It allows the operator to design and deploy the required infrastructure to perform attack and defences within a cyber scenario and allows designing and executing automated attacks and defences by using the deployed infrastructure.
The AA/DE system is made of four software modules: an administration module and web application; a deployment module; an automatic execution master module; automatic execution agents' modules.
The Administration module and web application allows the operator to manage the whole system through a web application, and to run and monitor attacks and defences within a training session.
The Deployment module interacts with the Exercise Manager & Orchestration (EMO) system, to deploy attack and defence tools into a training session, being ready to perform attack and defences tasks. The deployment module is accessed through a REST API.
Automatic Execution Master module.
The Master module is the automatic attack and defences execution manager. It coordinates the overall execution and instructs the Agent module about actions to perform and collects the result of each action. It also sends data to the awareness and scoring subsystems. The Master module is implemented as a daemon running outside the training session environment. It communicates with Agent modules on a dedicate channel, reserved for this use.
The Automatic Execution Agent modules - Agent modules are the automatic attack and defences actual executors. They are a software application (a daemon under Linux, a service under Windows) installed on the machines performing the attack/defence tasks inside the training scenario and impersonate the attack/defence users. When needed, they can also impersonate generic users performing common tasks, for example opening a mail message and an attachment within it. Agents receive instructions from the Master on a dedicated channel and translate them into whatever syntax is needed, to perform a given task using a specific attack or defence tool. They collect the result of the operation and return it to the Master. Multiple Agents can be active inside a training session at any given moment (one on each attack/defence system participating to the session). The communication between Master and Agents is fully asynchronous, and each Agent can also perform multiple tasks at the same time. The AA/DE currently has a catalogue of multiple cyber threats. In the IT / ICT Domain, cyber-attacks patterns have been developed to realize:

- Common and advanced/zero-day malware, including ransomware, across different operating systems;
- Brute-force attacks;
- Data leakage/exfiltration;

- Botnet command-and-control communications;
- Client-side and server-side vulnerabilities and exploits;
- Spam, phishing, and spear-phishing attacks;
- Malicious domains and websites, including phishing websites;
- Denial-of-service attacks with multiple variants.

In the ICS/SCADA/IoT domain, cyber-attacks patterns have been developed to realize:
- Cyber-attacks as DoS against SCADA for Metro Transport Power Traction via host exploitation;
- APT / Command and Control over SCADA server;
- Man in the Middle over Ethernet IP (Rockwell Controllers);
- Man in the Middle over IEC 104;
- DoS Attack over MT Regional Distribution SCADA for Control Room;
- Schneider PLC application integrity violation and recovery; MQTT misconfiguration and MITM attacks.

Other domains available for cyber-attacks are:
- Attacks to Cloud Services (AWS metadata server misconfiguration exploits);
- Attacks via malicious Android App to IT Web server (Server Side Request Forgery);
- Various TELCO attacks to various RAN 5G / CORE 5G implementations;
- Multi tactical domain cyber-attacks and effects emulation (air, land, maritime, space, EMS);
- 1553B Avionic Bus Spoofing Tests.

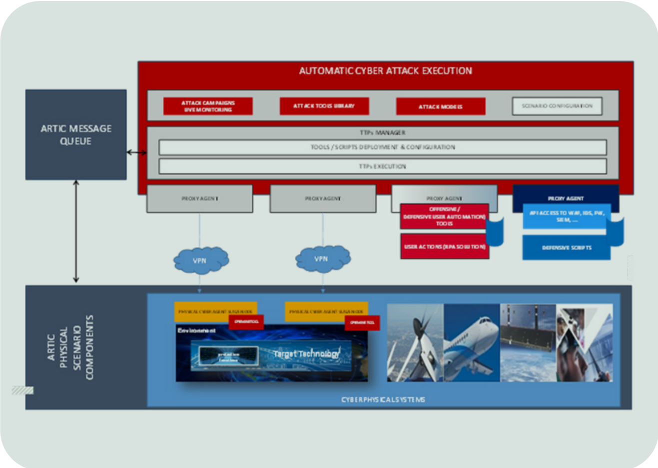

2–AA/DE Architecture for Cyber Physical Scenarios

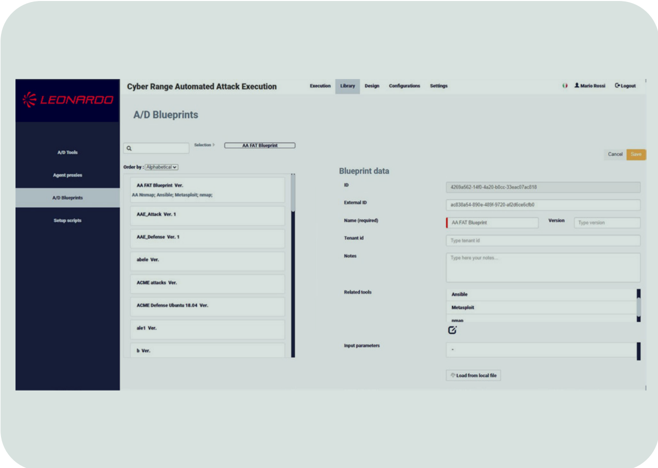

3–AA/DE User Interface

CYBER vs TRADITIONAL SIMULATION: THE ARTIC CYBER FEDERATION ADAPTER COMPONENT (CFA)

In the area of military operations simulation, the introduction of a cyber effect adaptation component which is compatible with various simulation frameworks that cover different level of details is key. This is true especially in face of the introduction of Cyberspace Operations into Multi Domain Operations scenarios, and when it is needed to validate their response in complex environmental condition (as in tactical tasks simulation sessions involving mission intelligence, recognition and execution tasks).

The Cyber Federation Adapter (CFA), as a bidirectional interoperability extension software component of the ARTIC Cyber emulation platform, might be developed with the aim of exposing its functionality towards a federated simulation. Its adoption would enable to simultaneously federate a desired number of cyber scenarios, to exchange events with multiple heterogeneous domains federates, including physical and cyber-physical simulators, visualization tools and hardware devices.

The solution adopted by the CFA to solve the known problem of semantic differences in integration among domains implements the SISO Cyber Data Exchange Model (DEM) standard, which uniquely defines messages of the *Cyber Objects* and *Cyber Events* types. The former ones represent the components of the infrastructure orchestrated by the range, while the latter represent the cyber actions undertaken during the scenario or their consequences on the cyber objects. The advantage of this solution derives also from the possibility of the dynamic configuration it offers, which allows to specify translation rules between the Cyber-DEM entities and the given Cyber Scenario API functionalities. In particular, Queries from the Cyber emulation side can be defined as the rules that detect cyber events, with the adapter publishing updates of the scenario status to the other simulation federates. From the CR side, the Cyber Procedures can be defined as the rules that interpret federation events to alter the state of the internal cyber scenario. In its first application, the adapter has bridged and enabled synchronized simulations from distinct domains.

SIMULATING THE EFFECT OF A CYBER ATTACK DURING AN AIR COMBAT MISSION

As a first example of application, the potential mix among physical, synthetic and cyber environments has been recently used in a demonstration that has simulated a cyber effect against an air combat platform (i.e. main fighter with controlled adjuncts), with the same detected cyber effect being reported to the pilot. Such an approach can enable accurate verification of adequacy of the logics devoted to the simultaneous management of a core fighter and of its various adjunct platforms, by evaluating the degree of autonomy of the latter and defining the characteristics necessary to avoid excessive pilot workload, thanks to the automatic selection of the most relevant information to present in each phase of the mission. The technological demonstration has been enabled by the collaboration of the Leonardo Cyber & Security Solutions Division's Cyber Fast Prototyping Lab (Genoa) with the Leonardo Aircraft Division's Simulation Battle LAB (in Turin). Thanks to such collaboration, the adjunct models have been refined with cyber elements provided by the CFA adaptation technology for translating cyber events, such as the simulated activation of dormant malware in an adjunct's electro-optical sensor in the Cyber emulation environment, into events useful for simulation of the overall tactical scenario. The activity has required a number of setup and preparation processes, like the integration of the Battle LAB TO network via the Leonardo Business Network, the preparation of the avionics bus emulation environment and its integration with the end system simulator (sensor) and cyber awareness modules, the preparation of the end system emulation environment (sensor) and its integration with AAE module, an avionics IDS simulation for traffic anomaly detection by the Cyber Agent software on board the sensor, and finally the simulation of the malware effect on the end system (injection of altered data onto the emulated avionics bus). The simulated scenario was a coordinated operation typical of a system of systems with allied and enemy forces, with the Objective to destroy the enemy command and control centre. The deployed Capabilities/Assets consisted of three formations called Packages, each being composed of an aircraft of the 6th-generation commanding four unmanned aircraft.

The unmanned aircraft were equipped on board with: electro-optical sensors and radar that provided information to the core platform (piloted aircraft) about the enemy infrastructure; an internal GPS detector; a cyber agent software boarded on the sensor's embedded operating system for the detection of cyber anomalies, with analysis and response capabilities. The Cyber Attack Target was the Electro-optical sensor of an unmanned aircraft. The attack trigger was an event generated by the remote activation of a malware via a rootkit embedded in the kernel of the sensor's embedded operating system. The local effect of the attack was the loss of capability/temporary operational use of the electro-optical sensor and therefore of the unmanned aircraft. In the developed main courses of action, the unmanned aircraft being no longer able to perform its task of monitoring the enemy C2, reports about its own degradation. The degradation is interpreted by the mission control as a possible "GPS denied" attack over the area of the unmanned aircraft. The mission is then rescheduled so that the manned aircraft that would release the missile could not enter the hypothetical GPS denied zone. The cyber agent software in the sensor's integrated operating system detects the altered data output from the sensor to an avionics virtual bus/link, identifies the presence of malware, identifies the corrupted file, generates a recovery hypothesis, evaluates it, and reports about the "ready to recover" status to the mission control.

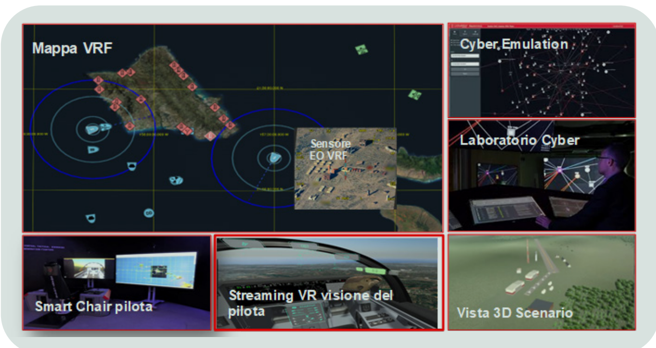

4 – Air Combat simulation with integrated Cyber Effects

SIMULATING THE EFFECT OF A CYBER ATTACK TO A WARSHIP

In this second example of application, a specific naval warfare cyber scenario with a mixed of integrated IT and OT/naval SCADA networks have been deployed by ARTIC services and has been hardened with a very high level of difficulty in penetrating networks and services. The Red Team evaluation rules have been defined, together with the configuration of additional KPIs and advanced functions, for evaluation and visibility of the possible attack effects. The preparation of the scenario and the experimentation has required design and/or modification of the exercise theatre of approximately 65 nodes, which included a simulated Internet (perimeter routers/firewalls), DMZ, Datacentre/Server Network, IoT, collaboration services, a simulated remote surveillance and security sensor network, a SIEM console, an IDS technology, and VPN connectivity. Such IT infrastructure was extended with a fully realistic naval SCADA systems with PLCs and simulated sensors and actuators of a warship.
The storyline had been focused on the crisis management phase within an imaginary “Dark Sea War”, where it had been assessed that a threat actor had laid several types of mines to jeopardize commercial routes, while preparing to adopt and realize a BLOCKSHIP as his maritime war tactic. Allied sea forces have planned to respond to the adversary threat mine laying, by sending units with MCM capabilities. The offensive cyber operation activity plan had been of injecting a malware in the SCADA network of a civilian CARGO SHIP crossing a contested sea strait.

The storyboard tells of the attacker sending a malicious payload via standard/fall back e-mail communication, requiring the Ship Officers to unwillingly execute the malware payload. When executed, the malware propagates and installs itself as a rootkit inside the naval SCADA server, reaching the steering and propulsion OT subsystems networks of the ship. It then instantiates a C2 communication with the remote attacker infrastructure and waits for a command to execute and wipe OT ship governance and control software.

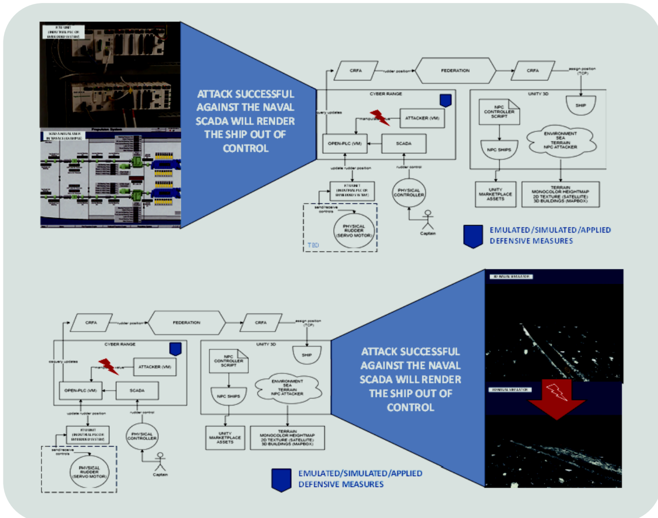

5 – Cyber OT attack effect on a Naval warship 3D simulator

SIMULATING THE EFFECT OF A CYBER SUPPLY ATTACK TO A CIVILIAN AIRPLANE

Cyber emulations scenarios can be connected to cyber-attack generation tools and procedures acting on target environments. Attacks and fault injections into the Supply chain can be realized by this platform against either virtual emulated environments or physical bus traffic.
In this third example of application, the storyboard tells that the Cyber Range Automated Attack/Defence Execution (AA/DE) module has been setup to use specialized tools for fault injection, and cyber exploits into an emulated (aircraft) scenario. Thanks to a sophisticated technology, the Cyber Federation Adapter (CFA) can intercept internal malicious events and can generate standard cyber events that are translated into safety critical effects that can thus be perceived by one or more federated flight simulation systems. In such a scenario, a widely known, open-source civilian aircraft flight simulation software has been federated to and ARTIC deployed scenario via the adapter (CFA). The adapter translates critical cyber conditions generated by the cyber-attack into effects that are “understandable” by external flight simulators. In a first run, a supply chain attack has been performed, using an instance of a virtualized avionic end system, without any specific cyber protection technology, causing an internal chain of both at rest and in transit violations of data integrity that result to be not correctly managed by the emulated avionic node and this affects its communication to the other bus terminal components. In this first run, a configuration file of the multi control display unit (MCDU) is altered by a malicious code that causes data integrity issues. The (enhanced) ARTIC platform monitoring and awareness dashboard shows the attack being in progress. The Cyber Federation Adapter then translates the ongoing data integrity violation into severe (extremized for the demo) availability issues on hydraulics of the airplane. Suddenly, in this example case the flight simulator features serious functional effects that affect the availability and thus the serviceability of critical hydraulic circuits of the simulated airplane. In the real world, the plane’s hydraulic energy provisioning is highly redundant in this type of platform, so the effect has been intentionally amplified in the simulation, for demonstration purposes.

A second run of the same attack is then performed, using a second replica of the target avionic end system, but this time, the Leonardo embedded Avionic Cyber Detection and Response (XDR) agent (LENS) is installed and actively runs in the same end-system that is the target of the attack. This second run is performed with Automated Attack/Defence Execution (AA/DE) using the same attack procedures of the first run, but this time the attack is performed against a replica of the avionic end system, where the XDR protection is active. The AA/DE executes the attack against the avionic end system application. The avionic cyber detection and response (XDR) agent detects the anomaly, as it can be seen from its own configuration and alarm analysis dashboard. While the attack is being executed in the Cyber Range, the cyber anomaly detection agent detects and classifies the anomaly, and the corresponding event appears on the ARTIC Dashboard. The avionic cyber detection and response (XDR) agent after the anomaly detection process, can highlight (via the CFA) through the pilot HMI of the airplane a running cyber alert condition. Triggered by the anomaly detection, via the CFA the Cyber Agent notifies the cyber alert condition.

Then, at the airplane pilot HMI level, the avionic cyber detection and response (XDR) agent responds and issues (again, via the CFA) a recovery notification message. Finally, the Cyber Agent neutralizes the threat and via the CFA notifies a recovery notification message to the pilot HMI.

6 – Cyber Supply Chain attack effect on a 3D Flight Simulator

SCENARIO DEFINITION LANGUAGE (SDL) AS THE NEAR FUTURE EVOLUTION FOR CYBER ATTACKS SIMULATION

The near future evolution for cyber-attacks simulation relies in the development of Scenario Definition Language (SDL) frameworks. A SDL framework improves the interoperability among cyber emulation environments and the re-use of scenarios. The intended SDL frameworks might be published as a reference under an open-source license for the use of several military actors, for developing and making their scenario to interoperate. The SDL component shall consider the features, capabilities and capacity of several existing federated cyber scenario emulators (cyber ranges), including their military use, network topology, services, workstations and servers with respect to operating systems (OSs), vendor patch levels and end-user applications, as well as their configurations. There are multiple target technologies that may benefit from the SDL, as this component allows various event types, (e.g., red team offensive training, blue team defensive training, blue team static digital forensic and investigation training), as well as evaluation and analysis of the training. It also allows for a structured presentation of incident response in playbooks.

The definition of natural language-driven scenarios offers a bridge between human intent and infrastructure-ready automation (IaaC/IaaS) for cyber emulators. The objective is of creating a unified pipeline from simple text to complete executable dynamic exercises. The current research is focusing on Virtual Scenario Description Language (VSDL) [1]-[3] that shall start from a meta-grammar that encodes domain concepts (nodes, networks, users, data, policies, and time) and compiles them into concrete automation artifacts for infrastructure deployment platforms. The core insight is that a human-centric, constraint-rich text can be parsed into a typed abstract syntax tree that captures both structure and semantics: computational resources, operating system families and versions, software inventories, connectivity, and firewall policies, among others. The VSDL supports comparative operators (equal, greater, not less, not greater), negations (is not), and temporal predicates (time-based rules) that control when the scenario elements change their state during the exercise (connections, updates, workload activation). Adopting VSDLs allows translating the model into orchestration blueprints (e.g., OpenStack images and flavors, TOSCA profiles, Ansible roles and playbooks, Terraform modules, Packer templates) with lifecycle hooks that provide hosts, configure services, and apply network rules. Another approach on cyber-attack and defence agents, complements the VSDL layer by introducing autonomous, role-driven entities capable of enacting adversarial and protective behaviours during exercises. Their attack–defence meta-model formalizes the interaction space: attacker goals, action primitives (scanning, exploitation, lateral movement, privilege escalation), defender controls (hardening, detection, containment, remediation), and the relationships tying actions to system states and observables.

Agents are integrated directly into the cyber range, embodying the red-team and blue-team roles that respond to scenario conditions. The value of this formalization lies in its operationalization: agents can reason over scenario graphs, capabilities, and vulnerabilities, select the best tactics in terms of reachability and payoff, and then produce traces that are useful both as training stimuli and as evaluation artifacts. Defence agents can monitor signals (IDS, SIEM, logs), match them against rules or learned patterns, and trigger controls such as isolating subnets, blocking ports, or rolling credentials. The meta-model's structure supports attaching arbitrary data to nodes and edges-credentials, exposure levels, attack surface attributes-which allow agents to operate on a rich semantic substrate. This approach enables repeatable, objective exercise scenarios involving comparable attack paths and defence responses.

CONCLUSION

The ARTIC project has paved the way to advanced research in the field of optimized, scalable, and declarative solutions based on cloud technologies and grey-box cyber-attack realization techniques. It has enabled the dynamic construction and deployment of automatic user and traffic simulation services, cyber-attacks scenario evaluation and composition, optimization of objectives, KPIs, participant access, and team role management. The platform, enriched with various artifacts and modules designed and developed by the Leonardo Cyber & Security Solutions Division Cyber Fast Prototyping Lab (CTO), enables the provisioning of federated military exercises setups towards Cyber Commands and Cyber Forces committed to conducting Cyberspace Operations, in support of the more traditional kinetic and/or recent non-kinetic (i.e. electro-magnetic, cognitive) domains. The solution also enables the generation of state-of-the-art cyber scenario contents supporting the integration of new domains (cloud technologies/applications, avionics architectures, critical terrestrial and space infrastructures), and it is going to be integrated soon in broader Scenario Definition Languages for complex Multi-Domain Operations, in which cyber-attacks must be taken into consideration as being part of the broader full-spectrum picture at both operational and tactical level.

Stefano Solari: stefano.solari@leonardo.com

REFERENCES

[1] Yamin, M. M., & Katt, B. (2022). Use of cyber attack and defense agents in cyber ranges: A case study."Virtual Scenario Description Language (VSDL)
[2] Costa, G., Russo, E., & Armando, A. (2023). Automating the Generation of Cyber Range Virtual Scenarios with VSDL
[3] Costa, G., Russo, E., & Armando, A. (2020). Building Next Generation Cyber Ranges with CRACK

The Leonardo Adaptive Training Ecosystem

Viola Mangiantini, Gianfranco Bellora
Leonardo - Aeronautics Division

The rapid evolution of operational scenarios underscores the pressing need for rethinking the military pilot training. Traditional training methods are often characterized by a one-size-fits-all approach. Thus, they might be not sufficient to fulfil the evolving demands and to face complexities of the modern aerial warfare. Given the need for development of more specialized skills in a variety of industries and for faster preparation times, organizations begin to pursue more adaptive and personalized learning. Artificial Intelligence can support this process by analysing the behaviour, identifying weaknesses and providing targeted feedback throughout the training continuum. The Leonardo Adaptive Training ecosystem aims to develop a machine learning–based approach applicable to all ground-based training devices, including computer-based training (CBT), scenario-based training (SBT), flight training devices (FTD) and full mission simulators (FMS). These platforms generate large volumes of performance data, the AI can process to detect patterns and deliver personalized recommendations. This paper describes how this approach benefits both instructors and students.

INTRODUCTION

In recent years, given the current geopolitical landscape, nations are under increasing push to maintain a ready and well-trained air force. Making pilots to reach faster the "ready to combat" level, while making their training more efficient and cost-effective, has become a strategic imperative. These developments have revealed limitations of the traditional training systems, which are no longer sufficient to meet the urgent demand for operationally ready pilots within the required timeframes and budgetary constraints.

As mission complexity increases and operational requirements evolve rapidly, there is no other way than rethinking the way pilots are trained. Conventional "one-size-fits-all" approaches that were effective in the past, are no longer adequate to ensure optimal learning outcomes for every individual [1]. Different trainees have different needs, learning speeds, and cognitive profiles. Therefore, what suits well to a student might be not suitable to another [2].

In this context, Adaptive Training emerges as an innovative instructional paradigm: a personalized, data-driven approach that dynamically adapts content, difficulty, and feedback of each training session according to the trainee's performance and cognitive state. By aligning training interventions to individual learning trajectories, Adaptive Training aims to optimize engagement, accelerate skill acquisition, and increase readiness of pilots.

Unlike the traditional training model, Adaptive Training adopts the concept of Optimized Learning, which is a data-driven strategy that leverages the vast and growing volume of information generated by digital training environments. The use of artificial intelligence algorithms and machine learning techniques enables to continuously assess the student's performance, identify skill gaps, and deliver tailored interventions aligned with each trainee's specific learning needs and progression rate.

The principle of maintaining the learner within the so-called Zone of Proximal Development (ZPD) is the key-feature of such approach.

The Zone of Proximal Development is the optimal range of difficulty whose tasks are just beyond the individual's current level of competence, but are still achievable with proper guidance.

Training within this zone has demonstrated to maximize the cognitive engagement, motivation, and learning efficiency of trainees. If the task were too easy, the trainee might feel disengaged due to lack of challenge.

If the task were too difficult, the trainee would likely to become frustrated and underperforming. Adaptive Training constantly calibrates the level of challenge, to ensure that each student is working at his/her most effective intensity according to his/her learning curve. In this way, the training process is optimized for delivering the right content at the right level at the right time, so ensuring it is neither overloading nor insufficient.
To operationalize this paradigm, a suite of AI-based tools has been developed and integrated into ground-based training systems. These tools are designed to provide objective and consistent evaluations of student's performance, based on a broad range of measurable inputs. By analysing telemetry, control inputs, and behavioural patterns, the system is able to detect learning trends, identify recurring errors, and generate personalized mission scenarios and targeted feedback.
Within modern training environments based on VR, XR, these adaptive systems leverage real-time data streams, such as eye-tracking metrics, to continuously monitor the pilot skill acquisition. The use of immersive technologies further enhances the realism and cognitive engagement of training sessions, while the AI models enable adjustments to scenario difficulty and instructional content based on the individual pilot's evolving performance profile. Machine learning algorithms are at the core of the adaptive training framework. Their primary function is to process high-dimensional behavioural and performance data to identify recurring patterns, detect skill gaps, and model individual learning trajectories over time. These insights allow the system to go beyond the static instruction, by generating dynamic training pathways that in real time adapt to the trainee's progress.
Based on the data collected during each training session, the system can automatically propose:

- Targeted feedback aimed at addressing specific weaknesses;
- Scenarios designed to strengthen underdeveloped competencies;
- Adaptive CGFs that evolve according to the student's proficiency.

This continuous adaptation process ensures that training remains within an optimal learning zone, by maintaining a balance between cognitive load and achievable challenge, which enhances the effectiveness of each session, and also accelerates the overall training efficiency and the mission readiness.
These developments mark the emerging of a new generation of intelligent training ecosystems, that are modular, interoperable, and capable of supporting instructors in designing and delivering more flexible, effective, and relevant training experiences. This transformation represents an innovation in training methodology and a concrete contribution to safety and sustainability of the future air operations.
Against this backdrop, the Italian Air Force is a concrete context in which the Adaptive Training methodologies can generate significant benefits. With a structured pilot training syllabus already in place, the challenge is how to evolve the existing training phases with data-driven tools and intelligent systems.

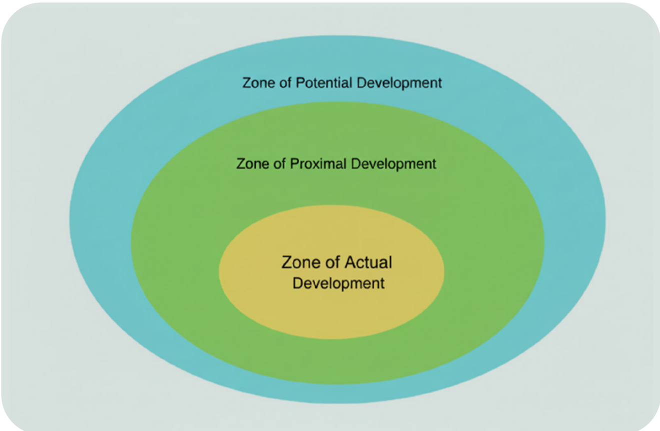

1 – Zone of Proximal Development

THE ITALIAN MILITARY AIR FORCE USE CASE

The Aeronautica Militare Italiana (Italian Air Force) applies a rigorous and structured pilot training model that is articulated in progressive phases of increasing complexity [3]. The training sequence begins with an initial selection process, after which pilots enter the Phase II that is a standardized stage attended by whom have successfully passed the access tests. At the end of this phase, every trainees are assigned to their future operational track.
Phase III represents a key instructional milestone. It consolidates the basic competencies acquired in the earliest stages, by combining academic instruction to practical flight training. Successful graduates are awarded the Military Pilot License (MPL) that formally qualifies them to serve as military pilots within the Italian Air Force.
Phase IV is also known as the "Lead-In to Fighter Training" program. It is specifically tailored to prepare pilots for integration into frontline operational Squadrons. This phase includes advanced training modules that simulate realistic mission profiles and operational environments, and represents a critical target for the introduction of Adaptive Training methodologies and technologies.

THE LEONARDO AERONAUTICS ADAPTIVE TRAINING ECOSYSTEM

Building upon the foundation of the Italian Air Force training pipeline, the Adaptive Training framework developed by Leonardo is designed to enhance Phase III and Phase IV training through the integration of advanced technological capabilities and intelligent evaluation tools. Its objective is to deliver a modular, data-driven training ecosystem that complements the existing tools with AI-enabled adaptivity, without replacing the essential role of the human instructor.

To support such a process, multiple training devices are employed that go ranging from e-learning platforms and procedural trainers to the Full Mission Simulator (FMS) of the M-346 aircraft. Each of these systems share data with a centralized evaluation engine that powers the adaptive features of the ecosystem.

ADVANCED ANALYTICS

The large volume of training data collected by the training device enables structuring personalized learning paths that offer more than just isolated performance evaluations. By analysing how each student progresses throughout manoeuvres and flight phases, the system identifies patterns of strengths and recurring difficulties, as well as atypical learning trajectories. This allows adapting dynamically the training intensity and its sequence, thus ensuring that each pilot spends longer time on his own critical skill gaps than on tasks that have already been mastered.

Clustering plays a central role in this process, by grouping students according to similarities in their performance profiles. Unlike rigid grading systems, clustering does not aim to assign a score, but rather to highlight natural patterns. For instance, it can reveal groups of students who have difficulty with a specific manoeuvre, or clusters of pilots who acquire certain skills more quickly than others do. These insights help instructors to anticipate learning needs and to design targeted interventions, while allowing the system to recommend differentiated learning paths for each cluster.

In addition to clustering, analysing the distribution of grades across missions provides another important perspective. The figure below shows how grades are spread for each mission, highlighting where students tend to perform better and where difficulties are more frequent. Such information allows instructors to understand which parts of the training program are more demanding and to better anticipate when additional support may be needed. Looking at data from many students rather than from one student, enables to see common problems and repeated patterns. For instance, if multiple students struggle with a particular manoeuvre, it suggests that additional practice or a different explanation is necessary. Data analysis can also identify early signs of slow progress or unusual behaviour, enabling instructors to provide timely assistance. Using data in this way creates a cycle of constant improvement, as it enables instructors to guide each student more effectively and, in case the same issues would arise for multiple students, provides valuable insights for adjusting the training program.

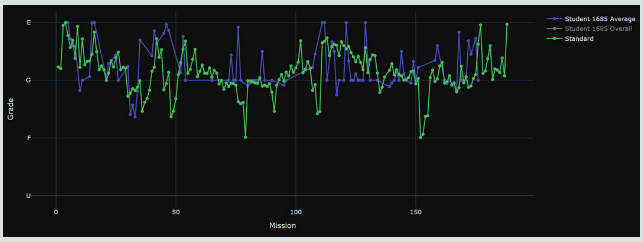

2 – Student Progress

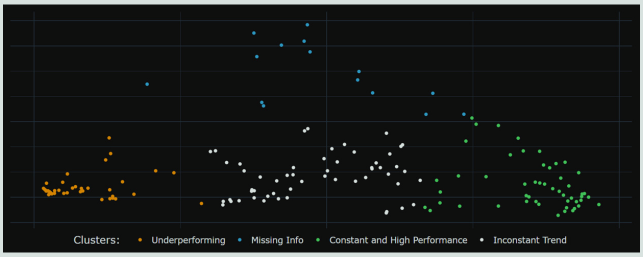

3 – Students Clustering

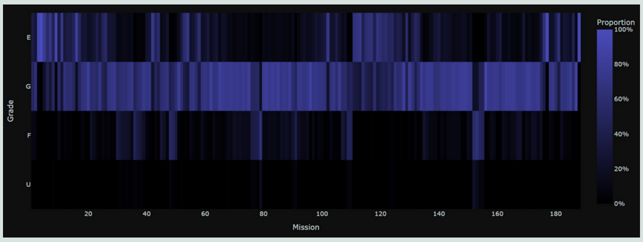

4 – Grade distribution per mission

AI-SUPPORTED EVALUATION

The core of the adaptive framework is a robust evaluation infrastructure, co-designed with Air Force instructors, which mimics the expert evaluation through transparent and validated artificial intelligence models. During each training session, data are collected at a frequency of 20 Hz, producing thousands of input variables per mission. These data points are categorized into:

- Flight parameters (e.g., angle of attack, pitch, roll, speed);
- Pilot control inputs (e.g., flap lever position, brake pedal use);
- System status and malfunctions (e.g., engine failure, degraded systems);
- Environmental conditions (e.g., weather, visibility).

Machine learning models process those variables to perform the mission analysis. A Random Forest Classifier is used for flight phase classification, incorporating feature selection and dimensionality reduction, via the Principal Component Analysis (PCA), to improve interpretability and performance. The pilot performance is then assessed by using a Gaussian Mixture Model (GMM) to detect anomalies and deviations from standard patterns. This is combined with statistical norms that define acceptable performance standards for each manoeuvre or flight phase. A weighted linear grading model turns such performance metrics into manoeuvre scores and the Bayesian Inference is applied to infer latent competencies from observable behaviours, linking the task execution to its relevant deeper skill acquisition. At the core of this ecosystem is an intuitive Human-Machine Interface (HMI) designed to support instructors in evaluation and debriefing.

The system displays flight data, segmented into significant sub-phases, for example:
Takeoff:

- Track alignment;
- Rotation;
- Climb;

Landing:

- Pre-flare;
- Flare;
- Touchdown.

Each segment is scored independently, and plots demonstrate deviations from expected parameters, delays in control inputs and anomaly patterns.

These insights allow instructors to focus on technically critical moments, enhancing the precision of feedback. Importantly, instructors maintain full authority over the final evaluation: while AI proposes scores and feedback, instructors can override or annotate any automated
assessment based on operational judgment and contextual factors.
This closed loop of human-AI collaboration ensures the system evolves with continuous supervision. As instructors validate or correct AI-generated outputs, the models learn from each interaction and gradually improve their own accuracy. From this perspective, data are the critical enabler of the entire adaptive framework. The continuous collection and curation of flight records allows the system to refine the existing models and to expand their scope, to include more complex manoeuvres and operational contexts. The richer and more diverse the data, the more effectively the framework can personalize training, identify early signs of proficiency or of difficulty, and provide actionable insights to support the improvement of each pilot.

5 – Adaptive Training HMI

6 – AI Feedback

COGNITIVE ASSESSMENT

The modular nature of the system also allows for the integration of additional biometric data streams, such as the eye-tracking embedded in VR headsets. These sensors can measure fixation duration, saccadic movements, and blink rate.

Recent scientific literature [4] has demonstrated that these physiological indicators are a valuable indicator of mental workload, which is a key component of pilot's performance that otherwise would be difficult to quantify. For example, prolonged fixation on a non-critical instrument may indicate cognitive tunnelling, and rapid saccades with inconsistent fixation duration may reflect information overload or loss of focus. An increased blink rate has also been associated with fatigue and cognitive strain [5].

When integrated into the Adaptive Training framework, these biometric signals are processed by using machine learning models to generate real-time estimations of the cognitive load.

These estimates can inform dynamic adjustments to training scenarios, such as modulating the mission complexity in response to elevated workload, or may guide post-mission debriefings, helping instructors to address not only what errors occurred but also why they occurred from the cognitive perspective. This approach extends the evaluation framework beyond the traditional task performance metrics, as it offers a more human-centred view of effectiveness of the pilot training.

7–Cognitive Assessment HMI

ADAPTIVE CGF

A key component of the Adaptive Training Ecosystem is the Adaptive Computer-Generated Force (CGF), an AI-driven module that controls the behaviour of friendly and adversary units during simulation-based training.

Unlike traditional CGFs, which follow pre-programmed behaviours, the Adaptive CGF continuously monitors the trainee's performance, cognitive load, and progression through the training mission. Based on this analysis, the CGF adapts its tactics, aggressiveness, and coordination levels, to keep the simulation within an optimal challenge zone, also avoiding any under-stimulation, which could lead to boredom and inefficiency, and any excessive difficulty, which could lead to frustration or failure [6][7].

For instance, if a trainee demonstrates a low cognitive load, the CGF may escalate the threat level by employing more sophisticated enemy manoeuvres or by coordinating multiple adversaries. Conversely, if the trainee struggles to maintain his situational awareness or if his task performance degrades, the CGF reduces the complexity to give the trainee more opportunity to focus on skill consolidation.

The adaptive CGF transforms simulations from static training tools into responsive and intelligent partners that can adapt to each trainee's evolving competence, in real time. This allows for targeted practice under adaptive pressure, which is a recognised accelerator of expertise development in demanding fields such as combat aviation.

ADAPTIVE SCENARIO

Complementing the real-time adaptivity of the CGF, the Adaptive Scenario engine operates at a broader instructional level, and acts as a recommender system for instructors during the planning and configuration of training missions.

8 – Adaptive Scenario (This image is a conceptual illustration generated with the aid of artificial intelligence. It does not represent a specific product or service and is not affiliated with any registered trademarks or commercial entities.)

Using historical performance data, syllabus objectives, and group benchmarks, the system suggests the most appropriate mission profile for each trainee at a given point along his progression. This ensures that the selected scenario addresses the learner's current weaknesses and his readiness for new challenges.
Importantly, the Adaptive Scenario module does not rigidly follow a one-size-fits-all syllabus timeline. It aligns with the same terminal objectives, allowing for flexibility in pacing and sequence. This enables the most proficient trainees to accelerate their progression and allows the others to receive additional exposure to targeted learning objectives.
Instructors remain in full control and can accept, modify, or reject the suggestions from the system, whose recommendations are supported by statistic insights from individual and collective pilot trajectories, which offer valuable assistance in customizing the learning path.
In brief, the Adaptive Scenario component shifts the focus from time-based training to competency-based progression, ensuring that pilots are effectively trained basing on demonstrated skills, instead of just on schedule.

ETHICAL CONSIDERATION AND PRIVACY

The integration of adaptive training systems based on artificial intelligence, biometric monitoring, and advanced analytics inevitably raises ethical and privacy concerns. On the one hand, the continuous collection of sensitive data enables personalization and efficiency. On the other hand, it requires robust safeguards to ensure data protection, transparency, and fairness. Ethical considerations include avoiding bias in algorithmic assessments, ensuring that AI-based evaluations remain supportive, and maintaining the central role of human instructors in the final decision-making. Privacy must also be addressed through strict data governance, anonymization protocols, and compliance with existing regulations. By embedding these principles, the adaptive training ecosystem can maintain compliance, protect individual rights, and provide instructors and trainees with a reliable tool for daily operations.

CONCLUSIONS

The development of the Adaptive Training ecosystem is a significant step in modernizing military pilot training. By integrating cutting-edge AI, real-time cognitive assessment, adaptive scenario generation, intelligent CGF, and advanced analytics, we create an innovative framework that addresses a critical operational challenge, by accelerating skill acquisition and maintaining rigorous safety, performance, and mission readiness.
This achievement stands out for its technological sophistication, as well as for how it has been grounded in rigorous methodology and operational reality. The result is a scientifically robust and mission-ready framework.

9 – Adaptive Training Ecosystem (This image is a conceptual illustration generated with the aid of artificial intelligence. It does not represent a specific product or service and is not affiliated with any registered trademarks or commercial entities.)

The integration of advanced algorithms, human-centred design, and close collaboration with experts has resulted in a system that can bridge the gap between research excellence and field applicability.
However, this is not the final picture. Let's think of the paradigm we have conceived and are implementing, as it were an evolving puzzle: several key pieces are in place, revealing a coherent and functional image, while some others are still being developed and refined. Each new component completes the design and enhances the performance, adaptability, and potential of the whole system, unlocking capabilities that today are only emerging on the horizon.

The vision, expertise, and ambition that have brought us this far will ensure that the next iterations complete the puzzle and set new global benchmarks for training excellence. The foundation is solid, and the evolution has only just begun.

Viola Mangiantini: viola.mangiantini@leonardo.com
Gianfranco Bellora: gianfranco.bellora@leonardo.com

REFERENCES

[1] Area9 Lyceum. (2023). How Adaptive Learning Optimizes Aviation Training. Retrieved from https://blog.area9lyceum.com/how-adaptive-learning-optimizes-aviation-training/
[2] Fuchs, S., Carpenter, A., Carroll, M., & Hale, K. (2011). A Hierarchical Adaptation Framework for Adaptive Training Systems. In D. D. Schmorrow & C. M. Fidopiastis (Eds.), Foundations of Augmented Cognition. Directing the Future of Adaptive Systems. Springer Berlin Heidelberg. https://books.google.mu/books?id=sXDtPB_7RAsC&hl=it&source=gbs_navlinks_s
[3] Schema Addestramento Volativo, https://www.accademiaaeronautica.it/addestramento-volativo-schema/
[4] Babu MD, JeevithaShree DV, Prabhakar G, Saluja KPS, Pashilkar A, Biswas P., Estimating Pilots' Cognitive Load from Ocular Parameters Through Simulation and In-Flight Studies https://pubmed.ncbi.nlm.nih.gov/33828735/
[5] Stern, J. A., Boyer, D., & Schroeder, D. (1994). Blink Rate: A Possible Measure of Fatigue. Human Factors: The Journal of the Human Factors and Ergonomics Society, 36(2), 285-297. https://doi.org/10.1177/001872089403600209
[6] Roessingh, J., van den Bosch, K., & Hofstede, G. J. (2016). Modeling behavior of Computer Generated Forces with Machine Learning Techniques: The NATO Task Group approach. NATO Modelling & Simulation Group. Retrieved from https://www.researchgate.net/publication/311486968_Modeling_behavior_of_Computer_Generated_Forces_with_Machine_Learning_Techniques_the_NATO_Task_Group_approach
[7] Reinisch, C., Bonner, M., & Vessella, J. (2021). Are We Machine Learning Yet? Computer Generated Forces with Learning Capabilities in Military Simulation. In Proceedings of the Interservice/Industry Training, Simulation, and Education Conference (I/ITSEC). Retrieved from https://www.researchgate.net/publication/356756043_Are_We_Machine_Learning_Yet_Computer_Generated_Forces_with_Learning_Capabilities_in_Military_Simulation

Battlefield Reconstruction Real Time

Adriano Tumino[1], Davide Bertelli[2], Luigi Rapetta[2]

[1]Leonardo – Electronics Division, [2]Alten Italia S.p.A.

Battlefield Reconstruction Real Time integrates advanced algorithms such as *Structure from Motion* (SfM) [1] [2] and *Gaussian Splatting* [3] to agnostically process a video data and turn it into high-fidelity, photorealistic digital environments. Primary objective of the system is to generate a *digital twin* of the recorded subject, within a predictable time window, while preserving its geometry, colour, and luminance properties. Although such reconstruction does not produce any interactive 3D mesh, the resulting model comes useful as a highly valuable asset for manual inspection and post-processing. Due to the substantial computational load and the need for deterministic processing times, the system exploits *High Performance Computing* (HPC) architectures, such as the Davinci-1 supercomputer, to reduce the total computation time and enhance accuracy of the reconstruction. The resulting point cloud can then be imported into advanced viewers and simulators, to take full advantage of its achieved photorealism.

INTRODUCTION

In recent decades, per-capita multimedia data production has been growing exponentially. Among the various data types, videos represent the most complex and informative resource, being them highly time-correlated and capable of capturing extremely heterogeneous subjects.

By leveraging such data, it is now possible to reconstruct environments and objects with remarkably high levels of detail. In particular, algorithms such as SfM and Gaussian Splatting enable indirect estimation of the scene geometry, while preserving its colour and luminance characteristics.

The outcome is a synthesized point cloud encoding the properties of the scene elements, suitable for use in simulation systems, game engines, and image generation frameworks, delivering photorealistic, high-fidelity reconstructions.

Battlefield Reconstruction Real Time uses these algorithms to generate a 3D copy of the scene within predictable time constraints, while preserving the chromatic, luminous, and reflective properties of the recorded environment.

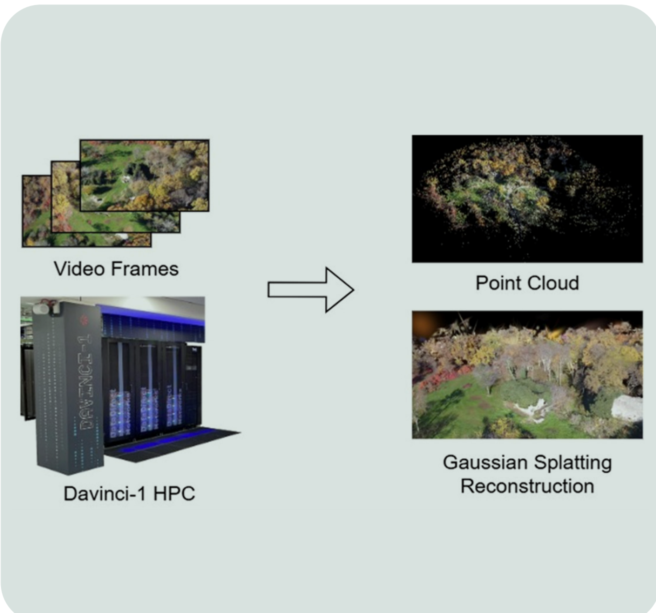

1-Battlefield Reconstruction Real Time Workflow

WORKFLOW

The process consists of multiple stages: video acquisition, frame extraction, SfM processing, Gaussian Splatting and scene visualization.

Data Acquisition

Reconstruction begins with the acquisition of raw video footage. This can be carried out by using consumer-grade devices - commercial drones, action cameras, or smartphones - or structured streams such as UDP or STANAG 4609, which may offer varying levels of quality and, in some cases, embedded geolocation metadata. Orbital camera trajectories around the subject yield the most accurate reconstructions. As shown in Figure 2 - Reconstruction results, from the left, at 25, 50 and 90 meters of altitude, the radius and altitude of the orbit directly influence the feature detection and detail preservation: distant objects appear with reduced texture sharpness, as high-frequency details are compressed into fewer pixels. Lighting conditions are crucial: strong illumination improves contrast and feature separability, while darker scenes compress pixel variability, making detection more challenging. Frames are then sampled to form the dataset for SfM. Video frames are highly time-correlated, thus exhibiting minimal visual changes among sequential images. Using all frames would unnecessarily inflate computational cost. When available, geolocation data (e.g., STANAG 4609) can assist both feature extraction and spatial alignment of the reconstructed scene within the viewer.

2-Reconstruction results, from the left, at 25, 50 and 90 meters of altitude

Structure from Motion

The reduced dataset is processed by using SfM, a photogrammetric technique that reconstructs a 3D structure from multiple images. It compares shared features between frames, to estimate scene geometry.
Outputs of the SfM:

- preliminary point cloud of the scene;
- the mathematical camera model for each frame: i.e. the estimated pose;
- refined subset of frames most relevant for matching.

A key advantage is that SfM is completely agnostic to the acquisition device, as it infers the camera position and orientation without requiring any prior hardware calibration.
The processing time increases exponentially along with the number of images, due to the *feature matching step*, where many frames are compared against one another. Therefore, frame sampling is essential for predictable execution times.

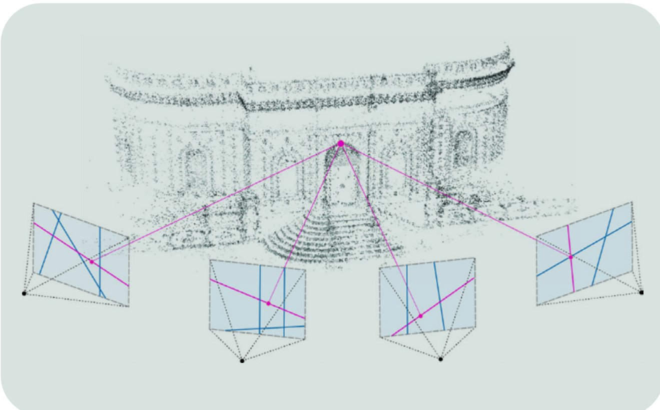

3-Structure-from-Motion example showing the tracking of a specific feature across the camera's different viewpoints

4 -Comparison between photogrammetry (left) and Gaussian Splatting reconstruction (right)

Gaussian Splatting

Once camera models, point cloud, and relevant frames are made available, the Gaussian Splatting is applied to recover colour, luminance, and reflectance.
Each point in the cloud is associated with a 3D Gaussian distribution, modelling the local light behaviour. An iterative optimization process adjusts the shape, orientation, and volume of the Gaussians to match the original frames, leveraging the differentiability of the representation. The algorithm progressively densifies the set of points to improve the visual fidelity.

Scene Visualization

The Gaussian Splatting requires an advanced rendering pipeline [4]: each *splat* is a semi-transparent 3D object, and all splats contribute additively to the final image.
Typical steps include:

- sorting splats by distance from the camera, for proper transparency handling;
- applying rotations and scaling;
- projecting the 3D Gaussian into a 2D Gaussian via billboarding.

Projection relies on the centroid and covariance matrix - derived from the scale and rotation matrices - to generate an elliptical footprint that represents the 3D distribution on screen.

OUTCOMES

Tests

Extensive tests have been conducted by varying video resolution, camera motion, and lighting conditions. Key findings include:

- orbital trajectories produce the most accurate reconstructions - vertical or linear movements often result in noisy, overlapping point clouds;
- low resolutions (e.g., 480×720) cause blurred reconstructions and sparse Gaussian distributions;
- higher resolutions produce better details but greatly increase the processing time:
  - 1600×1200 videos take ~20 minutes;
  - 4K videos take up to ~1 hour.

A good balance between detail and performance is achieved with 1920×1080 @ 30 FPS footage.

Limitations

Major limitations include:

- **rapidly increasing computational demand**, which conflicts with real-time constraints;
- **strong dependence on video quality**, especially resolution and contrast;
- **limited scalability**, due to hardware computations that using one GPU restrict the maximum reconstruction size;
- **transparency bottlenecks**: Gaussian sorting is computationally expensive and impacts the custom renderer's performance.

To mitigate these issues, it has been employed the Grendel-GS [5] implementation that improves the Gaussian correlation and allows batch processing, while exploiting multiple GPUs. Furthermore, both the SfM and Gaussian Splatting stages have been migrated to an HPC environment.

HPC COMPUTING

Given the exponential growth in computational requirements for reconstructing large, detailed environments, an HPC infrastructure is essential.

The Davinci-1 supercomputer, equipped with over 200 nodes and multi-petaflop performance, has proven itself ideal for handling demands of the Gaussian Splatting. Its high bandwidth ensured fast, reliable data transfers between workstations and the cluster throughout processing.

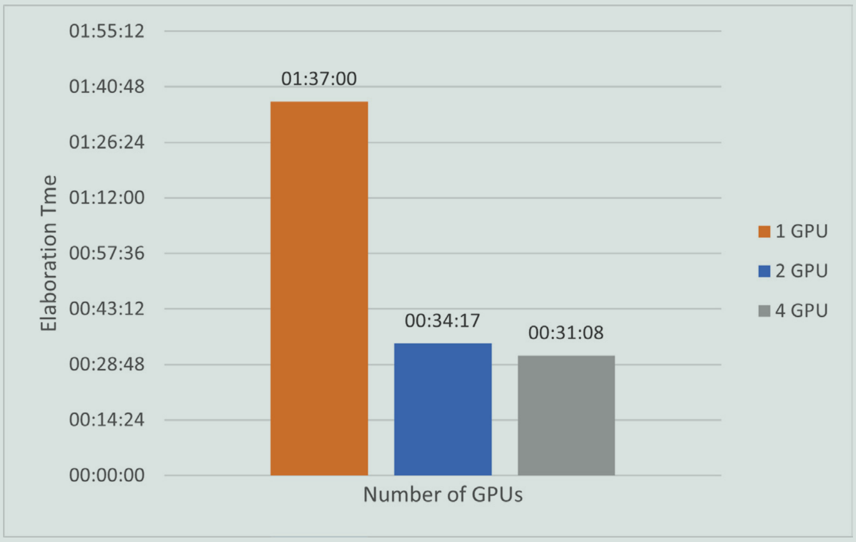

5-Processing time decreases as the number of GPUs increases

FUTURE DEVELOPMENTS

Planned improvements include:

- **cooperative data acquisition** using both aerial and ground agents to cover wide areas with high detail;
- **real-time processing and visualization**, enabling operators to receive live updates while reconstruction is ongoing;
- **fully scalable algorithms** capable of leveraging large GPU clusters;
- **Order-Independent Transparency (OIT)** [6][7][8] techniques to eliminate the need for expensive Gaussian sorting;
- **integration with VR and AR** to allow immersive interaction with the reconstructed environments.

CONCLUSIONS

Battlefield Reconstruction Real Time proves to be an extremely versatile technique for generating photorealistic 3D environments using low-cost hardware and heterogeneous video sources. Its applications span military, civil, and industrial sectors.

Key use cases include:

- pre-and post-event battlefield evaluation;
- risk assessment and post-disaster damage analysis;
- immersive training with realistic reconstructions.

Future integration with VR/AR and cooperative data collection will further enhance its value in operational planning for high-risk environments accessible only to robotic or expendable devices.

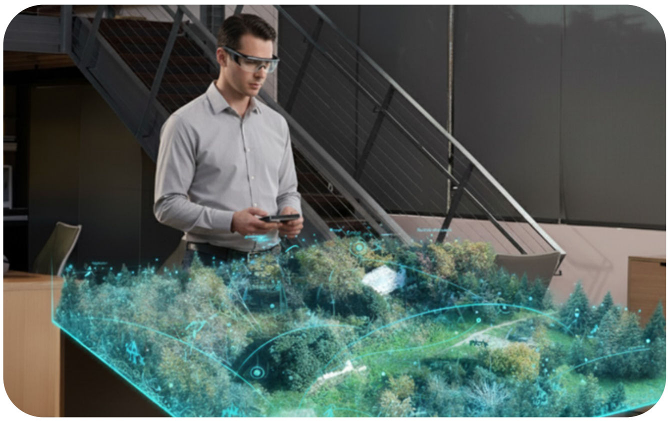

6-Concept of Battlefield Recontsruction Real Time AR application

Adriano Tumino: adriano.tumino@leonardo.com

REFERENCES

[1] J.L. Schönberger, Johannes Lutz and Frahm, Jan-Michael, "Structure-from-Motion Revisited", Conference on Computer Vision and Pattern Recognition (CVPR), 2016.
[2] J.L. Schönberger, Johannes Lutz and Zheng, Enliang and Pollefeys, Marc and Frahm, Jan-Michael, "Pixelwise View Selection for Unstructured Multi-View Stereo", European Conference on Computer Vision (ECCV), 2016.
[3] Kerbl, Bernhard and Kopanas, Georgios and Leimk ühler, Thomas and Drettakis, George, "3D Gaussian Splatting for Real-Time Radiance Field Rendering", ACM Transactions on Graphics, vol. 42, pp. 4, July, 2023.
[4] Matthias Zwicker, Hanspeter Pfister, Jeroen van Baar, and Markus Gross, "EWA volume splatting", Proceedings of the conference on Visualization '01 (VIS '01), pp. 29–36, 2001.
[5] Hexu Zhao and Haoyang Weng and Daohan Lu and Ang Li and Jinyang Li and Aurojit Panda and Saining Xie," On Scaling Up 3D Gaussian Splatting Training", https://arxiv.org/abs/2406.18533, 2024.
[6] Cass Everitt (2012), "Interactive Order-Independent Transparency", NVIDIA OpenGL Applications Engineering, available: https://web.archive.org/web/20110927071951/http://developer.nvidia.com/content/interactive-order-independent-transparency
[7] Louis Bavoil, Kevin Myers (2008), "Order Independent Transparency with Dual Depth Peeling", available: https://developer.download.nvidia.com/SDK/10/opengl/src/dual_depth_peeling/doc/DualDepthPeeling.pdf
[8] Morgan McGuire, Louis Bavoil (2013), "Weighted Blended Order-Independent Transparency", available: https://www.jcgt.org/published/0002/02/09/

Measuring and Verifying Electronic Warfare Mission Data Effectiveness using Synthetics

Romolo Gordini[1], Riccardo Dal Borgo[1], Miriam Chisari[1], Alessandro Moro[1], Matteo Ferfoglia[1], David French[2], Alun Riley[2], Tony Chandler[2], Henry Wong[2], Simon Capp[2]

[1]Leonardo – Electronics Division, [2]Leonardo UK – Electronics Division

Nowadays, technology dominates every aspect of military operations and decides the fate of conflicts. Electronic Warfare has become a fundamental component to battlefield success. In highly customizable Electronic Warfare systems, Mission Data are a critical factor in operational effectiveness of the military strikes and situational awareness. Good Mission Data allow users to positively identify and counteract threats in a timely and effective manner, but opportunities to test them during flight trials on aircraft are infrequent and extremely costly. High fidelity Electronic Warfare real-time simulation offers the ability to evaluate and verify Mission Data by using synthetic emitters in realistic scenarios, bringing flexibility and scalability into the Mission Data validation process.

INTRODUCTION

The defence sector is experiencing rapid transition toward the virtualization of testing and evaluation processes. Within the Radio Frequency (RF) domain, sensor manufacturers are increasingly adopting synthetic solutions to speed up the transformation of military requirements into deployable capabilities [1], enabling a modular, scalable, and repeatable validation process, especially when integrated into sophisticated modelling and simulation ecosystems, capable of testing and measuring the sensor performance in real time. Within the Electronic Warfare (EW) world, operational effectiveness is intrinsically tied to quality of the Mission Data/Intelligence Data (MD/ID) provided to the platform.

In military operations, good MD enables effective identification and neutralization of threats, ensuring safety and success of the missions. However, bad MD can severely undermine the system performance, compromising both the defensive and the offensive capabilities. The creation and validation of high-quality MD is an inherently difficult and resource-intensive task. MD files are often large and complex, and the infrequent nature of flight trials, especially those that are used to verify MD, poses significant challenges for on-field data validation.

This paper explores how high-fidelity real-time simulations can help bridging this gap by providing a synthetic yet realistic environment, in which MD can be tested and validated without any need for costly and infrequent flight trials, with support of the post-simulation analysis tools for verification of the EW systems against the expected behaviour.

ELECTRONIC WARFARE MISSION DATA IN SYNTHETICS

The importance of mission data in electronic warfare

Mission Data allow to deeply customize the EW system capabilities based on the environment in which it is intended to operate. An example of EW system is reported in Figure 1. Such customization affects the system performance under the following capabilities: threat detection, threat identification and threat countermeasure. Although some features are hardware dependent, e.g. the overall frequency bandwidth covered by the of EW system, MD allow to customize the way the system will exploit/optimize its own HW resources, to obtain the desired mission goals. Such configuration directly affects the *detection performance* of an EW receiver and must be properly tuned basing on the threat characteristics the system is expected to face, such as pulse signals, timing properties and emitters scanning patterns. For the correct *threat identification*, the intelligence data embedded in the MD emitter descriptors are crucial. This information can range from signal parameters such as frequency, pulse width, agilities and modulation type, to tactical information such as the emitter scanning patterns the and expected transmitted power. The goodness of these parameters and their coherence with real life RF threats the EW system is going to face, strongly impact the situational awareness and the system defence potential, since the *effectiveness of countermeasures* starts from a correct threat identification. The RF countermeasures library is another EW key component customizable through MD. Its proper configuration may determine the survival of assets deployed in a war scenario and the fate of the mission. It is clear that an accurate MD is essential for timely and accurate threat identification, which is crucial for countermeasures and tactical decision-making, and for platform protection against both traditional and advanced threats, such as radar or electronic jamming.

Any discrepancies in the MD, whether they arise from errors in data generation or incomplete intelligence, can negatively affect the system performance, leading to mission failure and loss of assets. This makes the verification of MD essential before it is used in live operations.

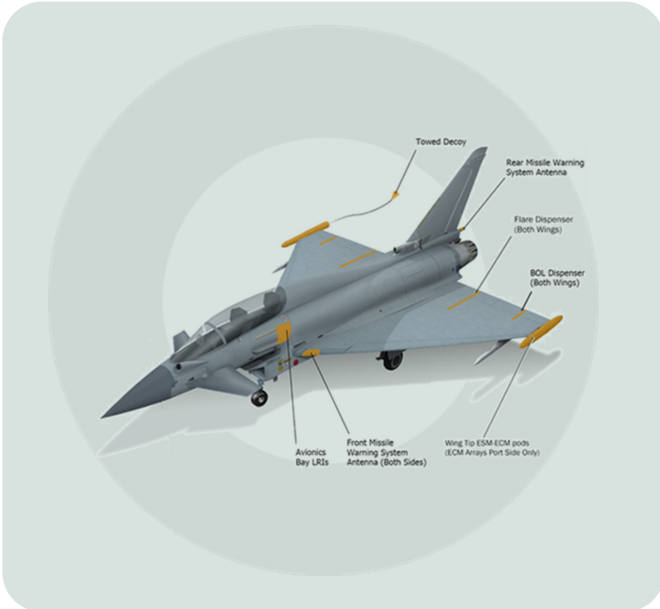

1-Example of EW System: PRAETORIAN DASS

Challenges in mission data verification

The traditional verification method of Mission Data foresees a long and arduous process mainly based on human operators, as depicted in Figure 2.
The traditional process poses the following challenges.

- *Size and Complexity*: MD files are typically large, comprising multiple layers of information that must be processed and validated. Ensuring the accuracy of this data can be a long and complex process, especially when working with multiple systems and threats.
- *Expensive Flight Trials*: the intrinsic expensiveness of flight trials leads to infrequent opportunities for MD verification in real world scenarios and, considering the logistical challenges involved, such trials cannot be relied upon as the primary method of verification.
- *Intelligence Data Leaking*: the Operational MD testing requires the usage of real intelligence data to characterize the environment emissions. These signals injected into the air can be intercepted by other entities and lead to diffusion of sensible data outside the defined perimeter.
- *Uncertainty in Emitting Platform Locations:* in most flight test scenarios, the exact locations and emitter descriptions of all emitting threats in the test zone are unknown. This makes difficult to compare the recorded data with the actual behaviour of the platform, and to fully trust that the results of the trials accurately reflect the system performance.

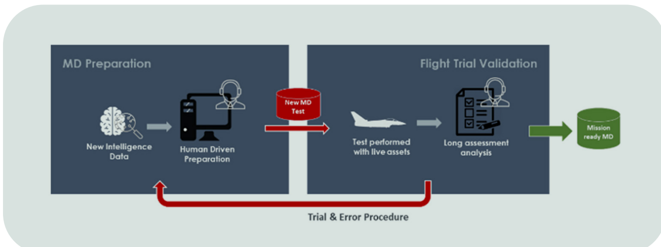

2-Traditional MD Verification Method

High-Fidelity simulation

A high-fidelity real-time simulation offers a solution to the aforementioned challenges by providing a realistic, credible and controlled environment, in which the MD can be tested without any need for expensive flight trials. One of the key components that make the high-fidelity simulation an effective tool for MD verification is the usage of *Real MD*: the simulated system can load the exact same MD that would be deployed on the real aircraft, ensuring that the test conditions in the simulation closely mirror those in a real operational environment.
This enables direct use of the same library between real and synthetics, saving operators time and avoiding possible translation discrepancies. Real MD are injected in the *Re-hosted Code* meaning that the avionic code used in the real EW system is embedded in the simulation environment, eliminating discrepancies that might arise from using different code or modelling approaches, hence ensuring that synthetics closely matches the real-world.

System components that are re-hosted into the synthetics are agnostic about the framework in which they operate, being it real or synthetic. A schematic solution is shown in Figure 3. Exploiting a high-fidelity representation of the electromagnetic environment and different programmable threats, the use of the simulation solution allows the operators to obtain high-fidelity data about the system performance based on the MD loaded into the simulated mission. The simulation can be defined as high-fidelity only after a *Performance Assessment*. This process allows the validation of the whole simulation chain, from synthetic environment and emitter description to EW system identification and countermeasures, against the real system behaviour. Real RIG and flight tests data are compared with the ones obtained in the synthetic solution, to assess that the system performances are coherent and to certify with the manufacturer the goodness of the solution. A pictorial representation of the Performance Assessment activity is illustrated in Figure 4.

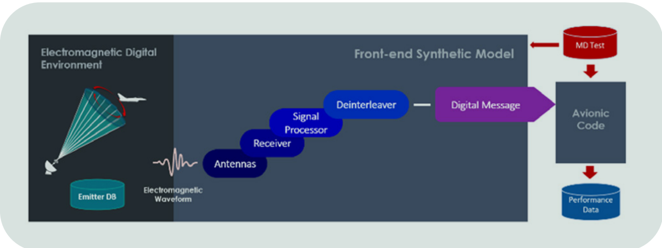

3-EW digital environment

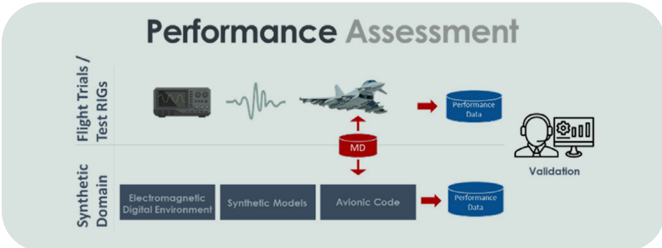

4-Performance Assessment procedure

Benefits of simulation

Synthetic emitters mimic the behaviour of real-world electronic threats. The high-fidelity model of a front-end EW receiver can be employed in a wide variety of scenarios for MD validation. The synthetic testing of EW systems provides the following benefits.

- *Flexibility*: synthetics allows to easily and quickly (re) configure any threat signals and the environment within the same test RIG allowing operators to test new threats or update them according to the latest intelligence data;
- *Scalability*: environmental complexity in terms of number of threat signals and their characterization can be arbitrarily increased and controlled in simulation, allowing more congested operative situations with respect to real flight trials.
- *Cost and Time effectiveness*: synthetics works on general purpose PC, thus eliminating the need for specific hardware.

Moreover, the use of synthetics for MD validation allows to streamline the trial-and-error procedure of defining new MD for threat identification or while testing the performance on different scenarios, avoiding logistic difficulties tied to the setup of traditional flight trials.

- *Repeatability*: synthetics offers a chance to automatically run test batches several times in different configurations, collecting a large amount of data usable for MD verification, overcoming the lack of system output data variety of real flight trials.
- *Low environmental impact*: synthetics allows avoiding non-negligible environmental impact of flight trials.
- *Data Driven Analysis*: the large amount of data available during the synthetic MD validation is an opportunity to bring the benefits of statistical analysis, machine learning and AI techniques into MD preparation workflow.

Once the simulation scenario has been executed, the Post-Simulation Analysis (PSA) tool provides detailed graphical representations of the mission, including visualizations of the emitter positions, platform dynamics and system responses. The PSA tool processes the simulation data, comparing batch of simulated EW measurements with known truths, helping to assess how well the EW system responds to the simulated threats with different MD configuration. The PSA allows to inspect the processing chain to understand why the tested MD set does not provide the desired performance and to investigate whether it could be the ambiguity in the emitter intelligence data that leads to a misidentification or the low-level configuration that leads to loss of threat detection opportunities. Figure 5 summarizes the workflow for MD validation through high-fidelity synthetics.

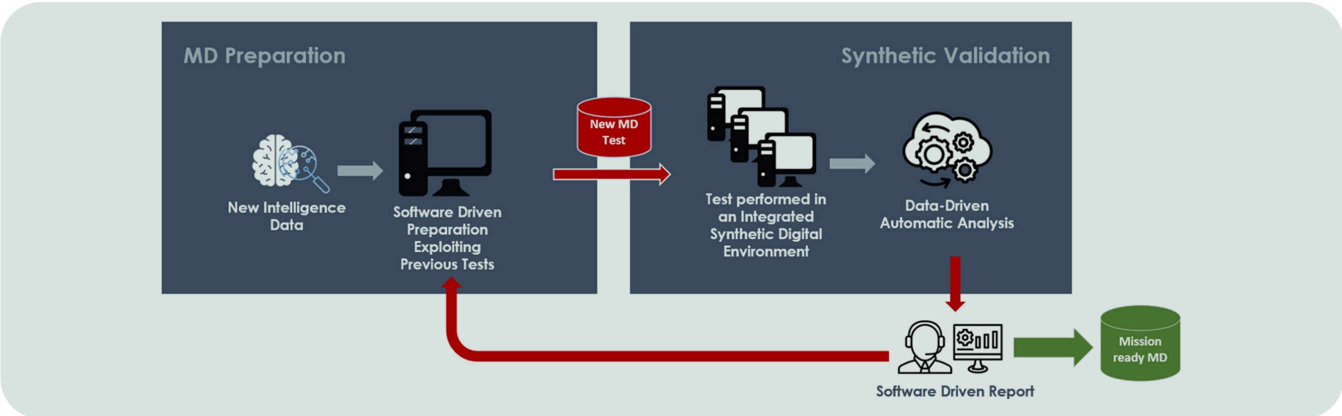

5–Synthetic MD Validation workflow

Applications

Within the Typhoon Eurofighter Program, ITST–Transition Package 2 at P2Eb standard, Leonardo Spa has developed for the Italian Ministry of Defence, a synthetic solution containing the Original Equipment Manufacturer ESM/ECM processor re-hosted code and real MD, as well as the simulated front-end based on the MARS-EW models. Offline tools have been developed in order to parse the real DASS MD, by using OEM IGS and MD SIDS. Performance Assessment activities have been carried out to compare the performance results and metrics of the real equipment against the synthetic solution.

CONCLUSIONS

A high-fidelity real-time simulation, combined with advanced MARS-EW tools, assures an effective mean to rigorously test and verify MD in a realistic setting and speeds up the process to reach operational-readiness at lower cost. Moreover, leveraging the synthetics throughout the entire development lifecycle, from design to operational support, military forces and industry can significantly enhance the performance of their own EW systems, while minimizing the risks associated with poor MD.

Romolo Gordini: romolo.gordini@leonardo.com

REFERENCES

[1] Rouffet, T., Poisson, J., Hottier, V., & Kemkemian, S. (2019). Digital twin: a full virtual radar system with the operational processing. International Radar Conference.

Revolutionizing Helicopter Design and Development through Virtual Reality

Maria Ester Augello[1], Alessandro Scandroglio[1], Francesco Di Cintio[1], Marco Ralli[2], Andrea Cristina[1], Fabio Iasi[1]
[1]Leonardo – Helicopters Division, [2]Leonardo – Corporate/Strategy and Innovation

Virtual Reality (VR) is revolutionizing the way we design, test, and refine rotorcraft at Leonardo Helicopters Division (LHD). By integrating VR into our processes, we are not only improving the pilot–machine interaction but also accelerating development cycles and reducing costs. Initially introduced to support ergonomic assessments, VR has progressively evolved into a versatile engineering enabler that reshapes development workflows and strengthens cross-department collaboration. The Phygital Mockup system represents a next-generation integration of physical components and high-fidelity virtual environments, which enables fast, data-driven evaluations throughout the design and prototyping phases. According to Leonardo's Industrial Plan, this approach strengthens engineering efficiency, improves quality of the early design validation and creates a scalable framework to extend to future applications across the product lifecycle.

UNDERSTANDING THE IMMERSIVE SIMULATION LANDSCAPE

Virtual Reality (VR) immerses users in a fully digital environment through the use of head-mounted displays and motion controllers, thus enabling natural interaction with virtual objects and surroundings. Augmented Reality (AR) overlays digital content onto the physical world, enhancing perception of the real-world without replacing it – for example through navigation overlays or interactive 3D models viewed via smartphones or tablets. Together with Mixed Reality (MR), in which virtual and physical elements coexist and interact in real time, these technologies form the broader domain of Extended Reality (XR), which spans a continuum from completely real to fully virtual environments:

- **Reality**: the physical world without augmentation;
- **Augmented Reality (AR)**: real-world view enhanced with digital overlays;
- **Mixed Reality (MR)**: integrated digital and physical components with real-time interaction;
- **Virtual Reality (VR)**: fully simulated digital environments.

Tools such as Unity, Unreal Engine, VRED and modern VR headsets like the Meta Quest family are fundamental enablers of XR development, as they allow engineers to create interactive 3D environments, immersive simulations, and advanced visualization tools used across training, engineering, design, and marketing.
Within this landscape, it is crucial – particularly for aeronautical applications – to distinguish between two main categories of XR solutions:

- simulation-environment applications with software-in-the-loop (SiL);
- **standalone VR** experiences.

SiL is directly coupled with engineering toolsets and real-time computational solvers, enabling high-accuracy assessment of pilot-training scenarios, rapid verification of software integrations, and in-depth exploration of piloting configurations. Standalone VR applications, by contrast, emphasize immersion and adaptability over computational coupling, enabling a broad range of human-centric analyses and scenario explorations that benefit from speed, accessibility, portability and ease of iteration.

1-Example of LHD mixed reality on AW189

Recognizing the distinction between these two XR paradigms allows different business functions within the company to identify the most suitable technology for their objectives – whether they require rigorous engineering validation or flexible, immersive visualization environments.

2 – Meta Quest Pro Headset

THE EVOLUTION OF VIRTUAL REALITY IN LHD

Traditionally, cockpit evaluations have relied only on large, static wooden mockups that are expensive to build, slow to modify, and inherently limited in their ability to support iterative design. A single physical-only mockup requires over 700 hours of engineering effort and significant logistical resources, often constraining the opportunity for human-centred adjustments once its structural constraints are fixed. These limitations motivated the exploration of a hybrid digital-physical approach capable of supporting rapid configuration changes and pilot-in-the-loop evaluations early in the design process.

The Phygital Mockup has been initially conceived as an ergonomics-driven evaluation tool introducing a breakthrough approach to helicopter design by combining selected physical interfaces with the immersive capabilities of virtual reality. Leveraging advanced VR platforms and lightweight headsets, the system enables precise assessment of cockpit layout, reachability, visibility, and human–machine interaction within hyper-realistic operational scenarios. The technology delivers immersive environments for **vehicle architecture and ergonomic analysis**, which allow pilots, co-pilots, and operators to evaluate configurations from multiple perspectives, anywhere in the world. Real-world geographic coordinates, weather conditions, and mission profiles can be replicated without any need for physical mockups, significantly improving flexibility while reducing cost and lead time.

Over the course of multiple programs, the system has progressively expanded its scope of application. Initially focused on cockpit and canopy development, VR enabled pilots to assess visibility in complex scenarios – such as shipboard operations or confined-area approaches – by switching among predefined aircraft attitudes and real-world camera positions.

Subsequently, the introduction of full-scale virtual walkaround environments extended the use of the Phygital Mockup to external evaluations. Pilots and maintainers were enabled to analyse accessibility, clearances, door and ramp ergonomics, and ground-operation workflows without requiring a physical prototype.

By addressing the limitations of traditional wooden mockups, the Phygital Mockup accelerates design timelines while improving test accuracy. This evolution has enabled to be performed ergonomic evaluations much earlier in the design process, supported by quantitative data such as gaze tracking and hand tracking, which contribute to provide more functional and intuitive cockpit configurations. As the helicopter market continues to grow – particularly in emergency medical services, search and rescue, and urban air mobility – Leonardo Helicopters Division continues to recognize the increasing demand for mission-ready aircraft that are validated through early, pilot-centred design methodologies.

In three years, the system has achieved remarkable efficiency milestones: a **Technology Readiness Level 8 (TRL8)** functioning prototype has been developed with negligible financial investment.

3-Virtual Reality applications directly from CAD models to transform development processes

Applied across multiple helicopter programs, including platforms with military requirements, the Phygital Mockup has delivered substantial return of investment by reducing both design time and iteration costs. The planned lightweight modular rig enhancements, available since mid-2025, include 3D-printed components, Bluetooth-integrated controls, and reconfigurable setups adaptable across different aircraft models, which further reinforce the system's scalability and long-term sustainability. LHD has progressively leveraged the Phygital Mockup beyond such pure ergonomic studies, to address program-specific needs related to Human–Machine Interface (HMI) and system integration. The inclusion of detailed virtual environments and advanced functionalities – such as haptic gloves and digital instrument panels – positions the platform as a comprehensive framework for multidisciplinary analysis, extending well beyond traditional prototyping.
In collaboration with the Flight Mechanics and Avionics departments, the virtual cockpit has been connected to software-in-the-loop avionics logics, to enable real-time updates of displays and controls. This integration supports rapid assessment of HMI changes, failure management behaviour, pilot workload, and pilot–machine interaction, significantly reducing dependency on hardware rigs. Originally conceived as an ergonomics-focused evaluation tool, the virtual cockpit has evolved into a versatile simulation environment that supports early design validation, system behaviour assessment, accident analysis, performance comparison of different aircraft in real manoeuvres, and cross-disciplinary collaboration throughout the development lifecycle.
As the platform matured, its capabilities have been extended to safety and operational analysis domains. The virtual reality environment has been successfully used to reconstruct accident scenarios in support of the Accident Analysis department, as it enables immersive replay of events based on available data and contributes both to internal investigations and dedicated to the SIA (Safety Investigation & Analysis) course.
This prototype integrates recorded flight parameters, control inputs, and environmental data into a synchronized three-dimensional reconstruction. The ability to visualise accident dynamics from multiple viewpoints – inside and outside the cockpit – has provided particularly valuable insights for understanding pilot actions, system responses, and environmental factors in a controlled and repeatable setting, beyond traditional time-series analysis.
In parallel, the same VR framework has been adopted within an engineering programme to reproduce real operational manoeuvres and conduct comparative evaluations across different aircraft flying together. By placing multiple rotorcraft in the same virtual scenario, engineers and pilots were able to assess relative performance, handling qualities, and mission effectiveness during coordinated flight profiles.
This demonstrates the flexibility of the tool in supporting complex, multi-vehicle analyses that would be difficult, costly, or impractical to perform in real flight.
This transformation results from a multidisciplinary effort spanning digital transformation and software development, vehicle architecture and ergonomics. The versatility of the Phygital Mockup fosters collaboration across Leonardo's teams and divisions, enabling rapid iterations, objective feedback, and continuous design refinement.
This technology is widely adopted across LHD activities, delivering immediate benefits. In particular, cockpit and canopy development applications enable test pilots and customer pilots to perform extensive evaluations and validations of both structural and avionics configurations, with the objective of optimising visibility and HMI. Even helicopters already in production – especially those belonging to the same Family line–benefit from VR-based validation of requested onboard system modifications, such as the central console layouts or the placement of emergency devices like the life-raft activation controls.
In addition, the technologies developed are successfully applied to rotorcraft still in the concept phase. Several programs have allowed test pilots to "fly" the aircraft from the earliest stages of its preliminary definition and to provide meaningful feedback on its overall configuration. Engineers have also gained early insight into aspects that previously could not be evaluated until a full-scale prototype was built, including maintenance accessibility, cabin access (doors, ramps, and walkways), system functionality, and overall HMI coherence. Moreover, this technology enables engineers to compare different configurations of the same aircraft through immersive evaluations of manoeuvres and flight paths, using both internal cockpit and external viewpoints. In the past year, VR played a crucial role in one of the most challenging regulation-related initiatives: the Piloting Standard project. Certification bodies, particularly EUROCAE, are currently assessing new flight-control solutions for emerging UAM applications. Virtual reality and phygital mockups formed the technological foundation of the simulation rig used by test pilots, certification authority pilots, and participating OEMs (Original Equipment Manufacturer), to evaluate alternative control strategies.

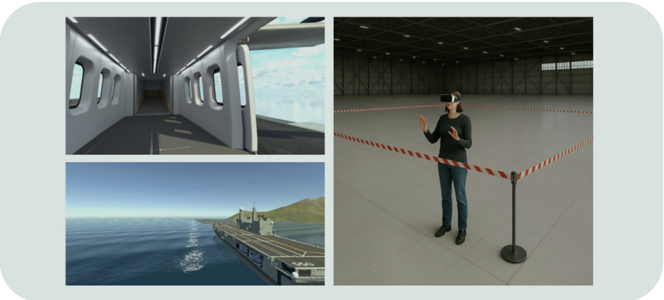

4 – Example of demo with clients

Finally, work is nearing completion on applying these technologies to the new LHD VXR flight mechanics simulator. The simulator reproduces the helicopter motion by using an actuated platform, and these VR applications enabled the rapid implementation of new aircraft configurations in just a few days, which extends the use of this certified training asset to intensive engineering and validation activities.
This initiative originated in 2022 within the "Digital Mindset and Business Analytics" course at Politecnico di Milano and evolved into a Proof of Concept through the LHD-sponsored **Digital Transformation Initiative**. As of 2025, the Phygital Mockup continues to mature as a strategic **digital enabler** within Leonardo Helicopters Division.
In essence, Virtual Reality enables the digitalization of traditional wooden mockups, which transforms static, resource-intensive assets into dynamic, reconfigurable environments that support continuous pilot-in-the-loop evaluation across the entire aircraft lifecycle.

5 – First RIG built to support fast integration between virtual reality and flight mechanics software and controls

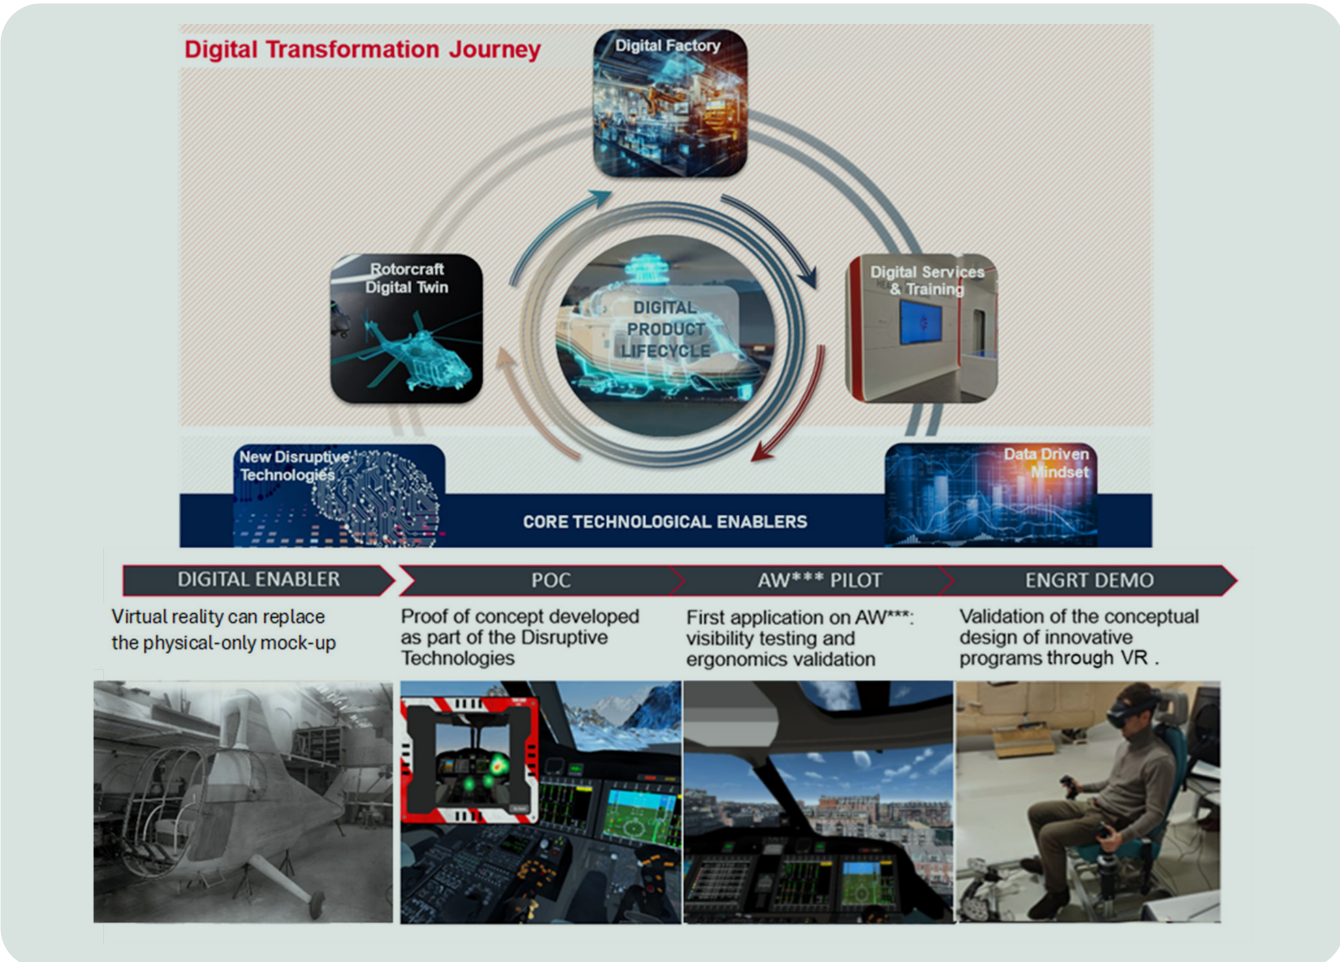

6 – Virtual Reality developed as digital enabler of the New Disruptive Technologies of Digital Transformation

FUTURE INDUSTRIAL APPLICATIONS

Built on activities already carried out within engineering – during the conceptual and experimental development phases of rotorcraft – this capability could be extended to support configuration definition for industrialized and commercialized products, into broader **industrial and commercial domains.**

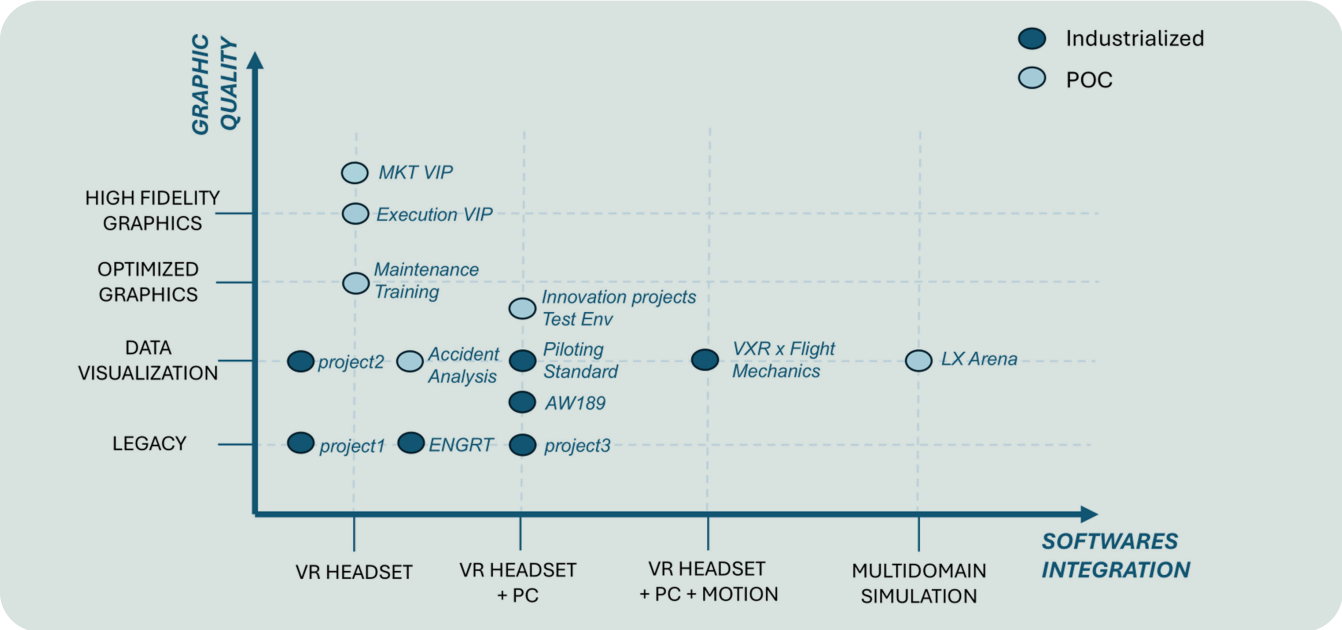

7-Virtual Reality technological roadmap

Customers actively customize conceptual cabin layouts, mission-role installations, and optional equipment within an immersive, data-driven environment aligned with the relevant certified product baseline. For Leonardo Helicopters Division, a VR framework could enable VIP and other customers to virtually inspect and interact with industrialized products during selected phases of contract execution. As part of a proof of concept, the solution would allow the assessment of alternative configurations, preview interior layouts prior to commitment, and better understand clearances, spatial constraints, and the functional implications of different design choices. If validated, these capabilities could be extended to use cases such as trade shows, virtual showrooms, or remote demonstrations, offering a scalable and efficient approach to future product presentation and customer engagement.

Preliminary proof-of-concept evaluations have also explored the application of VR to rotorcraft **maintenance training.** Although still at an early stage, the technology could provide maintainers with a controlled virtual workspace in which to rehearse inspection routines, practice removal and installation sequences, and diagnose system anomalies without requiring immediate access to a physical aircraft. Future developments may include the integration of procedural logic, step-guided workflows, remote instructor monitoring, and performance analytics, enabling VR to evolve into a complementary component of the technical training ecosystem. As the methodology matures, VR has the potential to become a valuable asset for safety investigations, operational debriefings, and for the iterative optimization of flight procedures.

Looking ahead, VR also could enable **multi-user and distributed battle simulation scenarios**, in which pilots, engineers, technicians, mission planners, and support personnel can operate simultaneously within the same virtual environment.

Such collaborative simulations could support mission planning iterations, cockpit–cabin coordination assessments, tactical decision-making exercises, thus enhancing the shared situational awareness and cross-functional collaboration.

Overall, virtual reality emerges as a versatile technological enabler capable of supporting the entire product lifecycle from the early design assessment and customer co-creation to training, marketing, after-sales support, safety analysis, and collaborative operational simulations. As XR technologies continue to mature, aeronautical organizations can increasingly leverage VR, AR, and MR solutions to enhance engineering workflows, operational capability development, and customer experience, across all major business domains.

8 - Example of LHD AW189 VIP Proof of Concept

CONCLUSIONS

Virtual Reality has evolved beyond its original role as an ergonomics-driven analysis tool to become a fully fledged, standalone framework supporting the entire helicopter product lifecycle. What initially emerged as a proof-of-concept focused on human–machine interaction has progressively matured into an agile, rapid, flexible, modular, and portable digital capability serving the whole helicopter value chain.
Today, Virtual Reality enables early and continuous interaction with aircraft concepts across design, engineering, maintenance, and operational domains, well before physical prototypes are available. Its effectiveness lies not only in immersive visualization, but in its ability to integrate heterogeneous data, users, and use cases within a single digital environment.
To fully unlock this potential, Virtual Reality is being formalized and reconstituted as a structured framework, tightly connected to shared corporate repositories and aligned with the development ecosystems common to multiple business areas. This integration ensures consistency, scalability, and reuse of data, tools, and methodologies across the organization.
The overarching objective is the definition of a common technological roadmap and guiding framework for Virtual Reality, capable of supporting the entire product value chain. By establishing shared principles and architectures, Virtual Reality becomes a strategic enabler for safer, faster, and more efficient helicopter design, development, and operation in the years ahead.

Maria Ester Augello: mariaester.augello@leonardo.com

REFERENCES

[1] Wikipedia, Reality–Virtuality Continuum, Wikimedia Foundation, https://en.wikipedia.org/wiki/Reality%E2%80%93virtuality_continuum
[2] Unity Technologies, *Unity Manual & Scripting API Documentation*, Unity Technologies, latest version.
[3] Unity Technologies, *High Definition Render Pipeline (HDRP) Documentation*, Unity Technologies, latest version.
[4] Meta Platforms Inc., *Meta XR All-In-One SDK Developer Documentation*, Meta Platforms Inc., latest version.
[5] Meta Platforms Inc., *Meta Quest Pro–Developer Guidelines and System Specifications*, Meta Platforms Inc., latest version.

Editor in Chief
Vincenzo Sabbatino

Editorial office
Emidio Di Pietro
Giovanni Cocca
Marco Morini
Patrizia Pozzoni

Published by:
Leonardo S.p.A.
GM Strategy & Innovation
Piazza Monte Grappa, 4
00195 Roma

Printed by:
Leonardo Global Solutions
BU Facility Management & Site Services
Centro Stampa - Via Tiburtina km 12,400
00131 Roma

The Editorial Team thanks Agatino Mursia
for serving as the Guest Editor, and Danilo Defant and Paolo Casanova for their contribution.

The POLARIS Innovation Journal is an editorial initiative of the General Management Strategy & Innovation.
Other initiatives of the POLARIS Innovation Journal are the Paperbacks and the Lunchtime Webinars.
The Journal invites questions and suggestions from readers.
Contact the Editorial Office at: polaris@leonardo.com

Scan this QR code to access the web version

https://www.leonardo.com/polaris_2026_54/

In compliance with the Leonardo sustainability policies,
and to contribute reducing the environmental footprint of the Company,
the POLARIS Innovation Journal is printed on certified paper (Xerox International Certificate).
The POLARIS Innovation Journal is published biannually.
Issue 54 – February 2026