Leveraging Digital Era Technologies for a new Synthetic Environment – a Panacea?

Philippe Perey
CAE Inc.
CANADA

philippe.perey@cae.com

Nick Giannias, Mátè Koch, Jean-Sebastien Dion, Jonathan Labonté, Dave Sexton
CAE Inc.
CANADA, GERMANY, UK

nick.giannias@cae.com mate.koch@cae.com jean-sebastien.dion@cae.com

jonathan.labonte@cae.com dsexton@cae.co.uk

ABSTRACT

Technology crossover from commercial markets to military applications is not new. Standard computers, programming languages and the use of graphics cards are examples of early adoption in military training devices. The digital era now offers a plethora of technologies. One area of promise is the combination of Gaming Engines, Cloud Computing and Artificial Intelligence to dramatically increase the digital immersion, fidelity, and scale of a next-generation Synthetic Environment (SE). With millions of users, the underlying technology would enable a shift from a bespoke and expensive SE, to a modular, robust, and affordable one.

This paper will summarize several years of research and numerous demonstration prototypes developed by CAE to exploit these technologies as it builds its next generation SE for tactical training, operational support, and concept development. The paper will highlight the strengths and pitfalls of these technologies applied to training as well as their use as a foundational component for validating platform mission capabilities in complex environments. The paper will also describe how new SE features may be combined to create a country-size digital twin, representing infrastructure, military assets, and citizens (including a trending view of sentiment and corresponding actions), as part of Assisted Decision Making, course of action analysis and government resiliency plans.

1.0 SYNTHETIC ENVIRONMENT OVERVIEW AND VISION

The ability of military forces to achieve mission-readiness in complex and large-scale scenarios increasingly relies on digitally recreating these environments for the purpose of doctrine definition, validation and ultimately as a realistic Synthetic Environment (SE) for training. Synthetic training environments have been exploited for decades, but the recent re-emergence of near-peer threats necessitates a step change in fidelity and scale to better train for the potential conflicts of tomorrow.

The simulation industry had developed two variants of synthetic environments, each focused on a specific use case. Lower temporal and spatial resolution SEs, often allowing faster than real time simulation are primarily used in constructive simulation for training military leaders. These systems allow commanders to control thousands of entities at a time, modelling up to brigade-level scenarios. To achieve scale, these SEs operate with simplified models, and lower temporal and spatial resolution (e.g., 15-meter grid and 5-second interval). Scenario visualization is also limited to 2D maps and moving icons, and have been used for at least 2 decades, training military leaders and strategic tasks such as doctrine development, decision support, and wargaming.

The second type of SE, providing both high-fidelity behavior and real-time movement of entities, is typically integrated with virtual training devices such as flight simulators. This type of SE, referred to as Computer-Generated Forces (CGF), is often limited to several hundred entities, especially when simulating fast-moving entities with high-resolution sensors and/or weapons.

The objective of the developments described in this paper is to provide the scale and the fidelity necessary to address both needs in a unified SE. It is CAE's belief that such a capability would not only enhance the fidelity of traditional training use cases but also expand its use to new applications described below.

1.1.1 Multi-Domain Operations

NATO allies are shifting their national defense strategies to address near-peer threats. This requires operation across all five domains of air, land, maritime, space and cyber in a more integrated and efficient way. To model the full spectrum of military assets (and the potential impact on civilians and infrastructure), SEs of the future must handle millions of dynamic entities. This enables more realistic and accurate simulation-based training but is also foundational to creating new decision support tools for military leaders. This could be done by running multiple simulations of different decisions and assessing outcomes and how they align to mission success criteria. Such capabilities could be embedded into future command and control systems, allowing commanders to assess course of action analysis.

1.1.2 Human Modelling

Realistic simulation of civilian populations at scale will require faithful representation of both individual and group behaviors. Unlike military interactions that can be modelled based on doctrines and orders, the behavior of a civilian crowd cannot be modelled with traditional rule-based logic. Cultural, religious and political views, along with the influence of social media are among the many factors that affect an individual's behavior. People follow daily routines, commute, work, and spend leisure time in a so-called pattern of life. Simulating this in future SEs, will enable modelling of complex group behaviors, supporting leaders and planners to better prepare and respond to a broad suite of potential events, spanning natural disasters (flooding, fires, drought, etc.), terrorist attacks, to large scale events (e.g. Olympic games). This could also include pandemic modelling (including spread of a virus based on specific public health policies), and urban planning (such as positioning future subway stations).

In the past, large scale population modelling has been performed using statistical techniques. While the predictive power of statistical models may be faster and marginally more accurate than causal models using techniques such as agent-based modelling [1], statistical models are theory-free black boxes that cannot be inspected or queried for insight into the cause of specific outcomes. Credible Agent-based, civilian population modelling enables the application of interventions from across the diplomatic, information, military and economic (DIME) spectrum and

[1] During Defense Advanced Research Projects Agency (DARPA) trials, Agent-based modelling achieved an accuracy of 88%, precision of 68% and a recall of 84% in the back-casting of events, a result that was close to traditional statistical modelling.

exploration of their impact upon the political, military, economic, social, information and infrastructure (PMESII) aspects of society in a way that can be queried using data collected during a simulation run.

Smart cities of the future could be connected to comprehensive SEs creating a true digital twin. This twin could be used both for planning and policy making, but also real-time decisions such as smart traffic management.

1.1.3 Development of Next-Generation Capabilities

A new SE could also accelerate development of next-generation weapon systems by stimulating operational software and hardware in the loop to validate capability. Let's take the example of a new missile system equipped with multiple sensors. The SE provides real-time FLIR and electro-optic images of a target to the missile's sensors, which in turn drives the device's image processing and target identification logic, which then commands the flight management system to re-align. Stimulating all sensors this way, combined with high fidelity flight models of the missile, would allow it to lay on the workbench and still believe it's closing in on a target.

Development of future weapon systems is also increasingly shifting from individual devices to a network of "system of systems". In tomorrow's battlespace, we foresee these manned and unmanned systems seamlessly connected through high bandwidth combat clouds, sharing tactical information, sensor data, and associated engagement logic. To develop and test such complex interaction, an interactive SE would be key to validating functionality and system effectiveness, while experimenting with new doctrines and tactics to exploit the capability to its maximum potential.

1.1.4 AI Validation and Certification

Artificial intelligence is playing an increasing role across both defense and civil markets. This includes self-driving cars, air taxies, and unmanned next-generation weapon systems to name a few. Deep learning AI systems require extensive data for algorithm training, reinforcement, and validation. SEs are already used today to train AI agents, but the scale, realism and representative interaction with a virtual environment must be enhanced to deal with complex and dense urban or military scenarios. Wide-ranging use cases are required to create a robust and operationally ready AI-based system. These SEs must include faster-than-real-time simulations to accelerate AI learning, simulating various weather scenarios, realistically dense city traffic and pedestrian activity and movement for emergency conditions.

1.2 Innovation Ecosystem

The broad availability and robustness of underlying SE technologies has the potential to further accelerate development and expand its use into new domains. This is typical of maturing technologies as it shifts from expensive and bespoke solutions in niche markets, to rapidly developing ecosystems available to the masses. The Apple App store and Google Play are perfect examples of ecosystems that have enabled the creation of over 5 million apps as of 2022.

The metaverse development is currently focused on entertainment, collaboration, and human interactions, but the underlying technology to build it will flow directly into future SEs. This may sound farfetched, but that is also the story of the Internet, initially connecting mainframe computers for scientists to share information progressing to driving the worldwide digital economy and connecting appliances and lights in our homes.

It is our belief that the fusion of digital era technologies in creating a realistic, dense and interactive Synthetic Environment may be the start of just such a journey.

2.0 SYNTHETIC ENVIRONMENT DEVELOPMENT

2.1 Introduction

In this paper, we define a Synthetic Environment (SE) as the compendium of datasets, models and simulations replicating the real world with sufficient fidelity and interactivity to support a training scenario; or assessing other mission or operational outcomes.

An SE includes geospatial data (terrain, water, buildings, roads, railways, etc…), an associated physical and electro-magnetic representation, dynamic behavior of scenario participants (such as ships, aircraft, vehicles, humans), weather, and their associated means of visualization across the electro-magnetic spectrum.

Current SE technologies used in military training typically involve a geospatial dataset supporting a Computer-Generated Forces (CGF) or Semi-Automated Forces (SAF) such as US Army's OneSAF, and a visualization system to present a real-time rendering and a sensor view of the scene to the trainee.

Over the past 5 years, Gaming Engines, Cloud Computing and Artificial Intelligence (AI) algorithms have matured. Through the COVID-19 pandemic, millions of users have exploited these capabilities, bringing a level of robustness and scale that could help redefine a new SE for tactical training, operational support, and concept development.

This paper summarizes the work and key observations that CAE has performed in prototyping and developing a new Synthetic Environment over the past several years, culminating in a new SE based on Epic Games' Unreal Engine 5 (released in the spring of 2022), as well as efforts to leverage Cloud Computing and AI to provide scale and realism to the simulation of dense urban environments. The paper is structured in four main sections: visualization, modelling & simulation, cloud computing implementation and AI-based modelling.

2.2 Visualization

In 2015, CAE began experimenting with different gaming engines to estimate the maturity and potential roadblocks of the technology as a new software foundation for professional-grade visualization for military simulators. CAE had been developing real-time image generators (IG's) since the early 1990's and launched the CAE Medallion-6000 IG family in the mid 2000's. Along with its civil-branded twin, that image generator software, in various revisions, has been fielded on over 1000 full-flight simulators, and numerous Naval and Land training devices.

The Medallion IG software is a robust and feature-rich system, but the software stack is aging, and it is becoming increasingly difficult to cost-effectively add new rendering capabilities and features. The rendering software, like many mature IG's systems on the market today, relies on the aging OpenGL graphics library. A migration to the Vulkan API – an OpenGL upgrade path – would be required to access new rendering capabilities.

CAE assessed different gaming engines and selected Epic Games' Unreal Engine as the basis for the further evaluation. Although other engines were equally capable, access to the source code meant that our engineers could resolve and/or extend capabilities in the underlying Unreal Engine according to our needs and specific use cases.

Figure 1: UE5 rendering of Eurofighter over ocean

An early 3-channel prototype was installed on a simulator to assess image quality and capabilities. At the time, several key issues remained, including anti-aliasing quality, rendering performance determinism, and large-scale numerical precision for whole-earth databases. The effort was shelved for several years.

By 2018, GPU improvements and gaming engine enhancements justified a large-scale development of a new IG based on Unreal Engine. The initial work was done on Unreal Engine 4, and then migrated to Unreal Engine 5 (UE5). A sample screenshot from that effort is displayed in Figure 1, demonstrating the significant increase in realism of advanced lighting of the UE5 engine.

Some of the key features and benefits of a gaming engine-based IG are listed in Table 1.

Table 1: Key Features and Benefits of Gaming Engine to SE Visualization

Capability/Feature of Game Engines	Value to Training/Digital Immersion
Lighting model	Improved material, contrast, and object perception compared to legacy IG's. Physics-based rendering (PBR)
Shadowing	Depth and height perception
Robust code base, extensively tested	High reliability and fewer deficiencies
Ray tracing / Global illumination	More realistic reflections and scene immersion
Highly optimized rendering pipeline	Higher scene density, higher mesh and texture Level of Detail (LODs)
Triple-A game rendering techniques	Increased scene realism; Access to future engine features
Marketplace of assets and plugins	Faster prototyping of new features and effects
Extensive Digital Content Creation (DCC) tools	Rapid development of special effects by artists
Integration of commercial DCC tools	Rapid 3D model content creation, lower artist training costs
Access to advanced graphics debugging tools	Quick turn around on low-level performance and stability issues
Comprehensive visual scripting language	Faster model animation and control by artists
On-going gaming engine evolution exploiting new GPU capabilities	Shorter delivery of new GPU rendering or performance enhancements to training

The quality of the rendering, the rich toolset, robust code base and worldwide marketplace of content were key drivers in the decision to proceed with the new image generator development. However, visual engineers had to assess each feature in detail to determine if alternate approaches were possible to circumvent any limitations. The following chapters discuss some of the key issues and the proposed solution that CAE is implementing to resolve them. UE5 code-level interventions have been minimised to ensure ongoing compatibility with future Unreal Engine releases.

2.2.1 Worldwide Scale

Modern game engines such as Unreal Engine are primarily designed and optimized for relatively small geographical areas. The launch of UE5 addresses this with 64-bit precision coordinates and a new geo-referencing plugin to help in converting coordinates between systems. UE5 also introduces a new “World Partition” system, designed to subdivide a large world into tiles, along with the capability to stream (in and out) content tiles as you move through the world.

These features go a long way in providing large world support, however not all UE5 features are updated to work properly with very large worlds. Several UE5 features still assume that the Z-axis is up. To maintain this, the run-time software needs to “rebase” the origin as you fly around the world, and such rebasing causes important performance glitches, affecting scene rendering determinism.

CAE’s solution is to use the 64-bit large world coordinates, fixing the origin once. A new custom tiling and paging system has been implemented, replacing the UE5 World Partition functionality.

2.2.2 Improved Anti-Aliasing (AA)

The UE5 engine supports many advanced anti-aliasing modes, optimized for the deferred rendering mode of the engine. This includes Fast-Approximate Anti-Aliasing (FXAA), Temporal Anti-Aliasing (TAA) and a new Temporal Super Resolution (TSR) technique. Some 3rd parties also offer plugins for UE5 that implement other advanced techniques, such as Deep Learning Anti-Aliasing, Deep Learning Super Sampling (DLAA/DLSS) as well as FidelityFX™ Super Resolution (FSR).

All these techniques are designed and optimized for the deferred rendering mode of the engine, and as such can NOT be used simultaneously with GPU-native Multi-Sample Anti-Aliasing (MSAA), the traditional go-to anti-aliasing of legacy (e.g., forward rendering) IG’s. UE5 can be configured in forward rendering, but this breaks many of the desirable UE5 features such as Nanite (a component of UE5 that allows virtualisation of the 3D geometry, enabling the rendering of orders of magnitude more polygons than before) and Lumen (UE5’s next generation dynamic global illumination and shadowing system).

These advanced anti-aliasing modes work well in a game context:

- Scenes move quickly which hide minor scintillation and AA issues on the screen,
- Players are close to the action (i.e., scene objects are closer in screen space),
- Players tend to use higher resolution display setups (a 24-inch 4K screen at 50cm from the player results in 71ppd, or 0.85 arc-minutes per pixel, typically twice the resolution of a flight trainer),
- Players are tolerant of minor artifacts.

In a full mission trainer, the use case is quite different:

- Pilots are looking at far away objects (small in screen space, such as runways, aircraft targets, etc.),
- The viewpoint is moving relatively slowly and fluidly, amplifying aliasing artifacts,

- Flight trainer displays typical have less optical resolution (despite having a much large field of view).

CAE's SE visualization solution is to retain the deferred rendering pipeline, but to add specially designed "per feature" anti-aliasing techniques to compensate. One such technique is an off-screen MSAA rendering pass for some key scene elements, such as runway markings. This allows us to benefit from MSAA on high-contrast markings geometry, while benefiting from total scene (TAA) anti-aliasing. This requires modification of the UE5 engine to add an additional pass to the rendering pipeline. CAE has also optimized the default TAA in UE5 to reduce sample jittering.

2.2.3 Latency and Rendering Performance Determinism

Deterministic rendering time is essential to all high-end training device image generators to ensure the scene presented to the trainee is smooth throughout the training event. Exceeding the frame time (typically 8.3msec or 16.6msec), causes the scene to momentarily freeze and can be very distracting to the trainee. Most game engines aren't designed for stringent real-time determinism. Surprisingly, the usual culprit is not the rendering, but rather data and memory management. The following issue is quite technical but worthy of discussion.

UE5 is well designed for streaming, but memory management is non-deterministic. It relies (like many systems) on a software programming task called "garbage collection" (GC) – configurable either on-demand or at fixed time intervals. Each time GC is run, a "marking phase" is done, going through all objects in the scene, marking those that are ready to be deleted. This marking phase is a synchronous task that blocks all critical UE threads, resulting in potential overruns. On a typical rendering IG computer, this marking phase can consume 10, 20, or even 100ms (!) in complex scenes. In this regard, other game engines such as Unity implement a better garbage collector that can run in incrementally, meaning it divides its work over multiple frames to avoid huge spikes.

By default, UE5 is configured with a "swap chain" of multiple buffers. This means that if we overrun while rendering one frame, other pre-rendered frames are still available for video scan out. This technique is good at filtering (small to medium) spikes from garbage collection. These multiple buffers add a lot of latency (about visual 2 frames), which most training devices can't afford.

For games that are well optimized and don't run at 100% capacity, and don't generate too much garbage (i.e., minimal paging of data, few dynamic objects, etc.) their marking phase time can be kept relatively low (e.g., < 50ms) and absorb spikes with the buffered frames. They can also select transition points in a game (such as entering an elevator, loading a new level, etc.) to execute GC and maintain a game's fluid experience.

CAE implemented several changes, starting with the default UE5 latency. The game engine's default "swap chain" is too deep and needed to be shortened. Several settings are available when using the UE5's VR mode, but thread synchronisation in VR mode does not synchronize rendering on the right event. It is unclear if this is a conscious choice from Epic Games or an oversight. To resolve this, we modified the UE5 source code, and corrected the swap chain and UE5 thread synchronization. The resulting system now has proper thread synchronisation with a total latency equivalent to a traditional IG (i.e. cull + draw + scan out).

Because now we cannot rely on a deep swap chain, the non-determinism of the garbage collection is a significant issue. We explored three alternatives to resolve this: the first option was to manage data to avoid creating many objects, pre-allocating and reusing the objects frequently. This requires budgeting of all needs and restricting functionality that did not fit the model. It was not considered to be viable.

The second option was to regroup data in "clusters" of objects. This is a feature supported natively in UE5. Using clusters, the marking phase is significantly reduced as the total number of clusters is lower. Due to database paging, significant computing is required to create these clusters, which must be done at runtime. Unfortunately, this also blocks the UE thread, creating overruns and stepping. This approach was abandoned.

A third option involving a partial rewrite of the UE5 garbage collector was explored. The intent is to convert the "marking phase" to an incremental task, allocating a fixed budget (e.g., 1ms) to execute marking at each frame. CAE is developing this new library and linking it to the UE5 engine, limiting the areas where code changes are required, and working closely with Epic Games technical support. Fortunately, this problem is well known to CAE as similar issues existed in the original code base of Medallion, and the GC was eventually re-written to achieve deterministic performance.

2.2.4 Sensor Simulation

Most game engines don't have professional-grade multi-spectral sensor scene rendering. The rendering pipeline is focused solely on the visible spectrum. Shortcuts are easily implemented to fake a FLIR or SWIR rendering with clever shaders or post-processing effects, but this would circumvent true physics calculations required for proper simulation, such as emitted and transmitted radiation, reflected radiance, radiation transfer in atmosphere, etc.

Given this gap, CAE turned to 3rd parties specialized in sensor simulation. There are two major providers offering high-fidelity sensor simulation plugins for UE5: JRM Technologies and Presagis Ondulus. Both products offer complete and exhaustive multi-spectral simulation add-on for UE5. CAE has teamed with Presagis for its sensor component of the new SE visualization.

2.2.5 Other Features

In the interest of paper length, we will not elaborate in detail on other UE5 features. The team identified certain limitations with the native UE5 weather simulation, multi-channel synchronization, and 3D ocean simulation. CAE has analyzed these and identified workarounds that are compatible with the needs of a next-generation high-end image generator.

2.3 Modelling & Simulation

Much like the earlier discussion on visualization, CAE has also developed and maintained its own computer-generated forces system for nearly three decades and delivered the capability on over 100 military training devices, primarily in the international market outside the U.S. The adoption of UE5 for visualization offered a unique opportunity to explore the potential of a new CGF based on the UE5 technology stack.

Over the past several years, the gap between the gaming and modelling & simulation industries has shrunk. There are now many high-quality, open-world, multi-player and immersive games vying to replicate a real-world environment. In parallel, the modelling & simulation industry is demanding larger, more realistic synthetic environments to enable nation-wide, coalition-based joint exercises as well as theatre-level mission operations and crisis management support.

It is CAE's belief that a new CGF built on UE5 technology has the potential to accelerate the technology crossover and provide the modelling and simulation industry a step-change in scale and realism, while facilitating access to 3rd party micro-services and plugins.

The following chapters describe some of our observations during prototyping efforts with the Unreal Engine over the past two years, organized along the themes of scalability, capabilities, and marketplace access.

2.3.1 Scalability

CAE has struggled to scale its own server-level CGF simulation beyond several thousand entities. CAE's CGF, is built on traditional object-oriented programming (OOP) principles. The gaming industry has introduced a software architectural pattern known as an Entity Component System (ECS) as a core foundation of many gaming engines. ECS uses a data-oriented approach to take full advantage of hardware acceleration using a contiguous memory data arrangement and parallel processing constructs (e.g., Single-Instruction Multiple-Data, or SIMD architectures).

Many Triple-A game engines use this ECS technique to offer gaming experiences with a significantly increased scale. Game studios are battling to offer bigger, richer and more dynamic environments, as well getting more players to coexist and evolve in the same game world. For example, Unity has its Data-Oriented Technology Stack (DOTS), and Epic Games recently introduced Mass Framework in Unreal Engine 5. Figure 2 illustrates graphically how a data fragment describing an entity (such as aircraft position) is organized contiguously in memory compared to traditional object-oriented data.

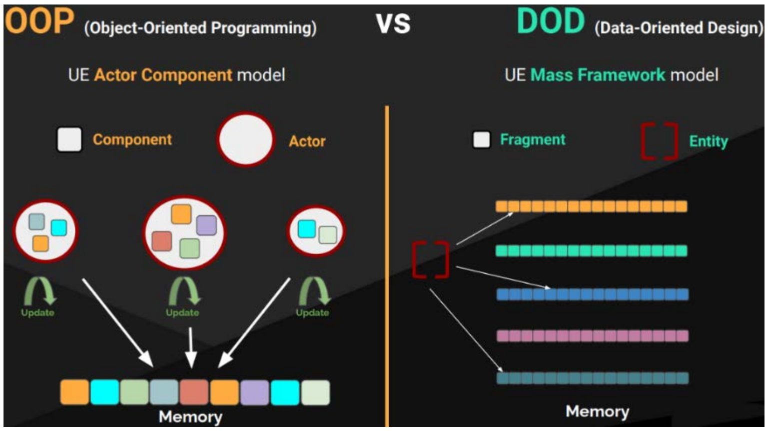

Figure 2: UE5 Mass Data Structure (courtesy of Epic Games)

This data restructuring is essential to achieve increased performance for a given instance of the CGF. In some cases, this is still insufficient to achieve the millions of entities that certain scenarios may require. To scale further, multiple coordinated instances are needed. This is the topic of the following chapter.

2.3.2 Scalability through Cloud Computing

There is a growing need for Synthetic Environments to realistically simulate vast crowds and so-called pattern of life behaviors in dense urban settings. Historically, crowd simulation involving millions of individuals used analytical, non-real-time models, often aggregated, with visualization limited to summary information (e.g., heat map, migration).

Some recent computer games have displayed crowds of 10,000 but employ significant over-simplification of each individual entity. This may be sufficient to generate a visually convincing picture, however such models could not be used to accurately simulate an evacuation or other large-scale crowd event. With the use of the scalable cloud technology, future SEs will be able to model and enable visualization of masses of entities.

CAE experimented with COTS spatial partitioning software to spread the CGF simulation over multiple cloud-based computing nodes. Notionally, these tools divide the virtual SE world into geographical areas (typically adjacent rectangular blocks) and each entity is assigned to its respective area. Each area is then assigned to a different computing node and is responsible for the simulation of all entities within its designated area.

Some partitioning software is also able to dynamically resize areas, splitting up denser ones and merging those that have less processing assignments. Experimentation confirms that this approach can effectively balance cloud resources. This task is notionally simple, but complicated to implement as an entity traverses a boundary and a handover is triggered between adjacent areas (and their related computing nodes), which requires both nodes to be aware of the respective entity during the transition.

Such spatial partitioning algorithms work best when there is little (or only short distance) interaction between the entities. In many computer games, the players of massive multiplayer games are usually on the ground, moving slowly and do not have long-range sensors looking at distant players. We foresee technical challenges in expanding this approach beyond civilians and vehicle traffic; however, CAE has not explored such use cases yet.

CAE successfully built a civil and a military focused demonstration, displayed in a map-like overview mode as well as an interactive, high framerate 3D environment. One of these demonstrations was in a city modelled with over 1 million buildings, 2 million dynamic entities across more than 100 Unreal Engine instances. Over 5000 entities were rendered within the field of view at 60Hz.

In another demonstration, depicted in Figure 3, an extremely dense crowd of civilians congregates at Trafalgar Square in London, as part of a simulated civil protest. Through improved rendering techniques, we were able to simultaneously visualize over 50 000 individually controlled entities

Figure 3: Dense crowd simulation leveraging cloud computing

Such large-scale demonstrators are compelling, confirming the effectiveness of cloud computer scaling, but they also push the limits of standard networking protocols and simulation interconnectivity interfaces. Specific solutions to these constraints have yet to be implemented, but we see new topologies emerging from the gaming industry, and we will assess these as part of on-going R&D at CAE.

2.3.3 Game Engine Capabilities

Beyond the architectural changes described earlier, other adaptations are aligning to the needs of military-focused SEs. The client-server architecture of game engines typically supports single player, local and networked multiplayer games. Similarly, our needs in Defense & Security are standalone, linked and networked training.

UE5 and other game engines natively include physics-based models and toolsets for modeling object dynamics, collisions, deformations, destructions, and fluid dynamics. AI modules are also used for complex tasks such as navigating in a dense environment. They offer a robust toolset for programmers, gameplay developers, artists, and animators to efficiently develop game logic, character behaviors, as well as 3D content. Our own experience building multiple demonstrations over the past two years is a 10-fold increase in modeling efficiency compared to existing CGF tools.

2.3.4 Marketplace Access

Smartphone App Stores and open-source communities are two well-known development ecosystems that have created a worldwide community of innovation, resulting in a plethora of widely available, powerful and inexpensive capabilities.

For the modelling & simulation industry, this brings significant opportunities and benefits. Some, like Unreal Engine, have created a community of millions of developers, contributing to technical forums, quality assurance, as well as to the development of plugins and 3D assets available publicly through an online marketplace.

CAE has purchased many assets and plugins as part of our evaluation and demonstrations. Our engineering and content development teams are impressed. As an example, we purchased a fully articulated 3D Eurofighter modelled with over one million polygons for less that 100 Euros and integrated it in a scenario in less than a day.

2.3.5 Limitations

Our experimentation confirms that some key areas of gaming engines and digital technologies are not yet fully adapted to the needs of the modelling & simulation industry. We believe that much like our visualization efforts described earlier, a combination of key partner collaboration, our workforce experience as well as existing CGF code base and content will enable solutions soon. Here is a short list of areas identified so far.

Table 2: Limitations requiring resolution for Modelling & Simulation industry needs

Game Engine	Multi-server support (scalability)
	Integration in a system of systems architecture (data exchange, networking, performances)
	Whole Earth terrain ellipsoid model (WGS 84)
	Representative dynamics and navigation for all domains (land, air, sea, space)
	Models of military equipment (sensors, weapons, countermeasures, communications, etc)
	Comprehensive library of entity models, behaviors, and scenarios
	Support of M&S standards (e.g. DIS, HLA/RPR-FOM, C2Sim, OGC CDB, OpenFlight)
Scalability Platform	Fast-moving entities supporting long-range interactions and complex state & behavior
	User interactions with a large number of entities, or with fast update rates
Cloud Infrastructure	Complex system of systems architecture with numerous interacting applications, exchanging large quantity of data with low latency and synchronization requirements
	Time for dynamic provisioning
	Controlled and classified data

2.4 AI Modelling of Social Media Behavior

Simulating large-scale civilian populations in a pattern of life invariably requires proper simulation of social media. Numerous events over the past decade show the influence of social media on public opinion and its direct impact on an individual's behavior and actions. A singular event may trigger a cascade of social media interactions, becoming viral and reaching millions worldwide (e.g., a diaspora reacting to events in their homeland). Such events have a direct and measurable impact on population sentiment, affecting the outcome of complex scenarios.

In 2021, CAE simulated the Greater London area as part of a concept demonstrator. The scenario simulated a cyber-attack on city transformer stations, affecting the power grid. The fictitious scenario had a small group spread information over a simulated social media network, culminating in a large-scale gathering in Trafalgar Square.

The simulation could have simplified the social media component and have people join the gathering after a certain time, but CAE was interested in a deeper analysis of the cascading events and associated causality. As such, CAE implemented a 2-stage approach:

1. Generation of social media posts
2. Analysis of social media posts

In the 1st stage, CAE took a sample of real social media messages, trained an AI based algorithm, and used it to generate a high number of realistic messages. These posts included messages including emoticons.

The synthetized social media messages were then analyzed and tagged by different classifiers according to sentiment, subjectivity, and emotion. Such techniques are also used on real social media content by marketing firms to evaluate the success of a brand marketing campaign. These three factors, presented as indicators for the population, were visualized in different ways. Analysis of the simulation outcome confirmed the sentiment shift towards fear and anger resulting in masses joining the protest.

CAE's demonstrator showed the value of including social media simulation in crowd behavior and sentiment, although the topic is worthy of significantly more research and validation to ensure the simulated outcomes are plausible and representative of real-world conditions. This topic is extensively discussed in a CAE paper by Harris et al, I/ITSEC 2022 paper 22175 (published in December 2022).

2.5 User Interface Considerations

Current SEs are typically delivered with editors and interfaces for creating scenarios, doctrines, models, and special databases for weapon and sensor systems. Developing and editing terrain databases and visual models require specialized tools as well. Most systems also provide a means of visualization through a map view, displaying entity movement and scenario control fields.

A future SE must include these, but also be able to edit vastly increased entity counts, levels of complexity, and new simulation models (such as human population attributes, social media, etc.). Controlling and editing the behavior of millions of scenario entities will invariably be done through higher-level controls to automate content creation parameters and AI. AI is already playing a role in world-scale 3D content creation, with hundreds of millions of 3D objects derived from satellite pictures that are analyzed through machine learning, detecting building footprints and creating region-specific 3D cities. Application of AI in cloud-enabled, scalable dataset generation pipelines is essential for operational decision support applications that must maintain currency between the real world and its synthetic representation in near real time. Examples include humanitarian relief missions to areas so changed/disrupted that an extant dataset may not be of value.

While game engine interfaces are attractive and efficient, we expect that operation of next generation SEs in our market will be addressed with a wider range of tailored interfaces for both users and developers.

3.0 SUMMARY

This paper has elaborated a comprehensive vision for a next-generation synthetic environment based on gaming engine technology, further enhanced through digital era capabilities of cloud computing and artificial intelligence. This new capability has the potential to unify earlier bespoke SEs that addressed large scale, aggregated entity representations for military command staff training, and those addressing real-time high-fidelity modeling associated with traditional military simulator needs. We believe the scale and realism of this new SE has the potential to fulfill upcoming needs of military multi-domain operations, population behavior modeling, next-generation system-of-systems development, as well as refinement of AI-based autonomous systems.

The paper has elaborated on key observations of the past 5 years, including experimentation, prototyping and detailed assessment of different gaming engines, including more recent efforts with UE5. When applied to real-time visualization, UE5 provides a compelling baseline of capabilities with advanced rendering, dynamic shadowing, efficient rendering pipeline as well as a massive – and relatively inexpensive – marketplace of quality 3rd party content and plugins. Several key limitations were also identified and described, including world scale positioning, anti-aliasing limitations, rendering performance determinism and sensor simulation. Each of these will require expert engineering teams to improve and/or rewrite sections of the UE5 software stack.

As a next-generation CGF, UE5 is also an attractive software baseline with an efficient and highly scalable data-oriented design with its Mass Framework. Model dynamics and realism have the potential to further improve by exploiting built-in physics-based models, character behaviors and AI-based navigation modules. Here again, a number of limitations have been identified that preclude out-of-the-box use, but solutions are foreseen as we continue product development.

Additionally, cloud computing and AI provide significant growth in capabilities and simulation complexity, and we have described use cases and experimentation results for each of these.

Our extensive evaluation and demonstration prototypes have confirmed that we are on the right path to offering our industry a more robust, cost-effective, and powerful synthetic environment for the needs of tomorrow.