

Interoperability Architectures

Mr. Jonathan Searle, Mr. John Brennan
Simulation and Synthetic Environment Laboratory (SSEL)
Defence Academy of the United Kingdom
Cranfield University
UK

J.R.Searle@cranfield.ac.uk / j.m.brennan@cranfield.ac.uk

INTRODUCTION

The first paper in this series addressed the general concepts of interoperability within the modelling and simulation (M&S) realm. The next stage, within the context of the lecture series as a whole, is to examine interoperability from an integration perspective. To accomplish this, one must look at the manner in which simulations (and their components) can be *bolted* together. Furthermore, to ensure a complete assessment is conducted one must consider more than just the technical aspects of interoperability and integration; it is critical that one gives due attention to the organisational, cultural and economic aspects of integration as well. As such, one can state that the main challenges of interoperability and integration within the M&S realm are composed of a mix of engineering, technical and conceptual issues.

The primary aim of this paper is to consider the **integration** and **interoperability** of M&S with particular emphasis on **federated** and **distributed** simulations.

To address this aim it is important for one to begin with a basic understanding of the significant terms within the statement, in the context of M&S. First of all, the term interoperability is meant to imply the ability of a given system to exchange data or information, and perform its required functions in concert with other distinctly separate systems. The term integration refers to the induction or amalgamation of M&S systems into existing programmes or studies such as training and acquisition. Therefore, in *simulation speak*, interoperability relates to the creation of *federations* whereas integration relates to the application of the federations. Finally, one must also explore the issues of composability and architectures – two fundamental elements underpinning interoperability and integration within the synthetic environment realm.

Prior to launching into the main element of this paper, it is worth taking the time to briefly explore the concepts of federated and distributed. These two terms are related, in some cases, but they do not necessarily imply mutual dependence within the M&S realm. Federated systems are systems that contribute to a whole while each element maintains self internal management – a federation is an organisation formed by merging several elements. The term distributed can have two meanings within the context of this paper. First, the term distributed computing generally refers to a collection of independent computers each performing partial elements of a greater task, which appears to any user as a single coherent system. Within the M&S realm, distributed has a similar meaning except that it normally also carries the implication of physical distribution (e.g. geographically disparate). The main reasons for employing distributed simulation are reliability, scalability, connectivity and extensibility (or composability). Essentially, it is important for one to understand that the term federation (or federated) does not imply distributed, or vice versa.

COMPOSABILITY

Composability as a concept is not new. The idea of building a system from individual (proven) components has been common within engineering for many years (e.g. the automotive industry). Similarly, the software engineering world adopted this concept, particularly with the advent of object oriented programming techniques. If one examines the implementation of a simulation, we can see that composability exists in many forms at a variety of levels. Take for instance, flight simulation. One can purchase a commercially available, PC-based flight simulator *game* (e.g. Microsoft’s Flight Simulator 2004) and see that it is in fact a composition of many individual elements, or components, of software code. A component of code handles the flight dynamics of the aircraft, while other components handle elements such as weather modelling, terrain modelling and computer generated air traffic. Not too dissimilar, is the typical civil aviation full flight simulator. These complex systems are constructed in a conceptually similar manner except that rather than simply a collection of software components, they are a collection of components, each comprised of hardware and software. Nonetheless, individual components will still perform specific functions as mentioned above, and the various hardware components will be networked to provide the user with the illusion that he is operating a single, coherent aircraft.

One can take this concept of composability to even higher levels wherein multiple platform simulators/simulations are linked together to provide a common mission space. This permits small team training, such as that necessary for flight elements or packages in the military fast jet arena. Extend yet again by bringing in land, maritime and C4ISR¹ assets, and one can create (compose) a joint simulation space with varied combinations of assets. Figure 1 attempts to capture this concept.




Figure 1: Scope of Composability.

Petty et al define composability as “... the capability to select and assemble simulation components in various combinations into simulation systems”^[1]. Their paper on composability goes on to describe two types of composability. The first is syntactic, or what can be considered engineering level composability. At this level one is most concerned with the technical aspects of connecting components through data interfaces or invocation mechanisms. For example, if one simulation provides position as an output, then this level of composability will be concerned with ensuring that other networked simulations or components can receive and interpret the position data packet when it is sent. The second type of composability mentioned is semantic, or modelling composability. At this level, consistent assumptions

¹ C4ISR – Command, Control, Communications, Computers, Intelligence, Surveillance, Reconnaissance.

and domain validity become important. That is, one is more concerned with whether or not the composition of two or more models is actually meaningful. For example, if one model issues an order to attack, do the other models that might receive this order *understand* what is meant by the attack order.

Essentially, the desire to assemble individual models or simulations as components of a larger system is certainly not a new concept. Nonetheless, one must be sure to address composability at both the technical level and the semantic (or modelling) level. For if there is no logical reason or meaning behind attempts to network individual components, no amount of assurance at the technical level will guarantee successful execution of the final system.

ARCHITECTURES FOR DISTRIBUTED SIMULATION

With the concept of building higher order systems by connecting multiple components having been examined, it is appropriate now to begin exploring the ways in which one can envision organising the components to facilitate a logical assembly. There are different conceptual approaches for categorising the elements of a distributed simulation (or federation). Two possible methods are to distribute by function and distribute by platform.

When distributing by function, one can begin by categorising the component models or simulations by military environments. For example, in a joint simulation system one could have a land battle component, a maritime battle component and an air battle component. Similarly, there would be components for such things as weather modelling, terrain modelling, ballistics modelling and sensor modelling. Figure 2 depicts this concept, which schematically represents the construct on which the Joint Training Confederation (JTC) is built. ^[ii]




Figure 2: Distributed Simulation by Function.

Another approach to distributed simulation, one that is likely more familiar, is by platform. This architecture is essentially a peer-to-peer distribution of equals (i.e. of similar components). Each station (or node) generates platform-level entities and all simulations/entities communicate with each other. One example of this approach is the IEEE 1278 standard known as Distributed Interactive Simulation (DIS). The DIS standard is based on a message passing protocol wherein ground truth state data is passed via protocol data units (PDUs) such that all participants potentially know all ground truth. This approach has benefits in the training domain but it is not very flexible beyond that.

Architectures for distributed simulation follow the conceptual examples above. First, the distribute by function concept is implemented as a network of functional nodes (see Figure 3). This architecture is optimal for applications other than training, such as the work and research done by industry or other like organisations. The primary advantage is that this architecture is relatively more easily reconfigured to

address a variety of studies or applications; as such, notwithstanding that it can be more complex, it is deemed more composable and has greater potential for reuse.

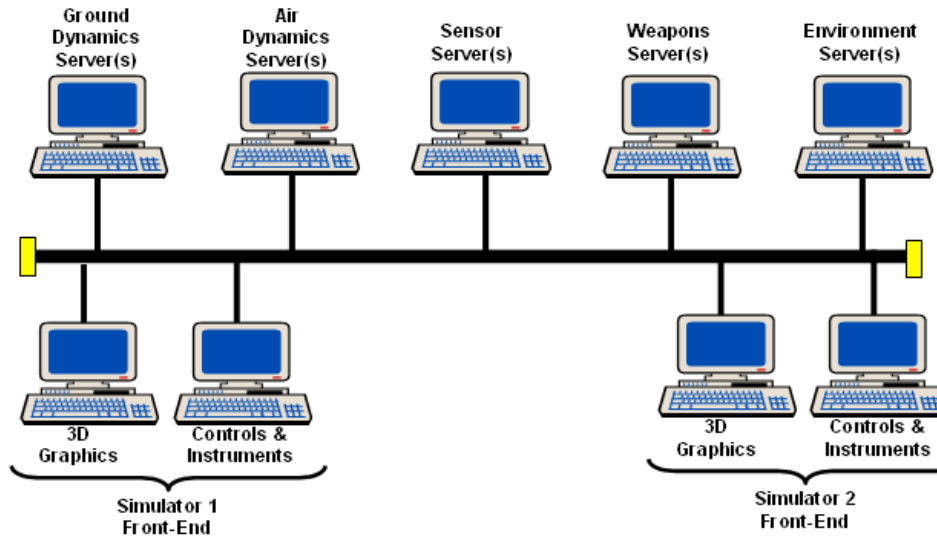


Figure 3: Functionally Distributed Architecture.

The distribute by platform concept is represented by a discrete, monolithic architecture (see Figure 4). Essentially a network of single computers or simulators, this architecture is good for training but not very flexible or scalable (i.e. one cannot go on adding more and more nodes *ad infinitum*). The systems that participate in this style of architecture are fairly homogenous and any modifications made to one participant may necessitate source code changes to all federates.

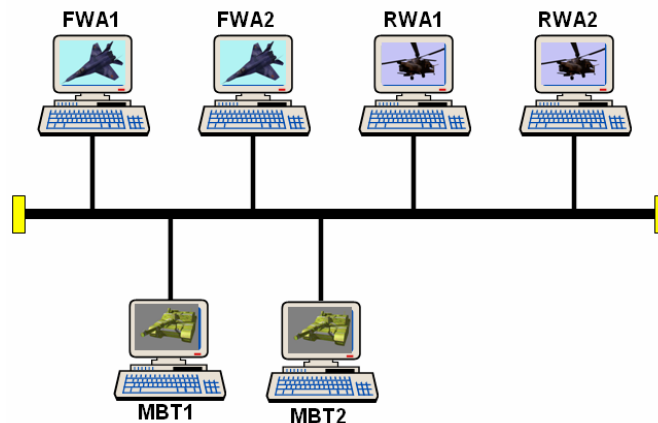


Figure 4: Discrete (Platform) Distributed Architecture.

INTEROPERABILITY AND INTEGRATION

When looking to build a federation, it is clearly a good starting point if each of the federates involved seemed to represent the same sort of things in a similar way. So, as an example, consider two hypothetical battle group level, ground combat simulations, which both represent individual vehicle platforms:

- An older legacy model, developed for stand-alone analysis use; and
- A newer model, developed from the outset as a training component.

One might ask the question, “How interoperable might they be?” Is it possible for the older simulation to be modified to make it interoperable with the newer system (i.e. to achieve reuse of the old model)? Furthermore, if this were possible, would it be sensible to do so?

The more traditional combat model was written in an older programming language, using traditional coding methods and styles. It was designed as a closed system, and therefore owns, controls and updates each and every entity within the environment on its own. Being a discrete event simulation², the speed of execution for the older system is complexity dependent. If executing in real-time, the simulation executes a small *battle slice* as quickly as possible (**much** faster than real-time), then simply pauses (sits idle) to let real-time catch up. At a lower level, the advances in processing actually occur in *bursts* (see Figure 5).




Figure 5: CPU Usage & Event Processing – Older System.

On the other hand, the newer model, like most current day systems, was designed and programmed in an object oriented manner. Although still event driven, it was written as proper real-time software, incorporating cyclic *callbacks* with very small time steps (see Figure 6). Consequently, the system executes consistently with relatively smooth processing and time advance.




Figure 6: CPU Usage & Event Processing – Newer System.

By way of example, the preceding two paragraphs begin to draw out some of the differences one can encounter when attempting to make two or more simulation systems interoperate. The following table identifies additional instances of the more common interoperability issues one might come upon.

² A discrete event simulation is one in which the state variables are updated periodically, in distinct steps. Discrete event simulations are typically time driven or event driven.

Typical Simulation Interoperability Issues		
Item	Legacy Simulation	Newer Simulation
Terrain	Matrix of cells	Often polygon based (but does not have to be)
Vehicle Orientation	Location & general orientation	Vehicle x, y, z and hdg, pitch, roll (i.e. 6 DoF)
Vehicle Space Occupation	Abstract (i.e. point mass)	Occupies physical space (i.e. bounding box)
Movement	Updated in discrete time slices	Updated more continuously (smoothly)
Engagement	Probability table (Hit & Kill); model determines all effects	Firing unit calculates ballistics & impact point; target calculates effects

The two simulation systems described are relatively similar in that they are both used to model ground combat; however, one can clearly see that there are issues which must be addressed. Furthermore, what if the simulations were fundamentally different in application; one would expect that the issues would be even more pronounced.

To be successful, interoperability requires due consideration to a range of issues. The initial problem is to establish how simulation state data will be shared when hindered by the physics, characteristics and topology of networks. Even with modern networked systems, one cannot underestimate the impact and costs of these issues. Compounded on this are the issues of security and intellectual property rights; these are particularly crucial in the military domain, when dealing across borders or with industry. But beyond this, there should also be concern regarding the meaningful interpretation and use of the data to be exchanged by each simulation and tool involved; technically two systems may be able to exchange data and even then understand it, but having done so, does it actually make sense? Essentially, in addition to considering the implementation level one must consider the conceptual level of interoperability that underpins the activity.

When considering the conceptual level, architecting interoperability must account for natural groupings of models and associated levels of interoperability. The typical hierarchical pyramid of models (see Figure 7) helps visualise natural groupings and the associated application areas of the wide expanse of models available today.




Figure 7: Hierarchical Pyramid of Models.

A second (or indeed alternative) way of establishing potential opportunities for interoperability is to use the approach that is now being termed *communities of interest* (COI). Essentially another way of classifying high level disciplines, COIs can represent groups of people or organisations that are focused on areas such as research and development, training or command and control (engineering, platform and aggregate foci as depicted in Figure 8).




Figure 8: Communities of Interest Concept (Graham Shanks, AMS, 2005).

One final element worth mentioning is the effort put forth by the European Cooperation for Long-term In Defence (EUCLID) research and Technology Program (RTP) 11.13 – the Synthetic Environment Development and Exploitation Process (SEDEP)ⁱⁱⁱ. Dubbed *Realising the Potential of networked Simulation in Europe* the SEDEP efforts aimed to provide tools and processes to mitigate obstacles to effective use of synthetic environments (SEs) within Europe. In short, the SEDEP program recommended

that, in order to enhance collaboration and reuse in a complex environment, one would need to provide processes and tools interfaced to repositories, and shared data environments and knowledge bases. Conclusions such as these tend to emphasise the fact that the use of SEs is problematic because of integration as much as because of interoperability.

SUMMARY

Networked and distributed simulation federations have come a long way since the development of SimNet in the early 1980s. The scope has increased dramatically necessitating much more complex architectures. The increased complexity has encouraged the practice of composability and reuse to help ease and facilitate the construct of complex federations. The architectures must account for not only the technical and engineering aspects but also the practical and conceptual characteristics of a federation. No single architecture can provide an all encompassing solution for all applications – each requirement must be approached independently in the first instance – then one can look to composability and reuse as means of providing efficiency.

REFERENCES

- [i] Taken from *Overview of a Theory of Composability*, Petty, Weisel & Mielke.
- [ii] See <http://alsp.ie.org/alsp/joint/index.html> for more information.
- [iii] See www.euclid1113.com for more details.