The Real-Time Data Company Ltd.

Introduction

This document's aim is to explain, at an introductory level, the differences between simulation techniques used for training purposes, and those needed to perform test and evaluation of systems. It uses visual and radar application as examples, but does not assume in-depth knowledge of either.

The defence sector, like many others, is facing challenges in the 21st Century. The challenges include increased complexity of demands and requirements. Limited funding seems not to reflect those increases. The complexity of requirements exacerbates the need for speed of delivery and flexibility. When a defence challenge is identified the population is at risk until the challenge can be countered. Complex equipments that take years to produce are not satisfactory solutions.

Examples of threats met by the military in recent years include 'surprise' invasion of one country on another, rapidly increasing cyber threats through to a disease pandemic. The effects of these ripple through to ordinary people by way of shortages, with food, electricity supplies and water, being amongst the most basic. Something needs to change and simulation is the 'something' that is being introduced.

Taking a long time to build and test a solution, then longer to train how to use it, runs contrary to the need for rapid in-Service response. Previous 'in the field' testing methods are generally expensive, slow to run and are limited in cases where undertaking them would place people or property at unacceptable risk. They find issues that would have been better identified much earlier. Simulation offers a continuous process of being able to test and evaluate the status of a project, against its requirements, at all lifecycle stages.

Synthetic environments can be understood from the world of gaming. Empires can be built and human players allowed to 'live' in these virtual worlds. They encourage and support experimentation in ways that interact with the virtual environment. Tools are made available in these virtual worlds to support interactions and these are not limited to those experienced in real-life. Clearly this has a strong parallel with training needs and exploring new concepts. It is also relatively quick to produce at relatively low cost – what's not to like!

Military training has exploited this, a prime example being VBS produced by Bohemia Interactive Solutions[1]. The virtual world is recognisably 'real' and the tools include representations of actual vehicles, weaponry and people. The software runs on a laptop computer and produces very credible graphics to immerse players into the scenarios being played out. The cooperative features of game playing are attractive to this type of training application.

Training aspects have had at least 15-20 years to get to this position. The Defence Chiefs have decided it is now time to consider how to introduce and test the new tools that the players in the virtual world can readily imagine. It is not too difficult to introduce new concepts to the training environment and 'test' them against threats. Their requirements and performance can be established and tweaked quite quickly. Then comes the need for implementation in the real world - and that is more difficult. Exactly why it is more difficult is expanded in this document.

[1] https://bisimulations.com/products/vbs4#:~:text=What%20is%20VBS4,fast%20and%20effective%20skills%20enhancement.

The Real-Time Data Company Ltd.

Training vs Real-World Implementation

To be clear, a real-world application needs testing of any mix of SW models through to physical implementations of the intended new product, as its lifecycle progresses. Typically, this requires raw signals to be generated by running scenarios with multiple models in a synthetic environment. In the early stages the signals stimulate a model of what is being proposed. In later stages the signals stimulate the physical entity (a radar system, a defensive aids suite etc.).

Training applications require only outcomes to be simulated.

> *Turn a steering wheel and the view changes. How quickly it changes is a matter of adjusting parameters in the simulation.*

In the real-world, many things happen in the vehicle to achieve that effect; engine load changes, drive gear loading changes, tyre-ground physical rules apply.

> *The outcome is achieved according to the application of physical laws and this is the challenge for simulation of a real-world implementation.*

Physical laws can often involve complex chain of sequences that are mutually inter-dependent. This presents a very different computing challenge to that required for games-based training. Instead of producing a simulation of outputs, it is necessary to simulate the inputs that will produce them, then applying the physical laws to see how the unit under test responds.

The outcome-only approach is referred to here as a training application and the physical modelling approach is referred to as a real-world application and can be used for test and evaluation of the physical model through to its real-world implementation. The concept of Digital Twins is derived from this, one twin being the model (usually implemented in software) and the other twin being its real-world implementation.

Simulation has two basic components, a model of the equipment to be evaluated and a model of the environment that generally will include other (friend and foe) models. In the most flexible implementations these are separated. This is true for training as well as real-world applications. Both also involve the use of scenarios that define how the equipment model interacts with the environment – usually to achieve an end that defines the operational requirement to be met. The iterative nature of this can be appreciated.

Some scenarios support human intervention, others are more like a replay of a pre-recorded sequence. Both have their place. The interactive option has exploratory and training benefits. The non-intervention option is better suited to repeatability for purposes of (for example) gathering statistical information, or changing such things as weather conditions to assess performance envelopes.

There is another significant difference between training and real-world applications. Clearly, training applications involve interactions with humans. Real-world applications may involve humans but the main interactions are between the equipment and the environment. In the case of sensor systems, the equipment may operate anywhere in the electromagnetic spectrum, not just the visual range. To illustrate the extent of differences, we will consider visual and radar sensor systems.

The nature of visual and radar imagery

Consider a simple scene as shown in Figure 1. The outer cube is the viewing area that contains two 3-D bodies. The viewing direction is from the front as shown by the arrow.

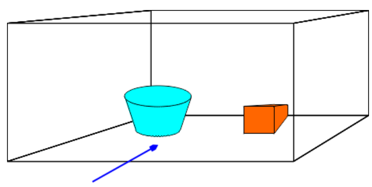

Figure 1 - A simple scene

The Real-Time Data Company Ltd.

The visual view of this scene would be as shown in Figure 2. The two bodies are each shown as a 2-D shape. The 3-D perspective (depth) is only achieved by changing the viewing angle and comparing the two pictures. The angle can be changed horizontally or vertically (azimuth or elevation).

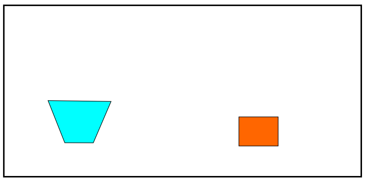

Figure 2 - Visual representation

Radar Sensing

To consider the radar view of this scene let us first understand the nature of radar illumination and the returns from the objects being illuminated.

Imagine a single very short radio frequency (RF) pulse being emitted from a point-source – although, in practice, it is more complex than this. The signal will radiate spherically, expanding with time and distance. Because distance is proportional to the time that the pulse has travelled (related by the speed of light), time and distance are interchangeable as units of measure. Radar systems have antennas to focus the emissions into a beam to increase the energy into the required direction instead of spherical transmission, although spreading is inevitable.

When the RF pulse encounters a solid, stationary, electrically conductive material, a current is induced and this results in an RF emission, at the same frequency but different amplitude and phase, emanating from the object. This is the 'return signal', or the 'radar echo return' from the object. It can be detected by a receiver tuned to the appropriate frequency. The frequency of the return will differ from the transmitted frequency if there is relative motion between the object and the emitter. There will be a delay between the time the pulse was emitted and the time the echo is received, proportional to the 'out and back' distance between the transmitter and the object that produced the echo. This is depicted in Figure 3 to Figure 5.

Figure 3 depicts a radar antenna and a display of the returned signal (right of picture) with the contact (a ship) to the left. The pulse is depicted as a spherically spreading wavefront shown as a green line in the centre of the picture.

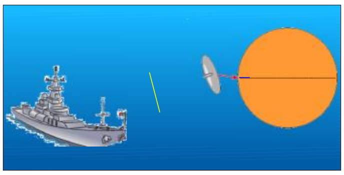

Figure 3 - Radar transmission

Figure 4 depicts a little while later. The transmitted pulse has progressed in time/distance and two return echoes are shown propagating, with some energy travelling towards the antenna.

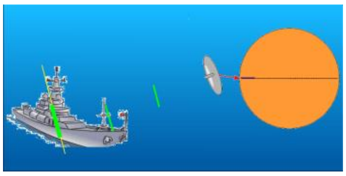

Figure 4 - Radar return 1

In Figure 5, time has progressed yet further and the outgoing pulse has moved beyond the contact and is out of sight. However, the 'last echo' is arriving at the antenna and the orange display shows the history of the amplitude of the echoes with time (distance). Observing the horizontal axis of the display (time/distance), the amplitudes (vertical direction) of the returns are a function of the 'radar cross-sections' (RCS - a measure of radar reflectivity and phase characteristics) of the contact and range. Clearly, the signal can be very weak on arrival at the receiving antenna because it is inversely proportional to the 4^{th} power of range (due to range-squared spreading of out-bound and inbound (echo) signals). Note that time of travel is out and back, so out-distance (range) is proportional to half the time on the display.

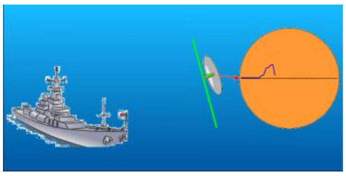

Figure 5 - Radar return 2

Radar Antennas and the Environment

Environments include many objects (building, trees, hills etc.) that also reflect in many parts of the electromagnetic spectrum – including radar. When used with radar systems, the many RCS components in the synthetic environment must also be implemented appropriately to contribute to the creation of the 'clutter' return. This is expanded in the following text to help establish the demands placed upon producing a synthetic environment to appropriate fidelity.

The two main types of radar antennas are 'traditional' rotating beam shapers (e.g. commercial maritime radars) and phased arrays that may rotate (e.g. Type 45 Frigate) or have limited movement relative to the body of the radar platform (e.g. Eurofighter Typhoon). For purposes here, the relevant factor is that they do not all have perfect narrow 3-D pencil beam shapes. A plan view might be as shown in Figure 6. The main beam is clear, the sidelobes add confusion to the interpretation of the echo returns. A strong echo at the angle of a sidelobe can be mistaken for a weaker signal in the main beam and this is an exploitable feature in some defence systems.

Figure 6 – Beam shape

The beam shape is shown overlaid on a map in Figure 7. The main beam is pointing east of north and the beam lobes can be seen.

Figure 7 - Beam View

Figure 8 - Sample View

The two arcs (actually, 2-D parts of a 3-D spherical shape) shown in Figure 8 intersect the 3-D beam shape. Each arc represents a range (time) sample following the transmission of a radar pulse. The return represented by the arc contains 'clutter' and 'object' return elements from the main beam and side lobes.

Each return step produces a signal that is a mathematical convolution of the actual transmitted waveform with each reflecting point at the relevant 3-D range. The gain of the radar varies with the position of reflecting points in the beam. The outline shape shows the beam's 3dB gain points. The phase will also be modified according to the design of the antenna.

In practice, there would be not two but many such arcs. Their number and their width being a function of the radar's analogue to digital sampling processes. It is an extension of the concept illustrated in Figure 5 above.

This is a complex calculation and the complexity increases as the required fidelity increases.

Radar Image Processing

A radar return is a single stream of amplitude-phase information, usually in the form of in-phase and quadrature (I-Q) signals. But this is not an image. 'Corner-turning' is the technique employed within the radar processor to generate images. To provide better understanding of the levels of processing needed for real-world simulations, the process is explained below.

Figure 9 depicts a radar return signal following the emission of a single pulse and is equivalent to that on the orange display in Figure 5. The horizontal axis (time-range) is as previously described. The vertical axis is also time (coarse), but on a longer timescale, a single vertical step being the time between transmitted pulses.

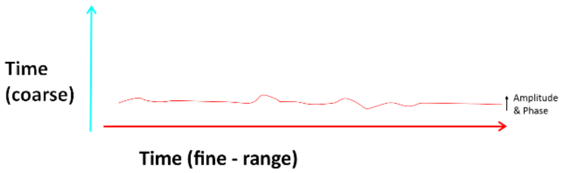

Figure 9 - Range - amplitude radar return

Figure 10 shows many returns and small variations in amplitude can be seen.

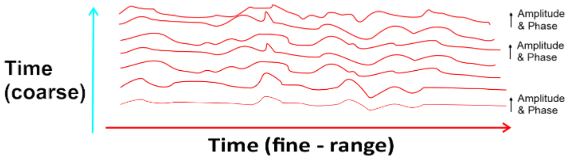

Figure 10 - Multiple returns

'Corner-turning' is the process depicted in Figure 11 where samples in range are taken from the returns of many pulses. If both the contact and the radar platform are stationary each return would be identical. However, if either or both are moving relative to each other there will be differences as shown. In this way, a series of 2-D images can be formed.

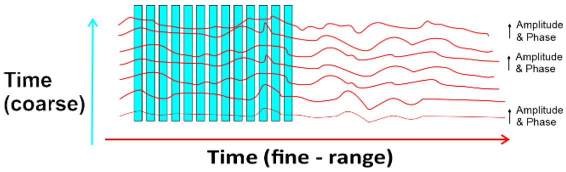

Figure 11 - Corner turning

Radar Imagery

We can now understand how the radar return from our simple scene would appear. The series of images would be as shown in the series of images in Figure 12, the sequence being from left to right.

Figure 12 - Series of radar images

The first image would be blank because the outgoing pulse has not yet reached the bodies. A little later, this will occur and the return pulse will arrive at the receiver causing the 'red' body to return its signal according to its RCS. The blue body has a sloping shape (varying with range) so only the RCS of the closest part (the top, in this case) will produce a return as depicted in the 2nd image. This process will continue with the 'red' body returning no more data because there is no outgoing pulse in-line with it and the 'blue' body return is from the RCS of its middle part. The last image shown is from the furthest-away part of the 'blue' body (the bottom part). After that there will be no further returns.

In the real world, the return range steps in this description would be a continuous process rather than steps. However, 'steps' are involved because the incoming continuous echo signal will be sampled and digitised by I-Q analogue to digital convertors. This raises two important points:

1. The sampling rate for digitisation defines the resolution of the image (time steps = range steps).
2. The RCS model of the image needs to match the resolution of the digitisation process – resulting in the model of the contact comprising many small elements of RCS data (amplitude and phase vs frequency).

Data Rates and Timing

The differences in processing needs between training and real-world applications result in very different data rates and timing requirements. Table 1 presents a comparison between a visual image generation for training applications and a radar simulation of raw data to be used to stimulate a real-world application.

Spectral Range	Performance	Data Rate (approx.)
Visual	• update rate of 100 frames/sec (10 msec/100Hz). • 4K video.	100Mbits/sec
Radar	• resolution 1 metre (an average performance). • 12-bit I-Q resolution.	14.4Gbits/sec

Table 1 - Data rates

The data rates are approximate but the difference (100 Mbits/sec vs 14,400 Mbits/sec) is clear. In both examples, the data rates are those resulting from the processing required to achieve the output streams. In addition, the processing loads to achieve data streams at these rates are high for both, but the radar processing load is much higher than for video.

There are two other important considerations:

1. Latency – the time required to have the return's first data point available following the transmission.
2. Availability of processing equipment.

Latency: A video image is required at around 60-100 Hz (16-10 msec) because a human eye is not able to distinguish changes faster than that. A 1-2 msec delay in producing the data would not be noticeable.

A radar pulse repetition frequency is measured in msec. That is the rate that a pulse will be emitted, after which the radar receiver will expect incoming echoes. Radar waves travel at the speed of light ($3x10^8$ M/sec) so a delay of 1 msec in receiving the first return data represents a range of 150 KM. That is the range for which no contacts would be reported. It is clearly unacceptable for a simulation because it does not occur in a real system where latency for radar returns is measured in nsec, if not psec. In terms of processing data for a simulation, that is a very short time to convolve the outgoing pulse shape with the sum of RCSs in the first range return's spatial volume.

Processing equipment: Training applications are a sub-set of the gaming applications and the market for these is large (projected by Statista to reach US$221.40bn in 2023[2]). This market-size has resulted in specialist video cards being available for PCs that are designed to meet video processing needs.

The real-world application market is miniscule by comparison and specialised cards are not available. However, video cards may be used or, for greater efficiency, Field Programmable Gate Array (FPGA) cards are available. FPGAs are specialist devices and have a cost-penalty because of their low-volume production, the required learning-curve and their complexity of use.

[2] https://www.statista.com/outlook/dmo/digital-media/video-games/worldwide#revenue

The Real-Time Data Company Ltd.

Spoof-Proof Issues

A real system (under-test) may be able to detect anomalies and decide to ignore a contact because it is a false target or a decoy. One way of doing this is that the system will create a signature of each contact and if it does not match a database entry it may be ignored or saved for post-contact deep analysis. That can lead to it being added to a database as a known decoy and for it to be ignored thereafter.

Some detection methods include:

1. Incorrect aspect ratios of a contact when viewed from different angles.
2. Inappropriate speed, acceleration or dimensions.
3. Inappropriate, or no, micro-Doppler where this is expected.

So, what does this mean for producing a synthetic environment or model for use in a test environment?

- **Aspect Ratio:** The models must be of sufficient fidelity to meet the requirements of the applications that use them. The sampling rate of the radar must be matched by the resolution of the synthetic environment and the models (contacts) contained in it. Also, the angular resolution of models must be such that these criteria are met – this can mean building RCS models with considerably less than 1° steps in azimuth and elevation viewing angles. That represents a lot of data for a large contact such as a ship.
- **Motion:** The contacts will generally be moving around in the environment and this allows their speeds and (3-D) accelerations to be calculated by the observing radar. These should fall within the envelope for the type of contact being represented. Position relative to ground can also be important. It would be inappropriate for many contact types to appear to fly, sink or rotate to some aspects (tanks do not fly or rotate to vertical positions whereas a helicopter might). These issues can arise with some representations of the environment with low-resolution position information.
- **Micro-Doppler** is a useful target identification parameter. For example, the rotation of propellers, turbines and helicopter blades have significant effect on radar returns and hence signatures generated from them. A flying helicopter with no micro-Doppler blade signature would not be realistic.

The impact of these issues on a synthetic test environment is to increase the processing load. So, each application needs to be understood to provide appropriate (value for money) processing power and fidelity.

SE/ Modelling Implementation Considerations

The use of simulation and synthetic environments for test and evaluation is an emerging topic. There are un-trodden paths to be taken. The following design considerations are amongst those to be considered, and satisfactorily resolved, when producing synthetic environments and the object models to be run within them:

1. Radar processing will require multiple processing streams for it to meet the total processing load. The architectural design to meet the requirements will determine the functional contributions of the streams and how many are needed. However, not all streams will produce their results simultaneously, or even at a consistent rate for all situations. This presents a synchronisation problem when summing results to produce the stimulation signals to the sensor under test.
2. The synthetic environments and models within them need to have user-friendly set-up and run features if they are to be accepted for general use. This topic includes the following aspects:
    a. Commonality of Application Program Interfaces (APIs) are one way to improve useability and reduce learning time. Work is required to identify what is required in this area to standardise the APIs.

b. Scenarios require models and synthetic environments to be merged. This raises issues of appropriate security levels being available, usually meaning fidelity levels, so that a compatible configuration can be built. Middleware is available to help manage these aspects e.g., Pitch[3] being one of the good examples for HLA and DIS architectures.

c. Consideration of how much real-time control is needed e.g., run-time changes in weather or manoeuvring of contacts in response to the unit-under-test activities. (This might include avoidance manoeuvres, jamming, or launch of weapons.)

d. Post-activity analysis may be required and this raises issues about what parameters need to be monitored during the running of scenarios, how they are captured and to what time resolution. It also raises issues of meta-data being generated during recording so that key activities can be located for replay.

3. Some military systems involve integration between multiple sensor systems, either at a sensor level or at a system level. Clearly, if these systems are to be tested in a synthetic environment it will be necessary for the environment and models to run in a coherent multi-spectral way. Imagine, for example, a radar signature indicated that a contact was likely to be one object but its RF or visual signature was seriously incompatible. The fusion processes on the platform would be likely to reject the object as a decoy and the platform's response (e.g. weapon release) would not be appropriate for the intended scenario.

4. Dual use training and T&E could save time and money as well as improving training realism. There are choices about how the lower-cost training environment might be integrated with data from a T&E environment. Consistency between the outcomes from the two applications is important.

5. Data Management issues:

    a. Meta data will be required so that models, scenarios, previous tests and their outcomes can be accessed. This is a huge topic involving wide stakeholder engagement to ensure all needs are covered in a cost-effective way.

6. It is likely that operation over several geographic locations may be needed for reasons of equipment availability, personnel availability as well as to address some security issues. This presents linking challenges that include compounded latency and data-access security issues.

7. The simulation equipment is likely to be replicated in part or in full so that multiple sites and multiple projects can be run simultaneously. Maintenance and updating needs consideration so that comparable results are obtained whenever and wherever a scenario is run.

Conclusions

There are big differences between models and synthetic environments used for training, and those needed for T&E. These differences strongly influence the processing power needed to support them. Training applications have advanced to a level that is well up the development curve, whereas T&E applications are at a much earlier stage.

Processing power has now advanced to the point where it is becoming practical to produce a T&E environment adequate for many purposes – including full lifecycle support – see http://www.thertdc.com/Downloads/DefDigThreadv1.pdf. There is some way to go to develop the techniques and skills needed to fully implement and support T&E applications.

This paper has compared visual and radar systems but electronic warfare, cyber, infra-red line-scan, sonar

[3] https://pitchtechnologies.com/developer-studio/

and other applications are also important.

There is increasing interest in applying reverse engineering techniques to estimate the nature of the body that produced responses to illuminating signals in any part of the electro-magnetic spectrum. This has application in auto-generating models of contact bodies for use in simulation environments and other applications.

This is an exciting and rapidly growing area with many challenges for practioners in commercial, data management, mathematical and technical fields, as well as educators. It offers career opportunities in areas that are currently not fully identified but are gradually emerging – some of these are introduced in http://www.thertdc.com/Downloads/DTSkillsv1.pdf.

Dave Murray – The Real-Time Data Co. Ltd.

Dave-M@TheRTDC.com

Jan 2023