

12th ICCRTS

“Adapting C2 to the 21st Century”

Paper: I-032

Title: Executable Architecture of Net Enabled Operations: State Machine of Federated Nodes

Topics: Modeling and Simulation
C2 Metrics and Assessment
Network-Centric Experimentation and Applications

Authors: Mark Ball
Joint Studies Operational Research Team
Centre for Operational Research and Analysis
Ottawa, Ontario, Canada

Ronald Funk
Joint Studies Operational Research Team
Centre for Operational Research and Analysis
Ottawa, Ontario, Canada

Richard Sorensen
Principal Systems Engineer
Vitech Corporation
Vienna, Virginia, USA

Point of Contact:

Mark Ball
Centre for Operational Research and Analysis
National Defence Headquarters
101 Colonel By Drive
K1A 0K2
Office: (613) 992-4539
Fax: (613) 992-3342
Email: ball.mg@forces.gc.ca

Executable Architecture of Net Enabled Operations: State Machine of Federated Nodes

By

Mark Ball, Ronald Funk and Richard Sorensen

The Defence Research and Development Canada (DRDC) Centre for Operational Research and Analysis (CORA) is developing capability-engineering analysis tools to support the building, demonstration, and analysis of executable architectures. Our paper to 11th ICCTS [1] described how to model workflows within an Operations Centre (OPCEN) employing a Net-Centric architecture. It used a State Machine (SM) model to simulate how multiple jobs can proceed in parallel when operators use Task, Post, Process, Use (TPPU) cycle to organize their work.

This paper extends the OPCEN SM model to track the interaction of work between OPCENs. The State Machine of Federated Nodes (SMOFN) model is organized around networked nodes that produce and consume products held in a virtual Repository. The data-driven simulation uses files to build customized job workflows and configure any combination of nodes without affecting the business logic. SMOFN also accounts for the following overhead activities:

- (1) Tracking consumer perception of product utility as it accrues and decays;
- (2) Consolidation of products into higher-level aggregated products; and
- (3) Triggering new jobs where needed whenever relevant products become available.

Customization of SMOFN is underway to account for the data and product flows between OPCENs in new Canadian Forces Command structure.

Introduction

Background

The Defence Research and Development Canada (DRDC) Centre for Operational Research and Analysis (CORA) is developing capability-engineering analysis tools to support the building, demonstration, and analysis of executable architectures. Our paper to 11th ICCTS [1] described how to model workflows within an Operations Centre (OPCEN) employing an executable Net-Centric architecture. It used a State Machine (SM) model to simulate how multiple jobs proceed in parallel when operators use a Task, Post, Process, Use (TPPU) cycle to organize their work. Further descriptions of previous SM work can be found in [2], [3], [4], and [5].

The OPCEN SM has been extended to account for Net-Enabled interactions between several OPCENs. In essence, the OPCEN SM accounted for the work done by a single OPCEN to produce several products based on data analysis. In the State Machine of Federated Nodes (SMOFN), not only are such production jobs tracked within several OPCENs, but the products they create are uploaded to a common, networked, Repository so that all products can be accessed and used by any OPCEN. The SMOFN is an attempt to capture the essential logic that governs the way work is conducted anywhere, but with particular emphasis on ensuring that it can be fully applied to networked OPCENs.

Primer on TPED vs. TPPU

Task, Process, Exploit, Disseminate (TPED) logic implies that each job is a serial process: jobs are worked on from start to finish without being interrupted. Flowchart diagrams (ie. in CORE) can simulate TPED by putting the job processes in a loop and repeating for each job. TPPU, on the other hand, allows jobs to be interrupted by higher priority jobs. As illustrated in Figure 1, updates of jobs processed through TPPU are posted at regular intervals and the utility of the job increases as more posts are made. This accrued utility is saved even if the job is interrupted. TPPU leads to concurrent behaviour (several partially processed jobs at a time) that is too complex to model using classical flowchart diagrams. The SM overcomes this by stopping time and executing the logic to assign operators to jobs in order of priority. It then steps forward in time and repeats the process. Progress within each job is tracked by recording the job’s state and updating it at each time step.

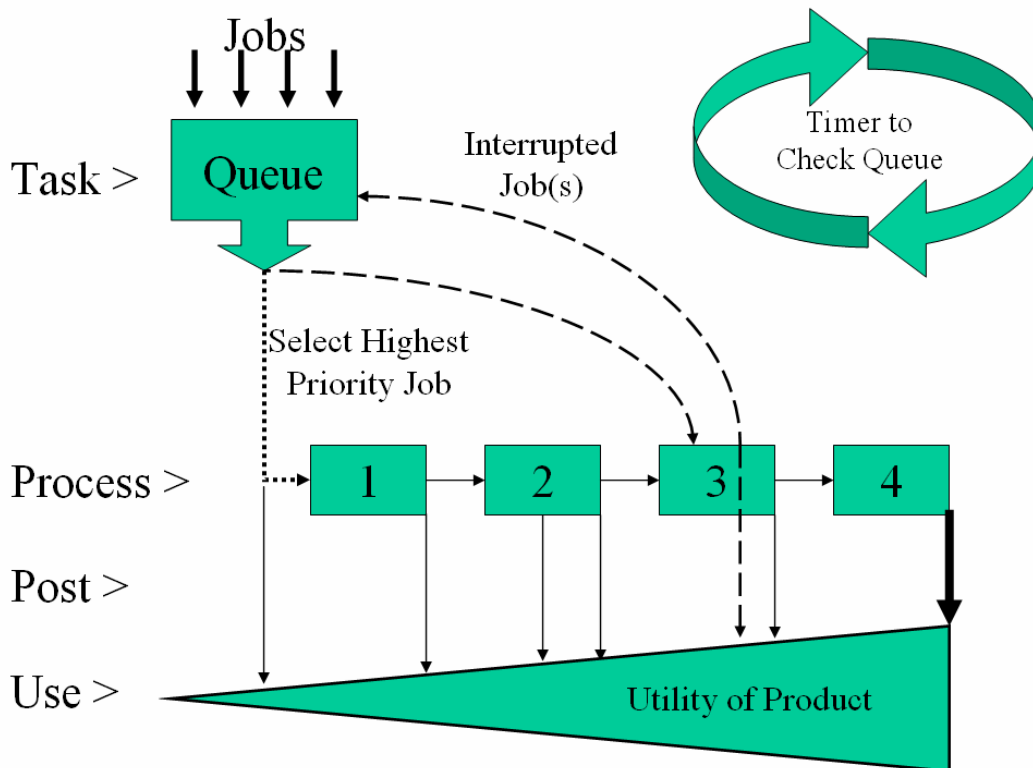


Figure 1: Task, Post, Process, Use (TPPU)

Conceptual Basis for SMOFN

Figure 2 illustrates the concept of SMOFN as a Department of Defense Architecture Framework (DoDAF) OV-1 (high-level operational concept graphic). It is a scale-free design in the sense that it is equally valid for representing interactions between nodes, or within a node, or even within a single operator's perspectives. The dashed lines show the transfer of Products, Questions, Queries, Results, and Responses, while the solid lines are a reminder that all of the transactions actually occur through the Repository, the interface to which is achieved using a Portal. This conceptual OV-1 diagram was also used as the basis for describing the logic that controls the interactions between the various nodes in the SMOFN. More specifically, it was used to define the interactions between each node and the Repository.

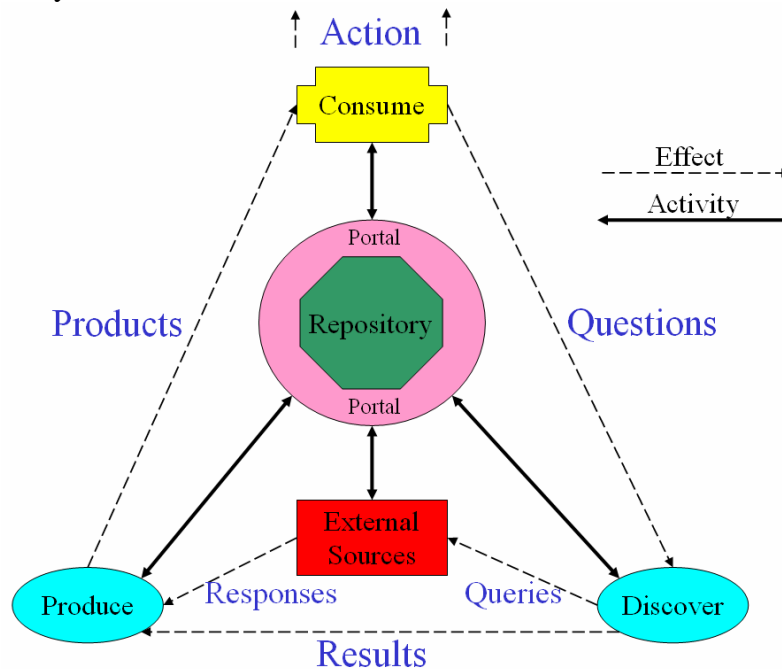


Figure 2: Scale-free OV-1

Producer-Repository-Consumer Model Logic

Figure 3 illustrates the ways in which Producer and Consumer nodes can communicate with each other and interact with the Repository. The most direct way for content to be transferred is for the Producer to push products directly to the Consumer. Note that during the actual execution the transfer will most often occur over a network with the Repository acting as the medium for accomplishing this. Alternately, the Producer may post content to the Repository where it waits to be pulled by interested Consumers. The logic options include allowing the Producer to notify the Consumer that content is available, or the Repository to use its business rules to decide when to notify the Consumer, or leave it for the Consumer to find when searching for new information.

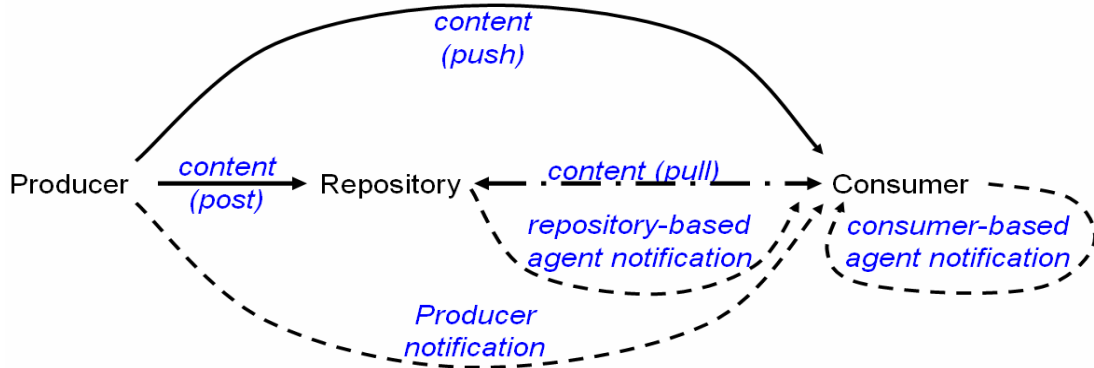


Figure 3: Producer-Repository-Consumer Paths

Operational Decision Making Logic

Extensions of OPCEN SM to SMOFN

In terms of the Producer node, very little has changed in the SMOFN model compared to the OPCEN SM described in [1]. Instead, the work concentrated on allowing users to create customized job threads. The OPCEN SM assumed every job involved 10 steps, and fewer steps were handled by assigning zero time to unneeded steps. Instead, SMOFN builds its job threads dynamically by assigning each step a unique name and linking any number of them together until it reaches a final step called “COMPLETE”. This ability to assemble building blocks also allows for a wider range of logic constructs to be used.

Producer Logic

The focus of the OPCEN SM was the coding of the logic controlling Producer threads. The SMOFN focus is to check for the arrival of new jobs at each OPCEN. Each OPCEN then uses the OPCEN SM logic rules to assign operators to jobs. First the jobs are examined in order of priority and assigned to available operators with the skills necessary to perform the job. Some jobs that are in progress from the previous time step can be interrupted if a higher priority job requires the same operator, or a job step considered to be generic can be interrupted to free a skilled operator so they can move to a job requiring their skill, leaving the generic job for someone else. Consumers determine the utility of any product, but Producers estimate the utility of jobs as they are worked upon. Jobs are abandoned when their utility decays faster than the operators can increase it. Jobs are also abandoned if they cannot be completed by their drop-dead time, which is specified separately for each job. If a job’s drop-dead time is omitted in the input files, it will instead be calculated based on the job’s arrival time at the queue and a standard allowable idle time for each job type.

Consumer Logic

In terms of SMOFN execution, the Consumer node’s job is very straightforward. Each time a product is received by the Consumer, there is a chance, based on the type of

product received, that a question will be generated. The Consumer starts by reviewing each product type and then generates questions after a certain delay time. These questions are sent to the Repository to be passed on to the Discover node.

Discover logic

The Discover node is responsible for answering questions generated by the Consumer. In real life terms, this tends to be done through a search for existing data. In modelling terms, it is handled through job threads similar to those used by the Producer. In fact, most of the logical scripting used by the Producer threads is shared with the Discover threads except for a few important differences. First, the Discover logic does not account for Utility because the Consumer has set the utility by asking the question. Similarly, the Discover thread always returns either a product that has some use to the Consumer or it tasks External Sources to collect it.

External Sources Logic

The External Sources node receives Requests for Information (RFIs), or Queries, from the Discover node. Exactly how the requested information is found is not within the scope of SMOFN. What is important is how much information is found (in terms of file size) and how long it takes to find it. The found information will typically be in the form of raw data, which is then sent to the Producer to trigger a new analysis job.

Repository Logic

The repository is the medium between all of the other nodes. Its job is to transfer products from Producers to Consumers, questions from Consumers to Discoverers, RFI's from Discoverers to External Sources, results from Discoverers to Producers, and responses from External Sources to Producers. The scripts to handle the logic of these data transfers are all similar, with nuances to account for differences between types of data or nodes, but more particularly for what type of information is relevant for the receiving node to execute its logic appropriately.

Implementation in COREsim

SMOFN Top Level Description

Figure 4 is the top-level diagram that controls SMOFN execution. It is colour-coded so each activity can be cross-referenced with the appropriate node in Figure 2. Unlike most diagrams in CORE, time does not increase as CORE steps through the process from the left to the right. In the case of SMOFN, time is actually stopped and the displayed left-to-right sequence represents the decision making process that takes place at each time step. Once the decision logic is completed SMOFN increments time another step and repeats the process.

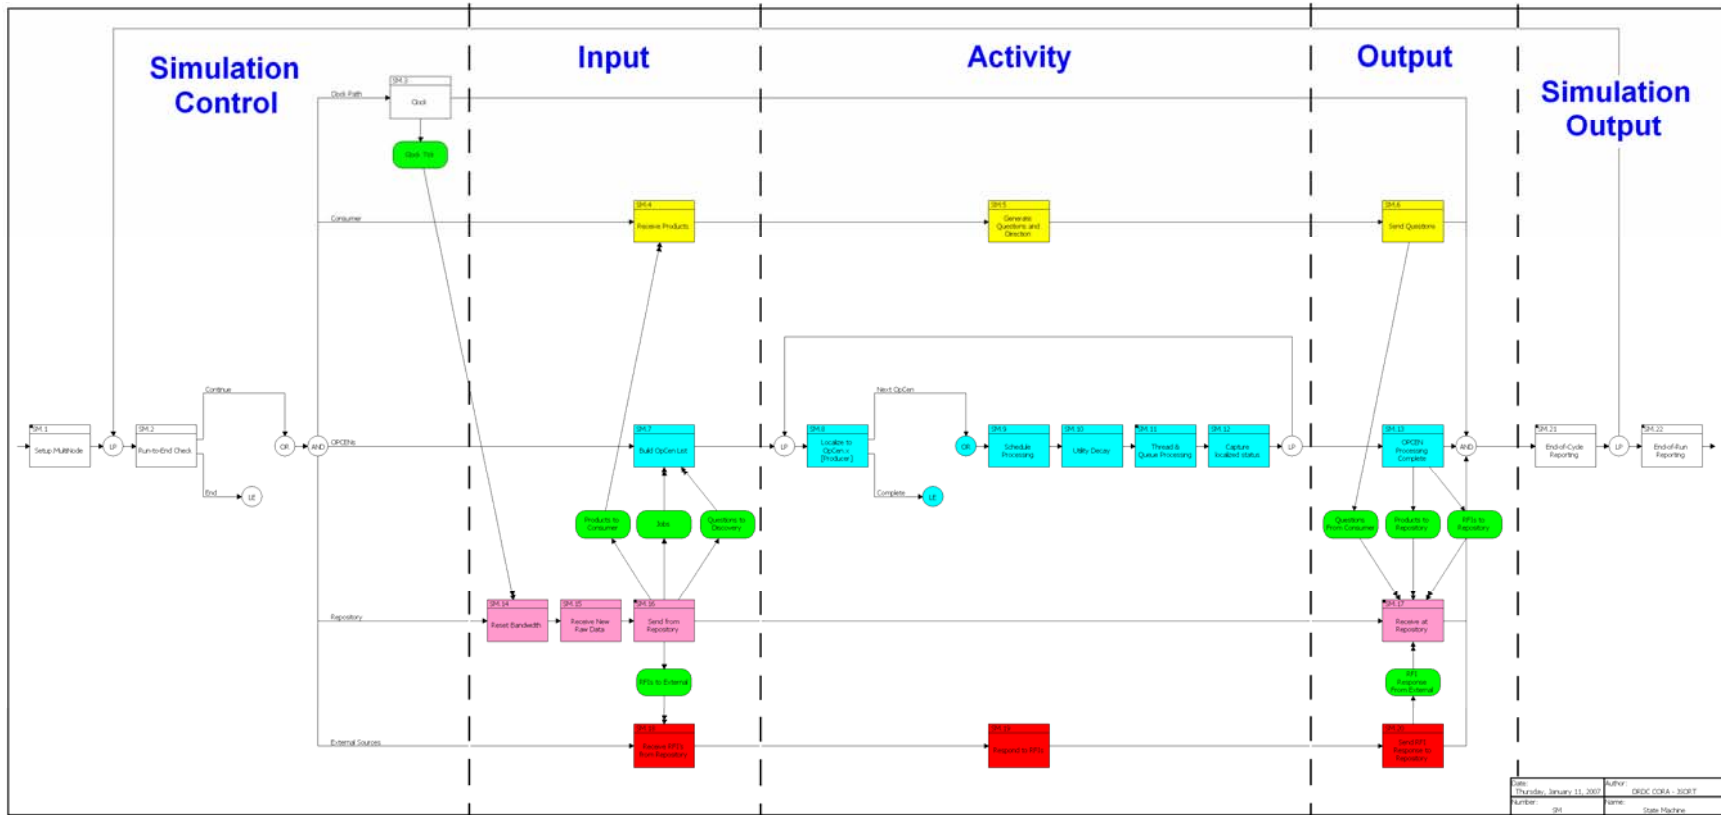


Figure 4: SMOFN Top-level Model

The Figure 4 diagram is also segmented into five phase of activity within each time step:

1. **Simulation Control** handles the process of incrementing the clock at the beginning of each time step and checking for the end of the simulation. It also includes the input of scenario data at the beginning of the simulation, before entering the time step loop;
2. **Input** simulates the delivery of information from the Repository to OPCENs (Producers, Consumers, and Discoverers) and to External Sources. Activity refers to the activity within those nodes, that is the generation of Questions by Consumers, the processing of jobs by Producers and Discoverers, and the production of new data by External Sources;
3. **Output** simulates the uploading of information from OPCENs and External Sources to the Repository; and
4. **Simulation Output** saves pertinent information to output variables at the end of each time step and saves variables to output files at the end of the simulation.

A key difference between Figure 4 and typical CORE behaviour diagrams is that the logic is not completely defined simply by the diagram itself. As mentioned earlier, the non-serial behaviour of TPPU logic is not scalable when modelled by classical flowchart diagrams, so SMOFN uses scripting embedded within each activity in the diagram to track the state of each job.

Logic Within Branches

The SMOFN elements logic shown as white boxes in Figure 4 is the simulation control and simulation output phases. When the simulation starts, the setup activity runs, reading data from input files and setting up global variables used to track data throughout the simulation. The SMOFN then enters a loop, which repeats for each time step. First the Run-to-End Check executes to see if the end of the simulation has been reached (this will occur at a pre-set time or when there are no jobs left to do). If the end of the simulation is reached then the loop exits and the SMOFN skips ahead to the End-of-Run Reporting which outputs the data that was tracked through the simulation. If the end is not reached, the clock increments a time step and each node's logic is again executed. Afterwards, the current state of every task and every operator is saved to the global variables that are the basis of the simulation output.

The simulation of the operational nodes is divided into three phases. Each node's activity is triggered by input received from the Repository, and that activity leads to an output that is sent back to the Repository.

Consume activity of OV-1 is represented by the yellow boxes in Figure 4. Each time step, the Consumer node begins by receiving products sent by the Repository. The Generate Questions activity then goes through all the products that have just been received and may generate questions based on those products. If any questions are generated, the delay before they are actually sent is determined and the time to send them is added to a schedule.

Produce and **Discover** logic elements of OV-1 are displayed by the blue boxes in Figure 4. The activity within both of these nodes is based on job threads where much of the scripting is the same. First a list of OPCENs is built and a loop is entered that is repeated for each OPCEN. The loop begins by choosing a new OPCEN and the Schedule Processing activity then looks at the job arrival schedule and then adds to the queue any job that arrives during this time step. Utility Decay then goes through all the jobs and checks to see how old their original data is. If the age of the data is such that the utility should decrease this time step, that is handled here. Thread & Queue Processing then goes through all jobs and assigns available operators to jobs in order of job priority.

It is here that the utility of jobs can increase as work progresses, or jobs can be abandoned if they cannot be completed on time. Produce and Discover jobs are allowed to draw from the same pool of operators so when the SMOFN sorts jobs in order of priority, it ignores whether the job is related to Production or Discovery. After these activities have executed, the Capture Localized Status element saves any data that must be tracked into the next time step. After each OPCEN loop has executed, the OPCEN Processing Complete element signals to the Repository that it can now execute its logic for receiving products and RFIs.

The pink boxes represent the work done by the **Repository** during each time step. First of all, the global variable tracking the bandwidth between the Repository and each OPCEN must be reset, as it is decremented whenever bandwidth is used to send or receive data. Bandwidth is tracked in terms of how much data can be transferred during each time step.

Next, the Repository checks for the arrival of new raw data. This check is based only on parameters read into the model during setup; raw data received from External Sources in response to RFIs is handled separately. Send From Repository contains the logic used to send products to the Consumer, jobs to the Producer, questions to the Discoverer, and RFIs to External Sources. This activity takes place before each of these other nodes executes during the time step so they can incorporate data as they receive it. Similarly, Receive At Repository waits until each node has completed their execution for the time step so it can receive data as soon as it is ready. This activity receives questions from Consumers, products from Producers, RFIs from Discoverers, and RFI responses from External Sources.

Finally, **External Sources** are simulated with the red boxes. They receive RFIs from the Repository, respond to those RFIs, and send the responses back to the Repository. As with the Consumer, the process required to respond to RFIs is not tracked in detail, only the amount of time required and the type of raw data returned are currently handled within the simulation. It is possible to model External Sources in more detail but it must be done in a way that does not detract from the OPCEN processes.

During each time step, items sent from the Repository must occur before any of the other nodes can execute their logic. Only after those nodes have executed can the Receive at Repository activity run. It should be noted that the Receive and Send activities of the Consumer and External Sources have no embedded logic scripts. Their logic is actually

handled within the Send and Receive activities of the Repository. They are shown to illustrate the Consumer and External Source perspectives of what is taking place.

Input Data Files

One advantage of the SMOFN is that the model itself is only concerned with the execution of the logic. All operating parameters, such as the number of operators at an OPCEN and their skill sets, or the number of jobs to be done, are read from data files during the setup stage of the simulation. This allows a single version of the model to be used across a wide spectrum of scenarios.

Simulation Setup

The first data file read into the SMOFN points to the data path where that scenario files reside. These files include a list of OPCENs and the data paths containing their setup information, a time at which the simulation will stop, and a switchboard to allow the analyst to select values for a list of generic business rule settings.

OPCEN Inputs

The parameters for each OPCEN is defined by five different files that are read into the SMOFN model:

1. An events list;
2. A summary of thread types;
3. A matrix of operator skills;
4. A set of utility decay curves; and
5. The step-by-step definitions of each thread.

Events List

The purpose of the events list is to define what jobs will be added to the OPCEN job queue, and when. The main entries for each job are the thread name, priority, and the time at which the job will be added to the queue. In the OPCEN SM, this list would include all jobs that the OPCEN would process over the course of the simulation. In the case of the SMOFN it only includes jobs that are based on OPCEN operating procedures. Other jobs, based on the arrival of new data are more appropriately listed in the schedule of new jobs that the Repository will send to each OPCEN (described in the next section).

The utility of the job will typically begin to decay the moment the job is added to the queue. To accomplish this, the file includes a field to specify when the data was collected so that the decay of utility can be calculated from that point in time. Finally, each job has a standard slack time for the particular job type is valid; there is also a field to set a specific deadline time.

Thread Types

The thread types file defines the overall characteristics of each job type that is processed at the particular OPCEN. The details specified here reflect the job as a whole, whereas the characteristics of each step within the job thread are specified in a separate file. Each entry is identified by thread name and priority (entries in the events list will have their priority rounded to those identified here if necessary). The data used to describe a thread includes the following fields to tailor the business rules to each job thread as appropriate:

1. Percentage of job will be declared as Nothing Significant To Report (NSTR);
2. Default slack time for the job;
3. Number of steps an operator can look ahead to estimate the expected utility gained over time spent on the job;
4. Whether the job is redundant with other jobs of its type (jobs that are redundant will be ignored if another job of the same type is in progress when a new one arrives in the queue); and
5. Utility of a job that returns NSTR, and the details of aggregating other jobs into this one.

For the SMOFN to properly account for the aggregation of several completed products into one new job, the following details need to be specified for the aggregating job thread:

1. Name of the job step where aggregation will take place,
2. Types of completed jobs to aggregate,
3. Amount of time added to the aggregation step for each job aggregated; and
4. Percentage of utility that carries over from the aggregated job to the aggregating job.

The time added and utility carried over can be different for each job type that is aggregated.

Operator Skills

The SMOFN refers to the operator skills matrix to assign operators in an OPCEN to queued jobs. Once assigned, this matrix also identifies how effectively operators do their work. The operator skills matrix identifies each operator in an OPCEN by a “name” that is a string containing only letters and numbers. The information tracked for each operator includes:

1. Their speed and quality of work (both expressed as percentages where 100% speed means jobs are done in the standard required time, and 100% quality means the operator adds the expected utility to jobs they work on);
2. Order in which they are assigned to generic work;

3. Percent chance they are able to take on generic work when they are not already assigned any specific job; and
4. List of their skills, beginning with their primary skill and allowing up to ten entries. When a queued job requires a particular skill, the SMOFN will look to each operator's first skill, then each operator's second skill, and so on until an available operator with the required skill is found.

Utility Decay Curves

The utility decay curves are used by the SMOFN to describe how the utility of products decreases as they age. Each line begins with a thread name and priority (the combination of which must match those in the thread types file) and a time expressed in units of the simulation's time steps (so far, we have used minutes for our simulations). Each line then ends with three numbers representing the lower bound, peak value, and upper bound of a triangular distribution. The logic of the utility decay is that after the product of a particular thread type has aged by the specified amount of time (age is counted from the moment the raw data was taken), its potential utility can be calculated by the specified triangular distribution. The SMOFN logic is defined so that potential utility never increases as products age, even if the distribution is constructed to allow it. The actual utility is calculated by multiplying the utility accrued from work done on the job times the potential utility based on product age.

Thread Definitions

The last file is used to describe the threads, in terms of how the SMOFN should handle each step within each job. This file uses one line of data for each step of each thread type (thread types must be the same, by name and priority, as those in the thread types file). The data for each line includes 12 fields covering the thread type, the name of the step, the time and skill required to complete it, the number of posts (utility updates) during the step, whether or not it can be interrupted, the lower bound, peak, and upper bound defining a triangular distribution used to calculate the utility achieved by the end of the step, the file size of the resulting product (used to track uploading to the Repository and subsequent distribution to Consumers), the type of step (in terms of which logical script to execute within the SMOFN) and the name of the next step to move to after the current one.

Seven types of steps are currently available to be used by SMOFN. The first three of which previously existed in the OPCEN SM:

1. No decision logic as to the following step, once the current step is complete the job will simply move on to the next;
2. NSTR decision: if the job is marked NSTR, it happens at the beginning of this step, after which all subsequent steps are skipped; and
3. Final step of any Production job, after which the job is marked complete and removed from the queue.

There are also four additional steps that are unique to Discovery jobs:

4. Decision as to the results of the Discovery process and has three possible exits.
 - a. No product exists,
 - b. Desired product is found, and
 - c. Some data is found but it is insufficient.
5. In the case that no product exists, the Discovery job will create a request for new data that will be sent to External Sources;
6. If the product is found, it must be analysed and put in the context of the originating question, so a new Production job is added to the appropriate OPCEN's schedule.
7. If insufficient data is found, then the Discovery will do a combination of the previous two steps. IT starts by sending a request to External Sources for more data, and in the mean time, a new Production job is added to the OPCEN schedule to analyse whatever data already exists.

Repository Inputs

A separate set of data files is used to describe the operating parameters of the Repository after the OPCEN initialization data. These files are:

1. Schedule for the arrival of new data;
2. Schedule for the arrival of new products created outside the simulation;
3. Matrix organizing the delivery of products to the appropriate Consumers;
4. List of each OPCEN's bandwidth; and
5. List controlling the generation of questions.

Data Arrival Schedule

The data arrival schedule is similar to the events list for each OPCEN. This file is used to simulate the fact that many jobs within an OPCEN are triggered by the arrival of new data in the Repository, rather than by the OPCEN's own operating procedures. Each line specifies the time that the new data was created, the time at which it will arrive at the Repository to be delivered to the Producer, the source of the data (such as a recce unit), the type and priority of the job thread that will be generated, the deadline for the data analysis job, the Producer OPCEN to which the data will be sent for analysis, the delivery method (i.e. push or pull), and the size of the file that must be delivered.

New Product Arrival Schedule

The schedule of outside products defines the arrival of products from Producers that are not tracked by the model. Product arrival is an important consideration for the SMOFN so we can simulate how one OPCEN (or several) reacts to products created by other OPCENs, without having to model those other OPCENs and their own product creation processes in detail. These products can then be forwarded to Consumer OPCENs as though they were created within the simulation. The required data includes:

1. Name of the producing OPCEN;
2. Type of job (including the priority rounded to the nearest appropriate thread type);
3. Specific priority;
4. File size to be transferred;
5. Time at which the original data was collected;
6. Time at which the product will arrive at the Repository;
7. Status of the job producing the product (i.e. complete, or the end of a specific job step, in which case the product corresponds to the most recent update);
8. Utility achieved due to work on the producing job; and
9. Decayed utility due to the data's age.

Delivery Matrix

The SMOFN uses the delivery matrix to identify what products are sent from specific Producers to specific Consumers. This applies equally to products created within the SMOFN simulation as well as to products identified in the new product arrival schedule. This file is the main source of information for the SMOFN to account for data sharing between OPCENs. A line of data contains the product type, the name of the Producer and Consumer OPCENs (one of each, products sent from one Producer to multiple Consumers must have one line for each Consumer), the percent probability that a product of the given type will be sent to the specified Consumer when created by the specified Producer, and the method of delivery (push or pull).

Bandwidth

The bandwidth file is very simple, containing only a list of OPCENs that connect to the repository along with their bandwidth. This identifies to the SMOFN the amount of data that can be transferred between each OPCEN and the Repository during a single time step. Aside from the Producer and Consumer OPCENs, the Discover and External Sources nodes should also be included here. Units are at the discretion of the analyst, but must be consistent with the file sizes for products identified in the thread definitions and the new product arrival schedule, as well as raw data in the data arrival schedule.

Question Generation

The questions file draws from the combinations of product types and Consumers that come out of the delivery matrix. This information is referred to whenever a product is sent to a Consumer so that questions are generated as appropriate. For each such combination, the questions file identifies to the SMOFN the percent chance that the Consumer will generate a question, the file size required to contain that question, and the amount of time that the Consumer will take between receiving the product and sending out the question.

Way Ahead

Work on the actual SMOFN model is nearing completion. The major issue now is to collect the data required to populate all of the initialization files described above. This data must be representative of the operating parameters of the OPCENs, as the goal is to accurately simulate data flow between these. Particular emphasis is being placed on describing job threads that involve multiple OPCENs, such as the process to report and respond to events in theatre or the preparation of high-level daily briefs. Another major component is the composition of various OPCENs, as far as the number of operators on shift and the skills they are trained in. Although much of this data will be classified, the intent is to keep the model unclassified as it only reads the data upon execution.

The planned process for building the future operational architecture of the Canadian Forces (CF) command structure with the help of SMOFN is illustrated in Figure 5. We begin by examining current C2 practices to capture them as executable threads, in light of the capabilities definition knowledge gained through previous work with the SMOFN model. These threads can then be converted into input to the SMOFN model. This, along with any developments of the SMOFN execution capability, leads to a more complete version of the SMOFN. The first payoff from the operator's perspective is that the practices captured as threads are documented and can be used to validate the Standard Operating Procedures (SOPs) in use. The second is that the inclusion of these threads into the SMOFN allows them to be analysed and used to potentially improve the SOPs. The payoff from the modeller's perspective is the creation of a more complete definition of C2 processes and lead to a better articulation of the target operational architecture.

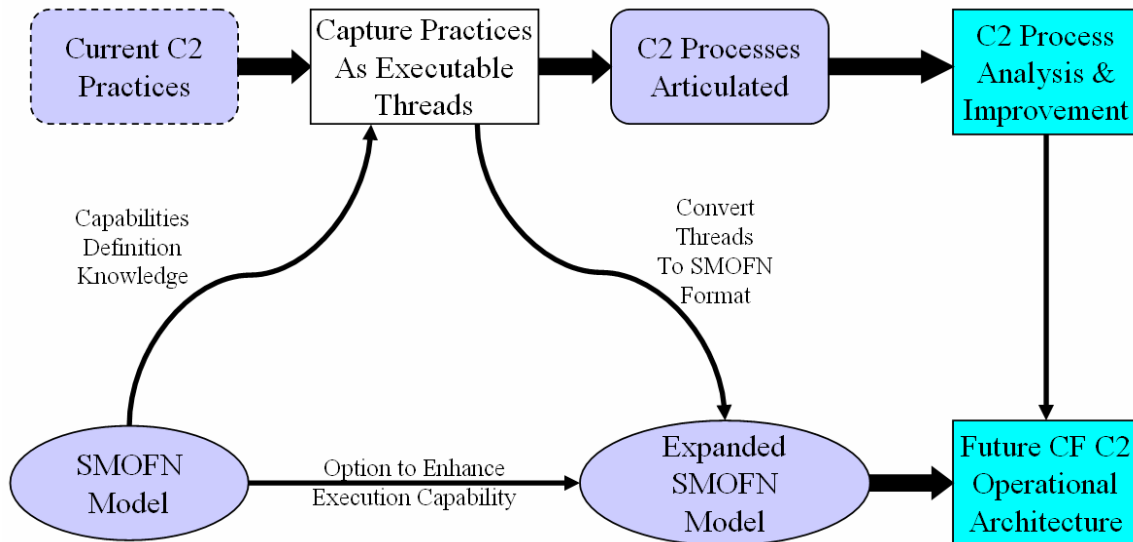


Figure 5: Planned Operational Architecture Creation Process

The above process is also an important part of the Verification, Validation, and Accreditation (VV&A) of the SMOFN model. If the executable threads captured from the operators are correctly transformed into inputs to the SMOFN, then the model should be able to replicate representative results of what is observed during operations. Once the VV&A process is complete for individual threads is it reasonable to start analyzing the SMOFN model parameters for cumulative effect of interactions between threads.

Results of Initial Trial

The process of capturing executable job threads in COREsim has recently been demonstrated as the final phase of SMOFN development. It started with capturing information about two job processes in the form of executable models:

1. The first case was an undocumented C2 practice used by the strategic level OPCEN night watch to prepare the daily electronic briefing for senior commanders. The data collection involved an operator subject matter expert (SME) describing process steps with personnel resources, timings and duration distributions while the analyst built and corrected the model logic. It took two days work to build a model of this C2 process that highlighted the effect of several process and resource bottlenecks; and
2. The second case started by building a quick model from a standard operating procedure (SOP) matrix that lists the actions required by several nodes when handling serious casualties during deployed operations. The analyst updated and tested the model logic with an operator SME while entering the associated data. It took five days part-time effort to build an executable model that works the way the process is made to work instead of what the existing SOP would suggest.

Both trial cases successfully demonstrated how to quickly build and validate detailed models with operator SMEs. The result was a much better articulated C2 process than existed before and it provided a means to test a broad range of process options. The added bonus of using COREsim was that formatted DoDAF product documents could be automatically generated for whatever data was entered into the model.

The final step in the modeling involves converting the C2 process threads to SMOFN data input file format. A portion of this has already been demonstrated using the above trial cases by decomposing the threads into simple sub-jobs and then linking them together. The project was still sorting out some technical aspects of this conversion process but it is anticipated these will be overcome in the near future.

The greatest challenge in implementing SMOFN will come from the need to have assured access to operator SMEs. They are the key to collecting, collating and validating the full range of C2 process threads. These operator SMEs are also the resource in greatest demand for ongoing operations. The staffs are investigating how to support the review their existing SOPs in order to document the C2 processes needed to support the future operational architecture.

Conclusions

A few significant conclusions have been drawn from the SMOFN modelling experience, as described in detail in this paper. The SMOFN has achieved its goal of extending the detailed OPCEN SM model logic across multiple nodes, each with their own jobs and resources but also with the ability to share information. The Net Centric virtual Repository business logic successfully controls the interaction of information between operator perspectives (i.e. producer, consumer and discovery) as well as between different OPCENs within the SMOFN. The SMOFN focus on business logic driven by data files for an instantiated C2 structure allows it to be both scalable and flexible.

The process of gathering data for process threads is both quick and effective; the major limitation noted so far is the availability of SMEs to provide the operational context. This is expected to continue to be the biggest implementation hurdle to building a faithful simulation and executable CF Command Structure operational architecture.

References

- [1] Funk, R.W., M.G. Ball, and R. Sorensen, "Building Executable Architectures of Net Enabled Operations Using State Machines to Simulate Concurrent Activities", presented to 11th ICCRTS, Cambridge, UK, September 28, 2006
http://www.dodccrp.org/events/11th_ICCRTS/html/papers/014.pdf
- [2] Funk, R.W. and Sorensen, R.L., "State Machine Modelling of TPED and TPPU", NDIA Proceedings, Oct 2005,
<http://www.dtic.mil/ndia/2005systems/thursday/sorensen.pdf>
- [3] Funk, R.W., Ball, M.G. and Sorensen, R.L., Building Executable Architectures to Simulate Concurrent Activities of Net Enabled Operations - Case of a Single Operations Centre, DRDC Technical Report TR 2006-24, Nov 2006.
- [4] Gauthier, S.M. and Funk, R.W., Estimating Partial Utility of Interim Products, DRDC CORA Technical Note TN 2006-02, April 2006.
- [5] Sorensen, R.L., Funk, R.W. and Ball, M.G., "Exploring Concurrent Activities: Using State Machines to Understand Net-Enabled Operations", Seventeenth Annual International Symposium of the International Council On Systems Engineering (INCOSE) 24 - 28 June 2007.