A DISCRETE EVENT SIMULATION MODEL FOR EXAMINING FUTURE SUSTAINABILITY OF CANADIAN FORCES OPERATIONS

Patricia Moorhead
Andrew Wind
Mira Halbrohr

Centre for Operational Research and Analysis
Defence Research and Development Canada
101 Colonel By Drive
Ottawa, ON K1A 0K2, CANADA

ABSTRACT

This paper presents a proof-of-concept discrete-event simulation model for examining the ability of the Canadian Forces (CF) to sustain operations, from a human resources perspective. Given a set of future operations for the CF, ranging from known ongoing domestic commitments to possible international missions, the goal is to identify potential shortages of deployable personnel by occupation, rank and unit, up to five years in advance of actual deployment. As a demonstration case, the proof-of-concept model was applied to a contingency analysis of the sustainability of Task Force Afghanistan over a three year planning horizon.

1 INTRODUCTION

The Government of Canada's Defence Policy Statement (GOC 2005) delineated a series of future tasks for the Canadian Forces (CF) in support of Canada's role in contributing to international peace and stability. Included in the list of specific tasks are: the ability to sustain for an indefinite period of time two land-based Mission Specific Task Forces of approximately 1,200 personnel each, in potentially different theatres of operation; and the ability to deploy a third task force of approximately 1,000 personnel for up to six months either to reinforce a current operation or mount a new short-term mission.

Given this direction by the Federal Government, the CF must be able to sustain up to three expeditionary operations, simultaneously, in potentially three different theatres of operation. The requirement to be able to sustain three distinct lines of operation at the same time, while continuing to meet domestic commitments, poses a challenge for the joint operational support community.

In order to assist with the operational planning process, the Canadian Operational Support Command requested a modeling tool be developed for projecting potential future shortages of joint operational support personnel, given a series of domestic and international task requirements over a three to five year planning horizon. This model would also be used to assess whether or not the distribution of deployable versus non-deployable joint operational support positions across the CF establishment is appropriate for meeting operational demands.

The set of future domestic and international operations to be considered will typically consist of a combination of tasks that occur on a known schedule (e.g. manning of instructor positions at training schools) and ones that may or may not occur during the planning horizon period (e.g. international disaster relief missions). To accommodate both deterministic and stochastic event generation mechanisms, a simulation model is required.

At the same time that the model development process began, a priority requirement to analyze the CF's ability to sustain Task Force Afghanistan (TFA) arose. As of the Fall of 2007, the mandate for TFA extended to February 2009. However, the possibility of the mandate timeframe being extended by the Canadian Government could not be ignored. Pro-active contingency analyses were required. Phase I of the project thus became the development of a model that considers international operations occurring on a known fixed schedule, in particular TFA. In future phases of the project, the model logic will be expanded to include stochastic mechanisms for task generation, and will consider both domestic and international operations.

This paper presents the operational sustainability simulation model developed in Phase I of the project. This model was used to investigate sustainment of TFA alone; all other international and domestic commitments were not considered. As such, the results derived are not a true reflection of reality but do serve to demonstrate the potential of modeling and simulation to project potential personnel

shortages and identify possible imbalances in the distribution of positions within the CF establishment.

The model inputs include: a set of positions, for domestic and/or international commitments, to be filled on a six-month rotating basis over a three to five year period; a static snapshot of CF Regular Force personnel demographics; a randomly generated health (the term "left out of battle" is used) status profile over time for each member; the most recent deployment history for each serving Regular Force member; and departmental personnel tempo policies. Given the inputs, the model determines where (e.g. occupation, rank) and when (e.g. operation and rotation) personnel shortages may occur. Based on the output data, it is possible to assess the number of deployable positions (by unit, occupation and rank) required to provide a desired ratio of deployed personnel abroad to replacements back home.

In Section 2 of this paper the data inputs and key assumptions underlying the Phase I simulation model are discussed. The model logic is explained in Section 3; verification and validation issues are also addressed. In Section 4 examples of the types of information that may be derived from the simulation results are shown. Finally, our conclusions and a discussion of model developments planned for future phases of the project are given in Section 5.

2 DATA INPUTS AND MODEL ASSUMPTIONS

2.1 CF Regular Force Population Data

While both Regular and Reserve Force members serve in domestic and international operations, only Regular Force demographic data is used in the Phase I operational sustainability model. As many Reserve Force members are part-time employees, the baseline population can change considerably from one month to the next. Obtaining Reserve Force demographic data of sufficiently quality to support the analyses being conducted is currently not feasible.

The first set of data inputs consists of demographic information on each deployable member of the CF Regular Forces. All members with appropriate minimum rank (Lieutenant and higher for Officers, Private-Trained or higher for Non-Commissioned Members) who are listed as being on manning in the departmental Human Resources Management System database are considered deployable. Members listed as non-effective strength (e.g. members on long-term disability, pursuing post-graduate studies, or on retirement leave) are not considered to be deployable. The population data input file contains the following snap-shot in time information for each deployable member: service number, rank, occupation, unit, whether or not the unit is deployable, and a series of dates indicating when the member is eligible to be tasked for international operations based on personnel tempo policy. The records are ordered

by rank, with the highest rank coming first. This is done to expedite the nomination process in the model.

The coding of units as deployable or not is done to reflect the concept that CF expeditionary operations should be sustained by personnel from deployable units only; infrastructure positions should not be emptied in order to fill operational billets. The coding derived remains to be validated by the various Force Generators and Force Employers within the Department of National Defence (DND).

The CF personnel tempo policy for international operations (NDHQ CDS 2007) is the key driver in the model. For deployments of 180 consecutive days or more, the policy states that a member is entitled to a 60 day respite period on return, and that an exemption period of 365 days (from the time of return) shall apply or else an operational waiver shall be required. During an exemption period members shall not normally be posted outside Canada or to an isolated post to which their families are not authorized to proceed at public expense. This does not preclude the possibility of a member beginning pre-deployment training within Canada before the exemption period has ended.

If operational requirements dictate that a member must be deployed internationally before the 365 day exemption period has ended, or an exempted member volunteers for another such deployment, a waiver document must be signed.

For deployments between 60 and 179 consecutive days, the personnel tempo policy recommends a two day exemption for each day deployed; operational waivers are not required for postings to international operations during this exemption period. For deployments of less than 60 days, the application of an exemption period on return is discretionary.

The personnel tempo policy is implemented in the operational sustainability model through a colour coding system, as follows:

- Red – the member cannot be tasked (he/she is currently deployed or in a 60 day respite period);
- Orange – the member is in the 365 day exemption period following a 180 day or longer deployment, but can be tasked (a waiver is required);
- Yellow – the member is in the recommended exemption period following deployments of between 60 and 179 days, but can be tasked (no waiver required); and
- Green – the member can be tasked, with no restrictions.

Within the model, exemption periods are not applied for deployments of less than 60 days.

The CF Taskings Plans and Operations (CFTPO) application captures deployment data for CF operations. For each member, "Go Green", "Go Yellow" and "Go Orange" dates are determined by examining the end date and duration of his/her most recent expeditionary operation deployment as recorded in CFTPO, and applying the person-

nel tempo policy as appropriate. For members who have not yet served in an international deployment, the "Go Green" date is set to be his/her hire date plus three or five years, for Non-Commissioned Members and Officers respectively.

The Phase I model assumes a static CF population throughout the simulation period: attrition, promotions, recruitment and posting cycles are not modeled.

2.2   Position Data

The second major data input file is a listing of all positions that are to be filled during the simulation period. CFTPO is mined to obtain position information from the table of organization for the current and next rotation of the operation(s) being examined. The current rotation is extracted as baseline information, while the position list for the next rotation is used as the template for all future rotations in the simulation period.

It is assumed in the Phase I model that all rotations for international operations are six months in duration and that all positions must be filled on a rotating six month basis. In reality some positions can span more than one rotation and be up to one year in duration. As a result, there may be some overlap in the manning of positions between the current (actual) rotation and the first simulated rotation.

For each position to be filled over the simulation period, the position data input file holds the following information: a unique record identifier; a task identifier (e.g. operation name); the task rotation number; a position identifier; the date by which the nomination must be filled (*NominateBy* date); the date by which a nominee must report for pre-deployment training (*ReportBy* date); the position start and end dates; the position occupation or occupational group requirement; the position rank requirement, which can be a single rank or a range of ranks; the nominee's service number (if the position has been filled); and the service component (Reserve or Regular Force) of the nominee.

At the time of the preliminary TFA analysis, the CFTPO database did not capture the *NominateBy* and *ReportBy* dates. For each position to be filled, the *NominateBy* date was set to be the position start date less 210 days. The *ReportBy* date was taken to be the position start date less 180 days.

Within CFTPO there is no data field in which the desired service component of a nominee is stipulated. Typically up to 15% of positions in expeditionary operations are filled by Reserve Force members. Ideally the distribution of Reserve Force members across positions in future rotations would be determined stochastically, based on historical Reserve Force augmentation rates. This is not done in the Phase I model, as the deadline of preliminary results did not allow sufficient time for this data mining exercise. Instead, positions in future rotations are hard-coded as to

be filled by either Reserve or Regular Force members, based on the pattern observed in the current rotation position list.

The position list in the data input file is ordered in the following manner. Firstly, the positions are listed chronologically by the *NominateBy* date. Secondly, the positions are ordered by the stringency of the occupation and rank requirements. Positions with a single occupation requirement are listed first, and within this subset positions with a single rank level requirement precede those with a range of eligible ranks. Positions with an occupational group requirement follow next, with the same division as before by rank requirements. Within each of these subsets, the positions are further ordered by occupation code and within occupation code by rank, with the lowest rank coming first.

In the data cleaning process, considerable effort is made to ensure that all position data is valid. Despite this, an occasional invalid occupation or rank code may not be detected. In such instances, the affected positions will not be filled during the nomination process.

2.3   Occupational Group Decodes

The personnel snap-shot data provides the occupation code for each member. Within the position list, the occupation requirement can be either a single occupation or an occupational group. A decode table was created to translate the occupational groups into their individual occupation components.

2.4   Left Out of Battle Data

The term left out of battle (LOB) refers to situations where deployable members cannot be tasked for deployment on any given day due to temporary personal circumstances, such as short-term illness or injury. To accommodate this reality in the model, a LOB probability distribution is used to randomly assign a LOB status profile to each deployable member.

The simulation period is divided into fixed six-month intervals, with the first interval starting on the simulation start date. The start dates of these consecutive intervals form one of the model input files. Within the model itself, the following LOB probability distribution is applied to each deployable member:

- 85% of members are healthy and "good-to-go" throughout the simulation period;
- 11% of members are LOB for one six-month interval during the simulation period;
- 3% of member are LOB for one one-year interval during the simulation period; and
- 1% of members are LOB throughout the entire simulation period.

This probability distribution has not been validated by military subject matter experts.

The above approach to modeling temporary unavailability of personnel for deployments accounts for pre-training losses only; it does not account for non-combat losses during training or any losses during the deployment itself.

The assignment of a LOB status profile to each member is the only stochastic mechanism implemented in the Phase I model. Task determination and Reserve augmentation patterns are strictly deterministic in nature; this will change in future phases of the project.

3 SIMULATION MODEL LOGIC

The proof-of-concept model was built using the Arena 11.0 simulation software by Rockwell Automation (www.arenasimulation.com). There are three segments to the model which are described below.

3.1 Rank Lookup Table Sub-Model

As the population data is being read into the model, a rank lookup table is created. Recall that the population data input file is ordered by rank, from highest to lowest. The purpose of this sub-model is to index the rank change points in the population data input file. The output of this model component is a lookup table linking each rank level to a record in the population data file where the listing of members with that rank begins.

3.2 Assign LOB Status Profile

As mentioned previously, the assignment of LOB status profiles is the only stochastic mechanism in the Phase I model. Each individual listed in the population data input file is assigned a LOB status profile for the duration of the simulation period, based on the LOB probability distribution discussed in Section 2.4. To begin, the LOB status array for each individual is first initialized to take on the value of "not LOB" in each of the six-month time intervals in the simulation period. The LOB probability distribution is then applied.

Let P be the number of six-month intervals in the simulation period. If the member is declared LOB for one six-month period, an interval is randomly chosen using a discrete uniform $[1, P]$ distribution. The corresponding entry in the LOB status array is then changed to "LOB".

If the member is declared LOB for a one-year period, a discrete uniform $[1, P-1]$ distribution is used to randomly select two consecutive six-month intervals. The two corresponding entries in the LOB status array are changed to "LOB".

For members declared LOB for the entire simulation period, all entries in the LOB status array are set to "LOB".

3.3 Nomination Process Sub-Model

The main component of the operational sustainability model is the nomination process sub-model. As each position to be filled is read into the model, the nomination process is activated.

The nomination process itself is deterministic; it is a straight-forward matching of personnel characteristics and position requirements. The stochastic functionality of the simulation environment is used to determine the values taken by some of the personnel characteristics, and in future model versions to determine which tasks are to be conducted when. The option existed to create the nomination process outside the simulation environment. For reasons of ease of use, and to minimize the number of file transfers between applications, the decision was made to code this sub-model inside the simulation environment.

The first step in the nomination process is to check if the nomination is already filled. This will be the case for all positions in the currently deployed rotation of the operation(s) being examined, and for nearly all positions in the next rotation to deploy. If the position is already filled, the nomination information is sent to the output file. If the position is not currently filled, the next step is to check if the position is designated for Reserve or Regular Force members. Recall that in Phase I of the project, these designations are done manually; in Phase II, Reserve Force augmentation patterns for deployments will be determined stochastically. The model attempts to fill Regular Force positions only. If the position is to be filled by a Reservist, this information is sent to the output file.

Before starting the search for members eligible to fill the Regular Force position under consideration, the list of occupations from which nominees must come is determined. The position requirement can be for a specific occupation or an occupational group. In the latter case, the occupational group decode table is used to obtain the list of acceptable occupations. If no acceptable occupations can be found to fill the position (i.e. the occupation code in the requirement is invalid), default nomination data indicating the inability to fill the position are written to the output file.

Once a list of occupations from which nominees must come is determined, the process of matching individuals to the position requirements begins. As shown in Figure 1, the first two criteria to be satisfied are: the individual must have the correct rank; and he/she must have a colour status other than Red before the *ReportBy* date (i.e. tasking the individual would not violate the personnel tempo policy). If no individual is found satisfying these criteria, default nomination information indicating so is written to the output file.

After an individual with the appropriate rank and colour status is found, the next question asked is whether or not this person is in a deployable position; only those indi-
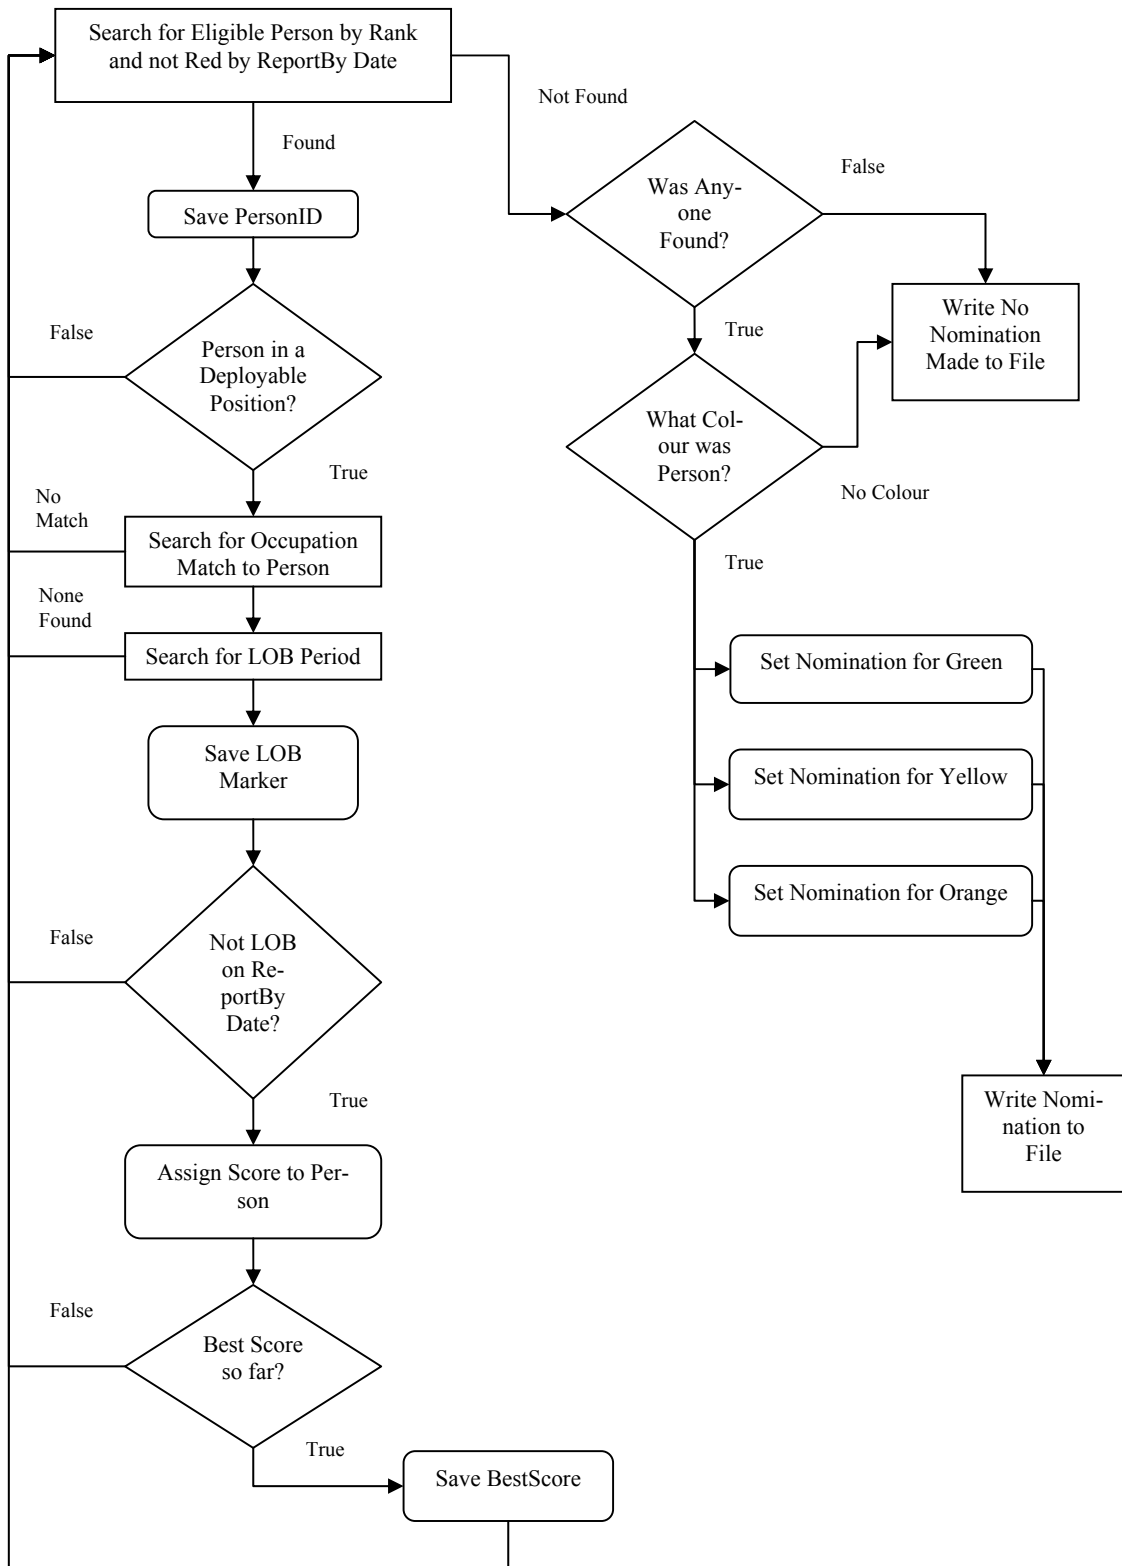
Figure 1: Model logic for the nomination process.

viduals in deployable positions will be considered for deployment to international operations. If the member is not in a deployable position, the model returns to searching the population list to find the next potential match.

At this stage, the model has checked for matches between the position requirement and individuals based upon component, rank, colour status and occupation. The next filter to be applied is the LOB status profile; members who are LOB on the *ReportBy* date cannot be nominated to fill the position, as they cannot report for pre-deployment training on time. The model does not check if the member is assigned an LOB status between the *ReportBy* and position end dates; such accounting for training and deployment losses will be incorporated in future versions of the model.

If the currently selected member is LOB on the *ReportBy* date, the model returns to the search block to find the next potential match in the population database.

Individuals who can report for pre-deployment training and match the rank and occupation requirements of the position to be filled, are next assigned a score. The scores are calculated based on how long the member has been coloured Green, Yellow or Orange, and his/her rank. Persons who have been Green the longest and are of the lowest required rank for the position to be filled are given the highest scores.

Let C represent the colour status (Green, Yellow or Orange) of the individual as of the position start date. The number of days the person has been Green, Yellow or Orange are denoted by D_G, D_Y and D_O respectively. Let R be the numeric value associated with the person's rank (ranging from 1 for General, to 18 for Private-Trained). The score assigned to an individual is then calculated in the following manner:

$$S_C = \begin{cases} 1{,}000{,}000{,}000 + 100 * D_G + R & C = Green \\ 1{,}000{,}000 + 100 * D_Y + R & C = Yellow \\ 1{,}000 + 100 * D_O + R & C = Orange . \end{cases}$$

If the member's score is the best (highest) observed so far, the member becomes the preferred potential nominee. The model then returns to the search block for finding eligible members by rank and colour match, to find the next potential nominee for the position.

This search process is repeated until all personnel eligible for the position to be filled have been scored. The person with the highest score is nominated for the position. Ties in scores are broken by taking the person who comes first in the population data listing.

There are two possible outcomes from the search for potential nominees: either a nominee was found, or there were no matches for the position. If no match was found, this information is written to the output file. When a nomination has been made, the next step is to update that indi-

vidual's colour change dates to reflect the new deployment.

Prior to updating the colour change dates, the old colour change dates are stored. This is done so that personnel tempo statistics can be calculated in follow-on analyses. As well, if the member was on an exemption period (i.e. had a colour status of either Orange or Yellow) at the time of nomination, the number of days remaining in that exemption period are recorded. Official personnel tempo policy does not require that unused exemption period time be returned to members after the second deployment. The operational sustainability model does, however, try to return this unused time to the member in the form of a non-mandatory exemption period after the new deployment finishes. The aim here is to distribute the operational burden as evenly as is practical among deployable personnel.

Let E represent the end date for the position being filled, and L represent the length of the deployment in days. Let T_Y represent the time that remained in a member's Yellow status when the member was nominated for the new deployment. This equals the number of days that remained between the old "Go Green" date (when the Yellow status would have ended) and the new position start date. Similarly T_O represents the "unused" Orange time for a member who was Orange at the time of nomination. Table 1 shows how the colour change dates are updated for the selected nominee, based upon the length of the new deployment and the member's colour status prior to the nomination.

Table 1: Formulae for determining the dates on which a nominee's colour status changes post-deployment, based upon pre-deployment status and deployment length.

Colour Status		Deployment Length		
Pre-Deploy	**Post-Deploy**	**< 60 days**	**60 to 179 days**	**180+ days**
Green	Green	E	E + 2*L	E+ 365
Yellow	Green	E+ T_Y	E + min[2*L + T_Y, 365]	E+ 365
Orange	Green	E+ T_O	E + min[2*L + T_O, 365]	E+ 365
Green Yellow Orange	Yellow	E	E	E+ 365
Green Yellow Orange	Orange	E	E	E + 60

With the nomination process now complete, the final step is to write the nomination information to the output file.

The model then returns to the position data input file to read in the next position to be filled. The process of

reading in positions and searching for nominees is repeated until all positions to be filled have been dealt with.

Multiple iterations of the model may be run. In each iteration, it is the LOB status profile for each member that may change, as this is determined via a probability distribution. In this manner the list of personnel eligible to fill each position may change from one iteration to another. The baseline population from which the eligibility list is drawn, and the set of positions to be filled remain constant across all iterations.

The information written to the output file for each position and its nominee consists of: simulation iteration number; the unique position identifier; the nominee's service number, rank, occupation and unit; the nominee's old "Go Green", "Go Yellow" and "Go Orange" dates; and the nominee's new "Go Green", "Go Yellow" and "Go Orange" dates.

In those instances where a nomination was not made (i.e. the position was already filled, designated as a Reserve Force position, or no match for the position requirements could be found), default values are inserted into the nominee information fields.

3.4 Model Verification and Validation

The coding of units as either deployable or non-deployable was conducted by a military subject matter expert. However, the designations have yet to be validated by the various Force Generators and Force Employers in DND.

Subject matter experts from the Directorate of Quality of Life were consulted to ensure the official personnel tempo policies are being accurately reflected in the operational sustainability model. The LOB probability distribution used has not yet been validated by subject matter experts.

All other model assumptions were briefed to and validated by the military sponsor.

4 EXAMPLE MODEL RESULTS

The Phase I operational sustainability model was used in the Fall of 2007 to project potential future personnel shortages associated with sustaining TFA. As stated previously, this was a pro-active contingency analysis: at the time the Canadian Government had not mandated TFA to extend beyond February 2009, but the possibility of an extension until 2010 or 2011 existed.

As the actual preliminary results obtained may reflect on the CF's operational capabilities in Afghanistan, they are not presented here. What can be said is that the initial results are very encouraging for the future utility of the operational sustainability model. While several of the model assumptions remain to be validated, the model framework developed to date is capable of addressing the issues it was designed to investigate. Questions such as when and where

the CF could potentially encounter shortages in joint operational support positions can be successfully examined.

To demonstrate the type of information that can be derived from the model results, some output data was created for a fictitious operation. Consider the situation wherein the CF is sustaining one expeditionary operation over a three year period, in which 700 positions must be repopulated every six months.

When running the operational sustainability model, there are two options: to restrict the pool of nominees to those in deployable units (the model default), or to allow all deployable members regardless of their unit type to be eligible for selection. The aim of running the model is to identify potential shortages of positions (and by association personnel) in the CF establishment by unit, occupation and rank. The shortages identified in the model can be of two types: the position was filled but an operational waiver had to be signed, or the position could not be filled at all. In the latter case this can be due to: there not being any personnel in the pool of potential nominees with the requisite characteristics (correct occupation and rank); or all personnel with the requisite characteristics were already deployed or in a mandatory 60 day respite period.

For the fictitious scenario above, suppose six occupations (labeled A through F) are identified as having the potential to experience personnel shortages during a three year deployment timeframe, when the pool of potential nominees is restricted to those in deployable units only. Figures 2 and 3 demonstrate the types of graphs that can be produced by the operational sustainability model in such a case. Figure 2 shows the percentages and counts of positions that could be left unfilled if personnel from deployable units only are tasked, and operational waivers are not signed. Personnel from deployable units could fill these positions if they signed operational waivers. Figure 3 shows the remaining percentages and counts of positions that could remain unfilled regardless of the use of operational waivers; the pool of eligible nominees within the deployable units has been exhausted.


Figure 2: Percentages and counts of positions in a fictional operation, by occupation and rotation, that could be left unfilled if personnel from deployable units only are tasked, and operational waivers are not signed.

Figure 3: Percentages and counts of positions in a fictional operation, by occupation and rotation, that could be left unfilled if personnel from deployable units only are tasked. The pool of eligible nominees has been exhausted.

The types of results shown in Figures 2 and 3 are based on the assumption that international operations are to be manned by personnel from deployable units only. Thus, any potential shortages identified to do not necessarily translate into shortages of people in the CF establishment; rather, they highlight potential shortages within the deployable units. The CF may very well have the necessary people to meet the operational requirements, but they could be currently employed in infrastructure positions (e.g. a Military Police Officer position at a headquarters building).

To help determine if the potential shortages shown above are due to a lack of personnel, or the distribution of deployable and non-deployable positions in the CF establishment, the operational sustainability model can be run with the "deployable units only" restriction removed. If opening up eligibility for nominations to all deployable members of the CF results in no potential shortages being identified, then the issue is the distribution of positions within the establishment, and not the number of personnel on manning. In this case, it may be advisable to re-examine the current distribution of positions to determine if changes can and should be made.

On the other hand, if allowing all deployable members to be considered for international deployments still results in the potential for shortages, one may conclude that there are insufficient numbers of personnel with the identified characteristics (occupation and rank) to satisfy the operational demands. Strategies for ameliorating the future impacts of such potential shortages may then be devised, examined and put into effect if the timeframe allows.

Another type of output from the operational sustainability model that may assist in the operational planning process is a projection of the apportionment of the nominations across the Environmental Chiefs of Staff (ECS) and Operational Commands (Op Cmds). Figure 4 shows an apportionment graph for a fictitious operation in which personnel from three ECS and three Op Cmds are tasked to deploy in the simulation model. Such information can be used to help determine which ECS or Op Cmd should be

tasked to generate the forces to fill certain positions in future rotations of an operation.


Figure 4: Apportionment of nominees across ECS and Op Cmds, by rotation, for a fictional operation.

5   CONCLUSION AND FUTURE WORK

In this paper, we have presented a proof-of-concept model for assessing the ability of the CF to sustain operations. The operational sustainability model can provide projections of potential personnel shortages several years in advance of deployments, and enables the assessment of whether or not the distribution of deployable versus non-deployable positions across the CF is appropriate for meeting operational demands. This study shows one way in which modeling and simulation can be used to inform the operational planning process, and the generation (assembly, training and deployment) of forces to meet the requirement.

The application of the proof-of-concept model to the analysis of sustaining TFA provided insight into several improvements that should be made to the model, in order to provide a more holistic and realistic analysis of the CF's ability to meet operational demands. In future phases of the project, the model logic will be enhanced to determine Reserve Force augmentation patterns stochastically based on historical rates, rather than the current hard-coding of specific positions. Linkages to existing models for projecting demographic changes should also be incorporated so the input population can be updated annually within the model to reflect the effects of attrition, promotions, postings and recruitment on the deployable population. As well, future versions of the model should incorporate both stochastic and deterministic mechanisms for deciding when future tasks and operations will occur.

REFERENCES

The Government of Canada. 2005. *Defence Policy Statement: A Role of Pride and Influence in the World*. Ottawa: The Government of Canada. Available via

<www.forces.gc.ca/site/reports/dps/ index_e.asp> [accessed March 13, 2008].

National Defence Headquarters, Chief of Defence Staff. 2007. *CANFORGEN 082/07: Perstempo Policy for CF International Operations*. Ottawa: National Defence Headquarters.

AUTHOR BIOGRAPHIES

PATRICIA MOORHEAD has a B.Sc. in Mathematics and Statistics, and a M.Sc. in Statistics from the University of Guelph. Since 2002, she has been employed as a Defence Scientist with the Centre for Operational Research and Analysis at the Canadian Department of National Defence. She is currently posted to the Canadian Operational Support Command Headquarters, and is conducting operational research studies on personnel and logistics issues related to joint operational support. Her e-mail address is <patricia.moorhead@drdc-rddc.gc.ca>.

MIRA HALBROHR is a Defence Scientist with the Centre for Operational Research and Analysis at the Canadian Department of National Defence. She is currently a member of the Personnel Operational Research Team under Director Strategic Military Operational Research and Analysis. Her e-mail address is <Mira.Halbrohr@drdc-rddc.gc.ca>.

ANDREW WIND has a B.Sc. in Physics from the University of Waterloo, and a M.Sc. in Medical Physics from Carleton University. Since 2004, he has been employed as a Defence Scientist with the Centre for Operational Research and Analysis at the Canadian Department of National Defence. Previously a member of the Personnel Operational Research Team under Director Strategic Military Operational Research and Analysis, he is now conducting operational research studies for the Canadian Navy. His e-mail address is <Andrew.Wind@drdc-rddc.gc.ca>.