Proceedings of the 2013 Winter Simulation Conference
R. Pasupathy, S.-H. Kim, A. Tolk, R. Hill, and M.E. Kuhl, eds

2 CANADIAN FORCES FLYING TRAINING SCHOOL (2 CFFTS) RESOURCE ALLOCATION SIMULATION TOOL

René Séguin
Charles Hunter

Defence Research and Development Canada
101 Colonel By Drive
Ottawa, ON K1A 0K2, CANADA.

ABSTRACT

2 Canadian Forces Flying Training School is responsible for the intermediate phases of all pilot training for the Royal Canadian Air Force. The school's operation is stochastic and dynamic in nature and a resource allocation planning tool has been built to simulate the interactions of its various components. For example, it takes into account weather, aircraft, simulator and instructor availability, and student failure. This paper gives an overview of the school's operation, describes how it is simulated with a custom built C++ application and shows how the tool has been used to estimate average course duration, to determine what resources are the most significant bottlenecks and to study the impacts of significant proposed changes to the way pilots are trained. The tool was instrumental in showing that one resource was clearly responsible for creating bottlenecks and was used to analyze a few mitigation options.

1 INTRODUCTION

Air Force Training is striving to increase pilot throughput to reach approved manning levels while insuring graduation dates are met. The resource allocation problem examined in this paper is closely related to the problem covered in Séguin (2011). The main difference being the type of occupations involved: pilots rather than air navigators and airborne electronic sensor operators. Even though the problems may seem quite similar on the surface, the two teaching programs are very different and use/combine/match their resources and students in very different ways. So much that two tools had to be developed. As discussed in Séguin (2011), both problems can be theoretically viewed as combinations of resource allocation (Ibaraki and Katoh 1988) and school time tabling (Burke and Petrovic 2002) without really being either in the classical sense. The objective here, as in Séguin (2011), is not to come up with an a priori optimal resource allocation and schedule as the stochastic and dynamic nature of this complex problem makes the task rather impossible. Unpredictable events happen almost daily and thus constant schedule adjustments are required. The goal is therefore finding a method to predict how long should each type of course typically last and determine what are the most significant factors that influence course duration and hence graduation dates.

2 SCHOOL OPERATION

2.1 Resources

The school uses the following resources: different types of instructors, simulators and aircraft. All resources have a maximum usage per day. Furthermore, several of the resources have random availability (for example, instructors get sick). Some resources may also be unavailable due to planned events (for example, periodic maintenance on aircraft (A/C) or instructors' parental leave).

978-1-4799-2076-1/13/$31.00 ©2013 IEEE

At 2 CFFTS there are two types of instructors: civilians for Flight Training Devices (FTD) and military pilots for flying missions (FM). The instructors' availability depends on several factors. First, as military personnel are posted to other assignments or locations on a regular basis (every three to four years), a portion of the instructors at the school have to be replaced every year as they are posted out and personnel getting transferred in require a certain amount of (re-) training by the school. Second, Air Force personnel are frequently required to attend specialized training courses, perform proficiency and secondary tasks, attend professional development activities, and the like. These activities are of various length and spread over the year; they reduce the instructors' availability for teaching. Third, instructors may occasionally be sick. Fourth, instructors may be on leave (all personnel have a certain amount of leave granted and some personnel may be guaranteed a certain amount of leave in the summer time). Last, instructors may become unavailable for extended periods of time; for example, maternity leave or prolonged sick leave.

The school has access to a certain number of FTDs and A/C of various types. A few time slots are available each morning and afternoon to do FMs but only one FTD session can be scheduled per AM or PM. The A/C are mainly used during daylight hours and, typically, more flights per day are flown during the summer as days are longer. Sometimes the training devices may be available outside regular hours when required. The school may also experience surge/loss episodes, where the available resources may be different for a certain number of days.

2.2 Teaching Activities

Activities are lessons that students attend; at 2 CFFTS, lessons consist of collective standard classroom lessons (ground school (GS)) followed by individual cockpit procedure training (CPT), FTD sessions or FMs. A complete series of specific activities form a course flow and 2 CFFTS teaches at least a dozen different courses, including those for its own instructors. An instance of a course is called a serial and usually more than one serial of each course is being taught at the school simultaneously, each being at a different stage of completion. For some course flows, a portion of the students may continue on an extended course once the common portion of the course is completed.

A typical course flow begins with several days of GS, both morning and afternoon. This may be followed by a series of days where half the day is spent in GS, and the rest of the day is spent following the course flow's sequence of CPT, FTD and FM activities. After a few days of half GS, only practical activities are performed. From now on, students can only be scheduled one single practical activity daily; in the morning or the afternoon.

Depending on the activity, the resources' requirements are quite different. For simulator sessions, there is usually one simulator and one instructor per student. For a FM, an instructor and an A/C are required. Some flights are multi-students meaning they have to be done simultaneously by two to four students requiring the appropriate amount of A/C and instructors. Some FMs may require a number of extra A/C and instructors; for example to act as simulated adversaries. Furthermore, FMs have varying weather requirements (see Section 2.3).

Individual students sometimes fail to complete FTD or FM activities. In this case, remedial actions may have to be performed. This implies that special activities are dynamically added to the scheduled curriculum. Students may also be unavailable to perform a training activity due to sickness or other duties and will thus be delayed in their progress.

If students are in the practical part of their course, they are allowed a refresher activity if they have not been in the simulator or if they have not been flying for a certain number of days. This usually happens after the Christmas break or when resources become limited and cause long delays. In this case a FTD or FM activity is dynamically added to the course flow.

For each course there is a general attrition rate: for example, 25% of the students fail the course or have to drop out for some reason. Historically, students attrite from courses at different sections of the curriculum with varying probabilities. For example, a large percentage of the attrition may happen at the first set of check rides ("marked" FMs) of a course flow.

2.3 Weather

Depending on the FM, different weather requirements will be used. These are defined using various thresholds for meteorological factors such as temperature, visibility, ceiling and wind. These elements are also combined to calculate wind-chill and runway orientation is used for cross-wind. Furthermore, runway friction index, A/C icing and the need for an alternate airfield have to be taken into consideration.

2.4 Training Interruptions

There is traditionally no instructional activity at the school on weekends, for two weeks at Christmas time, during sports days or when there is a graduation ceremony. For some courses, physical training needs to be taken into account as it limits scheduling of FTD sessions and FMs in a different way.

2.5 Scheduling Concepts

For large serials, or combinations of concurrent serials, there are often not enough resources to schedule every student's required activity when requested. Consequently and when added to mission failures, the students of a serial quickly begin to get out of synchronization with each other as they progress through the sequence of FTD and FM activities and careful monitoring and balancing of this behavior is mandatory. This is important as several activities require a certain number of students co-operating to perform the activity. Furthermore, students of the serial can graduate only when they are all done. When a resource shortage occurs, a decision has to be made as to what activities should be performed and which can be postponed, always making sure that they all eventually get accomplished. The overall objective is minimizing delays for course completion. Traditionally, the school has given priority to certain courses (e.g., instructor courses) and sometimes increased priority of courses when they get close to graduation. The options below have been proposed to improve student throughput and shorten course duration.

Students may be allowed to perform a combination of two activities per day: FTDs and/or FMs. Some serials may be deployed to another location with some of the school's resources for part of their curriculum. This deployed location would be selected such that weather would only minimally hinder students' progress. The school may also allow activities to happen during the weekend either in a planned regular way or as a reaction to observed delays.

The school has also thought of organizing course flows into blocks of similar activities in the hope of reducing failure rates and also allowing an easier control of resources usage. For each serial that executes the course flow, target end dates are given for each block. Rules governing when and how students may perform and combine two activities per day may then be defined for each block; the rules are one of the following: always, never, or only when students are late.

Finally, course flows have historically always been executed in a strictly sequential way: an activity always being a prerequisite for its successor. However, this does not strictly have to be the case and course flows could be scheduled in a more flexible way while insuring that certain rules are followed. Flexible scheduling is used to reduce idle periods caused by inclement weather or lack of resources. A student may thus be able to follow a few teaching "threads" simultaneously alternating from one to the other depending on resource availability and/or weather. That said, some activities do have prerequisites; others are tests, implying that a series of preceding activities are prerequisites; others may be barriers that prevent any subsequent activities until all preceding activities are completed; finally, some activities may be elements of a strict sequence portion which must be executed in strict succession with no other activities in between. All these constraints have to be considered when flexible scheduling is used.

3 METHODOLOGY

As explained earlier, the objective of this study is not to come up with an a priori optimal resource allocation and schedule but rather finding a method to predict how long should each type of course typically last and to determine what are the most significant factors that influence course duration and thus graduation dates. Due to the stochastic and dynamic nature of the underlying resource allocation and scheduling problems, a simulation approach has been adopted. A custom built C++ simulation software tool with a Visual Basic Graphical User Interface has been developed. Similar work has been done previously for the air navigator and airborne electronic sensor operator training school (Séguin 2011). However, the two training programs are so dissimilar that different algorithms and another simulation tool are being used. The pilot version has been built incrementally over several years and the first version was documented in Caron and Woodman (2003). The tool has grown in complexity and sophistication and this paper mainly documents work related to the last wave of improvements achieved primarily after implementation of all the optimization options mentioned above.

Even though all phases of pilot training are interconnected, they usually happen at different locations each using a single platform and are generally managed individually. For this reason the tool was designed to simulate a single training phase at a time. Furthermore, as students are only allowed to do at most one training activity per training period (morning or afternoon), this is the level of time discretization that has been adopted. The tool simulates the activity of the school every period for several years. It simulates weather, instructor sickness, leave and posting, student sickness, mission failures and attrition as well as take into account planned resource shortages or surges. In each period, decisions are made as to which students get scheduled and which ones get delayed. To ensure that each activity eventually gets scheduled, a priority is assigned to each student and, as delays are incurred, priority increases. A queue of students, sorted by priority is thus kept, and resources are allocated to each student in a greedy manner: in each period, the student with the highest priority is scheduled unless a needed resource is unavailable. If the student cannot be scheduled, he/she waits idly until the next period (his/her priority increasing). Whether the first student was scheduled or not, the scheduling algorithm moves on to the student with the next highest priority. This goes on until all active students for the period have been either scheduled or wait idly. Then, the process is repeated for the next period and so on.

3.1 Overview of the Main Algorithm

The program starts the simulation by creating master copies of the serial manager and calendar classes. These master copies will be used to start each single run of the simulation. The serial manager is responsible for queuing students of every serial and assigning resources. The calendar covers the span of days of the simulation and tracks weekends, holidays, resources availability (including surges and losses), deployments, etc. Then, the serials are sorted by dates, the earliest start date first.

Next, the main loop is begun with each pass being a run of the simulation. Copies of the calendar and serial manager are created for the current run, and a random seed is chosen. Two half-day 'sports days' training sessions are set for those serials who participate: sports days are scheduled the first and third Friday afternoons of the month. If one of those Fridays is a statutory holiday, the preceding Thursday is chosen. Then students in each serial are randomly selected to fail (i.e., not complete the course flow).

Another major loop nested inside the main loop is started. Each pass is a half-day training period. If the training period falls on a holiday, all students are marked idle. If it falls on the weekend, typically all students are marked idle. If it is a morning period, several activities are performed, including: 1) selecting a random number to determine whether the weather probabilities for the various weather conditions will be met today; 2) determining the number of available instructors (based on illness, surges/losses, deployment, etc.); and, 3) marking sports days in the schedules of affected students. Following this, the serial manager attempts to schedule all available students to perform their next mission.

3.2 Implementation Details

3.2.1 Remedial and Refresher Activities

Students who need to perform a refresher mission or remedial activities are diverted to the appropriate refresher or remedial course flow. Refresher missions occur when students of certain serials go several days without performing a mission. Under such conditions, students will spend a training period performing a simple FM or FTD mission before continuing their course flow. If the last mission performed was an FTD activity, or there is no A/C, or the weather prevents flying, it will be identified as an FTD refresher.

Remedial activities are the results of a student failing to complete an activity. If the reason is beyond the student's control, the mission is simply repeated. Otherwise, the precise remedial activities needed depend on the severity of the failure, how often the student has failed an activity before, which course is being followed, how far along the student is and the type of mission that was failed. The complexity of this process is far too great to be replicated closely by the simulation tool. However, as one is interested by the delays caused by the failures rather than the precise paths followed by individual students, it is possible to simulate the remedial delays.

The frequency of zero, one, two, three or more additional activities vary depending on the factors mentioned above but simple rules can generally be applied combined with the use of probabilities. The following activities are typically used as remedials: 1) assessment of the student's situation and classroom instruction; 2) one or two remedial simulation activities; and, 3) a remedial flight activity if the failed activity was a FM. Any of these will add a one period delay to the course duration and each will obviously require the appropriate resource. Finally, the original incomplete activity is always repeated.

Students who have failed a multi-student mission will perform it again after any required remedial course work but with instructors playing the roles of the other students.

3.2.2 Simulator Resource Assignment Rules

When verifying resource availability for the various training devices, the original requirement is first verified for a match and if no resource is available then other devices can be checked for availability while ensuring that the established hierarchy of devices is followed. This means, for example, that if a CPT is required but none is available then a FTD of class one can be used and if none is available then an FTD of class two can be used and so on. Furthermore, before checking for another type of device, off-hour availability on the original equipment is always checked first. If a match is attempted for a student during an AM period and only off-hour availability is found, then it is assumed that the student will perform the activity earlier in the morning. Similarly, for the PM period, a student would be assumed, in this case, to perform the activity after regular hours.

3.2.3 Aircraft, Simulator Devices and Instructor Availability

Aircraft availability is an issue that has not been included in the simulation itself but rather handled outside the simulation tools. In view of the significant number of A/C employed at the school, it was deemed sufficient to obtain average availability in the form of a number of flights that can be launched in each period. As for FTD, their number being much smaller, it is possible to include a probability of failure for each of them even though in practice they have been extremely reliable. If a probability is used, it is always assumed that the device would be unavailable for a full day when it fails. For instructors, average weekly availability is obtained a priori by taking into account posting, leave, training, professional development and secondary duties. From this, the number of instructors on sick leave on a given day is deducted. This number is taken from a Binomial distribution with parameters n and p where n is the number of instructors and p the probability of each instructor being sick on a given day.

3.2.4 Weather

Flying missions are affected by weather and delays are very frequently incurred due to adverse weather conditions. To be able to account for these delays the tool needs to be able to realistically simulate the weather effects. The tool examines weather the day of a flight and decides if the mission can be flown or not using historical weather data. To decide if a weather requirement is met for a specific mission, average monthly weather condition probabilities are used and compared to randomly generated numbers. All weather factors described in Section 2.3 are considered.

3.2.5 Scheduling

Each student is placed in the scheduling queue for a mission (e.g., CPT, FTD or FM) or scheduled directly for GS. The queue is self-sorting, according to the priority criteria described in the next subsection. The serial manager then attempts to empty the queue. It will try to match the first choice mission (the next non-completed activity in the ordered curriculum) of the top student in the queue to the available resources. If resources are exhausted or a detail of the mission prevents it from occurring (e.g., weather, multiple students required for the mission, student not allowed to combine FTD and FM activities in the same day, or to fly twice in one day) and flexible scheduling (FS) is used, the serial manager will attempt to match the student's second choice to resources and so on, if such choices are allowed by the rules of FS. Students who cannot perform any mission will be marked as idle and will be queued for the same mission again in the next training period with increased priority.

When FS is used, students do not have to strictly abide to course flow sequencing and a few alternative missions may be possible. These are determined using the flexibility rules mentioned in the previous section. Here is how FS work: first, the lowest index non-completed mission is always tried first (above, this mission was called the first choice mission); then, if this mission cannot be accomplished and it is not part of a strict sequence portion of the course flow, all higher-index missions for which all prerequisites are already accomplished (or for which there is no prerequisite) and that are not further ahead than a barrier mission are possible alternative missions.

3.2.6 Serial Priority Sorting

Below is the cascading ordered list of criteria used for sorting the students being queued in each training period. When students A and B are compared to see which one should be ranked higher,

1. If student A has been unable to attempt the mission in the previous period due to resource shortage, weather or other students not being ready for a multi-student mission, student A gets priority;
2. If student A is in a remedial course flow due to a previous failed mission, student A gets priority;
3. If student A's course flow has higher priority than student B's course flow, student A gets priority;
4. If student A's progress through his/her course flow grants him/her the "close to graduation" status and student B's progress through his/her course flow does not, student A gets priority; and,
5. Finally, the student who has been the longest without executing a mission (FM or FTD) gets priority.

3.2.7 Multi-student Missions

Whether or not FS is used, once a student is ready for a multi-student mission, a check needs to be made to ensure that a sufficient number of students still need to complete the same mission. If not, instructors will have to be used in place of the missing students. If FS is not used, the only option for the student ready to do the multi-student mission is to wait until a sufficient number of students catch up and are ready to do the mission. When FS is used, the situation is significantly more complex due to the varying priority reached by each student. Assuming that student A is the highest priority student and needs to be matched with other students to perform mission XYZ, the serial manager will initially search the queue for students who have XYZ as their first or primary choice mission (whatever their priority is, as student A is already the highest priority). If a sufficient number of students are found, student A and the other students are matched to resources and scheduled for the mission. If an insufficient number of students are found with XYZ as their primary mission, then mission XYZ is moved to the back of student A's possible mission list and the other potential missions for this student are examined and one will be scheduled if the resources are available. If not, mission XYZ may eventually returns to the front of the student's mission list; then, on this second try, the queue is searched for students for whom XYZ is a possible mission (not necessarily their primary). If enough students are identified (say students B, C and D), then student A is moved within the queue of students from the top of the list to a point just below student D, the lowest priority student of the set who might perform mission XYZ with student A. The decision to delay this assignment is based on the fact that for this assignment to be optimal at this point, the priorities of all students involved and those ranked between them would need to be assessed; this would require a very complex set of rules. It was deemed reasonable to delay this assignment as the highest priority student was already given a few chances for an assignment. At this point, the student next to A in the queue is examined for an assignment and so on. Eventually, the serial manager may reach student D who may be ready to perform mission XYZ, and in the queue below him/her we could have students A, B and C, each with mission XYZ as the top mission choice in their trimmed mission list (as their initial first choice may have been unschedulable). However, mission XYZ may not be student's D primary mission and/or students B and C may have been scheduled other missions. If the serial manager reaches student A again, because B, C and D each executed missions other than XYZ, student A will again search the queue to identify a sufficient number of students with which to perform the mission. This may identify new potential matches of lower priority than student D, or perhaps terminate the goal of scheduling a mission for student A in this training period.

3.2.8 Block Lateness Rules

If some course flows have been defined with blocks in an attempt to improve student throughput, lateness may have to be assessed to see if two missions per day may be allowed. This is done when scheduling is completed in the training period loop (the second major loop) if the training period is an afternoon and the day corresponds to the planned end date of a block for a serial. Then for subsequent periods, block lateness catch-up rules may apply to dictate how many missions per day students may be allowed to do. For example, students who have three missions outstanding on the date the block should have been completed may be allowed to perform two missions per day for their next six missions (note that this rule is less restrictive than allowing two missions a day for the next three days as the six missions can take more than three days to complete). These catch-up rules do not override but rather combine with the serial rules allowing specific combinations of FMs and FTD missions on a same day.

3.2.9 Dynamic Weekend Activities

With this improvement option, when a serial has a sufficient number of students who are late or critically late, these students may have training periods scheduled on Saturday and maybe Sunday. Lateness is assessed using the difference in percentage between actual progress and planned progress. A different percentage is used to define lateness and critical lateness. The comparison between actual and planned progress is made on each weekend day whenever dynamic weekend activities is allowed for a serial. On the day of the comparison, the planned progress of the serial is expressed as a percentage of the expected duration (i.e., divide the number of days passed since the start of the serial by the difference between the start and expected end date of the serial). For each student, actual progress is expressed as a similar percentage (i.e., divide the number of missions completed by the total number of course flow missions to perform). If the difference between the two percentages exceeds a user defined threshold, the student is late. However, the school will not open on the weekend if only a few students are behind, hence thresholds for the number of late and critically late students also need to be provided by the user.

3.2.10 Planned Saturday Activities

Under this improvement scenario, the school will set ranges of dates during which planned Saturday activities shall occur. During those periods, selected serials will be allowed to schedule certain activities on Saturdays unless they are part of a long weekend (i.e., including a statutory holiday). The type of activities allowed on Saturdays can be different for each serial.

3.2.11 Deployed Training

With this improvement option, A/C, instructors and students of selected serials travel to another location, typically to take advantage of better weather. The pool of A/C and instructors at the school is thus divided between the two locations, and the deployed serials take a day traveling to and returning from the location. While deployed, however, the serials will typically advance quickly; the students will typically work six days per week, and be allowed to perform two missions per day, which can include flying twice per day, with relatively few missions delayed by bad weather.

3.3 Validation

The simulation tool has been in use for ten years and its various versions have been validated in a number of ways. First, historical data has been used to assess the weather module by comparing actual serial duration with simulated ones and the tests showed that the simulated results were quite close to observed durations and much better than the school's a priori forecasts. Then, in 2007, results of simulating the school operations with the tool were compared with those obtained with the Arena model of one of the Air Force Training's contractors: both were consistent with each other and representative of actual school activity. In 2011, 12 months of history were used to revalidate the model. Finally, the fact that numerous Air Force Training military subject matter experts and contractors have trusted the tool to help them plan resource allocation and establish schedules for so many years is another testimony of the tool's quality.

4 ANALYSIS RESULTS

The objective of this study was to examine the implications of modifying the student Phase II course and blending it with the Phase III course so that both use the CT-156 Harvard II turboprop A/C whereas Phase III traditionally used the CT-155 Hawk jet. This is part of a complete reorganization of the pilot training program with the aim of significantly increasing the student throughput. At the same time Air Force Training wanted to validate the characteristics of a new training paradigm: use of flexible scheduling, only full day ground school with no half-day portion and course flows arranged in blocks of similar type activities (blocks of FMs alternating with blocks of FTDs).

Table 1 provides information about each of the different courses. For each course, the number of FM and FTD activities, the number of serials scheduled per year and the number of students per serial are provided. For the two Phase II course flows, two values are often given to provide information for the regular and the extended portion of the course. The first Phase II course flow is the traditional course that was planned to be taught until the end of 2011 and where the extended portion is followed only by future jet pilots. The Phase II Block course is the new proposed blended version where the extended portion is composed of former Phase III missions and also attended only by jet students. The Transition B course flow is identical to the A version with the exception that the course is interrupted in the summer to allow students to complete their posting move (Transition courses are taken by experienced career pilot transitioning between A/C type). The instructor pilots (IP) proficiency course flows are artificial courses containing only flying activities and dedicated to IP to maintain their proficiency. The flying instructors course flows (FIC) are for new IP posted at the school. They have to be included in the school schedule as the same A/C required for student is used. The difference between the two versions of the course is only to reduce the number of proficiency flights that will be flown during the shorter winter months. Table 1 also indicates the total number of students of each type enrolled annually. Even though values are provided for the IP Proficiency courses, they are artificial as the goal here is rather to generate approximately 1200 sorties annually for proficiency.

To ensure that all potential interactions between the serials are experienced, the tool is fed with a few years' worth of serials of each course type. It is fed data from late 2010 to the mid 2015 that roughly covers the course of 2011 to 2014. Older and future serials are used to "seed" the simulator, as meaningful course length and graduation dates for the years of interest (2012 and 2013) can only be obtained once a realistic set of previous, of-interest and succeeding serials are continuously competing for resources. The last column of Table 1 shows the number of serials of each type used in the simulation.

Table 1: Courses and serials characteristics.

Course Type	Number of FMs	Number of FTD Missions	Total Number of Missions	Number of Serials Per Year	Number of Students	Number of Students Per Year	Total Number of Serials for the Simulation
Ph II	75/98	24/28	89/126	8	14 or 16	120/28	8
Ph II Blk	54/105	32/48	86/153	8	14 or 16	124/32	24
Trans A	74	22	96	1	6 or 8	6-8	4
Trans B	74	22	96	1	8	8	3
FIC	45	21	66	3 or 4	2	6-8	13
IP Prof A	136	0	136	1	8	8	4
IP Prof B	32	0	32	1	4	4	5

Due to the stochastic nature of factors such as weather and resource availability, numerous runs (typically 500) are done to obtain average behavior and assessments are done using Measures of Effectiveness (MOEs) to determine whether a scenario is expected to be successful or not. Typically, Air Force Training plans a few years worth of staggered serials and each is expected to last a certain number of days depending on the time of the year and the number of students loaded on the serial. A scenario is deemed successful if MOE 1 is met and either MOE 2 or MOE 3 is also met where the MOEs are:

1. MOE 1: 75% of the serials of interest graduate on-time (or early) in at least 75% of the runs.
2. MOE 2: 90% of the serials of interest graduate no more than 5 working days late in 75% of the runs.
3. MOE 3: 75% of the serials of interest graduate no more than 5 working days late in 95% of the runs.

The MOEs were assessed for the two-year period 2012-2013, which means a total of 16 serials of interest since eight regular student serials are planned annually. In this case, the MOEs translate to:

1. MOE 1: 12 of 16 serials must graduate on-time in 75% of the runs.
2. MOE 2: 15 of 16 serials must graduate no more than 5 working days late in 75 % of the runs.
3. MOE 3: 12 of 16 serials must graduate no more than 5 working days late in 95% of the runs.

The first scenario was conducted with the legacy resource pool kept intact for the whole simulation: 21 FTD slots per day (three devices seven times a day), 70 IPs, 18 FTD instructors and 85 A/C launches per day (72 during winter). Apart from the differences noted previously, the blended course flow also has 50 full days of GS and no half days rather than 15 and 55 respectively.

Table 2 shows the results of the initial simulations. As can be seen, Scenario 1 was not successful as none of the MOEs were satisfied and only a few serials were on time. The three rightmost columns provide percentages of days where shortages of resources happened due to too many serials requesting them simultaneously. The two values are indications of extremes over all runs. It is clear the main bottleneck is created by a lack of FTDs. In Scenario 2, an extra slot was added for each device but that still was not sufficient to satisfy any of the MOEs. In Scenario 3, another slot was added. MOE 1 and MOE 3 were satisfied and the majority of the course serials were graduating five to ten days earlier than anticipated (Figure 1). However, all 27 slots would be utilized between 90 and 115 working days in a year (over all runs) with the rest of the days experiencing quite a spread in daily usage (see Figure 2 for a typical run with 105 days at maximum usage and a median of 16 slots when not at maximum usage). This significant variance implies staffing complications for the FTDs.

Table 2: Scenario Summaries and Results.

			Shortages		
Scenario Success	**FTD slots**	**MOE Evaluation**	**A/C**	**IP**	**FTD**
1 No	21	2 of 16 serials on-time in 62% of the runs 15 of 16 serials are no more than 30 days late 73% of the runs	0-2%	1-1%	27-65%
2 No	24	8 of 16 serials on-time 78% of the runs 9 of 16 serials are no more than 5 days late 96% of the runs 12 of 16 serials are no more than 5 days late 83% of the runs	0-2%	0-2%	22-40%
3 Yes	27	12 of 16 serials on-time 77% of the runs 13 of 16 serials are no more than 5 days late 95% of the runs	0-1%	0-1%	8-14%

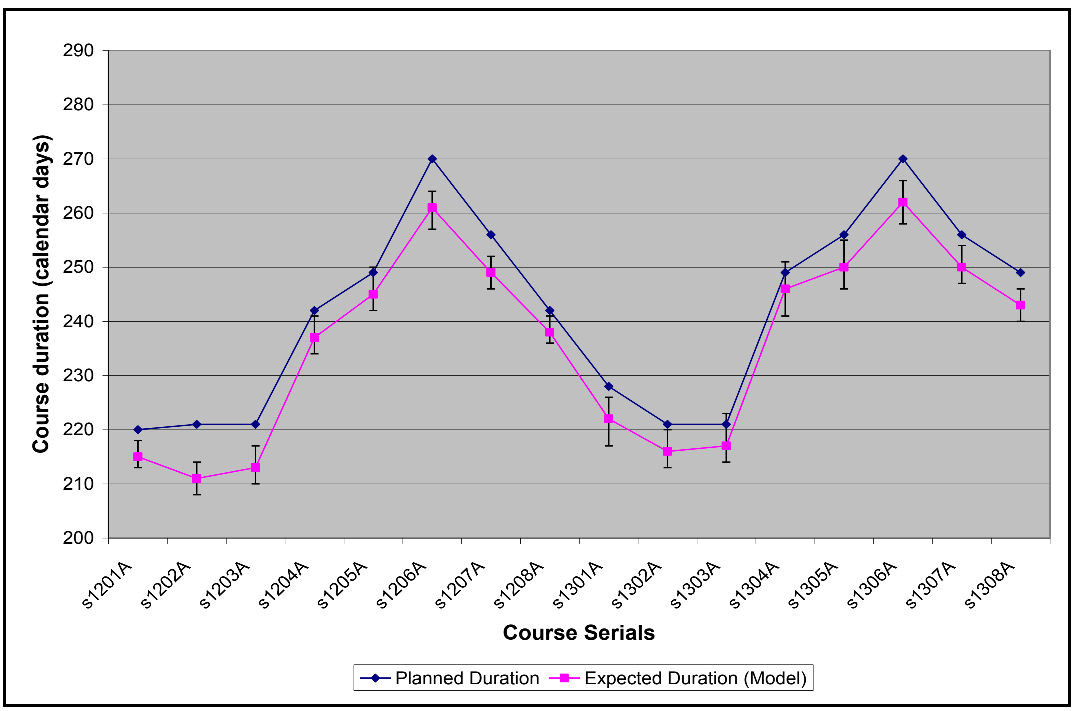

Figure 1: Scenario 3 2012-2013 Course Durations with 1st and 3rd quartiles.

Even though Air Force Training was already pleased with the results, other options were explored to find out if the situation could be improved. Scenarios 4 to 6 explore the effects of the number of full and half days of GS. Scenario 7 looks at the impact of course flow arranged in blocks and Scenario 8 looks at the effects of using FS. A full design of experiments was not run as the sponsor was quite reluctant to steer away from its original plan but we wanted to at least provide some insights on the impact these factors could have on the serials' chances of graduating on time. The results can be found in Table 3.

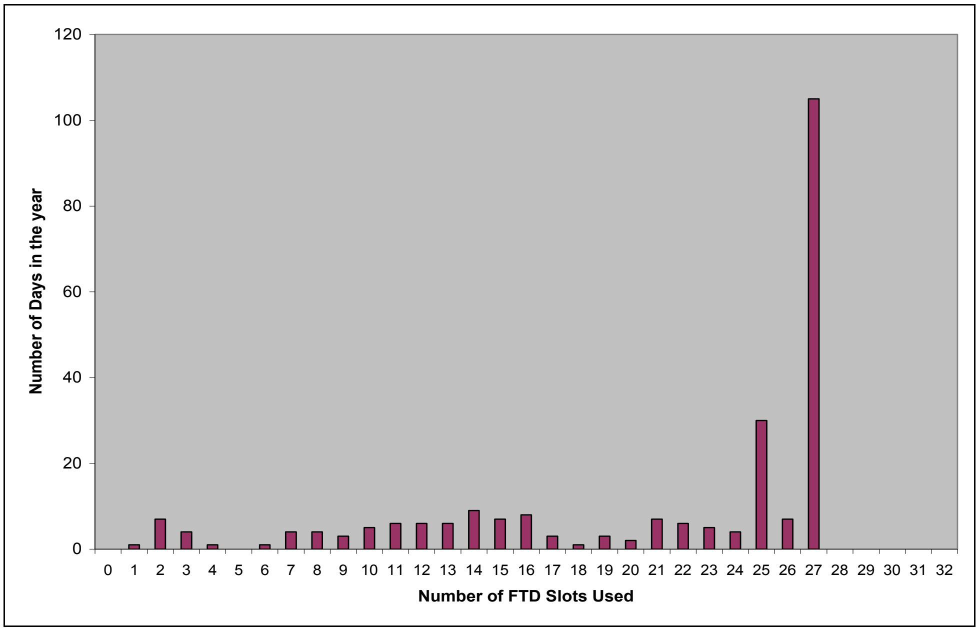

Figure 2: Scenario 3 Typical Daily FTD Slot Utilization.

Table 3: Scheduling Option Scenarios Summaries and Results.

Scenario Success	**FTD slots**	**Option**	**MOE Evaluation**	**FTD Short.**
4 Yes	27	35 Full & 30 Half Days	15 of 16 serials on-time in 80% of the runs 16 of 16 serials are no more than 5 days late 88% of the runs	5-9%
5 Yes	24	35 Full & 30 Half Days	12 of 16 serials on-time 76% of the runs 13 of 16 serials are no more than 5 days late 96% of the runs	15-27%
6 No	24	40 Full & 20 Half Days	11 of 16 serials on-time 75% of the runs 12 of 16 serials are no more than 5 days late 96% of the runs	13-28%
7 Yes	24	Scenario 2 but No Blocks	14 of 16 serials on-time 78% of the runs 15 of 16 serials are no more than 5 days late 96% of the runs	4-11%
8 No	27	Scenario 3 but No Flex. Sched.	1 of 16 serials on-time 37% of the runs 15 of 16 serials are no more than 25 days late 81% of the runs	4-9%

As can be seen in Scenario 4, results of the successful Scenario 3 were further improved when half-day GS was used rather than only full days. This was expected as allowing students to start practical activities early tends to spread out the FTD demand. This is particularly important for Phase II courses arranged in blocks as the first block consists exclusively of FTD sessions. In fact, as shown by Scenario 5, reducing the number of full days from 50 to 35 results in the same level of success as Scenario 3 but with only 24 FTD slots compared to 27. As Air Force Training wanted to keep FTD sessions as close as possible to their associated FMs in the next block, they wanted to see if the number of full day GS could be increased while keeping 24 FTD slots. Scenario 6 confirms that 7 weeks of full day GS is the upper limit as 8 weeks was not successful even though quite close. Scenario 7 explores the option of keeping the legacy mode of training with FTD sessions and FMs constantly alternating with one another for the Phase II course. This option provided the best results overall when 24 FTD slots are used: better graduation success, better FTD usage and close proximity of FTD sessions and associated FMs. The number of days at maximum usage varied between 113 and 141 days and the spread was much smaller with a typical median at 18 slots when not at maximum usage. A median of 18 slots with a maximum of 24 is significantly easier to staff than a median of 16 with a maximum of 27. Finally, Scenario 8 clearly demonstrates the huge impact FS is having: success rate is even lower than Scenario 1 where only 21 slots were used. Without FS, the planned schedule of serials using the new Phase II blended course would have virtually no chances of graduating on time unless even more resources are devoted to the program. Note that now weather, A/C and IP availability were responsible for delays as FTD slots were sufficient.

5 CONCLUSION

A tool has been built that allows for realistic simulation of the school operations. It has been used to obtain useful insights on causes of delays, course durations and usefulness of various scheduling options. The simulations clearly showed that the number of FTD slots is the most significant bottleneck for the specific phase of training examined. The results were instrumental in helping decision makers select an appropriate plan for implementation of the new program. As the only significant issue was the availability of the FTDs and extra capacity was relatively easy to add (requiring “only” additional funds), opening the school on weekends was not entertained. Similarly, as weather and A/C availability was not an issue, deployed training was not explored for this phase. In summary, Air Force Training decided to trial the course flow in blocks with no half-day GS even though these features seems slightly less effective as they wanted to evaluate their pedagogical impacts. The tool will be used next to study the impact of overhauling the rotary-wing and multi-engine programs as well as Phase IV for jet pilots. The mitigation options will undoubtedly be examined for those programs as well as when plans to significantly increase the pilot production are put forward.

REFERENCES

Burke, E.K., and S. Petrovic. 2002. “Recent Research Directions in Automated Timetabling.” *European Journal of Operational Research* 140/2: 266-280.

Caron, J.-D., and B. Woodman. 2003. “Resource Allocation Model for NFTC Phase IV.” Research Note 0303, Centre for Operational Research and Analysis, Defence R & D Canada, Ottawa, Canada.

Ibaraki, T., and N. Katoh. 1988. *Resource Allocation Problems (Algorithmic Approaches)*. Boston: The MIT Press.

Séguin, R. 2011. “1 Canadian Forces Flying Training School (1 CFFTS) Resource Allocation Simulation Tool.” In *Proceedings of the 2011 Winter Simulation Conference*, Edited by S. Jain, R.R. Creasey, J. Himmelspach, K.P. White, and M. Fu, 2495–2506. Piscataway, New Jersey: IEEE, Inc.

AUTHOUR BIOGRAPHIES

RENÉ SÉGUIN holds a Ph.D. in Operations Research from the University of Montréal. Since 2002, he has been a Defence Scientist in the Centre for Operational Research and Analysis of Defence Research and Development Canada. He is currently working on readiness issues pertaining to the Canadian Armed Forces. He is an associate editor of *INFOR*. His current research interests are resource allocation, scheduling and simulation. His email address is rene.seguin@drdc-rddc.gc.ca.

CHARLES J HUNTER holds an MSc in Statistics from the University of Manitoba. He joined the Operational Research and Analysis Establishment in 1983. During his 30 years with National Defence he has spent 26 of these directly supporting the Air Force at various headquarter locations. During the past decade he has focused his efforts on pilot production modeling and training issues. His email address is charles.hunter@forces.gc.ca.