



Interfaces

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Optimal Allocation of Students to Naval Nuclear-Power Training Units

Michael R. Miller, Robert J. Alexander, Vincent A. Arbige, Robert F. Dell, Steven R. Kremer, Brian P. McClune, Jane E. Oppenlander, Joshua P. Tomlin

To cite this article:

Michael R. Miller, Robert J. Alexander, Vincent A. Arbige, Robert F. Dell, Steven R. Kremer, Brian P. McClune, Jane E. Oppenlander, Joshua P. Tomlin (2017) Optimal Allocation of Students to Naval Nuclear-Power Training Units. *Interfaces* 47(4):320–335. <https://doi.org/10.1287/inte.2017.0905>

This article was written and prepared by U.S. government employee(s) on official time and is therefore in the public domain.

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes. For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Optimal Allocation of Students to Naval Nuclear-Power Training Units

Michael R. Miller,^a Robert J. Alexander,^a Vincent A. Arbige,^a Robert F. Dell,^b Steven R. Kremer,^a Brian P. McClune,^a Jane E. Oppenlander,^c Joshua P. Tomlin^a

^a Naval Nuclear Laboratory, Kesselring Site, Schenectady, New York 12301; ^b Operations Research Department, Naval Postgraduate School, Monterey, California 93943; ^c School of Business, Clarkson University, Schenectady, New York 12308

Contact: michaelr.miller@unnpp.gov (MRM); bobby.j.alexander@gmail.com (RJA); varbige@gmail.com (VAA); dell@nps.edu (RFD); skremer@nycap.rr.com (SRK); bpmcclune@gmail.com (BPM); joppenla@clarkson.edu,  http://orcid.org/0000-0001-8778-6461 (JEO); JoTomlin509@gmail.com (JPT)

Received: November 12, 2015

Revised: July 11, 2016; December 1, 2016

Accepted: March 24, 2017

Published Online in Articles in Advance:
June 27, 2017

<https://doi.org/10.1287/inte.2017.0905>

Copyright: This article was written and prepared by U.S. government employee(s) on official time and is therefore in the public domain.

Abstract. The U.S. Navy operates an impressive fleet of nuclear-powered submarines and aircraft carriers and has safely operated its nuclear fleet for more than 60 years, while steaming over 154 million miles. Rigorous training has been key to maintaining such an impressive record. The U.S. Naval Nuclear Propulsion Training Program develops, certifies, and delivers the nuclear-operator qualification training for enlisted and officer personnel operating its nuclear fleet. This training finishes at one of four nuclear-power training units (NPTUs), operates under a complex set of hard and soft constraints, varies depending on the type of student, and requires significant personnel and equipment resources. We developed and implemented a mixed-integer linear program (MILP) that prescribes how many students of each type to allocate to each NPTU at the start of each class (a group of students who train together) and how allocated students complete NPTU training. The use of MILP has improved student allocation by an estimated eight percent and led to significantly improved use of both NPTU personnel and equipment resources. In this paper, we describe this unique optimization application, the MILP formulation, its path to adoption, its user interface, and impacts from its development and use over the past three years.

History: This paper was refereed.

Funding: The submitted manuscript has been authored by contractor of the U.S. Government [Contract DE-NR-0000031]. Accordingly, the U.S. Government retains a non-exclusive, royalty-free license to publish or reproduce the published form of this contribution, or allow others to do so, for U.S. Government purposes.

Keywords: military • personnel: programming • integer • applications: education systems • planning: decision analysis • applications

Nuclear-powered submarines and aircraft carriers (Figure 1) are key elements for the defense of the United States and for the maintenance of free and open commerce across the world's oceans (Department of the Navy and Department of Energy 2014). These vessels are staffed by highly trained enlisted and officer personnel who operate and maintain the power-generation and propulsion systems capable of extended unsupported operations. The U.S. Naval Nuclear Propulsion Training Program develops, certifies, and delivers the nuclear-operator qualification training for enlisted and officer personnel who operate its nuclear fleet. This training finishes at one of four Nuclear Power Training Units (NPTUs). This paper describes the benefits achieved by using a MILP to prescribe the number of students of different types to allocate to each NPTU at

the start of each class and the activity sequence for allocated students to complete NPTU training.

Certifying Nuclear Operators

Certification as a naval nuclear operator requires rigorous training that lasts at least one year for each of five student types referred to by the name of the certification: electrician's mate; machinist's mate; electronics technician; engineering laboratory technician; and engineering officer of the watch. Each student type completes a unique training track consisting of knowledge and hands-on requirements. A student is certified (i.e., qualified) to operate a specific area of a naval nuclear-propulsion plant only after demonstrating mastery of propulsion plant equipment.

Figure 1. The Nuclear-Powered Aircraft Carrier USS JOHN C STENNIS (CVN 74), with Destroyer Escort (to Left), and the Nuclear-Powered Submarine USS SEAWOLF (SSN 21) Operates on Deployment in the Pacific Ocean



Notes. Operating the nuclear power plants on these ships requires highly trained enlisted and officer personnel. (Photo from <http://navy.mil>.)

Depending on the student type, students attend one or more schools prior to beginning NPTU training. While required schools vary by student type, all student types must satisfactorily complete a six-month program (i.e., nuclear-power school), consisting primarily of classroom instruction, prior to six additional months of NPTU training. Ideally, upon completion of this classroom instruction, students immediately begin training at one of four NPTUs at one of two training sites; each site has two units. During NPTU training, students engage in a mix of classroom, simulator, and hands-on training. Delays in starting NPTU training often occur due to limited resources; this produces a backlog of students waiting to begin NPTU training. Navy leadership carefully monitors this backlog.

Each NPTU is a self-contained training facility composed of a nuclear reactor, simulators, classrooms, staff instructors, and other training assets. Each NPTU class consists of a group of students who train together, which is designated by a sequential number based on the fiscal year. For example, class 1501 corresponds to the first class started in fiscal year 2015. The starting week of each NPTU class is known, and all plants start classes on the same day; therefore, each plant runs

classes with the same class number. With rare exceptions, a new class starts every eight weeks and three classes normally train simultaneously at each NPTU. Students are assigned (i.e., allocated) to a NPTU class and train together as a class.

NPTU training consists of a classroom phase (seven weeks) followed by a hands-on phase (17 weeks). Hands-on activities at the NPTU are critical because they teach students to perform the tasks that are required to safely operate nuclear reactors. Much of this hands-on training (referred to as “watchstanding”) takes place at “watchstations” located in either a plant that contains a nuclear reactor or at a simulator. Training conducted outside of the plant is referred to as “off-watch” training. Although extremely realistic, simulator training can only satisfy a fraction of the required watchstanding. A qualified staff instructor must be present at each watchstation in both the plant and a simulator to ensure proper plant operation. During operations, students are able to perform watchstanding. Simulators require staff time only when operating for training.

There are five types of staff instructors, each performing different training functions. Whereas students arrive in batches at defined intervals as part of a class,

staff instructors continuously flow in and out of a NPTU based on their individual military assignments. Typically, a staff instructor is assigned to a NPTU for three years consistent with the typical length of other Navy assignments. In addition to supporting hands-on training in the simulator and plant, staff instructors engage in a variety of other duties, including plant operation, providing classroom instruction, and administrative tasks. The amount of time devoted to each duty depends on staffing levels and the number of students being trained.

It is generally desirable to allocate as many students as possible to a class, while satisfying a variety of constraints and assuring the efficient use of limited staff instructors, equipment, and facilities. The number and type of students impacts the use of resources; each student type has a different set of qualification requirements, some that are unique and others that are common to other student types. Assigning the number and type of students to a training class is referred to as “student allocation” (or simply “allocation”). The training capability model (TCM), a MILP, prescribes how many students of each class and type to allocate to each NPTU; prescribes weekly staff-instructor assignments; and prescribes weekly student watchstanding and off-watch training.

Historical Approach

For more than 20 years prior to the adoption of the TCM, a single training analyst made student allocations using an iterative process, with expert judgment

applied at each iteration. Because of the importance of the allocation, a second person verified calculations to ensure potential errors were minimized. The analyst used a spreadsheet application to store data and assist with calculations. Over time, this spreadsheet application grew to more than 100 worksheets, which included numerous formulas and calculations, aided by Visual Basic for Applications (VBA) code (Microsoft 2015). The analyst required days to plan a single student allocation and the iterative effort was difficult to duplicate for any what-if analyses. In addition, several simplifying assumptions were employed for staff and simulator availability.


The Need for Optimization and an Expanded Model

Figure 2 shows the relationship between the number of nuclear operators qualifying each year and the available number of NPTUs. In recent years, the annual number of students requiring training has stayed near historic highs, while the number of NPTUs has decreased. This has necessitated the increased use of simulators and increased staffing levels. This in turn has complicated the task of determining student allocations. An improved student-allocation method capable of being used by more than a single analyst and capable of rapid what-if analysis was considered essential to best utilize the few remaining NPTUs.

Literature Review

Naval nuclear operator training is unique, but it shares much in common with other military training.

Figure 2. In the 1980s, the Number of NPTUs (Solid Line Graph) Was Eight; It Is Four Today and We Expect It to Decrease to Three After 2017, While the Number of Students Requiring Training, Which We Express as a Percentage of the Peak Number Trained in 1983 (Dotted Line Graph), Has Increased in Recent Years



Military training often involves the need to complete a sequence of qualification activities, which take place at “schools.” Each school lasts a number of weeks, sometimes occurs at different locations, and is only available periodically. Each military specialty typically has unique schools and schools common to other military specialties. Waiting often occurs between the end of one school and the start of another. Minimizing this waiting, or minimizing the backlog in the case of the TCM, is desirable. Grant (2000) minimizes waiting time for Marines by selecting the military occupational specialty (MOS) for each graduate of a common beginning school. Here, each MOS has its own series of schools, the timetable of when each school class begins is given, and capacity is simply the maximum number of students allowed in any class. Detar (2004) and Whaley (2001) also seek to minimize the waiting time of Marines between schools by prescribing a timetable of when each school course should start and how many students should be allocated to each class, with each capacity again simply the maximum number of students allowed in any class. Capacity constraints for the TCM are more complex because each student type impacts various training resources in different ways.

There is substantial operations research literature on military manpower planning as it relates to managing and growing military services. Early published work on hierarchical organizations can be found in Seal (1945) and Vajda (1947). The military services employ various models to determine recruiting, promotion rates, and retirement (Ginther 2006, Gibson 2007, Workman 2009). Wang (2005) provides a review of operations research applications in manpower planning, mostly with a focus on military training. His review includes applications that address optimization in the areas of: cost minimization for hiring and redeployment, personnel promotion, recruitment, and the mix and frequency of training modes (e.g., simulators, training aids) to maintain force proficiency. In general, these military manpower planning models have little in common with the TCM.

There is also substantial operations research literature on the related problem of course scheduling, where prescriptions assign students and instructors to classes, and classes to rooms and times; examples include de Werra (1985), Bonutti et al. (2012), and their extensive reference lists. The TCM


primarily differs from these course-scheduling applications because different student types train simultaneously and impact resources in different ways.

Similar resource-allocation problems can be found in healthcare literature. Caunhye et al. (2012) and Cardoen et al. (2010) provide reviews of the literature for emergency logistics and operating room planning, respectively. Examples of the allocation of operating room capacity using mixed-integer programming can be found in Zhang et al. (2009) and Blake and Donald (2002). These applications do not include prerequisite events, as required in the naval nuclear operator training environment. Integer programming and simulation are used in decision support systems that allocate medical assets during public health emergencies (Lee et al. 2009) and improving emergency department operations (Lee et al. 2015). Ernst et al. (2004) give an annotated bibliography of over 700 papers on personnel scheduling and rostering in a variety of application areas.

TCM Formulation

The TCM employs an elastic MILP to determine the number of students of each type to assign to each NPTU to comprise each training class over multiple years. Additionally, the TCM determines the number of simulator sessions each week at each NPTU. It does not explicitly plan the watchstanding sequence for each individual student. Instead, for each week, it plans the off-watch hours, plant watchstanding hours for each watch, and simulator watchstanding hours for each watch for all students of each student type and class at each NPTU. It establishes a preferred watchstanding window for each class at each NPTU with soft constraints that limit the number of watchstanding hours occurring very early, early, and late with respect to this preferred watchstanding window (Figure 3).

The TCM MILP models a number of practices that balance the competing needs for supplying qualified operators to the fleet and providing for plant maintenance periods and limits on staff working hours. An important objective is to minimize the number of students who must wait to receive NPTU training. Consequently, the first (and primary) term of the TCM objective function, which we show in Section A.5, Equation (A.1) in the appendix, imposes a penalty for each student in the training backlog (i.e., each student

Figure 3. The Preferred Watchstanding Window for a Class at a NPTU Occurs Between Its 14th and 21st Training Week

Notes. Up to five percent of the total watchstanding requirements can occur in weeks 8 and 9 and up to 50 percent of the requirements in weeks 10–13; however, any early watchstanding incurs a penalty. Late watchstanding is allowed, with increasing penalty severity, beyond week 21.

waiting to begin training). In some situations related to plant availability, a class of students is not assigned to a NPTU. This is referred to as “skipping a class” and should be avoided if at all possible. The second term of the objective function penalizes skipping classes. The third term of the objective function limits staff time to only what is needed to provide student training. Finally, a set of elastic penalties guide a solution to numerous goals.

Next, we give a summary of the primary prescriptions and constraints to provide a general understanding of the richness of the TCM MILP. Mathematical details can be found in the appendix.

The primary TCM variables are as follows.

- The integer number of each student type to start each class at each NPTU;
- The integer number of each student type waiting for training after the start of each class;
- The integer number of simulator sessions for each simulator at each NPTU each week;
- The number of hours each student type in each class performs each watchstanding requirement in its NPTU plant each week;
- The number of hours each student type in each class performs each watchstanding requirement in its NPTU simulator each week;
- The number of hours each student type in each class at each NPTU performs off-watch training each week; and

- The number of hours each staff-instructor type at each NPTU is assigned for off-watch and simulator watch instruction each week.

We organized the hard and soft constraints for assigning students to a training class and NPTU into five groups. In the following description, we use “goal” for a soft constraint (a constraint that can be violated at a cost), “limit” for a hard constraint, “each” for a constraint that exists for each permitted value of an index, and “all” when summing over all permitted index values. In Sections A.3 and A.5 in the appendix, we give the individual constraint sets and associated mathematical details for each group in the order presented here.

The class-composition constraint group establishes goals and limits on class size and distribution of student types; see constraints (A.2)–(A.7) in Sections A.3 and A.5 in the appendix. These constraints include:

- Bookkeeping to keep track of the backlog of each student type waiting for training after the start of each class;
- A lower goal and an upper goal for each student type in each class at each NPTU;
- A lower limit and an upper goal on all students in each class at each NPTU;
- An upper limit on all students at each NPTU across all simultaneous classes;
- An upper limit on all students in all simultaneous classes at each site; and
- An upper limit on all students of each student type in each class at each site.

The student-training constraint group establishes goals and limits on watchstanding in each NPTU simulator and plant; see Constraints (A.8)–(A.14) in Sections A.3 and A.5 in the appendix. These constraints include:

- A lower limit on required training hours for each student type in each class at each NPTU and each watchstation;
- An upper limit on the total training hours for each student type in each class at each NPTU for each watch in each week;
- An upper goal on total training hours for each student type in each class at each NPTU each week;
- An upper limit on total simulator and off-watch hours each week for all simultaneous classes for each student type at each NPTU; the upper limit is set by TCM decisions on staff-instructor assignments;
- An upper limit on total simulator watchstanding for each watch across all weeks for each student type in each class at each NPTU;
- An upper limit on simulation watchstanding for each watch across all classes and all student types for each week at each NPTU; and
- An upper limit on each NPTU plant's watch hours each week for all simultaneous classes and all student types.

The staff-instructor work constraint group establishes goals and limits on the assignment of staff hours; see constraints (A.15)–(A.17) in Sections A.3 and A.5 in the appendix. These constraints include:

- An upper goal on staff-instructor hours available for simulator and off-watch instruction for each staff type at each NPTU each week; both weekly control limits and cumulative sustained limits on staff-instructor availability are set;
- A lower limit on the number of staff hours required for each simulator session watch at each NPTU each week; the lower limit is set by a TCM decision on the number of simulator sessions; and
- A lower goal and an upper goal on simulator sessions for each simulator at each NPTU each week.

The watch placement constraint group establishes goals and limits on the pace of student training relative to the preferred watchstanding window; see constraints (A.18)–(A.21) in Sections A.3 and A.5 in the appendix. These constraints include:

- An upper goal on the percentage of watchstanding to be completed prior to a very early week (and an

early week) for each student type in each class at each NPTU and each watch; and

- An upper limit on the percentage of student off-watch hours relative to watchstanding hours for each student type in each class at each NPTU during early and late weeks.

Finally, the persistent (Brown et al. 1997) constraint group establishes goals and limits on adherence to a desired partial solution; see constraints (A.22) and (A.23) in Sections A.3 and A.5 in the appendix.

TCM Implementation Features

Personnel impacted by TCM prescriptions drive many of the features of TCM and its objective function. For any student allocation, the throughput of students (and the expedient reduction of any student backlog) are of primary concern, but secondary considerations, such as the efficient use of staff and plant training resources, are also considered. The TCM reflects this tiered priority structure with different penalty values for its objective function terms. The implementation also includes a time-based “reverse” discount factor applied to encourage students to complete training earlier rather than later in the planning horizon.

Implementers and approvers of the TCM prescriptions expect that small perturbations in inputs to the TCM (e.g., a small adjustment in staff-instructor hours available for weeks during a class) will yield no or at most only small changes in the TCM student-allocation prescriptions. We address this expectation by implementing persistence (Brown et al. 1997) as a feature, allowing a previous student allocation to be referenced as a preferred target.

Additional features are driven by practical considerations. Although a typical planning horizon spans two to four years, 10-year student allocations are often required to evaluate the expected long-term performance of the training program. For such long-term instances, the TCM employs another time-based discount that ensures near-term constraint violations (excluding those for classes already in training at the time of the allocation, which we designate as “fixed” in the appendix) are penalized more than violations that occur further in the future. Solutions over longer horizons can quickly stretch both the bounds of runtime practicality and the limits of desktop computer

resources. To ensure the TCM provides long-term prescriptions on standard desktop machines, we implement a solution cascade or receding-horizon solution (Brown et al. 1987, Baker and Rosenthal 1998). This method uses a rolling horizon to solve overlapping subsets of the planning horizon, thus reducing solution time. While this method has no guarantee of optimality, in practice, the TCM prescriptions are face valid using a solution cascade, and TCM users now prefer this solution method.

We implement the TCM using the General Algebraic Modeling System (GAMS) (GAMS 2015a) and solve it using CPLEX (GAMS 2015b). For a typical cascade subset of the planning horizon, about 1.5 years, an instance of the TCM has approximately 500,000 constraints and 950,000 variables (950 of them integer). The cascade subsets typically overlap by 0.5 years. Solution time per cascade is approximately five minutes using a Windows 7 workstation with two 3.33 GHz Intel Xeon X5680 processors and 8 GB RAM. Typical TCM planning horizons of 2.5 years, for example, require two cascade iterations and take approximately 10 minutes to solve, while long-term student allocations of 10 years typically require about 50 minutes.

TCM users find solution times without the use of a rolling horizon undesirable. For example, we recently conducted some experiments with real instances that have a planning horizon of 3.5 years. Using the preferred method of three cascade iterations, these took between 10 and 15 minutes to solve. Solving these instances without a rolling horizon took between 30 and 45 minutes. An examination of their respective solutions showed the results obtained using both methods were almost identical; the distributions of students to classes and weekly instructor staff workload varied only slightly. Empirical results such as these have led TCM users to almost exclusively use solution cascades. Despite this strong preference, we maintain an interface option for TCM users to easily disable solution cascades or adjust the subset of the planning horizon considered for each cascade subset.

For improved solution time, TCM users are also often willing to accept a solution that is only guaranteed to be within approximately one percent of optimality. With such a permitted gap, staff assigned hours (the third term of the objective function) may exceed

the number required for a given allocation. This cosmetic annoyance is corrected by solving a revised MILP that fixes the student allocation, thereby fixing the first two terms of the objective function, and minimizes the third term of the objective function.

TCM Testing Before Adoption

Several management layers were needed to approve the adoption of the TCM; therefore, we designed a rigorous test program. At the time of the TCM development, decision makers had no approved and universally applied criteria when evaluating allocations. Management judged new student allocations primarily based on the backlog and how they compared with historic allocations. Given this history, the first phase in the test program was to engage a broad range of stakeholders to develop objectives, criteria, and metrics for allocations using value-focused thinking techniques (Keeney 1994). Ewing et al. (2006) and Parnell and West (2011) provide examples of applying value-focused thinking. This resulted in the following overarching objective statement: The allocation process seeks to maximize student throughput with on-time training completion while efficiently utilizing staff and facilities. The characteristics of a good allocation were identified as equity in class assignments across NPTUs, on time completion, and full utilization of training resources. We then defined a set of 11 questions reflecting these characteristics for use by training experts in evaluating an allocation.

Simultaneously, training experts established 13 benchmark test cases representing a range of routine and abnormal scenarios. For each test case, several training experts were asked to independently evaluate TCM results using the 11 established questions, each rated on a four-point Likert scale (i.e., unsatisfactory, marginal, good, excellent). Evaluations, including open-ended comments, were submitted to a database created to enable long-term collection and analysis of TCM results. Once all test cases had been evaluated, statistics were computed quantifying the acceptability of the allocation and the associated inter-rater reliability using the method from Fleiss (1971).

Each testing round concluded with a meeting of the evaluators and model developers to discuss the results and agree on model refinements, if needed. For

example, one round resulted in changes to the staff-instructor workload model. The evaluators found it unacceptable to expect sustained periods of high staff-instructor workload, even in theoretical cases of low instructor staffing with high demand for qualified students. A new TCM constraint was implemented to control sustained staff-instructor workload. This feature allows a periodic surge of high workload if it is balanced by a period of lower workload.

In addition to benchmark testing, a series of retrospective tests were performed using the most recent two years of data. In the first test, the TCM prescribed student allocations, given actual plant availability and staffing levels. In the second test, actual student allocations were also fixed, with the objective of solving for staff workload. The TCM has elastic constraints that allow minor violations, and we experimented extensively with penalties for these violations to ensure that all penalties were scaled, such that each was meaningful, and to capture the trade-off between excessive staff workload and having a backlog of students waiting for training. Following this experimentation, default penalty values were established where the student throughput was less than a five percent difference from historical values and staff workload was within acceptable ranges.

Near the end of development, the TCM was run in parallel with the legacy model for several allocations. In the parallel operations, the primary training planner found the results of the TCM acceptable. In addition, the TCM provided insights that were previously unavailable. These insights, coupled with the results from the benchmark and retrospective testing, were reviewed with training management who requested immediate adoption of the TCM.

Interface Design

One principle analyst and several assistants are responsible for using the TCM to produce student allocations. Preparing TCM input requires considerable knowledge of student-training database systems and the TCM is run many times each week. A broad range of decision makers rely on its prescriptions to both determine the student allocation and to plan (e.g., staff-instructor schedules and plant maintenance periods). In addition, decision makers frequently request what-if analyses representing different training scenarios. All

these decision makers require graphical displays of the prescriptions.

A Microsoft Excel VBA application serves as the TCM interface. It contains all TCM documentation and input spreadsheets, serves to obtain input data from several independent databases, calls GAMS, produces all the TCM output reports, and displays all the graphics. The TCM produces 15 standard graphs; Figures 4–9 show examples. The TCM replicated the basic look and feel of all legacy graphics, while also providing new visualizations to display information not previously available in the legacy application.

Results and Impact

Over three years, the insights gained from the use of the TCM have increased the number of students trained by an estimated eight percent (when compared with the legacy model). This improvement stems primarily from a holistic understanding of how student training is impacted by the interrelationships between plant availability, staffing, facility availability, and simulator utilization. The ability to rapidly conduct what-if analysis that explicitly considers these interrelationships has led to these new insights; as a result, decision makers have altered their allocation decisions to better balance available resources.

The key to communicating the TCM prescriptions is effective visual displays. The primary display for

Figure 4. The Student Backlog (Number and Type Represented by the Different Shades in the Bar Graph) as a Function of Time (the Horizontal Axis) Provides Executive-Level Decision Makers with Key Information Concerning the Flow of Students through the Training Program, and Forms the Basis for Comparison of What-If Scenarios Involving the Allocation of Resources





Figure 5. Student Backlog Is Presented for Three Cases Based on the Placement of an Extended Maintenance Period, Which Would Result in Skipping Several Classes of Students at One NPTU



executive-level decision makers is the expected student backlog. Figure 4 shows an example of the projected backlog of four student types waiting for NPTU training over time. This figure shows that initially the student backlog is large because of resource constraints (i.e., qualified staff instructors, plant availability, or facilities). As student-training resources become available with each succeeding class, the backlog of students reduces. Subsequently, resources are again constrained resulting in a rise in the backlog. These fluctuations continue in response to changing resource profiles.

Figure 5 shows an example of a what-if scenario. The scenario is to determine the impact on student training based on the placement of a required maintenance period, which will make the plant unavailable

for student training for an extended period and result in several skipped classes. The baseline indicates the current schedule for the maintenance period; alternate 1 places the maintenance period in the schedule one year later; alternate 2 places it two years later. Schedules for the other plants are unaffected and all other resources are fixed; the only resource affected is the availability of the plant at which the maintenance is taking place. The comparison clearly shows that the baseline generates the fewest number of backlogged students; however, if the maintenance is moved from the current placement, alternate 2 would be the preferred option.

The TCM prescribes weekly staff-instructor variable workload across the planning horizon for each of the five types of staff. Figure 6 shows a staff-instructor workload display for one instructor type. It displays the fixed workload (duties required to keep the plants operational) and the variable workload for simulator

Figure 6. For Each Staff-Instructor Type, Workload Is Controlled Within the Desired Level, Sustained Limit, and Control Limit by Varying the Number and Type of Students Allocated to Each Class






Figure 7. Each Bar Represents the Total Number (and with Shading the Actual Number) of Students Allowed to Simultaneously Train at the Start of Class at a NPTU



Note. The student limits are a proxy for facility capacities, such as classrooms, study cubicles, and computer stations.

Figure 8. TCM Output Provides a Graphical Representation of Plant Availability in Conjunction with the Watchstanding Timeline Associated with Each Class



Notes. The figure presented is simplified and truncated for readability. Plant availability, student allocations, and watchstanding preferences by class are shown. A box in each plant's row indicates the week when a class completes watchstanding. For example, at Plant 1, Class 2 completes watchstanding in Week 65.

watchstanding and off-watch training as prescribed by the TCM. Figure 6 also displays the desired staff workload level, the maximum sustained workload limit, and the maximum control limit—a surge limit that must be offset by reduced hours a few weeks prior to or after the surge.

Insights provided by the TCM into the details of staff-instructor workload for each staff type yielded one of the greatest early benefits from adopting TCM, in part because the legacy model did not explicitly consider instructor workload. By using TCM, it became clear that some staff types were working at capacity (and thereby preventing additional student allocations), while other instructor types had hours available. Equipped with this new insight (and verification by staff instructors), changes were implemented to better balance workload among staff-instructor types and thereby increase student allocations. These changes included modifications of work assignments and adjustments in the number and type of instructors assigned to the NPTUs.


Optimally planning staff-instructor workload also allowed increased student allocations when part of training for a class coincided with a plant shutdown. Before the TCM, a simple rule of thumb dictated that only a few or no students would be allocated to a plant

when its availability prevented watchstanding completion before a standard number of weeks. Using the TCM, it became apparent that this rule of thumb was too restrictive. The TCM showed how to maintain near-normal student allocations by using increased staff-instructor hours during the shutdown for off-watch and simulator training. Sometimes this needs to be coupled with a modest extension beyond a normal training deadline, with the TCM providing the details used to obtain permission for such an exception to normal operations.

The TCM also ensures that student loading for each class fits within the capacity limits of the training facilities. This includes limits for individual classes (e.g., classroom number and size, study cubicles, computer stations), and cumulative NPTU and site-loading limits for all classes projected to be at a facility at one time. TCM does this by limiting the number of students in each NPTU and each site for all simultaneous classes. Figure 7 shows an example of student capacity of the training facilities over time.

Figure 8 shows a simplified summary of the TCM results for a student allocation at the four NPTUs; these results replicate much of a legacy report used by management. The timeline shows how the allocation fits within each plant's operating schedule, with no plant watchstanding occurring during maintenance periods.

Figure 9. TCM Provides Detailed Information for Each Plant and Class (with Only Plant 1 Shown for Readability)



Notes. For example, in Week 66, Class 2 has one simulation session and Class 3 has 20. The last row reports the percentage of plant watchstanding completion for each class. For example, in Week 64, Class 2 has 92 percent of its watchstanding complete and Class 3 has four percent.

The relationship between the preferred watchstanding periods for the three classes appears in the top portion of the timeline. The timeline view includes the size and composition of each class (represented in Figures 8 and 9 by the “# Students” in each of the class boxes), simulator utilization, and class completion. Figure 9 magnifies the timeline to show detailed information for each class. The timeline view is particularly useful to those making short-term training decisions.

In addition to the what-if scenario presented in Figure 5 (i.e., variation in the placement of planned maintenance periods), other typical scenarios involve changes to the number of staff instructors, additional facilities to expand student capacity, and the type and quantity of simulation equipment that could be integrated into the training program. In each of these what-if scenarios, the key element is the number of students that can be trained and the cost of training them, including personnel, facilities, and (or) equipment costs, as compared to the required number of trained students that must be transferred to the fleet to maintain the desired staffing levels onboard submarines and aircraft carriers. These what-if scenarios have driven changes to the number of assigned instructor staff, placement of maintenance periods, and strategic decisions concerning future investments in equipment and recapitalization.

The use of TCM, with its ability to conduct multiple what-if scenarios that produce accurate and repeatable results in the current resource-constrained environment, has enabled both tactical and strategic decisions to ensure fleet staffing needs are continually

met to satisfy U.S. Navy operational and strategic requirements.

Conclusions

The TCM provides an optimal use of the key resources needed to qualify naval nuclear operators in a challenging operational and budgetary environment. The TCM has helped increase the number of students trained by showing how to better employ these key resources, especially staff instructors. It eliminated the long-standing dependence on the few (i.e., only one or two) experts who can perform capacity analysis. Four analysts now routinely use the TCM.

The TCM’s prescriptions were quickly found to be superior to the legacy model. Today, decision makers frequently request a broad range of what-if analyses that rely on using the TCM.

Acknowledgments

The authors thank Martin Andrew, Fred Lanou, and Paul Zanella for their executive support throughout the development and deployment of the TCM, Scott Ciampa (Naval Nuclear Laboratory), Dan Arguello (Naval Nuclear Laboratory), Karina Rodriguez (Naval Nuclear Laboratory), and Anton Rowe (Naval Postgraduate School) for their technical contributions, Torre Bissell (retired) for his mentoring and software development expertise, and LT Kai Seglem (USN) for his many contributions to initial TCM development. Additionally, the authors thank the numerous Navy junior officers whose contributions at different stages of the project ensured a successful product delivery. Finally, the authors thank the *Interfaces* editors and referees for their thoughtful reviews that helped improve this article.

Appendix

We present a representative TCM formulation using $\overset{\circ}{\geq}$, $\overset{\circ}{\leq}$, and $\overset{i}{\leq}$ as shorthand for elastic constraints (i.e., constraints that can be violated at a cost per given unit of violation), as discussed in Brown et al. (1997). The elastic constraints represented by $\overset{\circ}{\geq}$ and $\overset{\circ}{\leq}$ can be violated in any period, while the elastic constraints represented by $\overset{i}{\leq}$ can only be violated for periods corresponding to the first τ classes, whose members have already begun training and are considered “fixed.” For example, if τ equals 2, with Class 1 allowed to stand watch for weeks $t \in \{1, 2\}$ and Class 2 weeks $t \in \{1, 2, \dots, 9\}$, elastic constraints represented by $\overset{i}{\leq}$ could only be violated where $c \in \{1, 2\}$ and $t \in \{1, 2, \dots, 9\}$. Let c (or alias c') be the index for class number; then the allocation for classes $c \in \{1, 2, \dots, \tau\}$ are fixed as starting conditions for the TCM and $c \in \{\tau + 1, \tau + 2, \dots\}$ are not fixed. Because these fixed classes may violate desired planning, we must have the ability to violate constraints for these fixed classes. Some of the terms in the constraints are applicable for only a single student type (referred to as rate $r = r3$).

In the following Sections A.1–A.6, we present the TCM formulation for indices, index sets, parameters, variables, objective function, and constraints, and then conclude with a brief description. Beale et al. (1974) originally documented this ordering, which was adhered to for decades under the label NPS format in hundreds of published theses and papers by Naval Postgraduate School (NPS) students and faculty, with acknowledgment of the original Beale et al. reference (Brown and Dell 2007), and recently reintroduced with additional guidance by Teter et al. (2016).

A.1. TCM Indices [~Cardinality]

- c, c' class [6 per year];
- d training site [2];
- j simulator type [2];
- p NPTU [4];
- r student type or rate [7];
- s staff-instructor type [5];
- t week of the planning horizon [208];
- w, w' watch to stand [25].

A.2. TCM Index Sets

- $p \in TS_d$ set of NPTU p at training site d ;
- $r \in RW_w$ set of all student types r that stands watch w ;
- $s \in SW_w$ set of staff s that can stand watch w ;
- $t \in AW_{cpr}$ set of all possible watch weeks t for class c , student type r , at NPTU p ;
- $t \in EW_{cr}$ set of all early watch weeks t for class c for student type r ;
- $t \in FW_{cpr}$ set of weeks t when class c can finish at NPTU p for student type r ;
- $t \in LW_{cpr}$ set of all late weeks t at the end of class c at NPTU p for student type r ;

- $t \in VW_{cr}$ set of all very early watch weeks t for class c for student type r ;
- $w \in AS$ set of all simulator watches w ;
- $w \in OW$ set of watches w satisfying off-watch requirements;
- $w \in PT$ set of watches w that require a plant;
- $w \in SS$ set of all watches w where a simulator can substitute;
- $w' \in SB_w$ set of watches w' that can substitute for watch w ;
- $w' \in SM_w$ set of simulator watch w' that can be substituted for watch w .

A.3. TCM Parameters [Units]

Parameters for the Objective Function

- pcr_{cr} penalty for student type r student waiting for training after the start of class c [penalty per student];
- pcp_{cp} penalty for not starting class c at NPTU p [penalty per class];
- psw_{sw} penalty for staff s standing watch w [penalty per hour].

Parameters for Constraints (A.2)–(A.7) (The Class-Composition Group)

- $roll_r$ number of student type r students waiting for training at the start of planning [students];
- new_{cr} number of newly arriving student type r for class c [students];
- $\underline{gcpr}_{cpr}, \overline{gcpr}_{cpr}$ lower and upper goal on the number of class c students of student type r desired at NPTU p [students];
- $\underline{lcp}_{cp}, \overline{gcp}_{cp}$ lower limit and upper goal on the number of class c students (excluding student type 3) at NPTU p [students];
- $\overline{lcp}_{cp}$ upper limit on the number of students NPTU p facilities can support during the first week class c is held [students];
- $\overline{lcd}_{cd}$ upper limit on number of students Site d facilities can support during the first week class c is held [students];
- $\overline{lcd}_{cdr}$ upper limit on number student type r site d facilities can support during the first week for class c [students].

Parameters for Constraints (A.8)–(A.14) (The Student-Training Group)

- req_{cprw} hours of class c watch w training required by student type r at plant p [hours per student];
- $done_{cprw}$ hours of watch w already completed at the start of planning by class c at NPTU p by student type r [hours];
- nww_{cprw} minimum number of weeks per class c , student type r , and watch w that must contain some watch-standing (divisor used to provide an upper bound on weekly training for each watch);

- $\overline{grt}_{rt}$ upper goal on the number of hours student type r can work in week t [hours per student];
- $\overline{lcrw}_{crw}$ upper limit on the number of class c watch hours allowed in simulators by student type r for watch w [hours per student];
- $\overline{ljp}_{jpw}$ hours available for simulator j , watch w for each simulator watch session at NPTU p [hours per session];
- $\overline{lptw}_{ptw}$ upper limit on hours of watch w available at NPTU p in week t [hours].

Parameters for Constraints (A.15)–(A.17) (The Staff-Instructor Group)

- $\overline{gpst}_{pst}$ upper goal on hours of staff s for assignment at NPTU p in week t [hours];
- $hfix_{pst}$ fixed hours for staff s at NPTU p in week t [hours];
- $hsim_{jpw}$ staff hours required for each j simulator watch w session at NPTU p [hours per session];
- $\underline{gjpt}_{jpt}, \overline{gjpt}_{jpt}$ lower and upper goal on the number of simulator j watch sessions in NPTU p in week t [sessions].

Parameters for Constraints (A.18)–(A.21) (The Watch Placement Group)

- $\overline{grv}_r$ upper limit on the fraction of student type r very early training allowed [hours per hours];
- $\overline{gre}_r$ upper limit on the fraction of student type r early training allowed [hours/hours].

Parameters for Constraints (A.22) and (A.23) (The Persistent Group)

- τ maximum class for index c for which allocations X_{cpr} are fixed for all p, r ;
- $fixX_{cpr}$ allocation of student type r from class c fixed to train at NPTU p [students];
- $oldX_{cpr}$ prior value of X_{cpr} used to guide a persistent solution [student].

A.4. TCM Variables

- X_{cpr} integer number of student type r to start class c at NPTU p ;
- K_{cr} integer number of student type r waiting for training after start of class c ;
- I_{jpt} integer number of simulator j sessions at NPTU p in week t ;
- G_{cp} binary variable with value one if class c is in session at NPTU p and zero otherwise;
- F_{cprtw} plant watch w training hours from fixed plant operations assigned; to student type r in class c at NPTU p in week t ;
- U_{cprtw} simulator watch w hours assigned to student type r in class c at NPTU p in week t ;
- V_{cprtw} Off-watch w hours assigned to student type r in class c at NPTU p in week t ;
- H_{pstw} staff s hours assigned at watch w (includes simulator and off-watch instruction only) at NPTU p in week t .

A.5. TCM Formulation

$$\text{Minimize } \left\{ \sum_{cr} pcr_{cr} K_{cr} + \sum_{cp} pcp_{cp} (1 - G_{cp}) + \sum_{ptw, s \in SW_w} psw_{sw} H_{pstw} + \text{elastic penalties} \right\} \quad (\text{A.1})$$

Subject to:

$$K_{cr} = roll_r |_{c=1} + K_{c-1, r} |_{c>1} + new_{cr} - \sum_p X_{cpr}, \quad \forall c, r, \quad (\text{A.2})$$

$$\underline{gcp}_{cpr} \leq X_{cpr} \leq \overline{gcp}_{cpr} G_{cp}, \quad \forall c, p, r, \quad (\text{A.3})$$

$$\underline{lcp}_{cp} G_{cp} \leq \sum_{r \neq r3} X_{cpr} \leq \overline{gcp}_{cp}, \quad \forall c, p, \quad (\text{A.4})$$

$$\sum_{c'=c-2, r \neq r3} X_{c'pr} + \sum_{c'=c-1, r=3} X_{c'pr} \leq \overline{lcp}_{cp}, \quad \forall c, p, \quad (\text{A.5})$$

$$\sum_{c'=c-2, r \neq r3, p \in TS_d} X_{c'pr} + \sum_{c'=c-1, r=3, p \in TS_d} X_{c'pr} \leq \overline{lcd}_{cd}, \quad \forall c, d, \quad (\text{A.6})$$

$$\sum_{p \in TS_d} X_{cpr} \leq \overline{lcd}_{cd}, \quad \forall c, d, r, \quad (\text{A.7})$$

$$\begin{aligned} & \sum_{t, w' \in SB_w} (F_{cprtw'} + V_{cprtw'} + U_{cprtw'}) \\ & \geq req_{cprw} X_{cpr} - done_{cprw}, \quad \forall c, p, r, w, \end{aligned} \quad (\text{A.8})$$

$$\begin{aligned} & \sum_{w' \in SB_w} (F_{cprtw'} + V_{cprtw'} + U_{cprtw'}) \\ & \leq \frac{(req_{cprw} X_{cpr} - done_{cprw})}{nw w_{crw}}, \quad \forall c, p, r, t, w, \end{aligned} \quad (\text{A.9})$$

$$\sum_w (F_{cprtw} + V_{cprtw} + U_{cprtw}) \leq \overline{grt}_{rt} X_{cpr}, \quad \forall c, p, r, t, \quad (\text{A.10})$$

$$\sum_{c, r \in RW_w} (V_{cprtw} + U_{cprtw}) \leq \sum_{s \in SW_w} H_{pstw}, \quad \forall p, t, w, \quad (\text{A.11})$$

$$\sum_{t \in AW_{cpr}, w' \in SM_w} U_{cprtw'} \leq \overline{lcrw}_{crw} X_{cpr}, \quad \forall c, p, r, w \in SS, \quad (\text{A.12})$$

$$\sum_{c | t \in AW_{cpr}, r \in RW_w} U_{cprtw} \leq \sum_j \overline{ljp}_{jpw} I_{jpt}, \quad \forall p, t, w \in AS, \quad (\text{A.13})$$

$$\sum_{c, r \in RW_w} F_{cprtw} \leq \overline{lptw}_{ptw}, \quad \forall p, t, w \in PT, \quad (\text{A.14})$$

$$\sum_{w \notin PT} H_{pstw} + hfix_{pst} \leq \overline{gpst}_{pst}, \quad \forall p, s, t, \quad (\text{A.15})$$

$$\sum_{s \in SW_w} H_{pstw} \geq \sum_j hsim_{jpw} I_{jpt}, \quad \forall p, t, w \in AS, \quad (\text{A.16})$$

$$\underline{gjpt}_{jpt} \leq I_{jpt} \leq \overline{gjpt}_{jpt}, \quad \forall j, p, t, \quad (\text{A.17})$$

$$\begin{aligned} & \sum_{t \in VW_{cr}} \left(F_{cprtw} + \sum_{w' \in SB_w \cap AS} U_{cprtw'} \right) \\ & \leq \overline{grv}_r req_{cprw} X_{cpr}, \quad \forall c, p, r, w \notin OW, \end{aligned} \quad (\text{A.18})$$

$$\begin{aligned} & \sum_{t \in EW_{cr}} \left(F_{cprtw} + \sum_{w' \in SB_w \cap AS} U_{cprtw'} \right) \\ & \leq \overline{gre}_r req_{cprw} X_{cpr}, \quad \forall c, p, r, w \notin OW, \end{aligned} \quad (\text{A.19})$$

$$\frac{\sum_{t' \geq t, w \in OW} V_{cprtw}}{\sum_{w \in OW} req_{cprw}} \geq \frac{\sum_{w \in PT} F_{cpr, t-1, w}}{\sum_{w \in PT} req_{cprw}} + \frac{\sum_{w \in PT, w' \in SB_w \cap AS} U_{cpr, t-1, w'}}{\sum_{w \in PT, w' \in SB_w \cap AS} req_{cprw'}}, \quad \forall c, p, r, t \in LW_{cpr}, \quad (A.20)$$

$$\frac{\sum_{w \in OW} V_{cprtw}}{\sum_{w \in OW} req_{cprw}} \geq \frac{\sum_{w \in PT} F_{cpr, t, w}}{\sum_{w \in PT} req_{cprw}} + \frac{\sum_{w \in PT, w' \in SB_w \cap AS} U_{cpr, t, w'}}{\sum_{w \in PT, w' \in SB_w \cap AS} req_{cprw'}}, \quad \forall c, p, r, t \in FW_{cpr}, \quad (A.21)$$

$$X_{cpr} = fix X_{cpr}, \quad \forall c \leq f c, p, r, \quad (A.22)$$

$$X_{cpr} \geq old X_{cpr}, \quad \forall c, p, r, \quad (A.23)$$

$$\begin{cases} X_{cpr} \geq 0 \text{ and integer, } & \forall c, p, r, \\ I_{jpt} \geq 0 \text{ and integer, } & \forall j, p, t, \\ K_{cr} \geq 0 \text{ and integer, } & \forall c, r, \end{cases} \quad (A.24)$$

$$G_{cp} \in \{0, 1\}, \quad \forall c, p, \quad (A.25)$$

$$\begin{cases} F_{cprtw} \geq 0, & \forall c, p, r, t, w \in PT, \\ U_{cprtw} \geq 0, & \forall c, p, r, t, w \in AS, \\ V_{cprtw} \geq 0, & \forall c, p, r, t, w \in OW, \\ H_{pstw} \geq 0 & \forall p, s, t, w \notin PT. \end{cases} \quad (A.26)$$

The objective function (A.1) expresses the total penalty value. The first term of the objective function is the weighted penalty of the student backlog; the second term is the weighted penalty for skipped classes; the third term is the weighted penalty for staff workload, and the last penalty term includes all elastic constraint violations (Section A.6 in this appendix). Constraint set (A.2) tracks student backlog (i.e., inventory of students by student type waiting to start training after the start of a class). For the first period ($c = 1$), the backlog is initial backlog ($roll_r$). The student backlog increases for any newly arriving students who cannot be trained in the current class, and decreases whenever more students may be trained than those newly arriving for the current class. Constraint set (A.3) measures deviation from desired lower and upper limits for each student type in each class at each plant, and ensures that no student of any student type is assigned to a skipped class. Constraint set (A.4) measures deviation from the desired upper and lower bounds for the total number of students in a class at each NPTU (excluding student type $r3$). Constraint set (A.4) removes the lower bound for skipped classes to ensure feasibility. Constraint set (A.5) sets an upper bound on the number of students in all simultaneous classes (classes $c - 2$ to c for all student types, excluding student type $r3$ and classes $c - 1$ to c for student type $r3$). Similarly, constraint set (A.6) sets an upper bound for total number of students a training site can support across all simultaneous classes. Constraint set (A.7) limits the number of each student type for a class at a site. Constraint set (A.8) ensures sufficient training hours are assigned for each student type at each watch station (watch station includes plant and off-watch requirements). Constraint set (A.9) limits weekly training hours for

each watch. Constraint set (A.10) restricts total weekly student work hours. Constraint set (A.11) limits student and off-watch hours to those with assigned staff. Constraint set (A.12) restricts simulation training to be no more than a user input fraction of the total for training that can be conducted using simulation. Constraint set (A.13) restricts simulator hours to those with assigned personnel. Constraint set (A.14) restricts plant hours. Constraint set (A.15) limits staff hours available for watch and off-watch duties. Constraint set (A.16) ensures adequate personnel for each simulation session. Constraint set (A.17) restricts weekly simulator sessions. Constraint set (A.18) allows no more than a user input fraction of the total plant-based training to be completed very early. Constraint set (A.19) allows no more than a user input fraction of the training of the total plant-based training to be completed early. Constraint sets (A.20) and (A.21) ensure sufficient off-watch training in the last weeks. Constraint (A.22) fixes the initial student allocation for the first classes. Constraint (A.23) measures negative deviation from a prior solution. Constraints (A.24)–(A.26) declare variable types.

A.6. Elastic Constraints

The TCM formulation has many elastic constraints that we have already introduced using notation shorthand. Here are the details for elastic constraint set (A.27), which controls the weekly and sustained staff-instructor workload by adding a new “elastic” variable O_{pst} for the overtime worked at NPTU p by staff type s in week t :

$$\sum_{w \notin PT} H_{pstw} + hfix_{pst} \leq \overline{gpst}_{pst} + O_{pst}, \quad \forall p, s, t. \quad (A.27)$$

In most cases, the addition of such an elastic variable (and its corresponding penalty term in the objective function) is all that is required to convert an elastic constraint from shorthand to more traditional notation. In this case, there are additional constraints on the elastic variable. Constraint set (F.1a) limits the extent of overtime allowed at NPTU p by staff type s in week t :

$$O_{pst} \leq \overline{over}_{pst}, \quad \forall p, s, t, \quad (F.1a)$$

and constraint set (F.1b) limits the sustained overtime allowed at NPTU p by staff type s over all watch weeks for each student type r and each class c :

$$\sum_{t \in AW_{cpr}} O_{pst} \leq \overline{overa}_{ps}, \quad \forall c, p, r, s. \quad (F.1b)$$

References

- Baker SF, Rosenthal RE (1998) A cascade approach for staircase linear programs. Technical Report NPS-OR-98-004, Naval Postgraduate School, Monterey, CA.
- Beale EML, Beare GC, Tatham PB (1974) The DOAE reinforcement and redeployment study: A case study in mathematical programming. Hammer PL, Zoutendijk G, eds. *Math. Programming Theory and Practice: Proc. NATO Adv. Study Inst., Figueira da Foz, Portugal* (Elsevier, New York), 417–442.

- Blake JT, Donald J (2002) Mount Sinai Hospital uses integer programming to allocate operating room time. *Interfaces* 32(2):63–73.
- Bonutti A, De Cescio F, Di Gaspero L, Schaerf A (2012) Benchmarking curriculum-based course timetabling: Formulations, data formats, instances, validation, visualizations, and results. *Ann. Oper. Res.* 194(1):59–70.
- Brown GG, Dell RF (2007) Formulating integer linear programs: A rogues' gallery. *INFORMS Trans. Ed.* 7(2):153–159.
- Brown GG, Dell RF, Wood RK (1997) Optimization and persistence. *Interfaces* 27(5):15–37.
- Brown GG, Graves GW, Ronen D (1987) Scheduling ocean transportation of crude oil. *Management Sci.* 33(3):335–346.
- Cardoen B, Demeulemeester E, Beliën J (2010) Operating room planning and scheduling: A literature review. *Eur. J. Oper. Res.* 201(3):921–932.
- Caunhye AM, Xiaofeng N, Pokharel S (2012) Optimization models in emergency logistics: A literature review. *Socio-Econom. Planning Sci.* 46(1):4–13.
- de Werra D (1985) An introduction to timetabling. *Eur. J. Oper. Res.* 19(2):151–162.
- Department of the Navy and Department of Energy (2014) *The United States Naval Nuclear Propulsion Program* (Department of the Navy and Department of Energy, Washington, DC).
- Detar PJ (2004) Scheduling Marine Corps entry-level MOS schools. Unpublished master's thesis, Naval Postgraduate School, Monterey, CA.
- Ernst AT, Jiang H, Krishnamoorthy M, Owens B, Sier D (2004) An annotated bibliography of personnel scheduling and rostering. *Ann. Oper. Res.* 127(1):21–144.
- Ewing PL Jr, Tarantino W, Parnell GS (2006) Use of decision analysis in the Army Base Realignment and Closure (BRAC) 2005 military value analysis. *Decision Anal.* 3(1):33–49.
- Fleiss JL (1971) Measuring nominal scale agreement among many raters. *Psych. Bull.* 76(5):378–382.
- GAMS (2015a) GAMS: Cutting edge modeling. Accessed July 15, 2015, <http://www.gams.com>.
- GAMS (2015b) CPLEX 12. Accessed July 15, 2015, <http://www.gams.com>.
- Gibson HO (2007) The total Army competitive category optimization model: Analysis of U.S. Army officer accessions and promotions. Unpublished master's thesis, Naval Postgraduate School, Monterey, CA.
- Ginther TA (2006) Army Reserve enlisted aggregate flow model. Unpublished master's thesis, Naval Postgraduate School, Monterey, CA.
- Grant JM (2000) Minimizing time awaiting training for graduates of the basic school. Unpublished master's thesis, Naval Postgraduate School, Monterey, CA.
- Keeney RL (1994) Creativity in decision making with value-focused thinking. *Sloan Management Rev.* 35(4):33–41.
- Lee EK, Chen C-H, Pietz F, Benecke B (2009) Modeling and optimizing the public-health infrastructure for emergency response. *Interfaces* 39(5):476–490.
- Lee EK, Atallah HY, Wright MD, Post ET, Thomas IV C, Wu DT, Haley LL Jr (2015) Transforming hospital emergency department workflow and patient care. *Interfaces* 45(1):58–82.
- Microsoft (2015) Introduction to the Visual Basic programming language. Accessed August 31, 2015, [https://msdn.microsoft.com/library/xk24xdbef\(v=vs.90\).aspx](https://msdn.microsoft.com/library/xk24xdbef(v=vs.90).aspx).
- Parnell GS, West PD (2011) Systems decision process overview. Parnell GS, Driscoll PJ, Henderson DL, eds. *Decision Making in Systems Engineering and Management* (John Wiley & Sons, New York), 275–295.
- Seal HL (1945) The mathematics of a population composed of k stationary strata each recruited from the stratum below and supported at the lowest level by a uniform number of annual entrants. *Biometrika* 33(3):226–230.
- Teter MD, Newman AM, Weiss M (2016) Consistent notation for presenting complex optimization models in technical writing. *Surveys Oper. Res. Management Sci.* 21(1):1–17.
- Vajda S (1947) The stratified semi-stationary population. *Biometrika* 34(3–4):243–254.
- Wang J (2005) A review of operations research applications in workforce planning and potential modeling of military training. Report DSTO-TR-1688, DSTO Systems Sciences Laboratory, Edinburgh, South Australia, Australia.
- Whaley DL (2001) Scheduling the recruiting and MOS training of enlisted Marines. Unpublished master's thesis, Naval Postgraduate School, Monterey, CA.
- Workman PE (2009) Optimizing security force. Unpublished master's thesis, Naval Postgraduate School, Monterey, CA.
- Zhang B, Murali P, Dessouky M, Belson D (2009) A mixed-integer programming approach for allocating operating room capacity. *J. Oper. Res. Soc.* 60(5):663–673.

Verification Letter

Geoffrey Guido, Manager, Training Programs and Technology Naval, Training and Simulation, Naval Nuclear Laboratory, Kesselring Site, Schenectady, NY 12301, writes:

"This is to verify that the claims made in the manuscript 'Optimal Allocation of Students to Naval Nuclear Power Training Units' are accurate. This manuscript describes the development and use of an application that optimizes the number of students assigned to Naval Nuclear Power Training Units. Since its adoption two years ago, we have realized a gain in student throughput of approximately eight percent and have significantly improved the deployment of our personnel and equipment."

Michael R. Miller is an engineer with the Naval Nuclear Laboratory. He spent 12 years in the Naval Nuclear Laboratory's Nuclear Operations Program training naval personnel on an operating reactor plant. He leads a software development group dedicated to providing analytical solutions for the Naval nuclear training community. He holds a Bachelor of Science in Mechanical Engineering from the University at Buffalo.

Robert J. Alexander works at the Naval Nuclear Laboratory developing analytic methods to support engineering and training efforts. Robert earned his Bachelor of Science degree in chemical engineering from Brigham Young University.

Vincent A. Arbige works at the Naval Nuclear Laboratory developing analytical data models to support allocation of training resources. Vincent earned his Bachelor of Science degree in chemistry from the University of Rhode Island and has worked in the Naval Nuclear Training Program for 36 years. He spent 10 years in various training positions and then built the first resource allocation models for the training program, which he has supported for the past 26 years.

Robert F. Dell is a professor of operations research (OR) at the Naval Postgraduate School (NPS). He joined NPS as

an OR assistant professor in 1990 and served as chairman of the OR Department from 2009 to 2015. During his tenure as chairman the department received the 2013 INFORMS Smith Prize. Professor Dell has been awarded the Barchi, Koopman, and Rist prizes for military operations research. He has also received two Department of the Army Payne Memorial Awards for Excellence in Analysis and two Department of the Navy Superior Civilian Service Awards. Professor Dell is editor-in-chief of the *Military Operations Research Journal*.

Steven R. Kremer works at the Naval Nuclear Laboratory developing analytic methods to support training. He is a retired Navy Captain who commanded USS ARCHERFISH (a nuclear powered submarine) and Naval Station Bremerton. He graduated from the United States Naval Academy with a Bachelor of Science degree in mechanical engineering and holds a master's degree in political science from Auburn University at Montgomery.

Brian P. McClune is a scientist at the Naval Nuclear Laboratory. Brian divides his time between development of optimization software for the nuclear training community and development of high performance computing applications in support of reactor design. He earned Bachelors of Science degrees in mathematics and physics and a Master's of Science degree in computer science from Clarkson University.

Jane E. Oppenlander teaches statistics in the School of Business and the Bioethics Program at Clarkson University. She earned her PhD in engineering and administrative systems from Union College. Jane recently retired after a 35-year career as a statistician at the Naval Nuclear Laboratory.

Joshua P. Tomlin works at the Naval Nuclear Laboratory programming systems to support engineering and training efforts. Along with his programming support, Joshua maintains and develops processes for the Learning Management System. He earned his master's degree in computer science from The College of Saint Rose.