informs
http://pubsonline.informs.org/journal/inte/

INTERFACES
Articles in Advance, pp. 1–18
ISSN 0092-2102 (print), ISSN 1526-551X (online)

Optimal Course Scheduling for United States Air Force Academy Cadets

Gerardo Gonzalez,[a] Christopher Richards,[b] Alexandra Newman[c]

[a] Department of Management, United States Air Force Academy, Colorado Springs, Colorado 80920; [b] Deputy Under Secretary of the Air Force, Management, 1790 Air Force Pentagon, Washington, DC 20330; [c] Department of Mechanical Engineering, Colorado School of Mines, Golden, Colorado 80401

Contact: gerry.gonzalez@usafa.edu (GG); christopher.richards@us.af.mil (CR); anewman@mines.edu, http://orcid.org/0000-0001-9886-9721 (AN)

Received: November 27, 2016
Revised: May 13, 2017; August 11, 2017
Accepted: August 21, 2017
Published Online in Articles in Advance:

https://doi.org/10.1287/inte.2017.0935

Abstract. Scheduling students and academic courses at the United States Air Force Academy (USAFA), a military commissioning source, has required unique software and considerable manual effort. The recent discontinuation of the Oracle-based student information system mandates that the USAFA superintendent invest in new software, the customization of which will incur millions in additional costs if USAFA continues to rely upon a fixed alternating-day schedule format. We present an integer program that generates a course schedule using the repeated-week format common to most commercial-off-the-shelf (COTS) systems. The integer program uses cadet registration information to determine the number of sections to be offered and how cadets should be assigned to them to ensure on-time graduation, while accomplishing mandatory military training. Hard constraints enforce institutional restrictions that require all athletes to attend practice, limit the number of cadets who delay required courses, keep classroom usage and number of sections to campus and faculty availability, and ensure cadets are assigned only to scheduled sections without overlapping time requirements. Flexible constraints reflect faculty and cadet preferences; their violation is minimized to honor teaching requests from each department, maintain minimum and maximum section sizes, restrict the number of evening sections, and meet cadet registrations. In contrast to the previous USAFA process, we generate schedules that reduce the number of unmet student registrations by more than 75 percent, use 21 percent fewer sections, and respect nearly 90 percent of faculty teaching preferences. Results from our methodology are easily reproducible and measurable in terms of time to adjudicate, desirability, and demand on faculty resources. By accommodating a standard repeated-week format, rather than adhering to the current alternating-day approach, our model integrates easily as a front end to a COTS system and avoids $120 million in customization costs. Our program reduces the reliance on manual manipulation and makes it possible to find feasible schedules that permit section length and patterns to vary according to pedagogy—a break from over 50 years of rigid time-blocking techniques that sacrifice desirability for feasibility and timeliness.

History: This paper was refereed.

Keywords: university timetabling • university scheduling • integer programming applications

Introduction

Established in 1954, the United States Air Force Academy (USAFA) is a public four-year university, comprised of 25 academic departments, with an approximate enrollment of 4,400 undergraduates pursuing over 30 possible degrees. As a military academy, however, USAFA differs from other four-year universities in four key aspects: (1) cadets have military jobs and duties that require a significant portion of their day to accomplish; however, some commitments may last only a portion of a semester, leaving a gap unsuitable for a standard academic class; (2) USAFA requires at least one airmanship course (e.g., flying gliders, parachuting out of aircraft, or flying motorized airplanes) that requires airfield access and a larger portion of the day than a standard academic course; (3) USAFA requires 100 percent cadet involvement in athletic training. Specifically, about 25 percent of the cadet body participates in an intercollegiate sport, while the remainder participates in intramural sports; and (4) with taxpayer-funded tuition and military quotas designed to be consistent with predetermined commissioning dates, cadets must graduate within four years.

For cadets to meet the competing requirements of academics, military training, and athletics, the registrar must necessarily consider cadet registrations each semester to simultaneously build course schedules and cadet assignments. When the registrar starts her process, she cannot gauge the number of cadet registrations her initial schedule will fail to fulfill, the work required to fix them, or if they can even be fixed at all. To plan for the worst case, the registrar must either collect information regarding cadet course preferences and department course offerings months prior to the semester being scheduled, or release cadet schedules for the next semester very late in the current semester. The former strategy precludes cadets and department staff from thoroughly considering the following semester's academic requirements, while the latter leaves insufficient time to resolve conflicts that may arise.

The majority of the information concerning cadet status and resource allocation is tracked by a collection of applications called the Cadet Administrative Management Information System (CAMIS), which is designed specifically to accommodate USAFA's unique scheduling requirements by implementing strict, uniform block-like timetables. This structure simplifies the manual manipulation in which departments and the registrar must engage each semester to resolve conflicts in a timely manner. Efforts to modernize CAMIS have consumed more than $30 million over the past decade, culminating in missed milestones and cost overruns (LaRivee 2017). As a result, Air Force headquarters has advised that all further attempts at modernization must rely heavily on commercial off-the-shelf (COTS) systems with minimal customization. Recently, the software that supports CAMIS has been discontinued, and commercial support has been withdrawn. As a result, USAFA must make the critical decision on how to replace its student information system. To adopt COTS software, USAFA must either have the system customized to accommodate the unique timetable structure it currently uses, or develop the ability to generate schedules that adapt to the standard repeated-week format common to COTS systems, while still meeting USAFA requirements.

Acting on the recommendations of a 2012 report (HQ USAFA 2012), Superintendent Lieutenant General Michelle Johnson mandated that a Pathways to Excellence Team develop a model capable of generating repeated-week schedules. To be implemented, such a model must maintain the emphasis on academic, military, and athletic requirements, but also introduce flexibility with additional section lengths and meeting patterns. Rather than remove departments and the registrar from the process, input from both should determine how to honor departmental preferences, while respecting institutional requirements. This approach allows departments to focus on arranging courses based on factors they deem pedagogically sound (e.g., whether a particular course is more suited as a once-a-week lecture or a thrice-a-week practicum). Because of the number and variety of stakeholders involved in such a model, USAFA leadership agrees with Schaerf (1999) that such a timetabling problem cannot and should not be completely automated. Any model, therefore, must remain part of an iterative process between departments and the registrar that can quickly identify, modify, and reevaluate proposed resolutions. Implementation must incur less cost than a COTS customization, currently projected at $220 million, based on the initial estimate completed by the government (LaRivee 2017).

To this end, we introduce a mixed-integer programming (MIP) model ($\mathcal{S}$) that enforces *hard* constraints (i.e., limits on classroom usage, the number of sections that can be taught, and the number of required courses that can be delayed, requirements that athletes attend practice, and logical constraints to ensure that cadets are assigned to actively scheduled sections without overlapping time requirements) and minimizes the violation of *flexible* constraints (i.e., teaching requests, minimum and maximum section sizes, restrictions on the number of evening sections, and meeting cadet registrations). Together, these constraints consider USAFA's policies, resources, and academic preferences. Penalties elicited from leadership reflect USAFA priorities regarding deviation from the desires of individual departments and result in a scheduling and assignment process that respects institutional rules, reduces the time spent resolving registration conflicts, and sufficiently honors concerns raised by departments. The goal is for ($\mathcal{S}$) to serve as the front end to a COTS scheduling system, building timetables and section rosters that comply with both USAFA requirements and the standard weekly framework common to such software. As a tool that faculty and administrators will use regularly, the input ($\mathcal{S}$) requires and the output it provides must be intuitive and straightforward. Additionally, rather than being black-box software, the formulation of ($\mathcal{S}$) should retain sufficient flexibility that future shifts in USAFA preferences or requirements can be easily accounted for by someone with a moderate understanding of integer programming.

The remainder of this paper is organized as follows: First, we provide a literature review on academic timetabling; we then summarize the criteria that compose the timetabling challenge at USAFA and define several necessary terms and parameters. Subsequently, we describe how the previous USAFA scheduling process attempts to incorporate military and athletic requirements before we detail the mixed-integer program that improves on that process. We next discuss the implementation of our work, and present results. We conclude with impacts and extensions.

Literature Review

Wren (1996) describes timetabling as "... the allocation, subject to constraints, of given resources to objects being placed in space-time, in such a way as to satisfy nearly as possible a set of desirable objectives" (p. 53). Each set of potential objects to be placed—nursing shifts (Jaumard et al. 1998), train crews (Caprara et al. 1999), athletic competitions (Easton et al. 2001)—come with equally unique requirements and constraints. Although the applications of timetabling remain broad, the growth in both number and size of secondary and postsecondary educational institutions (National Center for Education Statistics 2015) has led to an increased focus on academic applications. University schedules pose a particularly difficult challenge; with so many stakeholders (e.g., students, professors, deans), a model that generates a desirable schedule is often far more difficult to design than one that simply finds feasible solutions (Easton et al. 2001). Although new approaches (Babaei et al. 2015) that attempt to address these issues have been introduced in recent years, more than 40 percent of universities create large-scale schedules via manual processes, which can require as many as 17 weeks to complete (Burke et al. 2005). According to Schaerf (1999), many researchers agree that university timetabling cannot be completely automated because of the number of subjective factors involved. The complete process of scheduling students to courses, teachers to courses, courses to periods, and courses to classrooms has been termed the *population and course timetabling problem* (PCTP), and Boland et al. (2008) point out that few papers attempt to tackle it entirely. Instead, they tend to focus on one or two aspects. Feng et al. (2016), for example, develop both an integer program and a genetic algorithm, which schedule courses to specific times and rooms but consider only predicted class sizes without actually assigning students. Conversely, Cheng et al. (2003) assume a predetermined set of course offerings and attempt to create conflict-free student assignments that meet their registration requests. Researchers such as Aubin and Ferland (1989) address both course scheduling and student assignments in their models, but do so sequentially, holding either course schedules or student assignments constant at each iteration, and making small changes at each step until no further improvements can be made. Hertz (1991) develops a tabu search that improves upon this iterative method, but the size of the solution space forces both Hertz and Aubin and Ferland to focus on feasibility over desirability, minimizing the number of overlapping assignments rather than maximizing instructor and (or) student preferences. Müller and Murray (2010) address the issue of preference ex posteriori. After scheduling courses and assigning students using a similar local-search method, they share their initial solution online and allow students to individually view and edit their registrations within limits. One approach to more directly satisfy student and faculty preferences is to assign courses and students concurrently rather than sequentially. Bakir and Aksop (2008) and Boland et al. (2008) take such an approach, but they must reduce the number of decision variables in their models by partitioning students and courses and then scheduling these common "blocks," rather than individual students or sections. In addition to problem size, such concurrent consideration impacts the complexity of the problem. Dostert et al. (2016) prove that although students can be assigned to a specific number of sections within a given timetable in polynomial time, respecting student registrations immediately transforms the problem into an NP-complete one.

Because model complexity and associated instance size are both concerns, it is not surprising that a large emphasis has been placed on approximation techniques and heuristics capable of generating practical university schedules in a reasonable amount of time. Lewis (2008) provides both a summary and a taxonomy of many such heuristics that have been presented and (or) implemented in recent years. In addition to the integer programming and tabu search methodologies, a variety of computational techniques, including simulated annealing (Tuga et al. 2007, Aycan and Ayav 2009), genetic algorithms (Khonggamnerd and Innet 2009, Alsmadi et al. 2011), and hybrid colony optimization (Fong et al. 2014), have been applied to solving different aspects and instances of university timetabling. For the most part, however, although these approaches may differ in how they *search* for a solution, they are very similar in how they *rate* a solution. Babaei et al. (2015) explain that feasibility of a solution is normally dictated by hard constraints that cannot be violated, regardless of preference (e.g., the number of available seats in a classroom, that a student cannot be in two places at once), while desirability of a solution is measured by the degree to which soft constraints (e.g., no evening classes, even distribution of students among sections) are violated. Even with this structure, however, many universities remain unable to generate schedules without some form of manual manipulation. This is partially because, although sections and students can be scheduled at different times, the number of each is often fixed; that is, the only way to improve a given solution is to permute section start times and attempt to reassign students. In their work concerning conference seminar scheduling, Eglese and Rand (1987) introduce the flexibility to choose how many of a particular session to offer to better meet attendee interest. Although their approach has an obvious application to tailoring a course timetable to student registration, its practical success is heavily dependent upon the fact that the number of seminars and attendees to be scheduled is small, approximately 20 and 300, respectively, relative to the course offerings and class size of an average university.

Perhaps closest to our own challenge, Sampson and Weiss (1995) build on Eglase and Rand's work by developing an integer program that simultaneously considers section and student preferences while adding capacity limits. A single integer program that handles course and student scheduling would easily integrate into a COTS system, and, once developed, its parameters and penalties can be adjusted over time to reflect evolving USAFA priorities without major changes to the program's underlying structure. A practical implementation of their approach (Sampson et al. 1995), however, leaves more than 5 percent of student registrations unmet, requires the dean to predetermine the total number of sections, and applies a heuristic, because instances of more than 1,000 students and 40 sections make solving an integer program intractable. Although the computational power available to solve integer programs has grown substantially in the past two decades, researchers such as Lübbecke (2015) still hold that accommodating student sectioning while scheduling courses makes large-scale instances too complex to solve. Instead, many recent integer programming approaches to real-world university timetabling, such as Phillips et al. (2015), create models that focus solely on scheduling classes, and assume the student assignment is completed either a priori or ex posteriori.

Undergraduates at USAFA, however, represent not only students pursuing a degree, but cadets being trained for a commission; any feasible course schedule must necessarily consider both of these missions. For USAFA to realize the benefits of an integer program, we introduce several methods that increase the tractability of our own instance, which includes 4,000 cadets and 400 courses across 1,000 sections. Furthermore, in addition to capacity constraints, graduation deadlines require that our model limit allowable registration conflicts based on future class availability. Finally, many military duties do not begin and end according to academic semesters and require that courses and duties be deconflicted according to the fraction of the semester they occupy. We provide a detailed description of specific constraints related to our integer program in Appendix A.

Definitions and Problem Statement

Before describing the previous USAFA model and introducing ($\mathcal{S}$), we define the necessary technical terms used throughout this paper and the requirements for a feasible and desirable USAFA schedule.

Block: A specific fraction of a semester required by a course. For example, we divide our semester into three blocks; thus, Block 2 would refer to the time during the second one-third of the semester.

Period: A predetermined continuous span of time on a specific weekday (e.g., 0900–1030 on Monday).

Epoch: A continuous span of time during which a unique set of periods overlaps. For example, given periods on Monday from 0800–0900, 0800–0915, and 0800–1000, three epochs would exist: 0800–0900, 0900–0915, and 0915–1000.

Course: A specific academic, athletic, or military requirement for which cadets can register (e.g., calculus 101, intramural football, airmanship).

Pattern: A set of blocks, with an associated set of nonoverlapping periods (e.g., {{1, 2}, {(Tuesday 0800–1000), (Thursday 0900–1100)}}).

Section: A (*course*, *pattern*) instance to which cadets can be assigned. Courses with large numbers of registrations may require multiple sections.

Registration: A (*cadet*, *course*) requirement pair.

Cadet assignment: A designated (*cadet*, *section*) pair.

Course offering: A set of patterns for all courses that possess registrations.

Room type: A category assigned to potential locations at which a section can meet, defined based on seating capacity and amenities (e.g., academic versus athletic).

Registration conflict: A registration that cannot be matched to a cadet assignment within a feasible course offering.

An acceptable USAFA scheduling process must include a set of periods that departments can use to build a course offering. This preferred course offering (PCO), along with section enrollment minima and maxima, accommodates the total number of registrations for all academic courses and represents the preferred teaching schedule of a department's instructors in aggregate. (Departments make specific section-instructor assignments once an offering is finalized.) The PCO is combined with all registrations, and the room type and blocks required by each course. Using this input, a model must develop a revised course offering (RCO) and a set of cadet assignments. Patterns and sections within the RCO should match those of the PCO whenever possible and must adhere to limits concerning room-type availability within a given epoch and block. Cadet assignments should obey the preferred section minima and maxima, minimize the number of registration conflicts, and must not require a cadet to delay a course that cannot be delayed or to attend periods that overlap during a particular block.

Previous USAFA Scheduling Model

Cadets must carry an average of six courses per semester to accumulate 141 semester hours and meet their four-year commission deadline. Cadets construct a four-year plan and, each semester, meet with their advisers to review and (or) adjust registration for the following semester. The outcome of these meetings provides course enrollment numbers for the upcoming semester that, in turn, departments use to generate their PCOs.

Although the registrar has the ultimate responsibility for creating the upcoming semester's course offering and cadet assignments, individual departments are more keenly aware of their own staffing and course sequence requirements. To this end, creating an effective schedule has always been a lengthy and iterative dialog between departments and the registrar. Each department has a dedicated representative assigned to this process, and each is given an eight-page guideline for building its initial PCO. These rules are the outcome of years of information and best practices to create an initial course offering that is capable of meeting as many cadet registrations as possible. For example, courses with only a single section must use morning periods to minimize conflicts with afternoon athletics. Courses with more than 10 sections must utilize all periods evenly. The desired result of this process is a PCO that allows most registrations to be matched to feasible cadet assignments. After all possible assignments have been made, remaining registration conflicts are addressed individually.

To accommodate this process, the *previous USAFA model* divides the semester uniformly; each school day has seven 53-minute periods, and each course meets exactly 40 times. Rather than assigning courses to specific patterns, all courses simply meet on alternating weekdays. So, in a typical, two-week period, a cadet attends a course five days: Monday, Wednesday, and Friday, followed by Tuesday and Thursday. However, interruptions such as holidays may result in situations in which a course meets as few as three times in a two-week period.

This uniform scheduling structure exists largely to simplify the registrar's task of generating an RCO and resolving the subsequent registration conflicts. By not using patterns that depend on specific days and requiring that all periods be of equal length, the registrar can more easily move sections and modify cadet assignments during her process without excessive cascading effects. This rigid uniformity, however, significantly reduces the freedom departments possess to teach courses according to pedagogy and results in courses that meet either more or less frequently than is necessary. At present, a standard USAFA semester requires 18.5 weeks to meet the necessary contact time with cadets. The school day is also limited, starting at 0730 and ending at 1430, because USAFA deems it necessary to reserve a large portion of the afternoon to accommodate any possible military and (or) athletic training a cadet might require, rather than including it as part of the scheduling process.

Under the old regime, the registrar repeats the following steps until a satisfactory schedule is built:

1. Identify cadets and their registrations.
2. Solicit an initial PCO from departments.
3. Identify registrations that cannot be matched to assignments under the PCO.
4. Assign all nonathlete cadets to intramural sports and assign all athletes to their intercollegiate sport's practice requirements.
5. Find cadet assignments for as many registrations as possible using a simple greedy heuristic.
6. Use an exhaustive search to try to find cadet assignments for any remaining registrations.
7. Identify cadets with assignments that can be exchanged for another feasible assignment and determine if any such alteration permits the resolution of a lingering registration conflict.
8. Manipulate cadet assignments in an attempt to fill sections currently below the minimum preferred section size.
9. Generate a report that includes the algorithm's RCO and its remaining conflicts and deficiencies.

The above algorithm has three primary shortcomings: (1) subjective human decisions result in solutions that are not reproducible; with the exception of steps 5 and 6, all actions reflect a manual effort; (2) it is focused almost entirely on reducing registration conflicts and respecting section enrollment limits and largely ignores any measure of preference; and (3) the greedy heuristic used by the algorithm is shortsighted; in addition to being completely undocumented and severely outdated, it considers each registration, in no particular order, and matches it to the first available assignment, also evaluated in no particular order and with no consideration of cascading effects. This approach does not guarantee an optimal solution; in addition, the registrar's guidelines and conflict resolution techniques leave approximately 200 cadets with unassigned registrations for at least some of their required courses.

These shortcomings force the registrar and departments to invest weeks of effort, after the RCO has been generated, to resolve the remaining registration conflicts either by overfilling a current section, obtaining approval for additional sections to be taught, delaying the assignment of registrations to a later semester, and (or) requesting a complete or partial PCO resubmission from departments and restarting the process. This resolution must occur each semester to ensure that every cadet is able to receive his (her) commission in four years.

Proposed Mixed-Integer Programming Approach

To generate an RCO and cadet assignments that result in fewer conflicts and a shorter resolution process, $(\mathcal{S})$ uses three primary decision variables: Z_{cp}, Y_{ucp}, and D_{uc}. Like the RCO generated by the registrar's algorithm, the solutions produced by $(\mathcal{S})$ do not match sections to specific rooms or instructors. Instead, Z_{cp} determines how many of section (c, p), given as a course pattern pair, to schedule; Y_{ucp} assumes a value of one if registration (u, c) of a student to a course is assigned to such a section; and the binary D_{uc} reflects whether a registration (u, c) is delayed to a later semester. To keep USAFA decision makers apprised and to simplify the problem (without compromising optimality), departments consider specific instructor allocation during the creation of the PCO and make explicit assignments ex posteriori.

Our formulation produces a solution provable to within a five percent optimality gap and provides three unique advantages over similar mixed-integer programming approaches to academic timetabling (see MirHassani and Habibi 2013 for examples): (1) it generates a course offering and cadet assignments *concurrently*; (2) the solutions accommodate unique Academy challenges, such as the period and block requirements of military training, the specialized needs of airmanship and athletics, and the necessity of on-time commissioning; and (3) the problem remains tractable despite the magnitude of our effort, which includes thousands of cadets and hundreds of courses in contrast to instances found in the literature, which involve hundreds of students and only dozens of courses.

Model Description

Before departmental preferences can be considered, ($\mathcal{S}$) must include several *hard* constraints that cannot be violated. All intercollegiate athletic practice registrations must be assigned. The total number of sections in a single epoch can never exceed the total number of rooms of a particular type. In addition to department preferences, faculty size and capabilities enforce a hard limit on how many times a specific section can be repeated in a course offering. Logical constraints require that each cadet assignment be matched to a section that exists within the associated offering, and that no cadet be required to attend two sections with overlapping patterns. Finally, to achieve commissioning deadlines, each cadet may delay at most one registration per semester, and the number of delayed registrations for each course must be limited based on the future availability of that course.

With restrictions and mandates met, the remaining *flexible* constraints in ($\mathcal{S}$) promote adherence to the preferences submitted by departments. This includes using sections from the PCO wherever possible, minimizing the number of cadets assigned to evening periods, and seeking to assign cadets within section-specific minimum and maximum capacities. The penalties associated with violating these preferences are elicited from USAFA leadership and use registrations in conjunction with the faculty PCO to measure the impact of each infraction. For example, using only the sections requested by the PCO for a particular course is preferable; however, if additional sections must be used, USAFA leadership mandates that (1) the penalty be inversely proportional to the number of sections the PCO requests for the course, and (2) the penalty for using a section outside of the PCO be double that of repeating a requested section. Although leadership initially mandates the cost of violating each flexible constraint, we conduct a sensitivity analysis (Appendix C) to determine the impact of these values and recommend changes, if appropriate.

Finally, registrations that can be neither assigned nor delayed under any feasible course offering result in a registration conflict. Although a solution without registration conflicts is preferable, in practice this is not possible to achieve; thus, the constraint must remain flexible to preclude the model from being infeasible.

Registration conflicts can only be resolved through direct manual intervention, either by an adviser working with a cadet to modify his (her) four-year graduation plan, or by USAFA leadership approving an exemption or authorizing the additional manpower required for supplemental or larger sections to be taught. To minimize the number of conflicts and the associated resources necessary to do so, each violation carries a penalty that is high relative to those associated with other flexible constraints (i.e., six times, on average) and proportional to the credit hours associated with the course responsible for the conflict.

Tractability

The structure associated with our model and the number of variables and constraints in realistic instances puts significant strain on current hardware, software, and MIP solution techniques. The computational complexity of a constrained decision problem such as ($\mathcal{S}$) falls in the category of nondeterministic polynomial-time complete (NP-complete) problems, signifying that the problem is among the hardest to solve because there is no known polynomial-time solution algorithm. Moreover, the average USAFA enrollment includes approximately 4,000 cadets, each of whom possesses up to eight registrations per semester. For instances that contain 400 courses and 100 different periods, ($\mathcal{S}$) can consist of billions of variables, and nearly twice as many constraints, and can require days to find a feasible solution that adheres to the proposed PCO. In practice, we mitigate this via several techniques designed to increase problem tractability: (1) elminate unnecessary variables and constraints through set indexing and by combining certain penalties; (2) add

Figure 1. (Color online) In This Matrix of Available Weekly Periods, We Show That in Comparison to the Previous USAFA Model, Departments Are Allowed Eight Additional Times at Which Periods Can Start and Six Different Period Lengths, Which They May Combine Into Any Desired Weekly Pattern

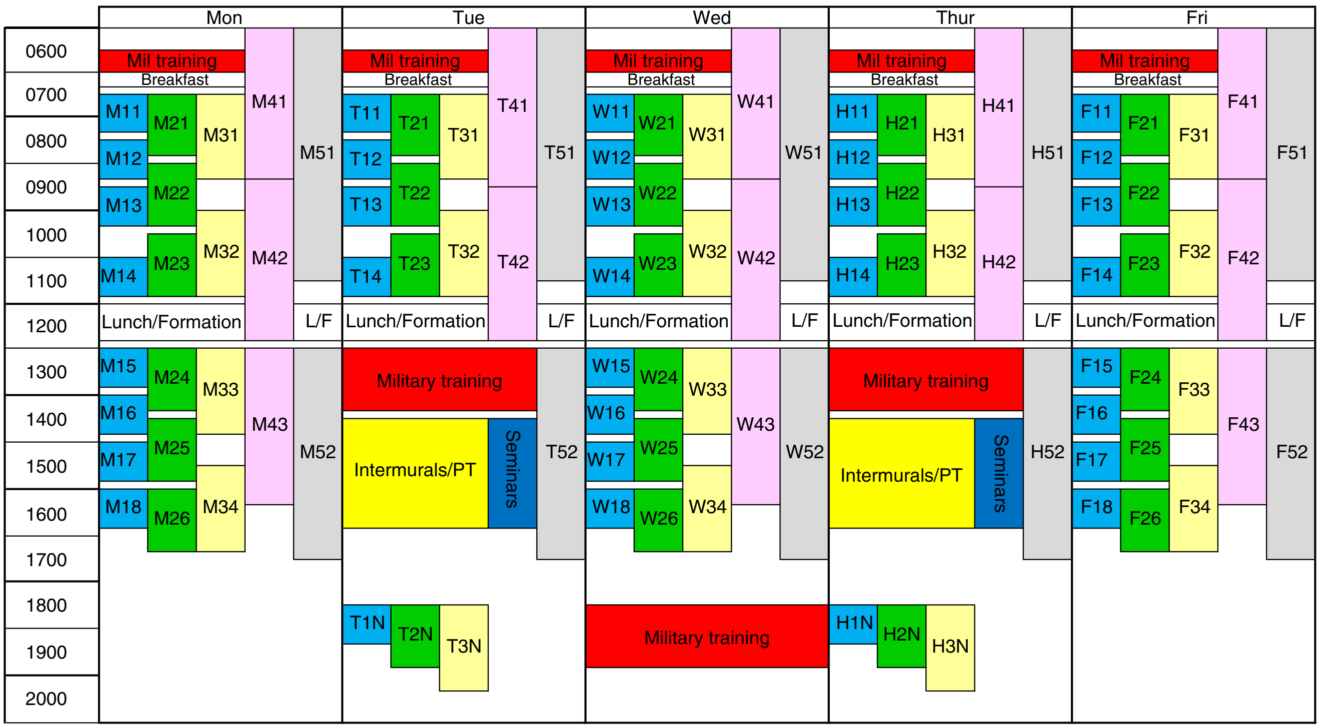

strengthening constraints; (3) tune the solver software; and (4) introduce a procedure to obtain an initial feasible solution. Appendix B provides specific details for the application of each technique.

New Scheduling Process

To facilitate the transition from the previous USAFA model to implementation of ($\mathcal{S}$), we present a simple four-step iterative procedure that closely mirrors the previous process but gives departments more options without involving them in behind-the-scenes optimization mechanics.

Request for PCO

During a semester, after cadets have met with their advisers and registered for courses, an initial PCO for each department can be formed. To this end, we present each department with the new available periods (Figure 1).

Because ($\mathcal{S}$) ensures military training and athletics requirements are met, each department is encouraged to build its PCO in a manner it deems pedagogically appropriate. Rather than attempting to adhere to the aforementioned complicated guidelines, ($\mathcal{S}$) determines the feasibility of meeting the department's preferences, and provides alternatives when necessary.

PCO Submission

A PCO is submitted from each department in the form of a spreadsheet that contains four pieces of information for each unique section a department plans to offer: (1) the associated course; (2) a set of periods from Figure 1; (3) the number of times a department would like to offer the section; and (4) a minimum and maximum number of cadets preferred in each section. Table 1 provides an example.

Table 1. PCO Input: This Sample Shows Departmental Preferences for Course Periods, Patterns (Which Can Be More Than Three Periods), and Sizes

Course number	First period	Second period	Third period	No. of sections	Minimum enrollment	Maximum enrollment
c1523	M12	W12	F12	3	10	20
c2317	F41	—	—	1	0	50
c8921	T26	H26	—	5	15	25
...	...	...	...	...	...	...

Registrations and additional course information from the registrar are added to the PCO data. This information includes which courses have special location requirements (e.g., airfield or physical education), which can be delayed, and which require the entire semester versus only certain blocks.

Data Additions and Processing

To facilitate the implementation of ($\mathcal{S}$), preprocessing code forms the PCO and registration data into sets, which it then expands upon. This includes simple rearrangements of data for constraint qualifications (e.g., constructing $\mathcal{U}_c$, undergrads registered for course c, from the registrar's $\mathcal{C}_u$, courses for which undergrad u has registered) or more convenient groupings (e.g., collecting all patterns for a particular course c into a single set $\mathcal{P}_c$). More importantly, in addition to rearranging the given data, code facilitates the following:

First, using the preferences from the PCO as a template, additional sections are included as alternatives to those requested by the PCO. For example, in Table 1, although course c2317 is preferred to be offered only as a pattern involving the single period on Friday (F41) by the *department*, this does not consider possible *registrar* complications (e.g., a cadet with a registration for another class whose only sections involve overlapping patterns). With this in mind, the code relaxes ($\mathcal{S}$) to consider additional sections that mimic the patterns requested in the PCO, each of which is assigned capacity parameters (i.e., enrollment minima and maxima) identical to those in the PCO, with the exception that their preferred number of sections is set to zero. In our example, 12 additional single-period sections of course c2317 are made available using the periods from Figure 1 that match the length of F41. The model ($\mathcal{S}$) now has more feasible sections in which to schedule c2317 but is penalized for choosing any pattern outside the department's PCO. Second, the code provides the option of generating an initial solution, which assigns seniors, juniors, and sophomores successively, fixing the assignments of each class before proceeding. By partitioning the problem, the number of variables is controlled at each step to provide a quick, feasible solution that prioritizes upperclass requirements and that ($\mathcal{S}$) can use as a starting point.

Output and Feedback

Once ($\mathcal{S}$) solves, two files are generated and returned to departments: (1) summary statistics of the solution, which highlight both the deviation from the original PCO and the number of registration conflicts remaining (Table 2); and (2) interactive course and cadet schedules under the current solution (Figures 2 and 3).

Figure 2 provides a measure of how closely the proposed solution matches the PCO, allowing the registrar and departments to gauge the extent to which

Table 2. PCO Deviation: Numbers Represent the Difference Between the PCO Request and the Solution Proposed by ($\mathcal{S}$), in Which Negative Numbers Signify PCO Requests That Were Unused, Positive Values Indicate Sections Added by ($\mathcal{S}$), and "—" Corresponds to Perfect Agreement Between the PCO and the Proposed Schedule

Summary statistics:
Registration conflicts: 24
Total PCO sections requested: 1,451
Unmet PCO requests: 325
Non-PCO sections used: 404

		Total deviation from PCO by period and block						
Blocks	Course	M11	M12	M13	M14	M15	M16	M17
1–3	c201	—	(−1)	1	3	(−2)	—	—
1–2	c218	—	—	—	—	—	(−2)	4
1–3	c223	—	1	—	1	—	—	—
1–3	c233	1	1	—	—	—	(−1)	—
1–3	c234	1	1	(−1)	1	—	—	—
2–3	c323	—	—	—	—	1	—	—
1–3	c335	(−2)	—	—	1	1	1	1

Figure 2. (Color online) In This Interactive Course Schedule, Enrollment Numbers Are Given by Period for an Arbitrary Course (Chemistry 110)

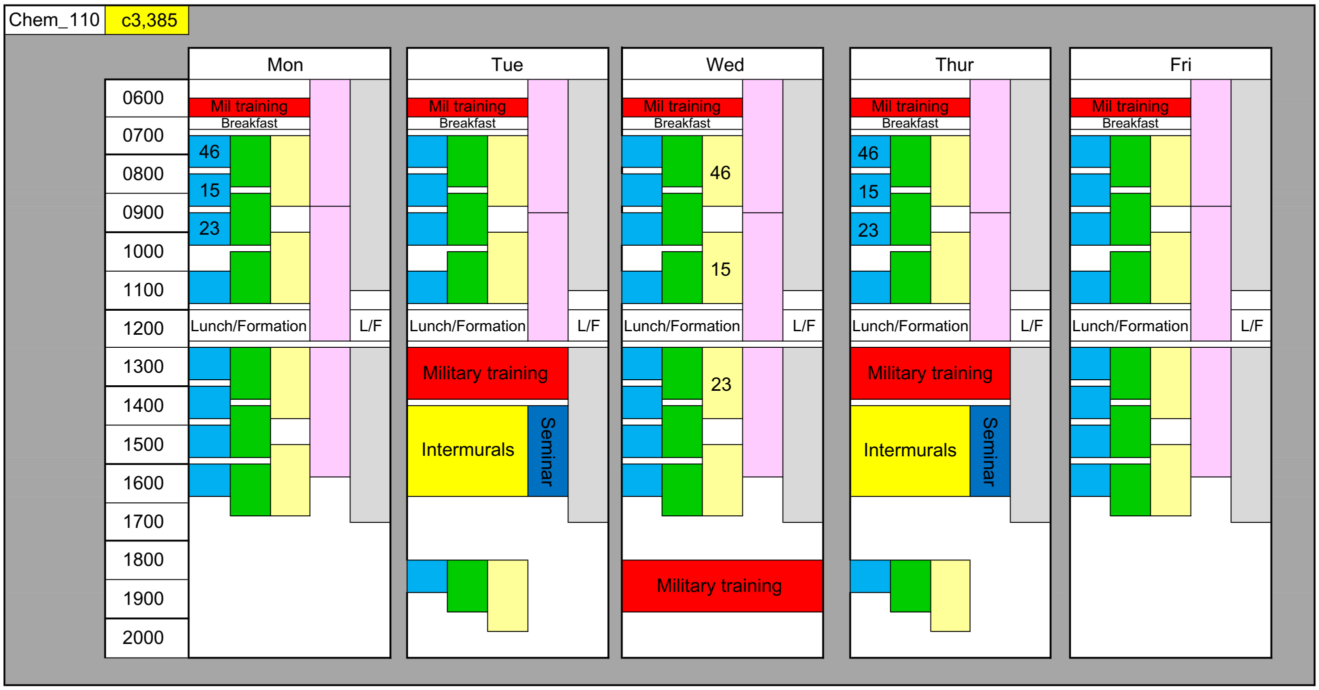

Figure 3. (Color online) This Interactive Cadet Schedule Gives a Weekly Schedule for an Arbitrary Cadet, Including a List of Registration Conflicts

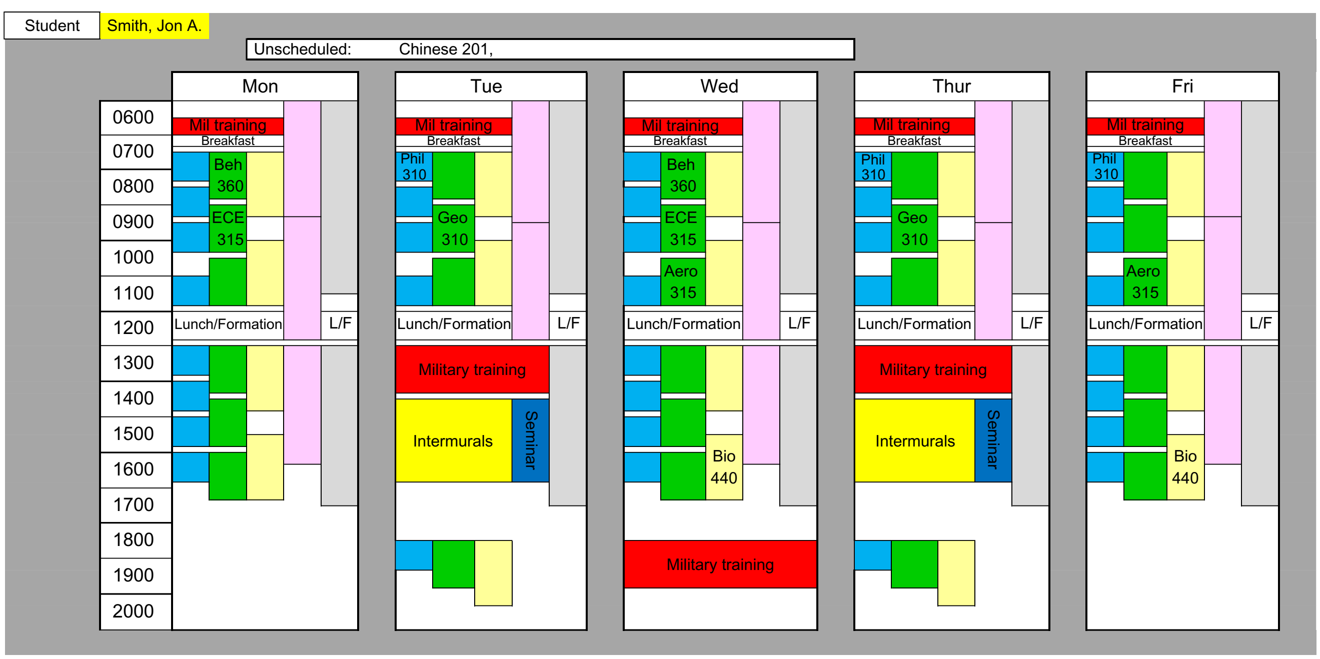

the initial PCO meets cadet requirements. The course schedule in Figure 2 serves as a tool with which departments can evaluate the proposed solution and decide whether the changes represented are feasible based on staffing and course material.

The cadet schedule in Figure 3 is intended to aid the registrar in focusing on those cadets with registration conflicts and evaluating the work required to resolve them. Once these evaluations have been made, the proposed solution can either be accepted as is or with one-off modifications. Individual departments can resubmit a modified PCO based on the registration conflicts and unmet requests associated with their original preference. The process is repeated until it produces a solution that is acceptable to both the registrar and the departments.

Results

The superintendent has contracted with a third-party vendor to work with the registrar and Information Technology Department at USAFA to develop the front end, which integrates ($\mathcal{S}$) into COTS scheduling software and facilitates information sharing and processing between departments and the registrar. To reflect the hardware available to the registrar, we use 32 Intel(R) Xeon(R) 2.60 GHz processors with 32 GB RAM running CPLEX 12.6.2 to obtain our results. A decision regarding the implementation of ($\mathcal{S}$) was based on USAFA registration data for the 2016 fall semester. A request for a PCO with the newly structured periods (Figure 1), along with appropriate instructions, was sent to all departments. The returned PCO includes 1,401 desired sections, 1,069 of which are unique. The registrar provided 30,986 registrations from 3,894 cadets for 459 courses (approximately 8 registrations per cadet).

Although the process of resolving all conflicts and publishing a finalized schedule requires multiple iterations between departments and the registrar, the time required to complete that process is largely defined by the quality of the initial attempt to merge PCO and registration requirements into an RCO and a set of cadet assignments. The quality of this solution is measured using four separate indicators:

Registration conflicts: The number of unmet student-course pairs. Fewer conflicts indicate more registrations have been met; thus, the solution requires that the registrar spend less time arbitrating resolutions with USAFA leadership and departments.

Unmet preferences: The number of sections from the PCO that were not used by the RCO; a smaller value indicates closer adherence to departmental preferences and a solution that the department will more readily accept.

Total RCO sections: The total number of sections required by the RCO; smaller values represent less strain on faculty resources and a solution USAFA leadership will more readily approve. Further, this measures how many alternative, nonpreferred sections were needed in lieu of any unmet preferences.

Time: Hours required to generate the RCO; shorter run times provide quicker feedback to departments concerning the feasibility of their PCOs and an estimate of the time required to evaluate subsequent PCO revisions.

Table 3 compares the initial RCO generated by four methods: (1) the *previous USAFA model*, an almost completely manual effort that uses the alternating-day format with fixed 53-minute periods; (2) a *basic MIP* based on an initial proposed formulation (Gonzalez 2011) and developed by contractors, it contains only fundamental USAFA constraints that focus on feasibility over desirability and lacks our tractability enhancements; (3) our proposed ($\mathcal{S}$) *formulation*, described in the previous section, which integrates faculty preferences along with tractability modifications to generate schedules that are viable and desirable in a reasonable amount of time; and (4) an implementation of ($\mathcal{S}$) using the initial feasible solution obtained from the algorithm we present in Appendix B. We solve all models to a five percent optimality gap. USAFA chose this gap because, in practice, we found that using a smaller gap greatly increases solution times without significantly improving solution quality.

The *previous USAFA model* generates a schedule that uses only PCO sections because it is not equipped to investigate alternatives. This comes at the price of 200 registration conflicts, which require a sizable effort on the part of the registrar to resolve; each adjustment must be made manually and can cause cascading conflicts. After generating the RCO reflected in Table 3, the *previous USAFA model* requires that the registrar work with departments to make an average of 2,715 adjustments per day, affecting more than 1,000 cadets.

Table 3. Results: Statistics for the Fall 2016 Semester (Solved to Within Five Percent of Optimality, Where Applicable); Solution Times Represent a Single Iteration

	Instance characteristics		Solution quality			
	Number of variables	Number of constraints	Registration conflicts	Unmet preferences	Total RCO sections	Time (hours)
Previous USAFA model†	NA	NA	200	0	1,682	0.5
Basic MIP††	349,932	219,353	(41)	(563)	(1,280)	(106)
($\mathcal{S}$) *formulation*	221,180	99,573	44	147	1,359	1.6
($\mathcal{S}$) + initial solution†††	222,609	99,224	43	74	1,430	1.3

†The previous model is primarily a manual effort; the time listed accounts only for the greedy algorithm and omits time spent during manual manipulation.
††The basic MIP is unable to reach a solution using the initial hardware; the parenthetical numbers reflect results obtained by increasing system memory.
†††Variable and constraint counts pertain to ($\mathcal{S}$) using the improved solution provided; time includes that spent finding an initial solution (Appendix B).

The *basic MIP* in Table 3 closes the optimality gap to only 23 percent before exhausting system memory. Run on a machine with increased memory, it reaches the desired optimality tolerance in just over 106 hours. The solution greatly reduces registration conflicts to 41 but ignores over 40 percent of the section preferences from the PCO. This solution requires fewer fixes to cadet assignments at the expense of teaching preferences, indicating that departments would likely resubmit a revised PCO and wait four additional days for it to be evaluated. Conversely, ($\mathcal{S}$) achieves a similar registration conflict count of 45, and does so in 1.6 hours; additionally, it respects 90 percent of department preferences. Although ($\mathcal{S}$) uses 79 more sections than the *basic MIP*, the total number is far below that used by the *previous USAFA model* and remains less than the 1,401 requested in the PCO. Figure 4 shows that reducing the total number of sections is possible, but only at the expense of additional unmet preferences. Prior to implementing ($\mathcal{S}$) into the new COTS software, the registrar set a benchmark of 60 registration conflicts, satisfied that this goal would significantly reduce the required person-hours to finalize a schedule. Solutions found by ($\mathcal{S}$) surpass this goal by 25 percent, and their close adherence to preferences within the PCO has led to an overwhelming recommendation from the USAFA faculty senate to adopt the new schedule of calls.

Using the initial solution provided by the algorithm in Appendix B, ($\mathcal{S}$) generates a solution in 19 percent less time (including time spent on the algorithm); the schedule incurs approximately half of the unmet preferences as ($\mathcal{S}$) does without the initial solution, but it uses 71 additional sections. The difference in quality metrics is consistent with the trade-offs illustrated in Figure 4 and can be controlled in subsequent iterations if the initial RCO and cadet assignments are not accepted.

Impacts and Future Work

Scheduling courses and cadets via ($\mathcal{S}$) represents key advantages over the *previous USAFA model* in flexibility

Figure 4. Number of Unmet PCO Preferences as a Function of Restricting the Total Number of Sections Offered Is Shown

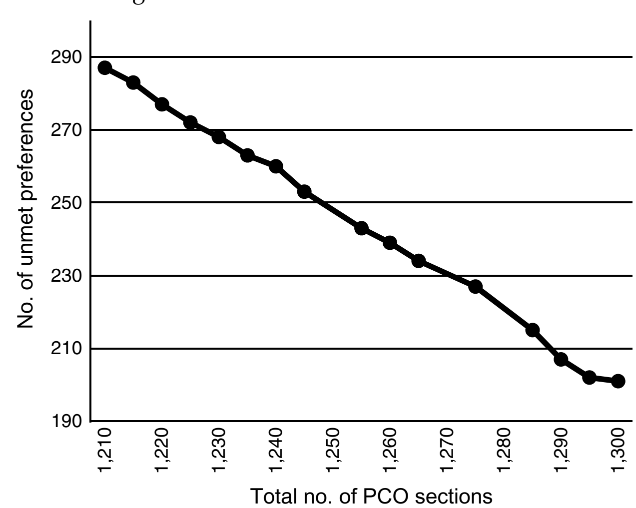

Note. Reducing the total number of sections required by the RCO is possible, but only at the expense of additional PCO deviation.

of use, quality of solution, and execution time. With the diversity of options provided in Figure 1, many instructors report that using a twice-a-week 80-minute pattern (not available under the *previous USAFA model*) greatly improves their effective contact time, while those who wish to retain their current lesson plans continue to do so using thrice-a-week 50-minute patterns. Replacing the previous alternating-day approach with a repeated-week pattern allows ($\mathcal{S}$) to accommodate both pattern preferences and generate a schedule in which the required contact hours for a semester are met in 15 weeks, 3.5 fewer than the *previous USAFA model*. Furthermore, with a total proposed cost of \$100 million, incorporating ($\mathcal{S}$) into a COTS system represents \$120 million in savings in comparison to maintaining the previous alternating-day model (LaRivee 2017). Satisfied that the analysis and successful application of ($\mathcal{S}$) meets or exceeds the conditions imposed on the Pathways to Excellence Team, the superintendent has approved the use of ($\mathcal{S}$) and has made the decision to obtain a request for proposal from industry to implement COTS software and replace the current Oracle system. This decision fundamentally changes the way in which the Academy has approached its schedule of calls for more than 50 years and has already provided benefits. During an early implementation, USAFA was able to combine two separate lessons of an economics course into a single night class. The new structure eliminates time spent resummarizing material and allows professors to get "... more than twice as much done in less than two full periods of time" (Branum 2014). By shortening the overall semester, ($\mathcal{S}$) is also opening new opportunities for military field training.

During the writing of this paper, the Pathways to Excellence Team modified the matrix in Figure 1 several times in response to requests from USAFA leadership. In each case, the core design of ($\mathcal{S}$) required no changes. Instead, only minor set and parameter updates were necessary to implement the revisions. The flexibility in the application of ($\mathcal{S}$) not only allows it to evolve with USAFA's needs, it introduces additional applications within the university. Work has begun to use ($\mathcal{S}$) to plan cadet flight schedules around their academic ones. Building these two types of schedules in concert maximizes flight time during good weather and reduces potential safety issues. ($\mathcal{S}$) has also proven useful as a what-if decision tool outside of the scheduling process to determine the cost-benefit analysis of hiring additional professors or permitting more sections during evening or weekend hours in future semesters.

One current limitation of our model is that it generates course schedules and cadet assignments for the current semester only. USAFA is interested in extending this time horizon to encompass an entire school year in future iterations. If cadet requirements and faculty preferences can be simultaneously identified for the spring, summer, and fall semesters, the additional clairvoyance could allow courses to be scheduled over multiple semesters, which would lead to fewer sections in a given semester and (or) a further reduction in the total number of registration conflicts; however, the model would have to be updated both to capture the unique characteristics of a summer semester (i.e., shortened length and reduced availability of faculty and cadets) and to weigh the benefits of assigning cadets to future semesters against the uncertainty of resource availability.

In addition to expanding the applicability of ($\mathcal{S}$), heuristics can reduce its run time. Tests show that our own process for finding an initial solution may help achieve this goal by using a different prioritization scheme. Although class year is an intuitive option, other partitioning choices, such as participation in athletics or degree path, may prove more beneficial.

Implementing new procedures in any organization presents challenges; it requires cooperation from a broad range of stakeholders. USAFA is no exception; with an annual budget in the hundreds of millions of dollars and a workforce in the thousands, overhauling a half century of dogma on how cadets should spend their time is not a simple task. However, most agencies on campus have finally agreed that "... Cadets are suffering ... the schedule of calls, as it currently exists, is an obstacle, ... " (Branum 2014). The flexibility offered by ($\mathcal{S}$) is the first step in refocusing on cadets and ensuring that USAFA remains competitive among four-year universities while accomplishing its military mission.

Appendix A. Formulation ($\mathcal{S}$)

Our formulation of ($\mathcal{S}$) does not include an explicit objective function. Instead, we use notation borrowed from Brown et al. (1997) in which the goal of the program is simply to minimize the amount by which elastic constraints (denoted by $\doteq$ or $\dot{\leq}$) are violated.

Sets:

$\mathcal{U}$: Undergraduate cadets
$\mathcal{C}$: Courses
$\mathcal{T}$: Time periods
$\mathcal{P}$: Patterns
$\mathcal{B}$: Blocks
$\mathcal{E}$: Epochs
$\mathcal{R}$: Classroom types
$\mathcal{T}^V \subset \mathcal{T}$: Evening time periods
$\mathcal{C}^D \subset \mathcal{C}$: Courses that can be delayed to a later semester
$\mathcal{C}^I \subset \mathcal{C}$: Intercollegiate athletics courses

Indexed Sets:

$\mathcal{C}_b \subset \mathcal{C}$: Courses that require block b
$\mathcal{C}_r \subset \mathcal{C}$: Courses that require room type r
$\mathcal{U}_c$: Cadets registered for course c
$\mathcal{C}_u$: Courses registered for by cadet u
$\mathcal{P}_c$: Available patterns for course c
$\mathcal{C}_p$: Courses that are compatible with pattern p
$\mathcal{M}_p$: Time periods included in pattern p
$\mathcal{F}_p$: Time periods that overlap with those in $\mathcal{M}_p$
$\mathcal{T}_e$: Time periods that overlap epoch e

Parameters:

σ_{cp} = Number of preferred (c, p) sections (from PCO)
$\bar{\sigma}_c$ = Maximum number of sections of course c that can be scheduled in a semester
$\underline{g}_{cp}$ = Minimum preferred number of cadets assigned to section (c, p)
$\bar{g}_{cp}$ = Maximum preferred number of cadets assigned to section (c, p)
$\bar{d}_c$ = Maximum number of cadets allowed to delay course c
$\bar{s}_r$ = Maximum number of rooms of type r available
M_{cp} = "Big-M" bounding the number of possible cadet assignments to section (c, p)
M_u = "Big-M" bounding the number cadet assignments possible for cadet u

Variables:

Z_{cp} = Number of (c, p) sections scheduled
Y_{ucp} = 1 if a cadet u is assigned to attend course c using pattern p, 0 otherwise
D_{uc} = 1 if cadet u delays course c, 0 otherwise

Constraints:

$$Z_{cp} \dot{\le} \sigma_{cp} \quad \forall\, c \in \mathcal{C},\, p \in \mathcal{P}_c, \tag{A.1}$$

$$\sum_{p\in\mathcal{P}_c} Z_{cp} \le \bar{\sigma}_c \quad \forall\, c \in \mathcal{C}, \tag{A.2}$$

$$\sum_{c\in\mathcal{C}_r} \sum_{p\in\mathcal{P}_c:\, \mathcal{M}_p\cap\mathcal{T}_e\neq\emptyset} Z_{cp} \le \bar{s}_r \quad \forall\, r \in \mathcal{R},\, e \in \mathcal{E}, \tag{A.3}$$

$$\sum_{u\in\mathcal{U}_c} Y_{ucp} \dot{\le} \bar{g}_{cp} Z_{cp} \quad \forall\, c \in \mathcal{C},\, p \in \mathcal{P}_c, \tag{A.4}$$

$$\sum_{c\in\mathcal{C}_p} \sum_{u\in\mathcal{U}_c} Y_{ucp} \doteq 0 \quad \forall\, p \in \mathcal{P}:\, \mathcal{T}^V \cap \mathcal{M}_p \neq \emptyset, \tag{A.5}$$

$$\sum_{u\in\mathcal{U}_c} Y_{ucp} \le M_{cp} Z_{cp} \quad \forall\, c \in \mathcal{C},\, p \in \mathcal{P}_c, \tag{A.6}$$

$$\sum_{p\in\mathcal{P}_c} Y_{ucp} = 1 \quad \forall\, c \in \mathcal{C}^I,\, u \in \mathcal{U}_c, \tag{A.7}$$

$$D_{uc} + \sum_{p\in\mathcal{P}_c} Y_{ucp} \doteq 1 \quad \forall\, u \in \mathcal{U},\, c \in \mathcal{C}_u, \tag{A.8}$$

$$\sum_{c\in\mathcal{C}^D\cap\mathcal{C}_u} D_{uc} \le 1 \quad \forall\, u \in \mathcal{U}, \tag{A.9}$$

$$\sum_{u\in\mathcal{U}_c} D_{uc} \le \bar{d}_c \quad \forall\, c \in \mathcal{C}^D, \tag{A.10}$$

$$\sum_{c\in(\mathcal{C}_b\cap\mathcal{C}_u)} \left(\sum_{p\in\mathcal{P}_c:\, t\in\mathcal{F}_p} Y_{ucp} + M_u \sum_{p\in\mathcal{P}_c:\, t\in\mathcal{M}_p} Y_{ucp} \right) \le M_u \quad \forall\, b \in \mathcal{B},\, u \in \mathcal{U},\, t \in \mathcal{T}, \tag{A.11}$$

$$Z_{cp} \in \{0\} \cup \mathbb{Z}^+ \quad \forall\, c \in \mathcal{C},\, p \in \mathcal{P}_c,$$
$$Y_{ucp} \in \{0, 1\} \quad \forall\, u \in \mathcal{U},\, c \in \mathcal{C}_u,\, p \in \mathcal{P}_c,$$
$$D_{uc} \in \{0, 1\} \quad \forall\, u \in \mathcal{U},\, c \in \mathcal{C}^D \cap \mathcal{C}_u.$$

In the absence of an explicit objective function, the constraints that define ($\mathcal{S}$) either ensure the solution is feasible under USAFA requirements or that said solution is as desirable as possible to departments and cadets. Constraint (A.1) encourages adherence to the PCO and its violation, that is, scheduling more sections in a particular pattern than requested by the PCO, incurs a penalty based on the total number of sections initially requested. (Larger penalties are assigned to courses with fewer sections.) Constraints (A.2) and (A.3) represent strict manpower and resource constraints on the number of sections and cannot be violated. Specifically, constraint (A.2) limits the number of sections of a particular course that can be taught per semester per departmental guidelines. Constraint (A.3) reflects physical resources by limiting all sections of all courses scheduled during a particular epoch to the number of rooms that are eligible and available across campus. Constraints (A.4)–(A.7) similarly control how cadets are assigned to course sections. Constraints (A.4) and (A.5) attempt to honor PCO requests with flexible limits on section enrollment and minimal assignment of cadets to evening sections, respectively, via penalties similar to those associated with constraint (A.1). Constraints (A.6) and (A.7) are hard constraints that enforce a logical restriction against assigning cadets to nonscheduled sections, and that require intercollegiate athletes to attend the necessary practice (offered only in a single pattern), respectively.

Ultimately, every registration must either be met or delayed to a later semester when the course is reoffered. Historically, finding a schedule that meets these requirements and does not deviate from the initial PCO is highly unlikely. Instead, the *previous USAFA model* uses guidelines and heuristics to try and meet as many course registrations as it can, and then begins a lengthy negotiation process with departments in an attempt to fix the remaining registration conflicts. Although ($\mathcal{S}$) has the flexibility to deviate from the PCO and generate a schedule with no registration conflicts, it does so at the cost of violating flexible constraints.

Completely eliminating registration conflicts inevitably results in a schedule that is unlikely to be endorsed by all departments. Thus, violation of constraint (A.8) remains flexible, with each violation translating into a registration conflict that must be resolved by USAFA in some fashion. Because these violations represent a measure of the man-hours necessary to adjudicate a final schedule, their penalty is high relative to those associated with other flexible constraints (i.e., six times, on average) and is proportional to the credit hours associated with the course c that is responsible for the conflict. Furthermore, although registration conflicts are definitely undesirable, delaying registrations only simplifies the current semester by complicating a later one. To this end, constraint (A.9) enforces a USAFA restriction that each cadet delay no more than a single registration in a given semester, and constraint (A.10) limits the number of registrations for course c that may be delayed. (The parameter $\bar{d}_c$ is based on the future availability of course c.)

Finally, constraint (A.11) ensures that a cadet cannot be in more than one place at a time and must account for all sections to which a cadet is assigned, the periods during which they meet, and all possible ways in which these periods may overlap within individual blocks. Initially, this requirement was modeled as two separate constraints, an equality and an inequality, to possibly generate a stronger formulation. This construction was ultimately discarded because any gains were overshadowed by the increase to the required number of variables and constraints.

The objective function is implicit in this formulation and consists of the sum of the following terms representing weighted deviations associated with the following: (1) scheduling a particular section more often than is requested by the PCO: constraint (A.1); (2) exceeding a maximum class size: constraint (A.4); (3) assigning cadets to evening periods: constraint (A.5); and (4) incurring registration conflicts: constraint (A.8).

Appendix B. Tractability Improvements

Size Reduction and Bound Tightening

The following tactics reduce the size of our monolith, tighten its bounds, and provide our instances with an initial feasible solution, thus expediting solutions (Klotz and Newman 2013). We reduce the number of assignment variables by using indexed sets (Appendix A). For example, without careful control, Y_{ucp} can include approximately 10^9 binary variables. However, by ignoring instances in which a cadet u does not require a course c, as well as those in which c is not offered in pattern p, this number is drastically reduced to just over 200,000. Controlling Y_{ucp} in this way also allows constraints involving the eliminated assignment possibility to be removed and reduces the number of constraints defined in constraints (A.5)–(A.8). Similar reductions apply to Z_{cp}, D_{uc}, and their associated constraints.

The size of the model is further limited by removing variables and constraints associated with penalties below a certain threshold and incorporating them elsewhere in the model. For example, constraint (A.4) originally included the lower bound $\underline{g}_{cp}$ on the average number of cadets assigned to each section. However, the difficulties associated with exceeding the upper bound (e.g., faculty size, room types, office hours) heavily outweigh the inconvenience of falling below the minimum. Therefore, we remove the constraints related to under-assignment and instead modify the penalties associated with constraint (A.1) to be proportional to $\underline{g}_{cp}$ (i.e., as $\underline{g}_{cp}$ increases, so does the penalty for opening an additional section). This modification results 16,000 fewer variables and constraints, and eschews solving conflicts simply by opening many small sections.

Although some variables cannot be explicitly removed from ($\mathcal{S}$), they can be fixed for a particular instance based on USAFA mandates. For example, in the fall 2016 PCO, USAFA identified several specific courses that cadets must delay to a summer semester, if possible. (Resources for the summer section had already been committed.) In addition to fixing many Y_{ucp} variables a priori, any constraint that involves Y_{ucp} is also strengthened with the removal of penalty variables.

The model is also tightened via the following means: (1) An upper bound on the sum of assignment binaries is equal to the total number of registered courses. (2) The magnitude by which flexible constraints can be violated is limited. At present, this upper bound is prescribed by USAFA leadership and is applied a priori across all sections and cadets. Future iterations of ($\mathcal{S}$) will allow input from departments to tailor individual limits to their particular needs. (3) Flexible constraints associated with penalties above a certain threshold are converted into hard constraints. An example of this latter methodology involves constraint (A.3); although not impossible, appropriating additional rooms for regular academic use is so difficult that violating constraint (A.3) is never optimal under any solution found by ($\mathcal{S}$). As a hard constraint, constraint (A.3) requires fewer variables, is stronger, and, as implemented now, reduces the time ($\mathcal{S}$) spends searching through feasible solutions that are never optimal in practice. The relevant node logs reveal that, in solving ($\mathcal{S}$), CPLEX never branches off the root node; all optimization is done through a combination of cuts and internal heuristics. Although our own cuts are effective, it should be possible to examine the postprocessed model to identify the most powerful cuts and apply them a priori.

We also tune our solver; probing fixes binary variables involved in packing and partitioning constraints. Setting a probe value of 2 reduces our solve time in practice. Furthermore, because the variables greatly outnumber the constraints, we employ options to ensure that our solver uses a primal versus dual simplex at the root node. Finally, we introduce a heuristic to produce an initial feasible solution. All tractability enhancements described in this section are also applied to the process of finding an initial solution, described next.

Initial Solution for ($\mathcal{S}$)

To find an initial solution, ($\mathcal{S}$) partitions the data it receives according to year group (i.e., sophomores, juniors, seniors) to solve smaller subproblems that schedule courses with unique

Figure B.1. (Color online) We Illustrate the Performance of $(\mathcal{S})$ With and Without an Initial Solution

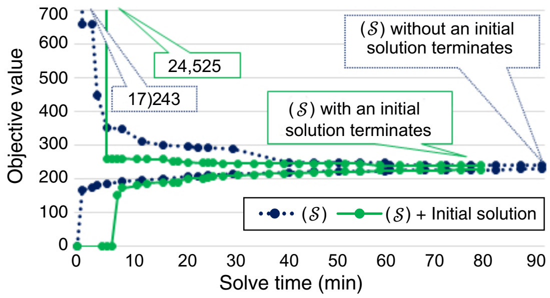

Note. Early on, $(\mathcal{S})$ closes its optimality gap more quickly without an initial solution; however, time spent deriving the initial solution is quickly recuperated and allows a 19 percent reduction in overall solve time.

year-group registrations. Figure B.1 demonstrates how the derived solution affects the overall solve time and the rate at which $(\mathcal{S})$ closes its optimality gap. The pseudocode that follows (Figure B.2) details how the initial solution is derived.

Figure B.2. The Pseudocode to Derive an Initial Solution for $(\mathcal{S})$ Is Shown

Initial Solution for $(\mathcal{S})$;
`//Partition Cadets (Freshmen schedules are preset and not included)`
$\mathcal{U}_1 \leftarrow$ Senior cadets $u \in U$;
$\mathcal{U}_2 \leftarrow$ Junior cadets $u \in U$;
$\mathcal{U}_3 \leftarrow$ Sophomore cadets $u \in U$;
`//Identify courses unique to each year-group`
for $n = 1, \ldots, 3$ **do**
  $\mathcal{Y}_n \leftarrow \left(\bigcup_{u \in \mathcal{U}_n} C_u\right) \setminus \left(\bigcup_{u' \in \mathcal{U}_{n'}, n' \neq n} C_{u'}\right)$
end
`//Make cadet assignments for all non-unique courses`
$(\mathcal{S})^{\mathcal{Y}_0} \leftarrow (\mathcal{S})$ without constraints or registrations related to courses $c \in \bigcup_{n=1}^{3} \mathcal{Y}_n$;
Solve $(\mathcal{S})^{\mathcal{Y}_0}$;
`//Save assignments`
foreach Y_{ucp} **in** $(\mathcal{S})^{\mathcal{Y}_0}$ **do**
  $y^0_{ucp} \leftarrow Y_{ucp}$
end
`//Assign each year-group individually;`
`//Retain assignments for non-unique courses`
for $n = 1, \ldots, 3$ **do**
  $(\mathcal{S})^{\mathcal{Y}_n} \leftarrow (\mathcal{S})$ with only $u \in U_n$ and the associated registrations and constraints;
  foreach y^0_{ucp} **such that** $u \in U_n$ **do**
    Add constraint $Y_{ucp} = y^0_{ucp}$ to $(\mathcal{S})^{\mathcal{Y}_n}$
  end
  Solve$(\mathcal{S})^{\mathcal{Y}_n}$;
  foreach Y_{ucp} **in** $(\mathcal{S})^{\mathcal{Y}_n}$ **do**
    $y^n_{ucp} = Y_{ucp}$
  end
end
`//Derive initial solution using assignments from each class`
$(\mathcal{S})^I \leftarrow (\mathcal{S})$;
for $n = 1, \ldots, 3$ **do**
  Add constraint $Y_{ucp} = y^n_{ucp}$ to $(\mathcal{S})^I$
end
Solve $(\mathcal{S})^I$;
Return solution.

Appendix C. Penalty Sensitivity Analysis

The formulation in Appendix A includes four flexible constraints: (A.1), (A.4), (A.5), and (A.8), the violation of which is controlled by three types of penalties. (Constraints (A.4) and (A.5) are both related to overfilling a section and, thus, use similar penalties.) Although USAFA leadership provides course-specific and pattern-specific penalties for the violation of each constraint, we conduct our own analysis to determine the appropriateness of the magnitude of the penalties. Using increments of 25 percent, the penalty for each type of violation (i.e., using additional sections, overfilling a section,

Figure C.1. (Color online) We Illustrate Quality Metrics, (a) Solve Time, (b) Total Sections, and (c) Unmet Preferences, as Functions of Adjustments to the Penalties for Flexible Constraint Violation

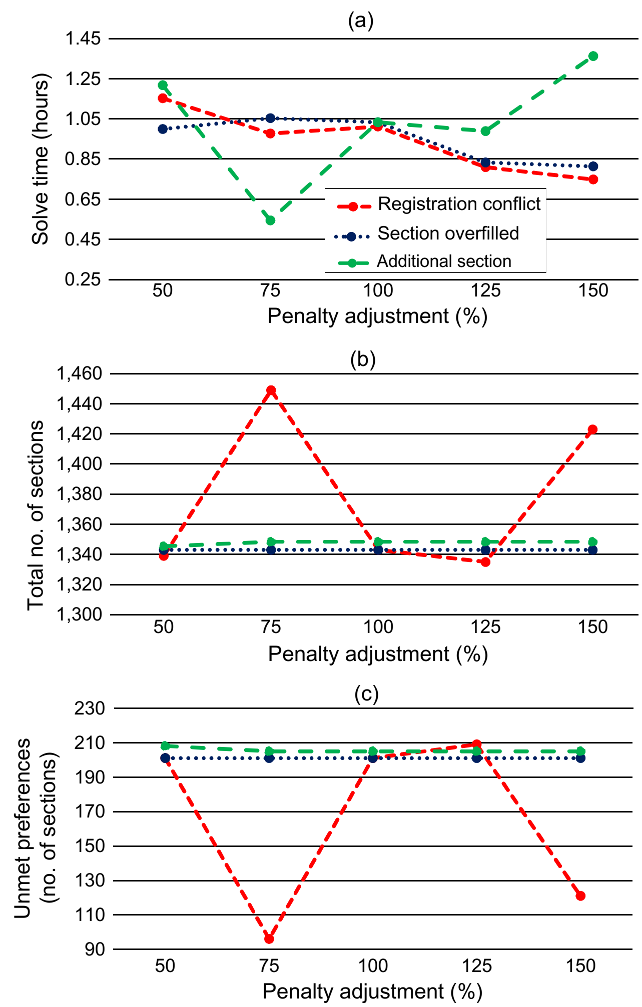

Notes. For each violation (i.e., registration conflict, overfilling a section, and using an additional section), the associated dashed line represents the result of adjusting its penalty while all others are held constant. The trend lines in each chart are slightly perturbed for ease of visualization.

and incurring a registration conflict) varies from 50 to 150 percent of the prescribed USAFA value. Using the fall 2016 PCO and registration data, our analysis solves ($\mathcal{S}$) for each possible combination. Because of the large number of requisite runs, we use an optimality gap of 8 percent to reduce the time necessary to test all 125 combinations. Figure C.1 summarizes some of the main comparative results of the analysis.

Solutions are evaluated using the four metrics presented in our results (i.e., registration conflicts, unmet preferences, total sections, and solve time). Because registration conflicts are assigned a high penalty relative to those associated with other violations, their occurrence varies only slightly within our analysis, between 44 and 48, and is not included in the summaries provided in Figure C.1. Concerning the remaining metrics, unmet preferences range from 96 to 222, total sections from 1,322 to 1,449, and run times from 0.5 to 1.9 hours.

The analysis shows that variation of the penalties associated with overfilling and adding sections has minimal effect on both the number of unmet preferences and the total number of sections. Furthermore, adjustments to registration conflict penalties result largely in simple trade-offs between unmet preferences and total number of sections, as predicted by Figure 4. Based on the results illustrated in Figure C.1(a), we increase the penalties associated with registration conflicts to 125 percent of their original values. Coincidentally, this increase reduces average run times by approximately 20 percent and changes the solution slightly: the eight additional unmet preferences are offset by the eight fewer total sections.

References

Alsmadi OMK, Za'er S, Abu-Al-Nadi DI, Algsoon A (2011) A novel genetic algorithm technique for solving university course timetabling problems. *Proc. 7th Internat. Workshop Systems Signal Processing Appl.* (*WOSSPA*) (Institute of Electrical and Electronics Engineers, Washington, DC), 195–198.

Aubin J, Ferland JA (1989) A large scale timetabling problem. *Comput. Oper. Res.* 16(1):67–77.

Aycan E, Ayav T (2009) Solving the course scheduling problem using simulated annealing. Accessed September 12, 2017, http://web.iyte.edu.tr/tolgaayav/papers/Aycan-CourseSch.pdf.

Babaei H, Karimpour J, Hadidi A (2015) A survey of approaches for university course timetabling problem. *Comput. Indust. Engrg.* 86:43–59.

Bakir A, Aksop C (2008) A 0-1 integer programming approach to a university timetabling problem. *Hacettepe J. Math. Stat.* 37(1): 41–55.

Boland N, Hughes B, Merlot L, Stuckey P (2008) New integer linear programming approaches for course timetabling. *Comput. Oper. Res.* 35(7):2209–2233.

Branum D (2014) Pathways to Excellence Team envisions changes to improve cadets' schedules. Accessed September 12, 2017, http://www.usafa.af.mil/News/Features/Display/Article/619693/pathways-to-excellence-team-envisions-changes-to-improve-cadets-schedules/.

Brown G, Dell R, Wood K (1997) Optimization and persistence. *Interfaces* 27(5):15–37.

Burke E, Elliman D, Ford P, Weare R (2005) Examination timetabling in British universities—A survey. Burke E, De Causmaecker P, eds. *Proc. Internat. Conf. Practice Theory Automated Timetabling*, LNCS 1153 (Springer, Berlin), 76–90.

Caprara A, Fischetti M, Toth P (1999) A heuristic method for the set covering problem. *Oper. Res.* 47(5):730–743.

Cheng E, Kruk S, Lipman M (2003) Flow formulations for the student scheduling problem. Burke E, De Causmaecker P, eds. *Proc. Internat. Conf. Practice Theory Automated Timetabling*, LNCS 2740 (Springer, Berlin), 299–309.

Dostert M, Politz A, Schmitz H (2016) A complexity analysis and an algorithmic approach to student sectioning in existing timetables. *J. Scheduling* 19(3):285–293.

Easton K, Nemhauser G, Trick M (2001) The travelling tournament problem: Description and benchmarks. Walsh T, ed. *Proc. 7th Internat. Conf. Principles Practice Constraint Programming*, LNCS 2239 (Springer, Berlin), 580–585.

Eglese R, Rand G (1987) Conference seminar timetabling. *J. Oper. Res. Soc.* 38(7):591–598.

Feng X, Lee Y, Moon I (2016) An integer program and a hybrid genetic algorithm for the university timetabling problem. *Optim. Methods Software* 32(3):1–25.

Fong CW, Asmuni H, McCollum B, McMullan P, Omatu S (2014) A new hybrid imperialist swarm-based optimization algorithm for university timetabling problems. *Inform. Sci.* 283:1–21.

Gonzalez G (2011) Heuristics and cuts for integer programming scheduling models with applications at the Air Force Academy and in Generator Dispatch. Doctoral dissertation, Colorado School of Mines, Golden, CO.

Hertz A (1991) Tabu search for large scale timetabling problems. *Eur. J. Oper. Res.* 54(1):39–47.

HQ USAFA (2012) Cadet administrative management information system assignment and scheduling algorithms: Final report. Report, United States Air Force Academy, Colorado Springs, CO.

Jaumard B, Semet F, Vovor T (1998) A generalized linear programming model for nurse scheduling. *Eur. J. Oper. Res.* 107(1):1–18.

Khonggamnerd P, Innet S (2009) On improvement of effectiveness in automatic university timetabling arrangement with applied genetic algorithm. Sohn S, Kwak KD, Um K, Lee GY, Ko F, eds. *Proc. Fourth Internat. Conf. Comput. Sci. Convergence Inform. Tech.* (*ICCIT'09*) (Institute of Electrical and Electronics Engineers, Washington, DC), 1266–1270.

Klotz E, Newman A (2013) Practical guidelines for solving difficult mixed integer linear programs. *Surveys Oper. Res. Management Sci.* 18(1):18–32.

LaRivee D (2017) Future schedule of calls response. Memorandum, United States Air Force Academy, Colorado Springs, CO.

Lewis R (2008) A survey of metaheuristic-based techniques for university timetabling problems. *OR Spectrum* 30(1):167–190.

Lübbecke ME (2015) Comments on: An overview of curriculum-based course timetabling. *TOP* 23(2):359–361.

MirHassani S, Habibi F (2013) Solution approaches to the course timetabling problem. *Artificial Intelligence Rev.* 39(2):1–17.

Müller T, Murray K (2010) Comprehensive approach to student sectioning. *Ann. Oper. Res.* 181(1):249–269.

National Center for Education Statistics (2015) Digest of educational statistics. Accessed September 12, 2017, https://nces.ed.gov/programs/digest/d15/.

Phillips AE, Waterer H, Ehrgott M, Ryan DM (2015) Integer programming methods for large-scale practical classroom assignment problems. *Comput. Oper. Res.* 53:42–53.

Sampson SE, Weiss EN (1995) Increasing service levels in conference and educational scheduling: A heuristic approach. *Management Sci.* 41(11):1816–1825.

Sampson SE, Freeland JR, Weiss EN (1995) Class scheduling to maximize participant satisfaction. *Interfaces* 25(3):30–41.

Schaerf A (1999) A survey of automated timetabling. *Artificial Intelligence Rev.* 13(2):87–127.

Tuga M, Berretta R, Mendes A (2007) A hybrid simulated annealing with Kempe chain neighborhood for the university timetabling problem. Lee R, Chowdhury MU, Ray S, Lee T, eds. *Proc. 6th Internat. Conf. Comput. Inform. Sci.* (*ICIS*) (Institute of Electrical and Electronics Engineers, Washington, DC), 400–405.

Wren A (1996) Practice and theory of automated timetabling: Scheduling, timetabling and rostering—A special relationship? Burke E, Ross P, eds. *Proc. Internat. Conf. Practice Theory Automated Timetabling*, LNCS 1153 (Springer, Berlin), 46–75.

Verification Letter

David R. Larivee, AD-25, Director, Pathways to Excellence Team, United States Air Force Academy, writes:

"Let me take this opportunity to endorse the outstanding work of Lt Col Gerry Gonzalez and Major Christopher Richards in support of developing a new scheduling algorithm for the United States Air Force Academy (USAFA). USAFA has operated utilizing an alternating day format since its founding in 1954. As education and technology have advanced, this structure now constrains our charge to find ways to improve the student experience and the cost to support and/or modernize the current system is unmanageable.

"An Air Force business process analysis recommends that the institution adopt a commercial off-the-shelf (COTS) Student Information System (SIS) for consolidating and managing the 85 current applications that assist USAFA personnel in tracking the progress of cadets from pre-candidacy through graduation. Leadership had previously rejected COTS products due to the institutional assumption that our military, flight, athletic, and academic requirements would be more expensive to fit into a COTS product than to build a custom SIS.

"My Pathways to Excellence team works directly for the USAFA Superintendent, Lt General Michelle Johnson to analyze military, flight, athletic, and academic requirements and generate a new schedule of calls concept for the institution. For this analysis we needed a new assignment and scheduling model that allowed us to rapidly evaluate business rule changes and allowed users to focus on results rather than mechanics. The work at the Colorado School of Mines by Lt Col Gerry Gonzalez and Maj Christopher Richards has resulted in a robust scheduling tool which enables us to evaluate various schedule alternatives. This scheduling tool has become our primary means for addressing both the SIS acquisition decision and for evaluating a range of complex policy options to inform strategic decisions about new schedule of calls possibilities. The immediate contribution was our first ever ability to unequivocally debunk the assumption that USAFA needed a custom SIS. This insight alone saved the institution about $120 million by choosing a $100 million COTS solution instead of a $220 million custom solution.

"Due to the rapid change and execution capability offered by the scheduling tool, the Pathways team has to-date evaluated nearly a hundred different scheduling scenarios to arrive at a greatly improved schedule of calls structure for USAFA. Our suggested construct to leadership shows the ability to offer more pedagogically-sound formats for our classes and training programs, with fewer scheduling conflicts, more efficient and effective use of cadet and faculty time, while still honoring the constraints of the various mission elements. Last week we took our analysis and recommendation to the USAFA Faculty Senate and it overwhelmingly recommended adoption of our new schedule of calls approach made possible by the aforementioned scheduling tool.

"Thanks to the success of this scheduling tool, we can test and evaluate alternatives and implement the changes necessary to remain at the forefront of developing future officers. We can now base decisions on known costs and benefits rather than gut feel. We are extremely pleased with this research effort."

Gerardo Gonzalez is an assistant professor in the Management Department at the United States Air Force Academy (USAFA). He obtained his BS in operations research at USAFA, his MBA at the Naval Postgraduate School, and his PhD in mineral and energy economics at the Colorado School of Mines. His interests are in deterministic optimization modeling, especially as it applies to scheduling, project management, and decision analysis.

Christopher Richards is an analytical scientist at the Pentagon within the United States Air Force. Prior to his doctoral research at Colorado School of Mines, he was an associate professor at the United States Air Force Academy. He obtained his BS in mathematics and computer science at Pepperdine University and his MS in operations research at the Air Force Institute of Technology. He specializes in decision analysis, assisting both the Department of Homeland Security and the Joint Improvised Explosive Device Defeat Organization in building course-of-action selection models.

Alexandra Newman is a professor in the Mechanical Engineering Department at the Colorado School of Mines (CSM). Prior to joining CSM, she was a research assistant professor at the Naval Postgraduate School in the Operations Research Department. She obtained her BS in applied mathematics at the University of Chicago and her PhD in industrial engineering and operations research at the University of California at Berkeley. She specializes in deterministic optimization modeling, especially as it applies to energy and mining systems, and to logistics, transportation, and routing. She received a Fulbright Fellowship to work with industrial engineers on mining problems at the University of Chile in 2010 and was awarded the Institute for Operations Research and the Management Sciences (INFORMS) Prize for the Teaching of Operations Research and Management Science Practice in 2013.