Resource Allocation Modeling Techniques Applied to Air Force Crew Scheduling Problem

Submitted by:

Kepler Research, Inc.

February 2012

1. Introduction and Background

The 20th Air Force maintains and operates the Air Force's Intercontinental Ballistic Missile (ICBM) force, providing on-alert, combat-ready ICBMs to the President of the United States. The ICBM mission requires staffing trained and certified personnel for alert shifts (alerts) at geographically separated operation centers (op centers). Personnel must be scheduled for daily alerts, as well as monthly training and annual certifications. Therefore, the schedules must include the op center staffing and the training events, training resources, and trainers necessary to maintain current certification for all personnel.

The current scheduling process is manual and time-consuming, with many tradeoffs, conflicting availability, and limited training capacity. Schedulers create the monthly schedule for the following month. Once the schedule is created, unanticipated events or changes often cause rework during the month. With so many moving parts and numerous constraints, the 20th Air Force needed a way to automate the scheduling tasks and optimize the allocation of their resources, while also being flexible and dynamic to respond to changing conditions.

Currently schedulers create the monthly schedule using TimePiece™, a software product developed by Kepler Research (Kepler) that facilitates scheduling and provides visibility throughout the organization. TimePiece™ has enhanced the effectiveness of the scheduling process and reduced the time required to schedule a single Wing's resources to 1-2 weeks; however it does not provide the ability to optimally allocate resources and is not dynamic to respond to last minute changes.

Kepler, together with graduate students from George Mason University's (GMU) Operations Research (OR) program, designed and developed an optimization model for the 20th Air Force to optimally allocate their resources to meet critical mission requirements. The goal of the model was to create the monthly schedule and dramatically improve the efficiency and performance over the existing scheduling process.

2. Methodology and Approach

The team utilized linear programming (LP) and integer programming (IP) techniques to develop a mixed-integer program (MIP) for assigning resources to requirements. To fulfill requirements, the Air Force Operation Centers Scheduling (AFOSC) model executes in four phases, shown below in Figure 1. Each execution phase consists of running a mathematical program to perform a series of assignments, and contains several parameter assignments to be used in subsequent mathematical program execution phases. The model initially creates the "original schedule" – the monthly schedule for alerts and training for the following month. During the month the original schedule is being executed, the model has two methods to address daily changes in personnel availability and creates a "new schedule", as necessary.

Figure 1 Execution Phases for AFOCS Model

The following sections briefly describe at a high level the model assignments, constraints, objective functions, and dynamic portions.

Model Assignments

The main assignments made in each phase assign a person *p* to an event *e* on day *d*. Events are one of the following: *alert*, *o-day, leave,* classroom training events *(T1, T3, T4),* or simulation training events *(TR, eval)*. Personnel assigned to alerts are assigned to a site within their squadron, and are paired together with one person serving as the crew commander and the other as the deputy crew commander. A backup crew is assigned each day in case a crew member is unable to pull alert. The model assigns instructors and students to classroom events and *TR* simulation events, and assigns evaluators and students to *eval* simulation events. In the final phase, the model assigns instructors, evaluators, students, and simulators to simulator time slots for executing the simulation training events.

Model Constraints

Each person must complete the required classroom training events and the *TR* simulation event monthly. If a person is up for their annual evaluation, then they must complete the *eval* simulation event. Classroom training events require 1 instructor, *TR* events require 2 instructors, and *eval* events require 3 evaluators. Personnel are grouped into functional roles (crewdog, instructor, evaluator, flight commander), and are limited in the number of alerts they can pull per month based on their functional role (8 for crewdogs, 2 for the rest). Personnel must have 3 days between alerts, and cannot perform anything the day after they pull alert, which is considered their off-day (o-day). There are 2 simulators and 5 simulator slots per simulator available each day; personnel pulling alert the next day can only complete their *TR* or *eval* event in the first 3 simulator slots due to pre-alert rest restrictions.

Model Objective Functions

The most important metric for the 20th Air Force is the *Integral Crew Rate (ICR)*. This metric tracks the percentage of time a person is paired with their "crew mate" in alerts and training. The higher the ICR, the better. Therefore, in assigning personnel to events, the main objective function was to maximize personnel ICR. In addition, as MIP solvers often get greedy, a penalty function was added to spread out the number of alerts each person pulls within their functional grouping. This ensured that a person, such as a crewdog, does not pull 8 alerts per month, while another one only pulls 3.

Dynamic Portions

The dynamic portions allowed the model to be flexible and responsive in dealing with unforeseen personnel unavailability. Sometimes a person cannot meet the rigorous requirements to pull an alert, and therefore cannot perform the duty on a particular day. The alert duty must still be filled, and thus the dynamic portions were created to address those changing requirements.

There are two dynamic portions: a *heuristic function* and a *mathematical program*. The *heuristic function* can be used if a single crew commander or deputy cannot pull alert, or both cannot pull alert. The *heuristic* essentially takes the backup crew and assigns them to the site where the person can no longer pull alert. The person who can no longer pull alert is put on leave, and a new backup crew is enlisted. This causes almost no disruption to the original schedule, and eliminates any secondary effects.

The *mathematical program* handles large changes due to many people suddenly becoming unavailable to pull alert. The program creates a new month schedule, with the days completed before the disruption remaining the same, and takes into account the new personnel unavailability. In creating the new schedule, the program minimizes the difference between the new schedule and the original schedule, while maintaining feasibility. The goal of the program is to minimize the disruption to the original schedule, rather than maximizing ICR, as this is most important to the 20th when such events occur.

Model Execution and Solution Times

To solve a MIP, the model is sent to a commercial solver, such as CPLEX or Gurobi, that uses various algorithms and techniques to solve the mathematical program. For this model, there were 4 mathematical programs to solve, one for each phase, each with a different set of decision variables. The solution from each phase was used as input for the next phase, which the solver would use in solving for the next phase's decision variables. This continued until all 4 phases were complete and the monthly schedule was created. If any personnel availability changed during the execution of the monthly schedule, the dynamic portions were executed based on the need.

Solving a MIP takes longer than solving a LP due to the model's integer and binary decision variables. For a LP, the decision variables are continuous and the solver can find the solution by moving along the borders of the decision space. With integer and binary variables, the solver initially solves the LP version of the model, called the LP-relaxation, and then performs various techniques to find the optimal integer solution within a grid-like decision space. The LP-relaxation solution is called the LP upper-bound (for maximization problems), which is the best any solution can be, and is referred to as *optimality*. The goal of the solver is to minimize the gap between the LP upper-bound and the current integer solution found.

For this model, the solver finds an integer solution after 10min that is within 15.4% of optimality. After 2 hours, within 3.2% of optimality. After 3 hours, within 1.97% of optimality. After 4 hours, the solution was still at 1.97% of optimality, which was deemed an acceptable tolerance, so the solver was stopped. In practice, a tolerance is set for the solver to prevent it from churning on an IP for an unknown amount of time. Initially for this model no tolerance was set to test how close to optimality the solver would get at various time intervals.

3. Conclusion

The Resource Allocation problem is a common problem seen in almost every organization. Its difficulty arises out of having numerous conflicting resource availabilities and countless constraints on when, where, and how resources can be assigned, not to mention interdependencies among the resources that also impact their assignment. The 20th Air Force historically dealt with this problem manually by getting all the various schedulers and resource managers into a room for 1-2 weeks and letting them hash out a schedule for the following month. This process was costly and very time consuming, and the 20th Air Force needed a better solution. The Kepler and GMU team developed a model that quickly optimizes the allocation of 20th Air Force resources, significantly reducing the time required to complete the monthly schedules from 1-2 weeks to about 3 hours and producing a solution within 2% of optimality. The model ensures the allocation of resources is balanced across the force, and allows schedulers to quickly address unforeseen changes in resource availability and still fulfill critical mission requirements.

For more information

To learn more about our services, contact:

Rebecca McCrabb

Business Development Lead, Advanced Analytics

(571) 319-8137

Rebecca.McCrabb@KeplerResearch.com

About Us

Kepler Research, Inc. is an employee-owned management consulting firm, specializing in acquisition services, information technology, and advanced analytics. We are committed to excellence, bringing innovative solutions combined with comprehensive in-depth analyses and practical expertise. We partner with our clients on their most difficult problems to provide holistic solutions and achieve superior, lasting results. At Kepler, challenge is welcome, work is passion, success is absolute.