Operations Research Applications in Canada's Department of National Defence

René Séguin

Defence Research and Development Canada
Centre for Operational Research and Analysis, Canada
rene.seguin@drdc-rddc.gc.ca

Keywords: Applications, scheduling, resource allocation, routing, timetabling, simulation

1 Introduction

Operations Research (OR) originated in defence applications during World War II. Nowadays, Canada's Department of National Defence and Defence departments around the world continue to rely on OR to provide them with the scientific support that enables them to make better decisions. The field is still a reservoir of complex applications and unsolved problems, as rich today as it has been in the past. This talk will present a selection of applications that the author has encountered in the last five years and for which most is still to be done.

2 Scheduling of Airborne Surveillance Resources for Complex Imagery Requirements

Any organization that must observe and gather information on a set of objects of interest is faced with the problem of selecting, assigning, scheduling and routing surveillance assets (combination of platforms and sensors). Fixed path assets such as satellites and manoeuvrable airborne assets are usually considered. Surveillance requests specify frequencies of observation, time constraints (due dates and dynamic time windows), combinations of imagery equipment type preferences (radar, electro-optic, infrared and/or hyperspectral), priorities, imagery quality and object locations and types (points, lines, areas). Requests that are not fulfilled immediately are rescheduled for another collection cycle as necessary, which implies a rolling horizon. The objective is to find an optimum collection strategy to maximize the utility of the requests collected in a given period of time or collection cycle. This defines a very complex version of the classical Selective Traveling Salesman Problem which is defined as the problem of finding a circuit through different locations that maximizes some utility function (such as bonus points accumulated or profit earned at visited locations) while respecting a traveling constraint (usually a maximum distance or a time limit).

3 Scheduling and Resource Allocation at Specialized Air Forces Schools

The Air Force needs to train highly specialized personnel such as pilots, navigators or electronic sensor operators. In essence, this is accomplished as in any other schools or colleges by sequentially matching instructors, students and physical resources to complete a required program. However, in this military environment several major characteristics make the problem quite different. The specialized resources used (pilots, planes, simulators) are very expensive and thus used to capacity if not lacking. The ability to use planes is dynamic due to failures, maintenance and weather. The availability of instructors is also dynamic due to sickness, leave and secondary duties, which delays individual classes. Coupled with the limited number of resources, the delays generate competition for the resources and sometimes for several consecutive days, which means that a priority scheme has to be used to decide who

gets the resources. Individual classes often require two or four instructors. For example, for each flight, specific number of instructors and students are required in addition to pilots and the aircraft. There are many failure points for the students but a limited number of retests are allowed. While a student is retested, the whole class waits. Similarly, if a student is sick, the class is idle until his return. About a quarter of the instructors are posted out of the school every year. They are replaced by students from one of the programs. If graduation and posting are not properly synchronized, a temporary shortage of instructors occurs. The objective here is to find out at what date to start each course and minimize their time span. Finding out the bottlenecks is a secondary objective.

4 Production, Absorption and Attrition of Air Force Pilots

The military pilot community can be viewed as a system in a rather fragile equilibrium. Tracking the community's health is a constant concern of the Air Force. Pilots generally go from undergraduate training to operational training units and on to squadrons. They go from learning how to fly to learning how to fly types of aircraft (helicopters, multi engines or jets) and finally how to fly specific aircraft. To become operationally qualified, pilots are required to acquire flying hours on specific aircraft and this is accomplished thru mentoring with experienced pilots. However, experienced pilots have various roles in the Air Force; they can be instructors at the undergrad school or in the training units, they can be operational pilots at squadrons or in exchange/foreign positions, they can be on a demonstration unit and finally be on a staff position. During their career, pilots may transition from one specific aircraft to another, sometimes even to a different type. Each transfer requires retraining. Furthermore, new fleets and newer models of a current aircraft are acquired and old ones retired. New fleets required more pilots from the community as a whole whereas new models call for transition of pilots. All this implies more training. The number of "flying" positions is strictly controlled and limited by a number of factors. In practice, this means that inexperience pilots push experience pilots out. Simulating the pilot community to give decision makers the ability to gauge the effect of factors such as attrition of experienced pilots, production level at the undergrad school or flying rates on the overall health of the pilot system is of paramount importance. This allows to find out, among others, at what rate the transition to a new model of aircraft or the retirement of an old fleet can be accomplished, how many new recruits can be absorbed in a squadron, how fast can experience be acquired thru mentoring.