Interfaces
Vol. 43, No. 3, May–June 2013, pp. 232–239
ISSN 0092-2102 (print) | ISSN 1526-551X (online)

http://dx.doi.org/10.1287/inte.2013.0683

Mathematical Programming Guides Air-Ambulance Routing at Ornge

Timothy A. Carnes
Link Analytics, Atlanta, Georgia 30338, tim.carnes@linkanalytics.com

Shane G. Henderson, David B. Shmoys
School of Operations Research and Information Engineering, Cornell University, Ithaca, New York 14853
{sgh9@cornell.edu, david.shmoys@cornell.edu}

Mahvareh Ahghari
Ornge, Mississauga, Ontario L4W 5H8, Canada, mahghari@ornge.ca

Russell D. MacDonald
Ornge, Mississauga, Ontario L4W 5H8, Canada; and Faculty of Medicine, University of Toronto, Toronto, Ontario, Canada, rmacdonald@ornge.ca

Ornge provides air-ambulance services to the province of Ontario. A major part of its services involves prescheduled transports between medical facilities. These transports almost exclusively require fixed-wing aircraft because of the distances involved and cost considerations. The requests are received in advance, scheduled overnight, and typically executed the following day. We describe our work in developing a planning tool that determines an assignment of requests to aircraft to minimize cost, subject to a range of complicating constraints. Ornge flight planners use the tool daily, resulting in substantial savings over the manual approach they used previously to develop schedules. We describe the problem, our formulation, its implementation, and the impact on operations at Ornge.

Key words: set partitioning; transport medicine; dial-a-ride.
History: This paper was refereed.

The province of Ontario, Canada has approximately 13 million residents spread over approximately one million square kilometers. Ontario ensures that all its residents receive a high level of medical care, partly through an advanced medical transport capability. Ornge, a not-for-profit company, provides these transport services primarily through air-ambulance services, although it also has several land ambulances. Each year, Ornge transports approximately 19,000 patients. These transports can be either scene calls or interfacility transports. Scene calls refer to calls in which paramedics respond to the location of an accident by helicopter, sometimes with assistance from land ambulances when, the sending or receiving facility does not have a heliport; these are the calls that immediately spring to mind when one thinks of air-ambulance service. Interfacility transfers refer to patient transfers between medical facilities. They can be emergent (42 percent), urgent (21 percent), or non-urgent (37 percent). Emergent and urgent interfacility transports require an immediate response. In this study, we focus on the non-urgent interfacility transports, which can be scheduled in advance.

Non-urgent interfacility transports use a subset of aircraft that are staffed and equipped for non-urgent requests. These aircraft are occasionally seconded for emergent transports in times of overwhelming demand, usually caused by patients with emergent, time-sensitive conditions; such situations result in disruptions in the non-urgent patient transfer schedules. Thus, there is a limited amount of sharing of aircraft between non-urgent and emergent/urgent calls. We ignore this sharing in our current work; however, we may consider it in future work, as we discuss in the *Impact and Next Steps* section.

Non-urgent interfacility transports (which we refer to as requests in the remainder of this paper) typically number 10–20 requests per day (although as many as 30 requests are sometimes made), and are almost exclusively handled by fixed-wing aircraft that are stationed around the province.

A key problem at the core of Ornge's mandate is to determine how to schedule and route available aircraft to handle these requests at minimal cost. The problem is complex principally because (1) aircraft can handle up to approximately four requests within a duty shift, (2) the requests can be handled in any order, subject to certain time restrictions, (3) although some aircraft can carry up to two patients at a time, others can carry only one, and (4) some patients (e.g., infectious patients) cannot be transported with another patient. Previously, Ornge flight planners manually determined the aircraft assignments and routes, based on their experience and personal practice.

Beginning in 2009, and in concert with several engineering master's and PhD students at Cornell University, we developed an optimization-based tool that determines the optimal assignment of requests to aircraft and the optimal routes to fly, taking into account a number of complicating side constraints. The tool uses an intuitive Excel interface and a C++ implementation, which assembles a set-partitioning formulation of the problem, invokes an integer programming (IP) solver, and returns the optimization results to the Excel tool. Flight planners now use this tool to develop plans for the next day's operations. They set up the problem in the Excel tool, obtain the solution, study the solution for its practicality and adherence to policies that are not easily encoded in a mathematical formulation, adjust certain parameters that relate to these policies, and solve the problem again. Although the first solution obtained is often the one implemented, a small number of iterations are generally required.

We tested the tool and assessed it for validity using retrospective Ornge data from randomly selected dates between July 2010 and February 2011 (MacDonald et al. 2011), and used the results to guide a live implementation. Our study showed that the tool would yield estimated decreases of 12 percent in flying hours, 13 percent in distance flown, and 16 percent in cost. In May 2011, we implemented the application in real time on a test basis, with a trial implementation on randomly selected days in June and August 2011. During the first eight weeks of fully implementing the tool, the application resulted in only a three percent decrease in average daily cost (MacDonald, Ahghari, et al. 2011, MacDonald, Walker, et al. 2012). However, the requests that were scheduled over that period had above-average origin-destination distances. Once one adjusts for the greater distances that needed to be flown, the actual cost savings appear to be on the order of eight percent. The residual difference between estimated and actual cost decreases likely resulted from deviations between optimized routings and actual routings. Without using the application, flight planners had changed the optimized routings to divert aircraft to pick up patients with emergency conditions. We are addressing this gap between fully optimized and actual use by incorporating a schedule repair option into a new version of the application.

The underlying mathematical problem Ornge faces is known in the operations research (OR) literature as a static dial-a-ride problem (Cordeau and Laporte 2007) with nonhomogeneous vehicles (i.e., aircraft) and multiple depots (i.e., aircraft bases). The problem is static in that requests are not scheduled as they are received, but rather collected and scheduled simultaneously. Cordeau and Laporte (2007), Cordeau et al. (2008), and Parragh et al. (2008) include comprehensive surveys on this problem and other vehicle routing problems with pickups and deliveries. Beaudry et al. (2010) provide an example of a challenging dynamic instance of a dial-a-ride problem in a healthcare setting.

A number of different approaches are available to solve static dial-a-ride problems. In addition to a variety of heuristics that can scale to very large instances, Cordeau and Laporte (2007) describe two exact methods based on IP formulations with decision variables that determine the arcs that vehicles transverse, and therefore the routes that they take. The formulations differ in that one determines vehicle-specific arcs; the other applies to homogenous vehicle fleets. Neither formulation can be applied in our setting because of the difficulty in effectively capturing our side constraints.

Ornge views the ability to capture complex constraints as vital, and the solution produced must require little or no additional modification by the end user. To capture side constraints, we adopt a set-partitioning formulation with enumerative pricing of feasible assignments. Parragh et al. (2012) give a similar formulation. In contrast to that work, we are able to solve our instances to optimality, perhaps because of the smaller scale of our problems. This may also be because we exploit the combinatorial structure of the problem to efficiently enumerate pricing and check the feasibility of routes. Espinoza et al. (2008a, 2008b) also considered the set-partitioning formulation as one of two formulations for an air-taxi problem. Even the linear programming relaxation of the set-partitioning problem could not be solved in these studies because of the size of the instances. Instead, their underlying formulation is a multicommodity network flow with side constraints, with special heuristic techniques for scaling to large problem instances.

Certain side constraints can be captured using dynamic column generation in time-space networks (Engineer et al. 2011); however, we can avoid the associated algorithmic complexity because our problem is small enough that we can enumerate and price all feasible columns in a manner that allows flight planners to use the application in real time. Therefore, unlike most (but not all) studies in the literature, we solve our instances to optimality.

To the best of our knowledge, the work we describe in this paper is the first to address such a problem within the air-medical transport industry. Commercial airlines use related strategies to develop flight schedules, routes, aircraft assignments, and crew schedules. Our application differs because the schedule differs each day.

We organized the remainder of this paper as follows: The *Problem* section defines the problem in detail, explaining the structure of requests, the costs associated with routes, and the primary complicating side constraints. The *Impact and Next Steps* section describes the results of an experiment to validate the tool and measure its performance relative to previous practice. It also covers the subsequent impact of this work at Ornge and outlines ongoing work. The appendix describes the set-partitioning problem formulation, and explains the recursion used to compute column costs. It also explains how we handle complicating side constraints.

Problem

Requests

On a given day, Ornge receives a number of requests that must be scheduled for the following day. Each request consists of the following information.

Origin. Airport at which the transport originates.

Destination. Airport at which the transport terminates.

Pickup after. The earliest time a pickup can be completed.

Drop off before. The latest time a dropoff can be completed.

Level of care. A representation of the needs of a patient, which can be primary, advanced, or critical care. These levels represent increasingly constraining requirements of personnel, scope of practice, and equipment needed on a transporting aircraft.

Stretchers. The number of stretchers required.

Escorts. The number of escorts required. Some patients, such as children, require escorts; this requirement may impact the number of patients an aircraft can transport.

Solitary. Whether a patient can be transported with other patients. This field also indicates if a patient is infectious; if so, extra time is required to disinfect a plane after the patient has been dropped off.

Maximum time. The maximum time that the patient can remain on the aircraft from pickup to drop-off time. We include this to partially consider patient convenience. Flight planners sometimes use it when an optimized plan calls for a patient to be kept on a plane for an inordinate amount of time.

Aircraft

Each aircraft that Ornge uses has a range of identifying characteristics, listed in the following.

Base. Where the plane begins and ends each route.

Level of care. The level of care that the plane can provide. A plane can carry a patient whose level of care is at most equal to the plane's level of care.

Stretchers. Number of stretchers (equivalently, patients) the plane can carry at any one time. A plane can usually carry two patients at a time; however, some planes can carry only one.

Escort capacity. The number of escorts that can be taken as a function of the number of stretcher-bound patients (one or two) on board.

Airspeed. The cruising speed of the aircraft.

Fuel cost. Fuel cost per hour flown.

Charter cost. The charter cost per hour flown. This is a charge for the use of an aircraft that is paid to the aircraft provider.

Advanced charter cost. The charter cost per hour flown when the plane is carrying advanced or critical patients.

Routes and Costs

Each plane that is used in one day flies a route that consists of multiple legs, where each leg consists of a takeoff and a landing with no intermediate stops. The route begins and ends at the plane's base. Calculating the cost of a route is complicated, but we can approximate it using the following three components.

Fuel cost. Charged on each leg based on the time spent in the air. We assume an average fuel price per liter across the province for all aircraft.

Charter cost. Charged on each leg based on the time spent in the air and the level of care provided.

Detention cost. Charged per hour spent waiting on the ground in excess of a certain ground-holding time.

Computing the fuel and charter costs requires that we know the time spent in the air, which depends primarily on the distance traveled. This distance depends on flight-path requirements and on avoiding dangerous weather systems. We ignore such complexities in our model, instead assuming that flights follow great-circle paths. Wind complicates the calculation. Even if wind speed and direction are constant over a route that begins and ends at the same place, the wind's effects do not cancel out. Instead, one must explicitly account for these effects. We do not have access to detailed location-and-time-specific wind information. Therefore, we assume a constant wind speed and direction across the province, and compute flying time accordingly.

Routes have two primary requirements.

1. A route cannot last longer than a specific government-mandated length of time. We assume this bound to be 12 hours. In practice, this time limit also limits the number of requests that a plane can handle in one day. We assume this limit to be four, in accordance with current Ornge practice. As we will see, this limit on the number of requests is vital to ensuring that the optimization problems we address are computationally tractable.

2. A route cannot be so long that a patient is kept on board for too long relative to the time the patient would have been on board if he or she had flown directly from origin to destination. Similarly, the number of legs on which a patient will travel during transport has a limit. Flight planners impose and adjust these soft constraints to varying degrees.

The problem involves selecting airplane routes and takeoff times for each leg on the routes that handle all requests at minimum total cost, while not violating any constraint discussed earlier.

Impact and Next Steps

A tool of this type needs careful validation. Therefore, we went through several phases in which we validated it, and adjusted requirements and functionality, before we reached the current implementation. A key step in this process was a retrospective study in which we randomly sampled 50 days between July 2010 and February 2011 (MacDonald et al. 2011). For each of the 50 days, we completed the following steps.

1. Retrieved the actual schedule flown and the information available when the schedule was constructed;

2. derived an optimized plan, as we describe herein;

3. calculated the total flying time, distance traveled, and number of legs for the optimized plan, assuming an empty plane;

4. calculated the cost of the optimized plan, based on data from financial records; and

5. used expert opinions to ensure the validity of the optimized plan.

Before discussing the results of the study, we acknowledge three limitations. First, the statistics for the plans the flight planners computed incorporated aviation factors, such as the need to fly

	Flight planners	Optimized plan	Difference
Flights	312	305	7
Hours	1,417	1,253	164
Distance (km)	481,381	417,156	64,225
Empty legs (%)	35.5	32.8	2.7

Table 1: Summary statistics show the differences in characteristics in plans that the flight planners developed with and without the optimization tool.

around weather systems, whereas those for the model did not. Second, the plans the flight planners computed included real-time disruptions that resulted from unscheduled urgent/emergent calls that required reorganization of the schedule; those for the optimized plan did not. Third, the costs modeled in the optimization tool only approximate the realized costs; however, we also used these estimated costs in evaluating the routing constructed by Ornge flight planners.

During the study period, 838 requests for transfers were received; the daily mean was 16.8 ± 5.4 requests. Based on the costs we computed in this study, we projected that the optimized plans would yield savings of approximately 16.5 percent. The summary statistics in Table 1 provide evidence for the substantial benefits of an optimized plan. The differences shown in the table, when computed as averages over the 50 days of the study, are statistically significant at the five percent level of an unpaired t-test.

Given the limitations discussed previously, we expected that the actual realized savings would be smaller than 16.5 percent; however, the benefits of the optimized plan are not in dispute, and Ornge management is extremely supportive of the tool. Ornge flight planners use the tool each day, except on days when the flight planners who are trained in its use are not available. However, because management recognizes the benefits of using the tool, the company is now training additional staff in its use.

Currently, we have two primary activities related to this project. First, schedule disruption is inevitable for several reasons, including weather and new calls that require urgent attention. Therefore, in conjunction with a team of Cornell master's of engineering students, we are adapting the tool to address the schedule-repair problem, which arises when a disruption during the day renders the planned schedule infeasible. One then attempts to reoptimize the remaining schedule, which is essentially a smaller version of the problem discussed previously. Second, the approximate costs used in the model have a number of limitations that lead to nontrivial discrepancies between the costs assumed in the model and those realized in practice; we are working to reduce these discrepancies.

Other research questions have arisen as a result of this study. For example, where should aircraft be based around the province to minimize expected daily cost? Carnes (2010) provides an approximation algorithm that addresses this question on a stylized model of air-ambulance operations. This question also relates to ambulance-location problems; Brotcorne et al. (2003) include a survey. However, those models are designed to address emergency responses one request at a time, rather than through routes that cover a set of requests. Can one answer this question with a single model that treats both non-urgent and urgent/emergent calls, as opposed to treating these classes of calls and their associated aircraft separately? Another issue involves designing a plan for the scheduled requests that considers the potential disruptions that could arise in executing the plan. The natural modeling framework is two-stage stochastic programming, although other methods might be appropriate. This issue relates to the dynamic dial-a-ride problem, as surveyed in Cordeau and Laporte (2007).

Appendix

We formulate the problem using a set-partitioning integer program. In this integer program, a binary variable is associated with each combination of a plane and a set of up to four requests. Given r requests and p planes, the number of variables n in the integer program is therefore

$$p\left(\binom{r}{1}+\binom{r}{2}+\binom{r}{3}+\binom{r}{4}\right).$$

In practice, this is an upper bound because some combinations of planes and sets of requests are infeasible; therefore, we omit such variables from the formulation. For example, for 30 requests and 40 airplanes, the formulation has 1,277,200 variables; however, this typically falls to approximately 750,000 in a number of test instances after we remove the infeasible combinations. Let the resulting variables be $x_1, x_2, \ldots, x_n$.

Define $A_{ij} = 1$ if request i is included in the jth column (i.e., the jth combination of both the plane and the set of requests it will serve), and 0 otherwise. Let $B_{ij} = 1$ if plane i is associated with the jth column, and 0 otherwise. Let $e^{(m)}$ denote the m-dimensional column vector with $e_i^{(m)} = 1$ for all i. Let c_j denote the cost of completing the set of requests in the jth column with the associated plane, as we discuss next. The optimization problem is therefore a set-partitioning problem of the form

$$\begin{aligned} \min_x \quad & c'x \\ \text{subject to} \quad & Ax = e^{(r)} \\ & Bx \le e^{(p)} \\ & x \text{ binary.} \end{aligned}$$

We must still discuss how we compute the cost coefficient vector c, which is the heart of the problem. We temporarily ignore many of the side constraints of the problem, such as the need to transport some patients alone, and we assume that all planes can carry two patients. We will discuss the additional complexities of the problem at the end of this appendix. For a more detailed and more general discussion of the recursion described next, see Carnes (2010, Chapter 3).

Fix a particular column, column j, in the formulation. The value c_j represents the minimal cost of handling the set of requests $\mathcal{R}_j$ using the plane associated with column j. Requests (i.e., patients) can be picked up and dropped off in any order, subject to the requirements listed in the *Problem* section. This complicates the computation of c_j. We enumerate all possible orderings of pickups and dropoffs of the requests $\mathcal{R}_j$, and choose the feasible sequence that minimizes costs. If no sequence is feasible (e.g., if the route would take too long and therefore violate the duty-day requirement), then we omit the column from the formulation.

If $\mathcal{R}_j$ consists of a single request, then c_j is the cost of flying (if necessary) from the plane's base to the request's origin, next to the destination, and then returning to the plane's base.

Suppose that $\mathcal{R}_j$ consists of $k > 1$ requests. We consider all $k!$ orderings of the requests in turn, where each ordering is defined by the order in which the requests are picked up. For each such sequence, assuming that planes can carry two patients, 3^{k-1} routes are possible, where a route consists of the sequence of pickups and dropoffs. To understand the reason, suppose that the first patient (the current patient) has already been picked up. The route can now take one of the following actions.

1. Drop off the current patient and pick up the next patient,
2. Pick up the next patient and drop off the current patient, or
3. Pick up the next patient and drop off the next patient.

In each of these three cases, we are left with one patient on board, so we can repeat the argument for each subsequent patient, obtaining a factor of three for each patient, except the first one.

This argument also provides a recursion that we use to enumerate all possible routes that can be used to complete the set of requests $\mathcal{R}_j$. We then compute the cost of each route and test it for adherence to the many requirements and side constraints; if it is feasible and cheaper than the best route found so far, we retain it as the incumbent route.

In performing the test for feasibility of a route, we consider the plane's characteristics (e.g., level of care, capacity). We enforce a minimum turnaround time at each stop on a route and also add time on the ground (e.g., to disinfect the aircraft after an infectious patient has been transported). These times are factored in when determining whether the time windows can be satisfied. We assume that the plane takes off at a time chosen by the flight planners, and we sequentially delay the takeoff times through the route to attempt to adhere to the time window constraints for requests. This will yield a feasible way to accommodate the time windows if one exists (Cordeau and Laporte 2007). However, it does not necessarily minimize the detention cost. One can apply an algorithm to minimize detention cost by delaying the initial takeoff (Savelsbergh 1992), but Ornge prefers to fix the initial takeoff times, partly for crew convenience. After the route's takeoff times have been fixed, computing the route's cost is then straightforward.

We complete the abovementioned steps for every combination of the set of requests and planes. Therefore, the total computation required to assemble the set-partitioning problem is approximately

$$p\sum_{k=1}^{4}\binom{r}{k}k!\,3^{k-1}, \tag{1}$$

where the upper bound (i.e., 4) in the summation arises because of the current Ornge practice of not assigning more than this number of requests to any aircraft. We impose it as a practical limit because of Ornge policies; however, it has the side benefit of ensuring that the recursion previously described completes quickly, because if this bound were much larger, then the computational effort, Equation (1), for larger numbers of requests r would be formidable.

One might consider using column generation to attempt to avoid the computational limitations we face in using complete enumeration of the columns. However, this is not possible because of the complicating side constraints that form an essential part of the formulation.

Acknowledgments

We thank master's of engineering students Jong-Yub Chae, Johannes Essl, Ying Xian, Ben Cheron, Anchal Dube, Lokesh Manohar, and Kevin Yu for their efforts in formulation and cost modeling, and PhD student Alex Fix for his programming assistance. Shane Henderson's work was partially supported by the National Science Foundation [Grants CMMI-0758441 and CMMI-1200315]. David Shmoys's work was partially supported by the National Science Foundation [Grants CCR-0635121, DMS-0732196, CCF-0832782, and CCF-1017688]. Tim Carnes's work was partially supported by the National Science Foundation [Grants CCR-0635121, DMS-0732196, and CCF-0832782].

References

Beaudry A, Laporte G, Melo T, Nickel S (2010) Dynamic transportation of patients in hospitals. *OR Spectrum* 32(1):77–107.

Brotcorne L, Laporte G, Semet F (2003) Ambulance location and relocation models. *Eur. J. Oper. Res.* 147(3):451–463.

Carnes TA (2010) Approximation algorithms via the primal-dual schema: Applications of the simple dual-ascent method to problems from logistics. PhD thesis, Operations Research and Information Engineering, Cornell University, Ithaca, NY.

Cordeau J-F, Laporte G (2007) The dial-a-ride problem: Models and algorithms. *Ann. Oper. Res.* 153(1):29–46.

Cordeau J-F, Laporte G, Ropke S (2008) Recent models and algorithms for one-to-one pickup and delivery problems. Golden B, Raghavan S, Wasil E, eds. *The Vehicle Routing Problem: Latest Advances and New Challenges* (Springer, New York), 327–357.

Engineer FG, Nemhauser GL, Savelsbergh MWP (2011) Dynamic programming-based column generation on time-expanded networks: Application to the dial-a-flight problem. *INFORMS J. Comput.* 23(1):105–119.

Espinoza D, Garcia R, Goycoolea M, Nemhauser GL, Savelsbergh MWP (2008a) Per-seat, on-demand air transportation part I: Problem description and an integer multicommodity flow model. *Transportation Sci.* 42(3):263–278.

Espinoza D, Garcia R, Goycoolea M, Nemhauser GL, Savelsbergh MWP (2008b) Per-seat, on-demand air transportation part II: Parallel local search. *Transportation Sci.* 42(3):279–291.

MacDonald RD, Ahghari M, Carnes TA, Henderson SG, Shmoys DB (2011) Use of a novel application to optimize aircraft utilization for non-urgent patient transfers. *Air Medical J.* 30(5):255.

MacDonald RD, Walker L, Ahghari M, Carnes TA, Henderson SG, Shmoys DB (2012) Prospective, real-time use of an optimization application for non-urgent patient transfers using fixed wing aircraft. *Air Medical J.* 31(5):230.

Parragh SN, Doerner KF, Hartl RF (2008) A survey on pickup and delivery problems. *J. Betriebswirtschaft* 58(2):81–117.

Parragh SN, Cordeau J-F, Doerner KF, Hartl RF (2012) Models and algorithms for the heterogenous dial-a-ride problem with driver-related constraints. *OR Spectrum* 34(3):593–633.

Savelsbergh MWP (1992) The vehicle routing problem with time windows: Minimizing route duration. *ORSA J. Comput.* 4(2):146–154.

Verification Letter

Julius Ueckermann, Vice President, Logistics, Ornge Transport Medicine, 5310 Explorer Drive, Mississauga, Ontario L4W-5H8, Canada writes:

"The above-named manuscript was recently submitted to your journal by Timothy Carnes et al. The manuscript describes a novel application that was developed to help optimize flight routing for non-emergency calls carried out at Ornge Transport Medicine in the Province of Ontario, Canada. This letter can serve as verification the application was implemented in our communication center, resulting in a decrease in flight hours, distance flown, empty legs flown, and cost. Based on our positive experience with the application, we intend to fully implement the latest version in our communication center in the near future."

Timothy A. Carnes is a senior analyst at Link Analytics where he develops optimization engines and designs solutions for a wide variety of problems in big data and business analytics. He received his BS in mathematics from Harvey Mudd College, and his MS and PhD in operations research from Cornell University. He served as a postdoctoral fellow at the Sloan School of Management at MIT, where he worked on healthcare applications on a collaborative project with Massachusetts General Hospital.

Shane G. Henderson is a professor in the School of Operations Research and Information Engineering at Cornell University. He received his PhD from Stanford University, and has held academic positions at the University of Michigan and the University of Auckland. His research interests include discrete-event simulation, simulation optimization, and emergency-services planning.

David B. Shmoys is a professor of operations research and information engineering as well as computer science at Cornell University. He obtained his PhD in computer science from the University of California at Berkeley, and held postdoctoral positions at MSRI in Berkeley and Harvard University, and a faculty position at MIT before joining the Cornell faculty. His research has focused on the design and analysis of efficient algorithms for discrete optimization problems, with applications including scheduling, inventory theory, computational biology, and most recently, computational sustainability.

Mahvareh Ahghari is a research coordinator at Ornge. She received her master's degree in industrial engineering from the University of Toronto. Her research and professional activities are in the areas of discrete-event simulation, resource-allocation optimization, process improvement, and performance analysis of emergency-medical services, call centers, and production systems.

Russell D. MacDonald is the medical director at Ornge Transport Medicine. He is also an associate professor in the Faculty of Medicine and holds an appointment at the Institute of Medical Sciences at the University of Toronto. He practices emergency medicine at Sunnybrook Health Sciences Center in Toronto, and is the co-director of the university's Emergency Medicine Fellowship Program. His expertise includes transport medicine, disaster preparedness, and international health. His academic interests include patient safety and adverse events in prehospital care, decision-making and optimization of resource allocation in transport medicine, and the relationship between public health and emergency services.