Interfaces
Vol. 37, No. 4, July–August 2007, pp. 334–341
ISSN 0092-2102 | EISSN 1526-551X | 07 | 3704 | 0334

DOI 10.1287/inte.1060.0265

Static and Dynamic Order Scheduling for Recreational Rental Vehicles at Tourism Holdings Limited

Andreas T. Ernst, Mark Horn, Mohan Krishnamoorthy
CSIRO Mathematical and Information Sciences, Canberra ACT 2601, Australia
{andreas.ernst@csiro.au, mark.horn@csiro.au, mohan.krishnamoorthy@csiro.au}

Philip Kilby
Research School of Information Sciences and Engineering, Australian National University, Canberra ACT 0200, Australia

Phil Degenhardt, Michael Moran
Platinum Ray Pty. Ltd., Hillside, Victoria 3037, Australia
{phil@platinumray.com, michael@platinumray.com}

Tourism Holdings Limited (THL), with 14 locations in Australia and New Zealand, operates a fleet of approximately 4,000 recreational rental vehicles of many types. It allocates vehicles to bookings centrally. If demand for a particular vehicle type at a location exceeds supply, THL may substitute vehicles of similar types or relocate vehicles from other locations to the location that needs the vehicles. The *static problem* that THL faces daily is to determine a vehicle schedule that minimizes the tangible and intangible costs of such substitutions and relocations. The *dynamic problem* is to determine—sometimes as the customer waits—whether a vehicle will be available to cover a potential booking and to incorporate that booking into the schedule. The Commonwealth Scientific and Industrial Research Organisation (CSIRO) has researched, developed, and supplied software, VASS and D-VASS, to solve the static and dynamic aspects of THL's schedule creation and maintenance. This paper describes the THL problem, the systems that CSIRO implemented, and how THL embedded these systems into its operations.

Key words: rental-scheduling applications; network-flow algorithms; assignment problems; revenue management.
History: This paper was refereed.

Scheduling at THL

Tourism Holdings Limited (THL) is a New-Zealand–based company with revenue in excess of NZ$100 million. It operates a fleet of more than 4,000 recreational vehicles (e.g., motor homes and camper vans) at 10 locations in Australia and four in New Zealand.

THL schedules these rental-vehicle operations centrally, on a day-to-day basis, using two criteria: Can it accept a given booking request? If so, which vehicle should it assign? The reputation, success, and even the company's viability depend on its ability to answer these questions.

Scientists from the Commonwealth Scientific and Industrial Research Organisation (CSIRO) in Australia worked with THL over the past five years to implement systems that can help answer these questions, and to assist THL's scheduling staff to obtain maximum efficiency from its vehicle fleet. This paper describes the business issues facing THL, and the systems CSIRO developed to address these issues.

In operations research terms, THL faced a scheduling problem. It needed to construct a schedule that assigns a suitable vehicle to each accepted booking. We call this the *static problem*. The question then arises, "Can a given booking be placed into the schedule, and if so, where?" This is the *dynamic problem*. In the remainder of this section, we give an overview of these two scheduling problems.

We can state the problem as follows. The characteristics of a booking are its start and end date, its start and end location, and the booked product. A product is related to a vehicle type but is more general. For example, a product may specify a four-berth motor home, which covers several different makes and models of vehicles. Products also differ with respect to the periods of hire and insurance conditions.

The main characteristics of a vehicle are its availability dates and its vehicle type. The vehicles include standard four-door saloon cars, four-wheel drive vehicles, and camper vans with two, four, or six berths. If we consider these and other distinguishing features, such as make and model, there are about 50 distinct vehicle types.

To allow the largest set of bookings to be accepted, we consider a large pool of vehicles and use the concepts of *substitution* and *relocation*, described below, to maximize the number of available vehicles.

Substitution. The product that a booking specifies maps to one or more suitable vehicle types. Some of the vehicle types are natural matches for the product, while some may be of higher quality or capacity. For example, it is permissible, but not desirable, to substitute a six-berth vehicle for a product that specifies a four-berth vehicle. A substitution cost table assigns costs to each combination of vehicle type and product.

Relocation. A relocation involves shifting a vehicle from one location to another to meet demand. A relocation table specifies the consequent costs for each vehicle type and each possible pair of locations between which relocations can be made. For a given vehicle and pair of locations, there may be several relocation scenarios that have different durations and cost. For example, two drivers might relocate a vehicle nonstop at high cost, or a single driver might perform the relocation more economically but more slowly, or the company might offer the relocation task to backpackers at nominal cost.

Other costs associated with scheduling of the rental fleet include the following.

Turnaround Cost. When a customer returns a vehicle to a THL depot, THL cleans it and makes it ready for the next booking. An allowance of one day is standard for this process; however, a faster turnaround is possible, at an additional cost.

Usage Cost. Because of higher maintenance costs, customers prefer some vehicles to others. For example, they prefer newer vehicles to older ones. This preference is reflected in a per-day usage cost for each day the customer uses the vehicle.

In addition to bookings, fleet management requires scheduling for two types of special events—*maintenance activities* and *disposals*. Each vehicle has its own routine maintenance schedule. As a maintenance activity approaches, a "booking" reserves the vehicle for this purpose during the specified dates. The process of removing a vehicle from the fleet is called a disposal. Disposals are concerned with the process of continuous renewal and replacement, i.e., adding new vehicles and decommissioning older vehicles.

We can treat both maintenance activities and disposals like normal bookings; they include a start and end location, and a start and end date (the end date is infinite for a disposal). Unlike a normal booking, however, a maintenance activity or disposal refers to a particular vehicle, rather than to a vehicle type. In addition, a booking (e.g., for a high-profile customer) may sometimes apply to a particular vehicle rather than to a product or vehicle type. These references to specific vehicles complicate the scheduling task, as we will see later.

A solution to the scheduling problem assigns a vehicle to each booking, maintenance, and disposal. An optimal solution must minimize the total substitution, relocation, and other costs, subject to the company's business rules. Standard scheduling constraints apply, such as a requirement that bookings assigned to the same vehicle do not overlap in time. We describe the scheduling model in more detail in the *Scheduling Model* section.

The dynamic scheduling problem poses the question, "Can a vehicle be provided to satisfy a booking for product X over period P from location A to location B?" It is necessary to answer such a question in the course of an operator's dialog with a customer, that is, in near real time.

The recreational focus of THL's business distinguishes it from "mainstream" rental-car operations in several important respects. First, THL's customers often book vehicles for holiday travel well in advance and for long durations. Second, unlike many franchise-based, car-rental companies, its locations operate in concert to maximize company profit, allowing more options with respect to vehicle relocation. Third, a large percentage of THL's market comprises bookings that start and end at different locations. While these characteristics do apply to some car-rental operations, they are not as important as they are in the recreational-market segment.

Relocation is an example of the flexibility available to THL that would not usually be available to a rental-car company. This flexibility is particularly important because of the large distances between major Australian cities. In addition, the availability of reasonably priced, one-way air travel makes one-way vehicle hire an attractive proposition. Seasonal patterns can also occur on a larger scale. For example, in summer months, when the monsoon season makes the tropical north unpleasant, the southern regions of Australia are popular with travelers. Imbalances of this kind lead to "migratory" patterns of demand, which THL must balance by relocations in the opposite direction. Relocations require unproductive work, which can be expensive for THL.

THL addresses the issue of long lead times by using an "active horizon" of 200 days. It accepts bookings beyond the horizon without planning their implementations, and addresses them when they enter the horizon. Each of the schedules for the New Zealand and Australian operations currently comprise several thousand bookings.

Another distinctive element is the operational setting, which poses both static and dynamic scheduling problems. The vehicle assignment and scheduling system (VASS) provides a solution to the static scheduling problem. It runs nightly and builds a schedule that assigns a vehicle to each booking, maintenance activity, and disposal.

For the dynamic scheduling problem, we have developed a software system, Dynamic VASS (D-VASS) to provide support for THL's front-end reservations system. Its main purpose is to answer availability queries by THL's reservation staff and to incorporate new bookings into the schedule. It also adjusts the schedule in response to operational contingencies, e.g., vehicle breakdowns and late returns.

In handling an availability query, D-VASS first ascertains if a vehicle is available to satisfy the request precisely. If not, it may suggest available alternatives, e.g., different dates or products similar to those requested. When a customer proceeds from an availability query to an actual booking, D-VASS updates the schedule by inserting the new booking. D-VASS always maintains a feasible, current schedule.

Because D-VASS is a dynamic system, it is important that schedule updates be accomplished with minimal disruption to other activities, e.g., booking enquiries. Therefore, D-VASS performs updates as quickly as possible, using a heuristic solution improvement method. Because a long sequence of suboptimal updates could result in a substantially inferior current schedule, the software invokes a background improvement process whenever there are no other current activities.

In this section, we have given a broad description of THL scheduling. We discuss some of the trade-offs that influence scheduling decisions in the *Scheduling Trade-Offs* section. In the *Scheduling in Practice* section, we show how THL uses the system. We present our model of the scheduling problem in the *Scheduling Model* section, and provide a mathematical outline in the appendix. We describe VASS in the *Vass—Static Scheduler* section, and D-VASS in the *D-Vass—Dynamic Scheduler* section. In the *Operating Environment and Performance* section, we discuss system performance. In the *Benefits* and *Conclusion* sections, we discuss the benefits that VASS and D-VASS have provided to THL, and indicate some prospects for further development. We will now describe how the project began.

Project History

CSIRO's involvement in the VASS project began in 1999. Britz:Australia, a vehicle-rental company, had developed in-house, vehicle-scheduling software that used a simple heuristic to assign bookings to vehicles in an incremental fashion. However, as Britz:Australia expanded, the size of the problem reached the limits of what this software could handle. After searching the market for a suitable software product to replace its existing system, Britz:Australia concluded that no existing commercial solution would meet its needs. It commissioned CSIRO to undertake a pilot study to demonstrate a capability to provide a suitable scheduling system. The pilot study involved the construction of a prototype, static-scheduling system, a series of tests to apply the prototype to several trial data sets, and a comparison of the results with those obtained using Britz:Australia's in-house system.

The pilot study's success led to commissioning VASS, in June 1999, for operations spanning the whole of Australia. A significant test of the robustness and scaleability of VASS came shortly afterwards when THL acquired Britz:Australia, thereby almost doubling the size of the fleet to be scheduled. THL encountered no significant problems in dealing with this increase in fleet size, and VASS has been in continuous use since that time.

THL was sufficiently pleased with the VASS performance that it subsequently commissioned the implementation of D-VASS as part of a complete rewrite of its computerized reservations, booking, and vehicle-management systems. It has used D-VASS in New Zealand since September 2002 and in Australia since November 2002. D-VASS receives every new booking and incorporates it into the schedule. It also aids THL in planning for unusual conditions and in responding to contingent situations, e.g., late returns and breakdowns. VASS and D-VASS work in tandem. D-VASS runs continuously during the day to maintain the schedule; VASS runs nightly to return the schedule to a near-optimal state.

Scheduling Trade-Offs

We provided an overview of THL scheduling in the *Scheduling at THL* section. Many subtle details contribute to the scheduling. At its heart, the scheduling task involves managing a series of multifaceted trade-offs. We can view the choice between a relocation and a substitution to satisfy a particular booking as a trade-off between customer satisfaction and cost minimization. However, the decision must also consider the probable demand for a substituted vehicle, which is usually of higher value than the requested vehicle. At the time of the initial scheduling decision, there may be no known competing demand for the substituted vehicle, but the opportunity cost of the decision is not negligible.

Another important question that follows a decision to plan a relocation is the amount of time allowed for the trip—the longer the period allowed for a relocation, the lower the cost. Therefore, there is an incentive to make the relocation decision early. On the other hand, other scheduling changes subsequent to the booking may eliminate the need for relocation altogether. Therefore, there is a counterincentive to push relocation decisions as far into the future as possible.

D-VASS handles these conflicts by using relocation in preference to substitution at the initial time of booking. Thus, it matches the booking with the product the customer has specified, as much as is possible. As the booking date nears, the system devotes increasing attention to the possibility of using a substitution instead of a relocation, considering the trade-off of potential customer dissatisfaction (because of the substitution) versus dollar cost (for the relocation).

Another factor is the possibility of delays, where a customer must wait beyond the expected time to pick up a vehicle. THL does not plan for delays at the time of a booking request; if no vehicle is available to meet the request, it does not allow the booking to proceed. However, subsequent events, such as breakdowns and late returns by other customers, may make delays difficult to avoid. THL handles this situation by assigning a high notional cost to any delay, thereby discouraging the scheduling algorithms from incorporating delays into the schedule.

A final trade-off is associated with *turnaround*—the process of cleaning and preparing a vehicle after its return. The system allows a standard time interval for a turnaround (depending on vehicle type), but it is possible to rush the turnaround time or perform it overnight, if required. It assigns a notional cost to any turnaround interval less than the standard interval. The system will consider short turnarounds to accommodate bookings when the schedule is tight.

Scheduling in Practice

In practice, we can view the scheduling system from three main perspectives. The first is that of THL's reservation staff who uses the system when responding to enquiries and booking customer requests. For this purpose, accessing D-VASS is by way of a front-end module that handles pricing, payment, and other ancillary details, as well as the core scheduling function.

The second group is scheduling specialists and other staff members responsible for managing THL's vehicle fleet. Staff members enter information about vehicle purchases, maintenance, decommissioning, and related activities into the database, and the computing system infrastructure makes this data available to D-VASS. This allows THL schedulers to focus their attention on less routine tasks, such as planning for busy periods and handling exceptional circumstances, such as allocating a particular vehicle to a high-profile customer. The scheduling staff must also identify extra vehicles that can be hurried through the maintenance or commissioning processes in response to breakdowns, late returns, and other contingencies.

A third point of view is that of the customer-service staff who handle dispatching and arrivals of vehicles. For example, when a customer arrives to commence a rental, the service personnel will query D-VASS to identify the vehicle allocated to the rental, and then notify the system when they have dispatched the vehicle.

Scheduling Model

VASS and D-VASS share a common model of the scheduling task. The time resolution is half a day, and the main task is to assign bookings to vehicles based on the following inputs.

Index Sets. Index sets represent the sets of all locations, products, vehicle types, and rental types relevant to the problem.

Vehicles. Each vehicle has a unique identifier, a vehicle type, and a specification of its availability, i.e., the location and date of the next time it will be available for rental.

Bookings. A set of bookings represents the work to be scheduled. Each booking specifies a start date and location, an end date and location, and a product identifier to indicate the range of vehicles that can satisfy the booking. A number representing the value or priority of the booking helps to determine which bookings to leave unscheduled when it is not feasible to schedule all bookings. For some specialized purposes, we may define a booking as a fixed assignment that requires a particular vehicle or one of a limited subset of vehicles.

Maintenance. A maintenance activity nominates a particular vehicle, with fixed start and end dates and a specific location.

Disposal. A disposal also nominates a particular vehicle, with fixed start and end dates and a specific location.

Relocation. A relocation defines all legal relocations in terms of the source and destination cities, time required, and cost. A given pair of cities may have several different relocations available.

Substitution. For each product, a substitution table lists the vehicle types that may feasibly be used. "Natural" vehicle types have zero cost, others have greater costs, and unacceptable substitutions have infinite cost.

Parameters. Additional parameters that influence the composition of a scheduling solution include the time horizon over which to schedule, an incentive factor to delay relocations as much as possible, the maximum allowable delay, a penalty cost for the delays, etc. Normally, these parameter values do not change between VASS or D-VASS runs. However, they give THL additional control over the schedules that the software produces (also see the appendix).

VASS—Static Scheduler

VASS uses a two-phase approach that is based on a network-flow formulation of the problem. The first phase involves the solution of a master problem using a commercial, integer-linear-programming code. In the absence of a requirement that specific vehicles must be allocated to bookings, the schedule that VASS produces is optimal with respect to the modeled costs. That is, the master problem exploits the fact that all vehicles of a particular type are essentially identical. The requirement that some specific vehicles be allocated to particular activities (e.g., maintenance, disposal) disrupts this structure. This aspect of the problem is similar to "tail assignment" in aircraft scheduling, i.e., assigning individual aircraft to previously determined routes (Gopalan and Talluri 1998). VASS uses a heuristic to assign specific vehicles to activities based on the solution of the master problem.

D-VASS—Dynamic Scheduler

The main task for D-VASS is to insert a new activity (e.g., a booking or a maintenance activity) into the schedule. In the case of an availability query, the insertion is temporary. However, in other cases, D-VASS accepts the revision and updates the current schedule accordingly. To maintain acceptable online performance, D-VASS ignores some of the details that VASS addresses. In this respect, the most important simplification is that D-VASS preserves all substitutions that VASS has already programmed. That is, it does not try to change the vehicle type assigned to an activity.

D-VASS applies a heuristic procedure to an assignment formulation of the scheduling problem. In summary, it regards the schedule as a time-ordered sequence of activities. This sequence implies a successor or "assignment" relationship between each pair of activities that a given vehicle carries out and a graph in which each node is an assignment. The task then is to find the shortest path connecting these nodes for each vehicle in the fleet. The method that D-VASS uses to incorporate a new activity into the schedule is based on the successive shortest-path method for solving the assignment problem as Engquist (1982) describes. Kilby (2004) provides a more complete description of the algorithm.

A key advantage of this approach is that it can also be used in a "continuous optimization" scheme. Improvements in the schedule appear as negative cycles in the graph used to solve the assignment problem. The continuous improvement operator simply looks for such negative cycles and implements them.

Operating Environment and Performance

We configured the scheduling system for a Microsoft Windows™ platform. It is currently driven by a 1.5 GHz Pentium processor. We wrote both VASS and D-VASS in C++ and implemented them as dynamic link libraries. VASS uses the CPLEX™ linear-programming library to solve the network-flow problem.

Access to VASS is via D-VASS, which in turn is invoked by a COM+ server in middleware that we wrote for this purpose. The use of COM+ facilitates distributed access to the scheduling system from the front-end clients that we mentioned in the *Scheduling in Practice* section. For efficiency, the system retains schedule information in memory between client calls rather than rereading it from the database each time it invokes the server. Implementing this arrangement required special measures to protect the integrity of the database and to ensure robust operation in other respects.

An automatic process activates VASS each evening when booking activity is low. VASS typically requires 5 to 20 minutes to produce a schedule. D-VASS is available for use on an essentially continuous basis, and includes a "background improvement" procedure that it runs when it is not attending to incoming calls (see the *D-VASS—Dynamic Scheduler* section above). To avoid excessive delays, each standard (non-VASS) call to D-VASS is subject to a five-second time limit. This provides a comfortable ceiling, even for booking queries, indicating an adequate level of online performance.

Benefits

The pilot study (see the *Project History* section) showed the potential for significant savings in the quality of schedules that VASS created, compared with the schedules that the existing system produced; VASS reduced substitution costs by over 30 percent and reduced relocations by two percent. At the same time, it decreased the number of bookings left unassigned in the schedule by 20 percent.

In addition to cost savings, the flexibility that the system delivered has been an important benefit. THL's scheduling staff formerly spent many hours each day developing a schedule manually. An internally developed scheduling system saved some time. The introduction of VASS delivered an even more reliable schedule, reducing the amount of time that staff spent in developing a schedule even further. The schedulers have gradually developed a trust in the system and its ability to produce effective solutions. They are now concentrating their efforts on handling "exceptions" to normal operations, e.g., late returns, breakdowns, and dispatching of incorrect vehicles. As a result, the schedulers handle such situations much more effectively than they did in the past. The system has also enabled THL to keep its business running with much less dependence on a scheduling staff that is on call seven days per week.

THL's operations have changed markedly since it commissioned VASS and D-VASS. The company has changed hands. World events, which have led to a downturn in tourism, have affected its business. Customer profiles have changed—some market segments have shrunk, while others have grown.

All these changes, e.g., the shifts in demand patterns that THL has encountered since September 11, 2001, have the potential to affect THL's bottom line.

Although changes in fleet composition cannot always match changes in demand exactly, VASS and D-VASS have assisted the company in taking an adaptive approach. In particular, the scheduling system has facilitated phasing out old types of vehicles and introducing new ones through fine-grained control over substitution policy and by extracting maximum utility from the fleet.

The form of the parameter tables allows THL's scheduling staff to respond quickly to unusual events. For example, a "special event," e.g., a music festival, can cause an unexpected spike in demand; specialized substitution or relocation rules applied during the event can handle these spikes. The result is better customer service—whenever possible, a customer obtains the vehicle he or she desires or obtains a close substitute.

The new systems have yielded savings in direct operating costs of about two percent per annum. There is some evidence that VASS and D-VASS have had significant additional impact on THL's bottom line, e.g., through an effective increase in fleet capacity. These impacts are difficult to quantify because of the change of company ownership that occurred shortly after the VASS introduction. However, it is clear that the new systems enable THL to accept many more bookings, more rapidly, and with increased fleet utilization.

VASS and D-VASS constitute an excellent platform for implementing predictive techniques. In particular, THL (or a company with similar operations) could build expected demand into its systems by modeling future vehicle dispositions based on predicted as well as currently known demand. Revenue management, the science of balancing pricing against expected demand to achieve the best possible result, is a related predictive approach. McGill and van Ryzin (1999) provide a review of its concepts and techniques. For instance, the general flow of bookings for THL in New Zealand is north to south; therefore, bringing vehicles north again requires many relocations. A reasonable question is, "What is a suitable discount to offer customers to provide them with an incentive to travel against the flow?" Revenue-management techniques may provide the key.

Conclusions

CSIRO has produced a software system incorporating two main modules—VASS and D-VASS—that assist a large, recreational-vehicle-rental company in its daily operations. The software handles every availability query and incorporates every new booking into the vehicle schedule. It answers difficult questions that are central to the efficient running of the company. These questions include:

—What relocations does the company need to do today to meet demand?

—When should the company use a substitution rather than a relocation?

—Does the company have a vehicle available to cover a requested booking?

In the future, THL and CSIRO may examine ways of incorporating revenue-management concepts into VASS and D-VASS as a basis for more informed pricing and discounting decisions. CSIRO and THL would like to develop their relationship further to examine these and other questions.

Appendix. Problem Formulation

In this appendix, we show how the problem referred to in the *Scheduling Model* section is specified mathematically. For definitions of the terms we used, please refer to the main text.

We will use the following terminology to define a rental fleet and its associated activities.

V—set of vehicles in the rental fleet.

B—set of current rental bookings to be scheduled.

F—set of activities with a fixed vehicle required, such as maintenances and preassigned bookings.

P—set of products available.

T—set of vehicle types.

V_t—set of vehicles of type $t \in T$.

F_v—set of fixed activities for vehicle $v \in V$.

We include a dummy activity representing the current activity for vehicle v in F_v. This activity implies the location and time the vehicle will be available next.

The following terms define other scheduling elements and conditions.

sl_a—location where activity a is to commence.

el_a—location where activity a is to terminate.

st_a—time when activity a is to commence.

et_a—time when activity a is to terminate.

v_b—profit obtainable from booking b.

reloc-t(v, m, n)—time to relocate vehicle v from m to n.
Time is zero if $m = n$.

reloc-c(v, m, n)—cost to relocate vehicle v from m to n. Cost is zero if $m = n$.

subst-c(v, b)—substitution cost of using vehicle v for booking b. This is determined by the product of b and the vehicle type of v.

turn-t(v)—turnaround time needed to prepare vehicle v for a rental.

delay-c(b, t)—cost of delaying booking b by a time t. *delay-c*$(b, 0) = 0$.

D_a—maximum delay for activity a.

d_a—actual delay for activity a.

We define a schedule by the list of activities $A_v \subseteq (B \cup F)$ assigned to each vehicle $v \in V$:

A_v—list of activities scheduled for vehicle v.

A—set of all activities assigned to vehicles $= \bigcup_{v\in V} A_v$.

U—unassigned bookings $= B \backslash A$.

a_v^i—ith activity scheduled for vehicle v.

With the provision for delays, the timing of a rental or other activity performed on a given vehicle v may be delayed beyond its stipulated commencement time $st_{a_v^i}$ by tardiness in completion of the activity's predecessor a_v^{i-1}. The delay is zero for $i = 1$, while for $i > 1$ it is defined as follows:

$$st'_{a_v^i} = et_{a_v^{i-1}} + turn\text{-}t(v) + reloc\text{-}c(v, el_{a_v^{i-1}}, sl_{a_v^i}), \quad (1)$$

hence,

$$d_a = \max(st'_a - st_a, 0). \quad (2)$$

Four main constraints apply to the construction of a schedule A. In order, these are:

— each vehicle must perform the fixed activities preassigned to it,

— each activity can be scheduled on at most one vehicle,

— the activities assigned to each vehicle must be arranged in temporal sequence, and

— an upper limit is placed on any delay in commencing an activity.

$$F_v \subseteq A_v \quad \forall\, v \in V, \quad (3)$$

$$A_v \cap A_w = \varnothing \quad \forall\, v \neq w \in V, \quad (4)$$

$$st_{a_v^{i-1}} \leq st_{a_v^i} \quad \forall\, v \in V, \forall i\colon i > 1,\, a_v^i \in A_v, \quad (5)$$

$$d_a \leq D_a \quad \forall\, a \in A. \quad (6)$$

We define the following cost components for a schedule A:

Profit foregone from rentals that are not included in the current schedule:

$$CP = \sum_{b\in U} v_b. \quad (7)$$

Relocation costs:

$$CR = \sum_{v\in V} \sum_{i:\, i>1,\, a_v^i \in A_v} reloc\text{-}c(v, el_{a_v^{i-1}}, sl_{a_v^i}). \quad (8)$$

Substitution costs:

$$CS = \sum_{v\in V} \sum_{a\in A_v \cap B} subst\text{-}c(v, a). \quad (9)$$

Costs due to delays beyond requested commencement times:

$$CD = \sum_{v\in V} \sum_{i:\, i>1,\, a_v^i \in A_v} delay\text{-}c(a_v^i, d_{a_v^i}). \quad (10)$$

The problem, then, is to construct schedule A to minimize the weighted sum of CP, CR, CS, and CD, subject to conditions (3)–(6).

References

Engquist, M. 1982. A successive shortest path algorithm for the assignment problem. *Infor* **20**(4) 370–384.

Gopalan, R., K. T. Talluri. 1998. Mathematical models in aircraft schedule planning: A survey. *Ann. Oper. Res.* **76** 155–185.

Kilby, P. 2004. An online schedule update algorithm for vehicle fleets. CMIS Technical Report 04/17, CSIRO Mathematical and Information Sciences, Canberra, Australia.

McGill, J. L., G. van Ryzin. 1999. Revenue management: Research overview and prospects. *Transportation Sci.* **33**(2) 233–256.