Interfaces
Vol. 35, No. 1, January–February 2005, pp. 49–60
ISSN 0092-2102 | EISSN 1526-551X | 05 | 3501 | 0049

DOI 10.1287/inte.1040.0113


# Bombardier Flexjet Significantly Improves Its Fractional Aircraft Ownership Operations

Richard Hicks, Richard Madrid, Chris Milligan, Robert Pruneau, Mike Kanaley
Bombardier Aerospace, 3400 Waterview Parkway, Suite 400, Richardson, Texas 75080
{jrhicksjr61@cs.com, rmadrid@flexjet.com, cmilligan@flexjet.com,
robert.pruneau@businessacft.bombardier.com, mkanaley@flexjet.com}

Yvan Dumas, Benoit Lacroix
AD OPT Technologies Inc., 3535, chemin Queen-Mary, Suite 650, Montréal, Quebec, Canada H3V 1H8
{ydumas@ad-opt.com, blacroix@ad-opt.com}

Jacques Desrosiers
HEC Montréal and GERAD, 3000, chemin de la Côte-Sainte-Catherine, Montréal, Quebec, Canada H3T 2A7,
jacques.desrosiers@hec.ca

François Soumis
École Polytechnique de Montréal and GERAD, C.P. 6079, Succursale Centre-ville, Montréal, Quebec, Canada H3C 3A7,
francois.soumis@gerad.ca

The fractional aircraft market is the fastest growing segment of the business aircraft industry. A fractional aircraft operation is complex—essentially an unscheduled airline in a constantly changing environment. Bombardier Flexjet implemented a comprehensive three-module optimization system to simultaneously maximize its use of aircraft, crews, and facilities. AD OPT Technologies designed the modules, using the GENCOL optimizer developed at GERAD, which employs a column-generation approach to decompose large-scale mixed-integer nonlinear programming problems. Since inception, the project has generated savings in excess of $54 million with projected additional savings of $27 million annually, primarily by lowering crew levels (20 percent), aircraft inventory (40 percent), and supplemental charter aircraft usage (five percent) while increasing aircraft utilization (10 percent). The quality of customer service has remained consistently high, with significant reduction in supply.

*Key words*: transportation: models, network; programming: nonlinear.

The fractional aircraft market is the fastest growing sector of the business aircraft industry, with a 32 percent compounded annual growth rate from December 1998 to December 2003. The potential market is estimated to be over 120,000 owners. Fractional ownership programs allow individuals or companies to buy shares in a business jet at a fraction of the full cost of ownership. Owners select the aircraft type appropriate to their needs, determine the number of hours per year they fly, and purchase shares of business jets starting at a share size of 1/16th (equal to 50 hours' flying per year). Firms in the fractional business guarantee them 24-hour access to a fleet with as little as four hours notice. Owners can fly into 10 times the number of airports available to commercial flights. They can also fly anywhere in the world on as little as 48 hours notice. They pay predictable monthly management and usage fees, and the operators manage aircraft maintenance, flight crews, hangars, fuel, and insurance on their behalf. A fractional aircraft ownership operation is complex; it is essentially an unscheduled airline in a constantly changing environment with large daily variations in demand. Operators must match owners' requests to aircraft and crews in such a way that they operate flights safely, satisfy the owners, and maximize profits. They must also schedule crew members and their travel, maintain the aircraft, comply with regulations, cope with bad weather, and provide flight-related services, such as ground transportation and catering.

Bombardier entered the fractional aircraft business in mid-1995 with its Flexjet program. It scheduled aircraft and crews manually for the first year when its operations were small. It could no longer handle scheduling manually when the fleet size reached about 25 aircraft. From 1996 to 2000, it optimized its operations using an integer multicommodity flow model with a CPLEX engine. By late 1999, with about 80 aircraft to schedule, it realized the model's limitations in considering constraints. It could obtain only rudimentary aircraft routing in reasonable time. Considering even a small subset of crew working rules and maintenance restrictions took too long. Therefore operators scheduled, swapped, and tracked daily crews manually in a suboptimal way. They also produced monthly crew schedules manually using a bid-line approach. Indeed, its inadequate scheduling system caused Flexjet to maintain extra business jets and crews, costing it on the order of several hundred dollars per live flight hour.

In early 2000, Bombardier contracted with AD OPT Technologies, a leading provider of workforce-management systems, and GERAD, an operations research center operated jointly by four Montreal universities. Together we began a three-phase scheduling and optimization enhancement project to maximize Bombardier Flexjet's use of its aircraft, crews, and facilities. Flexjet could benefit from features of the phases as we completed them, gain experience with them, and seek needed revisions as we proceeded with the later phases. The requirements for the new system were an integrated aircraft and crew optimizer with cold and warm starts and a repair mode for use on small portions of the scheduling horizon; a preferential bidding system for optimizing the monthly work, training, and vacation schedule for crew members; and a personalized crew-assignment module for reoptimization during daily operations. The overall goal was to provide the highest possible number of flawlessly executed flights every day, including days with heavy demand, and to reduce the total operational costs (including the costs for crews, core aircraft, empty positioning flights, and supplemental charter flights). We also had to integrate the new software with the Bombardier Flexjet graphical user interface (GUI) to make it widely available to our personnel.

## An Unscheduled Airline

The overall integrated system had to allocate resources and produce flight and crew assignments that made efficient use of them. We use the following terminology throughout the paper: A *crew* is a group of people who operate an aircraft together, that is, a combination of captain, first officer, and flight attendant (when required). A *duty* is a combination of flight legs, connections, and ground waiting times that form a legal workday for a crew member. A *rotation* (or *pairing*), starting and ending at the crew member home base, is a sequence of duties separated by rest periods. A *monthly block* is a sequence of rotations separated by days off. The crews and aircraft have to cover two types of obligations: point-to-point movements demanded by owners and point-in-time obligations for which aircraft must be at certain locations at certain times for maintenance, contract signings, stand-by duty, and so forth. Schedulers must consider many factors not present for scheduled airlines, including the following: Owners request approximately 80 percent of flights 48 hours or more prior to departure and 20 percent with as little lead time as four hours prior to departure and also ask for changes in 30 percent of the flights in the system at least once within 48 hours prior to departure. Additions to the fleet of supplemental (charter) aircraft with the proper cost structure and operating characteristics cover days when demand exceeds Flexjet's internal supply. Upgrades and downgrades of aircraft can reduce costs or improve travel service to owners, for example, to those who refer prospective owners to the firm. Airport limitations prohibit aircraft too large for an airport's runways or require extra time buffers between flights because of congestion. Departure time windows around the requested times can be adjusted globally or by individual flight leg. Multiple locations for maintenance vary in maximum capacities and costs. Multiple locations for signing contracts with owners purchasing shares can be tax friendly or satisfy owner preferences. Swaps among Flexjet crew members may require repositioning by commercial airlines or Flexjet aircraft.

Managing swaps is important in maximizing the utilization of crews. Hard swaps (coming off duty) and soft ones (changing aircraft) are possible. The crew-swap aspect is complex, given that Flexjet must schedule crew movements around owners' requests, which change constantly; frequently a change in one leg request can alter the schedule of three or four aircraft.

The scheduled airline's problem is completely different. For example, major airlines decide the flight schedule several months in advance while fractional operators know their flight requests only several hours beforehand. Furthermore, the scheduled airlines' optimization process for using aircraft and crews follows from a sequential division of the problem, often resulting in suboptimal decisions. Scheduled airlines (1) build the flight schedule; (2) assign the fleet, that is, match aircraft types to flights; (3) route the aircraft, taking into account required maintenance capacity; (4) solve the crew-pairing problem, that is, construct crew itineraries over several days separately for each type of aircraft; (5) prepare monthly schedules for crew members; and, (6) in case of disruptions, recover aircraft and crews sequentially. Some researchers (Cordeau et al. 2001, Mainville Cohn and Barnhart 2003) have tried to integrate aircraft-routing and crew-scheduling modules for commercial airlines, but they have not achieved the strong integration we reached for the Flexjet fractional aircraft business.

Literature on the fractional aircraft business is fairly recent. Keskinocak and Tayur (1998) and Martin et al. (2003) present similar multicommodity network-based models that create single schedules for both the aircraft and the crew, that is, they impose the crew rules on top of the aircraft rules. Such models are clearly suboptimal because they artificially impose crew rest periods on the aircraft. They allow for only one crew duty per aircraft per day and for swapping only at the end of a duty, and they do not optimize aircraft itineraries to select good alternative crew swaps. In our project, we modified the usual sequential process and strongly integrated many of its components. The system for the Flexjet fractional aircraft business has three modules:

(1) Monthly crew assignments based on days off, preferences, and forecasted owners' requests;

(2) For all aircraft, simultaneous integration of flight scheduling, aircraft-type assignment, aircraft routing, and crew scheduling;

(3) Each day, assignment of crew members to working periods on specific aircraft.

In reengineering the decision process, we started with the monthly assignments, a long-term planning problem. In the second module, we combined the first four stages for commercial airlines. The second module also replaces the sequential recovery stages for commercial airlines when it is used in the repair mode as it recovers aircraft and crews simultaneously. The third module links the results of the previous modules and provides personalized schedules.

## System Architecture

This system architecture is based on the three integrated optimization modules developed by AD OPT (Figure 1). Altitude is a registered trademark of AD OPT Technologies Inc.

Altitude VBS (vacation bidding system) and Altitude PBS (preferential bidding system) determine annual leaves and monthly schedules (lines), that is, they assign crew members duty days and vacations over four-week periods, trading off coverage requirements (determined by forecasting models) with crew members' preferences while considering such events as crew training. These long-term planning systems enable Flexjet planners to precisely match crewing levels to expected demand in terms of total crews on each day and crews on duty through each overnight period. Matching crew levels to our daily forecasted targets is important in optimizing aircraft and crews with TOSCA.

TOSCA (the optimizer for scheduling corporate aircraft) is the heart of the integrated system, the central short-term optimization module. It allows Flexjet planners to look out over a one- to three-day solving horizon and plan the movement of all types of aircraft in the fleet. TOSCA simultaneously assigns flights to aircraft and optimizes crew duty periods. It considers both hard constraints (based on legality or physical limitations, such as Federal Aviation Administration (FAA) regulations, company rules, or aircraft performance limits) and soft constraints (based on cost trade-offs). For each aircraft, TOSCA determines whether to use no crew (unused aircraft), one crew duty period or two crew duty periods, and when and where to swap crew members. TOSCA plans crew duty periods considering the crews already assigned and the number of crews available according to

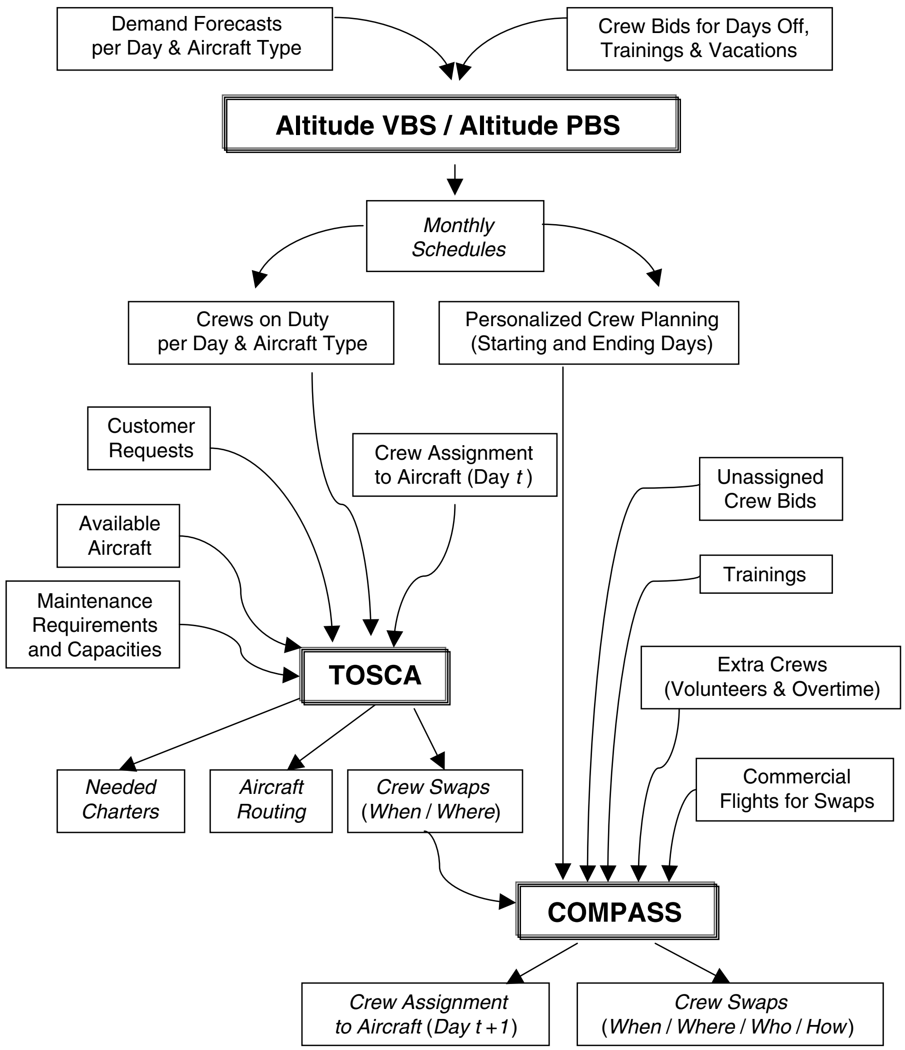


**Figure 1: The system architecture presents an integrated optimization system that provides routing of aircraft and crews. The vacation bidding (VBS) and the preferential bidding (PBS) systems determine annual leaves and monthly schedules, based on long-term forecasts. Once in operation on a day-to-day basis, TOSCA simultaneously optimizes aircraft and crew schedules over a one- to three-day horizon to cover owners' flight requests, and COMPASS assigns crew members to specific aircraft and minimizes crew-swap costs.**

monthly lines determined by Altitude PBS and Altitude VBS. The TOSCA optimizer can operate in three modes: a cold start in less than 30 minutes, a warm start in less than 10 minutes, and a repair mode in two or three minutes or less.

COMPASS (crew optimization module for paired assignment) assigns individual crew members each day to the crew duty periods that TOSCA has planned for specific aircraft. It considers any constraints that prohibit certain crew members from working on certain flights. It also matches first officers to captains based on rotation lengths and originating bases to minimize crew-swap costs determined by the TOSCA module. It also matches trainers to trainees.

## A Brief History

The existing Flexjet optimizer with a CPLEX solver (the legacy system) was unable to handle the growing complexity of the scheduling problem. It could not consider many aspects of the aircraft-routing model in a reasonable run time, in particular, many crew-duty rules and the crew-swap requirements.

Whatever capability the legacy system had to consider, crew constraints did not matter in practice because Flexjet planners need solutions within 30 minutes to insure that changes during the run time do not render the solutions useless. The legacy system strung legs together to minimize deadhead (empty) trips without considering which aircraft had crews that need to come off duty. Schedulers had to review the schedule aircraft by aircraft to find aircraft that required crew swaps to see if by chance they could devise a way to swap crews somewhere in the string of flight legs the optimizer assigned. Because on average a third of the fleet requires crew swaps each day, the schedulers frequently had to change assignments manually.

The legacy system did not consider all the basic crew rules: schedulers had to check all aircraft to make sure the optimizer did not schedule overlong duty days or violate overnight crew rules. Schedulers ran the optimizer each morning to schedule flights and crews for the following day, and they made all the changes after that—both changes to comply with owner requests and changes to make up for the shortcomings of their system. The legacy system was not solving the full aircraft-crew problem. Bombardier Flexjet needed a better system.

Bombardier representatives first met with Professors Jacques Desrosiers and François Soumis from GERAD in March 2000 to explain the fractional ownership business and the firm's optimization problem. Yvan Dumas and Benoit Lacroix from AD OPT Technologies later visited the Flexjet owner operations center in April and conducted a series of meetings in May and June. In these meetings, the Flexjet managers listed the new features they wanted, and the AD OPT representatives confirmed the technical feasibility. From the meetings, the Flexjet managers determined that enough new features would be technically feasible to justify the expense of the project. In August, they authorized the project. They awarded the following development contracts to AD OPT: simultaneous aircraft-and-crew-scheduling optimization software in September 2000; preferential-bidding-system and vacation-bidding-system software in March 2001; and daily-crew-assignment optimization software in November 2002.

## GENCOL: The Mathematics Behind the Solution

AD OPT Technologies designed the three modules using the GENCOL optimizer Desrosiers and Soumis developed at GERAD beginning in 1981. It employs a column-generation approach to decompose large-scale mixed-integer nonlinear programming problems. GENCOL is at the core of several scheduling systems that are widely used in urban, air, and rail transportation. For example, it is used to schedule bus drivers (Desrochers and Soumis 1989) in such cities as Tokyo, Singapore, New York, and Chicago. It is also used to solve scheduling problems by such airlines as UPS, FedEx, United, Air Canada, and Qantas.

The model used in GENCOL is based on an integer multicommodity network flow problem. It uses resource variables and constraints to model the various rules governing aircraft itineraries and crew members' schedules. The solution approach considers a master problem and subproblems alternatively, that is, it relies on the Dantzig-Wolfe (1960) decomposition method that has been adapted to nonlinear integer programs (Desaulniers et al. 1998). GENCOL mathematically decomposes the network-based model and places all the nonlinearities within constrained shortest-path problems solved by powerful dynamic programming algorithms (Desrosiers et al. 1995) to generate the columns sent to the master problem. Columns represent aircraft itineraries and crew-member schedules. At the coordinating level, GENCOL solves the master problem with a linear program that selects the best mix of columns. GENCOL generates only a small subset of the possible billions of these aircraft itineraries and crew-member schedules. A customized branch-and-bound search tree embeds this two-level process to obtain integer solutions. Modeling all constraints and cost factors for each module was the main research effort in this project.

## TOSCA

During the first phase of the project (April 2000 through June 2001), we developed TOSCA. This central module is a short-term optimizer that simultaneously routes and schedules aircraft and crew members. It also optimizes the flight schedule given departure-time windows. The optimizer objective function takes into account costs and penalties, including the following: Hourly operating costs for Flexjet and chartered aircraft, which influence their assignment to different leg types; Fleet assignment costs per flight leg, which differ by aircraft type and influence our hierarchies of assignment choices for the various legs; The charter-assignment penalty that accounts for the losses in quality of service resulting from placing owners on non-Flexjet aircraft; The early or late departure penalties that discourage our shifting departure times from those customers request; The crew-swap costs that include travel cost per mile and cost per minute in crew time; The cost for flight legs left uncovered in the solution that must be covered by outside aircraft; Re-assignment penalties when using the optimizer in warm start and repair modes that insure we focus only on near-term changes to the schedule that save enough money to justify using the limited capacity for setup rework; Threshold penalties on rest and duty length to equalize crew workloads by applying more restrictive limits to crew members that exceed a heavy workload score.

TOSCA uses column generation to separate the problem constraints into two groups: local and global constraints. The local constraints concern the validity of aircraft itineraries and crew members' schedules and are incorporated in the subproblems. These constraints govern aircraft maintenance times, aircraft connection times, time windows on flight departures, matching customers' requests that should be grouped together, aircraft initial and final positions, aircraft capabilities for international flights, runway limitations, crew rules, and crew members initial and final conditions. Global constraints are imposed on the master problem, a set-partitioning model. They specify that we must select one itinerary per aircraft, that we must fly flight-leg requests exactly once, that we must not exceed maintenance-location capacities and crew availabilities per fleet/base/day, and that we specify maximum aircraft upgrades and downgrades. In summary, TOSCA determines the type of aircraft to assign to each flight request, the aircraft positioning and itineraries, and when and where to perform maintenance on each aircraft. It optimizes flight starting times within time windows, when and where to swap crews, and from which base to get the new crew. It also determines crew rest periods and the need for aircraft double-duty crew assignments.

## Altitude VBS and Altitude PBS

In the second phase (April 2001 through April 2002), we added optimization for the yearly and monthly crew planning. Altitude VBS assigns annual leave to the crew members, based on available vacation slots and crew preferences. It also supports monthly rebids when new vacation slots become available to the crew members. For regular airlines, Altitude PBS assigns detailed schedules composed of sequences of flights and layovers (pairings) for pilots and flight attendants typically over a one-month period (Gamache et al. 1998). Fractional airlines requirements differ from commercial airlines; Altitude PBS assigns work days over a one-month period. These will be filled with specific owners' requested flights on a daily basis during operations. Bombardier Flexjet runs Altitude VBS and Altitude PBS prior to the other modules, rather than at the end as regular airlines do.

In general, constructing monthly schedules for crew members consists of assigning days on and off around annual leaves and training periods, maximizing satisfaction by considering weighted bids that reflect individual preferences. The assignments also favor matching the schedules of first officers and flight attendants to captains' schedules to reduce the costs of crew swaps while honoring seniority restrictions. Schedulers must never construct a maximum-score schedule for a particular crew member at the expense of a more senior employee. Altitude PBS determines all crew members' monthly schedules (lines), that is, the days they have on and off, and it lists those willing to work extra days. While satisfying all the crew legality rules in the subproblems (one for each crew member), the system ensures coverage of flight demand with crews, maintains continuity in crew rotations, and considers crew members' preferences. In GENCOL, each column of the master problem is a possible crew schedule that may or may not be selected, while the right-hand side is the forecasted flight demand for each day.

During this second phase, Bombardier Flexjet's internal business transformation and systems group developed an enhanced crew-scheduling-and-tracking system, FlexSchedule, which it integrated with Flexjet's other aircraft-scheduling software so the firm could track all flight crews in real time. It can examine daily crew availability while maintaining full compliance with FAA regulations for fractional operators. FlexSchedule has aircraft, crew, dispatch, owner, and maintenance displays that helped planners to accept and use all the new features of the system.

## COMPASS

In the third and final phase (June 2002 through December 2003), we further integrated the FlexSchedule GUI with the GENCOL optimization software suite and implemented the COMPASS module. We further improved TOSCA and fully integrated it with COMPASS, the third module. COMPASS assigns individual crew members to the crew duty periods TOSCA planned for specific aircraft while considering crew-transportation and overtime costs, reassignment penalties in the warm-start and repair modes, and assignment of overtime opportunities based on crew preferences. It also maximizes crew-member pairings by home base and rotation length (the consecutive days a crew member works), matches crew members' rotation ends with aircraft-maintenance schedules, matches trainees to trainers, and considers crew members' preferences to fly with or to avoid flying with others. COMPASS obeys any rules that prohibit a crew member from working on a certain flight leg. We again split the constraints into two categories. Global constraints on the master problem concern flights to be covered once, cockpit-qualification, crew-composition, and crew-swap limitations at particular airports or within regions. Local constraints on the subproblems restrict admissible crew members' assignments because of such factors as qualifications, reserve crew members on standby, connection times, crew domiciles, possible extra days, and crew repositioning. COMPASS also considers matching first officers to captains to reduce crew-swap costs in the TOSCA module. The COMPASS solution method is essentially a column-generation approach except that the size of the actual problems allows explicit generation of all possible variables (columns) of the master problem.

## Resistance to Change

As we implemented each new piece of software, we encountered resistance to change on the part of the schedulers who would be using the system and the flight crews.

The schedulers were skeptical about new systems and new processes; we overcame their objections with plenty of hands-on training and by implementing the feature-rich software in digestible portions. We trained schedulers individually. In the early stages of implementation, we monitored the system closely and asked the schedulers for feedback frequently. We gathered feedback from them at every stage of the project. For example, between flight segments, an aircraft requires a certain amount of turnaround time before it can be ready for the next flight leg. Based on schedulers' input, we designed the system to prevent aircraft originating outside the US from being assigned to land at noncustoms airports and also restricted these aircraft from assignments from early flights in the US. Gathering schedulers' comments and implementing their suggestions into an adaptive solution methodology was critical to a successful implementation. Flexjet's operating environment (fleet size, owner mix, owner flying trends, and information systems) changes constantly, and we frequently review the system to align it with the current operating environment. Flexjet relies on the schedulers to see new trends and to work with the developers enthusiastically to keep the system current. By empowering the schedulers, we enable them to drive the development of the system. Without their routine feedback, our results could quickly diverge from optimal. Because the costs for empty positioning and charter hours add up to thousands of dollars, any divergence from the optimal operating strategy can cause major losses. We seek feedback not just to avoid bad outcomes; we also want to foster an environment in which we can leverage operations research to exploit opportunities opened up by new strategies and features. We try to notice when users implement a strategy manually, and then we try to adjust the system to replicate that strategy automatically so that it is applied to the entire schedule and not just when a user notices the opportunity. Flexjet can thus use the system to disseminate good strategies among users on all shifts and move toward the goal of responding consistently to all situations regardless of time of day or scheduler.

The new automated bidding process replaced the prior manual bid-line system and thus affected the flight crews. We reassured them with communications and training. We provided complete instructions online for the pilots and included recurrent instruction in the crews' regularly scheduled training. We provided individual instruction in person and over the phone. Since implementation, we have continuously solicited suggestions for new features from crew members. We made a number of changes after implementation, for example, by allowing them to bid for lengths of time on or off without regard to particular dates and by allowing them to require four days off after a seven-day rotation. AD OPT developed the system with a great deal of crew input, and Flexjet continues to improve the system based on crew feedback.

## New Features

Because of limitations in the legacy system, Flexjet could not consider several elemental concepts. With the new integrated system, Flexjet now considers variables it previously could not.

The new optimizer considers all crew duty rules and produces itineraries that honor established constraints. Also, the system can assign multiple crews to one aircraft when it is optimal for it to travel longer than one crew's duty day. The new optimizer also considers the overnight rest requirements, preventing assignments that violate overnight rules. This means the setup is much less elaborate, providing schedulers the opportunity to fix changes by using optimizer warm starts rather than performing only manual changes.

A very important new feature controls crew swaps. Each day, about a third of crew members are in transit to or from their assigned aircraft. Using crew efficiently is the make-or-break factor in determining whether Flexjet can operate effectively with a reduced crew-per-aircraft ratio. The system establishes inputs that define the swap time windows for each swap location for crew members from each crew base. It determines which aircraft require swaps by considering the crew-rotation schedule (another input the legacy system lacked). In assigning strings of legs, it considers the swaps required; schedulers do not have to go through the schedule and make a lot of manual changes to accommodate swaps the leg assigner did not consider. It can also swap crews when it is optimal to do so in the middle of a crew's rotation.

The system uses the departure-time window in two ways depending on the type of day. Flexjet designates some days as peak days; on those days by contract it can move owners' departure times forward or back by up to three hours without penalty. The TOSCA optimizer can consider the six-hour time window in five-minute increments and still produce an answer in less than 30 minutes. This means that, for each leg, it considers departure on time, with a five-minute delay, with a 10-minute delay, with a 15-minute delay, and so forth up to three hours. For example, Flexjet's owners may request 190 live legs for a given day. For each of those legs the optimizer considers 36 possible delayed departures and 36 possible early departures, increasing the size of the problem. The legacy system could not evaluate delays in a timely manner, so for peak days, schedulers had to plan any use of the three-hour window manually. In effect, on Flexjet's busiest days when optimization was most important, schedulers would effectively build the schedule manually and suboptimally. This new feature benefits the owners as a group: by delaying some flight departures, fewer owners need undesirable supplemental charter aircraft on peak days.

Flexjet also uses the departure window on nonpeak days to allow up to 15-minute delays on any flights. During its years of operation, Flexjet discovered it could delay requested departure times slightly over the course of the day without noticeably affecting each owner's departure time. By allowing 15-minute delays, the optimizer can place three or four legs a day on more cost-efficient aircraft, usually by reducing the time required for positioning empty planes. By delaying three or four legs slightly, Flexjet could avoid using charter aircraft for these, which would in turn have bumped many legs off their assigned company aircraft onto the chartered aircraft. By allowing slight delays, the system can assign approximately 20 to 25 percent of the schedule more effectively. Otherwise, schedulers would have to make delays manually and would miss opportunities because they had too little time to scour the complex schedule for opportunities.

Another new feature that schedulers use every day is the multiple-location maintenance. Each day Flexjet typically has 10 to 15 maintenance requirements to be performed overnight in any of four locations. For each aircraft that requires maintenance, the legacy system could specify only one particular location. Maintenance personnel, who had no training in scheduling, decided which location to use without regard to an aircraft's location or the demand available to move it around. That would be 10 to 15 predetermined single-location requirements a day. With the multiple-location feature in TOSCA, the maintenance personnel can input maintenance events and the optimizer considers all maintenance bases. Thus, the new system replaces those 10 to 15 single-location events with a pool of multiple-location events. The optimizer can use the pool of flights available to get aircraft to maintenance locations, planning for the group of aircraft instead of sending each aircraft to a particular location. Associated with maintenance location are maintenance capacity constraints that limit the number of each aircraft type the location can handle per night. With the legacy system, maintenance personnel scheduled alternate locations manually when the ones entered initially were infeasible.

## Financial and Qualitative Impact

The integrated system and the new optimization strategies have greatly improved operations. Using the new system, Flexjet can alter the operating cost structure. It has improved cost measures that influence its scheduling decisions. Total operations costs dropped. Crew members per aircraft sold decreased by more than 20 percent, core fleet-support aircraft by more than 40 percent, and the percentage in supplemental charter aircraft used by more than five percent, while density (live hours per available aircraft per day) has increased by more than 10 percent. Most important, the key metric for customer service, percentage of flights executed flawlessly, has remained consistently high.

We based our figures on savings, a comparison of actual yearly performance until December 2003 with conditions before implementation of TOSCA, which was completed in July 2001. All figures were normalized by using a common reference point (for example, we based crew-level savings on the ratio of crew members employed to aircraft operated). The savings calculated are conservative due to the phased implementation.

For the three phases of the project, the total costs for the work AD OPT and Bombardier did can be estimated at a few million dollars. We estimated that the return is over $54 million in savings. We estimated savings at $27 million annually for approximately 80,000 flight hours, and we expect annual savings to grow as the fleet flight hours increase with the growth of the company. If the owner base doubles, the annual savings will double. Using linear proportionality in making our estimate is conservative given that the legacy system most likely would break down completely long before Bombardier Flexjet doubles its base. Given the projected rapid growth of the fractional airline business, our system should prove to be a major enabler of growth for Flexjet.

The project has produced many important qualitative benefits. For instance, crew members are getting better schedules than they would have by bidding on set lines. Flexjet needed to reduce its crewing levels compared to demand to compete in the fractional market. By reducing crew levels, the new system increased crew utilization (Figure 2). We had a few months' data under the legacy system after Flexjet reduced its crew and could see a very strong deterioration in the schedules the system was giving crew members. For example, senior crew members who usually received their first or second choice of schedules were beginning to receive their fifth or sixth choice. We implemented the preferential bidding system before most crew members noticed that trend.

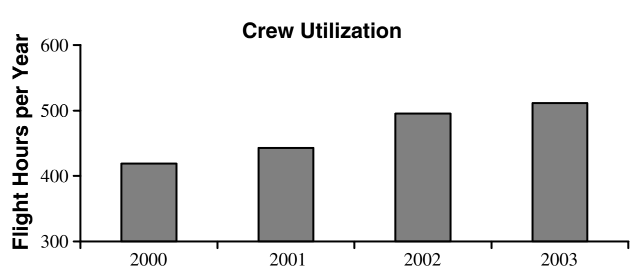


**Figure 2: Crew utilization increased by 22 percent between 2000 and 2003.**

The system helped Flexjet to avoid a major deterioration in crew morale.

The new system increased efficiency within the aircraft-scheduling group by 20 percent and within the crew-planning group by 50 percent. The Flexjet operations center could focus on other critical areas. For example, the long-term scheduler could use the time saved to coordinate crew movement with crew schedulers, to coordinate maintenance events with maintenance planners, and to answer requests made by owner services.

Flexjet's ability to match assets to varying daily demand has improved, increasing their utilization. As a result, Flexjet has expanded its product offerings and increased the flexibility of its owner program. For example, a one-sixteenth share in a Challenger 604 now gives owners 60 hours per year compared with the industry standard of 50 hours. This was offered without an increase in the purchase price or monthly management fees.

## Conclusions

In 2000, Flexjet was burdened with an uncompetitive cost structure. It was unlikely to be profitable at the asset levels it maintained. Its scheduling system was at capacity. Flexjet determined that it needed to overhaul its asset structure by scheduling more efficiently. Our three-year project with AD OPT Technologies helped Flexjet to achieve operational profitability and improve its scheduling process. Flexjet can now create new product offerings to continuously adapt to the rapidly changing fractional airline business. Flexjet has recognized OR/MS as a key strategic tool.

Flexjet has reduced its operations costs, decreasing crew levels by more than 20 percent, its fleet of aircraft by more than 40 percent, the charter aircraft percentage by more than five percent, and increasing density (live hours per available aircraft per day) by more than 10 percent. Its service level has remained high. Its total savings in 30 months exceed $50 million. Annual savings are about $27 million for 80,000 flight hours. Because these savings are proportional to fleet size, they should grow as the company grows. We designed the system so that computer speed increases faster than the problem complexity (appendix), so Flexjet will not face capacity limitations for many years.

## Appendix. A Flavor of the Powerful Decomposition Method Used in GENCOL

Desaulniers et al. (1998) describe the mathematics behind GENCOL. To save space, we describe only a simplified version of the decomposition method we use in the TOSCA model. Let $N$ be a set of requested flight legs and $K$ a heterogeneous fleet of aircraft located anywhere. For each aircraft, construct a graph $G^k = (N^k, A^k)$ over the planning horizon in which $N^k \subseteq N$ is the set of flight legs that can be served by aircraft $k$ and arc $(i, j) \in A^k \subset N^k \times N^k$ exists if it is possible to go from the end of flight $i$ to the beginning of flight $j$. For leg $i \in N^k$, $f_i^k$ is the flying time. We assume in $N^k$ an origin location $o^k$ and a fictitious destination $d^k$, and we let $m^k$ be the remaining flying time at $o^k$ before the next maintenance, that is, aircraft $k$ flying capacity. For each arc $(i, j) \in A^k$, let $c_{ij}^k$ be the cost incurred (including flight cost at node $i$ on the appropriate type of aircraft $k$) and $t_{ij}^k$ be the time to go from the beginning of flight $i$ to the beginning of flight $j$. The aim of the aircraft-routing problem is to cover each flight leg once while minimizing the total cost without exceeding the aircraft flying capacity.

Given that binary flow variables $x_{ij}^k$ take value 1 if aircraft $k$ uses arc $(i, j) \in A^k$, and 0 otherwise, a linear-programming formulation is given by the following multicommodity network-flow model:

$$\min \sum_{k\in K} \sum_{(i,j)\in A^k} c_{ij}^k x_{ij}^k \tag{1}$$

$$\text{subject to} \sum_{k\in K} \sum_{j:(i,j)\in A^k} x_{ij}^k = 1, \quad i \in N, \tag{2}$$

$$\sum_{j:(i,j)\in A^k} x_{ij}^k - \sum_{j:(j,i)\in A^k} x_{ji}^k = \Delta_i, \quad k \in K,\ i \in N^k, \tag{3}$$

$$\sum_{j:(i,j)\in A^k} t_{ij}^k x_{ij}^k \le m^k, \quad k \in K, \tag{4}$$

$$x_{ij}^k \in X^k, \quad k \in K,\ (i,j) \in A^k, \tag{5}$$

$$x_{ij}^k \text{ binary}, \quad k \in K,\ (i,j) \in A^k. \tag{6}$$

The objective function in (1) minimizes the total flying cost of the aircraft fleet. For each flight leg, a certain aircraft has to cover the leg in (2). Constraint sets (3) and (6) define a path structure for aircraft $k$ from its origin $o^k$ to its destination $d^k$: there is conservation of aircraft flow type when $\Delta_i = 0$ for each flight leg $i \in N^k$, while $\Delta_{o^k} = 1$ at the origin of aircraft $k$, and $\Delta_{d^k} = -1$ at the corresponding destination. In (4), the total flying time of any aircraft must not exceed its flying capacity. This is one of the restrictions that are local to the itinerary of aircraft $k$. Others are symbolized in (5) by sets $X^k$, $k \in K$. For small instances and disregarding (5), the above formulation can be solved directly as an integer linear program. However, constraints (3) through (6) and the objective function are separable by aircraft $k$ and form a block angular structure suitable for Dantzig-Wolfe decomposition. This suggests a natural decomposition by defining a constrained path subproblem for each aircraft while keeping the constraints on covering once each flight leg in the master, a set-partitioning problem.

Consider now imposing time window $a_i^k \le T_i^k \le b_i^k$ on the starting time $T_i^k$ of flight $i \in N^k$ possibly served by aircraft $k \in K$. These additional resource variables and constraints are just an example of the constraint sets that can be defined in (5). These are expressed in GENCOL in the form

$$\left(\sum_{j:(i,j)\in A^k} x_{ij}^k\right) a_i^k \le T_i^k \le \left(\sum_{j:(i,j)\in A^k} x_{ij}^k\right) b_i^k, \quad k \in K,\ i \in N^k, \quad (7)$$

$$x_{ij}^k (T_i^k + t_{ij}^k - T_j^k) \le 0, \quad k \in K,\ (i,j) \in A^k. \quad (8)$$

Time window $a_i^k \le T_i^k \le b_i^k$ in (7) is accessible only if flight $i$ is covered by aircraft $k$, that is, if $\sum_{j:(i,j)\in A^k} x_{ij}^k = 1$. Because flow variables are binary, a nonlinear constraint in (8) expresses that if $x_{ij}^k = 1$, then $T_j^k \ge T_i^k + t_{ij}^k$, that is, waiting is possible and the starting time of flight $j$ is at least $T_i^k + t_{ij}^k$. Combining with (7), $T_j^k$ is therefore computed as $\max\{a_j^k,\ T_i^k + t_{ij}^k\}$, a piecewise linear function, as long as it does not exceed $b_j^k$. Constraints in (8) can be linearized, but there are as many constraints as there are arcs in all the networks. For large instances, this prohibits the use of a direct solution approach by linear programming and branch and bound.

GENCOL performs mathematical decomposition and uses dynamic programming to solve constrained nonlinear shortest-path problems as integer programs. This structure is represented by (3) through (8) together with a modified cost function to account for the dual information provided by the master problem for the pricing of new columns. The generated path columns fill in the set-partitioning problem used to select the right mix of aircraft itineraries. Additional global constraints can be imposed on this master problem. These indicate, for example, that maintenance-location capacities and crew availabilities per fleet/base/day are not exceeded and that the maximum aircraft upgrade and downgrade assignments are specified. The size of model (1) through (8) follows: The number of variables is $O(n^2)$ times the number of aircraft, where $n$ is the number of flights. As the number of aircraft is a function of the number of flights, say $n/3$, the number of variables in that model is $O(n^3)$. The number of constraints is also $O(n^3)$, mainly depending on constraint set (8). Observe that the number of nonzero elements in the constraint matrix grows as $n^5$, and this is the reason why the direct solution of the linearized version of this model becomes harder to solve as the problem size increases. For the column-generation process used in GENCOL, we can separate the size of the master problem and that of the subproblems. For the master problem, a linear program, the number of constraints is approximately $n$, the number of flights to cover, while the number of generated path columns that are kept is only about three to five times the number of flights. Therefore, the number of nonzero elements only grows as $n^2$. For the subproblems, each of which represents a specific aircraft, most of the underlying network information is common. Hence, there is only one large network from which the system chooses the relevant data when solving a specific subproblem. The size of this network is $O(n^2)$. One has to add various linear and nonlinear functions to take into account all the rules to satisfy, but the number of nonzero elements also grows as $O(n^2)$. If we assume that the number of flights doubles in three years (at a growth rate of 25 percent), the computer speed increases at most by a factor of eight, while the number of nonzero elements of the linearized version of model (1) through (8) increases by a factor of 32. In GENCOL, the number of nonzero elements in the master problem and the subproblems increases only by a factor of four. In conclusion, the benefit of the decomposition approach is that it allows for substantial increase in problem size.

**Acknowledgments**
We gratefully acknowledge the contributions and support of the following people at Bombardier Flexjet: Tracy Ashberry, Ken Capps, Ryan Ford, Monica Fowler, Elton Gore, Arron Miller, Jason Miller, Cathy Santeiu, Robert Webb, the entire System Operations team, and the entire Flight Operations team. We also acknowledge the important contributions and efforts of John Allarie, vice-president of sales at AD OPT; François Lessard from GERAD and Fabrice Lavier from AD OPT, who developed TOSCA and COMPASS; Éric Gélinas, Norbert Lingaya, Gérald Marquis, and Yves Siko from AD OPT, who developed Altitude PBS and Altitude VBS. We extend our sincere thanks to Nicolas Bélanger and Daniel Villeneuve from AD OPT for their contributions to GENCOL and to our collaborators Guy Desaulniers and Michel Gamache at GERAD. Finally, special thanks go to Brian T. Denton, our coach for the Edelman presentation, for his valuable advice.

## References

Cordeau, Jean-François, François Soumis, Jacques Desrosiers. 2001. Simultaneous assignment of locomotives and cars to passenger trains. *Oper. Res.* **49**(4) 531–548.

Dantzig, George B., Philip Wolfe. 1960. Decomposition principle for linear programs. *Oper. Res.* **8**(1) 101–111.

Desaulniers, Guy, Jacques Desrosiers, Irina Ioachim, Marius M. Solomon, François Soumis, Daniel Villeneuve. 1998. A unified framework for deterministic time constrained vehicle routing and crew scheduling problems. T. Crainic, G. Laporte, eds. *Fleet Management and Logistics*. Kluwer, Norwell, MA, 57–93.

Desrochers, Martin, François Soumis. 1989. A column generation approach to the urban transit crew scheduling problem. *Transportation Sci.* **23**(1) 1–13.

Desrosiers, Jacques, Yvan Dumas, Marius M. Solomon, François Soumis. 1995. Time constrained routing and scheduling. M. O. Ball, T. L. Magnanti, C. L. Monma, G. L. Nemhauser, eds. *Handbooks in Operations Research and Management Science*, Vol. 8. *Network Routing*. Elsevier Science B.V., Amsterdam, the Netherlands, 35–139.

Gamache, Michel, François Soumis, Daniel Villeneuve, Jacques Desrosiers, Éric Gélinas. 1998. The preferential bidding system at Air Canada. *Transportation Sci.* **32**(3) 246–255.

Keskinocak, Pinar, Sridhar Tayur. 1998. Scheduling of time-shared jet aircraft. *Transportation Sci.* **32**(3) 277–294.

Mainville Cohn, Amy, Cynthia Barnhart. 2003. Improving crew scheduling by incorporating key maintenance routing decisions. *Oper. Res.* **51**(3) 387–396.

Martin, Chris, David Jones, Pinar Keskinocak. 2003. Optimizing on-demand aircraft schedules for fractional aircraft operators. *Interfaces* **33**(5) 22–35.