

UPS Optimizes Its Air Network

Andrew P. Armacost

Department of Management, US Air Force Academy, 2354 Fairchild Drive, Suite 6J100, Colorado Springs, Colorado 80920,
armacost@alum.mit.edu

Cynthia Barnhart

Department of Civil and Environmental Engineering, Massachusetts Institute of Technology, 77 Massachusetts Avenue,
Room 1-229, Cambridge, Massachusetts 02139, cbarnhar@mit.edu

Keith A. Ware, Alysia M. Wilson

UPS Operations Research Group, 1400 Hurstbourne Parkway, Louisville, Kentucky 40223
{air2kaw@ups.com, air2amm@ups.com}

Operations research specialists at UPS and the Massachusetts Institute of Technology (MIT) created a system to optimize the design of service networks for delivering express packages. The system simultaneously determines aircraft routes, fleet assignments, and package routings to ensure overnight delivery at minimal cost. It has become central to the UPS planning process, fundamentally transforming the process and the underlying planning assumptions. Planners now use the system's solutions and insights to improve plans. UPS management credits the system with identifying operational changes that have saved over \$87 million between 2000 and 2002. Anticipated future savings are expected to be in the hundreds of millions of dollars.

Key words: industries: transportation, shipping; transportation: models, network.

UPS is the world's leading package-delivery company, carrying an average of more than 13 million packages daily to nearly 8 million customers in over 200 countries and territories. With 256 aircraft and 78 more on order, UPS Airlines, a wholly owned subsidiary of UPS, is the 11th largest commercial airline in the world and the ninth largest in the United States. The airline is the key infrastructure that enables UPS to provide such expedited delivery services as same-day SonicAir, next day air, and second day air. The airline's 2002 next-day air operations produced over \$5 billion in revenue and averaged more than 1.1 million package deliveries a night.

To support next-day-air network planning and operations, our team from UPS and MIT developed and implemented Volume, Location, and Aircraft Network Optimizer (VOLCANO), an optimization-based planning system that is transforming the planning and business processes within UPS Airlines. This innovative modeling and algorithmic approach to an intractable network-design problem has been a tremendous success within the airline and the academic community. UPS credits the system with saving over \$87 million between its acceptance by UPS planners in late 2000 and the end of 2002, and senior UPS managers estimate saving approximately \$189 million in operating costs over the next decade. Even more important are the potential savings in aircraft acquisitions, as long-range planners use VOLCANO

to support fleet composition and acquisition. In addition, VOLCANO is cutting-edge technology consistently cited by senior UPS leaders, including Tom Weidemeyer, the president of UPS Airlines and the chief operating officer of UPS.

Using \$100 borrowed from a friend, 19-year-old Jim Casey founded the American Messenger Service in Seattle, Washington in 1907. Casey, with his brother George and a group of friends, provided private messenger and delivery services. The company grew to include retail package delivery, and in 1919, it adopted the name United Parcel Service (UPS). Ten years later, UPS began offering limited air-delivery service between a number of West Coast cities. By the mid 1950s, UPS was offering two-day delivery service between major cities from coast to coast. It was not until 1975, however, that the company received federal authorization to deliver packages between all 50 states. Today, UPS continues to grow globally, delivering more than 13 million packages and documents around the world every day.

To deliver packages efficiently to all states, UPS increasingly depended on air transportation. UPS's second-day-air service, originally called Blue Label Air, was available in all 50 states by 1978. In response to increasing demand for even more rapid delivery, UPS offered next-day-air service among the 48 contiguous states and the territory of Puerto Rico in 1985, and to Alaska, Hawaii, and international addresses soon thereafter. In the late 1980s, with permission

from the Federal Aviation Administration to operate its own airline, the company formed UPS Airlines in less than a year. Prior to establishing the airline, UPS leased space in the cargo holds of regularly scheduled commercial flights and leased entire planes, with the leasing company providing the aircraft and crews. Today, UPS Airlines operates its fleet of 256 aircraft and supplements its fleet with leased aircraft for the peak retail shipping period from Thanksgiving through December.

To ensure timely next-day-air service for its domestic customers, UPS Airlines operates a network with one all-points hub in Louisville, Kentucky; six regional hubs in Columbia, South Carolina, Dallas, Texas, Hartford, Connecticut, Ontario, California, Philadelphia, Pennsylvania, and Rockford, Illinois; and more than 100 airports. The airline operates nine aircraft types flying more than 1,800 flight legs daily. To provide efficient next-day-air service, the airline sets aggressive deadlines for package delivery. It establishes deadlines for customers to drop off packages, typically between 6:00 PM and 8:00 PM. UPS must then deliver the packages by 10:30 AM the following business day or, in the case of next-day-air early AM service, by 8:00 AM. To provide service within such tight time frames the airline must operate its airports, hubs, and aircraft efficiently.

In the UPS network (Figure 1), trucks carry packages to ground centers and from ground centers to airports. At the airports, workers load packages onto aircraft. Each aircraft transports its packages directly to an air hub or stops at one intermediate airport to pick up additional packages. At the hub, workers unload all packages from the inbound aircraft, sort them, and load them onto the outbound delivery aircraft. The aircraft performing pickup and delivery are the same, with each aircraft positioned at the air hub until it is fully loaded for its delivery route. The aircraft then fly to at most two airports, where workers

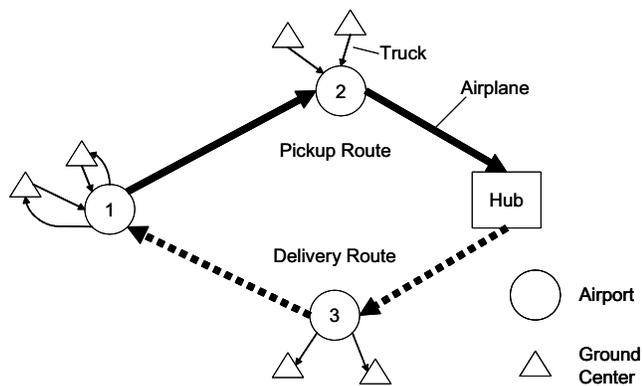


Figure 1: In this example of a next-day-air network, a two-leg pickup route runs from airport 1 to airport 2 to the hub and a two-leg delivery route runs from the hub to airport 3 to airport 1.

transfer the packages to trucks that carry them to a local ground center. At the ground centers, workers sort the packages again and load them on smaller trucks for delivery to their destinations. Thus, a next-day package typically travels from its origin to a ground center by truck, from the ground center to an airport by truck, from the airport to a hub by plane, from the hub to an airport by plane, from the airport to a ground center by truck, and from the ground center to its destination by truck.

The aircraft routes are carefully planned. Each type of aircraft has operating characteristics that determine which routes it can fly, including maximum flying range, effective speed, restrictions on the locations at which it can land, and cargo capacity. The number of airports a plane can visit on a pickup route or a delivery route, not including the hub, is two. Manual planners typically create route networks in which delivery routes simply reverse the order of a corresponding pickup route, while such mirror-image routes might not effectively use aircraft capacity. In addition, some of the aircraft used during the night for next-day-air deliveries are used during the day for second-day-air deliveries, so an important interface exists between the next-day-air network and the second-day-air network.

Manual Planning Prompts Research Effort

Planning an air network of this size takes many groups of planners and analysts. Within UPS Airlines' industrial engineering division, three groups of planners work on different next-day-air planning issues that share a common element: to determine the most cost-effective set of airplane routes and package movements that meet customer demands and timing requirements. First, the long-range planners develop network plans for two to 10 years in the future, specifying the capacity of the network by selecting its operating locations (airports), air hubs, and the mixture of different aircraft types needed to move projected volume. Network planners work on the plans for the current year, adjusting the existing plans to accommodate actual or anticipated changes in the system and enabling the airline to meet current demand. Peak planners focus on developing the network plan to enable operations during the busy retail season in November and December. Outside the industrial engineering group, financial analysts make similar plans regarding major capital outlays, for example, for expanding ground facilities and acquiring new aircraft. All planning groups participated in specifying requirements for the VOLCANO system and developing, implementing, and validating it. These groups focus on the next-day-air network; a separate corps of

planners work on similar planning issues for second-day air.

Prior to VOLCANO, expert planners could take up to nine months to manually produce a single plan. This process did not include analyzing sensitivity of the plan to key data, such as package volume levels. Planners were forced to plan to a single, conservative set of package volume projections. And the problem was continuing to grow. Rodger McLaughlin, a UPS planning manager with 17 years of experience, says, "The size and complexity of our operational system, and the amount of data available was so vast and so interdependent...that it became more than the human mind could process."

In 1994, the UPS Airlines operations research (OR) group launched an effort to develop optimization-based tools to facilitate long-range planning of its expedited service offerings. We focused on next-day-air network planning because the long-range aircraft acquisition strategy is driven by the requirements of the next-day-air system. We wanted to create a planning system that would allow planners to apply their expertise to more scenarios and to shift their focus from manually creating a feasible plan to quickly generating (near-) optimal plans for the next-day-air network. In 1995, the OR group initiated a joint research project with MIT to develop optimization methods for simultaneously determining aircraft movements and package flows that would minimize aircraft ownership and operating costs while considering numerous operating constraints on system capacity and customer service standards. No tractable optimization methods existed for designing a network of this size and complexity. We faced a tremendous technical hurdle in developing an optimization-based approach. Moving the technology from a theoretical proof of concept to the planners' desks and gaining their acceptance was equally daunting.

Technical Obstacles and Modeling Achievement

To develop VOLCANO, we had to overcome major modeling and computational obstacles. The most difficult piece of the system from an optimization standpoint was the network-design component, in which we simultaneously determine the minimum-cost set of routes, fleet assignment, and package flows that satisfy constraints on various operating issues, including limits on the number of aircraft of each fleet type; landing restrictions at airports; aircraft operating characteristics, such as range, speed, and load capacity; continuous aircraft flow requirements (that is, balance of flow); time windows for pickup and delivery; and sorting capacities and hours of operation for each hub. In addition, packages must arrive at

the hubs in a staggered manner to spread the package volume across the entire sorting period. Finally, the next-day-air network must interface with the daytime aircraft requirements used in the second-day-air network. For the second-day network, the number of aircraft of a given type at a particular location is known, and these requirements serve as boundary conditions for our aircraft flows in the next-day-air network.

We developed and tested three approaches for solving this problem (Kim et al. 1999, Barnhart et al. 2002b). Early attempts relied upon conventional network design methods (Magnanti and Wong 1984, Gendron et al. 1999, Crainic 2000) that are useful on smaller versions of network-design problems. Solution methods fell primarily into two categories. The first involved cutting planes (Van Roy and Wolsey 1985, Magnanti et al. 1995, Günlük 1999, Stallaert 2000) and the second involved approximation algorithms (Hochbaum and Naor 1996, Karger 1999).

Conventional formulations for our air-network-design problem include two types of decision variables: those modeled with integer-valued decision variables to represent the aircraft-routing decisions and those modeled with continuous decision variables to represent package flows. Package volume along a flight leg depends upon the capacity assigned on each flight leg. Enforcing this type of constraint for all possible flight legs leads to tractability issues. Namely, integer-valued solutions for the aircraft decision variables are notoriously difficult to achieve for network-design formulations. The constraints that enforce the continuity of aircraft flows further exacerbate this issue by propagating fractional aircraft through the network. The end result is that for a problem the size of the UPS next-day-air network, conventional methods of formulating and solving network-design problems are ineffective for what we are trying to accomplish.

To underscore the difficulty of solving this problem, we formulated the problem for a very small portion of the network, that portion involving UPS's smallest hub in Columbia, South Carolina. The results were that even after multiple-day run times, we were able to generate only poor quality solutions. If we were to solve this problem, we needed a new idea to fundamentally advance the solution of network-design problems. This idea began to take form as we asked the following question: *Can we reformulate this problem in terms of the aircraft-routing decisions and capture the package flows implicitly in the new formulation?*

The answer was a resounding Yes! As we describe in more detail elsewhere (Armacost et al. 2002), we introduced composite variable formulations that overcame the computational obstacles. Such formulations are unlike typical mathematical-programming formulations in that additional constraints (even nonlinear)

can reduce the problem size and enhance tractability. By exploiting special structure to identify and limit the set of decision variables over which to optimize, we obtained computational properties in the reformulation that allow us to obtain solutions to UPS’s next-day-air network in minutes rather than not at all, as was the case using conventional network-design formulations.

The primary computational obstacle associated with conventional network-design formulations is the inherent weakness of their linear-programming (LP) relaxations. To facilitate the process of identifying integer solutions, we redefined the decision variables to produce a network design formulation whose LP relaxation is provably stronger. In our new formulation approach, we combine aircraft routes and package flows into a single integer-valued decision variable.

We can illustrate this concept with an example (Figure 2), in which we have two possible aircraft flying the route from an airport to the air hub. The first aircraft has a capacity of 4,000 packages, while the second has a capacity of 10,000. Our goal is to provide sufficient capacity on the route to carry the total demand of 5,000 packages from the airport to the hub. The decision variables in conventional network-design formulations are the number of each aircraft to fly from the airport to the hub and the package volume to flow on each of those aircraft. The solution to the LP relaxation will use either 1.25 of aircraft type 1 or half of aircraft type 2, and searching for the optimal integer solution requires additional work.

Instead, we define new variables, called composites, which combine the original aircraft and package-flow decision variables to provide sufficient, yet minimal, capacity to move the required volume from the airport to the hub. In this example, one composite variable combines two aircraft of the first type, and another uses a single aircraft of the second type. The composite-variable formulation then is required to pick at least one composite variable, thereby guaranteeing sufficient capacity for the demand moving

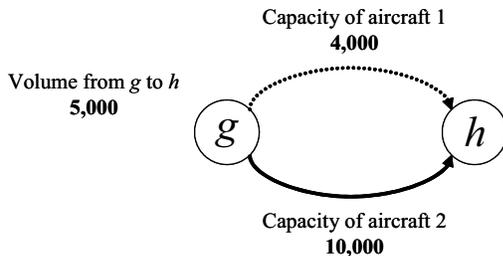


Figure 2: This two-location network consists of an airport, g , and an air hub, h . The objective is to move 5,000 packages from g to h using one of two aircraft types with different capacities.

from the airport to the hub. In the LP relaxation, the solution using only 1.25 of aircraft type 1 is not feasible because the composite variable is defined as two of aircraft type 1 and the value of the selected composite variables must be at least one. Similarly, the solution using half of aircraft type 2 is not feasible. The composite-variable formulation’s LP relaxation, therefore, removes some of the fractional solutions that exist in the conventional formulation’s LP relaxation, providing tighter bounds and requiring less work searching for an optimal or near-optimal integer solution.

Composites can be more complex than the simple combination of two aircraft routes (Figure 3). The first consists of two double-leg pickup routes bound for the same hub and stopping at the same intermediate location. Each plane has capacity to carry the entire package volume from its respective starting location. The residual capacity of the two planes, when combined, is sufficient to cover the package volume from the intermediate airport. In the second example, we have two planes converging on the same intermediate location but destined for different hubs. In this case, we might have capacity on the aircraft to transfer packages between planes at the intermediate location, thus allowing this combination of planes to cover even more gateway-to-hub package volumes.

In the composite-variable approach applied to the next-day-air network-design problem, we use time windows for pickup and delivery along with aircraft operating characteristics to generate all possible pickup and delivery routes (Figure 4). We also take advantage of predefined hub service territories, which include the airports whose packages will be sorted by each hub. Knowing the available capacity of the aircraft on a given route (which can depend upon the distance flown because of added fuel requirements) and the volume to be moved from the airports to the hub, we determine whether the aircraft route has sufficient capacity to carry all the airport-to-hub demand. If so, this aircraft route covers the demand

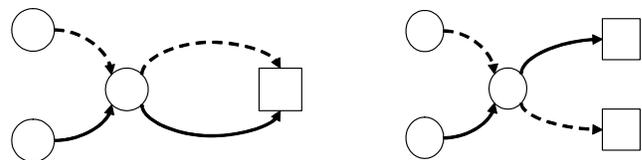


Figure 3: In these examples of composites, the left diagram shows two planes converging at the same intermediate stop en route to the hub (square). One aircraft is represented by the solid line, the other by the dashed line. On the right, we show a similar case in which two aircraft arrive at the intermediate location and travel to different hubs. If timing and capacity permit, packages can be transferred between aircraft at the intermediate location.

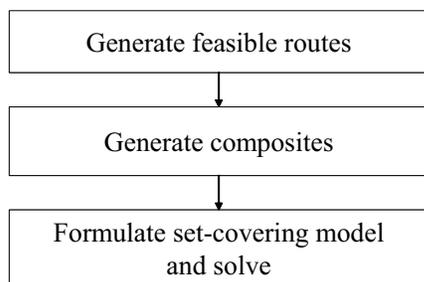


Figure 4: The composite-variable approach yields a set-covering formulation with appealing computational properties. It requires preparatory work to generate the feasible set of composite variables.

and is considered to be a composite. If not, we combine the aircraft route with capacity from other aircraft routes servicing some of the same airports. The composite is the combination of the aircraft routes and the corresponding flow of packages such that the airport-to-hub demands are covered. As we build these composites, we also look for opportunities to transfer packages from one aircraft to another, which depend on the timing of aircraft arrivals at common nonhub locations and the capacity of those aircraft.

Using composite variables, we solve a set-covering formulation in which we ensure a selection of composites that covers all origin-destination demands (Appendix). This formulation has several computational advantages, including improved LP bounds. Generating the set of composites requires additional work up front but typically takes only a few minutes. Moreover, we use the complexity of the planning problem to our advantage. We use the operating characteristics of the problem (for example, timing and aircraft capacity) to reduce the size of the candidate set of composite variables. The problem is very complex when we consider a network of seven hubs, over 100 airports, and more than 2,000 origin-destination pairs for package volume.

Early Test Results

Our early computational results for a small, single-hub example showed that the conventional network-design formulation requires much more effort than the equivalent composite-variable formulation to arrive at the optimal solution (Table 1). In this example, the LP relaxation of the conventional network-design formulation gives an optimal value that is 62.5 percent from the optimal integer value, requiring 781 nodes to be evaluated in the branch-and-bound tree. The composite-variable method, however, gives a tight LP relaxation producing the optimal integer solution at the root node of the branch-and-bound tree.

	Conventional Formulation	Composite-Variable Formulation
LP relaxation solution	10,663	28,474
IP solution	28,474	28,474
% difference (gap)	62.50	0
Branch-and-bound nodes	781	1

Table 1: For a simple, single-hub network, finding the optimal solution using a composite-variable formulation requires significantly less effort than with a conventional network-design formulation. The LP relaxation of the composite-variable formulation yields the optimal integer solution.

We obtain similar results when planning the entire next-day-air network. One scenario tested during the development phase included 101 airports, seven of which were hubs, and 160 aircraft available from seven fleet types. We conservatively estimated the nightly volume at 926,268 packages on the pickup side and 967,172 packages on the delivery side. Summary comparisons of our model’s solution with that of the planners (Table 2) revealed a seven percent reduction in operating cost, which translates into tens of millions of dollars in annual savings. The more significant savings come in ownership cost and the total number of aircraft required to operate the air network. For this scenario, the model achieved a 10.7 percent reduction in the number of aircraft required, or 16 fewer aircraft than originally planned. Moreover, the model required about two hours to run, compared to the six to nine months the planners took to design the next-day-air network.

In this example, the gap between the manual plan and the model’s plan is large. In practice, planners are unable to adopt an entire plan as suggested by the model. Instead, they have used VOLCANO to identify incremental changes to existing network plans. Such changes allow planners to reduce, but not close, this gap by implementing some, but certainly not all, of the suggested changes. As planners adopt more changes suggested by VOLCANO and recognize the economic potential, we expect UPS to realize an enormous payoff.

The model alters the fundamental design of the network. In their manual solutions, planners tend to create pickup and delivery routes that are mirror images.

	% Reduction Versus Planners’ Solution
Operating cost	7.00
Ownership cost	29.20
Number of aircraft	10.70

Table 2: The composite-variable formulation allows us to solve the complete next-day-air network-design problem. This allows significant reductions relative to the planners’ manual solution for operating cost, aircraft ownership cost, and number of aircraft required for the network.

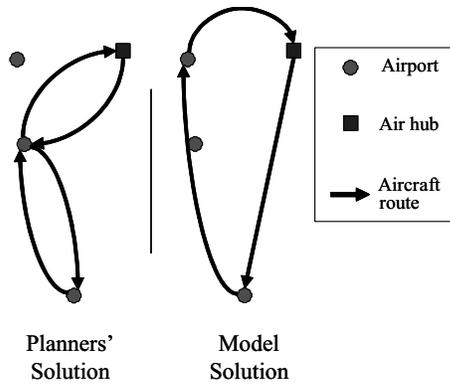


Figure 5: In this example, the planners' pickup route, originating at the southern airport, has one stop en route to the hub, and the delivery route from the hub mirrors the pickup route. The model tries to better match aircraft capacity with cargo volume, which results in asymmetry of the pickup and delivery routes. In this case, the pickup portion of the route has two pickup locations, originating at the southernmost airport. On the delivery side, the aircraft returns immediately to the originating airport, as the model has found a more efficient way to deliver the packages to the intermediate airports.

The model, however, tries to match the demand between the airports and hubs with the available capacity of the aircraft (Figure 5). Our approach also produces nonintuitive aircraft routes (Figure 6), often clever solutions that exploit time-zone changes and opportunities to transfer packages between planes at nonhub locations. Using methods to better match capacity to package volume, the model generates

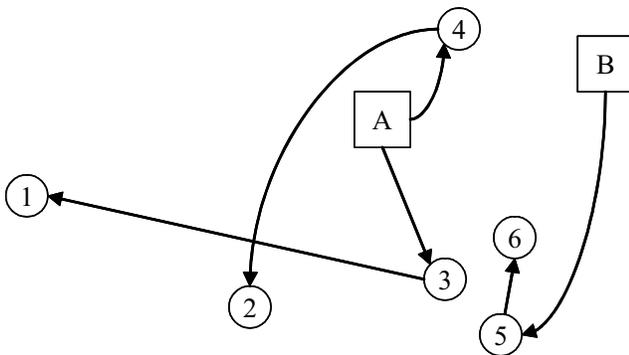


Figure 6: In these examples of nonintuitive delivery routes generated by the model, the model exploits time-zone changes and opportunities to transfer packages between planes at nonhub airports. The delivery route A-3-1 exploits the long westbound flight through time-zone changes to make up time after dropping off packages at airport 3. This allows the same aircraft to serve two locations, better utilizing its capacity. The backtracking in delivery route A-4-2 is done to ensure packages en route from A to another airport can be transferred at airport 4 to an aircraft en route to the other airport. In addition, the time-zone changes provide the extra time to allow this aircraft to satisfy the customer service delivery requirements at airport 2. A similar backtracking phenomenon is seen in route B-5-6, which uses extra capacity on its first leg to carry packages that will be transferred to another plane at airport 5.

solutions that rely on fewer aircraft and are less costly than manual solutions.

Organizational Obstacles

Although we overcame many technical hurdles, we faced equally challenging organizational hurdles throughout the project in gaining the acceptance of senior managers and the planning groups. Each year, senior managers, including the chief information officer and the chief financial officer of the airline, required us to justify continuing the UPS-MIT research effort. Although the research grant to MIT was only \$100,000 per year, the approval process was not simple. Our progress was monitored, and we had to document the details of our research. During the three years in which we struggled with difficult network-design aspects of the problem, the OR group actively championed the effort with senior management. The reputation of the researchers also carried weight in ensuring the continuation of the project. Most important, we maintained an open and honest dialogue with UPS managers about our progress, discussing both the ups and the downs. We took every opportunity to demonstrate empirically the payoff to the company, painting a picture of the system's possible uses and its organizational impact.

We also faced resistance from the planners. From the beginning of the effort, we sought the participation of network planners, peak planners, long-range planners, and financial analysts. Understandably, planners who had relied on their own expertise for years saw a threat in the new technology. We welcomed the skeptics over the life of the project and gradually earned their support in the annual review with senior management, which bolstered the acceptance of VOLCANO. We were also fortunate to gain the active participation of one planner with a background in OR. She worked with the OR group in validating the VOLCANO plans and served as an advocate by convincing fellow planners that they could become more effective using this new technology. Support for the system continued to build until the biggest skeptics became VOLCANO's strongest supporters. The system is now universally viewed as a useful planning tool.

Implementation

Implementation began in February 1999, and we rolled out the initial version of VOLCANO one year later. We continue to refine the system, with extensions forthcoming to cover additional air delivery service offerings, including second-day air, three-day select, and international delivery.

VOLCANO is more than an optimization model. It also includes modules that handle a variety of data,

including the operating characteristics of the aircraft, the time windows for customer service, the operating characteristics of the airports and hubs, and the volume to be moved through the air network. The core of the system—and the most difficult piece from an optimization standpoint—is the network-design problem solved by the composite-variable approach to produce an efficient set of aircraft fleet assignments, routes, and feasible package flows. Our modeling approach also ensures a feasible schedule for the aircraft routes. An additional module creates the precise timing required for the schedule.

The full-scale system required features to make it usable. To implement VOLCANO, we interfaced it with existing UPS databases, manipulated the data into usable form, built additional pre- and post-processing modules, and generated useful output for the planners. For the core analytical engine, we used the ILOG CPLEX Callable Library, a commercial off-the-shelf linear and integer programming solver. We used the C++ programming language to generate data structures and columns for the integer program. VOLCANO runs on a Hewlett Packard N4000 machine with the HPUNIX 11 operating system. Using a series of SQL calls and Perl scripts, we extract data from Oracle databases that reside on a separate DEC Alpha machine. We also developed a Visual Basic user interface that controls data input and output to the models and resides on the planner's personal workstation.

Impact

Since we delivered the initial version of VOLCANO in late 2000, it has changed the cost of operating the network and the way planners do their jobs. The three planning groups have accepted the system and routinely use it to support their planning processes. We categorize the system's widespread impact into five areas: quantifiable savings, potential future savings we cannot yet quantify, effects on the planning process, visibility of OR within UPS, and the portability of the modeling technology to a wide array of applications.

Quantifiable Cost Savings

The network planners and the peak planners have used VOLCANO since late 2000 to modify existing plans, looking for opportunities to reduce the costs of operating, lease, or ownership and to address operational problems in the network. Operating costs are primarily fuel and crew costs. Ownership cost is the cost of buying aircraft, often amortized over their lives. Lease costs are typically incurred during the peak season when UPS obtains additional capacity by leasing aircraft.

Operational network plan changes attributed to VOLCANO in 2002 included moving two planes from the network plan to the pool of spare aircraft, which reduced operating cost. The peak planners credit VOLCANO with downsizing three leased aircraft, replacing four UPS aircraft with two, and addressing several critical airlift-capacity problems. These changes reduced both operating and lease costs. Finally, in analyzing the 2005 plan, long-range planners used VOLCANO to defer the purchase of an aircraft for four years, saving both operating and ownership costs during that time. VOLCANO enables UPS to make such reductions in the number of aircraft by matching available capacity to the demand better than the manual planners. These published changes to the network, peak, and long-range plans reduce costs by more than \$87 million in present-day dollars. Using these early results as a baseline, planners and financial analysts have conservatively estimated savings in operating costs for the next decade to exceed \$189 million.

Future Savings

Planners typically use VOLCANO by taking an existing network plan as input and using the system to change the plan incrementally. These increments are the basis for official published changes to the network plan. UPS relies on this incremental approach for three reasons. First, VOLCANO is new technology that planners are just learning to use. Using the system for incremental changes helps the planning groups to accept it. Second, UPS had already published its network and peak plans for the years of initial use. Thus, planners were not willing to make dramatic changes to the existing plans. Third, UPS would need major resources and many personnel to execute the plans. It must make operational changes incrementally; it cannot support radical change during execution.

By using VOLCANO for incremental changes to existing network plans, we are realizing only a fraction of the potential savings. As planners use the system, they will become willing to accept larger changes to the network. Seventy-five percent of the \$87 million in cost savings between 2000 and 2002 are from changes published in 2002. By incrementally modifying the baseline plan over time, planners will eventually publish network plans close to VOLCANO's optimized solutions. It might take years for planners to achieve all potential savings. Because we based our estimates of savings over the next decade (\$189 million) on the initial incremental changes to existing plans, \$189 million represents only a portion of the savings in operating costs the airline can eventually realize.

Finally, the potential savings in operating costs are small relative to those in aircraft ownership costs. As

part of a recent long-range planning study, we executed 30 to 40 runs of VOLCANO to examine the ideal composition of the airline's future fleet. One unintended consequence of this study was a change in the way the long-range planners foresee planning future aircraft acquisitions. Ownership cost savings in this context is money UPS does not need to spend to acquire aircraft. The long-term effects can be enormous. With a single aircraft costing \$100 million (UPS Press Release 2001), the long-term savings in aircraft ownership might far exceed our estimate of savings in operating costs. In our early computational results, the model reduced by 16 the number of aircraft required to operate the 2001 plan.

Transforming the Planning Process

VOLCANO has become an integral part of the air-network-planning processes, changing the way planners plan and improving network, peak, and long-range plans. The system reduces the time it takes to evaluate a change to a plan, or to examine an operational problem, from months or weeks to days or hours. Using the system, planners can examine many scenarios for a single planning problem, whereas they could previously evaluate multiple scenarios only with difficulty, if at all. With VOLCANO's shortened planning cycle, planners can move from generating feasible plans to generating *optimal* plans.

The system is changing the nature of all planning groups' analyses. For example, peak planners used VOLCANO to find alternatives when three aircraft of a particular type were not available for lease. Prior to VOLCANO, they would have settled for leasing other aircraft at great cost to obtain equivalent or greater lift and accepted doing so as the cost of doing business. Instead, the planners and OR analysts used VOLCANO to modify routes and reassign aircraft to handle the demand at lower cost than that of the original plan.

In the wake of VOLCANO's implementation, planners are beginning to think differently about aircraft routing because of the routes generated by the system. When planners encounter resource problems similar to those they have previously analyzed with VOLCANO, they use the system-generated techniques for altering package movements and rerouting or reassigning aircraft without actually using VOLCANO to do so. For example, before they used the system, the planners would not have considered a route with an air hub as an intermediate stop. Routes originated or terminated at hubs, but did not pass through them; yet VOLCANO consistently designed such routes. In January 2002, network planners implemented a route of this nature without using VOLCANO.

Finally, planners can challenge the assumptions and methods they use in manual planning. They use

VOLCANO to examine scenarios from scratch, ignoring conventional ideas about how to solve them. By analyzing multiple scenarios, they establish the difference in cost between two solutions and quantify the effects of changing various operating practices. According to Mike Keenan, network next-day-air planning manager, "If we were able to change departure times at just one or two gateways (airports), significant savings could be realized. Before, we couldn't convince a gateway (airport) to change times because we couldn't quantify it." VOLCANO allows planners to quantify such changes.

Overall, planning groups have adopted VOLCANO as a critical element of what they do. Here is what they've had to say:

VOLCANO now gives us an opportunity to take a fresh look—a completely creative look—at how to potentially reroute airplanes or make changes to the year-round network that continue on through peak season.

Gary Graves, peak planning manager

There is no doubt that the VOLCANO product is an integral part of the planning process and a required tool for analyzing operations and developing acquisition strategies.

Rodger McLaughlin, long-range-planning manager

Project Visibility and Technology Portability

This project has been the subject of many write-ups in the airline's internal newsletter and has been mentioned in articles in *Information Week* (Barrett 2000) and *Business Week* (McGuire 2000). The savings generated by the system for the 2000 peak delivery season were mentioned at a quarterly UPS board of directors meeting, and they were used by the airline's president in his speech to investor analysts. The company's priority is to be a leader in technology applications, and the VOLCANO system has become a showcase project. The new system has had a dramatic impact on planning and operating the air network, giving the OR group great visibility within the airline and among senior decision makers.

The technology we developed is now used in other large-scale, complex planning-and-scheduling applications. We have successfully applied composite-variable formulations, the core element of the VOLCANO system, to extensions of the next-day-air network-design problem, including the second-day air and international networks. As part of an effort with UPS Service Parts Logistics, Cohn and Barnhart (2002) have used composite-variable models to determine the optimal warehousing and distribution of parts.

Beyond UPS, practitioners are using composite-variable methods to solve other large-scale planning problems. Passenger airlines face two such problems,

one to simultaneously determine fleet assignments and passenger flows and the other to simultaneously determine optimal crew schedules and aircraft routings. Previously, models designed to solve these problems were slow and often intractable. Composite-variable models greatly enhanced solvability and produced improved solutions over those generated using conventional approaches (Barnhart et al. 2002a, Cohn and Barnhart 2003).

Composite-variable formulations have also been developed for problems of large-scale mobility and logistics planning for the US Department of Defense. Of note are applications to intermodal movement of munitions by land, sea, and air (Clark et al. forthcoming); real-time air targeting within dynamic, hostile environments (Barth 2001); and designs of the peacetime network of aircraft for moving cargo between the United States and overseas locations (Nielsen et al. forthcoming). Outside of transportation, researchers at MIT are exploring the use of composite-variable methods for telecommunications planning. We laid the foundations for these efforts with our technical achievements on the UPS next-day-air network-design problem.

Concluding Remarks

UPS has saved over \$87 million from 2000 through 2002, and planners estimate UPS's savings over the next decade at \$189 million. The long-term impact might far exceed these estimates, as changes produced by VOLCANO affect aircraft acquisition decisions.

Tom Weidemeyer, the president of UPS Airlines and the chief operating officer of UPS, has summarized the overall future impact of our effort: "These changes in our planning process mean that VOLCANO will continually improve our operating plans and will affect the number of aircraft we will buy or retire as our company evolves. The potential cost savings associated with these moves are huge. More importantly, VOLCANO will make UPS more competitive by ensuring that we are running the most efficient network possible. We at UPS have good reason to believe in the power of operations research."

Appendix: Formulations

Conventional Formulation for Next-Day-Air Network Design

Sets

- F = set of fleet types.
- R^f = set of routes flown by fleet type, $f \in F$.
- A = set of all arcs in time-space network.
- G = set of gateways (airports).

H = set of hubs.

R_p = set of pickup routes.

R_D = set of delivery routes.

$R_p(\underline{h})$ = set of pickup routes terminating at the hub.

$R_p(\bar{g})$ = set of pickup routes originating at gateway g .

$R_D(\underline{g})$ = set of delivery routes terminating at gateway g .

$R_D(\bar{h})$ = set of delivery routes originating at hub h .

K = set of origin-destination commodities.

P^k = set of paths for commodity k .

Indicators

$\delta_{ij}^p = 1$ if path p includes arc (i, j) , 0 otherwise.

$\delta_{ij}^r = 1$ if route r includes arc (i, j) , 0 otherwise.

$\delta_h^p = 1$ if path p passes through hub h , 0 otherwise.

Data

d_r^f = cost of flying route r with fleet type f .

b^k = package demand associated with origin-destination commodity k .

u_r^f = capacity of aircraft type f flying route r .

e_h = package capacity of hub h .

n_f = number of aircraft of type f .

a_h = number of aircraft that can land at hub h .

Decision Variables

y_r^f = number of aircraft of type f assigned to fly route r .

x_p^k = fraction of commodity k 's demand flown along path p .

$$\begin{aligned} & \min \sum_{f \in F} \sum_{r \in R^f} d_r^f y_r^f \\ & \text{subject to } \sum_{k \in K} \sum_{p \in P^k} \delta_{ij}^p b^k x_p^k \leq \sum_{f \in F} \sum_{r \in R^f} \delta_{ij}^r u_r^f y_r^f \\ & \hspace{20em} \text{for all } (i, j) \in A, \\ & \sum_{p \in P^k} x_p^k = 1 \quad \text{for all } k \in K, \\ & \sum_{k \in K} \sum_{p \in P^k} \delta_h^p b^k x_p^k \leq e_h \quad \text{for all } h \in H, \\ & \sum_{r \in R_p^f(\bar{g})} y_r^f - \sum_{r \in R_D^f(\underline{g})} y_r^f = 0 \quad \text{for all } g \in G, f \in F, \\ & \sum_{r \in R_D^f(\bar{h})} y_r^f - \sum_{r \in R_p^f(\underline{h})} y_r^f = 0 \quad \text{for all } h \in H, f \in F, \\ & \sum_{r \in R_p^f} y_r^f \leq n_f \quad \text{for all } f \in F, \\ & \sum_{f \in F} \sum_{r \in R_p^f(\underline{h})} y_r^f \leq a_h \quad \text{for all } h \in H, \\ & x_p^k \geq 0 \quad \text{for all } p \in P^k, k \in K, \\ & y_r^f \geq 0 \quad \text{and integer for all } r \in R^f, f \in F. \end{aligned}$$

Composite Variable Formulation for Next-Day-Air Network Design

Sets

C = set of all composites.

C_P = set of composites constructed from pickup routes.

C_D = set of composites constructed from delivery routes.

Indicators

$\delta_c^{gh} = 1$ if composite c covers the demand between gateway g and hub h , 0 otherwise.

$\delta_{ij}^r = 1$ if route r includes arc (i, j) , 0 otherwise.

$\delta_h^p = 1$ if path p passes through hub h , 0 otherwise.

Data

γ_c^f = number of aircraft of type f included in composite c .

$\gamma_c^f(\bar{g})$ = number of aircraft of type f included in composite c originating at airport g .

$\gamma_c^f(\underline{g})$ = number of aircraft of type f included in composite c terminating at airport g .

$\gamma_c^f(\bar{g})$ and $\gamma_c^f(\underline{g})$ defined similarly for hubs.

d_c = cost of all aircraft routes in composite c .

b_p^{gh} = pickup demand volume from gateway g to hub h .

b_D^{gh} = delivery demand volume from hub h to gateway g .

Decision Variables

$v_c = 1$ if composite c is selected, 0 otherwise.

$$\min \sum_{c \in C} d_c v_c$$

$$\text{subject to } \sum_{c \in C_P} \delta_c^{gh} v_c \geq 1 \quad \text{for all } (g, h): b_p^{gh} > 0,$$

$$\sum_{c \in C_D} \delta_c^{gh} v_c \geq 1 \quad \text{for all } (g, h): b_D^{gh} > 0,$$

$$\sum_{c \in C_P} \gamma_c^f(\bar{g}) v_c - \sum_{c \in C_D} \gamma_c^f(\underline{g}) v_c = 0$$

$$\text{for all } g \in G, f \in F,$$

$$\sum_{c \in C_P} \gamma_c^f(\underline{h}) v_c - \sum_{c \in C_D} \gamma_c^f(\bar{h}) v_c = 0$$

$$\text{for all } h \in H, f \in F,$$

$$\sum_{c \in C_P} \gamma_c^f v_c \leq n_f \quad \text{for all } f \in F,$$

$$\sum_{f \in F} \sum_{c \in C_P} \gamma_c^f(\underline{h}) v_c \leq a_h \quad \text{for all } h \in H,$$

$$v_c \in \{0, 1\} \quad \text{for all } c \in C.$$

Acknowledgments

The success of this project is the direct result of a companywide commitment. Many people have been involved in ensuring the company's long-term commitment to the development and implementation of VOLCANO. They include Jerry Delgaudio, Roger Dingus, Phil Hunter, Linda MacLean, Sue Ratkowski, Greg Reinhardt, Jim Stickler, Larry Subacz, and, from the airline's finance group, Scott Childress, Mike Jones, and Jeff McCorstin. For the active support we have received from planning groups, we thank Scott Abell, Chris Beilefeld, Robert Brightwell, Paul Contois, Gary Graves, Bob Gunderson, Mike Keenan, Rodger McLaughlin, Bill Simpson, and Mark Perkins. The OR analysts and programmers who made incredible contributions during implementation include Chi-Yin Cheng, Wynand DePuy, Melody Doda, Kurt File, Kay Kessler, Melanie Minter, Lee Mitchell, Jerry Shaw, and Bob White. MIT graduate students providing invaluable contributions include Daeki Kim and Niranjana Krishnan. We also express our most sincere appreciation to our coaches in the Edelman Prize competition, Randy Robinson and Peter Norden. Finally, the first author, as a major in the US Air Force, is required to say that this paper does not represent the official policy of the US Air Force or the Department of Defense.

References

- Armacost, Andrew P., Cynthia Barnhart, Keith A. Ware. 2002. Composite variable formulations for express shipment service network design. *Transportation Sci.* 36(1) 1–20.
- Barnhart, Cynthia, Amr Farahat, Manoj Lohatepanont. 2002a. Airline fleet assignment: An enhanced revenue model. Working paper, Massachusetts Institute of Technology, Cambridge, MA.
- Barnhart, Cynthia, Niranjana Krishnan, Daeki Kim, Keith A. Ware. 2002b. Network design for express shipment delivery. *Comput. Optim. Appl.* 21(3) 239–262.
- Barrett, Alexandra. 2000. Optimize the enterprise. *Information Week Online*. Retrieved June 5 www.informationweek.com/789/optimize.htm.
- Barth, Christopher D. 2001. Composite mission variable formulation for real-time mission planning. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.
- Clark, Steven J., Cynthia Barnhart, Stephan E. Kolitz. 2004. Large-scale optimization planning methods for the distribution of United States Army munitions. *Math. Comput. Modeling*. Forthcoming.
- Cohn, Amy M., Cynthia Barnhart. 2002. Composite-variable modeling for service parts logistics. Working paper, Massachusetts Institute of Technology, Cambridge, MA.
- Cohn, Amy M., Cynthia Barnhart. 2003. Improving crew scheduling by incorporating key maintenance routing decisions. *Oper. Res.* 51(3) 387–396.
- Crainic, Teodor G. 2000. Service network design in freight transportation. *Eur. J. Oper. Res.* 122(2) 272–288.
- Gendron, Bernard, Teodor G. Crainic, Antonio Frangioni. 1999. Multicommodity capacitated network design. B. Sanso, P. Soriano, eds. *Telecommunications Network Planning*. Kluwer Academic Publishers, Dordrecht, The Netherlands, 1–19.
- Günlük, Oktay. 1999. A branch-and-cut algorithm for capacitated network design problems. *Math. Programming* 86(1) 17–39.
- Hochbaum, Dorit S., Joseph Naor. 1996. Approximation algorithms for network design problems. *J. Algorithms* 21(2) 403–414.
- Karger, David R. 1999. Random sampling in cut, flow, and network design problems. *Math. Oper. Res.* 24(2) 383–413.

- Kim, Daeki, C. Barnhart, Keith A. Ware, G. Reinhardt. 1999. Multimodal express package delivery: A service network design approach. *Transportation Sci.* **33**(4) 391–407.
- Magnanti, Thomas L., Richard T. Wong. 1984. Network design and transportation planning: Models and algorithms. *Transportation Sci.* **18**(1) 1–55.
- Magnanti, Thomas L., Prakash Mirchandani, Rita Vachani. 1995. Modeling and solving the two-facility capacitated network loading problem. *Oper. Res.* **43**(1) 142–157.
- McGuire, Patrick. 2000. Planning made simpler. *Business Week* (October 30) 174B–174D.
- Nielsen, Christopher A., Andrew P. Armacost, Cynthia Barnhart, Stephan E. Kolitz. 2004. Network design formulations for scheduling U.S. Air Force channel route missions. *Math. Comput. Modeling*. Forthcoming.
- Stallaert, Jan. 2000. Valid inequalities and separation for capacitated fixed charge flow problems. *Discrete Appl. Math.* **98**(3) 265–274.
- UPS Press Release. 2001. UPS Airlines places order for 60 Airbus freighters. UPS Press Room, pressroom.ups.com.
- Van Roy, Tony J., Laurence A. Wolsey. 1985. Valid inequalities and separation for uncapacitated fixed charge networks. *Oper. Res. Lett.* **4**(3) 105–112.